MIKE for Beginners

Michael Meyer[1] and Krijn Reijnders[2]

[1] University of Regensburg, Germany
[2] COSIC, KU Leuven, Belgium

Abstract. This work explains MIKE for beginners, i.e., those unfamiliar with isogenies. MIKE is both a fast and compact post-quantum NIKE and a beautiful scheme. However, this beauty is only visible once you are familiar with many more advanced concepts in isogeny-based cryptography. The beauty of MIKE then reveals itself in several different layers, each taking a step up in complexity. This work assumes no knowledge of isogenies, and explains three levels of understanding for MIKE, each level building on the previous one, preparing the reader for the full MIKE paper. Each level assumes a different level of knowledge; Level 1 assumes only familiarity with basic cryptography, Level 2 assumes a bit of familiarity with elliptic curves, Level 3 assumes a bit of familiarity with isogenies, which is provided in Level 2, and Level 3 ½ bridges the gap from Level 3 to the full MIKE paper [BDD+26].

Aim of this Document

MIKE [Rob24; BDD+26] is an exceptionally well-performing isogeny-based non-interactive key exchange (NIKE) and therefore an interesting scheme to many cryptographers beyond the discipline of isogeny-based cryptography. However, the literature in isogeny-based cryptography can be impenetrable, and in particular both papers describing MIKE require a serious investment in time as well as effort to understand. With this work, we hope to provide a simple introduction to MIKE which should be accessible even for those with no knowledge of isogenies.

1. Level 1 is aimed at those familiar with basic cryptographic knowledge, and aims to explain MIKE as the next evolution in Diffie–Hellman key exchange with decreasing amounts of structure. After Level 1, you should feel familiar with monoid actions.
2. Level 2 is aimed at those familiar with elliptic curves and cryptography, and aims to explain MIKE from the computational point of view: what is necessary to understand and implement the core computational protocol itself. After Level 2, you should feel somewhat familiar with isogenies.
3. Level 3 is aimed at those with a bit of knowledge of isogenies, and aims to explain the core underlying module action in MIKE, to show how the computational protocol is derived from the theory. After Level 3, you should feel somewhat familiar with module actions on abelian varieties, and may be ready to start reading the full paper [BDD+26].
4. Level 3½ aims to close the gap to the full paper [BDD+26] by giving a somewhat down-to-earth introduction to several of its core concepts. It assumes a bit more knowledge of the reader, points to the relevant sources to learn those concepts and the relevant sections of [BDD+26] to get a deeper understanding of many of the core concepts in the full module action framework.

E-mail: michael@random-oracles.org (Michael Meyer), krijn@cs.ru.nl (Krijn Reijnders)

This work is licensed under a "CC BY 4.0" license.
Date of this document: 2026-11-20.

Notation and Terminology. We always assume we are working with primes $p \equiv 3 \bmod 4$, and most of the times even $p \equiv 7 \bmod 8$, as in MIKE [BDD+26]. This allows us to think of $\mathbb{F}_{p^2}$ as $\mathbb{F}_p(i)$, with $i^2 = -1$. We will often use the commutative ring $R = \mathbb{Z}[\sqrt{-p}]$ of elements of the form $a + b\sqrt{-p}$, with $a, b \in \mathbb{Z}$, and similarly $R' = \mathbb{Z}[\frac{1+\sqrt{-p}}{2}]$ with elements of the form $a + b\frac{1+\sqrt{-p}}{2}$. We are aware that we write R' where [BDD+26] writes R. However, throughout this whole paper we have a strong identification between the symbols (E_0, R, W) and (E'_0, R', W'), justifying the change. In later levels, we may use $\omega = \frac{1+\sqrt{-p}}{2}$ to simplify the writing.

We will explain and work with abelian varieties of a certain dimension. In general, a variety of dimension 1 is called a *curve*, of dimension 2 a *surface*, of dimension 3 a *threefold*, and of dimension 4 a *fourfold*. We will sometimes write "abelian variety of dimension 4" and sometimes "abelian fourfold", but they are the same thing. We strictly use E or E_A to refer to elliptic curves, that is, abelian varieties of dimension 1. In dimension 2, almost all the surfaces we will see are products of elliptic curves or Jacobians of hyperelliptic curves of genus 2. Some computational steps require a bit more understanding of these Jacobians, but a reader unfamiliar with Jacobians can assume these steps as a black box. We say that a curve E or variety A is *defined over* $\mathbb{F}_p$, resp. $\mathbb{F}_{p^2}$, when the defining equation(s) are defined over $\mathbb{F}_p$, resp. $\mathbb{F}_{p^2}$.

In isogeny-based cryptography, abelian varieties are almost always considered together with a (principal) polarisation. To avoid having to introduce such technical machinery, in Levels 1 to 3 we implicitly assume that all our varieties are principally polarised. This is justified by the fact that all such varieties are obtained via polarised isogenies starting from products of elliptic curves, for which we assume the product polarisation.

We will explain and work with isogenies between abelian varieties of a certain dimension. In general, we use lowercase letters such as φ or ψ to denote (one-dimensional) isogenies between elliptic curves, and uppercase Φ or Ψ to denote higher-dimensional isogenies. A little bit of experience with two-dimensional isogenies, for which we recommend the explanations by Castryck, Decru, and Smith [CDS20] and Flynn and Ti [FT19], will be very helpful in computational examples in Level 3, but one can also safely assume these steps as a black box and still get a solid understanding of MIKE.

We apologise for overloading in notation, which was sadly unavoidable in some places. For example, x_P is used as the x-coordinate of a point P, but $x_A \in \mathbb{Z}_{2^e}$ is used as the secret value in key generation. Similarly, A is used both for the Montgomery coefficient of curves, as a shorthand for objects K_A or M_A associated to Alice, and sometimes for an abelian variety.

Examples and Code. We have tried to be generous with examples in this work. Most of these examples also have associated code (in SageMath) to demonstrate such situations in practice. As we are not Lorenz Panny, the use of LLMs has been very helpful in writing up this code nicely, although we stress that no LLMs were used in a single letter that is written in this text.

The associated Sage code will soon be available.

Acknowledgements. We thank the MIKE team and Maria Corte-Real Santos for their helpful feedback and improvements to an earlier version of this work. This work is supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement ISOCRYPT No. 101020788), by the Research Council KU Leuven grant C14/24/099, by CyberSecurity Research Flanders with reference number VOEWICS02, and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 577658121 (IMPACT).

AI Disclosure. Beyond the usage in writing up examples, an LLM was used as to check for correctness and as an editorial aid: given a final draft of the paper, it was used to identify inconsistencies, correct typographical errors, and improve grammar and presentation. This document is fully written by the authors and we take full responsibility for the accuracy, integrity, and final presentation of this work.

Level 1: MIKE as a monoid action NIKE

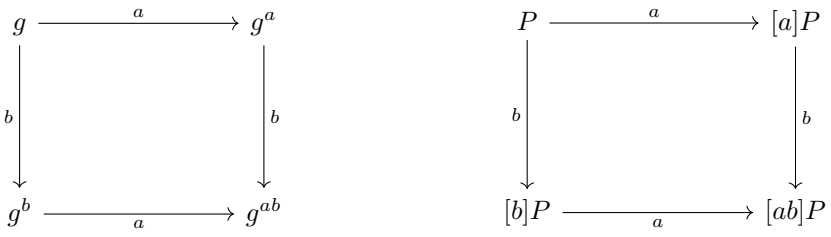
The first description of MIKE that we give is as a commutative *monoid action*, where a monoid is essentially a group, but without the requirement of the existence of an inverse for each element. To reach that description, we look at the evolution of Diffie–Hellman throughout recent decades, by peeling away more and more structure of the underlying objects. This gives a simple explanation for MIKE, which provides a nice intuition on why we believe it to be quantum secure. We start with pre-quantum Diffie–Hellman variants, before we discuss post-quantum variants.

1.1 Pre-quantum non-interactive key exchange

Diffie–Hellman key exchange. The origin of our path through the evolution of NIKEs is Diffie and Hellman's key exchange scheme—**the** classical NIKE every cryptography student learns about in their first course [DH76]. The Diffie–Hellman key exchange is remarkably simple: we pick a prime p, consider the full multiplicative group $\mathbb{F}_p^\times$, and fix some generator $g \in \mathbb{F}_p^\times$. Then Alice and Bob choose secret keys $a, b \in \{1, 2, \ldots, p-1\}$, and compute their public keys g^a and g^b, respectively. The shared secret is $(g^a)^b = g^{ab} = (g^b)^a$. We depict this in Fig. 1a.

The main security assumption behind Diffie–Hellman is the discrete logarithm problem (DLOG). Given g and g^a, it asks to recover a. The best classical attacks against Diffie–Hellman's DLOG problem are based on index calculus [Adl79]. This type of attack exploits the representation of the underlying field elements, and results in a subexponential classical attack. This is not devastating, but requires much larger parameters (that is, a large enough prime p), to withstand such attacks, which harms performance and key sizes.

Looking at this from an abstract point of view, the functionality of the key exchange only requires a commutative group, which Diffie–Hellman instantiates with the multiplicative group $\mathbb{F}_p^\times$. At the same time, the underlying rich structure of $\mathbb{F}_p$—that the key exchange does not use in its entirety—is available to the attacker, resulting in this attack that outperforms generic DLOG attacks.



(a) Classical Diffie–Hellman, using finite fields.

(b) ECDH, using cyclic groups.

Figure 1: Pre-quantum non-interactive key exchange.

Elliptic Curve Diffie–Hellman. This observation leads us to the next step in the evolution of NIKEs. Roughly a decade after the publication of Diffie–Hellman, Miller [Mil85] and Koblitz [Kob87] independently proposed cryptography based on elliptic curves. Given a prime p, an elliptic curve over $\mathbb{F}_p$ can be of the form

$$E : y^2 = x^3 + Ax^2 + x$$

with a fixed $A \in \mathbb{F}_p$. Then a point P on E is a tuple $(x, y) \in \mathbb{F}_p \times \mathbb{F}_p$ that satisfies the equation above. The set of points on an elliptic curve, together with the so-called point at

infinity $\mathcal{O}$ form a commutative group with an addition operation. Cryptographers love elliptic curves because of this abelian group structure! With addition, we can also define scalar multiplication by an integer $n \in \mathbb{Z}$, by adding together n copies of a given point P, written as $[n]P$. We call the smallest positive n such that $[n]P = \mathcal{O}$ the order of P.

Based on this, we can immediately describe an elliptic-curve analogue of Diffie–Hellman (ECDH): given a curve E and a base point P, we work with the cyclic (sub)group $G = \langle P \rangle$ of points generated by P, and denote its order by r. Alice and Bob choose positive integers $a, b < r$ as secret keys and compute $[a]P$ and $[b]P$ as their public keys, respectively. The shared secret is $[a]([b]P) = [ab]P = [b]([a]P)$. We depict this in Fig. 1b.

So far it looks like this makes the Diffie–Hellman scheme more complex by replacing the simple arithmetic of exponentiations in $\mathbb{F}_p^\times$ by the more complicated-looking elliptic curve arithmetic. But we gain more than we lose: by stripping off the rich underlying structure of $\mathbb{F}_p$, attackers lose an avenue for attacks. Indeed, for well-chosen curves E the best known classical attacks are generic exponential-time attacks that can't exploit the exact nature and features of the chosen group. In contrast, subexponential attacks like index calculus no longer work! In some sense, this means that ECDH is as secure as group-based Diffie–Hellman schemes can get, leading to much smaller parameters and keys for ECDH compared to Diffie–Hellman.

In a nutshell, by shaking off some structure that we don't need for the functionality of the scheme, we avoid certain attacks and obtain a more secure scheme—a strategy we will continue in order to reach a quantum-resistant NIKE.

⚠ Warning

We will not delve into the subtle discussion of what a *well-chosen* elliptic curve exactly is in the ECDH context. There are many pitfalls, and a bad choice of curve can render ECDH insecure. Our description above glosses over this and assumes the usage of established and standardised curves and parameters.

1.2 Post-quantum non-interactive key exchange

Isogeny-based group actions. We've already seen that with ECDH we basically mastered the art of defining a classically secure NIKE, but unfortunately this is only true for classical security. Problematically, Shor's quantum algorithm [Sho94] can solve the DLOG problem in groups in polynomial time on a sufficiently powerful quantum computer. Hence, both Diffie–Hellman and ECDH become completely insecure as soon as such quantum computers are available.[1] A potential way to resolve this predicament is—as you might have guessed—to get rid of some additional underlying structure, which leads us to the next step in the NIKE evolution.

As first proposed by Brassard and Yung [BY90], we can replace the involved commutative groups and their operations by a commutative *group action*, and still maintain the NIKE functionality. Abstractly speaking, this involves a commutative group $\mathcal{G}$, and a set $\mathcal{E}$ that is in general not equipped with a group structure. The group action of $\mathcal{G}$ on $\mathcal{E}$ then is a map $* : \mathcal{G} \times \mathcal{E} \to \mathcal{E}$.

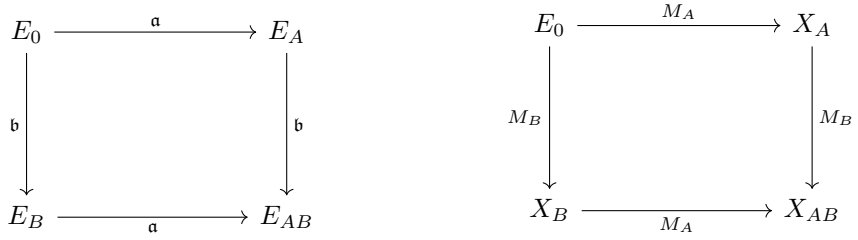
An isogeny-based instantiation and formalisation of this approach is due to Couveignes [Cou06], and independently to Rostovtsev and Stolbunov [RS06], which eventually led to the development of the CSIDH scheme [CLM+18] based on isogenies. In CSIDH, we still work with elliptic curves, but in a very different sense than in ECDH. The set $\mathcal{E}$ consists of a certain set of elliptic curves, while the group $\mathcal{G}$ is the so-called ideal class group. Its exact definition is not important for our high-level view on the CSIDH NIKE; it's enough to describe how this group acts on $\mathcal{E}$ via isogenies. At this point it suffices to

[1] We refrain from placing a bet on when the first successful quantum attack on ECDH will be published.

defer the details to Level 2 and think of an isogeny φ as some sort of *nice* map between two elliptic curves E and E' that preserves the group structure between these curves i.e., writing $\varphi : E \to E'$, we have $\varphi(P+Q) = \varphi(P) + \varphi(Q)$. We will omit all the magic tricks that explain how CSIDH works, and for now simply work with the viewpoint that the action of an ideal $\mathfrak{a} \in \mathcal{G}$ (or more precisely an ideal class) on some curve $E \in \mathcal{E}$ results in computing the codomain E' of the isogeny $\varphi_{\mathfrak{a}}$ corresponding to $\mathfrak{a}$.

Due to the commutativity of $\mathcal{G}$, we can again define a straightforward analogue of Diffie–Hellman: given a fixed starting curve $E_0 \in \mathcal{E}$, Alice and Bob choose secret ideal classes $\mathfrak{a}, \mathfrak{b} \in \mathcal{G}$, and compute their public keys $E_A = \mathfrak{a} * E_0$ and $E_B = \mathfrak{b} * E_0$ via the corresponding isogenies φ_A and φ_B, respectively. The shared secret is $\mathfrak{a} * E_B = (\mathfrak{a} \cdot \mathfrak{b}) * E_0 = \mathfrak{b} * E_A$, again computed via the corresponding isogenies φ'_A and φ'_B. We depict this in Fig. 2a.

Again it seems like we make earlier approaches like Diffie–Hellman and ECDH much more complex. But again we gain a lot: Shor's algorithm does not apply to such group action NIKEs, since it essentially requires a group-style operation on elements of $\mathcal{E}$, which is not available if $\mathcal{E}$ is not a group! So does that mean that CSIDH is the best post-quantum NIKE we can hope for? Well, unfortunately there is another quantum algorithm due to Kuperberg [Kup05] that applies to group actions. The good news is that it does not have polynomial-time complexity like Shor's algorithm, but subexponential complexity. So we are in a similar situation as for Diffie–Hellman: the attack is not devastating, but reaching a large enough security level requires scaling up parameters compared to when only considering generic attacks. But fortunately, we have one more step in the NIKE evolution up our sleeve.



(a) CSIDH NIKE, using group actions.

(b) MIKE, using monoid actions.

Figure 2: Post-quantum non-interactive key exchange.

Isogeny-based module actions. The strategy should be clear by now: Kuperberg's algorithm requires a *group* action, so let's get rid of the group structure! Looking at Fig. 2a and the resulting NIKE, we see that we never actually need inverses of group elements, but only the commutative action itself. So from an abstract point of view, it suffices to work with just a commutative *monoid* action, which exactly drops the requirement on the existence of inverses from the group setting before. This is believed to prevent the applicability of Kuperberg's algorithm.

But how do we instantiate this monoid action? As in the steps before, this unfortunately adds some complexity, but will result in an overview of how MIKE works. We will use the so-called *module action*, where modules act on *abelian varieties*; a framework first proposed by Robert [Rob24]. For our high-level overview in this section, it suffices to give a very rough description of these objects, and defer more details to later levels.

⚠ Warning

The descriptions in this section are deliberately vague and should not be confused with proper definitions. The aim of this section is solely to gain an initial intuition of the objects and computations in MIKE.

An abelian variety X is in some sense a generalisation of an elliptic curve. That is, it is defined by the zero set of some polynomials $F_i(x_0 : \cdots : x_m)$, and consists of points $P = (a_0 : \cdots : a_m)$ that are zeroes of the polynomials. Just like elliptic curves, they have a group structure, such that for points $P, Q \in X$, we get $P + Q \in X$. An important feature of an abelian variety is its *dimension* g, and elliptic curves are simply a special case of dimension $g = 1$. This generalisation also extends to isogenies: a g-dimensional isogeny $\Phi : X \to X'$ again is some *nice* map, now between g-dimensional abelian varieties X, X'.

To describe the modules involved in MIKE we first define a ring R. In our context we will for instance use $R = \mathbb{Z}[\sqrt{-p}]$ for a prime $p \equiv 3 \bmod 4$. For now, it suffices to abstractly think of a module M as a full-rank lattice $M \subseteq R^n$ with $n \geq 1$, so that an element $m \in M$ is of the form $m = (r_1, \ldots, r_n)$ with $r_i = a_i + b_i\sqrt{-p} \in \mathbb{Z}[\sqrt{-p}]$. We say that the module M in this case has *rank* n. Modules form a commutative monoid, and the operation is tensoring $M \otimes M'$, an analogue of multiplication for modules, although the details don't matter here. As the rank of this product is the product of their ranks, it is clear that there is generally no inverse for modules.

Putting things together, we want to see how a rank-n module M acts on an abelian variety X of dimension g. We will untangle this both in theory and practice in the following levels. For now, we content ourselves with this description: the action $Y = M * X$, where the module M has rank n and the variety X has dimension g, results in an abelian variety Y of dimension $n \cdot g$, and we can compute Y through an isogeny $\Phi : X^n \to Y$ of dimension $n \cdot g$. You will now have to believe us that (i) this works and (ii) this module action is commutative:

$$M_A * (M_B * X) = (M_A \otimes M_B) * X = (M_B \otimes M_A) * X = M_B * (M_A * X).$$

It's then easy to define a module-action variant of Diffie–Hellman. We choose to work with modules of rank two, and as starting object, we pick an abelian variety of dimension one, i.e., an elliptic curve E_0. Alice and Bob choose secret rank-2 modules M_A and M_B, and compute $X_A = M_A * E_0$ and $X_B = M_B * E_0$ as public keys, respectively. This means that X_A and X_B are two-dimensional abelian varieties. The shared secret is

$$X_{AB} = M_A * X_B = (M_A \otimes M_B) * E_0 = M_B * X_A,$$

a four-dimensional abelian variety.

This is how MIKE works, viewed from an abstract perspective. Note that if we worked with rank-1 modules M_A and M_B instead of rank-2 modules, and started from an elliptic curve E_0, the whole protocol would work in dimension one. This looks very similar to group-action NIKEs like CSIDH, which is not a coincidence, as we will see later on. This motivates the choice of rank-2 modules in MIKE, as the smallest possible choice to leave the group action setting.

In a nutshell, we again stripped off some of the underlying structure that is not strictly required to build a NIKE, and hopefully gain resistance against subexponential quantum attacks using Kuperberg's algorithm. So the currently known classical *and* quantum attacks both have exponential complexity.[2] But we also pay the price of an increased complexity of instantiating the original framework by Robert [Rob24], as impressively demonstrated by the need of a full paper plus 100 pages of appendices to introduce MIKE [BDD+26].

[2] We could therefore argue that MIKE is as secure as a NIKE can get both classically and quantumly. If you asked us if this is true, we'd answer with a clear ~~*Yes!*~~ *We don't know!* After all, MIKE is a very recent proposal and requires much more cryptanalysis effort before considering any kind of practical applications.

Level 2: MIKE from isogenies

In Level 1, we explained how MIKE, when viewed as simply a protocol, is essentially like Diffie–Hellman key exchange, but with much less structure in the underlying action. In this section, we describe MIKE from a computational point of view. This view is a bit different, as we find by comparing the core computational diagram in Fig. 3 with Fig. 2b.

The goal of this section is to explain the computational diagram Fig. 3 in more detail: how MIKE *in practice* is essentially the computation of *isogenies*, the 1-dimensional isogeny $\varphi_A : E'_0 \to E_A$, resp. φ_B, between elliptic curves E'_0 and E_A for key generation, and the four-dimensional isogeny Φ_A, resp. Φ_B for shared secret derivation to obtain the 4-dimensional abelian variety X_{AB}. We first explain what a (one-dimensional) isogeny is.[3]

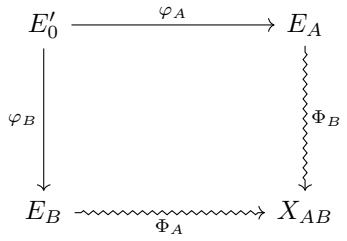


Figure 3: Computational diagram of MIKE.

⚠ Warning

It is clear that the dimensions between Fig. 2b and Fig. 3 are mismatched, and that the top left corner is different. We explain both differences in this section and Level 3, and apologise for introducing such confusion up front.

After explaining a bit more generally how one- and four-dimensional isogenies work in Section 2.1, we will need to clear up a few confusions that arise from Fig. 3. Namely, the four-dimensional isogeny Φ_A is properly defined as an isogeny between four-dimensional varieties $Y_B \to X_{AB}$. We explain how that four-dimensional variety Y_B is related to E_B, and how the kernels of φ_A and Φ_A are related in Section 2.2. Finally, we finish the shared secret computation in Section 2.3 by showing how Alice and Bob ensure they end up with the same shared key from X_{AB}.

Remark 1. We urge the reader to also read the excellent blog by Giacomo Pope [Pop26], available at

`https://giacomopope.com/mike/`.

This explains the isogenies in MIKE from an implementor's perspective, with a bit more focus on the underlying operations. We provide a complementary, high-level view of the computational diagram.

2.1 One-dimensional isogenies

Definition 1. A (one-dimensional) *isogeny* $\varphi : E \to E'$ is a map between elliptic curves such that

$$\varphi(P+Q) = \varphi(P) + \varphi(Q), \quad \text{for all } P, Q \in E,$$

[3] Costello [Cos19] has written an excellent introduction to SIDH and SIKE. Although SIKE is fundamentally different to MIKE, the key generation is almost the same, and his tutorial is therefore also very helpful for understanding the one-dimensional isogenies in MIKE.

that is surjective and has a finite kernel

$$\ker \varphi := \{P \in E \mid \varphi(P) = \mathcal{O}\}.$$

In general, an isogeny can be written as fractions of polynomials in the following way

$$\varphi : (x, y) \mapsto \left(\frac{f_1(x)}{f_2(x)},\ y \cdot \frac{g_1(x)}{g_2(x)} \right)$$

for some polynomials f_1, f_2, g_1, and g_2 in $\mathbb{F}_q[x]$. We may think of the *degree* of an isogeny as the maximal degree of the polynomials f_1 and f_2 that we need to write down the isogeny as a rational map. Let's look at three examples of isogenies, before we discuss the specific isogenies in MIKE's key generation.

