THÈSE PRÉSENTÉE

POUR OBTENIR LE GRADE DE

DOCTEUR
DE L'UNIVERSITÉ DE BORDEAUX

ÉCOLE DOCTORALE DE MATHÉMATIQUES ET INFORMATIQUE

Spécialité : Mathématiques pures

Par **Pierrick DARTOIS**

Calcul rapide d'isogénies en dimension supérieure pour les applications cryptographiques

Fast computation of higher dimensional isogenies for cryptographic applications

Sous la direction de : **Damien ROBERT**
Co-directeur : **Benjamin WESOLOWSKI**
Co-encadrant : **Luca DE FEO**

Soutenance prévue le 9 juillet 2025

Membres du jury :

M. Damien ROBERT	Directeur de Recherche	Université de Bordeaux	Directeur
M. Benjamin WESOLOWSKI	Chargé de Recherche	ENS de Lyon	Co-directeur
M. David KOHEL	Professeur	Aix-Marseille Université	Rapporteur
M. David LUBICZ	Directeur Scientifique (DGA)	Université de Rennes 1	Rapporteur
M. Pierrick GAUDRY	Directeur de Recherche	Université de Lorraine	Examinateur
Mme. Sabrina KUNZWEILER	Inria Starting Faculty Position	Université de Bordeaux	Examinatrice
Mme. Elisa LORENZO-GARCIA	Maîtresse Assistante	Université de Neuchâtel	Examinatrice
M. Frederik VERCAUTEREN	Professor	KU Leuven	Examinateur
M. Luca DE FEO	Research Staff Member	IBM Research Europe	Invité

Calcul rapide d'isogénies en dimension supérieure pour les applications cryptographiques

Résumé : Shor a découvert en 1995 un algorithme permettant à un ordinateur quantique suffisamment puissant d'attaquer tous les protocoles cryptographiques à clés publiques fondés sur le logarithme discret et la factorisation des nombres en produit de facteurs premiers, tels que RSA et les courbes elliptiques, largement utilisés aujourd'hui. Depuis, de gros efforts de recherche ont été accomplis pour proposer des protocoles résistants aux attaques quantiques et l'institut de standardisation et des technologies américain (NIST) a organisé deux compétitions internationales en ce sens. La cryptographie à base d'isogénies repose sur la difficulté à trouver des isogénies entre courbes elliptiques. En 2022, le protocole d'échange de clé Supersingular Isogeny Diffie-Hellman (SIDH) proposé au NIST fait l'objet d'attaques majeures. Loin d'entraver l'avenir de la cryptographie à base d'isogénies, ces attaques ont au contraire fortement dynamisé la recherche dans ce domaine. Les idées de ces attaques très efficaces ont en effet inspiré de nouveaux protocoles cryptographiques et l'amélioration de protocoles existants ne souffrant pas des faiblesses de sécurité de SIDH. Nos travaux de thèse ont notamment contribué à des améliorations de Short Quaternion Isogeny Signatures (SQIsign), un protocole de signature électronique également proposé au NIST. SQIsign s'appuie sur la correspondance de Deuring entre isogénies des courbes elliptiques supersingulières et idéaux d'une algèbre de quaternions. Bien que très compact, SQIsign était désavantagé par un algorithme de signature très lent consistant à traduire un idéal d'une algèbre de quaternions en l'isogénie qui lui correspond. Les attaques contre SIDH ont mené à de nouvelles idées pour améliorer cet algorithme de traduction de manière beaucoup plus efficace, tout en améliorant la preuve de sécurité et la compacité du protocole. Ces nouveaux algorithmes de traduction d'idéaux en isogénies ont aussi servi à améliorer le calcul de l'action du groupe des classes d'idéaux sur les courbes elliptiques supersingulières orientées, qui intervient par exemple dans le protocole Commutative Supersingular Isogeny Diffie Helmman (CSIDH). Les attaques contre SIDH et les nouvelles techniques de construction qui ont suivi s'appuient sur le calcul efficace d'isogénies en dimension supérieure à 2 (les courbes elliptiques étant de dimension 1). Cette thèse a ainsi contribué à l'élaboration et à l'implémentation d'algorithmes rapides de calcul d'isogénies en dimensions 2 et 4.

Mots-clés : cryptographie post-quantique, isogénies, isogénies en dimension supérieure, SQISign, lemme de Kani, actions de groupes cryptographiques.

Unité de recherche
Université de Bordeaux, CNRS, INRIA, IMB, UMR 5251, F-33400 Talence, France.

Fast computation of higher dimensional isogenies for cryptographic applications

Abstract: In 1995, Shor discovered an algorithm that would enable a sufficiently powerful quantum computer to attack all public-key cryptographic protocols based on discrete logarithm and prime factorisation, such as RSA and elliptic curves that are widely used today. Since then, major research efforts have been made to propose protocols that are resistant to quantum attacks, and the US National Institute of Standards and Technology (NIST) has organised two international competitions to standardise these quantum-resistant protocols. Isogeny-based cryptography is based on the difficulty of finding isogenies between elliptic curves. In 2022, the Supersingular Isogeny Diffie-Hellman (SIDH) key exchange protocol proposed to the NIST competition came under major attack. Far from hindering the future of isogeny-based cryptography, these attacks gave a major boost to research in this field. The ideas behind these highly effective attacks actually inspired new cryptographic protocols and improvements to existing ones that do not suffer from SIDH's security weaknesses. Our works contributed to improvements of Short Quaternion Isogeny Signatures (SQIsign), a digital signature scheme also proposed to the NIST competition. SQIsign is based on the Deuring correspondence between isogenies of supersingular elliptic curves and ideals in a quaternion algebra. Although very compact, SQIsign was disadvantaged by a very slow signature algorithm, which consists in translating an ideal of a quaternion algebra into its corresponding isogeny. Attacks on SIDH led to new ideas to improve this translation algorithm, while enhancing the protocol's security proof and compactness. These new algorithms for ideal-to-isogeny translation were also used to improve the computation of the ideal class group action on oriented supersingular elliptic curves, involved for instance in the Commutative Supersingular Isogeny Diffie Helmman (CSIDH) protocol. Attacks on SIDH and the new construction techniques that followed are based on the efficient calculation of isogenies in dimension greater than 2 (elliptic curves being of dimension 1). This thesis contributed to the design and implementation of fast algorithms to compute isogenies in dimensions 2 and 4.

Keywords: post-quantum cryptography, isogenies, higher dimensionsl isogenies, SQIsign, Kani's lemma, cryptographic group actions.

À maman, qui a été empêchée de me voir devenir docteur.

À mon père, pour son infaillible soutien pendant mes onze années d'études.

Remerciements - Thanks

Je tiens tout d'abord à remercier vivement mes encadrants de thèse. Merci à vous trois pour votre soutien constant tout au long des trois dernières années et pour les idées lumineuses dont vous m'avez fait part, sans lesquelles ce travail n'aurait pu aboutir. Toute thèse doit inévitablement beaucoup à l'aptitude des encadrants à créer un environnement propice à la progression et à l'épanouissement de leur étudiant. C'est une mission incontestablement réussie, malgré l'éclatement géographique.

Tout d'abord, merci Luca, car c'est grâce à toi que cette belle aventure a commencé. Merci de m'avoir introduit dans la communauté des isogénistes, et d'avoir bravé les épreuves bureaucratiques à mes débuts. C'est à tes suggestions judicieuses que je dois également des connaissances bien utiles et pratiques en programmation (les *makefiles* et les *command line interfaces* m'ont vraiment simplifié la vie). Tu n'as pas non plus perdu patience pour faire progresser mes présentations sur le plan de la clarté et de la concision. Je garderai toujours un très bon souvenir des randonnées que nous avons faites dans les montagnes suisses et des après-midi que j'ai passées à Zurich pour monologuer sur les fonctions thêta devant un public de jeunes isogénistes.

Merci Damien pour m'avoir transmis avec patience des idées théoriques ardues qui ont forgé mon savoir-faire mathématique, ont largement contribué à mes résultats et continueront certainement à ouvrir des voies de recherche fructueuses pour la suite de ma carrière. Les fonctions thêta jouent un rôle central dans ce manuscrit et il aurait été impossible de comprendre la théorie sous-jacente sans ton aide, tes références, suggestions, explications et corrections. Dernier point essentiel, tu as toujours su te montrer bienveillant et jamais condescendant, même face à mon ignorance naïve en quelques occasions.

Merci Benjamin d'avoir montré l'exemple de ce que doit être un travail de recherche de grande qualité, tant sur le fond que sur la forme. Outre la profondeur des résultats qu'ils présentent, tes papiers sont unanimement reconnus comme bien écrits, et tes présentations comme très claires. Les quelques compliments que mes travaux ont reçus doivent beaucoup à ton exemple que je continue à admirer et à tes exigences en matière de clarté et de rigueur qui m'ont largement fait progresser. Je te remercie de m'avoir donné le goût de « la bonne science ».

Antonin, avec le nombre incalculable de réunions hebdomadaires auxquelles tu as assisté, tu mériterais certainement de figurer parmi mes encadrants. Cela a été un grand plaisir de collaborer avec toi sur SQIsign, et j'espère pouvoir continuer à le faire à l'avenir. Tes astuces algorithmiques lumineuses sur les quaternions se sont révélées fort utiles à plusieurs reprises et tu m'as également bien inspiré pour expliquer la correspondance de Deuring de façon claire dans mes présentations.

I am also very grateful to all my other co-authors. Thank you Sarah Arpin, Andrea Basso, James Clements, Tako Boris Fouotsa, Arthur Herlédan Le Merdy, Riccardo Invernizzi, Jonathan Komada Eriksen, Péter Kutas, Luciano Maino, Giacomo Pope, Ryan Rueger and Frederik Vercauteren. It has been a pleasure to work with you over the course of this PhD.

I would also like to thank Pierrick Gaudry, David Kohel, Sabrina Kunzweiler, Elisa Lorenzo Garcia, David Lubicz and Federik Vercauteren for your participation in my PhD jury. Merci en particulier à David Kohel et David Lubicz pour votre travail de rapporteur, vos rapports très positifs et vos corrections fort utiles.

My experience as a PhD student owes a lot to the isogeny community, who have been very dynamic, kind and welcoming to me. I would like to thank in particular Maria, Jonathan and Krijn for all their organisational efforts with The Isogeny Club.

Si rétrospectivement, la thèse a été une période heureuse de ma vie, c'est aussi en grande partie grâce à mon équipe d'accueil, l'équipe Inria CANARI à l'Institut de Mathématiques de Bordeaux. Car travailler dans un environnement agréable est essentiel au bien-être. Merci aux permanents de

l'équipe qui m'ont très bien accueilli. J'ai une pensée particulière pour Andreas, Aurel, Alice, Bill, Henri, Elena, Razvan (et Damien bien évidemment). Merci également aux irréductibles doctorants du bureau 319, en particulier Fabrice, Jean et Nicolas, pour tout ce que nous avons partagé ensemble au travail comme en dehors, allant des discussions animées aux sorties cinéma, en passant par les pâtisseries *vegan*. Fabrice, la pause thé hebdomadaire n'aurait pas été la même sans le *soft power* de tes gâteaux sans produit d'origine animale. Je n'oublie pas non plus les autres doctorants que j'ai eu le plaisir de connaître : Agathe, Elie, Fabrice (l'autre Fabrice), Alfonso, Guilhem, Wouter et sans doute d'autres que j'oublie (mille excuses). I am also grateful to Sabrina, Wessel and Marc. It has been a pleasure to get to know you.

I would also like to thank all those who contributed one way or another to the pursuit of my research in the near future. Thank you Pierre-Alain Fouque, Aurore Guillevic, André Schrottenloher, Patrick Gros, Emmanuel Thomé, Jean Kieffer, Xavier Bonnetain, Pierre-Jean Spaenlehauer, Pierrick Gaudry, Frederik Vercauteren, Wouter Castryck, Péter Kutas, David Lubicz, and of course my PhD advisors. May your efforts come to fruition (very soon I hope...).

Je remercie finalement les organismes qui ont financé ma recherche, en premier lieu Inria, mais également l'ANR CIAO (ANR-19-CE48-0008) et le PEPR PQ-TLS sous l'égide du programme France 2030 (ANR- 22-PETQ-0008).

Table des matières

Introduction et résumé substantiel en français

Comme le prévoit l'article 21 alinéa 2 de l'arrêté du 25 mai 2016 fixant le cadre national de la formation et les modalités conduisant à la délivrance du diplôme national de doctorat, nous introduisons et résumons ici en français les principales contributions de cette thèse portant sur la cryptographie post-quantique à base d'isogénies. Le reste du document est écrit en anglais et l'introduction qui suit est également traduite en anglais.

La menace quantique

Au-delà des communications militaires et des actifs financiers de quelques technophiles enthousiastes, la cryptographie est essentielle à notre quotidien numérique. Il est impossible de naviguer sur internet, de faire une transaction bancaire ou de signer un document en ligne de manière sécurisée sans utiliser de protocole cryptographique. Par opposition aux protocoles *symétriques*, les protocoles *asymétriques* (dits aussi *à clés publiques*) utilisent des clés différentes à l'usage des parties impliquées, l'une étant publique et l'autre secrète. Les protocoles asymétriques permettent par exemple de signer des documents numériques ou encore d'échanger une clé secrète entre deux parties de manière sécurisée pour l'utiliser ensuite dans un protocole de chiffrement symétrique. Aujourd'hui, la sécurité de l'intégralité des protocoles asymétriques utilisés actuellement repose sur l'un des deux problèmes suivants.

Problème 1 (Factorisation RSA)**.** *Étant donné un entier $N := p \cdot q$, avec deux facteurs premiers p et q inconnus de taille comparable, trouver p et q.*

Problème 2 (Logarithme discret)**.** *Étant donnés un groupe cyclique d'ordre premier $G = \langle g \rangle$ de générateur g connu et un élément quelconque $h \in G$, trouver $x \in \{0, \cdots, \#G - 1\}$ tel que $h = g^x$.*

Ces deux problèmes sont réputés difficiles pour un ordinateur classique, les meilleures attaques génériques connues contre RSA et le logarithme discret étant respectivement de complexité sous exponentielle [BLP93] et exponentielle [Sha71] en la taille des objets ($\log(N)$ ou $\log(\#G)$). L'algorithme de Shor [Sho97] permet cependant de les résoudre en temps polynomial sur un ordinateur quantique. Par précaution, en prévision des progrès futurs de l'ordinateur quantique, d'importants efforts de recherche sont consacrés pour trouver des protocoles construits sur des problèmes mathématiques résistants aux attaques quantiques, que l'on appelle communément *protocoles post-quantiques*.

Selon les problèmes sous-jacents, les protocoles cryptographiques post-quantiques sont rangés par famille. Ainsi, la cryptographie à base de réseaux repose sur la difficulté à trouver, dans un réseau euclidien de grande dimension, des vecteurs courts ou des vecteurs proches d'un vecteur quelconque [Ajt98 ; DKRS03]. La cryptographie à base de codes repose sur la difficulté à décoder un message bruité auquel on a appliqué un code correcteur d'erreur linéaire [BMT78]. La cryptographie multivariée repose sur la difficulté à résoudre des systèmes d'équations polynomiales à plusieurs variables. La cryptographie à base d'isogénies, qui fait l'objet de cette thèse, repose sur la difficulté à trouver une isogénies entre deux courbes elliptiques. Parmi les alternatives, la cryptographie à base d'isogénies se distingue par sa compacité (les tailles de clés ou de signatures sont faibles) et par sa relative lenteur.

L'institut américain de standardisation et des technologies (NIST) a lancé deux compétitions pour standardiser des protocoles post-quantiques afin d'anticiper leur déploiement industriel. La première compétition a permis de standardiser deux protocoles de transport de clé et deux protocoles de

signatures électroniques, dont la moitié à base de réseaux euclidiens (voir Tableau 1). Le protocole d'échange de clés SIDH [JDF11] (Supersingular Isogeny Diffie Hellman) a survécu jusqu'au quatrième et dernier tour de la compétition mais a fait l'objet d'une attaque (classique) très rapide qui a eu d'importantes conséquences (positives) sur la cryptographie a base d'isogénies. Le NIST a lancé une deuxième compétition pour standardiser d'autres signatures électroniques compactes, rapides à vérifier et reposant sur des hypothèses de sécurité diversifiées[1], laissant donc une chance aux isogénies. Le protocole à base d'isogénies SQIsign [DFKLPW20] (Short Quaternion Isogeny Signature), qui n'est pas sujet à l'attaque contre SIDH (et en a même paradoxalement bénéficié), a été retenu pour le deuxième tour de cette compétition. C'est aujourd'hui un candidat sérieux pour un déploiement industriel que les travaux de cette thèse ont contribué à améliorer.

Fonction	Nom	Famille
Transport de clé	Crystals-Kyber HQC*	Réseaux euclidiens Codes correcteurs
Signature électronique	Crystals-Dilithium Sphincs+	Réseaux euclidiens Fonctions de hachage

TABLE 1 : Schémas cryptographiques standardisés ou voués à être standardisés à l'issue de la première compétition post-quantique du NIST.
*Standardisation en cours.

Enjeux de la cryptographie à base d'isogénies

Le problème de l'isogénie

Les courbes elliptiques sont des courbes projectives de dimension 1 d'équation affine de la forme $y^2 = x^3 + Ax + B$. Si E est une courbe elliptique définie sur un corps k, alors l'ensemble $E(k)$ des points k-rationnels de E est dotée d'une structure de groupe abélien dont la loi de groupe est algébrique et efficace à calculer. Lorsque k est fini, le groupe $E(k)$ est d'ordre fini. Pour cette raison, les courbes elliptiques définies sur les corps fini ont largement été utilisées pour construire des schémas cryptographiques à base de logarithme discret.

Les isogénies sont des morphismes non-nuls entre courbes elliptiques en tant que morphismes de groupes et en tant que morphismes de variétés algébriques, c'est-à-dire pouvant s'exprimer à l'aide de fractions rationnelles. La cryptographie à base d'isogénies repose sur la difficulté à résoudre le problème suivant (ou des problèmes proches) pour un ordinateur classique et quantique.

Problème 3 (Problème de l'isogénie)**.** *Étant données deux courbes elliptiques E_1 et E_2 définies sur un corps fini $\mathbb{F}_q$, trouver une isogénie $\varphi : E_1 \longrightarrow E_2$.*

La difficulté de ce problème a été largement étudiée et dépend de la structure de l'anneau des endomorphismes des courbes considérées. Étant donnée une courbe elliptique E, l'anneau des endomorphismes $\mathrm{End}(E)$ est constitué des isogénies $E \longrightarrow E$ et du morphisme nul. Si E est défini sur un corps fini, on distingue deux cas [Sil09, Theorem V.3.1] :

- $\mathrm{End}(E)$ est isomorphe à un ordre quadratique imaginaire, auquel cas E est dite *ordinaire* ;
- $\mathrm{End}(E)$ est isomorphe à un ordre maximal dans une algèbre de quaternions, auquel cas E est dite *supersingulière*.

Une courbe ordinaire ne peut pas être isogène à une courbe supersingulière, étant donné qu'une isogénie transporte en partie la structure de l'anneau des endomorphismes d'une courbe à l'autre. Si E_1 et E_2 sont ordinaires et que leurs anneaux d'endomorphismes sont isomorphes $\mathrm{End}(E_1) \simeq \mathrm{End}(E_2)$, alors on sait qu'il existe un algorithme quantique capable de trouver une isogénie entre E_1 et E_2 en temps sous-exponentiel en $\log(q)$, où q est la caractéristique du corps de base [CJS14]. Les meilleurs

[1] C'est-à-dire pas seulement des réseaux euclidiens.

algorithmes classiques et quantiques connus pour résoudre le problème de l'isogénie entre deux courbes elliptiques supersingulières sont tous de complexité exponentielle en $\log(q)$ [DG16 ; Gro96 ; BJS14], sauf dans le cas particulier des courbes orientées dont le graphe d'isogénies a des propriétés analogues à celui des courbes ordinaires. Les courbes supersingulières sont préférées aux courbes ordinaires en cryptographie en partie pour ces garanties de sécurité mais également pour des raisons d'efficacité qui seront vues en Section 1.1.4.

Attaque par interpolation en dimension supérieure

Le problème de l'isogénie devient facile lorsque l'on ajoute l'image de l'isogénie $\varphi : E_1 \longrightarrow E_2$ à trouver par certains points aux informations connues de l'attaquant. En effet, grâce à un résultat dû à Ernst Kani [Kan97], on peut construire des algorithmes d'interpolation utilisant des isogénies en dimension ≥ 2 capable d'évaluer φ en tout point en temps polynomial étant donnés E_1, E_2 et l'image des points connus de l'attaquant [CD23 ; MMPPW23 ; Rob23]. Ces algorithmes d'interpolation ont été utilisés pour attaquer le protocole d'échange de clés SIDH [JDF11] qui rendait public l'image de certains points par les isogénies servant à construire la clé secrète partagée.

On aurait pu penser que cette attaque était dévastatrice pour la cryptographie à base d'isogénies. Ce fut en réalité tout le contraire. Les algorithmes d'interpolation en dimension supérieure sont si efficaces qu'ils ont pu être utilisés pour construire de nouveaux protocoles et accélérer des protocoles existants. Les contributions les plus importantes de cette thèse découlent naturellement de cette idée. Ainsi, le protocole de signature SQIsignHD [DLRW24], dérivé de SQIsign, a été l'une des premières propositions. Il a été suivi de SQIsign2D-West [BDF+25] qui sert aujourd'hui de référence pour la soumission de SQIsign au deuxième tour de la compétition du NIST [AAA+25].

Accélération de la correspondance de Deuring effective

Correspondance de Deuring et applications cryptographiques

SQIsign est construit grâce à la correspondance de Deuring [Deu41] entre idéaux d'une algèbre de quaternions et isogénies entre courbes elliptiques supersingulières. Si E_1 et E_2 sont deux courbes elliptiques supersingulières définies sur $\overline{\mathbb{F}}_p$, alors leurs anneaux d'endomorphismes sont isomorphes à des ordres maximaux $\mathcal{O}_1 \simeq \mathrm{End}(E_1)$ et $\mathcal{O}_2 \simeq \mathrm{End}(E_2)$ dans une même algèbre de quaternions $\mathcal{B} = \mathcal{O}_1 \otimes \mathbb{Q} = \mathcal{O}_2 \otimes \mathbb{Q}$. À toute isogénie $\varphi : E_1 \longrightarrow E_2$ correspond un idéal I_φ *connectant* $\mathcal{O}_1$ et $\mathcal{O}_2$, c'est-à-dire un idéal à gauche pour $\mathcal{O}_1$ et à droite pour $\mathcal{O}_2$.

La correspondance de Deuring permet de trouver une isogénie entre E_1 et E_2 lorsque leurs anneaux d'endomorphismes sont connus. La première étape consiste à trouver un idéal I connectant $\mathcal{O}_1$ et $\mathcal{O}_2$. Cet idéal peut ensuite être traduit en une isogénie $\varphi_I : E_1 \longrightarrow E_2$ en temps polynomial [EHLMP18 ; DFKLPW20]. Cette méthode – qui constitue le principe de l'algorithme de signature dans SQIsign – est bien adaptée aux applications cryptographiques en raison de sa complexité polynomiale et de ses garanties de sécurité. En effet, sans connaissance des anneaux d'endomorphismes (qui constituent une information secrète), trouver une isogénie entre E_1 et E_2 revient à résoudre le problème de l'isogénie supersingulière réputé difficile. De plus, il a été démontré que le calcul de l'anneau des endomorphismes d'une courbe elliptique supersingulière est équivalent au problème de l'isogénie supersingulière [Wes22 ; MW25].

Des algorithmes plus rapides pour traduire les idéaux en isogénies

Toutefois, l'algorithme de signature de la version originelle de SQIsign est lent en pratique en raison des contraintes de l'algorithme de traduction d'un idéal I en une isogénie $\varphi_I : E_1 \longrightarrow E_2$. Avec les techniques standards précédant les attaques d'interpolation de SIDH, il n'était possible de traduire que des idéaux de normes friables. Un idéal de norme friable connectant $\mathcal{O}_1$ et $\mathcal{O}_2$ devait donc être calculé en utilisant un algorithme dû à Kohel, Lauter, Petit et Tignol (KLPT) [KLPT14], produisant des idéaux de norme certes friable mais très grande. Cet idéal de très grande norme était donc coûteux à traduire en isogénie. En outre, la nature non-uniforme des distributions d'idéaux obtenues grâce à l'algorithme KLPT rendaient plus difficile la preuve de sécurité de SQIsign.

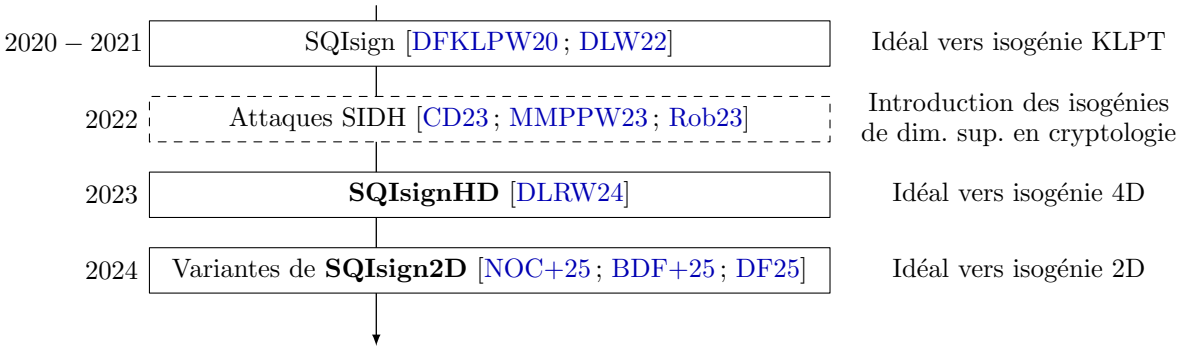

FIGURE 1 : Frise chronologique de SQIsign et ses variantes.

Les algorithmes d'interpolation en dimension supérieure ont permis de contourner totalement ces difficultés. En utilisant ces algorithmes pour la traduction d'idéal en isogénie, il n'est plus nécessaire d'imposer que la norme de l'idéal soit friable. Les idéaux connectants de norme non friable peuvent ainsi atteindre des normes plus petites et sont donc plus rapides à traduire (malgré le recours à la dimension supérieure) et ont de meilleures distributions que ceux de SQIsign. Un premier algorithme de traduction d'idéal en isogénie s'appuyant sur un calcul d'isogénie en dimension 4 a été proposé pour construire SQIsignHD. Des travaux ultérieurs [PR23 ; NO24] ont ensuite mené à la construction d'un algorithme de traduction d'idéal en isogénie avec un calcul en dimension 2. SQIsign2D-West s'appuie sur cet algorithme plus récent.

Nous présenterons ces nouveaux algorithmes de traduction en dimensions 2 et 4 dans le Chapitre 2 qui s'appuie notamment sur les travaux suivants menés au cours de la thèse :

- [DLRW24] Dartois, P., Leroux, A., Robert, D., Wesolowski, B. (2024). SQIsignHD : New Dimensions in Cryptography. Advances in Cryptology – EUROCRYPT 2024, Springer Nature Switzerland, 3-32.

- [BDF+25] Basso, A., Dartois, P., De Feo, L., Leroux, A., Maino, L., Pope, G., Robert, D., and Wesolowski, B. (2025). SQIsign2D-West : the Fast, the Small, and the Safer. Advances in Cryptology – ASIACRYPT 2024, Springer Nature Singapore, 339-370.

Améliorations successives de SQIsign

Ces améliorations algorithmiques ont permis d'accélérer considérablement la signature de SQIsign et d'obtenir des preuves de sécurité plus naturelles et moins heuristiques. Dans SQIsignHD, le calcul d'isogénie en dimension 4 est réalisé lors de la vérification mais pas lors de la signature. L'algorithme de signature obtenu est ainsi très rapide, au détriment de la vérification. Par rapport à la première implémentation NIST de SQIsign [CSSD+23], la signature de la dernière implémentation de SQIsignHD est respectivement 63,5 et 728,5 fois plus rapide aux niveaux de sécurité NIST-I et NIST-V (voir Tableau 2). Bien qu'il n'existe pas encore d'implémentation bas niveau de la vérification (il s'agit d'un travail en cours), la performance de la preuve de concept implémentée en SageMath est encourageante. La vérification termine en 600 ms au niveau de sécurité NIST-I sur une CPU Intel Core i5-1335U 4600MHz.

Dans SQIsign2D-West, la signature recourt à plusieurs calculs d'isogénies en dimension 2 et la vérification utilise un seul calcul d'isogénie en dimension 2. Cela permet d'accélérer considérablement la vérification, respectivement 3,1 et 24,3 fois plus rapide que la version originelle de SQIsign aux niveaux de sécurité NIST-I et NIST-V. La signature de SQIsign2D-West est moins rapide que SQIsignHD mais reste respectivement 9,5 et 99,9 fois plus rapide que SQIsign aux niveaux de sécurité NIST-I et NIST-V. D'autres versions de SQIsign en dimension 2, SQIsign2D-East [NOC+25] et SQIPrime [DF25], ont été proposées parallèlement à SQIsign2D-West (voir Figure 1). Bien qu'elles n'aient pas été implémentées en C ou en un langage bas niveau de performance comparable, ces variantes sont probablement plus performantes que SQIsign2D-West, au prix d'une preuve de sécurité plus heuristique. C'est pourquoi SQIsign2D-West a servi de référence pour la soumission de SQIsign au tour 2 du NIST, avec de

nouveaux gains de performance par rapport à la première implémentation de SQIsign2D-West (voir Tableau 2).

Niveau de sécurité		NIST-I	NIST-III	NIST-V
SQIsign NIST v 1.0	Génération de clé (ms)	355,72	5 625,72	22 445,3
	Signature (ms)	554,78	10 553,18	41 322,21
	Vérification (ms)	7,77	195,86	571,77
SQIsignHD	Génération de clé (ms)	14,85	48,5	112,31
	Signature (ms)	8,74	25,68	56,72
SQIsign2D-West	Génération de clé (ms)	16,53	52,24	113,18
	Signature (ms)	58,17	220,26	413,46
	Vérification (ms)	2,53	9,77	23,57
SQIsign NIST v 2.0	Génération de clé (ms)	10,63	32,05	51,37
	Signature (ms)	24,53	74,20	126,72
	Vérification (ms)	1,13	4,10	8,49

TABLE 2 : Durée en ms de signature, génération de clé et vérification de differentes versions de SQIsign implémentées en C sur une CPU Intel Core i5-1335U 4600MHz. La version NIST 2.0 dérivée de SQIsign2D-West a été implémentée avec une arithmétique des corps finis optimisée en assembleur pour les processeurs Intel. La vérification de SQIsignHD n'est pas affichée faute d'une implémentation en C.

Ces gains de performance et de sécurité ont également été associés à des gains de compacité. Ainsi, la taille des signatures est respectivement 40% et 15% plus faible dans SQIsignHD et SQIsign2D-West que dans la version originelle de SQIsign (voir Tableau 3). La description et l'analyse de sécurité des protocoles SQIsignHD et SQIsign2D-West seront présentés dans les Chapitres 3 et 4.

Niveau de sécurité	NIST-I	NIST-III	NIST-V
SQIsign (octets)	177	263	335
SQIsignHD (octets)	108	160	212
SQIsign2D-West (octets)	150	222	294

TABLE 3 : Comparaison des tailles de signatures (en octets) de SQIsign, SQisignHD et SQIsign2D-West.

Action non restreinte du groupe des classes sur les courbes supersingulières orientées

Si $\mathfrak{O}$ est un ordre quadratique imaginaire, une courbe supersingulière E est dite $\mathfrak{O}$-*orientée* lorsqu'elle est munie d'un plongement $\mathfrak{O} \hookrightarrow \mathrm{End}(E)$ qui ne peut être étendu à un ordre plus grand. Le groupe des classes d'idéaux $\mathrm{Cl}(\mathfrak{O})$ agit sur l'ensemble des (classes d'isomorphismes des) courbes $\mathfrak{O}$-orientées de telle sorte que pour toute courbe supersingulière $\mathfrak{O}$-orientée E, tout idéal $\mathfrak{a} \subseteq \mathfrak{O}$ corresponde à une isogénie $\varphi_{\mathfrak{a}} : E \longrightarrow \mathfrak{a} \cdot E$ dont le codomaine $\mathfrak{a} \cdot E$ est le résultat de l'action de $\mathfrak{a}$ sur E.

L'action de $\mathrm{Cl}(\mathfrak{O})$ sur les courbes $\mathfrak{O}$-orientées est une action de groupe cryptographique, c'est-à-dire notamment qu'elle est :

- Facile à calculer : on peut calculer calculer $\mathfrak{a} \cdot E$ à partir de $\mathfrak{a}$ et E en temps polynomial ;
- Difficile à inverser : il est difficile de calculer la classe de $\mathfrak{a}$ à partir de E et $\mathfrak{a} \cdot E$.

L'intérêt des actions de groupe cryptographiques est de pouvoir « traduire » dans le monde post-quantique des protocoles construits sur des logarithmes discrets. De telles actions de groupes ont d'abord été construites sur les courbes ordinaires [Cou06 ; RS06] avant leur introduction pour les courbes supersingulières, plus efficaces en pratique. Les protocoles d'échange de clé à base d'isogénies CSIDH [CLMPR18], OSIDH [CK20] et Scallop [FFK+23] ainsi que les signatures électroniques [BKV19] et protocoles plus avancés qui en découlent sont construits à partir d'actions de groupes sur les courbes supersingulières orientées.

En fait, l'action du groupe des classes d'idéaux $\mathrm{Cl}(\mathfrak{O})$ sur les courbes $\mathfrak{O}$-orientées ne définit pas une action de groupe cryptographique parfaite mais *restreinte*. Cela signifie que l'on peut calculer facilement l'action par des idéaux de petites normes mais pas par n'importe quel idéal. Heureusement, ces idéaux de petites normes engendrent $\mathrm{Cl}(\mathfrak{O})$ et on peut donc « mimer » l'action de l'ensemble du groupe des classes en considérant des produits d'idéaux.

Cependant, cette solution n'est pas sans inconvénients. En effet, il est utile dans certains protocoles d'agir par des classes d'idéaux ayant une distribution proche de l'uniforme dans $\mathrm{Cl}(\mathfrak{O})$. C'est le cas par exemple dans le schéma de signature électronique CSI-FiSh [BKV19] dérivé de CSIDH. Pour échantillonner des idéaux avec une telle distribution comme produit d'idéaux générateurs de petites normes, il convient de précalculer la structure de $\mathrm{Cl}(\mathfrak{O})$ et le réseau des relations des idéaux générateurs. Ce calcul est de complexité sous-exponentielle en $\log|\mathrm{disc}(\mathfrak{O})|$ et ne peut donc être réalisé pour de grands paramètres.

Or une attaque quantique connue due à Kuperberg [Kup05] contre le problème d'inversion des actions de groupes cryptographiques est elle-même de complexité sous-exponentielle. Des estimations prudentes conduisent donc à choisir de grands paramètres pour assurer un niveau de sécurité suffisant [BS20 ; Pei20 ; CSCDJRH21]. Pour CSIDH (et CSI-FiSh), il faudrait par exemple travailler sur un corps avec une caractéristique d'environ 4000 bits pour garantir un niveau de sécurité NIST-I [CSCDJRH21]. CSI-FiSh a donc un gros problème de passage à l'échelle.

Bien que Scallop [FFK+23] et son dérivé Scallop-HD [CLP24] aient été proposés pour résoudre ce problème, leur efficacité pratique n'est pas encore tout à fait satisfaisante. Une nouvelle approche algorithmique [PR23] qui a également inspiré SQIsign2D-West a permis de construire l'algorithme PEGASIS (practical effective group action using 4-dimensional isogenies) calculant l'action du groupes des classes par n'importe quel idéal. Cet algorithme qui utilise des techniques d'interpolation en dimension 4, se compare favorablement à l'approche concurrente KLaPoTi [PPS24] utilisant uniquement la dimension 2 ainsi qu'aux différentes implémentations d'actions de groupes cryptographiques restreintes de Scallop (voir Tableau 4). Implémenté en SageMath pour CSIDH, PEGASIS est le premier algorithme calculant des actions de groupes non-restreintes terminant en temps raisonnable même à niveau de sécurité élevé (4000 bits). Les premiers résultats sont donc déjà très encourageants et une implémentation en C devrait suivre (elle sera également utile à SQIsignHD et d'autres protocoles utilisant des isogénies de dimension 4).

Nous présenterons PEGASIS dans le Chapitre 2 consacré aux algorithmes de traduction des idéaux en isogénies en nous appuyant sur la prépublication suivante :

- [DEF+25] Dartois, P., Komada Eriksen, J., Fouotsa, T. B., Herlédan Le Merdy, A., Invernizzi, R., Robert, D., Rueger, R., Vercauteren, F., and Wesolowski, B. (2025). PEGASIS : Practical Effective Class Group Action using 4-Dimensional Isogenies. Preprint. Cryptology ePrint Archive, Paper 2025/401.

Calcul rapide d'isogénies en dimension supérieure

SQIsignHD et SQIsign2D-West n'auraient pu voir le jour sans algorithmes efficaces pour le calcul d'isogénies en dimension supérieure. La conception et l'implémentation de tels algorithmes est une contribution centrale de cette thèse. Nos travaux ont principalement porté sur le calcul de 2^e-isogénies (dont le degré est une puissance de 2) en dimensions 2 [DMPR25] et 4 [Dar24]. En dimension 1 comme en dimension supérieure, les 2^e-isogénies sont en effet les plus rapides à calculer car on peut les décomposer en chaînes d'isogénies élémentaires appelées 2-isogénies qui sont elles-mêmes rapides à calculer.

Papier	Langage	500	1000	1500	2000	4000
SCALLOP [FFK+23]*	C++	35s	12m30s	–	–	–
SCALLOP-HD [CLP24]*	Sage	88s	19m	–	–	–
PEARL-SCALLOP [ABE+24]	C++	30s	58s	12m	–	–
KLaPoTi [PPS24]	Sage	200s	–	–	–	–
	Rust	1.95s	–	–	–	–
PEGASIS (ce travail)	Sage	1.53s	4.21s	10.5s	21.3s	2m2s

TABLE 4 : Comparaison entre PEGASIS et d'autres actions de groupe cryptographiques dans la littérature sur une CPU Intel Core i5-1235U. Les 5 dernières colonnes donnent les temps correspondant aux différents niveaux de sécurité, où s/m donne le nombre de secondes/minutes. SCALLOP et SCALLOP-HD sont étoilés car ils ont été mesurés sur une configuration matérielle différente.

En dimension 2, les anciennes techniques pour calculer les 2-isogénies consistaient à utiliser les coordonnées de Mumford sur le modèle Jacobien (isogénies de Richelot) et ses dérivés. L'usage des coordonnées thêta de niveau 2 a permis de trouver des formules beaucoup plus rapides (cf. Tableau 5). Bien que les coordonnées thêta existent en d'autres niveaux que le niveau 2, ce choix est optimal en termes de complexité car il donne le nombre minimal de coordonnées thêta[2] pour obtenir une arithmétique non triviale (4 en dimension 2 et 16 en dimension 4). Les travaux précédents sur le calcul d'isogénies en coordonnées thêta [LR12 ; LR15 ; LR22] avaient une visée plus théorique (optimiser la complexité asymptotique des algorithmes) et ne permettaient pas de calculer des 2-isogénies en niveau 2.

$\log_2(p)$	e	**Thêta** Rust [DMPR25]	**Thêta** Sage [DMPR25]	Richelot Sage [OP22]	Richelot Sage [Kun22]	Kummer Sage [Kun22]
254	126	**2.13**	**108**	1028	760	467
381	208	**9.05**	**201**	1998	1478	858
1293	632	**463**	**1225**	12840	9196	5150

TABLE 5 : Temps d'exécution en ms du calcul d'une 2^e-isogénie en dimension 2 sur le corps de base $\mathbb{F}_{p^2}$ en Rust et en Sagemath sur une CPU Intel Core i7-9750H (2.6 GHz).

Une contribution ultérieure de cette thèse [Dar24] a permis de généraliser les algorithmes de calcul des 2-isogénies en coordonnées thêta de niveau 2 à toute dimension pour les implémenter en dimension 4 en SageMath[3]. Les résultats obtenus sont encourageants (cf. Tableau 6) et une implémentation C (en cours) devrait permettre de gagner un facteur 50 sur les temps d'exécution. À ce stade, cette implémentation en SageMath permet déjà de casser toutes les instances de SIDH en quelques secondes. À titre de comparaison, les implémentations précédentes[4] utilisant la dimension 2 et un certain nombre d'astuces couteuses pouvaient prendre plusieurs heures.

Après une introduction à la théorie des fonctions thêta due à Mumford [Mum66] au Chapitre 5, nous présentons au Chapitre 6 nos travaux sur le calcul d'isogénies en dimension supérieure en coordonnées thêta en nous appuyant sur les deux contributions suivantes :

- [DMPR25] Dartois, P., Maino, L., Pope, G., and Robert, D. (2025). An Algorithmic Approach to (2, 2)-Isogenies in the Theta Model and Applications to Isogeny-Based Cryptography. Advances in Cryptology – ASIACRYPT 2024, Springer Nature Singapore, 304-338.

[2] En général, sur une variété abélienne de dimension g, il y a n^g coordonnées thêta de niveau n.
[3] Voir ici https://github.com/Pierrick-Dartois/Theta_dim4.
[4] Voir https://github.com/GiacomoPope/Castryck-Decru-SageMath et https://github.com/Breaking-SIDH/direct-attack.

$\log_2(p)$	e	Dimension 4	Dimension 1 (G. Pope)
121	64	695	37
254	128	1428	83
365	192	2320	137

TABLE 6 : Temps d'exécution en ms du calcul d'une 2^e-isogénie en dimension 4 et en dimension 1 sur le corps de base $\mathbb{F}_{p^2}$ en SageMath sur une CPU Intel Core i5 double coeur (2.7 GHz).

- [Dar24] Dartois, P. (2024). Fast Computation of 2-Isogenies in Dimension 4 in the Theta model and Cryptographic Applications. Preprint, Cryptology ePrint Archive, Paper 2024/1180.

Une réduction entre problèmes à base d'isogénies

Nous avons choisi de ne pas présenter un article plus éloigné de la problématique centrale de la thèse que nos autres contributions :

- [ACD+24] Arpin, S., Clements, J., Dartois, P., Komada Eriksen, J., Kutas, P., and Wesolowski, B. (2024). Finding orientations of supersingular elliptic curves and quaternion orders. Designs, Codes and Cryptography 92(11), 3447–3493.

Nous discutons brièvement de sa portée ici.

La sécurité de la cryptographie reposant sur un ensemble de problèmes réputés difficiles, comprendre la difficulté relative de ces problèmes et donc les réductions entre ces problèmes est une démarche naturelle en cryptologie. L'une des contributions de l'article [ACD+24], due en grande partie à l'auteur de ces lignes, a porté sur une réduction du problème calculatoire consistant à trouver l'orientation d'une courbe elliptique $\mathfrak{O}$-orientable vers le problème décisionnel consistant à déterminer si une courbe est $\mathfrak{O}$-orientable ou non. L'article proposait aussi une solution au problème (analogue par correspondance de Deuring) consistant à trouver un plongement d'un ordre quadratique imaginaire dans un ordre d'une algèbre de quaternions.

Pour trouver une $\mathfrak{O}$-orientation d'une courbe elliptique, il suffit de trouver un endomorphisme pouvant s'identifier à un générateur de $\mathfrak{O}$. Un tel endomorphisme peut être trouvé en parcourant le graphe d'isogénies des courbes $\mathfrak{O}$-orientées à l'aide d'un oracle déterminant si une courbe est $\mathfrak{O}$-orientable ou non. On obtient ainsi une réduction sous-exponentielle du problème calculatoire vers le problème décisionnel d'$\mathfrak{O}$-orientation. Cela démontre qu'un oracle pour le problème décisionnel fournit une information non triviale puisque les meilleurs algorithmes connus pour trouver un endomorphisme d'une courbe elliptique supersingulière quelconque sont de complexité exponentielle.

Structure de la thèse

Dans le Chapitre 1, nous commençons par une série de préliminaires introduisant les notions et résultats fondamentaux sur les isogénies, la correspondance de Deuring, les courbes supersingulières orientées et les variétés abéliennes, qui seront utiles dans toute la suite.

Puis, dans une première partie, nous présentons les applications cryptographiques des algorithmes d'interpolation en dimension supérieure. Le Chapitre 2 est consacré aux algorithmes de traduction d'idéaux d'une algèbre de quaternions en isogénie et à PEGASIS, dont la démarche est proche. Le Chapitre 3 présente SQIsignHD et le Chapitre 4 présente SQIsign2D-West.

La deuxième partie de cette thèse est dédiée au calcul d'isogénies en dimension supérieure en coordonnées thêta. Le Chapitre 5 introduit la théorie des fonctions thêta due à Mumford [Mum66] et le Chapitre 6 présente plus spécifiquement nos travaux et les implémentations en dimensions 2 et 4 qui y sont associées.

Introduction

The quantum menace

Beyond military communications and the financial assets of a few enthusiastic technophiles, cryptography is essential to our digital daily lives. It has become impossible to browse the internet, make a bank transaction or sign an online document securely without using cryptography. As opposed to symmetric protocols, asymmetric protocols (also known as public key protocols) use different keys for the parties involved, one being public and the other secret. Asymmetric protocols can be used, for example, to sign digital documents or to securely exchange a secret key between two parties, which can then be used in a symmetric encryption protocol. Today, the security of all asymmetric protocols currently in use is based on one of the following two problems.

Problem 1 (RSA Factorisation)**.** *Given an integer $N := p \cdot q$, with two unknown prime factors p and q of comparable size, find p and q.*

Problem 2 (Discrete logarithm)**.** *Given a cyclic group of prime order $G = \langle g \rangle$ of known generator g and any element $h \in G$, find $x \in \{0, \cdots, \#G - 1\}$ such that $h = g^x$.*

These two problems are strongly assumed to be difficult for a classical computer, the best known generic attacks against RSA and the discrete logarithm problem being respectively of sub-exponential complexity [BLP93] and exponential complexity [Sha71] in the size of the objects ($\log(N)$ or $\log(\#G)$). However, Shor's algorithm [Sho97] can be used to solve them in polynomial time on a quantum computer. As a precaution, in anticipation of future advances in quantum computing, major research efforts have been devoted to finding protocols built on mathematical problems that are resistant to quantum attacks, also called *post-quantum protocols*.

Depending on the underlying problems, post-quantum protocols are classified into families. For instance, lattice-based cryptography relies on the difficulty of finding, in a high-dimensional Euclidean lattice, short vectors or vectors close to any vector [Ajt98 ; DKRS03]. Code-based cryptography relies on the difficulty of decoding a noisy message to which a linear error-correcting code has been applied [BMT78]. Multivariate cryptography is based on the difficulty of solving systems of polynomial equations with several variables. Isogeny-based cryptography, which is the main focus of this thesis, is based on the difficulty of finding an isogeny between two elliptic curves. Among the alternatives, isogeny-based cryptography is distinguished by its compactness (key or signature sizes are small) and its relative slowness.

The US National Institute of Standards and Technology (NIST) has launched two competitions to standardise post-quantum protocols in order to anticipate their industrial deployment. The first competition resulted in the standardisation of two key encapsulation mechanisms (KEM) and two digital signature protocols, half of which are lattice-based (see Table 7). The SIDH (Supersingular Isogeny Diffie Hellman) key exchange protocol made it to the fourth and final round of the competition but was subject to a very fast (classical) attack which had important (ironically positive) consequences for isogeny-based cryptography. The NIST launched a second competition to standardise other digital signatures that were compact, quick to verify and based on diversified security assumptions[5], thus giving isogenies a chance. The isogeny-based protocol SQIsign (Short Quaternion Isogeny Signature), which is not subject to the SIDH attack (and ironically has even benefited from it), has been selected for the second round of this competition. It is now a serious candidate for industrial deployment, which has been greatly improved by the works presented in this thesis.

[5] *i.e.* not only lattice-based.

Type of scheme	Name	Family
Key encapsulation (KEM)	Crystals-Kyber HQC*	Euclidean lattices Linear codes
Digital signature	Crystals-Dilithium Sphincs+	Euclidean lattices Hash functions

TABLE 7 : Cryptographic schemes already standardised or to be standardised after the first NIST post-quantum competition.
*Standard yet to be published.

Challenges in isogeny-based cryptography

The isogeny problem

Elliptic curves are 1-dimensional projective curves with affine equations of the form $y^2 = x^3 + Ax + B$. If E is an elliptic curve defined over a field k, then the set $E(k)$ of k-rational points of E has an abelian group structure whose group law is algebraic and efficient to compute. When k is finite, the group $E(k)$ is of finite order. For this reason, elliptic curves defined over finite fields have been widely used to construct cryptographic schemes based on discrete logarithms.

Isogenies are non-zero morphisms between elliptic curves as group homomorphisms and as morphisms of algebraic varieties, *i.e.* they can be expressed using rational fractions. Isogeny-based cryptography relies on the difficulty of solving the following problem (or similar problems) for a classical and a quantum computer.

Problem 3 (Isogeny Problem)**.** *Given two elliptic curves E_1 and E_2 defined over a finite field $\mathbb{F}_q$, find an isogeny $\varphi : E_1 \longrightarrow E_2$.*

The difficulty of this problem has been widely studied and depends on the structure of the endomorphism ring of the curves under consideration. Given an elliptic curve E, the endomorphisms ring $\mathrm{End}(E)$ is made of all isogenies $E \longrightarrow E$ and of the zero homomorphism. If E is defined over a finite field, then there are two cases [Sil09, Theorem V.3.1] :

- $\mathrm{End}(E)$ is isomorphic to a quadratic imaginary order, in which case E is called *ordinary* ;
- $\mathrm{End}(E)$ is isomorphic to a maximal order in a quaternion algebra, in which case E is called *supersingular*.

An ordinary curve cannot be isogenic to a supersingular curve, since an isogeny partly carries the structure of the ring of endomorphisms from one curve to the other. If E_1 and E_2 are ordinary and their endomorphism rings are isomorphic $\mathrm{End}(E_1) \simeq \mathrm{End}(E_2)$, then we know that there exists a quantum algorithm that can find an isogeny between E_1 and E_2 in sub-exponential time in $\log(q)$, where q is the characteristic of the base field [CJS14]. The best known classical and quantum algorithms for solving the isogeny problem between two supersingular elliptic curves are all of exponential complexity in $\log(q)$ [DG16 ; Gro96 ; BJS14], except in the special case of oriented curves, whose isogeny graph has properties very similar to the ordinary isogeny graph. Supersingular elliptic curves are preferred to ordinary elliptic curves in cryptography partly for these security guarantees but also for efficiency reasons which will be explained in Section 1.1.4.

Interpolation attack in higher dimension

The isogeny problem becomes easy when we add the image of the isogeny $\varphi : E_1 \longrightarrow E_2$ to be found by certain points to the information known by the attacker. Indeed, with a result due to Ernst Kani [Kan97], it is possible to construct interpolation algorithms using isogenies in dimension ≥ 2 to evaluate φ at any point in polynomial time given E_1, E_2 and the image of the points known to the attacker [CD23 ; MMPPW23 ; Rob23]. These interpolation algorithms were used to attack the SIDH key exchange protocol [JDF11] which made public the image of certain points by the isogenies used to construct the shared secret key.

One might think that this attack would be devastating for isogeny-based cryptography. It was actually the opposite. Higher-dimensional interpolation algorithms are so efficient that they have been used to build new protocols and speed up existing ones. The most important contributions of this thesis follow naturally from this idea. The SQIsignHD digital signature scheme, derived from SQIsign, was one of the first proposals. It was followed by SQIsign2D-West [BDF+25], which has been used as reference for SQIsign's submission to the second round of the NIST [AAA+25] competition.

Accelerating effective Deuring correspondence

The Deuring correspondence and cryptographic applications

SQIsign is constructed with the Deuring correspondence [Deu41] between quaternion ideals and isogenies between supersingular elliptic curves. If E_1 and E_2 are two supersingular elliptic curves defined over $\overline{\mathbb{F}}_p$, then their endomorphism rings are isomorphic to maximal orders $\mathcal{O}_1 \simeq \mathrm{End}(E_1)$ and $\mathcal{O}_2 \simeq \mathrm{End}(E_2)$ in the same quaternion algebra $\mathcal{B} = \mathcal{O}_1 \otimes \mathbb{Q} = \mathcal{O}_2 \otimes \mathbb{Q}$. To every isogeny $\varphi : E_1 \longrightarrow E_2$ corresponds an ideal I_φ *connecting* $\mathcal{O}_1$ and $\mathcal{O}_2$, *i.e.* a left $\mathcal{O}_1$-ideal which is also a right $\mathcal{O}_2$-ideal.

The Deuring correspondence can be used to find an isogeny between E_1 and E_2 when their endomorphism rings are known. The first step is to find an ideal I connecting $\mathcal{O}_1$ and $\mathcal{O}_2$. This ideal can then be translated into an isogeny $\varphi_I : E_1 \longrightarrow E_2$ in polynomial time [EHLMP18 ; DFKLPW20]. This method – which is the principle of SQIsign signature algorithm – is well suited to cryptographic applications because of its polynomial complexity and security guarantees. Indeed, without knowledge of the rings of endomorphisms (which are secret information), finding an isogeny between E_1 and E_2 is equivalent to solving the supersingular isogeny problem, which is strongly assumed to be difficult. Moreover, it has been proved that computing the endomorphism ring of a supersingular elliptic curve is equivalent to the supersingular isogeny problem [Wes22 ; MW25].

Faster algorithms for ideal-to-isogeny translations

However, the signature algorithm in the original version of SQIsign is slow in practice due to the constraints of the algorithm translating an ideal I into an isogeny $\varphi_I : E_1 \longrightarrow E_2$. With standard techniques known before the SIDH interpolation attacks, it was only possible to translate ideals of smooth norms. An ideal of smooth norm connecting $\mathcal{O}_1$ and $\mathcal{O}_2$ therefore had to be computed using an algorithm due to Kohel, Lauter, Petit and Tignol (KLPT) [KLPT14] that returned an output of very large norm. This ideal of very large and smooth norm was therefore costly to translate into isogeny. In addition, the non-uniform nature of the ideal distributions obtained using the KLPT algorithm made it more difficult to prove the security of SQIsign.

Higher-dimensional interpolation algorithms have made it possible to completely overcome these difficulties. By using these algorithms to translate ideals into isogenies, it is no longer necessary to require ideals to be of smooth norm. Non-smooth norm connecting ideals can thus reach smaller norms, are therefore faster to translate (despite the use of the higher dimension) and have better distributions than those of SQIsign. A first ideal-to-isogeny translation algorithm based on a 4-dimensional isogeny calculation, has been proposed for the SQIsignHD variant of SQIsign. Subsequent work [PR23 ; NO24] then led to the construction of an ideal-to-isogeny translation algorithm with 2-dimensional isogeny computations only. SQIsign2D-West is based on this more recent algorithm.

We shall present these new ideal-to-isogeny translation algorithms in dimensions 2 and 4 in Chapter 2 which is based, in particular, on the following works carried out for this thesis :

- [DLRW24] Dartois, P., Leroux, A., Robert, D., Wesolowski, B. (2024). SQIsignHD : New Dimensions in Cryptography. Advances in Cryptology – EUROCRYPT 2024, Springer Nature Switzerland, 3-32.

- [BDF+25] Basso, A., Dartois, P., De Feo, L., Leroux, A., Maino, L., Pope, G., Robert, D., and Wesolowski, B. (2025). SQIsign2D-West : the Fast, the Small, and the Safer. Advances in Cryptology – ASIACRYPT 2024, Springer Nature Singapore, 339-370.

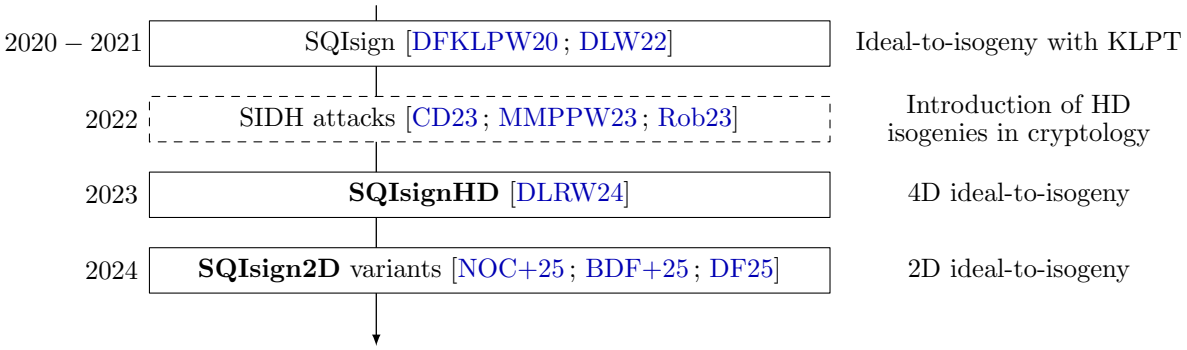

FIGURE 2 : Timeline of SQIsign and its variants.

Successive improvements of SQIsign

These algorithmic improvements have made it possible to speed up SQIsign's signature considerably and to obtain more natural and less heuristic security proofs. In SQIsignHD, the computation of a 4-dimensional isogeny is performed during the verification, but not during the signature. This results in a very fast signature algorithm, at the expense of the verification. Compared with the first NIST implementation of SQIsign [CSSD+23], the signature of the latest SQIsignHD implementation is respectively 63.5 and 728.5 times faster at NIST-I and NIST-V security levels (see Table 8). Although there is no low-level implementation of the verification for now (this is a work in progress), the proof-of-concept implementation in SageMath is encouraging. The verification takes 600 ms at NIST-I security level on an Intel Core i5-1335U 4600MHz CPU.

In SQIsign2D-West, signing uses several isogeny computations in dimension 2, and verification uses a single isogeny computation in dimension 2. This speeds up verification considerably, by a factor of 3.1 and 24.3 respectively, compared with the original version of SQIsign at NIST-I and NIST-V security levels. The SQIsign2D-West signature is slower than SQIsignHD, but remains 9.5 and 99.9 times faster than SQIsign at NIST-I and NIST-V security levels respectively. Other versions of SQIsign in dimension 2, SQIsign2D-East and SQIPrime, have been proposed alongside SQIsign2D-West (see Figure 2). Although they have not been implemented in C or a comparable low-level language, these variants are probably more powerful than SQIsign2D-West, at the expense of a more heuristic security proof. This is why SQIsign2D-West was selected as a reference for SQIsign's submission at the second round of the NIST competition, with further performance gains over the first SQIsign2D-West implementation (see Table 8).

These performance and security gains have also been combined with gains in compactness. For example, signature sizes are respectively 40% and 15% smaller in SQIsignHD and SQIsign2D-West than in the original version of SQIsign (see Table 9). The description and security analysis of the SQIsignHD and SQIsign2D-West protocols will be presented in Chapters 3 and 4.

Unrestricted action of the ideal class group on oriented supersingular elliptic curves

If $\mathfrak{O}$ is a quadratic imaginary order, a supersingular elliptic curve E is called $\mathfrak{O}$-*oriented* when there is an embedding $\mathfrak{O} \hookrightarrow \mathrm{End}(E)$ which cannot be extended to a superorder. The ideal class group $\mathrm{Cl}(\mathfrak{O})$ acts on (isomorphism classes of) $\mathfrak{O}$-oriented curves so that for any $\mathfrak{O}$-oriented supersingular curve E, any ideal $\mathfrak{a} \subseteq \mathfrak{O}$ corresponds to an isogeny $\varphi_{\mathfrak{a}} : E \longrightarrow \mathfrak{a} \cdot E$ whose codomain $\mathfrak{a} \cdot E$ is the result of the action of $\mathfrak{a}$ on E.

The action of $\mathrm{Cl}(\mathfrak{O})$ on $\mathfrak{O}$-oriented curves is a cryptographic group action, which means in particular that it is :

- Easy to compute : one can compute $\mathfrak{a} \cdot E$ from $\mathfrak{a}$ and E in polynomial time ;
- Difficult to invert : it is difficult to compute the class of $\mathfrak{a}$ from E and $\mathfrak{a} \cdot E$.

The main interest of cryptographic group actions is to "translate" protocols built on the discrete logarithm problem into the post-quantum world. Such group actions were first constructed on ordinary

Security level		NIST-I	NIST-III	NIST-V
SQIsign NIST v 1.0	Key generation (ms)	355.72	5 625.72	22 445.3
	Signature (ms)	554.78	10 553.18	41 322.21
	Verification (ms)	7.77	195.86	571.77
SQIsignHD	Key generation (ms)	14.85	48.5	112.31
	Signature (ms)	8.74	25.68	56.72
SQIsign2D-West	Key generation (ms)	16.53	52.24	113.18
	Signature (ms)	58.17	220.26	413.46
	Verification (ms)	2.53	9.77	23.57
SQIsign NIST v 2.0	Key generation (ms)	10.63	32.05	51.37
	Signature (ms)	24.53	74.20	126.72
	Verification (ms)	1.13	4.10	8.49

TABLE 8 : Key generation, signing and verification times of different versions of SQIsign on an Intel Core i5-1335U 4600MHz CPU. The NIST v 2.0 version based on SQIsign2D-West was implemented with an assembly optimised finite field arithmetic for Intel processors. SQIsignHD verification has not been implemented in C so verification times were not displayed for this scheme.

Security level	NIST-I	NIST-III	NIST-V
SQIsign (bytes)	177	263	335
SQIsignHD (bytes)	108	160	212
SQIsign2D-West (bytes)	150	222	294

TABLE 9 : Comparison of signature sizes (in bytes) in SQIsign, SQisignHD and SQIsign2D-West.

curves [Cou06 ; RS06] before being introduced for supersingular curves, which are more efficient in practice. The isogeny-based key exchange protocols CSIDH [CLMPR18], OSIDH [CK20] and Scallop [FFK+23] as well as the digital signature schemes [BKV19] and more advanced protocols derived from them are built from group actions on oriented supersingular elliptic curves.

Actually, the ideal class group action of $\mathrm{Cl}(\mathfrak{O})$ on $\mathfrak{O}$-oriented curves does not define a perfect cryptographic group action but a restricted one. This means that we can easily compute the action by ideals of small norms but not by any ideal. Fortunately, these ideals of small norms generate $\mathrm{Cl}(\mathfrak{O})$ and we can therefore "mimic" the action of the whole ideal class group by considering products of ideals.

However, this solution has drawbacks. Indeed, it is useful in some protocols to act by ideal classes with a distribution close to uniform in $\mathrm{Cl}(\mathfrak{O})$. This is the case, for instance, in the CSI-FiSh [BKV19] digital signature scheme derived from CSIDH. To sample ideals with such a distribution as a product of generating ideals of small norms, we need to precompute the structure of $\mathrm{Cl}(\mathfrak{O})$ and the relations lattice of the generating ideals. This computation is of sub-exponential complexity in $\log|\operatorname{disc}(\mathfrak{O})|$ and cannot therefore be carried out for large parameters.

However, a known quantum attack by Kuperberg [Kup05] against the inversion problem of cryptographic group actions is itself of sub-exponential complexity. Conservative estimates therefore lead to choosing large parameters to ensure a sufficient security level [BS20 ; Pei20 ; CSCDJRH21]. For CSIDH (and CSI-FiSh), we would have to work on a field of characteristic around 4000 bits to guarantee a NIST-I security level [CSCDJRH21]. CSI-FiSh therefore has a major scaling issue.

Although Scallop [FFK+23] and its derivative Scallop-HD [CLP24] have been proposed to solve this problem, their practical efficiency is not entirely satisfying. A new algorithmic approach [PR23] which also inspired SQIsign2D-West has made it possible to construct the PEGASIS algorithm (practical effective group action using 4-dimensional isogenies) which computes the ideal class group action

by any ideal. This algorithm using interpolation techniques in dimension 4, compares favourably with the competing approach KLaPoTi [PPS24] using dimension 2 only, and with the various implementations of Scallop's restricted cryptographic group actions (see Table 10). Implemented in SageMath for CSIDH, PEGASIS is the first algorithm to compute unrestricted group actions terminating in reasonable time even at high security level (4000 bits). These initial results are therefore already very encouraging and a C implementation should follow (it will also be useful for SQIsignHD and other protocols using dimension 4 isogenies).

We shall present PEGASIS in Chapter 2 devoted to ideal-to-isogeny translation algorithms, based on the following pre-publication :

- [DEF+25] Dartois, P., Komada Eriksen, J., Fouotsa, T. B., Herlédan Le Merdy, A., Invernizzi, R., Robert, D., Rueger, R., Vercauteren, F., and Wesolowski, B. (2025). PEGASIS : Practical Effective Class Group Action using 4-Dimensional Isogenies. Preprint. Cryptology ePrint Archive, Paper 2025/401.

Paper	Impl.	500	1000	1500	2000	4000
SCALLOP [FFK+23]*	C++	35s	12m30s	–	–	–
SCALLOP-HD [CLP24]*	Sage	88s	19m	–	–	–
PEARL-SCALLOP [ABE+24]	C++	30s	58s	12m	–	–
KLaPoTi [PPS24]	Sage	200s	–	–	–	–
	Rust	1.95s	–	–	–	–
PEGASIS (This work)	Sage	1.53s	4.21s	10.5s	21.3s	2m2s

TABLE 10 : Comparison between PEGASIS and other effective group actions in the literature. The last 5 columns gives the timings corresponding to the different security levels, where s/m gives the number of seconds/minutes in wall-clock time. SCALLOP and SCALLOP-HD are starred because they were measured on a different hardware setup.

Fast computation of higher-dimensional isogenies

SQIsignHD and SQIsign2D-West would not have been possible without efficient algorithms to compute higher-dimensional isogenies. The design and implementation of such algorithms is a central contribution of this thesis. Our work has mainly focused on the computation of 2^e-isogenies (whose degree is a power of 2) in dimensions 2 [DMPR25] and 4 [Dar24]. In dimension 1, as in higher dimensions, the 2^e-isogenies are indeed the fastest to compute because they can be decomposed into chains of elementary isogenies called 2-isogenies, which are themselves fast to compute.

In dimension 2, the former techniques for computing 2-isogenies used Mumford coordinates on the Jacobian model (Richelot isogenies) and its derivatives. The use of level 2 theta coordinates has made it possible to find much faster formulae (see Table 11). Although theta coordinates exist in levels other than 2, this choice is optimal in terms of complexity because it gives the minimum number of theta coordinates[6] to obtain non-trivial arithmetic (4 in dimension 2 and 16 in dimension 4). Previous works on computing isogenies in theta coordinates had a more theoretical focus (optimising the asymptotic algorithmic complexity) and did not allow to compute 2-isogenies with level 2 theta coordinates.

Another contribution of this thesis [Dar24] generalised the algorithms to compute 2-isogenies with level 2 theta coordinates in any dimension and proposed a 4-dimensional implementation in SageMath[7]. The obtained results are encouraging (cf. Table 12) and a C implementation (in progress) should make it possible to reduce execution times by a factor of 50. At this stage, the SageMath implementation can already break all SIDH instances in a few seconds. In comparison, previous implementations[8] using 2-dimensional isogenies and some expensive tricks could take several hours.

[6] In general, on an abelian variety of dimension g, there are n^g theta coordinates of level n.
[7] See https://github.com/Pierrick-Dartois/Theta_dim4.
[8] See https://github.com/GiacomoPope/Castryck-Decru-SageMath and https://github.com/Breaking-SIDH/

$\log_2(p)$	e	**Theta** Rust [DMPR25]	**Theta** Sage [DMPR25]	Richelot Sage [OP22]	Richelot Sage [Kun22]	Kummer Sage [Kun22]
254	126	**2.13**	**108**	1028	760	467
381	208	**9.05**	**201**	1998	1478	858
1293	632	**463**	**1225**	12840	9196	5150

TABLE 11 : Execution time in ms of a 2^e-isogeny computation in dimension 2 over the base field $\mathbb{F}_{p^2}$ in Rust and Sagemath on an Intel Core i7-9750H CPU (2.6 GHz).

$\log_2(p)$	e	Dimension 4	Dimension 1 (G. Pope)
121	64	695	37
254	128	1428	83
365	192	2320	137

TABLE 12 : Execution time in ms of a 2^e-isogeny computation in dimension 4 and in dimension 1 over the base field $\mathbb{F}_{p^2}$ in SageMath on a dual-core Intel Core i5 CPU (2.7 GHz).

After an introduction to the theory of theta functions due to Mumford [Mum66] in Chapter 5, we shall present in Chapter 6 our work on the computation of higher dimensional isogenies with theta coordinates based on the following two contributions :

- [DMPR25] Dartois, P., Maino, L., Pope, G., and Robert, D. (2025). An Algorithmic Approach to (2, 2)-Isogenies in the Theta Model and Applications to Isogeny-Based Cryptography. Advances in Cryptology – ASIACRYPT 2024, Springer Nature Singapore, 304-338.
- [Dar24] Dartois, P. (2024). Fast Computation of 2-Isogenies in Dimension 4 in the Theta model and Cryptographic Applications. Preprint, Cryptology ePrint Archive, Paper 2024/1180.

A reduction between isogeny-based problems

We have chosen not to present a paper further away from the central questions of the thesis than our other contributions :

- [ACD+24] Arpin, S., Clements, J., Dartois, P., Komada Eriksen, J., Kutas, P., and Wesolowski, B. (2024). Finding orientations of supersingular elliptic curves and quaternion orders. Designs, Codes and Cryptography 92(11), 3447–3493.

We briefly present its main results here.

Since cryptographic security is based on a set of problems that are known to be difficult, understanding the relative difficulty of these problems and hence the reductions between them is a natural approach in cryptology. One of the contributions of the article [ACD+24], largely due to the author of these lines, concerned a reduction from the computational problem of finding the orientation of an $\mathfrak{O}$-orientable elliptic curve to the decision problem of determining whether a curve is $\mathfrak{O}$-orientable or not. The paper also proposed a solution to the problem (analogous by Deuring correspondence) of finding an embedding of a quadratic imaginary order into an order of a quaternion algebra.

To find an $\mathfrak{O}$-orientation of an elliptic curve, all we need to do is find an endomorphism that can be identified with a generator of $\mathfrak{O}$. Such an endomorphism can be found by walking on the isogeny graph of $\mathfrak{O}$-oriented curves using an oracle that determines whether a curve is $\mathfrak{O}$-orientable or not. This results in a sub-exponential reduction of the computational problem to the $\mathfrak{O}$-orientation decision problem. This proves that an oracle for the decision problem provides non-trivial information since the best known algorithms for finding an endomorphism of any supersingular elliptic curve have exponential complexity.

`direct-attack`.

Structure of this thesis

In Chapter 1, we begin with a series of preliminaries introducing the fundamental notions and results on isogenies, the Deuring correspondence, oriented supersingular elliptic curves and abelian varieties, which will be useful throughout this thesis.

Then, in the first part, we present cryptographic applications of higher-dimensional interpolation algorithms. Chapter 2 is devoted to algorithms translating quaternion ideals into isogenies and to PEGASIS, whose approach is similar. Chapter 3 and 4 respectively introduce SQIsignHD and SQIsign2D-West.

The second part of this thesis is dedicated to the computation of higher dimensional isogenies with theta coordinates. Chapter 5 introduces the theory of theta functions due to Mumford [Mum66] and Chapter 6 more specifically presents our work and the associated implementations in dimensions 2 and 4.

Chapter 1

Preliminaries

1.1 Elliptic curves and isogenies

The reader of this thesis should be familiar with elliptic curves and isogenies but we recall some fundamental results in this section. We refer to the book by Silverman [Sil09] for a good introduction to this topic. In the following, we fix a field k.

1.1.1 Elliptic curves

Definition 1.1.1. An *elliptic curve* over k is a projective curve of the projective plane $E \subset \mathbb{P}^2(k)$ which is of genus 1 (as defined by the Riemann-Roch theorem [Har77, Theorem 1.3]) together with some base point $0_E \in E$. Note that such a curve is always smooth by definition.

In the following, we shall always assume that $\mathrm{char}(k) \neq 2, 3$, as this will be the case in our cryptographic applications. In that case, we can always translate the base point 0_E of an elliptic curve E to $(0 : 1 : 0)$ and rescale its equation to put it into (short) *Weierstrass form*:

$$Y^2Z = X^3 + AXZ^2 + BZ^3,$$

with $A, B \in k$. This equation admits an affine form

$$y^2 = x^3 + Ax + B,$$

with $x := X/Z$ and $y := Y/Z$. The base point $0_E = (0 : 1 : 0)$ cannot be represented on this affine curve and is called for that reason the *point a infinity*. Since E is smooth, the discriminant $\Delta(E) := 4A^3 + 27B^2$ is non-zero and we can define the *j-invariant* of E:

$$j(E) = 1728\frac{4A^3}{\Delta(E)}.$$

Two elliptic curves are isomorphic over $\overline{k}$ if and only if they have the same j-invariant [Sil09, Proposition III.1.4.b]. Besides, if $j_0 \in \overline{k}$, then there exists an elliptic curve defined over $k(j_0)$ with j-invariant j_0 [Sil09, Proposition III.1.4.c].

Elliptic curves admit an abelian group law [Sil09, § III.2] whose definition is very geometric. If E/k is an elliptic curve, and k'/k is a field extension, then $E(k')$ is a group given by the following law. Let $P, Q \in E(k')$ be distinct points and consider the line (PQ) in $\mathbb{P}^2(k')$. This line intersects E at a third point $R \in E(k')$ and we may consider the the symmetric S of R with respect to the axis $y = 0$. Then $P + Q = S$. When $P = Q$, we replace the line (PQ) by the tangent of E at P. With this group law, the point at infinity 0_E is the neutral element and the opposite of a point $P = (x, y)$ in affine coordinates is $-P = (x, -y)$. Also note that the group law is *algebraic*, which means formally that $(P, Q) \in E^2 \longmapsto P + Q \in E$ is a morphism of algebraic varieties and more concretely that the coordinates of $P + Q$ are expressed as explicit rational fractions in the coordinates of P and Q [Sil09, Algorithm III.2.3]. This group law have been exploited in pre-quantum cryptographic schemes based on the hardness of the discrete logarithm problem.

1.1.2 Montgomery elliptic curves and their arithmetic

A fast elliptic curve arithmetic is a crucial feature for both pre-quantum and isogeny based cryptography. Some models of elliptic curves are more appropriate than the Weierstrass model introduced above (*e.g.* Edwards, Jacobian, Montgomery...). In this thesis, we shall use the Montgomery model quite often. A *Montgomery elliptic curve* E over k has an affine equation of the form:

$$By^2 = x^3 + Ax^2 + x,$$

with $A, B \in k$. Unless explicitly stated otherwise, we shall assume that Montgomery curves have coefficient $B = 1$. When B is a quadratic non-residue in k, we say that E is a *quadratic twist* of the curve $E' : y^2 = x^3 + Ax^2 + x$ since it is isomorphic to E' over a quadratic extension of k but not over k. When k is a finite field and E/k is a Montgomery curve, it can be proved that $\#E(k)$ is always divisible by 4 [CS17, § 2.3], so E contains a k-rational point of 4-torsion or the full 2-torsion subgroup $E[2] \simeq (\mathbb{Z}/2\mathbb{Z})^2$ is k-rational. Indeed, the structure of torsion subgroups of elliptic curves is always given by the following.

Proposition 1.1.2. *[Sil09, Corollary III.6.4] Let E/k be an elliptic curve.*

(i) If $n \in \mathbb{N}^*$ is not divisible by $\mathrm{char}(k)$, then the n-th torsion subgroup $E[n]$ is isomorphic to $(\mathbb{Z}/n\mathbb{Z})^2$.

(ii) If $p = \mathrm{char}(k) > 0$, then we either have $E[p^n] \simeq (\mathbb{Z}/p^n\mathbb{Z})$ for all $n \in \mathbb{N}^*$ or $E[p^n] = \{0\}$ for all $n \in \mathbb{N}^*$.

For efficiency reasons, we often work on Montgomery curves with projective coordinates X and Z only (forgetting the Y). This allows for fast arithmetic formulas (see [CS17]) at the expense of the sign of points which is determined by the Y coordinate. Indeed, with *Montgomery $(X : Z)$-coordinates*, we no longer work on the elliptic curve E itself but on the Kummer line $E/\pm \simeq \mathbb{P}^1_k$. This means in particular that we can still perform point duplications $\pm P \longmapsto \pm[2]P$ but no longer perform simple point additions. If we want to add $\pm P$ and $\pm Q$ we can obtain either $\pm(P + Q)$ or $\pm(P - Q)$. To lift this sign ambiguity, we can obtain $\pm(P + Q)$ from $\pm P, \pm Q$ and the additional $\pm(P - Q)$. This is called *differential addition*.

1.1.3 Isogenies

Definition 1.1.3. An *isogeny* $\varphi : E_1 \longrightarrow E_2$ between elliptic curves over k is a non-constant morphism of k-varieties (in practice given by rational fractions in X, Y, Z over k) such that $\varphi(0_{E_1}) = 0_{E_2}$.

If E_1 and E_2 are elliptic curves defined over k, then for every algebraic extension k'/k, we can see E_1 and E_2 as defined over k' by extension of scalars and consider isogenies $E_1 \longrightarrow E_2$ defined over k' but not necessarily over k.

Theorem 1.1.4. *[Sil09, Theorem III.4.8] An isogeny $E_1 \longrightarrow E_2$ between elliptic curves over k induces a group homomorphism $E_1(\overline{k}) \longrightarrow E_2(\overline{k})$.*

If E_1 is an elliptic curve over k with equation $Y^2Z = f(X, Z)$, its field of functions is given by $\overline{k}(E_1) = \mathrm{Frac}(\overline{k}[X, Y, Z]/(Y^2Z - f(X, Z)))$. If $\varphi : E_1 \longrightarrow E_2$ is an isogeny, then it induces a map:

$$\varphi^* : s \in \overline{k}(E_2) \longmapsto s \circ \varphi \in \overline{k}(E_1).$$

The *degree* of φ denoted by $\deg(\varphi)$ is the degree of the extension $[\overline{k}(E_1) : \varphi^*\overline{k}(E_2)]$. By convention, the degree of 0 (which is a morphism but not an isogeny) is 0. The degree is a multiplicative map. An isogeny of degree d will also be called a *d-isogeny* in the following.

Theorem 1.1.5. *[Sil09, Theorem III.4.10] Let $\varphi : E_1 \longrightarrow E_2$ be an isogeny. Then:*

(i) φ is surjective with finite kernel.

(ii) $\#\ker(\varphi)$ divides $\deg(\varphi)$.

We say that an isogeny $\varphi : E_1 \longrightarrow E_2$ is *separable* when $\deg(\varphi) = \#\ker(\varphi)$. Most of the time the isogenies we shall consider will be separable. When $\mathrm{char}(k) = 0$, all isogenies between elliptic curves defined over k are separable. When $p := \mathrm{char}(k) > 0$, isogenies of degree coprime with p are separable. This is a general result also valid for isogenies between abelian varieties (Corollary 1.4.44).

An isogeny is called *cyclic* when its kernel is cyclic. In the literature, some authors generally assume that d-isogenies are cyclic isogenies of degree d. However, in the following, for consistency with abelian varieties (see Definition 2.2.1), we shall not assume that d-isogenies are cyclic.

Example 1.1.6. 1. If E is an elliptic curve over k, then for all $n \in \mathbb{Z} \setminus \{0\}$, the multiplication by $[n]$ is an isogeny $E \longrightarrow E$ of degree n^2 over k. It can be expressed explicitly coordinate-wise as $[n](x, y) = (f_n(x, y), g_n(x, y))$ with the use of *division polynomials*. It is separable when $\mathrm{char}(k) \nmid n$.

2. If k is of characteristic p (*e.g.* if k is finite) and E is an elliptic curve defined over k, we may consider the p-th *Frobenius* $\pi_p : (x, y) \in E \longrightarrow (x^p, y^p) \in E^{(p)}$, where $E^{(p)}$ is the elliptic curve of equation $y^2 = x^3 + A^p x + B^p$ obtained by action of the p-th Frobenius automorphism of k on the coefficients of the equation $y^2 = x^2 + Ax + B$ of E. The p-th Frobenius is an inseparable isogeny of degree p.

We define the p^n-th Frobenius for all $n \in \mathbb{N}^*$ similarly which is an inseparable isogeny of degree p^n. When the base field k is the finite field $\mathbb{F}_q$, the q-th Frobenius π_q is an endomorphism $E \longrightarrow E$. Actually, it can be proved that all inseparable isogenies are a product of a Frobenius and a separable isogeny [Sil09, Corollary II.2.12].

3. We terminate this series of examples with a concrete one. Let $E_1 : y^2 = x^3 + x + 4$ and $E_2 : y^2 = x^3 - x + 4$ be elliptic curves defined over $\mathbb{F}_7$. Then

$$\begin{array}{rcl} \varphi : E_1 & \longrightarrow & E_2 \\ (x, y) & \longmapsto & \left(\dfrac{x^2 - 2x - 1}{x - 2}, y \dfrac{x^2 + 3x - 2}{(x - 2)^2} \right) \end{array}$$

is a separable cyclic 2-isogeny with kernel $\ker(\varphi) = \langle (2, 0) \rangle$.

The following result ensures that separable isogenies are determined by their kernel up to post-composition by an isomorphism. This is particularly convenient for computational applications (see Section 1.1.5) as we can compute an isogeny from its kernel. The work of Vélu [Vé71] provides explicit formulas to express the rational fractions defining a cyclic separable isogeny φ coordinate-wise $\varphi(x, y) = (f(x, y), g(x, y))$ from the knowledge of its kernel generator in time $O(\deg(\varphi))$.

Theorem 1.1.7. *[Sil09, Proposition III.4.12] Let $K \subset E$ be a finite subgroup of an elliptic curve. Then there exists a separable isogeny $\varphi : E \longrightarrow E'$ of kernel K. If $\varphi' : E \longrightarrow E''$ is another separable isogeny with kernel K then there exists an isomorphism $\lambda : E' \xrightarrow{\sim} E''$ such that $\lambda \circ \varphi = \varphi'$.*

Proposition 1.1.8. *[Sil09, Corollary III.4.11] Let $\varphi : E \longrightarrow E'$ and $\varphi' : E \longrightarrow E''$ be two isogenies. Assume that $\ker(\varphi) \subseteq \ker(\varphi')$ and φ is separable. Then there exists an isogeny $\psi : E' \longrightarrow E''$ such that $\psi \circ \varphi = \varphi'$.*

The dual isogeny

Theorem 1.1.9. *[Sil09, Theorems III.6.1 and III.6.2] Let $\varphi : E_1 \longrightarrow E_2$ be an isogeny. Then there exists a unique isogeny $\widehat{\varphi} : E_2 \longrightarrow E_1$ such that $\widehat{\varphi} \circ \varphi = [\deg(\varphi)]_{E_1}$ called the* dual isogeny *of φ.*

The dual satisfies the following properties:

(i) $\varphi \circ \widehat{\varphi} = [\deg(\varphi)]_{E_2}$.

(ii) If $\psi : E_2 \longrightarrow E_3$ *then* $\widehat{\psi \circ \varphi} = \widehat{\varphi} \circ \widehat{\psi}$.

(iii) If $\varphi' : E_1 \longrightarrow E_2$ *then* $\widehat{\varphi + \varphi'} = \widehat{\varphi} + \widehat{\varphi'}$.

(iv) $\widehat{[n]} = [n]$ *for all* $n \in \mathbb{Z}$.

(v) $\deg(\widehat{\varphi}) = \deg(\varphi)$.

(vi) $\widehat{\widehat{\varphi}} = \varphi$.

Definition 1.1.10. Let E_1 and E_2 be two elliptic curves and consider $\mathrm{Hom}(E_1, E_2)$ the set of *homomorphisms* $E_1 \longrightarrow E_2$ which are either isogenies or the zero map. This set is naturally equipped with a $\mathbb{Z}$-module structure. When $E_1 = E_2$, $\mathrm{Hom}(E_1, E_1) = \mathrm{End}(E_1)$ has a ring structure for the composition of morphisms.

Theorem 1.1.9 admits an immediate and very useful corollary.

Corollary 1.1.11. *[Sil09, Corollary III.6.4] Let E_1 and E_2 be two elliptic curves. Then:*

(i) The degree map $\deg : \mathrm{Hom}(E_1, E_2) \longrightarrow \mathbb{Z}$ *is a positive definite quadratic form.*

(ii) If $\varphi \in \mathrm{End}(E_1)$*, then* $\varphi + \widehat{\varphi}$ *is the multiplication by an integer that we call the* trace *of* φ *and denote by* $\mathrm{Tr}(\varphi)$*.*

(iii) If $\varphi \in \mathrm{End}(E_1)$*, then* $\chi_\varphi := X^2 - \mathrm{Tr}(\varphi)X + \deg(\varphi)$ *is an annihilating polynomial of* φ *that we call the* characteristic polynomial *of* φ*.*

Proof. **(i)** Since $\deg(\varphi) \geq 0$ for all $\varphi \in \mathrm{Hom}(E_1, E_2)$ and $\deg(\varphi) = 0$ if and only if $\varphi = 0$, φ is positive definite. We only have to prove that

$$(\varphi, \psi) \in \mathrm{Hom}(E_1, E_2)^2 \longmapsto \langle \varphi, \psi \rangle := \deg(\varphi + \psi) - \deg(\varphi) - \deg(\psi) \in \mathbb{Z},$$

is bilinear. By Theorem 1.1.9, we obtain that for all $\varphi, \psi \in \mathrm{Hom}(E_1, E_2)$,

$$\begin{aligned}[\langle \varphi, \psi \rangle] &= [\deg(\varphi + \psi)] - [\deg(\varphi)] - [\deg(\psi)] = \widehat{\varphi + \psi} \circ (\varphi + \psi) - \widehat{\varphi} \circ \varphi - \widehat{\psi} \circ \psi \\ &= (\widehat{\varphi} + \widehat{\psi}) \circ (\varphi + \psi) - \widehat{\varphi} \circ \varphi - \widehat{\psi} \circ \psi = \widehat{\varphi} \circ \psi + \widehat{\psi} \circ \varphi,\end{aligned}$$

so the bilinearity follows immediately.

(ii) We notice that $[\mathrm{Tr}(\varphi)] = \varphi + \widehat{\varphi} = [\langle \varphi, [1] \rangle]$.

(iii) Finally,

$$\varphi^2 - [\mathrm{Tr}(\varphi)]\varphi + [\deg(\varphi)] = \varphi \circ \varphi - (\varphi + \widehat{\varphi}) \circ \varphi + \widehat{\varphi} \circ \varphi = 0.$$

□

Pushforward and pullbacks

Let $\varphi : E_1 \longrightarrow E_2$ be an a-isogeny and $\psi : E_1 \longrightarrow E_3$ be a b-isogeny, with $a, b \in \mathbb{N}^*$ coprime. Then we can construct the following commutative diagram of isogenies:

$$\begin{array}{ccc} E_3 & \xrightarrow{\varphi'} & E_4 \\ {\scriptstyle\psi}\uparrow & & \uparrow{\scriptstyle\psi'} \\ E_1 & \xrightarrow{\varphi} & E_2 \end{array}$$

with $\deg(\varphi') = a$ and $\deg(\psi') = b$. The commutative diagram above is sometimes called an (a, b)-*isogeny diamond* (see Definition 2.2.5). The isogenies φ' and ψ' are given by $\ker(\varphi') = \psi(\ker(\varphi))$ and $\ker(\psi') = \varphi(\ker(\psi))$.

Definition 1.1.12. We say that φ' is the *pushforward* of φ by ψ and denote $\varphi' := [\psi]_*\varphi$. Similarly, $\psi' = [\varphi]_*\psi$ is the pushforward of ψ by φ.

We also say that φ is the *pullback* of φ' by ψ and denote $\varphi = [\psi]^*\varphi'$. Also note that equivalently, φ is the pushforward of φ' by $\widehat{\psi}$: $\varphi = [\psi]^*\varphi' = [\widehat{\psi}]_*\varphi'$.

1.1.4 Elliptic curves over finite fields and supersingular elliptic curves

In the following, we consider elliptic curves defined over finite fields which are the only ones to be used in cryptographic applications.

Theorem 1.1.13 (Hasse-Weil)**.** *[Sil09, Theorem V.1.1] Let E be an elliptic curve defined over a finite field $\mathbb{F}_q$. Then the set of $\mathbb{F}_q$-rational points $E(\mathbb{F}_q)$ is finite and has cardinality:*

$$\#E(\mathbb{F}_q) = q + 1 - \mathrm{Tr}(\pi_q),$$

where π_q is the q-th Frobenius endomorphism. Furthermore, $|\mathrm{Tr}(\pi_q)| \leq 2\sqrt{q}$.

Proof. We notice that the points of $E(\mathbb{F}_q)$ are exactly points of E that are fixed by the q-th Frobenius, so $E(\mathbb{F}_q) = \ker([1]-\pi_q)$. Besides, it can be proved that $[1]-\pi_q$ is separable by computing its differential [Sil09, Corollary III.5.5], so that by Theorem 1.1.9 and Corollary 1.1.11.(ii),

$$\#E(\mathbb{F}_q) = \#\ker([1]-\pi_q) = \deg([1]-\pi_q) = 1 + \deg(\pi_q) - \mathrm{Tr}(\pi_q) = 1 + q - \mathrm{Tr}(\pi_q)$$

Besides, if we denote by $\langle , \rangle$ the bilinear map associated to the degree quadratic form, we have seen in the proof of Corollary 1.1.11 that $\mathrm{Tr}(\pi_q) = \langle \pi_q, [1] \rangle$. Then Cauchy-Shwarz inequality applies [Sil09, Lemma V.1.2] and ensures that:

$$\mathrm{Tr}(\pi_q) = \langle \pi_q, [1] \rangle \leq \sqrt{\langle \pi_q, \pi_q \rangle \langle [1], [1] \rangle} = 2\sqrt{\deg(\pi_q)\deg([1])} = 2\sqrt{q}.$$

□

Supersingular elliptic curves

Theorem 1.1.14. *[Voi21, Lemma 42.1.5 and Theorem 42.1.9] Let E be an elliptic curve defined over any field k. Then the endomorphism ring $\mathrm{End}(E)$ is either isomorphic to:*

(i) The ring of integers $\mathbb{Z}$.

(ii) An order in a quadratic imaginary field.

(iii) If $p := \mathrm{char}(k) > 0$, a maximal order an a quaternion algebra ramifying at p and ∞ (see Sections 1.2.1 and 1.2.2).

In the last two cases, we say that E has complex multiplication*.*

If E is defined over a finite field $\mathbb{F}_q$, the endomorphism ring $\mathrm{End}(E)$ not only contains $\mathbb{Z}$ but also the q-th Frobenius π_q, which, in most cases, is not a scalar. Actually, it can be proved that.

Theorem 1.1.15. *[Sil09, Theorem V.3.1] If E is an elliptic curve over a finite field $\mathbb{F}_q$, then E always has complex multiplication.*

Theorem 1.1.16. *[Sil09, Theorems V.3.1 and V.4.1.a] Let E be an elliptic curve over a field $\mathbb{F}_q$ of characteristic $p > 0$. Then the following are equivalent:*

(i) There exists $n \in \mathbb{N}^*$ such that $E[p^n] = \{0\}$.

(ii) For all $n \in \mathbb{N}^*$, $E[p^n] = \{0\}$.

(iii) $\mathrm{End}(E)$ is a maximal order in a quaternion algebra ramifying at p and ∞.

(iv) $\mathrm{Tr}(\pi_q) \equiv 0 \mod p$.

(v) The coefficient of x^{p-1} in $f(x)^{(p-1)/2}$ is zero, where $y^2 = f(x)$ is the equation of E.

If the above conditions are satisfied, we say that E is supersingular *and* ordinary *otherwise. If E is supersingular, we also have $j(E) \in \mathbb{F}_{p^2}$ so E is isomorphic (over $\overline{\mathbb{F}}_p$) to an elliptic curve defined over $\mathbb{F}_{p^2}$.*

In cryptographic applications, we mainly use supersingular elliptic curves for two reasons:

- **Security reasons** (see Section 2.1): hard problems regarding supersingular elliptic curves are assumed to be very hard even for quantum computers with well-known security reductions. Mainly, the problem of finding an isogeny between supersingular elliptic curves is equivalent to the problem of computing the endomorphism ring problem of a supersingular elliptic curve [Wes22; MW25]. All known algorithms to solve these problems are of exponential complexity in $\log(p)$ [DG16; EHLMP20; FIKMN25; BJS14] ($\widetilde{O}(\sqrt{p})$ classically and $\widetilde{O}(p^{1/4})$ quantumly).
- **Efficiency reasons:** we have some control over the number of points over $\mathbb{F}_{p^2}$ of a supersingular elliptic curve, which gives some control over torsion subgroups defined over $\mathbb{F}_{p^2}$ that are used to compute isogenies. Indeed, by Theorem 1.1.16 and Theorem 1.1.13, if $E/\mathbb{F}_{p^2}$ is supersingular, then the trace of the p^2-th Frobenius is divisible by p and satisfies $|\operatorname{Tr}(\pi_{p^2})| \leq 2p$, so $\operatorname{Tr}(\pi_{p^2}) \in \{-2p, -p, 0, p, 2p\}$ and $\#E(\mathbb{F}_{p^2}) = p^2 + 1 - \operatorname{Tr}(\pi_{p^2})$ can take only five distinct values. In some cases, we can even say more on $\#E(\mathbb{F}_{p^2})$.

Lemma 1.1.17. *Let p be a prime $\equiv 3 \mod 4$ and E be a supersingular Montgomery curve defined over $\mathbb{F}_{p^2}$. Then:*

(i) $E(\mathbb{F}_{p^2}) \simeq (\mathbb{Z}/(p+1)\mathbb{Z})^2$ or $E(\mathbb{F}_{p^2}) \simeq (\mathbb{Z}/(p-1)\mathbb{Z})^2$. We say that E is maximal *in the first case and* minimal *in the second case.*

(ii) If E has equation $y^2 = x^3 + Ax^2 + x$ where $A-2$ and $A+2$ are quadratic residue in $\mathbb{F}_{p^2}$ then E is maximal.

Proof. **(i)** By Theorem 1.1.13, we have $\#E(\mathbb{F}_{p^2}) = p^2+1-\operatorname{Tr}(\pi_{p^2})$ and $|\operatorname{Tr}(\pi_{p^2})| \leq 2p$ and since E is supersingular, $\operatorname{Tr}(\pi_{p^2}) \in \{-2p, -p, 0, p, 2p\}$ by Theorem 1.1.16. Since E is in Montgomery form, we have $\#E(\mathbb{F}_{p^2}) \equiv 0 \mod 4$ by [CS17, § 2.3]. Since $p \equiv 3 \mod 4$, it follows that $\operatorname{Tr}(\pi_{p^2}) \equiv 2 \mod 4$, so that $\operatorname{Tr}(\pi_{p^2}) = \pm 2p$ and $\#E(\mathbb{F}_{p^2}) = (p \mp 1)^2$.

Furthermore, the characteristic polynomial of the p^2-th Frobenius π_{p^2} is $\chi_{\pi_{p^2}} = X^2 \mp 2pX + p^2 = (X \mp p)^2$, so that $(\pi_{p^2} \mp [p])^2 = 0$ and $\pi_{p^2} = \pm[p]$ since $\operatorname{End}(E)$ is integral. We then have by Proposition 1.1.2:

$$E(\mathbb{F}_{p^2}) = \ker(\pi_{p^2} - [1]) = \ker(\pm[p] - [1]) = E[p \mp 1] \simeq (\mathbb{Z}/(p \mp 1)\mathbb{Z})^2.$$

(ii) When $A-2$ and $A+2$ are quadratic residue in $\mathbb{F}_{p^2}$, we have $8 | \#E(\mathbb{F}_{p^2})$ by [CS17, Table 1] so $\operatorname{Tr}(\pi_{p^2}) = p^2+1-\#E(\mathbb{F}_{p^2}) \equiv 2 \mod 8$ and $\operatorname{Tr}(\pi_{p^2})$ must equal $-2p$, so that $E(\mathbb{F}_{p^2}) \simeq (\mathbb{Z}/(p+1)\mathbb{Z})^2$ and E is maximal. This completes the proof. □

Remark 1.1.18. Note that every maximal elliptic curve defined over $\mathbb{F}_{p^2}$ *i.e.* with maximal number of $\mathbb{F}_{p^2}$-rational points $\#E(\mathbb{F}_{p^2}) = (p+1)^2$ is supersingular since its Frobenius endomorphism has trace $\operatorname{Tr}(\pi_{p^2}) = -2p \equiv 0 \mod p$. In that case, we also have $E(\mathbb{F}_{p^2}) \simeq (\mathbb{Z}/(p+1)\mathbb{Z})^2$ since $\operatorname{Tr}(\pi_{p^2}) = -2p$ ensures that $\pi_{p^2} = -[p]$.

As claimed previously, the torsion defined over $\mathbb{F}_{p^2}$ of maximal elliptic curves over $\mathbb{F}_{p^2}$ can be controlled. If $N | p+1$ then by Lemma 1.1.17.(i), any maximal elliptic curve E over $\mathbb{F}_{p^2}$ satisfies $E[N] \subseteq E(\mathbb{F}_{p^2})$. If we want the N-torsion to be defined over $\mathbb{F}_{p^2}$, we only have to choose a prime p such that $N | p+1$. This will be crucial for our cryptographic applications in the following. Indeed, we shall mainly work in the connected graph made of supersingular elliptic curves which are isogenous over $\mathbb{F}_{p^2}$ to a special supersingular elliptic curve curve $E_0 : y^2 = x^3 + x$ together with their isogenies.

Lemma 1.1.19. *Let $p \equiv 3 \mod 4$. Then the elliptic curve $E_0 : y^2 = x^3 + x$ defined over $\mathbb{F}_p$ is maximal, and every elliptic curve E isogenous to E_0 over $\mathbb{F}_{p^2}$ is maximal.*

Proof. We have

$$(x^3+x)^{(p-1)/2} = x^{(p-1)/2}(x^2+1)^{(p-1)/2} = x^{(p-1)/2} \sum_{k=1}^{(p-1)/2} \binom{(p-1)/2}{k} x^{2k}$$

and $(p-1)/2 + 2k = p-1$ if and only if $k = (p-1)/4$ but $p \equiv 3 \mod 4$ so $(p-1)/4$ is not an integer and the coefficient of x^{p-1} in $(x^3+x)^{(p-1)/2}$ is 0. By Theorem 1.1.16, it follows that E_0 is supersingular.

Since E_0 is a Montgomery curve with coefficient $A = 0$ and since $-2, 2 \in \mathbb{F}_p$ are quadratic residues in $\mathbb{F}_{p^2}$ (like every element of $\mathbb{F}_p$), Lemma 1.1.17.(ii) ensure that E_0 is maximal.

In addition, a classical theorem due to Tate [Tat66, Theorem 1] on abelian varieties ensure that a supersingular elliptic curve E isogenous to E_0 over $\mathbb{F}_{p^2}$ satisfies $\#E(\mathbb{F}_{p^2}) = \#E_0(\mathbb{F}_{p^2}) = (p+1)^2$ so is maximal. This completes the proof. □

The supersingular isogeny graph

Now, let us consider the *supersingular isogeny graph* whose vertices are isomorphism classes (over $\overline{\mathbb{F}}_p$) of supersingular elliptic curves *i.e.* j-invariants and edges are isogenies up to post-composition by an isomorphism. We can actually count the number of vertices with an easy formula.

Theorem 1.1.20. *[Sil09, Theorem V.4.1.c] If $p \geq 5$, the supersingular isogeny graph has the following number of vertices:*

$$\left\lfloor \frac{p}{12} \right\rfloor + \begin{cases} 0 & \text{if } p \equiv 1 \mod 12 \\ 1 & \text{if } p \equiv 5 \mod 12 \text{ or } p \equiv 7 \mod 12 \\ 2 & \text{if } p \equiv 11 \mod 12 \end{cases}$$

Theorem 1.1.21. *The supersingular isogeny graph is connected. If $E, E'/\overline{\mathbb{F}}_p$ are supersingular elliptic curves, then they are isogenous over $\overline{\mathbb{F}}_p$ and $\mathrm{Hom}(E, E')$ is even a $\mathbb{Z}$-module of rank 4.*

The connectivity properties of the supersingular isogeny graph are exceptional. It has been proved that this graph is Ramanujan even when we restrict to subgraphs with isogenies of certain degrees only [Piz90]. In particular, a random walk in this graph converges quickly to the uniform distribution (see Proposition 3.3.1 in particular). These properties offer good security guarantees as they ensure that it is easy "to get lost" in supersingular isogeny graphs.

1.1.5 Efficient representations of isogenies

In isogeny based cryptography, isogenies are often secret information that we want to hide from potential attackers. But what does it mean for a given party in a cryptographic protocol to *know* an isogeny? The knowledge of an isogeny can be formally defined naturally as some data called an *efficient representation* that gives the ability to evaluate the isogeny at any point in polynomial time.

Definition 1.1.22 (Efficient representation)**.** [Wes24, Definition 1.3] Let $\mathscr{A}$ be a polynomial time algorithm. We say that $\mathscr{A}$ is an *efficient isogeny evaluator* if for any sequence of bits $D \in \{0,1\}^*$ such that $\mathscr{A}(\mathsf{validity}, D)$ returns True, there exists an isogeny $\varphi : E \longrightarrow E'$ defined over a finite field $\mathbb{F}_q$ such that:

(i) $\mathscr{A}(\mathsf{curves}, D)$ returns (E, E').

(ii) $\mathscr{A}(\mathsf{degree}, D)$ returns $\deg(\varphi)$.

(iii) $\mathscr{A}(\mathsf{eval}, D, P)$ with $P \in E(\mathbb{F}_{q^r})$ returns $\varphi(P)$.

If D is of polynomial size in $\log(\deg(\varphi))$ and $\log(q)$, then we say that D is an *efficient representation* of φ (with respect to $\mathscr{A}$). In that case, $\mathscr{A}(\mathsf{curves}, D)$ and $\mathscr{A}(\mathsf{degree}, D)$ run in polynomial time in $\log(\deg(\varphi))$ and $\log(q)$ and $\mathscr{A}(\mathsf{eval}, D, P)$ runs in polynomial time in $\log(\deg(\varphi))$, $\log(q)$ and r.

This definition generalises to isogenies between abelian varieties that will be introduced in Definition 1.4.25.

When we mention efficient representations, the efficient isogeny evaluator $\mathscr{A}$ is rarely mentioned. As [Wes24] puts it, $\mathscr{A}$ may be seen as an algorithm implementing all the standard techniques to evaluate isogenies and we may require that efficient representations with respect to $\mathscr{A}$ specify which technique to use. We give some efficient isogeny representations that are widely used below.

Example 1.1.23. 1. If $\varphi : E \longrightarrow E'$ is a d-isogeny between elliptic curves defined over a finite field $\mathbb{F}_q$, where $d = \prod_{i=1}^{n} \ell_i$ is a smooth integer with prime factors ℓ_i bounded by a polynomial in $\log(d)$, then φ can be decomposed into a *chain of isogenies*:

$$E \xrightarrow{\varphi_1} E_1 \xrightarrow{\varphi_2} E_2 \quad \cdots \quad E_{n-2} \xrightarrow{\varphi_{n-1}} E_{n-1} \xrightarrow{\varphi_n} E', \tag{1.1}$$

where φ_i is an ℓ_i-isogeny for all $i \in [\![1\ ;\ n]\!]$ and $\varphi = \varphi_n \circ \cdots \circ \varphi_1$. Each φ_i can be expressed by rational fractions $\varphi_i(x,y) = (f_i(x,y), g_i(x,y))$ with numerators and denominators of linear degree in ℓ_i. To evaluate $P \in E(\mathbb{F}_{q^r})$, we can compute $\varphi(P) = \varphi_n \circ \cdots \circ \varphi_1(P)$ by evaluating the φ_i successively. Such an isogeny chain is an efficient representation of φ.

2. If furthermore, the d-isogeny φ from above is cyclic *i.e.* has cyclic kernel, a generator $P \in E(\mathbb{F}_{q^r})$ of $\ker(\varphi)$ with r bounded by a polynomial in $\log(d)$ is an efficient representation of φ. Indeed, we can obtain the isogeny chain described in Eq. (1.1) from the knowledge of P in polynomial time (in $\log(d)$ and $\log(q)$). Indeed, each isogeny φ_i can be expressed by rational fractions $\varphi_i(x,y) = (f_i(x,y), g_i(x,y))$ from the knowledge of its kernel $\ker(\varphi_i)$ in $O(\ell_i)$ operations over the field of definitions of $\ker(\varphi_i)$ with Vélu's formulas [Vé71], or even in $O(\sqrt{\ell_i})$ operations with further improvements from [BDFLS20]. Besides, for all $i \in [\![1\ ;\ n]\!]$, $\ker(\varphi_i)$ is generated by $[\ell_{i+1} \cdots \ell_n]\varphi_{i-1} \circ \cdots \circ \varphi_i(P)$ which can also be computed in polynomial time.

3. Prior to the introduction of higher dimensional isogeny interpolation algorithms, 1 and 2 were almost the only known efficient isogeny representations and we could only efficiently represent isogenies of smooth degrees. Now, we can efficiently represent isogenies of non smooth degree as follows. If $\varphi : E \longrightarrow E'$ is a d-isogeny defined over $\mathbb{F}_q$ and $N > d$ is a smooth integer (with a smoothness bound polynomial in $\log(d)$) and (P,Q) is a basis of $E[N]$ defined over $\mathbb{F}_{q^r}$ with r polynomial in $\log(d)$, then $(\varphi(P), \varphi(Q))$ is an efficient representation of φ. Basically, this means that from the knowledge of the interpolation data $(P, Q, \varphi(P), \varphi(Q))$, we can evaluate φ everywhere else. More details on the underlying higher dimensional isogeny interpolation algorithms will be given in Section 2.2.

By abuse of words, we say usually say that we *compute* an isogeny when we compute an efficient representation of this isogeny or translate an efficient representation from one to another. For instance, we say that we compute an isogeny when we translate a generator of its kernel into an isogeny chain, as explained in point 2 of Example 1.1.23. Similarly, from the interpolation data mentioned in point 3 of Example 1.1.23, we can compute an efficient representation of a higher dimensional isogeny embedding φ given as a chain similar to Eq. (1.1), from which φ can be evaluated everywhere.

1.2 Quaternion algebras and the Deuring correspondence

As we have seen in Theorem 1.1.16, the endomorphism rings of supersingular elliptic curves are isomorphic to maximal orders in the quaternion algebra ramifying at p and ∞. In this section, we define quaternion algebras properly and introduce some basic properties of these objects. We also introduce the Deuring correspondence relating quaternions and supersingular elliptic curves beyond Theorem 1.1.16. We shall see in particular how quaternion ideals and isogenies are related, a property which is crucial to this thesis.

1.2.1 Quaternion algebras, orders, ideals

Definition 1.2.1. A quaternion algebra $\mathcal{B}$ over a field F is a central simple algebra of dimension 4 over F. If $\mathrm{char}(F) \neq 2$, such a quaternion algebra is always isomorphic to a

$$\left(\frac{a,b}{F}\right) := F \oplus Fi \oplus Fj \oplus Fk$$

with $i^2 = a$, $j^2 = b$, $k = ij = -ji$ and $a, b \in F^*$ non-squares in F^* [Voi21, Lemma 2.2.5].

Definition 1.2.2. If $\mathcal{B} := (a,b/F)$ is a quaternion algebra over a field F with $\mathrm{char}(F) \neq 2$, we define the *conjugation* (also called *standard involution*), the *reduced trace* and *reduced norm* of $\mathcal{B}$ as follows:

$$\alpha = x + iy + jz + kt \in \mathcal{B} \longmapsto \overline{\alpha} = x - iy - jz - kt \in \mathcal{B},$$

$$\alpha = x + iy + jz + kt \in \mathcal{B} \longmapsto \mathrm{Tr}(\alpha) = \alpha + \overline{\alpha} = 2x \in F,$$

$$\alpha = x + iy + jz + kt \in \mathcal{B} \longmapsto \mathrm{nrd}(\alpha) = \alpha\overline{\alpha} = x^2 + ay^2 + bz^2 + abt^2 \in F.$$

In the following, we fix R a Dedekind domain and F its ring of fractions (*e.g.* $R = \mathbb{Z}$ and $F = \mathbb{Q}$) and we fix a quaternion algebra $\mathcal{B}$ over F.

Ideals and orders

Definition 1.2.3. A (R-)*lattice* $I \subset \mathcal{B}$ is an R-module of finite type such that $I \otimes_R F = \mathcal{B}$. An ($R$-)*order* $\mathcal{O} \subset \mathcal{B}$ is a (R-)lattice which is also a ring with unit ($1 \in \mathcal{O}$). A *maximal order* $\mathcal{O} \subset \mathcal{B}$ is an order which is maximal for the inclusion: any order $\mathcal{O}' \subset \mathcal{B}$ such that $\mathcal{O} \subseteq \mathcal{O}'$ must be equal to $\mathcal{O}$.

If $\mathcal{O} \subset \mathcal{B}$ is an order, a *fractional left $\mathcal{O}$-ideal* I is a lattice $I \subset \mathcal{B}$ such that $\mathcal{O} \cdot I \subseteq I$ and a *fractional right $\mathcal{O}$-ideal* J is a lattice $J \subset \mathcal{B}$ such that $J \cdot \mathcal{O} \subseteq J$. Conversely, if $I \subset \mathcal{B}$ is a lattice, we can construct the *left order* and the *right order* of I respectively as:

$$O_L(I) := \{\alpha \in \mathcal{B} \mid \alpha \cdot I \subseteq I\} \quad \text{and} \quad O_R(I) := \{\alpha \in \mathcal{B} \mid I \cdot \alpha \subseteq I\}.$$

By construction, I is a fractional left $O_L(I)$-ideal and a fractional right $O_R(I)$-ideal. We say that two lattices I and J are *compatible* when $O_R(I) = O_L(J)$, so that the intermediate orders of the product $I \cdot J$ "match-up".

Lemma 1.2.4. *[Voi21, Lemma 16.6.7] If $I \subset \mathcal{B}$ is a lattice, we have $O_L(I) = O_R(\overline{I})$ and $O_R(I) = O_L(\overline{I})$, where $\overline{I} := \{\overline{\alpha} \mid \alpha \in I\}$.*

A lattice $I \subset \mathcal{B}$ is principal if it is of the form $I = O_L(I)\alpha = \alpha O_R(I)$ for some $\alpha \in \mathcal{B}$ and *locally principal* if for all prime $\mathfrak{p}$ ideal of R, the localisation $I_{(\mathfrak{p})} = I \otimes_R R_{(\mathfrak{p})}$ is principal.

The reduced norm

The *reduced norm* of a lattice $I \subset \mathcal{B}$ denoted by $\mathrm{nrd}(I)$ is the R-submodule of F generated by $\{\mathrm{nrd}(\alpha) \mid \alpha \in I\}$. The reduced norm $\mathrm{nrd}(I)$ is a fractional ideal of R, so it is finitely generated over R [Voi21, Lemma 16.3.2]. By abuse, when $R = \mathbb{Z}$ (which is principal), we denote by $\mathrm{nrd}(I) \in \mathbb{Q}$ the unique non-negative generator of $\mathrm{nrd}(I)$. The norm is a multiplicative function.

Lemma 1.2.5. *[Voi21, Lemma 16.3.7] Let $I, J \subset \mathcal{B}$ be compatible lattices with either I or J locally principal. Then* $\mathrm{nrd}(I \cdot J) = \mathrm{nrd}(I)\,\mathrm{nrd}(J)$.

Lemma 1.2.6. *[Voi21, Lemma 16.3.8] Let $I \subset \mathcal{B}$ be a locally principal lattice. Then $\alpha \in I$ generates I if and only if* $\mathrm{nrd}(I) = \mathrm{nrd}(\alpha)R$.

Integrality

Definition 1.2.7. We say that a lattice $I \subset \mathcal{B}$ is *integral* if $I \subseteq O_R(I) \cap O_L(I)$.

Lemma 1.2.8. *[Voi21, Lemma 16.2.8] Let $I \subset \mathcal{B}$ be a lattice. Then the following are equivalent:*

(i) I *is integral.*

(ii) $I^2 \subseteq I$.

(iii) $I \subseteq O_L(I)$.

(iv) $I \subseteq O_R(I)$.

If I is integral, then every element of I is integral over R i.e. *is the root of a monic polynomial with coefficients in R.*

In particular, an order $\mathcal{O} \subset \mathcal{B}$ is an integral lattice and all of its elements are integral over R and the trace $\mathrm{Tr}(\alpha)$ and reduced norm $\mathrm{nrd}(\alpha)$ of any element $\alpha \in \mathcal{O}$ are defined over R, since $X^2 - \mathrm{Tr}(\alpha) + \mathrm{nrd}(\alpha)$ is the minimal polynomial of α over F, except when $\alpha \in R$.

Invertibility

Definition 1.2.9. A lattice $I \subset \mathcal{B}$ is *left invertible* if there exists a lattice $I' \subset \mathcal{B}$ such that I' is compatible with I and $I' \cdot I = O_R(I)$.

A lattice $I \subset \mathcal{B}$ is *right invertible* if there exists a lattice $I' \subset \mathcal{B}$ such that I is compatible with I' and $I \cdot I' = O_L(I)$.

A lattice $I \subset \mathcal{B}$ is *invertible* if there exists a lattice $I' \subset \mathcal{B}$ (two sided inverse) such that I' is compatible with I, I is compatible with I', $I' \cdot I = O_R(I)$ and $I \cdot I' = O_L(I)$.

Theorem 1.2.10. *[Voi21, Lemma 16.7.7] Let $I \subset \mathcal{B}$ be a lattice. Then the following are equivalent:*

(i) *I is invertible.*

(ii) *I is left invertible.*

(iii) *I is right invertible.*

(iv) *I is locally principal.*

If furthermore, $R = \mathbb{Z}$ and I is integral, the above conditions are equivalent to:

(v) $\mathrm{nrd}(I)^2 = [O_L(I) : I]$.

(vi) $\mathrm{nrd}(I)^2 = [O_R(I) : I]$.

Proposition 1.2.11. *[Voi21, Proposition 16.6.15] Let $I \subset \mathcal{B}$ be a lattice. Then:*

(i) *There exists two orders $\mathcal{O}, \mathcal{O}' \subset \mathcal{B}$ such that $O_L(I) \subset \mathcal{O}$, $O_R(I) \subset \mathcal{O}'$, $I \cdot \overline{I} = \mathrm{nrd}(I)\mathcal{O}$ and $\overline{I} \cdot I = \mathrm{nrd}(I)\mathcal{O}'$.*

(ii) *If either $O_L(I)$ or $O_R(I)$ is maximal, then both orders are maximal and I is invertible. In particular, $I \cdot \overline{I} = \mathrm{nrd}(I)O_L(I)$ and $\overline{I} \cdot I = \mathrm{nrd}(I)O_R(I)$.*

In the following, we shall only work with definite quaternion algebras over $\mathbb{Q}$ ($R = \mathbb{Z}$ and $F = \mathbb{Q}$), which are of the form $\mathcal{B} := (a, b/\mathbb{Q})$ with $a, b < 0$, so that the reduced norm nrd is a positive definite quadratic form. We shall only consider integral lattices $I \subset \mathcal{B}$ which are left $\mathcal{O}$-ideals for some maximal order $\mathcal{O}(= O_L(I))$. Such ideals are always invertible and are right ideals for a maximal order $\mathcal{O}' = O_R(I)$ by Proposition 1.2.11.(ii) so they satisfy the assumptions of Theorem 1.2.10.

1.2.2 The quaternion algebra ramified at p and ∞

Theorem 1.2.12. *[Voi21, Theorem 5.4.4] If F is a field with $\mathrm{char}(F) \neq 2$, then a quaternion algebra $\mathcal{B}$ over F is either a division algebra or isomorphic to $M_2(F)$, the algebra of 2×2 matrices with coefficients in F.*

Let $\mathcal{B}$ be a quaternion algebra over $\mathbb{Q}$. Then, for every place v of $\mathbb{Q}$ ($v = p$ a prime or $v = \infty$), $\mathcal{B} \otimes \mathbb{Q}_v$ is either a division algebra or isomorphic to the matrix algebra $M_2(\mathbb{Q}_v)$. Note that $\mathbb{Q}_p$ is the field of p-adic integers for all prime number p and that $\mathbb{Q}_\infty = \mathbb{R}$.

Definition 1.2.13. We say that a quaternion algebra $\mathcal{B}$ over $\mathbb{Q}$ *ramifies* at a place v of $\mathbb{Q}$ if $\mathcal{B} \otimes \mathbb{Q}_v$ is a division algebra and that $\mathcal{B}$ *splits* at v otherwise *i.e.* if $\mathcal{B} \otimes \mathbb{Q}_v \simeq M_2(\mathbb{Q}_v)$.

Theorem 1.2.14. *[Voi21, Theorem 14.1.3]*

(i) *A quaternion algebra over $\mathbb{Q}$ ramifies at a finite and even number of places.*

(ii) *Two quaternion algebras over $\mathbb{Q}$ are isomorphic if and only if they ramify at the same places.*

(iii) *For all finite set of places of $\mathbb{Q}$ with even cardinality, there exists a quaternion algebra ramifying at those places.*

We define the discriminant *of a quaternion algebra $\mathcal{B}$ denoted by* $\mathrm{disc}(\mathcal{B})$ *as the product of finite places where $\mathcal{B}$ ramifies. By the above, the discriminant determines the isomorphism class of $\mathcal{B}$.*

Corollary 1.2.15. *Let p be a prime. Then there exists a unique quaternion algebra $\mathcal{B}_{p,\infty}$ over $\mathbb{Q}$ up to isomorphism ramifying at p and ∞* i.e. *of discriminant p.*

(i) If $p = 2$, then $\mathcal{B}_{p,\infty} = (-1, -1/\mathbb{Q})$.

(ii) If $p \equiv 3 \mod 4$, then $\mathcal{B}_{p,\infty} = (-1, -p/\mathbb{Q})$.

(iii) If $p \equiv 1 \mod 4$, then $\mathcal{B}_{p,\infty} = (-q, -p/\mathbb{Q})$, where q is a prime $\equiv 3 \mod 4$ which is a not a quadratic residue modulo p.

Proof. See [Voi21, Example 14.2.13]. □

In the following, we shall see that the quaternion algebra $\mathcal{B}_{p,\infty}$ is crucial for the Deuring correspondence. Indeed, we have seen in Theorem 1.1.16 that endomorphism rings of supersingular elliptic curves are isomorphic to maximal orders in $\mathcal{B}_{p,\infty}$. The following gives a criteria to identify maximal orders.

Definition 1.2.16. If $\mathcal{O} \subset \mathcal{B}$ is an order in a quaternion algebra defined over $\mathbb{Q}$, and $(\alpha_1, \cdots, \alpha_4)$ is $\mathbb{Z}$-basis of $\mathcal{O}$, we define the *discriminant* of $\mathcal{O}$ as:

$$\mathrm{disc}(\mathcal{O}) = |\det(\mathrm{Tr}(\alpha_i \overline{\alpha}_j))_{1 \leq i,j \leq 4}|$$

This quantity is a well-defined invariant of $\mathcal{O}$ as does not depend on the $\mathbb{Z}$-basis. It can be proved that $\mathrm{disc}(\mathcal{O})$ is always a square in $\mathbb{N}^*$ and define the *reduced discriminant* as $\mathrm{discrd}(\mathcal{O}) := \sqrt{\mathrm{disc}(\mathcal{O})}$ [Voi21, Lemma 15.4.7].

Theorem 1.2.17. *[Voi21, Theorem 15.5.5] Let $\mathcal{B}/\mathbb{Q}$ be a quaternion algebra and $\mathcal{O} \subset \mathcal{B}$ be an order. Then $\mathcal{O}$ is maximal if and only if* $\mathrm{discrd}(\mathcal{O}) = \mathrm{disc}(\mathcal{B})$.

In particular, if $\mathcal{B} = \mathcal{B}_{p,\infty}$, then $\mathcal{O}$ is maximal if and only if $\mathrm{discrd}(\mathcal{O}) = p$.

1.2.3 Ideal equivalence

Definition 1.2.18. Let $\mathcal{B} := (a, b/\mathbb{Q})$ with $a, b < 0$. Two lattices $I, J \subset \mathcal{B}$ are (right) *equivalent* and we denote $I \sim J$ if there exists $\alpha \in \mathcal{B}^*$ such that $J = I\alpha$.

We can define a left equivalence similarly, but it will not be used. Only the right equivalence is useful for the Deuring correspondence. The reader will find below an explicit way to describe equivalent integral left ideals of a maximal order and how their right orders relate.

Lemma 1.2.19. *Let $\mathcal{B} := (a, b/\mathbb{Q})$ with $a, b < 0$ and $I \subset \mathcal{B}$ be an integral ideal such that $O_L(I)$ is maximal. Then:*

(i) The map:

$$\chi_I : \alpha \in I \setminus \{0\} \longmapsto I\frac{\overline{\alpha}}{\mathrm{nrd}(I)}$$

is a surjection to the set of integral ideals (right) equivalent to I.

(ii) For all $\alpha, \beta \in I \setminus \{0\}$, $\chi_I(\alpha) = \chi_I(\beta)$ if and only if there exists $\gamma \in O_R(I)^\times$ such that $\beta = \alpha\gamma$.

(iii) If $J = \chi_I(\alpha)$ is an integral ideal (right) equivalent to I, with $\alpha \in I \setminus \{0\}$, then we have

$$O_R(J) = \alpha \cdot O_R(I) \cdot \alpha^{-1}, \quad O_R(\alpha^{-1} \cdot J \cdot \alpha) = O_R(I), \quad \text{and} \quad O_L(\alpha^{-1} \cdot \overline{J} \cdot \alpha) = O_R(I).$$

Proof. Points (i) and (ii) are due to [DFKLPW20, Lemma 1]. Their proof being simple, we recall it here.

(i) If $\alpha \in I \setminus \{0\}$, then $J := \chi_I(\alpha) = I\overline{\alpha}/\mathrm{nrd}(I)$ is equivalent to a fractional ideal equivalent to I. Besides, $I\overline{\alpha} \subseteq I \cdot \overline{I} = \mathrm{nrd}(I)O_L(I)$ by Proposition 1.2.11, so that $J \subseteq O_L(I)$ and J is integral.

Conversely, if $J \subset \mathcal{B}$ is an integral ideal right equivalent to I, then $J = I \cdot \beta$, with $\beta \in \mathcal{B} \setminus \{0\}$. We then have $\overline{I} \cdot J = \overline{I} \cdot I \cdot \beta = \mathrm{nrd}(I) O_R(I) \cdot \beta$ by Proposition 1.2.11 and it follows that $\mathrm{nrd}(I)\beta \in \overline{I} \cdot J$. In addition, $J \subseteq O_L(J) = O_L(I) = O_R(\overline{I})$ since J is integral, so that $\overline{I} \cdot J \subseteq \overline{I}$ and $\mathrm{nrd}(I)\beta \in \overline{I} \setminus \{0\}$. Hence, there exists $\alpha \in I \setminus \{0\}$, such that $\beta = \overline{\alpha}/\mathrm{nrd}(I)$, so finally $J = \chi_I(\alpha)$.

(ii) If $\chi_I(\alpha) = \chi_I(\beta)$, then $I\overline{\alpha}/\mathrm{nrd}(I) = I\overline{\beta}/\mathrm{nrd}(I)$, so $\overline{I} \cdot I\overline{\alpha}/\mathrm{nrd}(I) = \overline{I} \cdot I\overline{\beta}/\mathrm{nrd}(I)$ *i.e.* $O_R(I)\overline{\alpha} = O_R(I)\overline{\beta}$ by Proposition 1.2.11. It follows that there exits $\gamma, \delta \in O_R(I)$ such that $\gamma\overline{\alpha} = \overline{\beta}$ and $\overline{\alpha} = \delta\overline{\beta}$ so that $\gamma\delta\overline{\beta} = \overline{\beta}$ and $\gamma\delta = 1$, so $\gamma \in O_R(I)^\times$ and finally $\beta = \alpha\overline{\gamma}$ with $\overline{\gamma} \in O_R(I)^\times$. The converse being trivial, this proves (ii).

(iii) We have:

$$\begin{aligned} O_R(J) &= \{\beta \in \mathcal{B}_{p,\infty} \mid J\beta \subseteq J\} = \left\{\beta \in \mathcal{B}_{p,\infty} \middle| I\frac{\overline{\alpha}\beta}{\mathrm{nrd}(I)} \subseteq I\frac{\overline{\alpha}}{\mathrm{nrd}(I)}\right\} = \left\{\beta \in \mathcal{B}_{p,\infty} \middle| I\frac{\overline{\alpha}\beta\alpha}{\mathrm{nrd}(\alpha)} \subseteq I\right\} \\ &= \{\beta \in \mathcal{B}_{p,\infty} \mid I\alpha^{-1}\beta\alpha \subseteq I\} = \{\beta \in \mathcal{B}_{p,\infty} \mid \alpha^{-1}\beta\alpha \in O_R(I)\} = \alpha \cdot O_R(I) \cdot \alpha^{-1}. \end{aligned}$$

Similarly,

$$\begin{aligned} O_R(\alpha^{-1} \cdot J \cdot \alpha) &= \{\beta \in \mathcal{B}_{p,\infty} \mid \alpha^{-1} \cdot J \cdot \alpha\beta \subseteq \alpha^{-1} \cdot J \cdot \alpha\} \\ &= \left\{\beta \in \mathcal{B}_{p,\infty} \middle| \alpha^{-1} \cdot I \cdot \frac{\overline{\alpha}\alpha\beta}{\mathrm{nrd}(I)} \subseteq \alpha^{-1} \cdot I \cdot \frac{\overline{\alpha}\alpha}{\mathrm{nrd}(I)}\right\} \\ &= \{\beta \in \mathcal{B}_{p,\infty} | \ I \cdot \beta \subseteq I\} = O_R(I). \end{aligned}$$

It follows that $O_L(\overline{\alpha^{-1} \cdot J \cdot \alpha}) = O_R(\alpha^{-1} \cdot J \cdot \alpha) = O_R(I)$. But

$$\overline{\alpha^{-1} \cdot J \cdot \alpha} = \overline{\alpha} \cdot \overline{J} \cdot \overline{\alpha}^{-1} = \mathrm{nrd}(\alpha)\alpha^{-1} \cdot \overline{J} \cdot \frac{\alpha}{\mathrm{nrd}(\alpha)} = \alpha^{-1} \cdot \overline{J} \cdot \alpha$$

This completes the proof. □

1.2.4 The Deuring correspondence

We now get to the main topic of this section: the Deuring correspondence which owes its name and discovery to Max Deuring [Deu41]. As we have seen in Theorem 1.1.16 endomorphism rings of supersingular elliptic curves defined over $\overline{\mathbb{F}}_p$ are isomorphic to a maximal order in the quaternion algebra $\mathcal{B}_{p,\infty}$ ramifying at p and ∞. We can actually not only connect endomorphism rings with quaternions but also isogenies as follows.

Let $E/\overline{\mathbb{F}}_p$ be a supersingular elliptic curve, $\mathcal{O} \subset \mathcal{B}_{p,\infty}$ be a maximal order isomorphic to $\mathrm{End}(E)$ via $\varepsilon : \mathcal{O} \xrightarrow{\sim} \mathrm{End}(E)$ and I be a left $\mathcal{O}$-ideal (integral by assumption as previously stated). We can consider the subgroup:

$$E[I] = \bigcap_{\alpha \in I} \ker(\varepsilon(\alpha)) \subset E(\overline{\mathbb{F}}_p),$$

If $\mathrm{nrd}(I)$ is coprime with p, by Theorem 1.1.7 we can associate to I a separable isogeny $\varphi_I : E \longrightarrow E_I$ of kernel $E[I]$. If $p | \mathrm{nrd}(I)$, we can factor $I = P^r I'$ with I' a left $\mathcal{O}$-ideal of norm coprime with p and P the unique two sided ideal of $\mathcal{O}$ of norm p (whose existence and uniqueness is ensured by [Voi21, Theorem 18.1.3]). Then we associate to I the isogeny $\varphi_I := \varphi_{I'} \circ \pi_{p^r}$, where $\pi_{p^r} : E \longrightarrow E^{(p^r)}$ is the p^r-th Frobenius corresponding to P^r. Note $E^{(p^r)}$ is isomorphic to E if r is even and $E^{(p)}$ if r is odd.

Remark 1.2.20. In [Voi21, Chapter 42], Voight defines $E[I]$ as a group scheme to deal with the inseparable case. We shall introduce group schemes in Section 1.4.4 but we prefer here the elementary point of view from above.

The map $I \longmapsto \varphi_I$ defines a correspondence between left $\mathcal{O}$-ideals and isogenies with domain E, since it has an inverse map associating a left $\mathcal{O}$-ideal:

$$I_\varphi = \{\alpha \in \mathcal{O} \mid \forall P \in \ker(\varphi), \varepsilon(\alpha)(P) = 0\}$$

to any isogeny $\varphi : E \longrightarrow E'$. The ideal I_φ is called the *kernel ideal* of φ. The correspondence satisfies the following properties.

Proposition 1.2.21.

(i) $\mathrm{nrd}(I) = \deg(\varphi_I)$.

(ii) $I_{\varphi_I} = I$, *which proves that* $I \longmapsto \varphi_I$ *has an inverse map.*

(iii) $\mathrm{Hom}(E_I, E)$ *is isomorphic to* I *as a left* $\mathcal{O}$*-module via the map:*

$$\begin{array}{rcl} \mathrm{Hom}(E_I, E) & \xrightarrow{\sim} & I \\ \psi & \longmapsto & \varepsilon^{-1}(\psi \circ \varphi_I) \end{array}$$

(iv) $\mathrm{End}(E_I)$ *is isomorphic to* $O_R(I)$ *as a ring via the map:*

$$\begin{array}{rcl} \mathrm{End}(E_I) & \xrightarrow{\sim} & O_R(I) \\ \psi & \longmapsto & \dfrac{1}{\mathrm{nrd}(I)}\varepsilon^{-1}(\widehat{\varphi}_I \circ \psi \circ \varphi_I) \end{array}$$

(v) *If* $J \sim I$ *is a left* $\mathcal{O}$*-ideal equivalent to* I*, then the respective codomains* E_I *and* E_J *of* φ_I *and* φ_J *are isomorphic.*

(vi) *If* J *is an integral ideal compatible with* I (i.e. *such that* $O_L(J) = O_R(I)$)*, then* $\varphi_{I\cdot J} = \varphi_J \circ \varphi_I$.

(vii) *For all* $\alpha \in \mathcal{O} \setminus \{0\}$*, the isogeny* $\varphi_{\mathcal{O}\alpha} : E \longrightarrow E_{\mathcal{O}\alpha}$ *associated to the principal ideal* $\mathcal{O}\alpha$ *is the endomorphism* $\varepsilon(\alpha) \in \mathrm{End}(E)$.

(viii) *The isogeny* $\varphi_{\overline{I}} : E_I \longrightarrow E_{\overline{I}}$ *associated to* $\overline{I}$ *is the dual* $\widehat{\varphi}_I : E_I \longrightarrow E$.

Proof. (i) and (ii) are [Voi21, Proposition 42.2.16], (iii) is [Voi21, Lemma 42.2.7], (iv) is [Voi21, Lemma 42.2.9], (v) is [Voi21, Lemma 42.2.13], (vi) is [Wat69, Proposition 3.12], (vii) follows from the definition of the Deuring correspondence and (viii) follows from (vi), (vii) and Proposition 1.2.11.(ii). □

From the above properties, we obtain the main theorem of the Deuring correspondence.

Theorem 1.2.22 (Deuring correspondence)**.**

(i) *For every isogeny* $\varphi : E \longrightarrow E'$*, there exists a left* $\mathcal{O}$*-ideal* I *and an isomorphism* $\lambda : E_I \xrightarrow{\sim} E'$ *such that* $\varphi = \lambda \circ \varphi_I$.

(ii) *For any maximal order* $\mathcal{O}' \subset \mathcal{B}_{p,\infty}$*, there exists a supersingular elliptic curve* $E/\mathbb{F}_{p^2}$ *such that* $\mathrm{End}(E') \simeq \mathcal{O}'$.

(iii) *The map* $[I] \longmapsto j(E_I)$ *induces a bijection between equivalence classes of left* $\mathcal{O}$*-ideals and isomorphism classes of supersingular elliptic curves over* $\overline{\mathbb{F}}_p$.

Proof. **(i)** We refer to the proof of [Voi21, Corollary 42.2.21].

(ii) We give the argument which is actually very instructive for isogeny based cryptography. Let $\mathcal{O}' \subset \mathcal{B}_{p,\infty}$ be a maximal order. Then we can construct a *connecting ideal* I between $\mathcal{O}$ and $\mathcal{O}'$ *i.e.* a left $\mathcal{O}$-ideal which is also a right $\mathcal{O}'$-deal as follows. For instance, we take $I := N\mathcal{O} \cdot \mathcal{O}'$ where $N := [\mathcal{O} \cdot \mathcal{O}' : \mathcal{O}]$ or any common denominator. We then consider $\varphi_I : E \longrightarrow E_I$ where $\mathrm{End}(E_I) \simeq O_R(I)$ by Proposition 1.2.21 and $O_R(I) = \mathcal{O}'$ by construction.

(iii) We refer to [Voi21, Corollary 42.3.7]. The definition of the map $[I] \longmapsto j(E_I)$ makes sense by Proposition 1.2.21.(v). The surjectivity is a consequence of (ii) and the injectivity comes from Proposition 1.2.21.(iii) and the fact that left $\mathcal{O}$-ideals that are isomorphic as left $\mathcal{O}$-modules are (right) equivalent [Voi21, Lemma 17.3.3]. □

In the spirit of Theorem 1.2.22.(iii), a natural question arises: is there a bijection between supersingular j-invariants and maximal orders? The answer is no, but this is close to be true.

Proposition 1.2.23. *[Voi21, Lemma 42.4.1] Let* $\mathcal{O} \subset \mathcal{B}_{p,\infty}$*. Then there exists at most two isomorphism classes of supersingular elliptic curves* E *such that* $\mathrm{End}(E) \simeq \mathcal{O}$*. These classes are given by* $j(E)$ *and its Galois conjugate* $j(E)^p = j(E^{(p)})$.

Supersingular elliptic curves	Quaternions
$j(E)$ or $j(E)^p$ supersingular	$\mathcal{O} \simeq \mathrm{End}(E)$ maximal order in $\mathcal{B}_{p,\infty}$
$\varphi : E \longrightarrow E'$	left $\mathcal{O}$-ideal and right $\mathcal{O}'$-ideal I_φ
$\varphi, \psi : E \longrightarrow E'$	$I_\varphi \sim I_\psi$
$\widehat{\varphi}$	$\overline{I}_\varphi$
$\varphi \circ \psi$	$I_\psi \cdot I_\varphi$
$\alpha \in \mathrm{End}(E)$	$\mathcal{O}\alpha$
$\deg(\varphi)$	$\mathrm{nrd}(I_\varphi)$

Table 1.1: The Deuring correspondence (credits to [DFKLPW20]).

The Deuring correspondence can be summarised in Table 1.1 due to [DFKLPW20].

We terminate this presentation of the properties of the Deuring correspondence with a technical lemma that will be used quite often.

Lemma 1.2.24. *Let E be a supersingular elliptic curve defined over $\mathbb{F}_{p^2}$, $\mathcal{O}$ be a maximal order of $\mathcal{B}_{p,\infty}$ isomorphic to $\mathrm{End}(E)$ via the isomorphism $\varepsilon : \mathcal{O} \xrightarrow{\sim} \mathrm{End}(E)$ and I, J be two (right) equivalent integral left-ideals of $\mathcal{O}$ such that $J = \chi_I(\alpha) = I \cdot \overline{\alpha}/\mathrm{nrd}(I)$ with $\alpha \in I \setminus \{0\}$. Then, we have:*

$$\widehat{\varphi}_J \circ \varphi_I = \varepsilon(\alpha).$$

Proof. By Lemma 1.2.19.(iii), we have $O_L(\alpha^{-1} \cdot \overline{J} \cdot \alpha) = O_R(I)$ so I and $\alpha^{-1} \cdot \overline{J} \cdot \alpha$ are compatible and we can consider their product:

$$I \cdot \alpha^{-1} \cdot \overline{J} \cdot \alpha = I \cdot \frac{\alpha^{-1} \cdot \alpha}{\mathrm{nrd}(I)} \cdot \overline{I} \cdot \alpha = \frac{1}{\mathrm{nrd}(I)} I \cdot \overline{I} \cdot \alpha = \mathcal{O} \cdot \alpha.$$

By the Deuring correspondence (Proposition 1.2.21.(vi) and (vii)), $I \cdot \alpha^{-1} \cdot \overline{J} \cdot \alpha$ corresponds to $\widehat{\varphi}_J \circ \varphi_I$, so that $\widehat{\varphi}_J \circ \varphi_I = \varepsilon(\alpha)$. □

An easy instance of the endomorphism ring problem

The Deuring correspondence is used to compute isogenies and construct cryptographic protocols (as we shall see in Section 2.1.1). The Deuring correspondence cannot be used effectively without knowing the endomorphism ring of the starting curve, and computing it is a hard problem in general. However, there are easy instances of this problem. The following is a classical easy instance that will be used without moderation in the rest of this thesis.

Lemma 1.2.25. *Let p be a prime $\equiv 3 \mod 4$ and E_0 be defined over $\mathbb{F}_p$ by the equation $y^2 = x^3 + x$. Then E_0 is supersingular (and even maximal) and its endomorphism ring is isomorphic to $\mathcal{O}_0 := \langle 1, i, (i+j)/2, (1+ij)/2 \rangle \subset \mathcal{B}_{p,\infty}$ via the isomorphism:*

$$\begin{array}{rcl} \varepsilon_0 : \mathcal{O}_0 & \xrightarrow{\sim} & \mathrm{End}(E_0) \\ i & \longmapsto & \iota : (x, y) \longmapsto (-x, \zeta_4 y) \\ j & \longmapsto & \pi_p : (x, y) \longmapsto (x^p, y^p) \end{array}$$

where $\zeta_4 \in \mathbb{F}_{p^2}$ is a square root of -1.

Proof. This endomorphism ring computation is very classical in the literature but there is no formal proof available, so we provide it here. By Lemma 1.1.19, E_0 is maximal so it is supersingular and its p-th Frobenius endomorphism satisfies $\pi_p^2 = \pi_{p^2} = -[p]$ by Remark 1.1.18.

We also verify immediately that $\iota^2 = [-1]$ and obtain that:

$$\pi_p \circ \iota(x, y) = ((-x)^p, \zeta_4^p y^p) = (-x^p, -\zeta_4^p y^p) = -(-x^p, \zeta_4^p y^p) = -\iota \circ \pi_p(x, y),$$

where we used the fact that $p \equiv 3 \mod 4$. Hence, ι and π_p corresponds to i and j in $\mathcal{B}_{p,\infty}$ respectively.

Finally, $E_0[2] = \{0, (0,0), (\zeta_4, 0), (-\zeta_4, 0)\}$ and

$$\pi_p(0,0) = (0,0) = \iota(0,0), \quad \iota \circ \pi_p(0,0) = (0,0),$$
$$\pi_p(\zeta_4,0) = (\zeta_4^p,0) = (-\zeta_4,0) = \iota(\zeta_4,0), \quad \iota \circ \pi_p(\zeta_4,0) = (-\zeta_4^p,0) = (\zeta_4,0),$$
$$\pi_p(-\zeta_4,0) = ((-\zeta_4)^p,0) = (\zeta_4,0) = \iota(-\zeta_4,0), \quad \iota \circ \pi_p(-\zeta_4,0) = ((-\zeta_4)^p,0) = (\zeta_4,0),$$

so that $\iota + \pi_p$ and $1 + \iota \circ \pi_p$ annihilate $E_0[2]$ and factor through $[2]$. This proves that ε_0 is well-defined.

To conclude, it suffices to prove that $\mathcal{O}_0$ is a maximal order in $\mathcal{B}_{p,\infty}$, which will ensure that ε_0 is surjective, so is an isomorphism (since it is trivially injective). We have:

$$\begin{aligned}
\operatorname{discrd}(\mathcal{O}_0) &= \begin{vmatrix} \operatorname{Tr}(1) & \operatorname{Tr}(\bar{i}) & \operatorname{Tr}\left(\overline{\frac{i+j}{2}}\right) & \operatorname{Tr}\left(\overline{\frac{1+ij}{2}}\right) \\ \operatorname{Tr}(i) & \operatorname{Tr}(i\bar{i}) & \operatorname{Tr}\left(i\overline{\frac{i+j}{2}}\right) & \operatorname{Tr}\left(i\overline{\frac{1+ij}{2}}\right) \\ \operatorname{Tr}\left(\frac{i+j}{2}\right) & \operatorname{Tr}\left(\frac{i+j}{2}\bar{i}\right) & \operatorname{Tr}\left(\frac{i+j}{2}\overline{\frac{i+j}{2}}\right) & \operatorname{Tr}\left(\frac{i+j}{2}\overline{\frac{i+j}{2}}\right) \\ \operatorname{Tr}\left(\frac{1+ij}{2}\right) & \operatorname{Tr}\left(\frac{1+ij}{2}\bar{i}\right) & \operatorname{Tr}\left(\frac{1+ij}{2}\overline{\frac{i+j}{2}}\right) & \operatorname{Tr}\left(\frac{1+ij}{2}\overline{\frac{1+ij}{2}}\right) \end{vmatrix}^{\frac{1}{2}} \\
&= \begin{vmatrix} 2 & 0 & 0 & 1 \\ 0 & 2 & 1 & 0 \\ 0 & 1 & \frac{p+1}{2} & 0 \\ 1 & 0 & 0 & \frac{p+1}{2} \end{vmatrix}^{\frac{1}{2}} = \begin{vmatrix} 1 & 0 & 0 & \frac{p+1}{2} \\ 0 & 1 & \frac{p+1}{2} & 0 \\ 0 & 2 & 1 & 0 \\ 2 & 0 & 0 & 1 \end{vmatrix}^{\frac{1}{2}} \\
&= \begin{vmatrix} 1 & 0 & 0 & \frac{p+1}{2} \\ 0 & 1 & \frac{p+1}{2} & 0 \\ 0 & 0 & -p & 0 \\ 0 & 0 & 0 & -p \end{vmatrix}^{\frac{1}{2}} = p = \operatorname{disc}(\mathcal{B}_{p,\infty})
\end{aligned}$$

It follows that $\mathcal{O}_0$ is a maximal order by Theorem 1.2.17. This completes the proof. □

1.2.5 Lattices of rank 4

Quaternion ideals can be seen as full-rank Euclidean lattices via an embedding into $\mathbb{R}^4$. Namely, $\mathcal{B} = (a, b/\mathbb{Q})$ can be embedded via the map $\iota : \mathcal{B} \hookrightarrow \mathbb{R}^4$,

$$1 \longmapsto (1,0,0,0),\ i \longmapsto (0, \sqrt{|a|}, 0, 0),\ j \longmapsto (0, 0, \sqrt{|b|}, 0),\ ij \longmapsto (0,0,0,\sqrt{|ab|}), \tag{1.2}$$

which is an isometry because $\|\iota(\alpha)\|^2 = \operatorname{nrd}(\alpha)$ for all $\alpha \in \mathcal{B}$. In the course of this thesis, for algorithmic applications of the Deuring correspondence, we shall study short vectors of quaternion ideals so we give some definitions and preliminary results that will be needed.

Definition 1.2.26. Let $\Lambda \subset \mathbb{R}^d$ be a lattice of rank d. For all $i \in [\![1\ ;\ d]\!]$, the i-th *minimum* of Λ denoted by $\lambda_i(\Lambda)$ is the quantity:

$$\lambda_i(\Lambda) = \min\left\{ \max_{1 \le j \le i} \|v_j\| \,\middle|\, v_1, \cdots, v_i \in \Lambda \text{ are linearly independent} \right\}.$$

In general, there is no basis reaching all successive minima of a lattice. For algorithmic applications and to prove some bounds on the successive minima (as defined above) of quaternion ideals, we use basis of short vectors that are Minkowski reduced in the following sense.

Definition 1.2.27. Let $\Lambda \subset \mathbb{R}^d$ be a lattice of rank d. We say that a basis $(b_1, \cdots, b_d)$ of Λ is *Minkowski reduced* if for all $i \in [\![1\ ;\ d]\!]$, b_i has minimal norm such that $(b_1, \cdots, b_i)$ can be extended into a basis of Λ. In particular, $\|b_1\| \le \cdots \le \|b_d\|$.

In low dimension, Minkowski reduced basis reach the successive minima, as desired.

Theorem 1.2.28. *[Wae56] If Λ is a lattice of rank d and $(b_1, \cdots, b_d)$ is a Minkowski reduced basis of Λ, then $\|b_i\| = \lambda_i(\Lambda)$ for all $1 \le i \le \min(d, 4)$. In particular, if $d \le 4$, then a Minkowski reduced basis of Λ reaches all successive minima of Λ.*

Additionally, in practice, Minkowski reduced basis can be computed efficiently in dimension up to four.

Theorem 1.2.29. *[NS09, Theorem 4.2.1] Let $\Lambda \subset \mathbb{R}^d$ be a full rank lattice with $d \leq 4$. Then, given $(b_1, \cdots, b_d)$, a basis of Λ with ordered vectors $\|b_1\| \leq \cdots \leq \|b_d\|$, a Minkowski reduced basis of Λ can be computed in time $O(\log \|b_d\|(1 + \log \|b_d\| - \log \lambda_1(\Lambda)))$.*

1.3 Oriented supersingular elliptic curves

In the previous section, we have seen the Deuring correspondence between quaternion ideals and isogenies between supersingular elliptic curves. There is also an analogue for ordinary elliptic curves with complex multiplication by some quadratic imaginary order $\mathfrak{O}$ *i.e.* whose endomorphism ring is isomorphic to $\mathfrak{O}$. The ideal class group $\mathrm{Cl}(\mathfrak{O})$ acts faithfully on these curves. Unlike the quaternion ideal action (where ideal equivalence classes do not even define a group), this is a commutative group action. This commutativity property has been leveraged to propose cryptographic schemes like Diffie-Hellman key exchange [Cou06; RS06]. Later, a commutative group action has been introduced for supersingular elliptic curves to define the CSIDH scheme (Commutative Supersingular Isogeny Diffie-Hellman) [CLMPR18] followed-up by more general works by Leonardo Colò, David Kohel [CK20] and Horoshi Onuki [Onu21].

1.3.1 Oriented supersingular elliptic curves and isogenies

The key idea is to restrict to a subring of the endomorphism ring of supersingular elliptic curves defining an *orientation*.

Definition 1.3.1. Let E be a supersingular elliptic curve and K be a quadratic imaginary field. A *K-orientation* of E is an embedding $\iota : K \hookrightarrow \mathrm{End}(E) \otimes \mathbb{Q}$. If $\mathfrak{O} \subset K$ is an order such that $\iota(\mathfrak{O}) \subset \mathrm{End}(E)$, we say that ι is an *$\mathfrak{O}$-orientation*. If furthermore, $\mathfrak{O} = \iota^{-1}(\mathrm{End}(E))$ *i.e.* if there is no superorder of $\mathfrak{O}$ mapping to $\mathrm{End}(E)$, we say that ι is a *primitive $\mathfrak{O}$-orientation*. We say that (E, ι) is a K or $\mathfrak{O}$-oriented supersingular elliptic curve.

In the following, we shall only consider primitive orientations and drop the mention of primitive when there is no ambiguity.

Definition 1.3.2. If K is a quadratic imaginary order, a *K-oriented isogeny* between two K-oriented elliptic curves $\varphi : (E, \iota) \longrightarrow (E', \iota')$ is a an isogeny satisfying:

$$\forall \alpha \in K, \quad \iota'(\alpha) = \varphi_*(\iota)(\alpha) := \frac{1}{\deg(\varphi)} \varphi \circ \iota(\alpha) \circ \widehat{\varphi}.$$

Given a K-orientation ι on E and an isogeny $\varphi : E \longrightarrow E'$, one can always define a K-orientation ι' on E' by the above formula.

If (E, ι) and (E', ι') are respectively (primitively) $\mathfrak{O}$ and $\mathfrak{O}'$-oriented elliptic curves then, we say that:

- φ is *descending* if $\mathfrak{O}' \subsetneq \mathfrak{O}$.
- φ is *horizontal* if $\mathfrak{O}' = \mathfrak{O}$.
- φ is *ascending* if $\mathfrak{O} \subsetneq \mathfrak{O}'$.

1.3.2 The ideal class group action

In the following, we fix $\mathfrak{O}$ an order of a quadratic imaginary field K. Let (E, ι) be a (primitively) $\mathfrak{O}$-oriented elliptic curve. Then, we can define the action of an ideal $\mathfrak{a} \subseteq \mathfrak{O}$ of norm $N(\mathfrak{a})$ coprime with p like in the Deuring correspondence as follows. Let:

$$E[\mathfrak{a}] = \bigcap_{\alpha \in \mathfrak{a}} \ker(\iota(\alpha)) \subset E(\overline{\mathbb{F}}_p),$$

and consider the separable isogeny $\varphi_{\mathfrak{a}} : E \longrightarrow E_{\mathfrak{a}}$ with kernel $E[\mathfrak{a}]$ and the induced K-orientation $\iota_{\mathfrak{a}} := \varphi_{\mathfrak{a}*}(\iota)$. This action by $\mathfrak{O}$-ideals satisfies natural properties of the Deuring correspondence.

Proposition 1.3.3. *(i) $\varphi_{\mathfrak{a}}$ is either horizontal or ascending. If $\mathfrak{a}$ is invertible in $\mathfrak{O}$, then $\varphi_{\mathfrak{a}}$ is horizontal.*

(ii) $\deg(\varphi_{\mathfrak{a}}) = N(\mathfrak{a})$.

(iii) If $\mathfrak{a}, \mathfrak{b} \subseteq \mathfrak{O}$ are invertible and of norms coprime with p, then $\varphi_{\mathfrak{ab}} = \varphi_{\mathfrak{b}} \circ \varphi_{\mathfrak{a}}$.

(iv) If $\mathfrak{a}, \mathfrak{b} \subseteq \mathfrak{O}$ are invertible and of norms coprime with p, then $\mathfrak{a} \cdot (E, \iota) \simeq \mathfrak{a} \cdot (E, \iota)$ via a K-oriented isomorphism if and only if $\mathfrak{a}$ and $\mathfrak{b}$ are equivalent i.e. *are equal up to multiplication by an element of K^*.*

(v) If (E_1, ι_1) and (E_2, ι_2) are $\mathfrak{O}$-oriented elliptic curves, then there exists an invertible ideal $\mathfrak{a} \subseteq \mathfrak{O}$ of norm coprime with p such that $(E_2, \iota_2) \simeq \mathfrak{a} \cdot (E_1, \iota_1)$ or $(E_2^{(p)}, \pi_{p*}(\iota_2)) \simeq \mathfrak{a} \cdot (E_1, \iota_1)$, the latter resulting from the application of the p-th Frobenius isogeny on (E_2, ι_2).

Proof. (i) is [Onu21, Proposition 3.5], (ii) is [Wat69, Proposition 3.15], (iii) is [Wat69, Proposition 3.12], (iv) have been proved in the proof of [Onu21, Theorem 3.4] and (v) is [Onu21, Proposition 3.3]. □

Note that $\mathfrak{a}$ is always invertible if $\mathfrak{O} = \mathfrak{O}_K$ is the maximal order or if the norm $N(\mathfrak{a})$ is coprime with the conductor $[\mathfrak{O}_K : \mathfrak{O}]$ of $\mathfrak{O}$ [Cox13, Lemma 7.18]. In that case, by Proposition 1.3.3.(i), $\mathfrak{a} \cdot (E, \iota) := (E_{\mathfrak{a}}, \iota_{\mathfrak{a}})$ is also $\mathfrak{O}$-oriented. Also, by Proposition 1.3.3, the isomorphism class of $\mathfrak{a} \cdot (E, \iota)$ only depends on the ideal class $[\mathfrak{a}]$ and we can compose this action by ideals. This defines a free action of the ideal class group $\mathrm{Cl}(\mathcal{O})$ on isomorphism classes of $\mathfrak{O}$-oriented elliptic curves denoted by $\mathrm{Ell}_p(\mathfrak{O})$.

Theorem 1.3.4. *[Onu21, Theorem 3.4] The map*

$$([\mathfrak{a}], [(E, \iota)]) \in \mathrm{Cl}(\mathfrak{O}) \times \mathrm{Ell}_p(\mathfrak{O}) \longmapsto [\mathfrak{a} \cdot (E, \iota)]$$

introduced above defines a free group action which is either transitive or admits two orbits which are conjugate of each other by application of the p-th Frobenius isogeny.

1.3.3 Commutative Supersingular Isogeny Diffie-Hellman (CSIDH)

In CSIDH [CLMPR18], we consider supersingular elliptic curves E defined over $\mathbb{F}_p$ and the orientations mapping to $\mathrm{End}_{\mathbb{F}_p}(E)$, the subring of $\mathrm{End}(E)$ (of rank 2) made of $\mathbb{F}_p$-rational endomorphisms. The ring $\mathrm{End}_{\mathbb{F}_p}(E)$ always contains the p-th Frobenius endomorphism $\pi_p : (x, y) \in E \longmapsto (x^p, y^p) \in E$. Since E is supersingular and defined over $\mathbb{F}_p$, $\mathrm{Tr}(\pi_p) \equiv 0 \mod p$ by Theorem 1.1.16 and $|\mathrm{Tr}(\pi_p)| \leq 2\sqrt{p}$ by Hesse-Weil's bound (Theorem 1.1.13) so $\mathrm{Tr}(\pi_p) = 0$ and $\pi_p^2 = -[p]$. As a consequence, π_p identifies with $\sqrt{-p}$ so that E is oriented by $\mathbb{Q}(\sqrt{-p})$ so it is primitively oriented either by the maximal order $\mathbb{Z}[(1+\sqrt{-p})/2]$ if $\mathrm{End}_{\mathbb{F}_p}(E) = \mathbb{Z}[(1+\pi_p)/2]$ or by the order of conductor 2 if $\mathbb{Z}[\sqrt{-p}]$ if $\mathrm{End}_{\mathbb{F}_p}(E) = \mathbb{Z}[\pi_p]$. Note that in either case, the orientation is naturally induced by the Frobenius $\iota : \sqrt{-p} \longmapsto \pi_p$ and needs not to be specified. This greatly simplifies the class group action computation.

The *supersingular isogeny graph over* $\mathbb{F}_p$ made of supersingular elliptic curves and their isogenies all defined over $\mathbb{F}_p$ has a *volcano* structure with two levels. On the *crater* or the *surface* of the volcano stand the curves oriented by the maximal order $\mathbb{Z}[(1+\sqrt{-p})/2]$. On the *floor* of the volcano stand the order $\mathbb{Z}[\sqrt{-p}]$ of conductor 2. The ideal class group action acts on each level of the volcano. Nonetheless, 2-isogenies can change levels. When $p \equiv 3 \mod 8$, 2 does not split in $\mathbb{Q}(\sqrt{-p})$ so there is no prime ideal of norm 2 and in particular, all 2-isogenies starting from the surface are descending isogenies. When $p \equiv 7 \mod 8$, 2 splits $\mathbb{Q}(\sqrt{-p})$ and the action of prime ideals lying above 2 yield horizontal isogenies. As a consequence, there are two horizontal and one descending 2-isogenies starting from each curve on the surface [CD20] (see Figures 1.1 and 1.2).

In CSIDH, curves on the floor *i.e.* oriented by $\mathbb{Z}[\sqrt{-p}]$ have been considered. The group action induced by $\mathrm{Cl}(\mathbb{Z}[\sqrt{-p}])$ on $\mathrm{Ell}_p(\mathbb{Z}[\sqrt{-p}])$ is a *restricted cryptographic group action* also called *restricted effective group action* [ADFMP20] since:

- It is free and transitive;

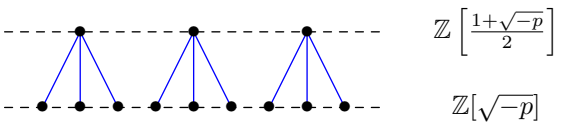

Figure 1.1: Volcano structure of the supersingular isogeny graph over $\mathbb{F}_p$ when $p \equiv 3 \mod 8$.

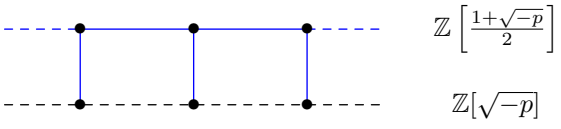

Figure 1.2: Volcano structure of the supersingular isogeny graph over $\mathbb{F}_p$ when $p \equiv 7 \mod 8$.

- It can be efficiently computed but only on ideals of smooth norms (which generate the ideal class group)[1];
- The *vectorisation problem* which requires to compute $[\mathfrak{a}] \in \mathrm{Cl}(\mathbb{Z}[\sqrt{-p}])$ from $E \in \mathrm{Ell}_p(\mathbb{Z}[\sqrt{-p}])$ and $\mathfrak{a} \cdot E$, is a hard problem even for a quantum computer.

This restricted cryptographic group action have been used to construct a Diffie-Hellman key exchange quite similar to the pre-quantum discrete logarithm based El Gamal as follows. A public curve $E_0 \in \mathrm{Ell}_p(\mathbb{Z}[\sqrt{-p}])$ is fixed. Alice chooses a secret ideal $\mathfrak{a} \subseteq \mathfrak{O}$ (of smooth norm) and sends $\mathfrak{a} \cdot E_0$ to Bob. Similarly, Bob chooses a secret ideal $\mathfrak{b} \subseteq \mathfrak{O}$ (of smooth norm) and sends $\mathfrak{b} \cdot E_0$ to Alice. Then Alice computes $\mathfrak{a} \cdot (\mathfrak{b} \cdot E_0) = (\mathfrak{ab}) \cdot E_0$ and Bob computes $\mathfrak{b} \cdot (\mathfrak{a} \cdot E_0) = (\mathfrak{ab}) \cdot E_0$. By commutativity of $\mathrm{Cl}(\mathfrak{O})$, the final key $(\mathfrak{ab}) \cdot E_0$ is shared between the two parties. These convenient cryptographic group action properties have also been used to build various cryptographic schemes including the digital signature scheme CSI-FiSh [BKV19], threshold [DFM20] and ring [BKP20] signatures, oblivious transfer [SGOPS20; LGSG21], oblivious pseudo-random functions [BKW20] and hash proof systems [ADFMP20]. Other cryptographic group actions based on other orientations than CSIDH and its analogue on the surface [CD20] also have been proposed [CK20; FFK+23; CLP24]. In all of these orientation based primitives, the restriction to smooth norm ideals can limit cryptographic applications in practice. A solution to overcome this limitation is a contribution of this thesis (see Section 2.5).

Another limitation to these schemes based on cryptographic group action is their vulnerability to Kuperberg's quantum subexponential attack (in $\log(p)$) on the verctorisation problem [Kup05]. There is currently no consensus on the efficiency of this attack on CSIDH (or similar schemes), hence on the security parameters that guarantee enough security [CSCDJRH21; BS20; Pei20]. According to the most conservative estimates, a base prime p of 4000 bits would be necessary to ensure a NIST-I level of security for CSIDH.

1.4 Polarised abelian varieties

In this section, we give the necessary algebraic geometry background to work with abelian varieties, following Milne's approach [Mil86]. This will be needed especially in Chapter 5 dedicated to theta functions. We do not require the reader to be a professional algebraic geometer but some familiarity with schemes and line bundles is recommended. We refer to [Har77, Chapter II, Sections 1-7] for an introduction to scheme theory. It will become clear in Chapter 5 that we only use some algebraic geometry to focus mainly on arithmetic aspects that are relevant for isogeny computations.

1.4.1 Abelian varieties

Informally, abelian varieties are a higher dimensional generalisation of elliptic curves. In particular, they have a group structure. However, we need some scheme theoretic notions to define them properly.

[1] This restriction of effectiveness to some group elements is why the cryptographic group action is *restricted*.

Definition 1.4.1 (Abelian variety)**.** Let k be a field. An *abelian variety* A over k is a *complete and geometrically integral group variety over* k.

We shall now give explicit definitions of those terms:

Variety: A k-variety V also called a variety defined over k is *k-scheme* that is *separated* and of *finite type*. This means that there is a structural morphism $\pi : V \longrightarrow \mathrm{Spec}(k)$ (k-scheme), that the diagonal map $\Delta : V \longrightarrow V \times_k V$ is a closed immersion (separateness), that V is quasi-compact as a topological space and for every open affine subset $U \subseteq V$, $\Gamma(U, \mathcal{O}_V)$ is a finite type k-algebra (finite type).

Complete: A k-scheme X is complete if for all k-scheme Y, the right projection $q : X \times_k Y \longrightarrow Y$ is closed.

Geometrically integral: A k-scheme X is *geometrically integral* if for every field extension k'/k, $X_{k'}$ is *integral*. We say that a scheme X is integral if for all open subset $U \subseteq X$, $\Gamma(U, \mathcal{O}_X)$ is an integral domain. Equivalently, X is integral if X is *reduced* and *irreducible* (as a topological space). A scheme X is reduced if for all $x \in X$, $\mathcal{O}_{X,x}$ is *reduced* i.e. does not contain non-zero nilpotent element. If X is integral, it is in particular connected (as a topological space).

Group variety: A *group scheme* G over k is a k-scheme together with two morphisms of k-schemes $m : G \times_k G \longrightarrow G$ and $\mathrm{inv} : G \longrightarrow G$ and an element $e \in G(k)$ such that m and inv induce a group structure on $G(\overline{k})$ with neutral element e. A *group variety* over k is a group scheme which is a k-variety.

This formal definition is almost never used in practice (the group structure excepted) and we usually rely on the following properties.

Theorem 1.4.2. *Let A be an abelian variety over a field k. Then:*

(i) The group law on $A(\overline{k})$ is commutative.

(ii) A is smooth.

(iii) A is a projective variety. In other words, there exists a closed immersion $A \hookrightarrow \mathbb{P}^n_k$.

Proof. **(i)** The proof follows [Mil86, Corollary 2.4] and is a consequence of the rigidity theorem [Mil86, Theorem 2.1]. We prove that the inversion map $\mathrm{inv} : A \longrightarrow A$ induces a group homomorphism on $A(\overline{k})$. Using the notations of Definition 1.4.1, consider $\varphi := m(\mathrm{inv} \circ m, \mathrm{inv} \circ m \circ (\mathrm{inv} \times \mathrm{inv}))$, which is a map $A \times A \longrightarrow A$ acting on points as $\varphi(a, a') = (a \cdot a')^{-1} \cdot (a^{-1} \cdot a'^{-1})^{-1}$. We have $\varphi(a, e) = e = \varphi(e, a)$ for all $a \in A(\overline{k})$. Since A is complete and $\varphi(A \times \{e\}) = \varphi(\{e\} \times A) = \{e\}$, the rigidity lemma below applies to φ so $\varphi = e$, which proves (i). As a consequence, we shall denote the group law on $A(\overline{k})$ additively and denote 0 or 0_A instead of e.

Lemma 1.4.3 (Rigidity Lemma [Mil86, Theorem 2.1])**.** *Let $f : V \times W \longrightarrow U$ be a morphism of k-varieties. Assume that V is complete and that there exist $u_0 \in U(k), v_0 \in V(k), w_0 \in W(k)$ such that $f(\{v_0\} \times W) = f(V \times \{w_0\}) = \{u_0\}$. Then $f(V \times W) = \{u_0\}$.*

(ii) The proof follows [Mil86, § 1]. By [GW10, Theorem 6.28], A is smooth if and only if $A_{\overline{k}}$ is smooth. Hence, we assume that k is algebraically closed ($A = A_{\overline{k}}$). Smooth points $x \in A$ admit an open affine neighbourhood $U \subseteq A$ such that there exists an open immersion

$$j : U \hookrightarrow \mathrm{Spec}\, k[T_1, \cdots, T_n]/(f_1, \cdots, f_{n-d})$$

with a Jacobian at x:

$$\left(\frac{\partial f_i}{\partial T_j}(j(x))\right)_{\substack{1 \leq i \leq n-d \\ 1 \leq j \leq n}} \in M_{n-d,n}(\kappa(x))$$

of rank $n - d$ [GW10, Definition 6.14]. This is an open condition. Since A is integral, it is reduced so the smooth locus A_{sm} of A is dense by [GW10, Theorem 6.19]. In particular, $A_{sm} \neq \emptyset$.

For all $x \in A$, consider the *translation* $t_x : A \longrightarrow A, a \longmapsto a + x \in A$ (well defined because the group law extends to $A(\kappa(x))$). For all $x \in A$, t_x is an isomorphism so the points of $t_x(A_{sm})$ are also

smooth i.e. $t_x(A_{sm}) \subseteq A_{sm}$. In particular, if $x_0 \in A_{sm}$, then for all $x \in A$, $x + x_0 \in A_{sm}$ so $A = A_{sm}$ i.e. A is smooth.

(iii) We refer to [Mil86, Theorem 7.1] for a complete proof. We give an overview of the proof in Theorem 1.4.21. □

The previous result confirms that abelian varieties generalise elliptic curves. Like elliptic curves, their group law is commutative and they are projective. This means that they can be described very concretely as a set of projective points annihilated by homogeneous polynomials. These polynomials can be explicitly computed with Riemann relations [Rob10, Theorem 4.7.1] due to Mumford [Mum66, pp. 336-349]. However, they are of limited use for our algorithmic applications.

1.4.2 Line bundles and divisors

Line bundles on abelian varieties are very useful because they provide a richer structure. The notion of polarised abelian variety follows naturally from the notion of line bundle. Line bundles are also used to construct maps from abelian varieties to projective spaces and to prove that abelian varieties are projective. As we shall see in Chapter 5, such maps to projective spaces define systems of coordinates which can be used to perform arithmetic operations on abelian varieties.

Line bundles: definition and basic properties

Let $(X, \mathcal{O}_X)$ be a locally ringed space. A *line bundle* also called *invertible sheaf* $\mathcal{L}$ on X is a locally free $\mathcal{O}_X$-module of rank 1. That is to say that X can be covered by open subsets $U \subseteq X$ such that $\mathcal{L}_{|U}$ is an $\mathcal{O}_{X|U}$-module of rank 1. In particular, $\mathcal{O}_X$ is a line bundle called the *trivial* line bundle. Every line bundle isomorphic to $\mathcal{O}_X$ is also called trivial.

Proposition 1.4.4. *Isomorphism classes of line bundles on X form an abelian group for the the tensor product over $\mathcal{O}_X$ with the trivial line bundle $\mathcal{O}_X$ as neutral element. This group is called the* Picard group *of X and denoted by* $\mathrm{Pic}(X)$.

Proof. We refer to [GW10, § 7.5] for a proof of this result. In particular, if $\mathcal{L}$ and $\mathcal{L}'$ are line bundles on X, $\mathcal{L} \otimes_{\mathcal{O}_X} \mathcal{L}'$ is also a line bundle on X. Besides, if $\mathcal{L}'' \simeq \mathcal{L}'$ is another line bundle on X isomorphic to $\mathcal{L}'$ then $\mathcal{L} \otimes_{\mathcal{O}_X} \mathcal{L}' \simeq \mathcal{L} \otimes_{\mathcal{O}_X} \mathcal{L}''$. Therefore, if $\mathcal{L}' \simeq \mathcal{O}_X$, then $\mathcal{L} \otimes_{\mathcal{O}_X} \mathcal{L}' \simeq \mathcal{L} \otimes_{\mathcal{O}_X} \mathcal{O}_X = \mathcal{L}$. Moreover, if $\mathcal{L}$ is a line bundle on X, then $\mathcal{L}^\vee := \mathrm{Hom}_{\mathcal{O}_X}(\mathcal{L}, \mathcal{O}_X)$ is a line bundle on X and $\mathcal{L}^\vee \otimes_{\mathcal{O}_X} \mathcal{L} \simeq \mathcal{O}_X$ [GW10, p. 7.5.8]. Hence, the isomorphism class $[\mathcal{L}]$ admits an inverse. □

In the following, we shall denote the group law multiplicatively. In particular, we shall denote $\mathcal{L}^{-1} := \mathcal{L}^\vee$ and $\mathcal{L}^n := \mathcal{L}^{\otimes n}$ for all $n \in \mathbb{N}^*$.

Correspondence with divisors

When X is a smooth and irreducible k-variety, line bundles can be seen as sheaves of functions on X or as divisors. This provides a "more concrete" way to look at them that is very useful in several proofs.

Consider a k-variety X. A *prime divisor* on X is a closed subvariety of codimension 1 in X. A *(Weil) divisor* D on X is a formal sum

$$D = \sum_{Z \in P(X)} n_Z [Z],$$

where $P(X)$ is the set of prime divisors on X and the n_Z are integers that are zero except on a finite number of prime divisors Z. The set of divisors on X denoted by $\mathrm{Div}(X)$ form an abelian group.

Example 1.4.5. On a plane or projective curve (e.g. an elliptic curve), prime divisors are points and divisors are formal finite sums of points.

Assume that X is a smooth and irreducible k-variety. X being integral, there exists a unique *generic point* $\eta \in X$ i.e. a point dense in X. This point corresponds to the zero-ideal in every open affine subvariety of X. Besides $\mathcal{O}_{X,\eta}$ is an integral domain so we can define its field of fractions

$k(X) := \mathrm{Frac}(\mathcal{O}_{X,\eta})$ called the *function field* of X. The generic point η being dense in X, we have injective maps $\Gamma(U, \mathcal{O}_X) \hookrightarrow \Gamma(V, \mathcal{O}_X) \hookrightarrow k(X)$ for all open subsets $V \subseteq U \subseteq X$ [GW10, Proposition 3.29] so sections of X can be viewed as functions of X.

Example 1.4.6. If $E := \mathrm{Proj}(k[x,y,z]/(y^2z - x^3 - Axz^2 - Bz^3))$ with $\Delta(A,B) := 4A^3 + 27B^2 \neq 0$ is an elliptic curve, then the function field of E is:

$$k(E) := \mathrm{Frac}(k[x,y,z]/(y^2z - x^3 - Axz^2 - Bz^3)).$$

Every prime divisor $Z \in P(X)$ of the k-variety X admits a generic point $\xi \in Z$ and $\mathcal{O}_{X,\xi}$ has Krull dimension 1 since Z has codimension 1 in X. We denote $\mathcal{O}_Z := \mathcal{O}_{X,\xi}$. The variety X being smooth, $\mathcal{O}_Z$ is a discrete valuation ring[2] for all $Z \in P(X)$ [GW10, Lemma 6.26 and Proposition 11.37]. We denote by ord_Z the valuation of $\mathcal{O}_Z$. Since $\mathrm{Frac}(\mathcal{O}_Z) = k(X)$ [GW10, Proposition 3.29], ord_Z defines a valuation on $k(X)$. Besides X is noetherian (as any k-variety), so for all $f \in k(X)^*$, $\mathrm{ord}_Z(f) \neq 0$ on a finite subset of prime divisors $Z \in P(X)$. Hence, we can define:

$$\begin{array}{rcl} \mathrm{div} : k(X)^* & \longrightarrow & \mathrm{Div}(X) \\ f & \longmapsto & \mathrm{div}(f) := \sum_{Z \in P(X)} \mathrm{ord}_Z(f)[Z] \end{array}$$

This is a group homomorphism. A divisor lying in the image of div is said *principal*. We denote $\mathrm{Princ}(X) := \mathrm{div}(k(X)^*)$.

If $D := \sum_{Z \in P(X)} n_Z[Z]$ is a divisor on X, we denote $D \geq 0$ if $n_Z \geq 0$ for all $Z \in P(X)$. For all open subset $U \subseteq X$, we also define the restriction:

$$D_{|U} := \sum_{\substack{Z \in P(X) \\ Z \cap U \neq \emptyset}} n_Z[Z \cap U].$$

We define the sheaf of $\mathcal{O}_X$-modules $\mathcal{L}(D)$ given by:

$$\Gamma(U, \mathcal{L}(D)) := \{f \in k(X)^* \mid \mathrm{div}(f)_{|U} + D_{|U} \geq 0\},$$

for all open subset $U \subseteq X$.

Theorem 1.4.7.

(i) For all $D \in \mathrm{Div}(X)$, $\mathcal{L}(D)$ is a line bundle on X.

(ii) The map $D \longmapsto \mathcal{L}(D)$ induces a group isomorphism $\mathrm{Div}(X)/\mathrm{Princ}(X) \xrightarrow{\sim} \mathrm{Pic}(X)$.

Proof. We refer to [GW10, Theorem 11.38]. The proof relies on the fact that divisors are locally principal on normal varieties, i.e. that for all $D \in \mathrm{Div}(X)$, there exists an open covering $(U_i)_{i \in I}$ of X such that for all $i \in I$, there exists $f_i \in k(X)$ such that $D_{|U_i} = \mathrm{div}(f_i)_{|U_i}$. In other words, Weil divisors are in bijection with Cartier divisors (as defined in [GW10, Definition 11.24]). □

Theorem 1.4.7 ensures that line bundles are in correspondence with divisors. This means that every result (e.g. in Chapter 5) using the language of line bundles could be translated in the language of divisors. We also obtained an explicit description of line bundles as sheaves of functions (up to isomorphism).

Ample line bundles and projective maps

On $\mathbb{P}^n_k = \mathrm{Proj}(k[T_0, \cdots, T_n])$, we consider the *twisted sheaf of Serre* defined by $\mathcal{O}(1) := \widetilde{M}$, where $M := \sum_{i=0}^n T_i k[T_0, \cdots, T_n]$. This is a sheaf of $\mathcal{O}_{\mathbb{P}^n_k}$-modules of rank 1 i.e. a line bundle. This sheaf is generated by global sections $T_0, \cdots, T_n \in \Gamma(\mathbb{P}^n_k, \mathcal{O}(1)) = M$ in the sense of the following definition.

Definition 1.4.8. A line bundle $\mathcal{L}$ on a k-scheme X is generated by global sections $s_0, \cdots, s_n \in \Gamma(X, \mathcal{L})$ if for all $x \in X$, the stalk $\mathcal{L}_x$ is generated by the images of $s_{0,x}, \cdots, s_{n,x}$ as an $\mathcal{O}_{X,x}$-module.

[2] The smoothness of X is a stronger assumption than necessary. By [GW10, Proposition 11.37], it suffices to assume that X is *normal i.e.* that $\mathcal{O}_{X,x}$ is normal for all $x \in X$ *i.e.* that the localisation of $\mathcal{O}_{X,x}$ at every prime is integrally closed.

By pulling-back $\mathcal{O}(1)$ on a scheme X, we obtain a line bundle generated by global sections which determines a projective map $X \longrightarrow \mathbb{P}^n_k$. This defines a correspondence between line bundles generated by global sections and projective maps.

Theorem 1.4.9. *[Har77, § II.7.1] Let X be a k-scheme.*

(i) If $\varphi : X \longrightarrow \mathbb{P}^n_k$ is a morphism of k-scheme, then $\varphi^*\mathcal{O}(1)$ is line bundle on X which is generated by global sections $s_i = \varphi^*(T_i)$ for $i \in [\![0\ ;\ n]\!]$.

(ii) Conversely, if $\mathcal{L}$ is a line bundle on X generated by global sections $s_0, \cdots, s_n \in \Gamma(X, \mathcal{L})$, then there is a unique morphism of k-schemes $\varphi : X \longrightarrow \mathbb{P}^n_k$ such that $\mathcal{L} \simeq \varphi^*\mathcal{O}(1)$ and $s_i = \varphi^*(T_i)$ for all $i \in [\![0\ ;\ n]\!]$.

Definition 1.4.10. Let X be a k-scheme of finite type and $\mathcal{L}$ be a line bundle on X.

(i) We say that $\mathcal{L}$ is *very ample* when there exists a close immersion $\varphi : X \hookrightarrow \mathbb{P}^n_k$ such that $\mathcal{L} \simeq \varphi^*\mathcal{O}(1)$.

(ii) We say that $\mathcal{L}$ is *ample* when there exists $n \in \mathbb{N}^*$ such that $\mathcal{L}^n$ is very ample.

(iii) Similarly, we say that a divisor D on X is (very) ample when $\mathcal{L}(D)$ is (very) ample.

Remark 1.4.11.

1. The definition of ample line bundles we provide is different from the one provided by [Har77, p. 153] but is equivalent over k-schemes of finite type by [Har77, Theorem 7.6].

2. If $\mathcal{L}$ is an ample line bundle on a complete k-variety X, then there exists $n \in \mathbb{N}^*$ such that $\mathcal{L}^n$ is generated by global sections so $\mathcal{L}$ has global sections. Hence, $\mathcal{L}^{-1}$ may not be ample, except if it is trivial $\mathcal{L} \simeq \mathcal{O}_X$. Indeed, $\mathcal{L}$ and $\mathcal{L}^{-1}$ have global sections if and only if $\mathcal{L}$ is trivial by [Mil86, Lemma 5.4].

Proposition 1.4.12. *If $\mathcal{L}$ and $\mathcal{M}$ are ample line bundles on a k-scheme of finite type X, then $\mathcal{L} \otimes \mathcal{M}$ is ample.*

Proof. See [Har77, Exercise II.7.5]. □

Lemma 1.4.13. *[The24, Tag 01PU] If X is a k-scheme of finite type and $Y \subseteq X$ is a closed subscheme then for every ample line bundle $\mathcal{L}$ on X, $\mathcal{L}_{|Y}$ is ample on Y.*

Line bundles on abelian varieties

Let A be an abelian variety over k. Then A is a smooth and irreducible k-variety so all the previous results on line bundle and divisors apply. If $f : X \longrightarrow A$ is a morphism of k-schemes and $\mathcal{L}$ is a line bundle of A, the pull back $f^*\mathcal{L}$ is a line bundle on X. The scheme theoretic definition of the pullback is provided in [Har77, § II.1] (for a general sheaf). Concretely, sections of $f^*\mathcal{L}$ can be seen as rational functions of A precomposed with f. In this paragraph, we present the properties of pullbacks of line bundles on A. The main result we shall prove is the theorem of the square (Theorem 1.4.19).

