MULTIPARTY NON-INTERACTIVE KEY EXCHANGE AND MORE FROM ISOGENIES ON ELLIPTIC CURVES

DAN BONEH[1], DARREN GLASS[2], DANIEL KRASHEN[3], KRISTIN LAUTER[4], SHAHED SHARIF[5], ALICE SILVERBERG[6], MEHDI TIBOUCHI[7], AND MARK ZHANDRY[8]

Abstract. We describe a framework for constructing an efficient non-interactive key exchange (NIKE) protocol for n parties for any $n \geq 2$. Our approach is based on the problem of computing isogenies between isogenous elliptic curves, which is believed to be difficult. We do not obtain a working protocol because of a missing step that is currently an open mathematical problem. What we need to complete our protocol is an efficient algorithm that takes as input an abelian variety presented as a product of isogenous elliptic curves, and outputs an isomorphism invariant of the abelian variety.

Our framework builds a *cryptographic invariant map*, which is a new primitive closely related to a cryptographic multilinear map, but whose range does not necessarily have a group structure. Nevertheless, we show that a cryptographic invariant map can be used to build several cryptographic primitives, including NIKE, that were previously constructed from multilinear maps and indistinguishability obfuscation.

1. Introduction

Let $\mathbb{F}_q$ be a finite field, let E be an ordinary elliptic curve over $\mathbb{F}_q$, and let X be the set of elliptic curves over $\mathbb{F}_q$ that are isogenous to E. The set X is almost always large (containing on the order of $\sqrt{q}$ elements). Moreover, under suitable conditions on E, the set X is endowed with a free and transitive action $*$ by a certain abelian group G, which is the ideal class group of the endomorphism ring of E. The action $*$ maps a given $g \in G$ and $E \in X$ to a curve $g * E \in X$.

This action, originally defined by Deuring [Deu41], has a number of properties that makes it useful in cryptography. First, for a fixed curve $E \in X$, the map $G \to X$ defined by $g \mapsto g * E$ is believed to be a one-way function. In other words, given a random curve $E' \in X$ it is difficult to find an element $g \in G$ such that $E' = g * E$. This suggests a Diffie–Hellman two-party key exchange protocol, proposed by Couveignes [Cou06] and Rostovtsev

[1] Stanford University
[2] Gettysburg College
[3] University of Georgia, Athens
[4] Microsoft Research
[5] California State University San Marcos
[6] University of California, Irvine
[7] NTT Corporation
[8] Princeton University

E-mail addresses: dabo@cs.stanford.edu, dglass@gettysburg.edu, daniel.krashen@gmail.com, klauter@microsoft.com, ssharif@csusm.edu, asilverb@uci.edu, tibouchi.mehdi@lab.ntt.co.jp, mzhandry@princeton.edu.

2010 *Mathematics Subject Classification.* Primary 14K02; Secondary 14Q20, 11Y16, 94A60.

Key words and phrases. Multilinear maps, Non-Interactive Key Exchange, Isogenies, Witness Encryption, Abelian Varieties.

and Stolbunov [RS06]: Alice chooses a random $a \in G$ and publishes $E_a := a * E$; Bob chooses a random $b \in G$ and publishes $E_b := b * E$. Their shared key is the curve

$$E_{ab} := (ab) * E = a * E_b = b * E_a,$$

which they can both compute. To ensure that both parties obtain the same key, their shared key is the j-invariant of the curve E_{ab}. More recently, De Feo, Jao, and Plût [FJP14], Galbraith [Gal18], Castryck et al. [CLM$^+$18], and De Feo, Kieffer, and Smith [FKS18] proposed variants of this protocol with better security and efficiency. Moreover, a supersingular version of the isogeny problem was introduced and proposed as the basis for a collision resistant hash function [CLG09]. Security of this one-way function was further studied in [EHL$^+$18].