Example 1.
Let $E : y^2 = x^3 + x$ and $E' : y^2 = x^3 + 5$ be two elliptic curves over $\mathbb{F}_{11}$. Then, there is an isogeny $\varphi : E \to E'$ of degree 3 given by

$$\varphi \quad : \quad (x, y) \quad \mapsto \quad \left(\frac{x^3 + x^2 + x + 2}{(x-5)^2}\ ,\ y \cdot \frac{x^3 - 4x^2 + 2}{(x-5)^3} \right).$$

To verify that this is actually an isogeny, one should check that for two points $P = (x_P, y_P)$ and $Q = (x_Q, y_Q)$ such that $R = P+Q = (x_R, y_R)$, we also have that $\varphi(P)+\varphi(Q) = \varphi(R)$. In practice, this is rather cumbersome, and we hope that you simply believe us on our blue and green eyes.

Example 2.
Let $E : y^2 = x^3 + Ax^2 + x$ be an elliptic curve over $\mathbb{F}_p$ for some prime p. Then there is a special isogeny *Frobenius* denoted by π, given by

$$\pi : E \to E, \qquad (x, y) \mapsto (x^p, y^p).$$

Clearly, π is of degree p, and acts trivially on a point $P = (x_P, y_P) \in E(\mathbb{F}_p)$, as we have $x_P^p = x_P$ and $y_P^p = y_P$ for $x_P, y_P \in \mathbb{F}_p$.
We could say that π captures which points of E are defined over $\mathbb{F}_p$, as π acts trivially, i.e. $\pi(P) = P$, only on points $P \in E(\mathbb{F}_p)$. The same holds for higher powers: π^k acts trivially on points $P \in E(\mathbb{F}_{p^k})$. More algebraically, we write that $\ker(\pi^k - 1) = E(\mathbb{F}_{p^k})$.

Example 3.
Anyone familiar with basic elliptic curve cryptography knows the *scalar multiplication map* $[n]$, which adds n copies of P together,

$$[n] : E \to E, \qquad P \mapsto \underbrace{P + \ldots + P}_{n \text{ times}}.$$

This map is also an isogeny! It is of degree n^2, which means that we could write this map in terms of fractions of polynomials of degree n^2. Silverman [Sil09] describes this in detail. We write the kernel of $[n]$ as $E[n]$, that is, the points $P \in E$ such that $[n]P = \mathcal{O}$. When E is defined over a finite field $\mathbb{F}_q$, this generally doesn't mean that the points $P \in E[n]$ are also $\mathbb{F}_q$-rational. Instead, we need to go to a (much!) larger extension field $\mathbb{F}_{q^k}$ to get $E[n] \subseteq E(\mathbb{F}_{q^k})$. More formally, we write

$$E[n] := \{P \in E(\overline{\mathbb{F}_q}) \ \mid \ [n]P = \mathcal{O}\}.$$

The first and last example[4] are what we call *separable* isogenies, and in these cases their degrees are equal to the size of their kernels. Frobenius π, the second example, is

[4] Assuming n is coprime to p.

inseparable: there are no non-zero points $P \in E$ that map to $\mathcal{O}$, but its degree is clearly p. From now on, we will only compute separable isogenies, but π will later reappear in crucial positions!

Isomorphisms. A certain class of isogenies $\varphi : E \to E'$ is very special: those of degree 1. They are *isomorphisms* of elliptic curves, and we write $E \cong E'$ whenever the curves are isomorphic. This concept shows up a lot because we usually think of curves up to isomorphism.

Example 4.
A good example that we need later is the curve $E_0' : y^2 = x^3 - x$ over a finite field $\mathbb{F}_p$. For practical purposes, we would like an isomorphic curve in Montgomery form, that is, $E_M : y^2 = x^3 + Mx^2 + x$ for some $M \in \mathbb{F}_p$. Let us pick $p = 47$ as an example, then $M = 34$ gives us such an isomorphic curve by the isomorphism

$$\psi : E_0' \to E_M, \qquad (x, y) \mapsto (27x + 20, 15y).$$

The field of definition of an isomorphism ψ is important for us, which can be stated as the field $\mathbb{F}_q$ so that we can write ψ in terms of $\mathbb{F}_q$-coefficients. Crucially, two curves E, E' can be defined over $\mathbb{F}_p$ with an isomorphism $\psi : E \to E'$ only defined over a field extension. We then say that E' is a *twist* of E. A specific case is when the isomorphism is defined over the (quadratic) extension $\mathbb{F}_{p^2}$, in which case E' is called the *quadratic twist* of E, which we usually simply call *the twist* of E and denote E^t. This twist is unique up to $\mathbb{F}_p$-isomorphism. We give two easy examples.

Example 5.
The easiest example of quadratic twists are $E_A : y^2 = x^3 + Ax^2 + x$ and $E_{-A} : y^2 = x^3 - Ax^2 + x$ over $\mathbb{F}_p$ with $p \equiv 3 \bmod 4$. Writing $\mathbb{F}_{p^2} = \mathbb{F}_p(i)$ with $i^2 = -1$, these curves are $\mathbb{F}_{p^2}$-isomorphic by

$$\psi : E_A \to E_{-A}, \quad (x, y) \mapsto (-x, iy),$$

and ψ cannot be written as a map over $\mathbb{F}_p$. Thus, $E_A^t = E_{-A}$.

Example 6.
From the previous example, one would say that $E_0 : y^2 = x^3 + x$ should have quadratic twist $E_0^t = E_{-0} = E_0$. And this is true! The reason is that E_0 is rather special as it has extra automorphisms, and the definition given above only really works for curves without extra autmorphisms. For example, the automorphism $\iota : (x, y) \mapsto (-x, iy)$ is an $\mathbb{F}_{p^2}$-isomorphism $E_0 \xrightarrow{\sim} E_0$ that cannot be defined over $\mathbb{F}_p$. In this sense, one should think of E_0 as being its own quadratic twist. The same happens for the curve $E_0' : y^2 = x^3 - x$.
Curves with extra automorphisms are rather special with regards to twists. For example, E_0 and E_0' are actually isomorphic too! One can see that the two curves cannot be isomorphic over $\mathbb{F}_p$, as E_0 only has one point of order 2, namely $(0, 0)$, but E_0' has three such points, namely $(0, 0)$, $(1, 0)$, and $(-1, 0)$. This can also be derived from their group structures

$$E_0(\mathbb{F}_p) = \mathbb{Z}_{p+1}, \qquad E_0'(\mathbb{F}_p) = \mathbb{Z}_{\frac{p+1}{2}} \times \mathbb{Z}_2.$$

It turns out that E_0' is a twist of E_0, but not a quadratic twist!

We will often write that something is unique 'up to isomorphism' when the result essentially holds for the isomorphism classes of the objects involved, where we consider isomorphisms over $\overline{\mathbb{F}}_p$.

Isogenies and Subgroups

The following result is crucial for isogeny-based cryptography, and should be remembered from now on until the passing of days.

Theorem 1 (General Mantra)**.** *There is a one-to-one correspondence between* finite subgroups $G \subseteq E$ *and separable* isogenies $\varphi : E \to E'$*, up to isomorphism of the codomain.*

More precisely, every such subgroup $G \subseteq E$ can be turned into an isogeny $\varphi_G : E \to E'$ with $\ker \varphi_G = G$. Vice versa, every isogeny $\varphi : E \to E'$ defines a unique subgroup $G \subseteq E$ by $G = \ker \varphi$, up to post-composition of φ with an isomorphism. Furthermore, as we have $\ker \varphi = G$, we have $\deg \varphi = |G|$.

The above theorem tells us nothing on how we would compute such an isogeny, even if we know the corresponding subgroup G. Luckily, Vélu [Vél71] gives us simple formulas to compute φ given the group G, in complexity $O(|G|)$. This means that if $\deg \varphi = |G|$ is quite large, computing such an isogeny is unfortunately rather slow. Luckily, any isogeny of degree $d = \prod \ell_i^{e_i}$ can be broken up into these smaller steps of prime degree ℓ_i. For example, an isogeny of degree 3^e can be computed as e steps of degree 3, where each step of degree 3 is a simple formula given by Vélu. We use a simplification of the simplified formulation from Costello and Hisil [CH17, Thm. 1] to demonstrate this.

Theorem 2. *Let $E_A : y^2 = x^3 + Ax^2 + x$ be an elliptic curve over $\mathbb{F}_p$ with $p > 2$. Let $K \in E_A(\overline{\mathbb{F}}_p)$ be a point of order $\ell = 2d + 1$, generating a subgroup $G = \langle K \rangle \subseteq E_A(\overline{\mathbb{F}}_p)$ of order ℓ. Then the degree-ℓ isogeny $\varphi : E_A \to E_{A'}$ with $\ker \varphi = G$ has codomain $E_{A'} : y^2 = x^3 + A'x^2 + x$, where*

$$A' = C^2 \cdot \left(6 \cdot \sum_{i=1}^{d} \frac{1 - x_{[i]K}^2}{x_{[i]K}} + A \right), \qquad C = \prod_{i=1}^{d} x_{[i]K}.$$

As a map, φ is given by

$$\varphi : (x, y) \mapsto (f(x), C \cdot y \cdot f'(x))$$

where

$$f(x) = x \cdot \prod_{i=1}^{d} \left(\frac{x \cdot x_{[i]K} - 1}{x - x_{[i]K}} \right)^2.$$

and $f'(x)$ its derivative.

Although this includes a few formulas, at its core it only requires the x-coordinates of the first $\frac{\ell-1}{2}$ multiples of K, so once we compute K, $[2]K$, ..., $[\frac{\ell-1}{2}]K$, the rest are simple formulas. We demonstrate this for $\ell = 5$.

Example 7 (Vélu isogeny)**.**

Take the curve $E_0 : y^2 = x^3 + x$ over the field $\mathbb{F}_p$ with $p = 419$. In $E_0(\mathbb{F}_p)$, we have a point $K = (349, 133)$ of order 5, which generates a subgroup

$$G = \{ \ (349, 133),\ (185, 346),\ (185, 73),\ (349, 286),\ \mathcal{O} \ \}.$$

Using the above Vélu's formulas, we can compute $\varphi : E_0 \to E_A$ with $\ker \varphi = G$. We get $d = \frac{\ell-1}{2} = 2$, so we just need $x_K = 349 = -70$ and $x_{[2]K} = 185$ to compute $C = -70 \cdot 185 = 39$ and

$$A = C^2 \cdot \left(6 \cdot \sum_{i=1}^{d} \frac{1 - x_{[i]K}^2}{x_{[i]K}} + 0 \right) = 221 \cdot 6 \cdot 39^2 = 199.$$

For the actual map in terms of (x, y)-coordinates, the core computation requires the

polynomial $f(x)$ which we compute as

$$f(x) = \frac{x \cdot (-70x-1)^2 \cdot (185x-1)^2}{(x+70)^2(x-185)^2}.$$

Thus, we get the isogeny $\varphi : E_0 \to E_{199}$ of degree 5 as the map that sends $P = (x, y)$ to

$$\left(\frac{x \cdot (-70x-1)^2 \cdot (185x-1)^2}{(x+70)^2(x-185)^2},\ -y \cdot \frac{179x^6 - 162x^5 - 107x^4 + 95x^3 - 104x^2 + 54x + 155}{(x+70)^3(x-185)^3} \right).$$

In practice, we only use one-dimensional isogenies φ with a kernel $G = \ker \varphi$ isomorphic to a cyclic group. This means that G is generated by a single point $K \in E$. This makes many steps in its computation simpler! We refer to such isogenies as *cyclic*. The situation in MIKE is even better: we essentially only use 2-isogenies[5] with kernel generators $K = (\alpha, 0)$ where $\alpha \neq 0$.[6] This is essentially the simplest isogeny to compute, and so the formulas, from Renes [Ren18], are remarkably easy: the isogeny $\varphi : E_A \to E_{A'}$ with $\ker \varphi = \{K, \mathcal{O}\}$ is given by

$$A' = 2(1-2\alpha^2), \qquad f(x) = x \cdot \frac{x \cdot \alpha - 1}{x - \alpha}, \qquad \varphi : (x, y) \mapsto (f(x),\ \sqrt{\alpha} \cdot y \cdot f'(x)). \quad (1)$$

The isogeny in MIKE's key generation

We pick the prime $p = 633 \cdot 2^{308} - 1$ and work mostly over $\mathbb{F}_p$ or $\mathbb{F}_{p^2}$. Alice's secret isogeny in key generation $\varphi_A : E_0' \to E_A$ will be a cyclic isogeny of degree 2^e, for the large value $e = 228$, for MIKE.[7] Thus, by Theorem 1, this corresponds to a (secret) subgroup $G_A = \langle K_A \rangle$ of E_0' of size 2^e generated by a (secret) point K_A of order 2^e. To make the computation efficient, we want K_A to be defined over $\mathbb{F}_{p^2}$.

The start. The starting curve in MIKE will be the very special curve $E_0' : y^2 = x^3 - x$. As almost always, in isogeny-based cryptography we like to work with curves that are *maximal supersingular*, which for now simply means that E_0' has the following group structure over $\mathbb{F}_{p^2}$,

$$E_0'(\mathbb{F}_{p^2}) \cong \mathbb{Z}_{p+1} \times \mathbb{Z}_{p+1}.$$

So, as $2^e \mid p+1$, the group $E_0'[2^e] \cong \mathbb{Z}_{2^e} \times \mathbb{Z}_{2^e}$ of all points such that $[2^e]P = \mathcal{O}$ is a subgroup of $E_0'(\mathbb{F}_{p^2})$. This means that we can find two points $P_0, Q_0 \in E_0'(\mathbb{F}_{p^2})$ of order 2^e such that every $K \in E_0'[2^e]$ can be written as a linear combination

$$K = [a]P_0 + [b]Q_0, \quad \text{for some } a, b \in \mathbb{Z}_{2^e}.$$

We compute these points P_0 and Q_0 once and for all in the implementation. Their precise description does not matter too much for now, although (P_0, Q_0) needs to satisfy some requirements that we will see in Level 3. We want K to be of order 2^e, so a and b can't both be even. As we are mostly interested in the subgroup generated by K, because of the General Mantra, assume for now that a is odd, so that its inverse mod 2^e exists. Then, replacing K by $[a^{-1}]K = P_0 + [a^{-1}b]Q_0$ doesn't change the subgroup $\langle K \rangle$.

[5] Although we often use 4-isogenies for computational efficiency.

[6] Such points $K \in E_A$ of order 2, where $E_A : y^2 = x^3 + Ax^2 + x$ have an x-coordinate that is a root of $x^3 + Ax^2 + x$, that is, $\alpha^3 + A\alpha^2 + \alpha = 0$. Clearly, $\alpha = 0$ is always one of the roots.

[7] This is the smallest value e such that finding φ_A given only E_0' and E_A should take about 2^{128} operations. Feel free to take a larger e up to 306 if you want. We could even go up to $e = 308$, but we will see in Level 3½ that it is beneficial to make sure that $2^e \mid \frac{p+1}{4}$, so $e \leq 306$.

⚠ Warning

MIKE uses a slightly different starting curve $E_0' : y^2 = x^3 - x$ than most isogeny-based cryptography such as SIDH, CSIDH, or SQIsign, which often uses $E_0 : y^2 = x^3 + x$. We will use both E_0 and E_0' in examples in this text, as E_0 has the benefit of already being in Montgomery form, and so makes a lot of examples easier to understand.

When working with E_0' over a given field $\mathbb{F}_p$, we can always precompute an isomorphism $\psi : E_0' \to E_M$ to a curve E_M in Montgomery form, and use the points $\psi(P_0), \psi(Q_0)$ for the basis of $E_M[2^e]$.

This sounds rather cumbersome and circumstantial, since MIKE *could* start with the curve E_0 in the first place. But there are good reasons to follow MIKE's approach, which will become clearer later in Level 3. In this section, we will sometimes use E_0 in examples for simplicity.

The secret. If Alice now picks a secret value $x_A \in \mathbb{Z}_{2^e}$, she can compute a secret point $K_A = P_0 + [x_A]Q_0$ of order 2^e, which corresponds to a secret subgroup $G_A = \langle K_A \rangle \subseteq E_0'(\mathbb{F}_{p^2})$. By Theorem 1 and Vélu's formulas, this allows her to compute a secret one-dimensional 2^e-isogeny $\varphi_A : E_0' \to E_A$ with $\ker \varphi_A = G_A$ in e steps of degree 2. Her *public key* is then the curve $E_A : y^2 = x^3 + Ax^2 + x$, for some value $A \in \mathbb{F}_{p^2}$. Of course, in practice, her public key is the value A, the *Montgomery coefficient* of E_A, of size $2 \log p$ bits. Later on, we will need the value $\frac{1+x_A}{2}$ to exist in $\mathbb{Z}_{2^e}$, so we should make sure Alice chooses an odd integer x_A here, instead of any $x_A \in \mathbb{Z}_{2^e}$. In practice, for deeper technical reasons, we should even restrict the secret scalar to $x_A \equiv 3 \bmod 8$, although these reasons are out of scope for this work.

x-only arithmetic. The only thing we really need to mention about this computation is that it can be done entirely given only the x-coordinate of K_A, which is also visible in Theorem 2. This is an analogue of computations in classical elliptic curve cryptography: by using *x-only arithmetic*, elliptic curve implementations often become much faster. The same holds for isogeny-based cryptography.

Example 8 (Key generation)**.**
We demonstrate key generation, using E_0 instead of E_0', for the small prime $p = 47 = 3 \cdot 2^4 - 1$, with $\mathbb{F}_{p^2} = \mathbb{F}_p(i)$ where $i^2 = -1$. We start on $E_0 : y^2 = x^3 + x$, and look for two points $P_0, Q_0 \in E_0(\mathbb{F}_{p^2})$ such that $E_0[2^4] = \langle P_0, Q_0 \rangle \cong \mathbb{Z}_{2^4} \times \mathbb{Z}_{2^4}$, for example

$$P_0 = (20 + 31i, 8 + 42i), \qquad Q_0 = (19, 43).$$

Alice chooses her favourite number, $x_A = 7$ as her secret key, and computes

$$K_A = P_0 + [7]Q_0 = (22 + 19i, 16 + 8i).$$

She could now compute the whole isogeny $\varphi_A : E_0 \to E_A$ with $\ker \varphi_A = \langle K_A \rangle$ in one go, but for larger degrees 2^e this would be entirely unfeasible. Instead, we break this up into $e = 4$ steps of 2-isogenies

$$\varphi_A = \varphi_4 \circ \varphi_3 \circ \varphi_2 \circ \varphi_1,$$

each of which has a kernel $\ker \varphi_i$ generated by some point of order 2. The first kernel is generated by the point of order 2 given by

$$K_1 = [8]K_A = (i, 0).$$

By Eq. (1), the isogeny $\varphi_1 : E_0 \to E_{A_1}$ is given by $A_1 = 2(1+2) = 6$ and $f_1(x) = \frac{x\cdot(ix-1)}{x-i}$. To get the next kernel, we first compute $K' = \varphi_1(K_A)$ of order 8, but we will only need its x-coordinate

$$x_{K'} = f_1(x_{K_A}) = \frac{(22+19i)\cdot(22i-19-1)}{22+18i} = 22+38i.$$

We can now repeat the previous step: the generator of the second kernel is the point of order 2

$$K_2 = [4]K' = (30, 0),$$

which gives $\varphi_2 : E_{A_1} \to E_{A_2}$ with $A_2 = 21$ and $f_2(x) = \frac{x\cdot(30x-1)}{x-30}$. We push K' through to get $K'' = \varphi_2(K')$ of order 4, whose x-coordinate is

$$x_{K''} = f_2(x_{K'}) = \frac{(22+38i)\cdot(30\cdot(22+38i)-1)}{-8+38i} = 24+41i,$$

giving the generator of the third kernel point K_3 and isogeny $\varphi_3 : E_{A_2} \to E_{A_3}$ as

$$K_3 = [2]K'' = (24, 0), \qquad A_3 = 1, \qquad f_3(x) = \tfrac{x\cdot(24x-1)}{x-24}.$$

One last time we push K'' through to get $K''' = \varphi_3(K'')$ of order 2 with x-coordinate

$$x_{K'''} = f_3(x_{K''}) = 23+41i.$$

This is also immediately the kernel generator of the fourth isogeny $\varphi_4 : E_{A_3} \to E_{A_4}$ given by

$$A_4 = 4+23i, \qquad f_4(x) = \tfrac{x\cdot((23+41i)x-1)}{x-(23+41i)}.$$

We have thus computed Alice's secret isogeny $\varphi_A : E_0 \to E_{A_4}$ as the composition

$$\varphi_A : E_0 \xrightarrow{\varphi_1} E_{A_1} \xrightarrow{\varphi_2} E_{A_2} \xrightarrow{\varphi_3} E_{A_3} \xrightarrow{\varphi_4} E_{A_4}.$$

For efficiency, in a real implementation, we represent the x-value x_K of K_A as a *projective coordinate* $(X_0 : X_1) \in \mathbb{P}^1(\mathbb{F}_{p^2})$ such that $x_K = \frac{X_0}{X_1}$. This allows us to keep track of the numerator and denominator of x_K separately, without the need to compute expensive inversions. However, the x-coordinate of a point K is the same as the x-coordinate of $-K$, and so, using x-only arithmetic we cannot distinguish such points. This causes many issues: when adding together P and Q in x-only coordinates, we can compute x_{P+Q} or x_{P-Q} given just x_P and x_Q, but we can only know which of the two we compute given the other value. In practice, we call this *differential addition*, where given x_P, x_Q, and x_{P-Q}, we can compute x_{P+Q}.

This is only one example of how x-only arithmetic often makes computations more cumbersome, but also faster. It is somewhat surprising that even with such obstacles of not being able to differentiate P and $-P$, we can still compute the isogeny $\varphi : E \to E'$ properly. Although all the formulas that are involved become slightly more complex, the resulting speed-up is tremendous and almost every implementation in isogeny-based cryptography thus uses this x-only arithmetic. Once more, we refer to the writings of Costello [Cos19] and Pope [Pop26] to emphasise their great introductions on these topics.

2.2 The isogeny in MIKE's shared secret

To properly describe Φ_A from Fig. 3, we first give a quick introduction to four-dimensional isogenies, before we describe the domain Y_B of Φ_A and its kernel.

Four-dimensional isogenies

⚠ Warning

This section describes four-dimensional isogenies as an analogue of one-dimensional isogenies for the sake of simplicity. This is an oversimplification, as in practice, going up in dimension significantly increases how difficult these objects are to work with, and how annoying it is to derive the right formulas, or to even write down the elements we are working with. Significantly, we skip explaining *maximal isotropy* for kernels of (polarised) isogenies, a technical requirement that ensures the codomain of an isogeny is in the shape we want it to be.

Broadly speaking, a four-dimensional isogeny $\Phi : X \to X'$ is again a surjective group homomorphism with a finite kernel, but now between four-dimensional abelian varieties X and X'.

Similar to the one-dimensional case, we still have a correspondence between subgroups $G \subseteq X$ and (separable) isogenies $\Phi : X \to X'$ with $\ker \Phi = G$. Similar to the one-dimensional case, MIKE uses only 2^e-isogenies, which can be split up into e steps of 2-isogenies. However, different to the one-dimensional case, our 2^e-isogeny φ has a cyclic kernel $\ker \varphi \cong \mathbb{Z}_{2^e}$, where our four-dimensional 2^e-isogeny Φ has a kernel

$$\ker \Phi \cong \mathbb{Z}_{2^e} \times \mathbb{Z}_{2^e} \times \mathbb{Z}_{2^e} \times \mathbb{Z}_{2^e},$$

and so has degree $\deg \Phi = 2^{4e}$. And so, where our one-dimensional isogeny could be described and computed given only the kernel generator K_A of order 2^e such that $\ker \varphi = \langle K_A \rangle$, in dimension four we need *four* points K_1, K_2, K_3, K_4 of order 2^e, linearly independent, to describe $\ker \Phi = \langle K_1, K_2, K_3, K_4 \rangle$, that is, every point K in the kernel of Φ is a linear combination

$$K = [a_1]K_1 + [a_2]K_2 + [a_3]K_3 + [a_4]K_4, \quad a_i \in \mathbb{Z}_{2^e}.$$

The analogue of x-only arithmetic. In dimension one, we used x-only arithmetic to speed up elliptic curve and isogeny computations, and we noticed that we could not differentiate P and $-P$ anymore. In more fancy terms, we were working on the elliptic curve E *up to* the automorphism $[-1]$. This idea generalises to higher dimensions as well: where we could represent $P \in E(\mathbb{F}_{p^2})$ up to $[-1]$, using x-only arithmetic, in 2^1 coordinates as $(X_0 : X_1) \in \mathbb{P}^1(\mathbb{F}_{p^2})$, we can also represent $P \in X(\mathbb{F}_p)$ up to $[-1]$ using an analogue of x-only arithmetic[8], properly called *Kummer arithmetic*, in 2^4 coordinates in $\mathbb{P}^{15}(\mathbb{F}_p)$. Thus, it is perfectly fine to think of points $P \in X(\mathbb{F}_p)$ as

$$P = (X_0 : \ldots : X_{15}), \quad X_i \in \mathbb{F}_p,$$

with the usual caveats that we cannot distinguish P and $-P$, and we need to use differential addition to compute $P + Q$ given P, Q and $P - Q$, and so on and so forth.