Theorem 1.4.14 (Seesaw principle)**.** *Let V be a complete k-variety and T be an integral scheme of finite type over k. Let $\mathcal{L}$ and $\mathcal{M}$ be line bundles over $V \times T$ such that $\mathcal{L}_{|V\times\{t\}} \simeq \mathcal{M}_{|V\times\{t\}}$ for all $t \in T$ and there exists $v \in V(k)$ such that $\mathcal{L}_{|\{v\}\times T} \simeq \mathcal{M}_{|\{v\}\times T}$. Then $\mathcal{L} \simeq \mathcal{M}$.*

Proof. See [Mil86, Theorem 5.1 and Corollary 5.2]. □

From the seesaw principle, one can obtain the theorem of the cube:

Theorem 1.4.15 (Theorem of the Cube [Mil86, Theorem 6.1])**.** *Let U, V, W be complete geometrically integral k-varieties, $\mathcal{L}$ a line bundle on $U \times V \times W$ and $u_0 \in U(k), v_0 \in V(k), w_0 \in W(k)$. Assume that $\mathcal{L}$ is trivial on $\{u_0\} \times V \times W$, $U \times \{v_0\} \times W$ and $U \times V \times \{w_0\}$. Then $\mathcal{L}$ is trivial on $U \times V \times W$.*

Corollary 1.4.16. *Let A be an abelian variety over k and $p_i : A^3 \longrightarrow A$ be the projection of the i-th component for all $i \in \{1,2,3\}$, $p_{ij} := p_i + p_j$ for all $i,j \in \{1,2,3\}$ and $p_{123} := p_1 + p_2 + p_3$. Let $\mathcal{L}$ be a line bundle on A. Then*

$$p_{123}^*\mathcal{L} \otimes p_{12}^*\mathcal{L}^{-1} \otimes p_{23}^*\mathcal{L}^{-1} \otimes p_{13}^*\mathcal{L}^{-1} \otimes p_1^*\mathcal{L} \otimes p_2^*\mathcal{L} \otimes p_3^*\mathcal{L}$$

is trivial.

Proof. Let $p, q : A^2 \longrightarrow A$ be the projections on the first and second component of A respectively and $m : A^2 \longrightarrow A$ be the addition map. Then, on $\{0\} \times A \times A$ identified with $A \times A$, the above line bundle becomes:

$$m^*\mathcal{L} \otimes p^*\mathcal{L}^{-1} \otimes m^*\mathcal{L}^{-1} \otimes q^*\mathcal{L}^{-1} \otimes p^*\mathcal{L} \otimes q\mathcal{L},$$

which is trivial. We prove similarly that the restrictions to $A \times \{0\} \times A$ and $A \times A \times \{0\}$ are trivial. The result follows from the theorem of the cube (Theorem 1.4.15). □

Corollary 1.4.17. *Let A be an abelian variety over k, V be a complete and geometrically integral k-variety, $f, g, h : V \longrightarrow A$ be three morphisms and $\mathcal{L}$ be a line bundle on A. Then*

$$(f+g+h)^*\mathcal{L} \otimes (f+g)^*\mathcal{L}^{-1} \otimes (g+h)^*\mathcal{L}^{-1} \otimes (f+h)^*\mathcal{L}^{-1} \otimes f^*\mathcal{L} \otimes g^*\mathcal{L} \otimes h^*\mathcal{L}$$

is trivial.

Proof. We apply the pullback of the $f \times g \times h : V^3 \longrightarrow A^3$ to the line bundle of Corollary 1.4.16. □

Corollary 1.4.18. *Let A be an abelian variety and $[n] : A \longrightarrow A$ be the the multiplication by $n \in \mathbb{Z}$ map. Then:*

$$[n]^*\mathcal{L} \simeq \mathcal{L}^{n(n+1)/2} \otimes [-1]^*\mathcal{L}^{n(n-1)/2}.$$

In particular, if $\mathcal{L}$ is symmetric ($[-1]^*\mathcal{L} \simeq \mathcal{L}$), then $[n]^*\mathcal{L} \simeq \mathcal{L}^{n^2}$ and if $\mathcal{L}$ is antisymmetric ($[-1]^*\mathcal{L} \simeq \mathcal{L}^{-1}$), then $[n]^*\mathcal{L} \simeq \mathcal{L}^n$.

Proof. We proceed by induction on $n \in \mathbb{N}$. The case $n = 0$ is $[0]^*\mathcal{L} \simeq \mathcal{O}_A$ which is true since the pullback by a constant map is always trivial. The case $n = 1$ is $[1]^*\mathcal{L} \simeq \mathcal{L}$, which is trivially true. For $n \geq 1$, we assume that the result holds for $n-1$ and n. We apply Corollary 1.4.17 with $f = [n], g = [1] = \mathrm{id}_A$ and $h = [-1]$ and we obtain that:

$$[n]^*\mathcal{L} \otimes [n+1]^*\mathcal{L}^{-1} \otimes [n-1]^*\mathcal{L}^{-1} \otimes [n]^*\mathcal{L} \otimes \mathcal{L} \otimes [-1]^*\mathcal{L} \simeq \mathcal{O}_A,$$

with $[n]^*\mathcal{L} \simeq \mathcal{L}^{n(n+1)/2} \otimes [-1]^*\mathcal{L}^{n(n-1)/2}$ and $[n-1]^*\mathcal{L} \simeq \mathcal{L}^{n(n-1)/2} \otimes [-1]^*\mathcal{L}^{(n-1)(n-2)/2}$, so that

$$\begin{aligned}
[n+1]^*\mathcal{L} &\simeq [n]^*\mathcal{L}^2 \otimes [n-1]^*\mathcal{L}^{-1} \otimes \mathcal{L} \otimes [-1]^*\mathcal{L} \\
&\simeq \mathcal{L}^{n(n+1)} \otimes [-1]^*\mathcal{L}^{n(n-1)} \otimes \mathcal{L}^{-n(n-1)/2} \otimes [-1]^*\mathcal{L}^{-(n-1)(n-2)/2} \otimes \mathcal{L} \otimes [-1]^*\mathcal{L} \\
&\simeq \mathcal{L}^{(n+1)(n+2)/2} \otimes [-1]^*\mathcal{L}^{n(n+1)/2}
\end{aligned}$$

This proves the result for $n+1$. Now, if $n \in \mathbb{Z} \setminus \mathbb{N}$, we apply $[-1]^*$ to the formula obtained for $-n$. This completes the proof. □

Theorem 1.4.19 (Theorem of the Square)**.** *Let A be an abelian variety over k and $\mathcal{L}$ be a line bundle on A. Then, for all $a, b \in A(\overline{k})$, we have:*

$$t_{a+b}^*\mathcal{L} \otimes \mathcal{L} \simeq t_a^*\mathcal{L} \otimes t_b^*\mathcal{L}.$$

Proof. We apply Corollary 1.4.17 with $f = \mathrm{id}_A$, g the constant map equal to a and h the constant map equal to b. □

Remark 1.4.20. If $D := \sum_{Z \in P(A)} n_X[Z]$ is a divisor on A and $t_a : A \longrightarrow A$ is the translation map $x \longmapsto x + a$ by $a \in A(\overline{k})$, then $t_a^*\mathcal{L}(D) = \mathcal{L}(D - a)$, where $D - a := \sum_{Z \in P(A)} n_X[Z - a]$. Hence, the theorem of the square ensures that $(D - a - b) + D \sim (D - a) + (D - b)$ for all $a, b \in A(\overline{k})$.

Why abelian varieties are projective

Theorem 1.4.21. *[Mil86, Theorem 7.1] Abelian varieties are projective. In other words, if A is an abelian variety over k, there exists a very ample line bundle on A.*

Proof. The proof from Milne [Mil86, Theorem 7.1] is made of two steps. We only give a very brief overview and refer to [Mil86] for details.

First, we assume that k is algebraically closed and find a line bundle $\mathcal{L}$ on A generated by global sections $s_0, \cdots, s_n \in \Gamma(A, \mathcal{L})$ forming a k-vector space $V = \mathrm{Span}_k(s_0, \cdots, s_n)$ that:

- *Separates points*: for all distinct closed points $x, y \in A(k)$, there exists $s \in V$ such that $s_x \in \mathfrak{m}_{A,x}\mathcal{L}_x$ and $s_y \notin \mathfrak{m}_{A,y}\mathcal{L}_y$.
- *Separates tangent vectors*: for all closed point $x \in A(k)$, $\{s \in V \mid s_x \in \mathfrak{m}_{A,x}\mathcal{L}_x\}$ spans $\mathfrak{m}_{A,x}\mathcal{L}_x/\mathfrak{m}^2_{A,x}\mathcal{L}_x$ as a k-vector space.

Then [Har77, Proposition II.7.3] ensures that $\mathcal{L}$ is very ample. To find such a very ample line bundle, we find a finite set of prime divisors $Z_1, \cdots, Z_n \subset A$ (subvarieties of codimension 1) such that $\bigcap_{i=1}^n Z_i = \{0\}$ and for all tangent vector $t \in T_0(A) = (\mathfrak{m}_{A,0}/\mathfrak{m}^2_{A,0})^*$, there exists $i \in [\![1\ ;\ n]\!]$ such that $t \notin T_0(Z_i) = (\mathfrak{m}_{Z_i,0}/\mathfrak{m}^2_{Z_i,0})^*$ (viewed naturally as a subspace of $T_0(A)$). Then, we consider $D := \sum_{i=1}^n [Z_i]$ and it can be proved that $\mathcal{L}(3D)$ is generated by global sections that separate points and tangent vectors i.e. is very ample. The main ingredient of the proof is the theorem of the square Theorem 1.4.19 (which translates local properties at 0 everywhere).

In the second step, we do not assume that k is algebraically closed. By the first step we obtain an ample divisor D on $A_{\overline{k}}$. Then D must be defined over a finite extension k'/k and we may consider $D' = \sum_{\sigma \in G} \sigma \cdot D$ with $G := \mathrm{Aut}_k(k')$. Then D' is fixed by the action of G, so is naturally defined over k'^G and it can be proved that D' is ample as a sum of ample divisors by Proposition 1.4.12. If k is perfect, then $k'^G = k$ and the proof is complete. Otherwise, $p := \mathrm{char}(k) > 0$ and k'^G is purely inseparable i.e. there exists $m \in \mathbb{N}$ such that $k'^{p^m} = k$. It can then be proved that $p^m D'$ is defined over k and ample, again as a sum of ample divisors. □

Theorem 1.4.21 ensures that (very) ample line bundles always exist on an abelian variety A over k and its proof relies on finding a specific ample line bundle $\mathcal{L}$ on A and proving that $\mathcal{L}^3$ is very ample. This is actually a general fact for any ample line bundle on A and its proof also relies on the theorem of the square.

Theorem 1.4.22. *[Mum74, Theorem p. 163] Let A be an abelian variety over k. Then for any ample line bundle $\mathcal{L}$ on A and $n \geq 3$, $\mathcal{L}^n$ is very ample.*

1.4.3 Isogenies

Definition 1.4.23. A *homomorphism* of abelian varieties $f : A \longrightarrow B$ over k is a morphism of k-schemes that induces a group homomorphism $A(\overline{k}) \longrightarrow B(\overline{k})$.

By abuse, we call homomorphism $A \longrightarrow B$ a homomorphism $f : A_{k'} \longrightarrow B_{k'}$ defined over an algebraic extension k'/k. By abuse, we also say that such a homomorphism f is *defined over k* or *k-rational* if f is the base change $f = g_{k'}$ of a homomorphism $g : A \longrightarrow B$ over k.

It is a widely known fact that an algebraic morphism between elliptic curves which maps 0 to 0 is a homomorphism. This fact generalises to abelian varieties.

Lemma 1.4.24. *Let $f : A \longrightarrow B$ be morphism of k-schemes between abelian varieties. Then $f = t_a \circ h$ where $h : A \longrightarrow B$ is a homomorphism of abelian varieties and $a := f(0) \in B(k)$. As a consequence, if $f(0) = 0$, then f is a homomorphism of abelian varieties.*

Proof. Replacing f by $t_{-a} \circ f$ with $a := f(0)$, we may assume that $f(0) = 0$. Let $m_A : A \times A \longrightarrow A$ and $m_B : B \times B \longrightarrow B$ be the addition maps and $\varphi := f \circ m_A - m_B \circ (f \times f) : A \times A \longrightarrow B$. Then $\varphi(0, x) = \varphi(x, 0) = f(x) - f(x) - f(0) = 0$ for all $x \in A(\overline{k})$ so $\varphi(\{0\} \times A) = \varphi(A \times \{0\}) = \{0\}$ and B is complete so the rigidity lemma (Lemma 1.4.3) applies and ensures that $\varphi = 0$. This completes the proof. □

Any homomorphism (of abelian varieties over k) $f : A \longrightarrow B$ has a kernel $\ker(f)$ defined as the scheme theoretic fiber of 0 (see [GW10, Definition 4.25]). This is a closed subgroup scheme of A [Sha86, p. 35].

Definition 1.4.25. An *isogeny* $f : A \longrightarrow B$ is a homomorphism which is surjective and has finite kernel. Following Definition 1.4.23, we say that such an isogeny f is defined over k when it is defined over k as a homomorphism.

Proposition 1.4.26. *[Mil86, Proposition 8.1] Let $f : A \longrightarrow B$ be a homomorphism. Then the following are equivalent:*

(i) *f is an isogeny.*

(ii) *$\dim(A) = \dim(B)$ and f is surjective.*

(iii) *$\dim(A) = \dim(B)$ and $\ker(f)$ is finite.*

Example 1.4.27. If $f : E_1 \longrightarrow E_2$ is a homomorphism between elliptic curves, then $\dim(E_1) = \dim(E_2) = 1$ so f is an isogeny if and only if it has finite kernel if and only if it is surjective. This was already a known fact.

When $f : A \longrightarrow B$ is an isogeny, its fibers are finite (from a set theoretic point of view). The associated structure sheaf morphism $f^{\#} : \mathcal{O}_B \longrightarrow f_*\mathcal{O}_A$ also satisfies a finiteness property.

Definition 1.4.28. Let $(X, \mathcal{O}_X)$ be a locally ringed space and $\mathcal{M}$ be an $\mathcal{O}_X$-module. Then $\mathcal{M}$ is *finite locally free* of rank $n \in \mathbb{N}$ if X is covered by open subsets $(U_i)_{i\in I}$ such that for all $i \in I$, $\mathcal{M}_{|U_i} \simeq \mathcal{O}^n_{X\,|U_i}$.

A morphism $f : X \longrightarrow Y$ of schemes is finite locally free (of rank n) if $f_*\mathcal{O}_X$ is a finite locally free $\mathcal{O}_Y$-module (of rank n) via the structure sheaf morphism $f^{\#} : \mathcal{O}_Y \longrightarrow f_*\mathcal{O}_X$ and if f is affine (preimages of affine subsets are affine subsets).

We introduce the notions of degree and separability of isogenies between abelian varieties, generalising these notions introduced for elliptic curve isogenies.

Proposition 1.4.29. *[Mil86, Proposition 8.1] Let $f : A \longrightarrow B$ be a homomorphism of abelian varieties. Then f is an isogeny if and only if it is finite locally free.*

Definition 1.4.30. The *degree* of a morphism $f : A \longrightarrow B$, denoted by $\deg(f)$, is the rank of $f_*\mathcal{O}_A$ as an $\mathcal{O}_B$-module.

It is easy to see that the degree is a multiplicative map. If $f : A \longrightarrow B$ and $g : B \longrightarrow C$ are isogenies, then $\deg(g \circ f) = \deg(f)\deg(g)$. It can also be proved that the degree is stable under base change (*i.e.* field extension).

As in differential geometry, we can define the tangent space of a k-scheme X at a point $x \in X$, as $T_x(X) := (\mathfrak{m}_{X,x}/\mathfrak{m}^2_{X,x})^*$ [GW10, § 6.2], which is a $\kappa(x)$-vector space. When X is smooth at $x \in X$, we have $\dim_{\kappa(x)} T_x(X) = \dim(X)$ and the converse also holds [GW10, Corollary 6.29]. If $f : X \longrightarrow Y$ is a morphism of k-schemes, then for all $x \in X$, $\kappa(x)$ is a field extension of $\kappa(f(x))$ induced by the local map $f^{\#} : \mathcal{O}_{Y,f(x)} \longrightarrow \mathcal{O}_{X,x}$ and we can define a differential map:

$$df_x : T_x(X) \longrightarrow T_{f(x)}(Y) \otimes_{\kappa(f(x))} \kappa(x)$$

as in [GW10, Remark 6.3]. When $\kappa(x) = \kappa(f(x))$, df_x is a map $T_x(X) \longrightarrow T_{f(x)}(Y)$. *Separable* morphisms are the analogue of local diffeomorphisms in differential geometry.

Definition 1.4.31. Let $f : A \longrightarrow B$ be an isogeny. We say that f is *separable* or *étale*[3] when the differential map df_x is an isomorphism of $\kappa(x)$-vector spaces for all $x \in A$.

The above definition generalises separable isogenies between elliptic curves.

[3] In general, a morphism of k-schemes $f : X \longrightarrow Y$ is étale when it is flat and separable/unramified. The invertibility of df_x means that the morphism is separable/unramified. If f is an isogeny, it is automatically flat because it is finite locally free [GW10, Proposition 12.19].

Proposition 1.4.32. *Let $f : A \longrightarrow B$ be an isogeny. Then $\#\ker(f) | \deg(f)$ and we have $\#\ker(f) = \deg(f)$ when f is separable.*

Proof. Let $K := \ker(f)$. The degree being stable under base change, we may assume that k is algebraically closed. Then, since f is finite locally free, we have by [GW10, Proposition 12.21]:

$$\deg(f) = \sum_{x \in K} e_x f_x,$$

where for all $x \in K$, e_x is the *ramification* of x, defined as the length of $\mathcal{O}_{K,x}$ as an $\mathcal{O}_{K,x}$-module ($K = \ker(f)$ being seen as the scheme theoretic fiber of 0) and $f_x := [\kappa(x) : k]$ is the *inertia* index of x. In particular, for all $x \in K$, $\kappa(x)/k$ is finite so $\kappa(x) = k$ since k is algebraically closed. Besides, since the translation map is an isomorphism, we obtain that $e_x = e_0$ for all $x \in K$, so that $\deg(f) = e_0 \# K$ and $\#K | \deg(f)$. When f is separable, it is in particular unramified by [The24, Tag 0B2G], so $e_0 = 1$ and $\deg(f) = \#K$. □

Proposition 1.4.33. *[Mil86, Proposition 8.2] Let A be an abelian variety over k of dimension $g > 0$ and $[n] : A \longrightarrow A$ be the multiplication by $n \in \mathbb{Z}^*$. Then:*

(i) *$[n]$ is an isogeny.*

(ii) *$\deg([n]) = n^{2g}$.*

(iii) *$[n]$ is separable if and only if $\mathrm{char}(k) \nmid n$.*

Proof. Except for (i), we only provide a sketch of the proof here. We refer to [Mil86, Proposition 8.2] for a precise proof.

(i) We prove that $\ker([n])$ is finite and conclude by Proposition 1.4.26. By Theorem 1.4.21, there exists a (very) ample line bundle $\mathcal{L}$ on A. Since $[-1]$ is an isomorphism, $[-1]^*\mathcal{L}$ is also ample and $\mathcal{L}' := \mathcal{L} \otimes [-1]^*\mathcal{L}$ is ample by Proposition 1.4.12. By construction, $\mathcal{L}'$ is symmetric so $[n]^*\mathcal{L}' \simeq \mathcal{L}'^{n^2}$ by Corollary 1.4.18. It follows that $\mathcal{M} := \mathcal{L}'^{n^2}$ is trivial on $\ker([n])$. Besides, $\mathcal{M}$ is ample by Corollary 1.4.18 and it is also ample on $\ker([n])$ as the restriction of an ample line bundle on a closed subscheme (Lemma 1.4.13). Since $\ker([n])$ admits a trivial ample line bundle, it must be of dimension 0 so finite. Indeed, if it was not of dimension 0, the intersection number $(\mathcal{M}^{\dim(\ker([n])} \cdot \ker([n]))$ would be positive by multilinearity of intersection numbers [The24, Tag 0BEV]. But $(\mathcal{M}^{\dim(\ker([n])} \cdot \ker([n])) = 0$ since $\mathcal{M}$ is trivial on $\ker([n])$. Contradiction. For a formal and abstract introduction to intersection numbers, we refer to [The24, Tag 0BEP] or alternatively [Sha13, § IV.1] for a more elementary one when k is algebraically closed.

(ii) It can be proved that $([n]^*\mathcal{L}'^g \cdot A) = \deg([n])(\mathcal{L}'^g \cdot A)$ (see [Mil86, Lemma 8.3] or [The24, Tag 0BET]) and that $([n]^*\mathcal{L}'^g \cdot A) = ((\mathcal{L}'^{n^2})^g \cdot A) = n^{2g}(\mathcal{L}'^g \cdot A)$ since $[n]^*\mathcal{L}' \simeq \mathcal{L}'^{n^2}$ and by [The24, Tag 0BER]. We also know that $([n]^*\mathcal{L}'^g \cdot A) > 0$ by [The24, Tag 0BEV] because A has positive dimension $g > 0$. Then, $\deg([n]) = n^{2g}$.

(iii) If $f, g : A \longrightarrow B$ are morphisms of abelian varieties, it can be proved that $d(f+g)_0 = df_0 + dg_0$ [Mum74, p. 42]. It follows that $d[n]_0 : T_0(A) \longrightarrow T_0(A)$ is the multiplication by n map which is an isomorphism if and only if $\mathrm{char}(k) \nmid n$. Using translation maps, we conclude that $[n]$ is separable if and only if $\mathrm{char}(k) \nmid n$. □

Corollary 1.4.34. *Let A be an abelian variety over k of dimension $g > 0$ and $n \in \mathbb{N}^*$. If $n \nmid \mathrm{char}(k)$, then $A[n](\overline{k}) \simeq (\mathbb{Z}/n\mathbb{Z})^{2g}$.*

Proof. The torsion subgroup $A[n]$ is the kernel of $[n]$ which is separable and of degree n^{2g} by Proposition 1.4.33. Then, Proposition 1.4.32 ensures that $\#A[n](\overline{k}) = n^{2g}$. We may assume $n \geq 2$ (the result being trivial for $n = 1$). By the finite abelian group structure theorem, there exists $d_1, \cdots, d_r \in \mathbb{N}^*$ such that $d_1 | \cdots | d_r$, $d_1 \geq 2$, $\prod_{i=1}^r d_i = n^{2g}$ and:

$$A[n](\overline{k}) \simeq \prod_{i=1}^{r} (\mathbb{Z}/d_i\mathbb{Z}).$$

We also have $d_r | n$ by the definition of $A[n](\overline{k})$ and $d_1 | d_r$, so $A[d_1](\overline{k}) \subseteq A[n](\overline{k})$, so that $A[d_1](\overline{k}) = A[n](\overline{k})[d_1] \simeq (\mathbb{Z}/d_1\mathbb{Z})^r$. It follows that $\#A[d_1](\overline{k}) = d_1^r$. But $\#A[d_1](\overline{k}) = d_1^{2g}$, again by Propositions 1.4.32 and 1.4.33. Hence, $r = 2g$. The equality $\prod_{i=1}^{2g} d_i = n^{2g}$ with $d_1 | \cdots | d_r | n$ then ensures that $d_1 = \cdots = d_r = n$, which completes the proof. □

Remark 1.4.35. Assume that k has characteristic $p > 0$ and that A is an abelian variety of dimension $g > 0$. Then, it can be proved that there exists an invariant $r \in [\![0\ ;\ g]\!]$ called the *p-rank* of A such that $A[p^n](\overline{k}) \simeq (\mathbb{Z}/p^n\mathbb{Z})^r$ for all $n \in \mathbb{N}$ (see [Mum74, p. 147]). As a group scheme, $A[p^n] \simeq (\mathbb{Z}/p^n\mathbb{Z})^r \times \mu_{p^n}^r \times G_n^0$, with $\mu_{p^n} := \mathrm{Spec}(\overline{k}[T]/(T^{p^n} - 1))$ and G_n^0 a local-local group scheme[4].

1.4.4 Isogenies as quotient maps

Between elliptic curves, separable isogenies can be defined as quotient maps by their kernel. In this section, we shall see that this point of view still holds for abelian varieties. We start by defining quotients by finite group schemes in a paragraph which is a bit abstract and that can be skipped at first reading. The results are a bit more general than the main statement we make in a second paragraph but become useful in formal proofs later.

Action of group schemes

Definition 1.4.36. Let G be a group scheme and X be a scheme, both defined over k. An *action* of G on X is a map $\mu : G \times_k X \longrightarrow X$ that induces a (group theoretic) action of $G(\overline{k})$ on $X(\overline{k})$. In other words, the composition

$$X \cong \mathrm{Spec}(k) \times_k X \xrightarrow{e \times \mathrm{id}_X} G \times_k X \xrightarrow{\mu} X$$

is the identity, where $e \in G(k)$ is the identity element; and the following diagram

$$\begin{array}{ccc} G \times_k G \times_k X & \xrightarrow{m \times \mathrm{id}_X} & G \times_k X \\ {\scriptstyle \mathrm{id}_G \times \mu}\downarrow & & \downarrow{\scriptstyle \mu} \\ G \times_k X & \xrightarrow[\mu]{} & X \end{array}$$

is commutative, where $m : G \times_k G \longrightarrow G$ is the multiplication map.

Definition 1.4.37. An action of a k-group scheme G on a k-scheme X, $\mu : G \times_k X \longrightarrow X$ is *free* (or *faithful*) when the map $(\mu, q) : G \times_k X \longrightarrow X \times_k X$ is a closed immersion, $q : G \times_k X \longrightarrow X$ being the projection on the right component.

Definition 1.4.38. Let $\mu : G \times_k X \longrightarrow X$ be an action of k-group scheme and $\mathcal{F}$ be a coherent sheaf on X. A *lift* of the action μ to $\mathcal{F}$ is an isomorphsim $\lambda : q^*\mathcal{F} \xrightarrow{\sim} \mu^*\mathcal{F}$ of sheaves on $G \times_k X$ (where $q : G \times_k X \longrightarrow X$ is the projection map on the second component), such that the following diagram of sheaves on $G \times_k G \times_k X$ commutes

$$\begin{array}{ccc} p_3^*\mathcal{F} & \xrightarrow{(p_2,p_3)^*\lambda} & \xi^*\mathcal{F} \\ & {\scriptstyle (m\times \mathrm{id}_X)^*\lambda}\searrow \quad \swarrow{\scriptstyle (\mathrm{id}_G \times \mu)^*\lambda} & \\ & \eta^*\mathcal{F} & \end{array}$$

where p_2 and p_3 are respectively the projection maps on the second and third components of $G \times_k G \times_k X$, $\xi := \mu \circ (p_1, p_2)$ and $\eta := \mu \circ (m \times \mathrm{id}_X)$.

[4] A local group scheme (i.e. contains a single point) whose Cartier dual (as defined in [Mum74, § 14]) is also local. See [Mum74, p. 136].

Now, we introduce two theorems which are used to prove several results in the theory of abelian varieties, including Theorem 1.4.41 (among others). Mumford [Mum74] proved them when the field k is algebraically closed but this hypothesis is actually not necessary and Mumford's proofs still hold otherwise.

Theorem 1.4.39. *[Mum74, Theorem 1.(A), p. 111] Let G be a finite k-group scheme acting on a k-scheme X such that the orbit of every point is contained in an affine open subset of X. Then there exists a k-scheme Y and a morphism of k-schemes $\pi : X \longrightarrow Y$ such that:*

(i) As a topological space, Y is the quotient X/G and π is the associated quotient map.

(ii) As a sheaf morphism, π induces an isomorphism $\mathcal{O}_Y \xrightarrow{\sim} (\pi_*\mathcal{O}_X)^G$, where $(\pi_*\mathcal{O}_X)^G$ is the subsheaf of $\pi_*\mathcal{O}_X$ made of G-invariant functions.

The couple (Y, π) is unique up to isomorphism as it satisfies the following universal property: for all G-invariant morphism of k-schemes $f : X \longrightarrow Z$, there exists a unique morphism $g : Y \longrightarrow Z$ such that $f = g \circ \pi$.

Theorem 1.4.40. *[Mum74, Theorem 1.(B), p. 111] Let G and X be as in Theorem 1.4.39. Assume furthermore that $G = \mathrm{Spec}(R)$ with $n := \dim_k(R)$. Then π is finite locally free of rank n.*

Besides, for all coherent $\mathcal{O}_Y$-module $\mathcal{F}$, $\pi^*\mathcal{F}$ has a natural G-action lifting the action of G on X and

$$\mathcal{F} \longmapsto \pi^*\mathcal{F}$$

is an equivalence of categories from coherent $\mathcal{O}_Y$-modules to coherent $\mathcal{O}_X$-modules with G-action, inducing an equivalence of categories from locally free $\mathcal{O}_Y$-modules of finite rank to locally free $\mathcal{O}_X$-modules of finite rank with G-action.

Correspondence between separable isogenies and their kernel

As desired, the following theorem generalises the well known correspondence between finite subgroups and separable elliptic curve isogenies. This is a consequence of Theorem 1.4.39 applied to subgroup schemes of abelian varieties acting by point translation. In this theorem, the point of view of finite subgroups K of an abelian variety defined over a field k seen as sets of points defined over $\overline{k}$ is equivalent to the point of view finite subgroup schemes G which are *étale i.e.* such that $\#G(\overline{k}) = \dim_k \Gamma(G, \mathcal{O}_G)$. It is in particular true for subgroups of cardinality non-divisible by $\mathrm{char}(k)$.

Theorem 1.4.41. *[Mum74, Theorem 4, p. 73] Let A be an abelian variety over k. Then, for every finite subgroup scheme $K \subset A$ defined over k, there exists an isogeny $f : A \longrightarrow B$ defined over k such that $\ker(f) = K$. This isogeny is unique up to post-composition by an isomorphism: if $g : A \longrightarrow C$ is another isogeny with kernel K, then there exists an isomorphism $\lambda : B \xrightarrow{\sim} C$ such that $g = \lambda \circ f$. Hence, we call B* the quotient of A by K *and denote $B := A/K$. When K is étale, f is separable.*

Quite often, we apply this theorem to subgroups $K \subset A_{\overline{k}}$ which are not defined over k. In that case, we obtain an isogeny $f : A_{\overline{k}} \longrightarrow B$ with kernel K. A natural question which is crucial for computational applications is whether this isogeny descends to an isogeny defined over k. The answer to this question is related to the action on K of automorphisms of $\overline{k}$ fixing k. Indeed, if A is an abelian variety over a non algebraically closed field k, then $\mathrm{Aut}_k(\overline{k})$ acts naturally on A.

Proposition 1.4.42. *[Mil08, Lemma IV.2.1] Let $f : A \longrightarrow B$ be a k-rational isogeny between abelian varieties defined over k. Then* $\ker(f)$ *is stable under the action of* $\mathrm{Aut}_k(\overline{k})$.

Conversely, if A is an abelian variety over k and if $K \subset A_{\overline{k}}$ is a finite group scheme stable under the action of $\mathrm{Aut}_k(\overline{k})$, *then the induced quotient isogeny $f : A_{\overline{k}} \longrightarrow B = A_{\overline{k}}/K$ via Theorem 1.4.41 is k-rational and the quotient $B = A/K$ is defined over k.*

Remark 1.4.43. Note that a finite subgroup $K \subset A_{\overline{k}}$ can be stable under the action of $\mathrm{Aut}_k(\overline{k})$ without being defined over k (the action by automorphisms might permute points but not fix them). Of course, when K is defined over k, K is automatically stable under the action of $\mathrm{Aut}_k(\overline{k})$.

When $k = \mathbb{F}_q$ is a finite field, the Galois group $\mathrm{Aut}_{\mathbb{F}_q}(\overline{\mathbb{F}}_q) = \mathrm{Gal}(\overline{\mathbb{F}}_q/\mathbb{F}_q)$ is generated by the Frobenius automorphism $\sigma : x \in \overline{\mathbb{F}}_q \longmapsto x^q \in \overline{\mathbb{F}}_q$. On an abelian variety $A/\mathbb{F}_q$, the action via σ

defines the (q-th) *Frobenius endomorphism* $\pi_q \in \mathrm{End}(A)$ which is an inseparable isogeny of degree q [EGM22, Proposition 5.15]. In order to apply Proposition 1.4.42 in this context, we only have to consider the stability of subgroups under the action of π_q.

Now, we give a factorisation result of separable isogenies following from Theorem 1.4.41 and its consequences.

Corollary 1.4.44.

(i) Let $f : A \longrightarrow B$ and $g : A \longrightarrow C$ be two isogenies between abelian varieties such that $\ker(g) \subseteq \ker(f)$. Then, there exists a unique isogeny $h : C \longrightarrow B$ such that $f = h \circ g$.

(ii) Let $f : A \longrightarrow B$ be an isogeny of degree d. Then there exists an isogeny $g : B \longrightarrow A$ such that $g \circ f = [d]$.

(iii) If $f : A \longrightarrow B$ be an isogeny between abelian varieties over k of degree d not divisible by $\mathrm{char}(k)$. Then f is separable.

Proof. **(i)** The inclusion $\ker(g) \subseteq \ker(f)$ ensures that f is $\ker(g)$-invariant. Therefore, the existence and uniqueness of h directly follows from Theorem 1.4.39. We simply have to justify that h is an isogeny. This is a consequence of Proposition 1.4.26. Since f and g are isogenies, $\dim(A) = \dim(B) = \dim(C)$ and f is surjective so h must be surjective, so it is an isogeny.

(ii) By Proposition 1.4.32, $\#\ker(f) | d$ so that $\ker(f) \subseteq A[d]$ (as groups). The inclusion $\ker(f) \subseteq A[d]$ also holds as group schemes (by [EGM22, Exercise 4.4]) so (i) applies and gives the desired result.

(iii) By (ii), there exists $g : B \longrightarrow A$ such that $g \circ f = [d]$. Furthermore, $\mathrm{char}(k) \nmid d$ so $[d]$ is separable by Proposition 1.4.33 and for all $x \in A$, $dg_{f(x)} \circ df_x = d[d]_x$ is invertible so df_x is invertible and f is separable. □

1.4.5 The dual abelian variety and polarisations

The dual abelian variety

Let A be an abelian variety over k. If $\mathcal{L}$ is a line bundle on A, then Theorem 1.4.19 ensures that the map:

$$\begin{array}{rccc} \varphi_{\mathcal{L}} : & A(\overline{k}) & \longrightarrow & \mathrm{Pic}(A_{\overline{k}}) \\ & x & \longmapsto & [t_x^*\mathcal{L} \otimes \mathcal{L}^{-1}] \end{array} \tag{1.3}$$

is a group homomorphism. When $\mathcal{L}$ is ample, we shall see that $\varphi_{\mathcal{L}}$ induces an isogeny. The dual abelian variety of A is constructed to be the codomain of $\varphi_{\mathcal{L}}$ that will be called a polarisation. Now, let us formalize more precise statements.

Definition 1.4.45. We define $\mathrm{Pic}^0(A) \subset \mathrm{Pic}(A)$, the subgroup of isomorphism classes of line bundles $\mathcal{M}$ of A such that $t_x^*\mathcal{M} \simeq \mathcal{M}$ for all $x \in A(\overline{k})$.

Proposition 1.4.46. *(i) $\varphi_{\mathcal{L}}$ defined in Eq. (1.3) maps to $\mathrm{Pic}^0(A_{\overline{k}})$.*

(ii) Assume that $\Gamma(A, \mathcal{L}) \neq \{0\}$. Then $K(\mathcal{L}) := \ker(\varphi_{\mathcal{L}}) = \{x \in A(\overline{k}) \mid t_x^*\mathcal{L} \simeq \mathcal{L}\}$ is finite if and only if $\mathcal{L}$ is ample.

(iii) If $\mathcal{L}$ is ample, the image of $\varphi_{\mathcal{L}}$ is $\mathrm{Pic}^0(A_{\overline{k}})$.

Proof. By the theorem of the square (Theorem 1.4.19), we have for all $x, y \in A(\overline{k})$,

$$t_y^*(t_x^*\mathcal{L} \otimes \mathcal{L}^{-1}) = t_{x+y}^*\mathcal{L} \otimes t_y^*\mathcal{L}^{-1} \simeq t_x^*\mathcal{L} \otimes t_y^*\mathcal{L} \otimes \mathcal{L}^{-1} \otimes t_y^*\mathcal{L}^{-1} = t_x^*\mathcal{L} \otimes \mathcal{L}^{-1},$$

so that $\varphi_{\mathcal{L}}(x) = [t_x^*\mathcal{L} \otimes \mathcal{L}^{-1}] \in \mathrm{Pic}^0(A_{\overline{k}})$. This proves (i). (ii) is proved in [Mum74, p. 60] and (iii) in [Mum74, Theorem 1, p. 77]. □

Theorem 1.4.47. *There exists a unique pair $(\widehat{A}, \mathcal{P})$ (up to isomorphism), where $\widehat{A}$ is an abelian variety over A and $\mathcal{P}$ is a line bundle over $A \times \widehat{A}$ satisfying the following universal property:*

1. *$\mathcal{P}_{|\{0\}\times\widehat{A}}$ is trivial.*

2. $\mathcal{P}_{|A\times\{x\}} \in \mathrm{Pic}^0(A_{\kappa(x)})$ *for all* $x \in \widehat{A}$.

3. *For every* k*-scheme* T *and line bundle* $\mathcal{Q}$ *on* $A \times T$ *such that* $\mathcal{Q}_{|\{0\}\times T}$ *is trivial and* $\mathcal{Q}_{|A\times\{t\}} \in \mathrm{Pic}^0(A_{\kappa(t)})$ *for all* $t \in T$*, there exists a unique morphism* $f : T \longrightarrow \widehat{A}$ *such that* $(id_A \times f)^*\mathcal{P} \simeq \mathcal{Q}$.

$\widehat{A}$ *is called the* dual abelian variety *and* $\mathcal{P}$ *the* Poincaré sheaf *of* A.

Proof. The uniqueness of $(\widehat{A}, \mathcal{P})$ is a consequence of the universal property (point 3). The construction of $(\widehat{A}, \mathcal{P})$ follows from [Mum74, § 13] and [Mil86, § 10]. We only give an overview in the following.

Let $\mathcal{L}$ be an ample line bundle of A and consider $\mathcal{L}^* := m^*\mathcal{L} \otimes p^*\mathcal{L}^{-1} \otimes q^*\mathcal{L}^{-1}$ where $m, p, q : A \times A \longrightarrow A$ are respectively the addition, the projection on the first component and the projection on the second component. Then $\mathcal{L}^*$ is a line bundle on $A \times A$ such that $\mathcal{L}^*_{|\{0\}\times A}$ is trivial and for all $x \in A$, $\mathcal{L}^*_{|A\times\{x\}} \simeq t_x^*\mathcal{L} \otimes \mathcal{L}^{-1} \in \mathrm{Pic}^0(A_{\kappa(x)})$ by Proposition 1.4.46.(i). Hence, if the Poincaré sheaf exists, we must have $(\mathrm{id}_A \times \varphi_{\mathcal{L}})^*\mathcal{P} = \mathcal{L}^*$.

In [Mum74, § 13], Mumford constructs $\widehat{A}$ as the quotient $A/K(\mathcal{L})$ (which is an abelian variety by Theorem 1.4.41) and $\mathcal{P}$ as the quotient of $\mathcal{L}^*$ by the action of $\{0\} \times K(\mathcal{L})$ lifting the action by translation of $A \times A$ (which exists by Theorem 1.4.40). Note that here, $K(\mathcal{L})$ is not seen as the group $K(\mathcal{L}) = \{x \in A(\overline{k}) \mid t_x^*\mathcal{L} \simeq \mathcal{L}\}$ defined in Proposition 1.4.46 but as a group scheme (with the same underlying topological space). Formally, $K(\mathcal{L})$ is the maximal subscheme of A such that $\mathcal{L}^*_{|A\times K(\mathcal{L})}$ is trivial. □

Remark 1.4.48. In the construction of $(\widehat{A}, \mathcal{P})$, we have seen that $(\mathrm{id}_A \times \varphi_{\mathcal{L}})^*\mathcal{P} = \mathcal{L}^*$ with $[\mathcal{L}^*_{|A\times\{x\}}] = [t_x^*\mathcal{L} \otimes \mathcal{L}^{-1}] = \varphi_{\mathcal{L}}(x)$ for all $x \in A$. It follows that $\mathcal{P}_{|A\times\{\varphi_{\mathcal{L}}(x)\}}$ corresponds to $\varphi_{\mathcal{L}}(x)$ for all $x \in A$ and since $\varphi_{\mathcal{L}}$ is surjective by Proposition 1.4.46.(iii), that for all $y \in \widehat{A}$, $\mathcal{P}_{|A\times\{y\}} \in \mathrm{Pic}^0(A_{\kappa(y)})$ corresponds to y canonically.

Remark 1.4.49. Let k'/k be a field extension and consider $T := \mathrm{Spec}(k')$. Then, we have $A_{k'} = A \times_k \mathrm{Spec}(k')$ and point 3 of Theorem 1.4.47 ensures that every line bundle $\mathcal{L} \in \mathrm{Pic}^0(A_{k'})$ uniquely determines a morphism $f : \mathrm{Spec}(k') \longrightarrow \widehat{A}$ i.e. a point of $\widehat{A}(k')$ such that $(\mathrm{id}_A \times f)^*\mathcal{P} \simeq \mathcal{L}$. Hence, we can identify $\widehat{A}(k')$ and $\mathrm{Pic}^0(A_{k'})$. We shall even write $\widehat{A}(k') = \mathrm{Pic}^0(A_{k'})$. In particular, $\widehat{A}(\overline{k}) = \mathrm{Pic}^0(A_{\overline{k}})$, which is the image of $\varphi_{\mathcal{L}} : A \longrightarrow \widehat{A}$ when $\mathcal{L}$ is an ample line bundle of A, as we previously expected.

Corollary 1.4.50.

(i) When $\mathcal{L}$ *is an ample line bundle on* A*, then* $\varphi_{\mathcal{L}}$ *is an isogeny* $A \longrightarrow \widehat{A}$.

(ii) $\dim(A) = \dim(\widehat{A})$.

Proof. (i) is an immediate consequence of Proposition 1.4.46 and Remark 1.4.49. (ii) follows from the fact that an ample line bundle always exist on A by Theorem 1.4.21, from (i) and from the fact that isogenies preserve the dimension by Proposition 1.4.26. □

Dual isogenies

Not only abelian varieties have duals but also isogenies, so dualisation is a (contravariant) functor.

Proposition 1.4.51. *Let* $f : A \longrightarrow B$ *be a homomorphism between abelian varieties over* k*. Let* $\mathcal{P}_A$ *and* $\mathcal{P}_B$ *be the Poincaré sheaves on* $A \times \widehat{A}$ *and* $B \times \widehat{B}$ *respectively. Then, there exists a unique homomorphism* $\widehat{f} : \widehat{B} \longrightarrow \widehat{A}$ *such that* $(id_A \times \widehat{f})^*\mathcal{P}_A \simeq (f \times id_{\widehat{B}})^*\mathcal{P}_B$.

This map $\widehat{f}$ *is called the* dual homomorphism *if* f *and is defined on points by* $[\mathcal{M}] \in \mathrm{Pic}^0(B) \longrightarrow [f^*\mathcal{M}] \in \mathrm{Pic}^0(A)$.

Proof. The line bundle $\mathcal{Q} := (f \times \mathrm{id}_{\widehat{B}})^*\mathcal{P}_B$ on $A \times \widehat{B}$ satisfies $\mathcal{Q}_{|\{0\}\times\widehat{B}} = \mathcal{P}_{B|\{0\}\times\widehat{B}} \simeq \mathcal{O}_{\widehat{B}}$ by point 1 of Theorem 1.4.47 and for all $y \in \widehat{B}$, $\mathcal{Q}_{|A\times\{y\}} = f^*(\mathcal{P}_{B|B\times\{y\}})$, with $\mathcal{P}_{B|B\times\{y\}} \in \mathrm{Pic}^0(B_{\kappa(y)})$ by point 2 of Theorem 1.4.47, so that $\mathcal{Q}_{|A\times\{y\}} \in \mathrm{Pic}^0(A_{\kappa(y)})$ by the following lemma.

Lemma 1.4.52. *Let* k'/k *be a field extension and* $\mathcal{M} \in \mathrm{Pic}^0(B_{k'})$*. Then* $f^*\mathcal{M} \in \mathrm{Pic}^0(A_{k'})$.

Proof. We have for all $x \in A(\overline{k})$,

$$t_x^* f^* \mathcal{M} = (f \circ t_x)^* \mathcal{M} = (t_{f(x)} \circ f)^* \mathcal{M} = f^* t_{f(x)}^* \mathcal{M} \simeq f^* \mathcal{M},$$

so $f^*\mathcal{M} \in \mathrm{Pic}^0(A_{k'})$. □

Hence, the universal property of the Poincaré sheaf $\mathcal{P}_A$ (point 3 of Theorem 1.4.47) ensures the existence of a unique map $\widehat{f} : \widehat{B} \longrightarrow \widehat{A}$ such that $(\mathrm{id}_A \times \widehat{f})^* \mathcal{P}_A \simeq (f \times \mathrm{id}_{\widehat{B}})^* \mathcal{P}_B$.

As we have seen in Remark 1.4.48, for all $x \in \widehat{A}$, $\mathcal{L}_x := \mathcal{P}_{A|A\times\{x\}}$ is the the line bundle of $\mathrm{Pic}^0(A_{\overline{k}})$ corresponding to x and for all $y \in \widehat{B}$, $\mathcal{M}_y := \mathcal{P}_{B|B\times\{y\}}$ is the the line bundle of $\mathrm{Pic}^0(B_{\overline{k}})$ corresponding to y. We then have on the one hand, for all $y \in \widehat{B}$, $((f \times \mathrm{id}_{\widehat{B}})^* \mathcal{P}_B)_{|A\times\{y\}} = f^*(\mathcal{P}_{B|B\times\{y\}}) = f^*\mathcal{M}_y$ and on the other hand $((\mathrm{id}_A \times \widehat{f})^* \mathcal{P}_A)_{|A\times\{y\}} = \mathcal{P}_{A|A\times\{\widehat{f}(y)\}} = \mathcal{L}_{\widehat{f}(y)}$. Hence, $\widehat{f} : \mathrm{Pic}^0(A) \longrightarrow \mathrm{Pic}^0(B), [\mathcal{M}] \longmapsto [f^*\mathcal{M}]$. This completes the proof. □

Proposition 1.4.53. *Let $f : A \longrightarrow B$ be a homomorphism between abelian varieties over k. Then f is an isogeny if an only if $\widehat{f}$ is an isogeny. In that case,* $\deg(f) = \deg(\widehat{f})$.

Proof. If f is an isogeny, then $\ker(f)$ is a finite group scheme. In [Mum74, Theorem 1, p. 143], Mumford proved that $\ker(\widehat{f})$ is isomorphic to the Cartier dual of $\ker(f)$, which is also a finite group scheme. Since $\dim(\widehat{A}) = \dim(A) = \dim(B) = \dim(\widehat{B})$ by Corollary 1.4.50 and Proposition 1.4.26, it follows that $\widehat{f}$ is an isogeny by Proposition 1.4.26. The converse will follow from Proposition 1.4.55.(ii).

The equality between degrees is proved in [Mum74, Corollary 4, p. 131] and follows from the computation of cohomology groups of Poincaré sheaves. □

The following proposition ensures that dualisation is an involution *i.e.* that we can canonically identify A with $\widehat{\widehat{A}}$ (identification we shall make in the future).

Proposition 1.4.54. *Let A be an abelian variety over k and $\mathcal{P}$ its associated Poincaré sheaf.*

(i) The map $x \longmapsto \mathcal{P}_{\{x\}\times\widehat{A}}$ induces a canonical isomorphism $\iota : A \xrightarrow{\sim} \widehat{\widehat{A}}$.

(ii) If $\mathcal{L}$ is an ample line bundle, then $\varphi_{\mathcal{L}} = \widehat{\varphi}_{\mathcal{L}} \circ \iota$.

Proof. To see that ι maps to $\widehat{\widehat{A}}$, we have to prove that $\mathcal{P}_{\{x\}\times\widehat{A}} \in \mathrm{Pic}^0(\widehat{A}_{\kappa(x)})$ for all $x \in A$. Let $\mathcal{L}$ be an ample line bundle and $\mathcal{L}^* := m^*\mathcal{L} \otimes p^*\mathcal{L}^{-1} \otimes q^*\mathcal{L}^{-1}$. Then, $\mathcal{P}$ is the quotient of $\mathcal{L}^*$ by the action of $\{0\} \times K(\mathcal{L})$ lifting the action by translation of $A \times A$ and for all $x \in A$, $\mathcal{L}^*_{|\{x\}\times A} = t_x^*\mathcal{L} \otimes \mathcal{L}^{-1} \in \mathrm{Pic}^0(A_{\kappa(x)})$ is invariant by translation over $\{x\} \times A$ so $\mathcal{P}_{\{x\}\times\widehat{A}}$ is invariant by translation over $\{x\} \times \widehat{A}$ so is an element of $\mathrm{Pic}^0(\widehat{A}_{\kappa(x)})$.

Now, we prove (ii). Consider the swap isomorphisms $s : (x, y) \in A \times A \longrightarrow (y, x) \in A \times A$ and $s' : (x, y) \in \widehat{A} \times A \longrightarrow (y, x) \in A \times \widehat{A}$. Let $x \in A$ that we see as a morphism $T := \mathrm{Spec}(\kappa(x)) \longrightarrow A$. Then, we have

$$\begin{aligned}
\varphi_{\mathcal{L}}(x) &= [t_x^*\mathcal{L} \otimes \mathcal{L}^{-1}] = [\mathcal{L}^*_{|A\times\{x\}}] = [(s^*\mathcal{L}^*)_{|A\times\{x\}}] = [(\mathrm{id}_A \times x)^* s^* \mathcal{L}^*] \\
&= [(\mathrm{id}_A \times x)^* s^* (\mathrm{id}_A \times \varphi_{\mathcal{L}})^* \mathcal{P}] = [((\mathrm{id}_A \times \varphi_{\mathcal{L}}) \circ s \circ (\mathrm{id}_A \times x))^* \mathcal{P}] \\
&= [(s' \circ (\varphi_{\mathcal{L}} \times \mathrm{id}_A) \circ (\mathrm{id}_A \times x))^* \mathcal{P}] = [(s' \circ (\mathrm{id}_{\widehat{A}} \times x) \circ (\varphi_{\mathcal{L}} \times \mathrm{id}_T))^* \mathcal{P}] \\
&= [(\varphi_{\mathcal{L}} \times \mathrm{id}_T)^* (s' \circ (\mathrm{id}_{\widehat{A}} \times x))^* \mathcal{P}] = [(\varphi_{\mathcal{L}} \times \mathrm{id}_T)^* (\mathcal{P}_{|\{x\}\times A})] \\
&= [(\varphi_{\mathcal{L}} \times \mathrm{id}_T)^* (\mathcal{P}_{|\{x\}\times A})] = \widehat{\varphi}_{\mathcal{L}}([\mathcal{P}_{|\{x\}\times A}]) = \widehat{\varphi}_{\mathcal{L}} \circ \iota(x)
\end{aligned}$$

This proves (ii).

Now, since $\mathcal{L}$ is ample, $\varphi_{\mathcal{L}}$ is an isogeny by Proposition 1.4.46. Hence, $\deg(\varphi_{\mathcal{L}}) = \deg(\widehat{\varphi}_{\mathcal{L}})$ by Proposition 1.4.53. Besides, (ii) implies that $\ker(\iota) \subseteq K(\mathcal{L})$ is finite so ι is an isogeny since $\dim(A) = \dim(\widehat{A}) = \dim(\widehat{\widehat{A}})$ by Corollary 1.4.50. (ii) also implies that $\deg(\varphi_{\mathcal{L}}) = \deg(\widehat{\varphi}_{\mathcal{L}}) \deg(\iota)$ so $\deg(\iota) = 1$ and ι is an isomorphism. This proves (i) and completes the proof. □

Proposition 1.4.55. *Let $f : A \longrightarrow B$ and $g : B \longrightarrow C$ be homomorphisms between abelian varieties over k. Then:*

(i) $\widehat{g \circ f} = \widehat{f} \circ \widehat{g}$.

(ii) $\widehat{\widehat{f}} = f$, *with the canonical identifications* $A = \widehat{\widehat{A}}$ *and* $B = \widehat{\widehat{B}}$.

Proof. **(i)** Let $\mathcal{P}_A, \mathcal{P}_B$ and $\mathcal{P}_C$ be the Poincaré sheaves of A, B and C respectively. Then, we have:

$$\begin{aligned}(\mathrm{id}_A \times \widehat{f} \circ \widehat{g})^*\mathcal{P}_A &= (\mathrm{id}_A \times \widehat{g})^*(\mathrm{id}_A \times \widehat{f})^*\mathcal{P}_A \simeq (\mathrm{id}_A \times \widehat{g})^*(f \times \mathrm{id}_{\widehat{B}})^*\mathcal{P}_B = (f \times \widehat{g})^*\mathcal{P}_B \\ &= (f \times \mathrm{id}_{\widehat{C}})^*(\mathrm{id}_B \times \widehat{g})^*\mathcal{P}_B \simeq (f \times \mathrm{id}_{\widehat{C}})^*(g \times \mathrm{id}_{\widehat{C}})^*\mathcal{P}_C = ((g \circ f) \times \mathrm{id}_{\widehat{C}})^*\mathcal{P}_C\end{aligned}$$

By unicity of the homomorphism satisfying $(\mathrm{id}_A \times \widehat{g \circ f})^*\mathcal{P}_A = ((g \circ f) \times \mathrm{id}_{\widehat{C}})^*\mathcal{P}_C$, it follows that $\widehat{g \circ f} = \widehat{f} \circ \widehat{g}$.

(ii) Let $\mathcal{P}_A$ and $\mathcal{P}_B$ be the Poincaré sheaves associated to A and B respectively. Let $\iota_A : A \xrightarrow{\sim} \widehat{\widehat{A}}$ and $\iota_B : B \xrightarrow{\sim} \widehat{\widehat{B}}$ be the canonical isomorphisms from Proposition 1.4.54. We prove that $\widehat{\widehat{f}} \circ \iota_A = \iota_B \circ f$. Let $x \in A$ that we identify with a morphism $T := \mathrm{Spec}(\kappa(x)) \longrightarrow A$. Then, we have:

$$\begin{aligned}\widehat{\widehat{f}} \circ \iota_A(x) &= [(\mathrm{id}_T \circ \widehat{f})^*(\mathcal{P}_{A|\{x\}\times\widehat{A}})] = [(\mathrm{id}_T \circ \widehat{f})^*(x \times \mathrm{id}_{\widehat{A}})^*\mathcal{P}_A] \\ &= [((x \times \mathrm{id}_{\widehat{A}}) \circ (\mathrm{id}_T \circ \widehat{f}))^*\mathcal{P}_A] = [((\mathrm{id}_A \circ \widehat{f}) \circ (x \times \mathrm{id}_{\widehat{B}}))^*\mathcal{P}_A] \\ &= [(x \times \mathrm{id}_{\widehat{B}})^*(\mathrm{id}_A \circ \widehat{f})^*\mathcal{P}_A] = [(x \times \mathrm{id}_{\widehat{B}})^*(f \times \mathrm{id}_{\widehat{B}})^*\mathcal{P}_B] \\ &= [((f \times \mathrm{id}_{\widehat{B}}) \circ (x \times \mathrm{id}_{\widehat{B}}))^*\mathcal{P}_B] = [(f(x) \times \mathrm{id}_{\widehat{B}})^*\mathcal{P}_B] = [\mathcal{P}_{B|\{f(x)\}\times\widehat{B}}] = \iota_B \circ f(x).\end{aligned}$$

This completes the proof. □

Polarisations

Definition 1.4.56. A *polarisation* $\lambda : A \longrightarrow \widehat{A}$ is an isogeny such that there exists an ample line bundle $\mathcal{L}$ on $A_{\overline{k}}$ such that $\lambda_{\overline{k}} = \varphi_{\mathcal{L}}$ (defined in Eq. (1.3)). We say that (A, λ) is a *polarised* abelian variety. When λ is an isomorphism, we say that λ is *principal* and that (A, λ) is *principally polarised*.

Example 1.4.57. If E/k is an elliptic curve, then it is principally polarised. We may consider $\mathcal{L} := \mathcal{L}((0_E))$. Then $\mathcal{L}$ is ample and

$$\begin{array}{rccc}\varphi_{\mathcal{L}} : & E & \longrightarrow & \widehat{E} \\ & P & \longmapsto & \mathcal{L}((-P) - (0_E))\end{array}$$

is a principal polarisation. It can be proved this is the only one [Mil86, Example 13.1].

Definition 1.4.58. Let (A, λ_A) and (B, λ_B) be polarised abelian varieties. An isogeny $f : A \longrightarrow B$ is a *polarised isogeny* $(A, \lambda_A) \longrightarrow (B, \lambda_B)$ when $\widehat{f} \circ \lambda_B \circ f = \lambda_A$.

Definition 1.4.59. Two line bundles $\mathcal{L}$ and $\mathcal{M}$ on an abelian variety A are *algebraically equivalent* if $\mathcal{L} \otimes \mathcal{M}^{-1} \in \mathrm{Pic}^0(A)$.

Lemma 1.4.60. *Let A and B be abelian varieties. Let $\mathcal{L}$ and $\mathcal{M}$ be ample line bundles on A and B respectively. Then $f : A \longrightarrow B$ is a polarised isogeny $(A, \varphi_{\mathcal{L}}) \longrightarrow (B, \varphi_{\mathcal{M}})$ if and only if $f^*\mathcal{M}$ and $\mathcal{L}$ are algebraically equivalent.*

Proof. Assume that $f^*\mathcal{M}$ and $\mathcal{L}$ are algebraically equivalent i.e. $f^*\mathcal{M} \otimes \mathcal{L}^{-1} \in \mathrm{Pic}^0(A)$. Then, Proposition 1.4.51 ensures that for all $x \in A(\overline{k})$,

$$\begin{aligned}\widehat{f} \circ \varphi_{\mathcal{M}} \circ f(x) &= [f^*(t^*_{f(x)}\mathcal{M} \otimes \mathcal{M}^{-1})] = [(t_{f(x)} \circ f)^*\mathcal{M} \otimes f^*\mathcal{M}^{-1}] = [(f \circ t_x)^*\mathcal{M} \otimes f^*\mathcal{M}^{-1}] \\ &= [t_x^* f^*\mathcal{M} \otimes f^*\mathcal{M}^{-1}] = [t_x^*\mathcal{L} \otimes \mathcal{L}^{-1}] = \varphi_{\mathcal{L}}(x),\end{aligned}$$

where the last equality comes from the fact that $f^*\mathcal{M} \otimes \mathcal{L}^{-1} \in \mathrm{Pic}^0(A)$. Hence, f is a polarised isogeny $(A, \varphi_{\mathcal{L}}) \longrightarrow (B, \varphi_{\mathcal{M}})$.

Conversely, if f is a polarised isogeny $(A, \varphi_{\mathcal{L}}) \longrightarrow (B, \varphi_{\mathcal{M}})$, then the equality $\widehat{f} \circ \varphi_{\mathcal{M}} \circ f = \varphi_{\mathcal{L}}$ ensures that $t_x^* f^*\mathcal{M} \otimes f^*\mathcal{M}^{-1} \simeq t_x^*\mathcal{L} \otimes \mathcal{L}^{-1}$ for all $x \in A(\overline{k})$ i.e. that $f^*\mathcal{M} \otimes \mathcal{L}^{-1} \in \mathrm{Pic}^0(A)$. This completes the proof. □

Polarisations and cohomology of invertible sheaves

Let A be an abelian variety over k. Then, we can define a cohomology of sheaves of abelian groups on A given by the right derived functor of $\Gamma(A, \cdot)$ (see [Har77, § 3.2]). A theorem due to Grothendieck [Har77, Theorem 2.7] ensures that if $g := \dim(A)$, then for any sheaf $\mathcal{F}$ of abelian group on A and $i > g$, we have $H^i(A, \mathcal{F}) = 0$.

Definition 1.4.61. In particular, if $\mathcal{L}$ is a line bundle on A, we have can define

$$\chi(\mathcal{L}) := \sum_{i=0}^{\infty} (-1)^i \dim_k H^i(A, \mathcal{L}) = \sum_{i=0}^{g} (-1)^i \dim_k H^i(A, \mathcal{L}).$$

Then, we have the following results proved in [Mum74, § 16, p.150].

Theorem 1.4.62 (Riemann-Roch)**.** *For any line bundle $\mathcal{L}$ on A we have:*

(i) $\chi(\mathcal{L}) = (\mathcal{L}^g \cdot A)/g!$ *where* $g := \dim(A)$ *and* $(\mathcal{L}^g \cdot A)$ *is the self intersection number of* $\mathcal{L}$ *(or of its associated divisor).*

(ii) $\deg(\varphi_{\mathcal{L}}) = \chi(\mathcal{L})^2$.

Theorem 1.4.63. *If $\mathcal{L}$ is a line bundle on A such that $K(\mathcal{L})$ is finite, then there exists an integer $i(\mathcal{L}) \in [\![0\ ;\ g]\!]$ such that $H^{i(\mathcal{L})}(A, \mathcal{L}) \neq \{0\}$ and $H^i(A, \mathcal{L}) = \{0\}$ for all $i \in \mathbb{N} \setminus \{i(\mathcal{L})\}$.*

Corollary 1.4.64. *If $\mathcal{L}$ is an ample line bundle on A, then $i(\mathcal{L}) = 0$ and $\chi(\mathcal{L}) = \dim_k \Gamma(A, \mathcal{L})$.*

Proof. If $\mathcal{L}$ is ample, then non trivial global sections exist on $\mathcal{L}$ so we have $\dim_k \Gamma(A, \mathcal{L}) > 0$. Since $H^0(A, \mathcal{L}) = \Gamma(A, \mathcal{L})$, and $K(\mathcal{L})$ is finite by Proposition 1.4.46.(ii), Theorem 1.4.63 applies with $i(\mathcal{L}) = 0$. We then have $\chi(\mathcal{L}) = \dim_k H^0(A, \mathcal{L}) = \dim_k \Gamma(A, \mathcal{L})$. □

1.4.6 The Weil pairing

In this section, we fix an abelian variety A over k. We construct Weil pairings on A following the proof from Silverman [Sil09, § III.8] over elliptic curves, but in the abelian varieties setting. Unlike in previous sections, most proofs are fairly elementary here so we give them *in extenso*.

Lemma 1.4.65. *Let $\mathcal{L}$ be a line bundle on A. Then the following are equivalent:*

(i) $[\mathcal{L}] \in \mathrm{Pic}^0(A)$.

(ii) $m^*\mathcal{L} \simeq p^*\mathcal{L} \otimes q^*\mathcal{L}$, *where $m, p, q : A \times A \longrightarrow A$ are respectively the addition, the projection on the first component and the projection on the second component.*

(iii) $[n]^*\mathcal{L} \simeq \mathcal{L}^n$ *for all $n \in \mathbb{Z}$.*

(iv) $\mathcal{L}$ *is antisymmetric i.e.* $[-1]^*\mathcal{L} \simeq \mathcal{L}^{-1}$.

(v) There exists $n \in \mathbb{N}^*$ such that $[\mathcal{L}^n] \in \mathrm{Pic}^0(A)$.

Proof. (i) $\iff$ (ii) If $\mathcal{L} \in \mathrm{Pic}^0(A)$, then, we have for all $x \in A(\overline{k})$,

$$m^*\mathcal{L}_{|A\times\{x\}} \simeq t_x^*\mathcal{L} \simeq \mathcal{L} \simeq (p^*\mathcal{L} \otimes q^*\mathcal{L})_{|A\times\{x\}},$$

and $m^*\mathcal{L}_{|\{0\}\times A} \simeq \mathcal{L} \simeq (p^*\mathcal{L} \otimes q^*\mathcal{L})_{|\{0\}\times A}$ so the seesaw principle (Theorem 1.4.14) applies and we have $m^*\mathcal{L} \simeq p^*\mathcal{L} \otimes q^*\mathcal{L}$. Conversely, if $m^*\mathcal{L} \simeq p^*\mathcal{L} \otimes q^*\mathcal{L}$, then for all $x \in A(\overline{k})$, we obtain that $t_x^*\mathcal{L} \simeq \mathcal{L}$ by applying the restriction to $A \times \{x\}$.

(ii) $\implies$ (iii) Assume (ii). We proceed by induction on $n \in \mathbb{N}$. For $n = 0$, the result is trivial. Now, let $n \in \mathbb{N}^*$ and assume the result at rank $n - 1$. Then, applying $([n-1] \times [1])^*$ to (ii), we obtain:

$$[n]^*\mathcal{L} \simeq [n-1]^*\mathcal{L} \otimes \mathcal{L} \simeq \mathcal{L}^{n-1} \otimes \mathcal{L} = \mathcal{L}^n,$$

We proceed similarly for $n \in \mathbb{Z} \setminus \mathbb{N}$.

(iii) $\Longrightarrow$ (iv) is trivial. (iv) $\Longrightarrow$ (v) Assume that $\mathcal{L}$ is antisymmetric. Then $\mathcal{L}^2 \simeq [-1]^*\mathcal{L}^{-1} \otimes \mathcal{L}$ and for all $x \in A(\overline{k})$,

$$\varphi_{[-1]^*\mathcal{L}^{-1}}(x) = [t_x^*[-1]^*\mathcal{L}^{-1} \otimes [-1]^*\mathcal{L}] = [[-1]^*(t_{-x}^*\mathcal{L} \otimes \mathcal{L}^{-1})^{-1}],$$

where $t_{-x}^*\mathcal{L} \otimes \mathcal{L}^{-1} \in \mathrm{Pic}^0(A_{\overline{k}})$ by Proposition 1.4.46.(iii) so $[-1]^*(t_{-x}^*\mathcal{L} \otimes \mathcal{L}^{-1})^{-1} \simeq t_{-x}^*\mathcal{L} \otimes \mathcal{L}^{-1}$ since (i) $\Longrightarrow$ (iv), and

$$\varphi_{[-1]^*\mathcal{L}^{-1}}(x) = [t_{-x}^*\mathcal{L} \otimes \mathcal{L}^{-1}] = \varphi_{\mathcal{L}}(-x) = -\varphi_{\mathcal{L}}(x).$$

Hence, $\varphi_{[-1]^*\mathcal{L}^{-1}\otimes\mathcal{L}}(x) = -\varphi_{\mathcal{L}}(x) + \varphi_{\mathcal{L}}(x) = 0$ for all $x \in A(\overline{k})$, so $[[-1]^*\mathcal{L}^{-1} \otimes \mathcal{L}] \in \mathrm{Pic}^0(A)$ and $[\mathcal{L}^2] \in \mathrm{Pic}^0(A)$.

(v) $\Longrightarrow$ (i) Let $n \in \mathbb{N}^*$ such that $[\mathcal{L}^n] \in \mathrm{Pic}^0(A)$. Then for all $x \in A(\overline{k})$, $0 = \varphi_{\mathcal{L}^n}(x) = [n]\varphi_{\mathcal{L}}(x) = \varphi_{\mathcal{L}} \circ [n](x)$ i.e. $\varphi_{\mathcal{L}} \circ [n] = 0$. Since $[n]$ is surjective (as an isogeny), we have $\varphi_{\mathcal{L}} = 0$ so $[\mathcal{L}] \in \mathrm{Pic}^0(A)$. This completes the proof. □

Lemma 1.4.66. *Let $f : A \longrightarrow B$ be a separable isogeny between abelian varieties over k and $s \in \overline{k}(A)$. Then, $t_x^*s = s$ for all $x \in \ker(f)$ if and only if there exists $t \in \overline{k}(B)$ such that $s = f^*t$.*

Proof. Let $K := \ker(f)$. Our goal is to prove that $\overline{k}(A)^K = f^*\overline{k}(B)$, where is the subfield of $\overline{k}(A)$ fixed by the action of K by translation. We easily prove that $f^*\overline{k}(B) \subseteq \overline{k}(A)^K$. Indeed, if $t \in \overline{k}(B)$, then for all $x \in K$, $t_x^*f^*t = (f \circ t_x)^*t = f^*t$.

Conversely, consider $G := \mathrm{Aut}(\overline{k}(A)/f^*\overline{k}(B))$. Then, by Artin's lemma [Lan04, Theorem VI.1.8] we know that $\overline{k}(A)/\overline{k}(A)^G$ is a finite Galois extension of degree $\#G$. Since f is separable, $\overline{k}(A)/f^*\overline{k}(B)$ is a finite separable extension of degree $\deg(f) = \#K$. Since, $f^*\overline{k}(B)$ is fixed by G, we have $f^*\overline{k}(B) \subseteq \overline{k}(A)^G$ and we only have to prove that $\#G = \#K$ to conclude.

We actually prove that $K \longrightarrow G, x \longmapsto t_x^*$ is a group isomorphism (where t_x^* is the action by translation on functions of $\overline{k}(A)$). Indeed, if $x \in K \setminus \{0\}$ then t_x^* is not trivial, i.e. there exists $s \in \overline{k}(A)$ such that $t_x^*s \neq s$. Indeed, A being projective, we may assume that $A_{\overline{k}} \subseteq \mathbb{P}^n_{\overline{k}} = \mathrm{Proj}(\overline{k}[T_0, \cdots, T_n])$ and write $x := (x_0 : \cdots : x_n)$ and $0_A := (e_0 : \cdots : e_n)$, with $x_i \neq e_i$ for some $i \in [\![1\ ;\ n]\!]$. Then, the function $s := T_i - x_i$ satisfies $s(0_A) = e_i - x_i \neq 0$ and $t_x^*s(0_A) = s \circ t_x(0_A) = s(x) = 0$ so $t_x^*s \neq s$. It follows that $x \longmapsto t_x^*$ is injective and that $\#K \leq \#G$ but $\#K \geq \#G$ since $f^*\overline{k}(B) \subseteq \overline{k}(A)^G$ so this homomorphism is also surjective. This completes the proof. □

Theorem 1.4.67. *For all $n \in \mathbb{N}^*$ such that $\mathrm{char}(k) \nmid n$, there exists a non-degenerate pairing $e_n : A[n](\overline{k}) \times \widehat{A}[n](\overline{k}) \longrightarrow k^*$ called the n-th* Weil pairing *of A.*

Proof. **Construction:** The construction of the Weil pairing on any abelian variety is very similar to elliptic curves. Let $n \in \mathbb{N}^*$ such that $\mathrm{char}(k) \nmid n$. Let $y \in \widehat{A}[n](\overline{k})$. Then y is represented canonically by a divisor $D_y \in \mathrm{Pic}^0(A_{\overline{k}})$ and we have by Lemma 1.4.65.(ii), $[n]^*D_y \sim nD_y$ and $nD_y \sim 0$ since $[n]y = 0$. Hence, there exists sections $f_y, g_y \in \overline{k}(A)$ such that $nD_y = \mathrm{div}(f_y)$ and $[n]^*D_y = \mathrm{div}(g_y)$. The abelian variety A being projective, f_y and g_y can be seen as functions of a subvariety of the projective space $\mathbb{P}^N_k$ for some $N \in \mathbb{N}^*$. Then we shall denote them as functions, using the composition instead of the pull-back notation. In particular, we may write

$$\mathrm{div}(f_y \circ [n]) = [n]^* \mathrm{div}(f_y) = [n]^*(nD_y) = n[n]^*D_y = n\,\mathrm{div}(g_y) = \mathrm{div}(g_y^n). \tag{1.4}$$

It follows that $\mathrm{div}(g_y^n/f_y \circ [n]) = 0$ so that $g_y^n/f_y \circ [n] \in \Gamma(A_{\overline{k}}, \mathcal{L}(0)) = \Gamma(A_{\overline{k}}, \mathcal{O}_{A_{\overline{k}}})$. But $\Gamma(A_{\overline{k}}, \mathcal{O}_{A_{\overline{k}}}) = \overline{k}$ since $A_{\overline{k}}$ is a projective variety defined over an algebraically closed field [Har77, Theorem I.3.4] and we may write $c := g_y^n/f_y \circ [n] \in \overline{k}$. Informally, a function without zero and poles is constant. Besides, if $x \in A[n](\overline{k})$ then we have,

$$g_y^n \circ t_x = c \cdot f_y \circ [n] \circ t_x = c \cdot f_y \circ t_{[n]x} \circ [n] = c \cdot f_y \circ [n] = g_y^n,$$

so $g_y/g_y \circ t_x$ is a constant (since $n\,\mathrm{div}(g_y/g_y \circ t_x) = 0$) and an n-th root of unity in $\overline{k}$. We can then define $e_n(x, y) := g_y/g_y \circ t_x$. This defines a map $A[n](\overline{k}) \times \widehat{A}[n](\overline{k}) \longrightarrow \overline{k}^*$. Now, we verify that e_n satisfies the desired properties.

Bilinearity: Let $x, x' \in A[n](\overline{k})$ and $y, y' \in \widehat{A}[n](\overline{k})$. Then, $g_y/g_y \circ t_x$ being constant, we have $g_y/g_y \circ t_x = g_y \circ t_{x'}/g_y \circ t_{x+x'}$, so that:

$$e_n(x+x', y) = \frac{g_y}{g_y \circ t_{x+x'}} = \frac{g_y \circ t_{x'}}{g_y \circ t_{x+x'}} \frac{g_y}{g_y \circ t_{x'}} = \frac{g_y}{g_y \circ t_x} \frac{g_y}{g_y \circ t_{x'}} = e_n(x,y)e_n(x',y).$$

Besides, the canonical isomorphism $A^0(\overline{k}) \cong \mathrm{Pic}^0(A_{\overline{k}})$ ensures that $D_{y+y'} \sim D_y + D_{y'}$ so there exists $h \in \overline{k}(A)$ such that $\mathrm{div}(h) = D_{y+y'} - D_y - D_{y'}$. It follows that:

$$\mathrm{div}(h \circ [n]) = [n]^* \mathrm{div}(h) = [n]^* D_{y+y'} - [n]^* D_y - [n]^* D_{y'} = \mathrm{div}(g_{y+y'}) - \mathrm{div}(g_y) - \mathrm{div}(g_{y'}),$$

so that $g_{y+y'} = g_y \cdot g_{y'} \cdot h \circ [n]$ up to a constant (that can be ignored). We then have:

$$e_n(x, y+y') = \frac{g_{y+y'}}{g_{y+y'} \circ t_x} = \frac{g_y \cdot g_{y'} \cdot h \circ [n]}{g_y \circ t_x \cdot g_{y'} \circ t_x \cdot h \circ [n] \circ t_x} = \frac{g_y \cdot g_{y'}}{g_y \circ t_x \cdot g_{y'} \circ t_x} = e_n(x,y)e_n(x,y'),$$

where we used the fact that $[n] \circ t_x = [n]$ since $[n]x = 0$. This proves that e_n is a pairing.

Non-degeneracy: Let $y \in \widehat{A}[n](\overline{k})$ such that for all $x \in A[n](\overline{k})$, $e_n(x,y) = 1$. Then for all $x \in A[n](\overline{k})$, $g_y = g_y \circ t_x$ so g_y factors through $[n]$ by the Lemma 1.4.66 ($[n]$ being a separable isogeny) i.e. we may write $g_y = h_y \circ [n]$ for some $h_y \in \overline{k}(A)$. It follows by Eq. (1.4) that $\mathrm{div}(f_y \circ [n]) = \mathrm{div}(h_y^n \circ [n])$ so that $f_y \circ [n] = c \cdot h_y^n \circ [n]$ for some $c \in \overline{k}^*$. Since $[n]$ is an isogeny, it is surjective so $f_y = c \cdot h_y^n$ and $nD_y = \mathrm{div}(f_y) = \mathrm{div}(h_y^n) = n\,\mathrm{div}(h_y)$, so that $D_y = \mathrm{div}(h_y) \sim 0$ and $y = 0$.

It follows that $y \in \widehat{A}[n](\overline{k}) \longmapsto e_n(.,y) \in \widehat{A[n](\overline{k})}$ is injective so it is a group isomorphism since $\#\widehat{A[n](\overline{k})} = \#A[n](\overline{k}) = n^{2g} = \#\widehat{A}[n]$ by Corollary 1.4.34. Hence, if $x \in A[n](\overline{k})\backslash\{0\}$, then there exists $\chi \in \widehat{A[n](\overline{k})}$ such that $\chi(x) \neq 1$ so there exists $y \in \widehat{A}[n](\overline{k})$ such that $e_n(.,y) = \chi$ and $e_n(x,y) \neq 1$. This proves that e_n is non-degenerate and completes the proof. □

Remark 1.4.68. If E is an elliptic curve over k, there is a unique canonical principal polarisation identifying E with its dual $\widehat{E}$ (Example 1.4.57). Therefore, the n-th Weil pairing of E defines a map $e_n : E[n](\overline{k}) \times E[n](\overline{k}) \longrightarrow k^*$. However, on an abelian variety A, we have to specify the polarisation relating A to $\widehat{A}$.

When $\lambda : A \longrightarrow \widehat{A}$ is a homomorphism and $n \in \mathbb{N}^*$ is coprime with $\mathrm{char}(k)$, we denote by e_n^λ, the pairing:

$$e_n^\lambda : A[n](\overline{k}) \times A[n](\overline{k}) \longrightarrow \overline{k}^*, (x,y) \longmapsto e_n(x, \lambda(x)).$$

When $\lambda = \varphi_{\mathcal{L}}$ for some ample line bundle $\mathcal{L}$ on $A_{\overline{k}}$, we denote $e_n^{\mathcal{L}} := e_n^\lambda$.

Proposition 1.4.69. *If $\lambda : A \longrightarrow \widehat{A}$ is a polarisation, then e_n^λ is skew-symmetric.*

Proof. We follow [Mum74, Theorem 1, p. 186]. It suffices to prove that for all $x \in A[n](\overline{k})$, we have $e_n^\lambda(x,x) = 1$. Let D be a divisor on $A_{\overline{k}}$ representing λ ($\lambda = \varphi_{\mathcal{L}(D)}$). Then $e_n^\lambda(x,x) = e_n(x, \lambda(x)) = g_{\lambda(x)}/g_{\lambda(x)} \circ t_x$ with $\mathrm{div}(g_{\lambda(x)}) = [n]^*(t_x^* D - D)$ since $t_x^* D - D$ is the divisor that canonically represents $\lambda(x) \in \widehat{A}_{\overline{k}} = \mathrm{Pic}^0(A_{\overline{k}})$. Let $y \in A(\overline{k})$ such that $[n]y = x$ (which exists since $[n]$ is surjective, as any isogeny). Then $[n] \circ t_y = t_x \circ [n]$, so that

$$\mathrm{div}(g_{\lambda(x)}) = [n]^*(t_x^* D - D) = t_y^* [n]^* D - [n]^* D,$$

and for all $i \in [\![0\ ;\ n-1]\!]$,

$$\mathrm{div}(g_{\lambda(x)} \circ t_{[i]y}) = t_{[i]y}^* \mathrm{div}(g_{\lambda(x)}) = t_{[i+1]y}^* [n]^* D - t_{[i]y}^* [n]^* D.$$

Hence,

$$\begin{aligned} \mathrm{div}\left(\prod_{i=0}^{n-1} g_{\lambda(x)} \circ t_{[i]y}\right) &= \sum_{i=0}^{n-1} (t_{[i+1]y}^* [n]^* D - t_{[i]y}^* [n]^* D) = t_{[n]y}^* [n]^* D - [n]^* D \\ &= t_x^* [n]^* D - [n]^* D = [n]^* t_{nx}^* D - [n]^* D = 0, \end{aligned}$$

since $nx = 0$. Hence, $h := \prod_{i=0}^{n-1} g_{\lambda(x)} \circ t_{[i]y}$ is constant and we have:

$$1 = \frac{h \circ t_y}{h} = \frac{\prod_{i=0}^{n-1} g_{\lambda(x)} \circ t_{[i+1]y}}{\prod_{i=0}^{n-1} g_{\lambda(x)} \circ t_{[i]y}} = \frac{g_{\lambda(x)} \circ t_{[n]y}}{g_{\lambda(x)}} = \frac{g_{\lambda(x)} \circ t_x}{g_{\lambda(x)}}.$$

Hence $e_n^\lambda(x,x) = g_{\lambda(x)}/g_{\lambda(x)} \circ t_x = 1$. This completes the proof. □

Proposition 1.4.70. *Let $f : A \longrightarrow B$ be a homomorphism of abelian varieties over k, $\lambda : A \longrightarrow \widehat{A}$ be a polarisation and $n, m \in \mathbb{N}^*$ such that $\mathrm{char}(k) \nmid nm$. Then:*

(i) $\forall x \in A[n](\overline{k}), y \in \widehat{B}[n](\overline{k})$, $e_n(x, \widehat{f}(y)) = e_n(f(x), y)$.

(ii) $\forall x, y \in A[n](\overline{k})$, $e_n^{\widehat{f} \circ \lambda \circ f}(x,y) = e_n^\lambda(f(x), f(y))$.

(iii) $\forall x \in A[mn](\overline{k}), y \in \widehat{A}[mn](\overline{k})$, $e_{mn}(x,y)^n = e_m([n]x, [n]y)$.

Proof. **(i)** Let $x \in A[n](\overline{k})$ and $y \in \widehat{B}[n](\overline{k})$. Let D_y be the divisor on B representing y and $g_y \in \overline{k}(B)$ such that $[n]^* D_y = \mathrm{div}(g_y)$. Then $f^* D_y$ represents $\widehat{f}(y)$ by Proposition 1.4.51 and $[n]^* f^* D_y = f^* [n]^* D_y = \mathrm{div}(g_y \circ f)$, so that

$$e_n(x, \widehat{f}(y)) = \frac{g_y \circ f}{g_y \circ f \circ t_x} = \frac{g_y \circ f}{g_y \circ t_{f(x)} \circ f} = \frac{g_y}{g_y \circ t_{f(x)}} = e_n(f(x), y),$$

since $g_y / g_y \circ t_{f(x)}$ is constant. This proves (i). (ii) follows immediately.

(iii) Let $x \in A[mn](\overline{k})$ and $y \in \widehat{A}[mn](\overline{k})$. Let D_y be the divisor of $A_{\overline{k}}$ associated to y and $g_y \in \overline{k}(A)$ such that $[nm]^* D_y = \mathrm{div}(g_y)$. Then nD_y represents $[n]y$ and there exists $g_{[n]y} \in \overline{k}(A)$ such that $[m]^*(nD_y) = \mathrm{div}(g_{[n]y})$. It follows that:

$$\mathrm{div}(g_{[n]y} \circ [n]) = [n]^* \mathrm{div}(g_{[n]y}) = [n]^*[m]^*(nD_y) = n[nm]^* D_y = \mathrm{div}(g_y^n),$$

so that $g_y^n = c \cdot g_{[n]y} \circ [n]$ for some constant $c \in \overline{k}^*$. Consequently,

$$e_{nm}(x,y)^n = \frac{g_y^n}{(g_y \circ t_x)^n} = \frac{g_{[n]y} \circ [n]}{g_{[n]y} \circ [n] \circ t_x} = \frac{g_{[n]y} \circ [n]}{g_{[n]y} \circ (t_{[n]x} \circ [n])} = \frac{g_{[n]y}}{g_{[n]y} \circ t_{[n]x}} = e_{nm}([n]x, [n]y).$$

This completes the proof. □

Part I

Cryptographic applications of higher dimensional isogenies

Chapter 2

Improving ideal-to-isogeny translation algorithms

In this chapter, we present several algorithms to translate quaternion ideals into isogenies between supersingular elliptic curves. In Section 2.1, we first give an overview of algorithms introduced in the first version of SQIsign [DFKLPW20] based on the KLPT method [KLPT14] due to Kohel, Lauter, Petit and Tignol to smoothen ideal norms. We then present Kani's embedding lemma [Kan97] in Section 2.2 that has been used for the first time to completely break the isogeny based protocol and NIST candidate SIDH (Supersingular Isogeny Diffie Hellman) [JDF11] before being used constructively in several cryptographic applications. These first two sections are mainly literature reviews without original result.

We then introduce several contributions of this PhD that led to dramatic improvements of state of the art ideal-to-isogeny translation algorithms with Kani's lemma. In Section 2.3, we present one of the first constructive applications of Kani's lemma using 4-dimensional isogenies to translate ideals into isogenies. This construction has been used in SQIsignHD [DLRW24], a variant of the digital signature scheme and NIST candidate SQIsign [DFKLPW20; DLW22], with significant improvements of the signing time and security proof at the expense of the verification. In Section 2.4, we present a new algorithm using 2-dimensional isogenies only based on the Clapoti method (class group action in polynomial time) introduced by Aurel Page and Damien Robert [PR23] after SQIsignHD. This algorithm has been used in the SQIsign2D-West variant of SQIsign [BDF+25] achieving better performance than SQIsignHD in terms of verification time, while preserving a competitive signing time and improving further the security proof. Finally, in Section 2.5 we present an algorithm for practical effective class group action on oriented elliptic curves using 4-dimensional isogenies (PEGASIS) [DEF+25], also based on the Clapoti method and with very good performance when applied to CSIDH.

Note that this chapter does not introduce any cryptographic protocol but only algorithms. We refer to Chapters 3 and 4 for a presentation of SQIsignHD and SQIsign2D-West.

2.1 KLPT based techniques of ideal-to-isogeny translation and applications

In this section, we motivate the use of the Deuring correspondence in cryptographic protocols. We then present techniques introduced for the original version of SQIsign [DFKLPW20; DLW22] to make this Deuring correspondence effective by translating ideals of smooth norm into isogenies.

2.1.1 A constructive use of the Deuring correspondence

Let E_1 and E_2 be supersingular elliptic curves over $\mathbb{F}_{p^2}$. Assume that we know their endomorphism rings $\mathrm{End}(E_1)$ and $\mathrm{End}(E_2)$ *i.e.* that we know a maximal order $\mathcal{O}_i \subset \mathcal{B}_{p,\infty}$ and an isomorphism $\iota_i : \mathcal{O}_i \xrightarrow{\sim} \mathrm{End}(E_i)$ for $i \in \{1, 2\}$. We can use this information to compute an isogeny $\varphi : E_1 \longrightarrow E_2$ as follows.

First, we find a *connecting ideal* I between $\mathcal{O}_1$ and $\mathcal{O}_2$, that is to say a left $\mathcal{O}_1$-ideal that is also a right $\mathcal{O}_2$-ideal. For instance, we may take $I := d \cdot \mathcal{O}_1 \cdot \mathcal{O}_2$, with $d \in \mathbb{N}^*$ such that $I \subseteq \mathcal{O}_1 \cap \mathcal{O}_2$. By the Deuring correspondence, this ideal determines an isogeny $\varphi_I : E_1 \longrightarrow E_2$. However, this ideal has no reason to be of smooth norm, so to be translatable into an isogeny with standard techniques. For that reason, we look for an equivalent ideal $J \sim I$ of smooth norm that can then be translated into an isogeny $\varphi_J : E_1 \longrightarrow E_2$ that is efficiently representable with one-dimensional techniques because it has smooth degree $\mathrm{nrd}(J)$. Such an equivalent ideal of smooth norm $J \sim I$ can be obtained by the KLPT algorithm [KLPT14] due to Kohel, Lauter, Petit and Tignol. The outputs of KLPT generally have very big norm (at least $\widetilde{O}(p^{9/2})$ in general, reduced to $\widetilde{O}(p^{15/4})$ in SQIsign [DFKLPW20, § 8.3]), which hampers efficiency in practice as we shall see in Section 2.1.4.

Nonetheless, all of the above operations take polynomial time (in $\log(p)$). In addition, they become very difficult when either one of the endomorphism rings $\mathrm{End}(E_1)$ or $\mathrm{End}(E_2)$ is unknown. Finding an isogeny $\varphi : E_1 \longrightarrow E_2$ in this context is essentially an instance of the supersingular isogeny problem below.

Problem 2.1.1 (Supersingular Isogeny Problem)**.** Given two supersingular elliptic curves E_1, E_2 defined over $\mathbb{F}_{p^2}$, find an efficient representation of an isogeny $\varphi : E_1 \longrightarrow E_2$.

This problem has been proved to be equivalent to the supersingular endomorphism ring problem below [Wes22; PW24; MW25]. The best known algorithms to solve the latter problem all have exponential complexity $\widetilde{O}(\sqrt{p})$ on a classical computer [DG16; EHLMP20; FIKMN25] and $\widetilde{O}(p^{1/4})$ on a quantum computer with a Grover search [Gro96; BJS14].

Problem 2.1.2 (Supersingular Endomorphism Ring Problem)**.** Given a supersingular elliptic curve E defined over $\mathbb{F}_{p^2}$, find a maximal order $\mathcal{O} \subset \mathcal{B}_{p,\infty}$ and an isomorphism $\iota : \mathcal{O} \xrightarrow{\sim} \mathrm{End}(E)$.

In practice, this equivalence between Problem 2.1.1 and Problem 2.1.2 is used, since the knowledge of the endomorphism rings $\mathrm{End}(E_i)$ is provided by the knowledge of an isogeny $\varphi_i : E_0 \longrightarrow E_i$ starting from a supersingular curve $E_0/\mathbb{F}_{p^2}$ of known endomorphism ring. For instance, when $p \equiv 3 \mod 4$, the elliptic curve $E_0 : y^2 = x^3 + x$ is an easy instance of the endomorphism ring problem. Indeed, $\mathrm{End}(E_0)$ is isomorphic to $\mathcal{O}_0 := \langle 1, i, (i+j)/2, (1+ij)/2 \rangle \subset \mathcal{B}_{p,\infty}$ (where $i^2 = -1$ and $j^2 = -p$) by Lemma 1.2.25.

Both features of this isogeny path problem with secret endomorphism ring information - being solvable in polynomial time and hard without the necessary secret information - are convenient for cryptographic applications. For instance, this problem can be used in identity based protocols like digital signatures where parties have to prove they know some secret identifying them without revealing it.

2.1.2 Piecewise ideal-to-isogeny translation

In the following, we explain how to translate a left ideal $I \subseteq \mathcal{O}_1$ of norm ℓ^e (where ℓ is a prime) into an isogeny $\varphi_I : E_1 \longrightarrow E_2$. We assume that I is *primitive* in the sense of Definition 2.1.3 below, so that φ_I is cyclic.

Definition 2.1.3. Let $\mathcal{O} \subseteq \mathcal{B}_{p,\infty}$ be a maximal order. We say that a left $\mathcal{O}$-ideal $I \subseteq \mathcal{O}$ is *primitive* if $I \not\subseteq n\mathcal{O}$ for any integer $n > 1$. The isogeny associated to such an ideal is cyclic.

The straight forward method (proposed in [GPS20] for instance) would be to compute:

$$\ker(\varphi_I) = E[I] = \{P \in E(\overline{\mathbb{F}}_{p^2}) \mid \forall \alpha \in I, \quad \alpha(P) = 0\}$$

and then to apply Vélu's formulas [Vé71] to obtain φ_I. However, $E[I]$ is a cyclic subgroup of order ℓ^e of the ℓ^e-torsion $E[\ell^e]$ that may be defined over a field extension of $\mathbb{F}_{p^2}$ of exponentially large degree, especially if I is an output of KLPT ($\ell^e = \Omega(p^{15/4})$). This would make the use of Vélu's formulas impractical.

The solution proposed in SQIsign [DFKLPW20] to circumvent this difficulty was to cut the ideal-to-isogeny translation into pieces. We write $I = I_1 \cdot \cdots \cdot I_n$, with ideals $I_1, \cdots, I_n$ of norm dividing ℓ^f so that the ℓ^f-torsion is defined over a small field extension of $\mathbb{F}_{p^2}$ or even on $\mathbb{F}_{p^2}$ itself (by requiring

that $\ell^f | p+1$). For all $i \in [\![1\ ;\ n]\!]$, let $\psi_i : E'_i \longrightarrow E'_{i+1}$ be the isogeny associated to I_i. Assuming that the knowledge of $\mathcal{O}_1 \simeq \mathrm{End}(E_1)$ is given by an ℓ^g-isogeny $\varphi_1 : E_0 \longrightarrow E_1$ (and its associated ideal J_1), as explained in Section 2.1.1, we are able to compute ψ_1 from the knowledge of φ_1 and ψ_{i+1} from the knowledge of $\psi_i \circ \cdots \circ \psi_1 \circ \varphi_1$ for all $i \in [\![1\ ;\ n-1]\!]$. This will be explained in the next subsection.

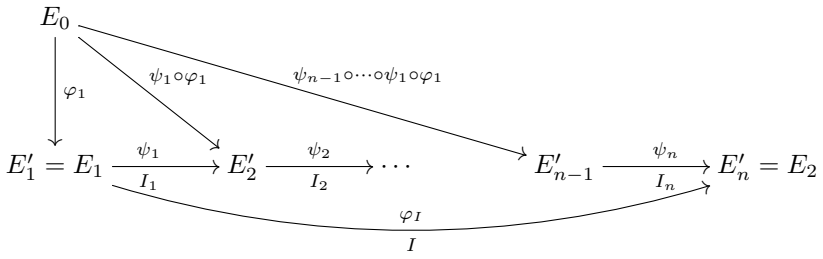

2.1.3 How to translate a piece of ideal

As we have seen, the translation of an ideal I of big norm ℓ^e reduces to the following problem.

Problem 2.1.4. Let $\varphi_J : E_0 \longrightarrow E$ be an ℓ^g-isogeny of associated left $\mathcal{O}_0$-ideal J. Let $\mathcal{O} := O_R(J) \simeq \mathrm{End}(E)$ and I' be a left $\mathcal{O}$-ideal of norm ℓ^f such that $E[\ell^f] \subseteq E(\mathbb{F}_{p^2})$. We want to compute the isogeny $\varphi_{I'} : E \longrightarrow E'$.

With the notations from Section 2.1.2, if we want to compute ψ_i, we solve this problem with $I' := I_i$, $J := J_1 \cdot I_1 \cdot \cdots \cdot I_{i-1}$ and $\varphi_J := \psi_{i-1} \circ \cdots \circ \psi_1 \circ \varphi_1$.

The naive method to solve Problem 2.1.4 would be to compute $E[I']$ directly, but this would mean evaluating elements of $I' \subseteq \mathcal{O}$ on $E[\ell^f]$. Unfortunately, this is not possible with standard techniques. Indeed, since we know $\mathrm{End}(E_0)$ and an ℓ^g-isogeny $\varphi_J : E_0 \longrightarrow E$, we can evaluate endomorphisms of E but only on points of order coprime with ℓ (*e.g.* with the techniques of [EHLMP18, Algorithm 4]).

The following trick was proposed in [DLW22] to circumvent this difficulty. The idea is to find an endomorphism $\theta \in \mathrm{End}(E)$ of degree coprime with ℓ decomposed into two isogenies $\theta := \widehat{\rho}_2 \circ \rho_1$, with $\rho_i : E_0 \longrightarrow E$ for $i \in \{1, 2\}$. The ρ_i can be evaluated on ℓ^f-torsion points, as well as θ. This endomorphism θ can then be leveraged to compute I' by the following lemma.

Lemma 2.1.5. *[DLW22, Lemma 8] Let K and I' be two left $\mathcal{O}$-ideals of norm ℓ^f not contained in $\ell\mathcal{O}$ and an let ε be an isomorphism $\mathcal{O} \xrightarrow{\sim} \mathrm{End}(E)$. Let $\theta \in \mathcal{O} \setminus (\mathbb{Z} + K + \ell\mathcal{O})$ have norm coprime with ℓ. Assume that $E[K]$ is generated by $P \in E[\ell^f]$. Then, for all $C, D \in \mathbb{Z}$, the following statements are equivalent:*

(i) $E[I'] = \langle [C]P + [D]\varepsilon(\theta)(P) \rangle$.

(ii) $\gcd(C, D, \ell) = 1$ *and there exists* $\alpha \in I$ *such that* $I = \mathcal{O} \cdot \alpha + \mathcal{O}\ell^f$ *and* $\alpha \cdot (C + D\theta) \in K$.

There exists an algorithm [DLW22, Algorithm 3], that returns $\theta \in \mathcal{O} \setminus (\mathbb{Z} + K + \ell\mathcal{O})$ of norm dividing T^2 when given a powersmooth integer $T = \Theta(p^{5/4})$ and a left $\mathcal{O}$-ideal K of norm ℓ^f not contained in $\ell\mathcal{O}$. Combining this algorithm with the above lemma, [DLW22] obtained Algorithm 2.1 that solves Problem 2.1.4 in polynomial time (in $\log(p)$).

2.1.4 On the practical efficiency of KLPT based techniques

In SQIsign [DFKLPW20; DLW22], to make Algorithm 2.1 efficient, the $\ell^f T$-torsion is required to be defined over $\mathbb{F}_{p^2}$ (for either the elliptic curve or its quadratic twist). For that reason, the prime p is selected so that $\ell^f | p+1$ and $\ell^f T | p^2 - 1$. Primes of this form can be found when $\ell^f T = O(p^{3/2})$ *i.e.* when $\ell^f = O(p^{1/4})$ but are non-trivial to find. Looking for such primes has motivated intensive research efforts [Cos20; CMN21; BCRSC+23; Ste23; SEMRH24].

Besides, this efficiency compromise bears a non-negligible cost. Indeed, on the one hand, with $\ell^f = O(p^{1/4})$ and an ideal I of norm $\mathrm{nrd}(I) = O(p^{15/4})$ to translate, we need to cut I into $n = 15$ smaller ideals $I_1, \cdots, I_n$ and to apply Algorithm 2.1 $n = 15$ times. On the other hand, with the hard to satisfy requirement $\ell^f T | p^2 - 1$, the powersmooth number T still has large factors, and consequently,

Algorithm 2.1: Ideal of small norm to isogeny [DLW22, Algorithm 4].

Data: A powersmooth integer $T = \Theta(p^{5/4})$ coprime with ℓ, a left $\mathcal{O}_0$-ideal J of norm ℓ^g (with $g \geq f$), its associated isogeny $\varphi_J : E_0 \longrightarrow E$ and a left $\mathcal{O}$-ideal I' of norm ℓ^f, where $\mathcal{O} := O_R(J) \simeq \mathrm{End}(E)$.

Result: The isogeny $\varphi_{I'} : E \longrightarrow E'$ associated to I'.

$K \leftarrow \overline{J} + \mathcal{O}\ell^f$;
Compute $\theta \in \mathcal{O} \setminus (\mathbb{Z} + K + \ell\mathcal{O})$ of norm dividing T^2 using [DLW22, Algorithm 3];
Select $\alpha \in I$ such that $I = \mathcal{O}\alpha + \mathcal{O}\ell^f$;
Compute $C, D \in \mathbb{Z}$ such that $\alpha \cdot (C + D\theta) \in K$ and $\gcd(C, D, \ell) = 1$;
Factor $\mathrm{nrd}(\theta) = N_1 N_2$ with $N_1 | T$ and $N_2 | T$;
$H_1 \leftarrow \mathcal{O}\theta + \mathcal{O}N_1$ and $H_2 \leftarrow \mathcal{O}\overline{\theta} + \mathcal{O}N_2$;
for $i \in \{1, 2\}$ **do**
    Compute $L_i \leftarrow [J]^* H_i$;
    Compute ρ_i of kernel $\varphi_J(E_0[L_i])$;
end
$Q \leftarrow \widehat{\rho}_2 \circ \rho_1(P)$;
Compute $\varphi_{I'}$ of kernel $\langle [C]P + [D]Q \rangle$;
return $\varphi_{I'}$;

isogenies ρ_1 and ρ_2 from Algorithm 2.1 are costly to compute. As p increases, these primes factors of T naturally get bigger, which also explains that SQIsign does not scale well to higher security parameters.

Alternatives have been proposed to reduce this ideal-to-isogeny translation cost while preserving the algorithmic approach based on KLPT and division into smaller ideals. In [CRSEMR24], the authors work over small field extensions of $\mathbb{F}_{p^2}$ instead of $\mathbb{F}_{p^2}$ to compute the isogenies ρ_1 and ρ_2 from Algorithm 2.1. In [ON25], it has been proposed to use 2-dimensional isogenies to replace the use of the intermediate endomorphism θ of degree coprime to ℓ. This improved the efficiency and scalability of the KLPT based algorithmic approach. However, we shall see that radical changes of approach without KLPT and division into smaller ideals lead to better improvements. As the idea proposed in [ON25], these new approaches rely on higher dimensional isogenies and Kani's lemma that will be presented in the next section.

2.2 Kani's embedding lemma and isogeny interpolation

In this section, we introduce our main tool to improve ideal-to-isogeny translation algorithm. With Kani's lemma, we are able to interpolate isogenies between elliptic curves (even when they have non-smooth degrees) by embedding them into higher dimensional isogenies of smooth degrees. This tool has first been introduced in key recovery attacks against SIDH [CD23; MMPPW23; Rob23], before being used constructively for our application and several other cryptographic applications.

2.2.1 Kani's embedding lemma

In this section, we present Kani's lemma that has been first used in a proof by Ernst Kani [Kan97, Theorem 2.3] in dimension 2 and later generalised to any dimension [Rob23, Lemma 3.6]. We first need to introduce more notions and preliminary results on isogenies between principally polarised abelian varieties. In the following, we fix a field k.

Polarised degree

Definition 2.2.1. Let (A, λ_A) and (B, λ_B) be principally polarised abelian varieties over k and $\varphi : A \longrightarrow B$ be an isogeny.

(i) We define the *polarised dual* of φ by $\widetilde{\varphi} := \lambda_A^{-1} \circ \widehat{\varphi} \circ \lambda_B : B \longrightarrow A$.

(ii) If $d \in \mathbb{N}^*$, we say that φ is a *d-isogeny* is $\widetilde{\varphi} \circ \varphi = [d]_A$, or equivalently if φ is a polarised isogeny $(A, [d]\lambda_A) \longrightarrow (B, \lambda_B)$ in the sense of Definition 1.4.58. We also say that φ has *polarised degree d*.

The notion of d-isogeny generalises the notion of d-isogeny between elliptic curves. Between elliptic curves, polarised dual and dual isogenies can be identified because principal polarisations are canonically determined (see Example 1.4.57). Between elliptic curves, every isogeny of degree d is a d-isogeny. However, not all isogenies are d-isogenies for some $d \in \mathbb{N}^*$ between abelian varieties.

Lemma 2.2.2. *Let (A, λ_A) and (B, λ_B) be principally polarised abelian varieties of dimension g over k and $\varphi : A \longrightarrow B$ be an isogeny. Then:*

(i) $\widetilde{\widetilde{\varphi}} = \varphi$.

(ii) If φ is a d-isogeny, then $\varphi \circ \widetilde{\varphi} = [d]_B$, so $\widetilde{\varphi}$ is a d-isogeny.

(iii) If φ is a d-isogeny, then $\deg(\varphi) = \deg(\widetilde{\varphi}) = d^g$.

Proof. **(i)** We have by definition and by Proposition 1.4.55.(i):

$$\widetilde{\widetilde{\varphi}} = \lambda_B^{-1} \circ \widehat{\widetilde{\varphi}} \circ \lambda_A = \lambda_B^{-1} \circ \widehat{\lambda_B} \circ \widehat{\widehat{\varphi}} \circ \widehat{\lambda_A^{-1}} \circ \lambda_A$$

By Proposition 1.4.54.(ii), polarisation and their dual can be identified $\lambda_B = \widehat{\lambda_B}$. Besides, the dual of the identity being the identity, we obtain by Proposition 1.4.55.(i) that the dual of the inverse is the inverse of the dual, so that $\widehat{\lambda_A^{-1}} = \widehat{\lambda_A}^{-1} = \lambda_A^{-1}$. Finally, $\widehat{\widehat{\varphi}} = \varphi$ by Proposition 1.4.55.(ii), so that $\widetilde{\widetilde{\varphi}} = \varphi$.

(ii) If φ is a d-isogeny, then $\widetilde{\varphi} \circ \varphi = [d]_A$ so $\varphi \circ \widetilde{\varphi} \circ \varphi = \varphi \circ [d]_A = [d]_B \varphi$ and $\varphi \circ \widetilde{\varphi} = [d]_B$ since φ is surjective (as any isogeny).

(iii) We have $\deg(\widetilde{\varphi}) = \deg(\lambda_A^{-1}) \deg(\widehat{\varphi}) \deg(\lambda_B) = \deg(\widehat{\varphi}) = \deg(\varphi)$ by Proposition 1.4.53. Besides, $\deg([d]_A) = d^{2g}$ by Proposition 1.4.33.(ii), so if φ is a d-isogeny, the equality $\widetilde{\varphi} \circ \varphi = [d]_A$ ensures that $\deg(\varphi)^2 = d^{2g}$ so that $\deg(\varphi) = d^g$. This completes the proof. □

Product of polarised abelian varieties

Kani's lemma involves products of polarised abelian varieties. In this paragraph, we give some properties of these objects.

Lemma 2.2.3. *Let A and B be two abelian varieties defined over k. Then:*

(i) $\widehat{A \times B} \cong \widehat{A} \times \widehat{B}$ *with the canonical identification* $([\mathcal{L}], [\mathcal{M}]) \in \mathrm{Pic}^0(A_{\overline{k}}) \times \mathrm{Pic}^0(B_{\overline{k}}) \longmapsto [p_1^*\mathcal{L} \otimes p_2^*\mathcal{M}] \in \mathrm{Pic}^0(A_{\overline{k}} \times B_{\overline{k}})$, *where $p_1 : A \times B \longrightarrow A$ and $p_2 : A \times B \longrightarrow B$ are the projection maps.*

(ii) Let λ_A and λ_B be polarisations on A and B respectively. Then $\lambda_A \times \lambda_B : A \times B \longrightarrow \widehat{A} \times \widehat{B}$ is a polarisation on $A \times B$ called the product polarisation *of (A, λ_A) and (B, λ_B).*

Proof. **(i)** If $\mathcal{L}$ and $\mathcal{M}$ are line bundles of $\mathrm{Pic}^0(A_{\overline{k}})$ and $\mathrm{Pic}^0(B_{\overline{k}})$ respectively, then $[-1]^*\mathcal{L} \simeq \mathcal{L}^{-1}$ and $[-1]^*\mathcal{M} \simeq \mathcal{M}^{-1}$ by Lemma 1.4.65, so we easily verify that $[-1]^*(p_1^*\mathcal{L} \otimes p_2^*\mathcal{M}) \simeq (p_1^*\mathcal{L} \otimes p_2^*\mathcal{M})^{-1}$, so that $p_1^*\mathcal{L} \otimes p_2^*\mathcal{M} \in \mathrm{Pic}^0(A_{\overline{k}} \times B_{\overline{k}})$.

Conversely, let $\mathcal{N} \in \mathrm{Pic}^0(A_{\overline{k}} \times B_{\overline{k}})$. Consider $\mathcal{L} := \mathcal{N}_{|A \times \{0_B\}}$, $\mathcal{M} := \mathcal{N}_{|\{0_A\} \times B}$ and $\mathcal{N}' := p_1^*\mathcal{M} \otimes p_2^*\mathcal{M}$. We prove that $\mathcal{N} \simeq \mathcal{N}'$ with the seesaw principle (Theorem 1.4.14). Indeed, if $y \in B$, then we may see y as a map $T := \mathrm{Spec}(\kappa(y)) \longrightarrow B$, and 0_A and 0_B as maps $\mathrm{Spec}(k) \longrightarrow A$ and $\mathrm{Spec}(k) \longrightarrow B$ respectively, so that

$$\begin{aligned}
\mathcal{N}'_{|A \times \{y\}} &= (\mathrm{id}_A \times y)^*\mathcal{N}' = (\mathrm{id}_A \times y)^*((p_1 \times \mathrm{id}_T)^*(\mathrm{id}_A \times 0_B)^*\mathcal{N} \otimes (\mathrm{id}_T \times p_2)^*(0_A \times \mathrm{id}_B)^*\mathcal{N}) \\
&= ((\mathrm{id}_A \times 0_B) \circ (p_1 \times \mathrm{id}_T) \circ (\mathrm{id}_A \times y))^*\mathcal{N} \otimes ((0_A \times \mathrm{id}_B) \circ (\mathrm{id}_T \times p_2) \circ (\mathrm{id}_A \times y))^*\mathcal{N} \\
&= (\mathrm{id}_A \times 0_B)^*\mathcal{N} \otimes ([0]_A \times y)^*\mathcal{N} = ((\mathrm{id}_A \times 0_B), ([0]_A \times y))^*(p^*\mathcal{N} \otimes q^*\mathcal{N}),
\end{aligned}$$

where $[0]_A : A \longrightarrow A$ is the zero morphism (to be distinguished from $0_A : \mathrm{Spec}(k) \longrightarrow A$) and p, q are the projections from $(A \times B)^2$ to the first and second component respectively. By Lemma 1.4.65, we know that $p^*\mathcal{N} \otimes q^*\mathcal{N} \simeq m^*\mathcal{N}$, where $m : (A \times B)^2 \longrightarrow A \times B$ is the addition map. It follows that

$$\mathcal{N}'_{|A\times\{y\}} \simeq ((\mathrm{id}_A \times 0_B), ([0]_A \times y))^* m^*\mathcal{N} = (\mathrm{id}_A \times y)^*\mathcal{N} = \mathcal{N}_{|A\times\{y\}}.$$

Similarly, we easily obtain that $\mathcal{N}'_{|\{0_A\}\times B} \simeq \mathcal{N}_{|\{0_A\}\times B}$. Then, the seesaw principle (Theorem 1.4.14) applies and ensures that $\mathcal{N} \simeq \mathcal{N}'$. This proves (i).

(ii) By (i), we can canonically identify $\widehat{A \times B}$ with $\widehat{A} \times \widehat{B}$ so $\lambda_A \times \lambda_B$ defines an isogeny $A \times B \longrightarrow \widehat{A \times B}$. Besides, there exists line bundles $\mathcal{L}$ and $\mathcal{M}$ over $A_{\overline{k}}$ and $B_{\overline{k}}$ respectively, such that $(\lambda_A)_{\overline{k}} = \varphi_{\mathcal{L}}$ and $(\lambda_B)_{\overline{k}} = \varphi_{\mathcal{M}}$. We can then consider $\mathcal{N} := p_1^*\mathcal{L} \otimes p_2^*\mathcal{M}$ and we have for all $(x, y) \in A \times B$,

$$\begin{aligned}\lambda_A \times \lambda_B(x, y) &= (\lambda_A(x), \lambda_B(y)) = [p_1^*(t_x^*\mathcal{L} \otimes \mathcal{L}^{-1}) \otimes p_2^*(t_y^*\mathcal{M} \otimes \mathcal{M}^{-1})] \\ &= [t_{(x,y)}^*(p_1^*\mathcal{L} \otimes p_2^*\mathcal{M}) \otimes (p_1^*\mathcal{L} \otimes p_2^*\mathcal{M})^{-1}] = \varphi_{\mathcal{N}}(x, y),\end{aligned}$$

so that $(\lambda_A \times \lambda_B)_{\overline{k}} = \varphi_{\mathcal{N}}$ and $\lambda_A \times \lambda_B$ is a polarisation on $A \times B$. This completes the proof. □

Let $F : A_1 \times A_2 \longrightarrow B_1 \times B_2$ be a morphism between products of abelian varieties. Then, F can be written as a matrix

$$F := \begin{pmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \end{pmatrix},$$

where $f_{ij} := \pi_i \circ F \circ \iota_j$, $\pi_i : B_1 \times B_2 \longrightarrow B_i$ being the projection and $\iota_j : A_i \longrightarrow A_1 \times A_2$ being the injection for all $i, j \in \{1, 2\}$. Naturally, F acts on points of $A_1 \times A_2$ by matrix multiplication on the left:

$$F(x_1, x_2) = (f_{11}(x_1) + f_{12}(x_2), f_{21}(x_1) + f_{22}(x_2)).$$

Lemma 2.2.4. *Let $F : A_1 \times A_2 \longrightarrow B_1 \times B_2$ be a morphism between products of abelian varieties as above. Then*

(i) The dual of F is

$$\widehat{F} = \begin{pmatrix} \widehat{f}_{11} & \widehat{f}_{21} \\ \widehat{f}_{12} & \widehat{f}_{22} \end{pmatrix} : \widehat{B}_1 \times \widehat{B}_2 \longrightarrow \widehat{A}_1 \times \widehat{A}_2.$$

(ii) If the (A_i, λ_{A_i}) and (B_i, λ_{B_i}) are principally polarised for $i \in \{1, 2\}$, then the polarised dual of F with respect to the principal polarisations $\lambda_{A_1} \times \lambda_{A_2}$ and $\lambda_{B_1} \times \lambda_{B_2}$ is

$$\widetilde{F} = \begin{pmatrix} \widetilde{f}_{11} & \widetilde{f}_{21} \\ \widetilde{f}_{12} & \widetilde{f}_{22} \end{pmatrix}$$

Proof. **(i)** Let $(y_1, y_2) \in \widehat{B}_1 \times \widehat{B}_2$ that we identify with $([\mathcal{M}_1], [\mathcal{M}_2]) \in \mathrm{Pic}^0((B_1)_{\overline{k}}) \times \mathrm{Pic}^0((B_2)_{\overline{k}})$, so that

$$\begin{aligned}\widehat{F}(y_1, y_2) &= [F^*(\pi_1^*\mathcal{M}_1 \otimes \pi_2^*\mathcal{M}_2)] = [(\pi_1 \circ F)^*\mathcal{M}_1 \otimes (\pi_2 \circ F)^*\mathcal{M}_2] \\ &= [(m_1 \circ (f_{11} \times f_{12}))^*\mathcal{M}_1 \otimes (m_2 \circ (f_{21} \times f_{22}))^*\mathcal{M}_2] \quad \text{(where } m_i \text{ is the addition map on } B_i\text{)} \\ &= [(f_{11} \times f_{12})^* m_1^*\mathcal{M}_1 \otimes (f_{21} \times f_{22})^* m_2^*\mathcal{M}_2] \\ &= [(f_{11} \times f_{12})^*(p_1^*\mathcal{M}_1 \otimes q_1^*\mathcal{M}_1) \otimes (f_{21} \times f_{22})^*(p_2^*\mathcal{M}_2 \otimes q_2^*\mathcal{M}_2)] \\ &\quad \text{(by Lemma 1.4.65 and with } p_i \text{ and } q_i \text{ the projection on the first and second} \\ &\quad \text{components of } B_i^2 \text{ respectively for all } i \in \{1, 2\}) \\ &= [\pi'^*_1 f_{11}^*\mathcal{M}_1 \otimes \pi'^*_2 f_{12}^*\mathcal{M}_1 \otimes \pi'^*_1 f_{21}^*\mathcal{M}_2 \otimes \pi'^*_2 f_{22}^*\mathcal{M}_2] \\ &\quad \text{(where } \pi'_i \text{ is the projection } A_1 \times A_2 \longrightarrow A_i \text{ for all } i \in \{1, 2\}) \\ &= [\pi'^*_1(f_{11}^*\mathcal{M}_1 \otimes f_{21}^*\mathcal{M}_2) \otimes \pi'^*_2(f_{12}^*\mathcal{M}_1 \otimes f_{22}^*\mathcal{M}_2)] \\ &= (\widehat{f}_{11}(y_1) + \widehat{f}_{21}(y_2), \widehat{f}_{12}(y_1) + \widehat{f}_{22}(y_2)).\end{aligned}$$

This proves (i).

(ii) By point (i) and Lemma 2.2.3.(ii), we have

$$\widetilde{F} = (\lambda_{A_1} \times \lambda_{A_2})^{-1} \circ \widehat{F} \circ (\lambda_{B_1} \times \lambda_{B_2}) = \begin{pmatrix} \lambda_{A_1}^{-1} & 0 \\ 0 & \lambda_{A_2}^{-1} \end{pmatrix} \circ \begin{pmatrix} \widehat{f}_{11} & \widehat{f}_{21} \\ \widehat{f}_{12} & \widehat{f}_{22} \end{pmatrix} \circ \begin{pmatrix} \lambda_{B_1} & 0 \\ 0 & \lambda_{B_2} \end{pmatrix}$$
$$= \begin{pmatrix} \lambda_{A_1}^{-1} \circ \widehat{f}_{11} \circ \lambda_{B_1} & \lambda_{A_1}^{-1} \circ \widehat{f}_{21} \circ \lambda_{B_2} \\ \lambda_{A_2}^{-1} \circ \widehat{f}_{12} \circ \lambda_{B_1} & \lambda_{A_2}^{-1} \circ \widehat{f}_{22} \circ \lambda_{B_2} \end{pmatrix} = \begin{pmatrix} \widetilde{f}_{11} & \widetilde{f}_{21} \\ \widetilde{f}_{12} & \widetilde{f}_{22} \end{pmatrix}.$$

This completes the proof. □

Kani's lemma and its converse

We are now ready to introduce Kani's lemma.

Definition 2.2.5 (Isogeny diamond)**.** Let $a, b \in \mathbb{N}^*$. An (a, b)-*isogeny diamond* is a commutative diagram of isogenies between principally polarized abelian varieties over k

$$\begin{array}{ccc} A' & \xrightarrow{\varphi'} & B' \\ {\scriptstyle\psi}\uparrow & & \uparrow{\scriptstyle\psi'} \\ A & \xrightarrow{\varphi} & B \end{array}$$

where φ and φ' are a-isogenies and ψ and ψ' are b-isogenies.

Lemma 2.2.6 (Kani)**.** *We consider an (a, b)-isogeny diamond over k, as above, with $d := a+b$ prime to the characteristic of the base field of abelian varieties. Then, the isogeny $F : A \times B' \longrightarrow B \times A'$ given in matrix notation by*

$$F := \begin{pmatrix} \varphi & \widetilde{\psi'} \\ -\psi & \widetilde{\varphi'} \end{pmatrix}$$

is a d-isogeny with $d = a + b$, for the product polarisations.

If a and b are coprime and d is not divisible by $\mathrm{char}(k)$*, the kernel of F is*

$$\ker(F) = \{(\widetilde{\varphi}(x), \psi'(x)) \mid x \in B[d]\} = \{([a]x, \psi' \circ \varphi(x)) \mid x \in A[d]\}.$$

Proof. By Lemma 2.2.4, we have

$$\tilde{F} = \begin{pmatrix} \tilde{\varphi} & -\widetilde{\psi} \\ \psi' & \varphi' \end{pmatrix}.$$

Hence

$$\tilde{F} \circ F = \begin{pmatrix} \widetilde{\varphi} \circ \varphi + \widetilde{\psi} \circ \psi & \widetilde{\varphi} \circ \widetilde{\psi'} - \widetilde{\psi} \circ \widetilde{\varphi'} \\ \psi' \circ \varphi - \varphi' \circ \psi & \psi' \circ \widetilde{\psi'} + \varphi' \circ \widetilde{\varphi'} \end{pmatrix} \tag{2.1}$$

with $\widetilde{\varphi} \circ \varphi + \widetilde{\psi} \circ \psi = [a]_A + [b]_A = [d]_A$ since φ is an a-isogeny and ψ is a b-isogeny. Similarly, we get that $\psi' \circ \widetilde{\psi'} + \varphi' \circ \widetilde{\varphi'} = [d]_{B'}$ after dualising (the dual being anti-commutative and the dual of an integer being an integer). Clearly, $\psi' \circ \varphi - \varphi' \circ \psi = 0$ since φ, φ', ψ and ψ' form an isogeny diamond. Then, we obtain $\widetilde{\varphi} \circ \widetilde{\psi'} - \widetilde{\psi} \circ \widetilde{\varphi'} = 0$ by applying the polarised dual to the previous equality (the polarised dual being anti-commutative, as the dual). Hence, $\tilde{F} \circ F = [d]_{A \times B'}$ so F is a d-isogeny.

If $x \in B[d]$, we have

$$\begin{aligned} F(\widetilde{\varphi}(x), \psi'(x)) &= (\varphi \circ \widetilde{\varphi}(x) + \widetilde{\psi'} \circ \psi'(x), -\psi \circ \widetilde{\varphi}(x) + \widetilde{\varphi'} \circ \psi'(x)) \\ &= ([a]x + [b]x, 0) = ([d]x, 0) = (0, 0) \end{aligned}$$

where we used the fact that $\psi \circ \widetilde{\varphi} = \widetilde{\varphi'} \circ \psi'$. Indeed, $\psi' \circ \varphi = \varphi' \circ \psi$, which implies that $[a]\psi \circ \widetilde{\varphi} = [a]\widetilde{\varphi'} \circ \psi'$ *i.e.* $\psi \circ \widetilde{\varphi} \circ [a] = \widetilde{\varphi'} \circ \psi' \circ [a]$ after multiplying on the right by $\widetilde{\varphi}$ and on the left by $\widetilde{\varphi'}$, so that $\psi \circ \widetilde{\varphi} = \widetilde{\varphi'} \circ \psi'$ since $[a]$ is surjective (as any isogeny).

It follows that $\ker(F)$ contains the subgroub:

$$S := \{(\widetilde{\varphi}(x), \psi'(x)) \mid x \in B[d]\}.$$

Since $\widetilde{\varphi}$ and ψ' are have polarised degrees a and b respectively, we have $\ker(\widetilde{\varphi}) \subseteq B[a]$ and $\ker(\psi') \subseteq B[b]$. It follows that $\ker(\widetilde{\varphi}) \cap \ker(\psi') = \{0\}$, when a and b are coprime, so that $x \in B \longmapsto (\widetilde{\varphi}(x), \psi'(x))$ is injective and $\#S = \#B[d]$. Since d is not divisible by $\mathrm{char}(k)$, we have $\#B[d] = d^{2g}$ with $g := \dim(B)$ by Corollary 1.4.34. A being isogenous to B and B', we also have $g = \dim(A) = \dim(B')$ and $\deg(F) = d^{\dim(A \times B')} = d^{2g}$ since F is a d-isogeny and by Lemma 2.2.2.(iii), so F is separable and $\#\ker(F) = \deg(F) = d^{2g}$. Hence, $\ker(F) = S$ when a and b are coprime.

When a and b are coprime, a is coprime with d, so $\ker(\varphi) \cap A[d] \subseteq A[a] \cap A[d] = \{0\}$, so φ induces an isomorphism from $A[d]$ to $B[d]$. It follows that

$$\begin{aligned}\ker(F) &= \{(\widetilde{\varphi}(x), \psi'(x)) \mid x \in B[d]\} = \{(\widetilde{\varphi} \circ \varphi(x), \psi' \circ \varphi(x)) \mid x \in A[d]\} \\ &= \{([a]x, \psi' \circ \varphi(x)) \mid x \in A[d]\}.\end{aligned}$$

This completes the proof. □

When a and b are coprime, if we know $\psi' \circ \varphi$ on $A[d]$, then we know $\ker(F)$ and we are able to compute F in polynomial time when d is smooth [LR12; LR15; LR22]. Knowing F, we can evaluate φ everywhere since $F(x, 0) = (\varphi(x), -\psi(x))$. This is why F is an "embedding" of φ and why F can be computed to interpolate φ from the image of some points generating $A[d]$ (as long as we also know an auxiliary b-isogeny ψ'). We say that F is an efficient representation of φ, in the sense of Definition 1.1.22.

Actually, we can also prove a weak converse of Kani's lemma.

Lemma 2.2.7 (Converse of Kani's lemma)**.** *Let $F : A \times B' \longrightarrow B \times A'$ be a d-isogeny (for the product principal polarizations), where d is prime to the characteristic of the base field. Write F as a matrix:*

$$F := \begin{pmatrix} \varphi & \widetilde{\psi'} \\ -\psi & \widetilde{\varphi'} \end{pmatrix}$$

and suppose φ is an a-isogeny. Then $\varphi, \varphi', \psi, \psi'$ form the following (a, b)-isogeny diamond with $b := d - a$:

$$\begin{array}{ccc} A' & \xrightarrow{\varphi'} & B' \\ {\scriptstyle\psi}\uparrow & & \uparrow{\scriptstyle\psi'} \\ A & \xrightarrow{\varphi} & B \end{array}$$

Proof. Since F is a d-isogeny, we have $\widetilde{F} \circ F = [d]$, so we get that $\psi' \circ \varphi = \varphi' \circ \psi$, $\widetilde{\varphi} \circ \varphi + \widetilde{\psi} \circ \psi = [d]$ and $\psi' \circ \widetilde{\psi'} + \varphi' \circ \widetilde{\varphi'} = [d]$ by Eq. (2.1). Since, φ is an a-isogeny, we have $\widetilde{\varphi} \circ \varphi = [a]$, so $\widetilde{\psi} \circ \psi = [b]$ and ψ is a b-isogeny.

We also have $F \circ \widetilde{F} = [d]$ by Lemma 2.2.2.(ii), so that $\varphi \circ \widetilde{\varphi} + \psi' \circ \widetilde{\psi'} = [d]$ and we get that ψ' is a b-isogeny. Then, the equality $\psi' \circ \widetilde{\psi'} + \varphi' \circ \widetilde{\varphi'} = [d]$ ensures that φ' is an a-isogeny. This completes the proof. □

Remark 2.2.8. We obtain the same result if we suppose that any of the isogenies φ', ψ and ψ' is an a or b-isogeny. If we consider isogenies between product of elliptic curves, then φ is always an a-isogeny so we can always apply the converse of Kani's lemma.

Lemma 2.2.9. *We keep the notations and assumptions from Lemma 2.2.6. Then, the polarised dual $\widetilde{F}$ is given by:*

$$\widetilde{F} := \begin{pmatrix} \widetilde{\varphi} & -\widetilde{\psi} \\ \psi' & \varphi' \end{pmatrix} : B \times A' \longrightarrow A \times B'$$

and has kernel:

$$\ker(\widetilde{F}) = \{(\varphi(x), -\psi(x)) \mid x \in A[d]\} = \{([a]x, -\psi \circ \widetilde{\varphi}(x)) \mid x \in B[d]\}.$$

Proof. The expression of $\widetilde{F}$ follows from Lemma 2.2.4. For the kernel computation, we consider the following (a, b)-isogeny diamond:

$$\begin{array}{ccc} B' & \xrightarrow{\widetilde{\varphi'}} & A' \\ {\scriptstyle\psi'}\uparrow & & \uparrow{\scriptstyle\psi} \\ B & \xrightarrow{\widetilde{\varphi}} & A \end{array}$$

We can either prove that the above diagram commutes and is associated to $\widetilde{F}$ manually or by applying Lemma 2.2.7 to $\widetilde{F}$. Applying Lemma 2.2.6 to the above (a, b)-isogeny diamond, we obtain the expression of $\ker(\widetilde{F})$. □

2.2.2 Isogeny interpolation

The main interest of Kani's lemma is to interpolate an isogeny φ whose images by some torsion points is known by embedding it into a higher dimensional isogeny F that can be computed in polynomial time. This higher dimensional isogeny F is an efficient representation of φ in the sense of Definition 1.1.22 as it provides a way to evaluate φ on any point in polynomial time. Since F can be computed in polynomial time from sufficiently many torsion point images via φ, these images themselves are an efficient representation of φ. In the following, we prove this result due to Damien Robert [Rob23, Theorem 1.1] with close but slightly different assumptions. Note that this theorem is purely theoretical and never used in practice for algorithmic applications. Nonetheless, its proof is instructive to understand the rest of this chapter (and beyond).

Theorem 2.2.10. *Let $\varphi : E_1 \longrightarrow E_2$ be an elliptic curve isogeny defined over a finite field $\mathbb{F}_q$ and of degree d coprime with q. Let $N_1, \cdots, N_r$ be pairwise coprime integers, all coprime with d and q. Assume that r and $N_1, \cdots, N_r$ are all bounded by a polynomial function in $\log(d)$ and that $\prod_{i=1}^{r} N_i > d$. For all $i \in [\![1\ ;\ r]\!]$, let (P_i, Q_i) be a basis of $E_1[N_i]$. Then the data made of $(P_i, Q_i)_{1\leq i\leq r}$, $(\varphi(P_i), \varphi(Q_i))_{1\leq i\leq r}$, E_1, E_2, d is an efficient representation of φ and its dual $\widehat{\varphi}$.*

Proof. Let $N := \prod_{i=1}^{r} N_i$. Then Lagrange's four squares theorem [Lag70] ensures the existence of $a_1, \cdots, a_4 \in \mathbb{Z}$ such that such that $a_1^2 + \cdots + a_4^2 + d = N$. Pollack and Treviño's algorithm [PT18, § 4] can provably find $a_1, \cdots, a_4 \in \mathbb{Z}$ in polynomial time in $\log(N)$. Since r and $N_1, \cdots, N_r$ are polynomial in $\log(d)$, $\log(N)$ is polynomial in $\log(d)$ so finding $a_1, \cdots, a_4 \in \mathbb{Z}$ takes polynomial time in $\log(d)$.

Now, for $i \in \{1, 2\}$, consider the following $(a_1^2 + \cdots + a_4^2)$-isogeny:

$$\alpha_i := \begin{pmatrix} a_1 & -a_2 & -a_3 & -a_4 \\ a_2 & a_1 & a_4 & -a_3 \\ a_3 & -a_4 & a_1 & a_2 \\ a_4 & a_3 & -a_2 & a_1 \end{pmatrix} \in \mathrm{End}(E_i^4),$$

and let Φ be the diagonal d-isogeny $\mathrm{Diag}(\varphi, \cdots, \varphi) : E_1^4 \longrightarrow E_2^2$. Then, Φ and the α_i form a $(a_1^2 + \cdots + a_4^2, d)$-isogeny diamond:

$$\begin{array}{ccc} E_2^4 & \xrightarrow{\alpha_2} & E_2^4 \\ {\scriptstyle\Phi}\uparrow & & \uparrow{\scriptstyle\Phi} \\ E_1^4 & \xrightarrow{\alpha_1} & E_1^4 \end{array}$$

Then, by Kani's lemma Lemma 2.2.6, we can consider the N-isogeny:

$$F := \begin{pmatrix} \alpha_1 & \widetilde{\Phi} \\ -\Phi & \widetilde{\alpha}_2 \end{pmatrix} \in \mathrm{End}(E_1^4 \times E_2^4),$$

whose kernel is:

$$\ker(F) = \{(\widetilde{\alpha}_1(x), \Phi(x)) \mid x \in E_1^4[N]\}, \tag{2.2}$$

since d and N are coprime.

The N_i being pairwise coprime, we shall see in Remark 6.3.2 that F can be decomposed as a chain of isogenies $F = F_r \circ \cdots \circ F_1$, where for all $i \in [\![1\ ;\ r]\!]$, $F_i : A_{i-1} \longrightarrow A_i$ is an N_i-isogeny of kernel

$$\ker(F_i) = F_{i-1} \circ \cdots F_1(\ker(F)) = F_{i-1} \circ \cdots F_1(\ker(F)[N_i]). \tag{2.3}$$

In particular, $\ker(F_1) = \ker(F)[N_1] = \ker(F) \cap (E_1^4 \times E_2^4)[N_1]$. Since φ is $\mathbb{F}_q$-rational, it commutes with the q-th Frobenius endomorphism as well as scalar multiplication maps. As a consequence, $E_1[N_1]$ is stable by the q-th Frobenius endomorphism, and by Eq. (2.2), we see that $\ker(F)[N_1]$ is stable by the q-th Frobenius endomorphism. Then Proposition 1.4.42 ensures that F_1 and its codomain A_1 are $\mathbb{F}_q$-rational. It follows that $\ker(F_2) = F_1(\ker(F)[N_2])$ is stable by the q-th Frobenius endomorphism and that F_2 and its codomain A_2 are $\mathbb{F}_q$-rational. By induction, we obtain that the isogeny F_i and its codomain A_i are $\mathbb{F}_q$-rational for all $i \in [\![1\ ;\ r]\!]$.

For all $i \in [\![1\ ;\ r]\!]$, $E_1[N_i]$ is defined over a finite field extension $\mathbb{F}_{q^{k_i}}/\mathbb{F}_q$ of degree $k_i \leq N_i^2$ (a fact that we easily obtain by considering the N_i-th division polynomial of E_1). So k_i is polynomial in $\log(d)$. This ensures in particular that $(\varphi(P_i), \varphi(Q_i))_{1 \leq i \leq r}$ takes a number of bits polynomial in $\log(q)$ and $\log(d)$ to store.

Furthermore, since the F_i are all $\mathbb{F}_q$-rational, $\ker(F_i)$ is defined over $\mathbb{F}_{q^{k_i}}$ for all $i \in [\![1\ ;\ r]\!]$. Hence, by [LR22] given a basis of $\ker(F_i)$ for $i \in [\![1\ ;\ r]\!]$, the computation of F_i costs $O(N_i^8)$ field operations over $\mathbb{F}_{q^{k_i}}$, so in polynomial time in $\log(d)$ and $\log(q)$. Such a basis of $\ker(F_i)$ can be obtained by applying $F_{i-1} \circ \cdots \circ F_1$ to the basis:

$$\mathscr{B}_i := ((\widetilde{\alpha}_1(x), \Phi(x)))_{x \in \{(P_i,0,0,0), \cdots, (0,0,0,P_i), (Q_i,0,0,0), \cdots, (0,0,0,Q_i)\}}$$

of $\ker(F)[N_i]$ by Eq. (2.3). Since $\varphi(P_i), \varphi(Q_i)$ is given by assumption, $\mathscr{B}_i$ can be computed in polynomial time in $\log(d)$ and $\log(q)$. Since $F_1, \cdots, F_{i-1}$ are $\mathbb{F}_q$-rational, evaluating $\mathscr{B}_i$ by $F_{i-1} \circ \cdots \circ F_1$ costs $O(N_1^8 + \cdots + N_{i-1}^8)$ operations over $\mathbb{F}_{q^{k_i}}$ which takes polynomial time in $\log(d)$ and $\log(q)$, since $i \leq r$ is also polynomial in $\log(d)$. As a consequence, the computation of F takes polynomial time in $\log(d)$ and $\log(q)$.

Finally, F is an efficient representation of φ and $\widehat{\varphi}$ as it takes a polynomial number of bits in $\log(d)$ and $\log(q)$ to store as a chain of isogenies (by [LR22]) and since we also have:

$$\begin{aligned} \forall P \in E_1, \quad F(P,0,0,0,0,0,0) &= ([a_1]P, [a_2]P, [a_3]P, [a_4]P, -\varphi(P), 0, 0, 0) \\ \forall Q \in E_2, \quad F(0,0,0,0,0,0,Q) &= (0,0,0,\widehat{\varphi}(Q), [a_4]Q, [a_3]Q, -[a_2]Q, [a_1]Q), \end{aligned}$$

so we can evaluate φ and $\widehat{\varphi}$ on any point in polynomial time. Since $(\varphi(P_i), \varphi(Q_i))_{1 \leq i \leq r}$ has polynomial size in $\log(q)$ and $\log(d)$ and yields F in polynomial time in $\log(q)$ and $\log(d)$, it is also an efficient representation of φ and $\widehat{\varphi}$. □

2.2.3 The SIDH protocol

SIDH (Supersingular Isogeny Diffie Hellman) is an isogeny based key exchange protocol. A key encapsulation protocol SIKE, based on SIDH has been proposed to the NIST post-quantum standardisation competition before being broken in 2022, closely after the beginning of round 4 of the competition. The main vulnerability of SIDH was the publication of torsion point images that could be leveraged by an attacker using Kani's lemma. In this section, we recall how this protocol is built.

In SIDH, Alice and Bob are given as public parameters:

- A prime of the form $p = c\ell_A^{e_A}\ell_B^{e_B} - 1$, with $c \in \mathbb{N}^*$ small, distinct small primes ℓ_A and ℓ_B and exponents $e_A, e_B \in \mathbb{N}^*$ such that $\ell_A^{e_A} \simeq \ell_B^{e_B} \simeq \sqrt{p}$;
- A starting supersingular elliptic curve E_0 defined over $\mathbb{F}_{p^2}$;
- Two basis (P_A, Q_A) and (P_B, Q_B) of $E_0[\ell_A^{e_A}]$ and $E_0[\ell_B^{e_B}]$ respectively.

Alice and Bob sample secret integers $s_A \in \mathbb{Z}/\ell_A^{e_A}\mathbb{Z}$ and $s_B \in \mathbb{Z}/\ell_B^{e_B}\mathbb{Z}$ respectively. Then Alice computes an $\ell_A^{e_A}$-isogeny $\varphi_A : E_0 \longrightarrow E_A$ of kernel $\ker(\varphi_A) = \langle P_A + [s_A]Q_A \rangle$ and Bob computes an $\ell_B^{e_B}$-isogeny $\varphi_B : E_0 \longrightarrow E_B$ of kernel $\ker(\varphi_B) = \langle P_B + [s_B]Q_B \rangle$. Alice sends $(E_A, \varphi_A(P_B), \varphi_A(Q_B))$ to Bob and Bob sends $(E_B, \varphi_B(P_A), \varphi_B(Q_A))$ to Alice. Alice can then compute $\psi_A := [\varphi_B]_* \varphi_A : E_B \longrightarrow$

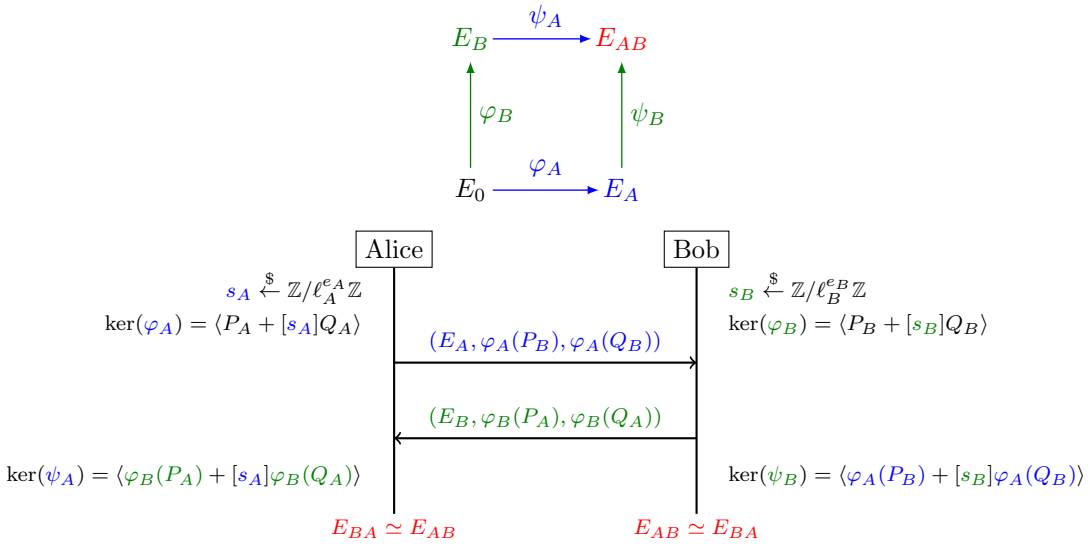

Figure 2.1: Principle of the SIDH key exchange protocol.

E_{BA} of kernel $\ker(\psi_A) = \langle \varphi_B(P_A) + [s_A]\varphi_B(Q_A) \rangle$ and Bob computes $\psi_B := [\varphi_A]_*\varphi_B : E_A \longrightarrow E_{AB}$ of kernel $\ker(\psi_B) = \langle \varphi_A(P_B) + [s_B]\varphi_A(Q_B) \rangle$. Then, they share knowledge of a secret elliptic curve $E_{AB} \simeq E_{BA}$.

The exchange of torsion points $(\varphi_A(P_B), \varphi_A(Q_B))$ and $(\varphi_B(P_A), \varphi_B(Q_A))$, though necessary to compute ψ_A and ψ_B, makes SIDH particularly vulnerable. In the next section, we present attacks that leverage this torsion point information and use Kani's lemma to completely recover one of the isogenies φ_A or φ_B, which is sufficient to recover the shared secret key E_{AB}. Indeed, if an attacker recovers φ_B, then they know $\ker(\varphi_B)$ so Bob's secret s_B and can use it to compute ψ_B whose kernel is $\ker(\psi_B) = \langle \varphi_A(P_B) + [s_B]\varphi_A(Q_B) \rangle$ and codomain is E_{AB}.

2.2.4 Attacks against SIDH

The first attacks introduced against SIDH in [CD23; MMPPW23] involved a 2-dimensional isogeny computation. Though reasonably efficient and practical at that time, these attacks involved several restrictions (*e.g.* on the starting curve E_0). Then Damien Robert proposed to apply Kani's lemma in dimension 4 or 8 to lift these restrictions [Rob23]. Even though no implementation was known at the time of publication, Robert's attack in dimension 8 was proved to run in polynomial time without restriction. Later, my efficient implementation of 4-dimensional 2-isogenies (see Section 6.6) made Robert's 4-dimensional attack very competitive for all SIKE parameters with any starting curve E_0.

Throughout this section, we keep the notations from Section 2.2.3 and assume that we want to recover φ_B from the knowledge of $(\varphi_B(P_A), \varphi_B(Q_A))$.

The first attacks in dimension 2

We present here the attack introduced by Maino, Martindale, Panny, Pope and Wesolowski in [MMPPW23]. We assume that $\ell_A^{e_A} > \ell_B^{e_B}$ and that $\mathrm{End}(E_0)$ is known. Indeed, in SIKE, the starting curve E_0 is either $y^2 = x^3 + x$ or $y^2 = x^3 + 6x^3 + x$ whose endomorphism rings are known. We consider the $(\ell_B^{e_B}, \ell_A^{e_A} - \ell_B^{e_B})$-isogeny diamond

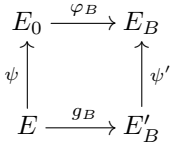

that yields by Kani's lemma (Lemma 2.2.6) a 2-dimensional $\ell_A^{e_A}$-isogeny

$$F := \begin{pmatrix} g_B & \widehat{\psi'} \\ -\psi & \widehat{\varphi}_B \end{pmatrix} : E \times E_B \longrightarrow E'_B \times E_0,$$

with kernel:

$$\ker(F) = \{([\ell_B^{e_B}]P, \varphi_B \circ \psi(P)) \mid P \in E[\ell_A^{e_A}]\} = \{([\ell_B^{e_B}]\widehat{\psi}(P), -[\ell_B^{e_B}]\varphi_B(P)) \mid P \in E_0[\ell_A^{e_A}]\}.$$

Hence, to compute $\ker(F)$, it suffices to find an f-isogeny $\widehat{\psi} : E_0 \longrightarrow E$ with $f := \ell_A^{e_A} - \ell_B^{e_B}$ and evaluate it on $E_0[\ell_A^{e_A}]$. Once this is done, we can compute F in polynomial time (in $\log(p)$) with various algorithms from the literature [KR09; LR12; LR15; CR15; LR22; DMPR25].

The main difficulty is to find this auxiliary f-isogeny $\widehat{\psi} : E_0 \longrightarrow E$. When $\mathrm{End}(E_0)$ is knwon, we can use the Deuring correspondence. We start by finding a left-ideal $I \subseteq \mathrm{End}(E_0)$ of norm f. This can be done with [MMPPW23, Algorithm 3]. Since this ideal is of non-smooth norm, we cannot translate it into an isogeny directly[1]. But we only need to find the codomain E of $\widehat{\psi} := \varphi_I : E_0 \longrightarrow E$ and to evaluate it on $E_0[\ell_A^{e_A}]$ so we can proceed as follows. We find an equivalent ideal $J \sim I$ of norm ℓ_B^e with e big enough via the KLPT algorithm [KLPT14] (in practice, $\ell_B^e \simeq p^3$). We can then translate it into an isogeny $\varphi_J : E_0 \longrightarrow E$ via the KLPT based techniques introduced in Section 2.1. By assumption, we know an isomorphism $\mathcal{O}_0 \xrightarrow{\sim} \mathrm{End}(E_0)$, where $\mathcal{O}_0 \subset\subseteq \mathcal{B}_{p,\infty}$ is a maximal order. Since $J \sim I$, we may write $J = I\overline{\theta}/\operatorname{nrd}(I)$ with $\theta \in I \subseteq \mathcal{O}_0$ and by Lemma 1.2.24, we have:

$$\widehat{\varphi}_J \circ \varphi_I = \varepsilon_0(\theta) \quad i.e. \quad [\ell_B^e]\varphi_I = \varphi_J \circ \varepsilon_0(\theta),$$

so to evaluate $P \in E_0[\ell_A^{e_A}]$, we find an inverse λ of ℓ_B^e modulo $\ell_A^{e_A}$ and compute $\varphi_J \circ \theta([\lambda]P) = [\ell_B^e\lambda]\varphi_I(P) = \varphi_I(P)$.

With this method, we can evaluate $\widehat{\psi} := \varphi_I$, hence compute F in polynomial time (in $\log(p)$) under some plausible heuristics. We can then evaluate $\widehat{\varphi}_B$ as follows:

$$F(0, P) = (\widehat{\psi'}(P), \widehat{\varphi}_B(P)),$$

and in particular, compute $\ker(\varphi_B) = \widehat{\varphi}_B(E_B[\ell_B^{e_B}])$ and find the secret $s_B \in \mathbb{Z}/\ell_B^{e_B}\mathbb{Z}$.

When $\mathrm{End}(E_0)$ is not known, the 2-dimensional attack proposed in [MMPPW23, § 3] is much more complex. We essentially have to tweak parameters to simplify the search for the auxiliary isogeny $\widehat{\psi} : E_0 \longrightarrow E$. We consider the following commutative diagram:

$$\begin{array}{ccccc} E_0 & \xrightarrow{\varphi'_B} & E' & \xrightarrow{\varphi_{\mathsf{guess}}} & E_B \\ {\scriptstyle\psi}\uparrow & & \uparrow{\scriptstyle\psi'} & & \\ E & \xrightarrow{g_B} & E'' & & \end{array} \qquad \text{(with } \varphi_B : E_0 \to E_B \text{ the composite)}$$

where φ_{guess} is a guess of the last j steps of the ℓ_B-isogeny chain φ_B, so that $\deg(\varphi_{\mathsf{guess}}) = \ell_B^j$, $\deg(\varphi'_A) = \ell_B^{e_B - j}$ and $\deg(\psi) := f := e\ell_A^{e_A - i} - \ell_B^{e_B - j}$. We look for small parameters e, i, j so that f is smooth enough to compute $\widehat{\psi}$ efficiently. The look for such parameters and the computation of $\widehat{\psi}$ takes subexponential time (in $\log(p)$). Then, we try several guesses φ_{guess} and apply Kani's lemma to the $(\ell_B^{e_B}, f)$-isogeny diamond on the left part of the diagram above to compute a 2-dimensional $e\ell_A^{e_A - i}$-isogeny embedding φ'_B. If we have chosen the correct guess, we expect the codomain of the 2-dimensional isogeny to split into a product of elliptic curves. Otherwise, we expect it to be a non-product principally polarised abelian surface (with overwhelming probability). Taking into account at most ℓ_B^j tries of φ_{guess}, the attack still takes subexponential time (in $\log(p)$).

These attacks have been successfully implemented in SageMath. The implementation by Castryck, Decru, Pope and Oudompheng[2] only works with a special starting curve of known endomorphism ring and broke SIKE with the biggest parameter (p751) in 1 h. The implementation by Maino, Panny, Pope and Wesolowski[3] works with any starting curve but only for small parameters.

[1] At least with the techniques known at that time, this is now possible as we shall see in Section 2.4.

[2] https://github.com/GiacomoPope/Castryck-Decru-SageMath

[3] https://github.com/Breaking-SIDH/direct-attack

An efficient attack in dimension 4

We first look for $e \in \mathbb{N}^*$ and $a_1, a_2 \in \mathbb{Z}$ such that

$$a_1^2 + a_2^2 + \ell_B^{e_B} = \ell_A^e. \tag{2.4}$$

If such an equation has a solution, we can find it by trying different values of $e \in \mathbb{N}^*$ until $\ell_A^e - \ell_B^{e_B}$ is a prime congruent to 1 mod 4. Then, we can apply Cornacchia's algorithm [Cor08] to find $a_1, a_2 \in \mathbb{Z}$ in polynomial time (in $\log(p)$).

Once we have solved Eq. (2.4), we consider the following $(a_1^2 + a_2^2, \ell_B^{e_B})$-isogeny diamond:

$$\begin{array}{ccc} E_B^2 & \xrightarrow{\alpha_B} & E_B^2 \\ {\scriptstyle \Phi_B}\uparrow & & \uparrow{\scriptstyle \Phi_B} \\ E_0^2 & \xrightarrow{\alpha_0} & E_0^2 \end{array}$$

where, for $i \in \{0, B\}$, α_i is an $(a_1^2 + a_2^2)$-isogeny

$$\alpha_i := \begin{pmatrix} a_1 & a_2 \\ -a_2 & a_1 \end{pmatrix} \in \mathrm{End}(E_i^2),$$

and Φ_B is a 2-dimensional diagonal $\ell_B^{e_B}$-isogeny $\Phi_B := \mathrm{Diag}(\varphi_B, \varphi_B) : E_0^2 \longrightarrow E_B^2$. Then, Kani's lemma (Lemma 2.2.6) ensures that:

$$F := \begin{pmatrix} \alpha_0 & \widetilde{\Phi}_B \\ -\Phi_B & \widetilde{\alpha}_B \end{pmatrix} \in \mathrm{End}(E_0^2 \times E_B^2),$$

is an ℓ_A^e-isogeny with kernel:

$$\begin{aligned} \ker(F) &= \{(\widetilde{\alpha}_0(P,Q), \Phi_B(P,Q)) \mid P, Q \in E_0[\ell_A^e]\} \\ &= \{([a_1]P - [a_2]Q, [a_2]P + [a_1]Q, \varphi_B(P), \varphi_B(Q)) \mid P, Q \in E_0[\ell_A^e]\}. \end{aligned}$$

Consequently, if $e \leq e_A$, $(\varphi_B(P_A), \varphi_B(Q_A))$ yields a basis of $\ker(F)$ that can be used to compute F in polynomial time with the algorithms from [LR12; LR15; LR22].

In general, we cannot expect to obtain a solution of Eq. (2.4) with $e \leq e_A$, so we divide the computation of F in two: $F = F_2 \circ F_1$, where F_i is an $\ell_A^{e_i}$-isogeny for $i \in \{1, 2\}$, $e = e_1 + e_2$ and $e_1 \simeq e_2$. We expect that $e_1, e_2 \leq e_A$, a condition much easier to satisfy. Computing $\widetilde{F}$ with Lemma 2.2.4.(ii), we see by Lemma 2.2.7 that $\widetilde{F}$ is obtained from an isogeny diamond. Applying Kani's lemma again, we obtain that

$$\begin{aligned} \ker(\widetilde{F}) &= \{(\alpha_0(P,Q), -\Phi_B(P,Q)) \mid P, Q \in E_0[\ell_A^e]\} \\ &= \{([a_1]P + [a_2]Q, -[a_2]P + [a_1]Q, -\varphi_B(P), -\varphi_B(Q)) \mid P, Q \in E_0[\ell_A^e]\}. \end{aligned}$$

By Lemma 6.3.1.(ii), we have $\ker(F_1) = \ker(F)[\ell_A^{e_1}]$ and $\ker(\widetilde{F}_2) = \ker(\widetilde{F})[\ell_A^{e_2}]$, so $\ker(F_1)$ and $\ker(\widetilde{F}_2)$ can be computed from the knowledge of $(\varphi_B(P_A), \varphi_B(Q_A))$ since $e_1, e_2 \leq e_A$. Hence, we can compute F_1 and $\widetilde{F}_2$ in polynomial time and then obtain $\widetilde{\widetilde{F}}_2 = F_2$ and $F = F_2 \circ F_1$. We refer to Section 6.4.2 for more details on this decomposition of F. Since F is an embedding of φ_B, we can evaluate φ_B everywhere as follows

$$F(P, 0, 0, 0) = ([a_1]P, -[a_2]P, -\varphi_B(P), 0),$$

so we can recover $\ker(\varphi_B)$ and s_B as previously.

With real NIST submission SIKE parameters, $\ell_A = 2$, $\ell_B = 3$ and $c = 1$ so that $p = 2^{e_A}3^{e_B} - 1$ and $(e_A, e_B) \in \{(216, 137), (250, 159), (305, 192), (372, 239)\}$ depending on the security level. Using our 4-dimensional 2^e-isogeny computation algorithms with level 2 theta coordinates [Dar24] (that will be presented in Section 6.6), we implemented a complete key recovery attack of SIDH running in a few seconds on a laptop for all SIKE parameters from a random starting curve E_0 (see Table 2.1). These very competitive timings are to be compared with the aforementioned ones in dimension 2 and

SIKE prime	e_A	e_B	e	Attack timing (s)
p434	216	137	225	3.82
p503	250	159	290	5.47
p610	305	192	407	8.61
p751	372	239	589	14.02

Table 2.1: Timings (in s) and parameters of the complete SIDH key recovery attack with a random starting curve in Python/SageMath for various NIST SIKE primes on a 2,7 GHz Intel Core i5 CPU.

indicate that 4-dimensional isogenies are efficient enough to be used constructively (as we shall see in Sections 2.3 and 2.5).

Note that Eq. (2.4) does not always admit a solution that can be found easily in polynomial time. In particular, when $\ell_A = 2$ and $\ell_B = 3$, e_B has to be odd. This is not the case for the p610 SIKE parameter ($(e_A, e_B) = (305, 192)$). In that case, we simply post-compose φ_B by a 3-isogeny $\psi : E_B \longrightarrow E'_B$ and apply the attack to $(\psi \circ \varphi_B(P_A), \psi \circ \varphi_B(Q_A))$ in order to change e_B into $e_B + 1$. More generally, the argument that we can solve Eq. (2.4) efficiently in polynomial time (in $\log(p)$) for any set of parameters is heuristic. There is no formal proof of this assumption. For that reason, Damien Robert also proposed an 8-dimensional attack which provably runs in polynomial time (in $\log(p)$).

A theoretical polynomial time attack in dimension 8

We have already seen this idea in the proof of Theorem 2.2.10. If $e \in \mathbb{N}^*$ is such that $\ell_A^e > \ell_B^{e_B}$, then Lagrange's four squares theorem [Lag70] ensures that we can always find $a_1, a_2, a_3, a_4 \in \mathbb{Z}$ such that:

$$a_1^2 + a_2^2 + a_3^2 + a_4^2 + \ell_B^{e_B} = \ell_A^e. \tag{2.5}$$

Pollack and Treviño's algorithm [PT18, § 4] can provably solve this equation in polynomial time (in $\log(p)$).

We can then apply Kani's lemma in dimension 8. For $i \in \{0, B\}$, consider the $(a_1^2 + a_2^2 + a_3^2 + a_4^2)$-isogeny:

$$\alpha_i := \begin{pmatrix} a_1 & -a_2 & -a_3 & -a_4 \\ a_2 & a_1 & a_4 & -a_3 \\ a_3 & -a_4 & a_1 & a_2 \\ a_4 & a_3 & -a_2 & a_1 \end{pmatrix} \in \mathrm{End}(E_i^4).$$

Consider also the 4-dimensional diagonal $\ell_B^{e_B}$-isogeny $\Phi_B := \mathrm{Diag}(\varphi_B, \varphi_B, \varphi_B, \varphi_B) : E_0^4 \longrightarrow E_B^4$. Then, we have an $(a_1^2 + a_2^2 + a_3^2 + a_4^2, \ell_B^{e_B})$-isogeny diamond

$$\begin{array}{ccc} E_B^4 & \xrightarrow{\alpha_B} & E_B^4 \\ {\scriptstyle \Phi_B}\uparrow & & \uparrow{\scriptstyle \Phi_B} \\ E_0^4 & \xrightarrow{\alpha_0} & E_0^4 \end{array}$$

By Kani's lemma, we have an 8-dimensional ℓ_A^e-isogeny:

$$F := \begin{pmatrix} \alpha_0 & \widetilde{\Phi}_B \\ -\Phi_B & \widetilde{\alpha}_B \end{pmatrix} \in \mathrm{End}(E_0^4 \times E_B^4),$$

with kernel:

$$\ker(F) = \{(\widetilde{\alpha}_0(P, Q, R, S), \Phi_B(P, Q, R, S)) \mid P, Q, R, S \in E_0[\ell_A^e]\}.$$

As previously, if $e \leq e_A$, we can compute a basis of $\ker(F)$ from $(\varphi_B(P_A), \varphi_B(Q_A))$ and then compute F in polynomial time (in $\log(p)$) using the algorithms from [LR12; LR15; LR22]. Otherwise, if $e/2 \leq e_A$, we can compute decompose the computation of F into $F = F_2 \circ F_1$, as previously. This approach takes polynomial time by [Rob23, Theorem 1.1] but has not been implemented and is purely theoretical.

2.2.5 Higher dimensional isogeny computation algorithms

In the following, we shall use constructively the isogeny interpolation techniques introduced previously for the SIDH attacks. These techniques rely on the efficient computation of higher dimensional isogenies. For efficiency, reasons we shall only use 2^e-isogenies in higher dimension. With the algorithms we shall introduce in Chapter 6 using level 2 theta coordinates, we are able to compute such a 2^e-isogeny from torsion points lying above its kernel and forming an isotropic subgroup of the 2^{e+2}-torsion, in the sense of the following definition. Hence, in the rest of this chapter, we shall use freely the following theorems that will be proved in Chapter 6.

Definition 2.2.11. If A is an abelian variety defined over a field k, $n \in \mathbb{N}^*$ is coprime with $\mathrm{char}(k)$, we say that a subgroup $K \subset A[n](\overline{k})$ is *isotropic* in $A[n]$ if the n-th Weil pairing is trivial on K: for all $x, y \in K$, $e_n(x, y) = 1$. We say K is *maximal isotropic* if it is maximal (as a subgroup of $A[n]$) for this property.

Theorem 2.2.12 (Section 6.3)**.** *Let k be a field of characteristic $\mathrm{char}(k) \neq 2$. Then there exists an algorithm that takes as input:*

(i) A principally polarised abelian variety A of dimension g defined over k;

(ii) Points $T_1, \cdots, T_g \in A[2^{e+2}]$ defined over k forming a maximal isotropic subgroup of $A[2^{e+2}]$;

And returns a 2^e-isogeny $f : A \longrightarrow B$ with kernel $\langle [4]T_1, \cdots, [4]T_g \rangle$ represented as a chain of 2-isogenies with a number of operations over k polynomial in e and 2^g.

In dimension 2, the 2^{e+2}-torsion requirement to compute a 2^e-isogeny can be relaxed, at the expense of 5 square root computations.

Theorem 2.2.13 (Section 6.5)**.** *Let k be a field of characteristic $\mathrm{char}(k) \neq 2$. Then there exists an algorithm that takes as input:*

(i) A principally polarised abelian surface A defined over k;

(ii) Points $T_1, T_2 \in A[2^e]$ defined over k forming a maximal isotropic subgroup of $A[2^e]$;

And returns a 2^e-isogeny $f : A \longrightarrow B$ with kernel $\langle T_1, T_2 \rangle$ represented as a chain of 2-isogenies with a number of operations over k polynomial in e.

2.3 Translating ideals of short norm with 4-dimensional isogenies

Recall the algorithmic problem from Section 2.1. We are given:

- A supersingular elliptic curve E_0 over $\mathbb{F}_{p^2}$ of known endomorphism ring $\mathrm{End}(E_0)$ *i.e.* such that we know a maximal order $\mathcal{O}_0 \subseteq \mathcal{B}_{p,\infty}$ and an isomorphism $\varepsilon_0 : \mathcal{O}_0 \xrightarrow{\sim} \mathrm{End}(E_0)$. For instance, E_0 may be the curve $y^2 = x^3 + x$ (with $p \equiv 3 \mod 4$).
- Two isogenies $\varphi_1 : E_0 \longrightarrow E_1$ and $\varphi_2 : E_0 \longrightarrow E_2$ of odd degrees d_1 and d_2 respectively, along with their respective corresponding ideals $I_1, I_2 \subseteq \mathcal{O}_0$ via the Deuring correspondence.
- An ideal $I \subset \mathcal{B}_{p,\infty}$ connecting $\mathcal{O}_1 := O_R(I_1) \simeq \mathrm{End}(E_1)$ and $\mathcal{O}_2 := O_R(I_2) \simeq \mathrm{End}(E_2)$, *i.e.* which is a left $\mathcal{O}_1$-ideal and a right $\mathcal{O}_2$-ideal.

We assume that the characteristic p is such that $2^f | p + 1$ with $2^f = \Omega(\sqrt{p})$, so that the 2^f-torsion of the supersingular elliptic curves E_0, E_1, E_2 is defined over $\mathbb{F}_{p^2}$. We want to compute the isogeny $\varphi_I : E_1 \longrightarrow E_2$ corresponding to I via the Deuring correspondence.

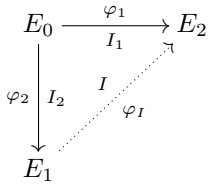

The KLPT based algorithms presented in Section 2.1 were costly because $\mathrm{nrd}(I)$ was required to be smooth, so I was expected to be an output of the KLPT algorithm [KLPT14] which returns ideals of very big norms. In the following, we no longer require I to be smooth but to be close to the smallest possible value $\mathrm{nrd}(I) \simeq \sqrt{p}$ and use Kani's lemma to embed φ_I in dimension 4. There are still restrictions on $\mathrm{nrd}(I)$ even if they are less constraining than previously.

2.3.1 Conditions on the ideal norm

We assume that we are given an ideal of norm $q := \mathrm{nrd}(I)$ that is (e, B)-*good* in the following sense.

Definition 2.3.1. Given $e, B \in \mathbb{N}^*$, we say that an integer q is (e, B)-*good* if $2^e - q = r \cdot s_1 \cdot s_3^2$, where r is a prime $\equiv 1 \mod 4$, s_1 is a B-smooth integer with prime factors $\equiv 1 \mod 4$ and s_3 is a B-smooth integer with prime factors $\equiv 3 \mod 4$. In particular, an (e, B)-good integer must be odd.

If q is (e, B)-good with B small enough ($B = 10^2$ to 10^3 in practice), then $2^e - q$ can be easily decomposed into a sum of two squares, so that:

$$a_1^2 + a_2^2 + q = 2^e, \tag{2.6}$$

with $a_1, a_2 \in \mathbb{Z}$. Indeed, we can find the factorisation $2^e - q = r \cdot s_1 \cdot s_3^2$ easily by trial and division of all primes $\leq B$ and testing that the quotient r is a prime $\equiv 1 \mod 4$. Then, for every prime $\ell | 2^e - q$ congruent to $1 \mod 4$ we easily decompose $\ell = a_\ell^2 + b_\ell^2$ via Cornacchia's algorithm [Cor08] and write

$$\prod_{\substack{\ell | 2^e - q \\ \ell \equiv 1 \mod 4}} (a_\ell + i b_\ell) = u_1 + i u_2,$$

with $u_1, u_2 \in \mathbb{Z}$ and $i = \sqrt{-1}$ in $\mathbb{C}$. We then return $a_1 := s_3 u_1$ and $a_2 := s_3 u_2$.

In the following, we fix $e \in \mathbb{N}^*$ (and $B \in \mathbb{N}^*$) such that there always exists a connecting ideal I between $\mathcal{O}_1$ and $\mathcal{O}_2$ of (e, B)-good norm. The probability for an integer q to be a prime congruent to $1 \mod 4$ is $\Omega(1/\log(q))$ so the probability for an integer to be (e, B)-good is $\Omega(1/e)$. Besides, the following lemma ensures that we can find ideal connecting ideals between $\mathcal{O}_1$ and $\mathcal{O}_2$ of norms bounded by $O(\sqrt{p})$. Hence, heuristically, we can expect to find ideals connecting $\mathcal{O}_1$ and $\mathcal{O}_2$ of (e, B)-good norm bounded by $O(e\sqrt{p})$, so we may select e such that $2^e = \Theta(\log(p)\sqrt{p})$.

Lemma 2.3.2. *Let $\mathcal{O} \subseteq \mathcal{B}_{p,\infty}$ be a maximal order and I be a left $\mathcal{O}$-ideal. Then there exists an equivalent ideal $J \sim I$ of norm $\leq 2\sqrt{2p}/\pi$.*

Proof. Consider the quadratic form $q_I : \alpha \in I \longmapsto \mathrm{nrd}(\alpha)/\mathrm{nrd}(I) \in \mathbb{N}$. We know that equivalent ideals $J \sim I$ are of the form $J = I\overline{\alpha}/\mathrm{nrd}(I)$ with $\alpha \in I$ and have norm $\mathrm{nrd}(J) = q_I(\alpha)$. Hence, we have to find an upper bound on the non-zero minimal value of q_I. Consider the canonical embedding $\iota : \mathcal{B}_{p,\infty} \hookrightarrow \mathbb{R}^4$ from Eq. (1.2):

$$1 \longmapsto (1,0,0,0), i \longmapsto (0,1,0,0), j \longmapsto (0,0,\sqrt{p},0), k \longmapsto (0,0,0,\sqrt{p}). \tag{2.7}$$

ι is an isometry in the following sense $\|\iota(\alpha)\|^2 = \mathrm{nrd}(\alpha)$ for all $\alpha \in \mathcal{B}_{p,\infty}$, where $\|.\|$ is the Euclidean norm of $\mathbb{R}^4$. By Minkowski's second theorem, the successive minima λ_i of the lattice $\iota(I)$ satisfy

$$\lambda_1 \cdots \lambda_4 \leq 2^4 \frac{\mathrm{Covol}(\iota(I))}{\mathrm{Vol}(B(0,1))} = \frac{32}{\pi^2} \mathrm{Covol}(\iota(I)). \tag{2.8}$$

But we have $\mathrm{Covol}(\iota(\mathcal{O})) = \mathrm{discrd}(\mathcal{O})/4$. Indeed, if $(\alpha_1, \cdots, \alpha_4)$ is a basis of $\mathcal{O}$, we get that

$$\begin{aligned}
\mathrm{Covol}(\iota(\mathcal{O})) &= |\det(\langle \iota(\alpha_i), \iota(\alpha_j) \rangle)_{1 \leq i,k \leq 4}|^{\frac{1}{2}} \\
&= \left| \det \left(\frac{1}{2}(\mathrm{nrd}(\alpha_i + \alpha_j) - \mathrm{nrd}(\alpha_i) - \mathrm{nrd}(\alpha_j)) \right)_{1 \leq i,k \leq 4} \right|^{\frac{1}{2}} \\
&= \left| \det \left(\frac{1}{2} \mathrm{Tr}(\alpha_i \overline{\alpha_j}) \right)_{1 \leq i,k \leq 4} \right|^{\frac{1}{2}} = \frac{1}{4} \mathrm{discrd}(\mathcal{O})
\end{aligned}$$

Besides, by Theorem 1.2.17, $\mathrm{discrd}(\mathcal{O}) = \mathrm{disc}(\mathcal{B}_{p,\infty}) = p$ since $\mathcal{O}$ is a maximal order in $\mathcal{B}_{p,\infty}$. We then have

$$\mathrm{Covol}(\iota(I)) = [\mathcal{O} : I]\,\mathrm{Covol}(\iota(\mathcal{O})) = [\mathcal{O} : I]\,\mathrm{discrd}(\mathcal{O})/4 = \mathrm{nrd}(I)^2 p/4 \tag{2.9}$$

It follows by Eq. (2.8) and Eq. (2.9) that the minimal value of q_I is

$$q_I(\alpha) = \frac{\lambda_1^2}{\mathrm{nrd}(I)} \leq \frac{2\sqrt{2p}}{\pi}.$$

This completes the proof. □

Remark 2.3.3. Selecting e such that $2^e = \Theta(\log(p)\sqrt{p})$ only offers a heuristic guarantee. Proving formally that we can always find a connecting ideal I of (e, B)-good norm would certainly require to increase 2^e by a lot. As [RT22] indicates, we should expect lower bounds close to $2^e = \omega(p^2)$, causing a huge efficiency loss and limiting competitive interest compared to approaches relying on KLPT [KLPT14].

2.3.2 Application of Kani's lemma

Assume $q := \mathrm{nrd}(I)$ is (e, B)-good and we have solved Eq. (2.6) *i.e.*, we have found $a_1, a_2 \in \mathbb{Z}$ such that $a_1^2 + a_2^2 + q = 2^e$. Then, we proceed as in the 4-dimensional attack against SIDH presented in Section 2.2.4. We consider the $(a_1^2 + a_2^2, q)$-isogeny diamond:

$$\begin{array}{ccc} E_2^2 & \xrightarrow{\alpha_2} & E_2^2 \\ {\scriptstyle \Phi_I}\uparrow & & \uparrow{\scriptstyle \Phi_I} \\ E_1^2 & \xrightarrow{\alpha_1} & E_1^2 \end{array}$$

where, for $i \in \{1, 2\}$, α_i is an $(a_1^2 + a_2^2)$-isogeny

$$\alpha_i := \begin{pmatrix} a_1 & a_2 \\ -a_2 & a_1 \end{pmatrix} \in \mathrm{End}(E_i^2),$$

and Φ_I is the 2-dimensional diagonal q-isogeny $\Phi_I := \mathrm{Diag}(\varphi_I, \varphi_I) : E_1^2 \longrightarrow E_2^2$. Then, Kani's lemma (Lemma 2.2.6) ensures that:

$$F_I := \begin{pmatrix} \alpha_1 & \widetilde{\Phi}_I \\ -\Phi_I & \widetilde{\alpha}_2 \end{pmatrix} \in \mathrm{End}(E_1^2 \times E_2^2), \tag{2.10}$$

is a 2^e-isogeny with kernel:

$$\begin{aligned} \ker(F_I) &= \{(\widetilde{\alpha}_1(P, Q), \Phi_I(P, Q)) \mid P, Q \in E_1[2^e]\} \\ &= \{([a_1]P - [a_2]Q, [a_2]P + [a_1]Q, \varphi_I(P), \varphi_I(Q)) \mid P, Q \in E_1[2^e]\}. \end{aligned}$$

We can then compute $\ker(F_I)$ provided we can evaluate φ_I on a basis of $E_1[2^e]$. In order to compute F_I efficiently with level 2 theta coordinates, we need to compute $T_1, \cdots, T_4$ forming an isotropic subgroup of $\ker(F_I)$ and such that $\ker(F_I) = \langle [4]T_1, \cdots, [4]T_4 \rangle$ (Theorem 2.2.12). These points T_i can be obtained from the image of φ_I on a basis of $E_1[2^{e+2}]$, so we need $e + 2 \leq f$. Details of this computation can be found in Section 6.6.

Remark 2.3.4. Our algorithmic approach still succeeds even when $e + 2 > f$. Indeed, in our setting, $2^e = \Theta(\log(p)\sqrt{p})$ and $2^f = \Omega(\sqrt{p})$ so we have $f \geq e/2 + 2$, so as explained in Section 2.2.4, we may decompose $F_I = F_2 \circ F_1$, where F_i is 2^{e_i}-isogeny for $i \in \{1, 2\}$, with $e = e_1 + e_2$ and $e_1, e_2 \leq f - 2$. Knowing φ_I on a basis of $E_1[2^f]$, we are able to compute F_1 and $\widetilde{F}_2$, then $F_2 = \widetilde{\widetilde{F}}_2$ and finally $F_I = F_2 \circ F_1$. We refer to Section 6.6.7 for more details.

2.3.3 Evaluation of torsion points

It remains to explain how we can evaluate φ_I on a basis of $E_1[2^f]$. Since I is connecting ideal between $\mathcal{O}_1$ and $\mathcal{O}_2$, we have $I_2 \sim I_1 \cdot I$ so $I_2 = I_1 \cdot I \cdot \overline{\theta}/\operatorname{nrd}(I_1 \cdot I)$ with $\theta \in I_1 \cdot I \subseteq \mathcal{O}_0$ and by Lemma 1.2.24, we have:

$$\widehat{\varphi}_2 \circ \varphi_I \circ \varphi_1 = \varepsilon_0(\theta) \quad i.e. \quad [d_1 d_2]\varphi_I = \varphi_2 \circ \varepsilon_0(\theta) \circ \widehat{\varphi}_1,$$

where ε_0 is an isomorphism $\mathcal{O}_0 \xrightarrow{\sim} \operatorname{End}(E_0)$. Since $d_1 d_2$ is odd, we can find its inverse λ modulo 2^f and for all $P \in E_1[2^f]$, we can compute

$$\varphi_2 \circ \varepsilon_0(\theta) \circ \widehat{\varphi}_1([\lambda]P) = [d_1 d_2 \lambda]\varphi_I(P) = \varphi_I(P).$$

This proves that we can evaluate φ_I and compute its 4-dimensional embedding F_I defined in Eq. (2.10). Algorithm 2.2 follows.

Algorithm 2.2: ideal-to-isogeny translation in dimension 4.

Data: A supersingular elliptic curve $E_0/\mathbb{F}_{p^2}$, a maximal order $\mathcal{O}_0 \subseteq \mathcal{B}_{p,\infty}$ and an isomorphism $\varepsilon_0 : \mathcal{O}_0 \xrightarrow{\sim} \operatorname{End}(E_0)$, two isogenies $\varphi_1 : E_0 \longrightarrow E_1$ and $\varphi_2 : E_0 \longrightarrow E_2$ of odd degrees d_1 and d_2 respectively, along with their respective corresponding ideals $I_1, I_2 \subseteq \mathcal{O}_0$ and an ideal $I \subset \mathcal{B}_{p,\infty}$ connecting $\mathcal{O}_1 := O_R(I_1)$ and $\mathcal{O}_2 := O_R(I_2)$ of (e, B)-good norm.

Result: A 4-dimensional 2^e-isogeny $F_I \in \operatorname{End}(E_1^2 \times E_2^2)$ embedding φ_I.

$q \leftarrow \operatorname{nrd}(I)$;
Find $a_1, a_2 \in \mathbb{Z}$ such that $a_1^2 + a_2^2 + q = 2^e$ with Cornacchia's algorithm [Cor08] (see Section 2.3.1);
Find $\theta \in \mathcal{O}_0$ such that $I_1 \cdot I \cdot \overline{I}_2 = \mathcal{O}_0\theta$;
Compute $\lambda \in \mathbb{Z}$ such that $\lambda d_1 d_2 \equiv 1 \mod 2^f$;
Generate a basis (P, Q) of $E_1[2^f]$;
$\varphi_I(P) \leftarrow \varphi_2 \circ \varepsilon_0(\theta) \circ \widehat{\varphi}_1([\lambda]P)$;
$\varphi_I(Q) \leftarrow \varphi_2 \circ \varepsilon_0(\theta) \circ \widehat{\varphi}_1([\lambda]Q)$;
Use $P, Q, \varphi_I(P), \varphi_I(Q), a_1, a_2$ and the algorithms from Section 6.6 to compute the 4-dimensional 2^e-isogeny $F_I \in \operatorname{End}(E_1^2 \times E_2^2)$ embedding φ_I defined in Eq. (2.10);
return F_I;

2.4 Translating any ideal from a special curve with isogenies in dimension 2

In this section, we explain how to translate any left-ideal I of $\mathcal{O}_0 \simeq \operatorname{End}(E_0)$, into an isogeny $\varphi_I : E_0 \longrightarrow E_I$, where E_0 is the special supersingular elliptic curve of equation $y^2 = x^3 + x$ using 2-dimensional isogenies only. The method we apply here is deeply inspired from an algorithmic approach called Clapoti (class group action in polynomial time) introduced by Aurel Page and Damien Robert [PR23] to compute ideal class group actions on oriented supersingular elliptic curves. It has been used in the SQIsign2D-West [BDF+25] variant of SQIsign, which inspired the round 2 NIST submission of SQIsign due to its performance and security properties. In Chapter 3, we also use it to improve SQIsignHD. In the following, we consider a prime of the form $p = c \cdot 2^e - 1$ with $c \in \mathbb{N}^*$ small and odd and $e \in \mathbb{N}$ of several hundreds. We shall work with supersingular elliptic curves defined over $\mathbb{F}_{p^2}$, which all have their 2^e-torsion defined over $\mathbb{F}_{p^2}$.

2.4.1 Computing an isogeny of arbitrary odd degree from a special curve

Let E_0 be the elliptic curve given by $y^2 = x^3 + x$, defined over $\mathbb{F}_p$. Given an odd integer $u < 2^e$, we explain how to compute a u-isogeny $\varphi_u : E_0 \longrightarrow E_u$ that will serve in the ideal-to-isogeny translation algorithm. The following algorithm is due to Nakagawa and Onuki [NO24, Algorithm 2].

By Lemma 1.2.25, $\operatorname{End}(E_0)$ is isomorphic to $\mathcal{O}_0 := \langle 1, i, (i+j)/2, (1+ij)/2 \rangle \subset \mathcal{B}_{p,\infty}$ so we may consider an isomorphism $\varepsilon_0 : \mathcal{O}_0 \xrightarrow{\sim} \operatorname{End}(E_0)$. In particular, it contains the suborder $\langle 1, i, j, ij \rangle$, where norm equations are easy to solve. If $N = \Omega(p \log(p))$, then we can solve the equation:

$$N = \operatorname{nrd}(x + iy + jz + tij) = x^2 + y^2 + p(z^2 + t^2), \tag{2.11}$$

with $x, y, z, t \in \mathbb{Z}$ as follows. We sample small values $z, t \in \mathbb{Z}$ until $N - p(z^2+t^2)$ is a prime $\equiv 1 \mod 4$. Then, we apply Cornacchia's algorithm [Cor08] to find $x, y \in \mathbb{Z}$ such that $N - p(z^2+t^2) = x^2 + y^2$. We refer to [Ler22, Algorithm 2] for more details. In our use case, we may apply this algorithm to $N = u(2^f - u)$ with $f \leq e$ as small as possible so that $u(2^f - u) = \Omega(p \log(p))$. We then obtain $\theta \in \mathcal{O}_0$ of norm $\mathrm{nrd}(\theta) = u(2^f - u)$.