Second, as alluded to above, the star operator satisfies the following useful property: for all $g_1, \ldots, g_n \in G$ the abelian varieties

$$A_1 := (g_1 * E) \times \cdots \times (g_n * E) \quad \text{and} \quad A_2 := (g_1 \cdots g_n) * E \times E^{n-1}$$

are isomorphic (see Appendix A.4). As we will see in the next section, this suggests an n-party non-interactive key exchange protocol, as well as many other cryptographic constructions. This property leads to a more general cryptographic primitive that we call a **cryptographic invariant map**, defined in the next section. This primitive has properties that are similar to those of cryptographic multilinear maps [BS03, GGH13a], which have found numerous applications in cryptography (e.g, [GGH$^+$13b, GGSW13, BW13, BLR$^+$15]). We discuss applications of cryptographic invariant maps in Section 3. In Remark 4.4 we explain why we use ordinary and not supersingular elliptic curves. Section 4 describes our approach to constructing cryptographic invariant maps from isogenies.

This work leads to the following question in algebraic geometry.

An open problem. To make the cryptographic applications discussed above viable we must first overcome an important technical challenge. While the varieties A_1 and A_2 defined above are isomorphic, they are presented differently. Our applications require an efficient way to compute an invariant that is the same for A_1 and A_2. In addition, the invariant must distinguish non-isomorphic varieties. We do not know any such computable isomorphism invariant, and we present this as an open problem. In Section 5 we explain why some natural proposals for isomorphism invariants do not seem to work. In Remarks 2.4 and 4.2 we show that a solution to this open problem, even for $n = 2$, would solve the isogeny decision Diffie–Hellman problem. Further, we give evidence that computing a particular isomorphism invariant might be equivalent to solving the elliptic curve isogeny problem, which is believed (or hoped) to be a quantum-resistant hard problem. Thus, Section 5 might be useful from the point of view of cryptanalysis of isogeny-based cryptography.

2. Cryptographic invariant maps

Definition 2.1. Let X be a finite set and let G be a finite abelian group. We say that G **acts efficiently on X freely and transitively** if there is an efficiently computable map $* : G \times X \to X$ such that:

- the map is a group action: $g * (h * x) = (gh) * x$, and there is an identity element $\mathrm{id} \in G$ such that $\mathrm{id} * x = x$, for all $x \in X$ and all $g, h \in G$;
- the action is transitive: for every $(x, y) \in X \times X$ there is a $g \in G$ such that $g * x = y$; and
- the action is free: if $x \in X$ and $g, h \in G$ satisfy $g * x = h * x$, then $g = h$.

Definition 2.2. By a **cryptographic invariant map** we mean a randomized algorithm MapGen that inputs a security parameter λ, outputs public parameters $\mathsf{pp} = (X, S, G, e)$, and runs in time polynomial in λ, where:

- X and S are sets, and X is finite,
- G is a finite abelian group that acts efficiently on X freely and transitively,
- e is a deterministic algorithm that runs in time polynomial in λ and n, such that for each $n > 0$, algorithm e takes λ as input and computes a map $e_n : X^n \to S$ that satisfies:
  - **Invariance property** of e_n: for all $x \in X$ and $g_1, \ldots, g_n \in G$,
  $$e_n(g_1 * x, \ldots, g_n * x) = e_n\big((g_1 \cdots g_n) * x, x, \ldots, x\big);$$
  - **Non-degeneracy** of e_n: for all i with $1 \le i \le n$ and $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n \in X$, the map $X \to S$ defined by $y \mapsto e_n(x_1, \ldots, x_{i-1}, y, x_{i+1}, \ldots, x_n)$ is injective.

In our candidate instantiation for cryptographic invariant maps the set X is a set of isogenous elliptic curves and the group G acting on X is a class group. The elements of S are isomorphism invariants of products of elliptic curves.

Definition 2.2 is quite ambitious in that it asks that e_n be defined for all $n > 0$ and run in polynomial time in n (and λ). A cryptographic invariant map that is defined even for a single $n > 2$, and satisfies the security assumptions in the next subsection, would still be quite interesting. We require a construction that works for all n because our framework using elliptic curve isogenies seems to support it. Similarly, we note that a construction that works for all $n > 0$, but runs in time exponential in n is still useful. It would limit our ability to evaluate e_n to relatively small n, but that is still of great interest. In the first three proposals in Section 5 we study candidates for e_n that run in time exponential in n, satisfy the non-degeneracy property, but do not satisfy the invariance property. It is an open problem to find a map that also satisfies the invariance property.