The domain of Φ_A

In the end, we want Φ_A to be a four-dimensional isogeny $Y_B \to X_{AB}$ where Y_B and X_{AB} are four-dimensional abelian varieties defined over $\mathbb{F}_p$. But key generation leaves us with a (one-dimensional) elliptic curve E_A, resp. E_B defined over $\mathbb{F}_{p^2}$. It turns out that we can simply do two rather simple things to derive first a two-dimensional X_A over $\mathbb{F}_p$ from E_A

[8] Although we can no longer state this idea as "the x-coordinate x_P of a point $P \in X$", there are good mathematical reasons why this is the correct analogy.

over $\mathbb{F}_{p^2}$, resp. X_B from E_B, and then a four-dimensional Y_A from X_A, resp. Y_B from X_B. Luckily, we can dress up both techniques in rather fancy mathematical terms.[9]

From dim-1 over $\mathbb{F}_{p^2}$ to dim-2 over $\mathbb{F}_p$. Given an elliptic curve $E_A : y^2 = x^3 + Ax^2 + x$ over $\mathbb{F}_{p^2}$, hence $A \in \mathbb{F}_{p^2}$, the simplest way to get something over $\mathbb{F}_p$ is to just pretend that we are working over $\mathbb{F}_p$. That is, write $x = x_0 + i \cdot x_1$ and $y = y_0 + i \cdot y_1$ for new variables x_0, x_1, y_0, y_1, and similarly write $A = A_0 + i \cdot A_1$ for $A_0, A_1 \in \mathbb{F}_p$. We can then rewrite the left-hand side as

$$y^2 = (y_0 + i \cdot y_1)^2 = \underbrace{(y_0^2 - y_1^2)}_{\text{real part}} + i \cdot \underbrace{(2y_0y_1)}_{\text{im. part}},$$

and the right-hand side $x^3 + Ax^2 + x$ as

$$\begin{aligned} x^3 + Ax^2 + x &= (x_0 + i \cdot x_1)^3 + (A_0 + i \cdot A_1) \cdot (x_0 + i \cdot x_1)^2 + (x_0 + i \cdot x_1) \\ &= (x_0^3 - 3x_0x_1^2 + A_0(x_0^2 - x_1^2) + x_0 - 2A_1x_0x_1) \\ &\quad + i \cdot (3x_0^2x_1 - x_1^3 + 2A_0x_0x_1 + A_1(x_0^2 - x_1^2) + x_1). \end{aligned}$$

When we collect the real and imaginary parts of these equations, we get two equations W_{Re} and W_{Im} that describe a surface (a 2-dimensional variety) W over $\mathbb{F}_p$, given by

$$\begin{aligned} W_{\text{Re}} : y_0^2 - y_1^2 &= x_0^3 - 3x_0x_1^2 + A_0(x_0^2 - x_1^2) + x_0 - 2A_1x_0x_1, \\ W_{\text{Im}} : 2y_0y_1 \quad &= 3x_0^2x_1 - x_1^3 + 2A_0x_0x_1 + A_1(x_0^2 - x_1^2) + x_1. \end{aligned}$$

Now, given a point $P = (x, y)$ over $\mathbb{F}_{p^2}$ satisfying $y^2 = x^3 + Ax^2 + x$, i.e. $P \in E_A(\mathbb{F}_{p^2})$, we get a point $W(P) = (x_0, x_1, y_0, y_1)$ over $\mathbb{F}_p$ satisfying both W_{Re} and W_{Im}, and so we say $W(P) \in W$. Furthermore, for three points satisfying $R = P + Q$ on E_A we can similarly give a group law on W such that $W(R) = W(P) + W(Q)$. All in all, we find that W is an *abelian surface*, an abelian variety of dimension 2, over $\mathbb{F}_p$ which is in many ways similar to E_A over $\mathbb{F}_{p^2}$, although also crucially different in a few ways we will see later.

The letter W here is not chosen by accident, as the abelian surface we have constructed is called the *Weil restriction* of E_A. We will therefore refer to this surface as $W(E_A)$. In more mathematical terms, the Weil restriction is *functorial* and therefore has a lot of nice properties to ensure that everything works as we want it to. The Weil restriction is in general a natural way to turn a variety of dimension n over $\mathbb{F}_{q^k}$ into a variety of dimension $n \cdot k$ over $\mathbb{F}_q$, but we only need this simplest variant in MIKE to turn E_A into a dimension-2 variety $X_A = W(E_A)$ over $\mathbb{F}_p$. It is this X_A that we also see in Fig. 2b.

⚠ Warning

In reality, X_A is a slightly different object from $W(E_A)$, namely $W'(E_A)$. This construction comes with a few subtleties, which we will see in Level 3. At this point, however, it is okay to think of X_A as $W(E_A)$. Pinky promise.

Another model for the Weil restriction. In practice, the above model for the Weil restriction, using points (x_0, x_1, y_0, y_1) that vanish on both W_{Re} and W_{Im} is cumbersome and slow. The practical way to use the Weil restriction is rather surprising: when we view the surface $W(E_A)$, defined over $\mathbb{F}_p$, as a variety over $\mathbb{F}_{p^2}$, one can show that we get the isomorphism

$$W(E_A) \cong E_A \times E_A^{(p)},$$

[9] We leave this mathematical charade[10] to the proper papers on MIKE.

[10] In this footnote, "mathematical charade" should be read as a compliment similar to "abstract nonsense".

where $E_A^{(p)} : y^2 = x^3 + A^p x^2 + x$ is the *conjugate curve* of E_A. Strangely enough, on $W(E_A)$ there is a different Frobenius map than usual, namely $\pi_W : (P, Q) \mapsto (\pi(Q), \pi(P))$ where $\pi(P) = (x^p, y^p) \in E_A^{(p)}$. We should think of $\mathbb{F}_p$-rational points as points on which Frobenius acts trivially, which in this case means that $\mathbb{F}_p$-rational points $W(P)$ on $W(E_A)$ corresponds to pairs of points $(P, \pi(P))$ on $E_A \times E_A^{(p)}$ with $P \in E_A(\mathbb{F}_{p^2})$, so that $\pi_W(P, \pi(P)) = (P, \pi(P))$. For us, this is very practical: we get a 2-dimensional surface $W(E_A)$ which for mathematical purposes is defined over $\mathbb{F}_p$, and we can simply think of the points on $W(E_A)(\mathbb{F}_p)$ as special elements $(P, \pi(P))$ for $P \in E_A(\mathbb{F}_{p^2})$ on the product of elliptic curves $E_A \times E_A^{(p)}$.

⚠ Warning

In the above, we use π to denote the map $\pi : E_A \to E_A^{(p)}$ defined by $(x, y) \mapsto (x^p, y^p)$. This should not be confused with the map $\pi : E \to E$ defined by $(x, y) \mapsto (x^p, y^p)$ from Example 2. The crucial difference being that the latter is defined for a curve E over $\mathbb{F}_p$, whereas above E_A is assumed to be over $\mathbb{F}_{p^2}$, so $E_A \neq E_A^{(p)}$.

From dim-2 over $\mathbb{F}_p$ to dim-4 over $\mathbb{F}_p$. So far, we have derived a 2-dimensional surface $X_A = W(E_A)$ over $\mathbb{F}_p$ from a 1-dimensional curve E_A over $\mathbb{F}_{p^2}$. But we are looking for a 4-dimensional variety Y_A over $\mathbb{F}_p$. As $2 \cdot 2 = 4$, a simple solution is to square $W(E_A)$. It turns out[11] that this is also the correct mathematical solution: we define $Y_A := W(E_A) \times W(E_A)$ as pairs of points on $W(E_A)$, and so Y_A is a 4-dimensional abelian variety over $\mathbb{F}_p$.

With the previous association between $W(E_A)$ and $E_A \times E_A^{(p)}$ we may then also think of general points on Y_A as 4-tuples (P, Q, R, S) with $P, R \in E_A$ and $Q, S \in E_A^{(p)}$. The $\mathbb{F}_p$-rational points $Y_A(\mathbb{F}_p)$ are the special 4-tuples $(P, \pi(P), R, \pi(R))$.

The kernel of Φ_A

From Section 2.1, we understand how a choice of $x_A \in \mathbb{Z}_{2^e}$ leads to Alice's secret isogeny $\varphi_A : E_0' \to E_A$. From the previous section we understand how we can associate a four-dimensional variety Y_A, resp. Y_B, to E_A, resp. E_B, and we have an intuitive understanding how 4-dimensional isogenies, such as $\Phi_A : Y_B \to X_{AB}$ work. Hence, if we can describe the four kernel generators K_1, K_2, K_3, K_4 of $\ker \Phi_A$ in terms of x_A, then we have all the ingredients we need to compute Φ_A. However, writing down points as tuples quickly becomes fairly unreadable. We therefore write the kernel points as linear combinations. Let P, Q be any basis[12] of $E_B[2^e]$ and define the following eight points on Y_B as

$$P_1 = (P, 0, 0, 0), \quad P_2 = (0, \pi(P), 0, 0), \quad P_3 = (0, 0, P, 0), \quad P_4 = (0, 0, 0, \pi(P)),$$

and

$$Q_1 = (Q, 0, 0, 0), \quad Q_2 = (0, \pi(Q), 0, 0), \quad Q_3 = (0, 0, Q, 0), \quad Q_4 = (0, 0, 0, \pi(Q)).$$

Informally, we may then rewrite the very crucial result of Theorem 4.1 [BDD+26] in the following manner.

Theorem 3. *If Alice's secret key is $x_A \in \mathbb{Z}_{2^e}$, giving $\varphi_A : E_0' \to E_A$, then the kernel of her isogeny $\Phi_A : Y_B \to X_{AB}$, where $Y_B := W(E_B)^2$, is generated by four points*

[11] We find out why in Level 3.

[12] In contrast to key generation, we no longer need a special basis (P_0, Q_0).

$K_1, K_2, K_3, K_4 \in Y_B$ *given by*

$$\begin{aligned}
K_1 &= \left[\tfrac{1-x_A}{2}\right](P_1 + P_2) + \left[\tfrac{1+x_A}{2}\right](P_3 + P_4),\\
K_2 &= \left[\tfrac{1-x_A}{2}\right](Q_1 + Q_2) + \left[\tfrac{1+x_A}{2}\right](Q_3 + Q_4),\\
K_3 &= \left[\tfrac{1+x_A}{2}\right](P_2 - P_1) + \left[\tfrac{1-x_A}{2}\right](P_3 - P_4),\\
K_4 &= \left[\tfrac{1+x_A}{2}\right](Q_1 - Q_2) + \left[\tfrac{1-x_A}{2}\right](Q_4 - Q_3).
\end{aligned}$$

Note that the kernel here only relies on Alice's secret x_A and Bob's public key E_B. This should feel miraculous to you, as it does to us. We hope to clear a bit of mystery in Level 3.

⚠ Warning

We call this map Φ_A to correspond correctly to Fig. 3. Although this is really the map that Alice should compute, we will see in Level 3 why the actual situation is a bit more subtle, and we refer to this map as F_A there.

Theorem 3 gives us the necessary mathematical detail to understand Alice's isogeny Φ_A. In practice, the computation requires us to translate those points K_i into points $(X_0 : \ldots : X_{15}) \in \mathbb{P}^{15}(\mathbb{F}_p)$ to apply the isogeny formulas for "x-only arithmetic" of dimension 4. This process, together with the right formulas for dimension-4 isogenies, is fairly cumbersome to derive. However, the good news is that the right steps are already derived in MIKE [BDD+26], so that any implementation can simply follow those formulas. With Fig. 4, we then have a more complete computational diagram for MIKE.

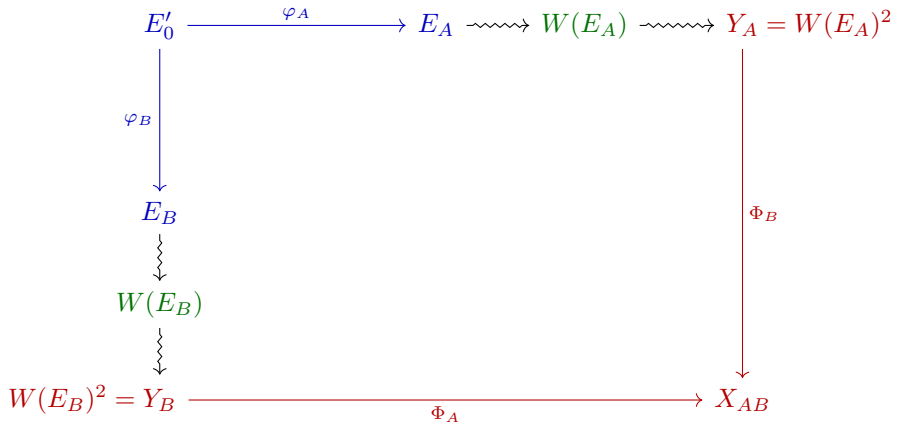


Figure 4: Detailed computational diagram of MIKE. Objects and isogenies in dimension one are shown in blue, in dimension two in green, and in dimension four in red. Squiggly arrows denote constructions, not isogenies.

Remark 2. We remark once more that Pope [Pop26] provides an excellent perspective on the viewpoint of implementing MIKE, when thinking of these elements simply as arrays of $\mathbb{F}_p$ and/or $\mathbb{F}_{p^2}$ elements.

2.3 Invariants in dimension four

Assuming you made it so far, you essentially know how to implement MIKE: you can compute φ_A and/or φ_B as one-dimensional isogenies, and you know the kernel of Φ_A and Φ_B as four-dimensional isogenies which you can implement if you have the right formulas

from (the implementation of) MIKE [BDD+26]. But we have one problem left: when both Alice and Bob compute Φ_A and Φ_B, how do they get a shared secret from their codomain X_{AB}? We represent such a four-dimensional abelian variety by its *theta null point*

$$t = (a_0 : \ldots : a_{15}) \in \mathbb{P}^{15}(\mathbb{F}_p).$$

We can view t as an analogue to the Montgomery A-coefficient for representing elliptic curves. But because of the way Alice and Bob compute these isogenies, they are extremely likely to end up on different theta null points representing the same variety X_{AB}.

The analogue for elliptic curves. An analogue from dimension 1 may help to illustrate the point. Take the curves $E_5 : y^2 = x^3 + 5x^2 + x$ and $E_{12} : y^2 = x^3 + 12x^2 + x$ over $\mathbb{F}_p$ with $p = 17$. These curves are *isomorphic*. How does this help us to derive the same shared secret? In dimension 1, we have a powerful friend, the j-invariant:

$$j(E_A) = \frac{256(A^2 - 3)^3}{A^2 - 4}.$$

The j-invariant is usually used to show that two curves are isomorphic, such as the two curves above which both have j-invariant $j(E_5) = j(E_{12}) = 10$, but can therefore also be used as a shared secret derived from the two curves.

For the fourfolds in MIKE. A slight inconvenience is that such an invariant, like the j-invariant, cannot exist in dimension four, as the space of abelian fourfolds is 10-dimensional, whereas a single invariant can only span one dimension. A second inconvenience is that finding a *set of invariants* $j_1, \ldots, j_{10}$ to establish if two theta null points $t, t' \in \mathbb{P}^{15}(\mathbb{F}_p)$ of abelian fourfolds are isomorphic, is an unsolved open problem.

We are in luck, however: the abelian fourfolds X_{AB} that we end up on in MIKE are not just any random abelian fourfold, but lie in a very specific (four-dimensional) subspace of abelian fourfolds.[13] One of the big contributions in MIKE [BDD+26] is the derivation of such invariants for those specific MIKE theta null points. We call these the MIKE invariants, and they are readily computable from a given theta null point.

Definition 2. Let $t = (a_0 : \ldots : a_{15}) \in \mathbb{P}^{15}(\mathbb{F}_p)$ be a theta null point of a MIKE fourfold X_{AB} as computed by the MIKE shared secret computation. Then the *relative MIKE invariants* $I_0, \ldots, I_9$ are computed as

$$\begin{aligned}
I_0 &= (a_0 + a_6 + a_9 + a_{15})^2 & \qquad I_5 &= 4(a_2 - a_{11})^2,\\
I_1 &= (a_0 - a_6 + a_9 - a_{15})^2 & \qquad I_6 &= 4(a_1 + a_7)^2,\\
I_2 &= (a_0 + a_6 - a_9 - a_{15})^2 & \qquad I_7 &= 4(a_1 - a_7)^2,\\
I_3 &= (a_0 - a_6 - a_9 + a_{15})^2 & \qquad I_8 &= 4(a_3 + a_5)^2,\\
I_4 &= 4(a_2 + a_{11})^2, & \qquad I_9 &= 4(a_3 - a_5)^2.
\end{aligned}$$

Definition 3. Given $I_0, \ldots, I_9$, we define the *normalised MIKE invariants* as $I'_k := I_k / I_0$ for $1 \leq k \leq 9$, and the *absolute MIKE invariants* as

$$J_2 := \sum_{k=1}^{9} {I'_k}^2, \quad J_3 := \sum_{k=1}^{9} {I'_k}^3, \quad J_4 := \sum_{k=1}^{9} {I'_k}^4, \quad J_5 := \sum_{k=1}^{9} {I'_k}^5.$$

We write $J_{\mathsf{MIKE}}(t) = (J_2(t), J_3(t), J_4(t), J_5(t)) \in \mathbb{F}_p^4$ to describe the computation of these invariants from $t = (a_0 : \ldots : a_{15})$. One can compute that $J_1 := \sum I'_k = 3$, so we drop it.

[13] This is due to the specific construction we use in MIKE, and not inherent to the module action framework in general.

With this, we may informally rephrase the key Theorem 4.3 from MIKE [BDD+26], to finally derive the shared secret between Alice and Bob!

Theorem 4. *Let $t = (a_0 : \ldots : a_{15})$ be the theta null point Alice computes from Φ_A, representing X_{AB}. Let $t' = (b_0 : \ldots : b_{15})$ be the theta null point Bob computes from Φ_B, representing X_{AB}. If both computations are done correctly, then*

$$J_{\mathsf{MIKE}}(t) = J_{\mathsf{MIKE}}(t') \in \mathbb{F}_p^4.$$

Conversely, under a mild non-vanishing assumption, the equality $J(t) = J(t')$ implies that t and t' represent the same fourfold. Hence, Alice and Bob may use this value as their shared secret.

Summary. In short, key generation is essentially just a one-dimensional 2^e-isogeny $\varphi_A : E_0' \to E_A$, generated by choosing a secret $x_A \in \mathbb{Z}_{2^e}$ and deriving $K_A = P_0 + [x_A]Q_0$ as the kernel generator $\ker \varphi_A = \langle K_A \rangle$. From Bob's public key E_B, Alice first derives the Weil restriction $W(E_B)$ using Section 2.2, and then uses her secret x_A to derive the kernel of $\Phi_A : W(E_B)^2 \to X_{AB}$ using Theorem 3. If Bob does the same, both Alice and Bob get theta null points $t_A, t_B \in \mathbb{P}^{15}(\mathbb{F}_p)$ as representations of X_{AB}. To ensure they derive the same shared secret, they compute the J-invariants using Definition 2 and Definition 3. Then, Theorem 4 will ensure equality of $J(t_A)$ and $J(t_B)$, as long as both performed their computations correctly.

We stress that the secret x_A is only involved in deriving the kernel generators K_A in dimension 1, and $K_1, \ldots, K_4$ in dimension 4, and not in the actual isogeny computations. This makes it much easier to implement MIKE in constant time, as explained also by Pope [Pop26].

Level 3: MIKE as a module action

So far we have presented two diagrams to explain MIKE: In Level 1, we described MIKE as a module action NIKE in Fig. 2b, and in Level 2, we described MIKE in terms of isogenies in Fig. 3. The goal of Level 3 is to unify these diagrams using the module action framework: why do the isogenies φ_A, φ_B, resp. Φ_A, Φ_B correspond to the module actions of M_A and M_B on E_0', resp. X_B, X_A?

To achieve this, we first look at a simpler example in Section 3.1: why (quadratic) ideals correspond to isogenies in CSIDH. This will give us the right intuition to understand why modules, as a generalisation of ideals, correspond to higher-dimensional isogenies in MIKE, which we explore in Section 3.2. However, this is still not enough to understand how the isogenies φ_A and φ_B relate to modules M_A and M_B, as those isogenies are significantly different from the isogenies we derive from ideals in CSIDH. We derive M_A and M_B from φ_A and φ_B in Section 3.3, and with this final bit of understanding, we will be able to derive Φ_A and Φ_B from M_A and M_B in Section 3.4 in terms of the description of their kernel as in Theorem 3. As a bonus, we then perform a concrete example of a MIKE key exchange over $\mathbb{F}_p$ for $p = 127$ in Section 3.5.

⚠ Warning

This level mostly assumes we are working with *supersingular* elliptic curves. For now, it is okay to think of these as curves over $\mathbb{F}_p$ with $p+1$ points, or as curves over $\mathbb{F}_{p^2}$ with $(p+1)^2$ points, the latter being *maximal* supersingular curves.

3.1 An isogeny from a quadratic ideal (CSIDH)

To understand how ideals give us isogenies, we first need to know about endomorphisms.

Definition 4. An endomorphism $\theta : E \to E$ is an isogeny from E to itself.

We've already seen two examples of endomorphisms in Level 2: the scalar multiplication map $[n] : E \to E$ for any $n \in \mathbb{Z}$, and Frobenius $\pi : (x, y) \mapsto (x^p, y^p)$ for curves over $\mathbb{F}_p$. There is a curious relation between π and such scalar multiplications. Take for example the starting curve from MIKE, $E_0' : y^2 = x^3 - x$ over $\mathbb{F}_p$, then $\pi^2(P) = [-p]P$ for every $P \in E_0'(\overline{\mathbb{F}}_p)$, so we can think of π as $\sqrt{-p}$. This can be made precise, and in this way, we can think of endomorphisms of elliptic curves as behaving like numbers: any combination $[a] + [b]\pi$ for $a, b \in \mathbb{Z}$ is an endomorphism of E_0' now, and we can identify this with $a + b\sqrt{-p} \in \mathbb{Z}[\sqrt{-p}]$, a ring of quadratic integers.

Remark 3. From now on, we will use the notation $[\cdot]$ to move between the world of integers and endomorphisms, for example $n \in \mathbb{Z}$ corresponds to $[n] : E \to E$, and similarly $\sqrt{-p}$ corresponds to $[\sqrt{-p}] = \pi : (x, y) \mapsto (x^p, y^p)$ on E_0 or E_0', and so $\gamma = a + b\sqrt{-p} \in \mathbb{Z}[\sqrt{-p}]$ acts as $[\gamma](P) = [a]P + [b]\pi(P)$. This notation hides one significant aspect that we should be aware of: there is an implicit *choice* made by choosing π to represent $\sqrt{-p}$, and it is possible that a curve has multiple valid endomorphisms $[\alpha]$ that could represent the same quadratic integer α. More information on this can be found in [CLM+18; CK20; DFK+23].

Quadratic ideals $\mathfrak{a}$ in this ring $R = \mathbb{Z}[\sqrt{-p}]$ now correspond to "ideals of endomorphisms" on curves defined over $\mathbb{F}_p$, such as E_0 or E_0'. We will show how these correspond to isogenies $\varphi_\mathfrak{a} : E_0 \to E$ with a kernel $E_0[\mathfrak{a}]$. The scheme CSIDH [CLM+18] is the easiest example using this phenomenon, based on the same idea from the older CRS scheme [Cou06; RS06; Sto10]. We first give an example on E_0 à la CSIDH, before we give a proper definition for general curves E.

Example 9.

Take the ring $R = \mathbb{Z}[\sqrt{-p}]$ and the curve $E_0 : y^2 = x^3 + x$ over the field $\mathbb{F}_p$ with $p = 419$.