The endomorphism θ can be written as $\varepsilon_0(\theta) = \psi_u \circ \varphi_u = \varphi'_u \circ \psi'_u$, where φ_u, φ'_u have degree u and ψ_u, ψ'_u have degree $2^f - u$. Now, consider the $(u, 2^f - u)$-isogeny diamond:

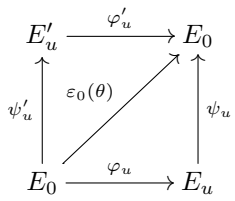

Applying Kani's Lemma (Lemma 2.2.6), we obtain a 2^f-isogeny:

$$\Phi_u = \begin{pmatrix} \varphi_u & \widehat{\psi_u} \\ -\psi'_u & \widehat{\varphi'}_u \end{pmatrix} : E_0 \times E_0 \longrightarrow E_u \times E'_u.$$

with kernel:

$$\ker(\Phi_u) = \{([u]P, \varepsilon_0(\theta)(P)) \mid P \in E_0[2^f]\}.$$

Then, by Theorems 2.2.12 and 2.2.13 there exists very efficient algorithms to compute Φ_u that will be presented in Section 6.5.

If $f \leq e - 2$, we can evaluate $\varepsilon_0(\theta)$ on a basis of $E_0[2^{f+2}]$ and then obtain 2^{f+2}-torsion points $T_1, T_2 \in E_0^2[2^{f+2}]$ such that $\ker(\Phi_u) = \langle [4]T_1, [4]T_2 \rangle$ and forming an isotropic 2^{f+2}-torsion subgroup *i.e.* such that $e_{2^{f+2}}(T_1, T_2) = 1$. With the algorithms from Section 6.5 taking T_1 and T_2 as input, we can compute Φ_u as a chain of 2-isogenies of length f in level 2 theta coordinates. These algorithms also convert image points by Φ_u from theta coordinates to Montgomery $(X : Z)$-coordinates on the codomain $E_u \times E'_u$ (see Section 6.5.4). In particular, we know that the component on E_u appears first in the image points by Φ_u and we can evaluate φ_u on any point $P \in E_0$ as follows $\Phi_u(P, 0) = (\varphi_u(P), -\psi'_u(P))$.

If $e - 1 \leq f \leq e$, we can evaluate $\varepsilon_0(\theta)$ on a basis of $E_0[2^f]$ and then obtain generators of $\ker(\Phi_u)$ that can also be used to compute Φ_u as a chain of 2-isogenies of length f in level 2 theta coordinates by Theorem 2.2.13. However, this costs 5 additional square root computations compared to using the 2^{f+2}-torsion as previously. Moreover, when converting codomain theta coordinates into product Montgomery $(X : Z)$-coordinates, the elliptic curves E_u and E'_u may be swapped. Image points can either be expressed in $E_u \times E'_u$ or $E'_u \times E_u$. To distinguish E_u from E'_u (and φ_u from ψ'_u), we may use Weil pairings as follows.

Let (P_0, Q_0) be a basis of $E_0[2^e]$, $(P, P') := \Phi_u(P_0, 0)$ and $(Q, Q') := \Phi_u(Q_0, 0)$. If image points lie in $E_u \times E'_u$, then $e_{2^e}(P, Q) = e_{2^e}(\varphi_u(P_0), \varphi_u(Q_0)) = e_{2^e}(P_0, Q_0)^u$. Otherwise, image points lie in $E'_u \times E_u$ and $e_{2^e}(P, Q) = e_{2^e}(-\psi'_u(P_0), -\psi'_u(Q_0)) = e_{2^e}(P_0, Q_0)^{2^f - u} \neq e_{2^e}(P_0, Q_0)^u$. Since image points are expressed in product Montgomery $(X : Z)$-coordinates, we only know their value on Kummer lines $(\pm P, \pm P')$ and $(\pm Q, \pm Q')$. To lift this sign ambiguity and be able to compute the Weil pairing, we may compute $\Phi_u(P_0 - Q_0, 0) = (\pm(P - Q), \pm(P' - Q'))$.

Finally, we can also easily obtain the left $\mathcal{O}_0$-ideal associated to φ_u via the Deuring correspondence which can be useful in some contexts (*e.g.* SQIsignHD).

Lemma 2.4.1. *The ideal associated to φ_u is $I_u := \mathcal{O}_0\theta + u\mathcal{O}_0$.*

Proof. We have $E_0[\mathcal{O}_0\theta + u\mathcal{O}_0] = \ker(\theta) \cap E_0[u]$. First, observe that $\ker(\varphi_u) \subseteq \ker(\varepsilon_0(\theta)) \cap E_0[u]$ since $\varepsilon_0(\theta)$ factors through φ_u, which has degree u.

Conversely, if $P \in \ker(\varepsilon_0(\theta)) \cap E_0[u]$, then $\psi_u \circ \varphi_u(P) = 0$. As a result, $\varphi_u(P) \in E_u[u] \cap \ker(\psi_u)$ and $E_u[u] \cap \ker(\psi_u) \subseteq E_u[u] \cap E_u[2^f - u] = \{0\}$, since u and $2^f - u$ are coprime. Thus, $P \in \ker(\varphi_u)$, proving that $E_0[\mathcal{O}_0\theta + u\mathcal{O}_0] = \ker(\varphi_u)$.

The claim $I_u = \mathcal{O}_0\theta + u\mathcal{O}_0$ follows from the Deuring correspondence. □

We summarise the procedure described above in Algorithm 2.3 that evaluates φ_u on (P_0, Q_0) (which will be enough for our use case in Section 2.4.2) and returns the associated ideal I_u as a bonus.

Algorithm 2.3: Evaluation of an isogeny of fixed odd degree.

Data: An odd positive integer $u < 2^e$ such that $u(2^e - u) = \Omega(p\log(p))$ and a basis (P_0, Q_0) of $E_0[2^e]$.
Result: $(E_u, \varphi_u(P_0), \varphi_u(Q_0), I_u)$, where $\varphi_u : E_0 \longrightarrow E_u$ is a u-isogeny and $I_u \subset \mathcal{O}_0$ is its associated ideal.

Find $f \leq e$ as small as possible so that $u(2^f - u) = \Omega(p\log(p))$;
Find $\theta \in \mathcal{O}_0$ of norm $u(2^f - u)$; // [Ler22, Algorithm 2]
$I_u \leftarrow \mathcal{O}_0\theta + u\mathcal{O}_0$;
if $f \leq e - 2$ **then**
  $P \leftarrow [2^{e-f-2}]P_0$, $Q \leftarrow [2^{e-f-2}]Q_0$;
  $T_1 \leftarrow ([u]P, \varepsilon_0(\theta)(P))$, $T_2 \leftarrow ([u - 2^f]Q, \varepsilon_0(\theta)(Q))$;
  Use T_1 and T_2 to compute $\Phi_u : E_0^2 \longrightarrow E_u \times E'_u$ of kernel $\langle [4]T_1, [4]T_2\rangle$; // see Section 6.5
  $(\pm\varphi_u(S), *) \leftarrow \Phi_u(S, 0)$ for $S \in \{P_0, Q_0, P_0 - Q_0\}$;
  Lift $(\varphi_u(P_0), \varphi_u(Q_0))$ from $\pm\varphi_u(P_0), \pm\varphi_u(Q_0), \pm\varphi_u(P_0 - Q_0)$;
  return $(E_u, \varphi_u(P_0), \varphi_u(Q_0), I_u)$;
else
  $P \leftarrow [2^{e-f}]P_0$, $Q \leftarrow [2^{e-f}]Q_0$;
  $T_1 \leftarrow ([u]P, \varepsilon_0(\theta)(P))$, $T_2 \leftarrow ([u]Q, \varepsilon_0(\theta)(Q))$;
  Compute $\Phi_u : E_0^2 \longrightarrow E \times E'$ of kernel $\langle T_1, T_2\rangle$; // see Section 6.5
  $(\pm I(S), \pm I'(S)) \leftarrow \Phi_u(S, 0)$ for $S \in \{P_0, Q_0, P_0 - Q_0\}$;
  Lift $(I(P_0), I(Q_0))$ from $\pm I(P_0), \pm I(Q_0), \pm I(P_0 - Q_0)$;
  Lift $(I'(P_0), I'(Q_0))$ from $\pm I'(P_0), \pm I'(Q_0), \pm I'(P_0 - Q_0)$;
  if $e_{2^e}(I(P_0), I(Q_0)) = e_{2^e}(P_0, Q_0)^u$ **then**
    $E_u \leftarrow E$, $(\varphi_u(P_0), \varphi_u(Q_0)) \leftarrow (I(P_0), I(Q_0))$;
  else
    $E_u \leftarrow E'$, $(\varphi_u(P_0), \varphi_u(Q_0)) \leftarrow (I'(P_0), I'(Q_0))$;
  end
  return $(E_u, \varphi_u(P_0), \varphi_u(Q_0), I_u)$;
end

2.4.2 The ideal-to-isogeny translation algorithm

Let I be a left $\mathcal{O}_0$-ideal. We want to efficiently represent the associated isogeny $\varphi_I : E_0 \longrightarrow E_I$. In this section, we give an algorithm that computes E_I and evaluates φ_I on a basis (P_0, Q_0) of $E_0[2^e]$. The general outline is as follows:

1. Find $I_1, I_2 \sim I$ of coprime norms $d_1, d_2 \approx \sqrt{p}$, and $u, v \in \mathbb{N}^*$ such that
$$d_1 u + d_2 v = 2^f \tag{2.12}$$
with $f \leq e$ and $\gcd(ud_1, vd_2) = 1$.

2. Evaluate isogenies $\varphi_u, \varphi_v : E_0 \longrightarrow E_u, E_v$ of degrees u and v on (P_0, Q_0) using Algorithm 2.3.

3. Apply Kani's Lemma to $\varphi_u \circ \widehat{\varphi}_1 : E_I \longrightarrow E_u$ and $\varphi_v \circ \widehat{\varphi}_2 : E_I \longrightarrow E_v$, where $\varphi_1, \varphi_2 : E_0 \longrightarrow E_I$ are the isogenies corresponding to I_1 and I_2 respectively, to compute $\Phi : E_u \times E_v \longrightarrow E_I \times E'$ that embeds the isogenies $\varphi_1 \circ \widehat{\varphi}_u$ and $\varphi_2 \circ \widehat{\varphi}_v$.

4. Evaluate Φ on well chosen points to obtain $(\varphi_I(P_0), \varphi_I(Q_0))$.

In the following, we give more details on each step and we summarise these steps in Algorithm 2.5.

Step 1: norm equation

We sample ideals $I_1, I_2 \sim I$ of odd coprime norms d_1 and d_2 until we find positive integers u, v such that $d_1 u + d_2 v = 2^e$. It has been proved that this coin problem admits a solution whenever $2^e \geq (d_1 - 1)(d_2 - 1)$ [Syl82]. Hence, the norms d_1 and d_2 should be as small as possible. To find equivalent ideals of such norms, we sample $\beta_i \in I$ with sufficiently small reduced norm and choose $I_i := I\overline{\beta_i}/\operatorname{nrd}(I)$, so that $\operatorname{nrd}(I_i) = \operatorname{nrd}(\beta_i)/\operatorname{nrd}(I)$. Lemma 2.3.2 ensures that the shortest vector in I has norm $O(\operatorname{nrd}(I)\sqrt{p})$ so we should expect to find $d_1, d_2 = O(\sqrt{p})$ so that $(d_1 - 1)(d_2 - 1) \approx p \approx 2^e$ in general. This is not enough to rigorously ensure the existence of a solution (u, v) to the coin problem and we have to rely on heuristic arguments. In the following, we provide such heuristic arguments to justify that we can find a solution I_1, I_2, u, v of Eq. (2.12) for an overwhelming majority of ideals I.

Consider an integer $N < p$ (in our application $N = 2^e$) and coprime positive integers $d_1, d_2 \leq N/2$ also coprime with N. Then, for any solution $(u, v) \in (\mathbb{N}^*)^2$ to the coin problem $ud_1 + vd_2 = N$, v is fixed by the value of u as $v := (N - ud_1)/d_2$. In particular, $d_2 | N - ud_1$. Heuristically, when u is a random integer between 1 and $N/d_1 - 1$, the probability that $d_2 | N - ud_1$ conditionally to d_1 and d_2 is $\approx 1/d_2$, so the probability that such an integer u exists is $\approx N/d_1 \cdot 1/d_2 = N/(d_1 d_2)$. Hence, we can make the following heuristic assumption: for any $1 \leq d \leq N/2$, the probability that a couple (d_1, d_2) selected uniformly at random among couples of coprime integers such that $d_1, d_2 \leq d$ has a solution $(u, v) \in (\mathbb{N}^*)^2$ to the associated coin problem $ud_1 + vd_2 = N$ is larger than N/d^2.

This heuristic assumption is still not sufficient because the quadratic form $q_I(\beta) := \operatorname{nrd}(\beta)/\operatorname{nrd}(I)$ has not the same distribution as a uniformly random integer when we sample $\beta \in I$ such that $q_I(\beta) \leq d$. We now give more detail on how we sample $\beta \in I$. First, we find a Minkowski reduced basis $\mathscr{B} := (\alpha_1, \cdots, \alpha_4)$ of I, so that $q_I(\alpha_1) \leq \cdots \leq q_I(\alpha_4)$ are the successive minima of q_I in the lattice I by Theorem 1.2.28. Then we sample $x_j \in [\![-B_j \ ; \ B_j]\!]$ uniformly at random with $B_j := \lfloor 1/4\sqrt{d/q_I(\alpha_j)} \rfloor$ for all $j \in [\![1 \ ; \ 4]\!]$ and we set $\beta := \sum_{j=1}^4 x_j \alpha_i$. By the triangular inequality (which is valid since q_I is a positive definite quadratic form), we have,

$$q_I(\beta) \leq \left(\sum_{j=1}^{4} |x_j| \sqrt{q_I(\alpha_j)} \right)^2 \leq d.$$

Hence, we want that $q_I(\beta_1)$ and $q_I(\beta_2)$ satisfy the aforementioned integer heuristic assumption when β_1 and β_2 are sampled uniformly at random in the set:

$$P_d(I, \mathscr{B}) := \left\{ \sum_{j=1}^{4} x_j \alpha_j \;\middle|\; \forall j \in [\![1 \ ; \ 4]\!], \quad x_j \in [\![-B_j \ ; \ B_j]\!] \right\}. \tag{2.13}$$

i.e. that a solution $(u, v) \in (\mathbb{N}^*)^2$ to the coin problem $uq_I(\beta_1) + vq_I(\beta_2) = N$ exits with probability $\Omega(N/d^2)$.

We notice that if $J \sim I$ is an equivalent left $\mathcal{O}_0$-ideal, then q_I and q_J are equivalent quadratic forms. In particular, they take the same values. Indeed, by Lemma 1.2.19.(i), we may write $J = I\overline{\beta}/\operatorname{nrd}(I)$ with $\beta \in I$ so that every $\alpha \in J$ can be written as $\alpha = \gamma\overline{\beta}/\operatorname{nrd}(I)$ with $\gamma \in I$ and:

$$q_J(\alpha) = \frac{\operatorname{nrd}(\alpha)}{\operatorname{nrd}(J)} = \frac{\operatorname{nrd}(\gamma)\operatorname{nrd}(\beta)}{\operatorname{nrd}(I)^2 q_I(\beta)} = \frac{\operatorname{nrd}(\gamma)\operatorname{nrd}(\beta)}{\operatorname{nrd}(I)\operatorname{nrd}(\beta)} = \frac{\operatorname{nrd}(\gamma)}{\operatorname{nrd}(I)} = q_I(\gamma).$$

Hence, the heuristic assumption is satisfied uniformly in a given equivalence class of left $\mathcal{O}_0$-ideals. By the Deuring correspondence, there is a finite number $h(\mathcal{O}_0)$ of such equivalence classes, bounded by the number of supersingular j-invariants. Actually, the Eichler mass formula [Voi21, Theorem 25.1.1] ensures that $h(\mathcal{O}_0) = p/24 + O(1)$. We make the following heuristic assumption ensuring that the desired heuristic holds for an overwhelming majority of left $\mathcal{O}_0$-ideal equivalence classes.

Heuristic 2.4.2. *For any integer $N \in \mathbb{N}^*$, left $\mathcal{O}_0$-ideal I, Minkowski reduced basis $\mathscr{B}$ of I and $1 \leq d \leq N/2$, let us denote:*

$$\begin{aligned} S_{d,N}(I, \mathscr{B}) := \{ &(\beta_1, \beta_2) \in P_d(\mathscr{B})^2 \mid \gcd(q_I(\beta_1), q_I(\beta_2)) = \gcd(q_I(\beta_1), N) = \gcd(q_I(\beta_2), N) = 1 \\ &\text{and} \quad \exists u, v \in \mathbb{N}^*, \ uq_I(\beta_1) + vq_I(\beta_2) = N \}. \end{aligned}$$

Then, for all $\varepsilon > 0$ and $0 < \rho < 1$, there exists a constant $C_{\varepsilon,\rho} > 0$ such that for all prime $p \equiv 3 \bmod 4$ and integer $\rho \cdot p \leq N \leq p - 1$:

(i) There exists a set $\mathcal{S}_\varepsilon(N, p)$ of equivalence classes of left ideals in $\mathcal{O}_0 \subseteq \mathcal{B}_{p,\infty}$ of cardinality $\#\mathcal{S}_\varepsilon(N, p) \geq (1 - \varepsilon)h(\mathcal{O}_0)$;

(ii) For all left $\mathcal{O}_0$-ideal I whose equivalence class lies in $\mathcal{S}_\varepsilon(N, p)$, for all Minkowski reduced basis $\mathscr{B}$ of I and for all $1 \leq d \leq N/2$, we have:

$$\#S_{d,N}(I, \mathscr{B}) \geq C_{\varepsilon,\rho} \cdot \frac{\#P_d(I, \mathscr{B})^2 N}{d^2} - 1. \tag{2.14}$$

Remark 2.4.3. Let us make a few comments on this heuristic assumption. The constant ρ is the minimal admissible value of N/p and should be close to 1 in our application ($N = 2^e$). The constant $\varepsilon > 0$ sets the proportion $\#\mathcal{S}_\varepsilon(N, p)/h(\mathcal{O}_0) \geq 1 - \varepsilon$ of left $\mathcal{O}_0$-ideal equivalence classes satisfying Eq. (2.14). The constant $C_{\varepsilon,\rho} > 0$ used to bound $\#S_{d,N}(I, \mathscr{B})$ depends on ρ and ε as follows. The lower ρ, the harder it is to find a solution to the coin equation $uq_I(\beta_1) + vq_I(\beta_2) = N$ (since ρ is the minimal value of N/p), and the lower $C_{\varepsilon,\rho}$. Similarly, the lower ε, the more ideals we require to satisfy Eq. (2.14) and the lower $C_{\varepsilon,\rho}$. When the probability of finding a coin solution N/d^2 or the box $P_d(I, \mathscr{B})$ are small, the set $S_{d,N}(I, \mathscr{B})$ may be empty. This explains the -1 in Eq. (2.14).

In practice, to solve the norm equation $d_1u + d_2v = M$ with $M|N$, $u, v > 0$ and $\gcd(ud_1, vd_2) = 1$, we proceed as follows. We sample $\beta_1, \beta_2 \in P_d(I, \mathscr{B})$ until $d_1 := q_I(\beta_1)$ and $d_2 := q_I(\beta_2)$ are coprime and coprime with N. Then, we apply extended Euclide's algorithm to find $u_0, v_0 \in \mathbb{Z}$ such that $u_0d_1 + v_0d_2 = 1$. Then, all solutions to the Diophantine equation $ud_1 + vd_2 = N$ are of the form $(u, v) = (Nu_0 + kd_2, Nv_0 - kd_1)$ for $k \in \mathbb{Z}$. To find $u, v > 0$ we select an integer $-Nu_0/d_2 \leq k \leq Nv_0/d_1$. If there is no integer in the interval $[-Nu_0/d_2, Nv_0/d_1]$, then we sample β_1 and β_2 again. Finally, to ensure $\gcd(d_1u, d_2v) = 1$, we divide u and v and N by $\gcd(u, v)$. Indeed, every common factor of ud_1 and vd_2 is a common factor of u and v, which is also a factor of N. Algorithm 2.4 follows.

Algorithm 2.4: Step 1: Finding suitable ideals $I_1, I_2 \sim I$ and a coin problem solution (u, v).

Data: An ideal $I \subseteq \mathcal{O}_0 \simeq \mathrm{End}(E_0)$, an integer $N = \Theta(p)$ and a bound $d \leq N/2$.
Result: $\beta_1, \beta_2 \in I$, $M|N$ and $u, v \in \mathbb{N}^*$ such that $\gcd(uq_I(\beta_1), vq_I(\beta_2)) = 1$ and $uq_I(\beta_1) + vq_I(\beta_2) = M$.

Compute a Minkowski reduced basis $(\alpha_1, \cdots, \alpha_4)$ of I ; // [NS09, Fig. 3-4]
$B_j \leftarrow \lfloor 1/4\sqrt{d/q_I(\alpha_j)} \rfloor$ for $j \in [\![1\ ;\ 4]\!]$;
while True **do**
  Sample $x_j, y_j \in [\![-B_j\ ;\ B_j]\!]^4$ for $j \in [\![1\ ;\ 4]\!]$ independently and uniformly at random;
  $\beta_1 \leftarrow \sum_{j=1}^4 x_j\alpha_j$, $\beta_2 \leftarrow \sum_{j=1}^4 y_j\alpha_j$;
  $d_1 \leftarrow q_I(\beta_1)$, $d_2 \leftarrow q_I(\beta_2)$;
  if $\gcd(d_1, d_2) = \gcd(d_1, N) = \gcd(d_2, N) = 1$ **then**
    Apply extended Euclide's algorithm to find $u_0, v_0 \in \mathbb{Z}$ such that $u_0d_1 + v_0d_2 = 1$;
    if $[-Nu_0/d_2, Nv_0/d_1] \cap \mathbb{Z} \neq \emptyset$ **then**
      Take $k \leftarrow \lceil -Nu_0/d_2 \rceil$;
      $u \leftarrow Nu_0 + kd_2$, $v \leftarrow Nv_0 - kd_1$;
      $d \leftarrow \gcd(u, v)$;
      $u \leftarrow u/d$, $v \leftarrow v/d$, $M \leftarrow N/d$;
      return $\beta_1, \beta_2, M, u, v$;
    end
  end
end

Lemma 2.4.4. *Assume Heuristic 2.4.2. Then for all $\varepsilon > 0$ and $0 < \rho < 1$, there exists constants $T_{\varepsilon,\rho}, n_{\varepsilon,\rho} > 0$ such that for all $\max(\rho \cdot p, n_{\varepsilon,\rho}) \leq N \leq p-1$, Algorithm 2.4 taking N as input terminates after less than $T_{\varepsilon,\rho}$ iterations on average on a proportion $\geq 1 - \varepsilon$ of left $\mathcal{O}_0$-ideal classes for a well chosen bound $d = \Theta_{\varepsilon,\rho}(\sqrt{p})$.*

Proof. Note that Algorithm 2.4 terminates if $S_{d,N}(I,\mathscr{B}) \neq \emptyset$. Let $1 \leq d \leq N/2$. Then we have:

$$\#P_d(I,\mathscr{B}) = \prod_{j=1}^{4}(2B_j+1) \geq \prod_{j=1}^{4}\frac{1}{4}\sqrt{\frac{d}{q_I(\alpha_j)}} = \frac{d^2}{256\sqrt{\prod_{j=1}^{4} q_I(\alpha_j)}}. \tag{2.15}$$

Besides, $\mathscr{B} = (\alpha_1, \cdots, \alpha_4)$ being Minkowski reduced, $q_I(\alpha_1), \cdots, q_I(\alpha_4)$ are the successive minima of the lattice I by Theorem 1.2.28 and we obtain by Minkowski's second theorem (see Eq. (2.8)) that $\prod_{j=1}^{4} q_I(\alpha_j) \leq 64p^2/\pi^4$, so that:

$$\#P_d(I,\mathscr{B}) \geq \frac{Dd^2}{p},$$

with $D := \pi^2/2048$. By Heuristic 2.4.2, there exists $C_{\varepsilon,\rho} > 0$ such that on a proportion $\geq 1-\varepsilon$ of left $\mathcal{O}_0$-ideal classes $[I]$, the ideal I satisfies:

$$\#S_{d,N}(I,\mathscr{B}) \geq C_{\varepsilon,\rho} \cdot \frac{\#P_d(I,\mathscr{B})^2 N}{d^2} - 1$$

Hence, when $\beta_1, \beta_2 \in P_d(I,\mathscr{B})$ are sampled independently and uniformly at random as in Algorithm 2.4, we have:

$$\begin{aligned}\mathbb{P}((\beta_1,\beta_2) \in S_{d,N}(I,\mathscr{B})) = \frac{\#S_{d,N}(I,\mathscr{B})}{\#P_d(I,\mathscr{B})^2} &\geq C_{\varepsilon,\rho}\cdot\frac{N}{d^2} - \frac{1}{\#P_d(I,\mathscr{B})^2} \\ &\geq C_{\varepsilon,\rho}\cdot\frac{N}{d^2} - \frac{p^2}{D^2 d^4}\end{aligned}$$

The lower bound on the right is maximised by the value $d = d^* := p/D\sqrt{2/(C_{\varepsilon,\rho}N)} = \Theta_{\varepsilon,\rho}(\sqrt{p})$. Since $d \leq N/2$, it follows that

$$\frac{N^{3/2}}{p} \geq \frac{2\sqrt{2}}{D\sqrt{C_{\varepsilon,\rho}}}$$

and this equality is always satisfied whenever $N \geq n_{\varepsilon,\rho} := 8/(D^2\rho^2 C_{\varepsilon,\rho})$. Besides, for the optimal value $d = d^*$, we obtain:

$$\mathbb{P}((\beta_1,\beta_2) \in S_{d,N}(I,\mathscr{B})) \geq \frac{C_{\varepsilon,\rho}^2 D^2 N^2}{4p^2} \geq \frac{C_{\varepsilon,\rho}^2 D^2 \rho^2}{4}$$

Hence, Algorithm 2.4 terminates after less than $T_{\varepsilon,\rho} := 4/(C_{\varepsilon,\rho}^2 D^2 \rho^2)$ iterations on average for all $\max(\rho\cdot p, n_{\varepsilon,\rho}) \leq N \leq p-1$ on a proportion $\geq 1-\varepsilon$ of left $\mathcal{O}_0$-ideal classes. □

Remark 2.4.5. Lemma 2.4.4 tells us that N should be big enough for Algorithm 2.4 to terminate and that the termination may not happen for all ideal classes. The more success probability we require (measured by $1-\varepsilon$) the more values β_1, β_2 we need to try and the more time we have to wait for the algorithm to terminate (as the bound d increases).

In most cases, the successive minima $q_I(\alpha_j)$ are close to each other and close to $\sqrt{p}$ so the B_j are very close and small. For that reason, in the implemented version of Algorithm 2.4 we fix a bound $B \in \mathbb{N}^*$ and sample all of the x_j and y_j in $[\![-B\ ;\ B]\!]$. Besides, we enumerate x_j and y_j in $[\![-B\ ;\ B]\!]$ instead of sampling them at random to avoid trying the same couples (β_1,β_2) twice. Also note that the reduced basis $(\alpha_1, \cdots, \alpha_4)$ may be computed with different algorithms than the Minkowski reduction algorithm. In dimension 4, a reduced basis obtained with LLL [LLL82] or BKZ [Sch87] is generally very close to a Minkowski reduced one.

With the SQIsign2D-West [BDF+25] implementation, we encountered (very rare) failure cases, that may have been explained for two reasons:

- The successive minima $q_I(\alpha_j)$ were not balanced and the uniform box approach $[\![-B\ ;\ B]\!]$ was unsuccessful.
- The ideal class $[I]$ was outside of the set of size $\geq (1-\varepsilon)h(\mathcal{O}_0)$ where Algorithm 2.4 terminates.

To avoid these (very rare) failure cases, a new method has been proposed for the round 2 NIST submission of SQIsign [AAA+25] where other starting curves than E_0 of known endomorphism rings were considered. This will be explained in Section 2.4.3.

Step 2: finding isogenies of degrees u and v

We can use Algorithm 2.3 to evaluate isogenies $\varphi_u : E_0 \longrightarrow E_u$ and $\varphi_v : E_0 \longrightarrow E_v$ of degrees u and v on $E_0[2^e]$. Note that since $u, v = O(\sqrt{p})$, Algorithm 2.3 will compute a 2^f-isogeny in dimension 2 with $f \simeq e/2$ to evaluate each isogeny φ_u and φ_v. This saves half of the cost of a full 2^e-isogeny computation.

Nonetheless, the cost of this step could be further improved if, for instance, u was a sum of two squares $u = x^2 + y^2$. Indeed, in that case, we may consider $\varphi_u := [x] + [y]\iota \in \mathrm{End}(E_0)$, with $\iota : (x, y) \in E_0 \mapsto (-x, \sqrt{-1}y) \in E_0$, which is immediate to evaluate. However, requiring that u or v from Eq. (2.12) is a sum of squares dramatically increases the cost and reduces the success probability of Step 1. For these reasons, this trick was deemed not optimal and not implemented in SQIsign2D-West [BDF+25].

Step 3: Applying Kani's lemma

Assume we have computed a solution to Eq. (2.12) in Step 1 and consider the following (d_1u, d_2v)-isogeny diamond:

$$\begin{array}{ccccc} E' & & \xrightarrow{\quad\widehat{\varphi'}_v\quad} & & E_v \\ \uparrow{\scriptstyle \varphi'_u} & & & & \uparrow{\scriptstyle \varphi_v} \\ & & & & E_0 \\ & & & & \uparrow{\scriptstyle \widehat{\varphi}_2} \\ E_u & \xrightarrow{\widehat{\varphi}_u} & E_0 & \xrightarrow{\varphi_1} & E_I \end{array} \tag{2.16}$$

where $\varphi'_u := [\varphi_u \circ \widehat{\varphi}_1]_*(\varphi_v \circ \widehat{\varphi}_2)$ and $\varphi'_v := [\varphi_v \circ \widehat{\varphi}_2]_*(\varphi_u \circ \widehat{\varphi}_1)$ (pushforward isogenies). By Kani's Lemma (Lemma 2.2.6), we have a 2^f-isogeny:

$$\Phi := \begin{pmatrix} \varphi_1 \circ \widehat{\varphi}_u & \varphi_2 \circ \widehat{\varphi}_v \\ -\varphi'_u & \varphi'_v \end{pmatrix} : E_u \times E_v \longrightarrow E_I \times E',$$

with kernel:

$$\ker(\Phi) = \{([d_1]\varphi_u(P), \varphi_v \circ \widehat{\varphi}_2 \circ \varphi_1(P)) \mid P \in E_0[2^f]\}.$$

By Lemma 1.2.24, if we write $I_1 := I\overline{\beta_1}/\operatorname{nrd}(I)$ and $I_2 := I\overline{\beta_2}/\operatorname{nrd}(I)$ with $\beta_1, \beta_2 \in I$, and $\theta := \beta_2\overline{\beta_1}/\operatorname{nrd}(I)$, so that $I_2 = I_1\overline{\theta}/\operatorname{nrd}(I_1)$, we have $\widehat{\varphi}_2 \circ \varphi_1 = \varepsilon_0(\theta)$ so that $\widehat{\varphi}_2 \circ \varphi_1$ can be evaluated easily.

Let (P_0, Q_0) be a basis of $E_0[2^e]$ and assume that the images of P_0 and Q_0 by φ_u and φ_v have been computed in Step 2. To evaluate $\varphi_v \circ \widehat{\varphi}_2 \circ \varphi_1(P_0, Q_0)$ we evaluate $\varepsilon_0(\theta)(P_0, Q_0)$ and find $a, b, c, d \in \mathbb{Z}/2^e\mathbb{Z}$ such that $\varepsilon_0(\theta)(P_0) = [a]P_0 + [b]Q_0$ and $\varepsilon_0(\theta)(Q_0) = [c]P_0 + [d]Q_0$ with discrete logarithm computations, so that:

$$\varphi_v \circ \widehat{\varphi}_2 \circ \varphi_1(P_0) = [a]\varphi_v(P_0) + [b]\varphi_v(Q_0) \quad \text{and} \quad \varphi_v \circ \widehat{\varphi}_2 \circ \varphi_1(Q_0) = [c]\varphi_v(P_0) + [d]\varphi_v(Q_0).$$

We can then compute $T_1 := ([d_1]\varphi_u(P_0), \varphi_v \circ \widehat{\varphi}_2 \circ \varphi_1(P_0))$ and $T_2 := ([d_1 - 2^f\mu]\varphi_u(Q_0), \varphi_v \circ \widehat{\varphi}_2 \circ \varphi_1(Q_0))$ with $\mu \equiv 1/u \mod 2^e$. If $f \leq e - 2$, then $([2^{e-f-2}]T_1, [2^{e-f-2}]T_2)$ form an isotropic subgroup of $(E_u \times E_v)[2^{f+2}]$ above $\ker(\Phi)$ that can directly be used to compute $\Phi : E_u \times E_v \longrightarrow E_I \times E'$ with the algorithms from Section 6.5 (Theorem 2.2.12).

However, we expect solutions to Eq. (2.12) to be tight so $e - 1 \leq f \leq e$. Then (T_1, T_2) or $([2]T_1, [2]T_2)$ is a basis of $\ker(\Phi)$ that can be used to compute Φ with 5 square roots and a possible swap of E_I and E' in the codomain as we have seen in Section 2.4.1. To distinguish E_I from E' in the codomain of Φ (and $\varphi_1 \circ \widehat{\varphi}_u$ from $\varphi_2 \circ \widehat{\varphi}_v$), we may evaluate Φ and compute Weil pairings, as explained in Section 2.4.1. In Algorithm 2.5, we always apply the latter method to $([2^{e-f}]T_1, [2^{e-f}]T_2)$ since the case $e - 1 \leq f \leq e$ is the most frequent and to facilitate a potential future constant time implementation.

Step 4: point evaluation

Finally, we explain how to evaluate $(\varphi_I(P_0), \varphi_I(Q_0))$ for a basis (P_0, Q_0) of $E_0[2^e]$. We first notice that we can evaluate $\varphi_1 \circ \widehat{\varphi}_u$ from the two-dimensional isogeny Φ. This implies we can evaluate φ_1 on $E_0[2^e]$ as follows: $\Phi(\varphi_u(P_0), 0) = ([u]\varphi_1(P_0), *)$ and $\Phi(\varphi_u(Q_0), 0) = ([u]\varphi_1(Q_0), *)$ and we can multiply by $\mu \equiv 1/u \mod 2^e$ to get $(\varphi_1(P_0), \varphi_1(Q_0))$.

Since $I_1 := I\overline{\beta_1}/\operatorname{nrd}(I)$, by Lemma 1.2.24 we have $\widehat{\varphi}_1 \circ \varphi_I = \varepsilon_0(\beta_1)$, so that $[d_1]\varphi_I = \varphi_1 \circ \varepsilon_0(\beta_1)$. We may compute discrete logarithms $a, b, c, d \in \mathbb{Z}/2^e\mathbb{Z}$ such that $\varepsilon_0(\beta_1)(P_0) = [a]P_0 + [b]Q_0$ and $\varepsilon_0(\beta_1)(Q_0) = [c]P_0 + [d]Q_0$, so that:

$$\varphi_I(P_0) = [a\delta_1]\varphi_1(P_0) + [b\delta_1]\varphi_1(Q_0) \quad \text{and} \quad \varphi_I(Q_0) = [c\delta_1]\varphi_1(P_0) + [d\delta_1]\varphi_1(Q_0),$$

with $\delta_1 \equiv 1/d_1 \mod 2^e$. We summarise Steps 1-4 in Algorithm 2.5.

Algorithm 2.5: ideal-to-isogeny from E_0.

Data: A left $\mathcal{O}_0$-ideal I and a basis (P_0, Q_0) of $E_0[2^e]$
Result: The image $(E_I, \varphi_I(P_0), \varphi_I(Q_0))$ of the isogeny $\varphi_I : E_0 \to E_I$ associated to I.

Call Algorithm 2.4 with input $I, 2^e$ to obtain $\beta_1, \beta_2 \in I$ and $u, v \in \mathbb{N}^*$ and $f \leq e$ such that $\gcd(uq_I(\beta_1), vq_I(\beta_2)) = 1$ and $uq_I(\beta_1) + vq_I(\beta_2) = 2^f$, with $d_1 := \operatorname{nrd}(\beta_1)/\operatorname{nrd}(I)$ and $d_2 := \operatorname{nrd}(\beta_2)/\operatorname{nrd}(I)$;
Call Algorithm 2.3 with input u to compute the image $(E_u, \varphi_u(P_0), \varphi_u(Q_0))$ of a u-isogeny $\varphi_u : E_0 \longrightarrow E_u$;
Call Algorithm 2.3 with input v to compute the image $(E_v, \varphi_v(P_0), \varphi_v(Q_0))$ of a v-isogeny $\varphi_v : E_0 \longrightarrow E_v$;
$\theta \leftarrow \beta_2\overline{\beta_1}/\operatorname{nrd}(I)$;
Find $a, b, c, d \in \mathbb{Z}/2^e\mathbb{Z}$ such that $\varepsilon_0(\theta)(P_0) = [a]P_0 + [b]Q_0$ and $\varepsilon_0(\theta)(Q_0) = [c]P_0 + [d]Q_0$;
$T_1 \leftarrow ([2^{e-f}d_1]\varphi_u(P_0), [2^{e-f}a]\varphi_v(P_0) + [2^{e-f}b]\varphi_v(Q_0))$;
$T_2 \leftarrow ([2^{e-f}d_1]\varphi_u(Q_0), [2^{e-f}c]\varphi_v(P_0) + [2^{e-f}d]\varphi_v(Q_0))$;
Compute $\Phi : E_u \times E_v \longrightarrow E_1 \times E_2$ of kernel $\langle T_1, T_2 \rangle$; // see Section 6.5
$(\pm I(S), \pm I'(S)) \leftarrow \Phi(\varphi_u(S), 0)$ for $S \in \{P_0, Q_0, P_0 - Q_0\}$;
Lift $(I(P_0), I(Q_0))$ from $\pm I(P_0), \pm I(Q_0), \pm I(P_0 - Q_0)$;
Lift $(I'(P_0), I'(Q_0))$ from $\pm I'(P_0), \pm I'(Q_0), \pm I'(P_0 - Q_0)$;
$\mu \equiv 1/u \mod 2^e$, $\delta_1 \leftarrow 1/d_1 \mod 2^e$;
if $e_{2^e}(I(P_0), I(Q_0)) = e_{2^e}(P_0, Q_0)^{u^2 d_1}$ **then**
  $E_I \leftarrow E_1$, $\varphi_1(P_0) \leftarrow [\mu]I(P_0)$, $\varphi_1(Q_0) \leftarrow [\mu]I(Q_0)$;
else
  $E_I \leftarrow E_2$, $\varphi_1(P_0) \leftarrow [\mu]I'(P_0)$, $\varphi_1(Q_0) \leftarrow [\mu]I'(Q_0)$;
end
Find $a, b, c, d \in \mathbb{Z}/2^e\mathbb{Z}$ such that $\varepsilon_0(\beta_1)(P_0) = [a]P_0 + [b]Q_0$ and $\varepsilon_0(\beta_1)(Q_0) = [c]P_0 + [d]Q_0$;
$\varphi_I(P_0) \leftarrow [a\delta_1]\varphi_1(P_0) + [b\delta_1]\varphi_1(Q_0)$, $\varphi_I(Q_0) \leftarrow [c\delta_1]\varphi_1(P_0) + [d\delta_1]\varphi_1(Q_0)$;
return $(E_I, \varphi_I(P_0), \varphi_I(Q_0))$;

2.4.3 Improving the norm equation step success probability

We have seen that Algorithm 2.5 may fail because of difficulties to solve Eq. (2.12) in some very rare cases. In the following, we explain a method proposed in [AAA+25] to overcome this issue.

Starting from other curves than E_0

In Algorithm 2.3, we construct isogenies $\varphi_u : E_0 \longrightarrow E_u$ of given odd degree u by exploiting the structure of $\mathcal{O}_0 \simeq \operatorname{End}(E_0)$. More precisely, we exploit the fact that $\mathcal{O}_0$ contains a suborder of the form $\mathbb{Z}[i] \oplus j\mathbb{Z}[i]$ to solve norm equations more easily with [Ler22, Algorithm 2]. As Antonin Leroux remarked in his PhD thesis, [Ler22, Algorithm 2] applies to all *special extremal orders* containing an order of small discriminant.

Definition 2.4.6. Let $\delta \in \mathbb{N}^*$. We say that a maximal order $\mathcal{O} \subset \mathcal{B}_{p,\infty}$ is a *δ-special extremal order* if it contains a suborder of the form $\mathfrak{O} + j\mathfrak{O}$, where $\mathfrak{O}$ is a quadratic imaginary order of discriminant $-\delta$.

If $E/\mathbb{F}_{p^2}$ is a supersingular elliptic curve of known endomorphism ring isomorphic to a δ-special extremal order $\mathcal{O}$, then for any odd integer u such that $u(2^e - u) = \Omega(p\log(p)h(-\delta))$, a u-isogeny $\varphi_u : E \longrightarrow E_u$ can be constructed with the techniques from Algorithm 2.3.

More freedom in Step 1

Let $n \in \mathbb{N}^*$. Assume that we have precomputed for $t \in [\![1\ ;\ n]\!]$:

- A δ_t-special extremal order $\mathcal{O}_t \subset \mathcal{B}_{p,\infty}$ for a small discriminant δ_t;
- A supersingular elliptic curve $E_t/\mathbb{F}_{p^2}$ whose endomorphism ring is isomorphic to $\mathcal{O}_t$;
- An isogeny $\varphi_t : E_0 \longrightarrow E_t$ of odd degree n_t and its associated ideal J_t connecting $\mathcal{O}_0$ and $\mathcal{O}_t$.

As in Section 2.4.2, we want to translate a left $\mathcal{O}_0$-ideal I into an isogeny $\varphi_I : E_0 \longrightarrow E_I$ and we follow similar steps. However, we use the additional freedom offered by our precomputations to enhance the success probability of Step 1. Namely, we still look for integral ideals I_1, I_2 of respective norms d_1, d_2, $f \leq e$ and $u, v \in \mathbb{N}^*$ such that $\gcd(ud_1, vd_2) = 1$ and

$$ud_1 + vd_1 = 2^f.$$

However, we no longer require I_1 and I_2 to be left $\mathcal{O}_0$-ideals equivalent to I. We allow I_1 to be a left $\mathcal{O}_s$-ideal equivalent to $\overline{J}_s \cdot I$ and I_2 to be a left $\mathcal{O}_t$-ideal equivalent to $\overline{J}_t \cdot I$ for some $s, t \in [\![0\ ;\ n]\!]$ (with the convention $J_0 = \mathcal{O}_0$). In practice in SQIsign NIST round 2 implementation, with $n = 6$ or 7 additional precomputed special extremal orders, the failure probability in Step 1 becomes completely negligible. The new method is summarised in Algorithm 2.6.

A variant of Algorithm 2.5

Once Step 1 is complete, we may evaluate a u-isogeny $\varphi_u : E_s \longrightarrow E_u$ on $(\varphi_s(P_0), \varphi_s(Q_0))$ and a v-isogeny $\varphi_v : E_t \longrightarrow E_v$ on $(\varphi_t(P_0), \varphi_t(Q_0))$, where (P_0, Q_0) is a basis of $E_0[2^e]$. We may then consider the following (ud_1, vd_1)-isogeny diamond inspired from Eq. (2.16):

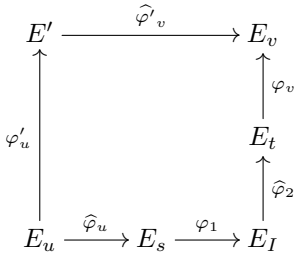

where $\varphi_1 : E_s \longrightarrow E_I$ and $\varphi_2 : E_t \longrightarrow E_I$ are the isogeny respectively associated to I_1 and I_2. Then, Kani's lemma (Lemma 2.2.6) yields a 2^f-isogeny:

$$\Phi := \begin{pmatrix} \varphi_1 \circ \widehat{\varphi}_u & \varphi_2 \circ \widehat{\varphi}_v \\ -\varphi'_u & \varphi'_v \end{pmatrix} : E_u \times E_v \longrightarrow E_I \times E',$$

with kernel:

$$\ker(\Phi) = \{([d_1]\varphi_u(P), \varphi_v \circ \widehat{\varphi}_2 \circ \varphi_1(P)) \mid P \in E_s[2^f]\}.$$

To find $I_1 \sim \overline{J}_s \cdot I$ and $I_2 \sim I$ in Step 1, we have found $\beta_1 \in \overline{J}_s \cdot I$ and $\beta_2 \in \overline{J}_t \cdot I$ such that $J_s \cdot I_1 = I\overline{\beta_1}/\operatorname{nrd}(I)$ and $J_t \cdot I_2 = I\overline{\beta_2}/\operatorname{nrd}(I)$. Then, by Lemma 1.2.24, we have $\widehat{\varphi}_s \circ \widehat{\varphi}_1 \circ \varphi_I = \varepsilon_0(\beta_1)$ and $\widehat{\varphi}_t \circ \widehat{\varphi}_2 \circ \varphi_I = \varepsilon_0(\beta_2)$. It follows that $[\operatorname{nrd}(I)n_s]\varphi_1 = \varphi_I \circ \varepsilon_0(\overline{\beta}_1) \circ \widehat{\varphi}_s$ and $[\operatorname{nrd}(I)n_t]\widehat{\varphi}_2 = \varphi_t \circ \varepsilon_0(\beta_2) \circ \widehat{\varphi}_I$, so that:

$$[n_s n_t]\widehat{\varphi}_2 \circ \varphi_1 = \varphi_t \circ \varepsilon_0(\theta) \circ \widehat{\varphi}_s,$$

with $\theta := \beta_2\overline{\beta}_1/\operatorname{nrd}(I)$. Since $n_s = \deg(\varphi_s)$ and $n_t = \deg(\varphi_t)$ are odd, we then obtain that $\ker(\Phi)$ is generated by T_1 and T_2 given by:

Algorithm 2.6: Step 1 modified: Finding suitable ideals I_1, I_2 and a coin problem solution (u, v).

Data: An ideal $I \subseteq \mathcal{O}_0 \simeq \mathrm{End}(E_0)$, an integer $N = \Theta(p)$, a bound $B = \Theta(\log(p))$ and precomputed data $(\mathcal{O}_t, J_t, n_t)_{0 \leq t \leq n}$.

Result: $s, t \in [\![0\ ;\ n]\!]$, $\beta_1 \in \overline{J}_s \cdot I$, $\beta_2 \in \overline{J}_t \cdot I$ such that $d_1 := \mathrm{nrd}(\beta_1)/(n_s \,\mathrm{nrd}(I))$ and $d_2 := \mathrm{nrd}(\beta_2)/(n_t \,\mathrm{nrd}(I))$ are coprime and coprime with N, along with $u, v \in \mathbb{N}^*$ and $M | N$ such that $\gcd(ud_1, vd_2) = 1$ and $ud_1 + vd_2 = M$.

```
for t = 0 to n do
    Compute a Minkowski reduced basis (α_{t,1}, ··· , α_{t,4}) of J̄_t · I ;        // [NS09, Fig. 3-4]
end
for s = 0 to n do
    for (x_1, ··· , x_4) ∈ [[−B ; B]]^4 do
        β_1 ← Σ_{j=1}^4 x_j α_{s,j};
        d_1 ← nrd(β_1)/(n_s nrd(I));
        if gcd(d_1, N) = 1 then
            for t = 0 to n do
                for (y_1, ··· , y_4) ∈ [[−B ; B]]^4 do
                    β_2 ← Σ_{j=1}^4 y_j α_{t,j};
                    d_2 ← nrd(β_2)/(n_t nrd(I));
                    if gcd(d_1, d_2) = gcd(d_2, N) = 1 then
                        Apply extended Euclide's algorithm to find u_0, v_0 ∈ ℤ such that
                         u_0 d_1 + v_0 d_2 = 1;
                        if [−N u_0/d_2, N v_0/d_1] ∩ ℤ ≠ ∅ then
                            Take k ← ⌈−N u_0/d_2⌉;
                            u ← N u_0 + k d_2, v ← N v_0 − k d_1;
                            d ← gcd(u, v, N);
                            u ← u/d, v ← v/d, M ← N/d;
                            return s, t, β_1, β_2, M, u, v;
                        end
                    end
                end
            end
        end
    end
end
```

$$T_1 := ([2^{e-f} d_1]\varphi_u \circ \varphi_s(P_0), [2^{e-f}\eta_t]\varphi_v \circ \varphi_t \circ \varepsilon_0(\theta)(P_0)),$$
$$\text{and} \quad T_2 := ([2^{e-f} d_1]\varphi_u \circ \varphi_s(Q_0), [2^{e-f}\eta_t]\varphi_v \circ \varphi_t \circ \varepsilon_0(\theta)(Q_0)),$$

with $\eta_t \equiv 1/n_t \mod 2^e$.

By Theorem 2.2.13, we can then apply the algorithms from Section 6.5 to compute $\Phi : E_u \times E_v \longrightarrow E_I \times E'$, and identify a potential swap between E_I and E' with Weil pairing computations, as explained previously. Finally, to evaluate $(\varphi_I(P_0), \varphi_I(Q_0))$, we evaluate $([u]\varphi_1 \circ \varphi_s(P_0), *) = \Phi(\varphi_u \circ \varphi_s(P_0), 0)$ and $([u]\varphi_1 \circ \varphi_s(Q_0), *) = \Phi(\varphi_u \circ \varphi_s(Q_0), 0)$ and use the fact that $[n_s d_1]\varphi_I = \varphi_1 \circ \varphi_s \circ \varepsilon_0(\beta_1)$ that follows from Lemma 1.2.24. Algorithm 2.7 follows.

How to precompute starting curves with special extremal endomorphism rings

Now, we explain how the data $(E_t, \mathcal{O}_t, J_t, n_t, \varphi_t(P_0), \varphi_t(Q_0))_{0 \leq t \leq n}$ is precomputed. We proceed as in [EPSV24; AAA+25]. We chose a sequence of small primes $q_1 < \cdots < q_n$ congruent to 3 $\mod 8$ and such that $(-p/q_t) = 1$ for $t \in [\![1\ ;\ n]\!]$. As in [Ibu82], we define for all $t \in [\![1\ ;\ n]\!]$ the $4q_t$-special extremal order:

$$\mathcal{O}_t := \left\langle 1, i', \frac{1+j}{2}, \frac{(r+j)i'}{2q_t} \right\rangle \subseteq \mathcal{B}_{p,\infty},$$

Algorithm 2.7: ideal-to-isogeny from E_0 with more freedom in Step 1.

Data: A left $\mathcal{O}_0$-ideal I and a basis (P_0, Q_0) of $E_0[2^e]$ and the precomputed data $(E_t, \mathcal{O}_t, J_t, n_t, \varphi_t(P_0), \varphi_t(Q_0))_{0\leq t\leq n}$.
Result: The image $(E_I, \varphi_I(P_0), \varphi_I(Q_0))$ of the isogeny $\varphi_I : E_0 \to E_I$ associated to I.

Call Algorithm 2.6 to find $s, t \in [\![0\ ;\ n]\!]$, $\beta_1 \in \overline{J}_s \cdot I$, $\beta_2 \in \overline{J}_t \cdot I$, $u, v \in \mathbb{N}^*$ and $f \leq e$ such that $\gcd(ud_1, vd_2) = 1$ and $ud_1 + vd_2 = 2^f$, with $d_1 := \mathrm{nrd}(\beta_1)/(n_s \,\mathrm{nrd}(I))$ and $d_2 := \mathrm{nrd}(\beta_2)/(n_t \,\mathrm{nrd}(I))$;
Compute the image $(E_u, \varphi_u \circ \varphi_s(P_0), \varphi_u \circ \varphi_s(Q_0))$ of a u-isogeny $\varphi_u : E_s \longrightarrow E_u$;
Compute the image $(E_v, \varphi_v \circ \varphi_t(P_0), \varphi_v \circ \varphi_t(Q_0))$ of a v-isogeny $\varphi_v : E_t \longrightarrow E_v$;
$\theta \leftarrow \beta_2\overline{\beta_1}/\,\mathrm{nrd}(I)$;
Find $a, b, c, d \in \mathbb{Z}/2^e\mathbb{Z}$ such that $\varepsilon_0(\theta)(P_0) = [a]P_0 + [b]Q_0$ and $\varepsilon_0(\theta)(Q_0) = [c]P_0 + [d]Q_0$;
$\eta_s \leftarrow 1/n_s \mod 2^e$, $\eta_t \leftarrow 1/n_t \mod 2^e$, $\mu \equiv 1/u \mod 2^e$, $\delta_1 \leftarrow 1/d_1 \mod 2^e$;
$T_1 \leftarrow ([2^{e-f}d_1]\varphi_u \circ \varphi_s(P_0), [2^{e-f}a\eta_t]\varphi_v \circ \varphi_t(P_0) + [2^{e-f}b\eta_t]\varphi_v \circ \varphi_t(Q_0))$;
$T_2 \leftarrow ([2^{e-f}d_1]\varphi_u \circ \varphi_s(Q_0), [2^{e-f}c\eta_t]\varphi_v \circ \varphi_t(P_0) + [2^{e-f}d\eta_t]\varphi_v \circ \varphi_t(Q_0))$;
Compute $\Phi : E_u \times E_v \longrightarrow E_1 \times E_2$ of kernel $\langle T_1, T_2\rangle$; // see Section 6.5
$(\pm I(S), \pm I'(S)) \leftarrow \Phi(\varphi_u \circ \varphi_s(S), 0)$ for $S \in \{P_0, Q_0, P_0 - Q_0\}$;
Lift $(I(P_0), I(Q_0))$ from $\pm I(P_0), \pm I(Q_0), \pm I(P_0 - Q_0)$;
Lift $(I'(P_0), I'(Q_0))$ from $\pm I'(P_0), \pm I'(Q_0), \pm I'(P_0 - Q_0)$;
if $e_{2^e}(I(P_0), I(Q_0)) = e_{2^e}(P_0, Q_0)^{u^2 d_1 n_s}$ **then**
  $E_I \leftarrow E_1$, $\varphi_1 \circ \varphi_s(P_0) \leftarrow [\mu]I(P_0)$, $\varphi_1 \circ \varphi_s(Q_0) \leftarrow [\mu]I(Q_0)$;
else
  $E_I \leftarrow E_2$, $\varphi_1 \circ \varphi_s(P_0) \leftarrow [\mu]I'(P_0)$, $\varphi_1 \circ \varphi_s(Q_0) \leftarrow [\mu]I'(Q_0)$;
end
Find $a, b, c, d \in \mathbb{Z}/2^e\mathbb{Z}$ such that $\varepsilon_0(\beta_1)(P_0) = [a]P_0 + [b]Q_0$ and $\varepsilon_0(\beta_1)(Q_0) = [c]P_0 + [d]Q_0$;
$\varphi_I(P_0) \leftarrow [a\delta_1\eta_s]\varphi_1 \circ \varphi_s(P_0) + [b\delta_1\eta_s]\varphi_1 \circ \varphi_s(Q_0)$;
$\varphi_I(Q_0) \leftarrow [c\delta_1\eta_s]\varphi_1 \circ \varphi_s(P_0) + [d\delta_1\eta_s]\varphi_1 \circ \varphi_s(Q_0)$;
return $(E_I, \varphi_I(P_0), \varphi_I(Q_0))$;

where $r^2 \equiv -p \mod 4q_t$ and $i'^2 = -q_t$. The quaternion i' is defined as $i' := i(x + yj)$ with $x, y \in \mathbb{Q}$ such that $x^2 + py^2 = q_t$ [EPSV24, Lemma 10]. To find such $x, y \in \mathbb{Q}$, we solve the Legendre equation $x'^2 + py'^2 - q_t z'^2 = 0$ with variables $x', y', z' \in \mathbb{Z}^*$ and set $x := x'/z'$ and $y := y'/z'$. Since $(-p/q_t) = 1$, Legendre's theorem below ensures this equation admits a solution. This solution can be found polynomial time in $\log(p)$ with Simon's algorithm [Sim05, Algorithm 3.4].

Theorem 2.4.7 (Legendre, 1785)**.** *Let $a, b, c \in \mathbb{Z}^*$. Assume that $-bc \mod a$, $-ac \mod b$ and $-ab \mod c$ are quadratic residue, then $ax^2 + by^2 + cz^2 = 0$ admits a non-zero solution $(x, y, z) \in \mathbb{Z}^3$.*

Proof. See [DH48]. □

Then, for all $t \in [\![1\ ;\ n]\!]$, we compute a connecting ideal J_t between $\mathcal{O}_0$ and $\mathcal{O}_t$ as $J_t := N \cdot \mathcal{O}_0 \cdot \mathcal{O}_t$, where N is a common denominator. We can then apply Algorithm 2.5 to J_t to compute $(E_t, \varphi_t(P_0), \varphi_t(Q_0))$.

Example 2.4.8. In SQIsign round 2 NIST submission, at NIST-I security level, $p = 5 \cdot 2^{248} - 1$ and six $4q_t$-special extremal orders with $q_t \in \{5, 17, 37, 41, 53, 97\}$ have been used [AAA+25, Appendix B].

2.5 Class group action with 4-dimensional isogenies

In this section, we present the PEGASIS algorithm (practical effective group action using 4-dimensional isogenies) introduced in [DEF+25]. The goal of this algorithm inspired by Clapoti [PR23] is to compute the class group action on oriented elliptic curves by any ideal. Indeed, with previous techniques, it was only possible to compute efficiently the action by products of ideals of small norms. For that reason, the class group action on oriented elliptic curves was called a *restricted effective (or cryptographic) group action* (REGA), unlike an *effective (or cryptographic) group action* (EGA) where the action by any group element is efficient.

For some cryptographic applications, an EGA is more suited than a REGA because the action by group elements with uniform distribution is required. A REGA can still be used in that case, at the expense of a precomputation of the group structure and the lattice of relations of group generators whose action can be computed efficiently. It is the case for instance, in the digital signature scheme CSI-FiSh [BKV19] relying on the action by $\mathrm{Cl}(\mathbb{Z}[\sqrt{-p}])$. The structure computation of $\mathrm{Cl}(\mathbb{Z}[\sqrt{-p}])$ may become impractical as p grows, and as a consequence, CSI-FiSh is not scalable to higher security levels. These problems do not appear with EGAs.

PEGASIS provides an EGA for any kind of orientation of supersingular elliptic curves but has only been implemented in the CSIDH/CSURF context. In the following, for the sake of clarity, we only consider the CSURF ideal class group action [CD20] but the algorithmic approach is general (we refer to [DEF+25] for the general case). As in Section 2.4, let p be a prime of the form $p = c \cdot 2^e - 1$ with $c \in \mathbb{N}^*$ a small and odd integer and $e \in \mathbb{N}$ of several hundreds. The CSURF group action is given by the maximal order $\mathfrak{O} := \mathbb{Z}[(1+\sqrt{-p})/2]$ of $\mathbb{Q}(\sqrt{-p})$. If E is a (primitively) $\mathfrak{O}$-oriented supersingular elliptic curve, then E admits a Montgomery model over $\mathbb{F}_p$ and $\mathrm{End}_{\mathbb{F}_p}(E) \simeq \mathfrak{O}$, where $\sqrt{-p}$ identifies with the p-th Frobenius endomorphism. We shall see that the Montgomery $(X:Z)$-coordinates of the 2^{e-1}-torsion of E are $\mathbb{F}_p$-rational so most of our arithmetic operations will take place over $\mathbb{F}_p$ in this context.

Let $\mathfrak{a} \subseteq \mathfrak{O}$ be an ideal and $E/\mathbb{F}_p$ be an $\mathfrak{O}$-oriented curve. We explain how the compute the action $E_\mathfrak{a} := \mathfrak{a} \cdot E$ of $\mathfrak{a}$ on E, which is the codomain of the isogeny $\varphi_\mathfrak{a} : E \longrightarrow E_\mathfrak{a}$ of kernel $E[\mathfrak{a}]$. As suggested in [PR23] and similarly to Section 2.4, we proceed as follows:

1. We solve the following problem.

Problem 2.5.1. Given an ideal $\mathfrak{a} \subseteq \mathfrak{O}$ and an integer f_{max}. Find ideals $\mathfrak{b}, \mathfrak{c} \subseteq \mathfrak{O}$ and integers N_1, N_2, u, v, f such that:

(i) The ideals $\mathfrak{b}, \mathfrak{c}$ are equivalent to $\mathfrak{a}$ and of the form $\mathfrak{b} := \mathfrak{b}_e \cdot \mathfrak{b}_k$ and $\mathfrak{c} := \mathfrak{c}_e \cdot \mathfrak{c}_k$, where $\mathfrak{b}_e$ and $\mathfrak{c}_e$ are product of small prime ideals and $N_1 := N(\mathfrak{b}_k)$ and $N_2 := N(\mathfrak{c}_k)$ are odd and coprime.

(ii) The integers $u, v \in \mathbb{N}^*$ are such that $\gcd(uN_1, vN_2) = 1$ and

$$uN_1 + vN_2 = 2^f. \tag{2.17}$$

(iii) We have $f \leq f_{max} - f_1 - f_2$, where $f_1 := v_2(N(\mathfrak{b}'_e))$, $f_2 := v_2(N(\mathfrak{c}'_e))$, v_2 being the 2-adic valuation, $\mathfrak{b}_e = \mathfrak{f}_0 \cdot \mathfrak{b}'_e$, $\mathfrak{c}_e = \mathfrak{f}_0 \cdot \mathfrak{c}'_e$ and $\mathfrak{f}_0$ is the greatest common factor of $\mathfrak{b}_e$ and $\mathfrak{c}_e$.

2. We compute and evaluate the isogenies $\varphi_{\mathfrak{b}_e} : E \longrightarrow E_1 := E_{\mathfrak{b}_e}$ and $\varphi_{\mathfrak{c}_e} : E \longrightarrow E_2 := E_{\mathfrak{c}_e}$ associated to $\mathfrak{b}_e$ and $\mathfrak{c}_e$ with standard techniques from Elkies [Elk98].

3. We compute and evaluate the 2-dimensional isogenies $\Phi_u : E_1^2 \longrightarrow A_u$ and $\Phi_v : E_2^2 \longrightarrow A_v$ of respective polarised degrees u and v.

4. Using the data from previous steps, we compute a 4-dimensional 2^f-isogeny $F : A_u \times A_v \longrightarrow E_\mathfrak{a}^2 \times A$ obtained from Kani's lemma and extract $E_\mathfrak{a}$ from the codomain $E_\mathfrak{a}^2 \times A$.

We now give more details on each one of the above steps.

2.5.1 Step 1: the norm equation

To simplify the computation of 2-dimensional isogenies of polarised degrees u and v in Step 3 (see Section 2.5.2), we require u and v to be sums of squares up to a small factor. In the following, we fix a set $\mathfrak{B}$ of small *Elkies primes i.e.* of primes that split completely in $\mathfrak{O} = \mathbb{Z}[(1+\sqrt{-p})/2]$ including 2. Instead of solving Problem 2.5.1, we solve the following problem.

Problem 2.5.2. Given an ideal $\mathfrak{a} \subseteq \mathfrak{O}$ and an integer f_{max}. Find ideals $\mathfrak{b}, \mathfrak{c} \subseteq \mathfrak{O}$ and integers $N_1, N_2, u, v, f, x_u, y_u, x_v, y_v, g_u, g_v$ such that:

(i) The ideals $\mathfrak{b}, \mathfrak{c}$ are equivalent to $\mathfrak{a}$ and of the form $\mathfrak{b} := \mathfrak{b}_e \cdot \mathfrak{b}_k$ and $\mathfrak{c} := \mathfrak{c}_e \cdot \mathfrak{c}_k$, where $\mathfrak{b}_e$ and $\mathfrak{c}_e$ are product of prime ideals lying above primes of $\mathfrak{B}$ and $N_1 := N(\mathfrak{b}_k)$ and $N_2 := N(\mathfrak{c}_k)$ are coprime and have no factor in $\mathfrak{B}$.

(ii) The integers $u, v \in \mathbb{N}^*$ and $f \leq f_{max}$ are such that $\gcd(uN_1, vN_2) = 1$ and

$$uN_1 + vN_2 = 2^f, \tag{2.18}$$

with $u := g_u(x_u^2 + y_u^2)$, $v := g_v(x_v^2 + y_v^2)$, $x_u, y_u, x_v, y_v \in \mathbb{Z}$ and $g_u, g_v \in \mathbb{N}^*$ are products of primes in $\mathfrak{B}$. We say that such integers u and v are $\mathfrak{B}$-*good*.

(iii) We have $f \leq f_{max} - f_1 - f_2$, where $f_1 := v_2(N(\mathfrak{b}'_e))$, $f_2 := v_2(N(\mathfrak{c}'_e))$, $\mathfrak{b}_e = \mathfrak{f}_0 \cdot \mathfrak{b}'_e$, $\mathfrak{c}_e = \mathfrak{f}_0 \cdot \mathfrak{c}'_e$ and $\mathfrak{f}_0$ is the greatest common factor of $\mathfrak{b}_e$ and $\mathfrak{c}_e$.

Remark 2.5.3. Since we need the 2^{f+2}-torsion to compute the 4-dimensional 2^f-isogeny in Step 4 (Theorem 2.2.12) and we want to work with $\mathbb{F}_p$-arithmetic as much as possible, we need $f + 2 \leq e - 1$ so we solve Problem 2.5.2 with $f_{max} := e - 3$.

The method to solve Problem 2.5.2 is quite similar to Step 1 in Section 2.4.2. We find suitable ideals $\mathfrak{b}, \mathfrak{c} \sim \mathfrak{a}$ by sampling $\beta, \gamma \in \mathfrak{a}$ of small norms and setting $\mathfrak{b} := \mathfrak{a}\overline{\beta}/N(\mathfrak{a})$ and $\mathfrak{c} := \mathfrak{a}\overline{\gamma}/N(\mathfrak{a})$. To sample β, γ of small norms, we find a Lagrange reduced basis (α_1, α_2) of $\mathfrak{a}$, sample at random x_1, x_2, y_1, y_2 in a small interval $[\![-m\ ;\ m]\!]$ and set $\beta := x_1\alpha_1 + x_2\alpha_2$ and $\gamma := y_1\alpha_1 + y_2\alpha_2$. By the following lemma (Lemma 2.5.4), we expect the norms of α_1, α_2 to be close to $\sqrt{p}N(\mathfrak{a})$ in most cases so we expect $N(\mathfrak{b}), N(\mathfrak{c}) = \Theta(\sqrt{p})$. Since $2^e = O(p)$, this would make the coin equation $uN(\mathfrak{b}) + vN(\mathfrak{c}) = 2^f$ with $f \leq e - 3$ tight to solve, and even more with the additional requirement that u and v are $\mathfrak{B}$-good.

Lemma 2.5.4. *If $\mathfrak{a}$ is an integral $\mathfrak{O}$-ideal and (α_1, α_2) is a Lagrange reduced basis of $\mathfrak{a}$, then*

$$\frac{pN(\mathfrak{a})^2}{\pi^2} \leq N(\alpha_1)N(\alpha_2) \leq \frac{4pN(\mathfrak{a})^2}{\pi^2}.$$

Proof. We proceed as in the proof of Lemma 2.3.2. Consider the canonical isomorphism $\iota : x + iy \in \mathbb{C} \longmapsto (x, y) \in \mathbb{R}^2$. When restricted to $\mathbb{Q}(\sqrt{-p})$, ι is an isometry in the following sense $\|\iota(\alpha)\|^2 = N(\alpha)$ for all $\alpha \in \mathbb{Q}(\sqrt{-p})$, where $\|.\|$ is the Euclidean norm of $\mathbb{R}^2$. By Minkowski's second theorem, the successive minima $\lambda_1 \leq \lambda_2$ of the lattice $\iota(\mathfrak{a})$ satisfy

$$\frac{2^2}{2!}\frac{\mathrm{Covol}(\iota(\mathfrak{a}))}{\mathrm{Vol}(B(0,1))} = \frac{2}{\pi}\mathrm{Covol}(\iota(\mathfrak{a})) \leq \lambda_1\lambda_2 \leq 2^2\frac{\mathrm{Covol}(\iota(\mathfrak{a}))}{\mathrm{Vol}(B(0,1))} = \frac{4}{\pi}\mathrm{Covol}(\iota(\mathfrak{a})). \tag{2.19}$$

But we have

$$\begin{aligned}\mathrm{Covol}(\iota(\mathfrak{O})) &= \begin{vmatrix} \langle \iota(1), \iota(1)\rangle & \left\langle \iota(1), \iota\left(\frac{1+i\sqrt{p}}{2}\right)\right\rangle \\ \left\langle \iota\left(\frac{1+i\sqrt{p}}{2}\right), \iota(1)\right\rangle & \left\langle \iota\left(\frac{1+i\sqrt{p}}{2}\right), \iota\left(\frac{1+i\sqrt{p}}{2}\right)\right\rangle \end{vmatrix}^{\frac{1}{2}} \\ &= \begin{vmatrix} 1 & \frac{1}{2} \\ \frac{1}{2} & \frac{1+p}{4} \end{vmatrix}^{\frac{1}{2}} = \left|\frac{1+p}{4} - \frac{1}{4}\right|^{\frac{1}{2}} = \frac{\sqrt{p}}{2}\end{aligned}$$

We then have

$$\mathrm{Covol}(\iota(\mathfrak{a})) = [\mathfrak{O} : \mathfrak{a}]\,\mathrm{Covol}(\iota(\mathfrak{O})) = \frac{N(\mathfrak{a})\sqrt{p}}{2}$$

It follows by Eq. (2.19) that

$$\frac{\sqrt{p}N(\mathfrak{a})}{\pi} \leq \lambda_1\lambda_2 \leq \frac{2\sqrt{p}N(\mathfrak{a})}{\pi}$$

Taking the square of this inequality, we obtain the desired result. □

To make the coin equation easier to solve, we factor out $\mathfrak{b}$ and $\mathfrak{c}$ by products of ideals lying above primes of $\mathfrak{B}$, that we denote $\mathfrak{b}_e$ and $\mathfrak{c}_e$ respectively, so that $\mathfrak{b} = \mathfrak{b}_e \cdot \mathfrak{b}_k$ and $\mathfrak{c} = \mathfrak{c}_e \cdot \mathfrak{c}_k$. We then set $N_1 := N(\mathfrak{b}_k)$, $N_2 := N(\mathfrak{c}_k)$ and solve the coin equation $uN_1 + vN_2 = 2^f$ (Eq. (2.18)) which is less tight. The method to solve Eq. (2.18) is very similar to Algorithms 2.4 and 2.6, except that we have to check that u and v are $\mathfrak{B}$-good and find their decompositions $u := g_u(x_u^2 + y_u^2)$, $v := g_v(x_v^2 + y_v^2)$. This is done by a trial and division method and Cornacchia's algorithm (as in Section 2.3.1). We summarise Step 1 in Algorithm 2.8.

Algorithm 2.8: Step 1: Finding a solution to Problem 2.5.2.

Data: An ideal $\mathfrak{a} \subseteq \mathfrak{O}$, two bounds $m, q = \Theta(\log(p))$ and a set of small primes $\mathfrak{B}$ splitting in $\mathfrak{O}$ (including 2).
Result: A solution $\mathfrak{b}, \mathfrak{c}, N_1, N_2, u, v, f, x_u, y_u, x_v, y_v, g_u, g_v$ to Problem 2.5.2 with respect to $\mathfrak{a}$ and $f_{max} := e - 3$.

Find a Lagrange reduced basis (α_1, α_2) of $\mathfrak{a}$;
for $(x_1, x_2) \in [\![-m\ ;\ m]\!]^2$ **do**
  $\beta_1 \leftarrow x_1\alpha_1 + x_2\alpha_2$, $\mathfrak{b} \leftarrow \mathfrak{a}\overline{\beta_1}/N(\mathfrak{a})$;
  Factor $N(\mathfrak{b}) := N(\beta_1)/N(\mathfrak{a})$ into $N(\mathfrak{b}) = M_1 \cdot N_1$, where all prime factors of M_1 lie in $\mathfrak{B}$ and none of the prime factors of N_1 lie in $\mathfrak{B}$;
  for $(y_1, y_2) \in [\![-m\ ;\ m]\!]^2$ **do**
    $\beta_2 \leftarrow y_1\alpha_1 + y_2\alpha_2$, $\mathfrak{c} \leftarrow \mathfrak{a}\overline{\beta_2}/N(\mathfrak{a})$;
    Factor $N(\mathfrak{c}) := N(\beta_2)/N(\mathfrak{a})$ into $N(\mathfrak{c}) = M_2 \cdot N_2$, where all prime factors of M_2 lie in $\mathfrak{B}$ and none of the prime factors of N_2 lie in $\mathfrak{B}$;
    if $\gcd(N_1, N_2) = 1$ **then**
      Apply extended Euclide's algorithm to find $u_0, v_0 \in \mathbb{Z}$ such that $u_0N_1 + v_0N_2 = 1$;
      $\mathfrak{d}_e \leftarrow \overline{\mathfrak{b}}_e\mathfrak{c}_e$, $d \leftarrow 0$;
      while $2\mathfrak{O}$ *is a factor of* $\mathfrak{d}_e$ **do**
        $d \leftarrow d + 1$;
        $\mathfrak{d}_e \leftarrow \mathfrak{d}_e/2$;
      end
      $f_1 \leftarrow v_2(M_1) - d$, $f_2 \leftarrow v_2(M_2) - d$;
      $f_m \leftarrow e - 3 - f_1 - f_2$;
      for $k \in [-2^{f_m}u_0/N_2, 2^{f_m}v_0/N_1] \cap \mathbb{Z}$ **do**
        $u \leftarrow 2^{f_m}u_0 + kN_2$, $v \leftarrow 2^{f_m}v_0 - kN_1$;
        $w \leftarrow \min(v_2(u), v_2(v))$;
        $u \leftarrow u/2^w$, $v \leftarrow v/2^w$, $f \leftarrow f_m - w$;
        Factor $u := g_u s_u^2 q_u r_u$ and $v := g_v s_v^2 q_v r_v$, where all their prime factors of g_u, g_v lie in $\mathfrak{B}$, all prime factors of s_u, s_v are $\leq q$ and all prime factors of q_u, q_v are $\equiv 1 \mod 4$ and $\leq q$;
        if r_u *and* r_v *are primes* $\equiv 1 \mod 4$ **then**
          Apply Cornacchia's algorithm to r_u and all prime factors of q_u to find $x'_u, y'_u \in \mathbb{Z}$ such that $q_u r_u = {x'}_u^2 + {y'}_u^2$;
          Apply Cornacchia's algorithm to r_v and all prime factors of q_v to find $x'_v, y'_v \in \mathbb{Z}$ such that $q_v r_v = {x'}_v^2 + {y'}_v^2$;
          $x_u \leftarrow s_u x'_u$, $y_u \leftarrow s_u y'_u$, $x_v \leftarrow s_v x'_v$, $y_v \leftarrow s_v y'_v$;
          return $\mathfrak{b}, \mathfrak{c}, N_1, N_2, u, v, f, x_u, y_u, x_v, y_v, g_u, g_v$;
        end
      end
    end
  end
end
return $\perp$;

Rerandomisation

For some ideals $\mathfrak{a}$, Problem 2.5.2 may not have a solution. Increasing the bound $\mathfrak{B}$ is always a possible solution, though this quickly becomes expensive and may be ineffective when $\mathfrak{a}$ is *unbalanced i.e.* when $N(\alpha_1) \ll \sqrt{p}N(\mathfrak{a})$ and $N(\alpha_2) \gg \sqrt{p}N(\mathfrak{a})$.

Another solution to this is to simply rerandomise the starting ideal $\mathfrak{a}$. Namely, If none of the sampled, equivalent ideals gives us a solution for Problem 2.5.2, we can multiply $\mathfrak{a}$ by a non-principal ideal $\mathfrak{l}$ for which the action is easy to evaluate (*e.g.* an ideal above $\ell = 2$, $\ell = 2$ being the smallest prime in $\mathfrak{B}$). We obtain a new ideal $\mathfrak{a}' = \mathfrak{a}\mathfrak{l}$ and try to solve Problem 2.5.2 for this ideal instead. For instance, if $\mathfrak{a}$ was unbalanced, we expect $\mathfrak{a}'$ to be balanced. If we manage to do so, we can simply

run Algorithm 2.8 on $\mathfrak{a}'$, using this solution, starting from the curve $E_{\bar{\mathfrak{l}}} = \bar{\mathfrak{l}} \cdot E$ instead of E, which is easy to compute with standard techniques by choice of $\mathfrak{l}$. Otherwise, we can rerandomise again by computing $\mathfrak{a}'' = \mathfrak{l}\mathfrak{a}'$ and running Algorithm 2.8 on $\mathfrak{a}''$, and so on. As soon as Problem 2.5.2 has a reasonable probability of being solved, this method takes care of failures in Algorithm 2.8 ($\perp$ returns) quite efficiently, making Step 1 faster at the cost of a few group action computations by $\bar{\mathfrak{l}}$.

2.5.2 Step 3: evaluating 2-dimensional isogenies of given polarised degree

In this section, we assume that we have computed $E_1 := E_{\mathfrak{b}_e}$ and $E_2 := E_{\mathfrak{c}_e}$ and we explain how to compute 2-dimensional isogenies $\Phi_u : E_1^2 \longrightarrow A_u$ and $\Phi_v : E_1^2 \longrightarrow A_v$ of respective polarised degrees u and v.

With u and v $\mathfrak{B}$-good

We assume that we have obtained a solution from Problem 2.5.2 in Step 1, so that u and v are $\mathfrak{B}$-good and can be written $u = g_u(x_u^2 + y_u^2)$ and $v = g_v(x_v^2 + y_v^2)$ with $x_u, y_u, x_v, y_v \in \mathbb{Z}$ and $g_u, g_v \in \mathbb{N}^*$ with prime factors in $\mathfrak{B}$. In that case, we may consider a g_u-isogeny $\varphi_u : E_1 \longrightarrow E_u$ given by the action of an ideal $\mathfrak{g}_u \subseteq \mathfrak{O}$ of norm g_u, and the 2-dimensional u-isogeny $\Phi_u := M_u \circ \mathrm{Diag}(\varphi_u, \varphi_u) : E_1^2 \longrightarrow E_u^2$, where M_u is the endomorphism:

$$M_u := \begin{pmatrix} x_u & -y_u \\ y_u & x_u \end{pmatrix} \in \mathrm{End}(E_u^2). \tag{2.20}$$

We may define $\Phi_v := M_v \circ \mathrm{Diag}(\varphi_v, \varphi_v) : E_2^2 \longrightarrow E_v^2$ similarly. Note that the abelian surfaces A_u and A_v are actually the elliptic products E_u^2 and E_v^2 respectively. In addition, the computation cost of Φ_u and Φ_v is restricted to 1-dimensional isogeny computations and elliptic curve arithmetic. This is not the case when u and v are not $\mathfrak{B}$-good.

With u or v not $\mathfrak{B}$-good

In practice, we always manage to solve Problem 2.5.2 in Step 1 so that u and v are $\mathfrak{B}$-good, by allowing more primes in $\mathfrak{B}$ as the parameter p grows. However, when p gets very big, solutions may be more difficult to find and Problem 2.5.2 may take more time to solve. Solutions of Problem 2.5.1 with either u or v (or both) not $\mathfrak{B}$-good may be worth considering to save time in Step 1 at the expense of Step 3. We now explain how to deal with this case, namely how to compute a u-isogeny $\Phi_u : E_1^2 \longrightarrow A_u$ when when u is an odd integer $< 2^{e-3}$ such that $u(2^{e-3} - u) = \Omega(p \log(p))$. The algorithmic approach is close to [NO24, Algorithm 2] (presented in Section 2.4.1) and involves the computation of a 4-dimensional isogeny, which is indeed more costly than the $\mathfrak{B}$-good approach presented above. Besides, the resulting codomain principally polarised abelian surface A_u is not an elliptic product in general.

The main idea follows from the following lemma.

Lemma 2.5.5. *Let $E/\mathbb{F}_p$ be an $\mathfrak{O}$-oriented curve, $a, b, c, d \in \mathbb{Z}$ and $N := a^2 + c^2 + p(b^2 + d^2)$. Assume that $N > 0$ and consider the endomorphism*

$$\alpha := \begin{pmatrix} [a] + [b]\pi_p & -[c] + [d]\pi_p \\ [c] + [d]\pi_p & [a] - [b]\pi_p \end{pmatrix} \in \mathrm{End}(E^2),$$

where $\pi_p : (x, y) \in E \longmapsto (x^p, y^p) \in E$ is the p-th Frobenius endomorphism. Then α is an N-isogeny.

Proof. By Lemma 2.2.4.(ii), the polarised dual of α (with respect to the naturally induced principal product polarisation on E^2) is

$$\widetilde{\alpha} = \begin{pmatrix} [a] + [b]\widehat{\pi}_p & [c] + [d]\widehat{\pi}_p \\ -[c] + [d]\widehat{\pi}_p & [a] - [b]\widehat{\pi}_p \end{pmatrix}.$$

Since E is supersingular and defined over $\mathbb{F}_p$, Hasse-Weil's bound implies that $\mathrm{Tr}(\pi_p) = 0$ so that $\pi_p^2 = -[p] = -\pi_p \circ \widehat{\pi}_p$ and $\widehat{\pi}_p = -\pi_p$. It follows that

$$\widetilde{\alpha} = \begin{pmatrix} [a] - [b]\pi_p & [c] - [d]\pi_p \\ -[c] - [d]\pi_p & [a] + [b]\pi_p \end{pmatrix}$$

and

$$\begin{aligned}\widetilde{\alpha}\circ\alpha &= \begin{pmatrix} [a]-[b]\pi_p & [c]-[d]\pi_p \\ -[c]-[d]\pi_p & [a]+[b]\pi_p \end{pmatrix} \circ \begin{pmatrix} [a]+[b]\pi_p & -[c]+[d]\pi_p \\ [c]+[d]\pi_p & [a]-[b]\pi_p \end{pmatrix} \\ &= \begin{pmatrix} [a^2]-[b^2]\pi_p^2+[c^2]-[d^2]\pi_p^2 & 0 \\ 0 & [c^2]-[d^2]\pi_p^2+[a^2]-[b^2]\pi_p^2 \end{pmatrix} \\ &= \begin{pmatrix} [a^2+c^2+p(b^2+d^2)] & 0 \\ 0 & [a^2+c^2+p(b^2+d^2)] \end{pmatrix} = [N]_{E^2}.\end{aligned}$$

This completes the proof. □

We fix an odd integer $u < 2^{e-3}$ and find $f \leq e-3$ as small as possible such that $u(2^f-u) = \Omega(p\log(p))$ (in practice $f \simeq e/2$). Then, we apply [Ler22, Algorithm 2] to find $a,b,c,d \in \mathbb{Z}$ such that

$$a^2+c^2+p(b^2+d^2) = \mathrm{nrd}(a+ci+bj+dij) = u(2^f-u),$$

and consider the endomorphism $\alpha \in \mathrm{End}(E^2)$ from Lemma 2.5.5 with $N := u(2^f-u)$. Then, the following lemma ensures that α can be decomposed as $\alpha = \Psi_u \circ \Phi_u = \Phi'_u \circ \Psi'_u$, where Φ_u, Φ'_u are u-isogenies and Ψ_u, Ψ'_u are (2^f-u)-isogenies, so that we have a $(u, 2^f-u)$-isogeny diamond

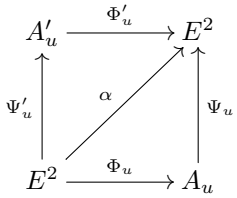

Lemma 2.5.6. *Let k be an algebraically closed field $d_1, d_2 \in \mathbb{N}^*$ coprime and not divisible by $\mathrm{char}(k)$ and $d := d_1d_2$, (A,λ_A) and (B,λ_B) be principally polarised abelian varieties and $f : A \longrightarrow B$ be a d-isogeny. Then:*

(i) There exists a principally polarised abelian variety (C,λ_C), a d_1-isogeny $f_1 : A \longrightarrow C$ and a d_2-isogeny $f_2 : C \longrightarrow B$ such that $f = f_2 \circ f_1$.

(ii) f_1 and f_2 are unique up to post or pre composition by an isomorphism and we have $\ker(f_1) = \ker(f)[d_1] = [d_2]\ker(f)$ and $\ker(f_2) = f_1(\ker(f))$.

Proof. This will be proved later. We refer to Lemma 6.3.1 and Remark 6.3.2. □

Applying Kani's lemma (Lemma 2.2.6) to the isogeny diamond above yields a 4-dimensional 2^f-isogeny:

$$F_u = \begin{pmatrix} \Phi_u & \widetilde{\Psi}_u \\ -\Psi'_u & \widetilde{\Phi'}_u \end{pmatrix} : E^4 \longrightarrow A_u \times A'_u,$$

with kernel:

$$\ker(F_u) = \{([u]P, [u]Q, \alpha(P,Q)) \mid P,Q \in E[2^f]\}.$$

From Lemma 6.4.3, we obtain 2^{f+2}-torsion points $T_1, \cdots, T_4$ forming an isotropic subgroup of $E^4[2^{f+2}]$ and such that $\ker(F_u) = \langle [4]T_1, \cdots, [4]T_4\rangle$. Using these points and the algorithmic approach from Sections 6.3 and 6.6, we can compute F_u as a chain of 2-isogenies in level 2 theta coordinates. With the techniques from Section 6.4.1, we can also extract A_u and A'_u from the codomain of F_u and express images $F_u(x,y) = (\Phi_u(x) + \widetilde{\Psi}_u(y), -\Psi'_u(x) + \widetilde{\Phi'}_u(y))$ in the product $A_u \times A'_u$. In particular, as desired, we can evaluate the u-isogeny $\Phi_u : E^2 \longrightarrow A_u$ since $F_u(x,0) = (\Phi_u(x), -\Psi'_u(x))$, so F_u is an efficient representation of Φ_u.

Since $f+2 \leq e-1$, note that most of the arithmetic operations performed to compute F_u take place over $\mathbb{F}_p$ instead of $\mathbb{F}_{p^2}$. Indeed, we have the following lemma which ensures that we can find a basis (P,Q) of $E[2^{f+2}]$ with $\mathbb{F}_p$-rational Montgomery $(X:Z)$-coordinates.

Lemma 2.5.7. *Let $E/\mathbb{F}_p$ be a supersingular Montgomery curve (primitively) $\mathfrak{O}$-oriented (with $\mathfrak{O} = \mathbb{Z}[(1+\sqrt{-p})/2]$). Then there exists a basis (P, Q) of $E[2^{e-1}]$ such that $\pi_p(P) = P$ and $\pi_p(Q) = -Q$, where π_p is the p-th Frobenius endomorphism. In particular, the Montgomery $(X : Z)$-coordinates of P and Q are $\mathbb{F}_p$-rational.*

Proof. Since E is supersingular and defined over $\mathbb{F}_p$, then $\#E(\mathbb{F}_p) = p + 1 = c2^e$. It follows that $\#E(\mathbb{F}_p)[2^e] = 2^e$. Since torsion subgroups of E have rank at most 2, we have $E(\mathbb{F}_p) \simeq (\mathbb{Z}/2^a\mathbb{Z}) \times (\mathbb{Z}/2^b\mathbb{Z})$ with $a \leq b$ and $a + b = e$. Then, $E[2^a] \subseteq E(\mathbb{F}_p) = \ker(\pi_p - 1)$ so $\pi_p - 1$ factors through $[2^a]$. The elliptic curve E being primitively $\mathfrak{O}$-oriented and $\mathfrak{O}$ being generated by $(\sqrt{-p}-1)/2$ which corresponds to $(\pi_p - 1)/2$, we must have $a = 1$ and $b = e - 1$, so that

$$E(\mathbb{F}_p)[2^e] = E(\mathbb{F}_p)[2^{e-1}] \simeq (\mathbb{Z}/2\mathbb{Z}) \times (\mathbb{Z}/2^{e-1}\mathbb{Z}). \tag{2.21}$$

This ensures the existence of $P \in E(\mathbb{F}_p)$ *i.e.* such that $\pi_p(P) = P$ of order 2^{e-1}. Let $A \in \mathbb{F}_p$ be the Montgomery coefficient of E, so that E is given by $y^2 = x^3 + Ax^2 + x$ and E^t be the quadratic twist over $\mathbb{F}_p$ given by $-y^2 = x^3 + Ax^2 + x$ isomorphic to E via $\tau : (x, y) \in E^t \longmapsto (x, \zeta_4 y) \in E$, with $\zeta_4^2 = -1$. Applying Eq. (2.21) to E^t, we obtain the existence of $P' \in E^t(\mathbb{F}_p)$ of order 2^{e-1}. Then, $Q := \tau(P') \in E(\mathbb{F}_{p^2}) \setminus E(\mathbb{F}_p)$ has order 2^{e-1}. Besides, since $p \equiv 3 \mod 4$ and $P' \in E^t(\mathbb{F}_p)$, we have:

$$\pi_p(Q) = \pi_p \circ \tau(P') = (x(P')^p, \zeta_4^p y(P')^p) = (x(P'), -\zeta_4 y(P')) = -\tau(P') = -Q.$$

To conclude, it suffices to prove that P and Q are linearly independent, *i.e.* that $[2^{e-2}]P \neq [2^{e-2}]Q$. But we have $\pi_p([2^{e-3}]P) = [2^{e-3}]P$ so $[2^{e-3}]P \in E(\mathbb{F}_p)$ and $\pi_p([2^{e-3}]Q) = -[2^{e-3}]Q$ so $[2^{e-3}]Q \in E(\mathbb{F}_{p^2}) \setminus E(\mathbb{F}_p)$. By Eq. (2.21), all 4-torsion points lying above $[2^{e-2}]P$ are defined over $\mathbb{F}_p$, so $[2^{e-3}]Q$ does not lie above $[2^{e-2}]P$ and we have $[2^{e-2}]P \neq [2^{e-2}]Q$. This completes the proof. □

We refer to [DEF+25, Appendix D.1] for an efficient algorithm to sample such a basis of $E[2^{e-1}]$ formed of eigenvectors for π_p.

Since conversion formulas from Montgomery $(X : Z)$-coordinates to level 2 theta coordinates on Montgomery elliptic curves are $\mathbb{F}_p$-rational (see Proposition 5.3.47 and Lemma 6.5.7), if we use a basis (P, Q) of $E[2^{f+2}]$ with $\mathbb{F}_p$-rational Montgomery $(X : Z)$-coordinates to express $T_1, \cdots, T_4$ we expect the product theta coordinates of these points to be $\mathbb{F}_p$-rational. Then, the computation of the first 2-isogenies of the chain F_u may involve operations over $\mathbb{F}_{p^2}$ (to apply changes of theta coordinates with Theorem 6.2.10 in particular) but the following 2-isogenies are computed over $\mathbb{F}_p$. Though, this convenient result is experimental and has not been proved over the course of this PhD. Further study is needed to obtain a formal proof.

2.5.3 Step 4: computing the 4-dimensional isogeny

Consider the following (uN_1, vN_2)-isogeny diamond

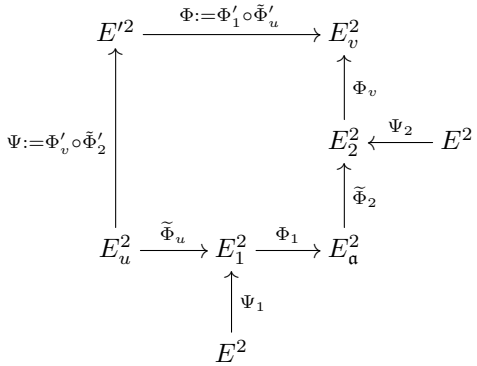

(2.22)

In the diagram above, we have:

- $\Psi_1 = \mathrm{Diag}(\varphi_{\mathfrak{b}_e}, \varphi_{\mathfrak{b}_e})$, where $\varphi_{\mathfrak{b}_e}$ is the isogeny corresponding to the action of $\mathfrak{b}_e$; analogously, $\Psi_2 = \mathrm{Diag}(\varphi_{\mathfrak{c}_e}, \varphi_{\mathfrak{c}_e})$;

- $\Phi_1 = \mathrm{Diag}(\varphi_{\mathfrak{b}_k}, \varphi_{\mathfrak{b}_k})$ and $\Phi_2 = \mathrm{Diag}(\varphi_{\mathfrak{c}_k}, \varphi_{\mathfrak{c}_k})$;
- Φ_u and Φ_v are isogenies of polarised degree u and v respectively; as explained in Section 2.5.2, we can write $\Phi_u = M_u \circ \mathrm{Diag}(\varphi_u, \varphi_u)$, with $\deg(\varphi_u) = g_u$, $\Phi_v = M_v \circ \mathrm{Diag}(\varphi_v, \varphi_v)$ with $\deg(\varphi_v) = g_v$, M_u and M_v being given by Eq. (2.20);

Furthermore, the isogenies Φ and Ψ are the isogenies of polarised degree uN_1 and vN_2 completing the square, whose existence is guaranteed by Lemma 2.5.6. In this setting we can also describe them explicitly, as the following lemma shows.

Lemma 2.5.8. *Φ is of the form $\Phi'_1 \circ \widetilde{\Phi}'_u$, where:*

- *$\Phi'_1 := \mathrm{Diag}(\varphi'_{\mathfrak{b}_k}, \varphi'_{\mathfrak{b}_k}) : {E'_1}^2 \longrightarrow E_v^2$ with $E'_1 := [\overline{\mathfrak{b}}_k] \cdot E_v$ and $\varphi'_{\mathfrak{b}_k}$, the isogeny associated to the action of $\mathfrak{b}_k$ on E'_1;*
- *$\Phi'_u := M_u \,\mathrm{Diag}(\varphi'_u, \varphi'_u) : {E'_1}^2 \longrightarrow {E'}^2$, M_u being given by Eq. (2.20) and φ'_u being the isogeny of degree g_u given by the action of the same ideal $\mathfrak{g}_u$ as φ_u (a product of prime ideals lying above primes in $\mathfrak{B}$).*

Ψ is of the form $\Phi'_v \circ \widetilde{\Phi}'_2$, where:

- *$\Phi'_2 := \mathrm{Diag}(\varphi'_{\mathfrak{c}_k}, \varphi'_{\mathfrak{c}_k}) : {E'_2}^2 \longrightarrow E_u^2$ with $E'_2 := [\overline{\mathfrak{c}}_k] \cdot E_u$ and $\varphi'_{\mathfrak{c}_k}$, the isogeny associated to the action of $\mathfrak{c}_k$ on E'_2;*
- *$\Phi'_v := M_v \,\mathrm{Diag}(\varphi'_v, \varphi'_v) : {E'_2}^2 \longrightarrow {E'}^2$, M_v being given by Eq. (2.20) and φ'_v being the isogeny of degree g_v given by the action of the same ideal $\mathfrak{g}_v$ as φ_v (a product of prime ideals lying above primes in $\mathfrak{B}$).*

In particular, the common codomain of $\widetilde{\Phi}$ and Ψ is a product of elliptic curves ${E'}^2$.

Proof. We have to verify that $\Phi'_1 \circ \tilde{\Phi}'_u$ and $\Phi'_v \circ \tilde{\Phi}'_2$ defined above are $N_1 u$ and $N_2 v$-isogenies respectively making the diagram commute, i.e. such that:

$$\Phi'_1 \circ \tilde{\Phi}'_u \circ \Phi'_v \circ \tilde{\Phi}'_2 = \Phi_v \circ \tilde{\Phi}_2 \circ \Phi_1 \circ \tilde{\Phi}_u. \tag{2.23}$$

First, we verify that the composition on the left makes sense i.e. that Φ'_u and Φ'_v have the same codomain. By definition, the codomain of Φ'_u is

$$E' = [\mathfrak{g}_u]E'_1 = [\mathfrak{g}_u \overline{\mathfrak{b}}_k]E_v = [\mathfrak{g}_u \overline{\mathfrak{b}}_k \mathfrak{g}_v]E_2 = [\mathfrak{g}_u \mathfrak{g}_v \overline{\mathfrak{b}}_k \mathfrak{c}_e]E$$

and the codomain of Φ'_v is

$$E'' = [\mathfrak{g}_v]E'_2 = [\mathfrak{g}_v \overline{\mathfrak{c}}_k]E_u = [\mathfrak{g}_v \overline{\mathfrak{c}}_k \mathfrak{g}_u]E_1 = [\mathfrak{g}_u \mathfrak{g}_v \overline{\mathfrak{c}}_k \mathfrak{b}_e]E.$$

But $\mathfrak{b} = \mathfrak{b}_e \mathfrak{b}_k$ and $\mathfrak{c} = \mathfrak{c}_e \mathfrak{c}_k$ are both equivalent to $\mathfrak{a}$. It follows that $\overline{\mathfrak{b}}_k \mathfrak{c}_e$ and $\overline{\mathfrak{c}}_k \mathfrak{b}_e$ are equivalent, so that $E' \simeq E''$.

By construction, $\Phi'_1 \circ \tilde{\Phi}'_u$ and $\Phi'_v \circ \tilde{\Phi}'_2$ are $N_1 u$ and $N_2 v$-isogenies respectively, so we only have to prove Eq. (2.23). On the one hand, we have:

$$\Phi'_1 \circ \widetilde{\Phi}'_u \circ \Phi'_v \circ \widetilde{\Phi}'_2 = \widetilde{M_u} \circ M_v \circ \mathrm{Diag}(\varphi'_{\mathfrak{b}_k} \circ \widehat{\varphi}'_u \circ \varphi'_v \circ \widehat{\varphi}'_{\mathfrak{c}_k}, \varphi'_{\mathfrak{b}_k} \circ \widehat{\varphi}'_u \circ \varphi'_v \circ \widehat{\varphi}'_{\mathfrak{c}_k}),$$

since $\widetilde{M_u}$ and M_v commute with diagonal isogenies. On the other hand:

$$\Phi_v \circ \widetilde{\Phi}_2 \circ \Phi_1 \circ \widetilde{\Phi}_u = M_v \circ \widetilde{M_u} \circ \mathrm{Diag}(\varphi_v \circ \widehat{\varphi}_{\mathfrak{c}_k} \circ \varphi_{\mathfrak{b}_k} \circ \widehat{\varphi}_u, \varphi_v \circ \widehat{\varphi}_{\mathfrak{c}_k} \circ \varphi_{\mathfrak{b}_k} \circ \widehat{\varphi}_u).$$

Both $\varphi'_{\mathfrak{b}_k} \circ \widehat{\varphi}'_u \circ \varphi'_v \circ \widehat{\varphi}'_{\mathfrak{c}_k}$ and $\varphi_v \circ \widehat{\varphi}_{\mathfrak{c}_k} \circ \varphi_{\mathfrak{b}_k} \circ \widehat{\varphi}_u$ correspond to the action of the ideal $\overline{\mathfrak{g}}_u \mathfrak{g}_v \overline{\mathfrak{c}}_k \mathfrak{b}_k$ on E_u so these isogenies must be equal. Besides, a simple matrix computation ensures that $\widetilde{M_u}$ and M_v commute. This proves Eq. (2.23) and the lemma. □

Eq. (2.22) and Kani's lemma (Lemma 2.2.6) yield a 2^f-isogeny:

$$F = \begin{pmatrix} \Phi_1 \circ \widetilde{\Phi}_u & \Phi_2 \circ \widetilde{\Phi}_v \\ -\Phi'_v \circ \tilde{\Phi}'_2 & \Phi'_u \circ \tilde{\Phi}'_1 \end{pmatrix} : E_u^2 \times E_v^2 \longrightarrow E_{\mathfrak{a}}^2 \times E'^2.$$

with kernel:

$$\begin{aligned} \ker(F) &= \{([uN_1]P, [uN_1]Q, \Phi_v \circ \widetilde{\Phi}_2 \circ \Phi_1 \circ \widetilde{\Phi}_u(P,Q)) \mid P, Q \in E_u[2^f]\} \\ &= \{([N_1]\Phi_u(P,Q), \Phi_v \circ \widetilde{\Phi}_2 \circ \Phi_1(\varphi_u(P), \varphi_u(Q))) \mid P, Q \in E_1[2^f]\} \\ &= \{([N_1]M_u(\varphi_u(P), \varphi_u(Q)), M_v(\varphi_v \circ \widehat{\varphi}_{\mathfrak{c}_k} \circ \varphi_{\mathfrak{b}_k}(P), \varphi_v \circ \widehat{\varphi}_{\mathfrak{c}_k} \circ \varphi_{\mathfrak{b}_k}(Q))) \mid P, Q \in E_1[2^f]\}. \end{aligned}$$

Let (P, Q) be a basis of $E_1[2^{f+2}]$, with $\mathbb{F}_p$-rational $(X : Z)$-coordinates (which does exist by Lemma 2.5.7 since $f \leq e - 3$), η_1, α be inverses of N_1, uN_1 modulo 2^{f+2} respectively, $(P_u, Q_u) := \varphi_u(P, Q)$ and $(P_v, Q_v) := \varphi_v \circ \widehat{\varphi}_{\mathfrak{c}_k} \circ \varphi_{\mathfrak{b}_k}(P, Q)$. Then, applying Lemma 6.4.1.(i), we obtain 2^{f+2}-torsion points:

$$\begin{aligned} T_1 &:= ([N_1 x_u]P_u, [N_1 y_u]P_u, [x_v]P_v, [y_v]P_v) \\ T_2 &:= ([-N_1 y_u]P_u, [N_1 x_u]P_u, [-y_v]P_v, [x_v]P_v) \\ T_3 &:= ([(1 - 2^e\alpha)x_u]Q_u, [(1 - 2^e\alpha)y_u]Q_u, [\eta_1 x_v]Q_v, [\eta_1 y_v]Q_v) \\ T_4 &:= ([-(1 - 2^e\alpha)y_u]Q_u, [(1 - 2^e\alpha)x_u]Q_u, [-\eta_1 y_v]Q_v, [\eta_1 x_v]Q_v) \end{aligned}$$

forming an isotropic subgroup $K' \subset (E_u^2 \times E_v^2)[2^{f+2}]$ such that $[4]K' = \ker(F)$ and we can compute F from these points by Theorem 2.2.12. More precisely, to compute F with the techniques from Sections 6.3 and 6.6, we need:

- The integers $N_1, N_2, g_u, x_u, y_u, g_v, x_v, y_v, e$ obtained from a solution to Problem 2.5.2 obtained in Step 1.
- The points P_u, Q_u, P_v, Q_v that needs to be computed in Step 2 (see Section 2.5.4).

Even though the approach to compute F is similar to Section 6.6, we refer to [DEF+25, Appendix B] for relevant specific details.

Remark 2.5.9. Since the $(X : Z)$-coordinates of P and Q are $\mathbb{F}_p$-rational and the isogenies $\varphi_u, \varphi_v, \varphi_{\mathfrak{b}_k}, \varphi_{\mathfrak{c}_k}$ associated to ideals of $\mathfrak{O}$ are also $\mathbb{F}_p$-rational (see Section 2.5.4), the points P_u, Q_u, P_v, Q_v have $\mathbb{F}_p$-rational $(X : Z)$-coordinates. As a consequence, we expect that most arithmetic operations involved in the computation of F take place over $\mathbb{F}_p$, for reasons explained in the end in Section 2.5.2.

2.5.4 Step 2: evaluating Elkies' isogenies

The goal of this section is to explain how to compute the isogenies associated to the ideals $\mathfrak{b}_e, \mathfrak{c}_e, \mathfrak{g}_u, \mathfrak{g}_v$ introduced previously and how to compute the points $(P_u, Q_u) := \varphi_u(P, Q)$ and $(P_v, Q_v) := \varphi_v \circ \widehat{\varphi}_{\mathfrak{c}_k} \circ \varphi_{\mathfrak{b}_k}(P, Q)$ introduced in the end of Section 2.5.3.

Computing $\mathbb{F}_p$-rational isogenies with Elkies' algorithm

Let $\mathfrak{l} \subseteq \mathfrak{O}$ be a prime ideal and $\varphi_{\mathfrak{l}} : (E, \iota) \longrightarrow (E_{\mathfrak{l}}, \iota_{\mathfrak{l}})$ be an associated $\mathfrak{O}$-oriented isogeny. By the definition of an $\mathfrak{O}$-oriented isogeny, we have $\varphi_{\mathfrak{l}} \circ \iota(\alpha) = \iota_{\mathfrak{l}}(\alpha) \circ \varphi_{\mathfrak{l}}$ for all $\alpha \in \mathfrak{O}$. In particular (with $\alpha = \sqrt{-p}$), we obtain that $\varphi_{\mathfrak{l}} \circ \pi_p^E = \varphi_{\mathfrak{l}} \circ \pi_p^{E_{\mathfrak{l}}}$, where π_p^E and $\pi_p^{E_{\mathfrak{l}}}$ are the p-th Frobenius endomorphisms of E and $E_{\mathfrak{l}}$ respectively. It follows that $\varphi_{\mathfrak{l}}$ is $\mathbb{F}_p$-rational.

However, its kernel $E[\mathfrak{l}]$ may not be defined over $\mathbb{F}_p$ or even $\mathbb{F}_{p^2}$, except if $\ell := N(\mathfrak{l})$ divides $p + 1$, which is not the case for all primes in $\mathfrak{B}$ in general. For that reason, we cannot compute $\varphi_{\mathfrak{l}}$ with Vélu's formulas [Vé71] without using field extensions of $\mathbb{F}_p$. To circumvent this difficulty, we use Elkie's algorithm introduced in [Elk98] instead to stay over $\mathbb{F}_p$.

Assume that $\ell = N(\mathfrak{l})$ is odd. Then, there are 2 horizontal ℓ-isogenies starting from E associated to $\mathfrak{l}$ and $\bar{\mathfrak{l}}$. The j-invariants $j(E_{\mathfrak{l}})$ and $j(E_{\bar{\mathfrak{l}}})$ defined over $\mathbb{F}_p$ may be found with the ℓ-th modular

polynomial. We compute $\Phi_\ell(x, j(E)) \in \mathbb{F}_p[x]$ and find roots of this polynomial defined over $\mathbb{F}_p$. These roots are $j(E_\mathfrak{l})$ and $j(E_{\bar{\mathfrak{l}}})$. We have no way of distinguishing between the two in general.

If $\varphi : E \longrightarrow E'$ is an $\mathbb{F}_p$-rational ℓ-isogeny, then its kernel polynomial $h(x)$ is a factor of degree $(\ell-1)/2$ of the division polynomial of E which is defined over $\mathbb{F}_p$. To find it knowing $j(E)$ and $j(E')$, we may use [BSS99, Algorithm VII.3]. Since we cannot distinguish between $j(E_\mathfrak{l})$ and $j(E_{\bar{\mathfrak{l}}})$, to check if the resulting kernel polynomial $h(x)$ corresponds to $\varphi_\mathfrak{l}$, we proceed as follows. The ideal $\mathfrak{l}$ is of the form $\mathfrak{l} = \ell\mathbb{Z} + (\sqrt{-p} - \lambda)/2\mathbb{Z}$, where $\lambda \in \mathbb{Z}$ is a square root of $-p \mod \ell$. It follows that the points of $E[\mathfrak{l}] = \ker(\varphi_\mathfrak{l})$ are eigenvectors of π_p for the eigenvalue λ. Since we do not work with ℓ-torsion points, we test the following equality symbolically

$$(x^p, y^p) \equiv [\lambda](x, y) \mod (h(x), y^2 - f(x)),$$

where $y^2 - f(x) = 0$ is the equation of E. If the above equality is not satisfied, it means that we picked the wrong j-invariant $j(E') = j(E_{\bar{\mathfrak{l}}})$ and we may apply [BSS99, Algorithm VII.3] again to the other j-invariant $j(E_\mathfrak{l})$ obtained as a root of $\Phi(X, j(E))$ to compute the kernel polynomial of $\varphi_\mathfrak{l}$. Once, we have found this kernel polynomial, we can express $\varphi_\mathfrak{l}$ with the formulas from [BMSS08, Proposition 4.1].

If $\ell = 2$, the computation of $\varphi_\mathfrak{l}$ is way simpler. Indeed, $\mathfrak{l}$ is of the form $2\mathbb{Z} + (\sqrt{-p} \pm 1)/2\mathbb{Z}$. Assume we know a basis (P, Q) of $E[4]$ such that $\pi_p(P) = P$ and $\pi_p(Q) = -Q$, so that the $(X : Z)$-coordinates of P and Q are defined over $\mathbb{F}_p$ (its existence is ensured by Lemma 2.5.7). Then, $E[\mathfrak{l}] = \langle [2]P \rangle$ if $\mathfrak{l} = 2\mathbb{Z} + (\sqrt{-p} - 1)/2\mathbb{Z}$ and $E[\mathfrak{l}] = \langle [2]Q \rangle$ if $\mathfrak{l} = 2\mathbb{Z} + (\sqrt{-p} + 1)/2\mathbb{Z}$. So we can easily compute $\varphi_\mathfrak{l}$ with Vélu's formulas.

Note that the isogeny formulas we use are adapted to Weierstrass models while our $\mathfrak{O}$-oriented supersingular elliptic curves are more naturally described by Montgomery models, so we apply conversion formulas between these two models. With the techniques presented above, we can rerandomise the starting curve E by computing an action $[\bar{\mathfrak{l}}^g] \cdot E$ as explained in the end of Section 2.5.1; and compute the isogenies $\varphi_{\mathfrak{b}_e} : E \longrightarrow E_1$, $\varphi_{\mathfrak{c}_e} : E \longrightarrow E_2$, $\varphi_u : E_1 \longrightarrow E_u$ and $\varphi_v : E_2 \longrightarrow E_v$ respectively associated to $\mathfrak{b}_e, \mathfrak{c}_e, \mathfrak{g}_u$ and $\mathfrak{g}_v$ which are products of prime ideals lying above primes of $\mathfrak{B}$.

Computing P_u, Q_u, P_v and Q_v

Now, let (P, Q) be a basis of $E_1[2^{f+2}]$ such that $\pi_p(P) = P$ and $\pi_p(Q) = -Q$ (Lemma 2.5.7). The computation of $(P_u, Q_u) := \varphi_u(P, Q)$ is straight forward once we have computed φ_u with the techniques presented above. We can even work with x-only or $(X : Z)$-only arithmetic to obtain $(\pm\varphi_u(P), \pm\varphi_u(Q))$ on the Kummer line $E_u/\pm$ and lift the result to $(\varphi_u(P), \varphi_u(Q))$ with a Weil pairing computation or by computing the additional Kummer line point $\pm\varphi_u(P - Q)$. We now explain how to compute $(P_v, Q_v) := \varphi_v \circ \widehat{\varphi}_{\mathfrak{c}_k} \circ \varphi_{\mathfrak{b}_k}(P, Q)$.

Let $\iota : \mathfrak{O} \xrightarrow{\sim} \mathrm{End}_{\mathbb{F}_p}(E)$ be the the $\mathfrak{O}$-orientation of E (mapping $\sqrt{-p}$ to π_p). Then, by construction, the ideals $\mathfrak{a}$ and $\mathfrak{b}$ obtained in Step 1 are given by $\mathfrak{b} = \mathfrak{a}\overline{\beta}_1/N(\mathfrak{a})$ and $\mathfrak{c} = \mathfrak{a}\overline{\beta}_2/N(\mathfrak{a})$, with $\beta_1, \beta_2 \in \mathfrak{a}$, so that $\bar{\mathfrak{c}}\mathfrak{b}$ is principal: $\bar{\mathfrak{c}}\mathfrak{b} = \mathfrak{O}\theta$ with $\theta := \overline{\beta}_1\beta_2/N(\mathfrak{a}) \in \mathfrak{O}$. It follows that $\widehat{\varphi}_\mathfrak{c} \circ \varphi_\mathfrak{b} = \iota(\theta)$, so that:

$$[N(\mathfrak{b}_e)N(\mathfrak{c}_e)]\widehat{\varphi}_{\mathfrak{c}_k} \circ \varphi_{\mathfrak{b}_k} = \varphi_{\mathfrak{c}_e} \circ \iota(\theta) \circ \widehat{\varphi}_{\mathfrak{b}_e}$$

Hence, when $N(\mathfrak{b}_e)$ and $N(\mathfrak{c}_e)$ are odd, we only have to evaluate $\varphi_{\mathfrak{c}_e} \circ \iota(\theta) \circ \widehat{\varphi}_{\mathfrak{b}_e}$ on (P, Q) and multiply the result by an inverse of $N(\mathfrak{b}_e)N(\mathfrak{c}_e) \mod 2^{f+2}$. The evaluation of $\widehat{\varphi}_{\mathfrak{b}_e}$ may be done by computing the action of $\bar{\mathfrak{b}}_e$ on E_1 or by the classical discrete logarithm computation presented in Algorithm 2.9[4]. To evaluate $\iota(\theta)$, we decompose $\theta = (a + b\sqrt{-p})/2$ with $a, b \in \mathbb{Z}$, so that $\iota(\theta) = ([a] + [b]\pi_p)/2$. A straightforward way to evaluate $\iota(\theta)$ on 2^{f+2}-torsion points is to use 2^{f+3}-torsion points that may not be defined over $\mathbb{F}_p$. An alternative involving Weil and Tate pairing computations has been introduced in [DEF+25, Algorithm 1].

Now, we assume that $N(\mathfrak{b}_e)$ or $N(\mathfrak{c}_e)$ is even. We write $\mathfrak{b}_e = \mathfrak{b}'_e \cdot \mathfrak{b}_2$ and $\mathfrak{b}_e = \mathfrak{c}'_e \cdot \mathfrak{c}_2$, where $\mathfrak{b}_2$ and $\mathfrak{c}_2$ are ideals of norms a power of 2. Since the action by common factors of $\mathfrak{b}_e, \mathfrak{c}_e$ can be computed at the beginning (as for the rerandomisation of $\mathfrak{a}$ in Step 1), we can assume that $\mathfrak{b}_2$ and $\mathfrak{c}_2$ are coprime *i.e.* that $\mathfrak{b}_2 = \mathfrak{l}^{f_1}$ and $\mathfrak{b}_2 = \bar{\mathfrak{l}}^{f_2}$, where $\mathfrak{l}$ is a prime ideal lying above 2. Writing $f_{12} = f_1 + f_2$, we get that

$$[N(\mathfrak{b}_{e'})N(\mathfrak{c}_{e'})2^{f_{12}}]\widehat{\varphi}_{\mathfrak{c}_k} \circ \varphi_{\mathfrak{b}_k} = \varphi_{\mathfrak{c}_e} \circ \iota(\sigma) \circ \widehat{\varphi}_{\mathfrak{b}_e}.$$

[4]That will be reused later.

The only added difficulty that needs to be accounted for is that we cannot find an inverse to $2^{f_{12}}$ modulo 2^{e+2}. We fix this by taking a basis (P, Q) of a larger torsion group $E[2^{f+2+f_{12}}]$ instead of $E[2^{f+2}]$ in the beginning. We remark that by construction, we have $f + 2 + f_{12} \leq e - 1$ since we have solved Problem 2.5.2 with $f_{max} = e - 3$ in Step 1.

Algorithm 2.9: Classical method to evaluate the dual of an isogeny.

Data: Coprime integers $N, d \in \mathbb{N}^*$, a basis (P_1, Q_1) of $E_1[N]$, its image $(\varphi(P_1), \varphi(Q_1))$ by a d-isogeny $\varphi : E_1 \longrightarrow E_2$ and a basis (P_2, Q_2) of $E_2[N]$.
Result: $(\widehat{\varphi}(P_2), \widehat{\varphi}(Q_2))$.
1 Find $\alpha, \beta, \gamma, \delta \in \mathbb{Z}/N\mathbb{Z}$ such that $\varphi_1(P_1) = [\alpha]P_2 + [\beta]Q_2$ and $\varphi_1(Q_1) = [\gamma]P_2 + [\delta]Q_2$;
2 $\Delta^{-1} \leftarrow (\alpha\delta - \beta\gamma)^{-1} \mod N$;
3 $\widehat{\varphi}_1(P_2) \leftarrow [d\Delta^{-1}\delta]P_1 - [d\Delta^{-1}\beta]Q_1$;
4 $\widehat{\varphi}_1(Q_2) \leftarrow -[d\Delta^{-1}\gamma]P_1 + [d\Delta^{-1}\alpha]Q_1$;
5 **return** $(\widehat{\varphi}(P_2), \widehat{\varphi}(Q_2))$;

2.5.5 Performance

The PEGASIS algorithm that we described for CSIDH/CSURF was implemented in SageMath[5] for primes of the form $p = c2^e - 1$ of different sizes. Our implementation covers primes from 500 bits to 4000 bits, the latter ensuring at least NIST-I level of security against the Kuperberg subexponential attack [Kup05; BS20; Pei20; CSCDJRH21]. For bigger primes p, we increased the size of the allowed Elkies' primes set $\mathfrak{B}$ to make Step 1 faster. The parameters may be found in Table 2.2.

Parameter set	e	c	$\mathfrak{B}$
500	503	33	3, 7, 11, 13
1000	1004	15	3, 5, 7, 11
1500	1551	9	3, 5, 11
2000	2026	51	3, 7, 11, 17
4000	4084	63	3, 7, 11, 17, 19

Table 2.2: Parameter sets used in our implementation. The prime p is of the form $p = c2^e - 1$ and $\mathfrak{B}$ is the set of small split primes used.

The results of our SageMath 10.5 implementation can be found in Table 2.3; timings for each steps are in seconds, and are obtained by averaging 100 runs on an Intel Core i5-1235U.

Parameter set	Step 1 (s)	Step 2 (s)	Step 3 (s)	Tot. Time (s)
500	0.097	0.477	0.960	1.534
1000	0.212	1.159	2.838	4.210
1500	1.186	2.853	6.491	10.530
2000	1.675	8.337	11.327	21.339
4000	15.606	52.808	53.463	121.876

Table 2.3: SageMath 10.5 timings in sec on Intel Core i5-1235U. Step 1 is the time to solve the norm equation, Step 2 the time to do the Elkies steps, and Step 3 the time to compute the dimension 4 isogeny

[5]The public repository can be found at https://github.com/pegasis4d/pegasis.

We conclude with a more detailed comparison with the other available isogeny based effective group actions (EGA) in the literature. The comparison is summarised in Table 2.4. The timings for SCALLOP [FFK+23] were reportedly measured on an Intel i5-6440HQ processor running at 3.5 GHz, while the timings for SCALLOP-HD [CLP24] were reportedly measured on an Intel Alder Lake CPU core clocked at 2.1 GHz. The timings for PEARL-SCALLOP [ABE+24] and the two versions of KLaPoTi [PPS24] were all re-measured on the same hardware setup as the timings presented in Table 2.3.

Paper	Impl.	500	1000	1500	2000	4000
SCALLOP [FFK+23]*	C++	35s	12m30s	–	–	–
SCALLOP-HD [CLP24]*	Sage	88s	19m	–	–	–
PEARL-SCALLOP [ABE+24]	C++	30s	58s	12m	–	–
KLaPoTi [PPS24]	Sage	200s	–	–	–	–
	Rust	1.95s	–	–	–	–
PEGASIS (This work)	Sage	1.53s	4.21s	10.5s	21.3s	2m2s

Table 2.4: Comparison between PEGASIS and other effective group actions in the literature. The last 5 columns gives the timings corresponding to the different security levels, where s/m gives the number of seconds/minutes in wall-clock time. SCALLOP and SCALLOP-HD are starred because they were measured on a different hardware setup.

PEGASIS is the first instantiation of an EGA at the 2000 bits and 4000 bits security level. In fact, it is even the first time that the full CSIDH group action can be computed at the 1000 bits level. However, as Table 2.4 shows, PEGASIS also significantly outperforms all earlier works at their security levels. The closest comparison comes with the Rust implementation of KLaPoTi. However, the fact that KLaPoTi was able to achieve a speedup by two orders of magnitude simply by switching from a high-level SageMath implementation to a low-level Rust implementation is also very promising for PEGASIS. Under the assumption that a comparable speedup would be possible for PEGASIS, we see that PEGASIS gives a highly-practical EGA, even at the highest security levels.

The main reason of this efficiency is that unlike all the other EGA instantiations, we are able to use the orientation given by Frobenius. This has three key benefits:

- Evaluating the orientation is very efficient.
- We do not need to represent, nor push forward the orientation.
- We can work over $\mathbb{F}_p$.

The price to pay is the need to go up to dimension 4 to compute the action, but as Table 2.3 shows that it is still reasonably efficient. We stress that our implementation is only a proof of concept, to explore whether level 4000 was feasible in practice. In particular, we are missing many standard implementation tricks; as a simple example, we work with affine coordinates instead of projective coordinates. Despite this, our timings are very encouraging, and we hope a low level optimised implementation could even be made comparable in efficiency to a state-of-the-art implementation of a restricted effective group action (REGA), like [CHMR25].

Chapter 3

SQIsignHD: faster and safer signatures with higher dimensional isogenies

With the algorithmic tools introduced in Chapter 2 to translate quaternion ideals into supersingular isogenies, we are ready to introduce the digital signature scheme SQIsignHD based on SQIsign, which was the first isogeny based scheme to use them. The presentation that follows is mainly based on an original contribution from this PhD [DLRW24] and on improvements from the SQIsign2D-West paper, also part of this PhD [BDF+25]. These improvements simplify the description and accelerate the fast version of SQIsignHD.

Recall that the bottleneck of SQIsign was the ideal-to-isogeny translation in the signing phase that relied on KLPT based techniques presented in Section 2.1. In SQisignHD, this costly ideal-to-isogeny translation is replaced by the evaluation of some points via an isogeny (which is fast in practice). To verify that these image points represent the expected isogeny, a higher dimensional isogeny computation (in dimension 4 in practice) takes place during the verification phase (using the techniques from Sections 2.2.4 and 2.3). This method enhances greatly the signing performance, allows for even more compact signatures than SQIsign and cleaner security proofs at the expense of the verification efficiency. SQIsign2D-West [BDF+25] that will be introduced in Chapter 4 certainly offers the best compromise on all fronts (even the verification), but SQIsignHD is still interesting as a trade-off between verification and signing time as it beats the record for fast isogeny based signatures, and even more with the most recent improvements.

Note that two variants of SQIsignHD will be introduced in the following: F-SQisignHD, with 4-dimensional verification, implemented and optimised for efficiency, and R-SQIsignHD, with 8-dimensional verification, then purely theoretical and optimised for clean security proofs. The description of the first variant is simple while the description of the latter is a bit more technical. We shall first give a high level description of SQIsignHD, present some algorithmic building blocks that were not already previously introduced (and that will also be useful in Chapter 4), then present each phase of both variants of SQIsignHD, before a presentation of their security analysis and a concluding section on the performance of F-SQIsignHD.

3.1 An overview of the SQIsign framework

3.1.1 An identification protocol

SQIsign is a digital signature scheme obtained via the Fiat-Shamir transform [FS87] (see Section 3.1.3) of an identification protocol where a party (called the *prover*) proves the knowledge of some secret information to another party (called the *verifier*) without revealing it. This protocol is built on the Deuring correspondence between quaternion ideals and isogenies. SQIsign and SQIsignHD mainly differ in the way of making the Deuring correspondence effective. While SQIsign only works with smooth degree isogenies between supersingular elliptic curves, SQIsignHD uses 4-dimensional isogenies

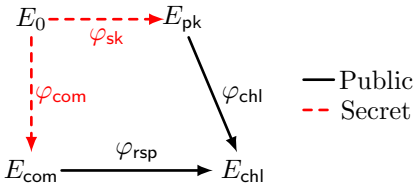

Figure 3.1: The SQIsign/SQIsignHD identification protocol. This diagram is non-commutative. Dashed red lines represent secrets.

in the verification process. In the following, we present the main building blocks of SQIsign (and SQIsignHD) identification protocol.

Public set-up. We choose a prime p and a supersingular elliptic curve $E_0/\mathbb{F}_{p^2}$ of known endomorphism ring $\mathcal{O}_0 \cong \mathrm{End}(E_0)$ such that E_0 has smooth torsion defined over a small extension of $\mathbb{F}_{p^2}$ (of degree 1 or 2). In practice, one may use the curve $E_0 : y^2 = x^3 + x$ (and $p \equiv 3 \mod 4$).

Key generation. The prover generates a random secret isogeny $\varphi_{\mathsf{sk}} : E_0 \longrightarrow E_{\mathsf{pk}}$ and publishes E_{pk} as its public key.

Commitment. The prover generates a random secret isogeny $\varphi_{\mathsf{com}} : E_0 \longrightarrow E_{\mathsf{com}}$ and sends E_{com} to the verifier as its commitment. For the identification protocol to be zero-knowledge (and the derived signature scheme to be secure), E_{com} has to be computationally indistinguishable from a uniformly random elliptic curve in the supersingular isogeny graph.

Challenge. The verifier generates and sends to the prover a random isogeny $\varphi_{\mathsf{chl}} : E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}$ of smooth degree sufficiently large for φ_{chl} to have high entropy. The challenge space should have size $\Omega(2^\lambda)$ to ensure λ bits of (soundness) security.

Response. The prover generates and transmits to the verifier an *efficient representation* (as defined in Definition 1.1.22) of an isogeny $\varphi_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$ which does not backtrack through φ_{chl} (i.e. $\widehat{\varphi}_{\mathsf{rsp}} \circ \varphi_{\mathsf{chl}}$ is cyclic).

Verification. The verifier checks that the response returned by the prover correctly represents an isogeny $\varphi_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$ and checks that this isogeny does not backtrack through φ_{chl}. The diagram in Fig. 3.1 illustrates the relationship between the various isogenies computed by the protocol.

To compute such an efficient representation of φ_{rsp} (that will be called φ_{rsp} by abuse of notations), the prover uses the Deuring correspondence. Returning $\varphi_{\mathsf{rsp}} = \varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}} \circ \widehat{\varphi}_{\mathsf{com}} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$ would make the scheme insecure. However, the prover can translate $\varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}} \circ \widehat{\varphi}_{\mathsf{com}}$ into an ideal I connecting the maximal quaternion orders isomorphic to $\mathrm{End}(E_{\mathsf{com}})$ and $\mathrm{End}(E_{\mathsf{chl}})$, find a random equivalent ideal $I_{\mathsf{rsp}} \sim I$ and translate I_{rsp} into φ_{rsp}.

3.1.2 From SQIsign to SQIsignHD

The ideal $I_{\mathsf{rsp}} \sim I$ is sampled to be relatively easy to translate into an isogeny and with a distribution that ensures one can simulate the response without secret knowledge (zero knowledge property). These two objectives are in tension and lead to a trade-off between efficiency and rigorous security proofs. As we have seen in Section 2.1, in SQIsign, $\mathrm{nrd}(I_{\mathsf{rsp}})$ had to be smooth to make the ideal-to-isogeny translation possible. The KLPT algorithm [KLPT14] was used to find I_{rsp}, resulting in big norms $\mathrm{nrd}(I_{\mathsf{rsp}}) \approx p^{15/4}$, slow ideal-to-isogeny translation and a very heuristic security proof.

In SQIsignHD [DLRW24], the smoothness condition on I_{rsp} is relaxed, allowing for smaller norms, a stronger security proof and a faster response at the expense of the verification time. To represent φ_{rsp}, we use the algorithmic approach from Section 2.3 directly inspired from SIDH attack techniques [CD23; MMPPW23; Rob23] presented in Section 2.2.4. Following the ideas of Section 2.3.3, the prover uses the secret knowledge of $\varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}} \circ \widehat{\varphi}_{\mathsf{com}}$ to evaluate φ_{rsp} on some torsion points. This torsion evaluation (along with $\deg(\varphi_{\mathsf{rsp}})$) is an efficient representation of φ_{rsp} that can be sent to the verifier. This makes SQIsignHD response considerably faster than SQIsign response. To verify the validity of this representation, the verifier computes a 4-dimensional (or 8-dimensional) isogeny that "embeds" φ_{rsp} by Kani's Lemma (Lemma 2.2.6). A low level implementation of 4-dimensional isogenies would be needed to make accurate comparisons, but SQIsignHD verification is expected to be slower than SQIsign verification, especially after the latest improvements of AprèsSQI [CRSEMR24].

This drawback has been solved by the 2-dimensional variants of SQIsign [BDF+25; NOC+25; DF25], including SQIsign2D-West that was proposed to the NIST and will be presented in Chapter 4.

Another advantage of SQIsignHD is the scalability of the prime parameter p to higher security levels. As we have seen in Section 2.1.4, to make the ideal-to-isogeny translation efficient in SQIsign, p was required to satisfy $2^f | p+1$ and $2^f T | p^2 - 1$ with $2^f = \Theta(p^{1/4})$ and $T = \Theta(p^{5/4})$ a powersmooth integer. The efficient search for such primes is still an active research question and scales badly as p grows, meaning that the search becomes harder and the smoothness bound we can impose on T grows as p grows, making the ideal-to-isogeny translation more costly. Among other alternative improvements [CRSEMR24; ON25], SQIsignHD completely solves this scalability problem by using primes of the form $p = c2^e - 1$ with $c \in \mathbb{N}^*$ small and $e \in \mathbb{N}^*$ growing linearly with the security level. Note that primes of the form $p = c2^e 3^{e'} - 1$ were initially used in SQIsignHD [DLRW24] but the power of 3 is no longer necessary with the new techniques presented in Section 2.4 that were introduced for SQIsign2D-West. Even though these were not the original techniques used in SQIsignHD, they greatly simplify the presentation while improving efficiency. We refer to [DLRW24] for a presentation of SQIsignHD with the former techniques.

As in [DLRW24], we introduce two versions of SQIsign:

- F-SQIsignHD, standing for Fast version of SQIsignHD, optimised for efficiency and requiring a 4-dimensional isogeny computation in the verification.

- R-SQIsignHD, standing for Rigorous version of SQIsignHD, optimised to make the security proof (of the zero-knowledge property) easier at the expense of efficiency and requiring an 8-dimensional isogeny computation in the verification.

Unlike F-SQIsignHD, R-SQIsignHD is purely theoretical and not implemented but demonstrates the potential of higher dimensional isogenies to obtain highly trustworthy security proofs.

3.1.3 The Fiat-Shamir transform

In this section, we explain how to transform our SQIsignHD identification protocol into a signature scheme using the Fiat-Shamir transform [FS87]. The method is analogous to the original SQIsign protocol.

In F-SQIsignHD, the challenge isogeny will be a 2^e-isogeny. In order to make the scheme non-interactive, we replace the call to the verifier by a hash function that generates it. The CGL hash function [CLG09] can be used to generate a cyclic 2^e-isogeny $\varphi := \mathsf{CGL}(E, s) : E \longrightarrow E'$ given a supersingular elliptic curve $E/\mathbb{F}_{p^2}$ and an integer $s \in [\![1 \ ; \ 2^e]\!]$. This isogeny is generated as a chain of 2-isogenies starting from E, where bits of s determine non-backtracking outgoing isogenies at every step. We also use another secure hash function $H : \{0,1\}^* \longrightarrow [\![1 \ ; \ 2^e]\!]$ to generate inputs s of CGL from messages and j-invariants. In R-SQIsignHD, the challenge isogeny will be of odd degree but a similar hashing method will apply, *e.g.* using the hash function from [DFSGF+21, § 3.1] generalising CGL.

Signature. To sign a message m with a secret key $\varphi_{\mathsf{sk}} : E_0 \longrightarrow E_{\mathsf{pk}}$, generate a random commitment $\varphi_{\mathsf{com}} : E_0 \longrightarrow E_{\mathsf{com}}$, let $s := H(j(E_{\mathsf{pk}}), j(E_{\mathsf{com}}), m)$ and $\varphi_{\mathsf{chl}} := \mathsf{CGL}(E_{\mathsf{pk}}, s) : E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}$. From the knowledge of φ_{sk}, φ_{chl} and φ_{com}, construct an efficient representation $\mathsf{rsp} = (\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}), \varphi_{\mathsf{rsp}}(Q_{\mathsf{com}}), q)$ given by the image of torsion points by a response isogeny $\varphi_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$ and return $(E_{\mathsf{com}}, \mathsf{rsp})$ as a signature.

Verification. A verifier receiving a signature $(E_{\mathsf{com}}, \mathsf{rsp})$ associated to the message m and public key E_{pk} computes $s = H(j(E_{\mathsf{pk}}), j(E_{\mathsf{com}}), m)$ and then $\varphi_{\mathsf{chl}} = \mathsf{CGL}(E_{\mathsf{pk}}, s) : E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}$. The verifier finally checks that rsp represents correctly an isogeny $\varphi_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$ by computing a higher dimensional isogeny embedding, as explained previously.

Once it is established that the SQIsignHD identification protocol is complete, sound, and honest verifier zero-knowledge, and assuming the hardness of the endomorphism ring problem (Problem 2.1.2), we obtain a universally unforgeable signature against chosen message attacks in the random oracle model [VV15, Theorem 7].

3.2 Algorithmic building blocks

In this section, we describe or recall the algorithms that will be used in the SQIsignHD protocol.

3.2.1 Ideal-to-isogeny translations and isogeny of fixed degree

Let us fix a prime of the form $p = c2^e - 1$, with $c \in \mathbb{N}^*$ odd and small and consider the starting supersingular curve $E_0/\mathbb{F}_{p^2}$ of equation $y^2 = x^3 + x$ whose endomorphism ring is isomorphic to a well known maximal order $\mathcal{O}_0 \subset \mathcal{B}_{p,\infty}$ (Lemma 1.2.25). We also fix a basis (P_0, Q_0) of $E_0[2^e]$. In this setting, we shall use the following algorithms:

- Algorithm 2.3 taking as input an odd integer $u < 2^e$ such that $u(2^e - u) = \Omega(p \log(p))$ and returning $(E_u, \varphi_u(P_0), \varphi_u(Q_0), I_u)$, where $\varphi_u : E_0 \longrightarrow E_u$ is a u-isogeny and $I_u \subset \mathcal{O}_0$ is its associated ideal.
- Algorithm 2.7 taking as input a left $\mathcal{O}_0$-ideal I (and some precomputed data) and returning the image $(E_I, \varphi_I(P_0), \varphi_I(Q_0))$ of the isogeny $\varphi_I : E_0 \longrightarrow E_I$ associated to I.
- The 4-dimensional techniques of Algorithm 2.2 to translate a connecting ideal I of (e, B)-good norm in the sense of Definition 2.3.1 between two maximal orders $\mathcal{O}_1$ and $\mathcal{O}_2$ into an isogeny $\varphi_I : E_1 \longrightarrow E_2$ when two isogenies $\varphi_1 : E_0 \longrightarrow E_1$ and $\varphi_2 : E_0 \longrightarrow E_2$ of odd degrees and their associated ideals $I_1, I_2 \subseteq \mathcal{O}_0$ (such that $O_R(I_1) = \mathcal{O}_1$ and $O_R(I_2) = \mathcal{O}_2$) are given. We shall adapt Algorithm 2.2 to our context: the evaluation of φ_I on 2^e-torsion (explained in Section 2.3.3) will take place during the response phase and the 4-dimensional embedding $F_I \in \mathrm{End}(E_1^2 \times E_2^2)$ of φ_I (as defined in Eq. (2.10)) will be computed during the verification phase.
- An algorithm due to Antonin Leroux [Ler22, Algorithm 19], Galbraith, Petit and Silva [GPS20, Algorithm 2] (described in Algorithm 3.1) taking as input a supersingular elliptic curve $E/\mathbb{F}_{p^2}$ of known endomorphism ring $\mathcal{O} \simeq \mathrm{End}(E)$ and a primitive $\mathcal{O}$-ideal I (in the sense of Definition 2.1.3) of smooth norm D and returning a kernel generator $P \in E[I]$. This direct ideal to kernel translation algorithm is efficient only when D is smooth and $E[D]$ is defined over a small extension of $\mathbb{F}_{p^2}$.

Algorithm 3.1: Ideal to kernel [Ler22, Algorithm 19].

Data: A supersingular elliptic curve $E/\mathbb{F}_{p^2}$, a basis $(\beta_1, \cdots, \beta_4)$ of a maximal order $\mathcal{O} \subset \mathcal{B}_{p,\infty}$ isomorphic to $\mathrm{End}(E)$, a left $\mathcal{O}$-ideal $I \subseteq \mathcal{O}$, a basis (P, Q) of $E[D]$ (where D is smooth), the images $(\varepsilon(\beta_i)(P), \varepsilon(\beta_i)(Q))_{1 \leq i \leq 4}$, where ε is an isomorphism $\mathcal{O} \xrightarrow{\sim} \mathrm{End}(E)$.

Result: A point $R \in E[D]$ of order D such that $E[I] = \langle P \rangle$.

Compute $\alpha \in I$ such that $I = \mathcal{O}\alpha + \mathcal{O}D$;
Write $\alpha := \sum_{i=1}^4 b_i \beta_i$, with $b_1, \cdots, b_4 \in \mathbb{Z}$;
$\varepsilon(\alpha)(P) \leftarrow \sum_{i=1}^4 [b_i]\varepsilon(\beta_i)(P)$, $\varepsilon(\alpha)(Q) \leftarrow \sum_{i=1}^4 [b_i]\varepsilon(\beta_i)(Q)$;
if *order*$(P) < D$ **then**
    Find $d \in \mathbb{Z}$ such that $\varepsilon(\alpha)(P) = [d]\varepsilon(\alpha)(Q)$;
    $R \leftarrow P - [d]Q$;
else
    Find $d \in \mathbb{Z}$ such that $\varepsilon(\alpha)(Q) = [d]\varepsilon(\alpha)(P)$;
    $R \leftarrow Q - [d]P$;
end
return R;

3.2.2 Isogeny to ideal

In addition to ideal-to-isogeny translation algorithms, we also need a converse algorithm translating an isogeny into an ideal (that will be applied to the challenge isogeny). Let $\varphi : E_1 \longrightarrow E_2$ be a cyclic

D-isogeny (with D smooth) between two supersingular elliptic curves over $\mathbb{F}_{p^2}$. We represent φ by a generator of its kernel P and suppose it is defined over a small extension of $\mathbb{F}_{p^2}$. We give an algorithm (Algorithm 3.2) also due to Antonin Leroux [Ler22, Algorithm 20], Galbraith, Petit and Silva [GPS20, Algorithm 3] to compute the ideal I_φ associated to φ when $\mathrm{End}(E_1)$ is known.

Algorithm 3.2: Kernel to ideal [Ler22, Algorithm 20].

Data: A point $P \in E_1$ of smooth order D, a basis $(\beta_1, \cdots, \beta_4)$ of a maximal order $\mathcal{O}_1 \subset \mathcal{B}_{p,\infty}$ isomorphic to $\mathrm{End}(E_1)$, a basis (R, S) of $E_1[D]$ and the images $(\varepsilon(\beta_i)(R), \varepsilon(\beta_i)(S))_{1\leq i\leq 4}$, where ε is an isomorphism $\mathcal{O}_1 \xrightarrow{\sim} \mathrm{End}(E_1)$.

Result: The left $\mathcal{O}_1$-ideal I_P associated to the isogeny of kernel $\langle P\rangle$.

Find $a, b \in \mathbb{Z}/D\mathbb{Z}$ such that $P = [a]R + [b]S$ (discrete logarithm);
for $i = 1$ **to** 4 **do**
  $Q_i \leftarrow [a]\varepsilon(\beta_i)(R) + [b]\varepsilon(\beta_i)(S) = \varepsilon(\beta_i)(P)$;
end
Find i, j such that (Q_i, Q_j) is a basis of $E_1[D]$;
For $k \neq i, j$, find $a, b \in \mathbb{Z}/D\mathbb{Z}$ such that $Q_k = aQ_i + bQ_j$ (discrete logarithm);
Let $\gamma := \beta_k - a\beta_i - b\beta_j$;
return $\mathcal{O}_1\gamma + \mathcal{O}_1 D$;

In order to apply Algorithm 3.2, one has to know the domain endomorphism ring $\mathrm{End}(E_1)$. In our application, we shall know the endomorphism ring of the public starting curve E_0 (of equation $y^2 = x^3 + x$) and an N-isogeny $\psi : E_0 \longrightarrow E_1$ along with its kernel ideal I_ψ. If N is coprime with D, we can then use ψ and I_ψ to compute a basis of $\mathrm{End}(E_1)$ and to evaluate it on $E_1[D]$ to obtain the required inputs of Algorithm 3.2 (see [DLRW24, Algorithm 8] or [EHLMP18, Algorithm 4]).

Instead, we use a simpler method also proposed in [AAA+25, § 4.4.3] that relies on pushforwards and pull-backs. The idea to compute the kernel ideal I_φ of $\varphi : E_1 \longrightarrow E_2$ given a kernel generator $P \in \ker(\varphi)$ is to use the equality $I_\varphi = [I_\psi]_*[I_\psi]^* I_\varphi$, valid when $N = \deg(\psi)$ and $D = \deg(\varphi)$ are coprime. Since $[I_\psi]^* I_\varphi$ is the kernel ideal of $\psi^*\varphi$ whose kernel is $\ker(\psi^*\varphi) = \langle\widehat{\psi}(P)\rangle$, we can simply apply Algorithm 3.2 to $\widehat{\psi}(P)$ (leveraging the knowledge of $\mathrm{End}(E_0)$) to obtain $[I_\psi]^* I_\varphi$ and then $I_\varphi = [I_\psi]_*[I_\psi]^* I_\varphi$. When N and D are coprime, $\widehat{\psi}(P)$ can easily be computed by evaluating a basis of $E_1[D]$ via ψ and computing a discrete logarithm as in Algorithm 2.9. We summarise this simpler method in Algorithm 3.3.

Algorithm 3.3: Kernel to ideal with a trapdoor isogeny.

Data: A point $P \in E_1$ of smooth order D, a basis $(\beta_1, \cdots, \beta_4)$ of a maximal order $\mathcal{O}_0 \subset \mathcal{B}_{p,\infty}$ isomorphic to $\mathrm{End}(E_0)$, a basis (R_0, S_0) of $E_0[D]$ and the images $(\varepsilon_0(\beta_i)(R_0), \varepsilon_0(\beta_i)(S_0))_{1\leq i\leq 4}$, where ε_0 is an isomorphism $\mathcal{O}_0 \xrightarrow{\sim} \mathrm{End}(E_0)$, and the image $(\psi(R_0), \psi(S_0))$ of an N-isogeny $\psi : E_0 \longrightarrow E_1$ with $\gcd(N, D) = 1$.

Result: The left $\mathcal{O}_1$-ideal I_P associated to the isogeny of kernel $\langle P\rangle$.

Find $a, b \in \mathbb{Z}/D\mathbb{Z}$ such that $P = [a]\psi(R_0) + [b]\psi(S_0)$ (discrete logarithm);
$\widehat{\psi}(P) \leftarrow [Na]R_0 + [Nb]S_0$;
Call Algorithm 3.2 on $\widehat{\psi}(P)$, D, $(\beta_1, \cdots, \beta_4)$ and $(\varepsilon_0(\beta_i)(R_0), \varepsilon_0(\beta_i)(S_0))_{1\leq i\leq 4}$ to obtain $[I_\psi]^* I_P$;
$I_P \leftarrow [I_\psi]_*[I_\psi]^* I_P$;
return I_P;

3.2.3 Sampling a uniformly random ideal of fixed norm

In the protocol, we shall need to uniformly sample at random cyclic isogenies $\varphi : E \longrightarrow E'$ of fixed degree N several times. When a maximal order $\mathcal{O} \subset \mathcal{B}_{p,\infty}$ isomorphic to $\mathrm{End}(E)$ is known, by the Deuring correspondence this reduces to sampling a primitive left $\mathcal{O}$-ideal I of norm N uniformly at random. I is then translated into an isogeny φ (*e.g.* using Algorithm 2.7 if $\mathcal{O} = \mathcal{O}_0$). For φ to be

cyclic, I has to be primitive in the sense of Definition 2.1.3, that is to say that $I \not\subseteq n\mathcal{O}$ for any integer $n > 1$.

Given a maximal quaternion order $\mathcal{O} \subseteq \mathcal{B}_{p,\infty}$ and an integer N coprime with p, we explain how to sample primitive left ideals $I \subseteq \mathcal{O}$ of norm N. The following analysis would be simpler if we assumed N prime but it still holds when N is composite and we found this of independent interest. We start by proving that primitive ideals of norm N are in bijection with

$$\mathbb{P}^1(\mathbb{Z}/N\mathbb{Z}) = \{(x : y) \in (\mathbb{Z}/N\mathbb{Z})^2 \mid \gcd(x, y, N) = 1\}.$$

Lemma 3.2.1. *[KV10, Lemma 7.2]*

(i) For all prime $\ell \neq p$, $\mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Z}_\ell \simeq M_2(\mathbb{Z}_\ell)$.

(ii) $\mathcal{O}/N\mathcal{O} \simeq M_2(\mathbb{Z}/N\mathbb{Z})$.

(iii) Let $\varphi_N : \mathcal{O}/N\mathcal{O} \xrightarrow{\sim} M_2(\mathbb{Z}/N\mathbb{Z})$ be an isomorphism and let us denote

$$\forall (x : y) \in \mathbb{P}^1(\mathbb{Z}/N\mathbb{Z}), \quad M_{x,y} := \begin{pmatrix} x & y \\ 0 & 0 \end{pmatrix}.$$

Then, the map:

$$\begin{array}{ccc} \mathbb{P}^1(\mathbb{Z}/N\mathbb{Z}) & \longrightarrow & \{\text{primitive left ideals } I \subseteq \mathcal{O} \text{ of norm } N\} \\ (x : y) & \longmapsto & I_{(x:y)} := \mathcal{O}\varphi_N^{-1}(\{M_{x,y}\}) + \mathcal{O}N \end{array}$$

is a bijection.

Proof. **(i)** Let ℓ be a prime $\neq p$. Then, $\mathcal{B}_{p,\infty}$ splits at ℓ so $\mathcal{B}_{p,\infty} \otimes \mathbb{Q}_\ell \simeq M_2(\mathbb{Q}_\ell)$. Hence, we have an embedding $\iota : \mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Z}_\ell \hookrightarrow M_2(\mathbb{Q}_\ell)$ mapping to an order $O = \iota(\mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Z}_\ell) \subset M_2(\mathbb{Q}_\ell)$. By Theorem 1.2.17, $\mathrm{discrd}(\mathcal{O}) = \mathrm{disc}(\mathcal{B}_{p,\infty}) = p$ and $1/p \in \mathbb{Z}_\ell$ so we easily obtain that O contains a suborder of reduced discriminant 1 so that $\mathrm{discrd}(O) = 1$ and O is maximal in $M_2(\mathbb{Q}_\ell)$ by [Voi21, Theorem 15.5.3]. It follows by [Voi21, Corollary 10.5.5] that O is conjugate to $M_2(\mathbb{Z}_\ell)$, so that $\mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Z}_\ell \simeq M_2(\mathbb{Z}_\ell)$, as desired.

(ii) By point (i), we have for all prime $\ell \neq p$ and $e \in \mathbb{N}^*$,

$$\mathcal{O}/\ell^e\mathcal{O} \simeq \mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Z}_\ell/(\ell^e\mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Z}_\ell) \simeq M_2(\mathbb{Z}_\ell/\ell^e\mathbb{Z}_\ell) = M_2(\mathbb{Z}/\ell^e\mathbb{Z}).$$

Then, if we decompose $N = \prod_{i=1}^r \ell_i^{e_i}$ as product of distinct prime factors, then the Chinese remainder theorem ensures that:

$$\mathcal{O}/N\mathcal{O} \simeq \prod_{i=1}^r \mathcal{O}/\ell_i^{e_i}\mathcal{O} \simeq \prod_{i=1}^r M_2(\mathbb{Z}/\ell_i^{e_i}\mathbb{Z}) \simeq M_2(\mathbb{Z}/N\mathbb{Z}).$$

(iii) We first prove that the map is well defined. Let $(x : y) \in \mathbb{P}^1(\mathbb{Z}/N\mathbb{Z})$. Then $I_{(x:y)} := \mathcal{O}\varphi_N^{-1}(\{M_{x,y}\}) + \mathcal{O}N$ is a left $\mathcal{O}$-ideal. If $I_{(x:y)} \subseteq n\mathcal{O}$ for some $n \in \mathbb{N}^*$, then $n|M_{x,y}$ and $n|N$ so $n = 1$ since $\gcd(x, y, N) = 1$ so $I_{(x:y)}$ is indeed primitive.

Now, we prove that $\mathrm{nrd}(I_{(x:y)}) = N$. By [Voi21, Lemma 3.4.2], we first observe that for all prime $\ell \neq p$, an isomorphism $\mathcal{B}_{p,\infty} \otimes \mathbb{Q}_\ell \xrightarrow{\sim} M_2(\mathbb{Q}_\ell)$ identifies the conjugation of $\mathcal{B}_{p,\infty}$ with the standard involution of $M_2(\mathbb{Q}_\ell)$ given by the transpose of the comatrix

$$A = \begin{pmatrix} a & c \\ b & d \end{pmatrix} \in M_2(\mathbb{Q}_\ell) \longmapsto A^\dagger := \begin{pmatrix} d & -c \\ -b & a \end{pmatrix} \in M_2(\mathbb{Q}_\ell)$$

By the Chinese remainder theorem, it follows that for all $\alpha \in \mathcal{O}$, $\varphi_N(\overline{\alpha}) = \varphi_N(\alpha)^\dagger$, so that $\overline{I}_{(x:y)} = \mathcal{O}\varphi_N^{-1}(\{M_{x,y}^\dagger\}) + \mathcal{O}N$. Since $M_{x,y} \cdot M_{x,y}^\dagger = 0$, we then have $I_{(x:y)} \cdot \overline{I}_{(x:y)} = \mathcal{O}\varphi_N^{-1}(\{0\}) = \mathcal{O}N$, so that $\mathrm{nrd}(I_{(x:y)})^2 = \mathrm{nrd}(I_{(x:y)} \cdot \overline{I}_{(x:y)}) = N^2$ and $\mathrm{nrd}(I_{(x:y)}) = N$. This proves that the map $(x : y) \longmapsto I_{(x:y)}$ is well-defined.

Now, this map is clearly injective. Indeed, if $(x:y),(x':y') \in \mathbb{P}^1(\mathbb{Z}/N\mathbb{Z})$ are such that $I_{(x:y)} = I_{(x':y')}$, then we have $M_2(\mathbb{Z}/N\mathbb{Z})M_{x,y} = \varphi_N(I_{(x:y)}) = \varphi_N(I_{(x':y')}) = M_2(\mathbb{Z}/N\mathbb{Z})M_{x',y'}$, so there exists $M \in M_2(\mathbb{Z}/N\mathbb{Z})$ such that $M_{x',y'} = M \cdot M_{x,y}$, so that $(x',y') = (M_{1,1}x, M_{1,1}y)$ and $(x:y) = (x':y')$ in $\mathbb{P}^1(\mathbb{Z}/N\mathbb{Z})$.

Finally, we prove that $(x:y) \longmapsto I_{(x:y)}$ is surjective, which is the most delicate point. Let I be a left $\mathcal{O}$-ideal of norm N. Then I is locally principal by Theorem 1.2.10 so we may write for all prime $\ell|N$, $I \otimes \mathbb{Z}_\ell = (\mathcal{O} \otimes \mathbb{Z}_\ell) \cdot \alpha_\ell$ with $\alpha_\ell \in \mathcal{O} \otimes \mathbb{Z}_\ell$. Let $\ell|N$, $e := v_\ell(N)$ and $\varphi_\ell : \mathcal{O} \otimes \mathbb{Z}_\ell \xrightarrow{\sim} M_2(\mathbb{Z}_\ell)$ be an isomorphism. Consider the $\mathbb{Z}/\ell^e\mathbb{Z}$-module:

$$V := \{v \in (\mathbb{Z}/\ell^e\mathbb{Z})^2 \mid \varphi_\ell(\alpha_\ell) \cdot v \equiv 0 \mod \ell^e\}.$$

We prove that V has rank 1. Since I is primitive, ℓ does not divide α_ℓ so $\varphi_\ell(\alpha_\ell) \not\equiv 0 \mod \ell$ and $V/\ell V$ does not fill the whole of $(\mathbb{Z}/\ell\mathbb{Z})^2$. Besides,

$$\mathrm{nrd}(I \otimes \mathbb{Z}_\ell) = \langle \mathrm{nrd}(\alpha) \mid \alpha \in I \otimes \mathbb{Z}_\ell \rangle = N\mathbb{Z}_\ell = \ell^e\mathbb{Z}_\ell,$$

so that $\mathrm{nrd}(\alpha_\ell) = \ell^e$ by Lemma 1.2.6 *i.e.* $\alpha_\ell\overline{\alpha}_\ell = \ell^e$. It follows that $\varphi_\ell(\alpha_\ell) \cdot \varphi_\ell(\alpha_\ell)^\dagger \equiv 0 \mod \ell^e$ so $\varphi_\ell(\alpha_\ell)$ is not invertible modulo ℓ so that $V/\ell V \neq 0$ and $\dim_{\mathbb{Z}/\ell\mathbb{Z}}(V/\ell V) = 1$. Besides, $\varphi_\ell(\alpha_\ell) \not\equiv 0 \mod \ell$ so $\varphi_\ell(\alpha_\ell)^\dagger \not\equiv 0 \mod \ell$ and one of the columns of $\varphi_\ell(\alpha_\ell)^\dagger$ yields a non-trivial element $v_0 \in V$ that generates $V/\ell V$.

Now, if $v \in V$, we prove by induction on $i \in [\![1\ ;\ e]\!]$ that $v \in \mathbb{Z}/\ell^i\mathbb{Z}v_0$, which will prove that V has rank 1, as desired. First, since $V/\ell V$ has dimension 1 as a $\mathbb{Z}/\ell\mathbb{Z}$-vector space, there exists $\lambda_1 \in \mathbb{Z}$ such that $v \equiv \lambda_1 v_0 \mod \ell$. Now, if $i \in [\![1\ ;\ e-1]\!]$ and $v \equiv \lambda_i v_0 \mod \ell^i$ for some $\lambda_i \in \mathbb{Z}$, we can write $v - \lambda_i v_0 = \ell^i w_i$ with $w_i \in (\mathbb{Z}/\ell^e\mathbb{Z})^2$ and we have $\ell^i\varphi_\ell(\alpha_\ell)\cdot w_i \equiv \varphi_\ell(\alpha_\ell)\cdot(v - \lambda_i v_0) \equiv 0 \mod \ell^e$, so that $\varphi_\ell(\alpha_\ell)\cdot w_i \equiv 0 \mod \ell^{e-i} \equiv 0 \mod \ell$ and $w_i \equiv \lambda' v_0 \mod \ell$ for some $\lambda' \in \mathbb{Z}$. Hence, $v \equiv (\lambda_i + \ell^i\lambda')v_0 \mod \ell^{i+1}$ and the induction follows.

Let us denote $v_0 := (-y, x) \in (\mathbb{Z}/\ell^e\mathbb{Z})^2$ the generator of V. We notice that for all $M \in G_2(\mathbb{Z}_\ell)$, $\varphi_\ell^{-1}(M)\alpha_\ell$ also generates $I \otimes \mathbb{Z}_\ell$ so we can make operations on the lines of $\varphi_\ell(\alpha_\ell)$. Since $\varphi_\ell(\alpha_\ell)$ has rank 1 modulo ℓ, by operations on the lines of $\varphi_\ell(\alpha_\ell)$, we may assume that

$$\varphi_\ell(\alpha_\ell) = \begin{pmatrix} a & b \\ 0 & 0 \end{pmatrix},$$

with $a, b \in \mathbb{Z}_\ell$, so that $-ay + bx \equiv 0 \mod \ell^e$. Since $v_0 \not\equiv 0 \mod \ell$, we have $\gcd(x, y, \ell) = 1$ so there exists $u, v \in \mathbb{Z}$ such that $ux + vy \equiv 1 \mod \ell^e$. Combining this equation with $-ay + bx \equiv 0 \mod \ell^e$ multiplied by u and v, we obtain that $b = (ua + vb)y$ and $a = (ua + vb)x$ so that $(a:b) = (x:y) \in \mathbb{P}^1(\mathbb{Z}/\ell^e\mathbb{Z})$. Since this is valid for all $\ell|N$, we obtain by the Chinese remainder theorem the existence of $(x:y) \in \mathbb{P}^1(\mathbb{Z}/N\mathbb{Z})$ such that $M_{x,y} \in \varphi_N(I)$ and $I_{(x:y)} = \mathcal{O}\varphi_N^{-1}(\{M_{x,y}\}) + \mathcal{O}N \subseteq I$. Since both $I_{(x:y)}$ and I have norm N, we conclude that $I_{(x:y)} = I$, which completes the proof. □

As a direct consequence of Lemma 3.2.1 above, we obtain:

Lemma 3.2.2. *The set of elements $\alpha \in \mathcal{O}$ invertible modulo N acts transitively (by multiplication on the right) on the set of primitive left $\mathcal{O}$-ideals of norm N. Those elements $\alpha \in \mathcal{O}$ invertible modulo N are those of norm coprime with N.*

Proof. Let I be a primitive left $\mathcal{O}$-ideal of norm N. Then, the ideal I corresponds to $(x:y) \in \mathbb{P}^1(\mathbb{Z}/N\mathbb{Z})$ via the bijection of Lemma 3.2.1 and is isomorphic to $M_2(\mathbb{Z}/N\mathbb{Z})\cdot M_{x,y}$ via the composition of the reduction modulo N and φ_N. For any representative $(x, y) \in \mathbb{Z}^2$ of $(x:y) \in \mathbb{P}^1(\mathbb{Z}/N\mathbb{Z})$, we have $\gcd(x, y, N) = 1$ so we may find $u, v \in \mathbb{Z}$ such that $xu + yv \equiv 1 \mod N$, so that:

$$M_{x,y}\begin{pmatrix} u & -y \\ v & x \end{pmatrix} \equiv M_{1,0} \mod N \quad \text{and} \quad \det\begin{pmatrix} u & -y \\ v & x \end{pmatrix} \equiv 1 \mod N$$

Hence, the ideal $M_2(\mathbb{Z}/N\mathbb{Z}) \cdot M_{x,y}$ is in the orbit of $M_2(\mathbb{Z}/N\mathbb{Z}) \cdot M_{1,0}$ under the right action of $GL_2(\mathbb{Z}/N\mathbb{Z})$, and as a consequence, $I/N\mathcal{O}$ is in the orbit of the ideal $I_0/N\mathcal{O} := \mathcal{O}\varphi_N^{-1}(\{M_{1,0}\})/N\mathcal{O}$ under the right action of $(\mathcal{O}/N\mathcal{O})^*$.

To conclude, it suffices to prove that the invertible elements of $\mathcal{O}$ modulo N are those of norm coprime with N. If $\alpha \in \mathcal{O}$ is invertible modulo N, there exists $\beta, \gamma \in \mathcal{O}$ such that $\alpha\beta = 1 + N\gamma$, so that

$$\operatorname{nrd}(\alpha)\operatorname{nrd}(\beta) = \operatorname{nrd}(1 + N\gamma) = 1 + N\operatorname{Tr}(\gamma) + N^2\operatorname{nrd}(\gamma) \equiv 1 \mod N,$$

so $\operatorname{nrd}(\alpha)$ is invertible modulo N. Conversely, if $\operatorname{nrd}(\alpha)$ is coprime with N, there exists $\lambda \in \mathbb{Z}$ such that $\operatorname{nrd}(\alpha)\lambda \equiv 1 \mod N$. Then, it follows that $\alpha\overline{\alpha}\lambda \equiv 1 \mod N$, so α is invertible modulo N. This completes the proof. □

Lemma 3.2.2 ensures that $(\mathcal{O}/N\mathcal{O})^*$ acts transitively on primitive left ideals of norm N by multiplication on the right. Hence, given a primitive left $\mathcal{O}$-ideal I_0 of norm N, if we sample $[\alpha] \in (\mathcal{O}/N\mathcal{O})^*$ uniformly at random, then $I_0\alpha + N\mathcal{O}$ is uniformly random among primitive left $\mathcal{O}$-ideals of norm N.

To obtain such an ideal I_0, we compute $\gamma \in \mathcal{O}$ of norm NM with $\gcd(N, M) = 1$ and without integral factor. This can be done with the algorithms of [Ler22, Section 3.3]. We then consider $I_0 := \mathcal{O}\gamma + \mathcal{O}N$ and sample $[\alpha] \in \mathcal{O}/N\mathcal{O}$ uniformly at random until it is invertible modulo N (which can be checked by computing $\operatorname{nrd}(\alpha)$). These operations are summarised in Algorithm 3.4.

Lemma 3.2.3. *Algorithm 3.4 terminates after $O(\log\log(N))$ iterations on average.*

Proof. Via the isomorphism $\varphi_N : \mathcal{O}/N\mathcal{O} \xrightarrow{\sim} M_2(\mathbb{Z}/N\mathbb{Z})$ from Lemma 3.2.1, we see that the invertible elements of $\mathcal{O}/N\mathcal{O}$ are in bijection with matrices of $GL_2(\mathbb{Z}/N\mathbb{Z})$. By the Chinese remainder theorem, if $N = \prod_{i=1}^{r} \ell_i^{e_i}$ is the prime decomposition of N, we have:

$$GL_2(\mathbb{Z}/N\mathbb{Z}) \simeq \prod_{i=1}^{r} GL_2(\mathbb{Z}/\ell_i^{e_i}\mathbb{Z}).$$

If ℓ is a prime, then we have

$$|GL_2(\mathbb{Z}/\ell\mathbb{Z})| = (\ell^2 - 1)(\ell^2 - \ell).$$

To find this formula, we enumerate all the possible values for the first column ($\ell^2 - 1$ non-zero columns) and the second ($\ell^2 - \ell$ columns linearly independent from the first one). Now, for $e \in \mathbb{N}^*$, we consider the surjective group homomorphism $GL_2(\mathbb{Z}/\ell^{e+1}\mathbb{Z}) \twoheadrightarrow GL_2(\mathbb{Z}/\ell^e\mathbb{Z})$ given by the reduction mod ℓ^e. We immediately see that any kernel element is of the form $I_2 + \ell^e M$ with $M \in M_2(\mathbb{Z}/\ell\mathbb{Z})$ so that $\det(I_2 + \ell^e M) \equiv \det(I_2) \equiv 1 \mod \ell^e$ and $I_2 + \ell^e M$ is invertible. Hence, the kernel is exactly $I_2 + \ell^e M_2(\mathbb{Z}/\ell\mathbb{Z})$, which has cardinality $|M_2(\mathbb{Z}/\ell\mathbb{Z}) = \ell^4$, so that

$$|GL_2(\mathbb{Z}/\ell^{e+1}\mathbb{Z})| = \ell^4 |GL_2(\mathbb{Z}/\ell^e\mathbb{Z})|$$

and by an easy induction:

$$\forall e \in \mathbb{N}^*, \quad |GL_2(\mathbb{Z}/\ell^e\mathbb{Z})| = \ell^{4(e-1)}(\ell^2 - 1)(\ell^2 - \ell).$$

We conclude that the probability that an element of $\mathcal{O}/N\mathcal{O}$ is invertible is:

$$\begin{aligned}\frac{|GL_2(\mathbb{Z}/N\mathbb{Z})|}{|M_2(\mathbb{Z}/N\mathbb{Z})|} &= \frac{1}{N^4}\prod_{i=1}^{r}|GL_2(\mathbb{Z}/\ell_i^{e_i}\mathbb{Z})| = \frac{1}{N^4}\prod_{i=1}^{r}\ell_i^{4(e_i-1)}(\ell_i^2 - 1)(\ell_i^2 - \ell_i) \\ &= \prod_{\ell | N}\left(1 - \frac{1}{\ell^2}\right)\left(1 - \frac{1}{\ell}\right)\end{aligned}$$

Since the series $\sum_\ell 1/\ell^2$ converges, the series $\sum_\ell \log(1 - 1/\ell^2)$ also converges and there is a universal lower bound for the product $\prod_{\ell|N}(1 - 1/\ell^2)$. By [HW75, Theorem 328], we also know that $\prod_{\ell|N}(1 - 1/\ell) = \Omega(1/\log\log(N))$. This completes the proof. □

Algorithm 3.4: Random ideal of fixed norm.

Data: A maximal order $\mathcal{O} \subseteq \mathcal{B}_{p,\infty}$ and an integer N such that $p \nmid N$.
Result: A primitive left $\mathcal{O}$-ideal I of norm N sampled uniformly at random.

Find $\gamma \in \mathcal{O}$ primitive of norm NM with $\gcd(N, M) = 1$; // Using [Ler22, Algorithm 2]
repeat
  Sample $u_1, \cdots, u_4 \in [\![0\ ;\ N-1]\!]$ uniformly at random;
  $\alpha \leftarrow \sum_{i=1}^{4} u_i \alpha_i$, where $(\alpha_1, \cdots, \alpha_4)$ is a basis of $\mathcal{O}$;
until $\gcd(\mathrm{nrd}(\alpha), N) = 1$;
$I \leftarrow \mathcal{O}\gamma\alpha + N\mathcal{O}$;
return I;

3.2.4 Sampling a uniformly random ideal of bounded small norm

During the response phase, the prover has to generate an ideal I_{rsp} of small norm equivalent to another ideal I. For security reasons, we require that I_{rsp} is uniformly random among ideals equivalent to I of norms bounded by some number $N = \Theta(\sqrt{p})$. By Lemma 1.2.19.(i), we can sample $I_{\mathsf{rsp}} \sim I$ by sampling an element of I uniformly at random among elements of norm $\leq N\,\mathrm{nrd}(I)$. In this section, to solve that problem, we propose an algorithm to sample an element uniformly at random in the intersection of a ball with a euclidean lattice. The approach follows from [DLRW24, Lemma 11], generalizing [Wes22, Lemma 3.3]. Another approach easier to implement has been proposed and implemented for the NIST submission [AAA+25, Algorithm 3.7] but the proofs are heuristic. The method we propose here is slower in practice but proven.

Sampling uniformly in the intersection of a ball and a lattice

Let $\Lambda \subset \mathbb{R}^d$ be a full-rank lattice and $(b_1, \cdots, b_d)$ be a reduced basis of Λ (*e.g.* Minkowski, LLL, or BKZ-reduced). Consider the canonical Euclidean norm $\|.\|$ and let $\rho > 0$ big enough. Our goal is to sample uniformly at random in the set:

$$B(0, \rho) \cap \Lambda \setminus \{0\} = \{x \in \Lambda \setminus \{0\} \mid \|x\| \leq \rho\}.$$

This can be done with Algorithm 3.5.

Algorithm 3.5: Sampling an element in the intersection of a ball with a lattice.

Data: A "reduced" basis $(b_1, \cdots, b_d)$ of a full-rank lattice $\Lambda \subset \mathbb{R}^d$ and a bound $\rho \geq \max(\|b_d\|, d(2\prod_{i=1}^{d} \|b_i\|)^{1/d})$.
Result: An element $x \in B(0, \rho) \cap \Lambda \setminus \{0\}$ sampled uniformly at random.

$\nu \leftarrow \sqrt{d}\|b_d\|/2$;
while *True* **do**
  Sample $u \in B(0, \rho + \nu)$ uniformly at random;
  Find a solution $\lambda(v) \in \Lambda$ of the closest vector problem (CVP) for v;
  if $\lambda(v) \in B(0, \rho) \setminus \{0\}$ **then**
    return $\lambda(v)$;
  end
end

Finding a solution to CVP takes exponential time in the dimension d but we apply Algorithm 3.5 to lattices of dimension $d = 4$ in practice. Hence, the limiting factor is the success probability of each iteration that we estimate in the following.

Lemma 3.2.4. *Let $B := (b_1, \cdots, b_d)$ be a basis of a full-rank lattice $\Lambda \subset \mathbb{R}^d$ and a bound $\rho \geq \max(\|b_d\|, d(2\prod_{i=1}^{d} \|b_i\|)^{1/d})$. Then with these inputs, Algorithm 3.5 returns elements of $B(0, \rho) \cap \Lambda \setminus$*

$\{0\}$ *with uniform distribution and terminates after an average number of iterations bounded by:*

$$\frac{2d^d\pi^{\frac{d}{2}}\left(1+\frac{\sqrt{d}}{2}\right)^d}{\Gamma\left(\frac{d}{2}+1\right)}\frac{\prod_{i=1}^d\|b_i\|}{\mathrm{Covol}(\Lambda)},$$

where Γ *is Euler's gamma function.*

Proof. First, we prove that Algorithm 3.5 returns elements with uniform distribution. Let $\mathcal{V} := \{v \in \mathbb{R}^d \mid \|v\| = \min_{\lambda\in\Lambda}\|v+\lambda\|\}$ be the Voronoi cell at the origin. Then, the closest vector $\lambda(v)$ satisfies $v \in \mathcal{V} + \lambda(v)$ and $\lambda(v)$ is unique when v is not at the border of a Voronoi cell, so it is unique with probability 1. Hence, for all $u \in B(0,\rho)\cap\Lambda$,

$$\mathbb{P}(\lambda(v)=u) = \frac{\mathrm{Vol}((\mathcal{V}+u)\cap B(0,\rho+\nu))}{\mathrm{Vol}(B(0,\rho+\nu))}.$$

Let $\mu := \inf\{r>0 \mid \forall v\in\mathbb{R}^d,\ \exists\lambda\in\Lambda, \|x-\lambda\|\le r\}$ be the covering radius of Λ. Then $\mathcal{V}\subseteq B(0,\mu)$ and a classical result [MG02, Theorem 7.9] ensures that $\mu \le \sqrt{d}\lambda_d/2$ where λ_d is the last minimum of Λ, so that $\mu \le \sqrt{d}\|b_d\|/2 = \nu$. It follows that $\mathcal{V}+u \subseteq B(0,\rho+\nu)$ for all $u\in B(0,\rho)\cap\Lambda$. Hence

$$\mathbb{P}(\lambda(v)=u) = \frac{\mathrm{Vol}(\mathcal{V}+u)}{\mathrm{Vol}(B(0,\rho+\nu))} = \frac{\mathrm{Vol}(\mathcal{V})}{\mathrm{Vol}(B(0,\rho+\nu))} = \frac{\mathrm{Covol}(\Lambda)}{\mathrm{Vol}(B(0,\rho+\nu))},$$

where the last equality holds because $\mathcal{V}$ is a fundamental domain. This quantity does not depend on u so the returned value $\lambda(v)$ has uniform distribution. In addition, the probability that each iteration of Algorithm 3.5 terminates is:

$$\begin{aligned}\mathbb{P}(\lambda(v)\in B(0,\rho)\setminus\{0\}) &= \#(B(0,\rho)\cap\Lambda\setminus\{0\})\frac{\mathrm{Covol}(\Lambda)}{\mathrm{Vol}(B(0,\rho+\nu))}\\ &= \#(B(0,\rho)\cap\Lambda\setminus\{0\})\frac{\mathrm{Covol}(\Lambda)\Gamma\left(\frac{d}{2}+1\right)}{\pi^{\frac{d}{2}}(\rho+\nu)^d}\\ &\ge \#(B(0,\rho)\cap\Lambda\setminus\{0\})\frac{\mathrm{Covol}(\Lambda)\Gamma\left(\frac{d}{2}+1\right)}{\pi^{\frac{d}{2}}\left(1+\frac{\sqrt{d}}{2}\right)^d\rho^d},\end{aligned} \tag{3.1}$$

where we used the fact that $\nu = \sqrt{d}\|b_d\|/2 \le \sqrt{d}\rho/2$.

As in Eq. (2.13), we consider

$$P_d(B) := \left\{\sum_{i=1}^d x_i b_i \;\middle|\; \forall i\in[\![1\ ;\ d]\!],\quad x_i\in[\![-B_i\ ;\ B_i]\!]\right\},$$

where for all $i\in[\![1\ ;\ d]\!]$, $B_i := \lfloor\rho/(d\|b_i\|)\rfloor$. Then, by Cauchy-Schwartz inequality $P_d(B)\subseteq B(0,\rho)\cap\Lambda$, so that:

$$\begin{aligned}\#(B(0,\rho)\cap\Lambda\setminus\{0\}) \ge \#P_d(B) - 1 &= \prod_{i=1}^d(2B_i+1) - 1 = \prod_{i=1}^d\left(2\left\lfloor\frac{\rho}{d\|b_i\|}\right\rfloor+1\right)-1\\ &\ge \prod_{i=1}^d\frac{\rho}{d\|b_i\|} - 1 = \frac{\rho^d}{d^d\prod_{i=1}^d\|b_i\|} - 1 \ge \frac{\rho^d}{2d^d\prod_{i=1}^d\|b_i\|},\end{aligned}$$

where the last equality follows from $\rho \ge d(2\prod_{i=1}^d\|b_i\|)^{1/d}$.

Finally, combining the above equality with Eq. (3.1), it follows that:

$$\mathbb{P}(\lambda(v)\in B(0,\rho)\setminus\{0\}) \ge \frac{\Gamma\left(\frac{d}{2}+1\right)}{2d^d\pi^{\frac{d}{2}}\left(1+\frac{\sqrt{d}}{2}\right)^d}\frac{\mathrm{Covol}(\Lambda)}{\prod_{i=1}^d\|b_i\|},$$

and the average number of iterations of Algorithm 3.5 is $1/\mathbb{P}(\lambda(v)\in B(0,\rho)\setminus\{0\})$. This completes the proof. □

The complexity bound in Lemma 3.2.4 depends on the quality of the reduced basis we give as input. The closer $\prod_{i=1}^d\|b_i\|$ approaches $\mathrm{Covol}(\Lambda)$, the better.

The four dimensional case

For our application, we work with lattices of rank $d = 4$ and we can efficiently compute a Minkowski reduced basis $B := (b_1, \cdots, b_4)$ of a full rank lattice $\Lambda \subset \mathbb{R}^4$ (Theorem 1.2.29). In general, the lattice can be unbalanced and have a very big last minimum so the condition $\rho \geq \|b_4\|$ of Algorithm 3.5 can be hard to satisfy. If $\|b_3\| \leq \rho < \|b_4\|$ we know that any element in $B(0,\rho) \cap \Lambda$ will be a linear combination of $b_1, \cdots, b_3$ only since B is Minkowski reduced so we can restrict the sampling to $\Lambda_3 := \langle b_1, b_2, b_3 \rangle$. If $\|b_2\| \leq \rho < \|b_3\|$, we restrict the sampling to $\Lambda_2 := \langle b_1, b_2 \rangle$ and if $\|b_1\| \leq \rho < \|b_2\|$, we restrict to $\Lambda_1 := \langle b_1 \rangle$. The question is whether the condition $\rho \geq d(2\prod_{i=1}^{d} \|b_i\|)^{1/d}$ is satisfied for $d \in [\![1\ ;\ 4]\!]$.

Lemma 3.2.5. *Let $\Lambda \subset \mathbb{R}^4$ be a full rank lattice and $B := (b_1, \cdots, b_4)$ be a Minkowski reduced lattice. Then, we have for all $d \in [\![1\ ;\ 4]\!]$,*

$$d\left(2\prod_{i=1}^{d}\|b_i\|\right)^{\frac{1}{d}} \leq \frac{12\cdot 3^{\frac{1}{12}}}{\sqrt{\pi}}\operatorname{Covol}(\Lambda)^{\frac{1}{4}}.$$

Proof. Since the $\|b_i\|$ are the successive minima of Λ by Theorem 1.2.28, we have by Minkowski's second theorem:

$$\frac{2^4}{4!}\frac{\operatorname{Covol}(\Lambda)}{\operatorname{Vol}(B(0,1))} = \frac{1}{3\pi^2}\operatorname{Covol}(\Lambda) \leq \prod_{i=1}^{4}\|b_i\| \leq 2^4\frac{\operatorname{Covol}(\Lambda)}{\operatorname{Vol}(B(0,1))} = \frac{32}{\pi^2}\operatorname{Covol}(\Lambda) \tag{3.2}$$

It follows that ($d = 4$):

$$4\left(2\prod_{i=1}^{4}\|b_i\|\right)^{\frac{1}{4}} \leq \frac{8\sqrt{2}}{\sqrt{\pi}}\operatorname{Covol}(\Lambda)^{\frac{1}{4}}.$$

Besides the upper bound of Eq. (3.2) ensures that:

$$\prod_{i=1}^{3}\|b_i\| \leq \frac{32\operatorname{Covol}(\Lambda)}{\pi^2\|b_4\|},$$

and the lower bound ensures that:

$$\|b_4\| \geq \left(\frac{1}{3\pi^2}\operatorname{Covol}(\Lambda)\right)^{\frac{1}{4}},$$

so that:

$$\prod_{i=1}^{3}\|b_i\| \leq \frac{32\cdot 3^{\frac{1}{4}}\operatorname{Covol}(\Lambda)^{\frac{3}{4}}}{\pi^{\frac{3}{2}}},$$

and ($d = 3$):

$$3\left(2\prod_{i=1}^{3}\|b_i\|\right)^{\frac{1}{3}} \leq \frac{12\cdot 3^{\frac{1}{12}}}{\sqrt{\pi}}\operatorname{Covol}(\Lambda)^{\frac{1}{4}}.$$

The upper bound of Eq. (3.2) also ensures that:

$$\|b_1\|\|b_2\| \leq \frac{32\operatorname{Covol}(\Lambda)}{\pi^2\|b_3\|\|b_4\|},$$

and the lower bound ensures that:

$$\|b_4\|\|b_3\| \geq \left(\frac{1}{3\pi^2}\operatorname{Covol}(\Lambda)\right)^{\frac{1}{2}},$$

so that:

$$\|b_1\|\|b_2\| \leq \frac{32\sqrt{3}\operatorname{Covol}(\Lambda)^{\frac{1}{2}}}{\pi},$$

and $(d = 2)$:

$$2(2\|b_1\|\|b_2\|)^{\frac{1}{2}} \leq \frac{8 \cdot 3^{\frac{1}{4}}}{\sqrt{\pi}} \operatorname{Covol}(\Lambda)^{\frac{1}{4}}.$$

Finally, the upper bound of Eq. (3.2) ensures that $(d = 1)$:

$$2\|b_1\| \leq \frac{2^{\frac{9}{4}}}{\sqrt{\pi}} \operatorname{Covol}(\Lambda)^{\frac{1}{4}}.$$

Since $\max(8\sqrt{2}, 12 \cdot 3^{1/12}, 8 \cdot 3^{1/4}, 2^{9/4}) = 12 \cdot 3^{1/12}$, the result follows. □

Hence, we can adapt Algorithm 3.5 to the rank 4 case and lower bound ρ as follows.