2.1. **Security assumptions.** Next, we define some security assumptions on cryptographic invariant maps. The notation $x \xleftarrow{R} X$ will denote an independent uniform random variable x over the set X. Similarly, we use $x' \xleftarrow{R} A(y)$ to define a random variable x' that is the output of a randomized algorithm A on input y.

The n-way computational Diffie–Hellman assumption states that, given only the public parameters and $(g_1 * x, \ldots, g_n * x) \in X^n$, it is difficult to compute $e_{n-1}\big((g_1 \cdots g_n) * x, x, \ldots, x\big)$. A precise definition is the following:

Definition 2.3. We say that MapGen satisfies the n-way **computational Diffie–Hellman assumption (n-CDH)** if for every polynomial time algorithm $\mathcal{A}$,

$$\Pr\Big[\mathcal{A}(\mathsf{pp},\ g_1 * x, \ldots, g_n * x) = e_{n-1}\big((g_1 \cdots g_n) * x, x, \ldots, x\big)\Big]$$

is a negligible function of λ, when $\mathsf{pp} \xleftarrow{R} \mathsf{MapGen}(\lambda)$, $g_1, \ldots, g_n \xleftarrow{R} G$, and $x \xleftarrow{R} X$.

Remark 2.4. The natural n-way decision Diffie–Hellman assumption on X does not hold when invariant maps exist. That is, for all $n > 0$ it is easy to distinguish $(g_1 \cdots g_n) * x \in X$ from a random element of X, given only $x, g_1 * x, \ldots, g_n * x$. Given a challenge $y \in X$, simply check if

$$e_n(y, x, \ldots, x) = e_n(g_1 * x, \ldots, g_n * x).$$

Equality holds if and only if $y = (g_1 \cdots g_n) * x$. However, in Definition 2.5 we define an n-way decision Diffie–Hellman assumption for e_{n-1}. It states that it is hard to distinguish

$e_{n-1}\big((g_1\cdots g_n)*x,x,\ldots,x\big)$ from a random element in the image of e_{n-1}, given only the public parameters, x, and $(g_1*x,\ldots,g_n*x)\in X^n$.

Definition 2.5. We say that MapGen satisfies the n-way **decision Diffie–Hellman assumption (n-DDH)** if the following two distributions, $\mathcal{P}_0$ and $\mathcal{P}_1$, are polynomially indistinguishable, when $\mathsf{pp}\overset{\mathrm{R}}{\leftarrow}\mathsf{MapGen}(\lambda)$, $g_1,\ldots,g_n\overset{\mathrm{R}}{\leftarrow}G$, and $x\overset{\mathrm{R}}{\leftarrow}X$:

- $\mathcal{P}_0$ is $(\mathsf{pp},\ g_1*x,\ldots,g_n*x,\ s_0)$ where $s_0=e_{n-1}\big((g_1\cdots g_n)*x,x,\ldots,x\big)$,
- $\mathcal{P}_1$ is $(\mathsf{pp},\ g_1*x,\ldots,g_n*x,\ s_1)$ where s_1 is random in $\mathrm{Im}(e_{n-1})\subseteq S$.

3. Applications

We show that suitable cryptographic invariant maps can be used to solve a number of important problems in cryptography.

n-way Non-Interactive Key Exchange (NIKE). We show how to use a cryptographic invariant map to construct a Non-Interactive Key Exchange (NIKE) protocol in which n parties create a shared secret key that only they can efficiently calculate, without any interaction among the n parties. Currently, secure n-party NIKE for $n>3$ is only known from general purpose indistinguishability obfuscation (e.g., [BZ17]). Our NIKE construction is similar to the one in [Jou04, BS03, GGH13a] and satisfies a “static” notion of security.