Associate π to $\sqrt{-p}$. Take the ideal $\mathfrak{a} = (5, \sqrt{-p}-1) \subseteq R$. The corresponding generators are then $[5]$ and $\pi - 1$ on E_0, and we can associate a subgroup $E_0[\mathfrak{a}] \subseteq E_0$ to $\mathfrak{a}$ by intersecting the kernels $E_0[5]$ and $E_0[\pi - 1]$ of these generators

$$E_0[\mathfrak{a}] = E_0[5] \cap E_0[\pi - 1].$$

As we have seen in Example 2, the subgroup $E_0[\pi-1]$ consists of those points $P \in E_0(\overline{\mathbb{F}}_p)$ such that $(\pi - 1)(P) = \mathcal{O}$, that is, $\pi(P) = P$, that is, $P \in E_0(\mathbb{F}_p)$. On the other hand, from Example 3, we know that $E_0[5]$ consists of those points such that $[5]P = \mathcal{O}$. So, we get for their intersection

$$\begin{aligned} E_0[\mathfrak{a}] &= \{P \in E_0(\mathbb{F}_p) \mid [5]P = \mathcal{O}\} \\ &= \{\mathcal{O}, (349, 133), (185, 346), (185, 73), (349, 286)\} \\ &= \langle K \rangle \qquad \text{for } K = (349, 133) \quad \text{(or choose any other non-zero} K \in E_0[\mathfrak{a}]) \end{aligned}$$

By the General Mantra, there is an isogeny $\varphi_{\mathfrak{a}}$ with kernel $E_0[\mathfrak{a}] = \langle K \rangle$. Entirely coincidentally, this is exactly the example we computed with Vélu's formulas in Example 7 as the map $\varphi_{\mathfrak{a}} : E_0 \to E_{199}$ that sends $P = (x, y)$ to

$$\left(\frac{x \cdot (-70x - 1)^2 \cdot (185x - 1)^2}{(x + 70)^2 (x - 185)^2}, \; -y \cdot \frac{179x^6 - 162x^5 - 107x^4 + 95x^3 - 104x^2 + 54x + 155}{(x + 70)^3 (x - 185)^3} \right).$$

To reflect the action of the ideal $\mathfrak{a}$, we write $E_{199} = [\mathfrak{a}] * E_0$.

The general definition of $E[\mathfrak{a}]$ and the associated isogeny $\varphi_{\mathfrak{a}}$ are now not so difficult. As we are extending the General Mantra, we capture this in another mantra.

Theorem 5 (Another Mantra)**.** *Let $R = \mathbb{Z}[\alpha]$ be a ring, where α is a quadratic integer. Let E be an elliptic curve together with a corresponding endomorphism $[\alpha] : E \to E$. Then, given a nonzero quadratic ideal $\mathfrak{a} \subseteq R$ with norm[14] coprime to p, whose elements $\theta \in \mathfrak{a}$ are of the form $\theta = a + b\alpha$, we get a unique subgroup $E[\mathfrak{a}]$ by*

$$E[\mathfrak{a}] = \bigcap_{\theta \in \mathfrak{a}} \ker[\theta].$$

Given $E[\mathfrak{a}]$, we get a unique isogeny $\varphi_{\mathfrak{a}} : E \to E'$ with $\ker \varphi_{\mathfrak{a}} = E[\mathfrak{a}]$, where we write $E' = [\mathfrak{a}] * E$.

Remark 4. It is very likely that there are multiple ideals $\mathfrak{a}, \mathfrak{b} \subseteq R$ that induce the same codomain $\mathfrak{a} * E = \mathfrak{b} * E$. In fact, if we take just a single endomorphism α of E, then the ideal $\mathfrak{a} = \langle \alpha \rangle$ has $E[\mathfrak{a}] = \ker \alpha$, and so just induces itself as the isogeny $\varphi_{\mathfrak{a}} = \alpha$. Such ideals generated by a single element are called *principal* ideals, and always correspond to endomorphisms. Therefore, the clean and precise resolution to this is to quotient out by principal ideals, which gives the *class group* $\mathcal{Cl}(R)$. The precise details are beyond the scope of this text, but a beautiful topic in itself as the class group acts faithfully on specific sets of elliptic curves. This is the crucial construction behind the NIKEs proposed by Couveignes [Cou06], Rostovtsev and Stolbunov [RS06; Sto10], and CSIDH [CLM+18]. In short, the 'proper' group action is that of ideal classes $[\mathfrak{a}] \in \mathcal{Cl}(R)$, which is how we will write this action from now on, but we can still compute this action using Theorem 5.

[14] If you do not know what the norm of such an ideal is, ignore this comment: we need this technicality so that the resulting isogeny is separable, which is what we want. The norm $N(\mathfrak{a})$ for an invertible ideal $\mathfrak{a}$ can be computed as the gcd of the norms $N(\beta) = \beta \cdot \bar{\beta}$ of the elements $\beta \in \mathfrak{a}$.

Key properties of the class group action

Let us take a closer look at the curve $E_0 : y^2 = x^3 + x$ over $\mathbb{F}_p$ before we look at E_0'. What curves can we reach when we act with $[\mathfrak{a}] \in \mathcal{Cl}(R)$? That is, what does the orbit

$$\{\ [\mathfrak{a}] * E_0 \mid [\mathfrak{a}] \in \mathcal{Cl}(R)\}$$

look like? The key element is the group structure of E_0 over $\mathbb{F}_p$,

$$E_0(\mathbb{F}_p) \cong \mathbb{Z}_{p+1}.$$

If $E_A : y^2 = x^3 + Ax^2 + x$ over $\mathbb{F}_p$ is any other curve with $E_A(\mathbb{F}_p) \cong \mathbb{Z}_{p+1}$, then there is an ideal $\mathfrak{a}$ such that $E_A = [\mathfrak{a}] * E_0$, and vice versa, as $E_0 = [\mathfrak{a}]^{-1} * E_A$, as shown in [CLM+18]. The crucial connection here is the ring $R = \mathbb{Z}[\sqrt{-p}]$: both curves have the Frobenius endomorphism $\pi : (x, y) \mapsto (x^p, y^p)$ that acts as $\sqrt{-p}$, and so we can say that both are *oriented* by R. But more importantly, there is no bigger ring $R' \supseteq R$ by which these curves are oriented, so we may say these are *primitively oriented* by R. The curve E_0' is now a crucial example, as E_0' is also oriented by R, but not primitively! Crucially, this curve E_0' has a different group structure[15]

$$E_0'(\mathbb{F}_p) \cong \mathbb{Z}_{\frac{p+1}{2}} \times \mathbb{Z}_2.$$

Example 10.

We look at the supersingular elliptic curve $E_0' : y^2 = x^3 - x$ over $\mathbb{F}_p$ with rational group structure $E_0'(\mathbb{F}_p) \cong \mathbb{Z}_{\frac{p+1}{2}} \times \mathbb{Z}_2$. This can be seen because $x^3 - x$ has three roots $0, 1, -1$ over $\mathbb{F}_p$, and so E_0' has three points of order 2, in contrast to E_0. This implies that $E_0'[2]$ is fully $\mathbb{F}_p$-rational, in contrast to $E_0[2]$.
We can again associate $\sqrt{-p}$ to $\pi : (x, y) \mapsto (x^p, y^p)$ on this curve, so that E_0' is oriented by $R = \mathbb{Z}[\sqrt{-p}]$. However, because E_0' has rational 2-torsion, there is a bigger ring $R' \supseteq R$ that orients E_0', so the orientation by R is not *primitive*. This bigger ring is

$$R' = \mathbb{Z}\left[\tfrac{1+\sqrt{-p}}{2}\right],$$

which we may associate to the endomorphism $\omega = \frac{1+\pi}{2}$ on E_0'.
To see how the fully rational 2-torsion is connected to the endomorphism $\frac{1+\pi}{2}$, we first observe that $(\pi - 1)P = \mathcal{O}$ for all $P \in E_0'[2]$. But since $E_0'[2]$ is also the kernel of the doubling map, we say that $\pi - 1$ *factors through* $[2]$, i.e., we can write $\pi - 1 = \alpha \circ [2]$ for some endomorphism α. Thus, we have that $\alpha = \frac{\pi-1}{2}$ is an endomorphism on E_0'. This looks very similar to what we wanted to achieve, and for the final step we notice that the action of $+1$ and -1 is identical on the 2-torsion, since $P = -P$ for $P \in E_0'[2]$. Therefore, we also get $(\pi + 1)P = \mathcal{O}$ for $P \in E_0'[2]$, and following the same argument as above, we get that $\frac{1+\pi}{2}$ is an endomorphism on E_0'.

With this example, we essentially grasp the full picture of the class group actions for curves over $\mathbb{F}_p$.

1. We either have curves like E_0, with group structure $\mathbb{Z}_{p+1}$ and so no fully rational 2-torsion. Such curves are primitively oriented by $R = \mathbb{Z}[\sqrt{-p}]$, and between two such curves E and E' there is an ideal class $[\mathfrak{a}] \in \mathcal{Cl}(R)$ such that $E' = [\mathfrak{a}] * E$. We say that these curves are on the *floor*.

[15] There are only two possible group structures for supersingular curves over $\mathbb{F}_p$, either $\mathbb{Z}_{p+1}$ such as E_0 or $\mathbb{Z}_{\frac{p+1}{2}} \times \mathbb{Z}_2$ such as E_0'.

2. Or, we have curves like E_0', with group structure $\mathbb{Z}_{\frac{p+1}{2}} \times \mathbb{Z}_2$ and thus fully rational 2-torsion. Such curves are primitively oriented by $R' = \mathbb{Z}[\frac{1+\sqrt{-p}}{2}]$, and between two such curves E and E' there is an ideal class $[\mathfrak{a}] \in \mathcal{Cl}(R')$ such that $E' = [\mathfrak{a}] * E$. We say that these curves live on the *surface*.

In summary, the class group $\mathcal{Cl}(R)$ is a beautiful abelian group which acts on R-oriented curves. This group action makes it easy to instantiate a NIKE, such as CSIDH in Fig. 2a, but also leads to Kuperberg's quantum attack, as stated before. In general, it is good to keep in mind that we want something *like* CSIDH, but generalized in such a way that it is no longer a *group* action.

The vertical ideal

There is a last bit of confusion we need to clear up before we go to the module action. Perhaps, you have wondered about the ideal $\mathfrak{f} = (2, \sqrt{-p} - 1) \subseteq R$ and its action on E_0. Let us compute this action in detail.

Example 11.
Let $E_0 : y^2 = x^3 + x$ over $\mathbb{F}_p$ and let $\mathfrak{f} = (2, \sqrt{-p} - 1) \subseteq R$. Then

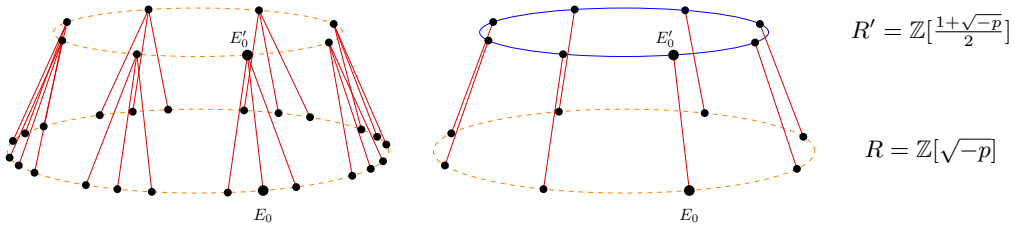
$$\begin{aligned} E_0[\mathfrak{f}] &= E_0[2] \cap E_0[\pi - 1] \\ &= E_0[2] \cap E_0(\mathbb{F}_p) \\ &= \{(0,0), \mathcal{O}\}. \end{aligned}$$

Thus, the associated isogeny $\varphi_\mathfrak{f}$ is the isogeny with kernel $\langle (0,0) \rangle$. It turns out that this is exactly the isogeny $\varphi_\mathfrak{f} : E_0 \to E_0'$, where $E_0' : y^2 = x^3 - x$ is the curve on the surface that we described before, and so we may write $E_0' = \mathfrak{f} * E_0$. However, this action is outside of the class group action we described before, because $\mathfrak{f}$ has no class in $\mathcal{Cl}(R)$. The crucial problem is that $\mathfrak{f}$ is not *invertible*: there is no ideal $\mathfrak{f}'$ such that $\mathfrak{f} \cdot \mathfrak{f}'$ is equivalent to $(1) \in \mathcal{Cl}(R)$. And so, while we can go 'up' from the floor to the surface, as showcased by $E_0' = \mathfrak{f} * E_0$, there is no ideal action $\mathfrak{f}' * E_0'$ that would take us back down to E_0. More precisely, there *is* an isogeny that goes down from E_0' to E_0, namely the dual isogeny, but it is not represented by an ideal action.

We can now summarise the full picture in an *isogeny volcano*, as in Fig. 5, although the details differ depending on the congruence of p modulo 8. On the *surface*, we have E_0' and all other curves that are primitively oriented by R'. On the *floor*, we have all curves primitively oriented by R, such as E_0. The curves on the floor (resp. on the surface) are connected through the action of ideals $[\mathfrak{a}] \in \mathcal{Cl}(R)$ (resp. $[\mathfrak{a}] \in \mathcal{Cl}(R')$). Since we stay on the same level, we call the corresponding isogenies *horizontal*. An example of this is shown via the dashed orange lines, although these should not be taken at face value.[16] Only when $p \equiv 7 \bmod 8$, we have horizontal 2-isogenies on the surface, depicted by the blue edges in Fig. 5b.

The connection between floor and surface is given in red through the action of $\mathfrak{f}$, where we can only go up, as described above. Hence, we call the corresponding 2-isogenies *vertical*. An in-depth analysis of the situation is beyond the scope of this text, but it turns out that there is a 3:1 relationship between curves on the floor and the surface when $p \equiv 3 \bmod 8$: there are triplets of curves on the floor that are connected to the same curve on the surface. In contrast, we get a 1:1 relationship when $p \equiv 7 \bmod 8$. For setting up CSIDH, we can then either decide to work on the floor or on the surface.

[16] Depending on the choice of $[\mathfrak{a}]$, we would get different cycles, or even multiple smaller cycles; the illustration here is somewhat inaccurately simplified. We refer to [CD20] for more details.



(a) $p \equiv 3 \bmod 8$ (b) $p \equiv 7 \bmod 8$

Figure 5: Isogeny volcanoes for different choices of p with curves on the surface primitively oriented by $R' = \mathbb{Z}[\frac{1+\sqrt{-p}}{2}]$ and curves on the floor primitively oriented by $R = \mathbb{Z}[\sqrt{-p}]$. Red edges depict vertical 2-isogenies corresponding to the action of $\mathfrak{f}$. Blue edges represent horizontal 2-isogenies. Orange dashed edges symbolise horizontal isogenies corresponding to the action of some ideals $[\mathfrak{a}] \in \mathcal{C}\ell(R)$ resp. $[\mathfrak{a}] \in \mathcal{C}\ell(R')$.

3.2 An isogeny from a quadratic module (MIKE)

Congratulations, you are now almost ready to begin reading MIKE [BDD+26], as this section aims to describe what happens in Section 2 of MIKE [BDD+26].

The core idea is simple: in the previous section, we took a ring $R = \mathbb{Z}[\alpha]$ and an ideal $\mathfrak{a} \subseteq R$ to get a subgroup $E[\mathfrak{a}]$ and thus an isogeny $\varphi_\mathfrak{a} : E \to E'$, and E' can be seen as the (group) action of $\mathfrak{a}$ on E, so $E' = \mathfrak{a} * E$. If we now take a *module* $M \subseteq R^n$, we show how to define a subgroup $E^n[M]$, which gives us an n-dimensional isogeny $\Phi_M : E^n \to B$ with kernel $\ker \Phi_M = E^n[M]$, and B can be seen as the *module action* $B = M * E$. Note in particular that B is the result of M acting on E, not on E^n. MIKE will mainly use modules of rank $n = 2$.

Table 1: The main dictionary to keep in mind, given a ring $R = \mathbb{Z}[\alpha]$.

Scheme	Type of Action	Acting by	Kernel	Isogeny	Result
CSIDH	group action	ideals $\mathfrak{a} \subseteq R$	$E[\mathfrak{a}] \subseteq E$	$\varphi_\mathfrak{a} : E \to E'$	$E' = \mathfrak{a} * E$
MIKE	monoid action	modules $M \subseteq R^n$	$E^n[M] \subseteq E^n$	$\Phi_M : E^n \to B$	$B = M * E$

Defining the module kernel

In MIKE, we are acting with modules $M \subseteq R^n$ of rank $n = 2$. To start with, we will use $R = \mathbb{Z}[\sqrt{-p}]$, and we work with elliptic curves and varieties over $\mathbb{F}_p$ where we have the easy association $\pi = [\sqrt{-p}]$. In this case, we say the curve or variety is R-oriented, and we can then think of elements $m \in M \subseteq R^n$ as $m = (r_1, \ldots, r_n)$, where each $r_i \in R$ is of the form $r_i = a_i + b_i\sqrt{-p}$, with $a_i, b_i \in \mathbb{Z}$. For an element $r_i \in R$, we remember that the associated endomorphism $[r_i] = [a_i] + [b_i]\pi$ on a curve E is simply

$$[r_i] \ : E \to E, \qquad P \ \mapsto \ [a_i]P + [b_i]\pi(P),$$

and so it is easy to define for $m = (r_1, \ldots, r_n) \in M$ the map $[m] : E^n \to E$ by

$$[m] \ : \ (P_1, \ldots, P_n) \ \mapsto \ \sum [r_i]P_i, \qquad r_i \in R.$$

Notice that we never actually use that E is one-dimensional, i.e., an elliptic curve! As π is also an endomorphism on our abelian variety $A/\mathbb{F}_p$, we may similarly define a map $[m] : A^n \to A$. That is, given $M \subseteq R^n$, we can describe a kernel $A^n[M]$ as follows, where we note that the same definition holds for R' instead of R.

Definition 5. Let $R = \mathbb{Z}[\sqrt{-p}]$ and let $M \subseteq R^n$ be a module of rank n. Let A be an abelian variety over $\mathbb{F}_p$ that is R-oriented. Then the *module kernel* $A^n[M] \subseteq A^n$ is the subgroup given by

$$A^n[M] := \{\ P = (P_1, \ldots, P_n) \in A^n \mid\ [m]P = \mathcal{O}_A \text{ for all } m \in M\}.$$

One should carefully compare this to Another Mantra, to see how this vastly generalises the previous result. Luckily, defining the right isogeny, and therefore module action, is now simple.

Definition 6. Let $R = \mathbb{Z}[\sqrt{-p}]$ and let $M \subseteq R^n$ be a module of rank n. Let A be an abelian variety over $\mathbb{F}_p$ that is R-oriented. Then M defines a subgroup $A^n[M]$, and there exists an isogeny $\Phi_M : A^n \to B$ with $\ker \Phi_M = A^n[M]$. We define the abelian variety $B := M * A$ as the codomain of Φ_M.

Let's do two simple non-trivial examples and one trivial example, using $R = \mathbb{Z}[\sqrt{-p}]$ and M of rank $n = 2$.

Example 12.
Let $R = \mathbb{Z}[\sqrt{-p}]$, take $E : y^2 = x^3 + 199 \cdot x^2 + x$ over $\mathbb{F}_p$ with $p = 419$, and take M the rank-2 module generated by $m_1 = (2, 0)$ and $m_2 = (\sqrt{-p}, -1)$. Then $E^2[M]$ are those points (P, Q) such that

$$[m_1](P, Q) = [2]P = \mathcal{O}, \qquad [m_2](P, Q) = \pi(P) - Q = \mathcal{O}.$$

From the first equation we get $P \in E[2]$, and from the second equation $Q = \pi(P)$, so

$$E^2[M] = \{\ (P, \pi(P)) \in E^2 \mid P \in E[2]\ \}.$$

The corresponding 2-dimensional isogeny $\Phi_M : E^2 \to B$ can be computed using Richelot isogenies [CDS20; FT19], using standard formulas from the literature [HLP00, Prop. 4], giving us that $B = M * E$ is the Jacobian of the hyperelliptic curve H of genus 2 defined by

$$H : y^2 = 312x^6 + 58x^4 + 58x^2 + 312.$$

Example 13.
In contrast, let us look at the action of the module M' generated by $m_1' = (1, 1)$ and $m_2' = (1, -1)$. Then $E^2[M']$ are those points (P, Q) such that

$$[m_1'](P, Q) = P + Q = \mathcal{O}, \qquad [m_2'](P, Q) = P - Q = \mathcal{O}.$$

These equations now give us that the kernel would be

$$E^2[M'] = \{\ (P, P) \in E^2 \mid P \in E[2]\ \},$$

and the corresponding 2-dimensional isogeny $\Phi_{M'} : E^2 \to B'$ is much more boring: the isogeny $\Phi_{M'}$ is simply the endomorphism $(P, Q) \mapsto (P + Q, P - Q)$ on E^2. So $B' = M' * E$ is simply $B' = E^2$.

Example 14.
Now let $M = R^2$, which is generated by $m_1 = (1, 0)$ and $m_2 = (0, 1)$, and act on any E over $\mathbb{F}_p$. So we need $[m_1](P, Q) = P = \mathcal{O}$ and $[m_2](P, Q) = Q = \mathcal{O}$, which implies

$$E^2[M] = \{(\mathcal{O}, \mathcal{O})\}.$$

But then the associated isogeny is just the identity $\Phi_M : E^2 \to E^2$, and we get $E^2 = R^2 * E$.

Key properties of the module action

Similar to the situation in Section 3.1, it is interesting to wonder what types of abelian varieties we reach depending on our starting curve. To simplify this a bit, we consider the action of rank-2 modules. Recall that in Section 3.1 for the class group action of different rings R, R', the main difference was the group structure of the starting curve. This will be very similar for the module action.

⚠ Warning

We implicitly assume we are working with projective modules $M \subseteq R^n$ for now, even though we haven't explained what these are. For us, they will be the right analogue to $[\mathfrak{a}] \in \mathcal{Cl}(R)$ in the previous section, and we go into more detail in Level 3½.

Case 1. Our usual starting curve is $E_0 : y^2 = x^3 + x$ over $\mathbb{F}_p$, and we recall that this curve is primitively oriented by $R = \mathbb{Z}[\sqrt{-p}]$. We already know that the objects we will reach by the action of a module of rank 2 are abelian surfaces A, and for these to be primitively oriented by R, we need such an A to be defined over $\mathbb{F}_p$.

Now, similar to the class group action case, when we act with M on E_0 we reach an abelian surface A over $\mathbb{F}_p$ with group structure

$$A(\mathbb{F}_p) \cong \mathbb{Z}_{p+1} \times \mathbb{Z}_{p+1},$$

and vice versa, we can reach all such A by a module M [Rob24; CLR26]. Note that this is also the group structure of $(E_0 \times E_0)(\mathbb{F}_p)$, which is reached by the trivial action of $M = R^2$ on E_0.

Example 15.

We have seen such abelian surfaces already! In Section 2.2, we took a curve $E_A : y^2 = x^3 + Ax^2 + x$ over $\mathbb{F}_{p^2}$ with $A \in \mathbb{F}_{p^2} \setminus \mathbb{F}_p$, and computed the Weil restriction $W(E_A)$ as an abelian surface over $\mathbb{F}_p$, which has the rational group structure

$$W(E_A)(\mathbb{F}_p) \cong \mathbb{Z}_{p+1} \times \mathbb{Z}_{p+1}.$$

We can associate $\sqrt{-p}$ to the Frobenius map π_W on $W(E_A)$, so that $W(E_A)$ is R-oriented. Thus, for Weil restrictions $W(E_A)$ of such curves E_A, we can always find a module $M_A \subseteq R^2$ such that $W(E_A) = M_A * E_0$. Similarly, in Example 12, we computed the action of such an M on E_{199} over $\mathbb{F}_p$ with $p = 419$ and got the Jacobian J of the hyperelliptic curve H of genus 2 defined by

$$H : y^2 = 312x^6 + 58x^4 + 58x^2 + 312,$$

and one can show that this Jacobian J has rational group structure

$$J(\mathbb{F}_p) \cong \mathbb{Z}_{p+1} \times \mathbb{Z}_{p+1}.$$

Case 2. We have also already seen the other curve $E_0' : y^2 = x^3 - x$ over $\mathbb{F}_p$ as an example of a curve primitively oriented by $R' = \mathbb{Z}[\frac{1+\pi}{2}]$ because of its group structure

$$E_0'(\mathbb{F}_p) \cong \mathbb{Z}_{\frac{p+1}{2}} \times \mathbb{Z}_2.$$

Acting by rank-2 modules $M \subseteq R'^2$, the resulting abelian surface $A = M * E_0'$ will then be an abelian surface over $\mathbb{F}_p$ primitively oriented by R' with group structure

$$A(\mathbb{F}_p) \cong \mathbb{Z}_{\frac{p+1}{2}} \times \mathbb{Z}_{\frac{p+1}{2}} \times \mathbb{Z}_2 \times \mathbb{Z}_2,$$

that is, the group structure of $(E_0' \times E_0')(\mathbb{F}_p)$.