Algorithm 3.6: Sampling an element in the intersection of a ball with a lattice of rank four.

Data: A Minkowski reduced basis $(b_1, \cdots, b_d)$ of a full-rank lattice $\Lambda \subset \mathbb{R}^4$ and a bound $\rho \geq 12 \cdot 3^{1/12}/\sqrt{\pi} \operatorname{Covol}(\Lambda)^{1/4}$.
Result: An element $x \in B(0, \rho) \cap \Lambda \setminus \{0\}$ sampled uniformly at random.

$d \leftarrow 4$;
if $\rho < \|b_4\|$ **then**
  Find the biggest $d \in [\![1\ ;\ 3]\!]$ such that $\rho \geq \|b_d\|$;
end
Call Algorithm 3.5 with $(b_1, \cdots, b_d)$ and ρ;

Lemma 3.2.6. *Algorithm 3.6 is correct and terminates with an average number of iterations bounded by* 131072.

Proof. The correctness of Algorithm 3.6 follows from Lemma 3.2.5, the correctness of Algorithm 3.5 proved in Lemma 3.2.4, and the fact that the basis given on entry is Minkowski reduced.

Now, if $d \leq 4$ and $\Lambda \subset \mathbb{R}^4$, we can without loss of generality work in $\operatorname{Span}(\Lambda) \simeq \mathbb{R}^d$ and assume that Λ is a full rank lattice of $\mathbb{R}^d$. Let $(b_1, \cdots, b_d)$ be a Minkowski reduced basis of Λ. Then by Lemma 3.2.4, Algorithm 3.5 terminates after at most

$$\frac{2d^d \pi^{\frac{d}{2}} \left(1 + \frac{\sqrt{d}}{2}\right)^d}{\Gamma\left(\frac{d}{2} + 1\right)} \frac{\prod_{i=1}^d \|b_i\|}{\operatorname{Covol}(\Lambda)}$$

iterations on average, and by Theorem 1.2.28 and by Minkowski's second theorem, we have

$$\prod_{i=1}^d \|b_i\| \leq \frac{2^d}{\operatorname{Vol}(B(0,1))} \operatorname{Covol}(\Lambda) = \frac{2^d \Gamma\left(\frac{d}{2} + 1\right)}{\pi^{\frac{d}{2}}} \operatorname{Covol}(\Lambda),$$

so the average number of iterations is bounded by

$$2^{d+1} d^d \left(1 + \frac{\sqrt{d}}{2}\right)^d \leq 2^{4+1} 4^4 \left(1 + \frac{\sqrt{4}}{2}\right)^4 = 131072.$$

This completes the proof. □

In our application, we consider a maximal order $\mathcal{O} \subset \mathcal{B}_{p,\infty}$ and a left $\mathcal{O}$-ideal I. As explained earlier, to find an equivalent ideal $I' \sim I$ sampled uniformly of bounded norm, we sample an element in I uniformly at random with bounded norm. Consider the isometry $\iota : \mathcal{B}_{p,\infty} \hookrightarrow \mathbb{R}^4$ from Eq. (1.2) and the lattice $\iota(I) \subset \mathbb{R}^4$. By Eq. (2.9), we have $\operatorname{Covol}(\iota(I)) = \operatorname{nrd}(I)^2 p/4$, so we can find a Minkowski reduced basis of I and apply Algorithm 3.6 to $\iota(I)$ with:

$$\rho^2 \geq 3^{13/6} \sqrt{p} \operatorname{nrd}(I)/\pi \tag{3.3}$$

to sample $\alpha \in I \setminus \{0\}$ such that $\operatorname{nrd}(\alpha) \leq \rho^2$ uniformly at random. In particular, we can sample ideals $I' \sim I$ of norm $O(\sqrt{p})$ uniformly at random, as desired.

3.3 Main phases of the SQIsignHD identification protocol

In this section we describe each phase of both versions of SQIsignHD: F-SQIsignHD optimised for efficiency, and R-SQIsignHD, theoretical and optimised for the security proof. We are given the following public parameters:

- A prime of the form $p = c2^e - 1$ with $c \in \mathbb{N}^*$ odd and small, and $e \simeq 2\lambda$ to grant a classical security level of λ bits.
- The supersingular elliptic curve E_0 of equation $y^2 = x^3 + x$ defined over $\mathbb{F}_{p^2}$ with an explicit isomorphism $\varepsilon_0 : \mathcal{O}_0 \xrightarrow{\sim} \mathrm{End}(E_0)$ between a maximal order of $\mathcal{B}_{p,\infty}$ and the endomorphism ring $\mathrm{End}(E_0)$ (given by Lemma 1.2.25).
- A basis (P_0, Q_0) of $E_0[2^e]$.
- A basis $(\beta_1, \cdots, \beta_4)$ of $\mathcal{O}_0$ (*e.g.* $(1, i, (i+j)/2, (1+ij)/2)$) and $(\varepsilon_0(\beta_i)(P_0), \varepsilon_0(\beta_i)(Q_0))_{1 \leq i \leq 4}$, the image of (P_0, Q_0).
- Some precomputed data to apply Algorithm 2.7 (as described in the last paragraph of Section 2.4.3).

We denote by pp these public parameters along with others that will be introduced in the following.

3.3.1 Key generation

In both F-SQIsignHD and R-SQIsignHD, to generate an asymmetric key pair, we fix a big enough prime N_{sk}. The prover samples a left $\mathcal{O}_0$-ideal I_{sk} of norm N_{sk} uniformly at random using Algorithm 3.4. Then, they apply Algorithm 2.7 to compute the image $(E_{\mathsf{pk}}, \varphi_{\mathsf{sk}}(P_0), \varphi_{\mathsf{sk}}(Q_0))$ of the isogeny $\varphi_{\mathsf{sk}} : E_0 \longrightarrow E_{\mathsf{pk}}$ associated to I_{sk}. The curve E_{pk} is published as the prover's public key and $\mathsf{sk} := (I_{\mathsf{sk}}, \varphi_{\mathsf{sk}}(P_0), \varphi_{\mathsf{sk}}(Q_0))$ is safely stored as their secret key. Algorithm 3.7 summarises the this key generation procedure in F-SQIsignHD.

Algorithm 3.7: Key Generation in F-SQIsignHD.

Result: A public key $\mathsf{pk} = E_{\mathsf{pk}}$ and the associated secret key $\mathsf{sk} = (I_{\mathsf{sk}}, \varphi_{\mathsf{sk}}(P_0), \varphi_{\mathsf{sk}}(Q_0))$.

1 Extract N_{sk} from the public parameters pp;
2 Call Algorithm 3.4 to sample an ideal I_{sk} of norm N_{sk} uniformly at random;
3 Call Algorithm 2.7 on I_{sk} and the public parameters pp to obtain the image $(E_{\mathsf{pk}}, \varphi_{\mathsf{sk}}(P_0), \varphi_{\mathsf{sk}}(Q_0))$ of the isogeny $\varphi_{\mathsf{sk}} : E_0 \longrightarrow E_{\mathsf{pk}}$ associated to I_{sk};
4 $\mathsf{pk} \leftarrow E_{\mathsf{pk}}$, $\mathsf{sk} \leftarrow (I_{\mathsf{sk}}, \varphi_{\mathsf{sk}}(P_0), \varphi_{\mathsf{sk}}(Q_0))$;
5 **return** pk, sk;

Let us discuss the choice of N_{sk}. Heuristically, with $N_{\mathsf{sk}} = \Theta(2^\lambda) = \Theta(\sqrt{p})$, the key space would be large enough. However, it would be safer to choose N_{sk} large enough to ensure that E_{pk} has a uniformly random distribution among supersingular elliptic curves defined over $\mathbb{F}_{p^2}$. Indeed, in that case a key recovery attack would require to break an average instance of the isogeny problem and the best known algorithms to solve this problem cost $\tilde{O}(\sqrt{p}) = \tilde{O}(2^\lambda)$, ensuring λ bits of security. By the following, choosing $N_{\mathsf{sk}} = \Theta(2^{4\lambda}) = \Theta(p^2)$ guarantees that E_{pk} has the desired uniformly random distribution.

Proposition 3.3.1. *Let $\varphi : E_0 \longrightarrow E$ be an isogeny of prime degree N sampled uniformly at random among the $N+1$ isogenies of degree N with domain E_0. Then, the distribution of E has statistical distance $O(\sqrt{p/N} + 1/p)$ to the uniform distribution in the supersingular isogeny graph.*

Proof. Let $\mathrm{SS}(p)$ be the set of supersingular elliptic curves over $\mathbb{F}_{p^2}$ (up to $\overline{\mathbb{F}}_p$-isomorphism) and S be the probability distribution on $\mathrm{SS}(p)$ given by $S(E) := K^{-1}/\#\operatorname{Aut}(E)$ for all $E \in \mathrm{SS}(p)$, with $K := \sum_{E \in \mathrm{SS}(p)} 1/\#\operatorname{Aut}(E)$. Let δ_0 be the Dirac distribution on E_0 and π the distribution obtained

from δ_0 after a single step N-isogeny walk from E_0. By [BCC+23, Theorem 11], the statistical distance between S and π satisfies

$$d_{TV}(S,\pi) := \frac{1}{2}\sum_{E\in \mathrm{SS}(p)} |S(E)-\pi(E)| \leq \frac{\sqrt{6KN}}{(N+1)}.$$

By the Eichler's mass formula [Voi21, p. 42.3.8], we know that $K = (p-1)/24$. Then, we get that $d_{TV}(S,\pi) = O(\sqrt{p/N})$.

Now, let U be the uniform distribution on $\mathrm{SS}(p)$. By [Sil09, Theorem III.10.1], we have $\#\operatorname{Aut}(E) = 2$ for all $E \in \mathrm{SS}(p)$ such that $j(E) \neq 0, 1728$, and $\#\operatorname{Aut}(E) \in \{4,6\}$ otherwise. By Theorem 1.1.20, there exists $C_p \in \mathbb{Z}$ small such that $\#\mathrm{SS}(p) = 2K + C_p = (p-1)/12 + C_p$. Hence, we have

$$\begin{aligned} d_{TV}(U,S) &= \frac{1}{2}\sum_{E\in \mathrm{SS}(p)} \left|\frac{1}{\#\mathrm{SS}(p)} - \frac{1}{K\#\operatorname{Aut}(E)}\right| = \frac{1}{2}\sum_{\substack{E\in \mathrm{SS}(p)\\ j(E)\neq 0,1728}} \left|\frac{1}{2K+C_p} - \frac{1}{2K}\right| + O(p^{-1}) \\ &= \frac{1}{2}\frac{C_p}{2K(2K+C_p)}(\#SS(p) + O(1)) + O(p^{-1}) = \frac{C_p}{4K} + O(p^{-1}) = O(p^{-1}). \end{aligned} \tag{3.4}$$

We finally get, by triangular inequality, that:

$$d_{TV}(U,\pi) \leq d_{TV}(U,S) + d_{TV}(S,\pi) = O(\sqrt{p/N} + 1/p).$$

□

Remark 3.3.2. Alternatively, we could fix N_{sk} as an odd prime power instead of a prime. Algorithm 3.4 would still apply and the security analysis would be comparable.

More image points required in R-SQIsignHD

The procedure is almost exactly the same in R-SQIsignHD, but some additional data is required to simplify the response phase. The image of φ_{sk} on points of odd orders of E_0 is also computed and joined to the secret key. We refer to Section 3.3.3 for more information on the points that needs to be evaluated.

Remark 3.3.3. Since N_{sk} is a big prime, we expect the torsion points to be evaluated in SQIsignHD to have order coprime with N_{sk} so Algorithm 3.3 applies to compute the kernel ideals of isogenies of smooth degrees with domain E_{pk}.

Note that Algorithm 2.7 only evaluates φ_{sk} on points of $E_0[2^e]$. Any point $P \in E_0[N]$ could be evaluated as in Algorithm 2.7 as long as N is coprime with all the parameters $(n_t)_{0\leq t\leq n}$ given on entry as precomputed data. Indeed, an inversion of some $n_t \mod N$ is needed to evaluate P. The integer ud_1 from Algorithm 2.7 also needs to be inverted modulo N. However, since $\gcd(ud_1, vd_2) = 1$, we can invert $vd_2 \mod N$ instead of $ud_1 \mod N$ if needed (and evaluate Φ on the second component instead of the first) to evaluate $\varphi_{\mathsf{sk}}(P)$.

3.3.2 Commitment

The key generation efficiency is less crucial than the commitment efficiency which is executed during the signature process. For that reason, we propose two distinct methods for F-SQIsignHD and R-SQIsignHD.

Commitment in F-SQIsignHD

In the commitment phase, the prover wants to generate an isogeny $\varphi_{\mathsf{com}} : E_0 \longrightarrow E_{\mathsf{com}}$ of big odd degree. To do that, they sample an odd integer $N_{\mathsf{com}} < 2^e$ at random such that $N_{\mathsf{com}}(2^e - N_{\mathsf{com}}) = \Omega(p\log(p))$ and apply Algorithm 2.3 to obtain the codomain E_{com}, the 2^e-torsion image basis $(\varphi_{\mathsf{com}}(P_0), \varphi_{\mathsf{com}}(Q_0))$ and the associated left $\mathcal{O}_0$-ideal I_{com} of an isogeny $\varphi_{\mathsf{com}} : E_0 \longrightarrow E_{\mathsf{com}}$ of degree N_{com}. The commitment E_{com} is published while the data $(I_{\mathsf{com}}, \varphi_{\mathsf{com}}(P_0), \varphi_{\mathsf{com}}(Q_0))$ is kept secret. Algorithm 3.8 follows.

This method is very efficient and yields very good signing times, as we shall see in Section 3.5.2. However, this choice is made at the expense of a heuristic security proof. Indeed, to prove the zero knowledge property, the commitment distribution E_{com} is expected to be uniform among supersingular elliptic curves. In theory, this method gives enough entropy: we have $\Theta(p)$ choices of $N_{\mathsf{com}} < 2^e$ and $\Theta(N_{\mathsf{com}})$ choices of isogenies φ_{com} for each choice of N_{com}, yielding $\Theta(p^2) = \Theta(2^{4\lambda})$ choices of isogenies φ_{com} in total. We would need bigger choices of N_{com} and more randomness in the generation of φ_{com} to make a perfectly rigorous argument but we can very well reasonably assume that E_{com} is computationally indistinguishable from a uniformly random supersingular elliptic curve.

Commitment in R-SQIsignHD

In R-SQIsignHD, we avoid the above heuristic assumption by proceeding as in the key generation, at the expense of efficiency. We fix a prime $N_{\mathsf{com}} = \Theta(2^{4\lambda}) = \Theta(p^2)$ ($N_{\mathsf{com}} = N_{\mathsf{sk}}$ would perfectly be reasonable). The prover samples a left $\mathcal{O}_0$-ideal I_{com} of norm N_{com} uniformly at random using Algorithm 3.4. Then, they apply Algorithm 2.7 to compute the image $(E_{\mathsf{com}}, \varphi_{\mathsf{com}}(P_0), \varphi_{\mathsf{com}}(Q_0))$ of the isogeny $\varphi_{\mathsf{com}} : E_0 \longrightarrow E_{\mathsf{com}}$ associated to I_{com}. As in F-SQIsignHD, the commitment E_{com} is published while the data $(I_{\mathsf{com}}, \varphi_{\mathsf{com}}(P_0), \varphi_{\mathsf{com}}(Q_0))$ is kept secret. Algorithm 3.9 follows.

Algorithm 3.8: Fast commitment (F-SQIsignHD).

Result: A commitment curve $\mathsf{com} = E_{\mathsf{com}}$ and the associated secret data $\mathsf{sc} = (I_{\mathsf{com}}, \varphi_{\mathsf{com}}(P_0), \varphi_{\mathsf{com}}(Q_0))$.

repeat
  Sample a random odd integer $N_{\mathsf{com}} \in [\![1\ ;\ 2^e - 1]\!]$;
until $N_{com}(2^e - N_{com}) = \Omega(p \log(p))$;
Call Algorithm 2.3 to compute the images $(E_{\mathsf{com}}, \varphi_{\mathsf{com}}(P_0), \varphi_{\mathsf{com}}(Q_0))$ and the associated ideal I_{com} of an N_{com}-isogeny $\varphi_{\mathsf{com}} : E_0 \longrightarrow E_{\mathsf{com}}$;
$\mathsf{com} \leftarrow E_{\mathsf{com}}$, $\mathsf{sc} \leftarrow (I_{\mathsf{com}}, \varphi_{\mathsf{com}}(P_0), \varphi_{\mathsf{com}}(Q_0))$;
return com, sc;

Algorithm 3.9: Rigorous commitment (R-SQIsignHD).

Result: A commitment curve $\mathsf{com} = E_{\mathsf{com}}$ and the associated secret data $\mathsf{sc} = (I_{\mathsf{com}}, \varphi_{\mathsf{com}}(P_0), \varphi_{\mathsf{com}}(Q_0))$.

Extract N_{com} from public parameters pp;
Call Algorithm 3.4 to sample an ideal I_{com} of norm N_{com} uniformly at random;
Call Algorithm 2.7 on I_{com} and the public parameters pp to obtain the image $(E_{\mathsf{com}}, \varphi_{\mathsf{com}}(P_0), \varphi_{\mathsf{com}}(Q_0))$ of the isogeny $\varphi_{\mathsf{com}} : E_0 \longrightarrow E_{\mathsf{com}}$ associated to I_{com};
$\mathsf{com} \leftarrow E_{\mathsf{com}}$, $\mathsf{sc} \leftarrow (I_{\mathsf{com}}, \varphi_{\mathsf{com}}(P_0), \varphi_{\mathsf{com}}(Q_0))$;
return com, sc;

3.3.3 Challenge

Challenge in F-SQIsignHD

If λ is the security level, then the challenge is an integer $\mathsf{chl} \in [\![0\ ;\ 2^\lambda - 1]\!]$. This way, we obtain a challenge space of size 2^λ. This integer describes the kernel of the challenge isogeny $\varphi_{\mathsf{chl}} : E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}$ as follows. One can generate a deterministic basis $(P_{\mathsf{pk}}, Q_{\mathsf{pk}})$ of $E_{\mathsf{pk}}[2^e]$ and compute the point $K_{\mathsf{chl}} := [2^{e-\lambda}]P_{\mathsf{pk}} + [\mathsf{chl} \cdot 2^{e-\lambda}]Q_{\mathsf{chl}}$ of order 2^λ that generates the kernel $\ker(\varphi_{\mathsf{chl}})$.

Challenge in R-SQIsignHD

In R-SQIsignHD, the challenge is still an integer $\mathsf{chl} \in [\![0\ ;\ 2^\lambda - 1]\!]$ representing an isogeny $\varphi_{\mathsf{chl}} : E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}$. However, φ_{chl} is assumed to have fixed odd degree N_{chl} in order to simplify the response

process. For security reasons (see Theorem 3.4.10 and Remark 3.4.11), we require $N_{\mathsf{chl}} = \Theta(p^{5/2})$ and for efficiency reasons, we require N_{chl} to be powersmooth.

Let us write the prime decomposition $N_{\mathsf{chl}} = \prod_{i=1}^{n} \ell_i^{e_i}$ with the prime powers $\ell_i^{e_i}$ small enough and distinct from the prime factors of the parameters $(n_t)_{0\leq t\leq n}$ given on entry of Algorithm 2.7 for technical reasons explained in Remark 3.3.3. Then, for all $i \in [\![1\ ;\ n]\!]$, $E_{\mathsf{pk}}[\ell_i^{e_i}]$ is defined over a small extension of $\mathbb{F}_{p^2}$, so we may efficiently generate a basis (P_i, Q_i) of $E_{\mathsf{pk}}[\ell_i^{e_i}]$ in a deterministic way. Then the challenge isogeny has kernel $\ker(\varphi_{\mathsf{chl}}) = \langle P + [\mathsf{chl}]Q\rangle$ with $P := \sum_{i=1}^{n} P_i$ and $Q := \sum_{i=1}^{n} Q_i$. It can be computed as a chain $\varphi_{\mathsf{chl}} = \varphi_n \circ \cdots \circ \varphi_1$, where for all $i \in [\![1\ ;\ n]\!]$, φ_i is an $\ell_i^{e_i}$-isogeny of kernel $\langle \varphi_{i-1} \circ \cdots \circ \varphi_1(P_i + [\mathsf{chl}]Q_i)\rangle$. This procedure is summarised in Algorithm 3.10 which also returns the associated ideal I_{chl} used by the prover in the response generation. Following standard ideas from [EHLMP18; DFKLPW20], the ideal I_{chl} is computed piecewise, as the isogeny chain φ_{chl}, applying Algorithm 3.3 to compute ideal pieces at every step.

As explained in Section 3.3.1, note that some additional data is joined in the secret key in order to compute I_{chl}. This additional data is the image $(\varphi_{\mathsf{sk}}(P_{0,i}), \varphi_{\mathsf{sk}}(Q_{0,i}))$ of deterministic basis $(P_{0,i}, Q_{0,i})$ of $E_0[\ell_i^{e_i}]$ for all $i \in [\![1\ ;\ n]\!]$, that we denote by $\varphi_{\mathsf{sk}|N_{\mathsf{chl}}} := (\varphi_{\mathsf{sk}}(P_{0,i}), \varphi_{\mathsf{sk}}(Q_{0,i}))_{1\leq i\leq n}$.

Algorithm 3.10: Challenge isogeny generation in R-SQIsignHD (during the response phase).

Data: The public parameters pp (including $N_{\mathsf{chl}} = \prod_{i=1}^{n} \ell_i^{e_i}$), the challenge $\mathsf{chl} \in [\![0\ ;\ 2^\lambda - 1]\!]$, the public key E_{pk} and the secret key sk.
Result: The corresponding isogeny $\varphi_{\mathsf{chl}} : E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}$ and its kernel ideal I_{chl}.

Extract $N_{\mathsf{chl}} = \prod_{i=1}^{n} \ell_i^{e_i}$ from pp;
Parse $I_{\mathsf{sk}}, \varphi_{\mathsf{sk}}(P_0), \varphi_{\mathsf{sk}}(Q_0), \varphi_{\mathsf{sk}|N_{\mathsf{chl}}} \leftarrow \mathsf{sk}$;
for $i = 1$ **to** n **do**
  Generate a deterministic basis (P_i, Q_i) of $E_{\mathsf{pk}}[\ell_i^{e_i}]$;
  $K_i \leftarrow P_i + [\mathsf{chl}]Q_i$;
  if $i \geq 2$ **then**
    $K_i \leftarrow \varphi_{i-1} \circ \cdots \circ \varphi_1(K_i)$;
  end
  Compute $\varphi_i : E_{i-1} \longrightarrow E_i$ of kernel $\langle K_i\rangle$;
  Generate a deterministic basis $(R_{0,i}, S_{0,i})$ of $E_0[\ell_i^{e_i}]$;
  $(U_i, V_i) \leftarrow (\varphi_{i-1} \circ \cdots \circ \varphi_1 \circ \varphi_{\mathsf{sk}}(R_{0,i}), \varphi_{i-1} \circ \cdots \circ \varphi_1 \circ \varphi_{\mathsf{sk}}(S_{0,i}))$;
  Extract from pp a basis $(\beta_1, \cdots, \beta_4)$ of $\mathcal{O}_0$ and compute $(\varepsilon_0(\beta_j)(R_{0,i}), \varepsilon_0(\beta_j)(S_{0,i}))_{1\leq j\leq 4}$;
  Call Algorithm 3.3 on K_i, $(\beta_1, \cdots, \beta_4)$, $(R_{0,i}, S_{0,i})$, $(\varepsilon_0(\beta_j)(R_{0,i}), \varepsilon_0(\beta_j)(S_{0,i}))_{1\leq j\leq 4}$ and (U_i, V_i) to obtain the kernel ideal I_i of φ_i;
end
return $\varphi_n \circ \cdots \circ \varphi_1, I_1 \cdots I_n$;

By the following lemma, the powersmooth integer $N_{\mathsf{chl}} = \Theta(p^{5/2})$ can be chosen so that Algorithm 3.10 terminates in polynomial time (in $\log(p)$).

Lemma 3.3.4. *The integer $N_{\mathsf{chl}} = \Theta(p^{5/2})$ can be chosen to be C-smooth with $C = O(\log(p))$. In that case, Algorithm 3.10 terminates in polynomial time in $\log(p)$.*

Proof. We consider

$$N = \prod_{\substack{\ell\leq C\\ \ell\notin S}} \ell,$$

where S is a finite set of inadmissible primes containing 2 (and the orders of points that cannot be evaluated by φ_{sk} for the technical reasons explained in Remark 3.3.3). By [HW75, Theorems 413 and 434], $\log(N_{\mathsf{chl}}) \sim C$ as $C \longrightarrow +\infty$. Hence, in order to achieve $N = \Theta(p^{5/2})$, taking $C = O(\log(p))$ is sufficient. This proves that C-powersmooth integers of size $\Theta(p^{5/2})$ with $C = O(\log(p))$ and prime factors not in the set S do exist. We select N_{chl} such an integer and write $N_{\mathsf{chl}} = \prod_{i=1}^{n} \ell_i^{e_i}$ its prime decomposition.

Let $i \in [\![1\ ;\ n]\!]$ and $E/\mathbb{F}_{p^2}$ be a supersingluar elliptic curve. Then $E[\ell_i^{e_i}]$ is defined over an extension $\mathbb{F}_{p^{2\delta_i}}/\mathbb{F}_{p^2}$ of degree $\delta_i = 6\ell_i^{e_i-1}(\ell_i - 1) = O(\log(p))$ by [ACD+24, Lemma 2.12]. By the most recent

improvements of Vélu's formulas [BDFLS20; Vé71], computing the isogeny $\varphi_i : E_{i-1} \longrightarrow E_i$ costs $O(\ell_i^{e_i/2})$ arithmetic operations over $\mathbb{F}_{p^{2\delta_i}}$ which takes polynomial time in $\log(p)$. Once φ_i has been expressed as rational fractions with these formulas, evaluating $\varphi_i(P)$ on $P \in E_{i-1}[\ell_j^{e_j}]$ for $j \geq i+1$ costs $O(\ell_i^{e_i})$ operations over $\mathbb{F}_{p^{2\delta_i\delta_j}}$ and the result is in $\varphi_i(P) \in E_i[\ell_j^{e_j}] \subseteq E_i(\mathbb{F}_{p^{2\delta_j}})$. Since $n = O(\log(p))$, it follows that point evaluations via the chain $\varphi_{i-1} \circ \cdots \circ \varphi_1$ in Algorithm 3.10 always take polynomial time in $\log(p)$, so the whole chain computation $\varphi_n \circ \cdots \circ \varphi_1$ runs in polynomial time in $\log(p)$. Besides, the cost of each call to Algorithm 3.3 is dominated by isogeny evaluations and discrete logarithms computations in subgroups of $\ell_i^{e_i}$-torsion. These operations also take polynomial time in $\log(p)$. It follows that the whole algorithm runs in polynomial time in $\log(p)$. □

3.3.4 Response

Given the secret key $\varphi_{\mathsf{sk}} : E_0 \longrightarrow E_{\mathsf{pk}}$, commitment $\varphi_{\mathsf{com}} : E_0 \longrightarrow E_{\mathsf{com}}$ and challenge $\varphi_{\mathsf{chl}} : E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}$ isogenies (*i.e.* data representing them efficiently), the prover has to return an isogeny $\varphi_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$ (*i.e.* some efficient representation of φ_{rsp}).

In both versions of SQIsignHD, the prover starts by computing the kernel ideal I_{chl} of φ_{chl} using the secret key data sk (and either Algorithm 3.3 in F-SQIsignHD or Algorithm 3.10 in R-SQIsignHD). With the secret data sk and sc obtained during the key generation and commitment phases, the prover also knows the ideals I_{sk} and I_{com} associated to the isogenies φ_{sk} and φ_{com} respectively so they can compute $I := \overline{I}_{\mathsf{com}} \cdot I_{\mathsf{sk}} \cdot I_{\mathsf{chl}}$. Then they sample a random ideal $I_{\mathsf{rsp}} \sim I$ using Algorithm 3.6 that can be translated into an isogeny $\varphi_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$. Note that the higher dimensional embedding of φ_{rsp} is not computed during the response but during the verification phase. The prover evaluates φ_{rsp} on torsion points instead using the techniques of Section 2.3.3. The main difference between F-SQIsignHD and R-SQIsignHD lie in the constraint imposed on the norm of I_{rsp} (smaller and with additional constraints in F-SQIsignHD than in R-SQIsignHD). This makes the description of the R-SQIsignHD response a bit more technical to handle cases when I_{rsp} has even norm.

Response in F-SQIsignHD

In F-SQIsignHD, we want make it possible for the verifier to embed φ_{rsp} in a 4-dimensional isogeny as in Section 2.3. For that reason, we fix $e/2 < f \leq e-2$ and a smoothness bound $B \in \mathbb{N}^*$ and require $q := \operatorname{nrd}(I_{\mathsf{rsp}})$ to be (f, B)-good in the sense of Definition 2.3.1. Indeed, if q is (f, B)-good then it is odd and it is easy to decompose $2^e - q$ as a sum of two squares and φ_{rsp} can be embedded in a 4-dimensional isogeny as defined in Eq. (2.10).

To sample such an ideal I_{rsp} of (f, B)-good norm q uniformly at random, the prover applies Algorithm 3.6 repeatedly to the lattice $I := \overline{I}_{\mathsf{com}} \cdot I_{\mathsf{sk}} \cdot I_{\mathsf{chl}}$ with a norm bound $\rho^2 = 2^f \operatorname{nrd}(I)$. Algorithm 3.6 will return an element $\alpha \in I \setminus \{0\}$ sampled uniformly at random among elements of norm $\leq 2^f \operatorname{nrd}(I)$. We stop the sampling once $q := q_I(\alpha) = \operatorname{nrd}(\alpha)/\operatorname{nrd}(I)$ is (f, B)-good and set $I_{\mathsf{rsp}} := I\overline{\alpha}/\operatorname{nrd}(I)$.

For this to succeed, we not only need f to satisfy Eq. (3.3) $2^f \geq 3^{13/6}\sqrt{p}/\pi$ which ensures that Algorithm 3.6 will terminate. We also need some margin to ensure that an element of (f, B)-good can be found. Heuristically, the probability for an integer $\simeq \sqrt{p}$ to be (f, B)-good is $\Omega(1/\log(p))$ so $2^f = \Omega(\sqrt{p}\log(p))$ should be sufficient. However, we have no formal proof that an element of (f, B)-good norm bounded by $O(\sqrt{p}\log(p))$ can be found.

Once I_{rsp} has been generated, we generate a deterministic basis $(P_{\mathsf{com}}, Q_{\mathsf{com}})$ of $E_{\mathsf{com}}[2^e]$ and evaluate $([2^\lambda]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}), [2^\lambda]\varphi_{\mathsf{rsp}}(Q_{\mathsf{com}}))$ with the techniques from Section 2.3.3 that we recall below. Recall that we have $I_{\mathsf{rsp}} = I\overline{\alpha}/\operatorname{nrd}(I)$ with $I = \overline{I}_{\mathsf{com}} \cdot I_{\mathsf{sk}} \cdot I_{\mathsf{chl}}$ and $\alpha \in I$, so that $I_{\mathsf{com}} \cdot I_{\mathsf{rsp}} = I_{\mathsf{sk}} \cdot I_{\mathsf{chl}} \cdot \overline{\alpha}/\operatorname{nrd}(I_{\mathsf{sk}} \cdot I_{\mathsf{chl}})$. By Lemma 1.2.24, it follows that:

$$\widehat{\varphi}_{\mathsf{com}} \circ \widehat{\varphi}_{\mathsf{rsp}} \circ \varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}} = \varepsilon_0(\alpha) \quad i.e. \quad [2^\lambda N_{\mathsf{com}} N_{\mathsf{sk}}]\varphi_{\mathsf{rsp}} = \varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}} \circ \varepsilon_0(\overline{\alpha}) \circ \widehat{\varphi}_{\mathsf{com}}, \tag{3.5}$$

where ε_0 is an isomorphism $\mathcal{O}_0 \xrightarrow{\sim} \operatorname{End}(E_0)$. Since $N_{\mathsf{com}}N_{\mathsf{sk}}$ is odd, we can find its inverse μ modulo $2^{e-\lambda}$ and for $P = P_{\mathsf{com}}, Q_{\mathsf{com}}$, we can compute:

$$\varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}} \circ \varepsilon_0(\overline{\alpha}) \circ \widehat{\varphi}_{\mathsf{com}}([\mu]P) = [2^\lambda N_{\mathsf{com}} N_{\mathsf{sk}} \mu]\varphi_{\mathsf{rsp}}(P) = [2^\lambda]\varphi_{\mathsf{rsp}}(P).$$

Note that the image of φ_{rsp} on the $2^{e-\lambda}$-torsion of E_{com} is sufficient to compute the 4-dimensional embedding in the verification phase as long as $e - \lambda \geq f/2 + 2$ (by Remark 2.3.4). This inequality

is largely satisfied with a margin close to $\lambda/2$ since $f \simeq \lambda \simeq e/2$. For compactness reasons (see Section 3.5.1) we use a tighter amount of torsion. Let $r := \lceil f/2 \rceil + 2$. Then we multiply $[2^\lambda]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}})$ and $[2^\lambda]\varphi_{\mathsf{rsp}}(Q_{\mathsf{com}})$ by $2^{e-r-\lambda}$ to return $([2^{e-r}]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}), [2^{e-r}]\varphi_{\mathsf{rsp}}(Q_{\mathsf{com}}), q)$ as a response. Algorithm 3.11 summarizes the whole F-SQIsignHD response procedure.

Algorithm 3.11: Fast response (F-SQIsignHD).

Data: The public parameters pp, the public key E_{pk}, the secret key sk, the commitment E_{com} and the secret commitment data sc and the challenge $\mathsf{chl} \in [\![0\ ;\ 2^\lambda - 1]\!]$.

Result: The response $\mathsf{rsp} = ([2^{e-r}]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}), [2^{e-r}]\varphi_{\mathsf{rsp}}(Q_{\mathsf{com}}), q)$, where $(P_{\mathsf{com}}, Q_{\mathsf{com}})$ is a deterministic basis of $E_{\mathsf{com}}[2^e]$ and $\varphi_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$ a q-isogeny.

Extract $e, \lambda, f, B, r, N_{\mathsf{com}}, N_{\mathsf{sk}}$ from pp;
Parse $I_{\mathsf{sk}}, \varphi_{\mathsf{sk}}(P_0), \varphi_{\mathsf{sk}}(Q_0) \leftarrow \mathsf{sk}$;
Parse $I_{\mathsf{com}}, \varphi_{\mathsf{com}}(P_0), \varphi_{\mathsf{com}}(Q_0) \leftarrow \mathsf{sc}$;
Compute a deterministic basis $(P_{\mathsf{pk}}, Q_{\mathsf{pk}})$ of $E_{\mathsf{pk}}[2^e]$;
$K_{\mathsf{chl}} \leftarrow [2^{e-\lambda}]P_{\mathsf{pk}} + [2^{e-\lambda}\mathsf{chl}]Q_{\mathsf{pk}}$;
Compute $\varphi_{\mathsf{chl}} : E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}$ of kernel K_{chl};
Extract from pp a basis $(\beta_1, \cdots, \beta_4)$ of $\mathcal{O}_0$ and its image $(\varepsilon_0(\beta_j)(P_0), \varepsilon_0(\beta_j)(Q_0))_{1 \le j \le 4}$;
Call Algorithm 3.3 on K_{chl}, $(\beta_1, \cdots, \beta_4)$, (P_0, Q_0) and $(\varepsilon_0(\beta_j)(P_0), \varepsilon_0(\beta_j)(Q_0))_{1 \le j \le 4}$ and $(\varphi_{\mathsf{sk}}(P_0), \varphi_{\mathsf{sk}}(Q_0))$ to compute the kernel ideal I_{chl} of φ_{chl};
$I \leftarrow \overline{I}_{\mathsf{com}} \cdot I_{\mathsf{sk}} \cdot I_{\mathsf{chl}}$;
repeat
  Call Algorithm 3.6 on I to sample $\alpha \in I$ of norm $\le 2^f \operatorname{nrd}(I)$ uniformly at random;
  $q \leftarrow \operatorname{nrd}(\alpha)/\operatorname{nrd}(I)$;
until q *is* (f, B)*-good*;
$\mu \leftarrow 1/(N_{\mathsf{com}} N_{\mathsf{sk}}) \mod 2^{e-\lambda}$;
Compute a deterministic basis $(P_{\mathsf{com}}, Q_{\mathsf{com}})$ of $E_{\mathsf{com}}[2^e]$;
Call Algorithm 2.9 on (P_0, Q_0), $(\varphi_{\mathsf{com}}(P_0), \varphi_{\mathsf{com}}(Q_0))$ and $(P_{\mathsf{com}}, Q_{\mathsf{com}})$ to compute $(\widehat{\varphi}_{\mathsf{com}}(P_{\mathsf{com}}), \widehat{\varphi}_{\mathsf{com}}(Q_{\mathsf{com}}))$;
Find $a, b, c, d \in \mathbb{Z}/2^e\mathbb{Z}$ such that $\varepsilon_0(\overline{\alpha}) \circ \widehat{\varphi}_{\mathsf{com}}(P_{\mathsf{com}}) = [a]P_0 + [b]Q_0$ and $\varepsilon_0(\overline{\alpha}) \circ \widehat{\varphi}_{\mathsf{com}}(Q_{\mathsf{com}}) = [c]P_0 + [d]Q_0$;
$[2^{e-r}]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}) \leftarrow [2^{e-r-\lambda}a\mu]\varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}}(P_0) + [2^{e-r-\lambda}b\mu]\varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}}(Q_0)$;
$[2^{e-r}]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}) \leftarrow [2^{e-r-\lambda}c\mu]\varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}}(P_0) + [2^{e-r-\lambda}d\mu]\varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}}(Q_0)$;
return $([2^{e-r}]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}), [2^{e-r}]\varphi_{\mathsf{rsp}}(Q_{\mathsf{com}}), q)$;

Response in R-SQIsignHD

In R-SQIsignHD, we embed φ_{rsp} in an 8-dimensional isogeny so we can relax the constraint on the ideal norm $q = \operatorname{nrd}(I_{\mathsf{rsp}})$. In addition, for security reasons that will be explained in Section 3.4.2, we require $q = \Theta(p^2)$ instead of $q = \Theta(\sqrt{p}\log(p))$. We select $f := 2(e-4)$, so that $e = f/2 + 2$ and a 2^f-isogeny of dimension 8 embedding φ_{rsp} can be inferred from the image of φ_{rsp} on a basis of $E_{\mathsf{com}}[2^e]$ by dividing the computation in two pieces as explained in Remark 2.3.4.

As in F-SQIsignHD, we start by computing the kernel ideal I_{chl} of the challenge. Since $\deg(\varphi_{\mathsf{chl}}) = N_{\mathsf{chl}} = \Theta(p^{5/2})$, the kernel of φ_{chl} is defined over a field extension of $\mathbb{F}_{p^2}$ of exponentially big degree so we cannot compute I_{chl} from a direct application of Algorithm 3.3. Instead, we use Algorithm 3.10 that computes the ideal I_{chl} piecewise while computing the challenge isogeny φ_{chl} as a chain of isogenies at the same time.

Then, in order to compute I_{rsp} equivalent to $I := \overline{I}_{\mathsf{com}} I_{\mathsf{sk}} I_{\mathsf{chl}}$, we sample $\alpha \in I \setminus \{0\}$ uniformly at random among elements of norm $\le 2^f \operatorname{nrd}(I)$ with Algorithm 3.6 and set $I_{\mathsf{rsp}} := I\overline{\alpha}/\operatorname{nrd}(I)$. The following steps are a bit more technical than in F-SQIsignHD in order to account for the odd norm factor of I_{rsp}. With a simple evaluation of 2^e-torsion points via φ_{rsp} similar to F-SQIsignHD, we would lose information and that would make the verification harder. Instead, we factor $I_{\mathsf{rsp}} = [2^a] I^{(1)}_{\mathsf{rsp}} \cdot I'_{\mathsf{rsp}} \cdot I^{(2)}_{\mathsf{rsp}}$, with $I^{(1)}_{\mathsf{rsp}}$ and $I^{(2)}_{\mathsf{rsp}}$ primitive and of norms $2^{b_1}, 2^{b_2} \le 2^e$ respectively and I'_{rsp} of odd

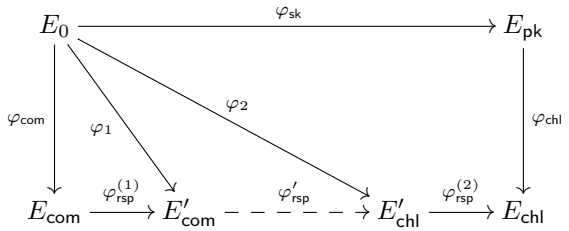

Figure 3.2: Non commutative diagram explaining the R-SQIsignHD response computation.

norm q'. We also consider the associated decomposition (see Fig. 3.2):

$$\varphi_{\mathsf{rsp}} = [2^a]\varphi^{(2)}_{\mathsf{rsp}} \circ \varphi'_{\mathsf{rsp}} \circ \varphi^{(1)}_{\mathsf{rsp}}.$$

To compute a response, we compute $a, \varphi^{(1)}_{\mathsf{rsp}}, \widehat{\varphi}^{(2)}_{\mathsf{rsp}}$ and the image of a deterministic 2^e-torsion basis via φ'_{rsp}. Then, since $\deg(\varphi'_{\mathsf{rsp}}) = q'$ is odd, the verifier will be able to embed φ'_{rsp} into an 8-dimensional 2^f-isogeny. The integer a is simply the greatest integer such that 2^a divides I_{rsp}. We can then compute $J_{\mathsf{rsp}} := I_{\mathsf{rsp}}/2^a = I^{(1)}_{\mathsf{rsp}} \cdot I'_{\mathsf{rsp}} \cdot I^{(2)}_{\mathsf{rsp}}$. The quaternion $\alpha \in I$ is also divisible by 2^a, so we can write $\alpha = 2^a\alpha'$, so that $J_{\mathsf{rsp}} = I\overline{\alpha'}/\operatorname{nrd}(I)$. To compute $I^{(1)}_{\mathsf{rsp}}$ and $I^{(2)}_{\mathsf{rsp}}$, we rely on the following lemma. We compute $\mathcal{O}_{\mathsf{com}} = O_R(I_{\mathsf{com}})$, $\mathcal{O}_{\mathsf{chl}} = O_R(I_{\mathsf{chl}})$ and $b := v_2(\operatorname{nrd}(J_{\mathsf{rsp}}))$. We then compute $b_1 := \min(b, e)$, $b_2 := b - b_1$ and set $I^{(1)}_{\mathsf{rsp}} := J_{\mathsf{rsp}} + 2^{b_1}\mathcal{O}_{\mathsf{com}}$ and $\overline{I}^{(2)}_{\mathsf{rsp}} := \overline{J}_{\mathsf{rsp}} + 2^{b_2}\mathcal{O}_{\mathsf{chl}}$, so that $\operatorname{nrd}(I^{(1)}_{\mathsf{rsp}}) = 2^{b_1}$ and $\operatorname{nrd}(I^{(2)}_{\mathsf{rsp}}) = 2^{b_2}$. Note that $I^{(2)}_{\mathsf{rsp}}$ is trivial when $v_2(\operatorname{nrd}(J_{\mathsf{rsp}})) \leq e$. The ideal $I^{(2)}_{\mathsf{rsp}}$ is only used when $v_2(\operatorname{nrd}(J_{\mathsf{rsp}})) > e$ in order to ensure that there is always enough accessible torsion to translate the even part of J_{rsp} into isogenies.

Lemma 3.3.5. *Let I be a left ideal of a maximal order $\mathcal{O} \subset \mathcal{B}_{p,\infty}$ of norm coprime with p. Let us write $I := m \cdot J$ with $m \in \mathbb{N}^*$ and $J \subseteq \mathcal{O}$ primitive. Let $d \in \mathbb{N}^*$ prime to m and $K := I + d\mathcal{O}$. Then* $\operatorname{nrd}(K) = \gcd(d, \operatorname{nrd}(J))$.

Proof. Let $E/\mathbb{F}_{p^2}$ be a supersingular elliptic curve of endomorphism ring isomorphic to $\mathcal{O}$. Then:

$$\begin{aligned} E[K] &= E[mJ + d\mathcal{O}] = E[mJ] \cap E[d] = \{P \in E \mid \forall \alpha \in J, [m]\alpha(P) = 0\} \cap E[d] \\ &= \{P \in E \mid [m]P \in E[J]\} \cap E[d] = [m]^{-1}(E[J]) \cap E[d] \end{aligned}$$

Since J is primitive, $E[J]$ is cyclic so we may consider a generator P of $E[J]$. Let $N' := \operatorname{nrd}(J)$ and $d' := \gcd(d, \operatorname{nrd}(J))$. Let $Q_0 := [N'/d']P$. Then, $[d]Q_0 = [d/d'][N']P = 0$ and $[m]Q_0 \in \langle P \rangle$ by construction, so that $Q_0 \in E[K]$. Conversely, let $Q \in E[K]$. Then $[m]Q = [k]P$ for some $k \in \mathbb{Z}$ and $[d]Q = 0$. In particular $[kd]P = [md]Q = 0$. Then, $N'|kd$ since P has order N', so that $N'/d'|k$, so we may write $k = k'N'/d'$ with $k' \in \mathbb{Z}$, so that $[m]Q = [k'N'/d']P = [k']Q_0$. Since m and d are coprime, there exists $u, v \in \mathbb{Z}$ such that $mu + dv = 1$ and we then have $Q = [mu + dv]Q = [um]Q = [uk']Q_0$. Hence, $E[K] = \langle Q_0 \rangle$ and finally

$$\operatorname{nrd}(K) = \#E[K] = \#\langle Q_0 \rangle = \#\langle [N'/d']P \rangle = d' = \gcd(d, \operatorname{nrd}(J)).$$

□

Once the computation of $I^{(1)}_{\mathsf{rsp}}$ and $\overline{I}^{(2)}_{\mathsf{rsp}}$ is done, we translate them into the isogenies $\varphi^{(1)}_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E'_{\mathsf{com}}$ and $\widehat{\varphi}^{(2)}_{\mathsf{rsp}} : E_{\mathsf{chl}} \longrightarrow E'_{\mathsf{chl}}$ respectively. We want to use Algorithm 3.1 to obtain their kernel but this algorithm requires to know the endomorphism ring of their respective domain. Instead, use the public knowledge of $\operatorname{End}(E_0)$ and apply Algorithm 3.1 to $[I_{\mathsf{com}}]^* I^{(1)}_{\mathsf{rsp}}$ and $[I_{\mathsf{sk}} \cdot I_{\mathsf{chl}}]^* \overline{I}^{(2)}_{\mathsf{rsp}}$ respectively, to obtain generators K_1 and K_2 of $\ker(\varphi^*_{\mathsf{com}}\varphi^{(1)}_{\mathsf{rsp}})$ and $\ker((\varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}})^* \widehat{\varphi}^{(2)}_{\mathsf{rsp}})$ respectively. We then have $\ker(\varphi^{(1)}_{\mathsf{rsp}}) = \langle \varphi_{\mathsf{com}}(K_1) \rangle$ and $\ker(\widehat{\varphi}^{(2)}_{\mathsf{rsp}}) = \langle \varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}}(K_2) \rangle$.

We can then finally try to evaluate φ'_{rsp} on $E'_{\mathsf{com}}[2^e]$. However, unlike in F-SQIsignHD, the method from Section 2.3.3 does not directly apply to $\varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}$ and $\widehat{\varphi}^{(2)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}}$ to evaluate φ'_{rsp} on

$E'_{\mathsf{com}}[2^e]$, since $\varphi^{(1)}_{\mathsf{rsp}}$ (and maybe $\varphi^{(2)}_{\mathsf{rsp}}$) have even degree. To circumvent this difficulty, we use isogenies $\varphi_1 : E_0 \longrightarrow E'_{\mathsf{com}}$ and $\varphi_2 : E_0 \longrightarrow E'_{\mathsf{chl}}$ of odd degrees. First, we set $I_1 := I_{\mathsf{com}} \cdot I^{(1)}_{\mathsf{rsp}}$ and $I_2 := I_{\mathsf{sk}} \cdot I_{\mathsf{chl}} \cdot \overline{I}^{(2)}_{\mathsf{rsp}}$, find $\alpha_1 \in I_1$ and $\alpha_2 \in I_2$ such that $\mathrm{nrd}(\alpha_1)/\mathrm{nrd}(I_1)$ and $\mathrm{nrd}(\alpha_2)/\mathrm{nrd}(I_2)$ are odd, so that the equivalent ideals $J_1 := I_1\overline{\alpha_1}/\mathrm{nrd}(I_1)$ and $J_2 := I_2\overline{\alpha_2}/\mathrm{nrd}(I_2)$ have odd norms N_1 and N_2. We define φ_1 and φ_2 to be the isogenies respectively associated to J_1 and J_2. Then, we use Algorithm 2.7 in order to evaluate $(\varphi_1(P_0), \varphi_1(Q_0))$ and $(\varphi_2(P_0), \varphi_2(Q_0))$. Finally, we use the following lemma to evaluate φ'_{rsp} on a deterministic basis $(P'_{\mathsf{com}}, Q'_{\mathsf{com}})$ of $E'_{\mathsf{com}}[2^e]$ as in Section 2.3.3.

Lemma 3.3.6. *We have:*

$$[N_{chl}N_{com}N_{\mathsf{sk}}N_1N_2]\varphi'_{rsp} = \varphi_2 \circ \varepsilon_0(\theta) \circ \widehat{\varphi}_1,$$

with $N_1 := \mathrm{nrd}(J_1)$ *and* $N_2 := \mathrm{nrd}(J_2)$ *and* $\theta := \alpha_2\overline{\alpha}'\overline{\alpha}_1/2^{b_1+b_2} \in \mathcal{O}_0$.

Proof. By Eq. (3.5), we have $[N_{\mathsf{com}}N_{\mathsf{sk}}N_{\mathsf{chl}}]\varphi_{\mathsf{rsp}} = \varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}} \circ \varepsilon_0(\overline{\alpha}) \circ \widehat{\varphi}_{\mathsf{com}}$, so that:

$$[N_{\mathsf{com}}N_{\mathsf{sk}}N_{\mathsf{chl}}2^{b_1+b_2}]\varphi'_{\mathsf{rsp}} = \widehat{\varphi}^{(2)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}} \circ \varepsilon_0(\overline{\alpha'}) \circ \widehat{\varphi}_{\mathsf{com}} \circ \widehat{\varphi}^{(1)}_{\mathsf{rsp}}, \tag{3.6}$$

since $\varphi_{\mathsf{rsp}} = [2^a]\varphi^{(2)}_{\mathsf{rsp}} \circ \varphi'_{\mathsf{rsp}} \circ \varphi^{(1)}_{\mathsf{rsp}}$ and $\alpha = 2^a\alpha'$. By Lemma 1.2.24, and by the definition of φ_1 and φ_2, we have:

$$\widehat{\varphi}_1 \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}} = \varepsilon_0(\alpha_1) \quad \text{and} \quad \widehat{\varphi}_2 \circ \widehat{\varphi}^{(2)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}} = \varepsilon_0(\alpha_2),$$

so that

$$[N_1N_{\mathsf{com}}]\widehat{\varphi}^{(1)}_{\mathsf{rsp}} = \varphi_{\mathsf{com}} \circ \varepsilon_0(\overline{\alpha}_1) \circ \widehat{\varphi}_1 \quad \text{and} \quad [N_2N_{\mathsf{sk}}N_{\mathsf{chl}}]\widehat{\varphi}^{(2)}_{\mathsf{rsp}} = \varphi_2 \circ \varepsilon_0(\alpha_2) \circ \widehat{\varphi}_{\mathsf{sk}} \circ \widehat{\varphi}_{\mathsf{chl}}.$$

Combining these equations with Eq. (3.6), we obtain that

$$[N_1N_2N_{\mathsf{com}}N_{\mathsf{sk}}N_{\mathsf{chl}}2^{b_1+b_2}]\varphi'_{\mathsf{rsp}} = \varphi_2 \circ \varepsilon_0(\alpha_2\overline{\alpha}'\overline{\alpha}_1) \circ \widehat{\varphi}_1.$$

Since $N_1 = \deg(\varphi_1)$ and $N_2 = \deg(\varphi_2)$ are odd, it follows that $\varepsilon_0(\alpha_2\overline{\alpha}'\overline{\alpha}_1)$ is divisible by $2^{b_1+b_2}$ in $\mathrm{End}(E_0)$, so that $\theta := \alpha_2\overline{\alpha}'\overline{\alpha}_1/2^{b_1+b_2} \in \mathcal{O}_0$. The result follows. □

The resulting response is of the form $(\varphi^{(1)}_{\mathsf{rsp}}, \widehat{\varphi}^{(2)}_{\mathsf{rsp}}, \varphi'_{\mathsf{rsp}}(P'_{\mathsf{com}}), \varphi'_{\mathsf{rsp}}(Q'_{\mathsf{com}}), q, a)$. The isogenies $\varphi^{(1)}_{\mathsf{rsp}}$ and $\widehat{\varphi}^{(2)}_{\mathsf{rsp}}$ can be represented as chains of 2-isogenies or, in a more compact way, by kernel generators (or even integers determining a kernel generator in a deterministic basis, as for the challenge isogeny). We summarise the R-SQIsignHD response process in Algorithm 3.12.

3.3.5 Verification

In F-SQIsignHD

Given a response $([2^{e-r}]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}), [2^{e-r}]\varphi_{\mathsf{rsp}}(Q_{\mathsf{com}}), q)$ along with the commitment E_{com} and challenge codomain E_{chl}, the prover proceeds exactly as in Section 2.3.2 to compute a 4-dimensional 2^f-isogeny $F \in \mathrm{End}(E^2_{\mathsf{com}} \times E^2_{\mathsf{chl}})$ embedding φ_{rsp}.

Since q is (f, B)-good, we can easily find $a_1, a_2 \in \mathbb{Z}$ such that $q + a_1^2 + a_2^2 = 2^f$, as explained in Section 2.3.1. We then consider the 4-dimensional 2^f-isogeny

$$F := \begin{pmatrix} \alpha_{\mathsf{com}} & \widetilde{\Phi}_{\mathsf{rsp}} \\ -\Phi_{\mathsf{rsp}} & \widetilde{\alpha}_{\mathsf{chl}} \end{pmatrix} \in \mathrm{End}(E^2_{\mathsf{com}} \times E^2_{\mathsf{chl}}), \tag{3.7}$$

where for $i \in \{\mathsf{com}, \mathsf{chl}\}$, α_i is an $(a_1^2 + a_2^2)$-isogeny

$$\alpha_i := \begin{pmatrix} a_1 & a_2 \\ -a_2 & a_1 \end{pmatrix} \in \mathrm{End}(E_i^2),$$

and Φ_{rsp} is the 2-dimensional diagonal q-isogeny $\Phi_{\mathsf{rsp}} := \mathrm{Diag}(\varphi_{\mathsf{rsp}}, \varphi_{\mathsf{rsp}}) : E^2_{\mathsf{com}} \longrightarrow E^2_{\mathsf{chl}}$.

Such an isogeny could be computed with 2^{f+2}-torsion points forming an isotropic subgroup above its kernel. However, the points $[2^{e-r}]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}), [2^{e-r}]\varphi_{\mathsf{rsp}}(Q_{\mathsf{com}})$ from the response have order 2^r with

Algorithm 3.12: Rigorous response (R-SQIsignHD).

Data: The public parameters pp, the public key E_{pk}, the secret key sk, the commitment E_{com}, the secret commitment data sc and the challenge $\mathsf{chl} \in [\![0\ ;\ 2^\lambda - 1]\!]$.

Result: The response $\mathsf{rsp} = (\varphi^{(1)}_{\mathsf{rsp}}, \widehat{\varphi}^{(2)}_{\mathsf{rsp}}, \varphi'_{\mathsf{rsp}}(P'_{\mathsf{com}}), \varphi'_{\mathsf{rsp}}(Q'_{\mathsf{com}}), q, a)$.

Extract $e, f, N_{\mathsf{com}}, N_{\mathsf{chl}}, N_{\mathsf{sk}}$ from pp;
Parse $I_{\mathsf{sk}}, \varphi_{\mathsf{sk}}(P_0), \varphi_{\mathsf{sk}}(Q_0), \mathcal{B}, \mathcal{B}_{|2^e}, \mathcal{B}_1, \varphi_{\mathsf{sk}|N_{\mathsf{chl}}} \leftarrow \mathsf{sk}$;
Parse $I_{\mathsf{com}}, \varphi_{\mathsf{com}}(P_0), \varphi_{\mathsf{com}}(Q_0) \leftarrow \mathsf{sc}$;
Call Algorithm 3.10 on pp, chl, E_{pk} and sk to compute the challenge isogeny φ_{chl} and its kernel ideal I_{chl};
$I \leftarrow \overline{I}_{\mathsf{com}} \cdot I_{\mathsf{sk}} \cdot I_{\mathsf{chl}}$;
Call Algorithm 3.6 on I to sample $\alpha \in I$ of norm $\leq 2^f \operatorname{nrd}(I)$ uniformly at random;
`// Factoring out the "even part" of the response`
Find the biggest $a \in \mathbb{N}$ such that 2^a divides α;
$\alpha' \leftarrow \alpha/2^a$, $J_{\mathsf{rsp}} \leftarrow I\overline{\alpha}'/\operatorname{nrd}(I)$;
$q \leftarrow \operatorname{nrd}(\alpha)/\operatorname{nrd}(I)$, $b \leftarrow v_2(q) - 2a$, $b_1 \leftarrow \min(b, e)$, $b_2 \leftarrow b - b_1$;
Compute deterministic basis $(P_{\mathsf{com}}, Q_{\mathsf{com}})$ and $(P_{\mathsf{chl}}, Q_{\mathsf{chl}})$ of $E_{\mathsf{com}}[2^e]$ and $E_{\mathsf{chl}}[2^e]$ respectively;
$\mathcal{O}_{\mathsf{com}} \leftarrow O_R(I_{\mathsf{com}})$, $\mathcal{O}_{\mathsf{chl}} \leftarrow O_R(I_{\mathsf{chl}})$;
$I^{(1)}_{\mathsf{rsp}} \leftarrow J_{\mathsf{rsp}} + 2^{b_1}\mathcal{O}_{\mathsf{com}}$, $\overline{I}^{(2)}_{\mathsf{rsp}} \leftarrow \overline{J}_{\mathsf{rsp}} + 2^{b_2}\mathcal{O}_{\mathsf{chl}}$;
Compute $[I_{\mathsf{com}}]^* I^{(1)}_{\mathsf{rsp}}$ and $[I_{\mathsf{sk}} \cdot I_{\mathsf{chl}}]^* \overline{I}^{(2)}_{\mathsf{rsp}}$;
Extract from pp a basis $(\beta_1, \cdots, \beta_4)$ of $\mathcal{O}_0$ and its image $(\varepsilon_0(\beta_j)(P_0), \varepsilon_0(\beta_j)(Q_0))_{1 \leq j \leq 4}$;
$(P_1, Q_1) \leftarrow ([2^{e-a_1}]P_0, [2^{e-a_1}]Q_0)$, $(P_2, Q_2) \leftarrow ([2^{e-a_2}]P_0, [2^{e-a_2}]Q_0)$;
Call Algorithm 3.1 on E_0, $(\beta_1, \cdots, \beta_4)$, $[I_{\mathsf{com}}]^* I^{(1)}_{\mathsf{rsp}}$, (P_1, Q_1) and $([2^{e-a_1}]\varepsilon_0(\beta_i)(P_0), [2^{e-a_1}]\varepsilon_0(\beta_i)(Q_0))_{1 \leq i \leq 4}$ to obtain a generator K_1 of $\ker(\varphi^*_{\mathsf{com}}\varphi^{(1)}_{\mathsf{rsp}})$;
Find $x, y \in \mathbb{Z}$ such that $K_1 = [x]P_1 + [y]Q_1$;
Apply Vélu's formulas to compute the 2-isogeny chain $\varphi^{(1)}_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E'_{\mathsf{com}}$ with kernel $\langle [2^{e-a_1}x]\varphi_{\mathsf{com}}(P_0) + [2^{e-a_1}y]\varphi_{\mathsf{com}}(Q_0) \rangle$;
Call Algorithm 3.1 on E_0, $(\beta_1, \cdots, \beta_4)$, $[I_{\mathsf{sk}} \cdot I_{\mathsf{chl}}]^* \overline{I}^{(2)}_{\mathsf{rsp}}$, (P_2, Q_2) and $([2^{e-a_2}]\varepsilon_0(\beta_i)(P_0), [2^{e-a_2}]\varepsilon_0(\beta_i)(Q_0))_{1 \leq i \leq 4}$ to obtain a generator K_2 of $\ker((\varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}})^* \widehat{\varphi}^{(2)}_{\mathsf{rsp}})$;
Find $z, t \in \mathbb{Z}$ such that $K_2 = [z]P_2 + [t]Q_2$;
Apply Vélu's formulas to compute the 2-isogeny chain $\widehat{\varphi}^{(2)}_{\mathsf{rsp}} : E_{\mathsf{chl}} \longrightarrow E'_{\mathsf{chl}}$ with kernel $\langle [2^{e-a_2}x]\varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}}(P_0) + [2^{e-a_2}y]\varphi_{\mathsf{chl}} \circ \varphi_{\mathsf{sk}}(Q_0) \rangle$;
`// Computing odd degree isogenies` $\varphi_1 : E_0 \longrightarrow E'_{\mathsf{com}}$ `and` $\varphi_2 : E_0 \longrightarrow E'_{\mathsf{chl}}$
$I_1 \leftarrow I_{\mathsf{com}} \cdot I^{(1)}_{\mathsf{rsp}}$, $I_2 \leftarrow I_{\mathsf{sk}} \cdot I_{\mathsf{chl}} \cdot \overline{I}^{(2)}_{\mathsf{rsp}}$;
for $i = 1, 2$ **do**
    Find $\alpha_i \in I_i$ such that $\operatorname{nrd}(\alpha_i)/\operatorname{nrd}(I_i)$ is odd;
    $J_i \leftarrow I_i\overline{\alpha}_i/\operatorname{nrd}(I_i)$, $N_i \leftarrow \operatorname{nrd}(\alpha_i)/\operatorname{nrd}(I_i)$;
    Call Algorithm 2.7 on J_i to obtain the image $(\varphi_i(P_0), \varphi_i(Q_0))$ of the associated isogeny;
end
`// Evaluating` φ'_{rsp}
$\mu \leftarrow 1/(N_1 N_2 N_{\mathsf{com}} N_{\mathsf{chl}} N_{\mathsf{sk}}) \mod 2^e$, $\theta \leftarrow \alpha_2\overline{\alpha}'\overline{\alpha}_1/2^{b_1+b_2}$;
Compute a deterministic basis $(P'_{\mathsf{com}}, Q'_{\mathsf{com}})$ of $E'_{\mathsf{com}}[2^e]$;
Call Algorithm 2.9 on (P_0, Q_0), $(\varphi_1(P_0), \varphi_1(Q_0))$ and $(P'_{\mathsf{com}}, Q'_{\mathsf{com}})$ to compute $(\widehat{\varphi}_1(P'_{\mathsf{com}}), \widehat{\varphi}_1(Q'_{\mathsf{com}}))$;
Find $a', b', c', d' \in \mathbb{Z}/2^e\mathbb{Z}$ such that $\varepsilon_0(\theta) \circ \widehat{\varphi}_1(P'_{\mathsf{com}}) = [a']P_0 + [b']Q_0$ and $\varepsilon_0(\theta) \circ \widehat{\varphi}_1(Q'_{\mathsf{com}}) = [c']P_0 + [d']Q_0$;
$\varphi'_{\mathsf{rsp}}(P'_{\mathsf{com}}) \leftarrow [a'\mu]\varphi_2(P_0) + [b'\mu]\varphi_2(Q_0)$;
$\varphi'_{\mathsf{rsp}}(P'_{\mathsf{com}}) \leftarrow [c'\mu]\varphi_2(P_0) + [d'\mu]\varphi_2(Q_0)$;
return $(\varphi^{(1)}_{\mathsf{rsp}}, \widehat{\varphi}^{(2)}_{\mathsf{rsp}}, \varphi'_{\mathsf{rsp}}(P'_{\mathsf{com}}), \varphi'_{\mathsf{rsp}}(Q'_{\mathsf{com}}), q, a)$;

$r = \lceil f/2 \rceil + 2$. Therefore, we may divide the computation of F in two as explained in Remark 2.3.4. We may decompose $F = F_2 \circ F_1$, where F_i is 2^{f_i}-isogeny for $i \in \{1, 2\}$, with $f = f_1 + f_2$ and $f_1, f_2 \leq r - 2$. Knowing φ_{rsp} on a basis of $E_{\mathsf{com}}[2^r]$, we are able to compute F_1 and $\widetilde{F}_2$, then $F_2 = \widetilde{\widetilde{F}}_2$. If the response is valid, the codomains of F_1 and $\widetilde{F}_2$ should correspond and we can finally compose $F = F_2 \circ F_1$. We refer to Section 6.6.7 for more details. As we shall explain in Sections 6.6 and 6.6.7, the knowledge of $a_1, a_2, q, f, E_{\mathsf{com}}, E_{\mathsf{chl}}$ and the points $[2^{e-r}]P_{\mathsf{com}}, [2^{e-r}]Q_{\mathsf{com}}, [2^{e-r}]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}), [2^{e-r}]\varphi_{\mathsf{rsp}}(Q_{\mathsf{com}})$ is sufficient to compute F_1 and $\widetilde{F}_2$ and then F.

Finally, as we shall see in Lemma 3.4.4, in order to verify that F represents a q-isogeny $\varphi_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$, we have to evaluate $F(P_{\mathsf{com}}, 0, 0, 0)$ and check the following equality:

$$F(P_{\mathsf{com}}, 0, 0, 0) = ([a_1]P_{\mathsf{com}}, -[a_2]P_{\mathsf{com}}, *, 0).$$

We summarise the F-SQIsignHD verification procedure in Algorithm 3.13.

Algorithm 3.13: Fast verification (F-SQIsignHD).

Data: The public parameters pp, the public key E_{pk}, commitment E_{com}, challenge $\mathsf{chl} \in [\![0\ ;\ 2^\lambda - 1]\!]$ and response rsp.
Result: A boolean value indicating if the response is valid.

Extract e, r, λ, f from pp;
Parse $[2^{e-r}]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}), [2^{e-r}]\varphi_{\mathsf{rsp}}(Q_{\mathsf{com}}), q \leftarrow \mathsf{rsp}$;
Compute a deterministic basis $(P_{\mathsf{pk}}, Q_{\mathsf{pk}})$ of $E_{\mathsf{pk}}[2^e]$;
$K_{\mathsf{chl}} \leftarrow [2^{e-\lambda}]P_{\mathsf{pk}} + [2^{e-\lambda}\mathsf{chl}]Q_{\mathsf{pk}}$;
Compute $\varphi_{\mathsf{chl}} : E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}$ of kernel K_{chl};
Compute a deterministic basis $(P_{\mathsf{com}}, Q_{\mathsf{com}})$ of $E_{\mathsf{com}}[2^e]$;
if *q is not (f, B)-good* **then**
  return False;
end
Find $a_1, a_2 \in \mathbb{Z}$ such that $a_1^2 + a_2^2 + q = 2^f$; // Section 2.3.1
$f_1 \leftarrow \lfloor f/2 \rfloor$, $f_2 \leftarrow f - f_1$;
Using $a_1, a_2, q, E_{\mathsf{com}}, E_{\mathsf{chl}}, [2^{e-r}]P_{\mathsf{com}}, [2^{e-r}]Q_{\mathsf{com}}, [2^{e-r}]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}), [2^{e-r}]\varphi_{\mathsf{rsp}}(Q_{\mathsf{com}})$, compute a 2^{f_1}-isogeny $F_1 : E_{\mathsf{com}}^2 \times E_{\mathsf{chl}}^2 \longrightarrow C_1$ and a 2^{f_2}-isogeny $\widetilde{F}_2 : E_{\mathsf{com}}^2 \times E_{\mathsf{chl}}^2 \longrightarrow C_2$ such that $F = F_2 \circ F_1$, where $F \in \mathrm{End}(E_{\mathsf{com}}^2 \times E_{\mathsf{chl}}^2)$ has been defined in Eq. (3.7) ; // Section 6.6
if $C_1 \neq C_2$ **then**
  return False;
end
Compute the polarised dual $F_2 := \widetilde{\widetilde{F}}_2$; // Section 6.2.3
$(T_1, T_2, T_3, T_4) \leftarrow F_2 \circ F_1(P_{\mathsf{com}}, 0, 0, 0)$;
if $T_1 = [a_1]P_{com}$ *and* $T_2 = -[a_2]P_{com}$ *and* $T_4 = 0$ **then**
  return True
else
  return False;
end

In R-SQIsignHD

In R-SQIsignHD, when given a response $(\varphi_{\mathsf{rsp}}^{(1)}, \widehat{\varphi}_{\mathsf{rsp}}^{(2)}, \varphi'_{\mathsf{rsp}}(P'_{\mathsf{com}}), \varphi'_{\mathsf{rsp}}(Q'_{\mathsf{com}}), q, a)$ along with the commitment E_{com} and challenge codomain E_{chl}, the prover computes an 8-dimensional 2^f-isogeny embedding φ'_{rsp}.

First, the prover verifies that $\varphi_{\mathsf{rsp}}^{(1)}$ and $\widehat{\varphi}_{\mathsf{rsp}}^{(2)}$ are (efficient representations of) isogenies with domains E_{com} and E_{chl} respectively and that the points $\varphi'_{\mathsf{rsp}}(P'_{\mathsf{com}}), \varphi'_{\mathsf{rsp}}(Q'_{\mathsf{com}})$ belong to the codomain E'_{chl} of $\widehat{\varphi}_{\mathsf{rsp}}^{(2)}$ and have order 2^e. Then, they compute $b := v_2(q) - 2a$, verify that $b \geq 0$ and that $\varphi_{\mathsf{rsp}}^{(1)}$ and $\widehat{\varphi}_{\mathsf{rsp}}^{(2)}$ have respective degrees 2^{b_1} and 2^{b_2} with $b_1 := \min(b, e)$ and $b_2 := b - b_1$.

The prover then finds $a_1, \cdots, a_4 \in \mathbb{Z}$ such that $a_1^2 + \cdots + a_4^2 + q' = 2^f$ where $q' := q/2^{a+2b}$, *e.g.* using Pollack and Treviño's algorithm [PT18, § 4] and consider the 8-dimensional 2^f-isogeny

$$F := \begin{pmatrix} \alpha_{\mathsf{com}} & \widetilde{\Phi}_{\mathsf{rsp}} \\ -\Phi_{\mathsf{rsp}} & \widetilde{\alpha}_{\mathsf{chl}} \end{pmatrix} \in \mathrm{End}(E'^4_{\mathsf{com}} \times E'^4_{\mathsf{chl}}), \tag{3.8}$$

where for $i \in \{\mathsf{com}, \mathsf{chl}\}$, α_i is an $(a_1^2 + \cdots + a_4^2)$-isogeny

$$\alpha_i := \begin{pmatrix} a_1 & -a_2 & -a_3 & -a_4 \\ a_2 & a_1 & a_4 & -a_3 \\ a_3 & -a_4 & a_1 & a_2 \\ a_4 & a_3 & -a_2 & a_1 \end{pmatrix} \in \mathrm{End}(E_i^4),$$

and Φ_{rsp} is the 4-dimensional diagonal q-isogeny $\Phi_{\mathsf{rsp}} := \mathrm{Diag}(\varphi^{(1)}_{\mathsf{rsp}}, \cdots, \varphi^{(1)}_{\mathsf{rsp}}) : E'^4_{\mathsf{com}} \longrightarrow E'^4_{\mathsf{chl}}$.

As previously, since $f/2 + 2 \leq e$, one can compute F in two parts and the knowledge of $a_1, \cdots, a_4, q', f$, $E'_{\mathsf{com}}, E'_{\mathsf{chl}}$, a deterministic basis $(P'_{\mathsf{com}}, Q'_{\mathsf{com}})$ of $E'_{\mathsf{com}}[2^e]$ and its image $(\varphi'_{\mathsf{rsp}}(P'_{\mathsf{com}}), \varphi'_{\mathsf{rsp}}(Q'_{\mathsf{com}}))$ is sufficient to compute F. Even though it has not been implemented and described in as much detail as in dimension 4, the algorithmic approach is close. We refer to Sections 6.3, 6.4.2 and 6.6 for more details.

As we shall see in Lemma 3.4.7, in order to verify that F represents a q'-isogeny $\varphi'_{\mathsf{rsp}} : E'_{\mathsf{com}} \longrightarrow E'_{\mathsf{chl}}$, we have to evaluate F on $(P'_{\mathsf{com}}, 0, \cdots, 0)$ and $(Q'_{\mathsf{com}}, 0, \cdots, 0)$ in order to check the following equalities:

$$F(P'_{\mathsf{com}}, 0, \cdots, 0) = ([a_1]P'_{\mathsf{com}}, [a_2]P'_{\mathsf{com}}, [a_3]P'_{\mathsf{com}}, [a_4]P'_{\mathsf{com}}, *, 0, 0, 0)$$
$$\text{and} \quad F(Q'_{\mathsf{com}}, 0, \cdots, 0) = ([a_1]Q'_{\mathsf{com}}, [a_2]Q'_{\mathsf{com}}, [a_3]Q'_{\mathsf{com}}, [a_4]Q'_{\mathsf{com}}, *, 0, 0, 0).$$

We summarise the R-SQIsignHD verification procedure in Algorithm 3.14.

3.4 Security analysis

As explained in Section 3.1, the digital signature scheme SQIsignHD is obtained from a 3-round identification scheme via the Fiat-Shamir transform. To obtain a scheme that is universally unforgeable against chosen massage attacks in the random oracle model, the identification scheme needs to be a Σ-protocol *i.e.* to satisfy the security assumptions of Theorem 3.4.2.

Consider a space of values K (playing the role of public keys) and a space of witnesses W (meant to be secret). An NP-*relation* $\mathcal{R}$ on $W \times K$ is a function $W \times K \longrightarrow \{0, 1\}$ that can be evaluated in polynomial time and we define the *language* associated to $\mathcal{R}$ as:

$$\mathcal{L}_{\mathcal{R}} = \{(w, x) \mid \mathcal{R}(w, x) = 1\}.$$

A prover knows some pair $(w, x) \in \mathcal{L}_{\mathcal{R}}$, where x is their public key and the witness w is secret. They want to prove knowledge of w to some verifier without revealing it by performing a 3-round identification protocol (commitment, challenge, response). Of course, such an identification protocol only makes sense if the relation $\mathcal{R}$ is hard.

Definition 3.4.1. The relation $\mathcal{R}$ is *hard* if we can generate pairs $(w, x) \in \mathcal{L}_{\mathcal{R}}$ in probabilistic polynomial time (meaning that key generation is efficient) but given $x \in K$, finding some witness $w \in W$ such that $\mathcal{R}(x, w) = 1$ cannot be done in probabilistic polynomial time with a non-negligible advantage (witness recovery is hard).

Theorem 3.4.2. *[VV15, Theorem 7] Consider a 3-round identification scheme between a prover and a verifier for an NP-relation $\mathcal{R} : W \times K \longrightarrow \{0, 1\}$ that satisfy the following:*

(i) *$\mathcal{R}$ is hard.*

(ii) ***Completeness:*** *The verifier always accepts a honest proof by the prover.*

(iii) ***Special soundness:*** *Given two protocol transcripts $(\mathsf{com}, \mathsf{chl}, \mathsf{rsp})$ and $(\mathsf{com}, \mathsf{chl}', \mathsf{rsp}')$ accepted by the verifier for a public key $\mathsf{pk} \in K$ with the same commitment com and distinct challenges $\mathsf{chl} \neq \mathsf{chl}'$, one can recover a witness $w \in W$ such that $\mathcal{R}(w, \mathsf{pk}) = 1$ in polynomial time.*

Algorithm 3.14: Rigorous verification (R-SQIsignHD).

Data: The public parameters pp, the public key E_{pk}, commitment E_{com}, challenge $\mathsf{chl} \in [\![0\ ;\ 2^\lambda - 1]\!]$ and response rsp.
Result: A boolean value indicating if the response is valid.

Extract $e, f, N_{\mathsf{chl}} = \prod_{i=1}^{n} \ell_i^{e_i}$ from pp;
Parse $\varphi_{\mathsf{rsp}}^{(1)}, \widehat{\varphi}_{\mathsf{rsp}}^{(2)}, \varphi'_{\mathsf{rsp}}(P'_{\mathsf{com}}), \varphi'_{\mathsf{rsp}}(Q'_{\mathsf{com}}), q, a \leftarrow \mathsf{rsp}$;
// Computation of E_{chl}
for $i = 1$ **to** n **do**
  Generate a deterministic basis (P_i, Q_i) of $E_{\mathsf{pk}}[\ell_i^{e_i}]$;
  $K_i \leftarrow P_i + [\mathsf{chl}]Q_i$;
  if $i \geq 2$ **then**
    $K_i \leftarrow \varphi_{i-1} \circ \cdots \circ \varphi_1(K_i)$;
  end
  Compute $\varphi_i : E_{i-1} \longrightarrow E_i$ of kernel $\langle K_i \rangle$;
end
$E_{\mathsf{chl}} \leftarrow E_n$;
// Verification of $\varphi_{\mathsf{rsp}}^{(1)}$ and $\widehat{\varphi}_{\mathsf{rsp}}^{(2)}$
if $\mathit{Domain}(\varphi_{rsp}^{(1)}) \neq E_{com}$ *or* $\mathit{Domain}(\widehat{\varphi}_{rsp}^{(2)}) \neq E_{chl}$ **then**
  return False;
end
$E'_{\mathsf{chl}} \leftarrow \mathsf{Codomain}(\widehat{\varphi}_{\mathsf{rsp}}^{(2)})$;
if $\varphi'_{rsp}(P'_{com}) \notin E'_{chl}$ *or* $\varphi'_{rsp}(Q'_{com}) \notin E'_{chl}$ **then**
  return False;
end
if $\mathit{order}(\varphi'_{rsp}(P'_{com})) \neq 2^e$ *or* $\mathit{order}(\varphi'_{rsp}(Q'_{com})) \neq 2^e$ **then**
  return False;
end
$b \leftarrow v_2(q) - 2a$, $q' \leftarrow q/2^{a+2b}$, $a_1 \leftarrow \min(a, e)$, $a_2 \leftarrow a - a_1$;
if $\deg(\varphi_{rsp}^{(1)}) \neq 2^{a_1}$ *or* $\deg(\varphi_{rsp}^{(1)}) \neq 2^{a_2}$ **then**
  return False;
end
// Computation of $F \in \mathrm{End}(E'^4_{\mathsf{com}} \times E'^4_{\mathsf{chl}})$
Generate a deterministic basis $(P_{\mathsf{com}}, Q_{\mathsf{com}})$ of $E_{\mathsf{com}}[2^e]$;
if $q \geq 2^f$ **then**
  return False;
end
Find $a_1, \cdots, a_4 \in \mathbb{Z}$ such that $a_1^2 + \cdots + a_4^2 + q' = 2^f$ using [PT18, § 4];
$f_1 \leftarrow \lfloor f/2 \rfloor$, $f_2 \leftarrow f - f_1$;
Using $a_1, \cdots, a_4, q', E'_{\mathsf{com}}, E'_{\mathsf{chl}}, P'_{\mathsf{com}}, Q'_{\mathsf{com}}, \varphi'_{\mathsf{rsp}}(P'_{\mathsf{com}}), \varphi'_{\mathsf{rsp}}(Q'_{\mathsf{com}})$, compute a 2^{f_1}-isogeny $F_1 : E'^4_{\mathsf{com}} \times E'^4_{\mathsf{chl}} \longrightarrow C_1$ and a 2^{f_2}-isogeny $\widetilde{F}_2 : E'^4_{\mathsf{com}} \times E'^4_{\mathsf{chl}} \longrightarrow C_2$ such that $F = F_2 \circ F_1$, where $F \in \mathrm{End}(E'^4_{\mathsf{com}} \times E'^4_{\mathsf{chl}})$ has been defined in Eq. (3.8) ; // Sections 6.3 and 6.4.2
if $C_1 \neq C_2$ **then**
  return False;
end
Compute the polarised dual $F_2 := \widetilde{\widetilde{F}}_2$; // Section 6.2.3
// Point evaluations
$T \leftarrow F_2 \circ F_1(P'_{\mathsf{com}}, 0, 0, 0, 0, 0, 0, 0)$;
$U \leftarrow F_2 \circ F_1(Q'_{\mathsf{com}}, 0, 0, 0, 0, 0, 0, 0)$;
if $T = ([a_1]P'_{com}, [a_2]P'_{com}, [a_3]P'_{com}, [a_4]P'_{com}, *, 0, 0, 0)$ *and* $U = ([a_1]Q'_{com}, [a_2]Q'_{com}, [a_3]Q'_{com}, [a_4]Q'_{com}, *, 0, 0, 0)$ **then**
  return True
else
  return False;
end

(iv) ***Honest verifier zero-knowledge (HVZK):*** *There exists a polynomial time simulator $\mathcal{S}$ that can generate transcripts $(\mathsf{com}, \mathsf{chl}, \mathsf{rsp})$ for any public key $\mathsf{pk} \in K$ with a distribution computationally indistinguishable from honest executions of the protocol for the same public key.*

Then this identification scheme yields a digital signature scheme via the Fiat-Shamir transform which is universally unforgeable under chosen message attacks in the random oracle model.

Intuitively, the special soundness ensures that an attacker cannot "guess" responses with a non-negligible advantage. Indeed, if they can produce valid responses with non negligible advantage, then with non-negligible probability, they are able to generate two valid transcripts $(\mathsf{com}, \mathsf{chl}, \mathsf{rsp})$ and $(\mathsf{com}, \mathsf{chl}', \mathsf{rsp}')$ for a public key $\mathsf{pk} \in K$ with the same commitment com and distinct challenges $\mathsf{chl} \neq \mathsf{chl}'$. By special soundness, they can recover a witness $w \in W$ such that $\mathcal{R}(w, \mathsf{pk}) = 1$ in polynomial time with non negligible probability, breaking the hardness of $\mathcal{R}$. Intuitively, the HVZK property means that transcripts do not leak any information on the secret key, since they can be simulated without knowing it.

Both versions of the SQIsignHD identification protocol is based on the NP-relation $\mathcal{R} : W \times K \longrightarrow \{0, 1\}$, where K is the space of supersingular elliptic curves over $\mathbb{F}_{p^2}$ and W is the space of endomorphisms of supersingular elliptic curves over $\mathbb{F}_{p^2}$ and:

$$\forall (\alpha, E) \in W \times K, \quad \mathcal{R}(\alpha, E) = 1 \iff \alpha \in \mathrm{End}(E) \setminus \mathbb{Z}. \tag{3.9}$$

This relation is hard provided that the *Supersingular One Endomorphism Problem* is hard.

Problem 3.4.3 (Supersingular One Endomorphism Problem)**.** Given a supersingular elliptic curve $E/\mathbb{F}_{p^2}$, find (an efficient representation of) a non-scalar endomorphism $\alpha \in \mathrm{End}(E) \setminus \mathbb{Z}$.

The fastest known algorithms for this problem have classical complexity in $\widetilde{O}(p^{1/2})$ [DG16; EHLMP20; FIKMN25] (see also [PW24, Theorem 8.8]) and a quantum complexity $\widetilde{O}(p^{1/4})$ using Grover's algorithm [Gro96; BJS14]. This problem is also proved to be equivalent to the Supersingular Endomorphism Ring Problem (Problem 2.1.2) and the Supersingular Isogeny Problem (Problem 2.1.1) [PW24; Wes22; MW25] which is the underlying problem of a key recovery attack. For these reasons, we believe the relation $\mathcal{R}$ to be hard and in order to ensure λ bits of classical and $\lambda/2$ bits of quantum security, we chose $p = \Theta(2^{2\lambda})$, so that $e \simeq 2\lambda$ as explained before.

The completeness of the protocol is basically trivial by construction: in Section 3.3, we described an identification protocol so that the verifier always accepts a honest execution (*i.e.* Algorithms 3.13 and 3.14 always return True if the prover executed their part honestly). It remains to prove the special soundness and HVZK properties that are discussed in the following sections, the latter being the most delicate. Note that in the original version of SQIsign, the HVZK argument was very ad-hoc and almost tautological: the response isogeny distribution was assumed to be computationally indistinguishable from uniform, even though it was known not to be uniform. The argument is stronger in both variants of SQIsignHD as it relies on more natural assumptions. We shall see that using an 8-dimensional isogeny of polarised degree $2^f = \Theta(p^2)$ in R-SQIsignHD simplifies greatly the proof of HVZK compared to F-SQIsignHD, at the expense of a more technical special soundness proof.

3.4.1 Special soundness

In F-SQIsignHD

We start by proving that the verification Algorithm 3.13 works as expected, meaning that a response that has been validated (the algorithm returns True) always efficiently represents an isogeny $\varphi : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$ of (f, B)-good degree.

Lemma 3.4.4. *Let $(E_{com}, \mathsf{chl}, \mathsf{rsp})$ be a transcript of the F-SQIsignHD identification protocol. If the response $\mathsf{rsp} = ([2^{e-r}]\varphi_{rsp}(P_{com}), [2^{e-r}]\varphi_{rsp}(Q_{com}), q)$ has been validated by the verifier* i.e. *if Algorithm 3.13 returned True, then q is (f, B)-good and rsp is an efficient representation of a q-isogeny $\varphi : E_{com} \longrightarrow E_{chl}$.*

Proof. Assume that Algorithm 3.13 has returned True. Then q is (f, B)-good by construction and Algorithm 3.13 has computed a 2^f-isogeny $F \in \mathrm{End}(E_{\mathsf{com}}^2 \times E_{\mathsf{chl}}^2)$ such that

$$F(P_{\mathsf{com}}, 0, 0, 0) = ([a_1]P_{\mathsf{com}}, -[a_2]P_{\mathsf{com}}, *, 0),$$

where $(P_{\mathsf{com}}, Q_{\mathsf{com}})$ is a deterministic basis of $E_{\mathsf{com}}[2^e]$. We may write $F := (f_{i,j})_{1\leq i,j\leq 4}$ in matrix notation. Then, since $\widetilde{F} \circ F = [2^f]$, we get that

$$\forall 1 \leq j \leq 4, \quad \sum_{i=1}^{4} \deg(f_{i,j}) = 2^f, \tag{3.10}$$

so the $f_{i,j}$ have degree $\leq 2^f$. By assumption, $f_{1,1}(P_{\mathsf{com}}) = [a_1]P_{\mathsf{com}}$, $f_{2,1}(P_{\mathsf{com}}) = -[a_2]P_{\mathsf{com}}$ and $f_{4,1}(P_{\mathsf{com}}) = 0$. Besides, by Cauchy-Schwarz inequality

$$\deg(f_{1,1} - [a_1]) \leq \left(\sqrt{\deg(f_{1,1})} + \sqrt{\deg([a_1])}\right)^2 < (2^{f/2} + 2^{f/2})^2 = 2^{f+2}.$$

Since $f \simeq e/2$, we have $e > f + 2$ so $f_{1,1} - [a_1]$ is annihilated by a subgroup of cardinality $2^e > \deg(f_{1,1} - [a_1])$ of E_{com}. It follows that $f_{1,1} = [a_1]$. We similarly obtain that $f_{2,1} = -[a_2]$ and $f_{4,1} = 0$. Hence, by Eq. (3.10), we get $\deg(f_{3,1}) = 2^f - a_1^2 - a_2^2 = q$. Besides, by construction $\varphi := f_{3,1}$ is a morphism $E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$, so it is a q-isogeny, as desired. Since F can be evaluated in polynomial time in $\log(p)$, so can φ. Hence, F is an efficient representation of φ. But F can be obtained from rsp in polynomial time in $\log(p)$ by Theorem 2.2.12. This proves that rsp is an efficient representation of φ and completes the proof. □

We also need to prove that an efficient representation of an isogeny is essentially the same as an efficient representation of its dual, a very intuitive and very useful result of independent interest that follows directly from the interpolation theorem (Theorem 2.2.10).

Lemma 3.4.5. *Let $\varphi : E_1 \longrightarrow E_2$ be a d-isogeny between two elliptic curves defined over $\mathbb{F}_q$. Then one can obtain an efficient representation of $\widehat{\varphi}$ in polynomial time in $\log(d)$ and $\log(q)$ from an efficient representation of φ.*

Proof. Similarly to the proof of Lemma 3.3.4, let $C \in \mathbb{N}^*$ and consider the products of primes:

$$N = \prod_{\substack{\ell \leq C \\ \ell \nmid d}} \ell \quad \text{and} \quad D = \prod_{\ell \leq C} \ell.$$

Then by [HW75, Theorems 413 and 434], $\log(D) \sim C$ as $C \longrightarrow +\infty$, so that $D \geq e^{C/2}$ for C big enough, so that $N \geq D/d \geq e^{C/2}/d$. Choosing $C = 4\log(d) + O(1)$, we obtain that $N > d$ as $d \longrightarrow +\infty$.

Let us write $N := \prod_{i=1}^{r} \ell_i$ with distinct primes $\ell_1, \cdots, \ell_r$. Then r and all the ℓ_i are in $O(\log(d))$. For all $i \in [\![1\ ;\ r]\!]$, we may generate a basis (P_i, Q_i) of $E_1[\ell_i]$. Since $E_1[\ell_i]$ is defined over an extension $\mathbb{F}_{q^{k_i}}/\mathbb{F}_q$ of degree $k_i \leq \ell_i^2$, this basis (P_i, Q_i) can be generated in polynomial time in $\log(d)$ and $\log(q)$ (*e.g.* using division polynomials or via point counting and scalar multiplication) and evaluated via φ in polynomial time in $\log(d)$ and $\log(q)$, for all $i \in [\![1\ ;\ r]\!]$ (with $r = O(\log(d))$). By Theorem 2.2.10, $(P_i, Q_i)_{1\leq i\leq r}$, $(\varphi(P_i), \varphi(Q_i))_{1\leq i\leq r}$ (together with d, E_1 and E_2) is an efficient representation of φ and $\widehat{\varphi}$. This completes the proof. □

We are now able to prove the:

Theorem 3.4.6. *The identification protocol F-SQIsignHD is special sound for the NP-relation defined in Eq. (3.9).*

Proof. Let $(E_{\mathsf{com}}, \mathsf{chl}, \mathsf{rsp})$ and $(E_{\mathsf{com}}, \mathsf{chl}', \mathsf{rsp}')$ be two transcripts for the same public key E_{pk} and commitment E_{com} but distinct challenges $\mathsf{chl} \neq \mathsf{chl}'$. Then we can access efficient representations of $\varphi_{\mathsf{chl}} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$ and $\varphi'_{\mathsf{chl}} : E_{\mathsf{com}} \longrightarrow E'_{\mathsf{chl}}$ from chl and chl'; and of $\varphi_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$ and $\varphi'_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E'_{\mathsf{chl}}$ from rsp and rsp' by Lemma 3.4.4 (since rsp and rsp' are valid). We may then consider the endomorphism

$$\alpha := \widehat{\varphi'}_{\mathsf{chl}} \circ \varphi'_{\mathsf{rsp}} \circ \widehat{\varphi}_{\mathsf{rsp}} \circ \varphi_{\mathsf{chl}} \in \mathrm{End}(E_{\mathsf{pk}})$$

We now prove that α is not scalar. Indeed, if it was, we would have $\alpha = [m]$ for some $m \in \mathbb{Z}$ and $qq'2^{2\lambda} = m^2$ where $q := \deg(\varphi_{\mathsf{rsp}})$ and $q' := \deg(\varphi'_{\mathsf{rsp}})$. By Lemma 3.4.4, q and q' are (f, B)-good so they are odd, so we may write $m = 2^f m'$ with $m' \in \mathbb{Z}$ odd ($m'^2 = qq'$). It follows that $[q']\widehat{\varphi}_{\mathsf{rsp}} \circ \varphi_{\mathsf{chl}} = [m']\widehat{\varphi'}_{\mathsf{rsp}} \circ \varphi'_{\mathsf{chl}}$. Since q, q' and m' are odd, we get that $\ker(\varphi_{\mathsf{chl}}) = \ker(\varphi'_{\mathsf{chl}})$ *i.e.* $\mathsf{chl} = \mathsf{chl}'$. Contradiction. This proves that α is not scalar.

Now, an efficient representation of φ'_{chl} yields an efficient representation of $\widehat{\varphi'}_{\mathsf{chl}}$ in polynomial time (in $\log(p)$) since we can compute $\ker(\widehat{\varphi'}_{\mathsf{chl}}) = \varphi'_{\mathsf{chl}}(E_{\mathsf{pk}}[2^\lambda])$. Besides, from an efficient representation of φ_{rsp}, we can obtain an efficient representation of $\widehat{\varphi}_{\mathsf{rsp}}$ by Lemma 3.4.5. We then have access to efficient representations of $\varphi_{\mathsf{chl}}, \widehat{\varphi'}_{\mathsf{chl}}, \widehat{\varphi}_{\mathsf{rsp}}$ and φ'_{rsp} that yield an efficient representation of α by composition. This completes the proof. □

In R-SQIsignHD

Lemma 3.4.7. *Let $(E_{com}, E_{chl}, \mathsf{rsp})$ be a transcript of the R-SQIsignHD identification protocol. If the response $\mathsf{rsp} = (\varphi_{rsp}(P_{com}), \varphi_{rsp}(Q_{com}), q)$ has been validated by the verifier* i.e. *if Algorithm 3.14 returned* True*, then rsp is an efficient representation of a q-isogeny $\varphi : E_{com} \longrightarrow E_{chl}$ with $q < 2^f$.*

Proof. Unsurprisingly, the proof is very similar to the proof of Lemma 3.4.4. Assume that Algorithm 3.13 has returned True. Then, keeping the notations from this algorithm, we have $q < 2^f$ and we have an efficient representation of a 2^{b_1}-isogeny $\varphi^{(1)} : E_{\mathsf{com}} \longrightarrow E'_{\mathsf{com}}$, a 2^{b_2}-isogeny $\widehat{\varphi}^{(2)} : E_{\mathsf{chl}} \longrightarrow E'_{\mathsf{chl}}$ and a 2^f-isogeny $F \in \mathrm{End}(E'^4_{\mathsf{com}} \times E'^4_{\mathsf{chl}})$ such that

$$F(P'_{\mathsf{com}}, 0, \cdots, 0) = ([a_1]P'_{\mathsf{com}}, [a_2]P'_{\mathsf{com}}, [a_3]P'_{\mathsf{com}}, [a_4]P'_{\mathsf{com}}, *, 0, 0, 0)$$
$$\text{and} \quad F(Q'_{\mathsf{com}}, 0, \cdots, 0) = ([a_1]Q'_{\mathsf{com}}, [a_2]Q'_{\mathsf{com}}, [a_3]Q'_{\mathsf{com}}, [a_4]Q'_{\mathsf{com}}, *, 0, 0, 0),$$

where $(P'_{\mathsf{com}}, Q'_{\mathsf{com}})$ is a deterministic basis of $E'_{\mathsf{com}}[2^e]$. As in the proof of Lemma 3.4.4, we may write $F := (f_{i,j})_{1 \le i,j \le 8}$ in matrix notation. Then, since $\widetilde{F} \circ F = [2^f]$, we get that

$$\forall 1 \le j \le 8, \quad \sum_{i=1}^{8} \deg(f_{i,j}) = 2^f, \tag{3.11}$$

so the $f_{i,j}$ have degree $\le 2^f$. By assumption, we have $f_{1,1}$ and $[a_1]$ coincide on $(P'_{\mathsf{com}}, Q'_{\mathsf{com}})$, so they coincide on a subgroup of cardinality 2^{2e} of E'_{com}. Applying Cauchy-Schwarz inequality again, we obtain:

$$\deg(f_{1,1} - [a_1]) \le \left(\sqrt{\deg(f_{1,1})} + \sqrt{\deg([a_1])}\right)^2 < (2^{f/2} + 2^{f/2})^2 = 2^{f+2},$$

where $f + 2 = 2(e-4) + 2 < 2e$. It follows that $f_{1,1} = [a_1]$. We obtain similarly that $f_{2,1} = [a_2]$, $f_{3,1} = [a_3]$, $f_{4,1} = [a_4]$, $f_{6,1} = f_{7,1} = f_{8,1} = 0$. By Eq. (3.11), we also obtain that $\varphi^{(1)} := f_{5,1}$ is a q'-isogeny $E'_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$, so that $\varphi := \varphi^{(1)} \circ \varphi^{(0)}$ is a q-isogeny $E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$. We complete the proof with the same arguments used in Lemma 3.4.4. □

In F-SQIsignHD, the special soundness argument easily follows from the fact that the response isogeny $\varphi_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$ dos not backtrack through the challenge isogeny $\varphi_{\mathsf{chl}} : E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}$ *i.e.* that $\ker(\widehat{\varphi}_{\mathsf{rsp}}) \cap \ker(\widehat{\varphi}_{\mathsf{chl}}) = 0$ since their degrees are coprime. In R-SQIsignHD, this assumption no longer holds but the challenge argument is saved by the fact that two distinct challenges φ_{chl} and φ'_{chl} starting from the same public curve E_{pk} con only have a *greatest common factor* of small degree in the following sense.

Lemma 3.4.8. *Let $\phi : E_1 \longrightarrow E_2$ and $\phi' : E_1 \longrightarrow E'_2$ be two cyclic isogenies. Let $\phi_0 : E_1 \longrightarrow E_3$ be the cyclic separable isogeny such that $\ker(\phi_0) = \ker(\phi) \cap \ker(\phi')$. Then, there exists two cyclic isogenies $\phi_1 : E_3 \longrightarrow E_2$ and $\phi'_1 : E_3 \longrightarrow E'_2$ such that $\phi = \phi_1 \circ \phi_0$, $\phi' = \phi'_1 \circ \phi_0$ and $\phi'_1 \circ \widehat{\phi_1}$ is cyclic. ϕ_0 will be called the* greatest common factor *of ϕ and ϕ'.*

Proof. We can always factor $\phi = \phi_1 \circ \phi_0$, $\phi' = \phi'_1 \circ \phi_0$, with $\phi_0 : E_1 \longrightarrow E_3$ such that $\ker(\phi_0) = \ker(\phi) \cap \ker(\phi')$.

First, we prove that $\ker(\phi_1) \cap \ker(\phi_1') = \{0\}$. If it was not the case, we could find a cyclic isogeny $\psi : E_2 \longrightarrow E_4$ with non-trivial kernel such that ϕ_1 and ϕ_1' factor through ψ, and both ϕ and ϕ' would factor through $\psi \circ \phi_0$, so we would have $\ker(\psi \circ \phi_0) \subseteq \ker(\phi) \cap \ker(\phi')$. But $\ker(\psi \circ \phi_0) = \phi_0^{-1}(\ker(\psi)) \supsetneq \ker(\phi_0)$ since $\ker(\psi)$ is non-trivial. Contradiction. Now, we prove that $\phi_1' \circ \widehat{\phi_1}$ is cyclic.

Step 1: Actually, it suffices to prove it when $\deg(\phi_1')$ and $\deg(\phi_1)$ are powers of the same prime ℓ. Indeed, if not, we can decompose $\phi_1 = \psi_2 \circ \psi_1$ and $\phi_1' = \psi_2' \circ \psi_1'$ with $\deg(\psi_1)\deg(\psi_1')$ coprime with $\deg(\psi_2)\deg(\psi_2')$. Then, we may write $\phi_1' \circ \widehat{\phi_1} = \psi_2' \circ \psi_3 \circ \widehat{\psi_2}$, with $\psi_3 = \psi_1' \circ \widehat{\psi_1}$. Using pushforward isogenies (as defined in Definition 1.1.12), we may write $[\psi_2]_*\psi_3 \circ \psi_2 = [\psi_3]_*\psi_2 \circ \psi_3$, so that $\psi_3 \circ \widehat{\psi_2} = \widehat{[\psi_3]_*\psi_2} \circ [\psi_2]_*\psi_3$ and

$$\phi_1' \circ \widehat{\phi_1} = \psi_2' \circ \widehat{[\psi_3]_*\psi_2} \circ [\psi_2]_*\psi_3.$$

If we assume that ψ_3 is cyclic then $[\psi_2]_*\psi_3$ is cyclic. Besides, ψ_2' and $[\psi_3]_*\psi_2$ are cyclic and $\ker(\psi_2') \cap \ker([\psi_3]_*\psi_2) = \{0\}$. Indeed, if $P \in \ker(\psi_2') \cap \ker([\psi_3]_*\psi_2)$, then $P \in \ker([\psi_3]_*\psi_2) = \psi_3(\ker(\psi_2))$ so we may write $P = \psi_3(Q)$ with $Q \in \ker(\psi_2)$. Let $R := \widehat{\psi_1}(Q)$. Then, $\phi_1(R) = \psi_2 \circ \psi_1 \circ \widehat{\psi_1}(Q) = [\deg(\psi_1)]\psi_2(Q) = 0$ and $\phi_1'(R) = \psi_2'(P) = 0$, so $R \in \ker(\phi_1) \cap \ker(\phi_1') = \{0\}$ and $R = 0$, so $P = \psi_1'(R) = 0$.

Besides, we have

$$\deg([\psi_2]_*\psi_3) = \deg(\psi_3) = \deg(\psi_1)\deg(\psi_1') \quad \text{and} \quad \deg(\psi_2' \circ \widehat{[\psi_3]_*\psi_2}) = \deg(\psi_2)\deg(\psi_2').$$

Since products of cyclic isogenies of coprime degrees are cyclic, if we can prove that $\psi_2' \circ \widehat{[\psi_3]_*\psi_2}$ is cyclic, then we can prove that $\phi_1' \circ \widehat{\phi_1}$ is cyclic.

Step 2: By Step 1, we can assume that $\deg(\phi_1')$ and $\deg(\phi_1)$ are powers of the same prime ℓ. We proceed by induction on $\deg(\phi_1')$. When $\deg(\phi_1') = 1$ it follows from the fact that the dual of a cyclic isogeny is cyclic. Now, we assume the result holds when $\deg(\phi_1') = \ell^n$ with $n \in \mathbb{N}$ and prove it holds when $\deg(\phi_1') = \ell^{n+1}$. We may factor $\phi_1' := \phi_2 \circ \phi_2'$ with $\deg(\phi_2) = \ell$ and $\deg(\phi_2') = \ell^n$. By assumption, $\phi_3 := \phi_2' \circ \widehat{\phi_1}$ is cyclic so we only have to prove that $\phi_2 \circ \phi_3$ is cyclic, *i.e.* that $\ker(\phi_2 \circ \phi_3) = \phi_3^{-1}(\ker(\phi_2))$ is cyclic.

Let Q be a generator of $\ker(\phi_2)$, P be a generator of $\ker(\phi_3)$ and $P' \in E_2$ such that $Q = \phi_3(P')$. Then

$$\ker(\phi_2 \circ \phi_3) = \phi_3^{-1}(\ker(\phi_2)) = \langle P, P' \rangle.$$

To conclude, it suffices to prove that $P \in \langle P' \rangle$. We have $P' \in \ker(\phi_2 \circ \phi_3) \subset E_2[\ell^{m+1}]$, with $\deg(\phi_3) := \ell^m$ and $[\ell^m]P' = \widehat{\phi_3} \circ \phi_3(P') = \widehat{\phi_3}(Q)$ and $\widehat{\phi_3}(Q) \neq 0$. Indeed, if $\widehat{\phi_3}(Q) = 0$, we may write $S := \widehat{\phi_2'}(Q)$, so that $\phi_1(S) = \widehat{\phi_3}(Q) = 0$ and $\phi_1'(S) = \phi_2 \circ \phi_2' \circ \widehat{\phi_2'}(Q) = [\ell^n]\phi_2(Q) = 0$, so that $S \in \ker(\phi_1) \cap \ker(\phi_1') = \{0\}$ and $\widehat{\phi_2'}(Q) = 0$. Hence, ϕ_2' factors through ϕ_2 and ϕ_1' factors through $[\ell] = \phi_2 \circ \widehat{\phi_2}$ and is not cyclic. Contradiction. Hence, $\widehat{\phi_3}(Q) \neq 0$ and P' has order ℓ^{m+1}.

Let $R \in E_2[\ell^m]$ such that $([\ell]P', R)$ is a basis of $E_2[\ell^m]$. Then, we may write $P := [a\ell]P' + [b]R$ for some $a, b \in \mathbb{Z}$ since $P \in \ker(\phi_3) \subset E_2[\ell^m]$. Since $Q \in \ker(\phi_2)$ has order ℓ, we get that

$$0 = \phi_3(P) = [a\ell]Q + [b]\phi_3(R) = [b]\phi_3(R),$$

and that $\phi_3(R)$ generates $\phi_3(E_2[\ell^m]) = \ker(\widehat{\phi_3})$, which is cyclic so it has order ℓ^m. It follows that $b \equiv 0 \mod \ell^m$, so that $P = [a\ell]P' \in \langle P' \rangle$. This completes the proof. □

Lemma 3.4.9. *Let $\varphi_{chl} : E_{\mathsf{pk}} \longrightarrow E_{chl}$ and $\varphi'_{chl} : E_{\mathsf{pk}} \longrightarrow E'_{chl}$ be two distinct challenges from the same public curve E_{pk}. Then, the largest integer dividing $\varphi'_{chl} \circ \widehat{\varphi}_{chl} \in \mathrm{Hom}(E_{chl}, E'_{chl})$ is $< 2^\lambda$.*

Proof. Recall that the challenge isogeny φ_{chl} has degree $N_{\mathsf{chl}} = \prod_{i=1}^n \ell_i^{e_i}$ and is defined by the kernel $\langle K_{\mathsf{chl}} \rangle$ with $K_{\mathsf{chl}} = P + [\mathsf{chl}]Q$ where $0 \leq \mathsf{chl} < 2^\lambda$, $P := \sum_{i=1}^n P_i$ and $Q := \sum_{i=1}^n Q_i$ and (P_i, Q_i) is a deterministic basis of $E_{\mathsf{pk}}[\ell_i^{e_i}]$ for all $i \in [\![1\ ;\ n]\!]$. In particular, (P, Q) form a basis of $E_{\mathsf{pk}}[N_{\mathsf{chl}}]$ (that cannot be represented without using exponentially large field extensions of $\mathbb{F}_{p^2}$).

The second challenge isogeny φ'_{chl} is defined similarly by its kernel generator $K_{\mathsf{chl}'} = P + [\mathsf{chl}']Q$, for some $\mathsf{chl} \neq \mathsf{chl}'$. Since φ_{chl} and φ'_{chl} are cyclic, by Lemma 3.4.8 there exists three cyclic isogenies $\phi_0 : E_{\mathsf{pk}} \longrightarrow E$, $\phi_1 : E \longrightarrow E_{\mathsf{chl}}$ and $\phi_1' : E \longrightarrow E'_{\mathsf{chl}}$ such that $\varphi_{\mathsf{chl}} = \phi_1 \circ \phi_0$, $\varphi'_{\mathsf{chl}} = \phi_1' \circ \phi_0$ and

$\phi_1' \circ \widehat{\phi_1}$ is cyclic. The *greatest cyclic factor* ϕ_0 of φ_{chl} and φ'_{chl} has kernel $\ker(\phi_0) = \ker(\varphi_{\mathsf{chl}}) \cap \ker(\varphi'_{\mathsf{chl}})$. Since $\varphi'_{\mathsf{chl}} \circ \widehat{\varphi}_{\mathsf{chl}} = [\deg(\phi_0)]\phi_1' \circ \widehat{\phi_1}$, we see that $\deg(\phi_0)$ is the largest integer dividing $\varphi'_{\mathsf{chl}} \circ \widehat{\varphi}_{\mathsf{chl}}$ in $\mathrm{Hom}(E_{\mathsf{chl}}, E'_{\mathsf{chl}})$, so we only have to prove that $\deg(\phi_0) < 2^\lambda$.

Let $R \in E_{\mathsf{pk}}$ be a generator of $\ker(\phi_0)$. Then, $R = [a]K_{\mathsf{chl}} = [b]K_{\mathsf{chl}'}$ for some $a, b \in [\![0\ ;\ N_{\mathsf{chl}} - 1]\!]$, *i.e.* :

$$[a-b]P + [a \cdot \mathsf{chl} - b \cdot \mathsf{chl}']Q = 0.$$

Since (P, Q) is a basis of $E_{\mathsf{pk}}[N_{\mathsf{chl}}]$, it follows that $a - b \equiv 0 \mod N_{\mathsf{chl}}$ so $a = b$ and $a(\mathsf{chl} - \mathsf{chl}') \equiv 0 \mod N_{\mathsf{chl}}$. It follows that N_{chl}/d divides a, where $d := \gcd(\mathsf{chl} - \mathsf{chl}', N_{\mathsf{chl}})$, so that $R \in E_{\mathsf{pk}}[d]$. Since $0 \leq \mathsf{chl} \neq \mathsf{chl}' < 2^\lambda$, it follows that $d < 2^\lambda$, and $\deg(\phi_0) \leq d < 2^\lambda$. This completes the proof. □

We are now able to finally prove the:

Theorem 3.4.10. *Assume that $N_{\mathsf{chl}} \geq 2^{f+\lambda}$. Then the identification protocol R-SQIsignHD is special sound for the NP-relation defined in Eq. (3.9).*

Proof. We proceed as in the proof of Theorem 3.4.6. Let $(E_{\mathsf{com}}, \mathsf{chl}, \mathsf{rsp})$ and $(E_{\mathsf{com}}, \mathsf{chl}', \mathsf{rsp}')$ be two transcripts for the same public key E_{pk} and commitment E_{com} but distinct challenges $\mathsf{chl} \neq \mathsf{chl}'$. Then we have seen that we can access efficient representations of $\varphi_{\mathsf{chl}} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$ and $\varphi'_{\mathsf{chl}} : E_{\mathsf{com}} \longrightarrow E'_{\mathsf{chl}}$ from chl and chl'; and of $\varphi_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$ and $\varphi'_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E'_{\mathsf{chl}}$ from rsp and rsp' by Lemma 3.4.7 (since rsp and rsp' are valid). As in the proof of Theorem 3.4.6, consider the endomorphism

$$\alpha := \widehat{\varphi'}_{\mathsf{chl}} \circ \varphi'_{\mathsf{rsp}} \circ \widehat{\varphi}_{\mathsf{rsp}} \circ \varphi_{\mathsf{chl}} \in \mathrm{End}(E_{\mathsf{pk}}).$$

Using similar arguments as previously (and Lemma 3.4.5 in particular), we can obtain an efficient representation of α in polynomial time in $\log(p)$ from efficient representations of φ_{chl}, φ'_{chl}, φ_{rsp} and φ'_{rsp}.

To conclude, we prove that α is non-scalar. Suppose by contradiction that $\alpha = [m]$ for some $m \in \mathbb{Z}$. We deduce

$$[m] \circ \varphi'_{\mathsf{chl}} \circ \widehat{\varphi}_{\mathsf{chl}} = [N_{\mathsf{chl}}^2] \circ \varphi'_{\mathsf{rsp}} \circ \widehat{\varphi}_{\mathsf{rsp}}. \tag{3.12}$$

Using Lemma 3.4.8, write $\varphi'_{\mathsf{chl}} \circ \widehat{\varphi}_{\mathsf{chl}} = [d] \circ \psi$ and $\varphi'_{\mathsf{rsp}} \circ \widehat{\varphi}_{\mathsf{rsp}} = [d'] \circ \nu$ where ψ and ν have cyclic kernel. We deduce from Eq. (3.12) that $dm = d' N_{\mathsf{chl}}^2$ is the largest integer dividing either side of the equality, and $\psi = \nu$ is the cyclic part of either side.

On one hand, we have $\deg(\nu) \leq \deg(\varphi_{\mathsf{rsp}}) \deg(\varphi'_{\mathsf{rsp}}) \leq 2^{2f}$ by Lemma 3.4.7. On the other hand, Lemma 3.4.9 implies

$$\deg(\psi) = \frac{N_{\mathsf{chl}}^2}{d^2} > \frac{N_{\mathsf{chl}}^2}{2^{2\lambda}} \geq 2^{2f}.$$

This contradicts the equality $\psi = \nu$. □

Remark 3.4.11. Since we have chosen $2^f = \Theta(p^2)$ for reasons that will be explained in Section 3.4.2 and $p = \Theta(2^{2\lambda})$ in order to ensure λ bits of classical hardness for the relation defined in Eq. (3.9), the condition $N_{\mathsf{chl}} \geq 2^{f+\lambda}$ implies that $N_{\mathsf{chl}} = \Omega(p^{5/2})$. This justifies the choice $N_{\mathsf{chl}} = \Theta(p^{5/2})$ proposed in Section 3.3.3.

3.4.2 The zero knowledge property

The proof of the zero-knowledge property of SQIsignHD uses an oracle generating isogenies of non-smooth degree. To our knowledge, there is no efficient algorithm implementing such an oracle. Nonetheless, it is believed that access to such an oracle does not affect the hardness of the relation defined in Eq. (3.9) (see Section 3.4.3). In R-SQIsignHD, the definition of such an oracle is very natural. In F-SQIsignHD, we add (mild) conditions on the degree to account for the computational constraints imposed by the method in dimension 4. These degree constraints are the main reason why the signatures are represented in dimension 8 instead of 4 in R-SQIsignHD. Note that the very efficient[1] isogeny based digital signature scheme PRISM [BBC+25] has been built on the assumption that implementing oracles used in F-SQIsignHD or R-SQIsignHD is a hard problem. Also note that another security paradigm for SQIsignHD has been proposed recently in [ABDPW25] (we shall briefly discuss it in Section 3.4.3).

[1] In PRISM, signature is 1.8 times faster and verification 1.4 times slower than all SQIsign variants.

In F-SQIsignHD

Definition 3.4.12. A *random uniform good degree isogeny oracle* (RUGDIO) is an oracle taking as input integers $f, B \in \mathbb{N}^*$ and a supersingular elliptic curve E defined over $\mathbb{F}_{p^2}$ and returning an efficient representation of a random isogeny $\varphi : E \longrightarrow E'$ of (f, B)-good degree such that:

(i) The distribution of E' is uniform in the supersingular isogeny graph.

(ii) The conditional distribution of φ given E' is uniform among isogenies $E \longrightarrow E'$ of (f, B)-good degree.

In addition to the constraint on the degree of the RUGDIO output, we add constraints on the distributions of isogenies. These constraints are necessary to construct a simulator of F-SQIsignHD. Using the Deuring correspondence, we heuristically justified in Section 3.3.4 that these constraints can be mathematically satisfied, namely that for all supersingular elliptic curves E and E', there exists $\varphi : E \longrightarrow E'$ of (f, B)-good norm. Taking 2^f slightly bigger than $\sqrt{p}$ is heuristically sufficient. Note that to prove the zero-knowledge property, we not only need access to a RUGDIO, but also to make a heuristic assumption on the distribution of the commitment E_{com}. This assumption is not necessary in R-SQIsignHD.

Theorem 3.4.13. *Assume that the commitments E_{com} resulting from a honest execution of F-SQIsignHD are computationally indistinguishable from elliptic curves chosen uniformly at random in the supersingular isogeny graph. Then, the F-SQIsignHD identification protocol is computationally honest verifier zero-knowledge in the RUGDIO model.*

In other words, under this assumption, there exists a random polynomial time simulator $\mathcal{S}$ with access to a RUGDIO that simulates transcripts $(E_{com}^{\mathcal{S}}, \mathsf{chl}_{\mathcal{S}}, \mathsf{rsp}_{\mathcal{S}})$ with a computationally indistinguishable distribution from the transcripts of the F-SQIsignHD identification protocol.

Proof. First, we explain how to construct the simulator $\mathcal{S}$. Given a public key E_{pk}, the simulator starts by generating a challenge by sampling $\mathsf{chl}_{\mathcal{S}} \in [\![0\ ;\ 2^\lambda - 1]\!]$ uniformly at random exactly as in F-SQIsignHD and computes the associated isogeny $\varphi_{\mathsf{chl}}^{\mathcal{S}} : E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}^{\mathcal{S}}$. Then, it applies the RUGDIO on entry (f, B) and $E_{\mathsf{chl}}^{\mathcal{S}}$ to get an efficient representation of a dual response isogeny $\widehat{\varphi}_{\mathsf{rsp}}^{\mathcal{S}} : E_{\mathsf{chl}}^{\mathcal{S}} \longrightarrow E_{\mathsf{com}}^{\mathcal{S}}$. From this efficient representation, we obtain $q_{\mathcal{S}} := \deg(\widehat{\varphi}_{\mathsf{rsp}}^{\mathcal{S}}) = \deg(\varphi_{\mathsf{rsp}}^{\mathcal{S}})$ and the image $(\varphi_{\mathsf{rsp}}^{\mathcal{S}}(P_{\mathsf{com}}^{\mathcal{S}}), \varphi_{\mathsf{rsp}}^{\mathcal{S}}(Q_{\mathsf{com}}^{\mathcal{S}}))$ of a deterministic basis $(P_{\mathsf{com}}^{\mathcal{S}}, Q_{\mathsf{com}}^{\mathcal{S}})$ of $E_{\mathsf{com}}^{\mathcal{S}}[2^e]$ (*e.g.* using Algorithm 2.9 since $q_{\mathcal{S}}$ is odd or Lemma 3.4.5 to obtain an efficient representation of $\varphi_{\mathsf{rsp}}^{\mathcal{S}}$ from an efficient representation of $\widehat{\varphi}_{\mathsf{rsp}}^{\mathcal{S}}$). Hence, we can compute $\mathsf{rsp}_{\mathcal{S}} := ([2^{e-r}]\varphi_{\mathsf{rsp}}^{\mathcal{S}}(P_{\mathsf{com}}^{\mathcal{S}}), [2^{e-r}]\varphi_{\mathsf{rsp}}^{\mathcal{S}}(Q_{\mathsf{com}}^{\mathcal{S}}), q_{\mathcal{S}})$ in polynomial time in $\log(p)$.

We now prove that the transcripts $(E_{\mathsf{com}}^{\mathcal{S}}, \mathsf{chl}_{\mathcal{S}}, \mathsf{rsp}_{\mathcal{S}})$ of $\mathcal{S}$ are computationally indistinguishable from the transcripts $(E_{\mathsf{com}}, \mathsf{chl}, \mathsf{rsp})$ of the F-SQIsignHD identification protocol. By construction, chl and $\mathsf{chl}_{\mathcal{S}}$ have exactly the same distribution. By the definition of the RUGDIO, $E_{\mathsf{com}}^{\mathcal{S}}$ is uniformly random in the supersingular isogeny graph conditionally to $E_{\mathsf{chl}}^{\mathcal{S}}$, so $E_{\mathsf{com}}^{\mathcal{S}}$ is uniformly random in the supersingular isogeny graph. Besides, E_{com} is computationally indistinguishable from a uniformly random supersingular elliptic curve by assumption.

Finally, conditionally to $E_{\mathsf{com}}^{\mathcal{S}}$ and $E_{\mathsf{chl}}^{\mathcal{S}}$, $\widehat{\varphi}_{\mathsf{rsp}}^{\mathcal{S}}$ (represented by $\mathsf{rsp}_{\mathcal{S}}$) is uniformly random among the isogenies $E_{\mathsf{chl}}^{\mathcal{S}} \longrightarrow E_{\mathsf{com}}^{\mathcal{S}}$ of (f, B)-good degree by the definition of the RUGDIO. The dual map $\phi \longmapsto \widehat{\phi}$ being a bijection preserving the degree, conditionally to $E_{\mathsf{com}}^{\mathcal{S}}$ and $E_{\mathsf{chl}}^{\mathcal{S}}$, $\varphi_{\mathsf{rsp}}^{\mathcal{S}}$ is also uniformly random among the isogenies $E_{\mathsf{com}}^{\mathcal{S}} \longrightarrow E_{\mathsf{chl}}^{\mathcal{S}}$ of (f, B)-good degree. By construction (see Sections 3.2.4 and 3.3.4), conditionally to E_{com} and E_{chl}, φ_{rsp} has the same distribution. Hence, since the basis of $E_{\mathsf{com}}[2^e]$ and $E_{\mathsf{com}}^{\mathcal{S}}[2^e]$ we consider are both deterministic, rsp and $\mathsf{rsp}_{\mathcal{S}}$ have the same distributions conditionally to their respective commitment and challenge curves. This completes the proof. □

We finally justify that E_{com} is computationally indistinguishable from a uniformly random supersingular elliptic curve is a reasonable assumption. We have seen in Section 3.3.2 that the challenge generation method outputs isogenies starting from E_0 of degrees close to p and that up to $\Theta(p^2)$ distinct isogenies can be generated in that way. The number of supersingular j-invariants is $\sim p/12$ and they are all linked to E_0 by an isogeny of degree $O(\sqrt{p})$ by Lemma 2.3.2. It follows that the commitment generation method has largely enough entropy to cover the whole supersingular isogeny graph, which weighs in favour of our heuristic assumption.

In R-SQIsignHD

In R-SQIsignHD, the oracle we use to simulate transcripts is much more natural than the RUGDIO used for F-SQIsignHD as it simply outputs uniformly random isogenies of bounded degree without any other condition on the degree or the codomain distribution.

Definition 3.4.14. A *random any degree isogeny oracle* (RADIO) is an oracle taking as input an integer $N \in \mathbb{N}^*$ and a supersingular elliptic curve E defined over $\mathbb{F}_{p^2}$ and returning an efficient representation of an isogeny $\varphi : E \longrightarrow E'$, which is uniformly random among the isogenies of degree $q < N$ with domain E.

Since R-SQIsignHD commitment isogenies have a fixed degree $N_{\mathsf{com}} = \Theta(p^2)$ and are generated uniformly at random by construction (see Algorithm 3.9), Proposition 3.3.1 ensures that the distribution of their codomain E_{com} has statistical distance $O(p^{-1/2})$ to the uniform distribution. Since $p = \Theta(2^{2\lambda})$, this ensures that an attacker with computing power bounded by 2^λ cannot distinguish E_{com} from a uniformly random supersingular elliptic curve. Recall that in the proof of Theorem 3.4.13, we used the fact that codomain outputs of the RUGDIO have a uniform distribution to prove that simulated and real commitments have close distributions. With a choice of bound $2^f = \Theta(p^2)$, we can prove that this is also the case for the RADIO (up to a negligible statistical distance). This is the following result which is of independent interest and that will be proved later.

Theorem 3.4.15. *Let $\varepsilon \in]0, 2]$. Let $E/\mathbb{F}_{p^2}$ be a supersingular elliptic curve and π be the probability distribution of codomains E' (up to $\overline{\mathbb{F}_p}$-isomorphism) of isogenies $\varphi : E \longrightarrow E'$ chosen uniformly at random among isogenies of degree $\deg(\varphi) \leq p^{1+\varepsilon}$. Let U be the uniform distribution in the supersingular isogeny graph over $\mathbb{F}_{p^2}$. Then, the statistical distance between U and π satisfies $d_{TV}(U, \pi) = O(p^{-\varepsilon/2})$.*

Then the desired result easily follows.

Theorem 3.4.16. *Assume that $2^f = \Theta(p^2)$ (*e.g.* $f = 2(e-4)$ as explained in Section 3.3.4). Then, the F-SQIsignHD identification protocol is statistically honest verifier zero-knowledge in the RADIO model.*

In other words, under this assumption, there exists a random polynomial time simulator $\mathcal{S}$ with access to a RADIO that simulates transcripts $(E^{\mathcal{S}}_{com}, \mathsf{chl}_{\mathcal{S}}, \mathsf{rsp}_{\mathcal{S}})$ with a statistically indistinguishable distribution from the transcripts of the R-SQIsignHD identification protocol (at statistical distance $O(p^{-1/2}) = O(2^{-\lambda})$).

Proof. The construction of $\mathcal{S}$ is exactly the same as in the proof of Theorem 3.4.13 but using the RADIO instead of the RUGDIO. Keeping the same notations, in the R-SQIsignHD context we still have to justify that an efficient representation of the dual simulated response $\widehat{\varphi}^{\mathcal{S}}_{\mathsf{rsp}} : E^{\mathcal{S}}_{\mathsf{chl}} \longrightarrow E^{\mathcal{S}}_{\mathsf{com}}$ outputted from the RADIO yields a response

$$\mathsf{rsp}_{\mathcal{S}} = (\varphi^{\mathcal{S}(1)}_{\mathsf{rsp}}, \widehat{\varphi}^{\mathcal{S}(2)}_{\mathsf{rsp}}, {\varphi'}^{\mathcal{S}}_{\mathsf{rsp}}(P'^{\mathcal{S}}_{\mathsf{com}}), {\varphi'}^{\mathcal{S}}_{\mathsf{rsp}}(Q'^{\mathcal{S}}_{\mathsf{com}}), q_{\mathcal{S}}, a) \tag{3.13}$$

in polynomial time in $\log(p)$. First, $q_{\mathcal{S}} = \deg(\widehat{\varphi}^{\mathcal{S}}_{\mathsf{rsp}})$ can easily be obtained from the efficient representation of $\widehat{\varphi}^{\mathcal{S}}_{\mathsf{rsp}}$. We can the factor $q_{\mathcal{S}} = 2^v q'$ with q' odd and consider the factorisation

$$\widehat{\varphi}^{\mathcal{S}}_{\mathsf{rsp}} = [2^a]\widehat{\varphi}^{\mathcal{S}(1)}_{\mathsf{rsp}} \circ \widehat{\varphi'}^{\mathcal{S}}_{\mathsf{rsp}} \circ \widehat{\varphi}^{\mathcal{S}(2)}_{\mathsf{rsp}},$$

with $2a \leq v$, $\widehat{\varphi}^{\mathcal{S}(1)}_{\mathsf{rsp}}$ cyclic of degree 2^{b_1} where $b_1 := \min(v - 2a, e)$, $\widehat{\varphi}^{\mathcal{S}(2)}_{\mathsf{rsp}}$ cyclic of degree 2^{b_2} where $b_2 := v - 2a - b_1$ and $\widehat{\varphi}^{\mathcal{S}(1)}_{\mathsf{rsp}}$ of degree q'. With the efficient representation, we can evaluate $\widehat{\varphi}^{\mathcal{S}}_{\mathsf{rsp}}$ on a basis (P, Q) of $E^{\mathcal{S}}_{\mathsf{chl}}[2^e]$ and compute the maximal $2^{e'}$ of the orders of $\widehat{\varphi}^{\mathcal{S}}_{\mathsf{rsp}}(P)$ and $\widehat{\varphi}^{\mathcal{S}}_{\mathsf{rsp}}(Q)$ to obtain $a = e - e'$. Then, we compute:

$$\ker(\widehat{\varphi}^{\mathcal{S}(2)}_{\mathsf{rsp}}) = \ker(\widehat{\varphi}^{\mathcal{S}(1)}_{\mathsf{rsp}} \circ \widehat{\varphi'}^{\mathcal{S}}_{\mathsf{rsp}} \circ \widehat{\varphi}^{\mathcal{S}(2)}_{\mathsf{rsp}}) \cap E^{\mathcal{S}}_{\mathsf{chl}}[2^{b_2}]$$

by computing a discrete logarithm problem involving $[2^{e-a-b_2}]\widehat{\varphi}^{\mathcal{S}}_{\mathsf{rsp}}(P)$ and $[2^{e-a-b_2}]\widehat{\varphi}^{\mathcal{S}}_{\mathsf{rsp}}(Q)$. Once $\widehat{\varphi}^{\mathcal{S}(2)}_{\mathsf{rsp}}$ has been computed, we obtain an efficient representation of $\varphi^{\mathcal{S}}_{\mathsf{rsp}}$ from an efficient representation of $\widehat{\varphi}^{\mathcal{S}}_{\mathsf{rsp}}$ in polynomial time in $\log(p)$ by Lemma 3.4.5. We can then compute:

$$\ker(\varphi^{\mathcal{S}(1)}_{\mathsf{rsp}}) = \ker(\varphi^{\mathcal{S}(2)}_{\mathsf{rsp}} \circ {\varphi'}^{\mathcal{S}}_{\mathsf{rsp}} \circ \varphi^{\mathcal{S}(1)}_{\mathsf{rsp}}) \cap E^{\mathcal{S}}_{\mathsf{chl}}[2^{b_1}],$$

as we computed $\ker(\widehat{\varphi}_{\mathsf{rsp}}^{\mathcal{S}(2)})$. We finally obtain an efficient representation of ${\varphi'}_{\mathsf{rsp}}^{\mathcal{S}}$ by embedding it into an 8-dimensional isogeny F embedding ${\varphi'}_{\mathsf{rsp}}^{\mathcal{S}} : {E'}_{\mathsf{com}}^{\mathcal{S}} \longrightarrow {E'}_{\mathsf{chl}}^{\mathcal{S}}$ of odd smooth and squarefree polarised degree N, as in the proof of Theorem 2.2.10. In order to compute F, we need to evaluate ${\varphi'}_{\mathsf{rsp}}^{\mathcal{S}}$ on points of ${E'}_{\mathsf{com}}^{\mathcal{S}}[\ell]$ for all prime $\ell | N$. From the knowledge of efficient representations of $\varphi_{\mathsf{rsp}}^{\mathcal{S}}$, $\varphi_{\mathsf{rsp}}^{\mathcal{S}(1)}$ and $\widehat{\varphi}_{\mathsf{rsp}}^{\mathcal{S}(2)}$, we can evaluate ${\varphi'}_{\mathsf{rsp}}^{\mathcal{S}}$ on such points using the tricks from Section 2.3.3 since N is odd. Finally, with an efficient representation of ${\varphi'}_{\mathsf{rsp}}^{\mathcal{S}}$, we can evaluate this isogeny on a basis $({P'}_{\mathsf{com}}^{\mathcal{S}}, {Q'}_{\mathsf{com}}^{\mathcal{S}})$ of ${E'}_{\mathsf{com}}^{\mathcal{S}}[2^e]$. We finally obtain a simulated response in the shape of Eq. (3.13) in polynomial time in $\log(p)$.

We finally justify that real R-SQIsignHD transcripts $(E_{\mathsf{com}}, \mathsf{chl}, \mathsf{rsp})$ have the same distribution as simulated transcripts $(E_{\mathsf{com}}^{\mathcal{S}}, \mathsf{chl}_{\mathcal{S}}, \mathsf{rsp}_{\mathcal{S}})$. By construction, chl and $\mathsf{chl}_{\mathcal{S}}$ have exactly the same distribution so E_{chl} and $E_{\mathsf{chl}}^{\mathcal{S}}$ have the same distribution. By construction of the RADIO and Algorithm 3.12, conditionally to their domains and codomains, both φ_{rsp} and $\varphi_{\mathsf{rsp}}^{\mathcal{S}}$ have uniform distribution among isogenies of degree $\leq 2^f$. Hence, rsp and $\mathsf{rsp}_{\mathcal{S}}$ share the same conditional distributions conditionally to the commitment and challenge. Besides, since $2^f = \Theta(p^2)$, by Theorem 3.4.15, $E_{\mathsf{com}}^{\mathcal{S}}$ is at statistical distance $O(p^{-1/2})$ from the uniform distribution. Since $N_{\mathsf{com}} = \Theta(p^2)$, by Proposition 3.3.1, E_{com} is also at statistical distance $O(p^{-1/2})$ to the uniform distribution. It follows that transcripts from R-SQIsignHD and $\mathcal{S}$ are at statistical distance $O(p^{-1/2})$. This completes the proof. □

On the codomain distribution of random isogenies with bounded degree (Theorem 3.4.15)

The goal of this section is to prove a bound on the statistical distance between codomains of random isogenies with bounded degrees and the uniform distribution on the supersingular isogeny graph. Similar results have been proved on fixed degree smooth isogeny walks [GPS20, Theorem 1] and non-bactracking ℓ-isogeny walks [BCC+23, Theorem 11]. We generalize these results to the case of non-fixed degree. Heuristically, we should expect lower degree bounds than in the fixed degree case to be as close to the uniform distribution, but this is not the case. In particular, as for non-bactracking ℓ-isogeny walks, we need to allow isogenies of degree $p^{1+\varepsilon}$ to reach a statistical distance of $O(p^{-\varepsilon/2})$ to the uniform distribution. However, we provide an elementary proof that does not require to study adjacency matrices of the supersingular isogeny graph and modular forms (unlike [BCC+23]). The main ingredients are the Deuring correspondence and a count of small quaternion ideal vectors (Corollary 3.4.19). We start by proving a classical bound on the last minimum of quaternion ideals claimed in [KLPT14, § 3.1] but never proved so far.

Lemma 3.4.17. *Let $\mathcal{O} \subset \mathcal{B}_{p,\infty}$ be a quaternion order and I be a left ideal of $\mathcal{O}$. Let $(\alpha_1, \cdots, \alpha_4)$ be a Minkowski reduced basis of I for the quadratic form $q_I : \alpha \in I \longmapsto \mathrm{nrd}(\alpha)/\mathrm{nrd}(I)$, so that $q_I(\alpha_i) \leq q_I(\alpha_{i+1})$ for all $i \in \{1, 2, 3\}$. Then*

$$q_I(\alpha_4) \leq \frac{8p}{\pi^2}.$$

Proof. As we saw in the proof of Lemma 2.3.2, we have by Minkowski's second theorem (Eq. (2.8)):

$$\prod_{i=1}^{4} q_I(\alpha_i) \leq \frac{64p^2}{\pi^4}.$$

This inequality is not sufficient to conclude (we only get $q_I(\alpha_4) = O(p^2)$ instead of $O(p)$). To complete the proof, we follow [BST+17, Theorem 3.1].

Let $(\beta_1, \cdots, \beta_4)$ be a Minkowski reduced basis of $\mathcal{O}$. As the α_i, the β_i satisfy

$$\prod_{i=1}^{4} \mathrm{nrd}(\beta_i) \leq \frac{64p^2}{\pi^4}.$$

Let $A := (a_{i,j})_{1 \leq i,j \leq 4} \in M_4(\mathbb{Z})$, where for all $1 \leq i, j \leq 4$, $a_{i,j}$ is the coefficient of α_4 in the decomposition of $\beta_i \alpha_j$ in the basis $(\alpha_1, \cdots, \alpha_4)$ (this is an integer since $\mathcal{O}I \subseteq I$). Then A is invertible. Indeed, if $x \in \mathbb{Z}^4$ satisfies $Ax = 0$ *i.e.* $\sum_{j=1}^{4} a_{i,j} x_j = 0$ for all $i \in [\![1\ ;\ 4]\!]$, then $\alpha := \sum_{j=1}^{4} x_j \alpha_j$ satisfies $\mathcal{O}\alpha \subseteq \langle \alpha_1, \alpha_2, \alpha_3 \rangle$. But $\mathcal{O}\alpha$ has rank 4 whenever $\alpha \neq 0$, so $\alpha = 0$ and $x = 0$.

A being invertible, there exists a permutation $\sigma \in \mathfrak{S}_4$ such that $a_{i,\sigma(i)} \neq 0$ for all $i \in [\![1\ ;\ 4]\!]$. It follows that for all $i \in [\![1\ ;\ 4]\!]$, $\beta_i\alpha_{\sigma(i)}$ completes $\langle \alpha_1, \alpha_2, \alpha_3\rangle$ as a full-rank sublattice of I, so that $\mathrm{nrd}(\alpha_4) \leq \mathrm{nrd}(\beta_i\alpha_{\sigma(i)})$ *i.e.* $q_I(\alpha_4) \leq \mathrm{nrd}(\beta_i)q_I(\alpha_{\sigma(i)})$, since $(\alpha_1, \cdots, \alpha_4)$ is Minkowski reduced. It follows that

$$q_I(\alpha_4)^4 \leq \prod_{i=1}^{4}(\mathrm{nrd}(\beta_i)q_I(\alpha_{\sigma(i)})) = \prod_{i=1}^{4}\mathrm{nrd}(\beta_i)\prod_{i=1}^{4}q_I(\alpha_i) \leq \left(\frac{64p^2}{\pi^4}\right)^2.$$

This completes the proof. □

Now, we introduce a generalization of [Wes22, Lemma 3.2] in every dimension, counting the elements of bounded norm in a lattice.

Lemma 3.4.18. *Let $\Lambda \subseteq \mathbb{R}^d$ be a full-rank lattice of last minimum λ_d and $\rho > \sqrt{d}/2\lambda_d$. Then*

$$\frac{\pi^{d/2}\left(\rho - \frac{\sqrt{d}\lambda_d}{2}\right)^d}{\Gamma\left(\frac{d}{2}+1\right)\mathrm{Covol}(\Lambda)} \leq \#\Lambda \cap B(0,\rho) \leq \frac{\pi^{d/2}\left(\rho + \frac{\sqrt{d}\lambda_d}{2}\right)^d}{\Gamma\left(\frac{d}{2}+1\right)\mathrm{Covol}(\Lambda)},$$

where $B(0,\rho)$ is the ball of center 0 and radius ρ for the Euclidean norm and Γ is Euler's gamma function.

Proof. Let $\mathcal{V} := \{v \in \mathbb{R}^d \mid \|v\| = \min_{\lambda\in\Lambda}\|v+\lambda\|\}$ be the Voronoi cell at the origin of Λ and $\mu := \sup_{v\in\mathcal{V}}\|v\|$ be the covering radius of Λ. Then, we have

$$B(0,\rho-\mu) \subseteq \bigsqcup_{\lambda\in\Lambda\cap B(0,\rho)}(\lambda+\mathcal{V}) \subseteq B(0,\rho+\mu),$$

so that

$$\mathrm{Vol}(B(0,\rho-\mu)) \leq (\#\Lambda\cap B(0,\rho))\cdot\mathrm{Vol}(\mathcal{V}) \leq \mathrm{Vol}(B(0,\rho+\mu)).$$

Since $\mathrm{Vol}(\mathcal{V}) = \mathrm{Covol}(\Lambda)$, $\mathrm{Vol}(B(0,\rho\pm\mu)) = \pi^{d/2}(\rho\pm\mu)^d/\Gamma(d/2+1)$ and $\mu \leq \sqrt{d}\lambda_d/2$ by [MG02, Theorem 7.9], the result follows. □

Corollary 3.4.19. *Let $\mathcal{O} \subset \mathcal{B}_{p,\infty}$ be a maximal order and I be an integral left $\mathcal{O}$-ideal. Then, for all $\varepsilon > 0$ the number of ideals of norm $\leq p^{1+\varepsilon}$ that are right-equivalent to I is*

$$N_{p^{1+\varepsilon}}([I]) := \#\{J \sim I \mid \mathrm{nrd}(J) \leq p^{1+\varepsilon}\} = \frac{2\pi^2}{\#\mathcal{O}_R(I)^\times}p^{1+2\varepsilon}(1+O(p^{-\varepsilon/2})).$$

Proof. By Lemma 1.2.19.(i), an ideal J is right-equivalent to I if and only if it is of the form $J := I\overline{\alpha}/\mathrm{nrd}(I)$ for some $\alpha \in I$. Furthermore, α is uniquely determined by J up to multiplication on the right by an element of $\mathcal{O}_R(I)^\times$ and we have $\mathrm{nrd}(J) = \mathrm{nrd}(\alpha)/\mathrm{nrd}(I) = q_I(\alpha)$. It follows that

$$N_{p^{1+\varepsilon}}([I]) := \#\{J \sim I \mid \mathrm{nrd}(J) \leq p^{1+\varepsilon}\} = \frac{1}{\#\mathcal{O}_R(I)^\times}\#\{\alpha \in I \mid q_I(\alpha) \leq p^{1+\varepsilon}\}.$$

Let $\iota : \mathcal{B}_{p,\infty} \hookrightarrow \mathbb{R}^4$ be the canonical embedding defined by Eq. (1.2), such that $\|\iota(\alpha)\|^2 = \mathrm{nrd}(\alpha)$ for all $\alpha \in \mathcal{B}_{p,\infty}$, where $\|.\|$ is the Euclidean norm on $\mathbb{R}^4$. Consider the lattice $\Lambda := \iota(I)$. We then have

$$N_{p^{1+\varepsilon}}([I]) = \frac{1}{\#\mathcal{O}_R(I)^\times}\#\Lambda\cap B\left(0, p^{(1+\varepsilon)/2}\sqrt{\mathrm{nrd}(I)}\right)$$

By Lemmas 3.4.17 and 3.4.18, we get

$$\#\Lambda\cap B\left(0, p^{(1+\varepsilon)/2}\sqrt{\mathrm{nrd}(I)}\right) \leq \frac{\pi^2\left(p^{(1+\varepsilon)/2}\sqrt{\mathrm{nrd}(I)} + \frac{2\sqrt{2p\,\mathrm{nrd}(I)}}{\pi}\right)^4}{2\,\mathrm{Covol}(\Lambda)},$$

with $\mathrm{Covol}(\Lambda) = p/4\,\mathrm{nrd}(I)^2$ by Eq. (2.9). It follows, that the right term of the inequality is $2\pi^2p^{1+2\varepsilon}(1+O(p^{-\varepsilon/2}))$. Applying the lower bound of Lemma 3.4.18, we also get that

$$\#\Lambda\cap B\left(0, p^{(1+\varepsilon)/2}\sqrt{\mathrm{nrd}(I)}\right) \geq 2\pi^2p^{1+2\varepsilon}(1+O(p^{-\varepsilon/2})).$$

The result follows. □

We denote by $\mathrm{SS}(p)$ the set of supersingular elliptic curves over $\mathbb{F}_{p^2}$ (up to $\overline{\mathbb{F}}_p$-isomorphism) and S the probability distribution on $\mathrm{SS}(p)$ given by $S(E) := 1/(K\#\operatorname{Aut}(E))$ for all $E \in \mathrm{SS}(p)$, with $K := \sum_{E\in\mathrm{SS}(p)} 1/\#\operatorname{Aut}(E) = (p-1)/24$ by Eichler mass formula [Voi21, Theorem 25.1.1]. Let U be the uniform distribution on $\mathrm{SS}(p)$. Recall that by Eq. (3.4), the statistical distance between S and U is $d_{TV}(S,U) = O(p^{-1})$. We can now finally prove our main result.

Theorem 3.4.15. Let $\varepsilon \in]0,2]$. Let $E/\mathbb{F}_{p^2}$ be a supersingular elliptic curve and π be the probability distribution of codomains E' (up to $\overline{\mathbb{F}}_p$-isomorphism) of isogenies $\varphi : E \longrightarrow E'$ chosen uniformly at random among isogenies of degree $\deg(\varphi) \leq p^{1+\varepsilon}$. Let U be the uniform distribution in the supersingular isogeny graph over $\mathbb{F}_{p^2}$. Then, the statistical distance between U and π satisfies $d_{TV}(U,\pi) = O(p^{-\varepsilon/2})$.

Proof. By the Deuring correspondence, it suffices to prove that given a maximal order $\mathcal{O} \subset \mathcal{B}_{p,\infty}$ and $\mathrm{Cl}(\mathcal{O})$ the set of right-equivalence classes of left-ideals of $\mathcal{O}$, the distribution π' of the ideal classes $[I] \in \mathrm{Cl}(\mathcal{O})$ when I is sampled uniformly at random among ideals of norm $\leq p^{1+\varepsilon}$ (which is the quaternion analogue of π) is at statistical distance $O(p^{-\varepsilon/2})$ from the uniform distribution U' on $\mathrm{Cl}(\mathcal{O})$. We also denote by S' the quaternion analogue of S, namely the distribution on $\mathrm{Cl}(\mathcal{O})$ given by $S([I]) := 1/(K\#O_R(I)^\times)$ for all $[I] \in \mathrm{Cl}(\mathcal{O})$, where $K := \sum_{[I]\in\mathrm{Cl}(\mathcal{O})} 1/\#O_R(I)^\times = (p-1)/24$ is the Eichler mass. By Eq. (3.4), $d_{TV}(U',S') = O(p^{-1})$, so it suffices to prove that $d_{TV}(S',\pi') = O(p^{-\varepsilon/2})$.

By Corollary 3.4.19, the number of left $\mathcal{O}$-ideals of norm $\leq p^{1+\varepsilon}$ is

$$N_{p^{1+\varepsilon}} = \sum_{[I]\in\mathrm{Cl}(\mathcal{O})} N_{p^{1+\varepsilon}}([I]) = 2\pi^2 K p^{1+2\varepsilon}(1+O(p^{-\varepsilon/2})),$$

so the distribution π' is given by

$$\forall [I] \in \mathrm{Cl}(\mathcal{O}), \quad \pi'([I]) = \frac{N_{p^{1+\varepsilon}}([I])}{N_{p^{1+\varepsilon}}} = \frac{1}{K}(1+O(p^{-\varepsilon/2})).$$

It follows immediately that

$$d_{TV}(S',\pi') = \frac{1}{2}\sum_{[I]\in\mathrm{Cl}(\mathcal{O})} |S'([I]) - \pi'([I])| = \frac{\#Cl(\mathcal{O})}{K} O(p^{-\varepsilon/2}) = O(p^{-\varepsilon/2}).$$

The result follows. □

3.4.3 On isogeny generation oracles

On hardness of the supersingular endomorphism problem with access to an auxiliary oracle

The F-SQIsignHD identification protocol is secure assuming the hardness of the supersingular endomorphism problem (Problems 2.1.2 and 3.4.3), and zero-knowledge with respect to a simulator that has access to a RUGDIO, as defined in Definition 3.4.12 (or a RADIO for R-SQIsignHD, as defined in Definition 3.4.14). For the resulting signature scheme to be secure, one therefore needs to assume that the supersingular endomorphism problem remains hard even when given access to a RUGDIO.

While it currently seems out of reach to prove that the supersingular endomorphism problem is equivalent to the variant with RUGDIO access, let us argue that the RUGDIO indeed does not help. We focus the following discussion on the RUGDIO, but the same arguments apply to the RADIO despite the slightly different distribution.

The RUGDIO allows to generate random isogenies with a chosen domain E. Note that this task is already known to be easy, with isogenies of smooth degree. The RUGDIO only lifts this smoothness restriction and replaces it with other restrictions ((f,B)-good): it allows to generate random isogenies whose degrees have large prime factors. It does not allow to reach more target curves, nor does it give more control on which specific target to hit: if f is big enough, the target curve is uniformly distributed in the supersingular graph (Theorem 3.4.15), which was already possible with smooth degree isogenies.

Smoothness of random isogenies has never been an inconvenience in finding endomorphisms. In fact, the current fastest algorithms for this problem only require very smooth degree isogenies, typically a power of 2. The reason is the following: the purpose of constructing a random isogeny from a fixed source is to reach a random target. As very smooth isogenies (even 2-smooth) are sufficient for optimal randomisation, there is no incentive to involve much larger prime factors. More specifically, the best known strategies to solve the supersingular endomorphism ring problem [DG16; EHLMP20; FIKMN25] have classical time complexity $\widetilde{O}(\sqrt{p})$ (and quantum time complexity $\widetilde{O}(p^{1/4})$ with a Grover argument [Gro96; BJS14]) and essentially perform a meet-in-the-middle search in the supersingular isogeny graph. Access to a RUGDIO would allow to use isogenies of a different shape in the search, but would not speed it up, as the probability to find isogenies with matching codomains stays the same. Another illustration that having access to non-smooth degree isogenies does not help is the fact that the discovery of the $\sqrt{\text{élu}}$ algorithm [BDFLS20] (which dramatically improved the complexity of computing prime degree isogenies) did not affect the state-of-the-art of the supersingular endomorphism problem.

The above arguments support that random isogenies of non-smooth degrees are not more helpful than random isogenies of smooth degrees. Now, one may be concerned that the encoding of the output of the RUGDIO may leak more information than it should. Non-smooth degree isogenies are represented as a component of a higher dimensional isogeny. This representation is universal, in the sense that any efficient representation of an isogeny can be efficiently rewritten in this form. In particular, this encoding contains no more information than any other efficient representation of the same isogeny.

Another security framework for the higher dimensional variants of SQIsign

In both variants of SQisignHD, the security proof is done in a non-standard model where the simulator has access to an auxiliary oracle, namely a RUGDIO or a RADIO. This is also the case for SQIsign2D-West (see Section 4.2.2) and other higher dimensional variants [NOC+25; DF25]. Unfortunately, such isogeny oracles cannot be constructed in practice and constructing such oracles is even believed to be a hard problem that can be used for cryptographic applications like PRISM [BBC+25]. The need for non-standard security models with oracles indicate that the usual Fiat-Shamir transform is not the appropriate framework to prove the security of SQIsign higher dimensional variants.

In [ABDPW25], the *Fiat-Shamir with hints* framework has been introduced to bypass completely the use of such oracles and prove the universal unforgeability under chosen message attacks (EUF-CMA) of SQIsign2D-West purely in the random oracle model under the hardness of the supersingular endomorphism ring with hints. The argument is easily adaptable to other SQIsign higher dimensional variants, including SQIsignHD. In this new framework, the oracle is replaced by *hints* given to the zero-knowledge simulator following the same distribution as the former oracle, hence the same distribution as response isogenies generated with honest protocol executions. The authors of [ABDPW25] reduce the EUF-CMA security of SQIsign2D-West to:

- The difficulty of the supersingular one endomorphism problem *with hints* when the hint distribution can be pushed through 2^n-isogenies from one supersingular elliptic curve to another.

- A hint distribution indistinguishability assumption between hints used by the simulator and pushable hints through 2^n-isogenies.

[ABDPW25] also reduce the supersingular endomorphism problem with hints to the supersingular one endomorphism problem with hints whose hardness has been discussed in the previous paragraph (where hints replace the oracle).

The Fiat-Shamir with hints framework to prove EUF-CMA security of SQIsign variants is a strong contribution enhancing the confidence in their security. We chose not to present the security proof of SQIsignHD and SQIsign2D-West in this framework as this contribution was not ours and simultaneous to the writing of this PhD thesis.

3.5 Instantiation of the SQIsignHD signature scheme

In this section, we finally discuss the performance of the F-SQIsignHD digital signature scheme (referred to as SQIsignHD by abuse) obtained via the Fiat-Shamir transform. The performance of R-SQIsignHD is not discussed as it is not implemented.

3.5.1 Parameter choices and compression techniques

As explained previously, we choose $p = \Theta(2^{2\lambda})$ in order to offer λ bits of classical and $\lambda/2$ bits of quantum soundness security which relies on the hardness of the endomorphism ring problem (Problems 2.1.2 and 3.4.3) with best known classical attacks in $\widetilde{O}(\sqrt{p})$ and best quantum attacks in $\widetilde{O}(p^{1/4})$. The prime parameter p for different security levels can be found in Table 3.1. In order to simplify the implementation, p was chosen to fit in respectively 4, 6 and 8 words of 64-bits for NIST levels I, III and V. Note that this choice of primes is exactly the same in SQIsign2D-West that will be presented in Chapter 4 but far different from the one proposed in the original version of SQIsignHD [DLRW24, § 3.1] (where $p = 13 \cdot 2^{126} \cdot 3^{78} - 1$ for NIST-I level).

As explained in Sections 3.1.3 and 3.3.4, a signature in SQIsignHD is of the form $(E_{\mathsf{com}}, \mathsf{rsp})$ with $\mathsf{rsp} = ([2^{e-r}]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}), [2^{e-r}]\varphi_{\mathsf{rsp}}(Q_{\mathsf{com}}), q)$ and $(P_{\mathsf{com}}, Q_{\mathsf{com}})$ a basis of $E_{\mathsf{com}}[2^e]$. This data can be represented and compressed as follows. The curve E_{com} can be represented as a scalar in $\mathbb{F}_{p^2}$ *i.e.* as two elements of $\mathbb{F}_p$ which is in bijection with $[\![0\ ;\ p-1]\!]$. For instance, working with a Montgomery equation $y^2 = x^3 + A_{\mathsf{com}}x^2 + x$, E_{com} is determined by its Montgomery coefficient $A_{\mathsf{com}} \in \mathbb{F}_{p^2}$. Since p has 2λ bits, E_{com} can be represented with 4λ bits. We also know that $q < 2^f$ with $2^f = \Omega(\sqrt{p}\log(p))$ for reasons explained in Section 3.3.4, so we may choose $f = \lambda + \lceil \log_2(2\lambda) \rceil$. Hence q may be represented with $\lambda + \lceil \log_2(2\lambda) \rceil$ bits.

Finally, $([2^{e-r}]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}), [2^{e-r}]\varphi_{\mathsf{rsp}}(Q_{\mathsf{com}}))$ can be represented by 3 integers of r bits. Indeed, we can always generate a deterministic basis $(P_{\mathsf{chl}}, Q_{\mathsf{chl}})$ of $E_{\mathsf{chl}}[2^e]$ and express

$$[2^{e-r}]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}) = [a2^{e-r}]P_{\mathsf{chl}} + [b2^{e-r}]Q_{\mathsf{chl}}, \quad \text{and} \quad [2^{e-r}]\varphi_{\mathsf{rsp}}(Q_{\mathsf{com}}) = [c2^{e-r}]P_{\mathsf{chl}} + [d2^{e-r}]Q_{\mathsf{chl}},$$

with $a, b, c, d \in [\![0\ ;\ 2^r - 1]\!]$. Each one of the integers a, b, c, d can be written with λ bits. Actually, only three of them are necessary. Indeed, we have on the one hand:

$$e_{2^r}([2^{e-r}]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}), [2^{e-r}]\varphi_{\mathsf{rsp}}(Q_{\mathsf{com}})) = e_{2^r}([2^{e-r}]P_{\mathsf{com}}, [2^{e-r}]Q_{\mathsf{com}})^q,$$

and on the other hand:

$$e_{2^r}([2^{e-r}]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}), [2^{e-r}]\varphi_{\mathsf{rsp}}(Q_{\mathsf{com}})) = e_{2^r}([2^{e-r}]P_{\mathsf{chl}}, [2^{e-r}]Q_{\mathsf{chl}})^{ad-bc}.$$

Since both $([2^{e-r}]P_{\mathsf{chl}}, [2^{e-r}]Q_{\mathsf{chl}})$ and $([2^{e-r}]P_{\mathsf{com}}, [2^{e-r}]Q_{\mathsf{com}})$ generate the 2^r-torsion of their respective curves, their Weil pairings are primitive 2^r-th roots of unity and there exists $k \in [\![0\ ;\ 2^r - 1]\!]$ odd such that $e_{2^r}([2^{e-r}]P_{\mathsf{com}}, [2^{e-r}]Q_{\mathsf{com}}) = e_{2^r}([2^{e-r}]P_{\mathsf{chl}}, [2^{e-r}]Q_{\mathsf{chl}})^k$, so that

$$ad - bc \equiv kq \mod 2^r.$$

If a is odd, then a, b and c determine d. Otherwise, b must be odd (since φ_{rsp} has odd degree, it preserves the order of 2^r-torsion points), and a, b and d determine c. As a consequence, $([2^{e-r}]\varphi_{\mathsf{rsp}}(P_{\mathsf{com}}), [2^{e-r}]\varphi_{\mathsf{rsp}}(Q_{\mathsf{com}}))$ can be represented with $3r$ bits and the whole signature can be represented with $4\lambda + f + 3r$ bits. Since $f = \lambda + \lceil \log_2(2\lambda) \rceil$ and $r = \lceil f/2 \rceil + 2$, the signature size is $13/2\lambda + 5/2\log_2(2\lambda) + O(1)$ bits. Note that the Weil pairing and discrete logarithm computation must take place during the verification, which is already quite costly. In order to save some verification time, it would be possible to use signatures of $4\lambda + f + 4r = 7\lambda + 3\log_2(2\lambda) + O(1)$ bits instead.

As the commitment, the public key E_{pk} is simply represented by its Montgomery coefficient $A_{\mathsf{pk}} \in \mathbb{F}_{p^2}$ which can be stored in 4λ bits. The signature and public key sizes can be found in Table 3.1 along with prime choices. We refer to Table 3.2 for a comparison with SQIsign original signature sizes. We find that F-SQIsignHD signatures are 40% more compact than SQIsign ones.

Security level λ	NIST-I 128	NIST-III 192	NIST-V 256
$p = c \cdot 2^e - 1$	$5 \cdot 2^{248} - 1$	$65 \cdot 2^{376} - 1$	$27 \cdot 2^{500} - 1$
Signature size (bytes*)	108	160	212
Public key size (bytes*)	64	96	128

Table 3.1: Prime parameter, signature and public key sizes for different security levels in F-SQIsignHD. *One byte contains 8 bits.

Security level	NIST-I	NIST-III	NIST-V
SQIsign (bytes)	177	263	335
F-SQIsignHD (bytes)	108	160	212

Table 3.2: Comparison of signature sizes in SQIsign and F-SQisignHD.

3.5.2 Performance

F-SQIsignHD key generation and signature have been implemented in C mostly by Antonin Leroux as part of the SQIsign2D-West code. The code repository may be found at https://github.com/SQISign/sqisign2d-west-ac24. Timings are displayed in Table 3.3. At NIST-I security level, signing takes around 9 ms and key generation takes around 15 ms on an Intel Core i5-1335U 4600MHz CPU. This compares favourably to the former implementation of SQIsignHD that did not rely on SQIsign2D-West new algorithms (especially Algorithm 2.3) where signing took 28 ms and key generation 70 ms at security level NIST-I on the same CPU[2].

Most importantly, in both implementations of SQIsignHD, signing is considerably faster than in SQIsign. Comparing with NIST first implementation of SQIsign [CSSD+23], F-SQIsignHD last implementation signing is respectively 63.5 and 728.5 times faster at NIST-I and NIST-V security levels (see Table 3.3). The gap between SQIsign and SQIsignHD grows with higher security parameters, highlighting the scalability issue of SQIsign that SQIsignHD solves completely. As we shall see, SQIsignHD signing time is also faster than SQIsign2D-West signing time by a factor 6.7 at NIST-I level (see Section 4.3.2).

Security level			NIST-I	NIST-III	NIST-V
SQIsign NIST v 1.0	Key generation	ms	355.72	5 625.72	22 445.3
		10^6 cycles	889.29	14 064.29	56 113.25
	Signing	ms	554.78	10 553.18	41 322.21
		10^6 cycles	1 386.95	26 382.94	103 305.53
	Verification	ms	7.77	195.86	571.77
		10^6 cycles	19.43	489.65	1 429.42
F-SQIsignHD	Key generation	ms	14.85	48.5	112.31
		10^6 cycles	37.12	121.29	280.92
	Signing	ms	**8.74**	**25.68**	**56.72**
		10^6 cycles	21.83	64.24	141.86

Table 3.3: Comparison of timings of F-SQIsignHD and SQIsign original NIST submission on an Intel Core i5-1335U 4600MHz CPU. F-SQIsignHD verification has not been implemented in C so verification times were not displayed in this table.

However, this spectacular signing and key generation performance is realised at the expense of the verification which requires a 4-dimensional isogeny computation. For now, the verification is

[2]This implementation can be found at https://github.com/Pierrick-Dartois/SQISignHD-lib.

implemented in SageMath and imported as a submodule of the old SQIsignHD implementation https://github.com/Pierrick-Dartois/SQISignHD-lib. For NIST-I security level, the old SQIsignHD verification takes around 600 ms on an Intel Core i5-1335U 4600MHz CPU. Currently, the time spent on verification is as follows: around 60 ms for the challenge computation[3], 510 ms for the 4-dimensional isogeny and 30 ms for the image of a point through F. A C implementation of 4-dimensional isogenies is ongoing and is expected to considerably accelerate this verification phase. Besides, SQIsign2D-West has been proposed to accelerate the verification with a 2-dimensional instead of a 4-dimensional isogeny computation while keeping competitive signing time.

[3]Which consisted in computing a chain of 3-isogenies in dimension 1 in the old SQIsignHD version, which is more costly than a chain of 2-isogenies of similar degree, which is used in F-SQIsignHD.

Chapter 4

SQIsign2D-West: faster verification with 2-dimensional isogenies

In this chapter, we reach one of the main goals of this PhD. We introduce SQIsign2D-West [BDF+25], a variant of SQIsign using only 2-dimensional isogenies to translate quaternion ideals into isogenies between supersingular elliptic curves. Using Algorithm 2.7 (introduced for SQIsign2D-West in the first place), we are able to completely bypass the use of 4-dimensional isogenies in the verification phase that was necessary in F-SQIsignHD. With this optimisation, we obtain a verification that is even faster than the original version of SQIsign while achieving security properties almost as rigorous as the purely theoretical R-SQIsignHD. The signing time is slower than in SQIsignHD but remains very competitive compared to SQIsign.

Note that SQIsign2D-West is not the only 2-dimensional variant of SQIsign. SQIsign2D-East [NOC+25] and SQIPrime [DF25] have been introduced at the same time as SQIsign2D-West. Though probably more efficient than SQIsign2D-West[1], their security proof relies on heuristic or *ad-hoc* assumptions. There is also a heuristic version of SQIsign2D-West which is comparable (and close to SQIsign2D-East in particular).

For these reasons, SQIsign2D-West is the optimal compromise that has been proposed as a reference to replace the original version of SQIsign at round 2 of the NIST competition [AAA+25]. The presentation in this chapter follows the original SQIsign2D-West paper [BDF+25] and also relies on some already introduced algorithmic ideas from NIST round 2 specification [AAA+25] (especially Algorithm 2.7 improving Algorithm 2.5).

4.1 The SQIsign2D-West identification protocol

As SQIsignHD, SQIsign2D-West is a variant of SQIsign using different methods for effective Deuring correspondence. In this section, we describe the identification protocol underlying the SQIsign2D-West digital signature scheme obtained via the Fiat-Shamir transform. We mostly present the main version of SQIsign2D-West as described in [BDF+25], which inspired the NIST round 2 SQIsign submission [AAA+25]. There is also a faster version with a heuristic security proof introduced in [BDF+25, Appendix B] and referred to as H-SQIsign2D-West, that we will not present in detail.

4.1.1 Setting and algorithmic building blocks

As in SQIsignHD, we are given the following public parameters:

- A prime of the form $p = c2^e - 1$ with $c \in \mathbb{N}^*$ odd and small, and $e \simeq 2\lambda$ to grant a classical security level of λ bits.
- The supersingular elliptic curve E_0 of equation $y^2 = x^3 + x$ defined over $\mathbb{F}_{p^2}$ with an explicit isomorphism $\varepsilon_0 : \mathcal{O}_0 \xrightarrow{\sim} \mathrm{End}(E_0)$ between a maximal order of $\mathcal{B}_{p,\infty}$ and the endomorphism ring $\mathrm{End}(E_0)$ (by Lemma 1.2.25).

[1] A low level implementation of SQIsign2D-East and SQIPrime is lacking to certify it.

- A basis (P_0, Q_0) of $E_0[2^e]$.
- A basis $(\beta_1, \cdots, \beta_4)$ of $\mathcal{O}_0$ (*e.g.* $(1, i, (i+j)/2, (1+ij)/2)$) and the image of (P_0, Q_0) by this basis $(\varepsilon_0(\beta_i)(P_0), \varepsilon_0(\beta_i)(Q_0))_{1 \leq i \leq 4}$.
- Some precomputed data to apply Algorithm 2.7 translating left $\mathcal{O}_0$-ideals into isogenies (as described in the last paragraph of Section 2.4.3).

We denote by pp these public parameters along with others that will be introduced later.

We shall use the following algorithms. Even though they were presented in Chapter 3 to construct SQIsignHD most of them were introduced specifically for SQIsign2D-West in the first place (which was designed afterwards).

- Algorithm 2.3 taking as input an odd integer $u < 2^e$ such that $u(2^e - u) = \Omega(p \log(p))$ and returning $(E_u, \varphi_u(P_0), \varphi_u(Q_0), I_u)$, where $\varphi_u : E_0 \longrightarrow E_u$ is a u-isogeny and $I_u \subset \mathcal{O}_0$ is its associated ideal.
- Algorithm 2.7 taking as input a left $\mathcal{O}_0$-ideal I (and some precomputed data) and returning the image $(E_I, \varphi_I(P_0), \varphi_I(Q_0))$ of the isogeny $\varphi_I : E_0 \longrightarrow E_I$ associated to I.
- Algorithm 3.1 taking as input a supersingular elliptic curve $E/\mathbb{F}_{p^2}$ of known endomorphism ring $\mathcal{O} \simeq \mathrm{End}(E)$ and a primitive $\mathcal{O}$-ideal I (in the sense of Definition 2.1.3) of smooth norm D and returning a kernel generator $P \in E[I]$. This direct ideal to kernel translation algorithm is efficient only when D is smooth and $E[D]$ is defined over a small extension of $\mathbb{F}_{p^2}$.
- Algorithm 3.3 to compute the kernel ideal of a cyclic isogeny $\varphi : E_1 \longrightarrow E_2$ of smooth degree D given a kernel generator $P \in E_1[D]$, a basis $(\beta_1, \cdots, \beta_4)$ of $\mathcal{O}_0$, a basis (R_0, S_0) of $E_0[D]$ and the images $(\varepsilon_0(\beta_i)(R_0), \varepsilon_0(\beta_i)(S_0))_{1 \leq i \leq 4}$, and the image $(\psi(R_0), \psi(S_0))$ of an N-isogeny $\psi : E_0 \longrightarrow E_1$ with $\gcd(N, D) = 1$.
- Algorithm 3.4 taking as input a maximal order $\mathcal{O} \subset \mathcal{B}_{p,\infty}$ and an integer N coprime with p and returning a primitive left $\mathcal{O}$-ideal of norm N sampled uniformly at random.
- Algorithm 3.6 sampling a uniform non-zero element in the intersection of a ball and a lattice of rank 4. We use it to sample at random quaternion ideals of small norm equivalent to a given ideal.

4.1.2 Key generation and commitment

The key generation phase in SQIsign2D-West is identical to F-SQIsignHD key generation summarised in Algorithm 3.7. We fix a big prime number $N_{\mathsf{sk}} = \Theta(p^2) = \Theta(2^{4\lambda})$, where λ is the security level. During the key generation, the prover samples uniformly at random a left $\mathcal{O}_0$-ideal I_{sk} of norm N_{sk} with Algorithm 3.4. Then they call Algorithm 2.7 to compute the image $(E_{\mathsf{pk}}, \varphi_{\mathsf{sk}}(P_0), \varphi_{\mathsf{sk}}(Q_0))$ of the isogeny $\varphi_{\mathsf{sk}} : E_0 \longrightarrow E_{\mathsf{pk}}$ associated to I_{sk}. The curve $\mathsf{pk} = E_{\mathsf{pk}}$ is published as a public key and $\mathsf{sk} := (I_{\mathsf{sk}}, \varphi_{\mathsf{sk}}(P_0), \varphi_{\mathsf{sk}}(Q_0))$ is safely stored as a secret key.

The commitment in SQIsign2D-West is identical to SQIsign2D-West key generation and R-SQIsignHD commitment summarised in Algorithm 3.9. We fix a big prime number $N_{\mathsf{com}} = \Theta(p^2)$ (*e.g.* $N_{\mathsf{com}} = N_{\mathsf{sk}}$). The prover samples uniformly at random a left $\mathcal{O}_0$-ideal I_{com} of norm N_{com} with Algorithm 3.4 and call Algorithm 2.7 to compute the image $(E_{\mathsf{com}}, \varphi_{\mathsf{com}}(P_0), \varphi_{\mathsf{com}}(Q_0))$ of the isogeny $\varphi_{\mathsf{com}} : E_0 \longrightarrow E_{\mathsf{com}}$ associated to I_{com}. The curve $\mathsf{com} = E_{\mathsf{com}}$ is published as the commitment and $\mathsf{sc} := (I_{\mathsf{com}}, \varphi_{\mathsf{com}}(P_0), \varphi_{\mathsf{com}}(Q_0))$ is stored as secret data.

Remark 4.1.1. In the heuristic version H-SQIsign2D-West, the commitment phase is identical to F-SQIsignHD commitment summarised in Algorithm 3.8. Algorithm 2.3 is used instead of 2.7, saving more than half the computational time.

4.1.3 Challenge

As in SQIsignHD, the challenge consists in an integer $\mathsf{chl} \in [\![0\ ;\ 2^{e_{\mathsf{chl}}} - 1]\!]$, where e_{chl} is a parameter denoting the size of the challenge space. This integer chl describes the kernel of the challenge isogeny $\varphi_{\mathsf{chl}} \colon E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}$:, *i.e.* $\ker(\varphi_{\mathsf{chl}}) = \langle P_{\mathsf{pk}} + [\mathsf{chl}]Q_{\mathsf{pk}} \rangle$, where $(P_{\mathsf{pk}}, Q_{\mathsf{pk}})$ is a deterministic basis of $E_{\mathsf{chl}}[2^e]$.

It is worth noting that, although $\deg(\varphi_{\mathsf{chl}}) = 2^e$, the challenge space contains only $2^{e_{\mathsf{chl}}} \ll 2^e$ possible challenges, i.e. we only allow $2^{e_{\mathsf{chl}}}$ possible kernels. As in R-SQIsignHD, the extra length of φ_{chl} is needed to deal with the fact that response isogenies may backtrack with φ_{chl}. In order to ensure λ bits of soundness security, we require $e_{\mathsf{chl}} \simeq \lambda$ but e_{chl} may be a slightly smaller in practice to prevent backtracking with the response.

Recall that in the signature scheme obtained via the Fiat-Shamir transform, the challenge is generated with a hash function applied to the commitment and the message m to be signed $\mathsf{chl} = H(E_{\mathsf{com}}, m)$. With m fixed, finding a colliding message m' such that $H(E_{\mathsf{com}}, m') = H(E_{\mathsf{com}}, m)$ would provide a forgery: an attacker could claim to have signed m' by using the prover's signature of m. This collision attack costs $\Theta(2^{e_{\mathsf{chl}}})$ which is smaller than the required cost $\Theta(2^\lambda)$ to ensure λ bits of security. The solution proposed in [AAA+25] to increase the cost of the attack to $\Theta(2^\lambda)$ is to iterate the hash function application $2^{\lambda - e_{\mathsf{chl}}}$ times. This solution is called *grinding*.

Remark 4.1.2. In H-SQIsign2D-West, the challenge space is the same but the challenge isogeny $\varphi_{\mathsf{chl}} : E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}$ associated to $\mathsf{chl} \in [\![0\ ;\ 2^{e_{\mathsf{chl}}} - 1]\!]$ is of degree $2^{e_{\mathsf{chl}}}$ instead of 2^e. Indeed, the extra length to prevent backtracking with the response becomes unnecessary.

4.1.4 Response

Unlike F-SQIsignHD, SQIsign2D-West exploits Algorithm 2.7 to embed the response isogeny $\varphi_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$ in dimension 2 instead of 4 (or 8). The general idea is to generate a random response ideal $I_{\mathsf{rsp}} \sim \overline{I}_{\mathsf{com}} \cdot I_{\mathsf{sk}} \cdot I_{\mathsf{chl}}$ of norm $q < 2^{e_{\mathsf{rsp}}}$ using Algorithm 3.6 and an auxiliary ideal I''_{aux} in $\mathcal{O}_0$ of norm $2^{e_{\mathsf{rsp}}} - q$ using Algorithm 3.4. Then, we compute the pushforward $I'_{\mathsf{aux}} := [I_{\mathsf{com}} \cdot I_{\mathsf{rsp}}]_* I''_{\mathsf{aux}}$ and apply Algorithm 2.7 to $I_{\mathsf{com}} \cdot I_{\mathsf{rsp}} \cdot I'_{\mathsf{aux}}$ to evaluate the corresponding isogeny $\varphi'_{\mathsf{aux}} \circ \varphi_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}$ on the basis (P_0, Q_0) of $E_0[2^e]$. Knowing φ_{com} on (P_0, Q_0), we obtain easily $\varphi'_{\mathsf{aux}} \circ \varphi_{\mathsf{rsp}}$ on a basis of $E_{\mathsf{com}}[2^e]$ and this data can be used to compute a 2-dimensional $2^{e_{\mathsf{rsp}}}$-isogeny embedding of φ_{rsp}. However, this natural method only applies when $q = \mathrm{nrd}(I_{\mathsf{rsp}})$ is odd, a condition that we do not impose in order to have more uniform response ideal distributions, as in R-SQIsignHD.

The complete SQIsign2D-West response algorithm is explained in the following. The description is unfortunately is a bit technical, as in R-SQIsignHD, in order to account for the even part of I_{rsp}. The diagram to keep in mind as we explain is the following one (see Fig. 4.1), where:

- $\varphi_{\mathsf{chl}} \colon E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}$ is the isogeny described by the challenge chl;
- $\varphi'_{\mathsf{chl}} \colon E_{\mathsf{pk}} \longrightarrow E^{(0)}_{\mathsf{chl}}$ is the portion of φ_{chl} that does not backtrack with the response isogeny;
- $\varphi^{(1)}_{\mathsf{rsp}} \colon E_{\mathsf{com}} \longrightarrow E'_{\mathsf{chl}}$ is the odd part of the response isogeny;
- $\varphi^{(0)}_{\mathsf{rsp}} \colon E'_{\mathsf{chl}} \longrightarrow E^{(0)}_{\mathsf{chl}}$ is the even, non-backtracking part of the response isogeny;
- $\varphi_{\mathsf{aux}} \colon E_{\mathsf{com}} \longrightarrow E_{\mathsf{aux}}$ is the auxiliary isogeny needed to embed the isogeny $\varphi^{(1)}_{\mathsf{rsp}}$ into a 2-dimensional isogeny;
- $\varphi'_{\mathsf{aux}} \colon E'_{\mathsf{chl}} \longrightarrow E'_{\mathsf{aux}}$ is the pushforward of φ_{aux} under $\varphi^{(1)}_{\mathsf{rsp}}$.

The first step is to compute the ideal I_{chl} corresponding to the isogeny $\varphi_{\mathsf{chl}} : E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}$ with kernel $\langle P_{\mathsf{pk}} + [\mathsf{chl}]Q_{\mathsf{pk}} \rangle$, where $(P_{\mathsf{pk}}, Q_{\mathsf{pk}})$ is a deterministic basis of $E_{\mathsf{pk}}[2^e]$. This is done via Algorithm 3.3 using the secret key $\mathsf{sk} = (I_{\mathsf{sk}}, \varphi_{\mathsf{sk}}(P_0), \varphi_{\mathsf{sk}}(Q_0))$.

Then the prover can compute $I := \overline{I}_{\mathsf{com}} \cdot I_{\mathsf{sk}} \cdot I_{\mathsf{chl}}$, apply Algorithm 3.6 to sample $\alpha \in I \setminus \{0\}$ of norm $\leq 2^{e_{\mathsf{rsp}}} \mathrm{nrd}(I)$ uniformly at random and set $I_{\mathsf{rsp}} := I\overline{\alpha}/\mathrm{nrd}(I)$. Note that to apply Algorithm 3.6, e_{rsp} should be chosen so that $2^{e_{\mathsf{rsp}}} \geq 3^{13/6}\sqrt{p}/\pi$ by Eq. (3.3), so e_{rsp} is very close to but slightly bigger than $e/2 \simeq \lambda$.

Once I_{rsp} has been generated, we factor out its even part as in R-SQIsignHD response. Let us write the norm of I_{rsp} as $\mathrm{nrd}(I_{\mathsf{rsp}}) = q = 2^n q' < 2^{e_{\mathsf{rsp}}}$ for an odd q'. We decompose φ_{rsp} as $\varphi_{\mathsf{rsp}} := \psi \circ \varphi^{(1)}_{\mathsf{rsp}}$

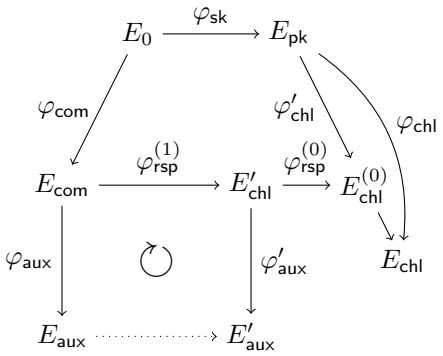

Figure 4.1: Response diagram.

with $\varphi^{(1)}_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E'_{\mathsf{chl}}$ of degree q' and $\psi : E'_{\mathsf{chl}} \longrightarrow E_{\mathsf{com}}$ of degree 2^n. It may happen that $\ker(\widehat{\psi}) \cap \ker(\widehat{\varphi_{\mathsf{chl}}})$ is not trivial. We say that φ_{rsp} *backtracks* through φ_{chl}. Let n_{bt} be the positive integer such that $2^{n_{\mathsf{bt}}} = \#\ker(\widehat{\psi}) \cap \ker(\widehat{\varphi_{\mathsf{chl}}})$. Equivalently, $2^{n_{\mathsf{bt}}}$ is the norm of the ideal $\overline{I}_{\mathsf{bt}} = \overline{I}_{\mathsf{rsp}} + \overline{I}_{\mathsf{chl}}$.