- Setup(λ): run $(X,S,G,e)\overset{\mathrm{R}}{\leftarrow}\mathsf{MapGen}(\lambda)$ and choose $x\overset{\mathrm{R}}{\leftarrow}X$. Output $\mathsf{pp}:=(X,S,G,e,x)$.
- For $i=1,\ldots,n$, party i chooses a random $g_i\overset{\mathrm{R}}{\leftarrow}G$, computes $x_i:=g_i*x\in X$, and publishes x_i on a public bulletin board.
- The shared key between the n-parties is

$$k:=e_{n-1}\big((g_1\cdots g_n)*x,x,\ldots,x\big)\in S.$$

Party $i\in\{1,\ldots,n\}$ computes k by obtaining $x_1,\ldots,x_n$ from the bulletin board, then choosing some $j\in\{1,\ldots,n\}$ where $j\neq i$, and computing

$$k=e_{n-1}(x_1,\ldots,x_{j-1},\ g_i*x_j,\ x_{j+1},\ldots,x_n)\in S,$$

where x_i is omitted from the input to e_{n-1}.

All n parties obtain the same key k by the invariance property of e_{n-1}. Static security follows from the n-way decision Diffie–Hellman assumption, as in [BS03]. Alternatively, we can rely on the weaker n-way computational Diffie–Hellman assumption by applying a hash function $H:S\to K$ to the key k. We model H as a random oracle in the security analysis. We leave the question of an adaptively-secure NIKE, in the sense of [FHKP13, Rao14], from an invariant map for future work.

Unique signatures and verifiable random functions (VRF). A digital signature scheme is made up of three algorithms: a key generation algorithm that outputs a public key and a secret key, a signing algorithm that signs a given message using the secret key, and a verification algorithm that verifies a signature on a given message using the public key. A signature scheme is a **unique signature scheme** if for every public key and every message, there is at most one signature that will be accepted as a valid signature for that message under the public key. While a number of unique signature schemes are known in the random oracle model (e.g., [BR96, BLS01]), it is quite hard to construct unique signatures without random oracles [Lys02, DY05]. Unique signatures are closely related to a simpler object called a verifiable random function, or VRF [MRV99]. Previous results show how to construct unique signatures and VRFs from multilinear maps without random oracles [BS03]. The same constructions work with a cryptographic invariant map. The unique signature scheme works as follows: The secret key is a random $(g_{1,0}, g_{1,1}, \ldots, g_{n,0}, g_{n,1}) \xleftarrow{\mathrm{R}} G^{2n}$. The public key is $(x, y_{1,0}, \ldots, y_{n,1}) \in X^{2n+1}$ where $x \xleftarrow{\mathrm{R}} X$ and $y_{i,b} := g_{i,b} * x$ for $i = 1, \ldots, n$ and $b = 0, 1$. The signature on an n-bit message $m \in \{0,1\}^n$ is $\sigma := (\prod_{i=1}^{n} g_{i,m_i}) * x \in X$. To verify a signature σ, check that $e_n(\sigma, x, \ldots, x) = e_n\big(y_{1,m_1}, \ldots, y_{n,m_n}\big)$. The security analysis of this construction is the same as in [BS03].

Constrained PRFs and broadcast encryption. We next describe how to construct *constrained pseudorandom functions* [BW13, KPTZ13, BGI14] for *bit-fixing constraints* from a cryptographic invariant map. Such constrained PRFs in turn can be used to build broadcast encryption with short ciphertexts [BW13].

A pseudorandom function (PRF) is a function $F : \mathcal{K} \times \mathcal{A} \to \mathcal{B}$ that is computable in polynomial time. Here, $\mathcal{K}$ is the key space, $\mathcal{A}$ is the domain, and $\mathcal{B}$ is the codomain. Intuitively, PRF security requires that, for a random key $k \in \mathcal{K}$, an adversary who obtains pairs $\big(a,\ F(k,a)\big)$, for $a \in \mathcal{A}$ of its choice, cannot distinguish these pairs from pairs $\big(a, f(a)\big)$ where f is a random function $\mathcal{A} \to \mathcal{B}$.

A *bit-fixing constrained PRF* is a PRF where a key $k \in \mathcal{K}$ can be constrained to only evaluate the PRF on a subset of the domain $\mathcal{A}$, where $\mathcal{A} = \{0,1\}^n$. Specifically, for $V \subseteq [n] = \{1, \ldots, n\}$ and a function $v : V \to \{0,1\}$, let $\mathcal{A}_v = \{a \in \mathcal{A} : \forall i \in V, a_i = v(i)\}$. A constrained key k_v enables one to evaluate $F(k,a)$ for all $a \in \mathcal{A}_v$, but reveals nothing about $F(k,a)$ for $a \notin \mathcal{A}_v$. We refer to [BW13] for the complete definition of this concept, and its many applications.