Example 16.
Take $p = 127 = 2^7 - 1$, write $\mathbb{F}_{p^2} = \mathbb{F}_p(i)$, and let's try to compute a 2-dimensional $(16, 16)$-isogeny $E_0'^2 \to J$ derived from a module action $M * E_0'$. This is a bit scary!
We take the module $M \subseteq R'^2$ generated by $m_1 = (16, 0)$, $m_2 = (0, 16)$, $m_3 = (-1 + \frac{1+\sqrt{-p}}{2}, 3 - 3\frac{1+\sqrt{-p}}{2})$, and $m_4 = (-3\frac{1+\sqrt{-p}}{2}, -\frac{1+\sqrt{-p}}{2})$. Although somewhat of a long computation, one can show that $K_1 = ((84 + 97i, 40 + 4i), (43 + 30i, 123 + 40i))$ and $K_2 = ((18, 90), (72, 47))$ are the two 16-torsion points that are kernel generators

$$E_0'^2[M] = \langle K_1, K_2 \rangle.$$

We will compute this $(16, 16)$-isogeny in four steps of $(2, 2)$-isogenies.[17]

Step 1. The first kernel is generated by $[8]K_1 = ((0, 0), (0, 0))$ and $[8]K_2 = ((1, 0), (1, 0))$. By closer inspection, this means that the kernel is (T, T) for $T \in E_0'[2]$. This gives us a well known $(2, 2)$-isogeny

$$\Phi_1 : E_0'^2 \to E_0'^2, \qquad (P, Q) \mapsto (P + Q, P - Q),$$

which just ends up back at $E_0'^2$. This is a bit strange, but not problematic.[18]

Step 2. If we push our kernel points K_1 and K_2 through Φ_1, we get 8-torsion points $K_1' = ((70 + 101i, 4 + 50i), (57 + 26i, 50 + 123i))$ and $K_2' = ((99, 38), (17, 18))$. The second $(2, 2)$-isogeny has kernel generators $[4]K_1' = ((1, 0), (-1, 0))$ and $[4]K_2' = ((1, 0), \mathcal{O})$. The $\mathcal{O}$ showing up shows us something funny is happening, in particular if we replace $[4]K_1'$ by $[4]K_1' + [4]K_2' = (\mathcal{O}, (-1, 0))$. Φ_2 is a *diagonal* isogeny, which essentially means that $\Phi_2 = \varphi_1 \times \varphi_2$, where the φ_i are two 1-dimensional 2-isogenies, performed at the same time. Here φ_1 has kernel generator $(1, 0)$, from $[4]K_2'$'s first argument, and φ_2 has kernel generator $(-1, 0)$ from the other generator's second argument. These we can easily compute with the given Vélu formulas from Eq. (1) as

$$\varphi_1 : E_0' \to E_{121}, \qquad \varphi_2 : E_0' \to E_6, \qquad \text{hence } \Phi_2 : E_0'^2 \to E_{121} \times E_6,$$

where we picked the right isomorphisms $E_0' \xrightarrow{\sim} E_M$ to compute these steps with the Montgomery formulas from Level 2.[19]

Step 3. Again, we push our previous kernel points K_1' and K_2' through Φ_2 to get the 4-torsion points $K_1'' = ((98+120i, 56+22i), (29+7i, 22+71i))$ and $K_2'' = ((44, 98), (1, 95))$, and we get the kernel generators $[2]K_1'' = ((98, 0), (29, 0))$ and $[2]K_2'' = ((35, 0), (0, 0))$ of our third $(2, 2)$-isogeny Φ_3. This is a *gluing* isogeny, a 2D-isogeny that takes a product to the Jacobian of a hyperelliptic curve C over $\mathbb{F}_p$. As before, we refer to [CDS20; CD23] for formulas to compute these, given $[2]K_1''$ and $[2]K_2''$. We get

$$\Phi_3 : E_{121} \times E_6 \to J_{C_3}, \qquad \text{where } C_3 : y^2 = 3x^6 + 5x^4 + 115x^2 + 52.$$

Step 4. Things now become a bit tricky if you have never worked with Jacobians. Pushing the kernel generators through Φ_3 gives us the kernel of the final Φ_4 as generated by the 2-torsion elements $K_1''' = \langle x^2 + 5x + 23, 0 \rangle$ and $K_2''' = \langle x^2 + 80x + 68, 0 \rangle$, using Mumford's representation. If you are aware of Richelot isogenies [CDS20; FT19], one can compute Φ_4 as the quadratic splitting of C_3 associated to $\langle K_1''', K_2''' \rangle$ given by

$$C_3 : y^2 = (x^2 + 5x + 23) \cdot (x^2 + 80x + 68) \cdot (3x^2 - x + 14),$$

which gives us $\Phi_4 : J_{C_3} \to J_{C_4}$ where C_4 is the hyperelliptic curve

$$C_4 : y^2 = 112x^6 + 67x^5 + 110x^4 + 53x^3 + 117x^2 + 10x + 121.$$

Altogether, this implies that $J_{C_4} = M * E_0'$. One can check in Magma that indeed

$$J_{C_4}(\mathbb{F}_p) \cong \mathbb{Z}_{64}^2 \times \mathbb{Z}_2^2 = \mathbb{Z}_{\frac{p+1}{2}}^2 \times \mathbb{Z}_2^2,$$

which one can also notice from the fact that the right-hand side of C_4 factors into linear terms over $\mathbb{F}_p$.

Case 3. Given cases 1 and 2, we are clearly missing the case of abelian surfaces A over $\mathbb{F}_p$ with

$$A(\mathbb{F}_p) \cong \mathbb{Z}_{p+1} \times \mathbb{Z}_{\frac{p+1}{2}} \times \mathbb{Z}_2.$$

Such surfaces exist, and are interesting to analyse from a module action perspective, but out of scope for this work as they are not relevant for an explanation of MIKE.

Again a vertical ideal

Similar to the vertical ideal $\mathfrak{f} \subseteq R$ that brought us upwards from E_0 to E_0', a similar situation happens more generally, taking us up from case 1 abelian surfaces to case 2 abelian surfaces. This is most clearly demonstrated using the following example.

Example 17.
Take a curve $E_A : y^2 = x^3 + Ax^2 + x$ with $j(E_A) \notin \mathbb{F}_p$ and $E_A[2] \subseteq E_A(\mathbb{F}_{p^2})$, represent its Weil restriction $W(E_A)$ as $E_A \times E_A^{(p)}$, where we should think of $\mathbb{F}_p$-rational points in $W(E_A)(\mathbb{F}_p)$ as the pairs $(P, \pi(P)) \in E_A \times E_A^{(p)}$ for $P \in E_A(\mathbb{F}_{p^2})$. We can define the kernel

$$G = \{(P, \pi(P)) \in W(E_A)(\mathbb{F}_p) \mid P \in E_A[2]\}.$$

The associated isogeny $\Phi_G : W(E_A) \to J_A$ has as codomain a Jacobian J_A over $\mathbb{F}_p$ whose rational group structure is

$$J_A(\mathbb{F}_p) \cong \mathbb{Z}_{\frac{p+1}{2}} \times \mathbb{Z}_{\frac{p+1}{2}} \times \mathbb{Z}_2 \times \mathbb{Z}_2,$$

in contrast to the group structure $W(E_A)(\mathbb{F}_p) \cong \mathbb{Z}_{p+1} \times \mathbb{Z}_{p+1}$. This vertical action is called *Scholten's construction*, and takes the R-oriented surface $W(E_A)$ to an R'-oriented surface J_A. By extending the analogy of E_0 to E_0' and R to R', we will denote J_A by $W'(E_A)$ in contrast to $W(E_A)$. This analogy goes further: note that the kernel G can be described as the $\mathbb{F}_p$-rational 2-torsion points on $W(E_A)$, that is,

$$\begin{aligned} G &= W(E_A)(\mathbb{F}_p) \cap W(E_A)[2] \\ &= W(E_A)[\pi_W - 1] \cap W(E_A)[2] \\ &= W(E_A)[\mathfrak{f}] \end{aligned}$$

for the vertical ideal $\mathfrak{f} = (2, \sqrt{-p} - 1) \subseteq R$.

[17] Those not familiar with $(2,2)$-isogenies will have to take some steps for granted, or read up on the literature [CDS20; FT19; CD23]. These steps will not matter too much for your understanding of MIKE, and can be taken as a black box, as long as you understand that the final codomain is an abelian surface over $\mathbb{F}_p$.

[18] With more isogeny experience, you will see that this must be the first step by how the kernel is chosen. This can be avoided, but makes other technical steps in MIKE a bit harder.

[19] If you have experience with SIKE, this is the two-dimensional analogue of the E_6 curve, as $121 \equiv -6 \bmod p$.

The analogy is therefore complete: for curves, we go up by $E_0' = \mathfrak{f} * E_0$ and similarly for Weil restrictions, we go up by $W'(E_A) = \mathfrak{f} * W(E_A)$, using the same vertical ideal $\mathfrak{f} = (2, \sqrt{-p} - 1) \subseteq R$. Remember that $\mathfrak{f}$ is not invertible, and so we cannot go down by such an ideal action. In general, given a curve E_A with $j(E_A) \notin \mathbb{F}_p$, the hyperelliptic curve associated to Scholten's construction is easy to compute.[20]

Theorem 6. *Let $E_A : y^2 = x^3 + Ax^2 + x$ be a Montgomery curve over $\mathbb{F}_{p^2}$ with $j(E_A) \notin \mathbb{F}_p$ and $E_A[4] \subseteq E_A(\mathbb{F}_{p^2})$. Let $r_1 \in \mathbb{F}_{p^2}$ be a nonzero root of $x^3 + Ax^2 + x$, take $P \in E_A[4]$ such that $[2]P = (r_1, 0)$, and let $r_2 = x_P - r_1$. Then Scholten's construction $W'(E_A)$ is the Jacobian of the genus-2 hyperelliptic curve C over $\mathbb{F}_p$ defined by*

$$C : y^2 = x \cdot (x - 1) \cdot (x - \lambda) \cdot (x - \mu) \cdot (x - \nu),$$

where

$$\lambda = \frac{N(r_2) + N(r_1) + 1}{2}, \qquad \mu = \frac{N(r_2) - N(r_1) + 1}{2}, \qquad \nu = \frac{\lambda\mu}{N(r_2)},$$

with $N : a + bi \mapsto a^2 + b^2$ the norm map $\mathbb{F}_{p^2} \to \mathbb{F}_p$.

Example 18.

We take again $p = 127 = 2^7 - 1$, with $\mathbb{F}_{p^2} = \mathbb{F}_p(i)$. Let us try the above for $A = 45 + 96i$, that is

$$E_A : y^2 = x^3 + (45 + 96i)x^2 + x.$$

With Sage, it's easy to find $P = (1 + 48i, 95 + 32i)$ of order 4 with $[2]P = (54 + 123i, 0)$, where $r_1 = 54 + 123i$ is indeed a non-zero root of $x^3 + (45 + 96i)x^2 + x$. This makes $r_2 = x_P - r_1 = 1 + 48i - 54 - 123i = 74 + 52i$. For their norms, we get $N(r_1) = 54^2 + 123^2 = 11$ as well as $N(r_2) = 74^2 + 52^2 = 52$, and so

$$\lambda = \frac{52 + 11 + 1}{2} = 32, \qquad \mu = \frac{52 - 11 + 1}{2} = 21, \qquad \nu = \frac{32 \cdot 21}{52} = 52.$$

And so, $W'(E_A)$ can be represented as the Jacobian J_C of the hyperelliptic curve C over $\mathbb{F}_p$ given by

$$C : y^2 = x \cdot (x - 1) \cdot (x - 32) \cdot (x - 21) \cdot (x - 52),$$

which expands to $C : y^2 = x^5 + 21x^4 + 104x^3 + 109x^2 + 19x$. One can check in Sage or Magma, or by the fact that the right-hand side factors over $\mathbb{F}_p$, that the Jacobian J_C must have the group structure

$$J_C(\mathbb{F}_p) \cong \mathbb{Z}_{\frac{p+1}{2}} \times \mathbb{Z}_{\frac{p+1}{2}} \times \mathbb{Z}_2 \times \mathbb{Z}_2.$$

3.3 Key generation as a module action

Our goal is now to find the right module M_A such that we interpret Alice's keygen $\varphi_A : E_0' \to E_A$ as a module action $X_A := M_A * E_0'$. This gets us to the heart of much confusion:

- **The bad:** We could have either started key generation on the floor from the R-oriented curve E_0 to get $\varphi_A : E_0 \to E_A$ for some curve E_A. Then $W(E_A)$ has the same group structure as E_0^2, and so there should be a projective module $M_A \subseteq R^2$ such that $W(E_A) = M_A * E_0$.

[20] Theorem 6 is a result obtained at Brainstorm Days of The Isogeny Club at Eurocrypt 2026. So far, it remains unpublished.

- **The good:** Or, we could start on the surface with the R'-oriented curve E_0' to get $\varphi_A : E_0' \to E_A$ for some curve E_A. Then $W'(E_A)$ has the same group structure as $E_0'^2$, and so there should be a projective module $M_A \subseteq R'^2$ such that $W'(E_A) = M_A * E_0'$.

- **The ugly:** Strangely enough, it looks like MIKE is combining both. It starts with E_0' for key generation, but the shared secret computation starts on $W(E_B)^2$.

We delay the ugly for Section 3.4, and look in detail into the first two cases.

Notation. For the rest of this section, we heavily use $\frac{1+\sqrt{-p}}{2}$, and we therefore will use the notation $\omega = \frac{1+\sqrt{-p}}{2}$. Thus, we get $R' = \mathbb{Z}[\omega]$.

3.3.1 MIKE on the floor or on the surface

In both cases, key generation gives us an $\mathbb{F}_{p^2}$-rational isogeny $\varphi_A : E_0 \to E_A$, resp. $E_0' \to E_A$, as in Fig. 3, and we want to find an associated 2D-isogeny $\Phi : E_0^2 \to X_A$, resp. $E_0'^2 \to X_A$, derived from a module M_A so that $\ker \Phi = E_0^2[M_A]$. This would justify $X_A := M_A * E_0$, resp. $M_A * E_0'$, as in Fig. 2b. So what is the big difference?

Note first that, over $\mathbb{F}_{p^2}$, the curves E_0 and E_0' are isomorphic, and so $W'(E_0) \cong W'(E_0')$, which we sometimes abuse in notation. So for the 1-dimensional $\mathbb{F}_{p^2}$-isogeny $\varphi_A : E_0 \to E_A$, it does not matter too much if we start on E_0 or E_0'. But from the module perspective, the module acts on E_0 or E_0' as an $\mathbb{F}_p$-curve, where they differ significantly! To associate a module to this 1-dimensional isogeny φ_A, it helps that we can associate two different 2-dimensional isogenies to φ_A:

- Applying the Weil restriction, we get an isogeny $W(\varphi_A) : W(E_0) \to W(E_A)$

- Applying Scholten's construction, we get an isogeny $W'(\varphi_A) : W'(E_0') \to W'(E_A)$.

Confusingly, both $W(E_0)$ and $W'(E_0)$ are derived from E_0 as an $\mathbb{F}_{p^2}$-curve, and so $W(E_0) \cong W(E_0')$ and $W'(E_0) \cong W'(E_0')$, which justifies the switch from E_0 to $W'(E_0')$ above.[21] It turns out that when deriving the associated M_A, it matters significantly if we work over the floor or on the surface.

Starting on the floor. On the R-oriented floor, we want to have $X_A := W(E_A)$ as the Weil restriction of E_A with group structure $\mathbb{Z}_{p+1} \times \mathbb{Z}_{p+1}$, that is, $W(E_A)$ is primitively oriented by $R = \mathbb{Z}[\sqrt{-p}]$. Then there is a rank-2 module $M_A \subseteq R^2$ such that $W(E_A) = M_A * E_0$. Computing key generation as an $\mathbb{F}_{p^2}$-rational isogeny $\varphi_A : E_0 \to E_A$ and deriving the correct module M_A associated to φ_A is much faster than computing the actual 2-dimensional isogeny $\Phi : E_0^2 \to W(E_A)$.

Unfortunately, it is not always easy or intuitive to derive the associated module M_A given the isogeny $E_0 \to E_A$. This is caused mainly by the fact that E_0 is (primitively) oriented by $R = \mathbb{Z}[\sqrt{-p}]$ instead of $R' = \mathbb{Z}[\frac{1+\sqrt{-p}}{2}]$. One significant difficulty is that $W(E_0)$ is only a twist of E_0^2 over $\mathbb{F}_p$, and so, while the isogeny $\varphi_A : E_0 \to E_A$ gives us a derived 2D isogeny $W(\varphi_A) : W(E_0) \to W(E_A)$, we cannot easily translate this into the isogeny $\Phi : E_0^2 \to W(E_A)$ that we want, so that we can derive M_A.[22]

[21] You need to carefully track if we are working over $\mathbb{F}_p$ or $\mathbb{F}_{p^2}$ in this section.

[22] As mentioned in the introduction, we implicitly assume principal polarisations in this level. To properly find this required module M_A requires a change of polarisation and therefore a deeper understanding of the module action. See Level 3½ and in particular Robert's in-depth description [Rob24].

Starting on the surface. Let's try this again on the R'-oriented surface: for $\varphi_A : E'_0 \to E_A$, if we want something that we can reach by a projective module, it must have the group structure of $E'^2_0(\mathbb{F}_p) = \mathbb{Z}_{\frac{p+1}{2}} \times \mathbb{Z}_{\frac{p+1}{2}} \times \mathbb{Z}_2 \times \mathbb{Z}_2$. This works for Scholten's construction $W'(E_A)$, so there should be a module $M_A \subseteq R'^2$ such that $W'(E_A) = M_A * E'_0$.

It turns out that this M_A is much easier to derive. Although there are deeper technical reasons, a key observation is that $W'(E'_0) \cong E'^2_0$. So if we apply the analogue of the above using W', we get that $W'(\varphi_A) : W'(E'_0) \to W'(E_A)$ is essentially an isogeny $E'^2_0 \to W'(E_A)$, using the above isomorphism. So, if we know how to describe $\ker W'(\varphi_A)$ given $\ker \varphi_A$, it will be much simpler to derive the module $M_A \subseteq R'^2$ such that $E'^2_0[M_A] \cong \ker W'(\varphi_A)$. This is exactly the M_A that gives us $W'(E_A) = M_A * E'_0$.

This gives us the best of both worlds: Alice can still perform key generation as a 1-dimensional isogeny $\varphi_A : E'_0 \to E_A$, yet have a simple translation to her secret module $M_A \subseteq R'^2$. We then get $X_A = W'(E_A) = M_A * E'_0$, and using M_A, we can compute the kernel $X^2_B[M_A]$ for the 4-dimensional isogeny $\Phi_A : X^2_B \to X_{AB}$. Having learned that we should be working on the surface, let us try to derive M_A.

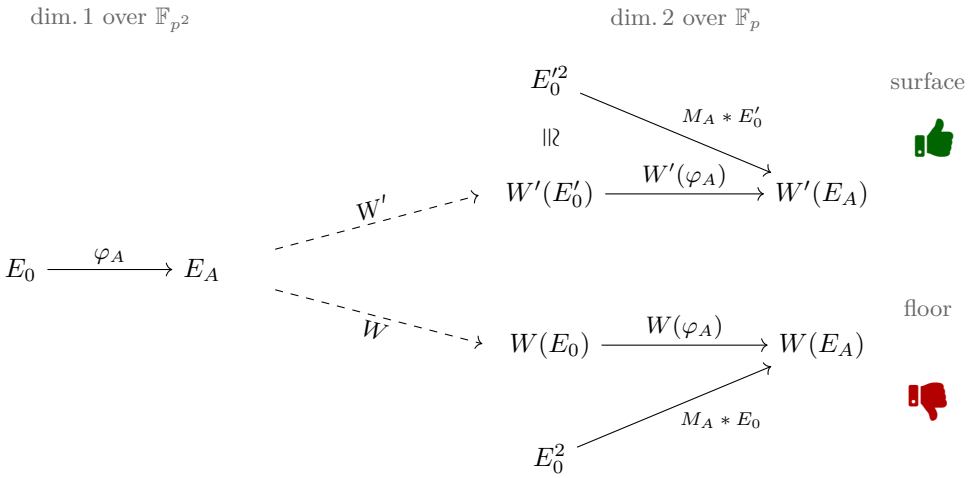


Figure 6: Applying W and W' give us 2-dimensional isogenies associated to $\varphi_A \colon E_0 \to E_A$, which help us to find an R-module, resp. R'-module, M_A so that $W(E_A) = M_A * E_0$ on the floor, resp. $W'(E_A) = M_A * E'_0$ on the surface. In both cases, such a module M_A exists, but it's cumbersome to find on the floor, whereas it's easy on the surface by the isomorphism $W'(E'_0) \to E'^2_0$.

3.3.2 Concrete derivation of M_A

To derive M_A in general, we will need three steps.

A. Given $\ker \varphi_A = \langle K_A \rangle$ with $K_A = P_0 + [x_A]Q_0$, we give a description of $\ker W'(\varphi_A) = \langle K_1, K_2 \rangle$ for some $K_1, K_2 \in E'^2_0$.

B. We finally explain what this special basis (P_0, Q_0) of $E'_0[2^e]$ is, which allows us to derive the right module elements $m \in R'^2$ that kill K_1 and K_2.

C. Using both **A** and **B**, we give a description of M_A in terms of x_A, such that $E'^2_0[M_A] = \ker W'(\varphi_A)$.

A. The Scholten kernel. Given an isogeny $\varphi : E \to E'$, we may apply W' to all arguments to get a 2D-isogeny $W'(\varphi) : W'(E) \to W'(E')$. The curve E'_0 is somewhat

special, as it is its own quadratic twist $E_0'^t$ by the map $\iota : (x, y) \mapsto (-x, iy)$. For any curve E over $\mathbb{F}_p$ with $E[2] \subseteq E(\mathbb{F}_p)$, we get $W'(E) = E \times E^t$, and so in the case of E_0' we get

$$W'(E_0') \cong E_0' \times E_0'^t \cong E_0'^2. \tag{2}$$

Following Corollary 3.1 [BDD+26], if $\ker \varphi_A = \langle K_A \rangle$, then the kernel of $W'(\varphi_A)$ is generated by $K_1 = (K_A, \iota(K_A))$ and $K_2 = (\omega(K_A), (\omega \circ \iota)(K_A)) = \omega(K_1)$. This interprets $W'(E_0')$ as $E_0'^2$, which is valid by the above isomorphisms.

B. The basis. Whenever $2^e \mid \frac{p+1}{4}$, we can choose a basis of $E_0'[2^e]$ in a rather special way. Namely, we can choose P_0 as $P_0 = [2]S$ for some $S \in E_0'(\mathbb{F}_p)$ of order 2^{e+1}, and this implies that

$$\omega P_0 = [1 + \sqrt{-p}]S = S + \pi(S) = [2]S = P_0.$$

Then, the action of $\alpha = a + b\omega \in R'$ on P_0 is simply

$$[\alpha](P_0) = [a]P_0 + [b]P_0 = [a + b]P_0.$$

If we then choose Q_0 as $\iota(P_0)$, we get that (P_0, Q_0) is a basis[23] for $E_0'[2^e]$, but also that $\iota(Q_0) = -P_0$. Writing $S = (x_S, y_S)$ with $x_S, y_S \in \mathbb{F}_p$ for a brief moment,

$$\pi(\iota(S)) = \pi(-x_S, iy_S) = (-x_S, -iy_S) = -\iota(S),$$

so that

$$\omega(Q_0) = [1 + \sqrt{-p}]\iota(S) = \iota(S) - \iota(S) = \mathcal{O},$$

and thus also $[\alpha](Q_0) = [a]Q_0$ for $\alpha = a + b\omega \in R'$. This special choice of basis will help us tremendously with deriving the right module elements to kill $W'(\varphi_A)$.

C. The right module. We now want to find the right $m_i \in R'^2$ such that for $M_A \subseteq R'^2$ generated by these m_i, we get $E_0'^2[M_A] = \langle K_1, K_2 \rangle$. An easy choice to start with is $m_1 = (2^e, 0)$ and $m_2 = (0, 2^e)$, to ensure that $E_0'^2[M_A] \subseteq E_0'^2[2^e]$, so that we now only have to look for elements m_i that cut out exactly K_1 and K_2 in $E_0'^2[2^e]$. We do this by a magical guess for such elements m_3 and m_4, which we secretly derived by rephrasing Corollary 3.1 from [BDD+26] in our terms.