Let $r := n - n_{\mathsf{bt}}$ and define $\varphi^{(0)}_{\mathsf{rsp}} : E'_{\mathsf{chl}} \longrightarrow E^{(0)}_{\mathsf{chl}}$ to be the isogeny with kernel $\ker(\psi)[2^r]$ – the isogeny $\varphi^{(0)}_{\mathsf{rsp}}$ coincides with the non-backtrack portion of φ_{rsp}. Now, let us factor I_{rsp} as $I^{(1)}_{\mathsf{rsp}} \cdot I^{(0)}_{\mathsf{rsp}} \cdot I_{\mathsf{bt}}$, where $\mathrm{nrd}(I^{(1)}_{\mathsf{rsp}}) = q'$ and $\mathrm{nrd}(I^{(0)}_{\mathsf{rsp}}) = 2^r$. The isogenies $\varphi^{(1)}_{\mathsf{rsp}}$ and $\varphi^{(0)}_{\mathsf{rsp}}$ correspond to $I^{(1)}_{\mathsf{rsp}}$ and $I^{(0)}_{\mathsf{rsp}}$, respectively. Actually, $\varphi^{(0)}_{\mathsf{rsp}}$ is a cyclic isogeny by the following lemma.

Lemma 4.1.3. *$I^{(0)}_{rsp}$ is a primitive left ideal of $\mathcal{O}'_{chl} := O_L(I^{(0)}_{rsp})$.*

Proof. Let us factor $I_{\mathsf{chl}} = I'_{\mathsf{chl}} \cdot I_{\mathsf{bt}}$. Then, we have

$$\overline{I}_{\mathsf{bt}} = \overline{I}_{\mathsf{rsp}} + \overline{I}_{\mathsf{chl}} = \overline{I}_{\mathsf{bt}}(\overline{I}^{(0)}_{\mathsf{rsp}} \cdot \overline{I}^{(1)}_{\mathsf{rsp}} + \overline{I}'_{\mathsf{chl}}).$$

Multiplying this equality on the left by I_{bt}, we get that $\overline{I}^{(0)}_{\mathsf{rsp}} \cdot \overline{I}^{(1)}_{\mathsf{rsp}} + \overline{I}'_{\mathsf{chl}} = 1/2^{n_{\mathsf{bt}}} I_{\mathsf{bt}} \cdot \overline{I}_{\mathsf{bt}} = \mathcal{O}'_{\mathsf{chl}}$. If $I^{(0)}_{\mathsf{rsp}}$ was not a primitive ideal, then we would have $I^{(0)}_{\mathsf{rsp}} \subseteq 2\mathcal{O}'_{\mathsf{chl}}$. In particular, there would exist $\alpha \in \mathcal{O}'_{\mathsf{chl}}$ and $\beta \in \overline{I}'_{\mathsf{chl}}$ such that $2\alpha + \beta = 1$, so that

$$\mathrm{nrd}(\beta) = \mathrm{nrd}(1 - 2\alpha) = 1 - 2\,\mathrm{Tr}(\alpha) + 4\,\mathrm{nrd}(\alpha) \equiv 1 \mod 2.$$

But $2^{e-n_{\mathsf{bt}}} = \mathrm{nrd}(\overline{I}'_{\mathsf{chl}}) \,|\, \mathrm{nrd}(\beta)$ and $n_{\mathsf{bt}} \leq e_{\mathsf{rsp}} < e$ so $2 \,|\, \mathrm{nrd}(\beta)$. Contradiction. □

To compute this ideal factorisation $I_{\mathsf{rsp}} = I^{(1)}_{\mathsf{rsp}} \cdot I^{(0)}_{\mathsf{rsp}} \cdot I_{\mathsf{bt}}$, the prover may compute $J_{\mathsf{rsp}} := I_{\mathsf{rsp}} \cdot I_{\mathsf{bt}}^{-1} = I_{\mathsf{rsp}} \cdot \overline{I}_{\mathsf{bt}}/2^{n_{\mathsf{bt}}}$. We then have $\overline{I}^{(0)}_{\mathsf{rsp}} = \overline{J}_{\mathsf{rsp}} + 2^r O_R(J_{\mathsf{rsp}})$ by Lemma 3.3.5 and since 2 does not divide J_{rsp} by Lemma 4.1.3. Finally, we have $I^{(1)}_{\mathsf{rsp}} = J_{\mathsf{rsp}} \cdot {I^{(0)}_{\mathsf{rsp}}}^{-1} = J_{\mathsf{rsp}} \cdot \overline{I}^{(0)}_{\mathsf{rsp}}/2^r$ which completes the factorisation.

Since $\varphi^{(1)}_{\mathsf{rsp}}$ has odd degree bounded by 2^f with $f := e_{\mathsf{rsp}} - n$, it can be represented via a 2^f-isogeny in dimension 2 by Kani's Lemma. This requires computing an auxiliary isogeny $\varphi'_{\mathsf{aux}} : E'_{\mathsf{chl}} \longrightarrow E'_{\mathsf{aux}}$ of degree $2^f - q'$.

For security reasons (see Section 4.2.2), we need the isogeny $\varphi'_{\mathsf{aux}} : E'_{\mathsf{chl}} \longrightarrow E'_{\mathsf{aux}}$ to be uniformly sampled among all the isogenies of degree $2^f - q'$. Hence, the prover samples a random left ideal I''_{aux} of $\mathcal{O}_0$ of norm $2^f - q'$ using Algorithm 3.4 and then computes I'_{aux} as the pushforward $I'_{\mathsf{aux}} := [I_{\mathsf{com}} \cdot I^{(1)}_{\mathsf{rsp}}]_* I''_{\mathsf{aux}}$. The prover can then evaluate $\varphi'_{\mathsf{aux}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}$ on (P_0, Q_0) by running Algorithm 2.7 on input $I_{\mathsf{com}} \cdot I^{(1)}_{\mathsf{rsp}} \cdot I'_{\mathsf{aux}}$.

While a representation of $\varphi'_{\mathsf{aux}} \circ \varphi^{(1)}_{\mathsf{rsp}}$ could act as a valid response, we want the identification protocol to be *commitment recoverable*, *i.e.* ensure it is possible to recompute the commitment curve from a the challenge and corresponding response. This eventually leads to a more compact signature (since the commitment does not need to be included). To achieve such a property, we want the isogeny $\varphi^{(0)}_{\mathsf{rsp}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \widehat{\varphi}_{\mathsf{aux}}$ connecting E_{aux} and E'_{chl}, passing through E_{com}. Thus, the prover has to

compute the isogeny $\varphi_{\mathsf{aux}} : E_0 \longrightarrow E_{\mathsf{aux}}$ of degree $2^f - q'$ with $f := e_{\mathsf{rsp}} - n$, fitting in the following $(q', 2^f - q')$-isogeny diamond:

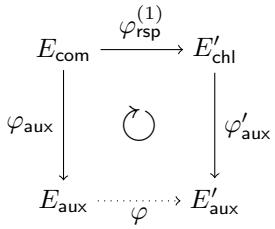

By Kani's lemma (Lemma 2.2.6), we may consider the 2^f-isogeny (with $f := e_{\mathsf{rsp}} - n$):

$$\Phi = \begin{pmatrix} \varphi^{(1)}_{\mathsf{rsp}} & \widehat{\varphi}'_{\mathsf{aux}} \\ -\varphi_{\mathsf{aux}} & \widehat{\varphi} \end{pmatrix} : E_{\mathsf{com}} \times E'_{\mathsf{aux}} \longrightarrow E'_{\mathsf{chl}} \times E_{\mathsf{aux}}. \tag{4.1}$$

whose kernel is:

$$\ker(\Phi) = \{([q']P, \varphi'_{\mathsf{aux}} \circ \varphi^{(1)}_{\mathsf{rsp}}(P)) \mid P \in E_{\mathsf{com}}[2^f]\}.$$

To compute Φ, the prover computes $T_1, T_2 \in (E_{\mathsf{com}} \times E'_{\mathsf{aux}})[2^{f+2}]$ forming an isotropic subgroup and such that $\ker(\Phi) = \langle [4]T_1, [4]T_2 \rangle$. The prover may define T_1, T_2 as follows:

$$T_1 := ([2^{e-f-2}q']\varphi_{\mathsf{com}}(P_0), [2^{e-f-2}]\varphi'_{\mathsf{aux}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(P_0))$$
$$\text{and} \quad T_2 := ([2^{e-f-2}(q' - 2^{e-f})]\varphi_{\mathsf{com}}(Q_0), [2^{e-f-2}]\varphi'_{\mathsf{aux}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(Q_0)),$$

thus exploiting the knowledge of $(\varphi_{\mathsf{com}}(P_0), \varphi_{\mathsf{com}}(P_0))$ and $(\varphi'_{\mathsf{aux}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(P_0), \varphi'_{\mathsf{aux}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(P_0))$ obtained previously. By Theorem 2.2.12, Φ can be computed from T_1, T_2 as a chain of 2-isogenies, with techniques that will be presented in Section 6.5.

To complete the response algorithm, we still need to compute the non-backtracking part of the response isogeny. Let $\varphi^{(0)}_{\mathsf{rsp}} : E'_{\mathsf{chl}} \longrightarrow E^{(0)}_{\mathsf{chl}}$ be such an isogeny, which corresponds to the ideal $I^{(0)}_{\mathsf{rsp}}$. From the evaluation of:

$$\Phi(\varphi_{\mathsf{com}}(P_0), 0) = (\varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(P_0), -\varphi_{\mathsf{aux}} \circ \varphi_{\mathsf{com}}(P_0))$$
$$\text{and} \quad \Phi(\varphi_{\mathsf{com}}(Q_0), 0) = (\varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(P_0), -\varphi_{\mathsf{aux}} \circ \varphi_{\mathsf{com}}(P_0)), \tag{4.2}$$

we extract $(\varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(P_0), \varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(Q_0))$. Then, $I^{(0)}_{\mathsf{rsp}}$ being primitive by Lemma 4.1.3, $\varphi^{(0)}_{\mathsf{rsp}}$ can be computed by applying Algorithm 3.3 to $I^{(0)}_{\mathsf{rsp}}$ and the image of $([2^{e-r}]\varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(P_0), [2^{e-r}]\varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(Q_0))$.

Let $\varphi'_{\mathsf{chl}} : E_{\mathsf{pk}} \longrightarrow E^{(0)}_{\mathsf{chl}}$ be the isogeny with kernel $\langle [2^{n_{\mathsf{bt}}}](P_{\mathsf{pk}} + [\mathsf{chl}]Q_{\mathsf{pk}}) \rangle$. In other words, φ'_{chl} is the portion of φ_{chl} that does not backtrack with the response isogeny. Even though φ'_{chl} and $\varphi^{(0)}_{\mathsf{rsp}}$ map onto the same elliptic curve, the curves obtained after an explicit computation of the two isogenies will only be equal up to isomorphism. Thus, the prover additionally has to compute an explicit isomorphism to let the two curves agree. The explicit computation of the isomorphism between the codomains of φ'_{chl} and $\varphi^{(0)}_{\mathsf{rsp}}$ is required to facilitate the verification.

Let $(P_{\mathsf{aux}}, Q_{\mathsf{aux}})$ be a deterministic basis of $E_{\mathsf{aux}}[2^{e_{\mathsf{rsp}} - n_{\mathsf{bt}} + 2}]$, $\mu \equiv (2^f - q')^{-1} \mod 2^{f+2}$, $\nu \equiv q'^{-1} \mod 4$ and define

$$P_{\mathsf{chl}} := [\mu]\varphi^{(0)}_{\mathsf{rsp}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \widehat{\varphi}_{\mathsf{aux}}(P_{\mathsf{aux}}), \quad Q_{\mathsf{chl}} := [\mu(1 - 2^f\nu)]\varphi^{(0)}_{\mathsf{rsp}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \widehat{\varphi}_{\mathsf{aux}}(Q_{\mathsf{aux}}). \tag{4.3}$$

$(P_{\mathsf{chl}}, Q_{\mathsf{chl}})$ can be computed by evaluating Φ and computing some discrete logarithms as follows. From Eq. (4.2) and the previous computation of $\varphi^{(0)}_{\mathsf{rsp}}$, we obtain $\widetilde{P}_{\mathsf{chl}} := \varphi^{(0)}_{\mathsf{rsp}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(P_0)$, $\widetilde{Q}_{\mathsf{chl}} := \varphi^{(0)}_{\mathsf{rsp}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(Q_0)$, $\widetilde{P}_{\mathsf{aux}} := \varphi_{\mathsf{aux}} \circ \varphi_{\mathsf{com}}(P_0)$ and $\widetilde{Q}_{\mathsf{aux}} := \varphi_{\mathsf{aux}} \circ \varphi_{\mathsf{com}}(Q_0)$. Now, $(\widetilde{P}_{\mathsf{aux}}, \widetilde{Q}_{\mathsf{aux}})$ is a basis of $E_{\mathsf{aux}}[2^e]$ so we can find $a, b, c, d \in \mathbb{Z}/2^{e_{\mathsf{rsp}} - n_{\mathsf{bt}} + 2}\mathbb{Z}$ such that:

$$P_{\mathsf{aux}} = [2^{e - e_{\mathsf{rsp}} + n_{\mathsf{bt}} - 2}]([a]\widetilde{P}_{\mathsf{aux}} + [b]\widetilde{Q}_{\mathsf{aux}}) \quad \text{and} \quad Q_{\mathsf{aux}} = [2^{e - e_{\mathsf{rsp}} + n_{\mathsf{bt}} - 2}]([c]\widetilde{P}_{\mathsf{aux}} + [d]\widetilde{Q}_{\mathsf{aux}}),$$

It then follows that:

$$\begin{aligned}P_{\mathsf{chl}} &= [\mu]\varphi^{(0)}_{\mathsf{rsp}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \widehat{\varphi}_{\mathsf{aux}}(P_{\mathsf{aux}}) = [\mu 2^{e-e_{\mathsf{rsp}}+n_{\mathsf{bt}}}]\varphi^{(0)}_{\mathsf{rsp}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \widehat{\varphi}_{\mathsf{aux}}([a]\widetilde{P}_{\mathsf{aux}} + [b]\widetilde{Q}_{\mathsf{aux}}) \\ &= [\mu 2^{e-e_{\mathsf{rsp}}+n_{\mathsf{bt}}}]([a]\varphi^{(0)}_{\mathsf{rsp}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \widehat{\varphi}_{\mathsf{aux}} \circ \varphi_{\mathsf{aux}} \circ \varphi_{\mathsf{com}}(P_0) + [b]\varphi^{(0)}_{\mathsf{rsp}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \widehat{\varphi}_{\mathsf{aux}} \circ \varphi_{\mathsf{aux}} \circ \varphi_{\mathsf{com}}(Q_0)) \\ &= [(2^f - q')^{-1}(2^f - q')2^{e-e_{\mathsf{rsp}}+n_{\mathsf{bt}}}]([a]\varphi^{(0)}_{\mathsf{rsp}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(P_0) + [b]\varphi^{(0)}_{\mathsf{rsp}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(Q_0)) \\ &= [2^{e-e_{\mathsf{rsp}}+n_{\mathsf{bt}}-2}]([a]\widetilde{P}_{\mathsf{chl}} + [b]\widetilde{Q}_{\mathsf{chl}})\end{aligned}$$

And similarly, $Q_{\mathsf{chl}} = [2^{e-e_{\mathsf{rsp}}+n_{\mathsf{bt}}-2}(1-2^f\nu)]([c]\widetilde{P}_{\mathsf{chl}}+[d]\widetilde{Q}_{\mathsf{chl}})$. Once $(P_{\mathsf{chl}}, Q_{\mathsf{chl}})$ is computed, the prover finally outputs the response that consists in $(E_{\mathsf{aux}}, P_{\mathsf{chl}}, Q_{\mathsf{chl}}, r, n_{\mathsf{bt}})$. We shall see in Section 4.1.5 that the verifier will be able to compute $\widetilde{\Phi}$ from this data. We summarise what has been explained in this section in Algorithm 4.1.

Remark 4.1.4. In H-SQIsign2D-West, I_{rsp} is selected so that $q = \mathrm{nrd}(I_{\mathsf{rsp}})$ is odd at the expense of some small but non-negligible failure probability to find such an ideal which complexifies the security analysis. However, the response procedure is greatly simplified from a technical point of view, as the factorisation of I_{rsp} and the computation of the non-backtracking even part $\varphi^{(0)}_{\mathsf{rsp}}$ can be avoided. The computation of the 2-dimensional isogeny $\Phi : E_{\mathsf{com}} \times E'_{\mathsf{aux}} \longrightarrow E_{\mathsf{chl}} \times E_{\mathsf{aux}}$ can also be avoided, which saves a significant amount of time. We refer to [BDF+25, Appendix B] for more details.

4.1.5 Verification

Given the public key E_{pk}, a challenge $\mathsf{chl} \in [\![0\ ;\ 2^{e_{\mathsf{chl}}} - 1]\!]$ and response $\mathsf{rsp} = (E_{\mathsf{aux}}, P_{\mathsf{chl}}, Q_{\mathsf{chl}}, r, n_{\mathsf{bt}})$, the verifier first generates a deterministic basis $(P_{\mathsf{pk}}, Q_{\mathsf{pk}})$ of $E_{\mathsf{pk}}[2^e]$ and computes the isogeny $\varphi'_{\mathsf{chl}} : E_0 \longrightarrow E^{(0)}_{\mathsf{chl}}$ with kernel $\langle [2^{n_{\mathsf{bt}}}](P_{\mathsf{pk}} + [\mathsf{chl}]Q_{\mathsf{pk}})\rangle$ – this corresponds to the non-backtracking portion of the challenge isogeny as in the previous paragraph. Additionally, they compute a deterministic basis $(P_{\mathsf{aux}}, Q_{\mathsf{aux}})$ of $E_{\mathsf{aux}}[2^{e_{\mathsf{rsp}}-n_{\mathsf{bt}}+2}]$

If $r > 0$, it means that the prover has chosen a response isogeny having an even, non-backtrack component. In this case, $[2^{e_{\mathsf{rsp}}-r-n_{\mathsf{bt}}+2}]P_{\mathsf{chl}}$ and $[2^{e_{\mathsf{rsp}}-r-n_{\mathsf{bt}}+2}]Q_{\mathsf{chl}}$ are linearly dependent, and $\langle [2^{e_{\mathsf{rsp}}-r-n_{\mathsf{bt}}+2}]P_{\mathsf{chl}}, [2^{e_{\mathsf{rsp}}-r-n_{\mathsf{bt}}+2}]Q_{\mathsf{chl}}\rangle$ is the kernel of the dual of the isogeny $\varphi^{(0)}_{\mathsf{rsp}}$ (see Fig. 4.1). The verifier then computes the isogeny $\widehat{\varphi}^{(0)}_{\mathsf{rsp}} : E^{(0)}_{\mathsf{chl}} \longrightarrow E'_{\mathsf{chl}}$ with kernel $\langle [2^{e_{\mathsf{rsp}}-r-n_{\mathsf{bt}}+2}]P_{\mathsf{chl}}, [2^{e_{\mathsf{rsp}}-r-n_{\mathsf{bt}}+2}]Q_{\mathsf{chl}}\rangle$ and computes $\widehat{\varphi}^{(0)}_{\mathsf{rsp}}(P_{\mathsf{chl}})$ and $\widehat{\varphi}^{(0)}_{\mathsf{rsp}}(Q_{\mathsf{chl}})$.

By Lemma 2.2.9, the polarised dual $\widetilde{\Phi} : E'_{\mathsf{chl}} \times E_{\mathsf{aux}} \longrightarrow E_{\mathsf{com}} \times E'_{\mathsf{aux}}$ of the 2-dimensional 2^f-isogeny defined in Eq. (4.1) has kernel:

$$\ker(\widetilde{\Phi}) = \{(\varphi^{(1)}_{\mathsf{rsp}}(P), -\varphi_{\mathsf{aux}}(P)) \mid P \in E_{\mathsf{com}}[2^f]\} = \{([\mu]\varphi^{(1)}_{\mathsf{rsp}} \circ \widehat{\varphi}_{\mathsf{aux}}(P), -P) \mid P \in E_{\mathsf{aux}}[2^f]\},$$

with $\mu \equiv (2^f - q')^{-1} \mod 2^{f+2}$. Besides, we have by Eq. (4.3):

$$\widehat{\varphi}^{(0)}_{\mathsf{rsp}}(P_{\mathsf{chl}}) = [2^r\mu]\varphi^{(1)}_{\mathsf{rsp}} \circ \widehat{\varphi}_{\mathsf{aux}}(P_{\mathsf{aux}}) \quad \text{and} \quad \widehat{\varphi}^{(0)}_{\mathsf{rsp}}(Q_{\mathsf{chl}}) = [2^r\mu(1 - 2^f\nu)]\varphi^{(1)}_{\mathsf{rsp}} \circ \widehat{\varphi}_{\mathsf{aux}}(Q_{\mathsf{aux}}),$$

with $\nu \equiv q'^{-1} \mod 4$. Since $f := e_{\mathsf{rsp}} - n_{\mathsf{bt}} - r$, it follows that $\ker(\widetilde{\Phi})$ is generated by $([4]T_1, [4]T_2)$, where:

$$T_1 = (\widehat{\varphi}^{(0)}_{\mathsf{rsp}}(P_{\mathsf{chl}}), -[2^r]P_{\mathsf{aux}}) \quad \text{and} \quad T_2 = (\widehat{\varphi}^{(0)}_{\mathsf{rsp}}(Q_{\mathsf{chl}}), -[2^r]Q_{\mathsf{aux}}),$$

form a maximal isotropic subgroup of $(E'_{\mathsf{chl}} \times E_{\mathsf{aux}})[2^{f+2}]$ by construction. Using the techniques from Section 6.5, the prover can compute $\widetilde{\Phi}$ from T_1 and T_2. Then the codomain of $\widetilde{\Phi}$ is expected to be of the form $E_{\mathsf{com}} \times E'_{\mathsf{aux}}$. If E_{com} is indeed the first component of the codomain, this proves that $\widetilde{\Phi}$ efficiently represents an isogeny connecting E_{com} and E'_{chl}, so that the response (together with the challenge) efficiently represents an isogeny connecting E_{com} and E_{chl}, as desired. Algorithm 4.2 summarises the whole verification procedure.

4.2 Security analysis

As SQIsignHD, SQIsign2D-West is a digital signature scheme obtained via the Fiat-Shamir transform so we rely on Theorem 3.4.2 to prove this scheme is universally unforgeable under chosen message

Algorithm 4.1: Response

Data: The public parameters pp, public key E_{pk}, the secret key sk, the commitment $(E_{\mathsf{com}}, \mathsf{com})$, the secret commitment data sc, and the challenge $\mathsf{chl} \in [\![0\ ;\ 2^{e_{\mathsf{chl}}} - 1]\!]$.

Result: A response $\mathsf{rsp} = (E_{\mathsf{aux}}, P_{\mathsf{aux}}, Q_{\mathsf{aux}}, r, n_{\mathsf{bt}})$

$I_{\mathsf{sk}}, \varphi_{\mathsf{sk}}(P_0), \varphi_{\mathsf{sk}}(Q_0) \leftarrow \mathsf{sk}$;

$I_{\mathsf{com}}, \varphi_{\mathsf{com}}(P_0), \varphi_{\mathsf{com}}(Q_0) \leftarrow \mathsf{sc}$;

Extract $e_{\mathsf{rsp}}, N_{\mathsf{com}}$, a basis $(\beta_1, \cdots, \beta_4)$ of $\mathcal{O}_0$ and $(\varepsilon_0(\beta_i)(P_0), \varepsilon_0(\beta_i)(Q_0))_{1 \le i \le 4}$ from pp;

`// Computation of` I_{chl}

Compute a deterministic basis $(P_{\mathsf{pk}}, Q_{\mathsf{pk}})$ of $E_{\mathsf{pk}}[2^e]$;

$K_{\mathsf{chl}} \leftarrow P_{\mathsf{pk}} + [\mathsf{chl}]Q_{\mathsf{pk}}$;

Call Algorithm 3.3 on K_{chl}, $(\beta_1, \cdots, \beta_4)$, $(\varepsilon_0(\beta_i)(R_0), \varepsilon_0(\beta_i)(S_0))_{1 \le i \le 4}$ and $(\varphi_{\mathsf{sk}}(P_0), \varphi_{\mathsf{sk}}(Q_0))$ to obtain I_{chl};

$I \leftarrow \overline{I}_{\mathsf{com}} \cdot I_{\mathsf{sk}} \cdot I_{\mathsf{chl}}$;

`// Computation and factorisation of` I_{rsp}

Call Algorithm 3.6 to sample $\alpha \in I \setminus \{0\}$ of norm $\mathrm{nrd}(\alpha) \le 2^{e_{\mathsf{rsp}}} \mathrm{nrd}(I)$ uniformly at random;

$I_{\mathsf{rsp}} \leftarrow I\overline{\alpha}/\mathrm{nrd}(I)$;

$q \leftarrow \mathrm{nrd}(I_{\mathsf{rsp}})$, $n \leftarrow v_2(q)$, $q' \leftarrow q/2^n$;

$\overline{I}_{\mathsf{bt}} \leftarrow \overline{I}_{\mathsf{rsp}} + \overline{I}_{\mathsf{chl}}$, $n_{\mathsf{bt}} \leftarrow \log_2(\mathrm{nrd}(I_{\mathsf{bt}}))$, $r \leftarrow n - n_{\mathsf{bt}}$;

$J_{\mathsf{rsp}} \leftarrow I_{\mathsf{rsp}} \cdot \overline{I}_{\mathsf{bt}}/2^{n_{\mathsf{bt}}}$, $\overline{I}^{(0)}_{\mathsf{rsp}} \leftarrow \overline{J}_{\mathsf{rsp}} + 2^r O_R(J_{\mathsf{rsp}})$,$I^{(0)}_{\mathsf{rsp}} \leftarrow \overline{\overline{I}}^{(0)}_{\mathsf{rsp}}$, $I^{(1)}_{\mathsf{rsp}} \leftarrow J_{\mathsf{rsp}} \cdot \overline{I}^{(0)}_{\mathsf{rsp}}/2^r$;

`// Evaluation of` $\varphi'_{\mathsf{aux}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}$

Call Algorithm 3.4 to sample a left $\mathcal{O}_0$-ideal I''_{aux} of norm $2^{e_{\mathsf{rsp}}-n} - q'$;

$I'_{\mathsf{aux}} \leftarrow [I_{\mathsf{com}} \cdot I^{(1)}_{\mathsf{rsp}}]_* I''_{\mathsf{aux}}$;

Call Algorithm 2.7 on $I_{\mathsf{com}} \cdot I^{(1)}_{\mathsf{rsp}} \cdot I'_{\mathsf{aux}}$ to compute the image $(E'_{\mathsf{aux}}, \varphi'_{\mathsf{aux}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(P_0), \varphi'_{\mathsf{aux}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(Q_0))$ of the associated isogeny $\varphi'_{\mathsf{aux}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}} : E_0 \longrightarrow E'_{\mathsf{aux}}$;

`// Computation and evaluation of` $\Phi : E_{\mathsf{com}} \times E'_{\mathsf{aux}} \longrightarrow E'_{\mathsf{chl}} \times E_{\mathsf{aux}}$

$T_1 \leftarrow ([2^{e-(e_{\mathsf{rsp}}-n)-2}q']\varphi_{\mathsf{com}}(P_0), [2^{e-(e_{\mathsf{rsp}}-n)-2}]\varphi'_{\mathsf{aux}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(P_0))$;

$T_2 \leftarrow ([2^{e-(e_{\mathsf{rsp}}-n)-2}(q' - 2^{e-(e_{\mathsf{rsp}}-n)})]\varphi_{\mathsf{com}}(Q_0), [2^{e-(e_{\mathsf{rsp}}-n)-2}]\varphi'_{\mathsf{aux}} \circ \varphi^{(1)}_{\mathsf{rsp}} \circ \varphi_{\mathsf{com}}(Q_0))$;

Apply the algorithms from Section 6.5 to (T_1, T_2) to compute $\Phi : E_{\mathsf{com}} \times E'_{\mathsf{aux}} \longrightarrow E'_{\mathsf{chl}} \times E_{\mathsf{aux}}$ with kernel $\langle [4]T_1, [4]T_2 \rangle$;

$(\pm I(S), \pm I'(S)) \leftarrow \Phi(\varphi_{\mathsf{com}}(S), 0)$ for $S \in \{P_0, Q_0, P_0 - Q_0\}$;

Lift $(\widetilde{P}_{\mathsf{chl}}, \widetilde{Q}_{\mathsf{chl}}) := (I(P_0), I(Q_0))$ from $\pm I(P_0), \pm I(Q_0), \pm I(P_0 - Q_0)$;

Lift $(\widetilde{P}_{\mathsf{aux}}, \widetilde{Q}_{\mathsf{aux}}) := (I'(P_0), I'(Q_0))$ from $\pm I'(P_0), \pm I'(Q_0), \pm I'(P_0 - Q_0)$;

`// Computation and evaluation of` $\varphi^{(0)}_{\mathsf{rsp}} : E'_{\mathsf{chl}} \longrightarrow E^{(0)}_{\mathsf{chl}}$ `and` $\varphi'_{\mathsf{chl}} : E_{\mathsf{pk}} \longrightarrow E'^{(0)}_{\mathsf{chl}}$

if $r > 0$ **then**

  Call Algorithm 3.3 on $I^{(0)}_{\mathsf{rsp}}$, $([2^{e-r}]\widetilde{P}_{\mathsf{chl}}, [2^{e-r}]\widetilde{Q}_{\mathsf{chl}})$, $I_{\mathsf{com}} \cdot I^{(1)}_{\mathsf{rsp}}$, $(\beta_1, \cdots, \beta_4)$ and $([2^{e-r}]\varepsilon_0(\beta_i)(P_0), [2^{e-r}]\varepsilon_0(\beta_i)(Q_0))_{1 \le i \le 4}$ to compute a generator $K^{(0)}_{\mathsf{rsp}} \in E'_{\mathsf{chl}}[2^r]$ of $\ker(\varphi^{(0)}_{\mathsf{rsp}})$;

  Compute the isogeny $\varphi^{(0)}_{\mathsf{rsp}} : E'_{\mathsf{chl}} \longrightarrow E^{(0)}_{\mathsf{chl}}$ of kernel $\langle K^{(0)}_{\mathsf{rsp}} \rangle$;

  $\widetilde{P}_{\mathsf{chl}}, \widetilde{Q}_{\mathsf{chl}} \leftarrow \varphi^{(0)}_{\mathsf{rsp}}(\widetilde{P}_{\mathsf{chl}}), \varphi^{(0)}_{\mathsf{rsp}}(\widetilde{Q}_{\mathsf{chl}})$;

end

Compute $\varphi'_{\mathsf{chl}} : E_{\mathsf{pk}} \to E'^{(0)}_{\mathsf{chl}}$ of kernel $\langle [2^{n_{\mathsf{bt}}}]K_{\mathsf{chl}} \rangle$;

Compute the isomorphism $\iota_{\mathsf{chl}} : E^{(0)}_{\mathsf{chl}} \to E'^{(0)}_{\mathsf{chl}}$;

$\widetilde{P}_{\mathsf{chl}}, \widetilde{Q}_{\mathsf{chl}} \leftarrow \iota_{\mathsf{chl}}(\widetilde{P}_{\mathsf{chl}}), \iota_{\mathsf{chl}}(\widetilde{Q}_{\mathsf{chl}})$;

`// Computation of` $(P_{\mathsf{chl}}, Q_{\mathsf{chl}})$

Compute a deterministic basis $(P_{\mathsf{aux}}, Q_{\mathsf{aux}})$ of $E_{\mathsf{aux}}[2^{e_{\mathsf{rsp}}-n_{\mathsf{bt}}+2}]$;

Compute $a, b, c, d \in \mathbb{Z}/2^{e_{\mathsf{rsp}}-n_{\mathsf{bt}}+2}\mathbb{Z}$ such that $P_{\mathsf{aux}} = [2^{e-e_{\mathsf{rsp}}+n_{\mathsf{bt}}-2}]([a]\widetilde{P}_{\mathsf{aux}} + [b]\widetilde{Q}_{\mathsf{aux}})$ and $Q_{\mathsf{aux}} = [2^{e-e_{\mathsf{rsp}}+n_{\mathsf{bt}}-2}]([c]\widetilde{P}_{\mathsf{aux}} + [d]\widetilde{Q}_{\mathsf{aux}})$;

$f \leftarrow e_{\mathsf{rsp}} - n$, $\nu \leftarrow q'^{-1} \mod 4$;

$P_{\mathsf{chl}}, Q_{\mathsf{chl}} \leftarrow [2^{e-e_{\mathsf{rsp}}+n_{\mathsf{bt}}-2}]([a]\tilde{P}_{\mathsf{chl}} + [b]\tilde{Q}_{\mathsf{chl}}), [2^{e-e_{\mathsf{rsp}}+n_{\mathsf{bt}}-2}(1 - 2^f\nu)]([c]\tilde{P}_{\mathsf{chl}} + [d]\tilde{Q}_{\mathsf{chl}})$;

return $(E_{\mathsf{aux}}, P_{\mathsf{chl}}, Q_{\mathsf{chl}}, r, n_{\mathsf{bt}})$;

Algorithm 4.2: Verify

Data: The public parameters pp, the public key E_{pk}, the commitment E_{com}, the challenge chl and the response $\mathsf{rsp} = (E_{\mathsf{aux}}, P_{\mathsf{chl}}, Q_{\mathsf{chl}}, r, n_{\mathsf{bt}})$.
Result: A boolean value indicating the validity of the response rsp.

Extract e, e_{rsp} from pp;
$E_{\mathsf{aux}}, P_{\mathsf{chl}}, Q_{\mathsf{chl}}, r, n_{\mathsf{bt}} \leftarrow \mathsf{rsp}$;
Compute a deterministic basis $(P_{\mathsf{pk}}, Q_{\mathsf{pk}})$ of $E_{\mathsf{pk}}[2^e]$;
Compute $\varphi'_{\mathsf{chl}} : E_0 \longrightarrow E^{(0)}_{\mathsf{chl}}$ with kernel $\langle [2^{n_{\mathsf{bt}}}](P_{\mathsf{pk}} + [\mathsf{chl}]Q_{\mathsf{pk}})\rangle$;
if $P_{chl} \notin E^{(0)}_{chl}$ *or* $Q_{chl} \notin E^{(0)}_{chl}$ **then**
  return False;
end
if $[2^{e_{rsp}-n_{\mathsf{bt}}+1}]Q_{chl} \neq 0$ **then**
  $R \leftarrow [2^{e_{\mathsf{rsp}}-n_{\mathsf{bt}}-r+2}]Q_{\mathsf{chl}}$;
else
  $R \leftarrow [2^{e_{\mathsf{rsp}}-n_{\mathsf{bt}}-r+2}]P_{\mathsf{chl}}$;
end
Compute $\widehat{\varphi}^{(0)}_{\mathsf{rsp}} : E^{(0)}_{\mathsf{chl}} \to E'_{\mathsf{chl}}$ of kernel $\langle R\rangle$;
Compute a deterministic basis $(P_{\mathsf{aux}}, Q_{\mathsf{aux}})$ of $E_{\mathsf{aux}}[2^{e_{\mathsf{rsp}}-n_{\mathsf{bt}}+2}]$;
$T_1 \leftarrow (\widehat{\varphi}^{(0)}_{\mathsf{rsp}}(P_{\mathsf{chl}}), -[2^r]P_{\mathsf{aux}}), T_2 \leftarrow (\widehat{\varphi}^{(0)}_{\mathsf{rsp}}(Q_{\mathsf{chl}}), -[2^r]Q_{\mathsf{aux}})$;
Apply the algorithms from Section 6.5 to (T_1, T_2) to compute $\widetilde{\Phi} : E'_{\mathsf{chl}} \times E_{\mathsf{aux}} \longrightarrow E_1 \times E_2$ with kernel $\langle [4]T_1, [4]T_2\rangle$;
if *the computation of* $\widetilde{\Phi}$ *fails* **then**
  return False;
else
  return $j(E_1) \stackrel{?}{=} j(E_{\mathsf{com}})$;
end

attacks in the random oracle model. SQIsign2D-West is also built on the NP-relation $\mathcal{R} : W \times K \longrightarrow \{0,1\}$, where K is the space of supersingular elliptic curves over $\mathbb{F}_{p^2}$ and W is the space of endomorphisms of supersingular elliptic curves over $\mathbb{F}_{p^2}$ and:

$$\forall(\alpha, E) \in W \times K, \quad \mathcal{R}(\alpha, E) = 1 \iff \alpha \in \mathrm{End}(E) \setminus \mathbb{Z},$$

as defined in Eq. (3.9). We have seen that $\mathcal{R}$ is hard provided the supersingular one endomorphism problem (Problem 3.4.3) is hard. The completeness follows by construction and we justify the special soundness and honest verifier zero knowledge (HVZK) property as in Section 3.4. The proof of the HVZK property we shall obtain relies on auxiliary isogeny generation oracles but is almost as rigorous as in R-SQIsignHD. As explained in Section 3.4.3, the Fiat-Shamir with hints framework introduced in [ABDPW25] gives an alternative and stronger security argument making the use of such oracles obsolete but we do not discuss it here.

4.2.1 Special soundness

We start by proving that the verification (Algorithm 4.2) works as expected, meaning that a response that has been validated (the algorithm returns True) always efficiently represents an isogeny $\varphi : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$ of degree $\deg(\varphi) < 2^{e_{\mathsf{rsp}}}$.

Lemma 4.2.1. *Let $(E_{com}, \mathsf{chl}, \mathsf{rsp})$ be a transcript of the SQIsign2D-West identification protocol. If the response $\mathsf{rsp} = (E_{aux}, P_{chl}, Q_{chl}, r, n_{\mathsf{bt}})$ has been validated by the verifier* i.e. *if Algorithm 4.2 returned* True, *then rsp is an efficient representation of an isogeny $\varphi : E_{com} \longrightarrow E_{chl}$ of degree $\deg(\varphi) < 2^{e_{rsp}}$.*

Proof. Consider the factorisation $\varphi_{\mathsf{chl}} = \varphi_{\mathsf{bt}} \circ \varphi'_{\mathsf{chl}}$ of the challenge isogeny, where $\varphi_{\mathsf{bt}} : E^{(0)}_{\mathsf{chl}} \longrightarrow E_{\mathsf{chl}}$ is a $2^{n_{\mathsf{bt}}}$-isogeny (the challenge part backtracking through the response). If Algorithm 4.2 has validated the response $\mathsf{rsp} = (E_{\mathsf{aux}}, P_{\mathsf{chl}}, Q_{\mathsf{chl}}, r, n_{\mathsf{bt}})$, then $P_{\mathsf{chl}}, Q_{\mathsf{chl}} \in E^{(0)}_{\mathsf{chl}}$ and these points yield a 2^r-isogeny

$\widehat{\varphi}^{(0)} : E_{\mathsf{chl}}^{(0)} \longrightarrow E'_{\mathsf{chl}}$ and a 2-dimensional 2^f-isogeny $\widetilde{\Phi} : E'_{\mathsf{chl}} \times E_{\mathsf{aux}} \longrightarrow E_{\mathsf{com}} \times E'_{\mathsf{aux}}$ (with $f := e_{\mathsf{rsp}} - n_{\mathsf{bt}} - r$).

By the converse of Kani's lemma (Lemma 2.2.7), $\widetilde{\Phi}$ can be written as a matrix

$$\widetilde{\Phi} = \begin{pmatrix} \phi_{1,1} & -\phi_{1,2} \\ \phi_{2,1} & \phi_{2,2} \end{pmatrix},$$

where $\widehat{\phi}_{1,1} \circ \phi_{1,1} = [a]$, $\widehat{\phi}_{2,2} \circ \phi_{2,2} = [a]$, $\widehat{\phi}_{2,1} \circ \phi_{2,1} = [b]$, $\widehat{\phi}_{1,2} \circ \phi_{1,2} = [b]$ and $a + b = 2^f$. Besides, we cannot have $a = 0$, otherwise we would have $\phi_{1,1} = 0$, $\phi_{2,2} = 0$ and:

$$\ker(\widetilde{\Phi}) = \ker(\phi_{2,1}) \times \ker(\phi_{1,2}).$$

Since $[4]T_1 = (*, [2^{r+2}]P_{\mathsf{aux}})$ and $[4]T_2 = (*, [2^{r+2}]Q_{\mathsf{aux}})$ belong to $\ker(\widetilde{\Phi})$, where T_1, T_2 have been defined in Line 15, with $(P_{\mathsf{aux}}, Q_{\mathsf{aux}})$ a basis of $E_{\mathsf{aux}}[2^{e_{\mathsf{rsp}}-n_{\mathsf{bt}}+2}]$, it follows that $\ker(\phi_{1,2})$ contains $E_{\mathsf{aux}}[2^f]$, so $b = \deg(\phi_{1,2}) \geq 2^{2f}$. Contradiction. So $a \neq 0$ and $\phi_{1,1}$ is an isogeny $E'_{\mathsf{chl}} \longrightarrow E_{\mathsf{com}}$ of degree $0 < a < 2^f$. By Lemma 3.4.5 an efficient representation of an isogeny yields an efficient representation of its dual. It follows that rsp yields an efficient representation of $\varphi := \varphi_{\mathsf{bt}} \circ \varphi^{(0)} \circ \widehat{\phi}_{1,1} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$. Its degree is $a2^{r+n_{\mathsf{bt}}} < 2^{f+r+n_{\mathsf{bt}}} = 2^{e_{\mathsf{rsp}}}$. This completes the proof. □

We also prove an analogue of Lemma 3.4.9 from which the special soundness argument directly follows.

Lemma 4.2.2. *Let $\varphi_{chl} : E_{\mathsf{pk}} \longrightarrow E_{chl}$ and $\varphi'_{chl} : E_{\mathsf{pk}} \longrightarrow E'_{chl}$ be two distinct challenges from the same public curve E_{pk}. Then, the largest integer dividing $\varphi'_{chl} \circ \hat{\varphi}_{chl} \in \mathrm{Hom}(E_{chl}, E'_{chl})$ is strictly smaller than $2^{e_{chl}}$.*

Proof. The proof is almost identical to Lemma 3.4.9. By Lemma 3.4.8, φ_{chl} and φ'_{chl} admit a greatest cyclic factor and we can use the structure of kernel generators of φ_{chl} and φ'_{chl} to bound its degree. □

Theorem 4.2.3. *Assume that $e_{chl} + e_{rsp} \leq e$. Then the SQIsign2D-West identification protocol is special sound for the NP-relation $\mathcal{R}$ defined in Eq. (3.9).*

Proof. This is the same argument used in the proof of Theorem 3.4.10. If $(E_{\mathsf{com}}, \mathsf{chl}, \mathsf{rsp})$ and $(E_{\mathsf{com}}, \mathsf{chl}', \mathsf{rsp}')$ are two transcripts for the same public key E_{pk} and commitment E_{com} but distinct challenges $\mathsf{chl} \neq \mathsf{chl}'$, then we consider the endomorphism:

$$\alpha := \widehat{\varphi'}_{\mathsf{chl}} \circ \varphi'_{\mathsf{rsp}} \circ \widehat{\varphi}_{\mathsf{rsp}} \circ \varphi_{\mathsf{chl}} \in \mathrm{End}(E_{\mathsf{pk}}).$$

By Lemma 4.2.1 and Lemma 3.4.5, the two transcripts give access to an efficient representation of α. Then, using Lemma 4.2.2, we can prove that α is non scalar when $e_{\mathsf{chl}} + e_{\mathsf{rsp}} \leq e$. □

4.2.2 The zero knowledge property

To prove that the protocol has the zero-knowledge property, we prove that there exists a simulator producing transcripts indistinguishable from an honest run of the protocol. Like in SQIsignHD (see Section 3.4.2), the simulator runs in polynomial time if it has access to an oracle producing random isogenies. This "random isogeny" oracle comes in two variants: the uniform target oracle (UTO) and the fixed degree isogeny oracle (FIDIO). Note that the zero knowledge property does not rely on any heuristic except the access to a UTO and a FIDIO (unlike F-SQIsignHD). Furthermore, in Section 4.2.3, we justify that a FIDIO can generate a RADIO as defined in Definition 3.4.14 and that the UTO is heuristically redundant. This means that the security proof in SQIsign2D-West is close to the rigorous security proof of the theoretical R-SQIsignHD variant optimised for security while being practical and efficient unlike the latter.

Definition 4.2.4. A *uniform target oracle* (UTO) is an oracle taking as input a supersingular elliptic curve E defined over $\mathbb{F}_{p^2}$ and an integer $N \geq 2\sqrt{2p}/\pi$, and outputs a random isogeny $\varphi : E \longrightarrow E'$ (in efficient representation) such that:

1. The distribution of E' is uniform among all the supersingular elliptic curves.
2. The conditional distribution of φ given E' is uniform among isogenies $E \longrightarrow E'$ of degree smaller or equal to N.

Remark 4.2.5. The condition $N \geq 2\sqrt{2p}/\pi$ ensures such an oracle exists: by Lemma 2.3.2, for any pair (E_1, E_2), the collection of isogenies $E_1 \longrightarrow E_2$ of degree smaller than N is non-empty.

Definition 4.2.6. A *fixed degree isogeny oracle* (FIDIO) is an oracle taking as input a supersingular elliptic curve E defined over $\mathbb{F}_{p^2}$ and an integer N, and outputting a uniformly random isogeny $\varphi : E \longrightarrow E'$ (in efficient representation) with domain E and degree N.

Theorem 4.2.7. *If $2^{e_{rsp}} \geq 2\sqrt{2p}/\pi$ and $N_{com} \geq 2^{4\lambda}$, then the SQIsign2D-West identification protocol is statistically honest-verifier zero-knowledge in the UTO and FIDIO model. In other words, there exists a polynomial time simulator $\mathcal{S}$ with access to a UTO and a FIDIO that produces random transcripts which are statistically indistinguishable from honest protocol transcripts.*

Remark 4.2.8. Note that we have chosen e_{rsp} in Section 4.1.4, so that $2^{e_{\mathsf{rsp}}} \geq 3^{13/6}\sqrt{p}/\pi$ in order to apply Algorithm 3.6. Consequently, the condition $2^{e_{\mathsf{rsp}}} \geq 2\sqrt{2p}/\pi$ is automatically satisfied.

Proof. The simulator proceeds as follows:

1. Generate an isogeny $\varphi_{\mathsf{chl}} : E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}$ according to the honest challenge distribution.
2. Call the UTO on input $(E_{\mathsf{chl}}, 2^{e_{\mathsf{rsp}}})$, resulting in the isogeny $\widehat{\varphi}_{\mathsf{rsp}} : E_{\mathsf{chl}} \longrightarrow E_{\mathsf{com}}$.
3. Decompose $\varphi_{\mathsf{rsp}} = \psi \circ \varphi_{\mathsf{rsp}}^{(1)}$ with $q' = \deg(\varphi_{\mathsf{rsp}}^{(1)})$ odd and $\deg(\psi) = 2^n$ a power of two. Let $2^{n_{\mathsf{bt}}} = \#(\ker(\hat{\psi}) \cap \ker(\hat{\varphi}_{\mathsf{chl}}))$. Let $r := n - n_{\mathsf{bt}}$ and $f := e_{\mathsf{rsp}} - n$.
4. Call the FIDIO on input $(E_{\mathsf{com}}, 2^f - q')$, resulting in the isogeny $\varphi_{\mathsf{aux}} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{aux}}$.

From the properties of the UTO and FIDIO, the above procedure returns transcripts with the same distribution as transcripts generated as follows:

1. Generate a uniformly random supersingular curve E_{com}
2. Generate an isogeny $\varphi_{\mathsf{chl}} : E_{\mathsf{pk}} \longrightarrow E_{\mathsf{chl}}$ according to the honest challenge distribution.
3. Generate a uniformly random isogeny $\varphi_{\mathsf{rsp}} : E_{\mathsf{com}} \longrightarrow E_{\mathsf{chl}}$, of degree at most $2^{e_{\mathsf{rsp}}}$.
4. Decompose $\varphi_{\mathsf{rsp}} = \psi \circ \varphi_{\mathsf{rsp}}^{(1)}$ with $q' = \deg(\varphi_{\mathsf{rsp}}^{(1)})$ odd and $\deg(\psi) = 2^n$ a power of two. Let $2^{n_{\mathsf{bt}}} = \#(\ker(\hat{\psi}) \cap \ker(\hat{\varphi}_{\mathsf{chl}}))$. Let $r := n - n_{\mathsf{bt}}$ and $f := e_{\mathsf{rsp}} - n$.
5. Generate a uniformly random isogeny φ_{aux} from E_{com} and of degree $2^f - q'$.

This is precisely the order in which an honest run of the protocol proceeds. The distribution for the first step matches the honest protocol run by Proposition 3.3.1. The distributions of following steps match the honest ones by construction. □

4.2.3 On the UTO and FIDIO oracles

Let us first argue that the UTO is essentially redundant: given a FIDIO, one can implement an oracle that is computationally indistinguishable from a UTO, at least when the bound N is sufficiently large. We proceed in two steps:

1. First, we use the FIDIO to build an oracle which outputs a uniform isogeny σ from E with $\deg(\sigma) \leq N$. In other words, one can turn a FIDIO into a RADIO, as defined in Definition 3.4.14.
2. Second, we argue that this distribution (the output of a RADIO) is indistinguishable from the output of a UTO.

Recall the definition of a RADIO.

Definition 4.2.9. A *random any-degree isogeny oracle* (RADIO) is an oracle taking as input a supersingular elliptic curve E defined over $\mathbb{F}_{p^2}$ and an integer N, and outputting a uniformly random isogeny $\varphi : E \longrightarrow E'$ (in efficient representation) with domain E and degree at most N.

Let us first explain how one can turn a FIDIO into a RADIO. Let f_N be the probability distribution of the degree of the output of a RADIO: for any integer q, let $f_N(q)$ be the probability that the degree of the output of a RADIO on input (E, N) is equal to q. Note that conditional on the degree of the output begin q, the FIDIO and the RADIO follow the same distribution: uniform among isogenies with domain E and degree q. Therefore, to simulate a RADIO, we can proceed as follows: on input (E, N),

1. sample an integer q following the distribution f_N;
2. call the FIDIO on input (E, q), and return the output.

To sample from the distribution f_N, observe that the value $f_N(q) = \widetilde{\Theta}(q/N^2)$ can be computed efficiently if the factorisation of q is known. Therefore, we can do rejection sampling by sampling uniformly random integers in $[\![1\ ;\ N]\!]$ *together with their factorisation* (see [Bac88]).

Now that we can turn a FIDIO into a RADIO, it remains to argue that a RADIO is indistinguishable from a UTO. For N large enough, it is indeed statistically indistinguishable: conditionally on the target curve, the two distributions are identical, and it is proven in Theorem 3.4.15 that when $N = \Theta(p^{1+\varepsilon})$ for $\varepsilon \in]0, 2]$, the distribution on the target curves are at statistical distance $O(p^{-\varepsilon/2})$. Therefore, when $N = \Theta(p^{1+\varepsilon})$, the RADIO and the UTO are at statistical distance $O(p^{-\varepsilon/2})$. The bound $N = O(p^{1/2})$ used in the protocol is not large enough for this theorem to apply, but we expect the distributions to remain computationally indistinguishable.

The conclusion of the above discussion is that in Theorem 4.2.7, the UTO is heuristically redundant. In other words, there is a (heuristic) simulator in the FIDIO model. It remains to argue that this FIDIO does not hurt the security assumption: access to a FIDIO does not help with solving the endomorphism ring problem. We refer to the analogous discussion about the security of SQIsignHD in Section 3.4.3. In essence, all a FIDIO does is compute a random walk from a source curve. We already know how to compute random walks of smooth degree (by taking a sequence of random isogeny steps of small prime degree), and a FIDIO extends this capability to random walks with potentially large prime steps.

4.3 Instantiation and performance

4.3.1 Parameter choices and signature sizes

The choice of prime $p = c2^e - 1$ is the same as in F-SQIsignHD (see Section 3.5.1 and Table 3.1 in particular) since SQIsign2D-West has the same security and torsion requirements. In particular, p has size close to but slightly smaller than 2λ bits in order to ensure λ bits of security against endomorphism ring attacks while fitting into $\lambda/32$ words of length 64 bits. We now explain in more detail the shape of the public key and signature.

As in F-SQIsignHD, the public key is a Montgomery curve $E_{\mathsf{pk}} : y^2 = x^3 + A_{\mathsf{pk}}x^2 + x$ represented by its Montgomery coefficient $A_{\mathsf{pk}} \in \mathbb{F}_{p^2}$. Since p has size 2λ, A_{pk} can be represented by 4λ bits. Some data can be added to the public key in order to speed up the verification. Indeed, the verifier has to generate a deterministic basis $(P_{\mathsf{pk}}, Q_{\mathsf{pk}})$ of $E_{\mathsf{pk}}[2^e]$. This basis is determined by some small integer $h \in \mathbb{N}$ called a *hint*. Giving this hint to the verifier speeds up the basis generation. We refer to [AAA+25, Algorithm 2.1 and 2.2] for more details on the deterministic basis generation with hints. It has been conjectured that the hint h can be represented with one byte (8 bits) with overwhelming probability. Hence, the total public key size is $4\lambda + 8$ bits.

Now, the signature is a bit different than in F-SQIsignHD where it was of the form $(E_{\mathsf{com}}, \mathsf{rsp})$, following the traditional Fiat-Shamir transform. Recall that SQIsign2D-West is a commitment recoverable scheme in order to allow for smaller signature. This means E_{com} does not have to be included in the signature and can be replaced by the challenge $\mathsf{chl} \in [\![0\ ;\ 2^{e_{\mathsf{chl}}} - 1]\!]$ which results from the application of a hash function $\mathsf{chl} = H(j(E_{\mathsf{pk}}), j(E_{\mathsf{com}}), m)$, where m is the message to be signed. The verifier will recover the commitment E_{com} during the verification process and will be able to check the

equality $\mathsf{chl} = H(j(E_{\mathsf{pk}}), j(E_{\mathsf{com}}), m)$. Given that E_{com} takes 4λ bits to store and e_{chl} is smaller but close to λ, this method saves 3λ bits on the signature.

Furthermore, since the verifier will also have to generate a deterministic basis of a torsion subgroup of E_{aux}, some hint h_{aux} has to be added to the signature in order to speed-up the verification. As a consequence, the signature is of the form $(\mathsf{chl}, E_{\mathsf{aux}}, P_{\mathsf{chl}}, Q_{\mathsf{chl}}, r, n_{\mathsf{bt}}, h_{\mathsf{aux}})$. As we have seen, chl takes $e_{\mathsf{chl}} \simeq \lambda$ bits to store, E_{aux} is stored as a Montgomery coefficient $A_{\mathsf{aux}} \in \mathbb{F}_{p^2}$ which takes 4λ bits. We also have $n_{\mathsf{bt}}, r \leq e_{\mathsf{rsp}}$ and by Remark 4.2.8, we may choose $e_{\mathsf{rsp}} = \lceil \log_2(3^{13/6}\sqrt{p}/\pi) \rceil \leq \lambda + 2$. Hence, n_{bt} and r take $\lceil \log_2(\lambda) \rceil$ bits to store each.

In addition, $P_{\mathsf{chl}}, Q_{\mathsf{chl}} \in E^{(0)}_{\mathsf{chl}}[2^{e_{\mathsf{rsp}}-n_{\mathsf{bt}}+2}]$ with $e_{\mathsf{rsp}} \leq \lambda + 2$. As explained in Section 3.5.1, $P_{\mathsf{chl}}, Q_{\mathsf{chl}}$ may be represented by 4 coefficients in $\mathbb{Z}/2^{e_{\mathsf{rsp}}-n_{\mathsf{bt}}+2}\mathbb{Z}$ determining their coordinates in a deterministic basis of $E^{(0)}_{\mathsf{chl}}[2^{e_{\mathsf{rsp}}-n_{\mathsf{bt}}+2}]$. Hence, $P_{\mathsf{chl}}, Q_{\mathsf{chl}}$ can be represented with $4(\lambda + 2 + 2) = 4\lambda + 16$ bits and some hint h_{chl} accelerating the deterministic basis generation, for a total of $4\lambda + 24$ bits. As in Section 3.5.1, we could gain $\lambda + 2$ bits by forgetting one integer coefficient at the expense of a Weil pairing computation during the verification. This optimisation was not considered in order to optimise verification time.

On the whole, the signature takes $9\lambda + 2\lceil \log_2(\lambda) \rceil + 32$ bits to store, including λ bits for chl, 4λ bits for E_{aux}, $4\lambda + 24$ bits for $(P_{\mathsf{chl}}, Q_{\mathsf{chl}})$, $2\lceil \log_2(\lambda) \rceil$ bits for r, n_{bt} and 8 bits for h_{aux}. The parameters, public key and signature sizes may be found in Table 4.1. Note that SQIsign2D-West is less compact than F-SQIsignHD but still 15% more compact than SQIsign (see Table 4.2).

Security level λ	NIST-I 128	NIST-III 192	NIST-V 256
$p = c \cdot 2^e - 1$	$5 \cdot 2^{248} - 1$	$65 \cdot 2^{376} - 1$	$27 \cdot 2^{500} - 1$
Signature size (bytes*)	150	222	294
Public key size (bytes*)	65	97	129

Table 4.1: Prime parameter, signature and public key sizes for different security levels in SQIsign2D-West. *One byte contains 8 bits.

Security level	NIST-I	NIST-III	NIST-V
SQIsign (bytes)	177	263	335
F-SQIsignHD (bytes)	108	160	212
SQIsign2D-West (bytes)	**150**	**222**	**294**

Table 4.2: Comparison of signature sizes (in bytes) in SQIsign, F-SQisignHD and SQIsign2D-West.

4.3.2 Performance

SQIsign2D-West and its heuristic version H-SQIsign2D-West have been implemented in C. The code repository may be found at https://github.com/SQISign/sqisign2d-west-ac24. This code was then used as reference and further optimised for the SQIsign round 2 NIST submission. The NIST v 2.0 repository may be found at https://github.com/SQIsign/the-sqisign/tree/nist-v2 and we refer to [AAA+25] for the associated specification. Timings on an Intel Core i5-1335U 4600MHz CPU are displayed in Table 4.3.

We observe that SQIsign2D-West verification times are very competitive and respectively 3.1 and 24.3 times faster than SQIsign original verification time at NIST-I and NIST-V levels (see Table 4.3). The NIST v 2.0 version with optimised finite field arithmetic even reaches 1 ms verification for NIST-I level. However, SQIsign2D-West signing is 6.7 times slower than F-SQIsignHD signing at NIST-I level. This gap is reduced by a factor two in the heuristic H-SQIsign2D-West and NIST v 2.0 optimised version. Furthermore, SQIsign2D-West signing is still respectively 9.5 and 99.9 times faster than the original SQIsign NIST submission at NIST-I and NIST-V levels. Even though it is three orders of magnitudes slower than lattice based or code based competitors, with such performance and

given its compactness, SQIsign NIST v 2.0 derived from SQIsign2D-West is a credible candidate for standardisation and industrial deployment.

Security level			NIST-I	NIST-III	NIST-V
SQIsign NIST v 1.0	Key generation	ms	355.72	5 625.72	22 445.3
		10^6 cycles	889.29	14 064.29	56 113.25
	Signing	ms	554.78	10 553.18	41 322.21
		10^6 cycles	1 386.95	26 382.94	103 305.53
	Verification	ms	7.77	195.86	571.77
		10^6 cycles	19.43	489.65	1 429.42
F-SQIsignHD	Key generation	ms	14.85	48.5	112.31
		10^6 cycles	37.12	121.29	280.92
	Signing	ms	8.74	25.68	56.72
		10^6 cycles	21.83	64.24	141.86
SQIsign2D-West	Key generation	ms	16.53	52.24	113.18
		10^6 cycles	41.37	130.65	283.04
	Signing	ms	58.17	220.26	413.46
		10^6 cycles	145.52	551.16	1 034.52
	Verification	ms	2.53	9.77	23.57
		10^6 cycles	6.32	24.45	58.93
H-SQIsign2D-West	Key generation	ms	14.84	47.44	107.51
		10^6 cycles	37.13	118.67	268.85
	Signing	ms	24.17	70.25	170.69
		10^6 cycles	60.48	175.64	426.78
	Verification	ms	2.52	8.61	23.16
		10^6 cycles	6.30	21.51	57.90
SQIsign NIST v 2.0	Key generation	ms	10.63	32.05	51.37
		10^6 cycles	26.59	80.13	128.43
	Signing	ms	24.53	74.20	126.72
		10^6 cycles	61.33	185.51	316.80
	Verification	ms	1.13	4.10	8.49
		10^6 cycles	2.83	10.26	21.22

Table 4.3: Key generation, signing and verification times of different versions of SQIsign on an Intel Core i5-1335U 4600MHz CPU. The NIST v 2.0 version based on SQIsign2D-West was implemented with an assembly optimised finite field arithmetic for Intel processors. F-SQIsignHD verification has not been implemented in C so verification times were not displayed for this scheme.

Part II

Fast computation of higher dimensional isogenies with the theta model

Chapter 5

Introduction to the algebraic theory of theta functions

Throughout this chapter, k will be an algebraically closed field, unless explicitly stated otherwise. In Section 1.4, we have seen that abelian varieties are projective. The goal of this chapter is to define systems of projective coordinates that are convenient to do arithmetic. We shall use them to compute isogenies in Chapter 6. The exposition of this chapter follows from Mumford's introductory article to the algebraic theory of theta functions [Mum66] and Damien Robert's PhD thesis [Rob10]. Essentially no original result will be presented here. Nonetheless, some proofs are more detailed than in the literature.

If $(A, \varphi_{\mathcal{L}})$ is a polarised abelian variety defined over k, and $\mathcal{L}$ is generated by global sections $s_0, \cdots, s_n$ (in the sense of Definition 1.4.8), these global sections define a map $A \longrightarrow \mathbb{P}^n_k$ Theorem 1.4.9. When $\mathcal{L}$ is very ample, this map is an embedding and $s_0, \cdots, s_n$ define coordinates on A. In Section 5.1, we define the *theta group* $G(\mathcal{L})$ and its action on global sections. We also define families of global sections called *theta functions* that behave nicely under the action of $G(\mathcal{L})$, ensuring convenient arithmetic properties. There are several choices of such families determined by choices of *theta structures*.

In Section 5.2, we study how theta groups, theta structures and theta functions are related via isogenies. We obtain an isogeny evaluation formula. In Section 5.3, we present *symmetric theta structures* that ensure even nicer arithmetic properties than generic ones. With such theta structures, we obtain differential addition and duplication formulas but also change of level that can be used to compute isogenies along with change of theta coordinate formulas (introduced later in Section 6.2).

5.1 Theta structures

Throughout this section, A will be an abelian variety over k.

5.1.1 The theta group

Let $\mathcal{L}$ be a line bundle on A. Recall the definition of the subgroup $K(\mathcal{L}) \subseteq A(k)$ formed by elements $x \in A(k)$ such that $t_x^*\mathcal{L} \simeq \mathcal{L}$.

Definition 5.1.1 (Theta group)**.** The *theta group* of $\mathcal{L}$ on A is the set of couples (x, ϕ_x), where $x \in A(k)$ and ϕ_x is an isomorphism $\mathcal{L} \xrightarrow{\sim} t_x^*\mathcal{L}$. It is denoted by $G(\mathcal{L})$.

The group structure on $G(\mathcal{L})$ is given by $(x, \phi_x) \cdot (y, \phi_y) := (x + y, t_x^*\phi_y \circ \phi_x)$, where $t_x^*\phi_y$ is the map $t_x^*\mathcal{L} \longrightarrow t_x^*(t_y^*\mathcal{L}) = t_{x+y}^*\mathcal{L}$ induced by $\phi_y : \mathcal{L} \longrightarrow t_y^*\mathcal{L}$.

Lemma 5.1.2. *If $\mathcal{L}$ is a line bundle on A, we have an exact sequence:*

$$1 \longrightarrow k^* \xrightarrow{\iota_{\mathcal{L}}} G(\mathcal{L}) \xrightarrow{\rho_{\mathcal{L}}} K(\mathcal{L}) \longrightarrow 0$$

where $\iota_{\mathcal{L}} : \lambda \in k^* \longmapsto (0, \lambda \cdot id_{\mathcal{L}})$ and $\rho_{\mathcal{L}} : (x, \phi) \in G(\mathcal{L}) \longmapsto x \in K(\mathcal{L})$ is the forgetful map.

Proof. $\rho_{\mathcal{L}}$ is surjective by the definition of $G(\mathcal{L})$ and $K(\mathcal{L})$. It remains to prove that the group of automorphisms $\phi : \mathcal{L} \xrightarrow{\sim} \mathcal{L}$ is isomorphic to k^*. Indeed, we have the classical result $\mathrm{Hom}_{\mathcal{O}_A}(\mathcal{L}, \mathcal{L}) \simeq \mathcal{O}_A$ [GW10, p. 7.5.7] and that $\Gamma(A, \mathcal{O}_A) = k$ since A is projective and k is algebraically closed [Har77, Theorem I.3.4]. It follows that $\mathrm{Aut}(\mathcal{L}) = \mathrm{Hom}_{\mathcal{O}_A}(\mathcal{L}, \mathcal{L})^* \simeq k^*$. □

5.1.2 Descending theta groups

In this section, we study how isogenies relate to theta groups. Let $f : A \longrightarrow B$ be a separable isogeny with kernel $K := \ker(f)$, $\mathcal{L}$ and $\mathcal{M}$ be separable ample line bundles on A and B respectively such that $\mathcal{L} \simeq f^*\mathcal{M}$ and α be an isomorphism $f^*\mathcal{M} \xrightarrow{\sim} \mathcal{L}$. Then, for all $x \in K$, $t_x^*\alpha$ induces an isomorphism

$$t_x^* f^*\mathcal{M} = (f \circ t_x)^*\mathcal{M} = f^*\mathcal{M} \xrightarrow{\sim} t_x^*\mathcal{L},$$

so $t_x^*\alpha \circ \alpha^{-1} : \mathcal{L} \longrightarrow t_x^*\mathcal{L}$ is well defined and is an isomorphism, so that $x \in K(\mathcal{L})$ and $(x, t_x^*\alpha \circ \alpha^{-1}) \in G(\mathcal{L})$. The subset $\widetilde{K} := \{(x, t_x^*\alpha \circ \alpha^{-1}) \mid x \in K\} \subseteq G(\mathcal{L})$ is a subgroup and the restriction of the forgetful map $\rho_{\mathcal{L}} : G(\mathcal{L}) \longrightarrow K(\mathcal{L})$ to $\widetilde{K}$ induces an isomorphism $\widetilde{K} \xrightarrow{\sim} K$.

Definition 5.1.3. For any subgroup $K \subseteq K(\mathcal{L})$, a *level subgroup above* K is a subgroup $\widetilde{K} \subseteq G(\mathcal{L})$ isomorphic to K via the forgetful map $\rho_{\mathcal{L}} : G(\mathcal{L}) \longrightarrow K(\mathcal{L})$.

Given a kernel $K = \ker(f)$, we have multiple choices of level subgroups $\widetilde{K} \subset G(\mathcal{L})$ which are determined by a choice of isomorphism $f^*\mathcal{M} \xrightarrow{\sim} \mathcal{L}$. Conversely, a choice of level subgroup determines $f^*\mathcal{M} \xrightarrow{\sim} \mathcal{L}$, so we have a correspondence.

Theorem 5.1.4 (Grothendieck's descent theorem [Mum74, Theorem 2, p. 231])**.** *Given a separable ample line bundle $\mathcal{L}$ on A and a separable isogeny $f : A \longrightarrow B$ of kernel $K \subseteq K(\mathcal{L})$, there is a one to one correspondence between level subgroups $\widetilde{K}$ of $G(\mathcal{L})$ above K and couples $(\mathcal{M}, \alpha)$, where $\mathcal{M}$ is a line bundle of B such that $f^*\mathcal{M} \simeq \mathcal{L}$ and $\alpha : f^*\mathcal{M} \xrightarrow{\sim} \mathcal{L}$ is an isomorphism of $\mathcal{O}_A$-modules. Explicitly, this correspondence maps $(\mathcal{M}, \alpha)$ to*

$$\widetilde{K} := \{(x, t_x^*\alpha \circ \alpha^{-1}) \mid x \in K\},$$

defined as above.

Proposition 5.1.5. *[Mum66, § 1, Proposition 2]Let $f : A \longrightarrow B$ be a separable isogeny with kernel $K := \ker(f)$, $\mathcal{L}$ and $\mathcal{M}$ be separable ample line bundles on A and B respectively such that $\mathcal{L} \simeq f^*\mathcal{M}$ and $\widetilde{K} \subset G(\mathcal{L})$ be a level subgroup above K. Then we have:*

(i) $f^{-1}(K(\mathcal{M})) \subseteq K(\mathcal{L})$.

(ii) The centralizer of $\widetilde{K}$ in $G(\mathcal{L})$ is $Z(\widetilde{K}) = \rho_{\mathcal{L}}^{-1} f^{-1}(K(\mathcal{M}))$.

(iii) $G(\mathcal{M})$ is canonically isomorphic to $Z(\widetilde{K})/\widetilde{K}$.

Proof. Let $\alpha : f^*\mathcal{M} \xrightarrow{\sim} \mathcal{L}$ be the isomorphism associated to $\widetilde{K}$ by Theorem 5.1.4. Let $x \in f^{-1}(K(\mathcal{M}))$ and $y = f(x) \in K(\mathcal{M})$. By the definition of $K(\mathcal{M})$, there exists an isomorphism $\psi : \mathcal{M} \xrightarrow{\sim} t_y^*\mathcal{M}$, so $f^*\psi$ is an isomorphism $f^*\mathcal{M} \xrightarrow{\sim} f^*t_y^*\mathcal{M}$. Furthermore,

$$f^*t_y^*\mathcal{M} = (t_y \circ f)^*\mathcal{M} = (f \circ t_x)^*\mathcal{M} = t_x^* f^*\mathcal{M}.$$

It follows that $t_x^*\alpha \circ f^*\psi \circ \alpha^{-1}$ is a well defined isomorphism $\mathcal{L} \xrightarrow{\sim} t_x^*\mathcal{L}$, and that $x \in K(\mathcal{L})$. This proves (i).

From this and (ii), we easily get (iii). Indeed, we have seen that $(x, t_x^*\alpha \circ f^*\psi \circ \alpha^{-1}) \in G(\mathcal{L})$ when $(f(x), \psi) \in G(\mathcal{M})$. Since f is surjective (as any isogeny), ψ is entirely determined by $f^*\psi$ (locally, hence globally), hence by $t_x^*\alpha \circ f^*\psi \circ \alpha^{-1}$, so we have a surjective map

$$\begin{array}{rccc} \alpha_f : H := \rho_{\mathcal{L}}^{-1} f^{-1}(K(\mathcal{M})) = \{(x, \phi_x) \mid f(x) \in K(\mathcal{M})\} & \longrightarrow & G(\mathcal{M}) \\ (x, t_x^*\alpha \circ f^*\psi \circ \alpha^{-1}) & \longmapsto & (f(x), \psi) \end{array} \tag{5.1}$$

which is a group homomorphism. If $\alpha_f(x,\phi_x) = (0, \mathrm{id}_{\mathcal{M}})$, then we have

$$\phi_x = t_x^*\alpha \circ f^*\mathrm{id}_{\mathcal{M}} \circ \alpha^{-1} = t_x^*\alpha \circ \mathrm{id}_{f^*\mathcal{M}} \circ \alpha^{-1} = t_x^*\alpha \circ \alpha^{-1}$$

so $(x,\varphi) \in \widetilde{K}$ and conversely, any element of $\widetilde{K}$ maps to $(0,\mathrm{id}_{\mathcal{M}})$ via α_f, so $\ker(\alpha_f) = \widetilde{K}$ and α_f induces a canonical isomorphism $G(\mathcal{M}) \simeq H/\widetilde{K}$. Hence, to prove (iii), it suffices to prove (ii) *i.e.* that $H = Z(\widetilde{K})$.

We obtain easily that $H \subseteq Z(\widetilde{K})$. Indeed, if $(x,\phi_x) \in H$ then we have seen that there exists an isomorphism $\psi : \mathcal{M} \xrightarrow{\sim} t_{f(x)}^*\mathcal{M}$ such that $\phi_x = t_x^*\alpha \circ f^*\psi \circ \alpha^{-1}$. Then, for all $(w,\phi_w) \in \widetilde{K}$, we have $\phi_w = t_w^*\alpha \circ \alpha^{-1}$ so, on the one hand

$$t_x^*\phi_w \circ \phi_x = t_{x+w}^*\alpha \circ t_x^*\alpha^{-1} \circ t_x^*\alpha \circ f^*\psi \circ \alpha^{-1} = t_{x+w}^*\alpha \circ f^*\psi \circ \alpha^{-1},$$

and on the other hand

$$t_w^*\phi_x \circ \phi_w = t_{x+w}^*\alpha \circ t_w^*f^*\psi \circ t_w^*\alpha^{-1}t_w^*\alpha \circ \alpha^{-1} = t_{x+w}^*\alpha \circ f^*t_{f(w)}^*\psi \circ \alpha^{-1} = t_x^*\phi_w \circ \phi_x$$

since $f(w) = 0$, so that (x,ϕ_x) and (w,ϕ_w) commute for all $(w,\phi_w) \in \widetilde{K}$ and $(x,\phi_x) \in Z(\widetilde{K})$.

Conversely if $(x,\phi_x) \in Z(\widetilde{K})$, then for all $(w,\phi_w) \in \widetilde{K}$, we have

$$(x+w, t_x^*\phi_w \circ \phi_x) = (x,\phi_x)\cdot(w,\phi_w) = (w,\phi_w)\cdot(x,\phi_x) = (x+w, t_w^*\phi_x \circ \phi_w).$$

Let D and E be divisors on A and B such that $\mathcal{L} \simeq \mathcal{L}(D)$ and $\mathcal{M} = \mathcal{L}(E)$ respectively. Then, for all $(y,\phi_y) \in G(\mathcal{L})$, $\phi_y : \mathcal{L} \xrightarrow{\sim} t_y^*\mathcal{L}$ is the multiplication by g_y^{-1} with $g_y \in k(A)$ such that $\mathrm{div}(g_y) = t_y^*D - D$. The equality $t_x^*\phi_w \circ \phi_x = t_w^*\phi_x \circ \phi_w$ is equivalent to

$$g_x \cdot (g_w \circ t_x) = g_w \cdot (g_x \circ t_w) \tag{5.2}$$

for all $(w,\phi_w) \in \widetilde{K}$.

Besides, we have $\mathcal{L} \simeq f^*\mathcal{M}$ so $D \sim f^*E$ and there exists $h \in k(A)$ such that $\mathrm{div}(h) = D - f^*E$, so that

$$\mathrm{div}(g_x) = t_x^*D - D = t_x^*f^*E - f^*E + t^*x\,\mathrm{div}(h) - \mathrm{div}(h) = f^*(t_{f(x)}^*E - E) + \mathrm{div}(h\circ t_x/h) \tag{5.3}$$

and for all $(w,\phi_w) \in \widetilde{K}$, $\mathrm{div}(g_w) = f^*(t_{f(w)}^*E - E) + \mathrm{div}(h\circ t_w/h) = \mathrm{div}(h\circ t_w/h)$, so there exists a constant $c_w \in k^*$ such that $g_w = c_w \cdot h \circ t_w/h$. Then Eq. (5.2) implies that $h_x := g_x \cdot h/h\circ t_x$ is invariant by translation by elements of K. Then, we can see h_x as a function $A \longrightarrow \mathbb{A}_k^1$ which is invariant by the action of K and since f is separable, we may apply Lemma 1.4.66 which ensures the existence of $s_x \in k(B)$ such that $h_x = s_x \circ f$. We then have $f^*\,\mathrm{div}(s_x) = f^*(t_{f(x)}^*E - E)$ by Eq. (5.3) so $\mathrm{div}(s_x) = t_{f(x)}^*E - E$ by surjectivity of f and the multiplication by s_x^{-1} is an isomorphism $\mathcal{M} \xrightarrow{\sim} t_{f(x)}^*\mathcal{M}$, so that $f(x) \in K(\mathcal{M})$ *i.e.* $(x,\phi_x) \in H$. This completes the proof. □

5.1.3 The commutator pairing

For $x,y \in K(\mathcal{L})$, let $\widetilde{x},\widetilde{y} \in G(\mathcal{L})$ be lifts of x,y respectively ($x = \rho_{\mathcal{L}}(\widetilde{x})$ and $y = \rho_{\mathcal{L}}(\widetilde{y})$). Since $\ker(\rho_{\mathcal{L}}) \simeq k^*$, $\widetilde{x}$ and $\widetilde{y}$ are defined up to a scalar. Hence,

$$e_{\mathcal{L}}(x,y) := \widetilde{x}\widetilde{y}\widetilde{x}^{-1}\widetilde{y}^{-1}$$

is independent of the choice of $\widetilde{x}$ and $\widetilde{y}$, so this formula defines a pairing $e_{\mathcal{L}} : K(\mathcal{L}) \times K(\mathcal{L}) \longrightarrow k^*$. This pairing is skew-symmetric and has values in k^* because its image is annihilated by $\rho_{\mathcal{L}}$ since $K(\mathcal{L})$ is abelian. We call it the *commutator pairing* of $\mathcal{L}$.

The commutator pairing satisfies the following convenient properties.

Proposition 5.1.6. *(i) If $f : A \longrightarrow B$ is a homomorphism of abelian varieties and $\mathcal{L}$ is a line bundle over B, then*

$$\forall x,y \in f^{-1}(K(\mathcal{L})), \quad e_{f^*\mathcal{L}}(x,y) = e_{\mathcal{L}}(f(x),f(y)).$$

(ii) For any line bundles $\mathcal{L}$ and $\mathcal{M}$ on A, we have:

$$\forall x, y \in K(\mathcal{L}) \cap K(\mathcal{M}), \quad e_{\mathcal{L}\otimes\mathcal{M}}(x,y) = e_{\mathcal{L}}(x,y)e_{\mathcal{M}}(x,y).$$

(iii) If $\mathcal{L} \in \mathrm{Pic}^0(A)$, then $e_{\mathcal{L}} = 1$.

(iv) If $\mathcal{L}$ and $\mathcal{M}$ are algebraically equivalent line bundles on A ($\mathcal{L}\otimes\mathcal{M}^{-1} \in \mathrm{Pic}^0(A)$), then $e_{\mathcal{L}} = e_{\mathcal{M}}$.

(v) If $\mathcal{L}$ is a line bundle on A, then for all $x \in K(\mathcal{L})$ and $y \in [n]^{-1}K(\mathcal{L})$,

$$e_{\mathcal{L}^n}(x,y) = e_{\mathcal{L}}(x,[n]y)$$

Proof. **(i)** Let $x, y \in f^{-1}(K(\mathcal{L}))$. Let $\tilde{z} := (f(x),\phi), \tilde{t} := (f(y),\psi) \in G(\mathcal{L})$ be lifts of x and y respectively. Then on the one hand, we have by the definition of $e_{\mathcal{L}}(f(x),f(y))$:

$$(0, e_{\mathcal{L}}(f(x),f(y))\mathrm{id}_{\mathcal{L}}) = \tilde{z}\tilde{t}\tilde{z}^{-1}\tilde{t}^{-1} = (0, \psi^{-1}\circ t^*_{f(y)}\phi^{-1}\circ t^*_{f(x)}\psi\circ\phi).$$

On the other hand, $\tilde{x} := (x, f^*\phi), \tilde{y} := (y, f^*\psi) \in G(f^*\mathcal{L})$ are lifts of x and y respectively, so we have by the definition of $e_{f^*\mathcal{L}}(x,y)$:

$$\begin{aligned}(0, e_{f^*\mathcal{L}}(x,y)\mathrm{id}_{f^*\mathcal{L}}) &= \tilde{x}\tilde{y}\tilde{x}^{-1}\tilde{y}^{-1} = (0, f^*\psi^{-1}\circ t^*_y f^*\phi^{-1}\circ t^*_x f^*\psi\circ f^*\phi)\\ &= (0, f^*(\psi^{-1}\circ t^*_{f(y)}\phi^{-1}\circ t^*_{f(x)}\psi\circ\phi)) = (0, f^*(e_{\mathcal{L}}(f(x),f(y))\mathrm{id}_{\mathcal{L}}))\\ &= (0, e_{\mathcal{L}}(f(x),f(y))\mathrm{id}_{f^*\mathcal{L}}).\end{aligned}$$

This proves (i).

(ii) Let $x, y \in K(\mathcal{L})\cap K(\mathcal{M})$ and $\tilde{x}_1 := (x,\phi_1), \tilde{y}_1 := (y,\psi_1)$ be lifts of x and y in $G(\mathcal{L})$ respectively and $\tilde{x}_2 := (x,\phi_2), \tilde{y}_2 := (y,\psi_2)$ be lifts of x and y in $G(\mathcal{M})$ respectively. Then $\tilde{x}_3 := (x, \phi_1\otimes\phi_2)$ and $\tilde{y}_3 := (y, \psi_1\otimes\psi_2)$ are lifts of x and y in $G(\mathcal{L}\otimes\mathcal{M})$ respectively. It follows that

$$\begin{aligned}(0, e_{\mathcal{L}\otimes\mathcal{M}}(x,y)\mathrm{id}_{\mathcal{L}\otimes\mathcal{M}}) &= (0, \psi_1^{-1}\otimes\psi_2^{-1}\circ t^*_y\phi_1^{-1}\otimes t^*_y\phi_2^{-1}\circ t^*_x\psi_1\otimes t^*_x\psi_2\circ\phi_1\otimes\phi_2)\\ &= (0, (\psi_1^{-1}\circ t^*_y\phi_1^{-1}\circ t^*_x\psi_1\circ\phi_1)\otimes(\psi_2^{-1}\circ t^*_y\phi_2^{-1}\circ t^*_x\psi_2\circ\phi_2))\\ &= (0, e_{\mathcal{L}}(x,y)e_{\mathcal{M}}(x,y)\mathrm{id}_{\mathcal{L}\otimes\mathcal{M}}).\end{aligned}$$

This proves (ii).

(iii) If $\mathcal{L} \in \mathrm{Pic}^0(A)$, then $K(\mathcal{L}) = A$ and $e_{\mathcal{L}}$ is a pairing $A\times A \longrightarrow k^*$. We can see $e_{\mathcal{L}}$ as a morphism of k-varieties $A\times A \longrightarrow \mathbb{G}_m := \mathrm{Spec}(k[T,T^{-1}])$. Since $A\times A$ is an abelian variety, it is complete and $\mathbb{G}_m$ is affine so $e_{\mathcal{L}}$ must be constant by [GW10, Corollary 12.67].

(iv) Let $\mathcal{L}$ and $\mathcal{M}$ be algebraically equivalent. Then $\mathcal{M} = \mathcal{L}\otimes\mathcal{N}$ with $\mathcal{N} \in \mathrm{Pic}^0(A)$ so $e_{\mathcal{M}} = e_{\mathcal{L}}\cdot e_{\mathcal{N}}$ by (ii) and $e_{\mathcal{N}} = 1$ by (iii), so $e_{\mathcal{M}} = e_{\mathcal{L}}$.

(v) By (ii), we get for all $x \in K(\mathcal{L})$ and $y \in [n]^{-1}K(\mathcal{L})$:

$$e_{\mathcal{L}^n}(x,y) = e_{\mathcal{L}}(x,y)^n = e_{\mathcal{L}}(x,[n]y).$$

□

Lemma 5.1.7. *Let $\mathcal{L}$ be an ample and separable line bundle and $n \in \mathbb{N}^*$ not divisible by $\mathrm{char}(k)$. Then $K(\mathcal{L}^n) = [n]^{-1}(K(\mathcal{L}))$ and $K(\mathcal{L}) = [n]K(\mathcal{L}^n)$.*

Proof. By the theorem of the square Theorem 1.4.19, we have for all $x \in A(k)$

$$\varphi_{\mathcal{L}^n}(x) = [t_x^*\mathcal{L}^n\otimes\mathcal{L}^{-n}] = [(t_x^*\mathcal{L}\otimes\mathcal{L}^{-1})^n] = [t^*_{[n]x}\mathcal{L}\otimes\mathcal{L}^{-1}] = \varphi_{\mathcal{L}}([n]x)$$

As a consequence, $K(\mathcal{L}^n) = \ker(\varphi_{\mathcal{L}^n}) = \ker(\varphi_{\mathcal{L}}\circ[n]) = [n]^{-1}\ker(\varphi_{\mathcal{L}}) = [n]^{-1}(K(\mathcal{L}))$. It follows that $[n]K(\mathcal{L}^n) \subseteq K(\mathcal{L})$ but the converse inclusion is also true since A is n-divisible (because $\mathrm{char}(k)\nmid n$). This completes the proof. □

Unsurprisingly, the commutator pairing is related to the Weil pairing.

Proposition 5.1.8. *If $\mathcal{L}$ is a line bundle on A and $\mathrm{char}(k) \nmid n$, then for all $x \in A[n]$ and $y \in [n]^{-1}(K(\mathcal{L}))$,*

$$e_n(x, \varphi_{\mathcal{L}}(y)) = e_{\mathcal{L}^n}(x, y),$$

where $e_n : A[n] \times \widehat{A}[n] \longrightarrow k^*$ is the n-th Weil pairing.

Proof. First, the equality makes sense because $[n]^{-1}K(\mathcal{L}) = \varphi_{\mathcal{L}}^{-1}(A[n])$, $A[n] \subseteq [n]^{-1}(K(\mathcal{L}))$ and $K(\mathcal{L}^n) = [n]^{-1}(K(\mathcal{L}))$ by Lemma 5.1.7. Now, let $x \in A[n]$ and $y \in [n]^{-1}(K(\mathcal{L}))$. Then, we may write $y = [n]z$ for some $z \in [n^2]^{-1}(K(\mathcal{L}))$ and we have by points (v) and (iv) of Proposition 5.1.6,

$$e_{\mathcal{L}^n}(x, y) = e_{\mathcal{L}^n}(x, [n]z) = e_{\mathcal{L}^{n^2}}(x, z) = e_{[n]^*\mathcal{L}}(x, z),$$

where the last equality comes from the fact that $[n]^*\mathcal{L}$ and $\mathcal{L}^{n^2}$ are algebraically equivalent. Indeed, by Corollary 1.4.18, we have $[n]^*\mathcal{L} \simeq \mathcal{L}^{n^2} \otimes ([-1]^*\mathcal{L} \otimes \mathcal{L}^{-1})^{n(n-1)/2}$ and by Lemma 1.4.65, we have $[-1]^*\mathcal{L} \otimes \mathcal{L}^{-1} \in \mathrm{Pic}^0(A)$ since it is antisymmetric. Hence, we have to prove that $e_{[n]^*\mathcal{L}}(x, z) = e_n(x, \varphi_{\mathcal{L}}(y))$.

Let D be a divisor such that $\mathcal{L} \simeq \mathcal{L}(D)$. By the definition of the Weil pairing (see Theorem 1.4.67), we then have $e_n(x, \varphi_{\mathcal{L}}(y)) = g_y/g_y \circ t_x$ with $g_y \in k(A)$ such that $\mathrm{div}(g_y) = [n]^*(t_y^*D - D)$.

Besides, $t_x^*[n]^*\mathcal{L} = [n]^*t_{[n]x}^*\mathcal{L} = [n]^*\mathcal{L}$ since $x \in A[n]$, so $(x, \mathrm{id}_{[n]^*\mathcal{L}}) \in G([n]^*\mathcal{L})$ lifts x. Let $(z, \phi_z) \in G([n]^*\mathcal{L})$ be a lift of z. Then, the composition map

$$[n]^*\mathcal{L} \overset{\mathrm{id}_{[n]^*\mathcal{L}}}{\longrightarrow} [n]^*\mathcal{L} \overset{\phi_z}{\longrightarrow} t_z^*[n]^*\mathcal{L} \overset{t_z^*\mathrm{id}_{[n]^*\mathcal{L}}}{\longrightarrow} t_z^*[n]^*\mathcal{L} \overset{t_x^*\phi_z^{-1}}{\longrightarrow} [n]^*\mathcal{L},$$

which equals $t_x^*\phi_z^{-1} \circ \phi_z$ is the multiplication by $e_{[n]^*\mathcal{L}}(x, z)$. But ϕ_z is the multiplication by h_z^{-1}, where $h_z \in k(A)$ satisfies

$$\mathrm{div}(h_z) = t_z^*[n]^*D - [n]^*D = [n]^*(t_{[n]z}^*D - D) = [n]^*(t_y^*D - D) = \mathrm{div}(g_y),$$

so that $h_z = c \cdot g_y$ for some $c \in k^*$ and $e_{[n]^*\mathcal{L}}(x, z) = h_z/h_z \circ t_x = g_y/g_y \circ t_x = e_n(x, \varphi_{\mathcal{L}}(y))$. This completes the proof. □

Lemma 5.1.9. *Let G be a finite abelian group and $e : G \times G \longrightarrow k^*$ be a non-degenerate skew-symmetric pairing. Then:*

(i) There exists a symplectic decomposition *of G,* i.e. *two subgroups $G_1, G_2 \subseteq G$ such that $G = G_1 \oplus G_2$ that are isotropic for e ($e(x, y) = 1$ for all $x, y \in G_i$ and $i = 1, 2$) and $G_2 \simeq \widehat{G}_1 := \mathrm{Hom}(G_1, k^*)$.*

(ii) There exists integers $d_1, \cdots, d_g \geq 2$ such that $d_1 | \cdots | d_g$ and for all tuple $(\zeta_1, \cdots, \zeta_g) \in (k^*)^g$ where ζ_i is a primitive d_i-th root of unity for all $i \in [\![1\ ;\ g]\!]$, there exists $x_1, \cdots, x_g, y_1, \cdots, y_g \in G$ such that

$$G = \langle x_1 \rangle \oplus \cdots \oplus \langle x_g \rangle \oplus \langle y_1 \rangle \oplus \cdots \oplus \langle y_g \rangle$$

and for all $i, j \in [\![1\ ;\ g]\!]$,

$$e(x_i, x_j) = e(y_i, y_j) = 1 \quad \text{and} \quad e(x_i, y_j) = \zeta_i^{\delta_{i,j}}.$$

*We say that $(x_1, \cdots, x_g, y_1, \cdots, y_g)$ is a $(\zeta_1, \cdots, \zeta_g)$-*symplectic basis *of G.*

Proof. (i) follows immediately form (ii) so we prove (ii). We proceed by strong induction on the cardinality of G. The result is immediate if $\#G = 1$.

Now, if $\#G > 1$ we may assume the result holds for finite abelian groups of cardinality $\leq \#G - 1$. Then G has a non-trivial exponent $d \geq 2$ and we may consider $x \in G$ of order d. Since e is non-degenerate, the character $e_x : y \in G \longmapsto e(x, y) \in k^*$ has order $k = d$. Indeed, we already know that $l | d$ since d is the exponent of G and $e_x^l = 1$, so that $e(lx, y) = 1$ for all $y \in G$, so $lx = 0$ since e is non-degenerate and $d | l$. It follows that there exists $y \in G$ such that $e(x, y)$ has order d, *i.e.* is a primitive d-th root of unity. Rescaling y by an integer coprime with d if necessary, we can choose the value $\zeta := e(x, y)$. We also immediately obtain that y has order d.

Consider the subgroup $H := \{z \in G \mid e(x, z) = e(y, z) = 1\}$. We prove that $G = \langle x \rangle \oplus H \oplus \langle y \rangle$. Then, applying the recursion hypothesis to H of cardinality $\#H = \#G/d^2 < \#G$ will complete the

proof. Let $\lambda, \mu \in \mathbb{Z}$ and $z \in H$ such that $\lambda x + \mu y + z = 0$. Then, we obtain that $e(x,y)^\mu = 1$ and $e(x,y)^\lambda = 1$, so $d|\mu$ and $d|\lambda$. Hence, $z = 0$. This proves that $\langle x \rangle, \langle y \rangle$ and H are in direct sum. Now, if $z \in G$, then $e(x,z)$ and $e(z,y)$ are d-th roots of unity so we may write $e(x,z) = e(x,y)^\mu$ and $e(z,y) = e(x,y)^\lambda$ for some $\lambda, \mu \in [\![0\ ;\ d-1]\!]$. Let $z' := z - \lambda x - \mu y$. Then,

$$e(x,z') = e(x,z)e(x,y)^{-\mu} = 1 \quad \text{and} \quad e(y,z') = e(y,z)e(y,x)^{-\lambda} = e(z,y)^{-1}e(x,y)^\lambda = 1,$$

so $z' \in H$. This completes the proof. □

Proposition 5.1.10. *Assuming $\mathcal{L}$ is ample and separable, $e_\mathcal{L}$ is non-degenerate.*

Proof. Let $\widetilde{K}$ be a maximal level subgroup of $G(\mathcal{L})$ and $K := \rho_\mathcal{L}(\widetilde{K})$. Then the centralizer of $\widetilde{K}$ is $Z(\widetilde{K}) = k^*\widetilde{K}$. Indeed, the center trivially contains $k^*\widetilde{K}$ and an element $z \notin k^*\widetilde{K}$ has image $\rho_\mathcal{L}(z) \notin K$ so if such an element commute with $\widetilde{K}$, the subgroup generated by z and $\widetilde{K}$ is a level subgroup above the group generated by $\rho_\mathcal{L}(z)$ and K, contradicting the maximality of $\widetilde{K}$.

We consider the isogeny $\pi : A \longrightarrow B := A/K$ of kernel K and $\mathcal{M} := \pi_*\mathcal{L}$ (so that $\mathcal{L} \simeq \pi^*\mathcal{M}$). Then, by Proposition 5.1.5, $G(\mathcal{M}) \simeq Z(\widetilde{K})/\widetilde{K} = k^*$ so $K(\mathcal{M}) = \{0\}$ by Lemma 5.1.2, so $\varphi_\mathcal{M}$ has degree 1. Since $\mathcal{L} \simeq \pi^*\mathcal{M}$, we get that $\varphi_\mathcal{L} = \widehat{\pi} \circ \varphi_\mathcal{M} \circ \pi$ by Lemma 1.4.60, so that

$$\#K(\mathcal{L}) = \deg(\varphi_\mathcal{L}) = \deg(\pi)^2 = (\#K)^2, \tag{5.4}$$

where we used the separability of $\mathcal{L}$ (hence of π), to obtain the first and last equalities.

Let $K_0 := \{x \in K(\mathcal{L}) \mid \forall y \in K(\mathcal{L}), \quad e_\mathcal{L}(x,y) = 1\}$ be the degenerate space of $e_\mathcal{L}$. Then, $e_\mathcal{L}$ induces a non-degenerate skew-symmetric pairing on $K(\mathcal{L})/K_0$ and by Lemma 5.1.9, it admits a symplectic decomposition $K(\mathcal{L})/K_0 = K_1 \oplus K_2$. We then have

$$\#(K(\mathcal{L})/K_0) = \#K_1 \cdot \#K_2 = (\#K_1)^2.$$

Besides, K_1 is maximal isotropic in $K(\mathcal{L})/K_0$, so its preimage K in $K(\mathcal{L})$ is maximal isotropic and we have by Eq. (5.4)

$$(\#K)^2 = \#K(\mathcal{L}) = \#(K(\mathcal{L})/K_0) \cdot \#K_0 = (\#K_1)^2 \cdot \#K_0 = (\#K/\#K_0)^2 \cdot \#K_0 = (\#K)^2/\#K_0$$

Hence, $\#K_0 = 1$ and $e_\mathcal{L}$ is non-degenerate. □

Corollary 5.1.11. *If $\mathcal{L}$ is ample and separable, the center of $G(\mathcal{L})$ is $\ker(\rho_\mathcal{L}) \simeq k^*$.*

Proof. If $g \in Z(G(\mathcal{L}))$ then, $e_\mathcal{L}(\rho_\mathcal{L}(g), y) = 1$ for all $y \in K(\mathcal{L})$ so $\rho_\mathcal{L}(g) = 0$ since $e_\mathcal{L}$ is non-degenerate. □

Not all subgroups $K \subseteq K(\mathcal{L})$ admit a level subgroup of $G(\mathcal{L})$ and define an isogeny. The commutator pairing gives a condition for that.

Proposition 5.1.12. *Let $K \subseteq K(\mathcal{L})$ be a subgroup. The following conditions are equivalent:*

(i) K admits a level subgroup of $G(\mathcal{L})$.

(ii) $\rho_\mathcal{L}^{-1}(K)$ is abelian.

(iii) K is isotropic for $e_\mathcal{L}$.

Proof. (ii) and (iii) are trivially equivalent by the definition of $e_\mathcal{L}$. Assume that (i) holds. Let $\widetilde{K} \subset G(\mathcal{L})$ be a level subgroup above K. Then, $\rho_\mathcal{L}^{-1}(K) = k^* \cdot \widetilde{K}$ and $\widetilde{K} \simeq K$ so $\widetilde{K}$ is abelian and $\rho_\mathcal{L}^{-1}(K)$ as well. Hence, (i) $\Longrightarrow$ (ii).

Conversely, assume that $\rho_\mathcal{L}^{-1}(K)$ is abelian. If we apply the theorem of finite abelian groups to K, we get a basis of K. $\rho_\mathcal{L}^{-1}(K)$ being abelian, it suffices to lift this basis to obtain a level subgroup, so we simply explain how to lift an element. Let $x \in K$ of order ℓ and $\widetilde{x} \in G(\mathcal{L})$ such that $\rho_\mathcal{L}(\widetilde{x}) = x$. Then, $\rho_\mathcal{L}(\widetilde{x}^\ell) = 1$ so $\alpha := \widetilde{x}^\ell \in k^*$. Let $\beta \in k^*$ be a ℓ-th root of α. Then, $\widetilde{x}/\beta$ has order ℓ and has image x via $\rho_\mathcal{L}$. Hence, (ii) $\Longrightarrow$ (i). This completes the proof. □

Using Grothendieck's descent theorem (Theorem 5.1.4) and the previous proposition, we obtain a characterisation of polarised isogenies by the isotropy of their kernel.

Corollary 5.1.13. *Let $f : (A, \mathcal{L}) \longrightarrow (B, \mathcal{M})$ be a polarised isogeny, as defined in Definition 1.4.58. Then $\ker(f)$ is isotropic for $e_{\mathcal{L}}$.*

Conversely, if $(A, \mathcal{L})$ is a polarised abelian variety and $f : A \longrightarrow B$ is an isogeny of kernel $\ker(f) \subseteq K(\mathcal{L})$ isotropic for $e_{\mathcal{L}}$, then there exists a line bundle $\mathcal{M}$ on B such that $f^*\mathcal{M} \simeq \mathcal{L}$, so that f is a polarised isogeny $(A, \mathcal{L}) \longrightarrow (B, \mathcal{M})$.

Proof. Assume that $f : (A, \mathcal{L}) \longrightarrow (B, \mathcal{M})$ is a polarised isogeny. Then Lemma 1.4.60 ensures that $f^*\mathcal{M} \simeq \mathcal{L} \otimes \mathcal{N}$ with $\mathcal{N} \in \mathrm{Pic}^0(A)$. By Grothendieck's descent theorem (Theorem 5.1.4), $\ker(f)$ admits a level subgroup in $G(\mathcal{L} \otimes \mathcal{N})$ so $\ker(f)$ is isotropic for $e_{\mathcal{L}\otimes\mathcal{N}}$ by Proposition 5.1.12. But $e_{\mathcal{L}\otimes\mathcal{N}} = e_{\mathcal{L}}$ by Proposition 5.1.6 so $\ker(f)$ is isotropic for $e_{\mathcal{L}}$.

Conversely, if $\ker(f) \subseteq K(\mathcal{L})$ is isotropic for $e_{\mathcal{L}}$, then there exists a level subgroup above $\ker(f)$ corresponding to a line bundle $\mathcal{M}$ on B and an isomorphism $\alpha : f^*\mathcal{M} \xrightarrow{\sim} \mathcal{L}$ by Grothendieck's descent theorem. This completes the proof. □

Definition 5.1.14 (Othogonality)**.** If $K \subseteq K(\mathcal{L})$ is a subgroup, we define the *orthogonal* of K by:

$$K^\perp := \{y \in K(\mathcal{L}) \mid \forall x \in K, \quad e_{\mathcal{L}}(x, y) = 1\}$$

This notion is different from the usual orthogonality because we may have $K \cap K^\perp \neq \{0\}$. It is especially the case when K is isotropic $K \subseteq K^\perp$.

Lemma 5.1.15. *Let $K \subseteq K(\mathcal{L})$ be a subgroup. Then:*

(i) $K^\perp \simeq \widehat{K(\mathcal{L})/K} := \mathrm{Hom}(K(\mathcal{L})/K, k^*)$.

(ii) $K^{\perp\perp} = K$.

(iii) K is maximal isotropic if and only if K is isotropic and $\#K(\mathcal{L}) = (\#K)^2$.

Proof. **(i)** Consider $y \in K^\perp \longmapsto e_{\mathcal{L}}(., y) \in \widehat{K(\mathcal{L})/K}$. This map is well-defined since $e_{\mathcal{L}}(x, y) = 1$ for all $x \in K$ and $y \in K^\perp$. If $y \in K^\perp$ satisfies $e_{\mathcal{L}}(x, y) = 1$ for all $x \in K(\mathcal{L})$, then $y = 0$ by non-degeneracy of $e_{\mathcal{L}}$ so the map is injective. Besides, if $\chi \in \widehat{K(\mathcal{L})/K}$, then χ induces a character $\widetilde{\chi} \in \widehat{K(\mathcal{L})}$ annihilating K and there exists $y \in K(\mathcal{L})$ such that $\widetilde{\chi} = e_{\mathcal{L}}(., y)$. Since $\widetilde{\chi}$ annihilates K, we must have $y \in K^\perp$, so the map is surjective and is an isomrophism $K^\perp \xrightarrow{\sim} \widehat{K(\mathcal{L})/K}$.

(ii) Since $\#\widehat{K(\mathcal{L})/K} = \#K(\mathcal{L})/K = \#K(\mathcal{L})/\#K$, we have $\#K^\perp = \#K(\mathcal{L})/\#K$ by (i). It follows that $\#K^{\perp\perp} = \#K(\mathcal{L})/\#K^\perp = \#K$. Furthermore, $K \subseteq K^{\perp\perp}$. Hence, $K^{\perp\perp} = K$.

(iii) Assume that K is isotropic. Then $K \subseteq K^\perp$ and K is maximal for this property if and only if $K = K^\perp$. But we have seen that $\#K(\mathcal{L}) = \#K \cdot \#K^\perp$ by (i), so $K = K^\perp$ if and only if $\#K(\mathcal{L}) = (\#K)^2$. This proves (iii). □

5.1.4 Theta structures

Previously, we have seen results on the structure of the theta group $G(\mathcal{L})$ (Lemma 5.1.2) and the subgroup $K(\mathcal{L})$ (Lemma 5.1.9) that give a purely canonical description of $G(\mathcal{L})$. In this section, we shall introduce this canonical description - called the *Heisenberg group* - and study isomorphisms between the Heisenberg group and the theta group - called *theta structures*.

The pairing $e_{\mathcal{L}}$ being non-degenerate, there exists by Lemma 5.1.9 a symplecitic decomposition $K(\mathcal{L}) = K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$, where $K_1(\mathcal{L})$ and $K_2(\mathcal{L})$ are maximal isotropic subgroups of $K(\mathcal{L})$ and $K_2(\mathcal{L}) \simeq \mathrm{Hom}(K_1(\mathcal{L}), k^*)$. By the finite abelian subgroups theorem, there exists $\delta := (d_1, \cdots, d_r) \in (\mathbb{N}^*)^r$ with $d_1 | \cdots | d_r$ such that

$$K_1(\mathcal{L}) \simeq K_2(\mathcal{L}) \simeq \prod_{i=1}^{r} \mathbb{Z}/d_i\mathbb{Z}.$$

It follows that $K(\mathcal{L}) \subseteq A[d_r]$. Since $A[d_r]$ has rank $2g$ by Corollary 1.4.34, we have $r \leq g$. We can assume $r = g$ and complete δ with ones if necessary.

Definition 5.1.16. When $K(\mathcal{L})$ can be decomposed as above, we say that $\mathcal{L}$ is *of type* δ. When $\delta = (n, \cdots, n)$ we say that $\mathcal{L}$ is of *level* n.

Let $K_1(\delta) := \prod_{i=1}^g \mathbb{Z}/d_i\mathbb{Z}$, $K_2(\delta) := \mathrm{Hom}(K_1(\delta), k^*)$ and $K(\delta) := K_1(\delta) \oplus K_2(\delta)$. Define a skew-symmetric pairing $e_\delta : K(\delta) \times K(\delta) \longrightarrow k^*$ by:

$$\forall (x_1, \chi_1), (x_2, \chi_2) \in K_1(\delta) \oplus K_2(\delta), \quad e_\delta((x_1, \chi_1), (x_2, \chi_2)) := \chi_2(x_1)\chi_1(x_2)^{-1}. \tag{5.5}$$

Proposition 5.1.17. *(i) $K(\mathcal{L})$ is symplectically isomorphic to $K(\delta)$,* i.e. *there exists an isomorphism $\phi : K(\delta) \xrightarrow{\sim} K(\mathcal{L})$ such that:*

$$\forall x, y \in K(\delta), \quad e_\delta(x, y) = e_{\mathcal{L}}(\phi(x), \phi(y)).$$

Such an isomorphism induces a symplectic decomposition $K(\mathcal{L}) = K_1(\phi) \oplus K_2(\phi)$ with $K_i(\phi) := \phi(K_i(\delta))$ for $i \in \{1, 2\}$.

(ii) Let $K(\mathcal{L}) = K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$ be a symplectic decomposition. Then there is a bijection between symplectic isomorphisms $K(\delta) \xrightarrow{\sim} K(\mathcal{L})$ mapping $K_i(\delta)$ to $K_i(\mathcal{L})$ for $i \in \{1, 2\}$ and isomorphisms $K_1(\delta) \xrightarrow{\sim} K_1(\mathcal{L})$. This bijection is given by the restriction to $K_1(\delta)$.

Proof. (i) follows easily from (ii) so we prove (ii). Let $\sigma : K_1(\delta) \xrightarrow{\sim} K_1(\mathcal{L})$ be an isomorphism. Let $(x_1, \cdots, x_g)$ be the canonical basis of $K_1(\delta)$ and $(\chi_1, \cdots, \chi_g)$ be its dual basis in $K_2(\delta)$ given by:

$$\forall i, j \in [\![1\ ;\ g]\!], \quad \chi_j(x_i) = e_\delta(x_i, \chi_j) = \zeta_i^{\delta_{i,j}},$$

where ζ_i is a δ_i-th primitive root of unity for all $i \in [\![1\ ;\ g]\!]$.

Consider the group homomorphism:

$$\Phi : y \in K_2(\mathcal{L}) \longmapsto (e_{\mathcal{L}}(\sigma(x_i), y))_{1 \leq i \leq g} \in \prod_{i=1}^g \mu_{d_i}(k),$$

where $\mu_{d_i}(k) \subset k^*$ is the subgroup of d_i-th roots of unity for all $i \in [\![1\ ;\ g]\!]$. If $y \in K_2(\mathcal{L})$ is such that $\Phi(y) = 1$, then $y \in \langle \sigma(x_1), \cdots, \sigma(x_g) \rangle^\perp = K_1(\mathcal{L})^\perp = K_1(\mathcal{L})$ since $(\sigma(x_1), \cdots, \sigma(x_g))$ is a basis of $K_1(\mathcal{L})$ and $K_1(\mathcal{L})$ is maximal isotropic. It follows that $y \in K_1(\mathcal{L}) \cap K_2(\mathcal{L}) = \{0\}$. Hence, Φ is injective. By cardinality, it is also surjective and $(\sigma(x_1), \cdots, \sigma(x_g))$ admits a unique dual basis $(y_1, \cdots, y_g)$ in $K_2(\mathcal{L})$ given by:

$$\forall i, j \in [\![1\ ;\ g]\!], \quad e_{\mathcal{L}}(\sigma(x_i), y_j) = \zeta_i^{\delta_{i,j}}.$$

Let $\tau : K_2(\delta) \longrightarrow K_2(\mathcal{L})$ such that $\tau(\chi_j) := y_j$ for all $j \in [\![1\ ;\ g]\!]$. Then τ is an isomorphism $K_2(\delta) \xrightarrow{\sim} K_2(\mathcal{L})$ such that

$$\sigma \times \tau : (x, \chi) \in K(\delta) \longmapsto \sigma(x) + \tau(\chi) \in K(\mathcal{L})$$

is symplectic and maps $K_i(\delta)$ to $K_i(\mathcal{L})$ for $i \in \{1, 2\}$ by construction. This proves the surjectivity.

Now, if ϕ and ψ are symplectic isomorphisms $K(\delta) \xrightarrow{\sim} K(\mathcal{L})$ mapping $K_i(\delta)$ to $K_i(\mathcal{L})$ for $i \in \{1, 2\}$ and such that $\phi_{|K_1(\delta)} = \psi_{|K_1(\delta)}$, then we have for all $(x, \chi) \in K_1(\delta) \times K_2(\delta)$,

$$\begin{aligned} e_{\mathcal{L}}(\phi(x, 1), \phi(0, \chi) - \psi(0, \chi)) &= e_{\mathcal{L}}(\phi(x, 1), \phi(0, \chi)) e_{\mathcal{L}}(\phi(x, 1), \psi(0, \chi))^{-1} \\ &= e_\delta((x, 1), (0, \chi)) e_{\mathcal{L}}(\psi(x, 1), \psi(0, \chi))^{-1} \\ &= \chi(x)\chi(x)^{-1} = 1. \end{aligned}$$

Since $e_{\mathcal{L}}$ is non-degenarate by Proposition 5.1.10, we conclude that $\phi(0, \chi) = \psi(0, \chi)$ for all $\chi \in K_2(\delta)$, so that $\phi = \psi$. This proves the injectivity and completes the proof. □

We now present, as promised, a description of $G(\mathcal{L})$ extending the description $K(\delta)$ of $K(\mathcal{L})$.

Definition 5.1.18. We define the *Heisenberg group* $\mathcal{H}(\delta)$ as $k^* \times K(\delta)$ with the following group law:

$$(\alpha, x_1, \chi_1) \cdot (\beta, x_2, \chi_2) := (\alpha\beta\chi_2(x_1), x_1 + x_2, \chi_1\chi_2).$$

Remark 5.1.19. e_δ is the equivalent of the commutator pairing $e_{\mathcal{L}}$. Indeed, it is also a commutator pairing in the Heisenberg group. Let $\rho_\delta : \mathcal{H}(\delta) \longrightarrow K(\delta), (\alpha, x, \chi) \longmapsto (x, \chi)$ be the forgetting map. Then, if $y, z \in K(\delta)$ and $\tilde{y}, \tilde{z} \in \mathcal{H}(\delta)$ are lifts of x, y respectively through ρ_δ, then we have $e_\delta(y, z) = \tilde{y}\tilde{z}\tilde{y}^{-1}\tilde{z}^{-1}$ (with the convention that scalars $\alpha \in k^*$ are identified with $(\alpha, 0, 1)$). Indeed, if we write $\tilde{y} := (\alpha, x_1, \chi_1)$ and $\tilde{z} := (\beta, x_2, \chi_2)$, then

$$\begin{aligned}\tilde{y}\tilde{z}\tilde{y}^{-1}\tilde{z}^{-1} &= (\alpha\beta\chi_2(x_1), x_1 + x_2, \chi_1\chi_2) \cdot (\alpha\beta\chi_1(x_2), x_1 + x_2, \chi_1\chi_2)^{-1} \\ &= \alpha\beta\chi_2(x_1)\alpha^{-1}\beta^{-1}\chi_1(x_2)^{-1}(1, x_1 + x_2, \chi_1\chi_2) \cdot (1, x_1 + x_2, \chi_1\chi_2)^{-1} \\ &= \chi_2(x_1)\chi_1(x_2)^{-1} = e_\delta((x_1, \chi_1), (x_2, \chi_2)) = e_\delta(y, z)\end{aligned}$$

Definition 5.1.20. A *theta structure* is an isomorphism of central extensions $\Theta_{\mathcal{L}} : \mathcal{H}(\delta) \xrightarrow{\sim} G(\mathcal{L})$, namely an isomorphism inducing a symplectic isomorphism $\overline{\Theta}_{\mathcal{L}} : K(\delta) \xrightarrow{\sim} K(\mathcal{L})$ such that the following diagram commutes:

$$\begin{array}{ccccccccc} 1 & \longrightarrow & k^* & \longrightarrow & \mathcal{H}(\delta) & \longrightarrow & K(\delta) & \longrightarrow & 0 \\ & & \| & & \downarrow{\scriptstyle \Theta_{\mathcal{L}}} & & \downarrow{\scriptstyle \overline{\Theta}_{\mathcal{L}}} & & \\ 1 & \longrightarrow & k^* & \longrightarrow & G(\mathcal{L}) & \longrightarrow & K(\mathcal{L}) & \longrightarrow & 0 \end{array}$$

Proposition 5.1.21. *There is a bijection between theta structures $\Theta_{\mathcal{L}} : \mathcal{H}(\delta) \longrightarrow G(\mathcal{L})$ and triples $(\overline{\Theta}_{\mathcal{L}}, s_1, s_2)$, where $\overline{\Theta}_{\mathcal{L}} : K(\delta) \longrightarrow K(\mathcal{L})$ is a symplectic isomorphism, $K_i(\overline{\Theta}_{\mathcal{L}}) := \overline{\Theta}_{\mathcal{L}}(K_i(\delta))$ and s_i is a section $K_i(\overline{\Theta}_{\mathcal{L}}) \longrightarrow G(\mathcal{L})$ for $i \in \{1, 2\}$. As a consequence, theta structures always exist.*

Proof. Let $\Theta_{\mathcal{L}}$ be a theta structure. Then it induces a symplectic isomorphism $\overline{\Theta}_{\mathcal{L}} : K(\delta) \xrightarrow{\sim} K(\mathcal{L})$ mapping the canonical symplectic decomposition $K(\delta) = K_1(\delta) \oplus K_2(\delta)$ to a symplectic decomposition $K(\mathcal{L}) = K_1(\overline{\Theta}_{\mathcal{L}}) \oplus K_2(\overline{\Theta}_{\mathcal{L}})$. Let $s_\delta : (x, \chi) \in K(\delta) \longmapsto (1, x, \chi) \in \mathcal{H}(\delta)$ and s_i be the composition

$$K_i(\overline{\Theta}_{\mathcal{L}}) \xrightarrow{\overline{\Theta}_{\mathcal{L}}^{-1}} K_i(\delta) \xrightarrow{s_\delta} \mathcal{H}(\delta) \xrightarrow{\Theta_{\mathcal{L}}} G(\mathcal{L})$$

for $i \in \{1, 2\}$. This defines sections $K_i(\overline{\Theta}_{\mathcal{L}}) \longrightarrow G(\mathcal{L})$, as desired.

Conversely, let $(\overline{\Theta}_{\mathcal{L}}, s_1, s_2)$ be a triple formed of a symplectic isomorphism and sections. Then the map $\mathcal{H}(\delta) \longrightarrow G(\mathcal{L})$ defined by the formula:

$$\Theta_{\mathcal{L}}(\alpha, x, \chi) := s_1 \circ \overline{\Theta}_{\mathcal{L}}(x, 0) + s_2 \circ \overline{\Theta}_{\mathcal{L}}(0, \chi)$$

is a theta structure. □

Remark 5.1.22. Similarly, the theta structure $\Theta_{\mathcal{L}}$ is determined by the couple $(\overline{\Theta}_{\mathcal{L}}, s_{\mathcal{L}})$, where $s_{\mathcal{L}} : K(\mathcal{L}) = K_1(\overline{\Theta}_{\mathcal{L}}) \oplus K_2(\overline{\Theta}_{\mathcal{L}}) \longrightarrow G(\mathcal{L})$ is the extension of the sections $s_i : K_i(\overline{\Theta}_{\mathcal{L}}) \longrightarrow G(\mathcal{L})$ for $i \in \{1, 2\}$.

Definition 5.1.23. The choice of the section $s_{\mathcal{L}} : K(\mathcal{L}) \longrightarrow G(\mathcal{L})$ lifting a symplectic decomposition $K(\mathcal{L}) = K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$ or equivalently of sections $s_i : K_i(\mathcal{L}) \longrightarrow G(\mathcal{L})$ or equivalently of the level subgroups $\widetilde{K_i}(\mathcal{L}) := s_i(K_i(\mathcal{L}))$ for $i \in \{1, 2\}$ is called a *level structure* of $G(\mathcal{L})$. It contains all the geometric information of a theta structure.

Lemma 5.1.24. *Two level structures $s_{\mathcal{L}}, s'_{\mathcal{L}} : K(\mathcal{L}) \longrightarrow G(\mathcal{L})$ lifting the same symplectic decomposition $K(\mathcal{L}) = K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$ differ by conjugation by a character $e_{\mathcal{L}}(c, \cdot) : K(\mathcal{L}) \longrightarrow k^*$, for some $c \in K(\mathcal{L})$.*

Proof. For $i \in \{1, 2\}$, $s'_{\mathcal{L}}$ differs from $s_{\mathcal{L}}$ on $K_i(\mathcal{L})$ by a character $\chi_i : x \in K_i(\mathcal{L}) \longrightarrow s'_{\mathcal{L}}(x) \cdot s_{\mathcal{L}}(x)^{-1} \in k^*$ (the result is indeed in k^* by Lemma 5.1.2). By non-degeneracy of $e_{\mathcal{L}}$, there exists $c_i \in K(\mathcal{L})$ such that $\chi_i = e_{\mathcal{L}}(c_i, \cdot)$. Since the $K_i(\mathcal{L})$ are isotropic and $K(\mathcal{L}) = K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$, we can assume that $c_1 \in K_2(\mathcal{L})$ and $c_2 \in K_1(\mathcal{L})$. Then $s'_{\mathcal{L}}$ differs from $s_{\mathcal{L}}$ by $e_{\mathcal{L}}(c, \cdot)$ with $c := c_1 + c_2$. □

5.1.5 Theta functions

By [Har77, Theorem II.7.1], if $\mathcal{L}$ is a line bundle on A generated by global sections, then these sections define a map $A \longrightarrow \mathbb{P}_k^n$. Such a map is a closed immersion when $\mathcal{L}$ is very ample. In this paragraph, we explain how to fix a basis of global sections given a theta structure $\Theta_{\mathcal{L}}$. Those will be the *theta functions*.

To find basis of global sections, we consider the action of $G(\mathcal{L})$ on the ring of global sections $\Gamma(A, \mathcal{L})$. This will lead to different choices of basis. Knowing the representations of the Heisenberg group $\mathcal{H}(\delta)$ better and fixing a theta structure will help to make this choice canonical.

Theorem 5.1.25. *[Mum66, Proposition 3]*

(i) $\mathcal{H}(\delta)$ has a unique irreducible representation $V(\delta)$ on which k^* acts naturally. $V(\delta)$ can be described explicitly as the space of functions $K_1(\delta) \longrightarrow k$ with the action:

$$\forall g \in V(\delta), (\alpha, x, \chi) \in \mathcal{H}(\delta), \quad (\alpha, x, \chi) \cdot g : y \longmapsto \alpha\chi(y)^{-1} g(y - x).$$

(ii) Any representation V of $\mathcal{H}(\delta)$ on which k^* acts naturally is isomorphic to $V(\delta)^r$ with $r := \dim_k(V^{\widetilde{K}})$, $V^{\widetilde{K}}$ being the subspace of $\widetilde{K}$-invariant elements of V for any maximal level subgroup $\widetilde{K} \subset \mathcal{H}(\delta)$.

Proof. **(i)** Let (V, ρ) be an irreducible representation on which $\mathcal{H}(\delta)$ acts naturally. Let $\widetilde{K} \subset \mathcal{H}(\delta)$ be any maximal level subgroup. Since $\rho(x) \in GL(V)$ is annihilated by $T^{\#\widetilde{K}} - 1 \in k[T]$ for all $x \in \widetilde{K}$, and $\mathrm{char}(k) \nmid \#\widetilde{K}$, we can diagonalize $\rho(x)$ for all $x \in \widetilde{K}$. Since $\widetilde{K}$ is abelian, we can co-diagonalize all of these endomorphisms with the same basis, so that V can be decomposed into

$$V = \bigoplus_{\chi \in \mathrm{Hom}(\widetilde{K}, k^*)} V_\chi,$$

where $V_\chi := \{v \in V \mid \forall x \in \widetilde{K}, \quad \rho(x)(v) = \chi(x) \cdot v\}$ is the eigenspace of weight χ for all character $\chi \in \mathrm{Hom}(\widetilde{K}, k^*)$.

Let $y \in \mathcal{H}(\delta)$. Then, by easy computations in the Heisenberg group, we get that

$$\forall z \in \widetilde{K}, \quad y^{-1} z y = e_\delta(z, y) \cdot z.$$

We denote by χ_y the character $z \in \widetilde{K} \longmapsto e_\delta(z, y) \in k^*$. Now, let χ_0 be the trivial character and $v \in V_{\chi_0}$. Then, we have for all $z \in \widetilde{K}$:

$$\begin{aligned}\rho(z)(\rho(y)(v)) &= \rho(y) \circ \rho(y^{-1} z y)(v) = \rho(y) \circ \rho(\chi_y(z) \cdot z)(v) = \chi_y(z) \cdot \rho(y) \circ \rho(z)(v) \\ &= \chi_y(z) \cdot \rho(y)(\chi_0(z) \cdot v) = \chi_y(z) \cdot \rho(y)(v)\end{aligned}$$

Hence, $\rho(y)(v) \in V_{\chi_y}$.

We consider the map

$$\begin{array}{rccc} \gamma : \mathcal{H}(\delta)/(k^* \cdot \widetilde{K}) & \longrightarrow & \mathrm{Hom}(\widetilde{K}, k^*) \\ y & \longmapsto & \chi_y \end{array}$$

This map is well defined since $\chi_y(z) = e_\delta(z, y) = z^{-1} y^{-1} z y = 1$ for all $y \in k^* \cdot \widetilde{K}$ and $z \in \widetilde{K}$. It is injective since $\chi_y(z) = 1$ for all $z \in \widetilde{K}$ implies that z commutes with y, so that $y \in k^* \cdot \widetilde{K}$ by maximality of $\widetilde{K}$. Finally, it is surjective because e_δ is non-degenerate. Hence γ is an isomorphism.

Now, since γ is surjective, all the characters of $\mathrm{Hom}(\widetilde{K}, k^*)$ are of the form χ_y for some $y \in \mathcal{H}(\delta)$ and since $\rho(y)$ maps V_{χ_0} to V_{χ_y}, we either have all the eigenspaces zero or non-zero. We suppose the latter. Furthermore, if we fix $v \in V_{\chi_0}$ and denote by W the subspace spanned by all the $\rho(y)(v)$ for $y \in \mathcal{H}(\delta)$ then $W \cap V_\chi$ must have dimension one for all $\chi \in \mathrm{Hom}(\widetilde{K}, k^*)$ by injectivity of γ. W being stable by ρ and V being irreducible, we must have $V = W$ so the V_χ all have dimension 1. This proves the uniqueness.

To conclude point (i), consider the space $V(\delta)$ of functions $K_1(\delta) \longrightarrow k$ with the action of $\mathcal{H}(\delta)$ given by:

$$\forall g \in V(\delta), (\alpha, x, \chi) \in \mathcal{H}(\delta), \quad (\alpha, x, f) \cdot g : y \longmapsto \alpha\chi(y)^{-1} g(y - x).$$

Taking $\widetilde{K} := s_\delta(K_1(\delta)) = \{(1, x, 1) \mid x \in K_1(\delta)\}$, we still have a decomposition

$$V(\delta) = \bigoplus_{\chi \in \mathrm{Hom}(\widetilde{K}, k^*)} V(\delta)_\chi.$$

With $V(\delta)_\chi = \{g \in V(\delta) \mid \forall x, y \in K_1(\delta), g(y - x) = \chi(x)g(y)\} = k \cdot \chi^{-1}$, so that $\dim_k V(\delta)_\chi = 1$ for all $\chi \in \mathrm{Hom}(\widetilde{K}, k^*)$. Consequently, $V(\delta)$ is an irreducible representation, and is the only one up to isomorphism. This proves (i).

(ii) Any representation V on which k^* acts naturally can be decomposed into a direct sum of irreducible representations, so V is isomorphic to $V(\delta)^r$ by (i). With this decomposition, we see that $V^{\widetilde{K}} = V_{\chi_0} \simeq V(\delta)^r_{\chi_0}$, χ_0 being the trivial character of a maximal level subgroup $\widetilde{K}$. Since $\dim_k V(\delta)_{\chi_0} = 1$, we conclude that $r = \dim_k(V^{\widetilde{K}})$. □

Corollary 5.1.26. *Let V be an irreducible representation of $\mathcal{H}(\delta)$ on which k^* acts naturally. Then for any maximal level subgroup $\widetilde{K} \subset \mathcal{H}(\delta)$, the subspace $V^{\widetilde{K}}$ of elements fixed by $\widetilde{K}$ has dimension* 1.

Proof. By point (i) of Theorem 5.1.25, we have $V \simeq V(\delta)$ so point (ii) of the same theorem ensures that $\dim_k(V^{\widetilde{K}}) = 1$ for any maximal level subgroup $\widetilde{K} \subset \mathcal{H}(\delta)$. □

Definition 5.1.27. We define the representation $V := \Gamma(A, \mathcal{L})$ of $G(\mathcal{L})$ given by:

$$\forall (x, \phi) \in G(\mathcal{L}), s \in \Gamma(A, \mathcal{L}), \quad U_{\mathcal{L}}(x, \phi)(s) := t^*_{-x}(\phi(s)).$$

Remark 5.1.28. This is well defined because ϕ maps $\mathcal{L}$ to $t^*_x\mathcal{L}$ so that $\phi(s) \in \Gamma(A, t^*_x\mathcal{L})$ and $t^*_{-x}(\phi(s)) \in \Gamma(A, t^*_{-x}t^*_x\mathcal{L}) \simeq \Gamma(A, \mathcal{L})$.

Theorem 5.1.29. *If $\mathcal{L}$ is ample and separable, then the representation $(\Gamma(A, \mathcal{L}), U_{\mathcal{L}})$ of $G(\mathcal{L})$ is irreducible.*

Proof. Let $\widetilde{K}$ be a maximal level subgroup of $G(\mathcal{L})$ and $K \subset K(\mathcal{L})$ its image via $\rho_{\mathcal{L}}$. Let $B := A/K$ and $\pi : A \longrightarrow B$ be the projection map (of kernel K). By Grothendieck's descent theorem (Theorem 5.1.4), $\mathcal{L}$ descends to a line bundle $\mathcal{M}$ such that $\mathcal{L} = \pi^*\mathcal{M}$. By maximality of $\widetilde{K}$, we get that the polarization $\varphi_{\mathcal{M}}$ associated to $\mathcal{M}$ has degree 1 (see the proof of Proposition 5.1.10). Hence, $\chi(\mathcal{M}) = \pm 1$ by Theorem 1.4.62. Since $\mathcal{L} = \pi^*\mathcal{M}$ is ample, $\mathcal{M}$ is also ample by [The24, Tag 0B5V] and it follows that $\dim_k \Gamma(A, \mathcal{M}) = \chi(\mathcal{M}) = 1$ by Corollary 1.4.64. Besides, π^* maps the sections of $\mathcal{M}$ to the sections of $\mathcal{L}$ invariant under the action of $\widetilde{K}$. It follows that $\dim_k \Gamma(A, \mathcal{L})^{\widetilde{K}} = \dim_k \Gamma(A, \mathcal{M}) = 1$, so that $(\Gamma(A, \mathcal{L}), U_{\mathcal{L}})$ is irreducible by Theorem 5.1.25. □

By Theorem 5.1.29 the action of the theta group $G(\mathcal{L})$ on global sections $\Gamma(A, \mathcal{L})$ is irreducible and by Theorem 5.1.25, this representation must be isomorphic to the canonical representation $V(\delta)$ of the Heisenberg group $\mathcal{H}(\delta)$. Hence, there exists an isomorphism $\beta : V(\delta) \xrightarrow{\sim} \Gamma(A, \mathcal{L})$ respecting the group actions of $\mathcal{H}(\delta)$ and $G(\mathcal{L})$, namely such that:

$$\forall v \in V(\delta), h \in \mathcal{H}(\delta), \quad \beta(h \cdot v) = \Theta_{\mathcal{L}}(h) \cdot \beta(v). \tag{5.6}$$

Since both $V(\delta)$ and $\Gamma(A, \mathcal{L})$ are irreducible, β is uniquely determined up to multiplication by a scalar, as a consequence of Schur's lemma [Lan04, Lemma XVIII.5.9]. $V(\delta)$, which is the space of functions $K_1(\delta) \longrightarrow k$ has a canonical basis $(\delta_i)_{i \in K_1(\delta)}$ given by the Kronecker delta functions:

$$\forall i, j \in K_1(\delta), \quad \delta_i(j) := \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$$

This gives a basis $(\theta_i)_{i \in K_1(\delta)}$ of $\Gamma(A, \mathcal{L})$ given by $\theta_i := \beta(\delta_i)$ for all $i \in K_1(\delta)$.

Definition 5.1.30. The basis $(\theta_i)_{i \in K_1(\delta)}$ is called the basis of *theta functions* associated to the theta structure $\Theta_{\mathcal{L}}$.

5.1.6 When theta functions become coordinates

When computing isogenies and performing arithmetic operations on abelian varieties, we shall naturally represent points in the projective space by evaluating a basis of theta functions at these points. The goal of this section is to define this evaluation properly and to determine when it fully represents points in the projective space *i.e.* when theta functions define a projective embedding and can be considered as coordinates. Let $\mathcal{L}$ be a separable ample line bundle on A.

Definition 5.1.31. Let $x \in A$. We define

$$\mathcal{L}(x) := \mathcal{L}_x \otimes_{\mathcal{O}_{A,x}} \kappa(x),$$

where $\mathcal{L}_x$ and $\mathcal{O}_{A,x}$ are respectively the stalks of $\mathcal{L}$ and $\mathcal{O}_A$ at x and $\kappa(x) := \mathcal{O}_{A,x}/\mathfrak{m}_{A,x}$ is the residue field of $\mathcal{O}_{A,x}$.

When x is a closed point ($x \in A(k)$), we have $\kappa(x) = k$ since k is algebraically closed, so $\mathcal{L}(x) \simeq k$. Consider an isomorphism $\lambda_x : \mathcal{L}(x) \xrightarrow{\sim} k$ and define for all $s \in \Gamma(A, \mathcal{L})$, $s(x) := \lambda_x(s_{(x)})$, where $s_{(x)}$ is the image of s in $\mathcal{L}(x)$. Since λ_x is determined up to a scalar in k^*, so is $s(x)$ but $[\theta_i(x)]_{i\in K_1(\delta)}$ defines a projective point (provided one of the coordinates is non-zero) independent of the choice of λ_x. When $\mathcal{L}$ is generated by global sections (so that $(\theta_i(x))_{i\in K_1(\delta)}$ is never zero), this defines a morphism

$$\begin{array}{rccc} J_{\mathcal{L}} : & A & \longrightarrow & \mathbb{P}_k^{d-1} \\ & x & \longmapsto & [\theta_i(x)]_{i\in K_1(\delta)} \end{array} \tag{5.7}$$

with $d := \#K_1(\delta) = \prod_{i=1} d_i$, which corresponds to the morphism induced by the basis of theta functions $(\theta_i)_{i\in K_1(\delta)}$ (Theorem 1.4.9). When $\mathcal{L}$ is very ample this map is a closed immersion.

Lemma 5.1.32. *Let $\mathcal{L}$ be an ample and separable line bundle and $n \in \mathbb{N}^*$ not divisible by* $\mathrm{char}(k)$*. If $A[n] \subseteq K(\mathcal{L})$, there exists an ample and separable line bundle $\mathcal{M}$ such that $\mathcal{L} = \mathcal{M}^n$.*

Proof. We assume that $A[n] \subseteq K(\mathcal{L})$. Then, for all $x, y \in A[n]$, we have by Proposition 5.1.6.(v)

$$e_{\mathcal{L}^n}(x, y) = e_{\mathcal{L}}(x, ny) = e_{\mathcal{L}}(x, 0) = 1.$$

Hence, $A[n]$ is isotropic in $K(\mathcal{L}^n)$, so by Proposition 5.1.12, we can lift $A[n]$ to a level subgroup of $G(\mathcal{L}^n)$ and by Grothendieck's descent theorem (Theorem 5.1.4), there exists an ample line bundle $\mathcal{M}'$ such that $\mathcal{L}^n = [n]^*\mathcal{M}'$. But, we know that $[n]^*\mathcal{M}' \simeq \mathcal{M}'^{n^2} \otimes (\mathcal{M}'^{-1} \otimes [-1]^*\mathcal{M}')^{(n^2-n)/2}$ by Corollary 1.4.18 and we know that $\mathcal{M}'^{-1} \otimes [-1]^*\mathcal{M}' \in \mathrm{Pic}^0(A)$ by Lemma 1.4.65, so that $(\mathcal{L} \otimes \mathcal{M}'^{-n})^n \in \mathrm{Pic}^0(A)$. It follows that $\mathcal{L} \otimes \mathcal{M}'^{-n} \in \mathrm{Pic}^0(A)$ by Lemma 1.4.65. Since $\mathrm{Pic}^0(A) = \widehat{A}(k)$ is n-divisible, we can write $\mathcal{L} \otimes \mathcal{M}'^{-n} = \mathcal{N}^n$ with $\mathcal{N} \in \mathrm{Pic}^0(A)$, so that $\mathcal{L} = \mathcal{M}^n$ with $\mathcal{M} := \mathcal{M}' \otimes \mathcal{N}$. $\mathcal{M}$ is ample and separable since $\mathcal{L}$ is and $\mathrm{char}(k) \nmid n$. This completes the proof. □

Theorem 5.1.33. *Let $\mathcal{L}$ be an ample separable line bundle of A of type δ. If there exists $n \geq 3$ such that $n|\delta$ then $\mathcal{L}$ is very ample and the map $J_{\mathcal{L}}$ from Eq. (5.7) is a closed immersion.*

Proof. Under the assumptions of the theorem, we have $A[n] \subseteq K(\mathcal{L})$ with $n \geq 3$ and there exists an ample and separable line bundle $\mathcal{M}$ on A such that $\mathcal{L} = \mathcal{M}^n$ by Lemma 5.1.32. Since $n \geq 3$ and $\mathcal{M}$ is ample, we then get that $\mathcal{L}$ is very ample by Theorem 1.4.22. □

In particular, by the above theorem, theta functions $(\theta_i^{\mathcal{L}})_i$ are theta coordinates *i.e.* define a projective embedding when $\mathcal{L}$ is of level 3 ($\delta = (3, \cdots, 3)$) or bigger. However, this theorem does not apply when $\mathcal{L}$ is of level 2 ($\delta = (2, \cdots, 2)$), an assumption that we shall make when computing isogenies to minimize the number of theta functions to handle. In that case, we still have a partial result that will be sufficient for our needs. We obtain that theta functions define embedding of the Kummer variety $K_A := A/\pm \hookrightarrow \mathbb{P}_k^{d-1}$. By abuse of language, we shall still call them theta *coordinates* in this case.

Theorem 5.1.34. *Let $\mathcal{L}$ be a symmetric line bundle on A ($[-1]^*\mathcal{L} \simeq \mathcal{L}$). Assume that* $\mathrm{char}(k) \neq 2$*. Then:*

(i) $\mathcal{L}^2$ is generated by global sections, so that $J_{\mathcal{L}^2} : A \longrightarrow \mathbb{P}_k^{d-1}$ introduced in Eq. (5.7) is well-defined.

(ii) $J_{\mathcal{L}^2}$ *factors through the projection* $\pi : A \longrightarrow K_A = A/\pm$.

(iii) *If* $(A, \mathcal{L})$ *is* irreducible i.e. *cannot be decomposed into a product of polarised abelian varieties, then* $J_{\mathcal{L}^2}$ *induces an embedding of the Kummer variety* $K_A \hookrightarrow \mathbb{P}_k^{d-1}$.

In other words, theta functions $(\theta_i^{\mathcal{L}^2})_i$ *induce projective coordinates on* K_A.

Proof. (i) has been proved in [PP02, Corollary 3.9.(ii)]. (ii) follows from Proposition 5.3.7 and point (ii) of Proposition 5.3.8. (iii) has been proved in [BL04, Theorem 4.8.1] over the complex numbers. The result is more general but a proof over any field could not be found in the literature. □

5.1.7 The theta null point

For algorithmic applications, it will be necessary to represent theta structures on a computer. In this section, we introduce the *theta null point* as the evaluation of theta functions at zero. We shall see that when $\mathcal{L}$ is very ample, the theta null point fully determines the theta-structure from which it is obtained and may be used to represent it.

Let $\mathcal{L}$ be an ample separable line bundle on A generated by global sections. In Section 5.1.6, we have seem how to evaluate theta functions at closed points. Since $0 \in A(k)$ is a closed point, we can evaluate the map defined in Eq. (5.7) $J_{\mathcal{L}}(0) := (\theta_i(0))_{i \in K_1(\delta)}$.

Definition 5.1.35. We call $J_{\mathcal{L}}(0) \in \mathbb{P}_k^{d-1}$ the *theta null point* attached to the theta structure $\Theta_{\mathcal{L}}$.

To prove our main theorem (Theorem 5.1.39), we shall see how a theta null point can determine the evaluation of theta functions on the whole of $K(\mathcal{L})$. We first explain how to evaluate sections of $\Gamma(A, \mathcal{L})$ at all points $x \in K(\mathcal{L})$ in a coherent way (if we work in affine space instead of projective space). Recall the notations of Section 5.1.6 and Definition 5.1.31 in particular. We define isomorphisms $\lambda_x : \mathcal{L}(x) \xrightarrow{\sim} k$ that factors through a fixed choice of isomorphism $\lambda_0 : \mathcal{L}(0) \xrightarrow{\sim} k$ for all $x \in K(\mathcal{L})$. As we have seen in Proposition 5.1.21, the theta structure $\Theta_{\mathcal{L}}$ defines a symplectic isomorphism $\overline{\Theta}_{\mathcal{L}} : K(\delta) \xrightarrow{\sim} K(\mathcal{L})$ and a section $s_{\mathcal{L}} : K(\mathcal{L}) \longrightarrow G(\mathcal{L})$ by the formula $s_{\mathcal{L}} := \Theta_{\mathcal{L}} \circ s_\delta \circ \overline{\Theta}_{\mathcal{L}}^{-1}$, with $s_\delta : (x, \chi) \in K(\delta) \longmapsto (1, x, \chi) \in \mathcal{H}(\delta)$. For all $x \in K(\mathcal{L})$, let us write $s_{\mathcal{L}}(x) := (x, \phi_x)$, with $\phi_x : \mathcal{L} \xrightarrow{\sim} t_x^*\mathcal{L}$. For all $x \in K(\mathcal{L})$, ϕ_x induces an isomorphism $\phi_x(0) : \mathcal{L}(0) \xrightarrow{\sim} t_x^*\mathcal{L}(0) = \mathcal{L}(x)$, so we can define $\lambda_x := \lambda_0 \circ \phi_x(0)^{-1} : \mathcal{L}(x) \xrightarrow{\sim} k$. As in Section 5.1.6, we can then define:

$$\forall s \in \Gamma(A, \mathcal{L}), \forall x \in K(\mathcal{L}), \quad s(x) := \lambda_x(s_{(x)}) = \lambda_0(\phi_x^{-1}(t_x^* s)_{(0)}),$$

where $s_{(x)}$ is the image of s in $\mathcal{L}(x)$. Up to a choice of symplectic isomorphism $\overline{\Theta}_{\mathcal{L}} : K(\delta) \longrightarrow K(\mathcal{L})$, this method to evaluate points of $K(\mathcal{L})$ yields a method to evaluate points of $K(\delta)$.

Definition 5.1.36. For all $(i, \chi) \in K(\delta)$ and $s \in \Gamma(A, \mathcal{L})$, we define the *value of* s *at* (i, χ) by:

$$s(i, \chi) := s(x) = \lambda_x(s_{(x)}) = \lambda_0(\phi_x^{-1}(t_x^* s)_{(0)}),$$

with $x := \overline{\Theta}_{\mathcal{L}}(i, \chi) \in K(\mathcal{L})$ and $s_{(x)}$ the image of s in $\mathcal{L}(x)$.

Lemma 5.1.37. *Let* $s \in \Gamma(A, \mathcal{L})$ *and* $g \in G(\mathcal{L})$. *Let* $(\alpha, i_1, \chi_1) := \Theta_{\mathcal{L}}^{-1}(g) \in \mathcal{H}(\delta)$. *Then, for all* $(i_2, \chi_2) \in K(\delta)$,

$$(g \cdot s)(i_2, \chi_2) = \alpha \chi_2(i_1)\chi_1(i_1)^{-1} s(i_2 - i_1, \chi_2 \chi_1^{-1}),$$

where $g \cdot s$ *is the result of the the theta group action of* g *on* s *defined in Definition 5.1.27.*

Proof. Let $(i_2, \chi_2) \in K(\delta)$ and $x := \overline{\Theta}_{\mathcal{L}}(i_2, \chi_2)$. Then

$$\begin{aligned} s(i_2, \chi_2) &= \lambda_0(\phi_x^{-1}(t_x^* s)_{(0)}) = \lambda_0((t_{-x}^* \phi_x^{-1})(s)_{(0)}) = \lambda_0(((-x, t_{-x}^* \phi_x^{-1}) \cdot s)_{(0)}) \\ &= \lambda_0((g'^{-1} \cdot s)_{(0)}), \end{aligned} \tag{5.8}$$

with $g' := s_{\mathcal{L}}(x) = (x, \phi_x) \in G(\mathcal{L})$. It follows that

$$(g \cdot s)(i_2, \chi_2) = \lambda_0((g'^{-1} g \cdot s))_{(0)}) = \lambda_0(((g^{-1} g')^{-1} \cdot s)_{(0)}). \tag{5.9}$$

Besides,

$$\begin{aligned}\Theta_{\mathcal{L}}^{-1}(g^{-1}g') &= \Theta_{\mathcal{L}}^{-1}(g)^{-1}\cdot\Theta_{\mathcal{L}}^{-1}(g')\\ &= (\alpha,i_1,\chi_1)^{-1}\cdot\Theta_{\mathcal{L}}^{-1}\circ s_{\mathcal{L}}\circ\overline{\Theta}_{\mathcal{L}}(i_2,\chi_2)\\ &= (\alpha,i_1,\chi_1)^{-1}\cdot s_\delta(i_2,\chi_2) = (\chi_1(i_1)\alpha^{-1},-i_1,\chi_1^{-1})\cdot(1,i_2,\chi_2)\\ &= (\chi_1(i_1)\chi_2(i_1)^{-1}\alpha^{-1},i_2-i_1,\chi_2\chi_1^{-1}),\end{aligned}$$

so that $g^{-1}g' = \gamma\cdot g''$ with $g'' := s_{\mathcal{L}}\circ\overline{\Theta}_{\mathcal{L}}(i_2-i_1,\chi_2\chi_1^{-1})$ and $\gamma := \chi_1(i_1)\chi_2(i_1)^{-1}\alpha^{-1}\in k$. Finally, we obtain by Eqs. (5.8) and (5.9):

$$\begin{aligned}(g\cdot s)(i_2,\chi_2) &= \lambda_0(((\gamma^{-1}\cdot g''^{-1})\cdot s)_{(0)}) = \gamma^{-1}\lambda_0((g''^{-1}\cdot s)_{(0)})\\ &= \gamma^{-1}s(i_2-i_1,\chi_2\chi_1^{-1}) = \alpha\chi_2(i_1)\chi_1(i_1)^{-1}s(i_2-i_1,\chi_2\chi_1^{-1}).\end{aligned}$$

□

Lemma 5.1.38. *For all $i\in K_1(\delta)$ and $(j,\chi)\in K(\delta)$,*

$$\theta_i(j,\chi) = \chi(i)\theta_{i-j}(0).$$

Proof. Let $i\in K_1(\delta)$, $(j,\chi)\in K(\delta)$ and $g := \Theta_{\mathcal{L}}(1,j,\chi)$. Recall the notations of the previous section: $\mathcal{H}(\delta)$ acts on $V(\delta)$, $\delta_i\in V(\delta)$ is the Kronecker function and β is an isomorphism $V(\delta)\xrightarrow{\sim}\Gamma(A,\mathcal{L})$ respecting the actions of $\mathcal{H}(\delta)$ on $V(\delta)$ and $G(\mathcal{L})$ on $\Gamma(A,\mathcal{L})$. Then

$$\begin{aligned}g\cdot\theta_i &= g\cdot\beta(\delta_i) = \beta((1,j,\chi)\cdot\delta_i) = \beta(\chi(i+j)^{-1}\delta_{i+j}) = \chi(i+j)^{-1}\beta(\delta_{i+j})\\ &= \chi(i+j)^{-1}\theta_{i+j}\end{aligned}$$

Combining this equality with Lemma 5.1.37, we get that

$$\chi(i+j)^{-1}\theta_{i+j}(0) = (g\cdot\theta_i)(0) = \chi(j)^{-1}\theta_i(-j,\chi^{-1}).$$

Replacing g by $g' := \Theta_{\mathcal{L}}(1,-j,\chi^{-1})$, we finally get the desired equality:

$$\theta_i(j,\chi) = \chi(i)\theta_{i-j}(0).$$

□

Theorem 5.1.39. *Let $\mathcal{L}$ be a separable line bundle generated by global sections.*

(i) If the symplectic isomorphism $\overline{\Theta}_{\mathcal{L}}: K(\delta)\xrightarrow{\sim}K(\mathcal{L})$ is fixed then the theta null point $(\theta_i(0))_{i\in K(\delta)}$ completely determines the section $s_{\mathcal{L}}: K(\mathcal{L})\longrightarrow G(\mathcal{L})$.

(ii) If $\mathcal{L}$ is very ample, the theta null point completely determines the full theta structure $\Theta_{\mathcal{L}}$.

Proof. **(i)** Assume $(\theta_i(0))_{i\in K(\delta)}\neq 0$ and let $y\in K(\mathcal{L})$ and $(j,\chi) := \overline{\Theta}_{\mathcal{L}}^{-1}(y)\in K(\delta)$. Let $i\in K_1(\delta)$ such that $\theta_{i-j}(0)\neq 0$. Then

$$\lambda_0(\phi_y^{-1}(t_y^*\theta_i)_{(0)}) = \theta_i(j,\chi) = \chi(i)\theta_{i-j}(0)\neq 0.$$

ϕ_y being determined up to a scalar, ϕ_y is entirely determined by the value $\theta_{i-j}(0)$ so $s_{\mathcal{L}}(y) = (y,\phi_y)$ is entirely determined by $\theta_{i-j}(0)$.

(ii) If $\mathcal{L}$ is very ample, then the map $J_{\mathcal{L}}: A\longrightarrow\mathbb{P}_k^{d-1}$ from Eq. (5.7) is an embedding so for all $(j,\chi)\in K(\delta)$ the projective point $(\theta_i(\overline{\Theta}_{\mathcal{L}}(j,\chi)))_{i\in K_1(\delta)}$ fully determines $\overline{\Theta}_{\mathcal{L}}(j,\chi)$. But by construction, $\theta_i(\overline{\Theta}_{\mathcal{L}}(j,\chi)) = \theta_i(j,\chi)$, which is determined by the theta null point by Lemma 5.1.38. It follows that the theta null point also determines $\overline{\Theta}_{\mathcal{L}}: K(\delta)\xrightarrow{\sim}K(\mathcal{L})$, so it determines the whole theta structure by (i) and Proposition 5.1.21 (see also Remark 5.1.22). □

Remark 5.1.40. When $4|\delta$ ($4|d_i$ for all $i\in[\![1\ ;\ g]\!]$), the theta null point also determines projective equations defining the image of the embedding $A\hookrightarrow\mathbb{P}_k^{d-1}$ (called *Riemann relations*), so it completely determines the triple $(A,\mathcal{L},\Theta_{\mathcal{L}})$ made of the polarised abelian variety with the theta structure. We refer to [Mum66, Corollary p. 340] for a proof of this result.

5.1.8 Action by translation of the theta group on theta functions

Let $\mathcal{L}$ be a separable ample line bundle on A generated by global sections. We have seen in Section 5.1.7 that the theta-group $G(\mathcal{L})$ acts on theta null point by translation (see Lemma 5.1.38 in particular). In this section, we prove that this fact is general. Let us define an action on $J_{\mathcal{L}}(A)$ as follows:

$$\forall g \in G(\mathcal{L}), \forall x \in A(k), \quad g \cdot J_{\mathcal{L}}(x) := [(g \cdot \theta_i^{\mathcal{L}})(x)]_{i \in K(\delta)}$$

Lemma 5.1.41. *This is a well-defined group action, in the sense that $J_{\mathcal{L}}(A)$ is stable by $G(\mathcal{L})$-action. More precisely, if $x \in A(k)$ and $(y, \phi_y) \in G(\mathcal{L})$, then $(y, \phi_y) \cdot J_{\mathcal{L}}(x) = J_{\mathcal{L}}(x + y)$.*

Proof. Let $x \in A(k)$, $(y, \phi_y) \in G(\mathcal{L})$ and $s \in \Gamma(A, \mathcal{L})$. Choose an isomorphism $\lambda_x : \mathcal{L}(x) \xrightarrow{\sim} k$ defining the evaluation at x (as in Section 5.1.7), so that

$$((y, \phi_y) \cdot s)(x) = \lambda_x(((y, \phi_y) \cdot s)_{(x)}) = \lambda_x(t_{-y}^* \phi_y(s)_{(x)}).$$

Now, define $\lambda_{x+y} : \mathcal{L}(x + y) \xrightarrow{\sim} k$ that factor through λ_x by $\lambda_{x+y} := \lambda_x \circ \phi_y(x)^{-1}$, where $\phi_y(x) : \mathcal{L}(x) \xrightarrow{\sim} t_y^* \mathcal{L}(x) = \mathcal{L}(x + y)$ is the image at x of $\phi_y : \mathcal{L} \xrightarrow{\sim} t_y^* \mathcal{L}$. We then define the evaluation of s at $x + y$ by

$$s(x + y) := \lambda_{x+y}(s_{(x+y)}) = \lambda_x(\phi_y(x)^{-1}(s_{(x+y)})) = \lambda_x(\phi_y(x)^{-1}(t_y^* s_{(x)})) = \lambda_x(\phi_y^{-1}(t_y^* s)_{(x)}).$$

But $s \longmapsto t_{-y}^* \phi_y(s)$ and $s \longmapsto \phi_y^{-1}(t_y^* s)$ are automorphisms of $\mathcal{L}$ and $\mathrm{Aut}(\mathcal{L}) \simeq k^*$ (as we have seen in the proof of Lemma 5.1.2) so they differ by a constant $\mu \in k^*$. It follows that $((y, \phi_y) \cdot s)(x) = \mu s(x+y)$ for all $s \in \Gamma(A, \mathcal{L})$, so that

$$(y, \phi_y) \cdot J_{\mathcal{L}}(x) = [((y, \phi_y) \cdot \theta_i^{\mathcal{L}})(x)]_{i \in K(\delta)} = [\mu \theta_i^{\mathcal{L}}(x + y)]_{i \in K(\delta)} = [\theta_i^{\mathcal{L}}(x + y)]_{i \in K(\delta)} = J_{\mathcal{L}}(x + y).$$

This completes the proof. □

Remark 5.1.42. If we fix a theta-structure $\Theta_{\mathcal{L}}$, this action on the basis of theta-functions is given by

$$\forall (\alpha, i, \chi) \in \mathcal{H}(\delta), \forall j \in K_1(\delta), \quad \Theta_{\mathcal{L}}(\alpha, i, \chi) \cdot \theta_j = \alpha \chi(i + j)^{-1} \theta_{i+j}. \tag{5.10}$$

If $x \in A(k)$ and $y = \overline{\Theta}_{\mathcal{L}}(j, \chi) \in K(\mathcal{L})$, we then have:

$$[\theta_i(x + y)]_i = [\Theta_{\mathcal{L}}(1, j, \chi) \cdot \theta_i(x)]_i = [\chi(i + j)^{-1} \theta_{i+j}(x)]_i = [\chi(i)^{-1} \theta_{i+j}(x)]_i. \tag{5.11}$$

The above result is consistent with Lemma 5.1.38 when $x = 0$.

5.2 Isogenies and theta structures

In this section, we present how the theory of theta functions applies to isogenies. Our goal is to prove a formula that relates theta functions on the domain and codomain of an isogeny called the *isogeny theorem* (Theorem 5.2.5). As expected, this formula will be crucial to our computational algorithms introduced in Chapter 6.

5.2.1 Compatible theta structures

Before introducing the isogeny theorem in Section 5.2.2, we give some compatibility conditions on theta structures defined on the domain and codomain of a polarised isogeny. This section is quite technical and the reader does not need to focus too much attention on details. The main goal is to characterise theta structures on the codomain that may be compatibles with a fixed theta structure on the domain. This is achieved in Proposition 5.2.4.

Let $f : (A, \mathcal{L}) \longrightarrow (B, \mathcal{M})$ be an isogeny between polarized abelian varieties (such that $f^* \mathcal{M} \simeq \mathcal{L}$). The goal of this paragraph is to relate the theta structures on $(A, \mathcal{L})$ and $(B, \mathcal{M})$ via f. For this, we first define a notion of compatibility between theta structures.

Let us recall some notations of Section 5.1.2. Let $K := \ker(f)$ and $\widetilde{K}$ be the level subgroup above K given by an isomorphism $\alpha : f^*\mathcal{M} \xrightarrow{\sim} \mathcal{L}$:

$$\widetilde{K} = \{(x, t_x^*\alpha \circ \alpha^{-1}) \mid x \in K\}.$$

Let $\alpha_f : Z(\widetilde{K}) \twoheadrightarrow G(\mathcal{M})$ be the surjective map defined by Equation 5.1. It induces an isomorphism $Z(\widetilde{K})/\widetilde{K} \xrightarrow{\sim} G(\mathcal{M})$ by Proposition 5.1.5.(iii).

Let $\Theta_\mathcal{L}$ and $\Theta_\mathcal{M}$ be theta structures on $(A, \mathcal{L})$ and $(B, \mathcal{M})$ respectively and $s_\mathcal{L} : K(\mathcal{L}) \longrightarrow G(\mathcal{L})$, $s_\mathcal{M} : K(\mathcal{M}) \longrightarrow G(\mathcal{M})$ be the sections induced by $\Theta_\mathcal{L}$ and $\Theta_\mathcal{M}$ respectively (see Remark 5.1.22). These sections define level structures on their respective theta groups given by $\widetilde{K_i}(s_\mathcal{L}) := s_\mathcal{L}(K_i(\overline{\Theta}_\mathcal{L}))$ and $\widetilde{K_i}(s_\mathcal{M}) := s_\mathcal{M}(K_i(\overline{\Theta}_\mathcal{M}))$ for $i \in \{1, 2\}$, given the symplectic decompositions $K(\mathcal{L}) = K_1(\overline{\Theta}_\mathcal{L}) \oplus K_2(\overline{\Theta}_\mathcal{L})$ and $K(\mathcal{M}) = K_1(\overline{\Theta}_\mathcal{M}) \oplus K_2(\overline{\Theta}_\mathcal{M})$ induced by $\overline{\Theta}_\mathcal{L}$ and $\overline{\Theta}_\mathcal{M}$ respectively.

Definition 5.2.1. We say that the level structures given by $s_\mathcal{L}$ and $s_\mathcal{M}$ are *compatible* (with respect to f) if:

(i) $\widetilde{K} = (\widetilde{K} \cap \widetilde{K_1}(s_\mathcal{L})) \oplus (\widetilde{K} \cap \widetilde{K_2}(s_\mathcal{L}))$.

(ii) α_f maps $\widetilde{K_i}(s_\mathcal{L}) \cap Z(\widetilde{K})$ to $\widetilde{K_i}(s_\mathcal{M})$ for $i \in \{1, 2\}$.

When $s_\mathcal{L}$ satisfies point (i) only, we say that $s_\mathcal{L}$ is compatible with $\widetilde{K}$.

We say that the theta structures $\Theta_\mathcal{L}$ and $\Theta_\mathcal{M}$ are compatible (with respect to f) if their induced level structures are compatible.

Lemma 5.2.2. *Assume that $s_\mathcal{L}$ is compatible with $\widetilde{K}$* i.e. *that $\widetilde{K} = s_\mathcal{L}(K) = (\widetilde{K} \cap \widetilde{K_1}(s_\mathcal{L})) \oplus (\widetilde{K} \cap \widetilde{K_2}(s_\mathcal{L}))$. Then:*

(i) $K = (K \cap K_1(\overline{\Theta}_\mathcal{L})) \oplus (K \cap K_2(\overline{\Theta}_\mathcal{L}))$ (we say that K is compatible with the decomposition $K(\mathcal{L}) = K_1(\overline{\Theta}_\mathcal{L}) \oplus K_2(\overline{\Theta}_\mathcal{L})$).

(ii) $K^\perp = (K^\perp \cap K_1(\overline{\Theta}_\mathcal{L})) \oplus (K^\perp \cap K_2(\overline{\Theta}_\mathcal{L}))$.

(iii) $s_\mathcal{L}(K^\perp \cap K_i(\overline{\Theta}_\mathcal{L})) = Z(\widetilde{K}) \cap \widetilde{K_i}(s_\mathcal{L})$ for $i \in \{1, 2\}$.

(iv) $Z(\widetilde{K}) = k^* \cdot (Z(\widetilde{K}) \cap \widetilde{K_1}(s_\mathcal{L})) \oplus (Z(\widetilde{K}) \cap \widetilde{K_2}(s_\mathcal{L}))$.

Proof. **(i)** Follows immediately from the assumption $\widetilde{K} = (\widetilde{K} \cap \widetilde{K_1}(s_\mathcal{L})) \oplus (\widetilde{K} \cap \widetilde{K_2}(s_\mathcal{L}))$.

(ii) Let $x \in K^\perp$ that we can write uniquely $x = x_1 + x_2$ with $x_i \in K_i(\overline{\Theta}_\mathcal{L})$ for $i \in \{1, 2\}$. Now, if $y \in K \cap K_1(\overline{\Theta}_\mathcal{L})$, we have

$$1 = e_\mathcal{L}(x, y) = e_\mathcal{L}(x_1, y)e_\mathcal{L}(x_2, y) = e_\mathcal{L}(x_2, y),$$

since $K_1(\overline{\Theta}_\mathcal{L})$ and K are isotropic, so x_2 is orthogonal to $K \cap K_1(\overline{\Theta}_\mathcal{L})$ and it is also orthogonal to $K \cap K_2(\overline{\Theta}_\mathcal{L})$ since $K_2(\overline{\Theta}_\mathcal{L})$ is isotropic. Hence $x_2 \in K^\perp$. Similarly, we get $x_1 \in K^\perp$. (ii) follows.

(iii) Let $y \in s_\mathcal{L}(K_i(\overline{\Theta}_\mathcal{L}))$. Then $y = s_\mathcal{L}(x)$ for some $x \in K_i(\overline{\Theta}_\mathcal{L})$. Let $y' \in \widetilde{K}$. Then there exists $x' \in K$ such that $y' = s_\mathcal{L}(x')$ (since $\widetilde{K} = s_\mathcal{L}(K)$ by assumption). We then have:

$$yy'y^{-1}y'^{-1} = e_\mathcal{L}(x, x'),$$

so $y \in Z(\widetilde{K})$ if and only if $x \in K^\perp$. (iii) follows.

(iv) If $y \in Z(\widetilde{K})$, then we have $\rho_\mathcal{L}(y) \in K^\perp$ as we saw in the proof of (iii). Since $ys_\mathcal{L}(\rho_\mathcal{L}(y))^{-1} \in \ker(\rho_\mathcal{L}) \simeq k^*$, we have $y \in k^* s_\mathcal{L}(K^\perp)$. Hence, the result follows from the equality $s_\mathcal{L}(K^\perp) = (Z(\widetilde{K}) \cap \widetilde{K_1}(s_\mathcal{L})) \oplus (Z(\widetilde{K}) \cap \widetilde{K_2}(s_\mathcal{L}))$ proved in point (iii). □

Lemma 5.2.3. *Assume that K is compatible with the decomposition of $K(\mathcal{L})$: $K = (K \cap K_1(\overline{\Theta}_\mathcal{L})) \oplus (K \cap K_2(\overline{\Theta}_\mathcal{L}))$. Then $s_\mathcal{L}$ is compatible with $\widetilde{K}$ if and only if $\alpha_f(\widetilde{K_i}(s_\mathcal{L}) \cap \rho_\mathcal{L}^{-1}(K)) = \{1\}$ for all $i \in \{1, 2\}$.*

Proof. Proposition 5.1.5 ensures that α_f induces an isomorphism $Z(\widetilde{K})/\widetilde{K} \xrightarrow{\sim} G(\mathcal{M})$. It follows that $\ker(\alpha_f) = \widetilde{K}$.

If $s_{\mathcal{L}}$ is compatible with $\widetilde{K}$, we have $\widetilde{K} = (\widetilde{K_1}(s_{\mathcal{L}}) \cap \rho_{\mathcal{L}}^{-1}(K)) \oplus (\widetilde{K_2}(s_{\mathcal{L}}) \cap \rho_{\mathcal{L}}^{-1}(K))$ so $\widetilde{K_i}(s_{\mathcal{L}}) \cap \rho_{\mathcal{L}}^{-1}(K) \subseteq \ker(\alpha_f)$ *i.e.* $\alpha_f(\widetilde{K_i}(s_{\mathcal{L}}) \cap \rho_{\mathcal{L}}^{-1}(K)) = \{1\}$ for all $i \in \{1,2\}$.

Conversely, if $\alpha_f(\widetilde{K_i}(s_{\mathcal{L}}) \cap \rho_{\mathcal{L}}^{-1}(K)) = \{1\}$ for all $i \in \{1,2\}$, then we have $(\widetilde{K_1}(s_{\mathcal{L}}) \cap \rho_{\mathcal{L}}^{-1}(K)) \oplus (\widetilde{K_2}(s_{\mathcal{L}}) \cap \rho_{\mathcal{L}}^{-1}(K)) \subseteq \ker(\alpha_f) = \widetilde{K}$. But $\rho_{\mathcal{L}}$ maps injectively $\widetilde{K}$ to K and we have $K = (K \cap K_1(\overline{\Theta}_{\mathcal{L}})) \oplus (K \cap K_2(\overline{\Theta}_{\mathcal{L}}))$, so the preceding inclusion must be an equality. This completes the proof. □

Proposition 5.2.4. *Let $(A, \mathcal{L}, \Theta_{\mathcal{L}})$ be a polarized abelian variety with a theta structure of type $\delta_{\mathcal{L}}$. Let $K \subseteq K(\mathcal{L})$ be a subgroup that we write $K = K_1 \oplus K_2$ with $K_i \subset K_i(\overline{\Theta}_{\mathcal{L}})$ for $i \in \{1,2\}$ and $f : A \longrightarrow B$ an isogeny of kernel K. Let $\widetilde{K}$ be the level subgroup above K given by:*

$$\widetilde{K} := s_{\mathcal{L}}(K) = (\widetilde{K_1}(s_{\mathcal{L}}) \cap \rho_{\mathcal{L}}^{-1}(K)) \oplus (\widetilde{K_2}(s_{\mathcal{L}}) \cap \rho_{\mathcal{L}}^{-1}(K)),$$

and $\mathcal{M}$ be the ample line bundle on B induced by $\widetilde{K}$ satisfying $f^*\mathcal{M} \simeq \mathcal{L}$ (see Theorem 5.1.4). We denote $\delta_{\mathcal{M}} \in \mathbb{Z}^g$ the type of $\mathcal{M}$.

Let $K^\perp = K^{\perp,1} \oplus K^{\perp,2}$ ($K^{\perp,i} = K^\perp \cap K_i(\overline{\Theta}_{\mathcal{L}})$ for $i \in \{1,2\}$) be the symplectic decomposition induced by the symplectic decomposition of $K(\mathcal{L})$. Then there is a bijection between theta structures of type $\delta_{\mathcal{M}}$ on $(B, \mathcal{M})$ compatible with $\Theta_{\mathcal{L}}$ and isomorphisms $\sigma : K^{\perp,1}/K_1 \xrightarrow{\sim} K_1(\delta_{\mathcal{M}})$.

Proof. Let $s_{\mathcal{L}} : K(\mathcal{L}) \longrightarrow G(\mathcal{L})$ be the section induced by $\Theta_{\mathcal{L}}$ (defining the level subgroups $\widetilde{K_i}(s_{\mathcal{L}})$). Consider the group $k^* \times K^{\perp,1} \times K^{\perp,2}$, equiped with the group law:

$$(\alpha, x_1, y_2) \cdot (\beta, x_2, y_2) := (\alpha\beta e_{\mathcal{L}}(x_1, y_2), x_1 + x_2, y_1 + y_2),$$

similar to the Heisenberg group law. Consider the map:

$$\begin{array}{ccc} k^* \times K^{\perp,1} \times K^{\perp,2} & \longrightarrow & Z(\widetilde{K}) \\ (\alpha, x, y) & \longmapsto & \alpha \cdot s_{\mathcal{L}}(x+y) \end{array}$$

which is a group homomorphism. This is actually an isomorphism by Lemma 5.2.2 (points (iii) and (iv)). Besides, it maps $\{1\} \times K_1 \times K_2$ to $s_{\mathcal{L}}(K) = \widetilde{K}$ so it induces an isomorphism

$$\Phi_{\mathcal{L}} : k^* \times K^{\perp,1}/K_1 \times K^{\perp,2}/K_2 \xrightarrow{\sim} Z(\widetilde{K})/\widetilde{K}.$$

If $x \in K^{\perp,1}$ and $y \in K^{\perp,2}$, then $e_{\mathcal{L}}(x,y)$ is invariant modulo K_1 on the left and modulo K_2 on the right. Hence

$$y \in K^{\perp,2} \longmapsto e_{\mathcal{L}}(.,y) \in \mathrm{Hom}(K^{\perp,1}, k^*)$$

induces a map $\varphi : K^{\perp,2}/K_2 \longrightarrow \mathrm{Hom}(K^{\perp,1}/K_1, k^*)$. If $\overline{y} \in K^{\perp,2}/K_2$ is in the kernel of this map then any representative $y \in K^{\perp,2}$ of $\overline{y}$ is orthogonal to $K^{\perp,1}$ and to $K^{\perp,2}$ since $K^{\perp,2} \subset K_2(\overline{\Theta}_{\mathcal{L}})$ is isotropic so $y \in K^{\perp\perp} = K$ (by Lemma 5.1.15.(ii)), and $\overline{y} = 0$. Hence, φ is injective. Now, let $\overline{\chi} \in \mathrm{Hom}(K^{\perp,1}/K_1, k^*)$. Then $\overline{\chi}$ induces a character $\chi \in \mathrm{Hom}(K^{\perp,1}, k^*)$ annihilating K_1. Let us extend χ to a charcter $\widetilde{\chi} \in \mathrm{Hom}(K(\mathcal{L}), k^*)$ annihilating $K_2(\overline{\Theta}_{\mathcal{L}})$ and write $\widetilde{\chi} = e_{\mathcal{L}}(.,y)$ for some $y \in K(\mathcal{L})$ (by non degeneracy of $e_{\mathcal{L}}$). Since $\widetilde{\chi}$ annihilates $K_2(\overline{\Theta}_{\mathcal{L}})$ which is maximal isotropic, $y \in K_2(\mathcal{L})$ and since $\widetilde{\chi}$ annihilates K_1 and $K_2 \subseteq K_2(\overline{\Theta}_{\mathcal{L}})$, we have $y \in K^\perp$, so $y \in K^{\perp,2}$. Hence φ is surjective, so it is an isomorphism.

Let σ be an isomorphism $K^{\perp,1}/K_1 \xrightarrow{\sim} K_1(\delta_{\mathcal{M}})$. Then $\widehat{\sigma} : y \longmapsto \varphi(y) \circ \sigma^{-1}$ is an isomorphism $K^{\perp,2}/K_2 \xrightarrow{\sim} K_2(\delta_{\mathcal{M}}) = \widehat{K_1}(\delta_{\mathcal{M}})$. We can then define a theta-structure $\Theta_{\mathcal{M}}$ on $\mathcal{M}$ making the following diagram commute:

$$\begin{array}{ccc} k^* \times K^{\perp,1}/K_1 \times K^{\perp,2}/K_2 & \xrightarrow{\Phi_{\mathcal{L}}} & Z(\widetilde{K})/\widetilde{K} \\ \downarrow{\scriptstyle \mathrm{id}_{k^*} \times \sigma \times \widehat{\sigma}} & & \downarrow{\scriptstyle \alpha_f} \\ \mathcal{H}(\delta_{\mathcal{M}}) & \xrightarrow{\Theta_{\mathcal{M}}} & G(\mathcal{M}) \end{array} \tag{5.12}$$

Now, we prove that $\Theta_{\mathcal{M}}$ is compatible with $\Theta_{\mathcal{L}}$. Indeed, $\widetilde{K} = s_{\mathcal{L}}(K)$ so point (i) of Definition 5.2.1 is satisfied. For point (ii), we have by construction of $\Theta_{\mathcal{M}}$:

$$\widetilde{K_1}(s_{\mathcal{M}}) = \Theta_{\mathcal{M}}(\{1\} \times K_1(\delta_{\mathcal{M}}) \times \{1\}) = \alpha_f(s_{\mathcal{L}}(K^{1,\perp})) = \alpha_f(\widetilde{K_1}(s_{\mathcal{L}}) \cap Z(\widetilde{K}))$$

the last equality following from Lemma 5.2.2.(iii). Similarly, we have $\widetilde{K_2}(s_{\mathcal{M}}) = \alpha_f(\widetilde{K_2}(s_{\mathcal{L}}) \cap Z(\widetilde{K}))$, so that $\Theta_{\mathcal{M}}$ and $\Theta_{\mathcal{L}}$ are indeed compatible.

Conversely, let $\Theta_{\mathcal{M}}$ be a theta-structure on $\mathcal{M}$ compatible with $\Theta_{\mathcal{L}}$. Then the canonical level subgroups defined by $\Theta_{\mathcal{M}}$ are $\widetilde{K_i}(s_{\mathcal{M}}) = \alpha_f(Z(\widetilde{K}) \cap \widetilde{K_i}(s_{\mathcal{L}}))$ for $i \in \{1, 2\}$. Hence, the level structure induced by $s_{\mathcal{M}}$ is completely determined and by Proposition 5.1.21 and $\Theta_{\mathcal{M}}$ only depends on a symplectic isomorphism $\overline{\Theta}_{\mathcal{M}} : K(\delta_{\mathcal{M}}) \longrightarrow K(\mathcal{M})$ respecting the symplectic decomposition $K(\mathcal{M}) = K_1(\overline{\Theta}_{\mathcal{M}}) \oplus K_2(\overline{\Theta}_{\mathcal{M}})$ with $K_i(\overline{\Theta}_{\mathcal{M}}) := \rho_{\mathcal{M}} \circ \alpha_f(Z(\widetilde{K}) \cap \widetilde{K_i}(s_{\mathcal{L}}))$ for $i \in \{1, 2\}$. By Proposition 5.1.17.(ii), $\overline{\Theta}_{\mathcal{M}}$ is determined by an isomorphism $K_1(\delta_M) \simeq K^{\perp,1}/K_1 \xrightarrow{\sim} K_1(\overline{\Theta}_{\mathcal{M}})$. □

5.2.2 The isogeny theorem

Theorem 5.2.5 (isogeny theorem)**.** *Let $f : (A, \mathcal{L}) \longrightarrow (B, \mathcal{M})$ be an isogeny between polarized abelian varieties of kernel K, $\Theta_{\mathcal{M}}$ and $\Theta_{\mathcal{L}}$ be compatible theta-structures on $(A, \mathcal{L})$ and $(B, \mathcal{M})$ respectively, and $\sigma : K^{\perp,1}/K_1 \xrightarrow{\sim} K_1(\delta_{\mathcal{M}})$ the isomorphism induced by $\Theta_{\mathcal{M}}$ as in Proposition 5.2.4.*

Let $(\theta_i^{\mathcal{L}})_{i \in K_1(\delta_{\mathcal{L}})}$ and $(\theta_i^{\mathcal{M}})_{i \in K_1(\delta_{\mathcal{M}})}$ the theta functions associated to $\Theta_{\mathcal{L}}$ and $\Theta_{\mathcal{M}}$ respectively. Then, there exists $\lambda \in k^*$ such that for all $i \in K_1(\delta_{\mathcal{M}})$,

$$f^*\theta_i^{\mathcal{M}} = \lambda \sum_{j \in \overline{\Theta}_{\mathcal{L}}^{-1}(\sigma^{-1}(\{i\}))} \theta_j^{\mathcal{L}}.$$

Proof. By Corollary 5.1.26, two sections of $\Gamma(A, \mathcal{L})$ differ by a constant if they are stable by the action of a maximal level subgroup of $G(\mathcal{L})$. We first need to prove that compatibility between theta-structures ensure the following compatibility between the $G(\mathcal{L})$ and $G(\mathcal{M})$-action (as defined in Definition 5.1.27):

$$\forall s \in \Gamma(A, \mathcal{M}), x \in K^{\perp}, \quad s_{\mathcal{L}}(x) \cdot f^*s = f^*(s_{\mathcal{M}}(f(x)) \cdot s).$$

Let $s \in \Gamma(A, \mathcal{M})$, $x \in K^{\perp}$ and $y := f(x)$. By compatibility of $\Theta_{\mathcal{L}}$ or $\Theta_{\mathcal{M}}$, we have $\alpha_f(s_{\mathcal{L}}(x)) = s_{\mathcal{M}}(y)$ (this is point (ii) of Definition 5.2.1 or a clear consequence of the construction of $\Theta_{\mathcal{M}}$ in the proof of Proposition 5.2.4). Let us write $s_{\mathcal{L}}(x) := (x, \phi_{\mathcal{L}}(x))$ and $s_{\mathcal{M}}(y) := (y, \phi_{\mathcal{M}}(y))$. Then, by the definition of α_f, we have (after identification $\mathcal{L} = f^*\mathcal{M}$)

$$\phi_{\mathcal{L}}(x)(f^*s) = f^*\phi_{\mathcal{M}}(y)(f^*s) = f^*(\phi_{\mathcal{M}}(y)(s)),$$

where we used the fact that $s_{\mathcal{L}}(x)$ is the unique element of $G(\mathcal{L})$ over x in $\alpha_f^{-1}(s_{\mathcal{M}}(y))$ (see the definition of α_f in the proof of Proposition 5.1.5). Hence, we have

$$s_{\mathcal{L}}(x) \cdot f^*s = t_{-x}^*(\phi_{\mathcal{L}}(x)(f^*s)) = t_{-x}^*(f^*(\phi_{\mathcal{M}}(y)(s))) = f^*(t_{-y}^*(\phi_{\mathcal{M}}(y)(s))) = f^*(s_{\mathcal{M}}(y) \cdot s).$$

We have in particular, for all $x \in K^{\perp}$,

$$s_{\mathcal{L}}(x) \cdot f^*\theta_0^{\mathcal{M}} = f^*(s_{\mathcal{M}}(f(x)) \cdot \theta_0^{\mathcal{M}}) = f^*\beta_{\mathcal{M}}(\Theta_{\mathcal{M}}^{-1}(s_{\mathcal{M}}(f(x))) \cdot \delta_0),$$

δ_0 is the Kronecker delta function at 0 on $K_1(\delta_{\mathcal{M}})$ and $\beta_{\mathcal{M}}$ is an isomorphism $V(\delta_{\mathcal{M}}) \xrightarrow{\sim} \Gamma(B, \mathcal{M})$ compatible with the actions of $G(\mathcal{M})$ and $\mathcal{H}(\delta_{\mathcal{M}})$ (defined in Section 5.1.5). Besides, by the definition of $\Theta_{\mathcal{M}}$,

$$\Theta_{\mathcal{M}}^{-1}(s_{\mathcal{M}}(f(x))) = (\mathrm{id}_{k^*} \times \sigma \times \widehat{\sigma}) \circ \Phi_{\mathcal{L}}^{-1} \circ \alpha_f^{-1}(s_{\mathcal{M}}(f(x))) = (1, \sigma(x_1), \widehat{\sigma}(x_2)),$$

where $x := x_1 + x_2$, with $x_i \in K^{\perp,i}$ for $i \in \{1, 2\}$. Since δ_0 is invariant by the action of $s_{\delta_{\mathcal{L}}}(K_2(\delta_{\mathcal{L}}))$ by Eq. (5.10) and σ annihilates K_1, $f^*\theta_0^{\mathcal{M}}$ is stable by the action of $s_{\mathcal{L}}(K_1 \oplus K^{\perp,2})$.

Let $x \in K_1 \oplus K^{\perp,2}$, that we write $x = \overline{\Theta}_{\mathcal{L}}(i_0, j_0)$, then we have

$$s_{\mathcal{L}}(x) \cdot \sum_{j \in \overline{\Theta}_{\mathcal{L}}^{-1}(K_1)} \theta_j^{\mathcal{L}} = \sum_{j \in \overline{\Theta}_{\mathcal{L}}^{-1}(K_1)} j_0(j + i_0)^{-1}\theta_{j+i_0}^{\mathcal{L}} = \sum_{j \in \overline{\Theta}_{\mathcal{L}}^{-1}(K_1)} j_0(j)^{-1}\theta_j^{\mathcal{L}}$$

with $j_0(j) = e_\delta((j,1),(0,j_0)) = 1$ for all $j \in \overline{\Theta}_{\mathcal{L}}^{-1}(K_1)$, since K_1 and $K^{\perp,2}$ are orthogonal. Hence, $\sum_{j\in\overline{\Theta}_{\mathcal{L}}^{-1}(K_1)} \theta_j^{\mathcal{L}}$ is stable by the action of $s_{\mathcal{L}}(K_1 \oplus K^{\perp,2})$.