We now explain how to construct bit-fixing constrained PRFs from cryptographic invariant maps. The construction and security proof are essentially the same as in Boneh and Waters [BW13], but translated to our setting. One complication is that the construction of Boneh and Waters requires a way to operate on invariants in S. We get around this by delaying the evaluation of the invariant to the very last step. We thus obtain the following bit-fixing constrained PRF:

- $\mathsf{Setup}(\lambda)$: run $(X, S, G, e) \xleftarrow{\mathrm{R}} \mathsf{MapGen}(\lambda)$ and choose $x \xleftarrow{\mathrm{R}} X$.
  Next choose $\alpha \xleftarrow{\mathrm{R}} G$ and $d_{i,b} \xleftarrow{\mathrm{R}} G$ for $i \in [n]$ and $b \in \{0,1\}$.
  Output the key $k = (X, S, G, e, \alpha, \{d_{i,b}\}_{i,b})$.
- The PRF is defined as: $F(k,a) = e_n\big(\ (\alpha \times \prod_{i=1}^{n} d_{i,a_i}) * x,\ x,\ \ldots,\ x\big)$.
  Here, $a \in \{0,1\}^n$ specifies a subset product of the set of $d_{i,b}$'s.
- $\mathsf{Constrain}(k, v)$: Let $V \subseteq [n]$ be the support of the function v, and assume V is not empty. The constrained key k_v is constructed as follows. Set $D_{i,b} = d_{i,b} * x$ for $i \notin V$. Let i_0 be the smallest element of V. Choose $|V| - 1$ random $g_i \in G$ for $i \in V \setminus \{i_0\}$, and set $g_{i_0} = \alpha \times \prod_{i \in V} d_{i,v_i} \times (\prod_{i \in V \setminus \{i_0\}} g_i)^{-1} \in G$. Let $h_i = g_i * x$ for $i \in V$.
  The constrained key is $k_v = \big(\{D_{i,b}\}_{i \notin V, b \in \{0,1\}},\ \{h_i\}_{i \in V}\big)$.
- $\mathsf{Eval}(k_v, a)$: To evaluate $F(k,a)$ using the constrained key k_v do the following. If $a \notin \mathcal{A}_v$, output $\diamond$. Otherwise, for $i = 1, \ldots, n$, let $C_i = D_{i,a_i}$ if $i \notin V$, and let $C_i = h_i$ otherwise. Output $e_n(C_1, \ldots, C_n)$. Then, by construction,

  $$e_n(C_1, \ldots, C_n) = e_n\left(\left(\prod_{i \notin V} d_{i,a_i} \prod_{i \in V} g_i\right) * x, x, \ldots, x\right) = F(k,a),$$

  as required.

The security proof for this construction is as in [BW13]. This construction can be further extended to a verifiable random function (VRF) by adapting Fuchsbauer [Fuc14] in a similar fashion.

Witness encryption. Witness encryption, due to Garg et al. [GGSW13], can be used to construct Identity-Based Encryption, Attribute-Based Encryption, broadcast encryption [Zha16], and secret sharing for NP statements. Witness encryption is a form of encryption where a public key is simply an NP statement, and a secret key is a witness for that statement. More precisely, a witness encryption scheme is a pair of algorithms:

- $\mathsf{Enc}(x, m)$ is a randomized polynomial-time algorithm that takes as input an NP statement x and a message m, and outputs a ciphertext c;
- $\mathsf{Dec}(x, w, c)$ is a deterministic polynomial-time algorithm that takes as input a statement x, supposed witness w, and ciphertext c, and attempts to produce the message m.

We require that if w is a valid witness for x, then for any message m, if $c \xleftarrow{R} \mathsf{Enc}(x, m)$, then $\mathsf{Dec}(x, w, c)$ outputs m with probability 1.