Lemma 1. *Let (P_0, Q_0) satisfy the above properties, and let $K_A = P_0 + [x_A]Q_0$ generate the kernel of $\varphi_A : E_0' \to E_A$. Let $\alpha, \beta \in R'$ be defined by $\alpha = -1 + \omega$ and $\beta = x_A - x_A \cdot \omega$. Then the module $M_A \subseteq R'^2$ generated by*

$$m_1 = (2^e, 0), \quad m_2 = (0, 2^e), \quad m_3 = (\alpha, \beta), \quad m_4 = (-\bar{\beta}, \bar{\alpha}),$$

satisfies $E_0'^2[M_A] = \ker W'(\varphi_A) = \langle K_1, K_2 \rangle$, with $K_1 = (K_A, \iota K_A)$ and $K_2 = \omega(K_1)$.

Hence, M_A is the module associated to Alice's secret $\varphi_A : E_0' \to E_A$, and $W'(E_A) = M_A * E_0'$.

These modules $M_A \subseteq R'^2$ are rather specific special modules, but they are enough to understand how MIKE works! We will refer to these specific modules as *MIKE modules*, to distinguish them from more general modules.

Remark 5. We promise that fairly little is lost by skipping the following proof, as it is rather computational. This is entirely because we do not want to introduce more advanced concepts here, and so we have to compute more. We derive a similar but slightly more streamlined version in Level 3½.

[23] This is because both are derived from an $S \in E_0'(\mathbb{F}_p)$ such that $[2^e]S \neq (0, 0)$, which holds as $p \equiv 7 \bmod 8$. As $Q_0 = \iota(P_0)$, we get $[2^{e-1}]P_0 = [2^e]S \neq [2^e]\iota(S) = [2^{e-1}]Q_0$, hence Q_0 is independent from P_0.

Proof. We essentially need to show that the elements $m_3 = (\alpha, \beta)$ and $m_4 = (-\bar{\beta}, \bar{\alpha})$ kill K_1 and K_2. As $\bar{\omega} = 1 - \omega$, we get

$$-\bar{\beta} = -x_A + x_A\bar{\omega} = -x_A\omega \quad \text{and} \quad \bar{\alpha} = -\omega.$$

Now take any $(S, T) \in E_0'^2[2^e]$, then we may write this element as

$$(S, T) = ([a]P_0 + [b]Q_0, [c]P_0 + [d]Q_0) \qquad \text{for some } a, b, c, d \in \mathbb{Z}_{2^e}.$$

By definition, such an (S, T) is in $E_0'^2[M_A]$ if both $[m_3]$ and $[m_4]$ kill it, that is, if

$$\begin{aligned}[m_3](S, T) &= [\alpha]([a]P_0 + [b]Q_0) + [\beta]([c]P_0 + [d]Q_0) \\ &= [-a + a]P_0 + [-b]Q_0 + [x_Ac - x_Ac]P_0 + [x_Ad]Q_0 \\ &= [-b + x_Ad]Q_0\end{aligned}$$

is 0, hence when $b = x_Ad$, and if

$$\begin{aligned}[m_4](S, T) &= [-\bar{\beta}]([a]P_0 + [b]Q_0) + [\bar{\alpha}]([c]P_0 + [d]Q_0) \\ &= [-x_A\omega]([a]P_0 + [b]Q_0) + [-\omega]([c]P_0 + [d]Q_0) \\ &= [-x_Aa - c]P_0\end{aligned}$$

is 0, hence when $c = -x_Aa$. Thus $E_0'^2[M_A]$ is generated by $L_1 = (P_0, [-x_A]P_0)$ and $L_2 = ([x_A]Q_0, Q_0)$. Now we observe that $K_1 = L_1 + L_2$ and

$$\begin{aligned}K_2 &= \omega(K_A, \iota K_A) \\ &= \omega((P_0 + [x_A]Q_0), (\iota P_0 + [x_A]\iota(Q_0))) \\ &= \omega((P_0 + [x_A]Q_0), (Q_0 - [x_A]P_0)) \\ &= (P_0, -[x_A]P_0) \\ &= L_1.\end{aligned}$$

We conclude that therefore $M_A \subseteq R'^2$ generated by m_1, m_2, m_3, m_4 indeed satisfies $E_0'^2[M_A] = \langle K_1, K_2 \rangle = \ker W'(\varphi_A)$. □

Example 19.

Let us translate a 2^4-isogeny $\varphi_A : E_0' \to E_A$ into a module M_A. We take again $p = 127 = 2^7 - 1$, with $\mathbb{F}_{p^2} = \mathbb{F}_p(i)$. A suitable 2^4-basis is $P_0 = (18, 90)$ and $Q_0 = (109, 90i)$. Alice picks the supersecure secret $x_A = 3$, which gives the 16-torsion point

$$K_A = P_0 + [3]Q_0 = (84 + 97i, 40 + 4i).$$

Precomposing with an isomorphism $\psi : E_0' \xrightarrow{\sim} E_{103}$, we can compute φ_A with $\ker \varphi_A = \langle K_A \rangle$ using the formulas from Level 2 as four 2-isogenies

$$\varphi_A : E_0' \xrightarrow{\sim} E_{103} \xrightarrow{\varphi_1} E_0 \xrightarrow{\varphi_2} E_6 \xrightarrow{\varphi_3} E_{67} \xrightarrow{\varphi_4} E_{45+96i}.$$

Using Lemma 1, we get $\alpha = -1 + \omega$ and $\beta = 3 - 3 \cdot \omega$, and so the module $M_A \subseteq R'^2$ is generated by

$$m_1 = (16, 0), \quad m_2 = (0, 16), \quad m_3 = (-1 + \omega, 3 - 3\omega), \quad m_4 = (-3\omega, -\omega).$$

Wait a minute!! This is exactly the module action we computed in Example 16, where we found $J_{C_4} = M_A * E_0'$ with J_{C_4} the Jacobian of

$$C_4 : y^2 = 112x^6 + 67x^5 + 110x^4 + 53x^3 + 117x^2 + 10x + 121.$$

If the theory is correct, we are saying that J_{C_4} is *also* isomorphic to Scholten's construction $W'(E_A)$. Wait a second minute!! We already computed $W'(E_A)$ in Example 18 as the Jacobian $J_{C'}$ where

$$C' : y^2 = x^5 + 21x^4 + 104x^3 + 109x^2 + 19x.$$

And indeed, with Sage or Magma, one can quickly find that C_4 and C' are $\mathbb{F}_p$-isomorphic. In fact, the path of $(2,2)$-isogenies given in Example 16 is an exact mirror of the W'-path of 2-isogenies given above, as shown in Fig. 7, making the examples full circle.

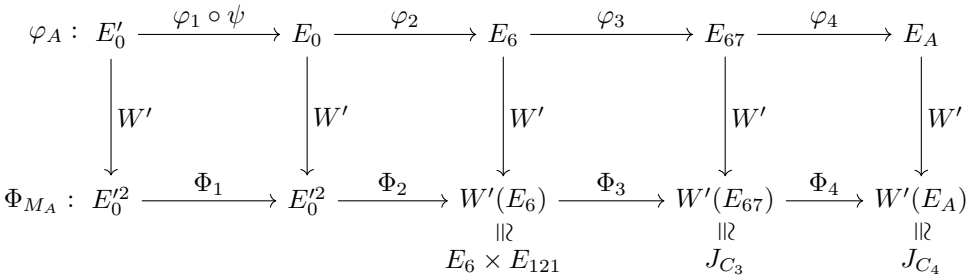


Figure 7: The 1D-isogeny $\varphi_A : E_0' \to E_A$ from Alice's key generation mirrored by the 2D-isogeny $\Phi_{M_A} : E_0'^2 \to W'(E_A)$ from Example 16, given by $W'(E_A) = M_A * E_0'$.

More general MIKE. One should note that MIKE derives M_A from φ_A purely for computational efficiency: Alice could take any secret rank-2 module $M_A \subseteq R'^2$ and define the public key as $X_A = M_A * E_0'$ which we may compute by the 2-dimensional isogeny $\Phi : E_0'^2 \to X_A$ with kernel $\ker \Phi = E_0'^2[M_A]$. If Bob does the same to obtain X_B, then both can still derive a shared secret X_{AB} by $M_A * X_B$, resp. $M_B * X_A$. The downside is twofold: this initial 2-dimensional isogeny Φ is slower to compute than the 1-dimensional isogeny $\varphi_A : E_0' \to E_A$, and the public key in this version is not a simple elliptic curve $E_A : y^2 = x^3 + Ax^2 + x$ which we can easily communicate by $A \in \mathbb{F}_{p^2}$, but instead an abelian surface X_A over $\mathbb{F}_p$, which is more annoying to communicate.

3.4 Shared secret derivation as a module action

We now have the necessary tools to resolve the remaining differences between Fig. 2b and Fig. 3. We understand that the rank-2 module action $X = M * E$ on an elliptic curve E over $\mathbb{F}_p$ corresponds to a 2-dimensional isogeny $\Phi_M : E^2 \to X$. When applied to E_0' and a secret module M_B, this corresponds to $W'(E_B) = M_B * E_0'$ and $\Phi_{M_B} : E_0'^2 \to W'(E_B)$. The squaring step in Level 2 to go from $W(E_B)$ to $W(E_B)^2$ derives exactly from the definition of the rank-2 module kernel $A^2[M]$ when M is of rank 2. In the previous section, we derived M_A as the module derived from $\varphi_A : E_0' \to E_A$ that gives us the isogeny $\Phi_{M_A} : E_0'^2 \to W'(E_A)$. Thus, the shared secret computation can really be seen as $M_A * W'(E_B)$ to get an isogeny $\Phi_A : W'(E_B)^2 \to X_{AB}$ of dimension $4 = \operatorname{rank}(M_A) \cdot \dim(W'(E_B))$, with

$$X_{AB} = M_A * W'(E_B) = (M_A \otimes M_B) * E_0' = (M_B \otimes M_A) * E_0' = M_B * W'(E_A).$$

Given that we now know how to translate $\varphi_A : E_0' \to E_A$ to a module $M_A \subseteq R'^2$, or more accurately, how to derive M_A given $x_A \in \mathbb{Z}_{2^e}$, one should wonder two things:

1. How do we describe the kernel of Φ_A given M_A in simple terms?
2. Why does Theorem 3 describe this kernel as elements of $W(E_B)$ instead of $W'(E_B)$?

The latter was referred to as "the ugly" before. We should apologise: this is actually rather cool.

3.4.1 The kernel of Φ_A

Properly speaking, $X_B = M_B * E'_0$ is now $W'(E_B)$, and not $W(E_B)$ or E_B, although we will get to that in a minute. So, the right kernel for Alice's shared secret derivation is defined by $X_B^2[M_A] = W'(E_B)^2[M_A]$. We will follow an approach similar to the proof of Lemma 1 for the MIKE modules M_A derived from x_A. The idea is rather simple: the elements $m_1 = (2^e, 0)$ and $m_2 = (0, 2^e)$ will always cut out $W'(E_B)^2[2^e]$, so we are trying to find a nice description of the elements $(S, T) \in W'(E_B)^2[2^e]$ such that $[m_3](S, T) = \mathcal{O}$ and $[m_4](S, T) = \mathcal{O}$. It turns out that with a bit of magic this description is not too bad. We apologise for the lengthy description of the following result, but by now you will know most of this already!

Lemma 2. *Let $x_A \in \mathbb{Z}_{2^e}$ with $x_A \equiv 3 \bmod 8$ be Alice's secret, with associated secret isogeny $\varphi_A : E'_0 \to E_A$. Let $\alpha, \beta \in R'$ be defined by $\alpha = -1 + \omega$ and $\beta = x_A - x_A \cdot \omega$, so that Alice's MIKE module $M_A \subseteq R'^2$ is generated by*

$$m_1 = (2^e, 0), \quad m_2 = (0, 2^e), \quad m_3 = (\alpha, \beta), \quad m_4 = (-\bar{\beta}, \bar{\alpha}).$$

Then M_A satisfies $E_0'^2[M_A] = \ker W'(\varphi_A)$. For a given public key E_B, denote with

$$\Phi_A : W'(E_B)^2 \to X_{AB}$$

Alice's 4D-isogeny associated to $M_A * W'(E_B)$. Let $(P_1, \ldots, P_4)$ denote any basis of $W'(E_B)[2^e]$, then the kernel of Φ_A is generated by $K_1, K_2, K_3, K_4 \in W'(E_B)^2[2^e]$ with

$$K_i := [\tfrac{1+x_A}{2}](P_i, -\pi(P_i)) + [\tfrac{1-x_A}{2}](\pi(P_i), P_i) \qquad \text{for all } 1 \le i \le 4.$$

Remark 6. In contrast to the situation we used to translate φ_A to M_A, which required a rather specific basis (P_0, Q_0) on E'_0, the above result really works for any basis $(P_1, \ldots, P_4)$ of $W'(E_B)[2^e]$. Furthermore, note that the coefficients $\frac{1 \pm x_A}{2}$ are exactly those of Theorem 3.

Proof. It is clear that m_1 and m_2 kill the K_i, so let's show that m_3 and m_4 kill the K_i. By now, you should know that this will be a bit computational. As an exercise, derive the following two identities which hold for ω and $\bar{\omega}$ on the 2^e-torsion, as $2^e \mid \frac{p+1}{4}$:

$$\omega\pi = \omega \quad \text{and} \quad \bar{\omega}\pi = -\bar{\omega}.$$

First apply m_3 to $(P_i, -\pi(P_i))$ to get $(\alpha - \beta\pi)P_i$ and to $(\pi(P_i), P_i)$ to get $(\alpha\pi + \beta)P_i$. With $\alpha = -1 + \omega = -\bar{\omega}$ and $\beta = x_A \cdot (1 - \omega) = x_A\bar{\omega}$, we get

$$\begin{aligned}
[m_3](K_i) &= \left[\left(\tfrac{1+x_A}{2}\right)(\alpha - \beta\pi) + \left(\tfrac{1-x_A}{2}\right)(\alpha\pi + \beta)\right] P_i \\
&= \left[\left(\tfrac{1+x_A}{2}\right)(-\bar{\omega} + x_A\bar{\omega}) + \left(\tfrac{1-x_A}{2}\right)(\bar{\omega} + x_A\bar{\omega})\right] P_i \\
&= \left[\left(\tfrac{1+x_A}{2}\right)(-1 + x_A) + \left(\tfrac{1-x_A}{2}\right)(1 + x_A)\right] \bar{\omega}(P_i) \\
&= \left[\left(\tfrac{-1+x_A^2}{2}\right) + \left(\tfrac{1-x_A^2}{2}\right)\right] \bar{\omega}(P_i) \\
&= \mathcal{O}.
\end{aligned}$$

To not keep any more computational delight from you, we leave $[m_4]K_i = \mathcal{O}$ as an exercise to the reader. This shows that $\langle K_1, \ldots, K_4 \rangle \subseteq W'(E_B)^2[M_A]$, so we are left to show that they indeed span the whole of $W'(E_B)^2[M_A]$ of size 2^{4e}. Clearly, each K_i has order 2^e, so we need to show linear independence of these K_i, which follows mainly from independence

of the basis P_i: If we would have $\sum \lambda_i K_i = \mathcal{O}$ for some $\lambda_i \in \mathbb{Z}_{2^e}$, then let $S = \sum \lambda_i P_i$, so that

$$\mathcal{O} = \sum \lambda_i K_i = [\tfrac{1+x_A}{2}](S, -\pi(S)) + [\tfrac{1-x_A}{2}](\pi(S), S).$$

We get an equation in S and $\pi(S)$ from each component:

$$\mathcal{O} = [\tfrac{1+x_A}{2}]S + [\tfrac{1-x_A}{2}]\pi(S) \quad \text{and} \quad \mathcal{O} = [-\tfrac{1+x_A}{2}]\pi(S) + [\tfrac{1-x_A}{2}]S.$$

We can use the two equations to eliminate the $\pi(S)$ part, and get

$$\mathcal{O} = [\tfrac{(1+x_A)^2}{4} + \tfrac{(1-x_A)^2}{4}]S = [\tfrac{1+x_A^2}{2}](S).$$

Because $x_A \equiv 3 \bmod 8$, and in fact we just need x_A odd, we get $1 + x_A^2 \equiv 2 \bmod 8$ hence $\frac{1+x_A^2}{2} \equiv 1 \bmod 4$. But then multiplying by S does not change its order, so S must have been $\mathcal{O}$ from the start. So $\mathcal{O} = S = \sum \lambda_i P_i$ implies $\lambda_i = 0$ for all i, which in turn proves the linear independence of the K_i, completing the proof. □

3.4.2 To $W(E_B)$ or not to $W(E_B)$?

We have done a lot of work to convince you that the proper isogeny Φ_A is a 4-dimensional isogeny $\Phi_A : W'(E_B)^2 \to X_{AB}$, derived in Lemma 2. Nevertheless, Theorem 3 and the MIKE paper both give you a kernel $G \subseteq W(E_B)^2$ associated to an isogeny $F_A : W(E_B)^2 \to X_{AB}$, instead of $W'(E_B)^2$.

Luckily, this is resolved rather easily: for simplicity, efficiency and compactness, it's much easier to give the public key as E_B rather than $W'(E_B)$, and as we model $W(E_B)$ as points $(P, Q) \in E_B \times E_B^{(p)}$ with the 'strange' Frobenius $(P, Q) \mapsto (\pi(Q), \pi(P))$, it is also easier to describe a kernel on $W(E_B)$ in terms of combinations of points on E_B and $E_B^{(p)}$. This avoids having to work with elements on the Jacobian $W'(E_B)$ to describe Φ_A, which has fairly annoying arithmetic in comparison to elliptic curves. So, instead of deriving the kernel $\ker \Phi_A = \langle K_1, K_2, K_3, K_4 \rangle \subseteq W'(E_B)^2$ using Lemma 2, in MIKE, following Appendix F [BDD+26], we instead derive the almost-equivalent isogeny $F_A : W(E_B)^2 \to X_{AB}$. The upside is that we can derive $\ker F_A$ from $\ker \Phi_A$ purely theoretically using some algebraic manipulation by the Scholten gluing

$$\Gamma : W(E_B) \to W'(E_B), \quad \ker \Gamma = \{(P, \pi(P)) \mid P \in E_B[2]\}$$

and (pages of) linear algebra, which we leave to Appendix F [BDD+26] for the meticulous reader. This bridges the gap between $\ker \Phi_A$ from Lemma 2 and $\ker F_A$ from Theorem 3. We stress that $\ker F_A \neq W'(E_B)^2[M_A] = \ker \Phi_A$. Instead, F_A is a clever trick to simplify the computation of Φ_A, and still derive the same codomain X_{AB}. We hope that this finally helps you understand the big MIKE diagram in Fig. 8.

Summary. To summarise, Alice wants to compute $\Phi_A : W'(E_B)^2 \to X_{AB}$, but instead, computes the almost-equivalent isogeny $F_A : W(E_B)^2 \to X_{AB}$ to derive X_{AB}. She does this as the kernel $\ker F_A$ can be easily described using combinations of points on E_B, and furthermore allows for easier computation. Importantly, the computation still gives the same codomain, and therefore computes the module action $M_A * W'(E_B)$. This module action should not be confused with the module action $M * W(E_B)$, which uses a different base ring R instead of R'.

3.5 A MIKE example over $p = 127$ for Alice and Bob

Let us throw this all together and compute an actual key exchange between Alice and Bob for the toy prime $p = 127 = 2^7 - 1$, using isogenies of degree $32 = 2^5$.[24] On $E_0' : y^2 = x^3 - x$,

[24] We apologise that we need to use $e = 5$, which seems to be the smallest proper example of MIKE that we could find. For rather technical reasons, $e < 5$ gives trivial key exchanges.

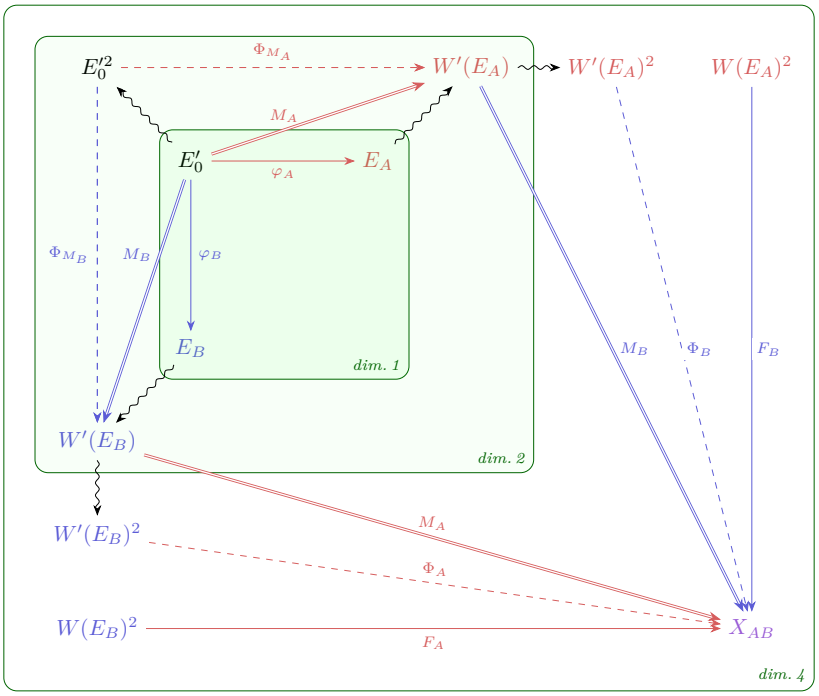


Figure 8: Adapted diagram from [BDD+26] of the MIKE protocol, with Alice in red and Bob in blue. Double arrows represent module actions; solid arrows are isogenies, with dashed arrows as implicit isogenies. Squiggly arrows represent dimension-upping via products or Scholten's construction. Note that F_A, resp. F_B, are only derived from Φ_A, resp. Φ_B, and do not factor through these isogenies: there is no map $G : W(E_B)^2 \to W'(E_B)^2$ such that $F_A = \Phi_A \circ G$, or vice versa.

we fix the 2^e-basis $P_0 = (16, 4)$ and $Q_0 = (111, 4i)$. Alice chooses her favourite[25] secret $x_A = 3$. We follow the same approach as in Example 19, and compute the 32-torsion point

$$K_A = P_0 + [3]Q_0 = (119 + 35i, 72 + 75i).$$

Precomposing with the same isomorphism $\psi : E_0' \xrightarrow{\sim} E_{103}$ as in Example 19, Alice can again compute φ_A with $\ker \varphi_A = \langle K_A \rangle$ using the formulas from Level 2, this time as five 2-isogenies

$$\varphi_A : E_0' \xrightarrow{\sim} E_{103} \xrightarrow{\varphi_{A,1}} E_0 \xrightarrow{\varphi_{A,2}} E_6 \xrightarrow{\varphi_{A,3}} E_{67} \xrightarrow{\varphi_{A,4}} E_{45+96i} \xrightarrow{\varphi_{A,5}} E_{86+77i},$$

which gives her the public key $A = 86 + 77i$.

Bob instead chooses $x_B = 19$, and gets

$$K_B = P_0 + [19]Q_0 = (59 + 46i, 33 + 11i).$$

Again precomposing with $\psi : E_0' \xrightarrow{\sim} E_{103}$, Bob can compute φ_B with $\ker \varphi_B = \langle K_B \rangle$ through five 2-isogenies

$$\varphi_B : E_0' \xrightarrow{\sim} E_{103} \xrightarrow{\varphi_{B,1}} E_0 \xrightarrow{\varphi_{B,2}} E_6 \xrightarrow{\varphi_{B,3}} E_{67} \xrightarrow{\varphi_{B,4}} E_{45+96i} \xrightarrow{\varphi_{B,5}} E_{115+34i},$$

giving him the public key $B = 115 + 34i$.

[25] We strongly recommend never choosing something like 3 or 4 as a private key, which unfortunately is still news to some organisations in 2026.

We immediately notice that actually Alice's and Bob's first few steps are the same! This is due to the setup: the restriction on secret keys that $x_A, x_B \equiv 3 \bmod 8$ fixes the first three steps. In our example, this extends to the fourth step since additionally $x_A \equiv x_B \bmod 16$, which is the main reason that we have to choose $e > 4$ for a meaningful example.[26]

⚠ Warning

The paper [BDD+26] chooses the lexicographically largest Montgomery coefficient to represent MIKE public keys E_A and E_B. For now, we skip this step just to simplify exposition, as it makes no significant difference for understanding MIKE. See Level 3½ for more details.