We have seen in the proof of Proposition 5.2.4 that $K^{\perp,2}/K_2 \simeq \widehat{K^{\perp,1}/K_1}$, so that

$$\#K^{\perp,2}/\#K_2 = \#K^{\perp,1}/\#K_1 \quad i.e. \quad \#K^{\perp,2}\cdot\#K_1 = \#K^{\perp,1}\cdot\#K_2.$$

But we also have $\#K\cdot\#K^{\perp} = \#K(\mathcal{L}) = \#K_1(\mathcal{L})^2$ by Lemma 5.1.15.(i), with $\#K = \#K_1\cdot\#K_2$ and $\#K^{\perp} = \#K^{\perp,1}\cdot\#K^{\perp,2}$, so that

$$\#K^{\perp,2}\cdot\#K_1 = \#K^{\perp,1}\cdot\#K_2 = \#K_1(\mathcal{L}).$$

Hence, $K_1 \oplus K^{\perp,2}$ is maximal isotropic, so we can apply Theorem 5.1.25.(ii) and conclude that there exists $\lambda \in k^*$ such that

$$f^*\theta_0^{\mathcal{M}} = \lambda \sum_{j\in\overline{\Theta}_{\mathcal{L}}^{-1}(K_1)} \theta_j^{\mathcal{L}}.$$

To conclude, we get by Eq. (5.10) that for all $i \in K_1(\delta_{\mathcal{M}})$, and $i_0 \in \overline{\Theta}_{\mathcal{L}}^{-1}(\sigma^{-1}(\{i\}))$,

$$f^*\theta_i^{\mathcal{M}} = f^*(\Theta_{\mathcal{M}}(1,i,1)\cdot\theta_0^{\mathcal{M}}) = \Theta_{\mathcal{L}}(1,i_0,1)\cdot f^*\theta_0^{\mathcal{M}} = \lambda\Theta_{\mathcal{L}}(1,i_0,1)\cdot\sum_{j\in\overline{\Theta}_{\mathcal{L}}^{-1}(K_1)}\theta_j^{\mathcal{L}}$$

$$= \lambda\sum_{j\in\overline{\Theta}_{\mathcal{L}}^{-1}(K_1)}\theta_{j+i_0}^{\mathcal{L}} = \lambda\sum_{j\in\overline{\Theta}_{\mathcal{L}}^{-1}(\sigma^{-1}(\{i\}))}\theta_j^{\mathcal{L}}.$$

□

5.3 Symmetric theta structures and arithmetic applications

In this section, we assume that $\mathrm{char}(k) \neq 2$. We introduce the theory of symmetric theta structures which are well suited for arithmetic applications. In particular, in Section 5.3.2, we obtain a duplication formula which yields algorithms not only to double points (as its name indicate) but also to perform differential addition of points. This formula will also be used to change the level of theta coordinates for isogeny evaluation in Section 6.1.1 and obtain change of theta coordinates in Section 6.2. Finally, we present formulas to compute level 2 symmetric theta structures on Montgomery curves and the associated theta coordinates.

5.3.1 The theory of symmetric theta structures

In this section that introduces the theory of *symmetric theta structures*, we start with the notion of *symmetric* and *totally symmetric line bundles*. We then define and study symmetric theta structures on theta groups of symmetric line bundles. When $\mathcal{L}$ is a totally symmetric line bundle, we introduce compatibility conditions between symmetric theta structures on $G(\mathcal{L})$ and $G(\mathcal{L}^2)$. Our main result is Theorem 5.3.30. In the next section, we shall use these compatibility conditions to obtain the duplication formula. Finally, we prove that isogenies map symmetric theta structures to symmetric theta structures and prove compatibility conditions on symmetric theta structures with respect to an isogeny. These properties will be essential in Chapter 6.

Symmetric line bundles

Definition 5.3.1. A line bundle $\mathcal{L}$ on A is symmetric if $[-1]^*\mathcal{L} \simeq \mathcal{L}$ and antisymmetric if $[-1]^*\mathcal{L} \simeq \mathcal{L}^{-1}$.

If $\mathcal{L}$ is a line bundle on A, then we immediately obtain that $\mathcal{L}\otimes[-1]^*\mathcal{L}$ is symmetric and $\mathcal{L}\otimes[-1]^*\mathcal{L}^{-1}$ is antisymmetric. We recall that being antisymmetric is equivalent to belonging to $\mathrm{Pic}^0(A)$ by Lemma 1.4.65.(iv), so that $\mathcal{L}\otimes[-1]^*\mathcal{L}^{-1} \in \mathrm{Pic}^0(A)$.

Proposition 5.3.2. *(i) The symmetric line bundles of* $\mathrm{Pic}^0(A)$ *correspond to the points of* $\widehat{A}[2]$.

(ii) *Every ample line bundle $\mathcal{L}$ on A is algebraically equivalent to a symmetric line bundle: there exists $\mathcal{L}_0 \in \mathrm{Pic}^0(A)$ such that $\mathcal{L} \otimes \mathcal{L}_0^{-1}$ is symmetric.*

(iii) *If $\mathcal{L}$ is an ample line bundle on A, then the set of symmetric line bundles algebraically equivalent to $\mathcal{L}$ is*

$$\{\mathcal{L}' \otimes \varphi_{\mathcal{L}}(x) \mid x \in \varphi_{\mathcal{L}}^{-1}(\widehat{A}[2]) = [2]^{-1}K(\mathcal{L})\},$$

where $\mathcal{L}'$ is any symmetric line bundle algebraically equivalent to $\mathcal{L}$. This set contains 2^{2g} elements (where $g := \dim(A)$).

Proof. Let $\mathcal{L}_0 \in \mathrm{Pic}^0(A)$. Then $\mathcal{L}_0$ is antisymmetric by Lemma 1.4.65.(iv), so it is antisymmetric if and only if $\mathcal{L}_0 \simeq [-1]^*\mathcal{L}_0 \simeq \mathcal{L}_0^{-1}$ *i.e.* $\mathcal{L}_0 \in \widehat{A}[2]$. This proves (i).

Let $\mathcal{L}$ be an ample line bundle on A. By Lemma 1.4.65.(iv), $\mathcal{L} \otimes [-1]^*\mathcal{L}^{-1} \in \mathrm{Pic}^0(A)$ so $[-1]^*\mathcal{L} \otimes \mathcal{L}^{-1} \in \mathrm{Pic}^0(A)$. $\mathcal{L}$ being ample, $\varphi_{\mathcal{L}} : x \in A(k) \longmapsto t_x^*\mathcal{L} \otimes \mathcal{L}^{-1} \in \mathrm{Pic}^0(A)$ is an isogeny by Proposition 1.4.46.(iii) so it is surjective and there exists $x \in A(k)$ such that $[-1]^*\mathcal{L} \otimes \mathcal{L}^{-1} \simeq t_x^*\mathcal{L} \otimes \mathcal{L}^{-1}$ *i.e.* $[-1]^*\mathcal{L} \simeq t_x^*\mathcal{L}$. Let $y \in A$ such that $x = 2y$. Then $t_y^*\mathcal{L}$ is symmetric. Indeed

$$[-1]^*t_y^*\mathcal{L} = t_{-y}^*[-1]^*\mathcal{L} \simeq t_{-y}^*t_x^*\mathcal{L} = t_{x-y}^*\mathcal{L} = t_y^*\mathcal{L}$$

and $t_y^*\mathcal{L} \otimes \mathcal{L}^{-1} = \varphi_{\mathcal{L}}(y) \in \mathrm{Pic}^0(A)$ so $\mathcal{L}$ is algebraically equivalent to $t_y^*\mathcal{L}$. This proves (ii).

Let $\mathcal{L}'$ be symmetric and algebraically equivalent to $\mathcal{L}$. Then any other symmetric line bundle equivalent to $\mathcal{L}$ is also equivalent to $\mathcal{L}'$ so it is of the form $\mathcal{L}' \otimes \mathcal{L}_0$ with $\mathcal{L}_0 \in \mathrm{Pic}^0(A)$ symmetric (since $\mathcal{L}'$ and $\mathcal{L}' \otimes \mathcal{L}_0$ are symmetric). We can then write $\mathcal{L}_0 = \varphi_{\mathcal{L}}(x)$ with $x \in \varphi_{\mathcal{L}}^{-1}(\widehat{A}[2]) = [2]^{-1}K(\mathcal{L})$ by (u). Conversely, any element of the form $\mathcal{L}' \otimes \varphi_{\mathcal{L}}(x)$ with $x \in \varphi_{\mathcal{L}}^{-1}(\widehat{A}[2])$ is symmetric and algebraically equivalent to $\mathcal{L}$. This proves that the set of symmetric line bundles algebraically equivalent to $\mathcal{L}$ is

$$\{\mathcal{L}' \otimes \varphi_{\mathcal{L}}(x) \mid x \in \varphi_{\mathcal{L}}^{-1}(\widehat{A}[2]) = [2]^{-1}K(\mathcal{L})\},$$

Since this set is in bijection with $\widehat{A}[2]$, its cardinality is $\#A[2] = 2^{2g}$ by Corollary 1.4.34 since $\mathrm{char}(k) \neq 2$. □

Normalization

Let $\mathcal{L}$ be a symmetric line bundle on A. Then, there exists an isomorphism $\Phi : \mathcal{L} \xrightarrow{\sim} [-1]^*\mathcal{L}$. For every closed point $x \in A(k)$, Φ induces an isomorphism

$$\Phi(x) : \mathcal{L}(x) \xrightarrow{\sim} [-1]^*\mathcal{L}(x) = \mathcal{L}(-x),$$

where $\mathcal{L}(x) := \mathcal{L}_x \otimes_{\mathcal{O}_{A,x}} \kappa(x)$, as defined in Definition 5.1.31.

Definition 5.3.3. We say that Φ is *normalized* when $\Phi(0) : \mathcal{L}(0) \xrightarrow{\sim} \mathcal{L}(0)$ is the identity.

Requiring Φ to be normalized uniquely determines Φ. Indeed, the composition

$$\mathcal{L} \xrightarrow{\Phi} [-1]^*\mathcal{L} \xrightarrow{[-1]^*\Phi} [-1]^*[-1]^*\mathcal{L} = \mathcal{L}$$

is an automorphism of $\mathcal{L}$ so it corresponds to the multiplication by a scalar $\lambda \in k^*$ (as we have seen in the proof of Lemma 5.1.2). Since Φ is normalized, we must have $\lambda = 1$ so $[-1]^*\Phi \circ \Phi$ is the identity. The equality $[-1]^*\Phi \circ \Phi = \mathrm{id}_{\mathcal{L}}$ determines Φ up to sign but, again, since Φ is normalised, Φ is actually uniquely determined.

Definition 5.3.4. Let $\mathcal{L}$ be a symmetric line bundle on A and $\Phi : \mathcal{L} \xrightarrow{\sim} [-1]^*\mathcal{L}$ a normalized isomorphism. If $x \in A[2]$, then $\Phi(x)$ is an automorphism $\mathcal{L}(x) \xrightarrow{\sim} \mathcal{L}(-x) = \mathcal{L}(x)$, so it corresponds to the multiplication by a scalar in k^* (since $\mathcal{L}(x) \simeq k$). We denote by $e_*^{\mathcal{L}}(x)$ this scalar.

Proposition 5.3.5. *1. The map $e_*^{\mathcal{L}} : A[2] \longrightarrow k^*$ takes values in $\{\pm 1\}$.*

2. If $\mathcal{L}$ and $\mathcal{M}$ are symmetric line bundles on A, then $e_*^{\mathcal{L}\otimes\mathcal{M}} = e_*^{\mathcal{L}} \cdot e_*^{\mathcal{M}}$.

3. If $f : A \longrightarrow B$ is a morphism, and $\mathcal{L}$ is a symmetric line bundle on B, then $e_*^{f^*\mathcal{L}}(x) = e_*^{\mathcal{L}}(f(x))$ for all $x \in A[2]$.

4. *If $\mathcal{L} \in \mathrm{Pic}^0(A)$ is a symmetric line bundle corresponding to $y \in \widehat{A}[2]$, then we have $e_*^{\mathcal{L}}(x) = e_2(x,y)$ for all $x \in A[2]$, where $e_2 : A[2] \times \widehat{A}[2] \longrightarrow \{\pm 1\}$ is the Weil pairing.*

5. *$e_*^{\mathcal{L}}$ is the quadratic form associated to the commutator pairing $e_{\mathcal{L}^2}$ on $A[2] \times A[2]$. In other words, we have:*
$$\forall x, y \in A[2], \quad e_{\mathcal{L}^2}(x,y) = e_*^{\mathcal{L}}(x+y)e_*^{\mathcal{L}}(x)^{-1}e_*^{\mathcal{L}}(y)^{-1}.$$

Proof. Since Φ is normalized, $[-1]^*\Phi \circ \Phi = \mathrm{id}_{\mathcal{L}}$ so for all $x \in A[2]$, $\Phi(-x) \circ \Phi(x) = \mathrm{id}_{\mathcal{L}(x)}$ *i.e.* $\Phi(x)^2 = \mathrm{id}_{\mathcal{L}(x)}$ *i.e.* $e_*^{\mathcal{L}}(x)^2 = 1$ *i.e.* $e_*^{\mathcal{L}}(x) \in \{\pm 1\}$. This proves (i). Points (ii) and (iii) immediately follow from the definition of $e_*^{\mathcal{L}}$.

The proof of point (iv) follows [Mum66, p. 305]. Let $\mathcal{L} \in \mathrm{Pic}^0(A)$ symmetric corresponding to $y \in \widehat{A}[2]$. Then, we have $[2]^*\mathcal{L} \simeq \mathcal{L}^2$ by Corollary 1.4.18 since $\mathcal{L}$ is antisymmetric (as any element of $\mathrm{Pic}^0(A)$) and $\mathcal{L}^2 \simeq \mathcal{O}_A$ since $\mathcal{L}$ has order 2. Then, we have an isomorphism $\psi : [2]^*\mathcal{L} \xrightarrow{\sim} \mathcal{O}_A$ that we can construct explicitly. Let $D_y \in \mathrm{Pic}^0(A)$ be the divisor on A representing $y \in \widehat{A}(k)$, so that we may identify $\mathcal{L}$ with $\mathcal{L}(D_y)$, and let $g_y \in k(A)$ such that $\mathrm{div}(g_y) = [2]^*D_y$. Then, we may define $\psi : [2]^*\mathcal{L} \xrightarrow{\sim} \mathcal{O}_A$ as the division by g_y. Indeed, for all open subset $U \subseteq A$, we have

$$\Gamma(U, [2]^*\mathcal{L}) \simeq \Gamma(U, \mathcal{L}([2]^*D_y)) = \{f \in k(U) \mid \mathrm{div}(f) + \mathrm{div}(g_y)_{|U} \geq 0\} = g_{y|U}^{-1}\Gamma(U, \mathcal{O}_A)$$

It follows that $t_x^*\psi : t_x^*[2]^*\mathcal{L} \xrightarrow{\sim} t_x^*\mathcal{O}_A$ is the multiplication by $g_y \circ t_x$. Besides, $t_x^*[2]^*\mathcal{L} = ([2] \circ t_x)^*\mathcal{L} = [2]^*\mathcal{L}$ since $x \in A[2]$ so the following diagram commutes

$$\begin{array}{ccc} [2]^*\mathcal{L} & \xrightarrow{\psi} & \mathcal{O}_A \\ \| & & \downarrow{\scriptstyle e_2(x,y)} \\ t_x^*[2]^*\mathcal{L} & \xrightarrow[t_x^*\psi]{} & t_x^*\mathcal{O}_A, \end{array}$$

since we have $e_2(x,y) = g_y/g_y \circ t_x$. Now, let $z \in A(k)$ such that $2z = x$. Then, by localizing the diagram at $-z$, we get that $\psi(z) = e_2(x,y)\psi(-z)$.

Besides, we have the following diagram

$$\begin{array}{ccc} [2]^*\mathcal{L} & \xrightarrow{\psi} & \mathcal{O}_A \\ \downarrow{\scriptstyle [2]^*\Phi} & & \| \\ [-2]^*\mathcal{L} & \xrightarrow[[-1]^*\psi]{} & [-1]^*\mathcal{O}_A \end{array}$$

which is commutative since it commutes at 0 (Φ being normalized). Localizing this diagram at z, we get that $\psi(-z) = \Phi(x) \circ [2]^*\Phi(z)$ where $[2]^*\Phi(z) = \Phi(x)$ is the multiplication by $e_*^{\mathcal{L}}(x)$. This proves (iv).

Point (v) requires to introduce more material. This will be proved in Corollary 5.3.24. □

Totally symmetric line bundles

Definition 5.3.6. Let $\mathcal{L}$ be a symmetric line bundle on A. We say that $\mathcal{L}$ is *totally symmetric* if $e_*^{\mathcal{L}}(x) = 1$ for all $x \in A[2]$.

Proposition 5.3.7. *Let $\mathcal{L}$ be a symmetric line bundle on A and $\pi : A \longrightarrow K_A := A/\pm$ be the projection to the Kummer variety. Then $\mathcal{L}$ is totally symmetric if and only if $\mathcal{L}$ descends to K_A via π, namely if and only if there exists a line bundle $\mathcal{M}$ on K_A such that $\mathcal{L} = \pi^*\mathcal{M}$.*

Proof. See [Mum66, Proposition 1, p. 305]. □

Proposition 5.3.8. *(i) If $\mathcal{L}$ is an ample and separable line bundle on A of type δ with $2|\delta$, then there exists a unique totally symmetric line bundle algebraically equivalent to $\mathcal{L}$.*

(ii) A line bundle $\mathcal{L}$ is totally symmetric if and only if there exists a symmetric line bundle $\mathcal{M}$ such that $\mathcal{L} = \mathcal{M}^2$.

Proof. Since $2 \mid \delta$, we have $A[2] \subseteq K(\mathcal{L})$ so there exists an ample and separable line bundle $\mathcal{M}$ such that $\mathcal{L} = \mathcal{M}^2$ by Lemma 5.1.32. We may write $\mathcal{L} = (\mathcal{M} \otimes [-1]^*\mathcal{M}) \otimes (\mathcal{M} \otimes [-1]^*\mathcal{M}^{-1})$, with $\mathcal{M} \otimes [-1]^*\mathcal{M}^{-1} \in \mathrm{Pic}^0(A)$ by Lemma 1.4.65. Hence, $\mathcal{L}$ is algebraically equivalent to $\mathcal{M} \otimes [-1]^*\mathcal{M}$. Besides, $\mathcal{M} \otimes [-1]^*\mathcal{M}$ is symmetric and by Proposition 5.3.5, we have:

$$\forall x \in A[2], \quad e_*^{\mathcal{M}\otimes[-1]^*\mathcal{M}}(x) = e_*^{\mathcal{M}}(x)e_*^{\mathcal{M}}(-x) = e_*^{\mathcal{M}}(x)^2 = 1.$$

Hence, $\mathcal{M} \otimes [-1]^*\mathcal{M}$ is totally symmetric.

Now, we prove that such a totally symmetric line bundle is unique in the algebraic class of $\mathcal{L}$. Without loss of generality, we assume that $\mathcal{L}$ is totally symmetric. Then, the other symmetric line bundles algebraically equivalent to $\mathcal{L}$ are the $\mathcal{L} \otimes \mathcal{M}$, where $\mathcal{M} \in \mathrm{Pic}^0(A)$ corresponds to $y \in \widehat{A}[2]$ by Proposition 5.3.2. We then have by Proposition 5.3.5 and since $\mathcal{L}$ is totally symmetric:

$$\forall x \in A[2], \quad e_*^{\mathcal{L}\otimes\mathcal{M}}(x) = e_*^{\mathcal{L}}(x)e_*^{\mathcal{M}}(x) = e_*^{\mathcal{M}}(x) = e_2(x, y).$$

The Weil pairing e_2 being non-degenerate, $\mathcal{L} \otimes \mathcal{M}$ is totally symmetric if and only if $y = 0$ *i.e.* $\mathcal{M} \simeq \mathcal{O}_A$. This proves (i).

If $\mathcal{L}$ is the square of a symmetric line bundle, it is clear by Proposition 5.3.5.(ii) that $\mathcal{L}$ is totally symmetric. Conversely, we assume that $\mathcal{L}$ is totally symmetric. Then Proposition 5.3.5.(v) ensures that $e_{\mathcal{L}^2}$ is trivial on $A[2] \subseteq K(\mathcal{L}^2)$. Let ϵ be the type of $\mathcal{L}^2$. Then $2|\epsilon$ (since $A[2] \subseteq K(\mathcal{L}^2)$) and we may identify $e_{\mathcal{L}^2}$ with $e_\epsilon : K_1(\epsilon) \times K_2(\epsilon) \longrightarrow k^*$ defined in Eq. (5.5) by:

$$\forall (x_1, \chi_1), (x_2, \chi_2) \in K_1(\epsilon) \times K_2(\epsilon), \quad e_\epsilon((x_1, \chi_1), (x_2, \chi_2)) = \chi_2(x_1)\chi_1(x_2)^{-1}.$$

For all $i \in [\![1\ ;\ g]\!]$, let $x_i \in K_1(\epsilon) = \prod_{i=1}^g \mathbb{Z}/\epsilon_i\mathbb{Z}$ be equal to 1 at index i and 0 everywhere else, ζ_i be a ϵ_i-th primitive root of unity and $\chi_i \in K_2(\epsilon)$ be the character given by $\chi_i(x_i) = \zeta_i$ and $\chi_i(x_j) = 1$ for all $1 \leq j \neq i \leq g$. Then, $\epsilon_i/2x_i$ and $\chi_i^{\epsilon_i/2}$ have order 2, so we have

$$\zeta_i^{\epsilon_i^2/4} = \chi_i^{\epsilon_i/2}(\epsilon_i/2x_i) = e_\delta((x_i, 1), (0, \chi_i)) = 1.$$

It follows that $\epsilon_i|\epsilon_i^2/4$, so that $4|\epsilon_i$. Hence, $A[4] \subseteq K(\mathcal{L}^2)$ and $A[2] = [2]A[4] \subseteq [2]K(\mathcal{L}^2) = K(\mathcal{L})$ by Lemma 5.1.7. By Lemma 5.1.32, there exists an ample and separable line bundle $\mathcal{M}$ such that $\mathcal{L} = \mathcal{M}^2$. Let $\mathcal{M}_0$ be a symmetric line bundle algebraically equivalent to $\mathcal{M}$ (it does exist by Proposition 5.3.2.(ii)). Then $\mathcal{M}_0^2$ is totally symmetric and algebraically equivalent to $\mathcal{L}$. By uniqueness, we get $\mathcal{L} = \mathcal{M}_0^2$. This proves (ii).

□

Symmetric theta structures

Definition 5.3.9. Let $\mathcal{L}$ be a symmetric line bundle on A and $\Phi : \mathcal{L} \xrightarrow{\sim} [-1]^*\mathcal{L}$ be a normalized isomorphism. For all $(x, \phi_x) \in G(\mathcal{L})$, we consider the composition:

$$\mathcal{L} \xrightarrow{\Phi} [-1]^*\mathcal{L} \xrightarrow{[-1]^*\phi_x} [-1]^*t_x^*\mathcal{L} = t_{-x}^*[-1]^*\mathcal{L} \xrightarrow{(t_{-x}^*\Phi)^{-1}} t_{-x}^*\mathcal{L}$$

and define $\Delta_{-1} : G(\mathcal{L}) \longrightarrow G(\mathcal{L})$ by:

$$\forall (x, \phi_x) \in G(\mathcal{L}), \quad \Delta_{-1}(x, \phi_x) := (-x, (t_{-x}^*\Phi)^{-1} \circ [-1]^*\phi_x \circ \Phi).$$

Δ_{-1} is a group homomorphism and makes the following diagram commute:

$$\begin{array}{ccccccccc} 1 & \longrightarrow & k^* & \longrightarrow & G(\mathcal{L}) & \xrightarrow{\rho_{\mathcal{L}}} & K(\mathcal{L}) & \longrightarrow & 0 \\ & & \| & & \downarrow{\scriptstyle \Delta_{-1}} & & \downarrow{\scriptstyle [-1]} & & \\ 1 & \longrightarrow & k^* & \longrightarrow & G(\mathcal{L}) & \xrightarrow{\rho_{\mathcal{L}}} & K(\mathcal{L}) & \longrightarrow & 0 \end{array}$$

We say that $g \in G(\mathcal{L})$ is *symmetric* when $\Delta_{-1}(g) = g^{-1}$. We say that a subgroup $G \subseteq G(\mathcal{L})$ is *symmetric* if all its elements are symmetric.

Remark 5.3.10. Since $[-1]^*\Phi \circ \Phi = \mathrm{id}_{\mathcal{L}}$, we easily check that Δ_{-1} is an involution.

Lemma 5.3.11. *Let $\mathcal{L}$ be a symmetric line bundle on A, $\Theta_{\mathcal{L}}$ be a theta structure on $G(\mathcal{L})$ and $(\widetilde{K_1}(\Theta_{\mathcal{L}}), \widetilde{K_2}(\Theta_{\mathcal{L}}))$ be the associated level structure. Then $\Theta_{\mathcal{L}}$ is compatible with itself with respect to the isogeny $[-1]$ (in the sense of Definition 5.2.1) if and only if $\widetilde{K_1}(\Theta_{\mathcal{L}})$ and $\widetilde{K_2}(\Theta_{\mathcal{L}})$ are symmetric.*

Proof. Since $\ker([-1]) = \{0\}$, the only level subgroup above $\{0\}$ is $\widetilde{K} = \{(0, \mathrm{id}_{\mathcal{L}})\}$ and $Z(\widetilde{K}) = G(\mathcal{L})$, Definition 5.2.1 ensures that $\Theta_{\mathcal{L}}$ is compatible with itself with respect to $[-1]$ if and only if the map $\alpha_{[-1]}$ defined by Equation 5.1 maps $\widetilde{K_i}(\Theta_{\mathcal{L}})$ to itself for $i \in \{1,2\}$. Since $\alpha_{[-1]} = \Delta_{-1}^{-1}$ by Equation 5.1, this is equivalent to $\Delta_{-1}(\widetilde{K_i}(\Theta_{\mathcal{L}})) = \widetilde{K_i}(\Theta_{\mathcal{L}})$ for $i \in \{1,2\}$. This means that for all $x \in K_i(\overline{\Theta}_{\mathcal{L}})$ ($i \in \{1,2\}$), the unique element $\widetilde{x} \in \widetilde{K_i}(\Theta_{\mathcal{L}})$ lying above x satisfies $\Delta_{-1}(\widetilde{x}) \in \widetilde{K_i}(\Theta_{\mathcal{L}})$ so $\Delta_{-1}(\widetilde{x}) = \widetilde{x}^{-1}$ since $\widetilde{x}^{-1}$ is the only element of $\widetilde{K_i}(\Theta_{\mathcal{L}})$ lying above $-x$. This means that the $\widetilde{K_i}(\Theta_{\mathcal{L}})$ ($i \in \{1,2\}$) are symmetric. □

Definition 5.3.12. If $\mathcal{L}$ has level δ, we define the analogue of Δ_{-1} on the Heisenberg group $\mathcal{H}(\delta)$ by:

$$\forall (\alpha, x, \chi) \in \mathcal{H}(\delta), \quad D_{-1}(\alpha, x, \chi) := (\alpha, -x, \chi^{-1}) = \frac{\alpha^2}{\chi(x)}(\alpha, x, \chi)^{-1}.$$

Similarly, we say that $h \in \mathcal{H}(\delta)$ is *symmetric* if $D_{-1}(h) = h^{-1}$.

The symmetric elements of $\mathcal{H}(\delta)$ are of the form (α, x, χ) with $\alpha^2 = \chi(x)$. It follows in particular that the canonical level subgroups $\widetilde{K_1}(\delta) := \{(1, x, 1) \mid x \in K_1(\delta)\}$ and $\widetilde{K_2}(\delta) := \{(1, 0, \chi) \mid \chi \in K_2(\delta)\}$ are symmetric.

Proposition 5.3.13. *Let $\Theta_{\mathcal{L}}$ be a theta structure associated to a symmetric line bundle $\mathcal{L}$ of level δ. Then, the following statements are equivalent:*

(i) The associated level subgroups $\widetilde{K_1}(\Theta_{\mathcal{L}})$ and $\widetilde{K_2}(\Theta_{\mathcal{L}})$ are symmetric.

(ii) $\Theta_{\mathcal{L}} \circ D_{-1} = \Delta_{-1} \circ \Theta_{\mathcal{L}}$.

If these assertions hold, we say that $\Theta_{\mathcal{L}}$ is a symmetric *theta structure.*

Proof. $(i) \implies (ii)$ Assume $\widetilde{K_1}(\Theta_{\mathcal{L}})$ and $\widetilde{K_2}(\Theta_{\mathcal{L}})$ are symmetric. Then, for all $(\alpha, x, \chi) \in \mathcal{H}(\delta)$, we have:

$$\begin{aligned}
\Theta_{\mathcal{L}} \circ D_{-1}(\alpha, x, \chi) &= \Theta_{\mathcal{L}}\left(\frac{\alpha^2}{\chi(x)}(\alpha, x, \chi)^{-1}\right) = \frac{\alpha^2}{\chi(x)}\Theta_{\mathcal{L}}(\alpha, x, \chi)^{-1} \\
&= \frac{\alpha^2}{\chi(x)}\left(\frac{\alpha}{\chi(x)} \cdot \Theta_{\mathcal{L}}(1, x, 1) \cdot \Theta_{\mathcal{L}}(1, 0, \chi)\right)^{-1} = \alpha \cdot \Theta_{\mathcal{L}}(1, 0, \chi)^{-1} \cdot \Theta_{\mathcal{L}}(1, x, 1)^{-1} \\
&= \alpha \cdot \Delta_{-1} \circ \Theta_{\mathcal{L}}(1, 0, \chi) \cdot \Delta_{-1} \circ \Theta_{\mathcal{L}}(1, x, 1) \quad \text{(since the } \widetilde{K_i}(\Theta_{\mathcal{L}}) \text{ are symmetric)} \\
&= \Delta_{-1} \circ \Theta_{\mathcal{L}}(\alpha, x, \chi),
\end{aligned}$$

so (ii) holds: $\Theta_{\mathcal{L}} \circ D_{-1} = \Delta_{-1} \circ \Theta_{\mathcal{L}}$.

$(ii) \implies (i)$ Assume $\Theta_{\mathcal{L}} \circ D_{-1} = \Delta_{-1} \circ \Theta_{\mathcal{L}}$. Then, for $i \in \{1,2\}$, we have:

$$\Delta_{-1}(\widetilde{K_i}(\Theta_{\mathcal{L}})) = \Delta_{-1} \circ \Theta_{\mathcal{L}}(\widetilde{K_i}(\delta)) = \Theta_{\mathcal{L}} \circ D_{-1}(\widetilde{K_i}(\delta)) = \Theta_{\mathcal{L}}(\widetilde{K_i}(\delta)) = \widetilde{K_i}(\Theta_{\mathcal{L}}),$$

where we used the fact that the canonical level subgroup $\widetilde{K_i}(\delta)$ is symmetric. Hence, $\widetilde{K_1}(\Theta_{\mathcal{L}})$ and $\widetilde{K_2}(\Theta_{\mathcal{L}})$ are symmetric and (i) holds. □

Lemma 5.3.14. *Let $\mathcal{L}$ be a symmetric line bundle on A and $x \in K(\mathcal{L})$. Then, there are only 2 symmetric elements $\pm\widetilde{x}$ lying above x.*

Proof. Let $x \in K$ and $\widetilde{x} \in G(\mathcal{L})$ be a lift of x. Then, $\Delta_{-1}(\widetilde{x})$ and $\widetilde{x}^{-1}$ both lie above $-x$ so there exists $\lambda \in k^*$ such that $\Delta_{-1}(\widetilde{x}) = \lambda\widetilde{x}^{-1}$. Let α be a square root of λ. Then, $\Delta_{-1}(\pm\widetilde{x}/\alpha) = \pm\alpha\lambda\widetilde{x}^{-1} = (\pm\widetilde{x}/\alpha)^{-1}$, so $\pm\widetilde{x}/\alpha$ are symmetric elements of $G(\mathcal{L})$ lying above x and they are the only ones. □

Lemma 5.3.15. *Let $\mathcal{L}$ be a symmetric line bundle on A and $x \in K(\mathcal{L})$ of order 2. Then, for every $\widetilde{x} \in G(\mathcal{L})$ lying above x, we have*

$$\Delta_{-1}(\widetilde{x}) = e_*^{\mathcal{L}}(x)\widetilde{x}.$$

Proof. Let $\Phi : \mathcal{L} \xrightarrow{\sim} [-1]^*\mathcal{L}$ be a normalized isomorphism. Let $x \in K(\mathcal{L})$ of order 2 and $\widetilde{x} := (x, \phi_x) \in G(\mathcal{L})$ lying above x. Then, $\Delta_{-1}(\widetilde{x}) = (x, (t_{-x}^*\Phi)^{-1} \circ [-1]^*\phi_x \circ \Phi)$ (since $-x = x$). Localizing $(t_{-x}^*\Phi)^{-1} \circ [-1]^*\phi_x \circ \Phi$ and ϕ_x at 0, we get maps $\mathcal{L}(0) \longrightarrow \mathcal{L}(-x) = \mathcal{L}(x)$ that differ by $\Phi(x)^{-1} : \mathcal{L}(x) \longrightarrow \mathcal{L}(-x) = \mathcal{L}(x)$ *i.e.* by $e_*^{\mathcal{L}}(x)$. The result follows. □

If $\mathcal{L}$ is symmetric, there does not always exist a symmetric theta structure on $G(\mathcal{L})$. It is the case, however, when $\mathcal{L}$ is totally symmetric by the following proposition. In general, the obstruction is a consequence of the preceding lemma.

Proposition 5.3.16. *Let $\mathcal{L}$ be a symmetric line bundle on A and K be an isotropic subgroup in $K(\mathcal{L})$. Then, the following conditions are equivalent:*

(i) There exists a symmetric level subgroup in $G(\mathcal{L})$ above K.

(ii) For all $x \in K[2]$, $e_*^{\mathcal{L}}(x) = 1$.

(iii) If $B = A/K$ and $f : A \longrightarrow B$ is the associated isogeny, there exists a symmetric line bundle $\mathcal{M}$ on B such that $f^*\mathcal{M} \simeq \mathcal{L}$.

In particular, if $\mathcal{L}$ is totally symmetric and if $K(\mathcal{L}) = K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$ is a symplectic decomposition, then we have a symmetric level structure $s_{\mathcal{L}} : K(\mathcal{L}) \longrightarrow G(\mathcal{L})$ lying above this decomposition. The other symmetric level structures are given by the action of the character $e_{\mathcal{L}}(c, \cdot)$ on $s_{\mathcal{L}}$, with $c \in A[2]$.

Proof. $(i) \iff (iii)$ Let $f : A \longrightarrow B$ be an isogeny of kernel K, $\mathcal{M}$ be a line bundle on B such that $f^*\mathcal{M} \simeq \mathcal{L}$ and α be an isomorphism $f^*\mathcal{M} \xrightarrow{\sim} \mathcal{L}$. Then, by Grothendieck's descent theorem (Theorem 5.1.4), the couple $(\mathcal{M}, \alpha)$ corresponds to a level subgroup $\widetilde{K}$ on $G(\mathcal{L})$ above K given by:

$$\widetilde{K} := \{(x, t_x^*\alpha \circ \alpha^{-1}) \mid x \in K\}.$$

Now, let $\Phi : \mathcal{L} \xrightarrow{\sim} [-1]^*\mathcal{L}$ be a normalized isomorphism and consider the isomorphism:

$$f^*[-1]^*\mathcal{M} = [-1]^*f^*\mathcal{M} \xrightarrow{[-1]^*\alpha} [-1]^*\mathcal{L} \xrightarrow{\Phi^{-1}} \mathcal{L}.$$

Then, the couple $([-1]^*\mathcal{M}, \Phi^{-1} \circ [-1]^*\alpha)$ corresponds to the level subgroup

$$\begin{aligned}\widetilde{K}' &:= \{(x, t_x^*(\Phi^{-1} \circ [-1]^*\alpha) \circ (\Phi^{-1} \circ [-1]^*\alpha)) \mid x \in K\} \\ &= \{(x, (t_x^*\Phi)^{-1} \circ [-1]^*(t_{-x}^*\alpha \circ \alpha) \circ \Phi) \mid x \in K\} \\ &= \{(-x, (t_{-x}^*\Phi)^{-1} \circ [-1]^*(t_x^*\alpha \circ \alpha) \circ \Phi) \mid x \in K\} = \Delta_{-1}(\widetilde{K}).\end{aligned}$$

Hence, $\widetilde{K}$ is symmetric if and only if $\widetilde{K}' = \Delta_{-1}(\widetilde{K}) = \widetilde{K}$ if and only if $\mathcal{M} \simeq [-1]^*\mathcal{M}$ (by Theorem 5.1.4) *i.e.* $\mathcal{M}$ is symmetric. This proves the equivalence $(i) \iff (iii)$.

$(i) \implies (ii)$ Let $\widetilde{K}$ be a symmetric level subgroup on $G(\mathcal{L})$ above K. Let $x \in K[2]$ and $\widetilde{x} \in \widetilde{K}$ the lift of x in $\widetilde{K}$. Then by Lemma 5.3.15, we have $\Delta_{-1}(\widetilde{x}) = e_*^{\mathcal{L}}(x)\widetilde{x}$ and $\Delta_{-1}(\widetilde{x}) = \widetilde{x}^{-1} = \widetilde{x}$ since $\widetilde{K}$ is symmetric and $\widetilde{x}$ has order 2. Hence, $e_*^{\mathcal{L}}(x) = 1$. This proves (ii).

$(ii) \implies (i)$ Assume that for all $x \in K[2]$, $e_*^{\mathcal{L}}(x) = 1$. To find a symmetric level subgroup $\widetilde{K}$ on $G(\mathcal{L})$ above K, we proceed as in the proof of Proposition 5.1.12. Since K is isotropic for $e_{\mathcal{L}}$, $\rho_{\mathcal{L}}^{-1}(K)$ is abelian by Proposition 5.1.12 so it suffices to lift a basis a basis of K with symmetric elements only.

Let $x \in K$ of order $n \in \mathbb{N}^*$ and $\widetilde{x} \in G(\mathcal{L})$ a symmetric lift of x (it does exist by Lemma 5.3.14). If n is even, write $n = 2m$ with $m \in \mathbb{N}^*$. Then, $\widetilde{x}^m$ is a symmetric element above $[m]x$, which has order 2, so Lemma 5.3.15 ensures that $\widetilde{x}^{-m} = \Delta_{-1}(\widetilde{x}^m) = e_*^{\mathcal{L}}(mx)\widetilde{x}^m = \widetilde{x}^m$, so $\widetilde{x}^m$ has order 2 and $\widetilde{x}$ has order $2m = n$. If n is odd, $\widetilde{x}^n$ is a symmetric element above $[n]x = 0$ which admits only 2 symmetric lifts $\pm 1 \in G(\mathcal{L})$ by Lemma 5.3.14. Hence, $\widetilde{x}^n = \pm 1$, so either $\widetilde{x}$ or $-\widetilde{x}$ is a symmetric lift of order n of x. This proves that we can find a symmetric lift of any element of K, which proves (i).

If $\mathcal{L}$ is totally symmetric and if $K(\mathcal{L}) = K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$ is a symplectic decomposition, then $K_1(\mathcal{L})$ and $K_2(\mathcal{L})$ satisfy (ii), so we have a symmetric level structure $s_{\mathcal{L}} : K(\mathcal{L}) \longrightarrow G(\mathcal{L})$ lying above this decomposition.

Let $s'_{\mathcal{L}} : K(\mathcal{L}) \longrightarrow G(\mathcal{L})$ be another symmetric level structure lying above $K(\mathcal{L}) = K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$. Then, by Lemma 5.1.24, $s'_{\mathcal{L}} = e_{\mathcal{L}}(c, \cdot) \cdot s_{\mathcal{L}}$ for some $c \in K(\mathcal{L})$. $s_{\mathcal{L}}$ and $s'_{\mathcal{L}}$ being symmetric, we must have $e_{\mathcal{L}}(c, \cdot) \in \{\pm 1\}$, so $c \in A[2]$. Conversely, since $\mathcal{L}$ is totally symmetric, we have $A[2] \subseteq K(\mathcal{L})$ by Proposition 5.3.8 and Lemma 5.1.7 so all conjugates of $s_{\mathcal{L}}$ by $c \in A[2]$ define symmetric level structures. □

Maps between theta groups of symmetric line bundles

In this technical paragraph, we introduce some maps relating theta groups $G(\mathcal{L}^n)$ when $\mathcal{L}$ is a symmetric line bundle. These maps will be useful to study the compatibility of theta-structures of $G(\mathcal{L}^n)$ when n varies. We shall also obtain Proposition 5.3.5.(v) as a corollary.

Definition 5.3.17. Let $\mathcal{L}$ be a separable line bundle on A and $n \in \mathbb{N}^*$ not divisible by $\mathrm{char}(k)$. Then we define the group homomorphism

$$\begin{array}{rcl} \varepsilon_n : G(\mathcal{L}) & \longrightarrow & G(\mathcal{L}^n) \\ (x, \phi_x) & \longmapsto & (x, \phi_x^{\otimes n}), \end{array}$$

where $\phi_x^{\otimes n}$ is the isomorphism $\mathcal{L}^n \xrightarrow{\sim} t_x^*\mathcal{L}^n$ induced by $\phi_x : \mathcal{L} \xrightarrow{\sim} t_x^*\mathcal{L}$, for all $(x, \phi_x) \in G(\mathcal{L})$.

ε_n makes the following diagram commute:

$$\begin{array}{ccccccccc} 1 & \longrightarrow & k^* & \longrightarrow & G(\mathcal{L}) & \xrightarrow{\rho_{\mathcal{L}}} & K(\mathcal{L}) & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle \lambda \mapsto \lambda^n} & & \downarrow{\scriptstyle \varepsilon_n} & & \downarrow & & \\ 1 & \longrightarrow & k^* & \longrightarrow & G(\mathcal{L}^n) & \xrightarrow{\rho_{\mathcal{L}^n}} & K(\mathcal{L}^n) & \longrightarrow & 0 \end{array}$$

Now, assume that $\mathcal{L}$ is a separable and symmetric line bundle on A and let $n \in \mathbb{N}^*$ not divisible by $\mathrm{char}(k)$. Then, we can define a map going the other way $G(\mathcal{L}^n) \longrightarrow G(\mathcal{L})$ as follows. By Corollary 1.4.18, since $\mathcal{L}$ is symmetric, there exists an isomorphism $\Psi : [n]^*\mathcal{L} \xrightarrow{\sim} \mathcal{L}^{n^2}$. For all $(x, \phi_x) \in G(\mathcal{L}^n)$, consider the isomorphism:

$$[n]^*\mathcal{L} \xrightarrow{\Psi} \mathcal{L}^{n^2} \xrightarrow{\phi_x^{\otimes n}} t_x^*\mathcal{L}^{n^2} \xrightarrow{t_x^*\Psi^{-1}} t_x^*[n]^*\mathcal{L} = [n]^* t_{nx}^*\mathcal{L}$$

and $\rho_x : \mathcal{L} \xrightarrow{\sim} t_{[n]x}^*\mathcal{L}$, the unique isomorphism such that $[n]^*\rho_x = t_x^*\Psi^{-1} \circ \phi_x^{\otimes n} \circ \Psi$. Note that this makes sense because $[n]x \in K(\mathcal{L})$ for all $x \in K(\mathcal{L}^n)$ by Lemma 5.1.7.

Definition 5.3.18. We denote $\eta_n(x, \phi_x) := ([n]x, \rho_x)$. η_n defines a map $G(\mathcal{L}^n) \longrightarrow G(\mathcal{L})$ that makes the following diagram commute:

$$\begin{array}{ccccccccc} 1 & \longrightarrow & k^* & \longrightarrow & G(\mathcal{L}^n) & \xrightarrow{\rho_{\mathcal{L}^n}} & K(\mathcal{L}^n) & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle \lambda \mapsto \lambda^n} & & \downarrow{\scriptstyle \eta_n} & & \downarrow{\scriptstyle [n]} & & \\ 1 & \longrightarrow & k^* & \longrightarrow & G(\mathcal{L}) & \xrightarrow{\rho_{\mathcal{L}}} & K(\mathcal{L}) & \longrightarrow & 0 \end{array}$$

Since $\Psi : [n]^*\mathcal{L} \xrightarrow{\sim} \mathcal{L}^{n^2}$ is defined up to multiplication by a scalar, η_n does not depend on it. Furthermore, it is a surjective homomorphism.

Definition 5.3.19. Let $\mathcal{L}$ be a symmetric line bundle on A and $n \in \mathbb{Z}$. We define for all $g \in G(\mathcal{L})$:

$$\Delta_n(g) := g^{n(n+1)/2} \cdot \Delta_{-1}(g)^{n(n-1)/2},$$

where Δ_{-1} has already been defined in Definition 5.3.9.

Lemma 5.3.20. *(i) Δ_n is a group homomorphism $G(\mathcal{L}) \longrightarrow G(\mathcal{L})$.*

(ii) Δ_n *makes the following diagram commute:*

$$\begin{array}{ccccccccc} 1 & \longrightarrow & k^* & \longrightarrow & G(\mathcal{L}) & \xrightarrow{\rho_{\mathcal{L}}} & K(\mathcal{L}) & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle \lambda\mapsto\lambda^{n^2}} & & \downarrow{\scriptstyle \Delta_n} & & \downarrow{\scriptstyle [n]} & & \\ 1 & \longrightarrow & k^* & \longrightarrow & G(\mathcal{L}) & \xrightarrow{\rho_{\mathcal{L}}} & K(\mathcal{L}) & \longrightarrow & 0 \end{array}$$

(iii) For all $n, m \in \mathbb{Z}$, $\Delta_n \circ \Delta_m = \Delta_{nm}$.

Proof. **(i)** Let $g := (x, \phi_x), h := (y, \phi_y) \in G(\mathcal{L})$. Then, by definition of the commutator pairing, we have $hg = e_{\mathcal{L}}(y,x)gh$, so that $(gh)^k = e_{\mathcal{L}}(y,x)^{k(k-1)/2} g^k h^k$ for all $k \in \mathbb{N}^*$ and:

$$\begin{aligned} \Delta_n(gh) &= (gh)^{\frac{n(n+1)}{2}} (\Delta_{-1}(gh))^{\frac{n(n-1)}{2}} \\ &= e_{\mathcal{L}}(y,x)^{\frac{n(n+1)(n^2+n-2)}{8}} g^{\frac{n(n+1)}{2}} h^{\frac{n(n+1)}{2}} \Delta_{-1}\left((gh)^{\frac{n(n-1)}{2}}\right) \\ &\qquad (\Delta_{-1} \text{ being a group homomorphism}) \\ &= e_{\mathcal{L}}(y,x)^{\frac{n(n+1)(n^2+n-2)+n(n-1)(n^2-n-2)}{8}} g^{\frac{n(n+1)}{2}} h^{\frac{n(n+1)}{2}} \Delta_{-1}(g)^{\frac{n(n-1)}{2}} \Delta_{-1}(h)^{\frac{n(n-1)}{2}} \end{aligned}$$

Besides, since $\Delta_{-1}(g)$ is a lift of $-x$, we have:

$$\begin{aligned} h^{\frac{n(n+1)}{2}} \Delta_{-1}(g)^{\frac{n(n-1)}{2}} &= e_{\mathcal{L}}\left(\frac{n(n+1)}{2}y, -\frac{n(n-1)}{2}x\right) \Delta_{-1}(g)^{\frac{n(n-1)}{2}} h^{\frac{n(n+1)}{2}} \\ &= e_{\mathcal{L}}(x,y)^{\frac{n^2(n-1)(n+1)}{4}} \Delta_{-1}(g)^{\frac{n(n-1)}{2}} h^{\frac{n(n+1)}{2}} \end{aligned}$$

Since $n(n+1)(n^2+n-2) + n(n-1)(n^2-n-2) = 2n^2(n-1)(n+1)$, we conclude that

$$\Delta_n(gh) = g^{\frac{n(n+1)}{2}} \Delta_{-1}(g)^{\frac{n(n-1)}{2}} h^{\frac{n(n+1)}{2}} \Delta_{-1}(h)^{\frac{n(n-1)}{2}} = \Delta_n(g)\Delta_n(h).$$

This proves (i).

(ii) This is trivial by the definition of Δ_n and Δ_{-1}.

(iii) Let $g \in G(\mathcal{L})$ and $n, m \in \mathbb{Z}$. Then, we have

$$\begin{aligned} \Delta_n \circ \Delta_m(g) &= \Delta_m(g)^{\frac{n(n+1)}{2}} \Delta_{-1} \circ \Delta_m(g)^{\frac{n(n-1)}{2}} \\ &= \left(g^{\frac{m(m+1)}{2}} \Delta_{-1}(g)^{\frac{m(m-1)}{2}}\right)^{\frac{n(n+1)}{2}} \left(\Delta_{-1}(g)^{\frac{m(m+1)}{2}} \Delta_{-1} \circ \Delta_{-1}(g)^{\frac{m(m-1)}{2}}\right)^{\frac{n(n-1)}{2}} \end{aligned}$$

But g and $\Delta_{-1}(g)$ commute because $e_{\mathcal{L}}(\rho_{\mathcal{L}}(g), \rho_{\mathcal{L}}(\Delta_{-1}(g))) = e_{\mathcal{L}}(\rho_{\mathcal{L}}(g), -\rho_{\mathcal{L}}(g)) = 1$ and Δ_{-1} is an involution by Remark 5.3.10, so that

$$\begin{aligned} \Delta_n \circ \Delta_m(g) &= g^{\frac{nm(n+1)(m+1)}{4}} \Delta_{-1}(g)^{\frac{nm(n+1)(m-1)}{4}} \Delta_{-1}(g)^{\frac{nm(n-1)(m+1)}{4}} g^{\frac{nm(n-1)(m-1)}{4}} \\ &= g^{\frac{nm(n+1)(m+1)+nm(n-1)(m-1)}{4}} + \Delta_{-1}(g)^{\frac{nm(n+1)(m-1)+nm(n-1)(m+1)}{2}} \\ &= g^{\frac{nm(nm+1)}{2}} \Delta_{-1}(g)^{\frac{nm(nm-1)}{2}} = \Delta_{nm}(g) \end{aligned}$$

This completes the proof. □

Proposition 5.3.21. *Assume that* $\mathcal{L}$ *is symmetric and that* $\mathrm{char}(k) \nmid n$. *For all* $m \in \mathbb{Z}$, *we denote by* $\Delta_m^{\mathcal{L}} : G(\mathcal{L}) \longrightarrow G(\mathcal{L})$ *and* $\Delta_m^{\mathcal{L}^n} : G(\mathcal{L}^n) \longrightarrow G(\mathcal{L}^n)$ *the morphisms defined in Definition 5.3.19. Then, we have:*

(i) $\Delta_{-1}^{\mathcal{L}} \circ \eta_n = \eta_n \circ \Delta_{-1}^{\mathcal{L}^n}$.

(ii) $\Delta_{-1}^{\mathcal{L}^n} \circ \varepsilon_n = \varepsilon_n \circ \Delta_{-1}^{\mathcal{L}}$.

(iii) $\eta_n \circ \varepsilon_n = \Delta_n^{\mathcal{L}}$.

(iv) $\varepsilon_n \circ \eta_n = \Delta_n^{\mathcal{L}^n}$.

Proof. See [Mum66, Proposition 5, p. 311]. □

Proposition 5.3.22. *Let $g \in G(\mathcal{L}^2)$ be an element of order 2 and $x := \rho_{\mathcal{L}^2}(g)$. Then $\eta_2(g) \in \ker(\rho_{\mathcal{L}^2}) \simeq k^*$ and $\eta_2(g)$ identifies with $e_*^{\mathcal{L}}(x)$.*

Proof. See [Mum66, Proposition 6, p. 312]. □

Corollary 5.3.23. *Let $\mathcal{L}$ be a symmetric line bundle on A. Then $\mathcal{L}$ is totally symmetric if and only if*

$$\ker(\eta_2) = \{g \in G(\mathcal{L}^2) \mid g^2 = 1\}.$$

Proof. $\Longleftarrow$ Assume that $\mathcal{L}$ is totally symmetric. Then $e_*^{\mathcal{L}}(x) = 1$ for all $x \in A[2]$, so by Proposition 5.3.22, $\ker(\eta_2)$ contains every element of order 2 of $G(\mathcal{L}^2)$.

Conversely, if $(x, \phi) \in \ker(\eta_2)$, then $[2]x = 0$ and $\phi^{\otimes 2} = t_x^*\Psi \circ \Psi^{-1}$, where Ψ is an isomorphism $[2]^*\mathcal{L} \xrightarrow{\sim} \mathcal{L}^4$. It follows that

$$t_x^*\phi^{\otimes 2} \circ \phi^{\otimes 2} = t_{2x}^*\Psi \circ t_x^*\Psi^{-1} \circ t_x^*\Psi \circ \Psi^{-1} = t_0^*\Psi \circ \Psi^{-1} = \mathrm{id}_{\mathcal{L}^4},$$

so that $t_x^*\phi \circ \phi = \pm \mathrm{id}_{\mathcal{L}^2}$. If $t_x^*\phi \circ \phi = -\mathrm{id}_{\mathcal{L}^2}$, then we may write $\phi := i\phi_0$ with $i^2 = -1$ and $t_x^*\phi_0 \circ \phi_0 = \mathrm{id}_{\mathcal{L}^2}$. Then, (x, ϕ_0) has order 2 so $\eta_2(x, \phi_0) = (0, \mathrm{id}_{\mathcal{L}})$ and $\eta_2(x, \phi) = (0, -\mathrm{id}_{\mathcal{L}}) \neq (0, \mathrm{id}_{\mathcal{L}})$. It follows that $t_x^*\phi \circ \phi = \mathrm{id}_{\mathcal{L}^2}$, so that (x, ϕ) has order 2.

$\Longrightarrow$ Assume that $\ker(\eta_2) = \{g \in G(\mathcal{L}^2) \mid g^2 = 1\}$. Then, for all $x \in A[2] \subseteq K(\mathcal{L}^2)$, we can find $\widetilde{x} \in G(\mathcal{L}^2)$ lying above x of order 2 and we have $e_*^{\mathcal{L}}(x) = \eta_2(\widetilde{x}) = 1$ by Proposition 5.3.22 (with canonical identifications). Hence, $\mathcal{L}$ is totally symmetric and the proof is complete. □

Corollary 5.3.24 (Proposition 5.3.5.(v))**.** *Let $\mathcal{L}$ be a symmetric line bundle on A. Then, we have:*

$$\forall x, y \in A[2], \quad e_*^{\mathcal{L}}(x+y) = e_*^{\mathcal{L}}(x)e_*^{\mathcal{L}}(y)e_{\mathcal{L}^2}(x, y).$$

Proof. Let $x, y \in A[2]$ and $g, h \in G(\mathcal{L}^2)$ be elements of order 2 lying above x and y respectively. Then $(gh)^2 = ghg^{-1}h^{-1} = e_{\mathcal{L}}(x, y)$. Let $\beta \in k^*$ be a square root of $e_{\mathcal{L}}(x, y)$. Then βgh has order 2 and lies above $x + y$, so we have by Proposition 5.3.22:

$$e_*^{\mathcal{L}}(x+y) = \eta_2(\beta gh) = \eta_2(\beta)\eta_2(g)\eta_2(h) = \beta^2 e_*^{\mathcal{L}}(x)e_*^{\mathcal{L}}(y) = e_*^{\mathcal{L}}(x)e_*^{\mathcal{L}}(y)e_{\mathcal{L}^2}(x, y).$$

The result follows. □

Compatible symmetric theta structures

Let $\mathcal{L}$ be a totally symmetric line bundle on A. In this section, we study theta structures on $G(\mathcal{L})$ and $G(\mathcal{L}^2)$ and their compatibility. The compatibility conditions we shall present are a requirement for the duplication formula we shall introduce in Section 5.3.2 for arithmetic applications.

Let $\delta := (d_1, \cdots, d_g)$ be the type of $\mathcal{L}$. Then, we have a natural embedding $K_1(\delta) \hookrightarrow K_1(2\delta)$ via the multiplication by 2 map and conversely, we have a natural surjective map $K_1(2\delta) \twoheadrightarrow K_1(\delta)$ mapping $x := (x_1 \mod 2d_1, \cdots, x_g \mod 2d_g)$ to $\overline{x} := (x_1 \mod d_1, \cdots, x_g \mod d_g)$. Looking at the dual, we also have a natural embedding $K_2(\delta) \hookrightarrow K_2(2\delta)$ mapping every $\chi \in K_2(\delta)$ to $2 \star \chi : x \in K_1(2\delta) \longmapsto \chi(\overline{x}) \in k^*$ and a surjective map $K_2(2\delta) \twoheadrightarrow K_2(\delta)$ mapping every $\chi \in K_2(2\delta)$ to $\overline{\chi} : x \in K_1(\delta) \longmapsto \chi(2x) \in k^*$.

Definition 5.3.25. We define $E_2 : \mathcal{H}(\delta) \longrightarrow \mathcal{H}(2\delta)$, $H_2 : \mathcal{H}(2\delta) \longrightarrow \mathcal{H}(\delta)$ and for all $n \in \mathbb{Z}$, $D_n : \mathcal{H}(\delta) \longrightarrow \mathcal{H}(\delta)$ respectively by:

$$\forall (\alpha, x, \chi) \in \mathcal{H}(\delta), \quad E_2(\alpha, x, \chi) := (\alpha^2, 2x, 2 \star \chi)$$

$$\forall (\alpha, x, \chi) \in \mathcal{H}(2\delta), \quad H_2(\alpha, x, \chi) := (\alpha^2, \overline{x}, \overline{\chi})$$

$$\forall (\alpha, x, \chi) \in \mathcal{H}(\delta), \quad D_n(\alpha, x, \chi) := (\alpha^{n^2}, nx, \chi^n)$$

E_2, H_2 and D_n are the analogues of ε_2, η_2 and Δ_n in Heisenberg groups. Indeed, we can prove that:

Lemma 5.3.26. *(i) E_2, H_2 and D_n are group homomorphisms;*

(ii) $E_2 \circ D_{-1}^{\mathcal{H}(\delta)} = D_{-1}^{\mathcal{H}(2\delta)} \circ E_2$ *and* $H_2 \circ D_{-1}^{\mathcal{H}(2\delta)} = D_{-1}^{\mathcal{H}(\delta)} \circ H_2$*;*

(iii) $E_2 \circ H_2 = D_2^{\mathcal{H}(2\delta)}$ *and* $H_2 \circ E_2 = D_2^{\mathcal{H}(\delta)}$*;*

(iv) For all $h \in \mathcal{H}(\delta)$, $D_n(h) = h^{n(n+1)/2} D_{-1}(h)^{n(n-1)/2}$*;*
where the exponents indicate the group of definition when there is an ambiguity.

Lemma 5.3.27. *Assume that* $2|\delta$*. Then:*

$$\ker(H_2) = \{h \in \mathcal{H}(2\delta) \mid h^2 = 1\}$$

Proof. By the definition of H_2,

$$\ker(H_2) = \{(\alpha, x, \chi) \in \mathcal{H}(2\delta) \mid \alpha^2 = 1, \overline{x} = 0 \text{ and } \overline{\chi} = 1\}.$$

We prove that these elements are exactly the elements of order 2 of $\mathcal{H}(2\delta)$. If $(\alpha, x, \chi) \in \ker(H_2)$, then $\overline{x} = 0$ so $d_i | x_i$ for all $i \in [\![1\ ;\ g]\!]$ (where $\delta := (d_1, \cdots, d_g)$). But $2|\delta$ so $2|x$. Besides, $\overline{\chi} = 1$ so $\chi(2y) = 1$ for all $y \in K_1(\delta)$. Now, if $y \in K_1(2\delta)$, we may write $y_i \equiv y_i' + \epsilon_i d_i \mod 2d_i$ with $0 \leq y_i' \leq d_i - 1$ and $\epsilon_i \in \{0, 1\}$ for all $i \in [\![1\ ;\ g]\!]$, so that $2y_i \equiv 2y_i' \mod 2d_i$, so that $2y = 2y'$ with $y' := (y_i')_{1 \leq i \leq g} \in K_1(\delta)$ and $\chi(2y) = \chi(2y') = 1$. Hence, $\chi^2 = 1$ and in particular $\chi(x) = 1$ since $2|x$. Hence, $(\alpha, x, \chi)^2 = (\alpha^2 \chi(x), 2x, \chi^2) = 1$. Conversely, if $(\alpha, x, \chi)^2 = 1$, then $2x = 0$ so $d_i | x_i$ for all $i \in [\![1\ ;\ g]\!]$ and $\overline{x} = 0$ and $2|x$. Furthermore, $\chi(2y) = 1$ for all $y \in K_1(2\delta)$ so $\chi(x) = 1$ and finally $\alpha^2 = 1$. Hence, $(\alpha, x, \chi) \in \ker(H_2)$. □

Definition 5.3.28. Let $\mathcal{L}$ be a totally symmetric line bundle. Let $\Theta_{\mathcal{L}}$ and $\Theta_{\mathcal{L}^2}$ be two theta-structures on $G(\mathcal{L})$ and $G(\mathcal{L}^2)$ respectively. We say they are *compatible* if both of the following diagrams commute:

$$\begin{array}{ccc} \mathcal{H}(\delta) & \xrightarrow{E_2} & \mathcal{H}(2\delta) \\ {\scriptstyle \Theta_{\mathcal{L}}}\downarrow & & \downarrow{\scriptstyle \Theta_{\mathcal{L}^2}} \\ G(\mathcal{L}) & \xrightarrow{\varepsilon_2} & G(\mathcal{L}^2) \end{array} \qquad \begin{array}{ccc} \mathcal{H}(2\delta) & \xrightarrow{H_2} & \mathcal{H}(\delta) \\ {\scriptstyle \Theta_{\mathcal{L}^2}}\downarrow & & \downarrow{\scriptstyle \Theta_{\mathcal{L}}} \\ G(\mathcal{L}^2) & \xrightarrow{\eta_2} & G(\mathcal{L}) \end{array}$$

We also say that $(\Theta_{\mathcal{L}}, \Theta_{\mathcal{L}^2})$ is a *pair of compatible theta structures* (for $(\mathcal{L}, \mathcal{L}^2)$).

Lemma 5.3.29. *If* $\Theta_{\mathcal{L}}$ *and* $\Theta_{\mathcal{L}^2}$ *are compatible theta-structures and* $\Theta_{\mathcal{L}^2}$ *is symmetric, then* $\Theta_{\mathcal{L}}$ *is symmetric.*

Proof. We assume that $\Theta_{\mathcal{L}^2}$ is symmetric. Then, by the compatibility of $\Theta_{\mathcal{L}}$ and $\Theta_{\mathcal{L}^2}$, Lemma 5.3.26.(ii) and Proposition 5.3.21, we have

$$\begin{aligned} \Delta_{-1}^{\mathcal{L}} \circ \Theta_{\mathcal{L}} \circ H_2 &= \Delta_{-1}^{\mathcal{L}} \circ \eta_2 \circ \Theta_{\mathcal{L}^2} = \eta_2 \circ \Delta_{-1}^{\mathcal{L}^2} \circ \Theta_{\mathcal{L}^2} = \eta_2 \circ \Theta_{\mathcal{L}^2} \circ D_{-1}^{\mathcal{H}(2\delta)} = \Theta_{\mathcal{L}} \circ H_2 \circ D_{-1}^{\mathcal{H}(2\delta)} \\ &= \Theta_{\mathcal{L}} \circ D_{-1}^{\mathcal{H}(\delta)} \circ H_2 \end{aligned}$$

and since H_2 is surjective, $\Delta_{-1}^{\mathcal{L}} \circ \Theta_{\mathcal{L}} = \Theta_{\mathcal{L}} \circ D_{-1}^{\mathcal{H}(\delta)}$ so $\Theta_{\mathcal{L}}$ is symmetric. □

Theorem 5.3.30. *(i) Every symmetric theta-structure* $\Theta_{\mathcal{L}^2}$ *on* $G(\mathcal{L}^2)$ *induces a unique symmetric theta-structure* $\Theta_{\mathcal{L}}$ *on* $G(\mathcal{L})$ *that is compatible with* $\Theta_{\mathcal{L}^2}$*.*

(ii) The resulting theta-structure $\Theta_{\mathcal{L}}$ *on* $G(\mathcal{L})$ *only depends on the symplectic isomorphism* $\overline{\Theta}_{\mathcal{L}^2} : K(2\delta) \xrightarrow{\sim} K(\mathcal{L}^2)$*.*

(iii) Every symmetric theta-structure on $G(\mathcal{L})$ *is induced by a symmetric theta-structure on* $G(\mathcal{L}^2)$*, or equivalently, by a symplectic isomorphism* $K(2\delta) \xrightarrow{\sim} K(\mathcal{L}^2)$*.*

Proof. **(i)** $\mathcal{L}$ is totally symmetric so $2|\delta$ by Proposition 5.3.8.(ii) and Lemma 5.1.7. Hence, by Lemma 5.3.27, we have:

$$\ker(H_2) = \{h \in \mathcal{H}(2\delta) \mid h^2 = 1\}.$$

Let $\Theta_{\mathcal{L}^2}$ be a symmetric theta-structure on $G(\mathcal{L}^2)$. Since $\mathcal{L}$ is totally symmetric, $\ker(\eta_2) = \{g \in G(\mathcal{L}^2) \mid g^2 = 1\}$ by Corollary 5.3.23 and $\Theta_{\mathcal{L}^2}$ is an isomorphism so $\ker(\eta_2 \circ \Theta_{\mathcal{L}^2}) = \ker(H_2)$ so $\eta_2 \circ \Theta_{\mathcal{L}^2}$ factors through H_2 and there exists an isomorphism $\Theta_{\mathcal{L}} : \mathcal{H}(\delta) \xrightarrow{\sim} G(\mathcal{L})$ such that $\eta_2 \circ \Theta_{\mathcal{L}^2} = \Theta_{\mathcal{L}} \circ H_2$.

We now verify that $\Theta_{\mathcal{L}}$ is a theta-structure. Since both H_2 and η_2 coincide with $\lambda \longmapsto \lambda^2$ on k^* and $\Theta_{\mathcal{L}^2}$ is the identity on k^*, $\Theta_{\mathcal{L}}$ is the identity on k^*. It remains to prove that the isomorphism $\overline{\Theta}_{\mathcal{L}} := \rho_{\mathcal{L}} \circ \Theta_{\mathcal{L}} \circ s_\delta : K(\delta) \xrightarrow{\sim} K(\mathcal{L})$ induced by $\Theta_{\mathcal{L}}$ is symplectic. Let $(x_1, \chi_1), (x_2, \chi_2) \in K(\delta)$. Then, for $i \in \{1, 2\}$, there exists $(x_i', \chi_i') \in K(2\delta)$ such that $\overline{x}_i' = x_i$ and $\overline{\chi}_i' = \chi_i$, so that $H_2 \circ s_{2\delta}(x_i', \chi_i') = s_\delta(x_i, \chi_i)$ and

$$\begin{aligned}\overline{\Theta}_{\mathcal{L}}(x_i, \chi_i) &= \rho_{\mathcal{L}} \circ \Theta_{\mathcal{L}} \circ s_\delta(x_i, \chi_i) = \rho_{\mathcal{L}} \circ \Theta_{\mathcal{L}} \circ H_2 \circ s_{2\delta}(x_i', \chi_i') = \rho_{\mathcal{L}} \circ \eta_2 \circ \Theta_{\mathcal{L}^2} \circ s_{2\delta}(x_i', \chi_i') \\ &= [2]\rho_{\mathcal{L}^2} \circ \Theta_{\mathcal{L}^2} \circ s_{2\delta}(x_i', \chi_i') = [2]\overline{\Theta}_{\mathcal{L}^2}(x_i', \chi_i').\end{aligned}$$

It follows that

$$\begin{aligned}e_{\mathcal{L}}(\overline{\Theta}_{\mathcal{L}}(x_1, \chi_1), \overline{\Theta}_{\mathcal{L}}(x_2, \chi_2)) &= e_{\mathcal{L}}([2]\overline{\Theta}_{\mathcal{L}^2}(x_1', \chi_1'), [2]\overline{\Theta}_{\mathcal{L}^2}(x_2', \chi_2')) \\ &= e_{\mathcal{L}^2}([2]\overline{\Theta}_{\mathcal{L}^2}(x_1', \chi_1'), \overline{\Theta}_{\mathcal{L}^2}(x_2', \chi_2')) \quad \text{(by Proposition 5.1.6.(v))} \\ &= e_{\mathcal{L}^2}(\overline{\Theta}_{\mathcal{L}^2}(x_1', \chi_1'), \overline{\Theta}_{\mathcal{L}^2}(x_2', \chi_2'))^2 \\ &= e_{2\delta}((x_1', \chi_1'), (x_2', \chi_2'))^2 = \chi_2'(x_1')^2 \chi_1'(x_2')^{-2} \\ &= \chi_2'(2x_1')\chi_1'(2x_2')^{-1} = \overline{\chi}_2'(\overline{x}_1')\overline{\chi}_1'(\overline{x}_2')^{-1} = \chi_2(x_1)\chi_1(x_2)^{-1} \\ &= e_\delta((x_1, \chi_1), (x_2, \chi_2)),\end{aligned}$$

so $\overline{\Theta}_{\mathcal{L}}$ is indeed symplectic and $\Theta_{\mathcal{L}}$ is a theta-structure.

We finally prove that $\Theta_{\mathcal{L}^2} \circ E_2 = \varepsilon_2 \circ \Theta_{\mathcal{L}}$. Let $h \in \mathcal{H}(\delta)$. Then, H_2 being surjective, there exists $h' \in \mathcal{H}(2\delta)$ such that $H_2(h') = h$. We then have

$$\begin{aligned}\varepsilon_2 \circ \Theta_{\mathcal{L}}(h) &= \varepsilon_2 \circ \Theta_{\mathcal{L}} \circ H_2(h') = \varepsilon_2 \circ \eta_2 \circ \Theta_{\mathcal{L}^2}(h') \\ &= \Delta_2^{\mathcal{L}^2} \circ \Theta_{\mathcal{L}^2}(h') \quad \text{(by Proposition 5.3.21.(iv))} \\ &= \Theta_{\mathcal{L}^2}(h')^3 \Delta_{-1}^{\mathcal{L}^2} \circ \Theta_{\mathcal{L}^2}(h') = \Theta_{\mathcal{L}^2}(h')^3 \Theta_{\mathcal{L}^2} \circ D_{-1}^{\mathcal{H}(2\delta)}(h') \quad \text{(since } \Theta_{\mathcal{L}^2} \text{ is symmetric)} \\ &= \Theta_{\mathcal{L}^2}(h'^3 D_{-1}^{\mathcal{H}(2\delta)}(h')) = \Theta_{\mathcal{L}^2} \circ D_2^{\mathcal{H}(2\delta)}(h') = \Theta_{\mathcal{L}^2} \circ E_2 \circ H_2(h') \quad \text{(by Lemma 5.3.26.(iii))} \\ &= \Theta_{\mathcal{L}^2} \circ E_2(h)\end{aligned}$$

We also obtain that $\Theta_{\mathcal{L}}$ is symmetric by Lemma 5.3.29. The uniqueness of $\Theta_{\mathcal{L}}$ follows from the equality $\eta_2 \circ \Theta_{\mathcal{L}^2} = \Theta_{\mathcal{L}} \circ H_2$ and the surjectivity of H_2. This proves (i).

(ii) Let $\Theta_{\mathcal{L}^2}$ and $\Theta'_{\mathcal{L}^2}$ be two symmetric theta symmetric theta-structures on $G(\mathcal{L}^2)$ inducing the same symplectic isomorphism $\overline{\Theta}_{\mathcal{L}^2} : K(2\delta) \xrightarrow{\sim} K(\mathcal{L}^2)$ and let $\Theta_{\mathcal{L}}$ and $\Theta'_{\mathcal{L}}$ be the compatible theta-structures on $G(\mathcal{L})$ they respectively induce. By Lemma 5.1.24, $\Theta'_{\mathcal{L}^2} = \xi \cdot \Theta_{\mathcal{L}^2}$, where ξ is a character $\xi : K(2\delta) \longrightarrow k^*$ and since both associated level structures $(\widetilde{K_1}(\Theta_{\mathcal{L}}), \widetilde{K_2}(\Theta_{\mathcal{L}}))$ and $(\widetilde{K_1}(\Theta'_{\mathcal{L}}), \widetilde{K_2}(\Theta'_{\mathcal{L}}))$ are symmetric, ξ maps to $\{\pm 1\}$ by Lemma 5.3.14. Then, $\xi^2 = 1$ so that $\eta_2 \circ \Theta'_{\mathcal{L}^2} = \eta_2 \circ \Theta_{\mathcal{L}^2}$ *i.e.* $\Theta'_{\mathcal{L}} \circ H_2 = \Theta_{\mathcal{L}} \circ H_2$ *i.e.* $\Theta_{\mathcal{L}} = \Theta'_{\mathcal{L}}$, since H_2 is surjective.

(iii) Let $\Theta_{\mathcal{L}}$ be a symmetric theta-structure on $G(\mathcal{L})$. Let $\overline{\Theta}_{\mathcal{L}} : K(\delta) \xrightarrow{\sim} K(\mathcal{L})$ be the induced symplectic isomorphism.

Step 1: First, we lift $\overline{\Theta}_{\mathcal{L}}$ to construct a symplectic isomorphism $\overline{\Theta}_{\mathcal{L}^2} : K(2\delta) \xrightarrow{\sim} K(\mathcal{L}^2)$ such that $\overline{\Theta}_{\mathcal{L}} \circ \overline{H}_2 = [2] \circ \overline{\Theta}_{\mathcal{L}^2}$ and $\overline{\Theta}_{\mathcal{L}} = \overline{\Theta}_{\mathcal{L}^2} \circ \overline{E}_2$, where

$$\overline{E}_2 : (x, \chi) \in K(\delta) \longmapsto (2x, 2 \star \chi) \in K(2\delta) \quad \text{and} \quad \overline{H}_2 : (x, \chi) \in K(2\delta) \longmapsto (\overline{x}, \overline{\chi}) \in K(\delta). \tag{5.13}$$

Let $(x_1, \cdots, x_g, \chi_1, \cdots, \chi_g)$ be a canonical basis of $K(2\delta) = K_1(2\delta) \oplus K_2(2\delta)$, given as follows: for all $i \in [\![1\ ;\ g]\!]$, let $x_i \in K_1(2\delta) = \prod_{i=1}^g \mathbb{Z}/2d_i\mathbb{Z}$ be equal to 1 at index i and 0 everywhere else and $\chi_i \in K_2(2\delta)$ be a character such that $\chi_i(x_i)$ is a $2d_i$-th primitive root of unity and $\chi_i(x_j) = 1$ for all $1 \leq j \neq i \leq g$. Let $(y_1, \cdots, y_g, z_1, \cdots, z_g)$ be a symplectic basis of $K(\mathcal{L}^2)$ such that $[2]y_i = \overline{\Theta}_{\mathcal{L}}(\overline{x}_i, 1)$

and $[2]z_i = \overline{\Theta}_{\mathcal{L}}(0, \overline{\chi}_i)$ for all $i \in [\![1\ ;\ g]\!]$. By Proposition 5.1.6.(v) and since $\overline{\Theta}_{\mathcal{L}}$ is symplectic, we then have for all $i \in [\![1\ ;\ g]\!]$:

$$e_{\mathcal{L}^2}(y_i, z_i)^2 = e_{\mathcal{L}^2}([2]y_i, z_i) = e_{\mathcal{L}}([2]y_i, [2]z_i) = e_\delta((\overline{x}_i, 1), (0, \overline{\chi}_i)) = \chi_i(x_i)^2,$$

so that $e_{\mathcal{L}^2}(y_i, z_i) = \pm\chi_i(x_i)$. But we may change the sign of $\chi_i(x_i)$ without affecting $\overline{\chi}_i$, so that $e_{\mathcal{L}^2}(y_i, z_i) = \chi_i$ for all $i \in [\![1\ ;\ g]\!]$. Then, we set $\overline{\Theta}_{\mathcal{L}^2}(x_i, 1) := y_i$ and $\overline{\Theta}_{\mathcal{L}^2}(0, \chi_i) := z_i$ for all $i \in [\![1\ ;\ g]\!]$. This defines a symplectic isomorphism $K(2\delta) \xrightarrow{\sim} K(\mathcal{L}^2)$ such that $\overline{\Theta}_{\mathcal{L}} \circ \overline{H}_2 = [2] \circ \overline{\Theta}_{\mathcal{L}^2}$.

But we still have some liberty on the choice of $\overline{\Theta}_{\mathcal{L}^2}$. Note that we could change y_i into $y_i + d_i z_i$ or z_i into $z_i + d_i y_i$ without changing the values of $e_{\mathcal{L}^2}$ on the new symplectic basis basis and without affecting the formula $\overline{\Theta}_{\mathcal{L}} \circ \overline{H}_2 = [2] \circ \overline{\Theta}_{\mathcal{L}^2}$.

Let $(x, \chi) \in K(\delta)$. Then, there exists $(x', \chi') \in K(2\delta)$ such that $\overline{H}_2(x', \chi') = (x, \chi)$. We then have, as desired:

$$\begin{aligned}\overline{\Theta}_{\mathcal{L}}(x, \chi) &= \overline{\Theta}_{\mathcal{L}} \circ \overline{H}_2(x', \chi') = [2] \circ \overline{\Theta}_{\mathcal{L}^2}(x', \chi') = \overline{\Theta}_{\mathcal{L}^2}(2x', \chi'^2) = \overline{\Theta}_{\mathcal{L}^2}(2x, 2 \star \chi) \\ &= \overline{\Theta}_{\mathcal{L}^2} \circ \overline{E}_2(x, \chi)\end{aligned}$$

Step 2: Now, we lift the level structure $s_{\mathcal{L}} : K(\mathcal{L}) \longrightarrow G(\mathcal{L})$ induced by $\Theta_{\mathcal{L}}$, namely, we define a level structure $s_{\mathcal{L}^2} : K(\mathcal{L}^2) \longrightarrow G(\mathcal{L}^2)$ such that $s_{\mathcal{L}} \circ [2] = \eta_2 \circ s_{\mathcal{L}^2}$ and $s_{\mathcal{L}^2|K(\mathcal{L})} = \varepsilon_2 \circ s_{\mathcal{L}}$. Let $i \in [\![1\ ;\ g]\!]$. Then, $[2]y_i = \overline{\Theta}_{\mathcal{L}}(\overline{x}_i, 1) \in K_1(\mathcal{L})$ by construction. Since $\Theta_{\mathcal{L}}$ is symmetric, we know by Proposition 5.3.13 that $s_{\mathcal{L}}([2]y_i)$ is symmetric. Now, η_2 being surjective, there exists $\widetilde{y}_i \in G(\mathcal{L}^2)$ such that $\eta_2(\widetilde{y}_i) = s_{\mathcal{L}}([2]y_i)$. Since $\widetilde{y}_i^{-1}$ and $\Delta_{-1}^{\mathcal{L}^2}(\widetilde{y}_i)$ both lift $-y_i$, we may write $\Delta_{-1}^{\mathcal{L}^2}(\widetilde{y}_i) = \lambda\widetilde{y}_i^{-1}$ with $\lambda \in k^*$. We then have by Proposition 5.3.21.(i):

$$\Delta_{-1}^{\mathcal{L}} \circ s_{\mathcal{L}}([2]y_i) = \Delta_{-1}^{\mathcal{L}} \circ \eta_2(\widetilde{y}_i) = \eta_2(\Delta_{-1}^{\mathcal{L}}(\widetilde{y}_i)) = \eta_2(\lambda\widetilde{y}_i^{-1}) = \lambda^2\eta_2(\widetilde{y}_i)^{-1} = \lambda^2 s_{\mathcal{L}}([2]y_i)^{-1},$$

so $\lambda = \pm 1$. If $\lambda = 1$, then $\widetilde{y}_i$ is symmetric and we set $s_{\mathcal{L}^2}(y_i) := \widetilde{y}_i$. If $\lambda = -1$, we may change y_i into $y_i + d_i z_i$ as suggested above. Indeed, if $g_i \in G(\mathcal{L}^2)$ is a lift of $d_i z_i$ of order 2, it is automatically symmetric by Lemma 5.3.15 and since $\mathcal{L}$ is totally symmetric. We then have

$$\Delta_{-1}^{\mathcal{L}^2}(\widetilde{y}_i g_i) = \Delta_{-1}^{\mathcal{L}^2}(\widetilde{y}_i)\Delta_{-1}^{\mathcal{L}^2}(g_i) = -\widetilde{y}_i^{-1}g_i^{-1} = -e_{\mathcal{L}^2}(y_i, d_i z_i)g_i^{-1}\widetilde{y}_i^{-1} = -e_{\mathcal{L}^2}(y_i, z_i)^{d_i}(\widetilde{y}_i g_i)^{-1},$$

with $e_{\mathcal{L}^2}(y_i, z_i)^{d_i} = -1$ since $e_{\mathcal{L}^2}(y_i, z_i) = \chi_i(x_i)$ is a primitive $2d_i$-th root of unity, so $\widetilde{y}_i g_i$ is symmetric and we may set $s_{\mathcal{L}^2}(y_i + d_i z_i) := \widetilde{y}_i g_i$. We proceed similarly to define $s_{\mathcal{L}^2}$ on $K_2(\mathcal{L}^2)$ (changing z_i into $z_i + d_i y_i$ if necessary).

By construction, $s_{\mathcal{L}} \circ [2] = \eta_2 \circ s_{\mathcal{L}^2}$. Now, we prove that $s_{\mathcal{L}^2|K(\mathcal{L})} = \varepsilon_2 \circ s_{\mathcal{L}}$. Indeed, by Proposition 5.3.21.(iv) and since the level structure $s_{\mathcal{L}^2}$ is symmetric, we have for all $x \in K(\mathcal{L}^2)$:

$$\begin{aligned}\varepsilon_2 \circ s_{\mathcal{L}}([2]x) &= \varepsilon_2 \circ \eta_2 \circ s_{\mathcal{L}^2}(x) = \Delta_2 \circ s_{\mathcal{L}^2}(x) = s_{\mathcal{L}^2}(x)^3 \cdot \Delta_{-1} \circ s_{\mathcal{L}^2}(x) = s_{\mathcal{L}^2}(x)^3 s_{\mathcal{L}^2}(-x) \\ &= s_{\mathcal{L}^2}([2]x).\end{aligned}$$

Since $K(\mathcal{L}) = [2]K(\mathcal{L}^2)$, we conclude that $s_{\mathcal{L}^2|K(\mathcal{L})} = \varepsilon_2 \circ s_{\mathcal{L}}$.

Step 3: We finally verify that the symmetric theta-structure $\Theta_{\mathcal{L}^2}$ is compatible with $\Theta_{\mathcal{L}}$. Let $(\alpha, x, \chi) \in \mathcal{H}(2\delta)$. Then:

$$\begin{aligned}\Theta_{\mathcal{L}} \circ H_2(\alpha, x, \chi) &= \Theta_{\mathcal{L}} \circ (\alpha^2, \overline{H}_2(x, \chi)) = \alpha^2 \cdot s_{\mathcal{L}} \circ \overline{\Theta}_{\mathcal{L}} \circ \overline{H}_2(x, \chi) = \alpha^2 \cdot s_{\mathcal{L}}([2]\overline{\Theta}_{\mathcal{L}^2}(x, \chi)) \\ &= \alpha^2 \cdot \eta_2 \circ s_{\mathcal{L}^2} \circ \overline{\Theta}_{\mathcal{L}^2}(x, \chi) = \eta_2(\alpha \cdot s_{\mathcal{L}^2} \circ \overline{\Theta}_{\mathcal{L}^2}(x, \chi)) = \eta_2 \circ \Theta_{\mathcal{L}^2}(\alpha, x, \chi),\end{aligned}$$

so $\Theta_{\mathcal{L}} \circ H_2 = \eta_2 \circ \Theta_{\mathcal{L}^2}$. Now, let $(\alpha, x, \chi) \in \mathcal{H}(\delta)$. Then:

$$\begin{aligned}\Theta_{\mathcal{L}^2} \circ E_2(\alpha, x, \chi) &= \Theta_{\mathcal{L}^2}(\alpha^2, \overline{E}_2(x, \chi)) = \alpha^2 \cdot s_{\mathcal{L}^2} \circ \overline{\Theta}_{\mathcal{L}^2} \circ \overline{E}_2(x, \chi) = \alpha^2 \cdot s_{\mathcal{L}^2} \circ \overline{\Theta}_{\mathcal{L}}(x, \chi) \\ &= \alpha^2 \cdot \varepsilon_2 \circ s_{\mathcal{L}} \circ \overline{\Theta}_{\mathcal{L}}(x, \chi) = \varepsilon_2(\alpha \cdot s_{\mathcal{L}} \circ \overline{\Theta}_{\mathcal{L}}(x, \chi)) = \varepsilon_2 \circ \Theta_{\mathcal{L}}(\alpha, x, \chi),\end{aligned}$$

so $\Theta_{\mathcal{L}^2} \circ E_2 = \varepsilon_2 \circ \Theta_{\mathcal{L}}$ and the proof is now complete. □

Remark 5.3.31. We have proved in Theorem 5.3.30.(iii) that any symmetric theta structure $\Theta_{\mathcal{L}}$ on $(A, \mathcal{L})$ is induced by a symplectic isomorphism $K(2\delta) \xrightarrow{\sim} K(\mathcal{L}^2)$. Equivalently, if $\delta := (d_1, \cdots, d_g)$, and if $\zeta_i \in k^*$ is a primitive $2d_i$-th root of unity for all $i \in [\![1\ ;\ g]\!]$, then $\Theta_{\mathcal{L}}$ is also induced by a

$(\zeta_1, \cdots, \zeta_g)$-symplectic basis $(x_1, \cdots, x_g, y_1, \cdots, y_g)$ of $K(\mathcal{L}^2)$ in the sense of Lemma 5.1.9. Indeed, if we denote by $e_i \in K_1(\delta)$ the element with 1 at index i and 0 everywhere else for all $i \in [\![1\ ;\ g]\!]$, and by $\chi_i \in K_2(\delta)$ the character given by $\chi_i(e_j) = \zeta_i^{\delta_{i,j}}$ for all $i, j \in [\![1\ ;\ g]\!]$, then we can set $\overline{\Theta}_{\mathcal{L}^2}(e_i, 1) := x_i$ and $\overline{\Theta}_{\mathcal{L}^2}(0, \chi_i) := y_i$ for all $i \in [\![1\ ;\ g]\!]$. This defines a symplectic isomorphism $K(2\delta) \xrightarrow{\sim} K(\mathcal{L}^2)$, hence a symmetric theta structure $\Theta_{\mathcal{L}}$ on $(A, \mathcal{L})$.

Symmetric theta structures and isogenies

In this paragraph, we relate symmetric theta structures to isogenies. In particular, we prove that an isogeny maps a symmetric theta structure to a symmetric theta structure compatible with the first one. We also obtain a partial converse of this result.

Lemma 5.3.32. *Let $f : (A, \mathcal{L}) \longrightarrow (B, \mathcal{M})$ be a polarised isogeny such that $f^*\mathcal{M} \simeq \mathcal{L}$ and $\mathcal{M}$ be a symmetric line bundle. Let $K := \ker(f)$ and $\widetilde{K}$ be the level subgroup induced by an isomorphism $\alpha : f^*\mathcal{M} \xrightarrow{\sim} \mathcal{L}$. Then $\mathcal{L}$ is symmetric and the canonical morphism $\alpha_f : Z(\widetilde{K}) \longrightarrow G(\mathcal{M})$ defined by 5.1 maps symmetric elements to symmetric elements.*

Proof. Since $\mathcal{M}$ is symmetric, we have $[-1]^*\mathcal{M} \simeq \mathcal{M}$ so $[-1]^*\mathcal{L} \simeq [-1]^*f^*\mathcal{M} = f^*[-1]^*\mathcal{M} \simeq f*\mathcal{M} \simeq \mathcal{L}$ and $\mathcal{L}$ is symmetric. More precisely, if $\Phi_{\mathcal{M}} : \mathcal{M} \xrightarrow{\sim} [-1]\mathcal{M}^*$ is an isomorphism, then we have an isomorphism $\Phi_{\mathcal{L}} : \mathcal{L} \xrightarrow{\sim} [-1]^*\mathcal{L}$ given by the composition

$$\mathcal{L} \xrightarrow{\alpha^{-1}} f^*\mathcal{M} \xrightarrow{f^*\Phi_{\mathcal{M}}} f^*[-1]^*\mathcal{M} = [-1]^*f^*\mathcal{M} \xrightarrow{[-1]^*\alpha} [-1]^*\mathcal{L}\ .$$

To conclude, we only have to prove that α_f "commutes" with Δ_{-1} *i.e.* that $\alpha_f \circ \Delta^{\mathcal{L}}_{-1} = \Delta^{\mathcal{M}}_{-1} \circ \alpha_f$, where $\Delta^{\mathcal{L}}_{-1}$ and $\Delta^{\mathcal{M}}_{-1}$ are the Δ_{-1} maps defined on $G(\mathcal{L})$ and $G(\mathcal{M})$ respectively. Let $g \in Z(\widetilde{K})$. Then, by Proposition 5.1.5, we may write $g := (x, t_x^*\alpha \circ f^*\psi \circ \alpha^{-1})$, with $(f(x), \psi) = \alpha_f(g) \in G(\mathcal{M})$. We then have:

$$\begin{aligned}
\alpha_f \circ \Delta^{\mathcal{L}}_{-1}(g) &= \alpha_f(-x, (t^*_{-x}\Phi_{\mathcal{L}})^{-1} \circ [-1]^*(t_x^*\alpha \circ f^*\psi \circ \alpha^{-1}) \circ \Phi_{\mathcal{L}}) \\
&= \alpha_f\big(-x, ([-1]^*t_x^*\alpha \circ f^*t^*_{-x}\Phi_{\mathcal{M}} \circ t^*_{-x}\alpha^{-1})^{-1} \circ ([-1]^*t_x^*\alpha \circ f^*[-1]^*\psi \circ [-1]^*\alpha^{-1}) \\
&\quad \circ ([-1]^*\alpha \circ f^*\Phi_{\mathcal{M}} \circ \alpha^{-1})\big) \\
&= \alpha_f\big(-x, t^*_{-x}\alpha \circ f^*((t^*_{-x}\Phi_{\mathcal{M}})^{-1} \circ [-1]^*\psi \circ \Phi_{\mathcal{M}}) \circ \alpha^{-1}\big) \\
&= (-f(x), (t^*_{-x}\Phi_{\mathcal{M}})^{-1} \circ [-1]^*\psi \circ \Phi_{\mathcal{M}}) = \Delta^{\mathcal{M}}_{-1}(f(x), \psi) = \Delta^{\mathcal{M}}_{-1} \circ \alpha_f(g),
\end{aligned}$$

This completes the proof. □

Remark 5.3.33. By the previous lemma, if $f : (A, \mathcal{L}) \longrightarrow (B, \mathcal{M})$ is a polarised isogeny, with $\mathcal{L}$ and $\mathcal{M}$ two symmetric line bundles satisfying $f^*\mathcal{M} \simeq \mathcal{L}$, $\Theta_{\mathcal{L}}$ and $\Theta_{\mathcal{M}}$ are compatible theta structures on $G(\mathcal{L})$ and $G(\mathcal{M})$ respectively (in the sense of Definition 5.2.1) and if $\Theta_{\mathcal{L}}$ is symmetric, then $\Theta_{\mathcal{M}}$ is automatically symmetric. Indeed, α_f maps $\widetilde{K_i}(\Theta_{\mathcal{L}}) \cap Z(\widetilde{K})$ (which is symmetric) to $\widetilde{K_i}(\Theta_{\mathcal{M}})$, which is then symmetric for all $i \in \{1, 2\}$.

It is natural to wonder when compatible theta structures with respect to an isogeny are symmetric and the above remark answers this question. Conversely, we can ask when symmetric theta structures are compatible with respect to isogeny. First, Proposition 5.3.34 below ensures that symmetric theta structures on the domain $(A, \mathcal{L})$ are automatically compatible with level subgroups $\widetilde{K}$ over the isogeny kernel (point (i) of Definition 5.2.1) when the codomain line bundle $\mathcal{M}$ is totally symmetric. Hence, to ensure that two symmetric theta structures are compatible with respect to an isogeny, we only have to verify that this isogeny maps the level structure on the domain to the one on the codomain (point (ii) of Definition 5.2.1). This condition is usually fastidious to verify but using Theorem 5.3.30 we obtain an equivalent condition which is much easier to verify in Proposition 5.3.36.

Proposition 5.3.34. *Let $f : (A, \mathcal{L}) \longrightarrow (B, \mathcal{M})$ be a polarised isogeny of kernel K, with $\mathcal{L}$ and $\mathcal{M}$ line bundles satisfying $f^*\mathcal{M} \simeq \mathcal{L}$ and $\mathcal{M}$ totally symmetric. Let $\widetilde{K}$ be the level subgroup above K associated to an isomorphism $\alpha : f^*\mathcal{M} \xrightarrow{\sim} \mathcal{L}$. Assume we have a symplectic decomposition $K(\mathcal{L}) = K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$ compatible with K: $K = (K \cap K_1(\mathcal{L})) \oplus (K \cap K_2(\mathcal{L}))$. Then any symmetric level structure on $G(\mathcal{L})$ respecting this decomposition is compatible with $\widetilde{K}$.*

Proof. Let $(\widetilde{K_1}(\mathcal{L}), \widetilde{K_2}(\mathcal{L}))$ be a symmetric level structure above $K(\mathcal{L}) = K_1(\mathcal{L}) \oplus K_2(\mathcal{L})$. By Lemma 5.2.3, it suffices to prove that $\alpha_f(\widetilde{K_i}(\mathcal{L}) \cap \rho_{\mathcal{L}}^{-1}(K)) = \{1\}$ for $i \in \{1,2\}$ to prove that this level structure is compatible with $\widetilde{K}$.

Let $i \in \{1,2\}$ and $\widetilde{x} \in \widetilde{K_i}(\mathcal{L}) \cap \rho_{\mathcal{L}}^{-1}(K)$ and $x := \rho_{\mathcal{L}}(\widetilde{x})$. Then $x \in K \cap K_i(\mathcal{L})$. Besides, $\mathcal{M}$ being totally symmetric, we have $B[2] \subseteq K(\mathcal{M})$ by Proposition 5.3.8 and Lemma 5.1.7. It follows by Proposition 5.1.5.(i) that $[2]^{-1}K = f^{-1}(B[2]) \subseteq f^{-1}(K(\mathcal{M})) \subseteq K(\mathcal{L})$, so there exists $y \in K(\mathcal{L})$ such that $2y = x$. Since $x \in K_i(\mathcal{L})$ and $K_1(\mathcal{L}) \cap K_2(\mathcal{L}) = \{0\}$, we can assume that $y \in K_i(\mathcal{L})$. Let $\widetilde{y} \in \widetilde{K_i}(\mathcal{L})$ be a lift of y. Then $\widetilde{y}^2$ and $\widetilde{x}$ both are lifts of x in $\widetilde{K_i}(\mathcal{L})$, so $\widetilde{y}^2 = \widetilde{x}$. In addition, $\alpha_f(\widetilde{y})$ is symmetric by Lemma 5.3.32 and is a lift of $f(y)$ in $G(\mathcal{M})$, which has order at most 2 since $2y = x \in K$. Since $\mathcal{M}$ is totally symmetric, Lemma 5.3.15 ensures that $\alpha_f(\widetilde{y}) = \Delta_{-1}(\alpha_f(\widetilde{y})) = \alpha_f(\widetilde{y})^{-1}$, so has order at most 2 *i.e.* that $\alpha_f(\widetilde{x}) = \alpha_f(\widetilde{y})^2 = 1$. This completes the proof. □

Lemma 5.3.35. *Let $f : (A, \mathcal{L}) \longrightarrow (B, \mathcal{M})$ be a polarised isogeny, where $\mathcal{L}$ and $\mathcal{M}$ are symmetric line bundles satisfying $f^*\mathcal{M} \simeq \mathcal{L}$. Let $\alpha : f^*\mathcal{M} \xrightarrow{\sim} \mathcal{L}$ be an isomorphism and $\alpha^{\otimes 2}$ be the induced isomorphism $f^*\mathcal{M}^2 \xrightarrow{\sim} \mathcal{L}^2$ and α_f and $\alpha_f^{\otimes 2}$ be the associated maps given by Eq. (5.1). Let $\eta_2^{\mathcal{L}} : G(\mathcal{L}^2) \longrightarrow G(\mathcal{L})$ and $\eta_2^{\mathcal{M}} : G(\mathcal{M}^2) \longrightarrow G(\mathcal{M})$ be the maps defined in Definition 5.3.17. Then, we have*

$$\eta_2^{\mathcal{M}} \circ \alpha_f^{\otimes 2} = \alpha_f \circ \eta_2^{\mathcal{L}}.$$

Proof. Let $\Psi_{\mathcal{M}} : [2]^*\mathcal{M} \xrightarrow{\sim} \mathcal{M}^4$ be an isomorphism (that does exist since $\mathcal{M}$ is symmetric). This isomorphism induces another isomorphism $\Psi_{\mathcal{L}} : [2]^*\mathcal{L} \xrightarrow{\sim} \mathcal{L}^4$ defined as the composition:

$$[2]^*\mathcal{L} \xrightarrow{[2]^*\alpha^{-1}} [2]^*f\mathcal{M} = f^*[2]^*\mathcal{M} \xrightarrow{f^*\Psi_{\mathcal{M}}} f^*\mathcal{M} \xrightarrow{\alpha^{\otimes 4}} \mathcal{L}^4$$

Let $\widetilde{K}^{\mathcal{L}} \subset G(\mathcal{L})$ and $\widetilde{K}^{\mathcal{L}^2} \subset G(\mathcal{L}^2)$ be level subgroups lying over $K := \ker(f)$. Let $g := (x, \phi_x) \in Z(\widetilde{K}^{\mathcal{L}^2})$. Then $\eta_2^{\mathcal{L}}(g) = ([2]x, \rho_x)$ with $[2]^*\rho_x^{\mathcal{L}} = t_x^*\Psi_{\mathcal{L}}^{-1} \circ \phi_x^{\otimes 2} \circ \Psi_{\mathcal{L}}$.

First, we prove that $\eta_2^{\mathcal{L}}(g) \in Z(\widetilde{K}^{\mathcal{L}})$. Recall that $Z(\widetilde{K}^{\mathcal{L}}) = \rho_{\mathcal{L}}^{-1}(K^{\perp,\mathcal{L}})$ and $Z(\widetilde{K}^{\mathcal{L}^2}) = \rho_{\mathcal{L}^2}^{-1}(K^{\perp,\mathcal{L}^2})$ where $K^{\perp,\mathcal{L}}$ and $K^{\perp,\mathcal{L}^2}$ are orthogonal of K for the commutator pairings $e_{\mathcal{L}}$ and $e_{\mathcal{L}^2}$ respectively. Hence, we have $x \in K^{\perp,\mathcal{L}^2}$ and we only have to prove that $[2]x \in K^{\perp,\mathcal{L}}$. But for all $y \in K$, we have by Proposition 5.1.6.(v),

$$e_{\mathcal{L}}([2]x, y) = e_{\mathcal{L}^2}(x, y) = 1,$$

so that $[2]x \in K^{\perp,\mathcal{L}}$ and $\eta_2^{\mathcal{L}}(g) \in Z(\widetilde{K}^{\mathcal{L}})$.

Hence, we can compose $\alpha_f \circ \eta_2^{\mathcal{L}}(g)$ and we have $\alpha_f \circ \eta_2^{\mathcal{L}}(g) = (f([2]x), \psi_{[2]x})$ with $\rho_x = t_{[2]x}^*\alpha \circ f^*\psi_{[2]x} \circ \alpha^{-1}$. We also have

$$\eta_2^{\mathcal{M}} \circ \alpha_f^{\otimes 2}(g) = \eta_2^{\mathcal{M}}(f(x), \psi_x) = ([2]f(x), \rho_{f(x)}),$$

with $\phi_x = t_x^*\alpha^{\otimes 2} \circ f^*\psi_x \circ (\alpha^{-1})^{\otimes 2}$ and $[2]^*\rho_{f(x)} = t_{f(x)}^*\Psi_{\mathcal{M}}^{-1} \circ \psi_x^{\otimes 2} \circ \Psi_{\mathcal{M}}$. Moreover, on the one hand

$$\begin{aligned}
[2]^*\rho_x &= t_x^*\Psi_{\mathcal{L}}^{-1} \circ \phi_x^{\otimes 2} \circ \Psi_{\mathcal{L}} \\
&= t_x^*([2]^*\alpha \circ f^*\Psi_{\mathcal{M}}^{-1} \circ (\alpha^{-1})^{\otimes 4}) \circ (t_x^*\alpha^{\otimes 4} \circ f^*\psi_x^{\otimes 2} \circ (\alpha^{-1})^{\otimes 4}) \circ (\alpha^{\otimes 4} \circ f^*\Psi_{\mathcal{M}} \circ [2]^*\alpha^{-1}) \\
&= t_x^*[2]^*\alpha \circ t_x^*f^*\Psi_{\mathcal{M}}^{-1} \circ f^*\psi_x^{\otimes 2} \circ f^*\Psi_{\mathcal{M}} \circ [2]^*\alpha^{-1} \\
&= [2]^*t_{[2]x}^*\alpha \circ (f^*t_{f(x)}^*\Psi_{\mathcal{M}}^{-1} \circ f^*\psi_x^{\otimes 2} \circ f^*\Psi_{\mathcal{M}}) \circ [2]^*\alpha^{-1} \\
&= [2]^*t_{[2]x}^*\alpha \circ f^*[2]^*\rho_{f(x)} \circ [2]^*\alpha^{-1}.
\end{aligned}$$

And on the other hand $[2]^*\rho_x = [2]^*t_{[2]x}^*\alpha \circ [2]^*f^*\psi_{[2]x} \circ [2]^*\alpha^{-1}$. It follows that $\rho_{f(x)} = \psi_{[2]x}$, so that $\alpha_f \circ \eta_2^{\mathcal{L}}(g) = \eta_2^{\mathcal{M}} \circ \alpha_f^{\otimes 2}(g)$. This completes the proof. □

Proposition 5.3.36. *Let $f : (A, \mathcal{L}) \longrightarrow (B, \mathcal{M})$ be a polarised isogeny of kernel K, with $\mathcal{L}$ and $\mathcal{M}$ line bundles satisfying $f^*\mathcal{M} \simeq \mathcal{L}$ and $\mathcal{M}$ totally symmetric. Let $\Theta_{\mathcal{L}}$ and $\Theta_{\mathcal{M}}$ be symmetric theta structures on $G(\mathcal{L})$ and $G(\mathcal{M})$ induced by symplectic isomorphisms $\overline{\Theta}_{\mathcal{L}^2} : K(2\delta_{\mathcal{L}}) \xrightarrow{\sim} K(\mathcal{L}^2)$ and $\overline{\Theta}_{\mathcal{M}^2} : K(2\delta_{\mathcal{M}}) \xrightarrow{\sim} K(\mathcal{M}^2)$ respectively (by Theorem 5.3.30.(iii)). Assume that:*

(i) $K = (K \cap K_1(\overline{\Theta}_{\mathcal{L}^2})) \oplus (K \cap K_2(\overline{\Theta}_{\mathcal{L}^2}))$;

(ii) and $f(K_i(\overline{\Theta}_{\mathcal{L}^2}) \cap f^{-1}(K(\mathcal{M}^2))) = K_i(\overline{\Theta}_{\mathcal{M}^2})$ *for* $i \in \{1,2\}$.

Then $\Theta_{\mathcal{L}}$ *and* $\Theta_{\mathcal{M}}$ *are compatible with respect to* f.

Proof. Let $\Theta_{\mathcal{L}^2}$ be a symmetric theta structure on $G(\mathcal{L}^2)$ lifting the symplectic isomorphism $\overline{\Theta}_{\mathcal{L}^2}$. Then $\Theta_{\mathcal{L}^2}$ is compatible with a level subgroup $\widetilde{K}^{\mathcal{L}^2} \subset G(\mathcal{L}^2)$ lying over $K := \ker(f)$ by Proposition 5.3.34 since $K = (K \cap K_1(\overline{\Theta}_{\mathcal{L}^2})) \oplus (K \cap K_2(\overline{\Theta}_{\mathcal{L}^2}))$. A theta structure $\Theta_{\mathcal{M}^2}$ on $G(\mathcal{M}^2)$ is compatible with $\Theta_{\mathcal{L}^2}$ if and only if $\alpha_f^{\otimes 2}(\widetilde{K_i}(\Theta_{\mathcal{L}^2}) \cap Z(\widetilde{K}^{\mathcal{L}^2})) = \widetilde{K_i}(\Theta_{\mathcal{M}^2})$ for $i \in \{1,2\}$, where $\alpha_f^{\otimes 2}$ is associated to an isomorphism $\alpha^{\otimes 2} : f^*\mathcal{M}^2 \xrightarrow{\sim} \mathcal{L}^2$. Since $f(K_i(\overline{\Theta}_{\mathcal{L}^2}) \cap f^{-1}(K(\mathcal{M}^2))) = K_1(\overline{\Theta}_{\mathcal{M}^2})$, the condition $\alpha_f^{\otimes 2}(\widetilde{K_i}(\Theta_{\mathcal{L}^2}) \cap Z(\widetilde{K}^{\mathcal{L}^2})) = \widetilde{K_i}(\Theta_{\mathcal{M}^2})$ can be satisfied for $i \in \{1,2\}$ and defines a theta structure $\Theta_{\mathcal{M}^2}$ on $G(\mathcal{M}^2)$ lifting the symplectic isomorphism $\overline{\Theta}_{\mathcal{M}^2} : K(2\delta_{\mathcal{M}}) \xrightarrow{\sim} K(\mathcal{M}^2)$. Then $\Theta_{\mathcal{M}^2}$ is symmetric by Lemma 5.3.32 and $\Theta_{\mathcal{M}}$ is compatible with $\Theta_{\mathcal{M}^2}$ in the sense of Definition 5.3.28 by Theorem 5.3.30.(iii).

We now prove that $\Theta_{\mathcal{L}}$ is compatible with $\Theta_{\mathcal{M}}$. By Proposition 5.3.34, $\Theta_{\mathcal{L}}$ is compatible with a level subgroup $\widetilde{K}^{\mathcal{L}} \subset G(\mathcal{L})$ lying over $K := \ker(f)$, since it is a symmetric theta structure. To prove the compatibility of $\Theta_{\mathcal{L}}$ and $\Theta_{\mathcal{M}}$, it suffices to prove that $\alpha_f(\widetilde{K_i}(\Theta_{\mathcal{L}}) \cap Z(\widetilde{K}^{\mathcal{L}})) = \widetilde{K_i}(\Theta_{\mathcal{M}})$ for all $i \in \{1,2\}$, where α_f is associated to the isomorphism $\alpha : f^*\mathcal{M} \xrightarrow{\sim} \mathcal{L}$ inducing $\alpha^{\otimes 2}$. We actually prove the equivalent property $s_{\mathcal{M}} \circ f = \alpha_f \circ s_{\mathcal{L}}$ on $f^{-1}(K(\mathcal{M}))$, where $s_{\mathcal{L}}$ and $s_{\mathcal{M}}$ are the level structures associated to $\Theta_{\mathcal{L}}$ and $\Theta_{\mathcal{M}}$ respectively. The compatibility of $\Theta_{\mathcal{L}^2}$ and $\Theta_{\mathcal{M}^2}$ already implies that $s_{\mathcal{M}^2} \circ f = \alpha_f^{\otimes 2} \circ s_{\mathcal{L}^2}$ on $f^{-1}(K(\mathcal{M}^2))$, where $s_{\mathcal{L}^2}$ and $s_{\mathcal{M}^2}$ are the level structures associated to $\Theta_{\mathcal{L}^2}$ and $\Theta_{\mathcal{M}^2}$ respectively. In addition, $(\Theta_{\mathcal{L}}, \Theta_{\mathcal{L}^2})$ and $(\Theta_{\mathcal{M}}, \Theta_{\mathcal{M}^2})$ being pairs of compatible symmetric theta structures in the sense of Definition 5.3.28, we have $s_{\mathcal{L}} \circ [2] = \eta_2^{\mathcal{L}} \circ s_{\mathcal{L}^2}$ and $s_{\mathcal{M}} \circ [2] = \eta_2^{\mathcal{M}} \circ s_{\mathcal{M}^2}$. Let $x \in f^{-1}(K(\mathcal{M}))$. Then, we may write $x = [2]y$ with $y \in f^{-1}(K(\mathcal{M}^2))$ since $[2]K(\mathcal{M}^2) = K(\mathcal{M})$ and we have by Lemma 5.3.35 that

$$\begin{aligned} s_{\mathcal{M}} \circ f(x) &= s_{\mathcal{M}} \circ [2] \circ f(y) = \eta_2^{\mathcal{M}} \circ s_{\mathcal{M}^2} \circ f(y) = \eta_2^{\mathcal{M}} \circ \alpha_f^{\otimes 2} \circ s_{\mathcal{L}^2}(y) \\ &= \alpha_f \circ \eta_2^{\mathcal{L}} \circ s_{\mathcal{L}^2}(y) = \alpha_f \circ s_{\mathcal{L}}([2]y) = \alpha_f \circ s_{\mathcal{L}}(x). \end{aligned}$$

This completes the proof. □

5.3.2 The duplication formula

In this section, we prove the duplication formula that relate theta coordinates associated to compatible symmetric theta structures $\Theta_{\mathcal{L}}$ and $\Theta_{\mathcal{L}^2}$ on $G(\mathcal{L})$ and $G(\mathcal{L}^2)$ respectively, when $\mathcal{L}$ is a totally symmetric line bundle over A. From this formula, we derive a differential addition algorithm to compute $x+y$ given closed points $x, y, x-y \in A(k)$ and also an algorithm to double points.

Product theta structures

The duplication formula involves product theta structures and product theta coordinates that we need to introduce first. Let $(A_1, \mathcal{L}_1), \cdots, (A_r, \mathcal{L}_r)$ be polarised abelian varieties, $A := \prod_{i=1}^r A_i$ and $\mathcal{L} := \bigotimes_{i=1}^r \pi_i^*\mathcal{L}_i$, where $\pi_i : A \longrightarrow A_i$ is the projection for all $i \in [\![1\ ;\ r]\!]$. Then $(A, \mathcal{L})$ is an abelian variety equipped with the *product polarization* $\varphi_{\mathcal{L}} := \mathrm{Diag}(\varphi_{\mathcal{L}_1}, \cdots, \varphi_{\mathcal{L}_r}) : A \longrightarrow \widehat{A}$ (as defined in Lemma 2.2.3). We have natural isomorphisms $K(\mathcal{L}) \simeq \bigoplus_{i=1}^r K(\mathcal{L}_i)$ and

$$G(\mathcal{L}) \simeq \prod_{i=1}^r G(\mathcal{L}_i)/\{(\lambda_1, \cdots, \lambda_r) \in k^* \mid \lambda_1 \cdots \lambda_r = 1\}.$$

Let $\delta^{(i)}$ be the type of $\mathcal{L}_i$ for all $i \in [\![1\ ;\ r]\!]$ and $\delta := \delta^{(1)} \vee \cdots \vee \delta^{(r)}$ be the concatenation of the $\delta^{(i)}$. Then, we also have:

$$\mathcal{H}(\delta) \simeq \prod_{i=1}^r \mathcal{H}(\delta^{(i)})/\{(\lambda_1, \cdots, \lambda_r) \in k^* \mid \lambda_1 \cdots \lambda_r = 1\}.$$

If $\Theta_{\mathcal{L}_1}, \cdots, \Theta_{\mathcal{L}_r}$ are theta-structures on $G(\mathcal{L}_1), \cdots, G(\mathcal{L}_r)$ respectively, the *product theta-structure* $\Theta_{\mathcal{L}} := \prod_{i=1}^r \Theta_{\mathcal{L}_i}$ is the isomorphism $\mathcal{H}(\delta) \xrightarrow{\sim} G(\mathcal{L})$ induced by $(h_1, \cdots, h_r) \longmapsto (\Theta_{\mathcal{L}_1}(h_1), \cdots, \Theta_{\mathcal{L}_r}(h_r))$. This theta-structure induces a natural symplectic decomposition of $K(\mathcal{L}) = K_1(\Theta_{\mathcal{L}}) \oplus K_2(\Theta_{\mathcal{L}})$, where $K_i(\Theta_{\mathcal{L}}) := \prod_{j=1}^r K_i(\Theta_{\mathcal{L}_j})$ for $i \in \{1,2\}$.

Lemma 5.3.37. *For all* $i := (i_1, \cdots, i_r) \in K_1(\delta^{(1)}) \times \cdots \times K_1(\delta^{(r)})$,

$$\theta_i^{\mathcal{L}} = \bigotimes_{j=1}^{r} \pi_j^* \theta_{i_j}^{\mathcal{L}_j}.$$

Proof. With the notations of Section 5.1.5, we have $V(\delta) \simeq \bigoplus_{j=1}^r V(\delta^{(j)})$ and $\Gamma(A, \mathcal{L}) = \bigotimes_{j=1}^r \pi_j^* \Gamma(A_j, \mathcal{L}_j)$. For all $j \in [\![1\,;\,r]\!]$, let $\beta_j : V(\delta^{(j)}) \xrightarrow{\sim} \Gamma(A_j, \mathcal{L}_j)$ be an isomorphism of representations satisfying Eq. (5.6) for $\Theta_{\mathcal{L}_j}$. Then, the isomorphism $\beta : V(\delta) \xrightarrow{\sim} \Gamma(A, \mathcal{L})$, $v_1 \otimes \cdots \otimes v_r \longmapsto \pi_1^*\beta_1(v_1) \otimes \cdots \otimes \pi_r^*\beta_r(v_r)$ is also an isomorphism of representations satisfying Eq. (5.6) for the product theta-structure $\Theta_{\mathcal{L}}$. The result follows. □

In concrete terms, Lemma 5.3.37 above ensures that if $(x_1, \cdots, x_r) \in A$ and $i := (i_1, \cdots, i_r) \in K_1(\delta)$, then we have:

$$\theta_i^{\mathcal{L}}(x_1, \cdots, x_r) = \prod_{j=1}^{r} \theta_{i_j}^{\mathcal{L}_j}(x_j). \tag{5.14}$$

Unsurprisingly, it suffices to multiply theta functions to obtain product theta functions. In the following, we denote $\theta_i^{\mathcal{L}} := \theta_{i_1}^{\mathcal{L}_1} \star \cdots \star \theta_{i_r}^{\mathcal{L}_r}$ instead of $\bigotimes_{j=1}^r \pi_j^* \theta_{i_j}^{\mathcal{L}_j}$ to lighten the notations.

The duplication formula

Theorem 5.3.38 (Duplication formula)**.** *Let* $\mathcal{L}$ *be a totally symmetric line bundle of type* δ *on* A *and* $(\Theta_{\mathcal{L}}, \Theta_{\mathcal{L}^2})$ *be a pair of compatible and symmetric theta structures for* $(\mathcal{L}, \mathcal{L}^2)$*. Consider the endomorphism:*

$$\begin{array}{rcl} \xi : A^2 & \longrightarrow & A^2 \\ (x, y) & \longmapsto & (x + y, x - y) \end{array}$$

Then:

(i) $\xi^*(\pi_1^*\mathcal{L} \otimes \pi_2^*\mathcal{L}) \simeq \pi_1^*\mathcal{L}^2 \otimes \pi_2^*\mathcal{L}^2$.

(ii) The product theta structures $\Theta_{\mathcal{L}^2} \times \Theta_{\mathcal{L}^2}$ *and* $\Theta_{\mathcal{L}} \times \Theta_{\mathcal{L}}$ *are compatible with respect to* ξ.

(iii) There exists $\lambda \in k^*$ *such that for all* $(i_1, i_2) \in K_1(\delta)^2$*, we have*

$$\xi^*\left(\theta_{i_1}^{\mathcal{L}} \star \theta_{i_2}^{\mathcal{L}}\right) = \lambda \sum_{\substack{(j_1,j_2) \in K_1(2\delta)^2 \\ j_1 + j_2 = 2i_1 \\ j_1 - j_2 = 2i_2}} \theta_{j_1}^{\mathcal{L}^2} \star \theta_{j_2}^{\mathcal{L}^2}.$$

Proof. **(i)** By the seesaw principle (Theorem 1.4.14), to prove that $\xi^*(\pi_1^*\mathcal{L} \otimes \pi_2^*\mathcal{L})$ and $\pi_1^*\mathcal{L}^2 \otimes \pi_2^*\mathcal{L}^2$ are isomorphic on A^2, it suffices to prove that they are isomorphic on $A \times \{a\}$ for all $a \in A$ and on $\{0\} \times A$. Let $a \in A$ that we see as a morphism $\mathrm{Spec}(\kappa(x)) \longrightarrow A$. Then, on the one hand

$$\begin{aligned} \xi^*(\pi_1^*\mathcal{L} \otimes \pi_2^*\mathcal{L})_{|A\times\{a\}} &= (\mathrm{id}_A \times a)^*\xi^*(\pi_1^*\mathcal{L} \otimes \pi_2^*\mathcal{L}) = (\pi_1 \circ \xi \circ (\mathrm{id}_A \times a))^*\mathcal{L} \otimes (\pi_2 \circ \xi \circ (\mathrm{id}_A \times a))^*\mathcal{L} \\ &= t_a^*\mathcal{L} \otimes t_{-a}^*\mathcal{L} \simeq t_0^*\mathcal{L} \otimes \mathcal{L} \quad \text{(by the theorem of the square, Theorem 1.4.19)} \\ &= \mathcal{L}^2, \end{aligned}$$

and on the other hand

$$\begin{aligned} (\pi_1^*\mathcal{L}^2 \otimes \pi_2^*\mathcal{L}^2)_{|A\times\{a\}} &= (\mathrm{id}_A \times a)^*(\pi_1^*\mathcal{L}^2 \otimes \pi_2^*\mathcal{L}^2) = (\pi_1 \circ (\mathrm{id}_A \times a))^*\mathcal{L}^2 \otimes (\pi_2 \circ (\mathrm{id}_A \times a))^*\mathcal{L}^2 \\ &= \mathrm{id}_A^*\mathcal{L}^2 \otimes a^*\mathcal{L}^2 \simeq \mathcal{L}^2, \end{aligned}$$

so that $\xi^*(\pi_1^*\mathcal{L} \otimes \pi_2^*\mathcal{L})_{|A\times\{a\}} \simeq (\pi_1^*\mathcal{L}^2 \otimes \pi_2^*\mathcal{L}^2)_{|A\times\{a\}}$. We obtain similarly that $\xi^*(\pi_1^*\mathcal{L} \otimes \pi_2^*\mathcal{L})_{|\{0\}\times A} \simeq (\pi_1^*\mathcal{L}^2 \otimes \pi_2^*\mathcal{L}^2)_{|\{0\}\times A} \simeq \mathcal{L}^2$. This proves (i).

(ii) Since $\Theta_{\mathcal{L}}$ and $\Theta_{\mathcal{L}^2}$ are symmetric, we easily obtain that their products $\Theta_{\mathcal{L}} \times \Theta_{\mathcal{L}}$ and $\Theta_{\mathcal{L}^2} \times \Theta_{\mathcal{L}^2}$ are symmetric. Since $\mathcal{L}$ is totally symmetric, we also obtain that $\pi_1^*\mathcal{L} \otimes \pi_2^*\mathcal{L}$ is totally symmetric by

Proposition 5.3.5 (points (ii) and (iii)). Hence, we can apply Proposition 5.3.36 provided its conditions (i) and (ii) are satisfied.

Let $\overline{\Theta}_{\mathcal{L}^4} : K(4\delta) \xrightarrow{\sim} K(\mathcal{L}^4)$ be a symplectic isomorphism extending $\overline{\Theta}_{\mathcal{L}^2} : K(2\delta) \xrightarrow{\sim} K(\mathcal{L}^2)$ and inducing $\Theta_{\mathcal{L}^2}$ by Theorem 5.3.30.(iii). The kernel of ξ is $K := \{(x,x) \mid x \in A[2]\}$ and we easily obtain that $A[2] = (K_1(\overline{\Theta}_{\mathcal{L}^4}) \cap A[2]) \oplus (K_2(\overline{\Theta}_{\mathcal{L}^4}) \cap A[2])$, so that

$$K = ((K_1(\overline{\Theta}_{\mathcal{L}^4}) \times K_1(\overline{\Theta}_{\mathcal{L}^4})) \cap K) \oplus ((K_2(\overline{\Theta}_{\mathcal{L}^4}) \times K_2(\overline{\Theta}_{\mathcal{L}^4})) \cap K),$$

which proves the first condition of Proposition 5.3.36.

For the second condition, we have to prove that for $i \in \{1,2\}$, we have:

$$\xi((K_i(\overline{\Theta}_{\mathcal{L}^4}) \times K_i(\overline{\Theta}_{\mathcal{L}^4})) \cap \xi^{-1}(K(\mathcal{L}^2) \times K(\mathcal{L}^2))) = K_i(\overline{\Theta}_{\mathcal{L}^2}) \times K_i(\overline{\Theta}_{\mathcal{L}^2}). \tag{5.15}$$

$\subseteq$ Let $i \in \{1,2\}$ and $(x,y) \in (K_i(\overline{\Theta}_{\mathcal{L}^4}) \times K_i(\overline{\Theta}_{\mathcal{L}^4})) \cap \xi^{-1}(K(\mathcal{L}^2) \times K(\mathcal{L}^2))$. Then $x+y$ and $x-y$ belong to $K(\mathcal{L}^2) \cap K_i(\overline{\Theta}_{\mathcal{L}^4}) = K_i(\overline{\Theta}_{\mathcal{L}^2})$ so that $\xi(x,y) \in K_i(\overline{\Theta}_{\mathcal{L}^2}) \times K_i(\overline{\Theta}_{\mathcal{L}^2})$.

$\supseteq$ Conversely, let $(s,t) \in K_i(\overline{\Theta}_{\mathcal{L}^2}) \times K_i(\overline{\Theta}_{\mathcal{L}^2})$. Since $K_i(\overline{\Theta}_{\mathcal{L}^2}) = [2]K_i(\overline{\Theta}_{\mathcal{L}^4})$, there exists $x, y \in K_i(\overline{\Theta}_{\mathcal{L}^4})$ such that $s+t = 2x$ and $s-t = 2y$. We then have $[2](s,t) - [2]\xi(x,y) = 0$ and $s - (x+y) = t - (x-y)$ so there exists $z \in A[2]$ such that $(s,t) = \xi(x,y) + (z,z) = \xi(x, y+z)$. Since $z = s - (x+y)$, we have $z \in K_i(\overline{\Theta}_{\mathcal{L}^4})$, so that $(x, y+z) \in K_i(\overline{\Theta}_{\mathcal{L}^4}) \times K_i(\overline{\Theta}_{\mathcal{L}^4})$. Also, $(s,t) \in K_i(\overline{\Theta}_{\mathcal{L}^2}) \times K_i(\overline{\Theta}_{\mathcal{L}^2}) \subseteq K(\mathcal{L}^2) \times K(\mathcal{L}^2)$. This proves Eq. (5.15) and the compatibility of $\Theta_{\mathcal{L}^2} \times \Theta_{\mathcal{L}^2}$ and $\Theta_{\mathcal{L}} \times \Theta_{\mathcal{L}}$ by Proposition 5.3.36.

(iii) We apply the isogeny theorem (Theorem 5.2.5). By Proposition 5.2.4, $\Theta_{\mathcal{L}} \times \Theta_{\mathcal{L}}$ is induced by $\Theta_{\mathcal{L}^2} \times \Theta_{\mathcal{L}^2}$ and an isomorphism $\sigma : K^{\perp,1}/K_1 \xrightarrow{\sim} K_1(\delta) \times K_1(\delta)$, where $K^\perp$ is the orthogonal of K for $e_{\mathcal{L}^2}$, $K^{\perp,1} := K^\perp \cap K_1(\overline{\Theta}_{\mathcal{L}^2})$ and $K_1 := K \cap K_1(\overline{\Theta}_{\mathcal{L}^2})$. Then Theorem 5.2.5 ensures the existence of $\lambda \in k^*$ such that for all $i_1, i_2 \in K_1(\delta)$,

$$\xi^*\left(\theta^{\mathcal{L}}_{i_1} \star \theta^{\mathcal{L}}_{i_2}\right) = \lambda \sum_{(j_1,j_2) \in (\overline{\Theta}_{\mathcal{L}^2} \times \overline{\Theta}_{\mathcal{L}^2})^{-1}(\sigma^{-1}(\{(i_1,i_2)\}))} \theta^{\mathcal{L}^2}_{j_1} \star \theta^{\mathcal{L}^2}_{j_2}. \tag{5.16}$$

By the diagram 5.12, we obtain that

$$\forall (x,y) \in K^{\perp,1}, \quad \overline{\Theta}_{\mathcal{L}} \times \overline{\Theta}_{\mathcal{L}}(\sigma(x,y),1) = \xi(x,y).$$

Besides, we know that $\varepsilon_2 \circ \Theta_{\mathcal{L}} = \Theta_{\mathcal{L}^2} \circ E_2$, so that $\overline{\Theta}_{\mathcal{L}}(i,1) = \overline{\Theta}_{\mathcal{L}^2}(2i,1)$ for all $i \in K(\delta)$. It follows that for all $i_1, i_2 \in K_1(\delta)$,

$$\begin{aligned}
(\overline{\Theta}_{\mathcal{L}^2} \times \overline{\Theta}_{\mathcal{L}^2})^{-1}(\sigma^{-1}(\{(i_1,i_2)\})) &= \{(j_1,j_2) \in K_1(2\delta)^2 \mid \exists (x,y) \in K^{\perp,1}, \quad \sigma(x,y) = (i_1,i_2) \\
&\qquad \text{and} \quad (\overline{\Theta}_{\mathcal{L}^2}(j_1), \overline{\Theta}_{\mathcal{L}^2}(j_2)) = \sigma(x,y)\} \\
&= \{(j_1,j_2) \in K_1(2\delta)^2 \mid \exists (x,y) \in K^{\perp,1}, \quad \sigma(x,y) = (i_1,i_2) \\
&\qquad \text{and} \quad (\overline{\Theta}_{\mathcal{L}^2}(j_1,1), \overline{\Theta}_{\mathcal{L}^2}(j_2,1)) = (x,y)\} \\
&= \{(j_1,j_2) \in K_1(2\delta)^2 \mid \exists (x,y) \in K^{\perp,1}, \quad (\overline{\Theta}_{\mathcal{L}}(i_1,1), \overline{\Theta}_{\mathcal{L}}(i_2,1)) \\
&\qquad = \xi(x,y) \quad \text{and} \quad (\overline{\Theta}_{\mathcal{L}^2}(j_1,1), \overline{\Theta}_{\mathcal{L}^2}(j_2,1)) = (x,y)\} \\
&= \{(j_1,j_2) \in K_1(2\delta)^2 \mid (\overline{\Theta}_{\mathcal{L}^2}(j_1+j_2,1), \overline{\Theta}_{\mathcal{L}^2}(j_1-j_2,1)) \\
&\qquad = (\overline{\Theta}_{\mathcal{L}}(i_1,1), \overline{\Theta}_{\mathcal{L}}(i_2,1))\} \\
&= \{(j_1,j_2) \in K_1(2\delta)^2 \mid (j_1+j_2, j_1-j_2) = (2i_1, 2i_2)\}
\end{aligned}$$

The above equality together with Eq. (5.16) complete the proof. □

A first arithmetic application of the duplication formula, is a very simple way to compute the opposite of a point $x \longmapsto -x$. In particular, in level 2 ($\delta = (2, \cdots, 2)$) we obtain that points and their opposite have the same theta coordinates, as they should on the Kummer variety $K_A = A/\pm$.

Corollary 5.3.39. *Let $\mathcal{L}$ be a totally symmetric line bundle of level δ and $\Theta_{\mathcal{L}}$ be a symmetric theta structure on $G(\mathcal{L})$. Let $(\theta^{\mathcal{L}}_i)_{i \in K_1(\delta)}$ be the associated basis of theta-functions of $\Gamma(A,\mathcal{L})$. Then, we have:*

$$\forall i \in K_1(\delta), \quad [-1]^*\theta^{\mathcal{L}}_i = \theta^{\mathcal{L}}_{-i}$$

Proof. Since $\Theta_{\mathcal{L}}$ is symmetric, it is compatible with itself with respect to $[-1]$ so we can apply the isogeny theorem 5.2.5 and obtain the existence of $\lambda \in k^*$ such that $[-1]^*\theta_i = \lambda\theta_{-i}$ for all $i \in K_1(\delta)$. Since $[-1]^2 = \mathrm{id}_A$, we must have $\lambda^2 = 1$, so $\lambda \in \{\pm 1\}$. We now prove that $\lambda = 1$.

Let $i \in K_1(\delta)$. Then, the duplication formula from Theorem 5.3.38 ensures the existence of $\mu \in k^*$ such that

$$\xi^*\left(\theta_i^{\mathcal{L}} \star \theta_i^{\mathcal{L}}\right) = \mu \sum_{\substack{(j_1,j_2)\in K_1(2\delta)^2 \\ j_1+j_2=2i \\ j_1-j_2=2i}} \theta_{j_1}^{\mathcal{L}^2} \star \theta_{j_2}^{\mathcal{L}^2} = \mu \sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \theta_{2i+t\delta}^{\mathcal{L}^2} \star \theta_{t\delta}^{\mathcal{L}^2},$$

and

$$\xi^*\left(\theta_i^{\mathcal{L}} \star \theta_{-i}^{\mathcal{L}}\right) = \mu \sum_{\substack{(j_1,j_2)\in K_1(2\delta)^2 \\ j_1+j_2=2i \\ j_1-j_2=-2i}} \theta_{j_1}^{\mathcal{L}^2} \star \theta_{j_2}^{\mathcal{L}^2} = \mu \sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \theta_{t\delta}^{\mathcal{L}^2} \star \theta_{2i+t\delta}^{\mathcal{L}^2}.$$

Applying these formulas at points $(0,x)$ and $(x,0)$ for $x \in A(k)$ respectively, we get

$$\theta_i^{\mathcal{L}}(x)\cdot\theta_i^{\mathcal{L}}(-x) = \mu \sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \theta_{2i+t\delta}^{\mathcal{L}^2}(0)\cdot\theta_{t\delta}^{\mathcal{L}^2}(x) \quad \text{and} \quad \theta_i^{\mathcal{L}}(x)\cdot\theta_{-i}^{\mathcal{L}}(x) = \mu \sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \theta_{t\delta}^{\mathcal{L}^2}(x)\cdot\theta_{2i+t\delta}^{\mathcal{L}^2}(0),$$

so that $\theta_i^{\mathcal{L}}(x)\cdot\theta_i^{\mathcal{L}}(-x) = \theta_i^{\mathcal{L}}(x)\cdot\theta_{-i}^{\mathcal{L}}(x)$. Hence it suffices to prove that there exists at least one section non-vanishing at 0 among the $\theta_i^{\mathcal{L}}$ to prove that $\lambda = 1$ so that $[-1]^*\theta_i = \theta_{-i}$. Since $\mathcal{L}$ is totally symmetric, it is the square of a symmetric line bundle by Proposition 5.3.8 and then Theorem 5.1.34.(i) ensures that $\mathcal{L}^2$ is generated by global sections, which implies the desired result. □

Differential addition formulas

In the following, $\mathcal{L}$ will always be a totally symmetric line bundle on A. The duplication formula is not very convenient because it involves a change of level (between theta coordinates for $\mathcal{L}$ and $\mathcal{L}^2$). A priori, if we want to compute $x+y$ knowing x, y and $x-y$ we might not know the $\theta_u^{\mathcal{L}^2}(x)$ and $\theta_v^{\mathcal{L}^2}(y)$ but only the $\theta_u^{\mathcal{L}}(x)$ and $\theta_v^{\mathcal{L}}(y)$. We circumvent this difficulty by considering "Fourier transforms" of theta coordinates.

Definition 5.3.40. For all $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$ and $i \in K_1(2\delta)$, define:

$$U_{\chi,i}^{\mathcal{L}^2} = \sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_{i+t\delta}^{\mathcal{L}^2}.$$

We call the $(U_{\chi,i})_{\chi,i}$ the *dual theta-coordinates* of the $(\theta_i^{\mathcal{L}^2})_{i\in K(2\delta)}$.

Theorem 5.3.41. *Let $x, y \in A(k)$. Then, there exists $\lambda_1, \lambda_2 \in k^*$ such that for all $i, j \in K_1(2\delta)$ such that $i \equiv j \mod \underline{2}$ and $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$, we have*

$$\theta_{(i+j)/2}^{\mathcal{L}}(x+y)\theta_{(i-j)/2}^{\mathcal{L}}(x-y) = \frac{\lambda_1}{2^g} \sum_{\chi\in\widehat{(\mathbb{Z}/2\mathbb{Z})^g}} U_{\chi,i}^{\mathcal{L}^2}(x)U_{\chi,j}^{\mathcal{L}^2}(y) \tag{5.17}$$

$$U_{\chi,i}^{\mathcal{L}^2}(x)U_{\chi,j}^{\mathcal{L}^2}(y) = \lambda_2 \sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_{(i+j+t\delta)/2}^{\mathcal{L}}(x+y)\theta_{(i-j+t\delta)/2}^{\mathcal{L}}(x-y) \tag{5.18}$$

Proof. Since $i \equiv j \mod \underline{2}$, we have $i+j \equiv 0 \mod \underline{2}$ and $i-j \equiv 0 \mod \underline{2}$ so $(i+j)/2$ and $(i-j)/2$ are well defined in $K_1(\delta)$. If $u, v \in K_1(2\delta)$ satisfy $u+v = i+j$ and $u-v = i-j$, then we have $2(u-i) = 0$ so $u = i+t\delta$ with $t \in (\mathbb{Z}/2\mathbb{Z})^g$ and $v = i+j-u = j+t\delta$. Conversely, $u = i+t\delta$ and $v = j+t\delta$ satisfy $u+v = i+j$ and $u-v = i-j$ for all $t \in (\mathbb{Z}/2\mathbb{Z})^g$. Hence, by Theorem 5.3.38, there exists $\lambda_1 \in k^*$ such that

$$\theta_{(i+j)/2}^{\mathcal{L}}(x+y)\theta_{(i-j)/2}^{\mathcal{L}}(x-y) = \lambda_1 \sum_{\substack{(u,v)\in K_1(2\delta)^2 \\ u+v=i+j \\ u-v=i-j}} \theta_u^{\mathcal{L}^2}(x)\theta_v^{\mathcal{L}^2}(y) = \lambda_1 \sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \theta_{i+t\delta}^{\mathcal{L}^2}(x)\theta_{j+t\delta}^{\mathcal{L}^2}(y)$$

Besides

$$\sum_{\chi\in(\widehat{\mathbb{Z}/2\mathbb{Z}})^g} U^{\mathcal{L}^2}_{\chi,i}(x)U^{\mathcal{L}^2}_{\chi,j}(y) = \sum_{s,t\in(\mathbb{Z}/2\mathbb{Z})^g}\left(\sum_{\chi\in(\widehat{\mathbb{Z}/2\mathbb{Z}})^g}\chi(s+t)\right)\theta^{\mathcal{L}^2}_{\chi,i+s\delta}(x)\theta^{\mathcal{L}^2}_{\chi,j+t\delta}(y).$$

And for all $s,t\in(\mathbb{Z}/2\mathbb{Z})^g$,

$$\begin{aligned}\sum_{\chi\in(\widehat{\mathbb{Z}/2\mathbb{Z}})^g}\chi(s+t) &= \sum_{(\chi_1,\cdots,\chi_g)\in(\widehat{\mathbb{Z}/2\mathbb{Z}})^g}\prod_{i=1}^g\chi_i(s_i+t_i) = \prod_{i=1}^g\left(\sum_{\chi_i\in(\widehat{\mathbb{Z}/2\mathbb{Z}})}\chi_i(s_i+t_i)\right)\\ &= \prod_{i=1}^g\left(1+(-1)^{s_i+t_i}\right) = \prod_{i=1}^g 2\delta_{s_i,t_i} = 2^g\delta_{s,t}.\end{aligned}$$

It follows that

$$\sum_{\chi\in(\widehat{\mathbb{Z}/2\mathbb{Z}})^g} U^{\mathcal{L}^2}_{\chi,i}(x)U^{\mathcal{L}^2}_{\chi,j}(y) = 2^g\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g}\theta^{\mathcal{L}^2}_{\chi,i+t\delta}(x)\theta^{\mathcal{L}^2}_{\chi,j+t\delta}(y).$$

Eq. (5.17) follows.

Now we prove Eq. (5.18):

$$\begin{aligned}U^{\mathcal{L}^2}_{\chi,i}(x)U^{\mathcal{L}^2}_{\chi,j}(y) &= \sum_{r,s\in(\mathbb{Z}/2\mathbb{Z})^g}\chi(r+s)\theta^{\mathcal{L}}_{i+r\delta}(x)\theta^{\mathcal{L}}_{j+s\delta}(y)\\ &= \sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g}\chi(t)\sum_{\substack{(u,v)\in K_1(2\delta)^2\\ u+v=i+j+t\delta\\ u-v=i-j+t\delta}}\theta^{\mathcal{L}^2}_u(x)\theta^{\mathcal{L}^2}_v(y)\\ &\qquad\text{(change of variables } t:=r+s, u:=i+r\delta, v:=j+s\delta)\\ &= \lambda_2\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g}\chi(t)\theta^{\mathcal{L}}_{(i+j+t\delta)/2}(x+y)\theta^{\mathcal{L}}_{(i-j+t\delta)/2}(x-y)\quad\text{(by Theorem 5.3.38),}\end{aligned}$$

where $\lambda_2\in k^*$. This proves Eq. (5.18). □

Differential addition and duplication algorithms

We call Eq. (5.17) and Eq. (5.18) differential addition formulas because they yield a differential addition algorithm. Assume that we want to compute the theta-coordinates $(\theta^{\mathcal{L}}_i(x+y))_i$ of $x+y$ knowing those of $x, y, x-y$ and $0\in A(k)$. Assume that $\theta^{\mathcal{L}}_i(x-y)\neq 0$ for all $i\in K_1(\delta)$. Then Eq. (5.17) ensures that for all $i\in K_1(\delta)$,

$$\theta^{\mathcal{L}}_i(x+y) = \frac{1}{\theta^{\mathcal{L}}_i(x-y)}\sum_{\chi\in(\widehat{\mathbb{Z}/2\mathbb{Z}})^g} U^{\mathcal{L}^2}_{\chi,2i}(x)U^{\mathcal{L}^2}_{\chi,0}(y)$$

up to a projective constant that we can ignore. To compute the $U^{\mathcal{L}^2}_{\chi,i}(x)U^{\mathcal{L}^2}_{\chi,0}(y)$, we use Eq. (5.18) twice to obtain:

$$U^{\mathcal{L}^2}_{\chi,2i}(x)U^{\mathcal{L}^2}_{\chi,0}(y) = \frac{1}{U^{\mathcal{L}^2}_{\chi,0}(0)^2}\left(\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g}\chi(t)\theta^{\mathcal{L}}_{i+t\delta/2}(x)^2\right)\left(\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g}\chi(t)\theta^{\mathcal{L}}_{t\delta/2}(y)^2\right).$$

Hence, we can compute the theta-coordinates of $x+y$ provided that $U^{\mathcal{L}^2}_{\chi,0}(0)\neq 0$. The $U^{\mathcal{L}^2}_{\chi,0}(0)^2$ can be precomputed, using Eq. (5.18) again:

$$U^{\mathcal{L}^2}_{\chi,0}(0)^2 = \sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g}\chi(t)\theta^{\mathcal{L}}_{t\delta/2}(0)^2.$$

This differential addition procedure is summarized in Algorithm 5.1. Note that we also derive easily a duplication procedure $x\longmapsto 2x$ summarized in Algorithm 5.2.

Remark 5.3.42. These algorithms work under the assumption that all dual theta constants $U^{\mathcal{L}^2}_{\chi,0}(0)$ do not vanish. Algorithm 5.2 requires in addition non vanishing theta constants $\theta^{\mathcal{L}}_i(0)$. These conditions are of course not always respected. In [Mum66, p. 339] and [Rob10, p. 81-82], it has been proved that when an even integer $n \geq 4$ divides δ, there are enough non vanishing dual theta constants $U^{\mathcal{L}^2}_{\chi,j}(0)$ to obtain $(\theta^{\mathcal{L}}_i(x+y))_i$ from the theta-coordinates of $x, y, x-y$ and Eqs. (5.17) and (5.18). This is unfortunately not the case for level 2 ($\delta = (2, \cdots, 2)$). However, we shall see in Chapter 6 that in practice the theta constants $\theta^{\mathcal{L}}_i(0)$ and the dual theta constants $U^{\mathcal{L}^2}_{\chi,0}(0)$ do not vanish when $(A, \mathcal{L})$ is neither a product of polarised abelian varieties nor 2-isogenous to such a product.

Remark 5.3.43. Algorithms 5.1 and 5.2 use several inversions which are much more costly over finite fields than multiplications. To reduce this cost, we can use a well known batch inversion method (Algorithm 5.3) to compute n inverses of field elements at the expense of only 1 inversion and $3(n-1)$ multiplications. Alternatively, since theta-coordinates are projective, we can remove all inversions from Algorithms 5.1 and 5.2 if we replace them by well chosen multiplications (*e.g.* by using Algorithm 5.3 without the inversion on Line 4, see Algorithm 6.3). In Algorithms 5.4 and 5.5, we propose inversion free fully projective versions of Algorithms 5.1 and 5.2 in dimension 2 with level 2 theta coordinates.

Notation 5.3.44. In the algorithms below, and throughout, we denote by $\mathbf{M}, \mathbf{S}, \mathbf{I}$ and $\mathbf{Sqrt}$ the computational cost of a multiplication, a squaring, an inversion or a square root computation over the field of definition of input data. The cost of additions and substractions, very low in comparison to other arithmetic operations, will be neglected.

Algorithm 5.1: Differential addition.

Data: $(\theta^{\mathcal{L}}_i(x))_{i\in K_1(\delta)}, (\theta^{\mathcal{L}}_i(y))_{i\in K_1(\delta)}, (\theta^{\mathcal{L}}_i(x-y))_{i\in K_1(\delta)}$ and $(\theta^{\mathcal{L}}_i(0))_{i\in K_1(\delta)}$ such that $\theta^{\mathcal{L}}_i(x-y) \neq 0$ for all $i \in K_1(\delta)$ and $U^{\mathcal{L}^2}_{\chi,0}(0) \neq 0$ for all $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$.

Result: $(\theta^{\mathcal{L}}_i(x+y))_{i\in K_1(\delta)}$.

Precompute: $U^{\mathcal{L}^2}_{\chi,0}(0)^{-2} \longleftarrow \left(\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta^{\mathcal{L}}_{t\delta/2}(0)^2\right)^{-1}$ for all $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$.

1 **for** $i \in K_1(\delta)$ **do**
2 **for** $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$ **do**
3 $U^{\mathcal{L}^2}_{\chi,i}(x)U^{\mathcal{L}^2}_{\chi,0}(y) \longleftarrow U^{\mathcal{L}^2}_{\chi,0}(0)^{-2}\left(\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta^{\mathcal{L}}_{i+t\delta/2}(x)^2\right)\left(\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta^{\mathcal{L}}_{t\delta/2}(y)^2\right)$;
4 **end**
5 $\theta^{\mathcal{L}}_i(x+y) \longleftarrow \theta^{\mathcal{L}}_i(x-y)^{-1}\sum_{\chi\in(\widehat{\mathbb{Z}/2\mathbb{Z}})^g} U^{\mathcal{L}^2}_{\chi,i}(x)U^{\mathcal{L}^2}_{\chi,0}(y)$;
6 **end**
7 **return** $(\theta^{\mathcal{L}}_i(x+y))_{i\in K_1(\delta)}$;

Algorithm 5.2: Duplication.

Data: $(\theta^{\mathcal{L}}_i(x))_{i\in K_1(\delta)}$ and $(\theta^{\mathcal{L}}_i(0))_{i\in K_1(\delta)}$ such that $\theta^{\mathcal{L}}_i(0) \neq 0$ for all $i \in K_1(\delta)$ and $U^{\mathcal{L}^2}_{\chi,0}(0) \neq 0$ for all $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$.

Result: $(\theta^{\mathcal{L}}_i(2x))_{i\in K_1(\delta)}$.

Precompute: $U^{\mathcal{L}^2}_{\chi,0}(0)^{-2} \longleftarrow \left(\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta^{\mathcal{L}}_{t\delta/2}(0)^2\right)^{-1}$ for all $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$.

1 **for** $i \in K_1(\delta)$ **do**
2 **for** $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$ **do**
3 $U^{\mathcal{L}^2}_{\chi,i}(x)U^{\mathcal{L}^2}_{\chi,0}(x) \longleftarrow U^{\mathcal{L}^2}_{\chi,0}(0)^{-2}\left(\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta^{\mathcal{L}}_{i+t\delta/2}(x)^2\right)\left(\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta^{\mathcal{L}}_{t\delta/2}(x)^2\right)$;
4 **end**
5 $\theta^{\mathcal{L}}_i(2x) \longleftarrow \theta^{\mathcal{L}}_i(0)^{-1}\sum_{\chi\in(\widehat{\mathbb{Z}/2\mathbb{Z}})^g} U^{\mathcal{L}^2}_{\chi,i}(x)U^{\mathcal{L}^2}_{\chi,0}(x)$;
6 **end**
7 **return** $(\theta^{\mathcal{L}}_i(2x))_{i\in K_1(\delta)}$;

Algorithm 5.3: Batch inversion.

Data: $a_1, \cdots, a_n \in k^*$.
Result: $1/a_1, \cdots, 1/a_n$.
1 $b_1 \longleftarrow a_0$ **for** $i = 2$ **to** n **do**
2 $\quad b_i \longleftarrow b_{i-1} \cdot a_i$; // $b_i = a_1 \cdots a_i$
3 **end**
4 $c_1 \longleftarrow 1/b_n$;
5 **for** $i = 2$ **to** n **do**
6 $\quad c_i \longleftarrow c_{i-1} \cdot a_{n-i+2}$; // $c_i = 1/(a_1 \cdots a_{n-i+1})$
7 **end**
8 $d_1 \longleftarrow c_n$;
9 **for** $i = 2$ **to** n **do**
10 $\quad d_i \longleftarrow b_{i-1} \cdot c_{n-i+1}$; // $d_i = (a_1 \cdots a_{i-1}) \cdot 1/(a_1 \cdots a_i) = 1/a_i$
11 **end**
12 **return** $d_1, \cdots, d_n$; // Total cost: $3(n-1)\mathbf{M} + \mathbf{I}$

The case of dimension 2

Algorithms 5.4 and 5.5 below are inversion free fully projective versions of Algorithms 5.1 and 5.2 in dimension 2 with level 2 theta coordinates ($\delta = (2, 2)$). These algorithms introduced in a paper I coauthored [Gau07] and improved in [LWZ24] optimise the number of multiplications better than the naive method replacing all inversions by a call to Algorithm 5.3 without inversion on Line 4.

In the algorithms below, we denote by $(x_P : y_P : z_P : t_P)$ the level 2 theta coordinates $(\theta_{00}(P) : \theta_{10}(P) : \theta_{01}(P) : \theta_{11}(P))$ of a point P in a 2-dimensional abelian variety A. We also denote by $(a : b : c : d)$ the theta null point $(\theta_{00}(0_A) : \theta_{10}(0_A) : \theta_{01}(0_A) : \theta_{11}(0_A))$. We denote by H the *Hadamard transform* given by the action by multiplication on the left of theta coordinates by the following matrix

$$\begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix}.$$

We also denote by $\mathbf{M}$ and $\mathbf{S}$ the cost of field multiplication and squaring respectively.

Algorithm 5.4: Differential addition in dimension 2 and level 2.

Data: The theta coordinates $(x_P : y_P : z_P : t_P)$, $(x_Q : y_Q : z_Q : t_Q)$ and $(x_{P-Q} : y_{P-Q} : z_{P-Q} : t_{P-Q})$ of P, Q and $P - Q$ respectively and the theta null point $(a : b : c : d)$. We assume that $x_{P-Q} \cdot y_{P-Q} \cdot z_{P-Q} \cdot t_{P-Q} \neq 0$ and $\alpha^2\beta^2\gamma^2\delta^2 \neq 0$, where $(\alpha^2 : \beta^2 : \gamma^2 : \delta^2) := H(a^2 : b^2 : c^2 : d^2)$.
Result: The theta coordinates $(x_{P+Q} : y_{P+Q} : z_{P+Q} : t_{P+Q})$ of $P + Q$.
Precompute: $(\beta^2\gamma^2\delta^2 : \alpha^2\gamma^2\delta^2 : \alpha^2\beta^2\delta^2 : \alpha^2\beta^2\gamma^2)$, where $(\alpha^2 : \beta^2 : \gamma^2 : \delta^2) := H(a^2 : b^2 : c^2 : d^2)$. // Cost: $6\mathbf{M} + 4\mathbf{S}$

1 $(x'_P, y'_P, z'_P, t'_P) \longleftarrow H(x_P^2, y_P^2, z_P^2, t_P^2)$;
2 $(x'_Q, y'_Q, z'_Q, t'_Q) \longleftarrow H(x_Q^2, y_Q^2, z_Q^2, t_Q^2)$;
3 $x_{P+Q} \longleftarrow \beta^2\gamma^2\delta^2 \cdot x'_P \cdot x'_Q$;
4 $y_{P+Q} \longleftarrow \alpha^2\gamma^2\delta^2 \cdot y'_P \cdot y'_Q$;
5 $z_{P+Q} \longleftarrow \alpha^2\beta^2\delta^2 \cdot z'_P \cdot z'_Q$;
6 $t_{P+Q} \longleftarrow \alpha^2\beta^2\gamma^2 \cdot t'_P \cdot t'_Q$;
7 $(x_{P+Q}, y_{P+Q}, z_{P+Q}, t_{P+Q}) \longleftarrow H(x_{P+Q}, y_{P+Q}, z_{P+Q}, t_{P+Q})$;
8 $xy_{P-Q} \longleftarrow x_{P-Q} \cdot y_{P-Q}$;
9 $zt_{P-Q} \longleftarrow z_{P-Q} \cdot t_{P-Q}$;
10 $x_{P+Q} \longleftarrow x_{P+Q} \cdot y_{P-Q} \cdot zt_{P-Q}$;
11 $y_{P+Q} \longleftarrow y_{P+Q} \cdot x_{P-Q} \cdot zt_{P-Q}$;
12 $z_{P+Q} \longleftarrow z_{P+Q} \cdot xy_{P-Q} \cdot t_{P-Q}$;
13 $t_{P+Q} \longleftarrow t_{P+Q} \cdot xy_{P-Q} \cdot z_{P-Q}$;
14 **return** $(x_{P+Q} : y_{P+Q} : z_{P+Q} : t_{P+Q})$; // Total cost: $18\mathbf{M} + 8\mathbf{S}$

Algorithm 5.5: Duplication in dimension 2 and level 2.

Data: The theta coordinates $(x_P : y_P : z_P : t_P)$ of P and the theta null point $(a : b : c : d)$. We assume that $abcd \neq 0$ and $\alpha^2\beta^2\gamma^2\delta^2 \neq 0$, where $(\alpha^2 : \beta^2 : \gamma^2 : \delta^2) := H(a^2 : b^2 : c^2 : d^2)$.
Result: The theta coordinates $(x_{[2]P} : y_{[2]P} : z_{[2]P} : t_{[2]P})$ of $[2]P$.
Precompute: $(bcd : acd : abd : abc)$ and $(\beta^2\gamma^2\delta^2 : \alpha^2\gamma^2\delta^2 : \alpha^2\beta^2\delta^2 : \alpha^2\beta^2\gamma^2)$, where $(\alpha^2 : \beta^2 : \gamma^2 : \delta^2) := H(a^2 : b^2 : c^2 : d^2)$. // Cost: $12\mathbf{M} + 4\mathbf{S}$

$(x_{[2]P}, y_{[2]P}, z_{[2]P}, t_{[2]P}) \longleftarrow H(x_P^2, y_P^2, z_P^2, t_P^2)$;
$x_{[2]P} \longleftarrow x_{[2]P}^2 \cdot \beta^2\gamma^2\delta^2$;
$y_{[2]P} \longleftarrow y_{[2]P}^2 \cdot \alpha^2\gamma^2\delta^2$;
$z_{[2]P} \longleftarrow z_{[2]P}^2 \cdot \alpha^2\beta^2\delta^2$;
$t_{[2]P} \longleftarrow t_{[2]P}^2 \cdot \alpha^2\beta^2\gamma^2$;
$(x_{[2]P}, y_{[2]P}, z_{[2]P}, t_{[2]P}) \longleftarrow H(x_{[2]P}, y_{[2]P}, z_{[2]P}, t_{[2]P})$;
$x_{[2]P} \longleftarrow x_{[2]P} \cdot bcd$;
$y_{[2]P} \longleftarrow y_{[2]P} \cdot acd$;
$z_{[2]P} \longleftarrow z_{[2]P} \cdot abd$;
$t_{[2]P} \longleftarrow t_{[2]P} \cdot abc$;
return $(x_{[2]P} : y_{[2]P} : z_{[2]P} : t_{[2]P})$; // Total cost: $8\mathbf{M} + 8\mathbf{S}$

5.3.3 Level 2 symmetric theta structures on Montgomery curves

For cryptographic applications, we mainly work with isogenies defined between products of elliptic curves. These elliptic curves are usually not given with a level 2 theta model, as we would require to compute product level 2 theta coordinates using Eq. (5.14), but with a Weierstrass or Montgomery model instead. In this section, we introduce conversion formulas between theta and Montgomery coordinates. The following presentation owes a lot to [RS24, § 2.2] and [Rob24, Chapter 7, Appendix A].

Definition 5.3.45. If E is an elliptic curve over k, a *Kummer line* of E is a degree 2 cover $\pi : E \longrightarrow \mathbb{P}^1_k$ with 4 ramification points that induces an isomorphism $E/\pm \xrightarrow{\sim} \mathbb{P}^1_k$. In other words, for all $P \in E(k)$, the fibre $\pi^{-1}(\{\pi(P)\}) = \{-P, P\}$ has cardinality 2, except when P is a point of 2-torsion (a ramification point), in which case it has cardinality 1.

In the following, we consider a Montgomery curve E be over k, whose equation is given by

$$BY^2Z = X^3 + AX^2Z + XZ^2 = X(X - \alpha Z)(X - 1/\alpha Z),$$

with $A, B, \alpha \in k$ and $A = -\alpha - 1/\alpha$. The *Montgomery Kummer line* of E is given by:

$$\begin{array}{rcl} \pi_M : E & \longrightarrow & \mathbb{P}^1_k \\ P := (X : Y : Z) & \longmapsto & \begin{cases} (X : Z) & \text{if } P \neq (0 : 1 : 0) \\ (1 : 0) & \text{if } P = (0 : 1 : 0) \end{cases} \end{array}$$

This Kummer line is used to represent points (up to sign) in $(X : Z)$-arithmetic. Its ramification points are

$$(1 : 0), (\alpha : 1), (0 : 1), (1/\alpha : 1),$$

the first one being the point at infinity.

Similarly, a level 2 theta structure defines a Kummer line on E, $\pi_\theta : P \longmapsto (\theta_0(P) : \theta_1(P))$ by Theorem 5.1.34. Its ramification points are given by Eq. (5.11):

$$(a : b), (b : a), (-a : b), (-b : a),$$

the first one being the theta null point. To obtain conversion formulas between Montgomery and theta coordinates, we use the following lemma.

Lemma 5.3.46. *Let $\pi_1, \pi_2 : E \longrightarrow \mathbb{P}^k_1$ be two Kummer lines. Then:*

(i) There exists a homography of $\mathbb{P}^k_1$, $h : (X : Z) \longmapsto (\alpha X + \beta Z : \gamma X + \delta Z)$ such that $\pi_2 = h \circ \pi_1$.

(ii) This homography is fully determined by the ramification points of π_1 and π_2.

Proof. Let $\overline{\pi}_1, \overline{\pi}_2 : E/\pm \xrightarrow{\sim} \mathbb{P}^k_1$ be the isomorphisms induced by π_1 and π_2 respectively. Then $h := \overline{\pi}_2 \circ \overline{\pi}_1^{-1}$ is an automorphism of $\mathbb{P}^1_k$ so it is a homography [Lan04, Exercise IV.10]. (i) follows immediately.

Now, h is determined by 4 coefficients over k so by its images on 4 distinct points, hence by the ramification points by π_1 and π_2. This proves (ii). □

Lemma 5.3.46 ensures that the change of coordinates map between Montgomery and a level 2 theta coordinates is a homography determined by the ramification points. This will be a key ingredient in the proof of our main result (Proposition 5.3.47).

Before, let us state some basic facts. Consider the line bundle $\mathcal{L}_0 := \mathcal{L}((0_E))$. It induces a principal polarisation as we have seen in Example 1.4.57 so $\mathcal{L}_0^2 \simeq \mathcal{L}(2(0_E))$ is of level 2. Besides, $\mathcal{L}_0$ is symmetric so $\mathcal{L}_0^2$ is totally symmetric by Proposition 5.3.8. Hence, Remark 5.3.31 ensures that a level 2 symmetric theta structure $\Theta_{\mathcal{L}_0^2}$ on $(E, \mathcal{L}_0^2)$ is determined by a basis of $E[4]$. With these facts in mind, we can state and prove the

Proposition 5.3.47. *Let E be a Montgomery curve over k of equation*

$$BY^2Z = X(X - \alpha Z)(X - 1/\alpha Z)$$

Let (P, Q) be a basis of $E[4]$ with $\pi_M(P) = (r : s)$ and $\pi_M(Q) = (-1 : 1)$, where π_M is the Montgomery Kummer line $(X : Y : Z) \longmapsto (X : Z)$. Then (P, Q) induces a level 2 symmetric theta structure $\Theta_{\mathcal{L}_0^2}$ on $(E, \mathcal{L}_0^2)$ with theta null point $(a : b) = (r + s : r - s)$ satisfying $\alpha = (a^2 + b^2)/(a^2 - b^2)$ and the change of coordinates from Montgomery to theta and theta to Montgomery coordinates are respectively given by

$$(X : Z) \longmapsto (\theta_0 : \theta_1) = (a(X - Z) : b(X + Z)) \quad \text{and} \quad (\theta_0 : \theta_1) \longmapsto (X : Z) := (b\theta_0 + a\theta_1 : -b\theta_0 + a\theta_1).$$

Proof. By [CS17, Eq. (10)], the duplication formula in Montgomery (X, Z) arithmetic is

$$[2](X : Z) = ((X + Z)^2(X - Z)^2 : 4XZ((X - Z)^2 + (A + 2)XZ)),$$

with $A := -\alpha - 1/\alpha$. It follows that $[2](-1 : 1) = (0 : 1)$ and $[2](r : s) \in \{(\alpha : 1), (1/\alpha : 1)\}$ since P and Q are linearly independent over $\mathbb{Z}/4\mathbb{Z}$. Without loss of generality, we can swap α and $1/\alpha$ and assume that $[2](r : s) = (\alpha : 1)$.

Let $(a : b)$ be the theta null point of the level 2 symmetric theta structure $\Theta_{\mathcal{L}_0^2}$ induced by (P, Q) and π_θ its associated theta Kummer line. Then $K_1(\overline{\Theta}_{\mathcal{L}_0^2}) = \langle [2]P \rangle$ and $K_2(\overline{\Theta}_{\mathcal{L}_0^2}) = \langle [2]Q \rangle$ so $\pi_\theta([2]P) = (b : a)$ and $\pi_\theta([2]Q) = (-a : b)$ by Eq. (5.11). By Lemma 5.3.46, we then have $\pi_\theta = h \circ \pi_M$ where h maps the ramification points $(1 : 0), (\alpha : 1), (0 : 1), (1/\alpha : 1)$ to $(a : b), (b : a), (-a : b), (-b : a)$ respectively.

Let us write $h : (X : Z) \longmapsto (\alpha' X + \beta Z : \gamma X + \delta Z)$ with $\alpha', \beta, \gamma, \delta \in k$ to be determined. Then, we have $(\alpha' : \gamma) = h(1 : 0) = (a : b)$ so we can assume (up to rescaling all coefficients) that $\alpha' = a$ and $\gamma = b$. Besides,

$$(\beta : \delta) = h(0 : 1) = (-a : b), \quad (\alpha'\alpha + \beta : \gamma\alpha + \delta) = h(\alpha : 1) = (b : a),$$
$$\text{and} \quad (\alpha' + \beta\alpha : \gamma + \delta\alpha) = h(1/\alpha : 1) = (-b : a),$$

so there exists $\lambda, \mu, \nu \in k^*$ such that $\beta = -\lambda a$, $\delta = \lambda b$, $\alpha'\alpha + \beta = \mu b$, $\gamma\alpha + \delta = \mu a$, $\alpha' + \beta\alpha = -\nu b$ and $\gamma + \delta\alpha = \nu a$. Combining the two first equations with the third and fourth, we obtain:

$$a \times a(\alpha - \lambda) = \mu ab = b \times b(\alpha + \lambda),$$

so that

$$\lambda = \frac{a^2 - b^2}{a^2 + b^2}\alpha,$$

where the division is legal because $a^2 \neq -b^2$, otherwise all 2-torsion theta points would not be distinct. Now, combining the two first equations with the two last, we obtain:

$$a \times a(1 - \lambda\alpha) = -\nu ab = -b \times b(a + \lambda\alpha),$$

so that

$$\lambda = \frac{a^2+b^2}{\alpha(a^2-b^2)},$$

where the division is legal because $a^2 \neq b^2$, otherwise all 2-torsion theta points would not be distinct and $\alpha \neq 0$, otherwise E would be a singular curve. Combining both expressions of λ, we obtain that:

$$\alpha = \pm\frac{a^2+b^2}{a^2-b^2} \quad \text{and} \quad \lambda = \pm 1,$$

so that $h : (X : Z) \longmapsto (a(X \mp Z) : b(X \pm Z))$.

By Algorithm 5.2, the duplication formula for theta coordinates is:

$$[2](\theta_0 : \theta_1) = (b(a^2-b^2)(\theta_0^2+\theta_1^2)^2 + b(a^2+b^2)(\theta_0^2-\theta_1^2)^2 : a(a^2-b^2)(\theta_0^2+\theta_1^2)^2 - a(a^2+b^2)(\theta_0^2-\theta_1^2)^2)$$

It follows that $[2](\pm 1 : 1) = (b(a^2-b^2) : a(a^2-b^2)) = (b : a)$, $[2](1 : 0) = (2ba^2 : -2ab^2) = (-a : b)$ and similarly, $[2](0 : 1) = (-a : b)$. Hence, $(1 : 1)$ and $(-1 : 1)$ are 4-torsion theta points lying above $(b : a)$ and $(1 : 0)$ and $(0 : 1)$ are 4-torsion theta points lying above $(-a : b)$. Hence, we may impose that h maps the Montgomery points $\pi_M(P) = (r : s)$ and $\pi_M(Q) = (-1 : 1)$ to $(1 : 1)$ and $(1 : 0)$ respectively. We then obtain $(a(r \mp s) : b(r \pm s)) = h(r : s) = (1 : 1)$ and $(a(-1 \mp 1) : b(-1 \pm 1)) = h(-1 : 1) = (1 : 0)$, so that $\pm = +$ and $a(r-s) = b(r+s)$ so we may assume (up to rescaling a and b) that $a = r+s$ and $b = r-s$, as desired. We conclude that $h : (X : Z) \longmapsto (a(X-Z) : b(X+Z))$.

Since $\pm = +$, we also have $\lambda = 1$ and $\alpha = (a^2+b^2)/(a^2-b^2)$. Finally, we also easily verify that $(\theta_0 : \theta_1) \longmapsto (X : Z) := (b\theta_0 + a\theta_1 : -b\theta_0 + a\theta_1)$ is the inverse of h. This completes the proof. □

Remark 5.3.48. In Proposition 5.3.47, we impose some constraint on the last point Q of the 4-torsion basis inducing the symmetric level 2 theta structure. The change of coordinates formulas are not valid for any symmetric level 2 theta structure, hence any 4-torsion basis. Using change of theta coordinates formulas (related to change of 4-torsion basis), we will be able to convert Montgomery coordinates to theta coordinates and conversely for any choice of 4-torsion basis (see Lemma 6.5.7).

Remark 5.3.49. So far we have always assumed that we work over an algebraically closed field k. However, if $k' \subset k$ is a subfield and the 4-torsion of E is defined over k', the level 2 theta null points we obtain from Proposition 5.3.47 are k'-rational and every k'-rational Montgomery point is converted into a k'-rational theta point. This fact is very useful for isogeny computations that take place over finite fields in practice.

Chapter 6

Computing 2-isogeny chains with level 2 theta coordinates

The goal of this chapter is to explain in detail how the higher dimensional 2^e-isogenies used in the cryptographic applications presented in Part I are computed. Namely, we prove Theorems 2.2.12 and 2.2.13. In the following, we shall work with level 2 theta coordinates that we introduced in Chapter 5. Unlike other models, level 2 theta coordinates make it possible to work efficiently in any dimension g (*e.g.* $g = 4$) while restricting the number of coordinates (to 2^g).

The following presentation is mainly based on an original publication on 2-dimensional 2^e-isogenies [DMPR25] and on a follow-up single author preprint [Dar24] generalizing the 2-dimensional approach to any dimension with a special focus on 4-dimensional 2^e-isogenies for SQIsignHD verification and SIDH attacks. Both papers [DMPR25; Dar24] are based on formulas introduced in a note by Damien Robert [Rob24] without detailed proofs and proved later in [Dar24].

Throughout this chapter, k will be an algebraically closed field of characteristic $\mathrm{char}(k) \neq 2$. All formulas will be proved over k but, in practice, computations will take place over subfields of k. All abelian varieties we shall consider will always be principally polarised with a polarisation induced by a symmetric line bundle. Fortunately, this is the case of elliptic products in particular. If A is a variety of dimension g over k and $\zeta \in k^*$ is a primitive d-th root of unity, we shall refer to $(\zeta, \cdots, \zeta)$-symplectic basis of $A[d]$ in the sense of Lemma 5.1.9 (ζ being repeated g times) as ζ-symplectic basis.

6.1 Computing 2-isogenies

The goal of this section is to describe algorithms using level 2 theta coordinates to compute 2-isogenies between principally polarised abelian varieties of any dimension, as elementary components of 2^e-isogenies decomposed into chains. In Section 6.1.1, we explain how to evaluate a 2-isogeny given the theta null point of its codomain. Then, we explain in section Section 6.1.2 how to compute the codomain theta null point given 8-torsion points lying above the kernel. In Section 6.1.3, we treat the special case when some theta constants (coordinates of the theta null point) vanish. This case may happen when we compute a gluing isogeny $A_1 \times A_2 \longrightarrow B$.

6.1.1 Change of level formula and isogeny evaluation

Consider two principally polarised abelian varieties $(A, \mathcal{L}_0)$ and $(B, \mathcal{M}_0)$ of dimension g with $\mathcal{L}_0$ and $\mathcal{M}_0$ symmetric. Let $f : (A, \mathcal{L}_0^d) \longrightarrow (B, \mathcal{M}_0)$ be a polarised isogeny with $f^*\mathcal{M}_0 \simeq \mathcal{L}_0^d$ for some integer $d \in \mathbb{N}^*$. Then f is a d-isogeny in the sense of Definition 2.2.1, *i.e.* it satisfies $\widehat{f} \circ \varphi_{\mathcal{M}_0} \circ f = \varphi_{\mathcal{L}_0^d} = [d]\varphi_{\mathcal{L}_0}$.

Now, consider $\mathcal{L} := \mathcal{L}_0^2$ and $\mathcal{M} := \mathcal{M}_0^2$. Those are totally symmetric line bundles of level 2 so theta structures on $G(\mathcal{L})$ and $G(\mathcal{M})$ induce level 2 theta coordinates. We also have $f^*\mathcal{M} \simeq \mathcal{L}^d$, so the isogeny theorem Theorem 5.2.5 relates theta coordinates of level $2d$ on the domain with theta coordinates on level 2 on the codomain. More precisely, we have the

Theorem 6.1.1. *Let $\zeta \in k^*$ be a primitive $4d$-th root of unity and $\mathscr{B} := (S_1, \cdots, S_g, T_1, \cdots, T_g)$ be a ζ-symplectic basis of $A[4d] = K(\mathcal{L}^{2d})$ such that $\ker(f) = [4]\langle T_1, \cdots, T_g\rangle$. Then, by Remark 5.3.31, $\mathscr{B}$ induces a symmetric theta structure $\Theta_{\mathcal{L}^d}$ on $(A, \mathcal{L}^d)$ and*

(i) $\mathscr{C} := ([d]f(S_1), \cdots, [d]f(S_g), f(T_1), \cdots, f(T_g))$ is a ζ^d-symplectic basis of $B[4]$ that induces a symmetric theta structure $\Theta_{\mathcal{M}}$ on $(B, \mathcal{M})$ compatible with $\Theta_{\mathcal{L}^d}$ with respect to f.

(ii) If $(\theta_i^{\mathcal{L}^d})_{i\in(\mathbb{Z}/2d\mathbb{Z})^g}$ and $(\theta_i^{\mathcal{M}})_{i\in(\mathbb{Z}/2\mathbb{Z})^g}$ denote the theta functions induced by $\Theta_{\mathcal{L}^d}$ and $\Theta_{\mathcal{M}}$ respectively, then we have

$$\forall i \in (\mathbb{Z}/2\mathbb{Z})^g, \quad f^*\theta_i^{\mathcal{M}} = \theta_{di}^{\mathcal{L}^d}.$$

Definition 6.1.2. We say that the symplectic basis $\mathscr{B}$ from Theorem 6.1.1 is *adapted* to f. We also say that the level $2d$ and 2 theta structures induced by $\mathscr{B}$ and $[d]\mathscr{B}$ respectively are *adapted* to f, along with their associated theta coordinates.

Proof of Theorem 6.1.1. **(i)** By Theorem 5.3.30.(iii), the symmetric theta structure $\Theta_{\mathcal{L}^d}$ is induced by a symmetric theta structure $\Theta_{\mathcal{L}^{2d}}$ on $(A, \mathcal{L}^{2d})$. Then, the isomorphism $\overline{\Theta}_{\mathcal{L}^{2d}} : (\mathbb{Z}/4d\mathbb{Z})^g \xrightarrow{\sim} K(\mathcal{L}^{2d}) = A[4d]$ induced by $\Theta_{\mathcal{L}^{2d}}$ is the one determined by $\mathscr{B}$ that induces $\Theta_{\mathcal{L}^d}$ (as in Remark 5.3.31).

Let $K := \ker(f)$. Then $K = [4]K_2(\overline{\Theta}_{\mathcal{L}^{2d}})$ so that $K = (K \cap K_1(\overline{\Theta}_{\mathcal{L}^{2d}})) \oplus (K \cap K_2(\overline{\Theta}_{\mathcal{L}^{2d}}))$ and by Proposition 5.3.34, $\Theta_{\mathcal{L}^{2d}}$ is compatible with the level subgroup $\widetilde{K}$ induced by a choice of isomorphism $f^*\mathcal{M}_0 \xrightarrow{\sim} \mathcal{L}_0^d$ that we fix. We can then apply Proposition 5.2.4 that ensures the existence of a theta structure $\Theta_{\mathcal{M}^2}$ on $(B, \mathcal{M}^2)$ compatible with $\Theta_{\mathcal{L}^{2d}}$, which is determined by a choice of isomorphism $\sigma : K^{\perp,1}/K_1 \xrightarrow{\sim} (\mathbb{Z}/4\mathbb{Z})^g$, where the orthogonal is defined within $K(\mathcal{L}^{2d})$, so that

$$\begin{aligned} K^\perp &= \{x \in K(\mathcal{L}^{2d}) \mid \forall y \in K,\ e_{\mathcal{L}^{2d}}(x, y) = 1\} = \{x \in K(\mathcal{L}^{2d}) \mid \forall y \in K_2(\overline{\Theta}_{\mathcal{L}^{2d}}),\ e_{\mathcal{L}^{2d}}(x, [4]y) = 1\} \\ &= \{x \in K(\mathcal{L}^{2d}) \mid [4]x \in K_2(\overline{\Theta}_{\mathcal{L}^{2d}})\} = [d]K_1(\overline{\Theta}_{\mathcal{L}^{2d}}) \oplus K_2(\overline{\Theta}_{\mathcal{L}^{2d}}) \end{aligned}$$

and $K^{\perp,1} = K^\perp \cap K_1(\overline{\Theta}_{\mathcal{L}^{2d}}) = [d]K_1(\overline{\Theta}_{\mathcal{L}^{2d}})$. We also have $K_1 = K \cap K_1(\overline{\Theta}_{\mathcal{L}^{2d}}) = \{0\}$, so $\sigma : [d]K_1(\overline{\Theta}_{\mathcal{L}^{2d}}) \xrightarrow{\sim} (\mathbb{Z}/4\mathbb{Z})^g$. Besides, Remark 5.3.33 ensures that $\Theta_{\mathcal{M}^2}$ is symmetric.

Now, we make a choice of σ and compute the resulting choice of $\Theta_{\mathcal{M}^2}$. By the diagram 5.12, we have

$$\overline{\Theta}_{\mathcal{M}^2} \circ (\sigma \times \widehat{\sigma}) = f_{|K^\perp},$$

where $\widehat{\sigma}$ is given by $y \in K^{\perp,2} = K_2(\overline{\Theta}_{\mathcal{L}^{2d}}) \longmapsto e_{\mathcal{L}^{2d}}(\sigma^{-1}(.), y) \in \widehat{(\mathbb{Z}/4\mathbb{Z})^g}$. In particular, we have for all $i \in [\![1\ ;\ g]\!]$, $\overline{\Theta}_{\mathcal{M}^2}(\sigma([d]S_i), 1) = [d]f(S_i)$, so we may naturally set $\sigma([d]S_i) := e_i$ where $e_i \in (\mathbb{Z}/4\mathbb{Z})^g$ equals 1 at index i and 0 everywhere else. This ensures that for all $i \in [\![1\ ;\ g]\!]$,

$$\overline{\Theta}_{\mathcal{M}^2}(e_i, 1) = [d]f(S_i),$$

as desired. Besides, for all $i, j \in [\![1\ ;\ g]\!]$, we have

$$\widehat{\sigma}(T_i)(e_j) = e_{\mathcal{L}^{2d}}(\sigma^{-1}(e_j), T_i) = e_{\mathcal{L}^{2d}}([d]S_j, T_i) = \zeta^{d\delta_{i,j}} = \chi_i(e_j),$$

where $\chi_i \in \widehat{(\mathbb{Z}/4\mathbb{Z})^g}$ is the character $e_j \longmapsto \zeta^{d\delta_{i,j}}$. So we also have for all $i \in [\![1\ ;\ g]\!]$,

$$\overline{\Theta}_{\mathcal{M}^2}(0, \chi_i) = \overline{\Theta}_{\mathcal{M}^2}(0, \widehat{\sigma}(T_i)) = f(T_i),$$

as desired. So $\overline{\Theta}_{\mathcal{M}^2}$ is determined by $\mathscr{C} = ([d]f(S_1), \cdots, [d]f(S_g), f(T_1), \cdots, f(T_g))$.

It remains to prove that the symmetric theta structure $\Theta_{\mathcal{M}}$ induced by $\overline{\Theta}_{\mathcal{M}^2}$ (hence by $\mathscr{C}$) is compatible with $\Theta_{\mathcal{L}^d}$. By Proposition 5.3.36, it suffices to prove that $K = (K \cap K_1(\overline{\Theta}_{\mathcal{L}^{2d}})) \oplus (K \cap K_2(\overline{\Theta}_{\mathcal{L}^{2d}}))$ (which already have been proved) and $f(K_i(\overline{\Theta}_{\mathcal{L}^{2d}}) \cap f^{-1}(K(\mathcal{M}^2))) = K_i(\overline{\Theta}_{\mathcal{M}^2})$ for $i \in \{1, 2\}$. But by construction, we have for all $i \in \{1, 2\}$,

$$\begin{aligned} f(K_i(\overline{\Theta}_{\mathcal{L}^{2d}}) \cap f^{-1}(K(\mathcal{M}^2))) &= f(K_i(\overline{\Theta}_{\mathcal{L}^{2d}}) \cap f^{-1}(A[4])) \\ &= f(K_i(\overline{\Theta}_{\mathcal{L}^{2d}}) \cap ([d]K_1(\overline{\Theta}_{\mathcal{L}^{2d}}) \oplus K_2(\overline{\Theta}_{\mathcal{L}^{2d}}))) = K_i(\overline{\Theta}_{\mathcal{M}^2}). \end{aligned}$$

This proves (i).