The basic notion of security for witness encryption is *soundness security*, which requires that if x is false, then $\mathsf{Enc}(x, m)$ hides all information about m. A stronger notion called *extractable security*, due to Goldwasser et al. [GKP+13], requires, informally, that if one can learn any information about m from $\mathsf{Enc}(x, m)$, then it must be the case that one "knows" a witness for x.

We briefly describe how to construct witness encryption from invariant maps. It suffices to give a construction from any NP-complete problem. There are at least two natural constructions from multilinear maps that we can use. One approach is to adapt the original witness encryption scheme of Garg et al. [GGSW13] based on the Exact Cover problem. This approach unfortunately also requires the same graded structure as needed by Boneh and Waters [BW13]. However, we can apply the same ideas as in our constrained PRF construction to get their scheme to work with invariant maps. Another is the scheme of Zhandry [Zha16] based on Subset Sum.[1]

As with the constructions of Garg et al. and Zhandry, the security of these constructions can be justified in an idealized attack model for the cryptographic invariant map, allowing only the operations explicitly allowed by the map—namely the group action and the map operation. Justification in idealized models is not a proof, but provides heuristic evidence for security.

4. Cryptographic invariant maps from isogenies

We begin by recalling some facts that are presented in more detail in Appendix A. Let E be an ordinary elliptic curve over a finite field $\mathbb{F}_q$ such that the ring $\mathbb{Z}[\pi]$ generated by its Frobenius endomorphism π is integrally closed. This implies in particular that $\mathbb{Z}[\pi]$ is the full endomorphism ring $\mathscr{O}$ of E. Let $\mathrm{Cl}(\mathscr{O})$ denote the ideal class group of this ring, and let $\mathrm{Ell}(\mathscr{O})$ denote the isogeny class of E. There exists a free and transitive action $*$ of $\mathrm{Cl}(\mathscr{O})$ on $\mathrm{Ell}(\mathscr{O})$, and there is a way to represent elements of $\mathrm{Cl}(\mathscr{O})$ (namely, as products of prime ideals of small norm) that makes this action efficiently computable. Moreover, one can efficiently sample close to uniform elements in $\mathrm{Cl}(\mathscr{O})$ under that representation. In addition, the "star operator" $*$ satisfies the following property: for any choice of ideal classes $\mathfrak{a}_1, \ldots, \mathfrak{a}_n, \mathfrak{a}'_1, \ldots, \mathfrak{a}'_n$ in $\mathrm{Cl}(\mathscr{O})$, the abelian varieties

$$(\mathfrak{a}_1 * E) \times \cdots \times (\mathfrak{a}_n * E) \quad \text{and} \quad (\mathfrak{a}'_1 * E) \times \cdots \times (\mathfrak{a}'_n * E) \tag{1}$$

[1] The basic scheme shown by Zhandry requires an "asymmetric" multilinear map, where the inputs to the map come from different sets. However, he also explains how to instantiate the scheme using symmetric multilinear maps. The symmetric scheme easily translates to use invariant maps.

are isomorphic over $\mathbb{F}_q$ if and only if $\mathfrak{a}_1 \cdots \mathfrak{a}_n = \mathfrak{a}'_1 \cdots \mathfrak{a}'_n$ in $\mathrm{Cl}(\mathscr{O})$. In particular:

$$(\mathfrak{a}_1 * E) \times \cdots \times (\mathfrak{a}_n * E) \cong (\mathfrak{a}_1 \cdots \mathfrak{a}_n) * E \times E^{n-1}. \tag{2}$$

Denote by $\mathrm{Ab}(E)$ the set of abelian varieties over $\mathbb{F}_q$ that are a product of the form (2), and assume that we can efficiently compute an isomorphism invariant for abelian varieties in $\mathrm{Ab}(E)$. In other words, assume that we have an efficiently computable map $\mathrm{isom}\colon \mathrm{Ab}(E) \to S$ to some set S that to any tuple $E_1, \dots, E_n$ of elliptic curves isogenous to E associates an element $\mathrm{isom}(E_1 \times \cdots \times E_n)$ of S such that $\mathrm{isom}(E_1 \times \cdots \times E_n) = \mathrm{isom}(E'_1 \times \cdots \times E'_n)$ if and only if the products $E_1 \times \cdots \times E_n$ and $E'_1 \times \cdots \times E'_n$ are isomorphic as abelian varieties. The curves E_i are given for example by their j-invariants, and in particular, the ideal classes $\mathfrak{a}_i$ such that $E_i \cong \mathfrak{a}_i * E$ are not supposed to be known.