The modules. Deriving the modules $M_A, M_B \subseteq R'^2$ is now super easy: we did this in Lemma 1 for the generic MIKE module $M \subseteq R'^2$ generated by

$$m_1 = (2^e, 0), \quad m_2 = (0, 2^e), \quad m_3 = (\alpha, \beta), \quad m_4 = (-\bar{\beta}, \bar{\alpha}),$$

with $\alpha = -1 + \omega$ and $\beta = x_A - x_A \cdot \omega$. So for M_A, Alice will use this same α, and $\beta = 3 \cdot (1 - \omega)$. For M_B, Bob will use the same α too, and $\beta = 19 \cdot (1 - \omega)$. Of course, in practice, neither actually computes M_A or M_B, both just throw their secrets x_A and x_B into Theorem 3.

Alice's kernel. First, Alice picks a basis for $E_B[2^7]$ instead of just $E_B[2^5]$ as it comes at no extra cost, but will improve the efficiency of the computation of F_A, see Level 3½. She computes

$$P_B = (68 + 35i, 22 + 92i), \quad Q_B = (107 + 88i, 66 + 70i).$$

It turns out that in practice, we don't even need all four kernel points of $\ker F_A \subseteq W(E_B)^2$, which is somewhat magical too. We explain this in Level 3½. First, Alice computes the points $T_3, T_4 \in E_B^2$ as

$$\begin{aligned} T_3 &= (-[\tfrac{1+x_A}{2}]P_B, [\tfrac{1-x_A}{2}]P_B) \\ &= ((125 + 23i, 52 + 79i), (68 + 35i, 105 + 35i)) \\ T_4 &= ([\tfrac{1+x_A}{2}]Q_B, -[\tfrac{1-x_A}{2}]Q_B) \\ &= ((38 + 112i, 70 + 28i), (107 + 88i, 66 + 70i)). \end{aligned}$$

The kernel points in Theorem 3 are given as 4-tuples of $W(E_B)^2$, whose points we modelled as (P, Q, R, S) for $P, R \in E_B$ and $Q, S \in E_B^{(p)}$. These $T_3, T_4 \in E_B^2$ play the role of (P, R), i.e., the 1st and 3rd entries of the 4-tuples in Theorem 3. The 2nd and 4th entries (Q, S) are simply $-\pi(T_3)$, resp. $-\pi(T_4)$. As Alice picked a basis (P_B, Q_B) for the 2^7-torsion, these translate to kernel generators for $\ker F_A$ via $K_3 = [4]T_3$ and $K_4 = [4]T_4$ and applying $-\pi$ similarly.

Bob's kernel. Similarly, on E_A, Bob picks a basis

$$P_A = (108 + 107i, 107 + 88i), \quad Q_A = (111 + 3i, 90 + 52i)$$

and derives the points

$$\begin{aligned} T_3' &= ((40 + 54i, 28 + 43i), (28 + 81i, 39 + 98i)) \\ T_4' &= ((113i, 37 + 8i), (114 + 45i, 116 + 116i)), \end{aligned}$$

[26] This is not a problem in terms of security: the MIKE parameter choices for the size of e consider the loss of a few bits due to the first steps being fixed.

with the similar warnings about $K_3' = [4]T_3'$ and $K_4' = [4]T_4'$, and the similar magic that we only need two instead of four of these points to compute the special MIKE 4D-isogenies. With that, we are ready to compute!

3.5.1 The 4D isogeny

As in Level 2, we oversimplify these computations a bit, by assuming that you just throw the right basis into formulas. This is also what we do here: we assume both Alice and Bob throw the points (T_3, T_4) and (T_3', T_4') into the right code from the `mike_py` repository.[27]

If Alice does that, she computes $F_A : W(E_B)^2 \to X_{AB}$ and gets as a theta null point for her codomain

$$t_A = (1 : 68 : 7 : 83 : 7 : 32 : 114 : 54 : 68 : 98 : 32 : 71 : 83 : 71 : 54 : 43).$$

Similarly, Bob computes $F_B : W(E_A)^2 \to X_{AB}$ and gets

$$t_B = (1 : 55 : 7 : 96 : 7 : 45 : 114 : 67 : 55 : 78 : 45 : 91 : 96 : 91 : 67 : 63).$$

Deriving the shared secret. We then only need to compute the MIKE invariants from Definitions 2 and 3. Alice computes these for t_A as

$$\begin{aligned}
I_0 &= (1+114+98+43)^2 = 4, & I_5 &= 4(7-71)^2 = 1,\\
I_1 &= (1-114+98-43)^2 = 62, & I_6 &= 4(68+54)^2 = 100,\\
I_2 &= (1+114-98-43)^2 = 41, & I_7 &= 4(68-54)^2 = 22,\\
I_3 &= (1-114-98+43)^2 = 30, & I_8 &= 4(83+32)^2 = 68,\\
I_4 &= 4(7+71)^2 = 79, & I_9 &= 4(83-32)^2 = 117.
\end{aligned}$$

By normalising by $I_0 = 4$, she gets

$$I_A' = (79, 42, 71, 115, 32, 25, 69, 17, 61).$$

By a similar computation, Bob gets $I_B = (4, 79, 41, 1, 62, 30, 100, 68, 22, 117)$, and normalised I_B' as

$$I_B' = (115, 42, 32, 79, 71, 25, 17, 69, 61).$$

As you can see, I_B' is a permutation of I_A', which is exactly Theorem 4.2 [BDD+26]. The goal of the absolute J-invariants is essentially to get a set of invariants covering all possible MIKE permutations of these normalised invariants I'. As the grand finale, the catharsis of this piece, Alice and Bob both compute these absolute J-invariants, and both obtain

$$J_{AB} = (115, 105, 117, 27).$$

Hence, the shared secret computation worked! Alice and Bob have shared a MIKE key!!

[27] https://github.com/tensor-mike/mike_py

Level 3½: Towards MIKE proper

We are sorry, but we lied: the scheme we have described so far is a simplified and perhaps slightly distorted version of MIKE. Although the core ideas are those presented in Levels 1 to 3, the actual scheme applies several clever optimisations to achieve significantly faster performance. Luckily, you should now be able to read the full MIKE paper [BDD+26] in all its complexities. To make it up to you, in this section we describe the differences and simplifications we applied, and why we did that.

3½.1 Not All Modules Are Created Equal

So far, we have been thinking of a module $M \subseteq R^n$ as a full-rank R-lattice generated by some set of elements $m = (r_1, \ldots, r_n) \in R^n$. For the purpose of understanding the module actions $M_A * E_0'$ and $M_A * X_B$, such an understanding is fine. The proper definition of a module, however, is much more general, and the associated module action is similarly more broad, derived elegantly from the categorical idea of the canonical power object construction on categories enriched in a closed symmetric monoidal category. This is far outside the scope of this work, and we therefore give a bit more of a hands-on description of three types of modules, and their associated module action. We then discuss why we avoided this discussion in the first place for MIKE.

3½.1.1 Types of modules

For a commutative ring R, a module M is essentially an abelian group with an action of R on it: $r \cdot m$ is well-defined for all $r \in R$ and $m \in M$, and we have some rather standard properties that this action satisfies for all $r, r' \in R$ and $m, m' \in M$:

$$\begin{aligned}(r + r') \cdot m &= r \cdot m + r' \cdot m \\ r \cdot (m + m') &= r \cdot m + r \cdot m' \\ (rr') \cdot m &= r \cdot (r' \cdot m) \\ 1 \cdot m &= m\end{aligned}$$

For this most general notion of an R-module M, we would not be able to define a meaningful module action $M * A$, where A is an (R-oriented) abelian variety.

Finitely presented. At the very least, we ask our module M to be *finitely presented*. This is a slightly stronger notion than *finitely generated*, which means that you can give a finite set $\{m_1, \ldots, m_n\}$ such that every element $m \in M$ can be written as $m = \sum r_i m_i$ with $r_i \in R$.

The definition is slightly categorical. We give it here for those category nerds that ask you if a construction is 'natural'.

Definition 7. Given a commutative ring R, an R-module M is *finitely presented* if there exists an exact sequence

$$R^k \xrightarrow{\varphi} R^n \xrightarrow{\psi} M \to 0,$$

where $k, n \in \mathbb{Z}_{>0}$.

A (finite) set of generators $(m_1, \ldots, m_n)$ can then be understood as the image under ψ of the elements $e_i = (0, \ldots, 1, \ldots, 0) \in R^n$ with a 1 in the i-th position, whereas φ gives us the relations between such generators m_i.

The associated module action for finitely presented modules is still rather horrible for cryptographers: for an abelian variety A and a finitely presented R-module M, we can now properly define an object $B = M * A$, but B will only be a *commutative proper group scheme*, which generalises the notion of abelian variety.

Projective. We need to avoid some nasty situations that arise when M is not torsion-free. For example, $M = R/n$ acts on A by $R/n * A = A[n]$, and the latter is not an abelian variety. To avoid this, we could ask that M is torsion-free, and then $M * A$ is an abelian variety under certain conditions on M and A. It turns out that a slightly stronger notion that we want is for M to be *projective*. Although for the ring R' we consider in MIKE, these notions are equivalent, this is still relevant for more general constructions. One particular case is the rank-1 torsion-free R-ideal $\mathfrak{f}$ from Example 17, which looks like any other module we discussed so far, which is torsion-free but not projective.[28] An easy way to view projective modules requires us to give intuition for a *free* module in this situation: if M is finitely presented and free, then there exists an exact sequence

$$0 \to R^n \xrightarrow{\psi} M \to 0.$$

In more concrete terms, such a free module has a generating set $\{m_1, \ldots, m_n\}$ with $m_i = \psi(e_i) \in M$ which is also linearly independent: for any $r_i \in R$ such that

$$\sum r_i m_i = 0, \qquad \text{we get} \qquad r_i = 0 \text{ for all } 1 \leq i \leq n.$$

The notion of projective is slightly weaker:

Definition 8. Let R be a commutative ring, and M a finitely presented R-module. Then M is *projective* if there exists another R-module M' such that $M \oplus M' \cong R^n$ for some $n \in \mathbb{N}$.

We now get a rather nice action! If A is an (R-oriented) abelian variety, and M is finitely presented and projective, then $B = M * A$ is again an abelian variety, which is precisely what we wanted. For the ring $R' = \mathbb{Z}[\frac{1+\sqrt{-p}}{2}]$ that we are using in MIKE, these finitely presented projective modules of a certain rank n come down to the full-rank R'-lattices $M \subseteq R'^n$ that we were using throughout Level 3, which we may always think of as generated by some set of $m = (r_1, \ldots, r_n)$ with $r_i \in R'$.

3½.1.2 Principal polarisations and Hermitian forms

So far, we have treated abelian varieties as vague generalisations of elliptic curves to higher dimensions $g > 1$, with only very few concrete examples that we have seen for dimension $g = 2$:

- we've seen the Jacobian of a genus-2 hyperelliptic curve C as an example of an abelian surface in Scholten's construction $E \to W'(E)$, for E (strictly) over $\mathbb{F}_{p^2}$,
- we've seen the product of elliptic curves $E_1 \times E_2$ as a (somewhat trivial) example of abelian surfaces,
- and we've seen the Weil restriction $W(E)$, for E (strictly) over $\mathbb{F}_{p^2}$ as an example of an abelian surface.

In general, abelian varieties are rather difficult and nuanced objects, and the main difficulty that is introduced in dimension $g > 1$ is that of a *polarisation*, which is a specific isogeny $\lambda : A \to \hat{A}$, where $\hat{A}$ is the dual of A. The difficulty is that explaining the notion of the dual of A is already beyond the scope of this text, which makes it particularly tricky to properly explain what a polarisation is. The only simple situation is elliptic curves, where $\hat{E} \cong E$ makes the situation so trivial that we never actually encounter polarisations.

Let us therefore informally try to shine some light on the situation: in practice, in cryptography, we should almost always think of the abelian varieties A that we encounter as a pair (A, λ) where λ is such a polarisation of degree 1, which we call a *principal polarisation*. The pair (A, λ) is then a *principally polarised abelian variety*, abbreviated p.p.a.v. As relevant for MIKE, let us look at the objects we have seen.

[28] This extends the CSIDH analogy, where $\mathfrak{f}$ is not *invertible*.

Dimension 2. Secretly, all of the three examples above were already p.p.a.v.'s of dimension 2, as their principal polarisations arise naturally in their definitions, which is why we never encountered them in Level 2 and Level 3:

- the Jacobian of a hyperelliptic curve C inherits a principal polarisation from the curve C,
- the product $E_1 \times E_2$ inherits the *product polarisation* $\lambda_1 \times \lambda_2$ from the natural principal polarisation λ_i on E_i,
- and a Weil restriction is isomorphic to a product over a field extension, which similarly gives it a product polarisation.

Dimension 4. Our 4-dimensional objects all start out from the product of two 2-dimensional objects, such as $W(E_B)^2$ or $W'(E_B)^2$, and therefore naturally get the product polarisations from $W(E_B)$, resp. $W'(E_B)$. Beyond that, the 4-dimensional isogenies $\Phi : A \to B$ we use ensure that the codomain B can be equipped with a principal polarisation λ_B by requiring that the kernel $\ker \Phi$ is maximally isotropic. This is guaranteed by our choice of modules and the description of our kernel elements, and something we have completely skipped over in Level 3 to keep it beginner-friendly. Therefore, the final object X_{AB} is in fact a principally polarised abelian variety.

Polarisations and the module action. The nice action we got before, when A is an abelian variety and M is finitely presented and projective, is therefore missing a crucial piece of information: the module M itself does not care for the (principal) polarisation of A! In order to keep track of polarisations, we have to introduce the notion of a Hermitian form H for a module M. The pair (M, H) will then mirror the pair (A, λ) for p.p.a.v.'s. By slight abuse of terminology, we define a Hermitian form as follows.

Definition 9. A *Hermitian form* on M is a map $H : M \times M \to R$ which is

- Sesquilinear: for all $r \in R$, we have $H(r \cdot m, m') = r \cdot H(m, m')$ and $H(m, r \cdot m') = \bar{r} H(m, m')$.
- Hermitian: for $m, m' \in M$, we have $H(m, m') = \overline{H(m', m)}$.
- Positive definite: for all non-zero $m \in M$, we have $H(m, m) > 0$.

A *Hermitian module* is a pair (M, H), with M an R-module and H a Hermitian form.

One can write many many pages about Hermitian modules, and their properties, as showcased by Robert [Rob24]. What we crucially need is that H mirrors the role of λ in the module world. The analogue of *principal* for a polarisation λ becomes *unimodular* for a Hermitian form H and is similarly defined in terms of a dual $\hat{M}$ of M: we want that $M \to \hat{M}$ given by $m \mapsto H(m, \cdot)$ is an isomorphism.

Example 20.
Although the above definitions are technical, in reality Hermitian forms can be very simple. For example, the map $H : R^n \times R^n \to R$ given by

$$H((r_1, \ldots, r_n), (r'_1, \ldots, r'_n)) = \sum r_i \bar{r'_i}$$

is a unimodular Hermitian form on R^n, and the pair (R^n, H) is a Hermitian module. This is the *standard* Hermitian form.

This is why the MIKE papers so often talk about Hermitian modules, as the resulting action now keeps track of polarisations: if we act with a Hermitian module (M, H) on

a polarised abelian variety (A, λ_A), the resulting $(B, \lambda_B) = (M, H) * (A, \lambda_A)$ is again a polarised abelian variety. Even better, if H is unimodular, and λ_A is principal, then the action takes a p.p.a.v. (A, λ_A) to a p.p.a.v. (B, λ_B).

In cryptography, we almost always want to work with p.p.a.v.'s and so the right module action is really the one by unimodular Hermitian modules (M, H). Luckily, if we start from E'_0, we get the natural principal polarisation λ_0, which ensures that the resulting codomains $(A, \lambda_A) = (M, H) * (E'_0, \lambda_0)$ are p.p.a.v.'s.

3½.1.3 MIKE modules

The context in MIKE is rather lucky: elliptic curves are naturally equipped with a principal polarisation, and by derivation from E, so are $W(E)$ and $W'(E)$. This makes it possible for us to describe the MIKE key exchange in terms of the module action without direct reference to principal polarisations: the modules and varieties used in the MIKE key exchange are of the precise form so that the principal polarisations on our objects are naturally defined, essentially because we start from an elliptic curve E over $\mathbb{F}_p$.

Nevertheless, we *also* ditched the (unimodular) Hermitian form H when we describe the modules M_A that we are acting with in Level 3. This is because the Hermitian forms are hidden in two places:

- On our starting curve E'_0, whenever we move to E'^2_0, we assume the product polarisation, which is equivalent to choosing the module R'^2 equipped with the standard Hermitian form from Example 20.
- Similarly, when we write M_A as generated by elements $m_i = (\alpha_i, \beta_i) \in R'^2$, we implicitly assumed a (Hermitian) norm form $H : (m_i, m_j) \mapsto \alpha_i \cdot \bar{\alpha_j} + \beta_i \cdot \bar{\beta_j}$, and this gives a unimodular Hermitian module by $(M_A, H/2^e)$.

A more correct version of MIKE would include these Hermitian forms, and the polarisations on the abelian varieties whenever we describe the action $(B, \lambda_B) = (M, H) * (A, \lambda_A)$, but this would hardly be beginner-friendly. Luckily, in the MIKE key exchange, we are in the simpler situation where a lot of these technical difficulties can be avoided by only a slight abuse of presentation. This is a core reason behind the jumps in levels of difficulty between Level 3, the appendices of [BDD+26], and the initial descriptions by Robert [Rob24].

3½.2 An isogeny from a quaternionic ideal

⚠ Warning

As it is outside of the main focus of this work, we will not give a proper in-depth introduction to quaternion arithmetic. Luckily, there are ample resources for those who want to dive deeper, such as Voight [Voi21] and Leroux [Ler22]. In particular, this section assumes a reasonable familiarity with quaternions, maximal orders, and the Deuring correspondence.

To prevent this from becoming a full introduction to quaternions, we mainly list the main facts that one should know to grasp this section.

- In general, endomorphisms of supersingular curves behave like quaternion elements γ in the quaternion algebra $B_{p,\infty}$. We write these as $\gamma = a + b\mathbf{i} + c\mathbf{j} + d\mathbf{k}$ with $a, b, c, d \in \mathbb{Q}$, with the multiplication rules $\mathbf{i}^2 = -1$, $\mathbf{j}^2 = -p$, and $\mathbf{k} = \mathbf{i} \cdot \mathbf{j} = -\mathbf{j} \cdot \mathbf{i}$. Note the non-commutativity!
- By the General Mantra, we can take a subgroup $G \subseteq E$ and find an isogeny $\varphi : E \to E'$ with $\ker \varphi = G$. For a general isogeny, by the Deuring correspondence, G

can be described by a quaternionic ideal I. This extends the situation in Section 3.1, where we could only do this for very specific subgroups G using quadratic ideals $\mathfrak{a}$.

- The main ingredient is a correspondence between (rank-4) endomorphism rings $\mathrm{End}(E)$ and maximal orders $\mathcal{O}$ in $B_{p,\infty}$. For example, $\mathrm{End}(E_0') = \langle 1, \iota, \frac{1+\pi}{2}, \frac{\iota-\kappa}{2}\rangle$, where $\kappa = \iota \circ \pi$, is isomorphic to $\mathcal{O}_0' = \langle 1, \mathbf{i}, \frac{1+\mathbf{j}}{2}, \frac{\mathbf{i}-\mathbf{k}}{2}\rangle$ through the isomorphism $[\cdot] : \mathcal{O}_0' \to \mathrm{End}(E_0')$ such that $[\mathbf{i}] = \iota$, and so on. Note that $\omega = \frac{1+\pi}{2}$ and that $\mathcal{O}_0'$ can be written as $\mathcal{O}_0' = R' \oplus R'\mathbf{i}$.

- In MIKE, we only need to understand this correspondence for isogenies $\varphi : E_0' \to E_A$, corresponding to *left* $\mathcal{O}_0'$*-ideals* of the form $I = \langle N, \gamma\rangle$ where N is a positive integer and $\gamma \in \mathcal{O}_0'$ is of the form

$$\gamma = a + b\cdot\mathbf{i} + c\cdot\tfrac{1+\mathbf{j}}{2} + d\cdot\tfrac{\mathbf{i}-\mathbf{k}}{2}, \qquad a,b,c,d \in \mathbb{Z},$$

  where $\gcd(a,b,c,d,N) = 1$. The association between φ and I goes through the kernel $\ker\varphi = E_0'[I]$ defined by

$$\begin{aligned} E_0'[I] &= E_0'[N] \cap E_0'[\gamma] \\ &= \{P \in E_0'(\overline{\mathbb{F}}_p) \mid [N]P = [\gamma]P = \mathcal{O}\} \\ &= \{P \in E_0'[N] \mid [\gamma]P = \mathcal{O}\}. \end{aligned}$$

- One should always keep in mind that isogenies $\varphi_{\mathfrak{a}} : E \to \mathfrak{a} * E$ derived from quadratic ideals $\mathfrak{a}$ are really a very special case of the association between 'general' isogenies $\varphi : E \to E'$ and quaternionic ideals I. Most significantly, the quadratic case can be seen as a group acting on elliptic curves over $\mathbb{F}_p$, and this group is commutative, which means that Kuperberg's quantum algorithm applies: given only E and $\mathfrak{a} * E$, we can recover the class $[\mathfrak{a}]$ of the secret ideal $\mathfrak{a}$ in subexponential time. This is why the underlying primes p in CSIDH explode to roughly 4000 bits [BDL+26], whereas the primes in SQIsign [AAA+26] and MIKE stay much smaller, around 320 bits.

Key generation as a quaternionic ideal

In Lemma 1, we derived the module M_A associated to Alice's secret isogeny $\varphi_A : E_0' \to E_A$. Quaternion ideals will help us to understand where this module is coming from, and in general strengthen our understanding of the link between the 1D isogeny φ_A and the derived 2D-isogeny $\Phi : E_0'^2 \to W'(E_A)$ associated to $M_A * E_0'$. This helps us understand the derivation of the kernel $\ker\Phi_A = W'(E_B)^2[M_A]$ of the 4D-isogeny $\Phi_A : W'(E_B)^2 \to X_{AB}$ too.

Remark 7. Although we need quaternions to *explain* MIKE, and in particular how M_A is derived from φ_A, it is crucial to remember that we do not need quaternions to *compute* MIKE. Instead, we use quaternions to theoretically derive how the kernel of Φ_A looks in terms of the secret $x_A \in \mathbb{Z}_{2^e}$, but the resulting kernel can be computed solely using x_A.

A crucial detail is that we compute an isogeny of degree $2^e \mid \frac{p+1}{4}$, which allows us to pick a special basis (P_0, Q_0) for $E_0'[2^e]$, namely such that

$$[\omega](P_0) = P_0, \qquad \text{and} \qquad Q_0 = [\mathbf{i}](P_0).$$

Lemma 3. *Let* $K_A = P_0 + [x_A]Q_0 \in E_0'[2^e]$. *Then* $\gamma = (\omega - 1)(1 - x_A\mathbf{i})$ *satisfies* $\gamma \in \mathcal{O}_0'$ *and* $[\gamma]K_A = 0$, *and for* $I = \langle 2^e, \gamma\rangle$ *we get* $E_0'[I] = \ker\varphi_A = \langle K_A\rangle$. *Furthermore, as* $\mathcal{O}_0' = R' \oplus R'\mathbf{i}$, *we can write* $\gamma = \alpha + \beta\mathbf{i}$ *for* $\alpha, \beta \in \mathbb{Z}[\frac{1+\sqrt{-p}}{2}]$, *where*

$$\alpha = (\omega - 1) \quad \textit{and} \quad \beta = x_A(1 - \omega).$$

Proof. This is exactly the first half of the proof of Corollary 3.1 [BDD+26], which comes down to computation. By writing out terms, we get

$$\begin{aligned}
[\gamma](K_A) &= [(\omega-1)(1-x_A\mathbf{i})](P_0 + x_A[\mathbf{i}]P_0) \\
&= [(\omega-1)(1-x_A\mathbf{i})(1+x_A\mathbf{i})]P_0 \\
&= [(\omega-1)(1+x_A^2)]P_0 \\
&= [(1+x_A^2)][(\omega-1)]P_0 \\
&= \mathcal{O}.
\end{aligned}$$

The crucial insight is that between line 3 and 4, the quaternion elements now commute (in contrast to ω and $\mathbf{i}$) as $(1+x_A^2)$ is now an integer. As γ is primitive and the reduced norm of I is 2^e, we are ensured that $E_0'[I] = E_0'[\gamma] \cap E_0'[2^e] = \langle K_A \rangle$. □

This finally explains our choices in Lemma 1 for $\alpha, \beta \in \mathbb{Z}[\frac{1+\sqrt{-p}}{2}]$, as derived from the associated quaternionic ideal! We rephrase that result in a more general form, which constitutes the second half of the proof of Corollary 3.1 [BDD+26].