(ii) By the isogeny theorem Theorem 5.2.5, we have for all $i \in (\mathbb{Z}/2\mathbb{Z})^g$,

$$f^*\theta_i^{\mathcal{M}} = \sum_{j \in \overline{\Theta}_{\mathcal{L}^d}^{-1}(\tau^{-1}(\{i\}))} \theta_j^{\mathcal{L}},$$

up to a projective constant that we can ignore (rescaling theta functions if necessary), where $\tau : K^{\perp,1}/K_1 = [d]K_1(\overline{\Theta}_{\mathcal{L}^d}) \xrightarrow{\sim} (\mathbb{Z}/2\mathbb{Z})^g$ is the isomorphism associated to $\Theta_{\mathcal{M}}$ by Proposition 5.2.4 (the orthogonal of K being taken in $K(\mathcal{L}^d) = A[2d]$). Since $\overline{\Theta}_{\mathcal{M}} = \overline{\Theta}_{\mathcal{M}^2} \circ [2]$, by the diagram Eq. (5.12), we have for all $l \in [\![1\ ;\ g]\!]$,

$$\overline{\Theta}_{\mathcal{M}^2}([2]\tau([2d]S_l), 1) = \overline{\Theta}_{\mathcal{M}}(\tau([2d]S_l), 1) = f([2d]S_l) = \overline{\Theta}_{\mathcal{M}^2}(\sigma([2d]S_l), 1),$$

so that $[2]\tau([2d]S_l) = \sigma([2d]S_l) = 2e_l \in (\mathbb{Z}/4\mathbb{Z})^g$ and $\tau([2d]S_l) = e_l \in (\mathbb{Z}/2\mathbb{Z})^g$. It follows that for all $i \in (\mathbb{Z}/2\mathbb{Z})^g$,

$$\begin{aligned}\overline{\Theta}_{\mathcal{L}^d}^{-1}(\tau^{-1}(\{i\})) &= \overline{\Theta}_{\mathcal{L}^d}^{-1}\left(\left\{\sum_{l=1}^g [2di_l]S_l\right\}\right) = (\overline{\Theta}_{\mathcal{L}^{2d}} \circ [2])^{-1}\left(\left\{\sum_{l=1}^g [2di_l]S_l\right\}\right) \\ &= \left\{\sum_{l=1}^g di_l \overline{\Theta}_{\mathcal{L}^{2d}}^{-1}(S_l)\right\} = \{di\}\end{aligned}$$

The result follows. □

In the following, we assume that $d = 2$, so that f is a 2-isogeny. Then Theorem 6.1.1 relates theta coordinates of level 2 on the codomain to theta coordinates on level 4 on the domain as follows:

$$\forall x \in A(k), \forall i \in (\mathbb{Z}/2\mathbb{Z})^g, \quad \theta_i^{\mathcal{M}}(f(x)) = \theta_{2i}^{\mathcal{L}^2}(x).$$

This is the simplest possible form of the isogeny theorem obtained by choice of a basis $\mathscr{B}$ adapted to f (hence justifying that choice). However, this is not yet sufficient to evaluate f since we only know level 2 coordinates $(\theta_i^{\mathcal{L}}(x))_{i \in (\mathbb{Z}/2\mathbb{Z})^g}$ and not level 4 coordinates $(\theta_{2i}^{\mathcal{L}^2}(x))_{i \in (\mathbb{Z}/2\mathbb{Z})^g}$. Hence, we need a *change of level* formula to relate these level 2 and level 4 theta coordinates. Fortunately, the differential addition formulas from Theorem 5.3.41 provide what we are looking for.

Corollary 6.1.3. *Under the assumptions of Theorem 6.1.1, for all $x \in A(k)$, there exists $\lambda \in k^*$ such that for all $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$,*

$$U_\chi^{\mathcal{M}}(f(x)) \cdot U_\chi^{\mathcal{M}}(0_B) = \lambda \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_t^{\mathcal{L}}(x)^2, \tag{6.1}$$

where the $U_\chi^{\mathcal{M}}$ are the dual theta coordinates *of the $\theta_i^{\mathcal{M}}$ as defined in Definition 5.3.40 by:*

$$U_\chi^{\mathcal{M}} := U_{\chi,0}^{\mathcal{M}} = \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_t^{\mathcal{M}},$$

where the second index is dropped because unused. The formula relating the $U_\chi^{\mathcal{M}}$ to the $\theta_i^{\mathcal{M}}$ is called a Hadamard transform.

Proof. Let $x \in A(k)$. Then Eq. (5.18) immediately ensures the existence of $\lambda \in k^*$ such that:

$$U_\chi^{\mathcal{L}^2}(x) \cdot U_\chi^{\mathcal{L}^2}(0_A) = \lambda \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_t^{\mathcal{L}}(x)^2,$$

where by Definition 5.3.40 and Theorem 6.1.1.(ii),

$$U_\chi^{\mathcal{L}^2}(x) = \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_t^{\mathcal{L}^2}(x) = \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_t^{\mathcal{M}}(f(x)) = U_\chi^{\mathcal{M}}(f(x)).$$

Similarly, $U_\chi^{\mathcal{L}^2}(0_A) = U_\chi^{\mathcal{M}}(f(0_A)) = U_\chi^{\mathcal{M}}(0_B)$. The result follows. □

When the codomain dual theta null point $(U^{\mathcal{M}}_{\chi}(0_B))_{\chi}$ is known, Corollary 6.1.3 yields an evaluation formula for f. Indeed, from $(\theta^{\mathcal{L}}_i(x))_i$, we obtain the dual theta point $(U^{\mathcal{M}}_{\chi}(f(x)))_{\chi}$ and a simple computation ensures that for all $i \in (\mathbb{Z}/2\mathbb{Z})^g$,

$$\sum_{\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}} \chi(i) U^{\mathcal{M}}_{\chi} = 2^g \theta^{\mathcal{M}}_i,$$

so we may obtain $(\theta^{\mathcal{M}}_i(f(x)))_{i \in (\mathbb{Z}/2\mathbb{Z})^g}$, as desired (up to a projective factor that can be ignored). Algorithm 6.1 follows. In the following section, we shall explain how the codomain theta null point and its dual can be computed. Note that Algorithm 6.1 only works if none of the dual theta constants $U^{\mathcal{M}}_{\chi}(0_B)$ vanish. The vanishing case will be treated in Section 6.1.3.

Algorithm 6.1: Generic 2-isogeny evaluation with level 2 theta coordinates.

Data: Theta coordinates $(\theta^{\mathcal{L}}_i(x))_i$ and inverse of the codomain dual theta null point $(1/U^{\mathcal{M}}_{\chi}(0_B))_{\chi}$.
Result: The theta coordinates $(\theta^{\mathcal{M}}_i(f(x)))_i$.

1 **for** $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$ **do**
2 $\quad U^{\mathcal{M}}_{\chi}(f(x)) \longleftarrow 1/U^{\mathcal{M}}_{\chi}(0_B) \cdot \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t) \theta^{\mathcal{L}}_t(x)^2$;
3 **end**
4 **for** $i \in (\mathbb{Z}/2\mathbb{Z})^g$ **do**
5 $\quad \theta^{\mathcal{M}}_i(f(x)) \longleftarrow \sum_{\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}} \chi(i) U^{\mathcal{M}}_{\chi}(f(x))$;
6 **end**
7 **return** $(\theta^{\mathcal{M}}_i(f(x)))_i$; // Total cost: $2^g\mathbf{M} + 2^g\mathbf{S}$

6.1.2 Computation of the codomain theta null point

In this section, we provide an algorithm to compute the codomain level 2 theta null point of a 2-isogeny $f : (A, \mathcal{L}^2) \longrightarrow (B, \mathcal{M})$ given 8-torsion points over its kernel. Recalling the notations from Section 6.1.1, we assume that $(A, \mathcal{L}_0)$ and $(B, \mathcal{M}_0)$ are principally polarised, that $\mathcal{L}_0$ and $\mathcal{M}_0$ are symmetric, that $f^*\mathcal{M}_0 \simeq \mathcal{L}_0^2$ and denote $\mathcal{L} := \mathcal{L}_0^2$ and $\mathcal{M} := \mathcal{M}_0^2$, which are totally symmetric. We also suppose the assumptions of Theorem 6.1.1 are satisfied, so that:

- We have a pair of compatible symmetric theta structures $(\Theta_{\mathcal{L}}, \Theta_{\mathcal{L}^2})$ for $(\mathcal{L}, \mathcal{L}^2)$;
- We have a symmetric theta structure $\Theta_{\mathcal{M}}$ on $(B, \mathcal{M})$ compatible with $\Theta_{\mathcal{L}^2}$ with respect to f;
- $\Theta_{\mathcal{L}^2}$ is induced by a ζ-symplectic basis $\mathscr{B} := (S_1, \cdots, S_g, T_1, \cdots, T_g)$ of $A[8]$ adapted to f *i.e.* such that $\ker(f) = [4]\langle T_1, \cdots, T_g \rangle$;
- $\Theta_{\mathcal{M}}$ is induced by the ζ^2-symplectic basis $\mathscr{C} := ([2]f(S_1), \cdots, [2]f(S_g), f(T_1), \cdots, f(T_g))$ of $B[4]$.

Then our algorithm returns the dual theta null point $(U^{\mathcal{M}}_{\chi}(0_B))_{\chi}$ given the 8-torsion points $T_1, \cdots, T_g$. Note that 8-torsion points are necessary to determine the theta null point without sign ambiguity. As we shall see in Section 6.5.2, we can work with 4-torsion or even 2-torsion in dimension 2 (and 3) at the expense of (costly) square root computations. Our dual theta null point computation algorithm relies on the following formula. For all $l \in [\![1\ ;\ g]\!]$, let us denote $\chi_l \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$, the character $i \in (\mathbb{Z}/2\mathbb{Z})^g \longmapsto (-1)^{i_l}$. Then, we have the

Lemma 6.1.4. *For all $l \in [\![1\ ;\ g]\!]$ and $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$,*

$$U^{\mathcal{M}}_{\chi \cdot \chi_l}(0_B) \left(\sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t) \theta^{\mathcal{L}}_t(T_l)^2 \right) = U^{\mathcal{M}}_{\chi}(0_B) \left(\sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\chi_l(t) \theta^{\mathcal{L}}_t(T_l)^2 \right).$$

Proof. Let $l \in [\![1\ ;\ g]\!]$. Then, by Corollary 6.1.3, there exists $\lambda \in k^*$ such that for all $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$,

$$U^{\mathcal{M}}_{\chi}(0_B) \cdot U^{\mathcal{M}}_{\chi}(f(T_l)) = \lambda \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t) \theta^{\mathcal{L}}_t(T_l)^2. \tag{6.2}$$

Since $[4]T_l \in \ker(f)$, $f(T_l)$ has order 4 so $f(T_l) \equiv f(T_l) + [2]f(T_l)$ in the Kummer variety $A/\pm$ and $\theta_i^{\mathcal{L}}(f(T_l)) = \theta_i^{\mathcal{L}}(f(T_l)+[2]f(T_l))$ for all $i \in (\mathbb{Z}/2\mathbb{Z})^g$. Besides, we have $[2]f(T_l) = \overline{\Theta}_{\mathcal{M}}(0, \chi_l) \in K_2(\overline{\Theta}_{\mathcal{M}})$ since $\Theta_{\mathcal{M}}$ is induced by the basis $\mathscr{C}$ previously introduced. Hence, Eq. (5.10) ensures that:

$$\forall i \in (\mathbb{Z}/2\mathbb{Z})^g, \quad \theta_i^{\mathcal{M}}(f(T_l)) = \theta_i^{\mathcal{M}}(f(T_l) + [2]f(T_l)) = \chi_l(i)^{-1}\theta_i^{\mathcal{M}}(f(T_l)) = \chi_l(i)\theta_i^{\mathcal{M}}(f(T_l)),$$

since $\chi_l^{-1} = \chi_l$, so that for all $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$,

$$\begin{aligned} U_\chi^{\mathcal{M}}(f(T_l)) &= \sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_t^{\mathcal{M}}(f(T_l)) = \sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\chi_l(t)\theta_t^{\mathcal{M}}(f(T_l)) \\ &= U_{\chi\cdot\chi_l}^{\mathcal{M}}(f(T_l)), \end{aligned} \tag{6.3}$$

Combining this with Eq. (6.2), we finally obtain that for all $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$,

$$\begin{aligned} U_{\chi\cdot\chi_l}^{\mathcal{M}}(0_B)\left(\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_t^{\mathcal{L}}(T_l)^2\right) &= U_{\chi\cdot\chi_l}^{\mathcal{M}}(0_B)\cdot\lambda^{-1}\cdot U_\chi^{\mathcal{M}}(0_B)\cdot U_\chi^{\mathcal{M}}(f(T_l)) \\ &= U_\chi^{\mathcal{M}}(0_B)\cdot\lambda^{-1}\cdot U_{\chi\cdot\chi_l}^{\mathcal{M}}(0_B)\cdot U_{\chi\cdot\chi_l}^{\mathcal{M}}(f(T_l)) \\ &= U_\chi^{\mathcal{M}}(0_B)\left(\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi_l(t)\chi(t)\theta_t^{\mathcal{L}}(T_l)^2\right). \end{aligned}$$

This completes the proof. □

Lemma 6.1.4 ensures that we can deduce $U_{\chi\cdot\chi_l}^{\mathcal{M}}(0_B)$ from $U_\chi^{\mathcal{M}}(0_B)$ provided the element

$$D_{\chi,l} := \sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_t^{\mathcal{L}}(T_l)^2$$

is non-zero. Since the χ_l generate $\widehat{(\mathbb{Z}/2\mathbb{Z})^g}$, this provides a method to compute the projective dual theta null point $(U_\chi^{\mathcal{M}}(0_B))_\chi$ step by step. We start by selecting $\chi_0 \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$ and set $U_{\chi_0}^{\mathcal{M}}(0_B) := 1$. Then we propagate the computation via the formula $U_{\chi\cdot\chi_l}^{\mathcal{M}}(0_B) = U_\chi^{\mathcal{M}}(0_B)D_{\chi\cdot\chi_l,l}/D_{\chi,l}$ assuming that $D_{\chi,l} \neq 0$ until we have computed $U_\chi^{\mathcal{M}}(0_B)$ for all $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$. From an algorithmic point of view, we fill in a computation tree whose vertices are characters of $\widehat{(\mathbb{Z}/2\mathbb{Z})^g}$, whose root is χ_0 and whose edges are quotients $D_{\chi\cdot\chi_l,l}/D_{\chi,l}$ relating a parent node χ and a child node $\chi\cdot\chi_l$. This tree does not contain duplicate edges and can be seen as a subgraph of the hypercube graph (since $\widehat{(\mathbb{Z}/2\mathbb{Z})^g} \simeq (\mathbb{Z}/2\mathbb{Z})^g$). To fill in the tree at every step, we iterate on leaves (*i.e* terminal nodes) χ and on $l \in [\![1\ ;\ g]\!]$. If $\chi\cdot\chi_l$ is not a node and $D_{\chi,l} \neq 0$ we append $\chi\cdot\chi_l$ as a child of χ with edge the $D_{\chi\cdot\chi_l,l}/D_{\chi,l}$. Algorithm 6.2 follows.

Note that in plain generality, the choice of root χ_0 matters. Indeed, if $U_{\chi_0}^{\mathcal{M}}(0_B) = 0$, we cannot expect the computation tree to cover the whole of $\widehat{(\mathbb{Z}/2\mathbb{Z})^g}$, otherwise $(U_\chi^{\mathcal{M}}(0_B))_\chi$ would be identically zero and would not be a projective point. For that reason, Algorithm 6.2 tests all roots $\chi_0 \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$ until the tree is filled in. However, if none of the dual theta constants $U_\chi^{\mathcal{M}}(0_B)$ vanish, then the initial choice of root χ_0 does not matter. If no suitable root is found, we expect f to be a gluing isogeny and we may apply a modified version of Algorithm 6.2 with additional entry points lying over $\ker(f)$ (see Section 6.1.3).

Lemma 6.1.5. *Assume that none of the dual theta constants $U_\chi^{\mathcal{M}}(0_B)$ vanish. Then, if Algorithm 6.2 returns the desired result, namely a computation tree covering the whole of $\widehat{(\mathbb{Z}/2\mathbb{Z})^g}$, it does with the first choice of root χ_0.*

Proof. Let $\chi_0 = 1 \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$ be the first choice of root. Let $\mathcal{T}$ be a tree returned by Algorithm 6.2 with root χ_0'. Then $\mathcal{T}$ covers $\widehat{(\mathbb{Z}/2\mathbb{Z})^g}$ and for every edge of $\mathcal{T}$ connecting χ to $\chi\cdot\chi_l$, we have

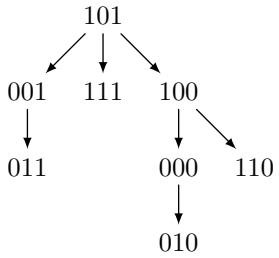

Figure 6.1: Example of tree $\mathcal{T}$ with $g = 3$ and root 101, where an index $i \in (\mathbb{Z}/2\mathbb{Z})^3$ identifies with the character $\chi^i : j \longmapsto (-1)^{\langle i|j \rangle}$.

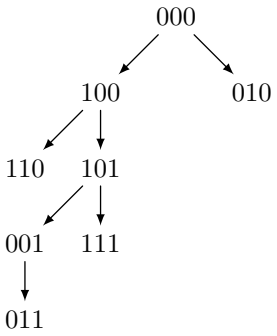

Figure 6.2: Tree $\mathcal{T}_0$ obtained from $\mathcal{T}$ with root 000.

$U^{\mathcal{M}}_{\chi\cdot\chi_l}(0_B)D_{\chi,l} = U^{\mathcal{M}}_{\chi,l}(0)D_{\chi\cdot\chi_l,0}$ by Lemma 6.1.4, where $U^{\mathcal{M}}_{\chi\cdot\chi_l}(0_B), U^{\mathcal{M}}_{\chi,l}(0), D_{\chi,l} \neq 0$ by assumption. It follows that $D_{\chi\cdot\chi_l,l} \neq 0$ so we can connect $\chi \cdot \chi_l$ to χ.

Exploiting this symmetry, we consider the tree $\mathcal{T}_0$ constructed from $\mathcal{T}$ as follows. We initialize $\mathcal{T}_0$ at χ_0. Then, for all leaf χ of $\mathcal{T}_0$, we consider all parents and children $\chi \cdot \chi_l$ of χ in $\mathcal{T}$ that do not belong to $\mathcal{T}_0$ yet and append them as children of χ in $\mathcal{T}_0$ (see Figs. 6.1 and 6.2). Such a tree $\mathcal{T}_0$ covers all nodes of $\mathcal{T}$, hence the whole of $\widehat{(\mathbb{Z}/2\mathbb{Z})^g}$ by construction.

Now, consider the undirected graph with vertex set $\widehat{(\mathbb{Z}/2\mathbb{Z})^g}$ and edges connecting χ and $\chi \cdot \chi_l$ if and only if $D_{\chi,l} \neq 0$ (if and only if $D_{\chi\cdot\chi_l,l} \neq 0$). By construction, Algorithm 6.2 explores this graph from χ_0 until it finds a tree covering the whole of $\widehat{(\mathbb{Z}/2\mathbb{Z})^g}$ if it does exist. Such a tree exists (we may consider $\mathcal{T}_0$ for instance) so Algorithm 6.2 terminates at root χ_0. This completes the proof. □

To save (costly) inversions in the execution of Algorithm 6.2, we save couples $(D_{\chi\cdot\chi_l,l}, D_{\chi,l})$ made of the numerator and denominator instead of quotients $D_{\chi\cdot\chi_l,l}/D_{\chi,l}$ on each edge of the tree. Once we have completed the computation tree $\mathcal{T}$, we can then inverse all the denominators $D_{\chi,l}$ by batch using Algorithm 5.3, then compute each quotient $D_{\chi\cdot\chi_l,l}/D_{\chi,l}$ on each edge and compute the theta constants step by step by descending the tree from the root: $U^{\mathcal{M}}_{\chi_0}(0_B) = 1$ and $U^{\mathcal{M}}_{\chi\cdot\chi_l}(0_B) = U^{\mathcal{M}}_{\chi,l}(0)D_{\chi\cdot\chi_l,0}/D_{\chi,l}$. However, this method still requires one inversion (on Line 4 of Algorithm 5.3) and we propose an inversion free method instead that computes the theta constants projectively.

This method uses Algorithm 6.3, a projective version of Algorithm 5.3 that returns projective inverses $\lambda/a_1 = a_2 \cdots a_n$, $\lambda/a_2 = a_1 a_3 \cdots a_n$, ..., $\lambda/a_n = a_1 \cdots a_{n-1}$ and $\lambda = a_1 \cdots a_n$ when $a_1, \cdots, a_n$ are given on entry. We start by a descending route on the computation tree $\mathcal{T}$ multiplying every numerator and denominator on the edges by the numerator and denominator of the parent edge (see Algorithm 6.4). Then, if $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$ is related to the root χ_0 by a branch $\chi_{l_1}, \cdots, \chi_{l_s}$ in $\mathcal{T}$, so that $\chi = \chi_0 \prod_{i=1}^{s} \chi_{l_i}$, then we have by Lemma 6.1.4 and a simple induction:

$$U^{\mathcal{M}}_{\chi}(0_B) \prod_{i=1}^{s} D_{\chi_0 \prod_{j=1}^{i-1} \chi_{l_j}, l_i} = U^{\mathcal{M}}_{\chi_0}(0_B) \prod_{i=1}^{s} D_{\chi_0 \prod_{j=1}^{i} \chi_{l_j}, l_i},$$

where $(\prod_{i=1}^{s} D_{\chi_0 \prod_{j=1}^{i} \chi_{l_j}, l_i}, \prod_{i=1}^{s} D_{\chi_0 \prod_{j=1}^{i-1} \chi_{l_j}, l_i})$ is the new couple (numerator, denominator) stored on the edge above χ. We can then apply Algorithm 6.3 to all denominators to obtain quotients of the form $\lambda / \prod_{i=1}^{s} D_{\chi_0 \prod_{j=1}^{i-1} \chi_{l_j}, l_i}$, where λ is the product of all denominators. We then set $U^{\mathcal{M}}_{\chi_0}(0_B) := \lambda$ and compute

$$U^{\mathcal{M}}_{\chi}(0_B) = \frac{\lambda}{\prod_{i=1}^{s} D_{\chi_0 \prod_{j=1}^{i-1} \chi_{l_j}, l_i}} \prod_{i=1}^{s} D_{\chi_0 \prod_{j=1}^{i} \chi_{l_j}, l_i}$$

for all non root character $\chi \in \mathcal{T}$. Algorithm 6.5 follows.

Algorithm 6.2: Tree filling algorithm for the codomain dual theta-null point computation.

Data: Theta-coordinates $\theta_i^{\mathcal{L}}$ of 8-torsion points $T_1, \cdots, T_g$ such that $\ker(f) = [4]\langle T_1, \cdots, T_g\rangle$.
Result: Full computation tree $\mathcal{T}$ as described above.

```
1  for χ_0 ∈ (Z/2Z)^g do
2      Initialize T at root χ_0;
3      while all terminal nodes of T are not not marked as leaves do
4          for every terminal node χ of T not marked as a leaf do
5              leaf ⟵ True;
6              for l = 1 to g do
7                  if χ · χ_l ∉ T then
8                      D_{χ,l} ⟵ Σ_{t∈(Z/2Z)^g} χ(t) θ_t^L(T_l)^2;
9                      if D_{χ,l} ≠ 0 then
10                         D_{χ·χ_l,l} ⟵ Σ_{t∈(Z/2Z)^g} χ_l(t)χ(t) θ_t^L(T_l)^2;
11                         Add χ · χ_l as the child of χ in T and store (D_{χ·χ_l,l}, D_{χ,l}) on the edge
                            from χ to χ · χ_l;
12                         leaf ⟵ False;
13                     end
14                 end
15             end
16             if leaf then
17                 Mark χ as a leaf;
18             end
19         end
20     end
21     if T covers (Z/2Z)^g then
22         return T;
23     end
24 end
25 return False;
```

Lemma 6.1.6. *Algorithm 6.5 uses $6 \cdot 2^g - 9$ multiplications and at most $g2^g(4^g - 2^g + 1)$ squarings. When $U_\chi^{\mathcal{M}}(0_B) \neq 0$ for all $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$, Algorithm 6.5 uses as many multiplications than in the general case but only $g4^g$ squarings.*

Proof. Let $\mathcal{T}$ be a (possibly incomplete) tree obtained at the end of the main loop of Algorithm 6.2 iterating on the root $\chi_0 \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$. Then to fill in $\mathcal{T}$, we had to compute $D_{\chi,l}$ for every (temporary) terminal node χ and every $l \in [\![1\ ;\ g]\!]$ such that $\chi \cdot \chi_l \notin \mathcal{T}$. This amounts to at most $g-1$ computations of $D_{\chi,l}$ if $\chi \neq \chi_0$ and g such computations otherwise. Each of these computations costs 2^g squarings so we have at most $2^g(\#\mathcal{T}(g-1)+1)$ squarings for the $D_{\chi,l}$ computations on Line 8 in total. The variable $D_{\chi\cdot\chi_l,l}$ on Line 10 is computed every time a tree node is found so $\#\mathcal{T} - 1$ times in total. Hence, Line 10 costs $2^g(\#\mathcal{T} - 1)$ squarings in total. On the whole, a main loop of Algorithm 6.2 costs at most $2^g g\#\mathcal{T}$ squarings. Besides, we have $\#\mathcal{T} < 2^g$ at the end of each iteration, except when $\mathcal{T}$ is full and the algorithm terminates. Hence, Algorithm 6.2 costs at most in total:

$$2^g g(2^g - 1)(2^g - 1) + 2^g g 2^g = 2^g g(4^g - 2^g + 1)$$

squarings. When $U_\chi^{\mathcal{M}}(0_B) \neq 0$ for all $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$, Algorithm 6.2 terminates at the first iteration by Lemma 6.1.5, so it costs at most $g4^g$ squarings in that case.

The call to Algorithm 6.4 in Algorithm 6.5 costs 2 multiplications per edge of the computation tree $\mathcal{T}$ returned by Algorithm 6.2, so $2(2^g - 1)$ multiplications in total. The projective batch inversion algorithm (Algorithm 6.3) costs $3(n-1)$ multiplications, where $n = 2^g - 1$ is the number of edges in $\mathcal{T}$. Finally, Line 7 of Algorithm 6.5 costs one multiplication per character $\chi \neq \chi_0$, so $2^g - 1$ multiplications in total. Hence, Algorithm 6.4 costs $6 \cdot 2^g - 9$ multiplications along with the number of squarings computed above. This completes the proof. □

Algorithm 6.3: Projective batch inversion.

Data: $a_1, \cdots, a_n \in k^*$.
Result: Projective inverses $a_2 \cdots a_n$, $a_1 a_3 \cdots a_n$, ..., $a_1 \cdots a_{n-1}$ and projective factor $a_1 \cdots a_n$.

$b_1 \longleftarrow a_0$ **for** $i = 2$ **to** n **do**
  $b_i \longleftarrow b_{i-1} \cdot a_i$; // $b_i = a_1 \cdots a_i$
end
$c_1 \longleftarrow 1$;
for $i = 2$ **to** n **do**
  $c_i \longleftarrow c_{i-1} \cdot a_{n-i+2}$; // $c_i = a_{n-i+2} \cdots a_n$
end
$d_1 \longleftarrow c_n$;
for $i = 2$ **to** n **do**
  $d_i \longleftarrow b_{i-1} \cdot c_{n-i+1}$; // $d_i = (a_1 \cdots a_{i-1}) \cdot (a_{i+1} \cdots a_n)$
end
return $d_1, \cdots, d_n, b_n$; // Total cost: $3(n-1)\mathbf{M}$

Algorithm 6.4: Recursive tree edge multiplication.

Data: A computation tree $\mathcal{T}$ (possibly non-full) with couples $(a, b) \in k^2$ on each edge and a couple $(\lambda, \mu) \in k^2$.
Result: A computation tree $\mathcal{T}'$ with the same structure as $\mathcal{T}$ and couples (a', b') on each edge made of products $a' = \lambda \prod a_i$ and $b' = \mu \prod b_i$ where the (a_i, b_i) lie on the parent edges of $\mathcal{T}$.

Let χ_0 be the root of $\mathcal{T}$;
for *every child* χ *of* χ_0 **do**
  Let (a, b) be the couple stored on the edge from χ_0 to χ;
  $(a, b) \longleftarrow (a\lambda, b\mu)$;
  Recurse on subtree $\mathcal{T}_\chi$ with root χ and initial couple value (a, b);
end
return $\mathcal{T}$;

Remark 6.1.7. Once the codomain dual theta null point $(U_\chi^{\mathcal{M}}(0_B))_\chi$ has been computed, we can apply Algorithm 6.3 directly to obtain its projective inverse $(1/U_\chi^{\mathcal{M}}(0_B))_\chi$ with $3(2^g - 1)$ multiplications. This precomputed result $(1/U_\chi^{\mathcal{M}}(0_B))_\chi$ can then be used as an entry of the evaluation algorithm (Algorithm 6.1).

6.1.3 The gluing case

Evaluation of a gluing isogeny

Let us keep all the notations from Section 6.1.2 and the assumptions from Theorem 6.1.1. If $f : A_1 \times A_2 \longrightarrow B$ is a gluing isogeny in the sense of Definition 6.1.8 below, we do not expect the evaluation algorithm (Algorithm 6.1) to work as expected because dual theta constants $U_\chi^{\mathcal{M}}(0_B)$ may vanish.

Definition 6.1.8. A *gluing isogeny* is a an isogeny $A_1 \times A_2 \longrightarrow B$, where A_1, A_2 and B are principally polarised abelian varieties and $A_1 \times A_2$ is considered with the product polarisation. Similarly, a *splitting isogeny* is an isogeny $B \longrightarrow A_1 \times A_2$.

If $f : A_1 \times A_2 \longrightarrow B$ is a gluing isogeny, we may not divide by $U_\chi^{\mathcal{M}}(0_B)$ in Eq. (6.1) as in Algorithm 6.1 but we may "twist" this equation to divide by a non-zero theta constant in order to obtain the coordinate $U_\chi^{\mathcal{M}}(f(x))$, $x \in A(k)$ being the point we want to evaluate. This twisting operation is given by translating x by 4-torsion points above $\ker(f)$.

Lemma 6.1.9. *For all* $l \in [\![1 \ ; \ g]\!]$ *and* $x \in A(k)$*, there exists* $\lambda_l \in k^*$ *such that for all* $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$*,*

Algorithm 6.5: Codomain dual theta-null point computation algorithm.

Data: Theta-coordinates $\theta_i^{\mathcal{L}}$ of 8-torsion points $T_1, \cdots, T_g$ such that $\ker(f) = [4]\langle T_1, \cdots, T_g\rangle$.
Result: The projective codomain dual theta-null point $(U_\chi^{\mathcal{M}}(0_B))_{\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g}$.

1 Call Algorithm 6.2 with entries $T_1, \cdots, T_g$ to get a computation tree $\mathcal{T}$;
2 Call Algorithm 6.4 on $\mathcal{T}$ and initial couple $(1, 1)$ to multiply the numerators and denominators on every edge recursively, going down the tree from its root;
3 Call Algorithm 6.3 all denominators D_χ lying on the edges of $\mathcal{T}$ to obtain projective inverses λ/D_χ and the projective factor λ;
4 Let χ_0 be the root of $\mathcal{T}$. Set $U_{\chi_0}^{\mathcal{M}}(0_B) \longleftarrow \lambda$;
5 **for** *every* $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g \setminus \{\chi_0\}$ **do**
6 Let (N_χ, D_χ) lying on the parent edge of χ;
7 $U_\chi^{\mathcal{M}}(0_B) \longleftarrow N_\chi \cdot \lambda/D_\chi$;
8 **end**
9 **return** $(U_{\chi,0}^{\mathcal{M}}(0_B))_{\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g}$;

we have:

$$U_\chi^{\mathcal{M}}(f(x)) \cdot U_{\chi\cdot\chi_l}^{\mathcal{M}}(0_B) = \lambda_l \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\chi_l(t)\theta_t^{\mathcal{L}}(x + [2]T_l)^2. \tag{6.4}$$

We also have for all $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$,

$$\left(\sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\chi_l(t)\theta_t^{\mathcal{L}}(x)^2\right) U_\chi^{\mathcal{M}}(0_B) = \lambda_l \left(\sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_t^{\mathcal{L}}(x + [2]T_l)^2\right) U_{\chi\cdot\chi_l}^{\mathcal{M}}(0_B). \tag{6.5}$$

Proof. Let $l \in [\![1\ ;\ g]\!]$ and $x \in A(k)$. Then by Corollary 6.1.3, there exists $\lambda_l \in k^*$ such that for all $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$,

$$U_\chi^{\mathcal{M}}(f(x + [2]T_l)) \cdot U_\chi^{\mathcal{M}}(0_B) = \lambda_l \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_t^{\mathcal{L}}(x + [2]T_l)^2. \tag{6.6}$$

Besides, by Theorem 6.1.1, we have $\overline{\Theta}_{\mathcal{M}}(0, \chi_l) = [2]f(T_l)$, so by Eq. (5.11), we have for all $i \in (\mathbb{Z}/2\mathbb{Z})^g$,

$$\theta_i^{\mathcal{M}}(f(x + [2]T_l)) = \theta_i^{\mathcal{M}}(f(x) + [2]f(T_l)) = \chi_l(i)^{-1}\theta_i^{\mathcal{M}}(x) = \chi_l(i)\theta_i^{\mathcal{M}}(x),$$

so that for all $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$,

$$\begin{aligned} U_\chi^{\mathcal{M}}(f(x + [2]T_l)) = \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_t^{\mathcal{M}}(f(x + [2]T_l)) = \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\chi_l(t)\theta_t^{\mathcal{M}}(f(x)) \\ = U_{\chi\cdot\chi_l}^{\mathcal{M}}(f(x)). \end{aligned} \tag{6.7}$$

It follows that for all $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$,

$$U_{\chi\cdot\chi_l}^{\mathcal{M}}(f(x)) \cdot U_\chi^{\mathcal{M}}(0_B) = \lambda_l \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_t^{\mathcal{L}}(x + [2]T_l)^2,$$

Changing χ into $\chi \cdot \chi_l$, we obtain Eq. (6.4).

Let $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$. Then applying Corollary 6.1.3 with index $\chi \cdot \chi_l$ and projective factor set to $\lambda = 1$ and combining with Eq. (6.7) and Eq. (6.6), we finally obtain:

$$\begin{aligned} \left(\sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\chi_l(t)\theta_t^{\mathcal{L}}(x)^2\right) U_\chi^{\mathcal{M}}(0_B) &= U_{\chi\cdot\chi_l}^{\mathcal{M}}(f(x)) \cdot U_{\chi\cdot\chi_l}^{\mathcal{M}}(0_B) \cdot U_\chi^{\mathcal{M}}(0_B) \\ &= U_\chi^{\mathcal{M}}(f(x + [2]T_l)) \cdot U_{\chi\cdot\chi_l}^{\mathcal{M}}(0_B) \cdot U_\chi^{\mathcal{M}}(0_B) \end{aligned}$$

$$= \lambda_l \left(\sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t) \theta_t^{\mathcal{L}}(x + [2]T_l)^2 \right) U^{\mathcal{M}}_{\chi \cdot \chi_l}(0_B)$$

This proves Eq. (6.5) and completes the proof. □

To compute $(U^{\mathcal{M}}_{\chi}(f(x)))_{\chi}$ given $(\theta_i^{\mathcal{L}}(x))_i$ when f is a gluing isogeny, we use Eq. (6.1) to obtain $U^{\mathcal{M}}_{\chi}(f(x))$ for all $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$ such that $U_\chi(0_B) \neq 0$. When $U_\chi(0_B) = 0$, we consider $l \in [\![1\ ;\ g]\!]$ such that $U_{\chi \cdot \chi_l}(0_B) \neq 0$ and compute $U^{\mathcal{M}}_{\chi}(f(x))$ via Eq. (6.4). We use Eq. (6.5) to compute the factor λ_l in Eq. (6.4). For each λ_l we need, we have to compute an inversion. The naive way to do it would be to apply the batch inversion algorithm (Algorithm 5.3).

Instead, we propose an inversion free method. Applying Algorithm 6.3, we obtain the $\lambda\lambda_l$ for some projective factor $\lambda \in k^*$. Then, for every $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$ such that $U_\chi(0_B) = 0$ and $U_{\chi \cdot \chi_l}(0_B) \neq 0$ we apply Eq. (6.4) with $\lambda\lambda_l$ instead of λ to obtain $\lambda U^{\mathcal{M}}_{\chi}(f(x))$. For all $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$ such that $U_\chi(0_B) \neq 0$, we apply Eq. (6.1) to obtain $U^{\mathcal{M}}_{\chi}(f(x))$ and multiply it by λ. Algorithm 6.6 follows.

Algorithm 6.6: Gluing isogeny evaluation algorithm.

Data: 4-torsion points $[2]T_1, \cdots, [2]T_g$ such that $\ker(f) = [4]\langle T_1, \cdots, T_g \rangle$, a subset of indices $L \subseteq [\![1\ ;\ g]\!]$, theta points $(\theta_i^{\mathcal{L}}(x))_i$ and $(\theta_i^{\mathcal{L}}(x + [2]T_l))_i$ for all $l \in L$, the codomain dual theta null point $(U^{\mathcal{M}}_{\chi}(0_B))_{\chi}$ and the inverses of the non-zero dual theta constants $1/U^{\mathcal{M}}_{\chi}(0_B)$ (up to a projective factor).

Result: $(\theta_i^{\mathcal{M}}(f(x)))_i$.

```
for χ ∈ (Z/2Z)^g hat do
    HS_χ ⟵ Σ_{t∈(Z/2Z)^g} χ(t) θ_t^L(x)^2;
    for l ∈ L do
        HS_{χ,l} ⟵ Σ_{t∈(Z/2Z)^g} χ(t) θ_t^L(x + [2]T_l)^2;
    end
end
for l ∈ L do
    Look for χ ∈ (Z/2Z)^g hat such that HS_{χ,l} ≠ 0 and U^M_{χ·χ_l}(0_B) ≠ 0;
    D_l ⟵ HS_{χ,l} · U^M_{χ·χ_l}(0_B);
    N_l ⟵ HS_{χ·χ_l} · U^M_χ(0_B);
end
Call Algorithm 6.3 on the D_l for all l ∈ L to obtain the λ/D_l for all l ∈ L and λ := ∏_{l∈L} D_l;
for l ∈ L do
    λλ_l ⟵ λ/D_l · N_l;
end
for χ ∈ (Z/2Z)^g hat do
    if U^M_χ(0_B) ≠ 0 then
        U^M_χ(f(x)) ⟵ λ · 1/U^M_χ(0_B) · HS_χ;
    else
        Find l ∈ L such that U^M_{χ·χ_l,0}(0_B) ≠ 0;
        U^M_χ(f(x)) ⟵ λλ_l · 1/U^M_{χ·χ_l}(0_B) · HS_{χ·χ_l,l};
    end
end
for i ∈ (Z/2Z)^g do
    θ_i^M(f(x)) ⟵ Σ_{χ∈(Z/2Z)^g hat} χ(i) U^M_χ(f(x))
end
return (θ_i^M(f(x)))_i ;        // Total cost: 2^g(#L + 1)S + (6#L − 3 + 2^{g+1})M
```

Remark 6.1.10. In practice, the translates $(\theta_i^{\mathcal{L}}(x + [2]T_l))_i$ for all $l \in L$ are not computed directly with the theta model using Algorithm 5.1 that requires to know $(\theta_i^{\mathcal{L}}(x - [2]T_l))_i$ and to divide by some (possibly zero) theta constants. For instance, if the domain is a product of elliptic curves, we may

use Weierstrass coordinates to perform additions and then translate the result into theta coordinates (using Proposition 5.3.47 for instance). This approach can be generalized to products of Jacobians of genus 2 curves.

We have no proof that any choice of L (even $L = [\![1\ ;\ g]\!]$) ensures the termination of Algorithm 6.6 *i.e.* that there are enough non-zero codomain dual theta constants for a suitable character $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$ to exist on Line 8. However, from our implementation in dimension 4 (that will be presented in Section 6.6) we were able to formulate the following conjecture:

Conjecture 6.1.11. *If f is a 4-dimensional gluing 2-isogeny $A_1 \times A_2 \longrightarrow B$, where $A_1 \times A_2$ is a product of principally polarised abelian surfaces that are not isomorphic to elliptic products (as polarised abelian varieties) and B is a 4-dimensional principally polarised abelian variety that is not isomorphic to a product of smaller dimensional abelian varieties (as a polarised abelian variety), then Algorithm 6.6 always terminates with any input set L such that $\#L \geq 2$.*

Algorithm 6.6 serves its purpose but is much more costly than the generic evaluation algorithm (Algorithm 6.1). Not only do we have to compute the translates $x + T_l$ for $l \in [\![1\ ;\ g]\!]$ prior to its application, but also to perform an amount of additional operations compared to Algorithm 6.1. This overhead is proportional to $\#L$. In the following, we propose an alternate algorithm due to Max Duparc and implemented in dimension 2 for the round 2 SQIsign submission [AAA+25, Algorithm 8.40]. This algorithm can be generalised to any dimension and its cost is close to Algorithm 6.1.

Lemma 6.1.12. *Let $x, T \in A(k)$. Then there exists $\lambda \in k^*$ such that for all $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$,*

$$U_\chi^{\mathcal{M}}(f(x)) \cdot U_\chi^{\mathcal{M}}(f(T)) = \lambda \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_t^{\mathcal{L}}(x+T)\theta_t^{\mathcal{L}}(x-T).$$

Proof. The proof is very similar to Corollary 6.1.3. Let $x, T \in A(k)$. Then Eq. (5.18) immediately ensures the existence of $\lambda \in k^*$ such that:

$$U_\chi^{\mathcal{L}^2}(x) \cdot U_\chi^{\mathcal{L}^2}(T) = \lambda \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_t^{\mathcal{L}}(x+T)\theta_t^{\mathcal{L}}(x-T),$$

where $U_\chi^{\mathcal{L}^2}(x) = U_\chi^{\mathcal{M}}(f(x))$ and $U_\chi^{\mathcal{L}^2}(T) = U_\chi^{\mathcal{M}}(f(T))$. The result follows. □

Consider T_1 the first point of the 8-torsion basis $\mathscr{B}$ lying above $\ker(f)$. Unlike the dual theta constants $U_\chi^{\mathcal{M}}(0_B)$, the dual theta coordinates $U_\chi^{\mathcal{M}}(f(T_1))$ of the 4-torsion point $f(T_1)$ do not vanish in practice, at least when f is a 2-dimensional gluing whose codomain is not a product or under the assumptions of Conjecture 6.1.11 in dimension 4. We may then apply Lemma 6.1.12 to $x \in A(k)$ and T_1 to obtain $(U_\chi^{\mathcal{M}}(f(x)))_\chi$ directly. This requires to have precomputed the projective inverse $(1/U_\chi^{\mathcal{M}}(f(T_1)))_\chi$ in the first place along with theta coordinates of $x + T_1$ and $x - T_1$. Algorithm 6.7 follows.

Algorithm 6.7: Generic 2-isogeny evaluation with level 2 theta coordinates.

Data: Theta coordinates $(\theta_i^{\mathcal{L}}(x))_i$, $(\theta_i^{\mathcal{L}}(x - T_1))_i$, $(\theta_i^{\mathcal{L}}(x + T_1))_i$ and inverse of the codomain dual theta coordinates $(1/U_\chi^{\mathcal{M}}(f(T_1)))_\chi$.
Result: The theta coordinates $(\theta_i^{\mathcal{M}}(f(x)))_i$.

1 **for** $\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g$ **do**
2 $\quad U_\chi^{\mathcal{M}}(f(x)) \longleftarrow 1/U_\chi^{\mathcal{M}}(f(T_1)) \cdot \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_t^{\mathcal{L}}(x+T_1) \cdot \theta_t^{\mathcal{L}}(x-T_1)$;
3 **end**
4 **for** $i \in (\mathbb{Z}/2\mathbb{Z})^g$ **do**
5 $\quad \theta_i^{\mathcal{M}}(f(x)) \longleftarrow \sum_{\chi \in (\widehat{\mathbb{Z}/2\mathbb{Z}})^g} \chi(i) U_\chi^{\mathcal{M}}(f(x))$;
6 **end**
7 **return** $(\theta_i^{\mathcal{M}}(f(x)))_i$; // Total cost: $2^{g+1}\mathbf{M}$

Remark 6.1.13. $(U_\chi^{\mathcal{M}}(f(T_1)))_\chi$ (and its projective inverse) can be precomputed using Algorithm 6.6, but we may save some computation time if we take advantage of a symmetry property of this dual theta

point. Indeed, by Eq. (6.3), we have, for all $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$, $U^{\mathcal{M}}_{\chi}(f(T_1)) = U^{\mathcal{M}}_{\chi\cdot\chi_1}(f(T_1))$. Combining this symmetry with Eq. (6.1), we may be able to evaluate $(U^{\mathcal{M}}_{\chi}(f(T_1)))_{\chi}$ without using Algorithm 6.6, provided that for all $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$, we either have $U^{\mathcal{M}}_{\chi}(0_B) \neq 0$ or $U^{\mathcal{M}}_{\chi\cdot\chi_1}(0_B) \neq 0$. This idea is already used to compute 2-dimensional isogenies in the SQIsign NIST round 2 implementation [AAA+25, Algorithm 8.39].

Codomain dual theta null point of a gluing isogeny

When we compute the dual theta null point $(U^{\mathcal{M}}_{\chi}(0_B))_{\chi}$ of the codomain of a gluing isogeny, the tree filling algorithm (Algorithm 6.2) may not have enough information to find a full computation tree. Namely, recalling the notation from Section 6.1.2:

$$D_{\chi,l} := \sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta^{\mathcal{L}}_t(T_l)^2,$$

there might be too many zero denominators $D_{\chi,l}$ for Algorithm 6.2 to return the desired result. To circumvent this difficulty, we generalise Lemma 6.1.4 to sums of T_l to add more possible edges to explore to fill in the computation tree.

Lemma 6.1.14. *For all multi-index $i \in (\mathbb{Z}/2\mathbb{Z})^g$, we denote $T_i := \sum_{l=1}^g [i_l]T_l$ and recall that $\chi^i := \prod_{l=1}^g \chi_l^{i_k}$. Then, for all $i \in (\mathbb{Z}/2\mathbb{Z})^g$ and $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$, we have:*

$$U^{\mathcal{M}}_{\chi\cdot\chi^i}(0_B)\left(\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta^{\mathcal{L}}_t(T_i)^2\right) = U^{\mathcal{M}}_{\chi}(0_B)\left(\sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\chi^i(t)\theta^{\mathcal{L}}_t(T_i)^2\right).$$

Proof. This is very similar to the proof of Lemma 6.1.4. □

Now, using Lemma 6.1.14, we are able to consider edges between characters χ and $\chi\cdot\chi^i$ for all $i \in (\mathbb{Z}/2\mathbb{Z})^g$ such that T_i has been given on entry of the tree filling algorithm (Algorithm 6.2). With this additional freedom, we can adapt Algorithms 6.2 and 6.5 very easily (including sums of points T_l on entry). With experimental results (see Section 6.6 in particular), we are able to state the following conjecture.

Conjecture 6.1.15. *(i) If f is not a gluing isogeny, the codomain dual theta constants $U^{\mathcal{M}}_{\chi}(0_B)$ never vanish and Algorithm 6.2 always returns a full computation tree (in one iteration by Lemma 6.1.5). Hence, Algorithm 6.5 terminates and computes a correct codomain dual theta null point $(U^{\mathcal{M}}_{\chi}(0_B))_{\chi}$ and the standard evaluation algorithm (Algorithm 6.1) can be used.*

(ii) If f is a 4-dimensional gluing 2-isogeny $A_1 \times A_2 \longrightarrow B$, where $A_1 \times A_2$ is a product of principally polarised abelian surfaces that are not isomorphic to elliptic products (as polarised abelian varieties) and B is a 4-dimensional abelian variety that is not isomorphic to a product of smaller dimensional abelian varieties (as a polarised abelian variety), then Algorithm 6.2 always returns a full computation tree so Algorithm 6.5 terminates and computes a correct codomain dual theta null point when 5 entry 8-torsion points $T_1, \cdots, T_4$ and $T_1 + T_2$ lying above $\ker(f)$ are given.

6.2 Change of theta coordinates

In order to compute a 2-isogeny $f : A \longrightarrow B$ in the last section, we made a choice of level 2 theta structure $\Theta_{\mathcal{L}}$ on the domain $(A, \mathcal{L})$ adapted to f that simplified the isogeny computation. In particular, we imposed the condition $K_2(\overline{\Theta}_{\mathcal{L}}) = \ker(f)$ in Theorem 6.1.1. Unfortunately, the domain theta structure $\Theta_{\mathcal{L}}$ obtained at the start (*e.g.* the product theta structure on an elliptic curve product) does not always satisfy this condition. For that reason, we need to find a new theta structure $\Theta'_{\mathcal{L}}$ adapted to f, satisfying $K_2(\overline{\Theta'}_{\mathcal{L}}) = \ker(f)$ in particular, and compute the change of theta coordinates from the former ones to the new ones. The goal of this section is to obtain explicit formulas to change theta coordinates.

In Section 6.2.1, we introduce some theory on the action of Heisenberg group automorphisms (relating different theta structures) on theta functions due to Faugère, Lubicz and Robert [FLR11, § 5.2-5.3] (see also [Rob10, § 3.5]). In Section 6.2.2, we follow the approach of [FLR11] and apply this theory to symmetric theta structures and obtain an explicit change of theta coordinates formula in that case (Theorem 6.2.10). This formula is the main contribution of this section. It was already known to Igusa [Igu72, Theorem V.2] and Cosset [Cos11, Proposition 3.1.24] but proved in the analytic setting of complex theta functions. Our proof only uses the algebraic setting of Mumford [Mum66] which is more suitable to isogeny computations over finite fields. Finally, in Section 6.2.3, we apply the change of theta coordinates formula in a very simple case to compute the polarised dual $\widetilde{f} : B \longrightarrow A$ of a 2-isogeny $f : A \longrightarrow B$.

6.2.1 Heisenberg group automorphisms

Let $(A, \mathcal{L})$ be a polarized abelian variety of type δ. Let $\mathrm{Aut}_{k^*}(\mathcal{H}(\delta))$ be the group of automorphisms of the Heisenberg group $\mathcal{H}(\delta)$ that are trivial on k^*.

Proposition 6.2.1. $\mathrm{Aut}_{k^*}(\mathcal{H}(\delta))$ *acts faithfully and transitively on the set of theta structures on* $(A, \mathcal{L})$ *by composition on the right.*

Proof. Let $\psi \in \mathrm{Aut}_{k^*}(\mathcal{H}(\delta))$. Then ψ induces a symplectic isomorphism $\overline{\psi} \in \mathrm{Aut}(K(\delta))$ such that the following diagram commutes:

$$\begin{array}{ccccccccc} 1 & \longrightarrow & k^* & \longrightarrow & \mathcal{H}(\delta) & \longrightarrow & K(\delta) & \longrightarrow & 0 \\ & & \| & & \downarrow{\scriptstyle \psi} & & \downarrow{\scriptstyle \overline{\psi}} & & \\ 1 & \longrightarrow & k^* & \longrightarrow & \mathcal{H}(\delta) & \longrightarrow & K(\delta) & \longrightarrow & 0 \end{array}$$

$\overline{\psi}$ is given by $\overline{\psi} := \rho_\delta \circ \psi \circ s_\delta$ where $\rho_\delta : \mathcal{H}(\delta) \longrightarrow K(\delta)$ and $s_\delta : K(\delta) \longrightarrow \mathcal{H}(\delta)$ are respectively the forgetful map and the canonical section. Now we justify that $\overline{\psi}$ is symplectic. Let $(i_1, \chi_1), (i_2, \chi_2) \in K(\delta)$. Then

$$\begin{aligned} e_\delta(\overline{\psi}(i_1, \chi_1), \overline{\psi}(i_2, \chi_2)) &= \psi(1, i_1, \chi_1)\psi(1, i_2, \chi_2)\psi(1, i_1, \chi_1)^{-1}\psi(1, i_2, \chi_2)^{-1} \\ &= \psi((1, i_1, \chi_1) \cdot (1, i_2, \chi_2))\psi((1, i_2, \chi_2) \cdot (1, i_1, \chi_1))^{-1} \\ &= \psi(\chi_2(i_1), i_1 + i_2, \chi_1\chi_2)\psi(\chi_1(i_2), i_1 + i_2, \chi_1\chi_2)^{-1} \\ &= \chi_2(i_1)\chi_1(i_2)^{-1}\psi(1, i_1 + i_2, \chi_1\chi_2)\psi(1, i_1 + i_2, \chi_1\chi_2)^{-1} \\ &= \chi_2(i_1)\chi_1(i_2)^{-1} = e_\delta((i_1, \chi_1), (i_2, \chi_2)), \end{aligned}$$

so $\overline{\psi}$ is symplectic.

Consequently, if $\Theta_\mathcal{L}$ is a theta structure, then $\Theta_\mathcal{L} \circ \psi$ is a theta structure so $\mathrm{Aut}_{k^*}(\mathcal{H}(\delta))$ acts on theta structures. Besides $\Theta_\mathcal{L} \circ \psi = \Theta_\mathcal{L}$ if and only if $\psi = \mathrm{id}_{\mathcal{H}(\delta)}$, so the action is faithful. Finally, if $\Theta_\mathcal{L}$ and $\Theta'_\mathcal{L}$ are theta structures, then $\Theta_\mathcal{L} \circ \Theta'^{-1}_\mathcal{L} \in \mathrm{Aut}_{k^*}(\mathcal{H}(\delta))$ so the action is transitive. □

Proposition 6.2.2. *Let* $\mathrm{Sp}(K(\delta))$ *be the group of symplectic morphisms of* $K(\delta)$*. Then, we have an exact sequence*

$$0 \longrightarrow K(\delta) \longrightarrow \mathrm{Aut}_{k^*}(\mathcal{H}(\delta)) \longrightarrow \mathrm{Sp}(K(\delta)) \longrightarrow 1,$$

where $\mathrm{Aut}_{k^*}(\mathcal{H}(\delta)) \longrightarrow \mathrm{Sp}(K(\delta))$ *is the forgetful map* $\psi \longmapsto \overline{\psi}$*.*

Proof. Let $\overline{\psi} \in \mathrm{Sp}(K(\delta))$. Then, $\overline{\psi}(K_i(\delta))$ for $i \in \{1, 2\}$ are (maximal) isotropic subgroups of $K(\delta)$. Hence, as we saw in the proof of Proposition Proposition 5.1.12 ((ii) $\Longrightarrow$ (i)), we can lift the $\overline{\psi}(K_i(\delta))$ to level subgroups of $\mathcal{H}(\delta)$, *i.e.* find sections $s_i : \overline{\psi}(K_i(\delta)) \longrightarrow \mathcal{H}(\delta)$ for $i \in \{1, 2\}$. We then define $\psi \in \mathrm{Aut}_{k^*}(\mathcal{H}(\delta))$ as follows:

$$\forall (\alpha, i, \chi) \in \mathcal{H}(\delta), \quad \psi(\alpha, i, \chi) := \alpha \cdot s_1(\overline{\psi}(i, 1)) \cdot s_2(\overline{\psi}(0, \chi)).$$

By construction, ψ reduces to $\overline{\psi}$, so the forgetful map $\mathrm{Aut}_{k^*}(\mathcal{H}(\delta)) \longrightarrow \mathrm{Sp}(K(\delta))$ is indeed surjective.

Now, let $\psi \in \mathrm{Aut}_{k^*}(\mathcal{H}(\delta))$ mapping to the identity in $\mathrm{Sp}(K(\delta))$. Then, there exists a character $\chi : K(\delta) \longrightarrow k^*$ such that

$$\forall (\alpha, y) \in \mathcal{H}(\delta), \quad \psi(\alpha, y) = (\alpha\chi(y), y).$$

Since e_δ is non-degenerate, there exists $c \in K(\delta)$ such that $\chi(y) = e_\delta(c, y)$ for all $y \in K(\delta)$. Hence, ψ is in the image of

$$\begin{array}{rcl} K(\delta) & \longrightarrow & \mathrm{Aut}_{k^*}(\mathcal{H}(\delta)) \\ c := (c_1, c_2) & \longmapsto & [\psi_c : (\alpha, i, \chi) \longmapsto (\alpha e_\delta(c, (i, \chi)), i, \chi) = (\alpha\chi(c_1)c_2(i)^{-1}, i, \chi)] \end{array} \tag{6.8}$$

and conversely, any morphism of the form ψ_c for some $c \in K(\delta)$ is trivial in $\mathrm{Sp}(K(\delta))$. Finally, $\psi_c = \mathrm{id}_{\mathcal{H}(\delta)}$ if and only if $c = 0$. This completes the proof. □

Remark 6.2.3. If $\overline{\psi} \in \mathrm{Sp}(K(\delta))$, then we can write a lift $\psi \in \mathrm{Aut}_{k^*}(\mathcal{H}(\delta))$ explicitly, as $\psi(\alpha, i, \chi) := (\alpha\xi(i, \chi), \overline{\psi}(i, \chi))$ for all $(\alpha, i, \chi) \in \mathcal{H}(\delta)$, where $\xi : K(\delta) \longrightarrow k^*$ is a *semi-character*, satisfying the following property:

$$\forall (i_1, \chi_1), (i_2, \chi_2) \in K(\delta), \quad \xi(i_1 + i_2, \chi_1 \cdot \chi_2) = \frac{\xi(i_1, \chi_2)\xi(i_2, \chi_2)\overline{\psi}_2(i_2, \chi_2)(\overline{\psi}_1(i_1, \chi_1))}{\chi_2(i_1)}. \tag{6.9}$$

In the following, we study the action of $\mathrm{Aut}_{k^*}(\mathcal{H}(\delta))$ on the canonical basis of theta functions $(\theta_i)_{i \in K_1(\delta)}$. Let $\Theta_{\mathcal{L}}$ be a theta structure, $\psi \in \mathrm{Aut}_{k^*}(\mathcal{H}(\delta))$ and $\Theta'_{\mathcal{L}} := \Theta_{\mathcal{L}} \circ \psi$ another theta structure. As we saw in Section 5.1.5, the basis of theta functions $(\theta_i)_{i \in K_1(\delta)}$ associated to $\Theta_{\mathcal{L}}$ is defined by $\theta_i := \beta(\delta_i)$, where the δ_i are the Kronecker delta functions for $i \in K_1(\delta)$ and $\beta : V(\delta) \xrightarrow{\sim} \Gamma(A, \mathcal{L})$ is the unique (up to scalar multiplication) isomorphism respecting the group action of $\mathcal{H}(\delta)$ and $G(\mathcal{L})$:

$$\forall h \in \mathcal{H}(\delta), v \in V(\delta), \quad \beta(h \cdot v) = \Theta_{\mathcal{L}}(h) \cdot \beta(v).$$

Similarly, we have a basis of theta functions $(\theta'_i)_{i \in K_1(\delta)}$ associated to $\Theta'_{\mathcal{L}}$, defined by $\theta'_i := \beta'(\delta_i)$ with $\beta' : V(\delta) \xrightarrow{\sim} \Gamma(A, \mathcal{L})$ such that:

$$\forall h \in \mathcal{H}(\delta), v \in V(\delta), \quad \beta'(h \cdot v) = \Theta'_{\mathcal{L}}(h) \cdot \beta'(v) = \Theta_{\mathcal{L}}(h) \circ \psi(h) \cdot \beta'(v).$$

We want to relate the θ_i to the θ'_i. By the definition of the representation $V(\delta)$ of $\mathcal{H}(\delta)$, we have for all $i \in K_1(\delta)$ and $j \in K_2(\delta)$:

$$\delta_i = (1, i, 1) \cdot \delta_0 \quad \text{and} \quad (1, 0, j) \cdot \delta_0 = \delta_0 \tag{6.10}$$

It follows that the θ'_i are fully determined by θ'_0 and that θ'_0 is invariant under the action of the maximal level subgroup $\widetilde{K}_2(\Theta'_{\mathcal{L}})$. By Corollary 5.1.26, a section invariant under the action of a maximal level subgroup is unique up to scalar multiplication. So we only need to find one section stabilized by $\widetilde{K}_2(\Theta'_{\mathcal{L}})$ to determine θ'_0 and then $(\theta'_i)_{i \in K_1(\delta)}$.

Proposition 6.2.4. *(i) There exists $i \in K_1(\delta)$ and $\lambda \in k^*$ such that:*

$$\theta'_0 = \lambda \sum_{j \in K_2(\delta)} \Theta_{\mathcal{L}} \circ \psi(1, 0, j) \cdot \theta_i.$$

(ii) When $\psi = \psi_c : (\alpha, i, \chi) \longmapsto (\alpha\chi(c_1)c_2(i)^{-1}, i, \chi)$, for some $c := (c_1, c_2) \in K(\delta)$ (ψ is trivial in $\mathrm{Sp}(K(\delta))$), we have $\theta'_0 = \lambda\theta_{c_1}$ and $\theta'_i = \lambda c_2(i)^{-1}\theta_{i+c_1}$ for all $i \in K_1(\delta)$ and for some $\lambda \in k^*$.

(iii) When $\psi(s_\delta(K_2(\delta))) = 1 \times K_1 \times K_2$ with $K_i \subset K_i(\delta)$ for $i \in \{1, 2\}$, then we have:

$$\theta'_0 = \lambda \sum_{i \in K_1} \theta_i,$$

with $\lambda \in k^*$.

Proof. **(i)** For all $i \in K_1(\delta)$, $T_i := \sum_{\chi \in K_2(\delta)} \Theta_{\mathcal{L}} \circ \psi(1,0,\chi) \cdot \theta_i$ is invariant under the action of $\widetilde{K}_2(\Theta'_{\mathcal{L}})$ by construction. Now we prove that at least one of the T_i is non-zero (so that $T_i = \theta'_0$ up to scalar multiplication). We write $\psi(1,0,\chi) := (\alpha(\chi), \sigma(\chi), \tau(\chi))$ for all $\chi \in K_2(\delta)$. Then

$$\forall i \in K_1(\delta), \chi \in K_2(\delta), \quad \psi(1,0,\chi) \cdot \delta_i = \alpha(\chi)\tau(\chi)(i + \sigma(\chi))^{-1} \delta_{i+\sigma(\chi)},$$

with

$$\tau(\chi)(\sigma(\chi)) = e_\delta((\sigma(\chi), 1), (0, \tau(\chi))) = 1.$$

Indeed, $(\sigma(\chi), 1), (0, \tau(\chi)) \in \overline{\psi}(K_2(\delta))$ since $\overline{\psi} = (\sigma, \tau)$ is an isomorphism. And $\overline{\psi}$ is symplectic so $\overline{\psi}(K_2(\delta))$ is isotropic.

It follows that for all $i \in K_1(\delta)$,

$$\begin{aligned} T_i &= \beta \left(\sum_{\chi \in K_2(\delta)} \psi(1,0,\chi) \cdot \delta_i \right) = \beta \left(\sum_{\chi \in K_2(\delta)} \alpha(\chi)\tau(\chi)(i)^{-1} \delta_{i+\sigma(\chi)} \right) \\ &= \beta \left[\sum_{l \in \mathrm{im}(\sigma)} \left(\sum_{\chi \in \sigma^{-1}(\{l\})} \alpha(\chi)\tau(\chi)(i)^{-1} \right) \delta_{i+l} \right] = \sum_{l \in \mathrm{im}(\sigma)} \left(\sum_{\chi \in \sigma^{-1}(\{l\})} \alpha(\chi)\tau(\chi)(i)^{-1} \right) \theta_{i+l} \end{aligned}$$

Hence, the θ_i forming a basis, it suffices to find $i \in K_1(\delta)$ such that:

$$\sum_{\chi \in \ker(\sigma)} \alpha(\chi)\tau(\chi)(i)^{-1} \neq 0$$

to get $T_i \neq 0$. But $\ker(\sigma) \cap \ker(\tau) = \{0\}$ since $\overline{\psi} = (\sigma, \tau)$ is an isomorphism so that $\#\,\mathrm{im}(\sigma) \cdot \#\,\mathrm{im}(\tau) = \#K_2(\delta)$ *i.e.* $\#\,\mathrm{im}(\tau) = \#K_2(\delta)/\#\,\mathrm{im}(\sigma) = \#\ker(\sigma)$ and τ induces an isomorphism $\ker(\sigma) \xrightarrow{\sim} \mathrm{im}(\tau)$. Hence, the $\tau(\chi)^{-1}$ for $\chi \in \ker(\sigma)$ are distinct characters so they are linearly independent and the result follows (the $\alpha(\chi)$ being non-zero for all $\chi \in \ker(\sigma)$).

(ii) We have, for all $\chi \in K_2(\delta)$,

$$\psi_c(1,0,\chi) \cdot \delta_{c_1} = (\chi(c_1), 0, \chi) \cdot \delta_{c_1} = \chi(c_1)\chi(c_1)^{-1} \delta_{c_1} = \delta_{c_1},$$

so θ_{c_1} is $\widetilde{K}_2(\Theta'_{\mathcal{L}})$-invariant and $\theta'_0 = \lambda \theta_{c_1}$ for some $\lambda \in k^*$. We then have

$$\theta'_i = \lambda\beta(\psi_c(1,i,1) \cdot \delta_0) = \lambda\beta((c_2(i)^{-1}, i, 1) \cdot \delta_0) = \lambda\beta(c_2(i)^{-1}\delta_{i+c_1}) = \lambda c_2(i)^{-1}\theta_{i+c_1}.$$

(iii) Let $T := \sum_{i \in K_1} \theta_i$ and let us write $\psi(1,0,\chi) := (1, \sigma(\chi), \tau(\chi))$ for all $\chi \in K_2(\delta)$. Then, for all $\chi \in K_2(\delta)$,

$$\Theta'_{\mathcal{L}}(\delta)(1,0,\chi) \cdot T = \beta \left(\sum_{i \in K_1} \psi(1,0,\chi) \cdot \delta_i \right) = \beta \left(\sum_{i \in K_1} \tau(\chi)(i + \sigma(\chi))^{-1} \delta_{i+\sigma(\chi)} \right),$$

with, for all $i \in K_1$

$$\tau(\chi)(i + \sigma(\chi)) = e_\delta((i + \sigma(\chi), 1), (0, \tau(\chi))) = 1$$

since $(i + \sigma(\chi), 1), (0, \tau(\chi)) \in \overline{\psi}(K_2(\delta))$. It follows that

$$\Theta'_{\mathcal{L}}(\delta)(1,0,\chi) \cdot T = \beta \left(\sum_{i \in K_1} \delta_{i+\sigma(\chi)} \right) = \beta \left(\sum_{i \in K_1} \delta_i \right) = T,$$

so that $\theta'_0 = \lambda T$ for some $\lambda \in k^*$.

□

6.2.2 Action of automorphisms on symmetric and compatible theta structures

In this section, we adapt the results of the previous section to the case of symmetric theta structures. We introduce symmetric Heisenberg automorphisms and study their action on symmetric theta structures. Our goal, achieved in Theorem 6.2.10, is to derive an explicit change of coordinates formula from Proposition 6.2.4 depending on the symplectic basis associated to the underlying symmetric theta structures (following Remark 5.3.31).

Definition 6.2.5. An automorphism $\psi \in \mathrm{Aut}(\mathcal{H}(\delta))$ is *symmetric* if $\psi \circ D_{-1} = D_{-1} \circ \psi$, where D_{-1} has been defined in Definition 5.3.12. We denote $\mathrm{Aut}^0(\mathcal{H}(\delta))$ the subgroup of symmetric automorphisms.

Lemma 6.2.6. *Let $\psi \in \mathrm{Aut}(\mathcal{H}(\delta))$. As in Remark 6.2.3, we may write $\psi(\alpha, i, \chi) := (\alpha\xi(i,\chi), \overline{\psi}(i,\chi))$ for all $(\alpha, i, \chi) \in \mathcal{H}(\delta)$, where $\overline{\psi} \in \mathrm{Sp}(K(\delta))$ is the symplectic automorphism induced by ψ and $\xi : K(\delta) \longrightarrow k^*$ is a semi-character. Then, ψ is symmetric if and only if*

$$\forall (i,\chi) \in K(\delta), \quad \xi(i,\chi)^2 = \chi(i)^{-1}\overline{\psi}_2(i,\chi)(\overline{\psi}_1(i,\chi)).$$

Proof. We have, for all $(\alpha, i, \chi) \in \mathcal{H}(\delta)$,

$$\psi \circ D_{-1}(\alpha, i, \chi) = \psi\left(\frac{\alpha^2}{\chi(i)}(\alpha, i, \chi)^{-1}\right) = \frac{\alpha^2}{\chi(i)}\psi(\alpha, i, \chi)^{-1}$$

$$\text{and} \qquad D_{-1} \circ \psi(\alpha, i, \chi) = \frac{\alpha^2\xi(i,\chi)^2}{\overline{\psi}_2(i,\chi)(\overline{\psi}_1(i,\chi))}\psi(\alpha, i, \chi)^{-1}$$

The result follows. □

Lemma 6.2.7. *The exact sequence of Proposition 6.2.2 yields an exact sequence*

$$0 \longrightarrow K(\delta)[2] \longrightarrow \mathrm{Aut}^0(\mathcal{H}(\delta)) \longrightarrow \mathrm{Sp}(K(\delta)) \longrightarrow 1.$$

Proof. Let $\overline{\psi} \in \mathrm{Sp}(K(\delta))$ and $\psi \in \mathrm{Aut}(\mathcal{H}(\delta))$ be a lift of $\overline{\psi}$. As in Remark 6.2.3, we may write $\psi(\alpha, i, \chi) := (\alpha\xi(i,\chi), \overline{\psi}(i,\chi))$ for all $(\alpha, i, \chi) \in \mathcal{H}(\delta)$, where $\xi : K(\delta) \longrightarrow k^*$ is a semi-character. By Lemma 6.2.6, ψ is symmetric if and only if

$$\forall (i,\chi) \in K(\delta), \quad \xi(i,\chi)^2 = \chi(i)^{-1}\overline{\psi}_2(i,\chi)(\overline{\psi}_1(i,\chi)).$$

It is sufficient that:

$$\forall (i,\chi) \in K(\delta), \quad \xi(i,1)^2 = \overline{\psi}_2(i,1)(\overline{\psi}_1(i,1)) \quad \text{and} \quad \xi(0,\chi)^2 = \overline{\psi}_2(0,\chi)(\overline{\psi}_1(0,\chi)). \tag{6.11}$$

We can then extend ξ by the semi-character formula (Eq. (6.9)):

$$\forall (i,\chi) \in K(\delta), \quad \xi(i,\chi) = \frac{\xi(i,1)\xi(0,\chi)\overline{\psi}_2(0,\chi)(\overline{\psi}_1(i,1))}{\chi(i)}.$$

Indeed, this defines a semi-character by construction, and we then have for all $(i,\chi) \in K(\delta)$,

$$\begin{aligned}
\xi(i,\chi)^2 &= \frac{\xi(i,1)^2 \cdot \xi(0,\chi)^2 \cdot \overline{\psi}_2(0,\chi)(\overline{\psi}_1(i,1))^2}{\chi(i)^2} \\
&= \frac{\overline{\psi}_2(i,1)(\overline{\psi}_1(i,1)) \cdot \overline{\psi}_2(0,\chi)(\overline{\psi}_1(0,\chi)) \cdot \overline{\psi}_2(0,\chi)(\overline{\psi}_1(i,1))^2}{\chi(i)^2} \\
&= \frac{\overline{\psi}_2(i,\chi)(\overline{\psi}_1(i,1)) \cdot \overline{\psi}_2(0,\chi)(\overline{\psi}_1(i,\chi))}{\chi(i)^2}
\end{aligned}$$

and

$$\chi(i) = e_\delta((i,\chi),(0,\chi)) = e_\delta(\overline{\psi}(i,\chi), \overline{\psi}(0,\chi)) = \overline{\psi}_2(0,\chi)(\overline{\psi}_1(i,\chi)) \cdot \overline{\psi}_2(i,\chi)(\overline{\psi}_1(0,\chi))^{-1},$$

so that, as expected,

$$\xi(i,\chi)^2 = \chi(i)^{-1}\overline{\psi}_2(i,\chi)(\overline{\psi}_1(i,\chi)).$$

In order to satisfy Eq. (6.11), we define ξ accordingly on basis of $K_1(\delta)$ and $K_2(\delta)$ (by taking square roots of the $\overline{\psi}_2(i,1)(\overline{\psi}_1(i,1))$ and $\overline{\psi}_2(0,\chi)(\overline{\psi}_1(0,\chi))$). This defines a symmetric lift $\psi \in \mathrm{Aut}(\mathcal{H}(\delta))$.

Now, if $\psi \in \mathrm{Aut}^0(\mathcal{H}(\delta))$ maps to $\mathrm{id}_{K(\delta)}$, we can write $\psi = \psi_c$ for some $c \in K(\delta)$, where $\psi_c(\alpha, y) := (\alpha e_{\mathcal{L}}(c,y), y)$ for all $(\alpha, y) \in \mathcal{H}(\delta)$ by Eq. (6.8). We then get that ψ_c is symmetric if and only if $e_\delta(c,\cdot)$ takes values in $\{\pm 1\}$, *i.e.* $c \in K(\delta)[2]$. This completes the proof. □

Let $\mathcal{L}$ be a totally symmetric line bundle on an abelian variety A of type δ. Let $\Theta_{\mathcal{L}}$ and $\Theta_{\mathcal{L}^2}$ be compatible symmetric theta-structures on $G(\mathcal{L})$ and $G(\mathcal{L}^2)$ respectively. We now explain how the action of a symmetric automorphism $\psi \in \mathrm{Aut}^0(\mathcal{H}(2\delta))$ on $\Theta_{\mathcal{L}^2}$ transforms $\Theta_{\mathcal{L}}$ and the associated theta-functions.

Proposition 6.2.8. *Recall the definitions of $E_2 : \mathcal{H}(\delta) \longrightarrow \mathcal{H}(2\delta)$ and $H_2 : \mathcal{H}(2\delta) \longrightarrow \mathcal{H}(\delta)$ from Definition 5.3.25 and $\overline{H}_2 : K(2\delta) \longrightarrow K(\delta)$ from Eq. (5.13).*

(i) For all $\psi \in \mathrm{Aut}^0(\mathcal{H}(2\delta))$, there exists a unique $\psi' \in \mathrm{Aut}^0(\mathcal{H}(\delta))$ such that $\psi' \circ H_2 = H_2 \circ \psi$. Then ψ and ψ' automatically satisfy $\psi \circ E_2 = E_2 \circ \psi'$.

(ii) Let $\overline{\psi}, \overline{\psi}', \xi, \xi'$ be respectively the symplectic automorphisms and semi-characters associated to ψ and ψ'. Then, we have $\overline{\psi}' \circ \overline{H}_2 = \overline{H}_2 \circ \overline{\psi}$ and

$$\forall (i,\chi) \in K(2\delta), \quad \xi'(\bar{i},\overline{\chi}) = \chi(i)^{-1}\overline{\psi}_2(i,\chi)(\overline{\psi}_1(i,\chi)).$$

(iii) Let $(\Theta_{\mathcal{L}}, \Theta_{\mathcal{L}^2})$ and $(\Theta'_{\mathcal{L}}, \Theta'_{\mathcal{L}^2})$ be two pairs of compatible symmetric theta structures for $(\mathcal{L}, \mathcal{L}^2)$. Then, there exists $\psi \in \mathrm{Aut}^0(\mathcal{H}(2\delta))$ such that $\Theta'_{\mathcal{L}^2} = \Theta_{\mathcal{L}^2} \circ \psi$ and $\Theta'_{\mathcal{L}} = \Theta_{\mathcal{L}} \circ \psi'$, where $\psi' \in \mathrm{Aut}^0(\mathcal{H}(\delta))$ is induced by ψ.

Proof. **(i)** By Lemma 5.3.27, $\ker(H_2) = \mathcal{H}(2\delta)[2]$, so for all $h \in \mathcal{H}(2\delta)$, $H_2 \circ \psi(h) = 1$ if and only if $\psi(h)^2 = 1$ if and only if $h^2 = 1$ (since ψ is an automorphism). Hence, $H_2 \circ \psi$ factors through H_2 and this defines an automorphism $\psi' : \mathcal{H}(\delta) \xrightarrow{\sim} \mathcal{H}(\delta)$ such that $\psi' \circ H_2 = H_2 \circ \psi$. This automorphism ψ' is trivial on k^* because ψ is and H_2 act as $\lambda \longmapsto \lambda^2$. Besides, ψ' is symmetric by Lemma 5.3.26.(ii) and since H_2 is surjective, so $\psi' \in \mathrm{Aut}^0(\mathcal{H}(2\delta))$. The uniqueness is a consequence of the surjectivity of H_2.

We now prove that $\psi \circ E_2 = E_2 \circ \psi'$. By surjectivity of H_2, it suffices to prove that $\psi \circ E_2 \circ H_2 = E_2 \circ \psi' \circ H_2$ *i.e.* that $\psi \circ D_2 = D_2 \circ \psi$ since $\psi' \circ H_2 = H_2 \circ \psi$ and $E_2 \circ H_2 = D_2$ by Lemma 5.3.26.(iii). Let $h \in \mathcal{H}(2\delta)$. Then, $D_2(h) = h^3 D_{-1}(h)$ by Lemma 5.3.26.(iv) and:

$$\psi \circ D_2(h) = \psi(h^3 D_{-1}(h)) = \psi(h)^3 \cdot \psi \circ D_{-1}(h) = \psi(h)^3 \cdot D_{-1} \circ \psi(h) = D_2 \circ \psi(h)$$

The result follows.

(ii) Let $(\alpha, i, \chi) \in \mathcal{H}(2\delta)$. Then

$$H_2 \circ \psi(\alpha, i, \chi) = H_2(\alpha\xi(i,\chi), \overline{\psi}(i,\chi)) = (\alpha^2 \xi(i,\chi)^2, \overline{H}_2 \circ \overline{\psi}(i,\chi))$$

$$\text{and} \quad \psi' \circ H_2(\alpha, i, \chi) = \psi'(\alpha^2, \bar{i}, \overline{\chi}) = (\alpha^2 \xi'(\bar{i},\overline{\chi}), \overline{\psi'} \circ \overline{H}_2(i,\chi))$$

Hence $\overline{\psi}' \circ \overline{H}_2 = \overline{H}_2 \circ \overline{\psi}$ and by Lemma 6.2.6:

$$\xi'(\bar{i},\overline{\chi}) = \xi(i,\chi)^2 = \chi(i)^{-1}\overline{\psi}_2(i,\chi)(\overline{\psi}_1(i,\chi)).$$

The result follows.

(iii) Let $\psi := \Theta_{\mathcal{L}^2}^{-1} \circ \Theta'_{\mathcal{L}^2}$. Then, $\Theta_{\mathcal{L}^2}$ and $\Theta'_{\mathcal{L}^2}$ being symmetric, we have:

$$\psi \circ D_{-1} = \Theta_{\mathcal{L}^2}^{-1} \circ \Theta'_{\mathcal{L}^2} \circ D_{-1} = \Theta_{\mathcal{L}^2}^{-1} \circ \Delta_{-1} \circ \Theta'_{\mathcal{L}^2} = D_{-1} \circ \Theta_{\mathcal{L}^2}^{-1} \circ \Theta'_{\mathcal{L}^2} = D_{-1} \circ \psi,$$

so $\psi \in \mathrm{Aut}^0(\mathcal{H}(2\delta))$. Besides by compatibility of the pairs $(\Theta_{\mathcal{L}}, \Theta_{\mathcal{L}^2})$ and $(\Theta'_{\mathcal{L}}, \Theta'_{\mathcal{L}^2})$:

$$\Theta'_{\mathcal{L}} \circ H_2 = \eta_2 \circ \Theta'_{\mathcal{L}^2} = \eta_2 \circ \Theta_{\mathcal{L}^2} \circ \psi = \Theta_{\mathcal{L}} \circ H_2 \circ \psi = \Theta_{\mathcal{L}} \circ \psi' \circ H_2,$$

so that $\Theta'_{\mathcal{L}} = \Theta_{\mathcal{L}} \circ \psi'$ since H_2 is surjective. This completes the proof. □

Let $\delta := (d_1, \cdots, d_g)$ with $d_1 | \cdots | d_g$ and $\zeta \in k^*$ be a d_g-th primitive root of unity. Let us fix a canonical symplectic basis of $K(\delta)$ as follows. For $l \in [\![1\ ;\ g]\!]$, let e_l be the vector of $K_1(\delta)$ with 1 at index l and 0 everywhere else. For all $l \in [\![1\ ;\ g]\!]$, let $\chi_l \in K_2(\delta)$ be the character such that $\chi_l(e_m) = \zeta^{d_g/d_l \delta_{l,m}}$ for all $m \in [\![1\ ;\ g]\!]$. Then $K_1(\delta)$ can be canonically identified with $K_2(\delta)$ via the map $i \in K_1(\delta) \longmapsto \chi^i := \prod_{l=1}^g \chi_i^{i_l}$. We then have

$$\forall i, j \in K_1(\delta), \quad \chi^i(j) = \zeta^{\langle i|j\rangle} \quad \text{with} \quad \langle i|j\rangle := \sum_{l=1}^{g} \frac{d_g}{d_l} i_l j_l.$$

Such a basis is called a *canonical ζ-symplectic basis.*

Lemma 6.2.9. *Let $\sigma : K(\delta) \xrightarrow{\sim} K(\delta)$ be an automorphism of $K(\delta)$ and M be its matrix in the canonical ζ-symplectic basis $(e_1, \cdots, e_g, \chi_1, \cdots, \chi_g)$. Then σ is symplectic if and only if*

$$^tM \cdot J_\Delta \cdot M \equiv J_\Delta \mod d_g, \qquad \textit{where} \qquad J_\Delta := \begin{pmatrix} 0 & \Delta \\ -\Delta & 0 \end{pmatrix}$$

and $\Delta := \mathrm{Diag}(d_g/d_1, \cdots, d_g/d_{g-1}, 1)$.

If we write

$$M := \begin{pmatrix} A & C \\ B & D \end{pmatrix},$$

this is equivalent to $^tB\Delta A \equiv {}^tA\Delta B$, $^tD\Delta C \equiv {}^tC\Delta B$ and $^tA\Delta D - {}^tB\Delta C \equiv \Delta$ modulo d_g.

Proof. Let $l, m \in [\![1\ ;\ g]\!]$. Then

$$\begin{aligned} e_\delta(\sigma(e_l, 1), \sigma(e_m, 1)) &= e_\delta((Ae_l, \chi^{Be_l}), (Ae_m, \chi^{Be_m})) = \chi^{Be_m}(Ae_l)\chi^{-Be_l}(Ae_m) \\ &= \zeta^{\langle Ae_l|Be_m\rangle - \langle Be_l|Ae_m\rangle} = \zeta^{{}^te_l({}^tA\Delta B - {}^tB\Delta A)e_m} \end{aligned}$$

and $e_\delta((e_l, 1), (e_m, 1)) = 1$. Besides

$$\begin{aligned} e_\delta(\sigma(0, \chi_l), \sigma(0, \chi_m)) &= e_\delta((Ce_l, \chi^{De_l}), (Ce_m, \chi^{De_m})) = \chi^{De_m}(Ce_l)\chi^{-De_l}(Ce_m) \\ &= \zeta^{\langle Ce_l|De_m\rangle - \langle De_l|Ce_m\rangle} = \zeta^{{}^te_l({}^tC\Delta D - {}^tD\Delta C)e_m} \end{aligned}$$

and $e_\delta((0, \chi_l), (0, \chi_m)) = 1$. Finally

$$\begin{aligned} e_\delta(\sigma(e_l, 1), \sigma(0, \chi_m)) &= e_\delta((Ae_l, \chi^{Be_l}), (Ce_m, \chi^{De_m})) = \chi^{De_m}(Ae_l)\chi^{-Be_l}(Ce_m) \\ &= \zeta^{\langle Ae_l|De_m\rangle - \langle Be_l|Ce_m\rangle} = \zeta^{{}^te_l({}^tA\Delta D - {}^tB\Delta C)e_m} \end{aligned}$$

and $e_\delta((e_l, 1), (0, \chi_m)) = \zeta^{d_g/d_l\delta_{l,m}}$. Hence, σ is symplectic if and only if $^tB\Delta A \equiv {}^tA\Delta B$, $^tD\Delta C \equiv {}^tC\Delta B$ and $^tA\Delta D - {}^tB\Delta C \equiv \Delta$ modulo d_g. The result immediately follows. □

Theorem 6.2.10 (Change of theta coordinates)**.** *Let $\Theta_{\mathcal{L}^2}$ be a symmetric theta-structure on $G(\mathcal{L}^2)$ and $\Theta_{\mathcal{L}}$ be the induced compatible theta-structure on $G(\mathcal{L})$. Let $\psi \in \mathrm{Aut}^0(\mathcal{H}(2\delta))$ and $\psi' \in \mathrm{Aut}^0(\mathcal{H}(\delta))$ be the induced symmetric automorphism (following Proposition 6.2.8.(i)). Let ζ be a primitive $2d_g$-th root of unitiy and*

$$M := \begin{pmatrix} A & C \\ B & D \end{pmatrix}$$

be the matrix of $\overline{\psi} \in \mathrm{Sp}(K(2\delta))$ in the canonical ζ-symplectic basis. Let $(\theta_i^{\mathcal{L}})_{i \in K_1(\delta)}$ and $(\theta'^{\mathcal{L}}_i)_{i \in K_1(\delta)}$ be respectively the basis of theta-functions for $\Theta_{\mathcal{L}}$ and $\Theta'_{\mathcal{L}} := \Theta_{\mathcal{L}} \circ \psi'$. Then, there exists $i_0 \in K_1(\delta)$ and $\lambda \in k^*$ such that for all $i \in K_1(\delta)$,

$$\theta'^{\mathcal{L}}_i = \lambda \sum_{j \in K_1(\delta)} \zeta^{\langle i|j\rangle - \langle Ai + Cj + 2i_0 | Bi + Dj\rangle} \theta^{\mathcal{L}}_{Ai+Cj+i_0}.$$

We can choose any value of $i_0 \in K_1(\delta)$ such that

$$\sum_{j \in K_1(\delta)} \zeta^{-\langle Cj + 2i_0 | Dj\rangle} \theta^{\mathcal{L}}_{i_0 + Cj} \neq 0.$$

Proof. By Proposition 6.2.4, there exists $i_0 \in K_1(\delta)$ and $\lambda \in k^*$ such that $\theta'^{\mathcal{L}}_0 = \lambda T_{i_0}$, where

$$T_{i_0} := \sum_{j \in K_2(\delta)} \Theta_{\mathcal{L}} \circ \psi'(1,0,j) \cdot \theta^{\mathcal{L}}_{i_0}$$

is non-zero.

As explained before, we can identify $K_1(\delta)$ with $K_2(\delta)$ via the map $j \longmapsto \chi^j$, where $\chi^j(i) := \zeta^{2\langle i|j\rangle}$ for all $i, j \in K_1(\delta)$ (ζ^2 being a primitive d_g-th root of unity). Similarly, we identify $K_1(2\delta)$ with $K_2(2\delta)$ via the map $j \longmapsto \widetilde{\chi}^j$, where $\widetilde{\chi}^j(i) = \zeta^{\langle i|j\rangle}$ for all $i, j \in K_1(2\delta)$. Now, by Proposition 6.2.8, we can express ψ' as follows: for all $i, j \in K_1(\delta)$, we have:

$$\psi'(1,i,\chi^j) = \left(\widetilde{\chi}^{j'}(i')^{-1}\overline{\psi}_2(i',\widetilde{\chi}^{j'})(\overline{\psi}_1(i',\widetilde{\chi}^{j'})), \overline{\psi}_1(i',\widetilde{\chi}^{j'}), \overline{\psi}_2(i',\widetilde{\chi}^{j'})\right),$$

with $i', j' \in K_1(2\delta)$ such that $\overline{i'} = i$ and $\overline{j'} = j$. It follows that for all $(i,j) \in K_1(\delta)$,

$$\begin{aligned}
\psi'(1,i,\chi^j) &= \left(\zeta^{-\langle i'|j'\rangle}\widetilde{\chi}^{Bi'+Dj'}(Ai'+Cj'), \overline{Ai'+Cj'}, \overline{\widetilde{\chi}^{Bi'+Dj'}}\right) \\
&= (\zeta^{-\langle i'|j'\rangle + \langle Bi'+Dj'|Ai'+Cj'\rangle}, Ai+Cj, \chi^{Bi+Dj}) \\
&= (\zeta^{-\langle i|j\rangle + \langle Bi+Dj|Ai+Cj\rangle}, Ai+Cj, \chi^{Bi+Dj})
\end{aligned}$$

For the last equality, we can easily check that $-\langle i'|j'\rangle + \langle Bi'+Dj'|Ai'+Cj'\rangle$ only depends on the values of i' and j' modulo d_g. Consequently,

$$\begin{aligned}
T_{i_0} &= \sum_{j \in K_1(\delta)} \zeta^{\langle Dj|Cj\rangle}\chi^{Dj}(Cj+i_0)^{-1}\theta^{\mathcal{L}}_{Cj+i_0} \\
&= \sum_{j \in K_1(\delta)} \zeta^{\langle Dj|Cj\rangle}\zeta^{-2\langle Cj+i_0|Dj\rangle}\theta^{\mathcal{L}}_{Cj+i_0} \\
&= \sum_{j \in K_1(\delta)} \zeta^{-\langle Cj+2i_0|Dj\rangle}\theta^{\mathcal{L}}_{Cj+i_0}
\end{aligned}$$

And, if $T_{i_0} \neq 0$, we have for all $i \in K_1(\delta)$,

$$\begin{aligned}
\theta'^{\mathcal{L}}_i &= \Theta'_{\mathcal{L}}(1,i,1) \cdot \theta'^{\mathcal{L}}_0 = \lambda\Theta_{\mathcal{L}} \circ \psi(1,i,1) \cdot T_{i_0} \\
&= \lambda \sum_{j \in K_1(\delta)} \zeta^{-\langle Cj+2i_0|Dj\rangle}\Theta_{\mathcal{L}} \circ \psi(1,i,1) \cdot \theta^{\mathcal{L}}_{Cj+i_0} \\
&= \lambda \sum_{j \in K_1(\delta)} \zeta^{-\langle Cj+2i_0|Dj\rangle}\zeta^{\langle Bi|Ai\rangle}\chi^{Bi}(Ai+Cj+i_0)^{-1}\theta^{\mathcal{L}}_{Ai+Cj+i_0} \\
&= \lambda \sum_{j \in K_1(\delta)} \zeta^{-\langle Cj+2i_0|Dj\rangle}\zeta^{\langle Bi|Ai\rangle}\zeta^{-2\langle Bi|Ai+Cj+i_0\rangle}\theta^{\mathcal{L}}_{Ai+Cj+i_0} \\
&= \lambda \sum_{j \in K_1(\delta)} \zeta^{-\langle Cj+2i_0|Dj\rangle}\zeta^{-\langle Bi|Ai+2Cj+2i_0\rangle}\theta^{\mathcal{L}}_{Ai+Cj+i_0} \\
&= \lambda \sum_{j \in K_1(\delta)} \zeta^{-\langle Cj+2i_0|Bi+Dj\rangle - \langle Bi|Ai+Cj\rangle}\theta^{\mathcal{L}}_{Ai+Cj+i_0} \\
&= \lambda \sum_{j \in K_1(\delta)} \zeta^{-\langle Ai+Cj+2i_0|Bi+Dj\rangle + \langle Ai|Bi+Dj\rangle - \langle Bi|Ai+Cj\rangle}\theta^{\mathcal{L}}_{Ai+Cj+i_0} \\
&= \lambda \sum_{j \in K_1(\delta)} \zeta^{-\langle Ai+Cj+2i_0|Bi+Dj\rangle + \langle Ai|Dj\rangle - \langle Bi|Cj\rangle}\theta^{\mathcal{L}}_{Ai+Cj+i_0} \\
&= \lambda \sum_{j \in K_1(\delta)} \zeta^{-\langle Ai+Cj+2i_0|Bi+Dj\rangle + {}^t i({}^t A\Delta D - {}^t B\Delta C)j}\theta^{\mathcal{L}}_{Ai+Cj+i_0} \\
&= \lambda \sum_{j \in K_1(\delta)} \zeta^{-\langle Ai+Cj+2i_0|Bi+Dj\rangle + {}^t i\Delta j}\theta^{\mathcal{L}}_{Ai+Cj+i_0} \\
&= \lambda \sum_{j \in K_1(\delta)} \zeta^{-\langle Ai+Cj+2i_0|Bi+Dj\rangle + \langle i|j\rangle}\theta^{\mathcal{L}}_{Ai+Cj+i_0}
\end{aligned}$$

This completes the proof. □

6.2.3 Computing the dual of a 2-isogeny

We keep the notations from Section 6.1.1. Assume we have computed a 2-isogeny $f : (A, \mathcal{L}^2) \longrightarrow (B, \mathcal{M})$ with the techniques from Sections 6.1.2 and 6.1.3, it is then easy to compute its dual $\widetilde{f} : (B, \mathcal{M}^2) \longrightarrow (A, \mathcal{L})$ with the data we already have. By the following lemma, we only have to precompute the inverse theta-constants $(1/\theta_i^{\mathcal{L}}(0_A))_i$ to be able to evaluate $\widetilde{f}$. Up to Hadamard transforms, the formula is similar to Eq. (6.1). It is obtained via the change of theta coordinates formula proved in Theorem 6.2.10, which is, in this particular case, a Hadamard transform.

Lemma 6.2.11. *Let $f : (A, \mathcal{L}^2) \longrightarrow (B, \mathcal{M})$ be a 2-isogeny, ζ be an 8-th primitive root of unity and $\mathscr{B} := (S_1, \cdots, S_g, T_1, \cdots, T_g)$ be a ζ-symplectic basis adapted to f as in Theorem 6.1.1. Let $\Theta_{\mathcal{L}^2}$ be the symmetric theta structure on $G(\mathcal{L}^2)$induced by $\mathscr{B}$, $\Theta_{\mathcal{L}}$ its induced compatible symmetric theta structure on $G(\mathcal{L})$ and $\Theta_{\mathcal{M}}$ be the theta structure on $G(\mathcal{M})$ induced by $\mathscr{C} := ([2]f(S_1), \cdots, [2]f(S_g), f(T_1), \cdots, f(T_g))$ compatible with $\Theta_{\mathcal{L}^2}$ with respect to f by Theorem 6.1.1. Then:*

(i) $\widetilde{f}$ is a polarised abelian variety $(B, \mathcal{M}^2) \longrightarrow (A, \mathcal{L})$ of kernel $\ker(\widetilde{f}) = K_1(\overline{\Theta}_{\mathcal{M}})$.

(ii) Let $y \in B(k)$. Then there exists $\lambda \in k^*$ such that for all $i \in (\mathbb{Z}/2\mathbb{Z})^g$,

$$\theta_i^{\mathcal{L}}(\widetilde{f}(y)) \cdot \theta_i^{\mathcal{L}}(0_A) = \lambda \sum_{\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}} \chi(i) U_\chi^{\mathcal{M}}(y)^2,$$

where the theta coordinates above are associated to $\Theta_{\mathcal{L}}$ and $\Theta_{\mathcal{M}}$.

Proof. **(i)** By assumption f is a 2-isogeny, so $\widetilde{f}$ is a 2-isogeny by Lemma 2.2.2.(ii), so $\widetilde{f}$ is a polarised isogeny $(A, \mathcal{L}_0^2) \longrightarrow (B, \mathcal{M}_0)$ (recall that $\mathcal{L} = \mathcal{L}_0^2$ and $\mathcal{M} = \mathcal{M}_0^2$). It follows that $\widetilde{f}$ is a polarised isogeny $(A, \mathcal{L}^2) \longrightarrow (B, \mathcal{M})$.

Besides, $f(A[2]) \subseteq \ker(\widetilde{f})$ since $\widetilde{f} \circ f = [2]$ and $\widetilde{f}$ is separable since $\mathrm{char}(k)$ is odd so $\#\ker(\widetilde{f}) = \deg(\widetilde{f}) = \deg(f) = 2^g$. We also have $f(A[2]) = f(K_1(\overline{\Theta}_{\mathcal{L}})) = K_1(\overline{\Theta}_{\mathcal{M}})$ by Theorem 6.1.1.(i) and $\#K_1(\overline{\Theta}_{\mathcal{M}}) = 2^g$ so the inclusion $f(A[2]) \subseteq \ker(\widetilde{f})$ is an equality, which proves (i).

(ii) Our goal here is to find a well-chosen ζ-symplectic basis adapted to $\widetilde{f}$ in order to apply Theorem 6.1.1 and its corollary (Corollary 6.1.3). Let $\zeta' \in k^*$ such that $\zeta'^2 = \zeta$ and $\mathscr{B}' := (S_1', \cdots, S_g', T_1', \cdots, T_g')$ be a ζ'-symplectic basis of $A[16]$ such that $[2]\mathscr{B}' = \mathscr{B}$. Then, by Theorem 5.3.30 (points (ii) and (ii)), $\mathscr{B}'$ induces a symmetric theta structure $\Theta_{\mathcal{L}^4}$ that is compatible with $\Theta_{\mathcal{L}^2}$. Following the same reasoning as in the proof of Theorem 5.3.30.(i), we obtain that $\mathscr{C}' := ([2]f(S_1'), \cdots, [2]f(S_g'), f(T_1'), \cdots, f(T_g'))$ induces a symmetric theta structure $\Theta_{\mathcal{M}^2}$ compatible with $\Theta_{\mathcal{M}}$ and with $\Theta_{\mathcal{L}^4}$ with respect to f.

We now apply the Hadamard Heisenberg automorphism to $\Theta_{\mathcal{M}^2}$ and $\Theta_{\mathcal{L}}$ to obtain compatible theta structures with respect to $\widetilde{f}$ and apply Theorem 6.1.1. Let $\psi \in \mathrm{Aut}^0(\mathcal{H}(8, \cdots, 8))$ such that $\overline{\psi}$ has matrix

$$M_\psi := \begin{pmatrix} 0 & -I_g \\ I_g & 0 \end{pmatrix} \in \mathrm{Sp}_{2g}(\mathbb{Z}/8\mathbb{Z}),$$

in the canonical ζ-symplectic basis. Let $\psi' \in \mathrm{Aut}^0(\mathcal{H}(4, \cdots, 4))$ be the symmetric Heisenberg automorphism induced by ψ (Proposition 6.2.8.(i)). Let $\Theta'_{\mathcal{M}^2} := \Theta_{\mathcal{M}^2} \circ \psi$, $\Theta'_{\mathcal{M}} := \Theta_{\mathcal{M}} \circ \psi'$ and $\Theta'_{\mathcal{L}} := \Theta_{\mathcal{L}} \circ \psi'$. Then $\Theta'_{\mathcal{M}^2}$ is induced by $M_\psi^T \cdot \mathscr{C}' = (f(T_1'), \cdots, f(T_g'), -[2]f(S_1'), \cdots, -[2]f(S_g'))$ and $\Theta'_{\mathcal{L}}$ is induced by

$$\begin{aligned} M_\psi^T \cdot \mathscr{B} &= ([2]T_1, \cdots, [2]T_g, -[2]S_1, \cdots, -[2]S_g) \\ &= ([2]\widetilde{f}(f(T_1')), \cdots, [2]\widetilde{f}(f(T_g')), \widetilde{f}(-[2]f(S_1')), \cdots, \widetilde{f}(-[2]f(S_g'))). \end{aligned}$$

By (i), $\ker(\widetilde{f}) = K_1(\overline{\Theta}_{\mathcal{M}})$ so $M_\psi^T \cdot \mathscr{C}'$ is adapted to $\widetilde{f}$ and $\Theta'_{\mathcal{L}}$ is the theta structure compatible with $\Theta'_{\mathcal{M}^2}$ with respect to $\widetilde{f}$ defined in Theorem 6.1.1.(i). We can then apply Theorem 6.1.1.(ii) and then Corollary 6.1.3 to the coordinates $(\theta_i'^{\mathcal{M}})_i$ and $(\theta_i'^{\mathcal{L}})_i$ respectively associated to $\Theta'_{\mathcal{M}}$ and $\Theta'_{\mathcal{L}}$. If $y \in B(k)$ is fixed, we obtain the existence of $\lambda_1 \in k^*$ such that for all $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$,

$$U_\chi'^{\mathcal{L}}(\widetilde{f}(y)) \cdot U_\chi'^{\mathcal{L}}(0_A) = \lambda_1 \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t) \theta_t'^{\mathcal{M}}(y)^2,$$

where the $U'^{\mathcal{L}}_{\chi}$ are the dual theta coordinates of the $\theta'^{\mathcal{L}}_i$. Applying Theorem 6.2.10, we obtain the existence of $\lambda_2 \in k^*$ such that for all $i \in (\mathbb{Z}/2\mathbb{Z})^g$,

$$\theta'^{\mathcal{M}}_i = \lambda_2 \sum_{j\in(\mathbb{Z}/2\mathbb{Z})^g} \zeta^{4\langle i|j\rangle}\theta_{-j} = \lambda_2 \sum_{j\in(\mathbb{Z}/2\mathbb{Z})^g} (-1)^{\langle i|j\rangle}\theta_j = \lambda_2 \cdot U^{\mathcal{M}}_{\chi^i}.$$

Similarly, there exists $\lambda_2 \in k^*$ such that $\theta^{\mathcal{L}}_i = \lambda_3 \cdot U'^{\mathcal{L}}_{\chi^i}$. It follows that for all $i \in (\mathbb{Z}/2\mathbb{Z})^g$,

$$\theta^{\mathcal{L}}_i(\widetilde{f}(y)) \cdot \theta^{\mathcal{L}}_i(0_A) = \lambda \sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi^i(t) U^{\mathcal{M}}_{\chi^t}(y)^2 = \lambda \sum_{t\in(\mathbb{Z}/2\mathbb{Z})^g} \chi^t(i) U^{\mathcal{M}}_{\chi^t}(y)^2 = \lambda \sum_{\chi\in\widehat{(\mathbb{Z}/2\mathbb{Z})^g}} \chi(i) U^{\mathcal{M}}_{\chi}(y)^2,$$

with $\lambda := \lambda_1\lambda_2^2/\lambda_3^2$. This completes the proof. □

6.3 Computing a chain of 2-isogenies

The goal of this section is to explain how we compute a 2^e-isogeny between principally polarised abelian varieties $f : A \longrightarrow B$ with kernel of rank $g := \dim(A)$ *i.e.* admitting a basis with g elements. The following lemma ensures that f can be decomposed into a chain of 2-isogenies

$$A \xrightarrow{f_1} A_1 \xrightarrow{f_2} A_2 \quad \cdots \quad A_{e-2} \xrightarrow{f_{e-1}} A_{e-1} \xrightarrow{f_e} B. \tag{6.12}$$

Hence, if we know how to compute each 2-isogeny of the chain (as explained in Section 6.1), we can apply well known strategies used for elliptic curve isogeny computations [JDF11].

Lemma 6.3.1. *Let $d_1, d_2 \in \mathbb{N}^*$ not divisible by $\mathrm{char}(k)$ and $d := d_1d_2$. Let $(A, \mathcal{L}_0)$ and $(B, \mathcal{M}_0)$ be principally polarised abelian varieties of dimension g and $f : A \longrightarrow B$ be a d-isogeny. Assume that $\ker(f)$ has rank g,* i.e. *admits a basis with g elements. Then:*

(i) There exists a principally polarised abelian variety $(C, \mathcal{N}_0)$, a d_1-isogeny $f_1 : A \longrightarrow C$ and a d_2-isogeny $f_2 : C \longrightarrow B$ such that $f = f_2 \circ f_1$.

(ii) f_1 and f_2 are unique up to post or pre composition by an isomorphism respectively, and we have $\ker(f_1) = \ker(f)[d_1] = [d_2]\ker(f)$ and $\ker(f_2) = f_1(\ker(f))$.

Proof. **(i)** This is point (i) of Corollary 1.4.44 in the context of polarised isogenies. Consider $f_1 : A \longrightarrow C$ of kernel $\ker(f)[d_1]$ and $f_2 : C \longrightarrow B'$ of kernel $f_1(\ker(f))$. These isogenies are well defined by Theorem 1.4.41 and besides $\ker(f_2 \circ f_1) = \ker(f)$ by construction so Theorem 1.4.41 again ensures that $f_2 \circ f_1$ and f are equal up to post composition by an isomorphism $B' \xrightarrow{\sim} B$, so we may assume that $f = f_2 \circ f_1$.

We now prove that $\ker(f_1)$ is isotropic for $e_{\mathcal{L}_0^{d_1}}$ in order to apply Corollary 5.1.13 and prove the existence of a principal polarisation of C. Since f is a d-isogeny, we have $\widetilde{f} \circ f = [d]$ so $\ker(f) \subseteq A[d]$. Besides, f is separable since $\mathrm{char}(k)$ does not divide d, so that $\#\ker(f) = \deg(f) = \sqrt{\deg([d])} = d^g$ by Proposition 1.4.33 and Proposition 1.4.53. Besides, $\ker(f)$ has rank g so we may write $\ker(f) := \langle x_1\rangle \oplus \cdots \oplus \langle x_g\rangle$ with $x_1, \cdots, x_g \in A[d]$ of order d. We then have

$$\ker(f_1) = \ker(f)[d_1] = \langle [d_2]x_1\rangle \oplus \cdots \oplus \langle [d_2]x_g\rangle = [d_2]\ker(f). \tag{6.13}$$

Hence, for all $x, y \in \ker(f)[d_1]$, we may write $y = [d_2]y'$ with $y' \in \ker(f)$ and we then have by Proposition 5.1.6.(iv),

$$1 = e_{\mathcal{L}_0^d}(x, y') = e_{\mathcal{L}_0^{d_1}}(x, [d_2]y') = e_{\mathcal{L}_0^{d_1}}(x, y),$$

so that $\ker(f_1)$ is isotropic for $e_{\mathcal{L}_0^{d_1}}$.

Then, Corollary 5.1.13 ensures the existence of a line bundle $\mathcal{N}_0$ on C such that $f_1^*\mathcal{N}_0 \simeq \mathcal{L}_0^{d_1}$. We then have $\widehat{f_1} \circ \varphi_{\mathcal{N}_0} \circ f_1 = \varphi_{\mathcal{L}_0^{d_1}} = [d_1]\varphi_{\mathcal{L}_0}$. It follows that $\deg(f_1)^2 \deg(\varphi_{\mathcal{N}_0}) = d_1^{2g}$ with $\deg(f_1) = \#\ker(f_1) = d_1^g$ by separability of f_1 and by Eq. (6.13). Hence, $\deg(\varphi_{\mathcal{N}_0}) = 1$ so $\varphi_{\mathcal{N}_0}$ is a principal polarisation and f_1 is a d_1-isogeny with respect to $\varphi_{\mathcal{L}_0}$ and $\varphi_{\mathcal{N}_0}$.

Finally, we have on the one hand,

$$[d]\varphi_{\mathcal{L}_0} = \widehat{f} \circ \varphi_{\mathcal{M}_0} \circ f = \widehat{f_1} \circ \widehat{f_2} \circ \varphi_{\mathcal{M}_0} \circ f_2 \circ f_1$$

And on the other hand, $\widehat{f_1} \circ \varphi_{\mathcal{N}_0} \circ f_1 = [d_1]\varphi_{\mathcal{L}_0}$, so that

$$\widehat{f_1} \circ ([d_2]\varphi_{\mathcal{N}_0} - \widehat{f_2} \circ \varphi_{\mathcal{M}_0} \circ f_2) \circ f_1 = 0$$

Since f_1 is surjective, as any isogeny, we have $\widehat{f_1} \circ ([d_2]\varphi_{\mathcal{N}_0} - \widehat{f_2} \circ \varphi_{\mathcal{M}_0} \circ f_2) = 0$ so applying the dual to the last equality and using the surjectivity of f_1 again, we obtain that $[d_2]\varphi_{\mathcal{N}_0} - \widehat{f_2} \circ \varphi_{\mathcal{M}_0} \circ f_2 = 0$ *i.e.* that f_2 is a d_2-isogeny.

(ii) We now prove the uniqueness of f_1 and f_2 up to post or pre composition by an isomorphism. Indeed, assume $f = f_2' \circ f_1'$ where f_i' is a d_i-isogeny for $i \in \{1, 2\}$. Then $\ker(f_1') \subseteq A[d_1]$ since f_1' is a d_1-isogeny, so that $\ker(f_1') \subseteq \ker(f)[d_1]$. But $\#\ker(f_1') = \deg(f_1') = d_1^g$ by separability of f (hence of f_1') and $\#\ker(f)[d_1] = d_1^g$ by Eq. (6.13), so that $\ker(f_1') = \ker(f)[d_1]$. By Theorem 1.4.41, there exists an isomorphism λ such that $\lambda \circ f_1 = f_1'$, and we have $f = f_2' \circ \lambda \circ f_1 = f_2 \circ f_1$ so $f_2' \circ \lambda = f_2$ by surjectivity of f_1, which completes the proof. □

Remark 6.3.2. Note that the hypothesis $\ker(f)$ has rank g in Lemma 6.3.1 may be replaced by $\gcd(d_1, d_2) = 1$. Indeed, in that case we have

$$\ker(f) = \ker(f)[d_1] \oplus \ker(f)[d_2], \tag{6.14}$$

which ensures that $\ker(f_1) = \ker(f)[d_1] = [d_2]\ker(f)$, which is the key ingredient to prove that $\ker(f_1)$ is isotropic for $e_{\mathcal{L}_0^{d_1}}$ in point (i) and obtain the existence of a principal polarisation on the codomain C of f_1. Eq. (6.14) also ensures the uniqueness in point (ii).

Throughout this section, $(A, \mathcal{L}_0)$ and $(B, \mathcal{M}_0)$ will be principally polarised abelian varieties and $f : A \longrightarrow B$ will be a 2^e-isogeny with kernel of rank g, decomposed as a chain of 2-isogenies, as in Eq. (6.12). We assume that we are given $T_1, \cdots, T_g \in A[2^{e+2}]$ forming an isotropic subgroup (for $e_{\mathcal{L}_0^{2^{e+2}}}$, or equivalently, for the 2^{e+2}-th Weil pairing) and such that $\ker(f) = \langle [4]T_1, \cdots, [4]T_g \rangle$. As we have seen in Sections 6.1.2 and 6.1.3, this additional torsion requirement (2^{e+2}-torsion instead of 2^e-torsion) is necessary to compute each 2-isogeny f_i of the chain which requires 8-torsion points above its kernel.

6.3.1 Computing an adapted theta structure on the domain

In practice, when we start the isogeny chain computation, the domain $(A, \mathcal{L}_0^2)$ is not equipped with a level 2 theta structure. We need to determine a level 2 theta structure and to compute the associated theta coordinates. Quite often, $(A, \mathcal{L}_0)$ is a product of principally polarised abelian varieties of smaller dimension.

When $(A, \mathcal{L}_0)$ is a product of elliptic curves $E_1 \times \cdots \times E_g$ (equipped with their respective canonical principal polarisation), we may first convert each E_i to a Montgomery model and then use Proposition 5.3.47 to convert Montgomery $(X : Z)$-coordinates into level 2 theta coordinates $(\theta_0^{E_i}, \theta_1^{E_i})$. Then, Eq. (5.14) applies to compute the product theta coordinates associated to the product theta structure on $(A, \mathcal{L}_0^2)$:

$$\theta_i^{\mathcal{L}_0^2}(x_1, \cdots, x_g) = \prod_{j=1}^{g} \theta_{i_j}^{E_j}(x_j),$$

for all $i \in (\mathbb{Z}/2\mathbb{Z})^g$ and $(x_1, \cdots, x_g) \in A(k)$.

When $(A, \mathcal{L}_0)$ is the Jacobian of a hyperelliptic curve of genus 2, we can convert Mumford coordinates into theta coordinates using Thomae's formulas [Mum84, § III.a.8] (see also [CR15, § 5.1 and A.4]). If $(A, \mathcal{L}_0)$ is a product of such Jacobians, we can use Thomae's formulas for each component and then Eq. (5.14) to compute the product theta coordinates.

Usually, the level 2 theta structure $\Theta_{\mathcal{L}_0^2}$ on $(A, \mathcal{L}_0^2)$ (often obtained as a product) is not adapted to f. We do not expect $K_2(\overline{\Theta}_{\mathcal{L}_0^2}) = \ker(f_1)$ which is a necessary condition to apply the algorithms

from Section 6.1 to compute the first 2-isogeny f_1. We need to change the theta structure on $(A, \mathcal{L}_0^2)$ and to compute the new theta coordinates with the formulas of Theorem 6.2.10.

Let $\zeta \in k^*$ be a 2^{e+2}-th primitive root of unity. Consider $S_1, \cdots, S_g \in A[2^{e+2}]$ such that $\mathscr{B} := (S_1, \cdots, S_g, T_1, \cdots, T_g)$ is a ζ-symplectic basis of $A[2^{e+2}]$ adapted to f in the sense of Definition 6.1.2. Recall that we already assumed from the start that $\ker(f) = \langle [4]T_1, \cdots, [4]T_g \rangle$. Then $[2^{e-1}]\mathscr{B}$ is adapted to f_1 and we may consider the level 2 theta structure $\Theta'_{\mathcal{L}_0^2}$ associated to $[2^e]\mathscr{B}$ and the associated theta coordinates $(\theta_i'^{\mathcal{L}_0^2})_i$. If we know the theta points $(\theta_i'^{\mathcal{L}_0^2}([2^{e-1}]T_l))_i$ for $l \in [\![1\ ;\ g]\!]$, Theorem 6.1.1 ensures that we can compute the codomain theta null-point of f_1 using Algorithm 6.5 and then evaluate f_1 using Algorithm 6.1, 6.6 or 6.7. To proceed, we need to apply a change of coordinates formula from $(\theta_i^{\mathcal{L}_0^2})_i$ to $(\theta_i'^{\mathcal{L}_0^2})_i$ associated to $\Theta_{\mathcal{L}_0^2}$ and $\Theta'_{\mathcal{L}_0^2}$ respectively.

We can consider a ζ^{2^e}-symplectic basis $\mathscr{B}_0$ of $A[4]$ associated to $\Theta_{\mathcal{L}_0^2}$. For instance, if $(A, \mathcal{L}_0) = (A_1, \mathcal{N}_1) \times (A_2, \mathcal{N}_2)$ is a product of principally polarised abelian varieties and $\Theta_{\mathcal{L}_0^2} = \Theta_{\mathcal{N}_1} \times \Theta_{\mathcal{N}_2}$ is the product theta structure, where $\Theta_{\mathcal{N}_i}$ is induced by a ζ^{2^e}-symplectic basis $\mathscr{B}_i = (x_1^{(i)}, \cdots, x_{g_i}^{(i)}, y_1^{(i)}, \cdots, y_{g_i}^{(i)})$ of $A_i[4]$ for $i \in \{1, 2\}$, then

$$\mathscr{B}_0 := \mathscr{B}_1 \times \mathscr{B}_2 = ((x_1^{(1)}, 0), \cdots, (x_{g_1}^{(1)}, 0), (0, x_1^{(2)}), \cdots, (0, x_{g_2}^{(2)}), \\ (y_1^{(1)}, 0), \cdots, (y_{g_1}^{(1)}, 0), (0, y_1^{(2)}), \cdots, (0, y_{g_2}^{(2)})). \quad (6.15)$$

Such a basis is called a *product symplectic basis*. We can then compute the symplectic change of basis matrix $M \in \mathrm{Sp}_{2g}(\mathbb{Z}/4\mathbb{Z})$ from $\mathscr{B}_0$ to $[2^e]\mathscr{B}$ and obtain apply Theorem 6.2.10 to obtain the change of coordinates formula from $(\theta_i^{\mathcal{L}_0^2})_i$ to $(\theta_i'^{\mathcal{L}_0^2})_i$.

Remark 6.3.3. In plain generality, M has 4 blocks of size $g \times g$:

$$M = \begin{pmatrix} A & C \\ B & D \end{pmatrix}.$$

The blocks C and D can be computed directly by expressing the points $[2^e]T_1, \cdots, [2^e]T_g$ in the basis $\mathscr{B}_0$ (via multiple discrete logarithms). The blocks A and B giving the expression of $[2^e]S_1, \cdots, [2^e]S_g$ (which are unknown) in $\mathscr{B}$ can be found by solving the equations ${}^tBA \equiv {}^tAB$ and ${}^tAD - {}^tBC \equiv I_g \mod 4$. This is linear algebra over $\mathbb{Z}/4\mathbb{Z}$. Note however that in practice when f is obtained via Kani's lemma, we can obtain an expression for M directly and avoid the need for discrete logarithm computations and linear algebra over $\mathbb{Z}/4\mathbb{Z}$ (see Section 6.4).

6.3.2 How the adapted theta structure propagates along the chain

The change of theta coordinates needed to compute f_1 with the Algorithms from Section 6.1 does not need to be repeated at every step of the chain computation as the choice of adapted theta structure propagates along the chain.

Lemma 6.3.4. *If $\mathscr{B} = (S_1, \cdots, S_g, T_1, \cdots, T_g)$ is a ζ-symplectic basis of $A[2^{e+2}]$ adapted to f decomposed into a chain of 2-isogenies $f = f_e \circ \cdots \circ f_1$, then:*

(i) $[2^{e-1}]\mathscr{B}$ is a $\zeta^{2^{e-1}}$-symplectic basis of $A[8]$ adapted to f_1.

(ii) For all $i \in [\![1\ ;\ e-1]\!]$,

$$\mathscr{C}_i := ([2^{e-1}]f_i \circ \cdots \circ f_1(S_1), \cdots, [2^{e-1}]f_i \circ \cdots \circ f_1(S_g), \\ [2^{e-1-i}]f_i \circ \cdots \circ f_1(T_1), \cdots, [2^{e-1-i}]f_i \circ \cdots \circ f_1(T_g))$$

is a $\zeta^{2^{e-1}}$-symplectic basis of $A_i[8]$ adapted to $f_{i+1} : A_i \longrightarrow A_{i+1}$.

Proof. **(i)** This is a trivial consequence of the fact that $\ker(f_1) = [2^{e-1}]\ker(f)$, which follows from Lemma 6.3.1.(ii).

(ii) If $i \in [\![1\ ;\ e-1]\!]$, Lemma 6.3.1.(ii). implies again that $\ker(f_{i+1}) = [2^{e+1-i}]f_i \circ \cdots \circ f_1(\ker(f))$. Using Proposition 5.1.8 and the standard properties of the Weil pairing (Proposition 1.4.70), we see that $\mathscr{C}_i$ is a $\zeta^{2^{e-1}}$-symplectic basis of $A_i[8]$ and that:

$$\ker(f_{i+1}) = [2^{e+1-i}]f_i \circ \cdots \circ f_1(\ker(f)) = [4]\langle [2^{e-1-i}]f_i \circ \cdots \circ f_1(T_1), \cdots, [2^{e-1-i}]f_i \circ \cdots \circ f_1(T_g)\rangle,$$

so that $\mathscr{C}_i$ is adapted to f_{i+1}. The result follows. □

Lemma 6.3.4 above ensures that if we have computed f_1 using the theta structure induced by $[2^e]\mathscr{B}$ as explained in Section 6.3.1, then for all $i \in [\![1\ ;\ e-1]\!]$, Theorem 6.1.1 applies to the level 4 theta structure induced by $\mathscr{C}_i$. Again by Theorem 6.1.1, the image theta points of $f_i \circ \cdots \circ f_1$ are expressed in level 2 theta coordinates associated to the theta structure naturally induced by $[2]\mathscr{C}_i$, which is compatible with the one induced by $\mathscr{C}_i$. Hence, all the algorithms from Section 6.1 apply to compute $f_{i+1} : A_i \longrightarrow A_{i+1}$ from the theta coordinates of $[2^{e-i}]f_i \circ \cdots \circ f_1(T_1), \cdots, [2^{e-i}]f_i \circ \cdots \circ f_1(T_g)$.

6.3.3 Quasi-linear computational strategies

Similarly to elliptic curve isogeny chain computations, the computation of the isogeny chain f in dimension g requires to compute $[2^{e-i}]f_i \circ \cdots \circ f_1(T_1), \cdots, [2^{e-i}]f_i \circ \cdots \circ f_1(T_g)$ in order to obtain f_{i+1} for all $i \in [\![1\ ;\ g]\!]$. Hence, theta point duplications (*e.g.* using Algorithm 5.2) and isogeny evaluations are involved.

A naive method would consist in evaluating $f_i \circ \cdots \circ f_1(T_1), \cdots, f_i \circ \cdots \circ f_1(T_g)$ iteratively and then multiplying these theta points by 2^{e-1-i} in order to obtain f_{i+1} and proceed further. This would require $g(e-1)$ 2-isogeny evaluations and

$$g\sum_{i=0}^{e-1}(e-1-i) = \frac{ge(e-1)}{2}$$

duplications. Alternatively, we can compute the duplicates $T_l, [2]T_l, \cdots, [2^{e-1}]T_l$ for $l \in [\![1\ ;\ g]\!]$ and then evaluate $f_i \circ \cdots \circ f_1([2^{e-i-1}]T_1), \cdots, f_i \circ \cdots \circ f_1([2^{e-i-1}]T_g)$ iteratively to compute f_{i+1} for all $i \in [\![1\ ;\ e-1]\!]$. This costs $g(e-1)$ duplications and $ge(e-1)/2$ 2-isogeny evaluations.

Both of the above naive methods have a quadratic cost in the length e of the chain. As it was done for elliptic curves isogenies, we can propose alternate computation strategies minimizing the total cost depending on the relative cost of duplications and 2-isogeny evaluations. These divide and conquer strategies usually lead to a quasi-linear $O(e\log(e))$ number of duplications and 2-isogeny evaluations.

Strategies to compute a 2-isogeny chain

Computing the 2-isogeny chain f reduces to computing the leaves $\mathscr{C}_{i-1,e-i}$ of the binary computation tree whose:

- vertices are the basis $\mathscr{C}_{i,j} := [2^j]f_i \circ \cdots \circ f_1(T_1, \cdots, T_g)$ for all $i, j \in \mathbb{N}$ such that $i + j \leq e-1$;
- left edges are duplications $\mathscr{C}_{i,j-1} \xrightarrow{[2]} \mathscr{C}_{i,j}$;
- right edges are 2-isogeny evaluations $\mathscr{C}_{i-1,j} \xrightarrow{f_i} \mathscr{C}_{i,j}$.

Such a tree is displayed in Fig. 6.3 for $e = 5$. The computation tree can only be evaluated depth first and left first since the leaf $\mathscr{C}_{i-1,e-i}$ has to be computed prior to any evaluation by f_i. However, evaluating all the vertices $\mathscr{C}_{i,j}$ would be a waste of computational resources leading to a quadratic complexity $O(e^2)$. Optimal strategies consist in navigating the computation tree depth first and left first with a minimal number of duplications and evaluations to evaluate the leaves $\mathscr{C}_{i-1,e-i}$.

As in [JDF11], we can represent the computation tree as a *discrete equilateral triangle* T_e formed by points of the unit triangular equilateral lattice delimited by the x axis and the straight lines $y = \sqrt{3}x$ and $y = \sqrt{3}(e-1-x)$:

$$T_e := \left\{ \left(r + \frac{s}{2}, \frac{s\sqrt{3}}{2} \right) \middle| \ r, s \in \mathbb{N}, \quad r + s \leq e-1 \right\} \tag{6.16}$$

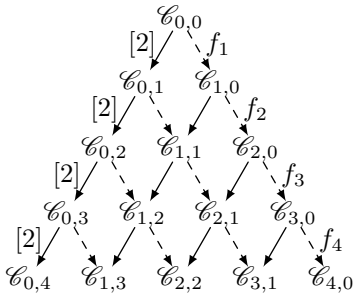

Figure 6.3: Computational structure of the 2^e-isogeny f with $e = 5$.

In T_e, *edges* are unit segments connecting two points of T_e. A *left edge* is a segment of positive slope and a *right edge* is a segment of negative slope. Edges are oriented in the direction of decreasing y coordinates. This defines an oriented graph structure on T_e. Vertices on $x, y \in T_e$ are ordered $x \to y$ if there exists a path from x to y. On a subgraph of T_e, the *root* is the initial point and *leaves* are final points.

Definition 6.3.5. A *strategy* S of T_e is a subgraph of T_e having a unique root. In the following, we only consider strategies that are:

1. *full*, meaning that S contains all leaves of T_e.
2. *well-formed*, meaning that there is only one path going through any interior point of S and no leaf in S distinct from the leaves of T_e.

Such a (full and well formed) strategy of T_e is also called a *strategy of depth* $e-1$. We denote $|S| = e$ its number of leaves.

Computing an optimal strategy

To compare strategies, we fix a measure μ parametrised by $(\alpha, \beta) \in \mathbb{R}_+^2$ on them, where α is the cost of a left edge (accounting for duplication cost) and β is the cost of a right edge (accounting for evaluation cost). Given such a measure, an *optimal strategy* of depth $e-1$ is a strategy of T_e with minimal cost.

We define the *tree topology* of a strategy S of depth $e-1$ as the binary tree with e leaves obtained by forgetting internal vertices of out degree less than two and keeping the same connectivity structure. Conversely, to any binary tree T with e leaves we associate a *canonical strategy* S_T of depth $e-1$ recursively as follows. If $e = 1$, we take $S_T := T_1$. If $e \geq 2$, we consider the left and right branches T' and T'' of T respectively and consider the canonical strategies $S' := S_{T'}$ and $S_{T''}$ associated to them. Let S'' be the translate of $S_{T''}$ by $|S'|$ to the right. Let r' and r'' be the roots of S' and S'' in T_e respectively and r be the root of T_e. Then the shortest paths rr' and rr'' from r to r' and r'' respectively are respectively made of left edges only and right edges only. We can then consider the strategy $S_T := rr' \cup rr'' \cup S' \cup S''$.

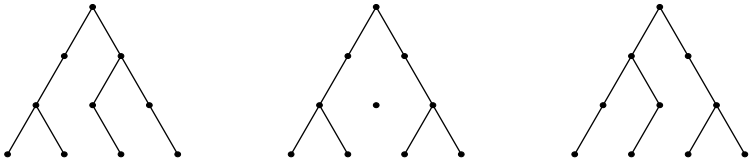

Figure 6.4: Three strategies of depth 3 sharing the same tree topology. The middle one is canonical.

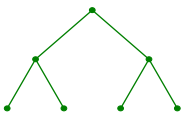

Figure 6.5: Tree topology of the strategies on the left.

The following result has been proved in [JDF11]:

Lemma 6.3.6. *[JDF11, Lemma 4.3] The canonical strategy is minimal, with respect to any measure, among all the strategies sharing the same tree topology.*

It follows that we can restrict to canonical strategies to find optimal strategies in the following. If S is a canonical strategy, we can consider its left and right branches S' and S'' as follows. If S has i leaves to the left of its root, we define $S' := S \cap T_i$ and $S'' := S \cap ((i,0) + T_{|S|-i})$.

Lemma 6.3.7. *[JDF11, Lemma 4.5] Let S be an optimal (canonical) strategy and let S' and S'' be its left and right branches respectively. Then, S' and S'' translated by $-|S'|$ are optimal strategies of $T_{|S'|}$ and $T_{|S''|}$ respectively.*

Proof. The proof is very natural. By Lemma 6.3.6, we know that S is a canonical strategy, so S' and S'' are well defined. If S' were not optimal, then by substituting an optimal strategy for S' inside S, we obtain a strategy with measure lower than $\mu(S)$. Contradiction. The same argument holds for S''. □

As pointed out in [JDF11], this suggests a dynamic programming approach to compute optimal strategies. For $e = 1$, the only optimal strategy is trivially $S = T_1$. Now, if we assume that we have computed optimal strategies $S_1, \cdots, S_{e-1}$ of $T_1, \cdots, T_{e-1}$ of respective measures $\mu(S_1), \cdots, \mu(S_{e-1})$, then the optimal strategy S_e will have left branch S_i and right branch S_{e-i} where:

$$i := \underset{1 \leq j \leq e-1}{\operatorname{argmin}} \left(\mu(S_j) + \mu(S_{e-j}) + (e-j)\alpha + j\beta\right).$$

In practice, we look for optimal strategies taking into account the higher cost of gluing isogenies that appear in the beginning of the isogeny chain. We refer to Section 6.6.5 for more details on these special optimal strategies in dimension 4.

Applying a strategy to compute a 2-isogeny chain

As suggested in [JAC+20, § 1.3.8], we can represent any strategy S in a unique way as a sequence of integers $(s_1, \cdots, s_{t-1})$ by considering the tree topology T_S of S. To establish this sequence $(s_1, \cdots, s_{t-1})$, we write down for every internal node of the tree T_S the number of leaves to its right and walk on it depth-first left-first.

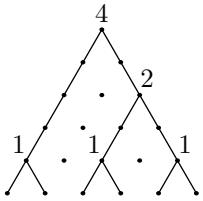

Figure 6.6: Strategy of depth 5 represented by $(4, 1, 2, 1, 1)$.

Given a strategy and a basis of the kernel, it is natural to compute the isogeny chain recursively, as proposed in [JAC+20, § 1.3.8]. An iterative version of the same algorithm derived from [CDPMR23, Algorithm 2] is presented in Algorithm 6.8.

Doubling points on domains of splitting isogenies

The computation of the isogeny chain in Algorithm 6.8 may involve point duplications on the domain A_i of an isogeny $f_{i+1} : A_i \longrightarrow A_{i+1}$ where Algorithm 5.2 may not be applied because of zero theta constants on A_i or zero dual theta constants on A_{i+1}.

Recall that if $(A, \mathcal{L})$ is a polarised abelian variety with a theta structure $\Theta_{\mathcal{L}}$ of level 2, and $(\theta_i^{\mathcal{L}}(x))_i$ is a theta point, then the computation of the double $(\theta_i^{\mathcal{L}}([2]x))_i$ in Algorithm 5.2 requires the computation of the inverse of the squared dual theta constants

$$U_\chi^{\mathcal{L}^2}(0_A)^2 = \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t)\theta_t^{\mathcal{L}}(0_A)^2$$

Algorithm 6.8: Computing an isogeny chain with a strategy.

Data: The level 2 theta coordinates (adapted to f) of $T_1, \cdots, T_g$ such that $\ker(f) = [4]\langle T_1, \cdots, T_g\rangle$ and a strategy $S = (s_1, \cdots, s_{t-1})$ of depth $e-1$.
Result: The 2-isogenies of the chain $f_1, \cdots, f_e$ such that $f = f_e \circ \cdots \circ f_1$.

$k \longleftarrow 1$;
$L_{levels} \longleftarrow [0]$;
$L_{basis} \longleftarrow [(T_1, \cdots, T_g)]$;
for $i = 1$ **to** e **do**
  $\mathscr{B}_K \longleftarrow$ last element of L_{basis};
  while $\sum_{x \in L_{levels}} x \neq e - k$ **do**
    Append s_k to L_{levels};
    $\mathscr{B}_K \longleftarrow [2^{s_k}]\mathscr{B}_K$;
    Append $\mathscr{B}_K$ to L_{basis};
    $k \longleftarrow k + 1$;
  end
  Use Algorithm 6.5 with input $\mathscr{B}_K$ to compute the isogeny f_i of kernel $[4]\langle\mathscr{B}_K\rangle$;
  Remove the last elements of L_{levels} and L_{basis};
  $L_{basis} \longleftarrow [f_i(\mathscr{C}) \mid \mathscr{C} \in L_{basis}]$ (Algorithm 6.1);
end
return $f_1, \cdots, f_e$;

for all $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}$. As a consequence of Theorem 6.1.1.(ii), $(U_\chi^{\mathcal{L}^2}(0_A))_\chi$ is the dual theta null point $(U_\chi^{\mathcal{M}}(0_B))_\chi$ of the codomain of the 2-isogeny $f : (A, \mathcal{L}^2) \longrightarrow (B, \mathcal{M})$ of kernel $\ker(f) = K_2(\overline{\Theta}_{\mathcal{L}})$.

If B is a product of abelian varieties *i.e.* if f is a splitting isogeny, some dual theta constants $U_\chi^{\mathcal{L}^2}(0_A) = U_\chi^{\mathcal{M}}(0_B)$ may be zero and Algorithm 5.2 may not be applied. In this case, to perform point duplications on A, we may apply a Hadamard transform on $\Theta_{\mathcal{L}}$ (as in the proof of Lemma 6.2.11) to obtain a theta structure $\Theta'_{\mathcal{L}}$ with associated theta coordinates

$$\theta'^{\mathcal{L}}_i = \sum_{j \in (\mathbb{Z}/2\mathbb{Z})^g} (-1)^{\langle i|j\rangle} \theta^{\mathcal{L}}_j \tag{6.17}$$

In that case, to perform point duplications, we need to invert the squared dual theta constants

$$U'^{\mathcal{L}^2}_\chi(0_A)^2 = \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t) \theta'^{\mathcal{L}}_t(0_A)^2,$$

where $(U'^{\mathcal{L}^2}_\chi(0_A))_\chi$ is the dual theta null point of the codomain of the 2-isogeny $f' : (A, \mathcal{L}^2) \longrightarrow (B', \mathcal{M}')$ of kernel $\ker(f') = K_2(\overline{\Theta}'_{\mathcal{L}}) = K_1(\overline{\Theta}_{\mathcal{L}})$. In an extreme majority of cases, we do not expect f' to be a splitting isogeny so the point duplication may be feasible. Once we have computed dual theta coordinates $(\theta'^{\mathcal{L}}_i([2]x))_i$, we obtain the desired theta coordinates $(\theta^{\mathcal{L}}_i([2]x))_i$ by applying another Hadamard transform as in Eq. (6.17). Indeed, the Hadamard transform is an involution up to the projective factor 2^g with $g = \dim(A)$.

Algorithm 5.2 cannot be applied either when some theta constants $\theta^{\mathcal{L}}_i(0_A)$ are zero. Then the Hadamard transform may also be a solution. If Algorithm 5.2 still cannot be applied because some theta constants $\theta'^{\mathcal{L}}_i(0_A)$ and $U'^{\mathcal{L}^2}_\chi(0_A)$ are still zero, we may apply a random change of theta coordinates obtained from Theorem 6.2.10.

6.3.4 Assumptions on the base field

So far we always have assumed that our base field k is algebraically closed. This assumption was meant to simplify the proofs in the formalism of algebraic theta functions. Note that Mumford also made this assumption in [Mum66]. However, in practice our base field k will be a finite field such that $A[2^{e+2}] \subseteq A(k)$ and the points $T_1, \cdots, T_g \in A[2^{e+2}]$ will be k-rational, so their theta coordinates will also be defined over k.

For instance, if A is a product of elliptic curves $E_1 \times \cdots \times E_g$, then the components of $T_1, \cdots, T_g$ in the E_i will have Weierstrass coordinates $(X : Y : Z)$ defined over k. Using Proposition 5.3.47 and Eq. (5.14), we shall obtain product theta coordinates defined over k. The change of theta coordinates from Section 6.3.1 involving a 2^{e+2}-th Weil pairing $\zeta \in k^*$ will lead to theta coordinates defined over k again. Then, duplications and 2-isogeny computations along the chain will take place over k. Hence, the computation of f will only require arithmetic operations over k. **In the following we shall only assume that k is a field such that the $A[2^{e+2}] \subseteq A(k)$.**

6.4 Isogenies obtained from Kani's lemma

The higher dimensional 2^e-isogenies considered in applications studied in Part I were all obtained from Kani's lemma. They are defined between products of abelian varieties and their kernel have a very specific form. For that reason, finding adapted symplectic basis on the domain to start the 2-isogeny chain computation (as explained in Section 6.3.1) is easier than in the general case.

In Section 6.4.1, we explain how such an adapted symplectic basis can be computed when the 2^{e+2}-torsion on the domain is available (*e.g.* rational over our base field k) and how we can recover the product theta structure on the codomain. In Section 6.4.2, we explain how to proceed when the codomain is known but only the 2^f-torsion is available with $f \geq e/2+2$. This last case is relevant for SQIsignHD, as explained in Sections 2.3.2 and 3.3.4.

6.4.1 Change of theta coordinates on the domain and codomain with full available torsion

Lemma 6.4.1. *Let a and b be odd and coprime integers and $d := a + b$ not divisible by $\mathrm{char}(k)$. Consider an (a, b)-isogeny diamond between principally polarised abelian varieties of dimension g:*

$$\begin{array}{ccc} A' & \xrightarrow{\varphi'} & B' \\ {\scriptstyle\psi}\uparrow & & \uparrow{\scriptstyle\psi'} \\ A & \xrightarrow{\varphi} & B \end{array}$$

and the associated d-isogeny by Kani's lemma:

$$F := \begin{pmatrix} \varphi & \widetilde{\psi'} \\ -\psi & \widetilde{\varphi'} \end{pmatrix} : A \times B' \longrightarrow B \times A'.$$

Let $\zeta \in k^*$ be a $4d$-th primitive root of unity and $(x_1, \cdots, x_g, y_1, \cdots, y_g)$ be a ζ-symplectic basis of $B[4d]$. Let α and β be modular inverses of a and b modulo $4d$ respectively.

(i) For all $i \in [\![1 \ ; \ g]\!]$, we denote:

$$S_i := ([-\alpha]\widetilde{\varphi}(y_i), 0), \quad S_{i+g} := (0, [\beta]\psi'(x_i)),$$

$$T_i := (\widetilde{\varphi}(x_i), \psi'(x_i)), \quad T_{i+g} := ([1 - \alpha d]\widetilde{\varphi}(y_i), \psi'(y_i)).$$

Then $\mathscr{B} := (S_1, \cdots, S_{2g}, T_1, \cdots, T_{2g})$ is a ζ-symplectic basis of $(A \times B')[4d]$ adapted to F i.e. *such that $\ker(F) = [4]\langle T_1, \cdots, T_{2g}\rangle$.*

(ii) For all $i \in [\![1 \ ; \ g]\!]$, we denote:

$$U_i := ([d]x_i, 0), \quad U_{i+g} := (0, [d]\psi \circ \widetilde{\varphi}(x_i)),$$

$$V_i := ([d]y_i, 0), \quad V_{i+g} := (0, [d\alpha\beta]\psi \circ \widetilde{\varphi}(y_i)).$$

Then $\mathscr{C}_0 := (U_1, \cdots, U_{2g}, V_1, \cdots, V_{2g})$ is a product ζ^d-symplectic basis of $(B \times A')[4]$.

(iii) *The ζ^d-symplectic basis $\mathscr{C} := ([d]F(S_1), \cdots, [d]F(S_{2g}), F(T_1), \cdots, F(T_{2g}))$ of $(B \times A')[4]$ naturally induced by F via Theorem 6.1.1 is related to $\mathscr{C}_0$ by the following formulas. For all $i \in [\![1\ ;\ g]\!]$:*

$$[d]F(S_i) = -V_i + [b]V_{i+g}, \quad [d]F(S_{i+g}) = U_i + [\beta]U_{i+g},$$
$$F(T_i) = U_i, \quad F(T_{i+g}) = [b]V_{i+g},$$

so that:

$$U_i = F(T_i), \quad U_{i+g} = [bd]F(S_{i+g}) - [b]F(T_i),$$
$$V_i = F(T_{i+g}) - [d]F(S_i), \quad V_{i+g} = [\beta]F(T_{i+g}).$$

To prove the above lemma, we use the following results on commutator pairings.

Lemma 6.4.2.

(i) *Let $(A, \mathcal{L})$ and $(B, \mathcal{M})$ be two polarised abelian varieties and consider $\mathcal{L} \star \mathcal{M} := \pi_1^*\mathcal{L} \otimes \pi_2^*\mathcal{M}$, the product of the line bundles $\mathcal{L}$ and $\mathcal{M}$ on $A \times B$, where π_i is the projection on the i-th component of $A \times B$ (for $i \in \{1, 2\}$). Then, for all $(x, y), (x', y') \in K(\mathcal{L}) \times K(\mathcal{M})$, we have:*

$$e_{\mathcal{L}\star\mathcal{M}}((x, y), (x', y')) = e_{\mathcal{L}}(x, x')e_{\mathcal{M}}(y, y').$$

(ii) *Let $(A, \mathcal{L}_0)$ and $(B, \mathcal{M}_0)$ be two principally polarised abelian varieties and $f : A \longrightarrow B$ be a d-isogeny. Then, for all $n \in \mathbb{N}^*$ not divisible by $\mathrm{char}(k)$ and $x, y \in A[n]$, we have:*

$$e_{\mathcal{M}_0^n}(f(x), f(y)) = e_{\mathcal{L}_0^n}(x, y)^d$$

Proof. **(i)** Let $(x, y), (x', y') \in K(\mathcal{L}) \times K(\mathcal{M})$. Then, by points (i) and (ii) of Proposition 5.1.6, we have:

$$\begin{aligned} e_{\mathcal{L}\star\mathcal{M}}((x, y), (x', y')) &= e_{\pi_1^*\mathcal{L}\otimes\pi_2^*\mathcal{M}}((x, y), (x', y')) \\ &= e_{\pi_1^*\mathcal{L}}((x, y), (x', y'))e_{\pi_2^*\mathcal{M}}((x, y), (x', y')) \\ &= e_{\mathcal{L}}(\pi_1(x, y), \pi_1(x', y'))e_{\mathcal{M}}(\pi_2(x, y), \pi_2(x', y')) \\ &= e_{\mathcal{L}}(x, x')e_{\mathcal{M}}(y, y'). \end{aligned}$$

(ii) Let $n \in \mathbb{N}^*$ not divisible by $\mathrm{char}(k)$ and $x, y \in A[n]$. Since f is a d-isogeny, it is a polarised isogeny $(A, \mathcal{L}_0^d) \longrightarrow (B, \mathcal{M}_0)$ so $f^*\mathcal{M}_0$ is algebraically equivalent to $\mathcal{L}_0^d$ by Lemma 1.4.60 and $e_{f^*\mathcal{M}_0^n} = e_{\mathcal{L}_0^{nd}}$ by points (ii) and (iv) of Proposition 5.1.6. Then, by points (i) and (v) of Proposition 5.1.6, we have:

$$e_{\mathcal{M}_0^n}(f(x), f(y)) = e_{f^*\mathcal{M}_0^n}(x, y) = e_{\mathcal{L}_0^{nd}}(x, y) = e_{\mathcal{L}_0^{nd}}(x, y) = e_{\mathcal{L}_0^n}(x, [d]y) = e_{\mathcal{L}_0^n}(x, y)^d.$$

This completes the proof. □

Proof of Lemma 6.4.1. **(i)** For all principally polarised abelian varieties C and D, we denote by $\mathcal{L}_C$ the line bundle associated to the principal polarisation on C and $\mathcal{L}_C \star \mathcal{L}_D := \pi_1^*\mathcal{L}_C \otimes \pi_2^*\mathcal{L}_D$ the product line bundle yielding a product principal polarisation on $C \times D$. Then, since $(x_1, \cdots, x_g, y_1, \cdots, y_g)$ is a ζ-symplectic basis of $A[4d]$, Proposition 5.1.6 and Lemma 6.4.2 yield the following results for the commutator pairing associated to the product polarisation $\mathcal{L}_A^{4d} \star \mathcal{L}_{B'}^{4d}$:

$$e_{\mathcal{L}_A^{4d}\star\mathcal{L}_{B'}^{4d}}(S_i, S_j) = e_{\mathcal{L}_A^{4d}}([-\alpha]\widetilde{\varphi}(y_i), [-\alpha]\widetilde{\varphi}(y_j)) = e_{\mathcal{L}_A^{4d}}(y_i, y_j)^{a\alpha^2} = 1$$

$$e_{\mathcal{L}_A^{4d}\star\mathcal{L}_{B'}^{4d}}(S_i, S_{j+g}) = e_{\mathcal{L}_A^{4d}}([-\alpha]\widetilde{\varphi}(y_i), 0)e_{\mathcal{L}_{B'}^{4d}}(0, [\beta]\psi'(x_j)) = 1$$

$$e_{\mathcal{L}_A^{4d}\star\mathcal{L}_{B'}^{4d}}(T_i, T_j) = e_{\mathcal{L}_A^{4d}}(\widetilde{\varphi}(x_i), \widetilde{\varphi}(x_j))e_{\mathcal{L}_{B'}^{4d}}(\psi'(x_i), \psi'(x_j)) = e_{\mathcal{L}_B^{4d}}(x_i, x_j)^{a+b} = 1$$

$$e_{\mathcal{L}_A^{4d}\star\mathcal{L}_{B'}^{4d}}(T_i, T_{j+g}) = e_{\mathcal{L}_A^{4d}}(\widetilde{\varphi}(x_i), [1 - \alpha d]\widetilde{\varphi}(y_j))e_{\mathcal{L}_{B'}^{4d}}(\psi'(x_i), \psi'(y_j))$$

$$= e_{\mathcal{L}_B^{4d}}(x_i, y_j)^{a(1-\alpha d)+b} = e_{\mathcal{L}_B^{4d}}(x_i, y_j)^0 = 1$$

$$e_{\mathcal{L}_A^{4d}\star\mathcal{L}_{B'}^{4d}}(S_i, T_j) = e_{\mathcal{L}_A^{4d}}([-\alpha]\widetilde{\varphi}(y_i), \widetilde{\varphi}(x_j))e_{\mathcal{L}_{B'}^{4d}}(0, \psi'(x_j)) = e_{\mathcal{L}_B^{4d}}(x_j, y_i)^{a\alpha} = \zeta^{\delta_{i,j}}$$

$$\begin{aligned} e_{\mathcal{L}_A^{4d}\star\mathcal{L}_{B'}^{4d}}(S_i, T_{j+g}) &= e_{\mathcal{L}_A^{4d}}([-\alpha]\widetilde{\varphi}(y_i), [1-\alpha d]\widetilde{\varphi}(y_j))e_{\mathcal{L}_{B'}^{4d}}(0, \psi'(y_j)) \\ &= e_{\mathcal{L}_B^{4d}}(y_i, y_j)^{-(1-\alpha d)\alpha} = 1 \end{aligned}$$

$$e_{\mathcal{L}_A^{4d}\star\mathcal{L}_{B'}^{4d}}(S_{i+g}, T_j) = e_{\mathcal{L}_A^{4d}}(0, \widetilde{\varphi}(x_j))e_{\mathcal{L}_{B'}^{4d}}([\beta]\psi'(x_i), \psi'(x_j)) = e_{\mathcal{L}_B^{4d}}(x_i, x_j)^{b\beta} = 1$$

$$\begin{aligned} e_{\mathcal{L}_A^{4d}\star\mathcal{L}_{B'}^{4d}}(S_{i+g}, T_{j+g}) &= e_{\mathcal{L}_A^{4d}}(0, [1-\alpha d]\widetilde{\varphi}(y_j))e_{\mathcal{L}_{B'}^{4d}}([\beta]\psi'(x_i), \psi'(y_j)) \\ &= e_{\mathcal{L}_B^{4d}}(x_i, y_j)^{b\beta} = e_{\mathcal{L}_B^{4d}}(x_i, y_j) = \zeta^{\delta_{i,j}}, \end{aligned}$$

for all $i, j \in [\![1\ ;\ g]\!]$. Hence, $\mathscr{B}$ is a ζ-symplectic basis of $(A \times B')[4d]$. Since $([4]x_1, \cdots, [4]x_g, [4]x_1, \cdots, [4]y_g)$ generates $A[d]$, we clearly have by Lemma 2.2.6:

$$[4]\langle T_1, \cdots, T_{2g}\rangle = \{([a]x, \psi' \circ \varphi(x)) \mid x \in A[d]\} = \ker(F).$$

This proves (i).

(ii) $\mathscr{C}_0 := (U_1, \cdots, U_{2g}, V_1, \cdots, V_{2g})$ is the product $\mathscr{C}_1 \times \mathscr{C}_2$ (as defined in Eq. (6.15)) with

$$\mathscr{C}_1 := ([d]x_1, \cdots, [d]x_g, [d]y_1, \cdots, [d]y_g)$$

$$\mathscr{C}_2 := ([d]\psi \circ \widetilde{\varphi}(x_1), \cdots, [d]\psi \circ \widetilde{\varphi}(x_g), [d\alpha\beta]\psi \circ \widetilde{\varphi}(y_1), \cdots, [d\alpha\beta]\psi \circ \widetilde{\varphi}(y_g)).$$

And we have for all $i, j \in [\![1\ ;\ g]\!]$:

$$e_{\mathcal{L}_B^4}([d]x_i, [d]x_j) = e_{\mathcal{L}_B^{4d}}(x_i, x_j)^d = 1$$

$$e_{\mathcal{L}_B^4}([d]y_i, [d]y_j) = e_{\mathcal{L}_B^{4d}}(y_i, y_j)^d = 1$$

$$e_{\mathcal{L}_B^4}([d]x_i, [d]y_j) = e_{\mathcal{L}_B^{4d}}(x_i, y_j)^d = \zeta^{d\delta_{i,j}},$$

so $\mathscr{C}_1$ is a ζ^d-symplectic basis of $B[4]$. Similarly, we verify that $\mathscr{C}_2$ is a ζ^d-symplectic basis of $A'[4]$. It follows by Lemma 6.4.2.(i) that $\mathscr{C} = \mathscr{C}_1 \times \mathscr{C}_2$ is a product ζ^d-symplectic basis of $(B \times A')[4]$, as desired.

(iii) The fact that $([d]F(S_1), \cdots, [d]F(S_{2g}), F(T_1), \cdots, F(T_{2g}))$ is a ζ^d-symplectic basis of $(B \times A')[4]$ follows from Theorem 6.1.1. It remains to compute for all $i \in [\![1\ ;\ g]\!]$:

$$\begin{aligned} [d]F(S_i) &= [d]F([-\alpha]\widetilde{\varphi}(y_i), 0) = ([-d\alpha]\varphi \circ \widetilde{\varphi}(y_i), [d\alpha]\psi \circ \widetilde{\varphi}(y_i)) \\ &= ([-da\alpha]y_i, [d\alpha]\psi \circ \widetilde{\varphi}(y_i)) = -V_i + [b]V_{i+g} \end{aligned}$$

$$\begin{aligned} [d]F(S_{i+g}) &= [d]F((0, [\beta]\psi'(x_i)) = ([d\beta]\widetilde{\psi'} \circ \psi'(x_i), [d\beta]\widetilde{\varphi'} \circ \psi'(x_i)) \\ &= ([d\beta b]x_i, [d\beta]\psi \circ \widetilde{\varphi}(x_i)) = ([d]x_i, [d\beta]\psi \circ \widetilde{\varphi}(x_i)) = U_i + [\beta]U_{i+g} \end{aligned}$$

$$\begin{aligned} F(T_i) &= F(\widetilde{\varphi}(x_i), \psi'(x_i)) = (\varphi \circ \widetilde{\varphi}(x_i) + \widetilde{\psi'} \circ \psi'(x_i), -\psi \circ \widetilde{\varphi}(x_i) + \widetilde{\varphi'} \circ \psi'(x_i)) \\ &= ([a+b]x_i, -\psi \circ \widetilde{\varphi}(x_i) + \psi \circ \widetilde{\varphi}(x_i)) = ([d]x_i, 0) = U_i \end{aligned}$$

$$F(T_{i+g}) = F([1-\alpha d]\widetilde{\varphi}(y_i), \psi'(y_i))$$

$$= ([1-\alpha d]\varphi \circ \widetilde{\varphi}(y_i) + \widetilde{\psi'} \circ \psi'(y_i), -[1-\alpha d]\psi \circ \widetilde{\varphi}(y_i) + \widetilde{\varphi'} \circ \psi'(y_i))$$
$$= ([a(1-\alpha d)+b]y_i, [\alpha d - 1 + 1]\psi \circ \widetilde{\varphi}(y_i)) = (0, [d]\psi \circ \widetilde{\varphi}(y_i)) = [b]V_{i+g},$$

where we used twice the fact that $\psi' \circ \varphi = \varphi' \circ \psi$ implies that $\widetilde{\varphi'} \circ \psi' = \psi \circ \widetilde{\varphi}$. The inverse equations:

$$U_i = F(T_i), \quad U_{i+g} = [bd]F(S_{i+g}) - [b]F(T_i),$$
$$V_i = F(T_{i+g}) - [d]F(S_i), \quad V_{i+g} = [\beta]F(T_{i+g}),$$

follow immediately. This completes the proof. □

In some contexts, it is easier to write kernel points of F as $([a]x, \psi' \circ \varphi(x))$ instead of $(\widetilde{\varphi}(x), \psi'(x))$ as in Lemma 6.4.1. We can adapt the result to this context.

Lemma 6.4.3. *Consider an (a,b)-isogeny diamond and an isogeny $F : A \times B' \longrightarrow B \times A'$ as in Lemma 6.4.1. Let $\zeta \in k^*$ be a $4d$-th primitive root of unity and $(x_1, \cdots, x_g, y_1, \cdots, y_g)$ be a ζ-symplectic basis of $A[4d]$. Let α and β be modular inverses of a and b modulo $4d$ respectively.*

(i) For all $i \in [\![1\ ;\ g]\!]$, we denote:

$$S_i := ([-\alpha]y_i, 0), \quad S_{i+g} := (0, [\beta]\psi' \circ \varphi(x_i)),$$
$$T_i := ([a]x_i, \psi' \circ \varphi(x_i)), \quad T_{i+g} := ([1-\alpha d]y_i, [\alpha]\psi' \circ \varphi(y_i)).$$

Then $\mathscr{B} := (S_1, \cdots, S_{2g}, T_1, \cdots, T_{2g})$ is a ζ-symplectic basis of $(A \times B')[4d]$ adapted to F i.e. *such that* $\ker(F) = [4]\langle T_1, \cdots, T_{2g}\rangle$.

(ii) For all $i \in [\![1\ ;\ g]\!]$, we denote:

$$U_i := ([d]\varphi(x_i), 0), \quad U_{i+g} := (0, [ad]\psi(x_i)),$$
$$V_i := ([\alpha d]\varphi(y_i), 0), \quad V_{i+g} := (0, [d\alpha\beta]\psi(y_i)).$$

Then $\mathscr{C}_0 := (U_1, \cdots, U_{2g}, V_1, \cdots, V_{2g})$ is a product ζ^d-symplectic basis of $(B \times A')[4]$.

(iii) The ζ^d-symplectic basis $\mathscr{C} := ([d]F(S_1), \cdots, [d]F(S_{2g}), F(T_1), \cdots, F(T_{2g}))$ of $(B \times A')[4]$ naturally induced by F via Theorem 6.1.1 is related to $\mathscr{C}_0$ by the following formulas. For all $i \in [\![1\ ;\ g]\!]$:

$$[d]F(S_i) = -V_i + [b]V_{i+g}, \quad [d]F(S_{i+g}) = U_i + [\beta]U_{i+g},$$
$$F(T_i) = U_i, \quad F(T_{i+g}) = [b]V_{i+g},$$

so that:

$$U_i = F(T_i), \quad U_{i+g} = [bd]F(S_{i+g}) - [b]F(T_i),$$
$$V_i = F(T_{i+g}) - [d]F(S_i), \quad V_{i+g} = [\beta]F(T_{i+g}).$$

Proof. We apply Lemma 6.4.1 to the ζ-symplectic basis $(\varphi(x_1), \cdots, \varphi(x_g), [\alpha]\varphi(y_1), \cdots, [\alpha]\varphi(y_g))$ of $B[4d]$. □

When $d = 2^e$, Lemmas 6.4.1 and 6.4.3 yield a way to compute the change of theta coordinates on the domain to obtain level 2 theta coordinates adapted to F. Indeed, using point (i) of this lemma, we can compute the symplectic change of basis matrix mentioned in Remark 6.3.3 and then apply Theorem 6.2.10 to compute the adapted level 2 theta coordinates.

In general, once F has been computed as a chain of 2-isogenies, the resulting theta coordinates on the codomain $B \times A'$ in which image theta points of F are expressed are non-product theta coordinates of $B \times A'$. The theta null point of $B \times A'$ is also not the product of theta null points of B and A'. Nonetheless, for applications, it is convenient to:

- Extract the theta null points of A' and B from the theta null point of $B \times A'$ since they give (almost) enough information to recover A' and B. This is in particular useful when A' and B are unknown at the start.

- Express images points of $F(P) = (F_B(P), F_{A'}(P))$ in the product $B \times A'$.

Fortunately, point (iii) of Lemma 6.4.1 also expresses the symplectic basis of the 4-torsion $\mathscr{C}$ naturally induced by F on its codomain in a product symplectic basis $\mathscr{C}_0$ of the 4-torsion. By Theorem 6.2.10, we obtain change of theta coordinates on the codomain towards product theta coordinates (induced by $\mathscr{C}_0$). After this change of theta coordinates, points $(x, y) \in B(k) \times A'(k)$ are then expressed as

$$\theta^{B\times A'}_{i,j}(x, y) = \theta^B_i(x) \cdot \theta^{A'}_j(y)$$

for all $(i, j) \in (\mathbb{Z}/2\mathbb{Z})^{2g}$. Hence, we can use Algorithm 6.9 on the product theta null point of $B \times A'$ to extract the theta null points of A' and B and on an image product theta point $F(P)$ to decompose it as $(F_B(P), F_{A'}(P))$ in $B \times A'$.

Algorithm 6.9: Extract components of a product theta point.

Data: Product level 2 theta coordinates $(\theta^{A\times B}_{i,j}(x, y))_{\substack{i\in(\mathbb{Z}/2\mathbb{Z})^{g_1}\\ j\in(\mathbb{Z}/2\mathbb{Z})^{g_2}}}$ of a point $(x, y) \in A(k) \times B(k)$.

Result: Projective level 2 theta points $(\theta^A_i(x))_{i\in(\mathbb{Z}/2\mathbb{Z})^{g_1}}$ and $(\theta^B_j(y))_{j\in(\mathbb{Z}/2\mathbb{Z})^{g_2}}$ in each component.

1 Find $(i_0, j_0) \in (\mathbb{Z}/2\mathbb{Z})^{g_1} \times (\mathbb{Z}/2\mathbb{Z})^{g_2}$ such that $\theta^{A\times B}_{i_0,j_0}(x, y) \neq 0$;
2 $(\theta^A_i(x))_{i\in(\mathbb{Z}/2\mathbb{Z})^{g_1}} \longleftarrow (\theta^{A\times B}_{i,j_0}(x, y))_{i\in(\mathbb{Z}/2\mathbb{Z})^{g_1}}$;
3 $(\theta^B_j(y))_{j\in(\mathbb{Z}/2\mathbb{Z})^{g_2}} \longleftarrow (\theta^{A\times B}_{i_0,j}(x, y))_{j\in(\mathbb{Z}/2\mathbb{Z})^{g_2}}$;
4 **return** $(\theta^A_i(x))_{i\in(\mathbb{Z}/2\mathbb{Z})^{g_1}}$ and $(\theta^B_j(y))_{j\in(\mathbb{Z}/2\mathbb{Z})^{g_2}}$;

6.4.2 Change of theta coordinates on the domain and codomain with half available torsion

In this section, we assume that we want to compute a 2^e-isogeny $F : \mathcal{A} \longrightarrow \mathcal{B}$ obtained from Kani's lemma but we can only access the 2^f-torsion of $\mathcal{A}$ and $\mathcal{B}$ with $f \geq e/2 + 2$, as in SQIsignHD. In that case, we assume that both $\mathcal{A}$ and $\mathcal{B}$ are known and we decompose $F = F_2 \circ F_1$ where $F_1 : \mathcal{A} \longrightarrow \mathcal{C}$ and $F_2 : \mathcal{C} \longrightarrow \mathcal{B}$ are respectively a 2^{e_1} and 2^{e_2}-isogeny, with $e = e_1 + e_2$ and $e_1, e_2 \leq f - 2$. We compute F_1 using 2^{e_1+2}-torsion points lying above $\ker(F_1)$ and $\widetilde{F}_2$ using 2^{e_2+2}-torsion points lying above $\ker(\widetilde{F}_2)$. Then we easily infer $\widetilde{\widetilde{F}}_2 = F_2$ from Section 6.2.3 and we are able to evaluate $F = F_2 \circ F_1$.

However, to be able to compose F_1 with F_2, we need the level 2 theta structures on $\mathcal{C}$ seen as the codomain of F_1 and as the domain of F_2 to be the same. Lemma 6.2.11 and its proof ensure that the latter theta structure is the dual of the theta structure induced by $\widetilde{F}_2$. Hence, we need the level 2 theta structures induced by F_1 and $\widetilde{F}_2$ to be dual of each other *i.e.* we need their associated symplectic 4-torsion basis to be related by the symplectic matrix

$$\begin{pmatrix} 0 & -I_g \\ I_g & 0 \end{pmatrix} \in \mathrm{Sp}_{2g}(\mathbb{Z}/4\mathbb{Z}),$$

so the associated theta coordinates are dual of one another in the sense of Corollary 6.1.3. Hence, we need two symplectic basis on the domain and codomain of F *adapted to the decomposition* $F = F_2 \circ F_1$ in the sense of the following definition.

Definition 6.4.4. Let $d_1, d_2 \in \mathbb{N}^*$ not divisible by $\mathrm{char}(k)$, $d := d_1 d_2$ and let $F : \mathcal{A} \longrightarrow \mathcal{B}$ be a d-isogeny between principally polarised abelian varieties of dimension g written as $F = F_2 \circ F_1$ where $F_1 : \mathcal{A} \longrightarrow \mathcal{C}$ and $F_2 : \mathcal{C} \longrightarrow \mathcal{B}$ are d_1 and d_2-isogenies respectively. Let $\zeta_i \in k^*$ be a primitive $4d_i$-th root of unity for $i \in \{1, 2\}$. We say that a ζ_1-symplectic basis $\mathscr{B}_1 := (S_1, \cdots, S_g, T_1, \cdots, T_g)$ of $\mathcal{A}[4d_1]$ and a ζ_2-symplectic basis $\mathscr{B}_2 := (S'_1, \cdots, S'_g, T'_1, \cdots, T'_g)$ of $\mathcal{B}[4d_2]$ are *adapted to the decomposition* $F = F_2 \circ F_1$ if:

(i) $\mathscr{B}_1$ is adapted to F_1 *i.e.* $\ker(F_1) = \langle [4]T_1, \cdots, [4]T_g \rangle$.

(ii) $\mathscr{B}_2$ is adapted to $\widetilde{F}_2$ *i.e.* $\ker(\widetilde{F}_2) = \langle [4]T'_1, \cdots, [4]T'_g \rangle$.

(iii) For all $i \in [\![1\ ;\ g]\!]$, $[d_2]\widetilde{F}_2(S'_i) = F_1(T_i)$ and $\widetilde{F}_2(T'_i) = -[d_1]F_1(S_i)$.

Lemma 6.4.5. *Let us keep the notations from Definition 6.4.4. Let $d', c_1, c_2 \in \mathbb{N}^*$ such that $d' = c_1d_1 = c_2d_2$ and let $\zeta \in k^*$ be a primitive $4d'$-th root of unity. Consider a ζ-symplectic basis $\mathscr{B}_1 := (S_1, \cdots, S_g, T_1, \cdots, T_g)$ of $\mathcal{A}[4d']$ and a ζ-symplectic basis $\mathscr{B}_2 := (S'_1, \cdots, S'_g, T'_1, \cdots, T'_g)$ of $\mathcal{B}[4d']$ such that:*

(i) $[c_1]\mathscr{B}_1$ is adapted to F_1 i.e. $\ker(F_1) = \langle [4c_1]T_1, \cdots, [4c_1]T_g\rangle$.

(ii) $[c_2]\mathscr{B}_2$ is adapted to $\widetilde{F}_2$ i.e. $\ker(\widetilde{F}_2) = \langle [4c_2]T'_1, \cdots, [4c_2]T'_g\rangle$.

(iii) For all $i \in [\![1\ ;\ g]\!]$, $T'_i = -F(S_i)$ and $T_i = \widetilde{F}(S'_i)$.

Then, $[c_1]\mathscr{B}_1$ and $[c_2]\mathscr{B}_2$ are adapted to the decomposition $F = F_2 \circ F_1$.

Proof. Point (iii) ensures that for all $i \in [\![1\ ;\ g]\!]$,

$$\widetilde{F}_2([c_2]T'_i) = -[c_2]\widetilde{F}_2 \circ F(S_i) = -[c_2]\widetilde{F}_2 \circ F_2 \circ F_1(S_i) = -[c_2d_2]F_1(S_i) = -[d_1]F_1([c_1]S_i),$$

and

$$F_1([c_1]T_i) = [c_1]F_1 \circ \widetilde{F}(S'_i) = [c_1]F_1 \circ \widetilde{F}_1 \circ \widetilde{F}_2(S'_i) = [c_1d_1]\widetilde{F}_2(S'_i) = [d_2]\widetilde{F}_2([c_2]S'_i).$$

Hence, $[c_1]\mathscr{B}_1$ and $[c_2]\mathscr{B}_2$ satisfy conditions (i), (ii) and (iii) of Definition 6.4.4 so they are compatible with the decomposition $F = F_2 \circ F_1$. □

Lemma 6.4.6. *Let us keep the notations from Lemma 6.4.1. We consider the d-isogeny*

$$F := \begin{pmatrix} \varphi & \widetilde{\psi'} \\ -\psi & \widetilde{\varphi'} \end{pmatrix} : A \times B' \longrightarrow B \times A',$$

obtained from an (a, b)-isogeny diamond with $d = a + b$. Let $d_1, d_2 \in \mathbb{N}^*$ not divisible by $\mathrm{char}(k)$ such that $d = d_1d_2$ and let $d', c_1, c_2 \in \mathbb{N}^*$ such that $d' = c_1d_1 = c_2d_2$. Let $\alpha, \beta \in \mathbb{N}^*$ be modular inverses of a and b modulo $4d'$, $\zeta \in k^*$ be a $4d'$-th primitive root of unity and $(x_1, \cdots, x_g, y_1, \cdots, y_g)$ be a ζ-symplectic basis of $B[4d']$.

Define $\mathscr{B}_1 := (S_1, \cdots, S_{2g}, T_1, \cdots, T_{2g})$ as

$$S_i := (\widetilde{\varphi}(y_i), 0), \quad S_{i+g} := (0, \psi'(x_i))$$

$$T_i := (-[\alpha]\widetilde{\varphi}(x_i), -[\alpha]\psi'(x_i)), \quad T_{i+g} := (-[\alpha]\widetilde{\varphi}(y_i), [\beta]\psi'(y_i)),$$

for all $i \in [\![1\ ;\ g]\!]$, and $\mathscr{B}_2 := (S'_1, \cdots, S'_{2g}, T'_1, \cdots, T'_{2g})$ as

$$S'_i := (-[\alpha]x_i, 0), \quad S'_{i+g} := (0, [\alpha\beta]\psi \circ \widetilde{\varphi}(y_i))$$

$$T'_i := (-[a]y_i, \psi \circ \widetilde{\varphi}(y_i)), \quad T'_{i+g} := (-[b]x_i, -\psi \circ \widetilde{\varphi}(x_i)),$$

for all $i \in [\![1\ ;\ g]\!]$. Then:

(i) $[c_1]\mathscr{B}_1$ is a ζ^{c_1}-symplectic basis of $(A \times B')[4d_1]$ compatible with F_1.

(ii) $[c_2]\mathscr{B}_2$ is a ζ^{c_2}-symplectic basis of $(B \times A')[4d_2]$ compatible with $\widetilde{F}_2$.

(iii) $[c_1]\mathscr{B}_1$ and $[c_2]\mathscr{B}_2$ are compatible with the decomposition $F = F_2 \circ F_1$.

Proof. **(i)** We can prove that $\mathscr{B}_1$ is a ζ-symplectic basis of $(A \times B')[4d']$ with similar pairing computations as in the proof of Lemma 6.4.1.(i). It follows that $[c_1]\mathscr{B}_1$ is a ζ^{c_1}-symplectic basis of $(A \times B')[4d_1]$. Then by Kani's lemma (Lemma 2.2.6) and Lemma 6.3.1.(ii), we have:

$$\ker(F_1) = [d_2]\ker(F) = \{[d_2](\widetilde{\varphi}(x), \psi'(x)) \mid x \in B[d]\} = \{(\widetilde{\varphi}(x), \psi'(x)) \mid x \in B[d_1]\}.$$

Besides, for all $i \in [\![1\ ;\ g]\!]$,

$$[4c_1]T_i = (\widetilde{\varphi}(-[4\alpha c_1]x_i), \psi'(-[4\alpha c_1]x_i)) \quad \text{and} \quad [4c_1]T_{i+g} := (\widetilde{\varphi}(-[4\alpha c_1]y_i), \psi'([4\beta c_1]y_i)).$$

Since $b = d - a \equiv -a \mod d_1$, we also have $\beta \equiv -\alpha \mod d_1$ so for all $i \in [\![1\ ;\ g]\!]$, we have $[4c_1]T_{i+g} := (\widetilde{\varphi}(-[4\alpha c_1]y_i), \psi'(-[4\alpha c_1]y_i))$. In addition, $-[4\alpha c_1]x_1, \cdots, -[4\alpha c_1]x_g, -[4\alpha c_1]y_1, \cdots, -[4\alpha c_1]y_g$ generate $[-4\alpha c_1]B[4d'] = B[d_1]$ so that,

$$\ker(F_1) = \langle [4c_1]T_1, \cdots, [4c_1]T_{2g} \rangle,$$

and $[c_1]\mathscr{B}_1$ is adapted to F_1.

(ii) We can prove that $\mathscr{B}_2$ is a ζ-symplectic basis of $(B \times A')[4d']$ with similar pairing computations as in the proof of Lemma 6.4.1.(i). It follows that $[c_2]\mathscr{B}_2$ is a ζ^{c_2}-symplectic basis of $(B \times A')[4d_2]$. By Lemma 2.2.9, we have

$$\ker(\widetilde{F}) = \{([a]x, -\psi \circ \widetilde{\varphi}(x)) \mid x \in B[d]\},$$

so that

$$\ker(\widetilde{F}_2) = [d_1]\ker(\widetilde{F}) = \{([a]x, -\psi \circ \widetilde{\varphi}(x)) \mid x \in B[d_2]\}.$$

Since $b = d - a \equiv -a \mod d_2$. Since $[4c_2]x_1, \cdots, [4c_2]x_g, -[4c_2]y_1, \cdots, -[4c_2]y_g$ generate $[-4c_2]B[4d'] = B[d_2]$, we obtain that

$$\ker(\widetilde{F}_2) = \langle [4c_2]T'_1, \cdots, [4c_2]T'_{2g} \rangle,$$

and $[c_2]\mathscr{B}_2$ is adapted to $\widetilde{F}_2$.

(iii) For all $i \in [\![1\ ;\ g]\!]$,

$$F(S_i) = F(\widetilde{\varphi}(y_i), 0) = (\varphi \circ \widetilde{\varphi}(y_i), -\psi \circ \widetilde{\varphi}(y_i)) = ([a]y_i, -\psi \circ \widetilde{\varphi}(y_i)) = -T_i$$

$$F(S_{i+g}) = F(0, \psi'(x_i)) = (\widetilde{\psi'} \circ \psi'(x_i), \widetilde{\varphi'} \circ \psi'(y_i)) = ([b]x_i, \psi \circ \widetilde{\varphi}(x_i)) = -T_{i+g},$$

where we used that $\psi' \circ \varphi = \varphi' \circ \psi$ since $\varphi, \varphi', \psi, \psi'$ form an (a, b)-isogeny diamond, so that $\widetilde{\varphi'} \circ \psi' = \psi \circ \widetilde{\varphi}$. For all $i \in [\![1\ ;\ g]\!]$, we also have:

$$\widetilde{F}(S'_i) = \widetilde{F}(-[\alpha]x_i, 0) = (-[\alpha]\widetilde{\varphi}(x_i), -[\alpha]\psi'(x_i)) = T_i$$

$$\begin{aligned}\widetilde{F}(S'_{i+g}) &= \widetilde{F}(0, [\alpha\beta]\psi \circ \widetilde{\varphi}(y_i)) = (-[\alpha\beta]\widetilde{\psi} \circ \psi \circ \widetilde{\varphi}(y_i), [\alpha\beta]\varphi' \circ \psi \circ \widetilde{\varphi}(y_i)) \\ &= (-[b\alpha\beta] \circ \widetilde{\varphi}(y_i), [a\alpha\beta]\psi'(y_i)) = (-[\alpha] \circ \widetilde{\varphi}(y_i), [\beta]\psi'(y_i)) = T_{i+g},\end{aligned}$$

where we use that $\psi' \circ \varphi = \varphi' \circ \psi$ implies $\varphi' \circ \psi \circ \widetilde{\varphi} = [a]\psi'$. Hence, by Lemma 6.4.5, $[c_1]\mathscr{B}_1$ and $[c_2]\mathscr{B}_2$ are compatible with the decomposition $F = F_2 \circ F_1$. □

6.5 Implementation in dimension 2

In this section, we explain specifically how to compute a 2^e-isogeny $F : E_1 \times E_2 \longrightarrow E_3 \times E_4$ between elliptic products in dimension 2 as a chain of 2-isogenies:

$$E_1 \times E_2 \xrightarrow{f_1} A_1 \xrightarrow{f_2} A_2 \quad \cdots \quad A_{e-2} \xrightarrow{f_{e-1}} A_{e-1} \xrightarrow{f_e} E_3 \times E_4.$$

This presentation follows from two works that I coauthored: a paper [DMPR25] and the NIST round 2 SQIsign specification [AAA+25]. We specialize the algorithms from Sections 6.1 to 6.3 to this case with several optimisations. In particular, we are able to relax the need to access 2^{e+2}-torsion points above $\ker(F)$. Generators of $\ker(F) \subset (E_1 \times E_2)[2^e]$ are sufficient, at the expense of square root computations. We also propose algorithms to compute the codomain theta null point of each 2-isogeny of the chain and to evaluate them with less arithmetic operations than the fully general algorithms from Section 6.1. Implementation results are presented in Section 6.5.5.

Assume that we are given $T_1, T_2 \in (E_1 \times E_2)[2^{e+2}]$ forming an isotropic subgroup such that $\ker(F) = \langle [4]T_1, [4]T_2 \rangle$. Then, the 2-isogeny chain F can be computed as follows:

1. Convert Montgomery $(X : Z)$-coordinates of the product $E_1 \times E_2$ into level 2 theta coordinates adapted to f_1, as we shall explain in Section 6.5.1.

2. Compute the gluing isogeny $f_1 : E_1 \times E_2 \longrightarrow A_1$ from 8-torsion points $[2^{e-1}]T_1$ and $[2^{e-1}]T_2$. This will be explained in Section 6.5.3.

3. For all $i \in [\![1\ ;\ e-1]\!]$, compute $f_{i+1} : A_i \longrightarrow A_{i+1}$ from 8-torsion points $[2^{e-i-1}]f_i \circ \cdots \circ f_1(T_1)$ and $[2^{e-i-1}]f_i \circ \cdots \circ f_1(T_2)$ with generic isogeny computation algorithms from Section 6.5.2. As explained in Section 6.3.3, a divide and conquer strategy can be used to minimize the number of 2-isogeny evaluations and duplications.

4. Convert codomain level 2 theta coordinates into Montgomery $(X : Z)$-coordinates of $E_3 \times E_4$. This will be explained in Section 6.5.4.

When we are given kernel generators $T_1, T_2 \in (E_1 \times E_2)[2^e]$ only, we proceed as above up to the following modifications:

- In step 2 above, f_1 is obtained from the 8-torsion points $[2^{e-3}]T_1$ and $[2^{e-3}]T_2$.
- In step 3 above, for all $i \in [\![1\ ;\ e-3]\!]$, f_{i+1} is obtained from the 8-torsion points $[2^{e-i-3}]f_i \circ \cdots \circ f_1(T_1)$ and $[2^{e-i-3}]f_i \circ \cdots \circ f_1(T_2)$.
- In step 3 above, f_{e-1} is obtained from 4-torsion points $f_{e-2} \circ \cdots \circ f_1(T_1)$ and $f_{e-2} \circ \cdots \circ f_1(T_2)$ with one square root computation. This will be explained in Section 6.5.2.
- In step 3 above, f_e is obtained from 3 square root computations. This will be explained in Section 6.5.2.

6.5.1 Computing an adapted theta structure on the domain

The general method

We can apply the general method explained in Section 6.3.1. We choose a 4-th primitive root of unity $\zeta_4 \in k^*$ and for $i \in \{1, 2\}$, we convert E_i into Montgomery form and find a basis (P_i, Q_i) of $E_i[4]$ with $x(Q_i) = -1$ and $e_4(P_i, Q_i) = \zeta_4$. Such a basis will be called $(\zeta_4$-*special*. We then use Proposition 5.3.47 to convert Montgomery $(X : Z)$-coordinates into level 2 theta coordinates associated to the theta structure induced by (P_i, Q_i). We can then compute product level 2 theta coordinates on $E_1 \times E_2$ given by:

$$\forall (P, Q) \in E_1 \times E_2, \forall i, j \in \mathbb{Z}/2\mathbb{Z}, \quad \theta_{i,j}^{E_1 \times E_2}(P, Q) = \theta_i^{E_1}(P)\theta_j^{E_2}(Q),$$

associated to the product theta structure induced by the product ζ_4-symplectic basis of $(E_1 \times E_2)[4]$:

$$\mathscr{B}_0 := ((P_1, 0), (0, P_2), (Q_1, 0), (0, Q_2)).$$

The conversion formula to product theta coordinates $(\theta_{00}, \theta_{10}, \theta_{01}, \theta_{11})$ can be written as a matrix $N_0 \in M_4(k)$ acting by left multiplication on the product $(X_1X_2, X_1Z_2, Z_1X_2, Z_1Z_2)$ of Montgomery $(X : Z)$-coordinates of $E_1 \times E_2$. We compute the level 2 theta null points of E_i given by $(a_i : b_i) := (X_i(P_i) + Z_i(P_i) : X_i(P_i) - Z_i(P_i))$ for $i \in \{1, 2\}$. Then, we have:

$$N_0 := \begin{pmatrix} a_1a_2 & -a_1a_2 & -a_1a_2 & a_1a_2 \\ b_1a_2 & -b_1a_2 & b_1a_2 & -b_1a_2 \\ a_1b_2 & a_1b_2 & -a_1b_2 & -a_1b_2 \\ b_1b_2 & b_1b_2 & b_1b_2 & b_1b_2 \end{pmatrix}. \tag{6.18}$$

To obtain theta coordinates $(\theta'^{E_1 \times E_2}_{i,j})_{i,j \in \mathbb{Z}/2\mathbb{Z}}$ adapted to F, we compute the symplectic matrix $M \in \mathrm{Sp}_4(\mathbb{Z}/4\mathbb{Z})$ from $\mathscr{B}_0$ to a ζ_4-symplectic basis $\mathscr{B}$ of $(E_1 \times E_2)[4]$ adapted to F. This matrix M can be obtained from kernel generators of F and discrete logarithm computations as explained in Remark 6.3.3. Alternatively, when F is obtained from Kani's lemma, Lemma 6.4.1 can be used. From M, ζ_4 and Theorem 6.2.10, we then obtain a matrix N_1 converting product theta coordinates $(\theta_{i,j}^{E_1 \times E_2})_{i,j \in \mathbb{Z}/2\mathbb{Z}}$ into theta coordinates $(\theta'^{E_1 \times E_2}_{i,j})_{i,j \in \mathbb{Z}/2\mathbb{Z}}$ adapted to F by multiplication on the left. The conversion formula directly from product $(X_1X_2, X_1Z_2, Z_1X_2, Z_1Z_2)$ of Montgomery $(X : Z)$-coordinates of $E_1 \times E_2$ is then given by the left multiplication by $N := N_1 \cdot N_0$.

Robert's method: on an elliptic curve

In [DMPR25, § 2.3], another method due to Damien Robert has been introduced to compute the conversion matrix N from product of Montgomery $(X : Z)$-coordinates on $E_1 \times E_2$ to level 2 theta coordinates adapted to F. The general idea is to use the properties of the theta group action on global sections, as in the proof of Proposition 6.2.4.

We first explain how this method works on a Montgomery curve E defined over k. Let (T_1', T_2') be a basis of $E[4]$. This basis determines a symmetric level 2 theta structure on $(E, \mathcal{L}_0^2)$ (with $\mathcal{L}_0 := \mathcal{L}((0_E))$) by Remark 5.3.31. Recall how this theta structure is constructed explicitly (specifying the proof of Theorem 5.3.30). Given a point of 2-torsion T, a 4-torsion point $T' \in E[4]$ such that $[2]T' = T$ determines a *symmetric element* (as defined in Definition 5.3.9) $\mathfrak{g}_T$ over T in the theta group $G(\mathcal{L}_0^2)$, as follows. By Lemma 5.3.14, there are two symmetic elements $\pm\mathfrak{g}_{T'} \in G(\mathcal{L}_0^4)$ over T'. They both have the same image $\mathfrak{g}_T := \eta_2(\pm\mathfrak{g}_{T'}) \in G(\mathcal{L}_0^2)$ (where η_2 has been defined in Definition 5.3.17), which is symmetric and lies over T. Let $T_1 = [2]T_1'$, $T_2 = [2]T_2'$. Then T_1', T_2' determine symmetric elements $\mathfrak{g}_1, \mathfrak{g}_2$ above T_1, T_2 respectively. Then $\mathfrak{g}_1, \mathfrak{g}_2$ form a symmetric level structure and determine a symmetric theta structure $\Theta_{\mathcal{L}_0^2}$ on $G(\mathcal{L}_0^2)$.

As we have seen in the proof of Proposition 6.2.4, the associated theta coordinate θ_0 is invariant under the action of $\mathfrak{g}_2$, so it is equal to the trace $\theta_0 := 1 \cdot s + \mathfrak{g}_2 \cdot s$, for any global section $s \in \Gamma(E, \mathcal{L}_0^2)$ such that the trace is non-zero. We also have $\theta_1 = \mathfrak{g}_1 \cdot \theta_0$. Hence, to determine θ_0 we need to study the action of symmetric elements on the global sections X and Z generating $\Gamma(E, \mathcal{L}_0^2)$.

Lemma 6.5.1. *Let $T' := (x_{T'} : y_{T'} : z_{T'})$ be a point of 4-torsion and $T := [2]T' := (x_T : y_T : z_T)$. Let $\mathfrak{g}_T \in G(\mathcal{L}_0^2)$ be the symmetric element determined by T'. Then*

$$\mathfrak{g}_T \cdot X = \frac{x_T z_{T'} \cdot X + (z_T x_{T'}^2 / z_{T'} - 2x_T x_{T'}) \cdot Z}{x_{T'} z_T - z_{T'} x_T} \quad \text{and} \quad \mathfrak{g}_T \cdot Z = \frac{z_T z_{T'} \cdot X - x_T z_{T'} \cdot Z}{x_{T'} z_T - z_{T'} x_T}.$$

Proof. We first treat the case when $T' := (1 : * : 1)$ and $T := (0 : 0 : 1)$. By Lemma 5.1.41, $\mathfrak{g}_T$ acts on global sections of $\Gamma(E, \mathcal{L}_0^2)$ as the translation by T up to a projective constant *i.e.* there exists $\lambda \in k^*$ such that $\mathfrak{g}_T \cdot s = \lambda t_T^* s = \lambda s \circ t_T$ for all $s \in \Gamma(E, \mathcal{L}_0^2)$. Since the translation by T maps $(X : Z)$ to $(Z : X)$, we have $\mathfrak{g}_T \cdot X = \lambda Z$ and $\mathfrak{g}_T \cdot Z = \lambda X$ and $\lambda^2 = 1$ because $\mathfrak{g}_T$ has order 2. Hence, $\lambda = \pm 1$, $\mathfrak{g}_T \cdot X = \pm Z$ and $\mathfrak{g}_T \cdot Z = \pm X$. It can be proved that the right sign choice is $\lambda = 1$ (see Remark 6.5.2) but we shall admit this result for now.

In plain generality, we map the Montgomery point $(x_{T'} : z_{T'})$ to $(1 : 1)$ and $(x_T : z_T)$ to $(0 : 1)$ via the homography $(X : Z) \longmapsto (X' : Z') := (z_{T'} z_T \cdot X - z_{T'} x_T \cdot Z : (x_{T'} z_T - z_{T'} x_T) \cdot Z)$. Using this change of variables, we obtain:

$$\begin{cases} \mathfrak{g}_T \cdot X' = Z' & \textit{i.e.} \quad z_{T'} z_T \cdot \mathfrak{g}_T \cdot X - z_{T'} x_T \cdot \mathfrak{g}_T \cdot Z = (x_{T'} z_T - z_{T'} x_T) \cdot Z \\ \mathfrak{g}_T \cdot Z' = X' & \textit{i.e.} \quad (x_{T'} z_T - z_{T'} x_T) \cdot \mathfrak{g}_T \cdot Z = z_{T'} z_T \cdot X - z_{T'} x_T \cdot Z \end{cases}$$

The result follows. □

Remark 6.5.2. As in Proposition 5.3.47, consider a special basis (T_1', T_2') of $E[4]$ with $T_2' = (-1 : * : 1)$ of double $T_2 = (0 : 0 : 1)$ and $T_1' := (a + b : * : a - b)$ of double $T_1 = (a^2 + b^2 : 0 : a^2 - b^2)$. Then the symmetric element above T_2 induced by T_2' acts by $\mathfrak{g}_2 \cdot X = -Z$ and $\mathfrak{g}_2 \cdot Z = -X$ by Lemma 6.5.1. Taking the trace of X under this action we get $\theta_0 = 1 \cdot X + \mathfrak{g}_2 \cdot X = X - Z$.

We then compute by Lemma 6.5.1 again:

$$\theta_1 = \mathfrak{g}_1 \cdot \theta_0 = \mathfrak{g}_1 \cdot (X - Z) = \frac{z_{T_1'}(x_{T_1} - z_{T_1})}{z_{T_1} x_{T_1'} - x_{T_1} z_{T_1'}} X + \frac{z_{T_1} x_{T_1'}^2 / z_{T_1'} - 2x_{T_1} x_{T_1'} + x_{T_1} z_{T_1'}}{z_{T_1} x_{T_1'} - x_{T_1} z_{T_1'}} Z = \frac{b}{a} X + \frac{b}{a} Z.$$

We recover the same conversion formula between Montgomery and theta coordinates as in Proposition 5.3.47. This proves in particular that the sign choice $\lambda = 1$ in the proof of Lemma 6.5.1 was the right one.

Robert's method: on a product of two elliptic curves

Now, we explain how this method generalizes to an elliptic product $E_1 \times E_2$. Let $\mathcal{L}_i := \mathcal{L}((0_{E_i}))$ be the line bundle inducing the principal polarisation of E_i for $i \in \{1, 2\}$ and consider the product

$\mathcal{L}_1 \star \mathcal{L}_2 := \pi_1^*\mathcal{L}_1 \otimes \pi_2^*\mathcal{L}_2$, where π_i is the projection on the i-th component of $E_1 \times E_2$ for $i \in \{1,2\}$. We have seen in Section 5.3.2 that $G(\mathcal{L}_1^2 \star \mathcal{L}_2^2) \simeq G(\mathcal{L}_1^2) \times G(\mathcal{L}_2^2)/\{(\lambda_1,\lambda_2) \in (k^*)^2 \mid \lambda_1\lambda_2 = 1\}$ via the map:

$$\begin{array}{rcl} G(\mathcal{L}_1^2) \times G(\mathcal{L}_2^2) & \longrightarrow & G(\mathcal{L}_1^2 \star \mathcal{L}_2^2) \\ (\mathfrak{g}_P, \mathfrak{g}_Q) = ((P,\phi_P),(Q,\phi_Q)) & \longmapsto & \mathfrak{g}_P \star \mathfrak{g}_Q = ((P,Q), \pi_1^*\phi_P \otimes \pi_2^*\phi_Q) \end{array}$$

By Definition 5.1.27, we easily see that an element $\mathfrak{g}_1 \star \mathfrak{g}_2 \in G(\mathcal{L}_1^2 \star \mathcal{L}_2^2)$ acts on a product of sections $s_1 \star s_2 \in \Gamma(E_1 \times E_2, \mathcal{L}_1^2 \star \mathcal{L}_2^2)$ as

$$(\mathfrak{g}_1 \star \mathfrak{g}_2) \cdot (s_1 \star s_2) = (\mathfrak{g}_1 \cdot s_1) \star (\mathfrak{g}_2 \cdot s_2).$$

Let T_1', T_2' forming an isotropic subgroup of $(E_1 \times E_2)[4]$ such that $\ker(f_1) = \langle T_1, T_2\rangle$, where $T_1 := [2]T_1'$ and $T_2 := [2]T_2'$. We may write $T_i' := (P_i', Q_i')$ and $T_i := (P_i, Q_i)$ for $i \in \{1,2\}$. Consider $S_1' := (0, Q_2')$ and $S_1' := (P_1', 0)$.

Lemma 6.5.3. *Either one of the following conditions is satisfied:*

(i) $\mathscr{B} := (S_1', S_2', T_1', T_2')$ *is a* ζ_4*-symplectic basis of* $(E_1 \times E_2)[4]$ *adapted to* f_1 i.e. *such that* $\ker(f_1) = \langle [2]T_1', [2]T_2'\rangle$, *where* $\zeta_4 := e_{\mathcal{L}_1^4}(P_1', P_2') = e_4(P_1', P_2')$.

(ii) f_1 *is a diagonal isogeny* $\mathrm{Diag}(\varphi_1, \varphi_2) : E_1 \times E_2 \longrightarrow E_1' \times E_2'$.

Proof. When f_1 is not diagonal, we prove that P_1', P_2', Q_1' and Q_2' have order 4. Assume that P_1' does not have order 4. Then, $P_1 = [2]P_1' = 0$ so $T_1 = (0, Q_1)$ and Q_1 has order 2 (as T_1). Hence, we may write $E_2[2] = \langle Q_1, R\rangle$ and $Q_2 = [\lambda]Q_1 + [\mu]R$ with $\lambda, \mu \in \{0,1\}$. Since (T_1, T_2) is isotropic (as the kernel of f_1), we have by Lemma 6.4.2.(i),

$$1 = e_{\mathcal{L}_1^2 \star \mathcal{L}_2^2}(T_1, T_2) = e_{\mathcal{L}_2^2}(Q_1, Q_2) = e_{\mathcal{L}_2^2}(Q_1, R)^{\mu},$$

with $e_{\mathcal{L}_2^2}(Q_1, R) = -1$ since Q_1 and R generate $E_2[2]$ so $\mu = 0$. Hence, we may write

$$\ker(f_1) = \langle T_1, T_2\rangle = \langle T_1, T_2 - \lambda T_1\rangle = \langle (P_2, 0), (0, Q_1)\rangle,$$

so that $f_1 = \mathrm{Diag}(\varphi_1, \varphi_2)$ with $\ker(\varphi_1) = \langle P_2\rangle$ and $\ker(\varphi_2) = \langle Q_1\rangle$. Similarly, if P_2', Q_1' or Q_2' does not have order 4, we obtain that $f_1 = \mathrm{Diag}(\varphi_1, \varphi_2)$ with $\ker(\varphi_1) = \langle P_1\rangle, \langle P_1\rangle$ or $\langle P_2\rangle$ and $\ker(\varphi_2) = \langle Q_2\rangle, \langle Q_2\rangle$ or $\langle Q_1\rangle$ respectively.

Now, we assume that f_1 is diagonal so that P_1', P_2', Q_1' and Q_2' have order 4. Besides, we have

$$1 = e_{\mathcal{L}_1^4 \star \mathcal{L}_2^4}(T_1', T_2') = e_{\mathcal{L}_1^4}(P_1', P_2') e_{\mathcal{L}_2^4}(Q_1', Q_2'),$$

so that $e_{\mathcal{L}_1^4}(P_1', P_2') = e_{\mathcal{L}_2^4}(Q_1', Q_2')^{-1}$. We prove that $\zeta_4 := e_{\mathcal{L}_1^4}(P_1', P_2')$ is a square root of -1. Indeed, if it was not the case, we would have $e_{\mathcal{L}_1^2}(P_1, P_2) = e_{\mathcal{L}_1^4}(P_1', P_2')^2 = 1$ so that $P_1 = P_2$ since both points have order 2. Similarly, we would have $e_{\mathcal{L}_2^2}(Q_1, Q_2) = e_{\mathcal{L}_2^4}(Q_1', Q_2')^2 = 1$ so $Q_1 = Q_2$ and $T_1 = T_2$. Hence, we would have $2^2 = \deg(f_1) = \#\ker(f_1) = 2$. Contradiction. It follows that $\zeta_4^2 = -1$ and that $\mathscr{B} := (S_1', S_2', T_1', T_2')$ is a ζ_4-symplectic basis of $(E_1 \times E_2)[4]$ adapted to f_1. □

In the following, we assume that f_1 is not diagonal. Otherwise, one dimensional isogenies computations would be sufficient to compute it and f_2 would be our starting gluing isogeny. Then, the above lemma ensures that $\mathscr{B} := (S_1', S_2', T_1', T_2')$ is a ζ_4-symplectic basis of $(E_1 \times E_2)[4]$ inducing a level 2 symmetric theta structure $\Theta_{\mathcal{L}_1^2 \star \mathcal{L}_2^2}$. As we have seen in the previous paragraph, any point $T' := (P', Q') \in (E_1 \times E_2)[4]$ determines a symmetric element $\mathfrak{g}_T \in G(\mathcal{L}_1^2 \star \mathcal{L}_2^2)$ above $T := [2]T := (P, Q)$ given by $\mathfrak{g}_T := \mathfrak{g}_P \star \mathfrak{g}_Q$, where $\mathfrak{g}_P$ and $\mathfrak{g}_Q$ are symmetric elements above P and Q determined by P' and Q' respectively. To find the coordinate θ_{00} associated to $\Theta_{\mathcal{L}_1^2 \star \mathcal{L}_2^2}$, we have to find a section stable by the action of $\widetilde{K}_2(\Theta_{\mathcal{L}_1^2 \star \mathcal{L}_2^2}) = \langle \mathfrak{g}_{T_1}, \mathfrak{g}_{T_2}\rangle$. Denoting $X_i, Z_i \in \Gamma(E_i, \mathcal{L}_i^2)$ the usual Montgomery coordinates for $i \in \{1,2\}$, we obtain that θ_{00} is the trace:

$$\begin{aligned} \theta_{00} &= \sum_{\mathfrak{g} \in \widetilde{K}_2(\Theta_{\mathcal{L}_1^2 \star \mathcal{L}_2^2})} \mathfrak{g} \cdot x_1 \star x_2 \\ &= x_1 \star x_2 + (\mathfrak{g}_{P_1} \cdot x_1) \star (\mathfrak{g}_{Q_1} \cdot x_2) + (\mathfrak{g}_{P_2} \cdot x_1) \star (\mathfrak{g}_{Q_2} \cdot x_2) + (\mathfrak{g}_{P_1+P_2} \cdot x_1) \star (\mathfrak{g}_{Q_1+Q_2} \cdot x_2). \end{aligned}$$

We then obtain that $\theta_{10} = \mathfrak{g}_{S_1} \cdot \theta_{00}$, $\theta_{01} = \mathfrak{g}_{S_2} \cdot \theta_{00}$ and $\theta_{11} = \mathfrak{g}_{S_1+S_2} \cdot \theta_{00}$. Theta group actions on global sections of the $\Gamma(E_i, \mathcal{L}_i)$ can be computed by Lemma 6.5.1 (see Algorithm 6.10). We then obtain the matrix $N \in M_4(k)$ whose action by multiplication on the left maps the product coordinates $(X_1X_2, X_1Z_2, Z_1X_2, Z_1Z_2)$ to the theta coordinates $(\theta_{00}, \theta_{10}, \theta_{01}, \theta_{11})$ by computing tensor products of matrices. Algorithm 6.11 details how this computation is done.

Algorithm 6.10: EC_action: elliptic curve theta group action.

Data: A point P' of order 4 in the Kummer line of an elliptic curve E/k.
Result: The matrix $M \in M_2(k)$ describing the action $(X, Z) \longmapsto (\mathfrak{g}_P \cdot X, \mathfrak{g}_P \cdot Z)$ by multiplication on the left, where $\mathfrak{g}_P \in G(\mathcal{L}(2(0_E)))$ is the symmetric element above $P := [2]P'$ determined by P'.

1 $P \leftarrow [2]P'$; // Cost: $2\mathbf{S} + 3\mathbf{M}$
2 $z_P x_{P'},\ z_P z_{P'},\ x_P x_{P'},\ x_P z_{P'} \leftarrow z_P \cdot x_{P'},\ z_P \cdot z_{P'},\ x_P \cdot x_{P'},\ x_P \cdot z_{P'}$;
3 $\delta \leftarrow z_P x_{P'} - x_P z_{P'}$;
4 Compute δ^{-1}, $z_{P'}^{-1}$ via batched inversions (Algorithm 5.3) ; // Cost: $3\mathbf{M} + 1\mathbf{I}$
5 $M_{00} \leftarrow x_P z_{P'} \cdot \delta^{-1}$;
6 $M_{10} \leftarrow z_P z_{P'} \cdot \delta^{-1}$;
7 $M_{01} \leftarrow x_{P'} \cdot z_{P'}^{-1} - x_P x_{P'} \cdot \delta^{-1}$;
8 $M_{11} \leftarrow -M_{00}$;
9 **return** $M := (M_{i,j})_{i,j \in \{0,1\}}$; // Total cost: $2\mathbf{S} + 14\mathbf{M} + 1\mathbf{I}$

Remark 6.5.4. Grouping together the batched inversions in all calls of Algorithm 6.10, we can save 3 inversions at the expense of 9 more multiplications, changing the total cost of Algorithm 6.11 to $8\mathbf{S} + 109\mathbf{M} + 1\mathbf{I}$.

General method vs. Robert's method: a comparison

To apply the general method, we first need to compute the change of basis matrix $M \in \mathrm{Sp}_4(\mathbb{Z}/4\mathbb{Z})$ from a product symplectic basis $\mathscr{B}_0$ of $(E_1 \times E_2)[4]$ to a symplectic basis $\mathscr{B}$ of $(E_1 \times E_2)[4]$ adapted to F. Recall that $\mathscr{B}_0$ is of the form $((U_1, 0), (0, U_2), (V_1, 0), (0, V_2))$, where (U_i, V_i) is a *special* basis of $E_i[4]$ with $\zeta_4^2 = -1$ and $x(Q_i) = -1$ for $i \in \{1, 2\}$. Hence, to compute M, we first have to find such basis (U_i, V_i) and express $\mathscr{B}$ in $\mathscr{B}_0$. If the the U_i and V_i are already known, any point $P \in E_i[4]$ can be written as $P = [a]U_i + [b]V_i$ with bi-discrete logarithms $a, b \in \mathbb{Z}/4\mathbb{Z}$ computed with Weil pairings as follows:

$$e_4(U_i, P) = e_4(U_i, V_i)^b \quad \text{and} \quad e_4(P, V_i) = e_4(U_i, V_i)^a.$$

However, when we need to decompose a basis (P, Q) into a special basis (U_i, V_i) while looking for U_i at the same time, we rely on the following lemma.

Lemma 6.5.5. *Let E/k be an elliptic curve, (P, Q) a basis of $E[4]$, $V \in E[4]$ a point of order 4 and $\zeta_4 \in k^*$ such that $\zeta_4^2 = -1$. Then, we can find $U \in E[4]$ such that (U, V) is a basis of $E[4]$ with $e_4(U, V) = \zeta_4$ and compute the change of basis matrix from (U, V) to (P, Q) with three 4-th Weil pairing computations $e_4 : E[4] \times E[4] \longrightarrow k^*$.*

Proof. We compute $e_4(P, V)$, $e_4(Q, V)$ and $e_4(P, Q)$. Assume that $e_4(P, V)$ has order 4. Then, $e_4(P, V) = \zeta_4^{\varepsilon}$ with $\varepsilon \in \{-1, 1\}$ so we may set $U := [\varepsilon]P$. Let $b_1, b_2 \in \mathbb{Z}/4\mathbb{Z}$ such that $e_4(Q, V) = \zeta_4^{b_1}$ and $e_4(P, Q) = \zeta_4^{b_2}$. Then, $Q = [\varepsilon b_1]U + [\varepsilon b_2]V$ and the change of basis matrix from (U, V) to (P, Q) is

$$\begin{pmatrix} \varepsilon & \varepsilon b_1 \\ 0 & \varepsilon b_2 \end{pmatrix}.$$

If $e_4(P, V)$ does not have order 4, then $e_4(Q, V)$ must have order 4, otherwise P and Q would be linearly dependent in $E[4]$. Then, $e_4(Q, V) = \zeta_4^{\varepsilon}$ with $\varepsilon \in \{-1, 1\}$ so we may set $U := [\varepsilon]Q$. Let $b_1, b_2 \in \mathbb{Z}/4\mathbb{Z}$ such that $e_4(P, V) = \zeta_4^{b_1}$ and $e_4(P, Q) = \zeta_4^{b_2}$. Then, $P = [\varepsilon b_1]U - [\varepsilon b_2]V$ and the change of basis matrix from (U, V) to (P, Q) is

$$\begin{pmatrix} \varepsilon b_1 & \varepsilon \\ -\varepsilon b_2 & 0 \end{pmatrix}.$$

Algorithm 6.11: Change of coordinates matrix.

Data: Points $T'_1 = (P'_1, P'_2)$ and $T'_2 = (Q'_1, Q'_2)$ forming an isotropic subgroup of $(E_1 \times E_2)[4]$ such that $\ker(f_1) = \langle [2]T'_1, [2]T'_2 \rangle$.

Result: The matrix $N \in M_4(k)$ describing the change of coordinates $(X_1X_2, X_1Z_2, Z_1X_2, Z_1Z_2) \longmapsto (\theta_{00}, \theta_{01}, \theta_{10}, \theta_{11})$ by multiplication on the left.

$G_1 \leftarrow \mathsf{EC\_action}(P'_1)$; // Algorithm 6.10. Cost: $2\mathbf{S} + 14\mathbf{M} + 1\mathbf{I}$
$G_2 \leftarrow \mathsf{EC\_action}(P'_2)$;
$H_1 \leftarrow \mathsf{EC\_action}(Q'_1)$;
$H_2 \leftarrow \mathsf{EC\_action}(Q'_2)$;
$t_{1,00} \leftarrow G_{1,00} \cdot H_{1,00} + G_{1,01} \cdot H_{1,10}$; // Action of $\mathfrak{g}_{P'_1+Q'_1}$ on x_1
$t_{1,01} \leftarrow G_{1,00} \cdot H_{1,01} + G_{1,01} \cdot H_{1,11}$;
$t_{2,00} \leftarrow G_{2,00} \cdot H_{2,00} + G_{2,01} \cdot H_{2,10}$; // Action of $\mathfrak{g}_{P'_2+Q'_2}$ on x_2
$t_{2,01} \leftarrow G_{2,00} \cdot H_{2,01} + G_{2,01} \cdot H_{2,11}$;
$N_{00} \leftarrow G_{1,00} \cdot G_{2,00} + H_{1,00} \cdot H_{2,00} + t_{1,00} \cdot t_{2,00} + 1$; // Trace for the first row (θ_{00})
$N_{01} \leftarrow G_{1,00} \cdot G_{2,01} + H_{1,00} \cdot H_{2,01} + t_{1,00} \cdot t_{2,01}$;
$N_{02} \leftarrow G_{1,01} \cdot G_{2,00} + H_{1,01} \cdot H_{2,00} + t_{1,01} \cdot t_{2,00}$;
$N_{03} \leftarrow G_{1,01} \cdot G_{2,01} + H_{1,01} \cdot H_{2,01} + t_{1,01} \cdot t_{2,01}$;
$N_{10} \leftarrow H_{2,00} \cdot N_{00} + H_{2,10} \cdot N_{01}$; // Action of $\mathfrak{g}_{(0,Q'_2)}$ for the second row (θ_{10})
$N_{11} \leftarrow H_{2,01} \cdot N_{00} + H_{2,11} \cdot N_{01}$;
$N_{12} \leftarrow H_{2,00} \cdot N_{02} + H_{2,10} \cdot N_{03}$;
$N_{13} \leftarrow H_{2,01} \cdot N_{02} + H_{2,11} \cdot N_{03}$;
$N_{20} \leftarrow G_{1,00} \cdot N_{00} + G_{1,10} \cdot N_{02}$; // Action of $\mathfrak{g}_{(P'_1,0)}$ for the third row (θ_{01})
$N_{21} \leftarrow G_{1,00} \cdot N_{01} + G_{1,10} \cdot N_{03}$;
$N_{22} \leftarrow G_{1,01} \cdot N_{00} + G_{1,11} \cdot N_{02}$;
$N_{23} \leftarrow G_{1,01} \cdot N_{01} + G_{1,11} \cdot N_{03}$;
$N_{30} \leftarrow G_{1,00} \cdot N_{10} + G_{1,10} \cdot N_{12}$; // Action of $\mathfrak{g}_{(P'_1,Q'_2)}$ for the fourth row (θ_{11})
$N_{31} \leftarrow G_{1,00} \cdot N_{11} + G_{1,10} \cdot N_{13}$;
$N_{32} \leftarrow G_{1,01} \cdot N_{10} + G_{1,11} \cdot N_{12}$;
$N_{33} \leftarrow G_{1,01} \cdot N_{11} + G_{1,11} \cdot N_{13}$;
return $N := (N_{i,j})_{0 \le i,j \le 3}$; // Total cost: $8\mathbf{S} + 100\mathbf{M} + 4\mathbf{I}$

Note that discrete logarithms computations within the group of 4-th roots of unity is essentially free. We only have to compute $\zeta_4^3 = -\zeta_4$ and look up in the set $\{1, \zeta_4, -1, -\zeta_4\}$. This completes the proof. □

Now, using the notations from Lemma 6.5.3, if the basis adapted to F is of the form $\mathscr{B} := (S'_1, S'_2, T'_1, T'_2)$ with $T'_1 := (P'_1, Q'_1)$, $T'_2 := (P'_2, Q'_2)$, $S'_1 := (0, Q'_2)$ and $S'_1 := (P'_1, 0)$, we may use the basis (P'_1, P'_2) of $E_1[4]$ and (Q'_1, Q'_2) of $E_2[4]$ and Lemma 6.5.5 to obtain ζ_4-special basis (U_i, V_i) of $E_i[4]$ for $i \in \{1, 2\}$ (where $\zeta_4 = e_4(P'_1, P'_2)$) along with the related change of basis matrices. Obtaining the product ζ_4-symplectic basis $\mathscr{B}_0$ and the change of basis matrix $M \in \mathrm{Sp}_4(\mathbb{Z}/4\mathbb{Z})$ from $\mathscr{B}_0$ to $\mathscr{B}$ is then immediate. By Lemma 6.5.5, this costs 6 Weil pairing computations in total and two square roots $2\mathbf{Sqrt}$ for the computation of the y-coordinate of V_1 and V_2.

Remark 6.5.6 (On the cost of Weil pairings)**.** The standard technique to compute a Weil pairing is via Miller's algorithm. We use the formula from [Mil04, Proposition 8]:

$$e_n(P, Q) = (-1)^n \frac{f_{n,P}(Q)}{f_{n,Q}(P)},$$

where Miller's function $f_{n,P}$ can be computed iteratively via the formula

$$f_{n+m,P} = f_{n,P} \cdot f_{n,P} \cdot \frac{L_{[m]P,[n]P}}{L_{[m+n]P,-[n+m]P}},$$

where $\mathrm{div}(L_{P,Q}) = (P) + (Q) + (-(P+Q)) - 3(0)$. The computation of the 4-th Weil pairing requires 4 Miller's iterations (corresponding to point duplication) and each one costs $14\mathbf{M} + 10\mathbf{S}$ by

[CJL+17, § 4.1]. It follows that the computation of $M \in \mathrm{Sp}_4(\mathbb{Z}/4\mathbb{Z})$ costs $24(14\mathbf{M} + 10\mathbf{S}) + 2\mathbf{Sqrt} = 336\mathbf{M} + 240\mathbf{S} + 2\mathbf{Sqrt}$ in total. As suggested in [CJL+17] when E_1 and E_2 are supersingular, we can save some computations by using Tate pairings instead of Weil pairings. Formulas for pairings can also be optimised depending on the curves E_1 and E_2 and the base field.

Then, computing the matrix N_0 from Eq. (6.18) only costs 4 multiplications. Besides, computing the change of coordinates matrix N_1 from product theta coordinates to theta coordinates adapted to F has a negligible cost. Indeed, powers of a 4-th primitive root of unity $\zeta_4 \in k^*$ can be computed without any multiplication, using the fact that $\zeta^2 = -1$. Hence, the computation of N_1 with the formulas from Theorem 6.2.10 only costs 16 additions or subtractions over the base field k and some (negligible) operations over $\mathbb{Z}/4\mathbb{Z}$. Finally, using the redundancies in N_0, the product $N = N_1 \cdot N_0$ can be computed with 16 multiplications only. Hence, computing N once the change of basis matrix $M \in \mathrm{Sp}_4(\mathbb{Z}/4\mathbb{Z})$ is known only costs $20\mathbf{M}$ and the total computation cost is $356\mathbf{M} + 240\mathbf{S} + 2\mathbf{Sqrt}$ by Remark 6.5.6, to be compared with the cost of Robert's method $(8\mathbf{S} + 109\mathbf{M} + 1\mathbf{I})$, which is much cheaper. Even improvements suggested in Remark 6.5.6 would make it hard for the general method to beat Robert's method. However, we shall see in Section 6.5.4 that when the 2^{e+2}-torsion is accessible and F is obtained from Kani's lemma, Lemma 6.4.1 ensures that the general method simplifies the splitting of the codomain of F into $E_3 \times E_4$.

An efficient combination of the two methods

The general method can be dramatically improved if combined with Robert's method over elliptic curves to convert Montgomery coordinates into theta coordinates associated to a non-special basis of the 4-torsion, and hence, relaxing the specialness condition to form the basis $\mathscr{B}_0$. This would essentially avoid any costly discrete logarithm computation.

Lemma 6.5.7. *Let E/k be an elliptic curve in Montgomery form, (P', Q') be a basis of $E[4]$, $P := [2]P'$ and $Q := [2]Q'$. Let X and Z be the usual Montgomery coordinates on E. Then, the level 2 theta coordinates associated to the symmetric theta structure induced by (P', Q') are given by $\theta_0 = \delta_{P'}(\lambda X + \mu Z)$ and $\theta_1 = \delta_{Q'}(\lambda' X + \mu' Z)$, with:*

$$\lambda := z_{Q'} x_{Q'} z_Q, \quad \mu := x_{Q'}(x_{Q'} z_Q - 2 x_Q z_{Q'}),$$

$$\lambda' := z_{P'}^2(\lambda x_P + \mu z_P), \quad \mu' := \lambda x_{P'}(z_P x_{P'} - 2 x_P z_{P'}) - \mu z_{P'}^2 x_P,$$

$\delta_{P'} := z_{P'}(x_{P'} z_P - z_{P'} x_P)$ and $\delta_{Q'} := z_{Q'}(x_{Q'} z_Q - z_{Q'} x_Q)$.

Proof. Let $\mathfrak{g}_P, \mathfrak{g}_Q \in G(\mathcal{L}(2(0_E)))$ be symmetric elements determined by P', Q' respectively. Then, θ_0 is invariant under the action of $\mathfrak{g}_Q$ so it can be defined as the trace $\theta_0 := \nu(x + \mathfrak{g}_Q \cdot x)$ for any projective constant $\nu \in k^*$, provided this trace is non-zero. By Lemma 6.5.1, we have

$$\begin{aligned}
\theta_0 &= \nu\left(x + \frac{x_Q z_{Q'} \cdot X + (z_Q x_{Q'}^2 / z_{Q'} - 2 x_Q x_{Q'}) \cdot Z}{x_{Q'} z_Q - z_{Q'} x_Q}\right) \\
&= \nu \frac{z_{Q'}(x_{Q'} z_Q - z_{Q'} x_Q) \cdot X + x_Q z_{Q'}^2 \cdot X + (z_Q x_{Q'}^2 - 2 x_Q x_{Q'} z_{Q'}) \cdot Z}{z_{Q'}(x_{Q'} z_Q - z_{Q'} x_Q)} \\
&= \nu \frac{z_{Q'} x_{Q'} z_Q \cdot X + (z_Q x_{Q'}^2 - 2 x_Q x_{Q'} z_{Q'}) \cdot Z}{z_{Q'}(x_{Q'} z_Q - z_{Q'} x_Q)} \\
&= \nu \frac{\lambda \cdot X + \mu \cdot Z}{\delta_{Q'}},
\end{aligned}$$

and

$$\begin{aligned}
\theta_1 = \mathfrak{g}_P \cdot \theta_0 &= \nu \frac{\lambda \mathfrak{g}_P \cdot X + \mu \mathfrak{g}_P \cdot Z}{\delta_{Q'}} \\
&= \frac{\nu\lambda(x_P z_{P'} \cdot X + (z_P x_{P'}^2 / z_{P'} - 2 x_P x_{P'}) \cdot Z)}{\delta_{Q'}(x_{P'} z_P - z_{P'} x_P)} + \frac{\nu\mu(z_P z_{P'} \cdot X - x_P z_{P'} \cdot Z)}{\delta_{Q'}(x_{P'} z_P - z_{P'} x_P)} \\
&= \frac{\nu\lambda(x_P z_{P'}^2 \cdot X + x_{P'}(z_P x_{P'} - 2 x_P z_{P'}) \cdot Z)}{\delta_{Q'}\delta_{P'}} + \frac{\nu\mu z_{P'}(z_P z_{P'} \cdot X - x_P z_{P'} \cdot Z)}{\delta_{Q'}\delta_{P'}}
\end{aligned}$$

$$= \frac{\nu z_{P'}^2(\lambda x_P + \mu z_P)\cdot X + \nu(\lambda x_{P'}(z_P x_{P'} - 2x_P z_{P'}) - \mu z_{P'}^2 x_P)\cdot Z}{\delta_{Q'}\delta_{P'}}$$

Taking $\nu := \delta_{Q'}\delta_{P'}$, we obtain the desired result. Note that $\lambda \neq 0$ since $z_{Q'}z_Q \neq 0$, otherwise Q or Q' would be zero and $x_{Q'} \neq 0$, otherwise Q' would be of order 2. This completes the proof. □

Now, we explain how to adapt the general method. Given an adapted symplectic basis $\mathscr{B} := (S_1', S_2', T_1', T_2')$ of $(E_1 \times E_2)[4]$ with $T_1' := (P_1', Q_1')$, $T_2' := (P_2', Q_2')$, $S_1' := (0, Q_2')$ and $S_1' := (P_1', 0)$, we consider the level 2 theta structures $\Theta_{\mathcal{L}_1^2}$ and $\Theta_{\mathcal{L}_2^2}$ on $(E_1, \mathcal{L}_1^2)$ and $(E_2, \mathcal{L}_2^2)$ induced by (P_1', P_2') and (Q_2', Q_1') respectively. The associated theta coordinates are given by

$$(\theta_0^{\mathcal{L}_i^2}, \theta_1^{\mathcal{L}_i^2}) = (m_{i,00}X_i + m_{i,01}Z_i, m_{i,10}X_i + m_{i,11}Z_i),$$

for $i \in \{1,2\}$, where the matrix $m_i := (m_{i,r,s})_{r,s\in\{0,1\}} \in M_2(k)$ is obtained from Lemma 6.5.7. Without any clever optimisation, computing each matrix costs $20\mathbf{M} + 1\mathbf{S}$. Then, the conversion matrix $N_0 \in M_4(k)$ to product theta coordinates $(\theta_{00}, \theta_{10}, \theta_{01}, \theta_{11})$ acting by left multiplication on the product $(X_1X_2, X_1Z_2, Z_1X_2, Z_1Z_2)$ can be written as:

$$N_0 := \begin{pmatrix} m_{1,00}m_{2,00} & m_{1,00}m_{2,01} & m_{1,01}m_{2,00} & m_{1,01}m_{2,01} \\ m_{1,10}m_{2,00} & m_{1,10}m_{2,01} & m_{1,11}m_{2,00} & m_{1,11}m_{2,01} \\ m_{1,00}m_{2,10} & m_{1,00}m_{2,11} & m_{1,01}m_{2,10} & m_{1,01}m_{2,11} \\ m_{1,10}m_{2,10} & m_{1,10}m_{2,11} & m_{1,11}m_{2,10} & m_{1,11}m_{2,11} \end{pmatrix}. \tag{6.19}$$

Then, as explained previously, the computation of the change of theta coordinate matrix $N_1 \in M_4(k)$ from the product theta coordinates induced by $\mathscr{B}_0 := ((P_1', 0), (0, Q_2'), (P_2', 0), (Q_1', 0))$ to theta coordinates induced by $\mathscr{B}$ adapted to F can be computed in negligible time. The product $N := N_1 \cdot N_0$ costs $64\mathbf{M}$. It follows that the total cost with this new method is $2\mathbf{S} + 104\mathbf{M}$, which is significantly less than Robert's method ($8\mathbf{S} + 109\mathbf{M} + 1\mathbf{I}$), given the absence of inversion.

6.5.2 Computing and evaluating a generic 2-isogeny

In this section, we adapt the algorithms from Sections 6.1.1 and 6.1.2 to compute a generic (*i.e.* non-gluing) 2-dimensional 2-isogeny $f : A \longrightarrow B$ in order to minimise arithmetic operations. We begin with the point evaluation algorithm taking as input the inverse of the dual codomain theta null point of f. We then explain how to compute this data when 8-torsion points above $\ker(f)$ are given, but also when 4-torsion points or only 2-torsion points are given, at the expense of square root computations. The possibility of relaxing 8-torsion point requirements is specific to dimension 2 and is not a general fact (as we shall see in Remark 6.5.13).

Evaluation

Let $(A, \mathcal{L}_0)$ and $(B, \mathcal{M}_0)$ be principally polarised abelian surfaces, $\mathcal{L} := \mathcal{L}_0^2$ and $\mathcal{M} := \mathcal{M}_0^2$ and $f : (A, \mathcal{L}^2) \longrightarrow (B, \mathcal{M})$ be a 2-isogeny. We assume we are given a level 2 theta structure $\Theta_{\mathcal{L}}$ on $G(\mathcal{L})$ adapted to f, a level 4 theta structure $\Theta_{\mathcal{L}^2}$ on $G(\mathcal{L}^2)$ compatible with $\Theta_{\mathcal{L}}$ and a level 2 theta structure $\Theta_{\mathcal{M}}$ on $G(\mathcal{M})$ induced by f and $\Theta_{\mathcal{L}^2}$, as in Theorem 6.1.1.

For all $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^2}$ and $i \in (\mathbb{Z}/4\mathbb{Z})^2$, recall the definition of $U_{\chi,i}^{\mathcal{L}^2}$ from Definition 5.3.40:

$$U_{\chi,i}^{\mathcal{L}^2} = \sum_{t\in(\mathbb{Z}/2\mathbb{Z})^2} \chi(t)\theta_{i+2t}^{\mathcal{L}^2}.$$

Recall that for all $i \in (\mathbb{Z}/2\mathbb{Z})^2$, we denote by χ^i the character $j \in (\mathbb{Z}/2\mathbb{Z})^2 \longmapsto (-1)^{\langle i|j\rangle}$. We can then define for all $i, j \in (\mathbb{Z}/2\mathbb{Z})^2$, $U_{i,j}^{\mathcal{L}^2} := U_{\chi^i,j}$, where the second index j is identified with an element of $\{0,1\}^2 \subseteq (\mathbb{Z}/4\mathbb{Z})^2$.

Definition 6.5.8. We say that a dual theta coordinate $U_{i,j}^{\mathcal{L}^2}$ is even (respectively odd) if $\langle i|j\rangle = 0 \mod 2$ (respectively $\langle i|j\rangle = 1 \mod 2$).

As a consequence of Theorem 6.1.1.(iii), we have seen that the dual theta null point of B satisfies:

$$(U_\chi^{\mathcal{M}}(0_B))_{\chi\in(\widehat{\mathbb{Z}/2\mathbb{Z}})^2} = (U_{\chi,0}^{\mathcal{M}}(0_B))_{\chi\in(\widehat{\mathbb{Z}/2\mathbb{Z}})^2} = (U_{\chi,0}^{\mathcal{L}^2}(0_A))_{\chi\in(\widehat{\mathbb{Z}/2\mathbb{Z}})^2}.$$

In particular, the dual theta null point is made of even theta constants. It has been proved that odd theta constants are always zero [Dup06, Proposition 5.1]. Fortunately, this is not the case for even theta constants. We assume that f is not a gluing isogeny, *i.e.* that A is not a product of elliptic curves. Then, the following result ensures that all even dual theta constants of B are non-zero, so that Algorithm 6.12 can be applied.

Proposition 6.5.9. *[Dup06, Proposition 6.5 & Corollary 6.1]*

(i) If $(A,\mathcal{L}_0)$ is not a product, then all even dual theta constants are non-zero.

(ii) If $(A,\mathcal{L}_0)$ is a product, then only one even dual theta constant $U_{i,j}^{\mathcal{L}^2}(0_A)$ is zero and the underlying theta structure $\Theta_{\mathcal{L}^2}$ is a product if and only if $U_{11,11}^{\mathcal{L}^2}(0_A)=0$.

Proof. In [Dup06], Dupont proved this result over $\mathbb{C}$ in the framework of analytic theta functions. To extend it to finite fields (our main case of interest), we would need to lift the abelian surface to a number field and then reduce it modulo some prime ideal lying above the characteristic. □

We denote by $(x_P : y_P : z_P : t_P)$ the associated level 2 theta coordinates $(\theta_{00}^{\mathcal{L}}(P) : \theta_{10}^{\mathcal{L}}(P) : \theta_{01}^{\mathcal{L}}(P) : \theta_{11}^{\mathcal{L}}(P))$ of a point $P\in A$ and similarly for a point in B. We denote by H the Hadamard transform given by the action by multiplication on the left of theta coordinates by the matrix

$$H := \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix}.$$

We denote by $(\alpha:\beta:\gamma:\delta)$ the dual theta null point of B and by $(\alpha^{-1}:\beta^{-1}:\gamma^{-1}:\delta^{-1})$ its projective inverse.

Algorithm 6.12: Generic 2-dimensional 2-isogeny evaluation.

Data: Theta coordinates $(x_P : y_P : z_P : t_P)$ of P and the projective inverse of the dual codomain theta null point $(\alpha^{-1}:\beta^{-1}:\gamma^{-1}:\delta^{-1})$ on B.
Result: The image theta coordinates $(x_{f(P)} : y_{f(P)} : z_{f(P)} : t_{f(P)})$ of $f(P)$.

$(x',y',z',t') \leftarrow H(x_P^2, y_P^2, z_P^2, t_P^2)$;
$x' \leftarrow \alpha^{-1}\cdot x'$;
$y' \leftarrow \beta^{-1}\cdot y'$;
$z' \leftarrow \gamma^{-1}\cdot z'$;
$t' \leftarrow \delta^{-1}\cdot t'$;
$x_{f(P)}, y_{f(P)}, z_{f(P)}, t_{f(P)} \leftarrow H(x',y',z',t')$;
return $(x_{f(P)} : y_{f(P)} : z_{f(P)} : t_{f(P)})$; // Total cost: $4\mathbf{M}+4\mathbf{S}$

Computation from 8-torsion lying above the kernel

Algorithm 6.13 explains how to compute the codomain theta null point $(a':b':c':d')$ of f, its dual $(\alpha:\beta:\gamma:\delta)$ and the projective inverse of this dual from points $T_1, T_2 \in A[8]$ forming a maximal isotropic subgroup and such that $\ker(f) = \langle [4]T_1, [4]T_2\rangle$. Even though the theta null point $(a':b':c':d')$ is not directly used in isogeny computations, it is still useful to perform point duplications on B (see Algorithm 5.5), as required in an isogeny chain computation. Algorithm 6.13 is much simpler than Algorithm 6.5 working in any dimension, while optimising the number of arithmetic operations. In particular, unlike the latter, it does not require computation trees.

As in Algorithm 6.5, the dual theta null point computation relies on Lemma 6.1.4.

Lemma 6.5.10. *There exists $x,y,z,t\in k$ such that*

$$H(x_{T_1}^2, y_{T_1}^2, z_{T_1}^2, t_{T_1}^2) = (x\alpha, x\beta, y\gamma, y\delta) \quad \text{and} \quad H(x_{T_2}^2, y_{T_2}^2, z_{T_2}^2, t_{T_2}^2) = (z\alpha, t\beta, z\gamma, t\delta).$$

Proof. Let us write $H(x^2_{T_1}, y^2_{T_1}, z^2_{T_1}, t^2_{T_1}) := (x'_1, y'_1, z'_1, t'_1)$, $H(x^2_{T_2}, y^2_{T_2}, z^2_{T_2}, t^2_{T_2}) := (x'_2, y'_2, z'_2, t'_2)$, $\chi^i : j \in (\mathbb{Z}/2\mathbb{Z})^2 \longmapsto (-1)^{\langle i|j\rangle}$ for all $i \in (\mathbb{Z}/2\mathbb{Z})^2$, $\chi_1 = \chi^{10}$ and $\chi_2 = \chi^{01}$. Then, applying Lemma 6.1.4 to $(\chi, l) = (\chi^{00}, 1)$, $(\chi^{01}, 1)$, $(\chi^{00}, 2)$ and $(\chi^{10}, 2)$ respectively, we obtain:

$$\beta \cdot x'_1 = U^{\mathcal{M}}_{\chi^{10}}(0_B) \cdot x'_1 = U^{\mathcal{M}}_{\chi^{00}}(0_B) \cdot y'_1 = \alpha \cdot y'_1, \quad \delta \cdot z'_1 = U^{\mathcal{M}}_{\chi^{11}}(0_B) \cdot z'_1 = U^{\mathcal{M}}_{\chi^{01}}(0_B) \cdot t'_1 = \gamma \cdot t'_1$$

$$\gamma \cdot x'_2 = U^{\mathcal{M}}_{\chi^{01}}(0_B) \cdot x'_2 = U^{\mathcal{M}}_{\chi^{00}}(0_B) \cdot z'_2 = \alpha \cdot z'_2, \quad \delta \cdot y'_2 = U^{\mathcal{M}}_{\chi^{11}}(0_B) \cdot y'_2 = U^{\mathcal{M}}_{\chi^{10}}(0_B) \cdot t'_2 = \beta \cdot t'_2$$

Hence, we may write $x'_1 = \alpha x$, $y'_1 = \beta x$, $z'_1 = \gamma y$, $t'_1 = \delta y$, $x'_2 = \alpha z$, $y'_2 = \beta t$, $z'_2 = \gamma z$, $t'_2 = \delta t$ with $x := x'_1/\alpha = y'_1/\beta$, $y := z'_1/\gamma = t'_1/\delta$, $z := x'_2/\alpha = z'_2/\gamma$ and $t := y'_2/\beta = t'_2/\delta$. This completes the proof. □

Algorithm 6.13 follows from Lemma 6.5.10 and returns a correct result under the following assumption.

Conjecture 6.5.11. *If f is not a gluing isogeny, then the values x, y, z, t from Lemma 6.5.10 are all non-zero.*

Algorithm 6.13: Generic 2-dimensional 2-isogeny codomain computation from 8-torsion points.

Data: Theta coordinates of $T_1, T_2 \in A[8]$ forming a maximal isotropic subgroup, such that $\ker(f) = \langle [4]T_1, [4]T_2 \rangle$.
Result: Codomain theta null point $(a' : b' : c' : d')$ of f, its dual $(\alpha : \beta : \gamma : \delta)$ and the projective inverse $(\alpha^{-1} : \beta^{-1} : \gamma^{-1} : \delta^{-1})$.

$(x\alpha, x\beta, y\gamma, y\delta) \leftarrow H(x^2_{T_1}, y^2_{T_1}, z^2_{T_1}, t^2_{T_1})$;
$(z\alpha, t\beta, z\gamma, t\delta) \leftarrow H(x^2_{T_2}, y^2_{T_2}, z^2_{T_2}, t^2_{T_2})$;
$x\alpha t\beta \leftarrow x\alpha \cdot t\beta$;
$z\alpha x\beta \leftarrow z\alpha \cdot x\beta$;
$\alpha \leftarrow z\alpha \cdot x\alpha t\beta$;
$\beta \leftarrow t\beta \cdot z\alpha x\beta$;
$\gamma \leftarrow z\gamma \cdot x\alpha t\beta$;
$\delta \leftarrow t\delta \cdot z\alpha x\beta$;
$z\gamma t\delta \leftarrow z\gamma \cdot t\delta$;
$\alpha^{-1} \leftarrow x\beta \cdot z\gamma t\delta$;
$\beta^{-1} \leftarrow x\alpha \cdot z\gamma t\delta$;
$\gamma^{-1} \leftarrow \delta$;
$\delta^{-1} \leftarrow \gamma$;
$(a', b', c', d') \leftarrow H(\alpha, \beta, \gamma, \delta)$;
return $(a' : b' : c' : d'), (\alpha : \beta : \gamma : \delta), (\alpha^{-1} : \beta^{-1} : \gamma^{-1} : \delta^{-1})$; // Total cost: $9\mathbf{M} + 8\mathbf{S}$

Computation from 4-torsion lying above the kernel

Now, we assume we are given 4-torsion points $T_1, T_2 \in A[4]$ such that $\ker(f) = \langle [2]T_1, [2]T_2 \rangle$. Then, we are still able to compute the codomain theta null point, its dual and the projective inverse of its dual at the expense of 2 square root computations. Actually, only T_1 is needed and T_2 is not used. The computation relies on the following fact.

Lemma 6.5.12. *There exists $x \in k^*$ such that $H(x^2_{T_1}, y^2_{T_1}, z^2_{T_1}, t^2_{T_1}) = (\alpha\beta x, \alpha\beta x, \gamma\delta x, \gamma\delta x)$.*

Proof. Corollary 6.1.3 ensures the existence of $\lambda \in k^*$ such that for all $\chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^2}$,

$$U^{\mathcal{M}}_{\chi}(f(T_1)) \cdot U^{\mathcal{M}}_{\chi}(0_B) = \lambda \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^2} \chi(t) \theta^{\mathcal{L}}_t(T_1)^2. \tag{6.20}$$

But as we have seen in the proof of Lemma 6.2.11.(ii), the dual theta coordinates $U^{\mathcal{M}}_{\chi}$ are associated to a level 2 theta structure $\Theta'_{\mathcal{M}}$ related to the theta structure $\Theta_{\mathcal{M}}$ induced by f by the symplectic change of basis matrix

$$\begin{pmatrix} 0 & -I_2 \\ I_2 & 0 \end{pmatrix} \in \mathrm{Sp}_4(\mathbb{Z}/4\mathbb{Z}),$$

so that $\overline{\Theta}'_{\mathcal{M}}((1,0),1) = f(T_1)$. It follows by Eq. (5.11) that

$$(U^{\mathcal{M}}_{\chi^{00}}(f(T_1)) : U^{\mathcal{M}}_{\chi^{10}}(f(T_1)) : U^{\mathcal{M}}_{\chi^{01}}(f(T_1)) : U^{\mathcal{M}}_{\chi^{11}}(f(T_1))) \\ = (U^{\mathcal{M}}_{\chi^{10}}(0_B) : U^{\mathcal{M}}_{\chi^{00}}(0_B) : U^{\mathcal{M}}_{\chi^{11}}(0_B) : U^{\mathcal{M}}_{\chi^{01}}(0_B)) = (\beta : \alpha : \delta : \gamma),$$

so there exists $\mu \in k^*$ such that

$$(U^{\mathcal{M}}_{\chi^{00}}(f(T_1)), U^{\mathcal{M}}_{\chi^{10}}(f(T_1)), U^{\mathcal{M}}_{\chi^{01}}(f(T_1)), U^{\mathcal{M}}_{\chi^{11}}(f(T_1))) = \mu(\beta, \alpha, \delta, \gamma)$$

It follows by Eq. (6.20) that

$$H(x^2_{T_1}, y^2_{T_1}, z^2_{T_1}, t^2_{T_1}) = \lambda^{-1}\mu(\alpha\beta, \alpha\beta, \gamma\delta, \gamma\delta).$$

This completes the proof. □

Besides, if $(a : b : c : d)$ is the theta null point of A, Corollary 6.1.3 also ensures that $(\alpha^2, \beta^2, \gamma^2, \delta^2) = H(a^2, b^2, c^2, d^2)$, up to a projective factor that we can ignore. From this formula, we can extract $\alpha\beta = \pm\sqrt{\alpha^2 \cdot \beta^2}$, $\alpha\gamma = \pm\sqrt{\alpha^2 \cdot \gamma^2}$ and finally $(\alpha : \beta : \gamma : \delta)$ from Lemma 6.5.12. As we shall see in Remark 6.5.13, the sign indetermination in square root computations does not impact the correctness of the result. Algorithm 6.14 follows.

Algorithm 6.14: Generic 2-dimensional 2-isogeny codomain computation from 4-torsion points.

Data: Theta coordinates of T_1 of order 4 such that $[2]T_1 \in \ker(f)$ and the theta null point $(a : b : c : d)$ of A.
Result: Codomain theta null point $(a' : b' : c' : d')$ of f, its dual $(\alpha : \beta : \gamma : \delta)$ and the projective inverse $(\alpha^{-1} : \beta^{-1} : \gamma^{-1} : \delta^{-1})$.

$(x\alpha\beta, \_, x\gamma\delta, \_) \leftarrow H(x^2_{T_1}, y^2_{T_1}, z^2_{T_1}, t^2_{T_1})$;
$(\alpha^2, \beta^2, \gamma^2, \delta^2) \leftarrow H(a^2, b^2, c^2, d^2)$;
$\alpha\beta \leftarrow \sqrt{\alpha^2 \cdot \beta^2}$;
$\alpha\gamma \leftarrow \sqrt{\alpha^2 \cdot \gamma^2}$;
$\beta \leftarrow \alpha\beta \cdot \alpha\gamma$;
$\delta^{-1} \leftarrow \beta \cdot x\gamma\delta$;
$\beta \leftarrow \beta \cdot x\alpha\beta$;
$\delta \leftarrow x\gamma\delta \cdot \alpha\beta \cdot \alpha^2$;
$\alpha \leftarrow x\alpha\beta \cdot \alpha^2$;
$\gamma \leftarrow \alpha \cdot \gamma^2$;
$\alpha \leftarrow \alpha \cdot \alpha\gamma$;
$\alpha^{-1} \leftarrow x\alpha\beta \cdot \delta^2$;
$\gamma^{-1} \leftarrow \alpha^{-1} \cdot \beta^2$;
$\alpha^{-1} \leftarrow \alpha^{-1} \cdot \gamma^2$;
$\beta^{-1} \leftarrow \alpha^{-1} \cdot \alpha\beta$;
$\alpha^{-1} \leftarrow \alpha^{-1} \cdot \beta^2$;
$\gamma^{-1} \leftarrow \gamma^{-1} \cdot \alpha\gamma$;
$\delta^{-1} \leftarrow \delta^{-1} \cdot \beta^2$;
$(a', b', c', d') \leftarrow H(\alpha, \beta, \gamma, \delta)$;
return $(a' : b' : c' : d')$, $(\alpha : \beta : \gamma : \delta)$, $(\alpha^{-1} : \beta^{-1} : \gamma^{-1} : \delta^{-1})$;
// Total cost: $8\mathbf{S} + 17\mathbf{M} + 2\mathbf{Sqrt}$

Computation from the domain theta null point

If $(a : b : c : d)$ is the theta null point of A, we have seen that

$$(\alpha^2, \beta^2, \gamma^2, \delta^2) = H(a^2, b^2, c^2, d^2), \tag{6.21}$$

up to a projective factor that we can ignore. When only kernel generators $T_1, T_2 \in A[2]$ are available, we rely on this formula to obtain the codomain dual theta null points $(\alpha : \beta : \gamma : \delta)$ with 3 square

Algorithm 6.15: Generic 2-dimensional 2-isogeny codomain computation from the domain theta null point.

Data: Theta coordinates of the theta null point $(a : b : c : d)$ of A.
Result: Codomain theta null point $(a' : b' : c' : d')$ of f, its dual $(\alpha : \beta : \gamma : \delta)$ and the projective inverse $(\alpha^{-1} : \beta^{-1} : \gamma^{-1} : \delta^{-1})$.

$(\alpha^2, \beta^2, \gamma^2, \delta^2) \leftarrow H(a^2, b^2, c^2, d^2)$;
$\alpha \leftarrow \alpha^2$;
$\beta \leftarrow \alpha^2 \cdot \beta^2$;
$\gamma \leftarrow \alpha^2 \cdot \gamma^2$;
$\delta \leftarrow \alpha^2 \cdot \delta^2$;
$\beta \leftarrow \sqrt{\beta}$;
$\gamma \leftarrow \sqrt{\gamma}$;
$\delta \leftarrow \sqrt{\delta}$;
$\alpha^{-1} \leftarrow \gamma^2 \cdot \delta^2$;
$\beta^{-1} \leftarrow \alpha^{-1} \cdot \beta$;
$\alpha^{-1} \leftarrow \alpha^{-1} \cdot \beta^2$;
$\gamma^{-1} \leftarrow \delta^2 \cdot \beta^2 \cdot \gamma$;
$\delta^{-1} \leftarrow \gamma^2 \cdot \beta^2 \cdot \delta$;
$(a', b', c', d') \leftarrow H(\alpha, \beta, \gamma, \delta)$;
return $(a' : b' : c' : d')$, $(\alpha : \beta : \gamma : \delta)$, $(\alpha^{-1} : \beta^{-1} : \gamma^{-1} : \delta^{-1})$;
// Total cost: $4\mathbf{S} + 10\mathbf{M} + 3\mathbf{Sqrt}$

root computations. Algorithm 6.15 follows. The kernel generators T_1 and T_2 are not directly used, but the correctness of Eq. (6.21) follows from the fact that the domain theta structure $\Theta_{\mathcal{L}}$ is adapted to f *i.e.* that $K_2(\overline{\Theta}_{\mathcal{L}}) = \langle T_1, T_2 \rangle = \ker(f)$.

Remark 6.5.13 (On the (im)possibility to generalise this approach to higher dimensions)**.** Unlike in the rest of this section, we assume that A and B can be of any dimension g. We assume that 8-torsion points $T_1, \cdots, T_g \in A[8]$ above $\ker(f)$ are not known. We then use the analogue of Eq. (6.21) in dimension g:

$$\forall \chi \in \widehat{(\mathbb{Z}/2\mathbb{Z})^g}, \quad U_\chi^{\mathcal{M}}(0_B)^2 = \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^g} \chi(t) \theta_i^{\mathcal{L}}(0_A)^2,$$

up to a projective factor that we can ignore. Hence, we can compute the codomain dual theta null-point $(U_\chi^{\mathcal{M}}(0_B))_\chi$ with $2^g - 1$ square root computations and sign choices. Unfortunately, this method is not sufficient to determine $(U_\chi^{\mathcal{M}}(0_B))_\chi$ in general because all sign choices may not be valid.

As Robert did in [Rob24, Chapter 7, Appendix B.2], we can prove that we can make at least $g(g+1)/2$ arbitrary sign choices among $2^g - 1$. Indeed, we may act on a symplectic basis of $B[4]$ inducing the theta-structure $\Theta_{\mathcal{M}}$ via the symplectic matrix:

$$M := \begin{pmatrix} I_g & 0 \\ B & I_g \end{pmatrix} \in \mathrm{Sp}_{2g}(\mathbb{Z}/4\mathbb{Z}),$$

which fixes the g last elements of the basis. By Theorem 6.2.10, the new resulting theta-coordinates are $\theta'^{\mathcal{M}}_i = \zeta_4^{-\langle i|Bi\rangle} \theta_i^{\mathcal{M}}$ for all $i \in (\mathbb{Z}/2\mathbb{Z})^g$ (up to a projective factor), where $\zeta_4^2 = -1$. Since M is symplectic, we have ${}^tB = B$ by Lemma 6.2.9 so we have $g(g+1)/2$ values to choose. We have:

$$\forall i \in (\mathbb{Z}/2\mathbb{Z})^g, \quad \langle i|Bi\rangle = \sum_{l=1}^{g} i_k^2 B_{l,l} + 2 \sum_{1 \le l < m \le g} i_l i_m B_{l,m},$$

so the $\langle i|Bi\rangle$ are determined by the $B_{l,l}$ and the $B_{l,m} \mod 2$. For all $l \in [\![1\ ;\ g]\!]$ and $m \in [\![l+1\ ;\ g]\!]$, we may fix $B_{l,l} \in \{0, 2\}$ and $B_{l,m} \in \{0, 1\}$ which fixes $\langle i|Bi\rangle \in 2\mathbb{Z}/4\mathbb{Z}$ so that $\theta'^{\mathcal{M}}_i = (-1)^{\langle i|Bi\rangle/2} \theta_i^{\mathcal{M}}$ for all $i \in (\mathbb{Z}/2\mathbb{Z})^g$, and in particular $\theta'^{\mathcal{M}}_{e_l} = (-1)^{-B_{l,l}/2} \theta^{\mathcal{M}}_{e_l}$, $\theta'^{\mathcal{M}}_{e_l+e_m} = (-1)^{-(B_{l,l}+B_{m,m}+2B_{l,m})/2} \theta^{\mathcal{M}}_{e_l+e_m}$, where e_l is the vector of $(\mathbb{Z}/2\mathbb{Z})^g$ with 1 at index l and 0 everywhere else. This amounts to choosing $g(g+2)/2$ signs among $2^g - 1$ and fixing the others.

In dimension $g = 2$, all $g(g+1)/2 = 3 = 2^g - 1$ arbitrary sign choices are valid, which proves the correctness of Algorithms 6.14 and 6.15. This is no longer true in dimension $g > 2$. In dimension $g = 3$, only 6 among 7 sign choices determine the last one with an explicit formula [KMM+24, Theorem 8]. In dimension $g \geq 4$, we have no such explicit formulas so the theta null-point is harder to guess.

6.5.3 Computing and evaluating a gluing 2-isogeny

In this section, we explain how to compute a non-diagonal gluing 2-isogeny $f : E_1 \times E_2 \longrightarrow B$. With the usual notations, let $\zeta_8 \in k^*$ be an 8-th primitive square root of unity and $\mathscr{B} := (S_1, S_2, T_1, T_2)$ be a ζ_8-symplectic basis of $(E_1 \times E_2)[8]$ adapted to f. By Lemma 6.5.3, we may write $T_1 := (P_1, Q_1)$, $T_2 := (P_2, Q_2)$, $S_1 := (0, Q_2)$ and $S_2 := (P_1, 0)$, where $P_1, P_2 \in E_1[8]$ and $Q_1, Q_2 \in E_2[8]$ have order 8. We may consider the product ζ_8-symplectic basis of $(E_1 \times E_2)[8]$:

$$\mathscr{B}_0 := ((P_1, 0), (0, Q_2), (P_2, 0), (0, Q_1)).$$

Let $\mathcal{L}_i := \mathcal{L}((0_{E_i}))$ for $i \in \{1, 2\}$ and $\Theta_{\mathcal{L}_1^2 \star \mathcal{L}_2^2}$, $\Theta'_{\mathcal{L}_1^2 \star \mathcal{L}_2^2}$, $\Theta'_{\mathcal{L}_1^4 \star \mathcal{L}_2^4}$ be the symmetric theta structure induced by $[2]\mathscr{B}_0$, $[2]\mathscr{B}$ and $\mathscr{B}$ respectively. As in Theorem 6.1.1, let $\Theta_{\mathcal{M}}$ be the level 2 theta structure on $(B, \mathcal{M})$ induced by $\mathscr{C} := ([2]f(S_1), [2]f(S_2), f(T_1), f(T_2))$ and let $(\alpha : \beta : \gamma : \delta)$ be the dual of the theta null point for this theta structure.

Lemma 6.5.14. *We have $\alpha\beta\gamma \neq 0$ and $\delta = 0$.*

Proof. The symplectic change of basis matrix from $[2]\mathscr{B}_0$ to $[2]\mathscr{B}$ is

$$\begin{pmatrix} 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

Then, by Theorem 6.2.10, the change of coordinates matrix from product theta coordinates associated to $\Theta_{\mathcal{L}_1^2 \star \mathcal{L}_2^2}$ to non-product theta coordinates associated to $\Theta'_{\mathcal{L}_1^2 \star \mathcal{L}_2^2}$ acting by left multiplication

$$(\theta_{00}^{\mathcal{L}_1^2 \star \mathcal{L}_2^2}, \theta_{10}^{\mathcal{L}_1^2 \star \mathcal{L}_2^2}, \theta_{01}^{\mathcal{L}_1^2 \star \mathcal{L}_2^2}, \theta_{11}^{\mathcal{L}_1^2 \star \mathcal{L}_2^2}) \overset{N}{\longmapsto} (\theta'^{\mathcal{L}_1^2 \star \mathcal{L}_2^2}_{00}, \theta'^{\mathcal{L}_1^2 \star \mathcal{L}_2^2}_{10}, \theta'^{\mathcal{L}_1^2 \star \mathcal{L}_2^2}_{01}, \theta'^{\mathcal{L}_1^2 \star \mathcal{L}_2^2}_{11})$$

is given by:

$$N := \begin{pmatrix} 1 & 1 & 1 & -1 \\ 1 & -1 & 1 & 1 \\ 1 & 1 & -1 & 1 \\ -1 & 1 & 1 & 1 \end{pmatrix}.$$

The product theta null point associated to $\Theta_{\mathcal{L}_1^2 \star \mathcal{L}_2^2}$ is of the form $(a_1a_2 : b_1a_2 : a_1b_2 : b_1b_2)$ and by Corollary 6.1.3, there exists $\lambda \in k^*$ such that

$$(\alpha^2, \beta^2, \gamma^2, \delta^2) = \lambda \cdot H(S(N(a_1a_2, b_1a_2, a_1b_2, b_1b_2))),$$

where $S : (x, y, z, t) \longmapsto (x^2, y^2, z^2, t^2)$. The computation then shows that $\delta = 0$, and that $\alpha\beta\gamma \neq 0$ since only one dual theta constant can be zero by Proposition 6.5.9. □

Evaluation

Assume we have computed the codomain dual theta null point $(\alpha : \beta : \gamma : 0)$ and its projective inverse $(\alpha^{-1} : \beta^{-1} : \gamma^{-1} : 0)$. Then, if $P \in E_1 \times E_2$, we can evaluate the dual theta coordinates $(x'_{f(P)} : y'_{f(P)} : z'_{f(P)} : t'_{f(P)})$ of $f(P)$, as follows. We use translates by 4-torsion points above the kernel as in Section 6.1.3. For the first three coordinates, we use Eq. (6.1), as for a generic evaluation:

$$H(x_P^2, y_P^2, z_P^2, t_P^2) = \lambda \cdot (\alpha x'_{f(P)}, \beta y'_{f(P)}, \gamma z'_{f(P)}, 0),$$

with $\lambda \in k^*$. For the last coordinate, we use Eq. (6.4) with $l = 1$:

$$H(x^2_{P+[2]T_1}, y^2_{P+[2]T_1}, z^2_{P+[2]T_1}, t^2_{P+[2]T_1}) = \mu \cdot (\alpha y'_{f(P)}, \beta x'_{f(P)}, \gamma t'_{f(P)}, 0),$$

with $\mu \in k^*$. Let us denote by (x'_P, y'_P, z'_P, t'_P) and $(x''_P, y''_P, z''_P, t''_P)$ the values of $H(x_P^2, y_P^2, z_P^2, t_P^2)$ and $H(x^2_{P+[2]T_1}, y^2_{P+[2]T_1}, z^2_{P+[2]T_1}, t^2_{P+[2]T_1})$ respectively. Then, if $x'_P \neq 0$, we have

$$\begin{aligned}(x'_P \cdot y''_P : \beta^{-1} \cdot y'_P \cdot \alpha \cdot y''_P : \gamma^{-1} \cdot z'_P \cdot \alpha \cdot y''_P : \gamma^{-1} \cdot z''_P \cdot \beta \cdot x'_P)& \\ = (\lambda\mu\alpha\beta x'^2_{f(P)} : \lambda\mu\alpha\beta x'_{f(P)} y'_{f(P)} : \lambda\mu\alpha\beta x'_{f(P)} z'_{f(P)} : \lambda\mu\alpha\beta x'_{f(P)} t'_{f(P)})& \\ = (x'_{f(P)} : y'_{f(P)} : z'_{f(P)} : t'_{f(P)}).&\end{aligned}$$

Alternatively, if $y'_P \neq 0$, we also have

$$(\alpha^{-1} \cdot x'_P \cdot \beta \cdot x''_P : y'_P \cdot x''_P : \gamma^{-1} \cdot z'_P \cdot \beta \cdot x''_P : \gamma^{-1} \cdot z''_P \cdot \alpha \cdot y'_P) = (x'_{f(P)} : y'_{f(P)} : z'_{f(P)} : t'_{f(P)}).$$

If both $x'_P = 0$ and $y'_P = 0$, we translate P by $[2]T_2$ instead and use Eq. (6.4) with $l = 2$. Algorithm 6.16 follows.

Remark 6.5.15. As explained in Section 6.1.3, another method has been introduced by Max Duparc for 2-dimensional gluing 2-isogeny computations in [AAA+25, Algorithm 8.40]. In plain generality, this method will require to compute both translates $P + [2]T_1$ and $P - [2]T_1$ (see Algorithm 6.7) but only $P + [2]T_1$ is actually necessary. This enables to save 4 multiplications. Another method has also been introduced in [AAA+25, Algorithm 8.41] to evaluate points of the form $(P_1, 0)$ and $(0, P_2)$ at a much lower cost.

Codomain

The codomain theta null point $(a' : b' : c' : d')$, its dual $(\alpha : \beta : \gamma : 0)$ and the projective inverse $(\alpha^{-1} : \beta^{-1} : \gamma^{-1} : 0)$ of the gluing isogeny $f : E_1 \times E_2 \longrightarrow A$ can be computed from $T_1, T_2 \in (E_1 \times E_2)[8]$ such that $\ker(f) = \langle [4]T_1, [4]T_2 \rangle$, as in the generic case. Hence, we can simply adapt Algorithm 6.13 to the gluing case. Algorithm 6.17 follows.

6.5.4 Computing product (theta) coordinates on the codomain

In this section, we explain how to obtain a product theta structure on the codomain of the 2^e-isogeny $F : E_1 \times E_2 \longrightarrow E_3 \times E_4$ computed as we previously explained. We are then able to split the codomain into a product $E_3 \times E_4$ and express image points $F(P)$ as $(F_3(P), F_4(P))$ into systems of coordinates on the Kummer lines associated to E_3 and E_4 (*e.g.* Montgomery $(X : Z)$-coordinates).

With Kani's lemma

We assume that the 2^e-isogeny $F : E_1 \times E_2 \longrightarrow E_3 \times E_4$ is obtained from Kani's lemma. By the converse of Kani's lemma (Lemma 2.2.7), we can actually always assume this is the case in dimension 2. Then, using the notations from Lemma 6.4.1, we may assume that

$$F := \begin{pmatrix} \varphi & \widehat{\psi'} \\ -\psi & \widehat{\varphi'} \end{pmatrix} : E_1 \times E_2 \longrightarrow E_3 \times E_4.$$

is obtained from the (a, b)-isogeny diamond:

$$\begin{array}{ccc} E_4 & \xrightarrow{\varphi'} & E_2 \\ {\scriptstyle\psi}\uparrow & & \uparrow{\scriptstyle\psi'} \\ E_1 & \xrightarrow{\varphi} & E_3 \end{array}$$

Assume that we used Lemma 6.4.1 to find an adapted theta structure on the domain of F. Then, we can keep track of the theta structure on the codomain. Indeed, by Theorem 6.1.1, the level 2 theta structure we obtain on $E_3 \times E_4$ from the computation of F is induced by the symplectic basis $\mathscr{C} := ([2^e]F(S_1), [2^e]F(S_2), F(T_1), F(T_2))$ of $(E_3 \times E_4)[4]$, where $\mathscr{B} := (S_1, S_2, T_1, T_2)$ is a symplectic basis of $(E_1 \times E_2)[2^{e+2}]$ adpated to F, where $\mathscr{B}$ is given by Lemma 6.4.1.(i). Points (ii) and (iii)

Algorithm 6.16: Gluing 2-dimensional 2-isogeny evaluation.

Data: Points expressed in Jacobian coordinates $P := (R_1, R_2) \in E_1 \times E_2$, $[2]T_1 := (P'_1, Q'_1), [2]T_2 := (P'_2, Q'_2) \in (E_1 \times E_2)[4]$ such that $\ker(f) = \langle [4]T_1, [4]T_2 \rangle$, a change of coordinates matrix N obtained in Section 6.5.1 (*e.g.* from Algorithm 6.11), the codomain dual theta null point $(\alpha : \beta : \gamma : 0)$ and its projective inverse $(\alpha^{-1} : \beta^{-1} : \gamma^{-1} : 0)$.

Result: The image theta coordinates $(x_{f(P)} : y_{f(P)} : z_{f(P)} : t_{f(P)})$ of $f(P)$.

$x'_P, y'_P, z'_P, t'_P \leftarrow N(X(R_1) \cdot X(R_2), X(R_1) \cdot Z(R_2), Z(R_1) \cdot X(R_2), Z(R_1) \cdot Z(R_2))$;
$x'_P, y'_P, z'_P, t'_P \leftarrow H(x'^2_P, y'^2_P, z'^2_P, t'^2_P)$;
if $x'_P \neq 0$ **then**
  $R'_1 \leftarrow R_1 + P'_1$; // Jacobian addition cost: $11\mathbf{M} + 5\mathbf{S}$
  $R'_2 \leftarrow R_2 + Q'_1$;
  $x''_P, y''_P, z''_P, t''_P \leftarrow N(X(R'_1) \cdot X(R'_2), X(R'_1) \cdot Z(R'_2), Z(R'_1) \cdot X(R'_2), Z(R'_1) \cdot Z(R'_2))$;
  $x''_P, y''_P, z''_P, t''_P \leftarrow H(x''^2_P, y''^2_P, z''^2_P, t''^2_P)$;
  $\alpha y''_P \leftarrow \alpha \cdot y''_P$;
  $\beta x'_P \leftarrow \beta \cdot x'_P$;
  $x'_{f(P)} \leftarrow x'_P \cdot y''_P$;
  $y'_{f(P)} \leftarrow \beta^{-1} \cdot y'_P \cdot \alpha y''_P$;
  $z'_{f(P)} \leftarrow \gamma^{-1} \cdot z'_P \cdot \alpha y''_P$;
  $t'_{f(P)} \leftarrow \gamma^{-1} \cdot z''_P \cdot \beta x'_P$;
else if $y'_P \neq 0$ **then**
  $R'_1 \leftarrow R_1 + P'_1$;
  $R'_2 \leftarrow R_2 + Q'_1$;
  $x''_P, y''_P, z''_P, t''_P \leftarrow N(X(R'_1) \cdot X(R'_2), X(R'_1) \cdot Z(R'_2), Z(R'_1) \cdot X(R'_2), Z(R'_1) \cdot Z(R'_2))$;
  $x''_P, y''_P, z''_P, t''_P \leftarrow H(x''^2_P, y''^2_P, z''^2_P, t''^2_P)$;
  $\beta x''_P \leftarrow \beta \cdot x''_P$;
  $\alpha y'_P \leftarrow \alpha \cdot y'_P$;
  $x'_P \leftarrow \alpha^{-1} \cdot x'_P \cdot \beta x''_P$;
  $y'_{f(P)} \leftarrow y'_P \cdot x''_P$;
  $z'_{f(P)} \leftarrow \gamma^{-1} \cdot z'_P \cdot \beta x''_P$;
  $t'_{f(P)} \leftarrow \gamma^{-1} \cdot z''_P \cdot \alpha y'_P$;
else
  $R'_1 \leftarrow R_1 + P'_2$;
  $R'_2 \leftarrow R_2 + Q'_2$;
  $x''_P, y''_P, z''_P, t''_P \leftarrow N(X(R'_1) \cdot X(R'_2), X(R'_1) \cdot Z(R'_2), Z(R'_1) \cdot X(R'_2), Z(R'_1) \cdot Z(R'_2))$;
  $x''_P, y''_P, z''_P, t''_P \leftarrow H(x''^2_P, y''^2_P, z''^2_P, t''^2_P)$;
  $\gamma x''_P \leftarrow \gamma \cdot x''_P$;
  $\alpha z'_P \leftarrow \alpha \cdot z'_P$;
  $x'_P \leftarrow \alpha^{-1} \cdot x'_P \cdot \gamma x''_P$;
  $y'_{f(P)} \leftarrow \beta^{-1} \cdot y'_P \cdot \gamma x''_P$;
  $z'_{f(P)} \leftarrow z'_P \cdot x''_P$;
  $t'_{f(P)} \leftarrow \beta^{-1} \cdot y''_P \cdot \alpha z'_P$;
end
$x_{f(P)}, y_{f(P)}, z_{f(P)}, t_{f(P)} \leftarrow H(x'_{f(P)}, y'_{f(P)}, z'_{f(P)}, t'_{f(P)})$;
return $(x_{f(P)} : y_{f(P)} : z_{f(P)} : t_{f(P)})$; // Total cost: $71\mathbf{M} + 18\mathbf{S}$

Algorithm 6.17: Gluing 2-dimensional 2-isogeny codomain computation from 8-torsion points.

Data: Theta coordinates of $T_1, T_2 \in (E_1 \times E_2)[8]$ forming a maximal isotropic subgroup, such that $\ker(f) = \langle [4]T_1, [4]T_2 \rangle$.
Result: Codomain theta null point $(a' : b' : c' : d')$ of f, its dual $(\alpha : \beta : \gamma : 0)$ and the projective inverse $(\alpha^{-1} : \beta^{-1} : \gamma^{-1} : 0)$.

$(x\alpha, x\beta, y\gamma, 0) \leftarrow H(x_{T_1}^2, y_{T_1}^2, z_{T_1}^2, t_{T_1}^2)$;
$(z\alpha, t\beta, z\gamma, 0) \leftarrow H(x_{T_2}^2, y_{T_2}^2, z_{T_2}^2, t_{T_2}^2)$;
$\alpha \leftarrow x\alpha \cdot z\alpha$;
$\beta \leftarrow x\beta \cdot z\alpha$;
$\gamma \leftarrow x\alpha \cdot z\gamma$;
$\alpha^{-1} \leftarrow x\beta \cdot z\gamma$;
$\beta^{-1} \leftarrow \gamma$;
$\gamma^{-1} \leftarrow \beta$;
$(a', b', c', d') \leftarrow H(\alpha, \beta, \gamma, 0)$;
return $(a' : b' : c' : d')$, $(\alpha : \beta : \gamma : 0)$, $(\alpha^{-1} : \beta^{-1} : \gamma^{-1} : 0)$; // Total cost: $4\mathbf{M} + 8\mathbf{S}$

of Lemma 6.4.1 ensure the existence of a product symplectic basis $\mathscr{C}_0$ of $(E_3 \times E_4)[4]$ such that the symplectic change of basis matrix from $\mathscr{C}$ to $\mathscr{C}_0$ is given by:

$$M := \begin{pmatrix} 0 & 0 & -1 & 0 \\ 0 & b & 0 & 0 \\ 1 & -b & 0 & 0 \\ 0 & 0 & 1 & \beta \end{pmatrix},$$

with $\beta b \equiv 1 \mod 4$. From M and Theorem 6.2.10, we obtain a change of theta coordinates matrix N to product theta coordinates

$$\begin{aligned}(\theta'_{00}(P,Q), \theta'_{10}(P,Q), \theta'_{01}(P,Q), \theta'_{11}(P,Q)) \\ \overset{N}{\longmapsto} (\theta_0^{E_3}(P) \cdot \theta_0^{E_4}(Q), \theta_1^{E_3}(P) \cdot \theta_0^{E_4}(Q), \theta_0^{E_3}(P) \cdot \theta_1^{E_4}(Q), \theta_1^{E_3}(P) \cdot \theta_1^{E_4}(Q))\end{aligned}$$

and we can obtain theta coordinates $(\theta_0^{E_i} : \theta_1^{E_i})$ on E_i for $i \in \{3, 4\}$ with Algorithm 6.9. In particular, we are able to extract the theta null points $(a_i : b_i)$ of E_i for $i \in \{3, 4\}$ and the Montgomery equations up to quadratic twist:

$$E_i : \quad BY^2Z = X(X - \alpha_i Z)(Z - 1/\alpha_i Z),$$

with $\alpha_i := (a_i^2 + b_i^2)/(a_i^2 - b_i^2)$ for $i \in \{3, 4\}$ by Proposition 5.3.47.

Without Kani's lemma

When we do not use Kani's lemma, we no longer keep track of the codomain theta structure. Consequently, to find a product theta structure on the codomain $E_3 \times E_4$, we can enumerate change of theta via symmetric Heisenberg automorphisms until we find a product theta structure. In practice, we can consider symplectic matrices in $M \in \mathrm{Sp}_4(\mathbb{Z}/4\mathbb{Z})$ and the associated change of theta coordinate matrices $N \in M_4(k)$ computed with Theorem 6.2.10. By Proposition 6.5.9.(ii) one of the even theta constants always $U_{i,j}^{\mathcal{L}^2}(0_A)$ vanishes and we know that we have found a product theta structure when the associated dual theta constant $U_{11,11}^{\mathcal{L}^2}(0_A)$ vanishes. Hence, we do not have to try each of the 979200 symplectic matrices $M \in \mathrm{Sp}_4(\mathbb{Z}/4\mathbb{Z})$, we can stop when we have found the right associated change of theta coordinates N so that $U_{11,11}^{\mathcal{L}^2}(0_A) = 0$. Actually, we can even predict in advance which change of theta coordinate N to choose depending on the even index (i, j) of the vanishing even dual theta constant $U_{i,j}^{\mathcal{L}^2}(0_A)$ for the original theta structure. Only 10 different matrices N cover all cases. Algorithm 6.18 follows.

Algorithm 6.18: Finding the product theta structure on the codomain.

Data: The codomain (non-product) theta null point $(a : b : c : d)$ of $E_3 \times E_4$ obtained at the end of the 2-isogeny chain computation F.

Result: A change of theta coordinates matrix N from the original (non-product) theta coordinates to product theta coordinates.

$x_0, x_1, x_2, x_3 \leftarrow a, b, c, d$;

$S \leftarrow \{(00,00), (00,10), (00,10), (00,11), (10,00), (10,01), (01,00), (01,10), (11,0), (11,11)\}$;

for $(i,j) \in S$ **do**

  Compute $U_{i,j} \leftarrow \sum_{t \in (\mathbb{Z}/2\mathbb{Z})^2} (-1)^{\langle i|t \rangle} x_{i+t} x_t$;

end

Let $(i,j) \in S$ such that $U_{i,j} = 0$;

if $(i,j) = (00,00)$ **then**

$$N \leftarrow \begin{pmatrix} 1 & \sqrt{-1} & 1 & \sqrt{-1} \\ 1 & -\sqrt{-1} & -1 & \sqrt{-1} \\ 1 & \sqrt{-1} & -1 & -\sqrt{-1} \\ -1 & \sqrt{-1} & -1 & \sqrt{-1} \end{pmatrix};$$

else if $(i,j) = (01,00)$ **then**

$$N \leftarrow \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & -1 & 1 \\ 1 & 1 & -1 & -1 \\ -1 & 1 & -1 & 1 \end{pmatrix};$$

else if $(i,j) = (10,00)$ **then**

$$N \leftarrow \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & -1 & -1 & 1 \\ -1 & -1 & 1 & 1 \end{pmatrix};$$

else if $(i,j) = (11,00)$ **then**

$$N \leftarrow \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ -1 & 1 & 1 & -1 \end{pmatrix};$$

else if $(i,j) = (00,01)$ **then**

$$N \leftarrow \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & -1 & 0 & 0 \end{pmatrix};$$

else if $(i,j) = (10,01)$ **then**

$$N \leftarrow \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & -1 & -1 & 1 \\ 1 & 1 & -1 & -1 \end{pmatrix};$$

else if $(i,j) = (00,10)$ **then**

$$N \leftarrow \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \end{pmatrix};$$

else if $(i,j) = (01,10)$ **then**

$$N \leftarrow \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix};$$

else if $(i,j) = (00,11)$ **then**

$$N \leftarrow \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix};$$

else if $(i,j) = (11,11)$ **then**

  $N \leftarrow I_4$;

return N ; // Total cost: 40**M**

6.5.5 Performance results

The algorithms described in this section have been implemented in SageMath and Rust and the code[1] has been published with the paper [DMPR25]. In Tables 6.1 and 6.2, we compare our timings with algorithms based on Jacobian or Kummer models to compute and evaluate a 2^e-isogeny over the base field $k = \mathbb{F}_{p^2}$ for different field sizes $\log_2(p)$ and isogeny chain lengths e. We improve timings by a factor ~ 5 at least for both computation and evaluation. This efficient implementation significantly contributed to the dynamic of isogeny based cryptography following the SIDH attacks. It has been used in various isogeny based protocols using Kani's lemma, and recently, a variant of this implementation in C has been used in SQIsign2D-West [BDF+25] and the subsequent round 2 SQIsign NIST submission [AAA+25].

Field size $\log_2(p)$	Length e	Theta Rust [DMPR25]	Theta Sage [DMPR25]	Richelot Sage [OP22]	Richelot Sage [Kun22]	Kummer Sage [Kun22]
254	126	**2.13**	**108**	1028	760	467
381	208	**9.05**	**201**	1998	1478	858
1293	632	**463**	**1225**	12840	9196	5150

Table 6.1: Execution times in ms for a 2-dimensional 2^e-isogeny (intermediate codomains) computation over the base field $\mathbb{F}_{p^2}$ in Rust and SageMath on an Intel Core i7-9750H (2.6 GHz) CPU. Source: [DMPR25, Tables 2 & 3].

Field size $\log_2(p)$	Length e	Theta Rust [DMPR25]	Theta Sage [DMPR25]	Richelot Sage [OP22]	Richelot Sage [Kun22]	Kummer Sage [Kun22]
254	126	**0.161**	**5.43**	114	66.7	18.4
381	208	**0.411**	**8.68**	208	119	31.4
1293	632	**17.8**	**40.8**	1203	593	170

Table 6.2: Execution times in msfor a 2-dimensional 2^e-isogeny point evaluation over the base field $\mathbb{F}_{p^2}$ in Rust and SageMath on an Intel Core i7-9750H (2.6 GHz) CPU. Source: [DMPR25, Table 2].

[1] Available at https://github.com/ThetaIsogenies/two-isogenies.

6.6 Implementation in dimension 4

In this section, we apply the general algorithms from Sections 6.1 to 6.3 to the computation of a 4-dimensional 2^e-isogeny $F \in \mathrm{End}(E_1^2 \times E_2^2)$ obtained from Kani's lemma with applications to SIDH attacks (Section 2.2.4) and ideal-to-isogeny translations in the context of SQIsignHD (Section 2.3). Let $\sigma : E_1 \longrightarrow E_2$ be an isogeny of odd degree q that we want to interpolate. We assume we are given a basis (P_1, P_2) of $E_1[2^f]$ (where $f \geq e/2 + 2$) and its image $(\sigma(P_1), \sigma(P_2))$ and $a_1, a_2 \in \mathbb{Z}$ such that $a_1^2 + a_2^2 + q = 2^e$. We compute the 2^e-isogeny:

$$F := \begin{pmatrix} \alpha_1 & \widetilde{\Sigma} \\ -\Sigma & \widetilde{\alpha}_2 \end{pmatrix} \in \mathrm{End}(E_1^2 \times E_2^2), \tag{6.22}$$

where $\Sigma := \mathrm{Diag}(\sigma, \sigma) : E_1^2 \longrightarrow E_2^2$ and for $i \in \{1, 2\}$,

$$\alpha_i := \begin{pmatrix} a_1 & a_2 \\ -a_2 & a_1 \end{pmatrix} \in \mathrm{End}(E_i^2).$$

We keep these notations in the following.

6.6.1 Locating gluings

As we have seen in Section 6.1.3, gluing isogenies are computed differently and may be more costly than generic ones to compute. Hence, for practicality and efficiency reasons, gluing isogenies (including diagonal isogenies) should be located in advance in the 2-isogeny chain. The following lemma ensures that gluing and diagonal isogenies in dimension 2 appear in the first steps of the 2-isogeny chain F.

Lemma 6.6.1. *Assume that $2|a_2$ and let $m := v_2(a_2)$ be its 2-adic valuation. Then $F := f_e \circ \cdots \circ f_1$, with*

$$E_1^2 \times E_2^2 \xrightarrow{f_1} A_1^2 \quad \cdots \quad A_{m-1}^2 \xrightarrow{f_m} A_m^2 \xrightarrow{f_{m+1}} B \longrightarrow \quad \cdots,$$

a chain of 2-isogenies, where the A_i are abelian surfaces and B is an abelian variety of dimension 4. We have:

(i) $f_1 = \mathrm{Diag}(\varphi_1, \varphi_1) \circ S$, with $S : (R_1, S_1, R_2, S_2) \in E_1^2 \times E_2^2 \longmapsto (R_1, R_2, S_1, S_2) \in (E_1 \times E_2)^2$ and $\varphi_1 : E_1 \times E_2 \longrightarrow A_1$ a gluing isogeny.

(ii) For all $i \in [\![2\ ;\ m]\!]$, f_i is a diagonal isogeny $\mathrm{Diag}(\varphi_i, \varphi_i)$, with $\varphi_i : A_{i-1} \longrightarrow A_i$.

(iii) f_{m+1} is a gluing isogeny.

(iv) $\ker(\varphi_m \circ \cdots \circ \varphi_1) = \{([a_1]P, \sigma(P)) \mid P \in E_1[2^m]\}$.

Proof. Kani's lemma ensures that:

$$\begin{aligned} \ker(F) &= \{(\widetilde{\alpha}_1(P, Q), \Sigma(P, Q)) \mid P, Q \in E_1[2^e]\} \\ &= \{([a_1]P - [a_2]Q, [a_2]P + [a_1]Q, \sigma(P), \sigma(Q)) \mid P, Q \in E_1[2^e]\}. \end{aligned}$$

Let $f_1, \cdots, f_{m+1}$ be the $m+1$ first elements of the 2-isogeny chain F. Then, since $a_2 \equiv 0 \mod 2^m$, we have

$$\ker(f_m \circ \cdots \circ f_1) = [2^{e-m}] \ker(F) = K_1 \oplus K_2,$$

where $K_1 := \{([a_1]P, 0, \sigma(P), 0) \mid P \in E_1[2^m]\}$ and $K_2 := \{(0, [a_1]P, 0, \sigma(P)) \mid P \in E_1[2^m]\}$. This proves the chain $f_m \circ \cdots \circ f_1$ has the desired form. This completes the proof. □

6.6.2 An overview of the isogeny chain computation

To compute the 2^e-isogeny $F \in \mathrm{End}(E_1^2 \times E_2^2)$ defined in Eq. (6.22), we assume that we are given a basis (P, Q) of $E_1[2^{e+2}]$ and its image $(\sigma(P), \sigma(Q))$ by σ. We also assume that a_2 is even (swapping a_1 and a_2 if necessary) in order to apply Lemma 6.6.1. We can decompose F as a 2-isogeny chain:

$$E_1^2 \times E_2^2 \xrightarrow{f_1} A_1^2 \quad \cdots \quad A_{m-1}^2 \xrightarrow{f_m} A_m^2 \xrightarrow{f_{m+1}} B_{m+1} \cdots B_{e-2} \xrightarrow{f_{e-1}} B_{e-1} \xrightarrow{f_e} E_1^2 \times E_2^2,$$

where $f_1, \cdots, f_{m+1}$ have been described in Lemma 6.6.1. To compute this 2-isogeny chain, we proceed as follows:

- We compute the chain of 2-dimensional 2-isogenies $\Phi := \varphi_m \circ \cdots \circ \varphi_1$ introduced in Lemma 6.6.1 with the algorithms from Section 6.5, as we shall explain in Section 6.6.3.
- We compute the 4-dimensional gluing isogeny of abelian surfaces $f_{m+1} : A_m^2 \longrightarrow B_{m+1}$ introduced in Lemma 6.6.1 with the algorithms from Section 6.1.3, as we shall see in Section 6.6.4.
- For all $i \in [\![m+2\ ;\ e]\!]$, we compute the 4-dimensional 2-isogeny $f_i : B_{i-1} \longrightarrow B_i$ as a generic isogeny, using the algorithms from Sections 6.1.1 and 6.1.2 and a computational strategy (that we shall introduce in Section 6.6.5) avoiding point duplications on the domain of splitting isogenies.
- We recover a product theta structure on $B_e = E_1^2 \times E_2^2$ in order to express point images by F in $(X : Z)$-Montgomery coordinates in the product $E_1^2 \times E_2^2$. This will be explained in Section 6.6.6.

In Section 6.6.7, we explain how to adapt this method when a basis of the 2^{e+2}-torsion of E_1 and its image by σ is not available, but only the a basis of the 2^f-torsion and its image, with $f \geq e/2+2$. In this case, we have to decompose F into $F = F_2 \circ F_1$ and compute F_1 and $\tilde{F}_2$, as explained in Section 6.4.2.

6.6.3 The first gluing in dimension 2

In this section, we explain how to compute the 2-dimensional 2^m-isogeny $\Phi := \varphi_m \circ \cdots \circ \varphi_1 : E_1 \times E_2 \longrightarrow A_m$ from Lemma 6.6.1. We start by finding a symplectic basis of $(E_1 \times E_2)[2^{m+2}]$ adapted to Φ, in the sense of Theorem 6.1.1. Recall that we are given a basis (P, Q) of $E_1[2^{e+2}]$ and its image $(\sigma(P), \sigma(Q))$. Let us denote $\zeta := e_{2^{e+2}}(P, Q)$.

Lemma 6.6.2. *Let r and b_1 be a modular inverses of q and a_1 modulo 2^{m+2} respectively. Consider $\mathscr{B}_1 := (J_1, J_2, K_1, K_2)$ with:*

$$J_1 := (-[2^{e-m}b_1]Q, 0), \quad J_2 := (0, [2^{e-m}r]\sigma(P)),$$

$$K_1 := ([2^{e-m}a_1]P, [2^{e-m}]\sigma(P)), \quad K_2 := ([2^{e-m}a_1(1+a_2^2b_1^2)]Q, [2^{e-m}]\sigma(Q)).$$

Then $\mathscr{B}_1$ is a $\zeta^{2^{e-m}}$-symplectic basis of $(E_1 \times E_2)[2^{m+2}]$ adapted to Φ.

Proof. As in Section 6.5, we consider the canonical line bundle $\mathcal{L}_i := \mathcal{L}((0_{E_i}))$ inducing the principal polarisation on E_i for all $i \in \{1, 2\}$ and the product principal polarisation induced by $\mathcal{L}_1 \times \mathcal{L}_2$ on $E_1 \times E_2$. Then, Proposition 5.1.8 ensures that we can identify the commutator pairing $e_{\mathcal{L}_1^n \star \mathcal{L}_2^n}$ with the Weil pairing e_n for all $n \in \mathbb{N} \setminus p\mathbb{N}$. Hence, we have to prove that $\mathscr{B}_1$ is $\zeta^{2^{e-m}}$-symplectic with respect to $e_{2^{m+2}}$ in the following. By Lemma 6.4.2 and the usual properties of the Weil pairing (see Section 1.4.6), we have:

$$\begin{aligned} e_{2^{m+2}}(J_1, J_2) &= e_{2^{m+2}}((-[2^{e-m}b_1]Q, 0), (0, [2^{e-m}r]\sigma(P))) = e_{2^{e+2}}((-[b_1]Q, 0), (0, [r]\sigma(P)))^{2^{e-m}} \\ &= e_{2^{e+2}}(-[b_1]Q, 0)^{2^{e-m}} e_{e^{e+2}}(0, [r]\sigma(P))^{2^{e-m}} = 1 \end{aligned}$$

By similar computations, we obtain that:

$$\begin{aligned} e_{2^{m+2}}(K_1, K_2) &= e_{2^{e+2}}([a_1]P, [a_1(1+a_2^2b_1^2)]Q)^{2^{e-m}} e_{2^{e+2}}(\sigma(P), \sigma(Q))^{2^{e-m}} \\ &= e_{2^{e+2}}(P, Q)^{2^{e-m}(a_1^2(1+a_2^2b_1^2)+q)} = e_{2^{e+2}}(P, Q)^{2^{e-m}(a_1^2+a_2^2+q)} = e_{2^{e+2}}(P, Q)^{2^{2e-m}} = 1, \end{aligned}$$

where we used the fact that $2e - m \geq e + 2$ *i.e.* $e \geq m + 2$, and that:

$$e_{2^{m+2}}(J_1, K_1) = e_{2^{e+2}}(Q, P)^{-2^{e-m}a_1b_1} = e_{2^{e+2}}(P, Q)^{2^{e-m}} = \zeta^{2^{e-m}}$$

$$e_{2^{m+2}}(J_2, K_2) = e_{2^{e+2}}(\sigma(P), \sigma(Q))^{2^{e-m}r} = e_{2^{e+2}}(P, Q)^{2^{e-m}qr} = \zeta^{2^{e-m}}$$

$$e_{2^{m+2}}(J_1, K_2) = e_{2^{e+2}}(Q, Q)^{2^{e-m}b_1a_1(1+a_2^2b_1^2)} = 1$$

$$e_{2^{m+2}}(J_2, K_1) = e_{2^{e+2}}(\sigma(P), \sigma(P))^{2^{e-m}ra_1} = 1$$

This proves that $\mathscr{B}_1$ is indeed a $\zeta^{2^{e-m}}$-symplectic basis of $(E_1 \times E_2)[2^{m+2}]$.

Now, since $2^m | a_2$, we have $2^{e+2-m}a_2^2 \equiv 0 \mod 2^{e+2}$, so that

$$\begin{aligned}\langle [4]K_1, [4]K_2\rangle &= \langle ([2^{e+2-m}a_1]P, [2^{e+2-m}]\sigma(P)), ([2^{e+2-m}a_1]Q, [2^{e+2-m}]\sigma(Q))\rangle \\ &= \{([a_1]R, \sigma(R)) \mid R \in E_1[2^m]\} = \ker(\Phi)\end{aligned}$$

by Lemma 6.6.1 and $\mathscr{B}_1$ is adapted to Φ. This completes the proof. □

Using Lemma 6.6.2, we can obtain level 2 theta coordinates on $E_1 \times E_2$ adapted to Φ from Montgomery $(X : Z)$-coordinates on E_1 and E_2, as explained in Section 6.5.1. We can compute Φ as a chain of 2-isogenies from the knowledge of J_1, J_2 (derived from the knowledge of $(\sigma(P), \sigma(Q))$) with the general approach introduced in the beginning of Section 6.5 for 2-dimensional isogenies. The first isogeny φ_1 is expected to be a non-diagonal gluing so Algorithms 6.17 and 6.16 should be used to compute and evaluate it. The following isogenies $\varphi_2, \cdots, \varphi_m$ are expected to be generic isogenies so Algorithms 6.13 and 6.12 should be used to compute and evaluate them. Algorithm 5.5 should also be used for point duplications. A computational strategy may be useful if m is big enough, which is generally not the case.

6.6.4 The second gluing in dimension 4

Assume that we have computed $\Phi := \varphi_m \circ \cdots \circ \varphi_1$, so that we know $f_m \circ \cdots \circ f_1 = \mathrm{Diag}(\Phi, \Phi) \circ S$. In this section, we explain how to compute the 4-dimensional gluing $f_{m+1} : A_m^2 \longrightarrow B_{m+1}$.

Consider a ζ-symplectic basis $\mathscr{B}_2 := (S_1, \cdots, S_4, T_1, \cdots, T_4)$ of $(E_1^4 \times E_2^4)[2^{e+2}]$ adapted to F given by Lemma 6.4.1.(i):

$$S_1 := (-[s]\widetilde{\alpha}_1(Q, 0), 0, 0) = (-[sa_1]Q, -[sa_2]Q, 0, 0),$$

$$S_2 := (-[s]\widetilde{\alpha}_1(0, Q), 0, 0) = ([sa_2]Q, -[sa_1]Q, 0, 0),$$

$$S_3 := (0, 0, [r]\sigma(P), 0), \quad S_4 := (0, 0, 0, [r]\sigma(P)),$$

$$T_1 := ([a_1]P, [a_2]P, \sigma(P), 0), \quad T_2 := (-[a_2]P, [a_1]P, 0, \sigma(P)),$$

$$T_3 := ([(1 - s2^e)a_1]Q, [(1 - s2^e)a_2]Q, \sigma(Q), 0), \quad T_4 := (-[(1 - s2^e)a_2]Q, [(1 - s2^e)a_1]Q, 0, \sigma(Q)),$$

where r and s are respectively the inverse of q and $a_1^2 + a_2^2$ modulo 2^{e+2}. Then Lemma 6.3.4.(ii) ensures that for all $i \in [\![m\ ;\ e-1]\!]$,

$$\begin{aligned}\mathscr{C}_i := ([2^{e-1}]f_i \circ \cdots \circ f_1(S_1), \cdots, [2^{e-1}]f_i \circ \cdots \circ f_1(S_g), \\ [2^{e-1-i}]f_i \circ \cdots \circ f_1(T_1), \cdots, [2^{e-1-i}]f_i \circ \cdots \circ f_1(T_g))\end{aligned}$$

is a $\zeta^{2^{e-1}}$-symplectic basis of $B_i[8]$ adapted to $f_{i+1} : B_i \longrightarrow B_{i+1}$. In particular, $\mathscr{C}_m$ is adapted to f_{m+1}.

However, the level 2 theta coordinates naturally induced by the computation of $f_m \circ \cdots \circ f_1$ are not the ones induced by $[2]\mathscr{C}_m$. Indeed, the domain of f_{m+1} is the product of the principally polarised abelian surface $(A_m, \mathcal{L}_m)$ with itself and by Theorem 6.1.1, we have a level 2 theta structure on $(A_m, \mathcal{L}_m^2)$ induced by the ζ^{2^e}-symplectic basis of $A_m[4]$:

$$\mathscr{D}_0 := ([2^m]\Phi(J_1), [2^m]\Phi(J_1), \Phi(K_1), \Phi(K_2)),$$

where (J_1, J_2, K_1, K_2) has been defined in Lemma 6.6.2. Images by Φ are expressed in level 2 theta coordinates induced by $\mathscr{D}_0$. Hence, we can easily express images by $f_m \circ \cdots \circ f_1$ in the level 2 product theta coordinates induced by $\mathscr{D}_0 \times \mathscr{D}_0$. To express these images in theta coordinates adapted to f_{m+1}, we need to apply the change of coordinates formula from Theorem 6.2.10 which depends on the symplectic change of basis matrix from $\mathscr{D}_0 \times \mathscr{D}_0$ to $[2]\mathscr{C}_m$.

Lemma 6.6.3. *The change of coordinates matrix from $\mathscr{D}_0 \times \mathscr{D}_0$ to $[2]\mathscr{C}_m$ is given by:*

$$\begin{pmatrix}
sa_1^2 & -sa_1a_2 & 0 & 0 & 0 & 0 & a_1a_2^2/2^m & a_1a_2 \\
0 & 0 & 1 & 0 & 0 & b_1qa_2/2^m & 0 & 0 \\
sa_1a_2 & sa_1^2 & 0 & 0 & 0 & 0 & -a_1a_2 & a_1a_2^2/2^m \\
0 & 0 & 0 & 1 & -b_1qa_2/2^m & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & -b_1a_2 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & 0 & b_1a_2 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1
\end{pmatrix}$$

Proof. Let us denote $G := f_m \circ \cdots \circ f_1$. Then we have:

$$\begin{aligned}
[2^e]G(S_1) &= (\Phi(-[2^e sa_1]Q, 0), \Phi(-[2^e sa_2]Q, 0)) = (\Phi(-[sa_1^2 2^e b_1]Q, 0), \Phi(-[sa_1a_2 2^e b_1]Q, 0)) \\
&= ([sa_1^2][2^m]\Phi(J_1), [sa_1a_2][2^m]\Phi(J_1))
\end{aligned}$$

$$\begin{aligned}
[2^e]G(S_2) &= (\Phi([2^e sa_2]Q, 0), \Phi(-[2^e sa_1]Q, 0)) = (\Phi([sa_1a_2 2^e b_1]Q, 0), \Phi(-[sa_1^2 2^e b_1]Q, 0)) \\
&= (-[sa_1a_2][2^m]\Phi(J_1), [sa_1^2][2^m]\Phi(J_1))
\end{aligned}$$

$$[2^e]G(S_3) = (\Phi(0, [2^e r]\sigma(P)), 0) = ([2^m]\Phi(J_2), 0)$$

$$[2^e]G(S_4) = (0, \Phi(0, [2^e r]\sigma(P))) = (0, [2^m]\Phi(J_2))$$

$$[2^{e-m}]G(T_1) = (\Phi([2^{e-m}a_1]P, [2^{e-m}]\sigma(P)), \Phi([2^{e-m}a_2]P, 0)) = (\Phi(K_1), \Phi([2^{e-m}a_2]P, 0)),$$

with:

$$\begin{aligned}
\Phi([2^{e-m}a_2]P, 0) &= \Phi([b_1a_2 2^{e-m}a_1]P, [b_1a_2 2^{e-m}]\sigma(P)) - \Phi(0, [b_1a_2 2^{e-m}]\sigma(P)) \\
&= [b_1a_2]\Phi(K_1) - \Phi(0, [b_1a_2q2^{e-m}r]\sigma(P)) = [b_1a_2]\Phi(K_1) - [b_1qa_2/2^m][2^m]\Phi(J_2),
\end{aligned}$$

so that:

$$[2^{e-m}]G(T_1) = (\Phi(K_1), [b_1a_2]\Phi(K_1) - [b_1qa_2/2^m][2^m]\Phi(J_2))$$

It follows that:

$$\begin{aligned}
[2^{e-m}]G(T_2) &= (\Phi(-[2^{e-m}a_2]P, 0), \Phi([2^{e-m}a_1]P, [2^{e-m}]\sigma(P))) \\
&= (-[b_1a_2]\Phi(K_1) + [b_1qa_2/2^m][2^m]\Phi(J_2), \Phi(K_1)).
\end{aligned}$$

Besides,

$$\begin{aligned}
[2^{e-m}]G(T_3) &= (\Phi([2^{e-m}(1 - s2^e)a_1]Q, [2^{e-m}]\sigma(Q)), \Phi([2^{e-m}(1 - s2^e)a_2]Q, 0)) \\
&= (\Phi([2^{e-m}a_1]Q, [2^{e-m}]\sigma(Q)), \Phi([a_1a_2 2^{e-m}b_1]Q, 0)) \\
&= (\Phi([2^{e-m}a_1]Q, [2^{e-m}]\sigma(Q)), -[a_1a_2]\Phi(J_1)),
\end{aligned}$$

where we used the fact that $m + 2 \leq e$ and where:

$$\begin{aligned}
\Phi([2^{e-m}a_1]Q, [2^{e-m}]\sigma(Q)) &= \Phi([2^{e-m}a_1(1 + a_2^2b_1^2)]Q, [2^{e-m}]\sigma(Q)) - \Phi([2^{e-m}a_2^2b_1]Q, 0) \\
&= \Phi(K_2) - \Phi([a_1a_2^2 2^{e-m}b_1]Q, 0) = \Phi(K_2) + [a_1a_2^2/2^m][2^m]\Phi(J_1),
\end{aligned}$$

so that:

$$[2^{e-m}]G(T_3) = (\Phi(K_2) + [a_1a_2^2/2^m][2^m]\Phi(J_1), -[a_1a_2]\Phi(J_1)).$$

Finally,

$$\begin{aligned}
[2^{e-m}]G(T_4) &= (\Phi(-[2^{e-m}(1 - s2^e)a_2]Q, 0), \Phi([2^{e-m}(1 - s2^e)a_1]Q, \sigma(Q))) \\
&= ([a_1a_2]\Phi(J_1), \Phi(K_2) + [a_1a_2^2/2^m][2^m]\Phi(J_1)).
\end{aligned}$$

This completes the proof. □

Once we have computed the change of theta coordinates to obtain theta coordinates adapted to f_{m+1}, we express the images of $[2^{e-m-1}]T_1, \cdots, [2^{e-m-1}]T_4$ and $[2^{e-m-1}](T_1+T_2)$ by $f_m \circ \cdots \circ f_1$ with these new theta coordinates. In practice (see Conjecture 6.1.15.(ii)), this data is enough to recover the codomain dual theta null point of f_{m+1} and we can apply (a gluing version of) Algorithm 6.5 with these 8-torsion points as input.

Then, to evaluate a point $R \in E_1^2 \times E_2^2$ via $f_{m+1} \circ \cdots \circ f_1$, we proceed as follows. We compute $R + [2^{e-m}]T_1$ and $R + [2^{e-m}]T_2$ with Weierstrass or Jacobian coordinates on the elliptic product $E_1^2 \times E_2^2$. Then, we evaluate $R, R + [2^{e-m}]T_1$ and $R + [2^{e-m}]T_2$ through the chain:

$$f_m \circ \cdots \circ f_1 = \mathrm{Diag}(\varphi_m \circ \cdots \circ \varphi_1, \varphi_m \circ \cdots \circ \varphi_1) \circ S.$$

The image points can then be used in Algorithm 6.6 to finally obtain $f_{m+1} \circ \cdots \circ f_1(R)$. In practice, the evaluation succeeds (Conjecture 6.1.11).

Alternatively, we may use the faster Algorithm 6.7 with $R + [2^{e-m-1}]T_1$ and $R - [2^{e-m-1}]T_2$ instead of $R + [2^{e-m}]T_1$ and $R + [2^{e-m}]T_2$.

6.6.5 Computing the generic 2-isogenies in the chain

Once we have computed the gluing isogenies $f_1, \cdots, f_{m+1}$, we compute the remaining 2-isogenies of the chain f_i (for $i \in [\![m+2\ ;\ e]\!]$) as generic isogenies with Algorithm 6.5. This algorithm takes as input the 8 torsion points $[2^{e-i-2}]f_{i-1} \circ \cdots \circ f_1(T_l)$ for $l \in [\![1\ ;\ 4]\!]$ to output the codomain dual theta null point of f_i. Hence, we have to compute the leaves of the computation tree mentioned in Section 6.3.3.

Quasi-linear divide and conquer strategies introduced in Section 6.3.3 also apply but they have to account for the already computed gluing isogeny chain $f_{m+1} \circ \cdots \circ f_1$, so the strategy:

- Must be of length $e - m$ instead of e (to start at $f_{m+1} \circ \cdots \circ f_1$, considered as *one* isogeny).
- Must take into account the higher cost of evaluation by the "first" isogeny $f_{m+1} \circ \cdots \circ f_1$.

Recall the definition of a strategy from Eq. (6.16) as a subgraph of the tree T_e defined in Eq. (6.16), where nodes are basis of points, left branches of the tree represent point duplications and right branches 2-isogeny evaluations. We now consider a different measure for strategies than the one introduced in Section 6.3.3 taking into account the additional cost of the "first" isogeny evaluation. We denote by μ' this measure parametrized by $(\alpha, \beta, \gamma) \in \mathbb{R}_+^3$ where α is the cost of a left edge (accounting for duplication cost), β is the cost of a right edge not starting from the root (accounting for a generic evaluation cost) and γ is the cost of a right edge starting from the root (accounting for the evaluation by the "first" isogeny).

Lemmas 6.3.6 and 6.3.7 ([JDF11, Lemma 4.3 and 4.5]) can easily be generalised to strategies with the measure μ'. Namely, we obtain that optimal strategies for μ' are canonical (hence determined by their tree topology) and that their left and right branches are optimal strategies.

Therefore, a dynamic programming approach is still valid here. Assuming we have computed optimal strategies $S'_1, \cdots, S'_{n-1}$ and $S_1, \cdots, S_{n-1}$ within the trees $T_1, \cdots, T_{n-1}$ for the measures μ parametrised by (α, β) from Section 6.3.3 and μ' parametrised by (α, β, γ) respectively, we can compute an optimal strategy S'_n for μ' within T_n with left branch S'_i and right branch S_{n-i}, where $i \geq 1$ is given by:

$$i := \underset{1 \leq j \leq n-1}{\mathrm{argmin}} \left(\mu'(S'_j) + \mu(S_{n-j}) + (n-j)\alpha + (j-1)\beta + \gamma\right).$$

6.6.6 Computing product (theta) coordinates on the codomain

Once we have computed the whole 2-isogeny chain $F = f_e \circ \cdots \circ f_1$, we want to be able to evaluate points $R \in E_1^2 \times E_2^2$ as:

$$F(R) = (F_1(R), \cdots, F_4(R)),$$

where the $F_i(R)$ are expressed in Montgomery $(X : Z)$-coordinates in $E_{\lceil i/2 \rceil}$ for $i \in [\![1\ ;\ 4]\!]$. In order to do that, we need to compute product theta coordinates on $E_1^2 \times E_2^2$ and to convert them into Montgomery $(X : Z)$-coordinates.

By Theorem 6.1.1, the theta coordinates naturally obtained after the isogeny chain computation are associated to the the theta structure induced by the ζ^{2^e}-symplectic basis $\mathscr{C} := ([2^e]F(S_1), \cdots, [2^e]F(S_4), F(T_1), \cdots, F(T_4))$ of $(E_1^2 \times E_2^2)[4]$, where $\mathscr{B}_2 := (S_1, \cdots, S_4, T_1, \cdots, T_4)$ has been defined in Section 6.6.4. Let $\mathscr{C}_0$ be the product ζ^{2^e}-symplectic basis of $(E_1^2 \times E_2^2)[4]$ given by:

$$\mathscr{C}_0 := (([2^e]P, 0, 0, 0), (0, [2^e]P, 0, 0), (0, 0, [2^e]\sigma(P), 0), (0, 0, 0, [2^e]\sigma(P)),$$
$$([2^e]Q, 0, 0, 0), (0, [2^e]Q, 0, 0), (0, 0, [2^e r]\sigma(Q), 0), (0, 0, 0, [2^e r]\sigma(Q))),$$

where r is a modular inverse of $q \mod 2^{e+2}$.

Lemma 6.6.4. *The change of basis matrix from $\mathscr{C}$ to $\mathscr{C}_0$ is given by:*

$$\begin{pmatrix} 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & -a_1 & -a_2 & 0 & 0 & 0 & 0 \\ 0 & 0 & a_2 & -a_1 & 0 & 0 & 0 & 0 \\ 1 & 0 & a_1 & a_2 & 0 & 0 & 0 & 0 \\ 0 & 1 & -a_2 & a_1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & ra_1 & ra_2 \\ 0 & 0 & 0 & 0 & 0 & 1 & -ra_2 & ra_1 \end{pmatrix}$$

Proof. Note that this is not a straightforward application of Lemma 6.4.1 given the shape of $\mathscr{C}_0$. We have:

$$\begin{aligned}[2^e]F(S_1) &= F(-[2^e s a_1]Q, -[2^e s a_2]Q, 0, 0) = (\alpha_1(-[2^e s a_1]Q, -[2^e s a_2]Q), [2^e s a_1]\sigma(Q), [2^e s a_2]\sigma(Q)) \\ &= (-[2^e s(a_1^2 + a_2^2)]Q, 0, [2^e s a_1]\sigma(Q), [2^e s a_2]\sigma(Q)) = (-[2^e]Q, 0, [2^e s a_1]\sigma(Q), [2^e s a_2]\sigma(Q))\end{aligned}$$

$$\begin{aligned}[2^e]F(S_2) &= F([2^e s a_2]Q, -[2^e s a_1]Q, 0, 0) = (\alpha_1([2^e s a_2]Q, -[2^e s a_1]Q), -[2^e s a_2]\sigma(Q), [2^e s a_1]\sigma(Q)) \\ &= (0, -[2^e s(a_2^2 + a_1^2)]Q, -[2^e s a_2]\sigma(Q), [2^e s a_1]\sigma(Q)) = (0, -[2^e]Q, -[2^e s a_2]\sigma(Q), [2^e s a_1]\sigma(Q))\end{aligned}$$

$$\begin{aligned}[2^e]F(S_3) &= F(0, 0, [2^e r]\sigma(P), 0) = (\widetilde{\Sigma}([2^e r]\sigma(P), 0), \widetilde{\alpha}_2([2^e r]\sigma(P), 0)) \\ &= ([2^e r]\widehat{\sigma} \circ \sigma(P), 0, [2^e r a_1]\sigma(P), [2^e r a_2]\sigma(P)) = ([2^e]P, 0, [2^e r a_1]\sigma(P), [2^e r a_2]\sigma(P))\end{aligned}$$

$$\begin{aligned}[2^e]F(S_4) &= F(0, 0, 0, [2^e r]\sigma(P)) = (\widetilde{\Sigma}(0, [2^e r]\sigma(P), 0), \widetilde{\alpha}_2(0, [2^e r]\sigma(P))) \\ &= (0, [2^e r]\widehat{\sigma} \circ \sigma(P), -[2^e r a_2]\sigma(P), [2^e r a_1]\sigma(P)) = (0, [2^e]P, -[2^e r a_2]\sigma(P), [2^e r a_1]\sigma(P))\end{aligned}$$

$$\begin{aligned}F(T_1) &= F([a_1]P, [a_2]P, \sigma(P), 0) = (\alpha_1([a_1]P, [a_2]P) + \widetilde{\Sigma}(\sigma(P), 0), -\Sigma([a_1]P, [a_2]P) + \widetilde{\alpha}_2(\sigma(P), 0)) \\ &= ([a_1^2 + a_2^2]P + \widehat{\sigma} \circ \sigma(P), 0, 0, 0) = ([2^e]P, 0, 0, 0)\end{aligned}$$

$$\begin{aligned}F(T_2) &= F(-[a_2]P, [a_1]P, 0, \sigma(P)) = (\alpha_1(-[a_2]P, [a_1]P) + \widetilde{\Sigma}(0, \sigma(P)), -\Sigma(-[a_2]P, [a_1]P) + \widetilde{\alpha}_2(0, \sigma(P))) \\ &= (0, [a_2^2 + a_1^2]P + \widehat{\sigma} \circ \sigma(P), 0, 0) = (0, [2^e]P, 0, 0)\end{aligned}$$

$$\begin{aligned}F(T_3) &= F([(1 - s2^e)a_1]Q, [(1 - s2^e)a_2]Q, \sigma(Q), 0) \\ &= (\alpha_1([(1 - s2^e)a_1]Q, [(1 - s2^e)a_2]Q) + \widetilde{\Sigma}(\sigma(Q), 0), \\ &\quad - \Sigma([(1 - s2^e)a_1]Q, [(1 - s2^e)a_2]Q) + \widetilde{\alpha}_2(\sigma(Q), 0)) \\ &= ([(1 - s2^e)(a_1^2 + a_2^2)]Q + \widehat{\sigma} \circ \sigma(Q), 0, [s2^e a_1]\sigma(Q), [s2^e a_2]\sigma(Q)) \\ &= (0, 0, [s2^e a_1]\sigma(Q), [s2^e a_2]\sigma(Q))\end{aligned}$$

$$F(T_4) = F(-[(1 - s2^e)a_2]Q, [(1 - s2^e)a_1]Q, 0, \sigma(Q))$$

$$\begin{aligned}
&= (\alpha_1(-[(1-s2^e)a_2]Q, [(1-s2^e)a_1]Q) + \widetilde{\Sigma}(0, \sigma(Q)),\\
&\quad - \Sigma(-[(1-s2^e)a_2]Q, [(1-s2^e)a_1]Q) + \widetilde{\alpha}_2(0, \sigma(Q)))\\
&= (0, [(1-s2^e)(a_2^2+a_1^2)]Q + \widehat{\sigma}\circ\sigma(Q), -[s2^e a_2]\sigma(Q), [s2^e a_1]\sigma(Q))\\
&= (0, 0, -[s2^e a_2]\sigma(Q), [s2^e a_1]\sigma(Q))
\end{aligned}$$

Inverting these relations, we obtain:

$$([2^e]P, 0, 0, 0) = F(T_1), \quad (0, [2^e]P, 0, 0) = F(T_2),$$

and:

$$\begin{aligned}
[2^e]F(S_3) - F(T_1) &= (0, 0, [2^e r a_1]\sigma(P), [2^e r a_2]\sigma(P)),\\
[2^e]F(S_4) - F(T_2) &= (0, 0, -[2^e r a_2]\sigma(P), [2^e r a_1]\sigma(P)),
\end{aligned}$$

so that:

$$\begin{aligned}
(0, 0, [2^e]\sigma(P), 0) &= -[a_1]([2^e]F(S_3) - F(T_1)) + [a_2]([2^e]F(S_4) - F(T_2)),\\
(0, 0, 0, [2^e]\sigma(P)) &= -[a_2]([2^e]F(S_3) - F(T_1)) - [a_1]([2^e]F(S_4) - F(T_2)),
\end{aligned}$$

where we used the fact that $r(a_1^2 + a_2^2) = r(2^e - q) \equiv -qr \equiv -1 \mod 4$. We also obtain:

$$([2^e]Q, 0, 0, 0) = -[2^e]F(S_1) + F(T_3), \quad (0, [2^e]Q, 0, 0) = -[2^e]F(S_2) + F(T_4)$$

$$(0, 0, [2^e r]\sigma(Q), 0) = [r a_1]F(T_3) - [r a_2]F(T_4), \quad (0, 0, 0, [2^e r]\sigma(Q)) = [r a_2]F(T_3) + [r a_1]F(T_4).$$

This completes the proof. □

Using the formulas from Theorem 6.2.10 with the matrix from Lemma 6.6.4 as input, we can compute the product theta coordinates associated to the product theta structure induced by $\mathscr{C}_0$ on $E_1^2 \times E_2^2$ from the non-product theta-coordinates associated to the theta structure naturally induced by $\mathscr{C}$. Applying Algorithm 6.9 twice, we can then decompose images of F as:

$$F(R) = (F_1(R), \cdots, F_4(R)),$$

with $F_1(R)$, $F_2(R)$ written in level 2 theta coordinates $(\theta_0^{E_1} : \theta_1^{E_1})$ associated to the theta structure induced by $([2^e]P, [2^e]Q)$ on E_1 and $F_3(R)$, $F_4(R)$ written in level 2 theta coordinates $(\theta_0^{E_2} : \theta_1^{E_2})$ associated to the theta structure induced by $([2^e]\sigma(P), [2^e r]\sigma(Q))$ on E_2. Using the change of coordinate formulas from Lemma 6.5.7, we can then express these points in Montgomery $(X : Z)$-coordinates, as desired.

6.6.7 Adaptations when only half of the torsion is available

Unlike previously, we assume that we want to compute the 2^e-isogeny F defined in Eq. (6.22) but we can only access the 2^f-torsion of E_1 and E_2 with $f \geq e/2 + 2$, as in SQIsignHD. As explained in Section 6.4.2, we decompose $F = F_2 \circ F_1$ where $F_1 : E_1^2 \times E_2^2 \longrightarrow C$ and $F_2 : C \longrightarrow E_1^2 \times E_2^2$ are respectively a 2^{e_1} and 2^{e_2}-isogeny, with $e = e_1 + e_2$ and $e_1, e_2 \leq f - 2$. We compute F_1 using 2^{e_1+2}-torsion points lying above $\ker(F_1)$ and $\widetilde{F}_2$ using 2^{e_2+2}-torsion points lying above $\ker(\widetilde{F}_2)$. Then we easily infer $\widetilde{\widetilde{F}}_2 = F_2$ from Section 6.2.3 and we are able to evaluate $F = F_2 \circ F_1$.

The computation of F_1 follows easily from what we previously explained. We follow the same steps, except that we stop the computation after e_1 2-isogenies computation when we reach C. A notable exception applies nonetheless: in order to obtain theta structures on C that are dual of each other from the computation of F_1 and $\widetilde{F}_2$, unlike previously, the 2^{e_1+2}-torsion points lying above $\ker(F_1)$ that we use are given by Lemma 6.4.6. Let (P, Q) be a basis of $E_1[2^f]$. Consider:

$$T_1 := (-[c_1 s a_1]P, -[c_1 s a_2]P, -[c_1 s]\sigma(P), 0), \quad T_2 := ([c_1 s a_2]P, -[c_1 s a_1]P, 0, -[c_1 s]\sigma(P)),$$

$$T_3 := (-[c_1 s a_1]Q, -[c_1 s a_2]Q, [c_1 r]\sigma(Q), 0), \quad T_4 := ([c_1 s a_2]Q, -[c_1 s a_1]Q, 0, [c_1 r]\sigma(Q)),$$

where $c_1 := 2^{f-e_1-2}$ and r and s are inverses of q and $a_1^2 + a_2^2$ modulo 2^f respectively. Then $T_1, \cdots, T_4$ form a maximal isotropic subgroup of $(E_1^2 \times E_2^2)[2^{e_1+2}]$ lying above $\ker(F_1)$ that we use to compute F_1.

Similarly, by Lemma 6.4.6, the points

$$T_1' := (-[c_2(a_1^2+a_2^2)]Q, 0, [c_2a_1]\sigma(Q), [c_2a_2]\sigma(Q)), \quad T_2' := (0, -[c_2(a_1^2+a_2^2)]Q, -[c_2a_2]\sigma(Q), [c_2a_1]\sigma(Q)),$$

$$T_3' := (-[c_2q]P, 0, -[c_2a_1]\sigma(P), -[c_2a_2]\sigma(P)), \quad T_4' := (0, -[c_2q]Q, [c_2a_2]\sigma(P), -[c_2a_1]\sigma(P)),$$

where $c_2 := 2^{f-e_2-2}$, form a maximal isotropic subgroup of $(E_1^2 \times E_2^2)[2^{e_2+2}]$ lying above $\ker(\widetilde{F}_2)$ that we use to compute $\widetilde{F}_2$ in order to obtain a level 2 theta structure on C which is the dual of the one induced by F_1.

The computation of $\widetilde{F}_2$ follows similar steps to F_1: a gluing in dimension 2, then a gluing in dimension 4 and the computation of the following generic 2-isogenies. Indeed, we have the following lemma.

Lemma 6.6.5. *Assume that $2|a_2$ and let $m := v_2(a_2)$ be its 2-adic valuation. Then $\widetilde{F} = g_e \circ \cdots \circ g_1$, with*

$$E_1^2 \times E_2^2 \xrightarrow{g_1} A_1'^2 \quad \cdots \quad A_{m-1}'^2 \xrightarrow{g_m} A_m'^2 \xrightarrow{g_{m+1}} B' \longrightarrow \quad \cdots,$$

a chain of 2-isogenies, where the A_i' are abelian surfaces and B' is an abelian variety of dimension 4. We have:

(i) $g_1 = \mathrm{Diag}(\psi_1, \psi_1) \circ S$, *with* $S : (R_1, S_1, R_2, S_2) \in E_1^2 \times E_2^2 \longmapsto (R_1, R_2, S_1, S_2) \in (E_1 \times E_2)^2$ *and* $\psi_1 : E_1 \times E_2 \longrightarrow A_1'$ *a gluing isogeny.*

(ii) For all $i \in [\![2\ ;\ m]\!]$, g_i *is a diagonal isogeny* $\mathrm{Diag}(\psi_i, \psi_i)$, *with* $\psi_i : A_{i-1}' \longrightarrow A_i'$.

(iii) g_{m+1} *is a gluing isogeny.*

(iv) $\ker(\psi_m \circ \cdots \circ \psi_1) = \{([a_1]P, -\sigma(P)) \mid P \in E_1[2^m]\}$.

Proof. We proceed as in the proof of Lemma 6.6.1. By Lemma 2.2.9, we obtain that:

$$\ker(\widetilde{F}) = \{([a_1]P + [a_2]Q, -[a_2]P + [a_1]Q, -\sigma(P), -\sigma(Q)) \mid P, Q \in E_1[2^e]\}.$$

Let $g_1, \cdots, g_{m+1}$ be the $m+1$ first elements of the 2-isogeny chain $\widetilde{F}$. Then, since $a_2 \equiv 0 \mod 2^m$, we have

$$\ker(g_m \circ \cdots \circ g_1) = [2^{e-m}] \ker(\widetilde{F}) = K_1 \oplus K_2,$$

where $K_1 := \{([a_1]P, 0, -\sigma(P), 0) \mid P \in E_1[2^m]\}$ and $K_2 := \{(0, [a_1]P, 0, -\sigma(P)) \mid P \in E_1[2^m]\}$. This proves the chain $g_m \circ \cdots \circ g_1$ has the desired form. □

6.6.8 Performance results

The computation and evaluation algorithms of F defined in Eq. (6.22) have been implemented in SageMath for the needs of SQIsignHD. In addition to the timings already presented in Section 3.5.2 with the fast Intel Core i5-1335U 4600MHz CPU, this computation has been tested on various parameters on the slower 2.7 GHz Intel Core i5 CPU with random supersingular elliptic curves E_1 defined over finite fields $\mathbb{F}_{p^2}$ of characteristic p between 30 and 378 bits. Primes are of the form $p = c \cdot 2^f \ell^{f'} - 1$ with $\ell = 3$ or 7, $f \geq e+2$ and c small. The isogeny $\sigma : E_1 \longrightarrow E_2$ "embedded" in $F \in \mathrm{End}(E_1^2 \times E_2^2)$ as defined in Eq. (6.22) is always a random cyclic isogeny of degree $q|\ell^{f'}$ and integers $a_1, a_2 \in \mathbb{Z}$ such that $q + a_1^2 + a_2^2 = 2^e$ are precomputed. In SQIsignHD verification, q is not smooth and may vary and a_1, a_2 are computed at runtime, however we have chosen $q|\ell^{f'}$ here to be able to verify that point images of F are correct. For every set of parameters, we compared the computation and evaluation of a 2^e-isogeny $F \in \mathrm{End}(E_1^2 \times E_2^2)$ in dimension 4 as defined in Eq. (6.22) with the computation and evaluation of a cyclic 2^e-isogeny in dimension 1 with domain E_1 (using x-only arithmetic code due to Giacomo Pope[2]). To compute F, both full torsion and half torsion cases (Section 6.6.7) were tested[3]. Computations were repeated 100 times and averaged.

[2] https://github.com/GiacomoPope/KummerIsogeny

[3] The 2^{e+2}-torsion is always available but we only used "half" of it to test the algorithmic approach introduced in Section 6.6.7. Note that we observe a small asymptotic performance gain in this case due to the quasi-linear complexity of isogeny chain computations ($e\log(e) > 2 \cdot (e/2)\log(e/2)$).

Results are displayed in Tables 6.3 and 6.4. We found that computing a 2^e-isogeny in dimension 4 is $16-18$ times more costly than in dimension 1 over a large base field $\mathbb{F}_{p^2}$, with a slight advantage to the half torsion algorithms (due to the quasilinear complexity of an isogeny chain computation). Timings for evaluation are ≈ 20 times faster in dimension 1 than in dimension 4. This suggests that our algorithmic approach is promising and can be made cryptographically relevant with a low level implementation (e.g. in C or Rust).

Table 6.3: Comparison of timings (in ms) for 2^e-isogeny computations in dimension 4 with full available torsion, half available torsion and in dimension 1 with G. Pope's code for various parameters in Python/Sagemath on a 2.7 GHz Intel Core i5 CPU.

				Dimension 4		Dimension 1
e	$\log_2(p)$	p	$\deg(\sigma)$	Full tors.	Half tors.	G. Pope
16	33	$2^{19}\cdot 3^9-1$	3^9	139	164	6
32	55	$2^{34}\cdot 3^{13}-1$	3^{13}	366	384	12
64	121	$11\cdot 2^{68}\cdot 3^{31}-1$	3^{31}	741	695	37
64	125	$5\cdot 2^{66}\cdot 3^{36}-1$	3^{35}	678	674	36
128	254	$2^{131}\cdot 3^{78}-1$	3^{75}	1519	1428	83
128	261	$5^2\cdot 2^{131}\cdot 3^{79}-1$	3^{79}	1586	1484	87
192	365	$2^{199}\cdot 3^{105}-1$	3^{105}	2447	2320	137
192	371	$239\cdot 2^{194}\cdot 3^{107}-1$	3^{107}	2459	2309	137
17	30	$3\cdot 2^{20}\cdot 7^3-1$	7^3	142	168	6
17	35	$2^{21}\cdot 7^5-1$	7^5	131	164	6
33	52	$3^2\cdot 2^{35}\cdot 7^5-1$	7^5	256	261	12
33	71	$2^{37}\cdot 7^{12}-1$	7^{11}	352	351	18
65	110	$109\cdot 2^{67}\cdot 7^{13}-1$	7^{13}	691	685	37
65	137	$5\cdot 2^{70}\cdot 7^{23}-1$	7^{23}	723	708	39
129	249	$261\cdot 2^{131}\cdot 7^{39}-1$	7^{39}	1559	1449	86
129	257	$15\cdot 2^{132}\cdot 7^{43}-1$	7^{43}	1612	1517	91
193	359	$3^2\cdot 2^{196}\cdot 7^{57}-1$	7^{57}	2499	2354	137
193	378	$97\cdot 2^{195}\cdot 7^{63}-1$	7^{63}	2488	2370	142

Table 6.4: Comparison of timings (in ms) for 2^e-isogeny evaluations in dimension 4 with full available torsion, half available torsion and in dimension 1 with G. Pope's code for various parameters in Python/Sagemath on a 2.7 GHz Intel Core i5 CPU.

				Dimension 4		Dimension 1
e	$\log_2(p)$	p	$\deg(\sigma)$	Full tors.	Half tors.	G. Pope
16	33	$2^{19} \cdot 3^9 - 1$	3^9	7.1	6.8	0.6
32	55	$2^{34} \cdot 3^{13} - 1$	3^{13}	14.2	13.9	0.8
64	121	$11 \cdot 2^{68} \cdot 3^{31} - 1$	3^{31}	27.5	26.8	1.8
64	125	$5 \cdot 2^{66} \cdot 3^{36} - 1$	3^{35}	25.9	26.1	1.8
128	254	$2^{131} \cdot 3^{78} - 1$	3^{75}	59.3	59.4	3.5
128	261	$5^2 \cdot 2^{131} \cdot 3^{79} - 1$	3^{79}	64.1	64.2	3.7
192	365	$2^{199} \cdot 3^{105} - 1$	3^{105}	107.7	109.9	5.4
192	371	$239 \cdot 2^{194} \cdot 3^{107} - 1$	3^{107}	106.6	106.9	5.4
17	30	$3 \cdot 2^{20} \cdot 7^3 - 1$	7^3	7.1	6.9	0.6
17	35	$2^{21} \cdot 7^5 - 1$	7^5	7.2	6.9	0.6
33	52	$3^2 \cdot 2^{35} \cdot 7^5 - 1$	7^5	10.0	9.7	0.8
33	71	$2^{37} \cdot 7^{12} - 1$	7^{11}	15.9	15.5	1.2
65	110	$109 \cdot 2^{67} \cdot 7^{13} - 1$	7^{13}	26.4	26.3	1.8
65	137	$5 \cdot 2^{70} \cdot 7^{23} - 1$	7^{23}	29.0	28.8	1.9
129	249	$261 \cdot 2^{131} \cdot 7^{39} - 1$	7^{39}	60.2	59.3	3.6
129	257	$15 \cdot 2^{132} \cdot 7^{43} - 1$	7^{43}	66.3	65.2	3.8
193	359	$3^2 \cdot 2^{196} \cdot 7^{57} - 1$	7^{57}	108.5	107.4	5.4
193	378	$97 \cdot 2^{195} \cdot 7^{63} - 1$	7^{63}	108.1	108.9	5.6

Bibliography

[AAA+25] M. A. Aardal, G. Adj, D. F. Aranha, A. Basso, I. A. C. Martínez, J. Chávez-Saab, M. C.-R. Santos, P. Dartois, L. De Feo, M. Duparc, J. K. Eriksen, T. B. Fouotsa, D. L. G. Filho, B. Hess, D. Kohel, A. Leroux, P. Longa, L. Maino, M. Meyer, K. Nakagawa, H. Onuki, L. Panny, S. Patranabis, C. Petit, G. Pope, K. Reijnders, D. Robert, F. R. Henríquez, S. Schaeffler, and B. Wesolowski. *SQIsign: Algorithm specifications and supporting documentation*. Tech. rep. Version 2. National Institute of Standards and Technology, 2025.

[ABDPW25] M. A. Aardal, A. Basso, L. De Feo, S. Patranabis, and B. Wesolowski. *A Complete Security Proof of SQIsign*. Cryptology ePrint Archive, Paper 2025/379. 2025. URL: https://eprint.iacr.org/2025/379.

[Ajt98] M. Ajtai. "The Shortest Vector Problem in L2 is NP-Hard for Randomized Reductions (Extended Abstract)". In: *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*. STOC '98. Dallas, Texas, USA: Association for Computing Machinery, 1998, 10–19. ISBN: 0897919629. DOI: 10.1145/276698.276705. URL: https://doi.org/10.1145/276698.276705.

[ADFMP20] N. Alamati, L. De Feo, H. Montgomery, and S. Patranabis. "Cryptographic Group Actions and Applications". In: *Advances in Cryptology – ASIACRYPT 2020*. Ed. by S. Moriai and H. Wang. Cham: Springer International Publishing, 2020, pp. 411–439. ISBN: 978-3-030-64834-3.

[ABE+24] B. Allombert, J.-F. Biasse, J. K. Eriksen, P. Kutas, C. Leonardi, A. Page, R. Scheidler, and M. T. Bagi. *PEARL-SCALLOP: Parameter Extension Applicable in Real Life for SCALLOP*. 2024.

[ACD+24] S. Arpin, J. Clements, P. Dartois, J. K. Eriksen, P. Kutas, and B. Wesolowski. "Finding orientations of supersingular elliptic curves and quaternion orders". In: *Des. Codes Cryptography* 92.11 (June 2024), 3447–3493. ISSN: 0925-1022. DOI: 10.1007/s10623-024-01435-5. URL: https://doi.org/10.1007/s10623-024-01435-5.

[Bac88] E. Bach. "How to generate factored random numbers". In: *SIAM Journal on Computing* 17.2 (1988), pp. 179–193. DOI: 10.1137/0217012.

[BBC+25] A. Basso, G. Borin, W. Castryck, M. C.-R. Santos, R. Invernizzi, A. Leroux, L. Maino, F. Vercauteren, and B. Wesolowski. *PRISM: Simple And Compact Identification and Signatures From Large Prime Degree Isogenies*. Cryptology ePrint Archive, Paper 2025/135. 2025. URL: https://eprint.iacr.org/2025/135.

[BCC+23] A. Basso, G. Codogni, D. Connolly, L. De Feo, T. B. Fouotsa, G. M. Lido, T. Morrison, L. Panny, S. Patranabis, and B. Wesolowski. "Supersingular Curves You Can Trust". In: *Advances in Cryptology – EUROCRYPT 2023*. Ed. by C. Hazay and M. Stam. Cham: Springer Nature Switzerland, 2023, pp. 405–437. ISBN: 978-3-031-30617-4.

[BDF+25] A. Basso, P. Dartois, L. D. Feo, A. Leroux, L. Maino, G. Pope, D. Robert, and B. Wesolowski. "SQIsign2D–West". In: *Advances in Cryptology – ASIACRYPT 2024*. Ed. by K.-M. Chung and Y. Sasaki. Singapore: Springer Nature Singapore, 2025, pp. 339–370. ISBN: 978-981-96-0891-1.

[BMT78] E. Berlekamp, R. McEliece, and H. van Tilborg. "On the inherent intractability of certain coding problems (Corresp.)" In: *IEEE Transactions on Information Theory* 24.3 (1978), pp. 384–386. DOI: 10.1109/TIT.1978.1055873.

[BDFLS20] D. J. Bernstein, L. De Feo, A. Leroux, and B. Smith. "Faster computation of isogenies of large prime degree". In: *Open Book Series, Proceedings of the Fourteenth Algorithmic Number Theory Symposium – ANTS XIV* 4.1 (2020), pp. 39–55.

[BKP20] W. Beullens, S. Katsumata, and F. Pintore. "Calamari and Falafl: Logarithmic (Linkable) Ring Signatures from Isogenies and Lattices". In: *Advances in Cryptology – ASIACRYPT 2020*. Ed. by S. Moriai and H. Wang. Cham: Springer International Publishing, 2020, pp. 464–492. ISBN: 978-3-030-64834-3.

[BKV19] W. Beullens, T. Kleinjung, and F. Vercauteren. "CSI-FiSh: Efficient Isogeny Based Signatures Through Class Group Computations". In: *Advances in Cryptology – ASIACRYPT 2019*. Ed. by S. D. Galbraith and S. Moriai. Cham: Springer International Publishing, 2019, pp. 227–247.

[BST+17] M. Bhargava, A. Shankar, T. Taniguchi, F. Thorne, J. Tsimerman, and Y. Zhao. "Bounds on 2-torsion in class groups of number fields and integral points on elliptic curves". In: *Journal of the American Mathematical Society* 33 (Jan. 2017). DOI: 10.1090/jams/945.

[BJS14] J.-F. Biasse, D. Jao, and A. Sankar. "A Quantum Algorithm for Computing Isogenies between Supersingular Elliptic Curves". In: *Progress in Cryptology – INDOCRYPT 2014*. Ed. by W. Meier and D. Mukhopadhyay. Cham: Springer International Publishing, 2014, pp. 428–442. ISBN: 978-3-319-13039-2.

[BL04] C. Birkenhake and H. Lange. *Complex Abelian Varieties*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004. ISBN: 978-3-662-06307-1. DOI: 10.1007/978-3-662-06307-1. URL: https://doi.org/10.1007/978-3-662-06307-1.

[BSS99] I. Blake, G. Seroussi, and N. Smart. *Elliptic Curves in Cryptography*. London Mathematical Society Lecture Note Series. Cambridge University Press, 1999.

[BKW20] D. Boneh, D. Kogan, and K. Woo. "Oblivious Pseudorandom Functions from Isogenies". In: *Advances in Cryptology – ASIACRYPT 2020*. Ed. by S. Moriai and H. Wang. Cham: Springer International Publishing, 2020, pp. 520–550. ISBN: 978-3-030-64834-3.

[BS20] X. Bonnetain and A. Schrottenloher. "Quantum Security Analysis of CSIDH". In: *Advances in Cryptology – EUROCRYPT 2020*. Ed. by A. Canteaut and Y. Ishai. Cham: Springer International Publishing, 2020, pp. 493–522. ISBN: 978-3-030-45724-2.

[BMSS08] A. Bostan, F. Morain, B. Salvy, and E. Schost. "Fast algorithms for computing isogenies between elliptic curves". In: *Mathematics of Computations* 77 (2008), pp. 1755–1778.

[BCRSC+23] G. Bruno, M. Corte-Real Santos, C. Costello, J. K. Eriksen, M. Meyer, M. Naehrig, and B. Sterner. "Cryptographic Smooth Neighbors". In: *Advances in Cryptology – ASIACRYPT 2023*. Ed. by J. Guo and R. Steinfeld. Singapore: Springer Nature Singapore, 2023, pp. 190–221. ISBN: 978-981-99-8739-9.

[BLP93] J. P. Buhler, H. W. Lenstra, and C. Pomerance. "Factoring integers with the number field sieve". In: *The development of the number field sieve*. Ed. by A. K. Lenstra and H. W. Lenstra. Berlin, Heidelberg: Springer Berlin Heidelberg, 1993, pp. 50–94. ISBN: 978-3-540-47892-8.

[CHMR25] F. Campos, A. Hellenbrand, M. Meyer, and K. Reijnders. *dCTIDH: Fast & Deterministic CTIDH*. Cryptology ePrint Archive, Paper 2025/107. 2025. URL: https://eprint.iacr.org/2025/107.

[CD20] W. Castryck and T. Decru. "CSIDH on the Surface". In: *Post-Quantum Cryptography*. Ed. by J. Ding and J.-P. Tillich. Cham: Springer International Publishing, 2020, pp. 111–129. ISBN: 978-3-030-44223-1.

[CD23] W. Castryck and T. Decru. "An Efficient Key Recovery Attack on SIDH". In: *Advances in Cryptology – EUROCRYPT 2023*. Ed. by C. Hazay and M. Stam. Cham: Springer Nature Switzerland, 2023, pp. 423–447. ISBN: 978-3-031-30589-4.

[CLMPR18] W. Castryck, T. Lange, C. Martindale, L. Panny, and J. Renes. *CSIDH: An Efficient Post-Quantum Commutative Group Action*. Cryptology ePrint Archive, Report 2018/383. https://eprint.iacr.org/2018/383. 2018.

[CLG09] D. X. Charles, K. E. Lauter, and E. Z. Goren. "Cryptographic Hash Functions from Expander Graphs". In: *Journal of Cryptology* 22.1 (2009), pp. 93–113. DOI: 10.1007/s00145-007-9002-x.

[CSCDJRH21] J. Chávez-Saab, J.-J. Chi-Domínguez, S. Jaques, and F. Rodríguez-Henríquez. "The SQALE of CSIDH: sublinear Vélu quantum-resistant isogeny action with low exponents". In: *Journal of Cryptographic Engineering* 12 (2021), pp. 349 –368. URL: https://api.semanticscholar.org/CorpusID:239668631.

[CSSD+23] J. Chávez-Saab, M. C.-R. Santos, L. De Feo, J. K. Eriksen, B. Hess, D. Kohel, A. Leroux, P. Longa, M. Meyer, L. Panny, S. Patranabis, C. Petit, F. R. Henríquez, S. Schaeffler, and B. Wesolowski. *SQIsign: Algorithm specifications and supporting documentation*. Tech. rep. Version 1. National Institute of Standards and Technology, 2023.

[CLP24] M. Chen, A. Leroux, and L. Panny. "SCALLOP-HD: Group Action from 2-Dimensional Isogenies". In: *Public-Key Cryptography – PKC 2024*. Ed. by Q. Tang and V. Teague. Cham: Springer Nature Switzerland, 2024, pp. 190–216. ISBN: 978-3-031-57725-3.

[CDPMR23] J.-J. Chi-Dominguez, A. Pizarro-Madariaga, and E. Riquelme. *Computing Isogenies of Power-Smooth Degrees Between PPAVs*. Cryptology ePrint Archive, Paper 2023/508. https://eprint.iacr.org/2023/508. 2023. URL: https://eprint.iacr.org/2023/508.

[CJS14] A. Childs, D. Jao, and V. Soukharev. "Constructing elliptic curve isogenies in quantum subexponential time". In: *Journal of Mathematical Cryptology* 8.1 (2014), pp. 1–29. DOI: doi:10.1515/jmc-2012-0016. URL: https://doi.org/10.1515/jmc-2012-0016.

[CK20] L. Colò and D. Kohel. "Orienting supersingular isogeny graphs". In: *Journal of Mathematical Cryptology* 14.1 (2020), pp. 414–437. DOI: doi:10.1515/jmc-2019-0034.

[Cor08] G. Cornacchia. "Su di un metodo per la risoluzione in numeri interi dell'equazione $\sum_{h=0}^{n} C_h x^{n-h} y^h = P$". In: *Giornale di matematiche di Battaglini* 46 (1908), pp. 33–90.

[CRSEMR24] M. Corte-Real Santos, J. K. Eriksen, M. Meyer, and K. Reijnders. "AprèsSQI: Extra Fast Verification for SQIsign Using Extension-Field Signing". In: *Advances in Cryptology – EUROCRYPT 2024*. Ed. by M. Joye and G. Leander. Cham: Springer Nature Switzerland, 2024, pp. 63–93. ISBN: 978-3-031-58716-0.

[Cos11] R. Cosset. "Applications des fonctions thêta à la cryptographie sur courbes hyperelliptiques". PhD thesis. Université Henri-Poincarré, Nancy 1, France, Nov. 2011. URL: http://docnum.univ-lorraine.fr/public/SCD_T_2011_0145_COSSET.pdf.

[CR15] R. Cosset and D. Robert. "Computing (ℓ, ℓ)-isogenies in polynomial time on Jacobian on genus 2 curves". In: *Mathematics of Computation* 84.294 (2015), pp. 1953–1975. ISSN: 00255718, 10886842. URL: http://www.jstor.org/stable/24489183 (visited on 02/19/2025).

[Cos20] C. Costello. "B-SIDH: Supersingular Isogeny Diffie-Hellman Using Twisted Torsion". In: *Advances in Cryptology – ASIACRYPT 2020*. Ed. by S. Moriai and H. Wang. Cham: Springer International Publishing, 2020, pp. 440–463. ISBN: 978-3-030-64834-3.

[CJL+17] C. Costello, D. Jao, P. Longa, M. Naehrig, J. Renes, and D. Urbanik. "Efficient Compression of SIDH Public Keys". In: *Advances in Cryptology – EUROCRYPT 2017*. Ed. by J.-S. Coron and J. B. Nielsen. Cham: Springer International Publishing, 2017, pp. 679–706. ISBN: 978-3-319-56620-7.

[CMN21] C. Costello, M. Meyer, and M. Naehrig. "Sieving for Twin Smooth Integers with Solutions to the Prouhet-Tarry-Escott Problem". In: *Advances in Cryptology – EUROCRYPT 2021*. Ed. by A. Canteaut and F.-X. Standaert. Cham: Springer International Publishing, 2021, pp. 272–301. ISBN: 978-3-030-77870-5.

[CS17] C. Costello and B. Smith. "Montgomery curves and their arithmetic: The case of large characteristic fields". In: *Journal of Cryptographic Engineering* 8 (Mar. 2017). DOI: 10.1007/s13389-017-0157-6.

[Cou06] J.-M. Couveignes. *Hard Homogeneous Spaces*. Cryptology ePrint Archive, Report 2006/291. https://eprint.iacr.org/2006/291. 2006.

[Cox13] D. A. Cox. *Primes of the form $x^2 + ny^2$*. Wiley, 2013, p. 349.

[Dar24] P. Dartois. *Fast computation of 2-isogenies in dimension 4 and cryptographic applications*. Cryptology ePrint Archive, Paper 2024/1180. 2024. URL: https://eprint.iacr.org/2024/1180.

[DEF+25] P. Dartois, J. K. Eriksen, T. B. Fouotsa, A. H. L. Merdy, R. Invernizzi, D. Robert, R. Rueger, F. Vercauteren, and B. Wesolowski. *PEGASIS: Practical Effective Class Group Action using 4-Dimensional Isogenies*. Cryptology ePrint Archive, Paper 2025/401. 2025. URL: https://eprint.iacr.org/2025/401.

[DLRW24] P. Dartois, A. Leroux, D. Robert, and B. Wesolowski. "SQIsignHD: New Dimensions in Cryptography". In: *Advances in Cryptology – EUROCRYPT 2024*. Ed. by M. Joye and G. Leander. Cham: Springer Nature Switzerland, 2024, pp. 3–32. ISBN: 978-3-031-58716-0.

[DMPR25] P. Dartois, L. Maino, G. Pope, and D. Robert. "An Algorithmic Approach to (2, 2)-Isogenies in the Theta Model and Applications to Isogeny-Based Cryptography". In: *Advances in Cryptology – ASIACRYPT 2024*. Ed. by K.-M. Chung and Y. Sasaki. Singapore: Springer Nature Singapore, 2025, pp. 304–338. ISBN: 978-981-96-0891-1.

[DH48] H. DAVENPORT and M. HALL. "ON THE EQUATION $ax^2 + by^2 + cz^2 = 0$". In: *The Quarterly Journal of Mathematics* os-19.1 (Jan. 1948), pp. 189–192. ISSN: 0033-5606. DOI: 10.1093/qmath/os-19.1.189. eprint: https://academic.oup.com/qjmath/article-pdf/os-19/1/189/4572198/os-19-1-189.pdf.

[DFKLPW20] L. De Feo, D. Kohel, A. Leroux, C. Petit, and B. Wesolowski. "SQISign: Compact Post-quantum Signatures from Quaternions and Isogenies". In: *Advances in Cryptology – ASIACRYPT 2020*. Ed. by S. Moriai and H. Wang. Cham: Springer International Publishing, 2020, pp. 64–93. ISBN: 978-3-030-64837-4.

[DLW22] L. De Feo, A. Leroux, and B. Wesolowski. *New algorithms for the Deuring correspondence: SQISign twice as fast*. Cryptology ePrint Archive, Paper 2022/234. https://eprint.iacr.org/2022/234. 2022. URL: https://eprint.iacr.org/2022/234.

[DFM20] L. De Feo and M. Meyer. "Threshold Schemes from Isogeny Assumptions". In: *Public-Key Cryptography – PKC 2020*. Ed. by A. Kiayias, M. Kohlweiss, P. Wallden, and V. Zikas. Cham: Springer International Publishing, 2020, pp. 187–212. ISBN: 978-3-030-45388-6.

[DFSGF+21] L. De Feo, C. Delpech de Saint Guilhem, T. B. Fouotsa, P. Kutas, A. Leroux, C. Petit, J. Silva, and B. Wesolowski. "Séta: Supersingular Encryption from Torsion Attacks". In: Singapore, Singapore: Springer-Verlag, 2021. ISBN: 978-3-030-92067-8. DOI: 10.1007/978-3-030-92068-5_9.

[DG16] C. Delfs and S. D. Galbraith. "Computing isogenies between supersingular elliptic curves over $\mathbb{F}_p$". In: *Designs, Codes and Cryptography* 78.2 (2016), pp. 425–440. DOI: 10.1007/s10623-014-0010-1. URL: https://doi.org/10.1007/s10623-014-0010-1.

[Deu41] M. Deuring. "Die Typen der Multiplikatorenringe elliptischer Funktionenkörper". In: *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 14.1 (1941), pp. 197–272. DOI: 10.1007/BF02940746.

[DKRS03] I. Dinur, G. Kindler, R. Raz, and S. Safra. "Approximating CVP to Within Almost-Polynomial Factors is NP-Hard". In: *Combinatorica* 23 (Apr. 2003), pp. 205–243. DOI: 10.1007/s00493-003-0019-y.

[DF25] M. Duparc and T. B. Fouotsa. "SQIPrime: A Dimension 2 Variant of SQISignHD with Non-smooth Challenge Isogenies". In: *Advances in Cryptology – ASIACRYPT 2024*. Ed. by K.-M. Chung and Y. Sasaki. Singapore: Springer Nature Singapore, 2025, pp. 396–429. ISBN: 978-981-96-0891-1.

[Dup06] R. Dupont. "Moyenne arithmético-géométrique, suites de Borchardt et applications". PhD thesis. PhD thesis, École polytechnique, 2006.

[EGM22] B. Edixhoven, G. van der Geer, and B. Moonen. *Abelian Varieties*. http://van-der-geer.nl/~gerard/AV.pdf. 2022.

[EHLMP18] K. Eisenträger, S. Hallgren, K. Lauter, T. Morrison, and C. Petit. "Supersingular Isogeny Graphs and Endomorphism Rings: Reductions and Solutions". In: *Advances in Cryptology – EUROCRYPT 2018*. Ed. by J. B. Nielsen and V. Rijmen. Cham: Springer International Publishing, 2018, pp. 329–368. ISBN: 978-3-319-78372-7.

[EHLMP20] K. Eisenträger, S. Hallgren, C. Leonardi, T. Morrison, and J. Park. "Computing endomorphism rings of supersingular elliptic curves and connections to path-finding in isogeny graphs". In: *Open Book Series, Proceedings of the Fourteenth Algorithmic Number Theory Symposium – ANTS XIV* 4.1 (2020), pp. 215–232.

[Elk98] N. D. Elkies. "Elliptic and modular curves over finite fields and related computational issues". In: *Computational perspectives on number theory (Chicago, IL, 1995)*. Vol. 7. AMS/IP Stud. Adv. Math. Amer. Math. Soc., Providence, RI, 1998, pp. 21–76. ISBN: 0-8218-0880-X. DOI: 10.1090/amsip/007/03. URL: https://doi.org/10.1090/amsip/007/03.

[EPSV24] J. Eriksen, L. Panny, J. Sotáková, and M. Veroni. "Deuring for the people: Supersingular elliptic curves with prescribed endomorphism ring in general characteristic". In: Jan. 2024, pp. 339–373. ISBN: 978-1-4704-7609-0. DOI: 10.1090/conm/796/16008.

[FLR11] J.-C. Faugère, D. Lubicz, and D. Robert. "Computing modular correspondences for abelian varieties". In: *Journal of Algebra* 343.1 (2011), pp. 248–277. ISSN: 0021-8693. DOI: https://doi.org/10.1016/j.jalgebra.2011.06.031. URL: https://www.sciencedirect.com/science/article/pii/S0021869311003802.

[FFK+23] L. D. Feo, T. B. Fouotsa, P. Kutas, A. Leroux, S.-P. Merz, L. Panny, and B. Wesolowski. "SCALLOP: Scaling the CSI-FiSh". In: *Public-Key Cryptography – PKC 2023*. Ed. by A. Boldyreva and V. Kolesnikov. Cham: Springer Nature Switzerland, 2023, pp. 345–375.

[FS87] A. Fiat and A. Shamir. "How To Prove Yourself: Practical Solutions to Identification and Signature Problems". In: *Advances in Cryptology — CRYPTO' 86*. Ed. by A. M. Odlyzko. Berlin, Heidelberg: Springer Berlin Heidelberg, 1987, pp. 186–194. ISBN: 978-3-540-47721-1.

[FIKMN25] J. Fuselier, A. Iezzi, M. Kozek, T. Morrison, and C. Namoijam. "Computing supersingular endomorphism rings using inseparable endomorphisms". In: *Journal of Algebra* 668 (2025), pp. 145–189. ISSN: 0021-8693. DOI: https://doi.org/10.1016/j.jalgebra.2025.01.012. URL: https://www.sciencedirect.com/science/article/pii/S0021869325000353.

[GPS20] S. D. Galbraith, C. Petit, and J. Silva. "Identification protocols and signature schemes based on supersingular isogeny problems". In: *Journal of Cryptology* 33.1 (2020), pp. 130–175.

[Gau07] P. Gaudry. "Fast genus 2 arithmetic based on Theta functions". In: *J. Math. Cryptol.* 1.3 (2007), pp. 243–265. DOI: 10.1515/JMC.2007.012.

[GW10] U. Görtz and T. Wedhorn. *Algebraic Geometry I Schemes with examples and exercises*. Advanced lectures in mathematics. Wiesbaden, Germany: Vieweg Teubner, 2010, pp. vii+615. ISBN: 3-8348-0676-5.

[Gro96] L. K. Grover. "A Fast Quantum Mechanical Algorithm for Database Search". In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. STOC '96. Philadelphia, Pennsylvania, USA: Association for Computing Machinery, 1996, 212–219. ISBN: 0897917855. DOI: 10.1145/237814.237866. URL: https://doi.org/10.1145/237814.237866.

[HW75] G. H. Hardy and E. M. Wright. *An Introduction to the Theory of Numbers*. Sixth. Oxford, 1975. DOI: 10.1093/oso/9780199219858.001.0001.

[Har77] R. Hartshorne. *Algebraic geometry*. Graduate Texts in Mathematics, No. 52. New York: Springer-Verlag, 1977, pp. xvi+496. ISBN: 0-387-90244-9.

[Ibu82] T. Ibukiyama. "On maximal orders of division quaternion algebras over the rational number field with certain optimal embeddings". In: *Nagoya Mathematical Journal* 88 (1982), 181–195. DOI: 10.1017/S002776300002016X.

[Igu72] J.-I. Igusa. *Theta functions*. Die Grundlehren der mathematischen Wissenschaften, Band 194. New York: Springer-Verlag, 1972, pp. x+232.

[JAC+20] D. Jao, R. Azarderakhsh, M. Campagna, C. Costello, L. De Feo, B. Hess, A. Jalali, B. Koziel, B. LaMacchia, P. Longa, M. Naehrig, J. Renes, V. Soukharev, D. Urbanik, G. Pereira, K. Karabina, and A. Hutchinson. *SIKE*. Tech. rep. available at https://csrc.nist.gov/projects/post-quantum-cryptography/round-3-submissions. National Institute of Standards and Technology, 2020.

[JDF11] D. Jao and L. De Feo. "Towards Quantum-Resistant Cryptosystems from Supersingular Elliptic Curve Isogenies". In: *Post-Quantum Cryptography*. Ed. by B.-Y. Yang. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 19–34. ISBN: 978-3-642-25405-5.

[Kan97] E. Kani. "The number of curves of genus two with elliptic differentials". In: *Journal für die reine und angewandte Mathematik* 1997.485 (1997), pp. 93–122. DOI: doi:10.1515/crll.1997.485.93. URL: https://doi.org/10.1515/crll.1997.485.93.

[KR09] T. Katsuyuki and Y. Reo. "An algorithm for computing a sequence of Richelot isogenies". In: *Bulletin of the Korean Mathematical Society* 46.4 (July 2009), pp. 789–802.

[KV10] M. Kirschmer and J. Voight. "Algorithmic Enumeration of Ideal Classes for Quaternion Orders". In: *SIAM Journal on Computing* 39.5 (2010), pp. 1714–1747. DOI: 10.1137/080734467.

[KLPT14] D. Kohel, K. Lauter, C. Petit, and J.-P. Tignol. *On the quaternion ℓ-isogeny path problem*. 2014. arXiv: 1406.0981 [math.NT].

[Kun22] S. Kunzweiler. *Efficient Computation of $(2^n, 2^n)$-Isogenies*. Cryptology ePrint Archive, Paper 2022/990. 2022. URL: https://eprint.iacr.org/2022/990.

[KMM+24] S. Kunzweiler, L. Maino, T. Moriya, C. Petit, G. Pope, D. Robert, M. Stopar, and Y. B. Ti. *Radical 2-isogenies and cryptographic hash functions in dimensions 1, 2 and 3*. Cryptology ePrint Archive, Paper 2024/1732. 2024. URL: https://eprint.iacr.org/2024/1732.

[Kup05] G. Kuperberg. "A Subexponential-Time Quantum Algorithm for the Dihedral Hidden Subgroup Problem". In: *SIAM Journal on Computing* 35.1 (2005), pp. 170–188. DOI: 10.1137/S0097539703436345. eprint: https://doi.org/10.1137/S0097539703436345. URL: https://doi.org/10.1137/S0097539703436345.

[Lag70] J. L. de Lagrange. "Démonstration d'un théoreme d'arithmétique". In: *Nouveau Mémoire de l'Académie Royale des Sciences de Berlin* (1770), pp. 123–133.

[LGSG21] Y.-F. Lai, S. D. Galbraith, and C. Delpech de Saint Guilhem. "Compact, Efficient and UC-Secure Isogeny-Based Oblivious Transfer". In: *Advances in Cryptology – EUROCRYPT 2021*. Ed. by A. Canteaut and F.-X. Standaert. Cham: Springer International Publishing, 2021, pp. 213–241. ISBN: 978-3-030-77870-5.

[Lan04] S. Lang. *Algèbre*. Springer, 2004, p. 934.

[LLL82] A. K. Lenstra, H. W. Lenstra, and L. Lovász. "Factoring polynomials with rational coefficients". In: *Mathematische Annalen* 261.4 (Jan. 1982), pp. 515–534. DOI: 10.1007/BF01457454.

[Ler22] A. Leroux. *Quaternion algebras and isogeny-based cryptography*. http://www.lix.polytechnique.fr/Labo/Antonin.LEROUX/manuscrit_these.pdf. 2022.

[LWZ24] J. Lin, S. Wang, and C.-A. Zhao. "A note on (2,2)-isogenies via theta coordinates". In: *Finite Fields and Their Applications* 99 (2024), p. 102496. ISSN: 1071-5797. DOI: https://doi.org/10.1016/j.ffa.2024.102496. URL: https://www.sciencedirect.com/science/article/pii/S1071579724001357.

[LR12] D. Lubicz and D. Robert. "Computing isogenies between abelian varieties". In: *Compositio Mathematica* 148.5 (Sept. 2012), pp. 1483–1515. DOI: 10.1112/S0010437X12000243.

[LR15] D. Lubicz and D. Robert. "Computing separable isogenies in quasi-optimal time". In: *LMS Journal of Computation and Mathematics* 18.1 (2015), 198–216. DOI: 10.1112/S146115701400045X.

[LR22] D. Lubicz and D. Robert. "Fast change of level and applications to isogenies". In: *Research in Number Theory* 9.7 (2022). DOI: 10.1007/s40993-022-00407-9.

[MMPPW23] L. Maino, C. Martindale, L. Panny, G. Pope, and B. Wesolowski. "A Direct Key Recovery Attack on SIDH". In: *Advances in Cryptology – EUROCRYPT 2023*. Ed. by C. Hazay and M. Stam. Cham: Springer Nature Switzerland, 2023, pp. 448–471. ISBN: 978-3-031-30589-4.

[MW25] A. H. L. Merdy and B. Wesolowski. *Unconditional foundations for supersingular isogeny-based cryptography*. Cryptology ePrint Archive, Paper 2025/271. 2025. URL: https://eprint.iacr.org/2025/271.

[MG02] D. Micciancio and S. Goldwasser. *Complexity of Lattice Problems: A Cryptographic Perspective*. Springer Science+Business Media, 2002, p. 220.

[Mil04] V. Miller. "The Weil Pairing, and Its Efficient Calculation". In: *The Journal of Cryptology* 17.4 (2004), pp. 235–261. DOI: 10.1007/s00145-004-0315-8. URL: https://doi.org/10.1007/s00145-004-0315-8.

[Mil86] J. S. Milne. "Abelian Varieties". In: *Arithmetic Geometry*. Ed. by G. Cornell and J. H. Silverman. New York, NY: Springer New York, 1986, pp. 103–150. ISBN: 978-1-4613-8655-1. DOI: 10.1007/978-1-4613-8655-1_5. URL: https://doi.org/10.1007/978-1-4613-8655-1_5.

[Mil08] J. S. Milne. *Abelian Varieties (v2.00)*. Available at www.jmilne.org/math/. 2008.

[Mum66] D. Mumford. "On the equations defining abelian varieties 1". In: *Inventiones mathematicae* 1.4 (1966), pp. 287–354. DOI: 10.1007/BF01389737. URL: https://doi.org/10.1007/BF01389737.

[Mum74] D. Mumford. *Abelian varieties*. Second Edition. Tata Institute of fundamental research studies in mathematics. London: Oxford University Press, 1974, pp. x+279.

[Mum84] D. Mumford. *Tata lectures on theta II*. With the collaboration of C. Musili, M. Nori, E. Previato, M. Stillman and H. Umemura. Boston: Birkbäuser, 1984, pp. xiv+272. ISBN: 0-8176-4569-1.

[NO24] K. Nakagawa and H. Onuki. "QFESTA: Efficient Algorithms and Parameters for FESTA Using Quaternion Algebras". In: *Advances in Cryptology – CRYPTO 2024*. Ed. by L. Reyzin and D. Stebila. Cham: Springer Nature Switzerland, 2024, pp. 75–106. ISBN: 978-3-031-68388-6.

[NOC+25] K. Nakagawa, H. Onuki, W. Castryck, M. Chen, R. Invernizzi, G. Lorenzon, and F. Vercauteren. "SQIsign2D-East: A New Signature Scheme Using 2-Dimensional Isogenies". In: *Advances in Cryptology – ASIACRYPT 2024*. Ed. by K.-M. Chung and Y. Sasaki. Singapore: Springer Nature Singapore, 2025, pp. 272–303. ISBN: 978-981-96-0891-1.

[NS09] P. Q. Nguyen and D. Stehlé. "Low-dimensional lattice basis reduction revisited". In: *ACM Trans. Algorithms* 5.4 (Nov. 2009). ISSN: 1549-6325. DOI: 10.1145/1597036.1597050. URL: https://doi.org/10.1145/1597036.1597050.

[Onu21] H. Onuki. "On oriented supersingular elliptic curves". In: *Finite Fields and Their Applications* 69 (2021), p. 101777. ISSN: 1071-5797. DOI: https://doi.org/10.1016/j.ffa.2020.101777.

[ON25] H. Onuki and K. Nakagawa. "Ideal-to-Isogeny Algorithm Using 2-Dimensional Isogenies and Its Application to SQIsign". In: *Advances in Cryptology – ASIACRYPT 2024*. Ed. by K.-M. Chung and Y. Sasaki. Singapore: Springer Nature Singapore, 2025, pp. 243–271. ISBN: 978-981-96-0891-1.

[OP22] R. Oudompheng and G. Pope. *A Note on Reimplementing the Castryck-Decru Attack and Lessons Learned for SageMath*. Cryptology ePrint Archive, Paper 2022/1283. 2022. URL: https://eprint.iacr.org/2022/1283.

[PR23] A. Page and D. Robert. *Introducing Clapoti(s): Evaluating the isogeny class group action in polynomial time*. Cryptology ePrint Archive, Paper 2023/1766. 2023. URL: https://eprint.iacr.org/2023/1766.

[PW24] A. Page and B. Wesolowski. "The Supersingular Endomorphism Ring and One Endomorphism Problems are Equivalent". In: *Advances in Cryptology – EUROCRYPT 2024*. Ed. by M. Joye and G. Leander. Cham: Springer Nature Switzerland, 2024, pp. 388–417. ISBN: 978-3-031-58751-1.

[PPS24] L. Panny, C. Petit, and M. Stopar. "KLaPoTi: An asymptotically efficient isogeny group action from 2-dimensional isogenies". In: *IACR Cryptol. ePrint Arch.* (2024), p. 1844. URL: https://eprint.iacr.org/2024/1844.

[PP02] G. Pareschi and M. Popa. "Regularity on abelian varieties I". In: *Journal of the American Mathematical Society* 16 (Nov. 2002), pp. 285–302. DOI: 10.1090/S0894-0347-02-00414-91.

[Pei20] C. Peikert. "He Gives C-Sieves on the CSIDH". In: *Advances in Cryptology – EUROCRYPT 2020*. Ed. by A. Canteaut and Y. Ishai. Cham: Springer International Publishing, 2020, pp. 463–492. ISBN: 978-3-030-45724-2.

[Piz90] A. K. Pizer. "Ramanujan graphs and Hecke operators". In: *Bulletin of the American Mathematical Society* 23 (1990), pp. 127–137. URL: https://api.semanticscholar.org/CorpusID:19789543.

[PT18] P. Pollack and E. Treviño. "Finding the Four Squares in Lagrange's Theorem". In: *Integers* 18A (2018), A15.

[Rob10] D. Robert. "Theta functions and cryptographic applications". PhD thesis. Université Henri-Poincarré, Nancy 1, France, July 2010. URL: http://www.normalesup.org/~robert/pro/publications/academic/phd.pdf.

[Rob23] D. Robert. "Breaking SIDH in Polynomial Time". In: *Advances in Cryptology – EUROCRYPT 2023*. Ed. by C. Hazay and M. Stam. Cham: Springer Nature Switzerland, 2023, pp. 472–503. ISBN: 978-3-031-30589-4.

[Rob24] D. Robert. *A note on optimising 2^n-isogenies in higher dimension*. Cryptology ePrint Archive, Paper 2024/406. https://eprint.iacr.org/2024/406. 2024.

[RS24] D. Robert and N. Sarkis. *Halving differential additions on Kummer lines*. Cryptology ePrint Archive, Paper 2024/1582. 2024. URL: https://eprint.iacr.org/2024/1582.

[RS06] A. Rostovtsev and A. Stolbunov. *Public-Key Cryptosystem Based On Isogenies*. Cryptology ePrint Archive, Report 2006/145. https://eprint.iacr.org/2006/145. 2006.

[RT22] J. Rouse and K. Thompson. *Quaternary quadratic forms with prime discriminant*. 2022. arXiv: 2206.00412 [math.NT].

[SGOPS20] C. D. de Saint Guilhem, E. Orsini, C. Petit, and N. P. Smart. "Semi-commutative Masking: A Framework for Isogeny-Based Protocols, with an Application to Fully Secure Two-Round Isogeny-Based OT". In: *Cryptology and Network Security*. Ed. by S. Krenn, H. Shulman, and S. Vaudenay. Cham: Springer International Publishing, 2020, pp. 235–258. ISBN: 978-3-030-65411-5.

[SEMRH24] M. C.-R. Santos, J. K. Eriksen, M. Meyer, and F. Rodríguez-Henríquez. "Finding Practical Parameters for Isogeny-based Cryptography". In: *IACR Communications in Cryptology* 1.3 (Oct. 7, 2024). ISSN: 3006-5496. DOI: 10.62056/ayojbhey6b.

[Sch87] C. P. Schnorr. "A Hierarchy of Polynomial Time Lattice Basis Reduction Algorithms". In: *Theoretical Computer Science* 53 (1987), 201–224. ISSN: 1088-6842. DOI: 10.1016/0304-3975(87)90064-8. URL: https://doi.org/10.1016/0304-3975(87)90064-8.

[Sha13] I. R. Shafarevich. *Basic Algebraic Geometry I*. Third Edition. Springer-Verlag Berlin, 2013, p. 310. ISBN: 978-3-642-37956-7. DOI: https://doi.org/10.1007/978-3-642-37956-7.

[Sha71] D. Shanks. "Class number, a theory of factorization, and genera". In: *Proceedings of Symposia in Pure Mathematics*. Vol. 20. 1971. ISBN: 978-0-8218-9306-7. URL: https://doi.org/10.1090/pspum/020.

[Sha86] S. S. Shatz. "Group Schemes, Formal Groups, and p-Divisible Groups". In: *Arithmetic Geometry*. Ed. by G. Cornell and J. H. Silverman. New York, NY: Springer New York, 1986, pp. 29–78. ISBN: 978-1-4613-8655-1. DOI: 10.1007/978-1-4613-8655-1_5. URL: https://doi.org/10.1007/978-1-4613-8655-1_5.

[Sho97] P. W. Shor. "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer". In: *SIAM Journal on Computing* 26.5 (1997), 1484–1509. ISSN: 1095-7111. DOI: 10.1137/s0097539795293172. URL: http://dx.doi.org/10.1137/S0097539795293172.

[Sil09] J. H. Silverman. *The Arithmetic of Elliptic Curves*. Springer, 2009, p. 522.

[Sim05] D. Simon. "Solving quadratic equations using reduced unimodular quadratic forms". In: *Math. Comput.* 74 (July 2005), pp. 1531–1543. DOI: 10.1090/S0025-5718-05-01729-1.

[Ste23] B. Sterner. *Towards Optimally Small Smoothness Bounds for Cryptographic-Sized Twin Smooth Integers and their Isogeny-based Applications*. Cryptology ePrint Archive, Paper 2023/1576. 2023. URL: https://eprint.iacr.org/2023/1576.

[Syl82] J. J. Sylvester. "On Subvariants, i.e. Semi-Invariants to Binary Quantics of an Unlimited Order". In: *American Journal of Mathematics* 5.1 (1882), pp. 79–136. ISSN: 00029327, 10806377. URL: http://www.jstor.org/stable/2369536 (visited on 03/13/2025).

[Tat66] J. Tate. "Endomorphisms of Abelian Varieties over Finite Fields". In: *Inventiones mathematicae* 2 (1966), pp. 134–144.

[The24] The Stacks project authors. *The Stacks project*. https://stacks.math.columbia.edu. 2024.

[VV15] D. Venturi and A. Villani. *Zero-Knowledge Proofs and Applications*. 2015. URL: http://danieleventuri.altervista.org/files/zero-knowledge.pdf.

[Voi21] J. Voight. *Quaternion Algebras*. Jan. 2021, p. 877. ISBN: 978-3-030-56692-0. DOI: 10.1007/978-3-030-56694-4.

[Vé71] J. Vélu. “Isogénies entre courbes elliptiques”. In: *Comptes-rendus de l’Académie des Sciences* 273 (1971). Available at https://gallica.bnf.fr, pp. 238–241.

[Wae56] B. L. V. der Waerden. “Die Reduktionstheorie der positiven quadratischen Formen”. In: *Acta Mathematica* 96 (1956), pp. 265–309. DOI: 10.1007/BF02392364.

[Wat69] W. C. Waterhouse. “Abelian varieties over finite fields”. eng. In: *Annales scientifiques de l’École Normale Supérieure* 2.4 (1969). http://eudml.org/doc/81852, pp. 521–560.

[Wes22] B. Wesolowski. “The supersingular isogeny path and endomorphism ring problems are equivalent”. In: *FOCS 2021 - 62nd Annual IEEE Symposium on Foundations of Computer Science*. Denver, Colorado, United States, Feb. 2022. URL: https://hal.archives-ouvertes.fr/hal-03340899.

[Wes24] B. Wesolowski. *Random walks in number-theoretic cryptography*. École Normale Supérieure de Lyon, Aug. 2024. URL: https://bweso.com/hdr.pdf.