Lemma 4. *Let $\gamma = \alpha + \beta\mathbf{i} \in \mathcal{O}_0'$ such that $[\gamma]K_A = 0$, and let $m = (\alpha, \beta) \in R'^2$ and $m' = (-\bar{\beta}, \bar{\alpha}) \in R'^2$. Let $K_1 := (K_A, \iota K_A)$ and $K_2 = \omega K_1$ generate $\ker W'(\varphi_A)$. Then for $i = 1, 2$ we get*

$$[m]K_i = 0, \qquad \textit{and} \qquad [m']K_i = 0.$$

Proof. This proof is now almost trivial! Let us compute $[m]K_1$:

$$\begin{aligned}
[m]K_1 &= \alpha K_A + \beta \iota K_A \\
&= [\alpha + \beta\mathbf{i}]K_A \\
&= [\gamma]K_A \\
&= 0,
\end{aligned}$$

and similarly, using that $r\mathbf{i} = \mathbf{i}\bar{r}$ for any $r \in R'$, we get

$$\begin{aligned}
[m']K_1 &= -\bar{\beta}K_A + \bar{\alpha}\iota K_A \\
&= [\mathbf{i}\beta\mathbf{i} + \mathbf{i}\alpha]K_A \\
&= [\mathbf{i}(\beta\mathbf{i} + \alpha)]K_A \\
&= [\mathbf{i}\gamma]K_A \\
&= 0.
\end{aligned}$$

To show the same holds for $K_2 = \omega K_1$, we only have to remember that R' is commutative so that $[m](\omega K_1) = [\omega][m]K_1 = 0$, and vice versa for m'. □

Lemma 4 therefore implies $\ker W'(\varphi_A) \subseteq E_0'^2[M_A]$. As γ is primitive and the reduced norm of $I = \langle 2^e, \gamma \rangle$ is $\mathrm{nrd}(I) = 2^e$, this implies that for the module M_A we get $|E_0'^2[M_A]| = \mathrm{nrd}(I)^2 = 2^{2e}$, following Example 2.4 [BDD+26]. One can show that K_1 and K_2 are linearly independent, and so we get the following result.

Corollary 1. *The MIKE module $M_A \subseteq R'^2$ generated by $m_1 = (2^e, 0)$, $m_2 = (0, 2^e)$, $m_3 = (\alpha, \beta)$, and $m_4 = (-\bar{\beta}, \bar{\alpha})$ ensures $E_0'^2[M_A] = \ker W'(\varphi_A)$.*

Remark 8. The proof of Lemma 4 models $W'(E_0') = E_0' \times E_0'^t$ as $E_0'^2$ using the twisting isomorphism $\iota : E_0'^t \to E_0'$. More generally for any E over $\mathbb{F}_p$ on the surface, we may pick any twisting isomorphism $t : E^t \to E$ which satisfies $\alpha t = t\bar{\alpha}$ for $\alpha \in R'$.

The above results crucially rely on our ability to write γ as $\alpha + \beta\mathbf{i}$, which is derived from the fact that $\mathcal{O}_0' = R' \oplus R'\mathbf{i}$. This is *the* crucial reason to work on the surface with R' and E_0' instead of on the floor with R and E_0. In fact, several key simplifications in MIKE rely on this crucial fact.

3½.3 Computing 4D-isogenies is actually hard

In Level 2, we've described four-dimensional isogenies as simply slapping some formulas on the given kernel points. Although this is true, this is also false. A more accurate statement is that for a given secret $x_A \in \mathbb{Z}_{2^e}$, where $2^e \mid \frac{p+1}{4}$, and choosing $x_A \equiv 3 \bmod 8$, the formulation of Theorem 4.1 [BDD+26] gives us a *level-2 theta structure* and a basis $\mathcal{B}$ adapted to the isogeny $F_A : W(E_B)^2 \to X_{AB}$. In theory, this is enough to derive the whole computation of F_A.

Theory versus practice. In practice, once you are given this setup $\mathcal{B}$, it is far from easy to derive the right formulas to compute F_A. Luckily, smarter people than us have already done this, in Appendix F [BDD+26]. The crucial difficulties in such a computation are a) figuring out the right formulation of $\mathcal{B}$ in terms of x_A in the first place, and b) figuring out *gluing* steps, i.e., the crucial steps $A \times B \to C$ where we go from two 2-dimensional varieties A, B to a four-dimensional C, or similarly when we go from two 1-dimensional varieties to a two-dimensional variety. Although both steps can be resolved with enough theta coordinates, linear algebra, ink, and determination, as demonstrated in Appendix F, the resulting algorithms are far from easily digestible. In summary, once you are given $\mathcal{B}$ in terms of x_A, it is only a matter of applying the right formulas and algorithms to obtain X_{AB}, but deriving those formulas and algorithms is a Herculean task, and the formulas and algorithms are neither simple, nor intuitive. Alas.

Theta coordinates. The crucial tools used in the computation of such four-dimensional isogenies are level-2 theta structures, briefly summarised as analogues of x-only arithmetic in Level 2. For principally polarised abelian varieties A that are not products, these coordinates give an embedding of A up to $[-1]$, that is, the Kummer variety $A/[-1]$, into $\mathbb{P}^{2^g-1}$ projective space. Hence, in dimension four, this gives an embedding

$$\theta : A/[-1] \hookrightarrow \mathbb{P}^{15}(\mathbb{F}_p), \qquad P \mapsto (\theta_0(P), \ldots, \theta_{15}(P)).$$

The previously mentioned theta null point $t = (a_0 : \ldots : a_{15})$ is then the image of 0_A, i.e., $t = \theta(0_A)$. The derivation of theta coordinates and their arithmetic is a classical subject initiated by Mumford [Mum66] and heavily used in modern isogeny-based cryptography to compute two- and four-dimensional isogenies [DMPR24; DLRW24; Dar25; DD26].

3½.4 The MIKE fourfolds X_{AB}

The fourfolds X_{AB} that we obtain as shared secrets from a key exchange are not generic 4-dimensional (principally polarised) abelian varieties, but live in a 4-dimensional space of rather special fourfolds. In general, the space of fourfolds is 10-dimensional, which means that if one would (and could) pick an abelian fourfold at random, the probability of getting a MIKE fourfold is negligible.

This is a direct result of how X_{AB} is derived in MIKE, and does not seem to pose any threats to the security of MIKE. On the contrary, this 'special locus' of fourfolds contains extra structure that resolves a few crucial problems in the construction of MIKE. The full paper describes this as an abelian fourfold *of Weil type*, and analyses their role in MIKE in Section 4.3 and Appendices G and H. We give a short introduction.

An extra automorphism. The special fourfolds X of Weil type have a surprising amount of extra structure. The main ingredient we need in MIKE is the existence of an extra automorphism $[i] : X \to X$ that behaves a bit like the imaginary unit i, e.g. $[i]^2 = [-1]$, on X. The existence of $[i]$ ensures a significant amount of structure on X: for

example, for the theta null point $t = (a_0 : \ldots : a_{15}) \in \mathbb{P}^{15}(\mathbb{F}_p)$ of a MIKE fourfold X_{AB}, when computed via the isogeny formulas in MIKE, we get the beautiful symmetries:

$$a_1 = a_8,\ a_2 = a_4,\ a_3 = a_{12},\ a_5 = a_{10},\ a_7 = a_{14},\ a_{11} = a_{13}.$$

Invariants. This symmetry not only explains the absence of $a_4, a_8, a_{10}, a_{12}, a_{13}$, and a_{14} in the invariants described in Definitions 2 and 3, but in fact, the extra automorphism $[i]$ on these fourfolds is what allows us to compute invariants in the first place! For general (principally polarised) abelian fourfolds, we do not know of such invariants.

The derivation of these invariants for the MIKE fourfolds, performed in Appendix H, is a spectacular piece of mathematics and fairly digestible if one is familiar with generic techniques to derive invariants.

Computation. The existence of $[i]$ can also improve computations of the four-dimensional isogenies, as demonstrated in Appendix E.2. The crucial observation here is that the symmetries induced by $[i]$ can simplify and reduce the required computational steps, mostly because it allows us to only keep track of two kernel generators, whereas a generic 4D-isogeny would require us to keep track of four kernel generators.

It seems possible to improve further upon these ideas in a way that is relatively simple to explain: we have explained that our computations use the analogue of x-only arithmetic by using level-2 theta coordinates. In a more abstract description, x-only arithmetic can be explained as working with the *Kummer variety* $E/[-1]$, in which we identify P with $-P$ for any $P \in E$. That is, we work modulo the automorphism $[-1] \in \mathrm{Aut}(E)$, which any elliptic curve or abelian variety has. Level-2 theta coordinates are similarly tools to do computations on the Kummer variety $A/[-1]$ for higher-dimensional abelian varieties A. This trick generalises for more automorphisms than just $[-1]$. In particular, for the MIKE fourfolds X, it should be possible to do all computations on $X/[i]$ using specialised arithmetic, instead of on $X/[-1]$ using level-2 theta arithmetic.

3½.5 The missing torsion

In MIKE, we define a prime $p = c2^f - 1$ of appropriate size, with a small odd cofactor c, and compute 2^e-isogenies in dimension one and four, respectively, with the assumption that $2^e \mid \frac{p+1}{4}$. That is, we are restricted to $e \leq f - 2$. This is very helpful in several directions:

- Firstly, we saw that writing down formulas for the four-dimensional isogenies in MIKE was a heroic effort. It's not surprising that we had to offer a sacrifice to the gods for making this possible, in the form of additional 2-torsion. More concretely, to compute 2-isogeny steps in dimension four in theta arithmetic, we actually need points of order 8: in addition to kernel points P_i of order 2, we also keep 8-torsion points above these, i.e., points P_i' with $[4]P_i' = P_i$. This additional information allows us to recover the codomain of each step without expensive square-roots and ambiguous sign choices. To have this extra torsion available all the way until the last step, we therefore want a setup that guarantees us to have 2^{e+2}-torsion available. We won't go into detail here, as we don't want to spoil the fun of reading Appendices E and F of [BDD+26].

- Similarly, during key generation (i.e., for the one-dimensional isogeny φ), Alice first computes a point K such that $[4]K$ is the kernel generator of φ of order 2^e. Pushing this point through φ, she obtains a point $\varphi(K) \in E_A[4]$. Such a point of order four allows us to write down the Montgomery coefficients of all equivalent public

keys essentially for free,[29] and pick a normalised representative, for instance the lexicographically largest one. Without such a point, we would require expensive square-root computations for this step.

Note that normalising E_A is not required for the key exchange to work, but is useful in practice, for example to obtain key binding properties. Since for the reasons above we require that extra 2-torsion is available anyway, it is worthwhile to add this normalisation step at such little cost.

- In several parts in the proofs of Level 3, we used that $\omega \cdot \bar{\omega} = \frac{p+1}{4}$, and so, that $[\omega \cdot \bar{\omega}](P) = \mathcal{O}$ and $[\omega^2]P = [\omega - \frac{p+1}{4}](P) = [\omega]P$ for any P of order $2^e \mid \frac{p+1}{4}$. This translates again to $e \leq f - 2$. Thus, such a choice for e allows us to choose a 2^e-basis (P, Q) with the very nice properties that $[\omega](P) = P$ and $[\omega](Q) = \mathcal{O}$. This significantly simplifies a lot of the computational steps involving ω, P, and Q.

Besides these computational advantages, there is another layer to the discussion on how to choose the exponent e: the two best known avenues to attack MIKE depend on the size of the field characteristic p and the parameter e, respectively. The former determines the hardness of the endomorphism ring problem—the main computational problem underlying isogeny-based cryptography. The latter determines the length of the secret isogeny chains, and therefore the complexity of recovering these isogenies through meet-in-the-middle approaches. This is very similar to the situation in SIKE, where meet-in-the-middle attacks were extensively studied to pick suitable isogeny lengths.

In summary for the choice of e, this tells us that for a given prime $p = c2^f - 1$, we get the upper bound of choosing $e \leq f - 2$, but we can go lower than that, as long as we maintain sufficient security against meet-in-the-middle attacks. For a detailed discussion on choices of p and e in the light of the recent improvements on solving the endomorphism ring problem by Wesolowski [Wes26], we refer to Appendix J of [BDD+26].

[29] In MIKE, any of the six possible A-coefficients of E_A, and the six possible A-coefficients of its conjugate $E_A^{(p)}$ are equivalent.

References

[AAA+26] Marius A. Aardal, Gora Adj, Diego F. Aranha, Andrea Basso, Giacomo Borin, Isaac Andrés Canales Martínez, Jorge Chávez-Saab, Maria Corte-Real Santos, Pierrick Dartois, Luca De Feo, Max Duparc, Thomas Espitau, Jonathan Komada Eriksen, Tako Boris Fouotsa, Décio Luiz Gazzoni Filho, Basil Hess, Riccardo Invernizzi, David Kohel, Antonin Leroux, Patrick Longa, Luciano Maino, Michael Meyer, Marzio Mula, Kohei Nakagawa, Hiroshi Onuki, Lorenz Panny, Sikhar Patranabis, Christophe Petit, Giacomo Pope, Krijn Reijnders, Damien Robert, Francisco Rodríguez-Henríquez, Sina Schaeffler, Frederik Vercauteren, Alexandre Wallet, Benjamin Wesolowski, and Wessel van Woerden. SQIsign. Technical report, National Institute of Standards and Technology, 2026. URL: https://sqisign.org.

[Adl79] Leonard M. Adleman. A subexponential algorithm for the discrete logarithm problem with applications to cryptography. In *20th Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 29-31 October 1979*, pages 55–60. IEEE Computer Society, 1979. DOI: 10.1109/SFCS.1979.2. URL: https://doi.org/10.1109/SFCS.1979.2.

[BDD+26] Andrea Basso, Pierrick Dartois, Max Duparc, Jonathan Komada Eriksen, Sabrina Kunzweiler, Michael Meyer, Giacomo Pope, Krijn Reijnders, Damien Robert, Ryan Rueger, and Sina Schaeffler. MIKE: a fast and compact post-quantum NIKE. Cryptology ePrint Archive, Paper 2026/2377, 2026. URL: https://eprint.iacr.org/2026/2377.

[BDL+26] Xavier Bonnetain, Max Duparc, Dania Lazzarini, Luciano Maino, Chloe Martindale, Christophe Petit, Sina Schaeffler, André Schrottenloher, and Alessandro Sferlazza. Quantum security analysis of unrestricted isogeny-based group actions. Cryptology ePrint Archive, Paper 2026/2215, 2026. URL: https://eprint.iacr.org/2026/2215.

[BY90] Gilles Brassard and Moti Yung. One-way group actions. In Alfred Menezes and Scott A. Vanstone, editors, *Advances in Cryptology - CRYPTO '90, 10th Annual International Cryptology Conference, Santa Barbara, California, USA, August 11-15, 1990, Proceedings*, volume 537 of *Lecture Notes in Computer Science*, pages 94–107. Springer, 1990. DOI: 10.1007/3-540-38424-3_7. URL: https://doi.org/10.1007/3-540-38424-3_7.

[CD20] Wouter Castryck and Thomas Decru. CSIDH on the surface. In Jintai Ding and Jean-Pierre Tillich, editors, *Post-Quantum Cryptography - 11th International Conference, PQCrypto 2020, Paris, France, April 15-17, 2020, Proceedings*, volume 12100 of *Lecture Notes in Computer Science*, pages 111–129. Springer, 2020. DOI: 10.1007/978-3-030-44223-1_7. URL: https://doi.org/10.1007/978-3-030-44223-1_7.

[CD23] Wouter Castryck and Thomas Decru. An efficient key recovery attack on SIDH. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology - EUROCRYPT 2023 - 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23-27, 2023, Proceedings, Part V*, volume 14008 of *Lecture Notes in Computer Science*, pages 423–447. Springer, 2023. DOI: 10.1007/978-3-031-30589-4_15. URL: https://doi.org/10.1007/978-3-031-30589-4_15.

[CDS20] Wouter Castryck, Thomas Decru, and Benjamin Smith. Hash functions from superspecial genus-2 curves using richelot isogenies. *Journal of Mathematical Cryptology*, 14(1):268–292, 2020.

[CH17] Craig Costello and Hüseyin Hisil. A simple and compact algorithm for SIDH with arbitrary degree isogenies. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part II*, volume 10625 of *Lecture Notes in Computer Science*, pages 303–329. Springer, 2017. DOI: 10.1007/978-3-319-70697-9_11. URL: https://doi.org/10.1007/978-3-319-70697-9_11.

[CK20] Leonardo Colo and David Kohel. Orienting supersingular isogeny graphs. *Journal of Mathematical Cryptology*, 14(1):414–437, 2020.

[CLM+18] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: an efficient post-quantum commutative group action. In Thomas Peyrin and Steven D. Galbraith, editors, *Advances in Cryptology - ASIACRYPT 2018 - 24th International Conference on the Theory and Application of Cryptology and Information Security, Brisbane, QLD, Australia, December 2-6, 2018, Proceedings, Part III*, volume 11274 of *Lecture Notes in Computer Science*, pages 395–427. Springer, 2018. ISBN: 978-3-030-03331-6. DOI: 10.1007/978-3-030-03332-3_15. URL: https://doi.org/10.1007/978-3-030-03332-3_15.

[CLR26] Maria Corte-Real Santos, Gioella Lorenzon, and Krijn Reijnders. Supersingularity and superspeciality verification for abelian surfaces. Cryptology ePrint Archive, Paper 2026/2259, 2026. URL: https://eprint.iacr.org/2026/2259.

[Cos19] Craig Costello. Supersingular isogeny key exchange for beginners. In Kenneth G. Paterson and Douglas Stebila, editors, *Selected Areas in Cryptography - SAC 2019 - 26th International Conference, Waterloo, ON, Canada, August 12-16, 2019, Revised Selected Papers*, volume 11959 of *Lecture Notes in Computer Science*, pages 21–50. Springer, 2019. DOI: 10.1007/978-3-030-38471-5_2. URL: https://doi.org/10.1007/978-3-030-38471-5_2.

[Cou06] Jean-Marc Couveignes. Hard homogeneous spaces. Cryptology ePrint Archive, Paper 2006/291, 2006. URL: https://eprint.iacr.org/2006/291.

[Dar25] Pierrick Dartois. *Fast computation of higher dimensional isogenies for cryptographic applications*. PhD thesis, Université de Bordeaux, France, July 2025. URL: https://theses.fr/s364682.

[DD26] Pierrick Dartois and Max Duparc. Chasing rabbits through hypercubes: better algorithms for higher dimensional 2-isogeny computations. Cryptology ePrint Archive, Paper 2026/114, 2026. URL: https://eprint.iacr.org/2026/114.

[DFK+23] Luca De Feo, Tako Boris Fouotsa, Péter Kutas, Antonin Leroux, Simon-Philipp Merz, Lorenz Panny, and Benjamin Wesolowski. SCALLOP: scaling the CSI-FiSh. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *Public-Key Cryptography - PKC 2023 - 26th IACR International Conference on Practice and Theory of Public-Key Cryptography, Atlanta, GA, USA, May 7-10, 2023, Proceedings, Part I*, volume 13940 of *Lecture Notes in Computer Science*, pages 345–375. Springer, 2023. DOI: 10.1007/978-3-031-31368-4_13. URL: https://doi.org/10.1007/978-3-031-31368-4_13.

[DH76] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Trans. Inf. Theory*, 22(6):644–654, 1976. DOI: 10.1109/TIT.1976.1055638. URL: https://doi.org/10.1109/TIT.1976.1055638.

[DLRW24] Pierrick Dartois, Antonin Leroux, Damien Robert, and Benjamin Wesolowski. SQIsignHD: new dimensions in cryptography. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology - EUROCRYPT 2024 - 43rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Zurich, Switzerland, May 26-30, 2024, Proceedings, Part I*, volume 14651 of *Lecture Notes in Computer Science*, pages 3–32. Springer, 2024. DOI: 10.1007/978-3-031-58716-0_1. URL: https://doi.org/10.1007/978-3-031-58716-0_1.

[DMPR24] Pierrick Dartois, Luciano Maino, Giacomo Pope, and Damien Robert. An algorithmic approach to (2, 2)-isogenies in the theta model and applications to isogeny-based cryptography. In Kai-Min Chung and Yu Sasaki, editors, *Advances in Cryptology - ASIACRYPT 2024 - 30th International Conference on the Theory and Application of Cryptology and Information Security, Kolkata, India, December 9-13, 2024, Proceedings, Part III*, volume 15486 of *Lecture Notes in Computer Science*, pages 304–338. Springer, 2024. DOI: 10.1007/978-981-96-0891-1_10. URL: https://doi.org/10.1007/978-981-96-0891-1_10.

[FT19] E Victor Flynn and Yan Bo Ti. Genus two isogeny cryptography. In *International Conference on Post-Quantum Cryptography*, pages 286–306. Springer, 2019.

[HLP00] Everett W. Howe, Franck Leprévost, and Bjorn Poonen. Large torsion subgroups of split Jacobians of curves of genus two or three. *Forum Mathematicum*, 12(3):315–364, 2000. DOI: 10.1515/form.2000.008. URL: https://doi.org/10.1515/form.2000.008.

[Kob87] Neal Koblitz. Elliptic curve cryptosystems. *Mathematics of computation*, 48(177):203–209, 1987.

[Kup05] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *SIAM J. Comput.*, 35(1):170–188, 2005. DOI: 10.1137/S0097539703436345. URL: https://doi.org/10.1137/S0097539703436345.

[Ler22] Antonin Leroux. *Quaternion Algebra and isogeny-based cryptography*. PhD thesis, Ecole doctorale de l'Institut Polytechnique de Paris, 2022.

[Mil85] Victor S. Miller. Use of elliptic curves in cryptography. In Hugh C. Williams, editor, *Advances in Cryptology - CRYPTO '85, Santa Barbara, California, USA, August 18-22, 1985, Proceedings*, volume 218 of *Lecture Notes in Computer Science*, pages 417–426. Springer, 1985. DOI: 10.1007/3-540-39799-X_31. URL: https://doi.org/10.1007/3-540-39799-X_31.

[Mum66] David Mumford. On the Equations Defining Abelian Varieties. I. *Inventiones Mathematicae*, 1:287–354, 1966. DOI: 10.1007/BF01389737.

[Pop26] Giacomo Pope. A String Theorist's Guide to Implementing MIKE. https://giacomopope.com/mike/, 2026. [Accessed 09-10-2026].

[Ren18] Joost Renes. Computing Isogenies Between Montgomery Curves Using the Action of (0, 0). In Tanja Lange and Rainer Steinwandt, editors, *Post-Quantum Cryptography - 9th International Conference, PQCrypto 2018, Fort Lauderdale, FL, USA, April 9-11, 2018, Proceedings*, volume 10786 of *Lecture Notes in Computer Science*, pages 229–247. Springer, 2018. DOI: 10.1007/978-3-319-79063-3_11. URL: https://doi.org/10.1007/978-3-319-79063-3_11.

[Rob24] Damien Robert. The module action for isogeny based cryptography. Cryptology ePrint Archive, Paper 2024/1556, 2024. URL: https://eprint.iacr.org/2024/1556.

[RS06] Alexander Rostovtsev and Anton Stolbunov. Public-key cryptosystem based on isogenies. Cryptology ePrint Archive, Paper 2006/145, 2006. URL: https://eprint.iacr.org/2006/145.

[Sho94] Peter W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *35th Annual Symposium on Foundations of Computer Science, Santa Fe, New Mexico, USA, November 20-22, 1994*, pages 124–134. IEEE Computer Society, 1994. DOI: 10.1109/SFCS.1994.365700. URL: https://doi.org/10.1109/SFCS.1994.365700.

[Sil09] Joseph H. Silverman. *The Arithmetic of Elliptic Curves*, volume 106 of *Graduate Texts in Mathematics*. Springer, 2nd edition, 2009. DOI: 10.1007/978-0-387-09494-6.

[Sto10] Anton Stolbunov. Constructing public-key cryptographic schemes based on class group action on a set of isogenous elliptic curves. *Adv. Math. Commun.*, 4(2):215–235, 2010.

[Vél71] Jacques Vélu. Isogénies entre courbes elliptiques. *Comptes-Rendus de l'Académie des Sciences*, 273:238–241, 1971.

[Voi21] John Voight. *Quaternion Algebras*, volume 288 of *Graduate Texts in Mathematics*. Springer, 2021. DOI: 10.1007/978-3-030-56694-4.

[Wes26] Benjamin Wesolowski. The supersingular isogeny problem in time and memory $p^{1/3+o(1)}$. Cryptology ePrint Archive, Paper 2026/1486, 2026. URL: https://eprint.iacr.org/2026/1486.