Based on such an isomorphism invariant isom, we construct a cryptographic invariant map as follows. The algorithm $\mathsf{MapGen}(\lambda)$ computes a sufficiently large base field $\mathbb{F}_q$, and an elliptic curve E over $\mathbb{F}_q$ such that the ring $\mathbb{Z}[\pi]$ generated by its Frobenius endomorphism is integrally closed (this can be done efficiently: see again Appendix A). The algorithm then outputs the public parameters $\mathsf{pp} = (X, S, G, e)$ where:

- $X = \mathrm{Ell}(\mathscr{O})$ is the isogeny class of E over $\mathbb{F}_q$;
- S is the codomain of the isomorphism invariant isom;
- $G = \mathrm{Cl}(\mathscr{O})$ is the ideal class group of $\mathscr{O}$; and
- the map $e_n\colon X^n \to S$ is given by $e_n(E_1, \dots, E_n) = \mathrm{isom}(E_1 \times \cdots \times E_n)$.

The facts recalled at the beginning of this section show that G acts efficiently on X freely and transitively in the sense of Definition 2.1, and that the properties of Definition 2.2 are satisfied. In particular, the invariance property follows from (2), and the non-degeneracy from the fact that the abelian varieties in (1) are isomorphic *only if* the corresponding products of ideal classes coincide. Thus, this approach does provide a cryptographic invariant map assuming isom exists.

Remark 4.1. In the 2-party case, the NIKE protocol obtained from this construction coincides with the isogeny key exchange protocols over ordinary curves described by Couveignes [Cou06] and Rostovtsev–Stolbunov [RS06].

Remark 4.2. The existence of isom breaks the isogeny decision Diffie–Hellman problem. Indeed, given three elliptic curves $(\mathfrak{a} * E, \mathfrak{b} * E, \mathfrak{c} * E)$ isogenous to E, one can check whether $\mathfrak{c} = \mathfrak{a}\mathfrak{b}$ in $\mathrm{Cl}(\mathscr{O})$ by testing whether the surfaces

$$(\mathfrak{a} * E) \times (\mathfrak{b} * E) \quad \text{and} \quad (\mathfrak{c} * E) \times E$$

are isomorphic. This does not prevent the construction of secure NIKE protocols (as those can be based on the computational isogeny Diffie–Hellman problem by applying a hash function: see Section 3), but currently, no efficient algorithm is known for this isogeny decision Diffie–Hellman problem.

Remark 4.3. For certain applications, it would be interesting to be able to hash to the set $X = \mathrm{Ell}(\mathscr{O})$, i.e., construct a random-looking curve E' in the isogeny class of E without knowing an isogeny walk from E to E'. An equivalent problem is to construct a random-looking elliptic curve with exactly $\#E(\mathbb{F}_q)$ points over $\mathbb{F}_q$. This seems difficult, however; the normal way of doing so involves the CM method, which is not efficient when the discriminant is large.

Remark 4.4. One can ask whether this construction extends to the supersingular case. Over $\mathbb{F}_{p^2}$ with p prime, the answer is clearly no, as the isogeny class of a supersingular elliptic curve is not endowed with a natural free and transitive group action by an abelian group. More importantly, isomorphism classes of products of isogenous supersingular elliptic curves over $\mathbb{F}_q$ are essentially trivial at least in a geometric sense. Indeed, according to a result of Deligne (see [Shi78, Theorem 3.5]), if $E_1, \ldots, E_n, E_1', \ldots, E_n'$ are all isogenous to a supersingular elliptic curve E, then

$$E_1 \times \cdots \times E_n \cong E_1' \times \cdots \times E_n' \quad \text{over } \overline{\mathbb{F}_q}$$

as soon as $n \geq 2$. In fact, the result holds over any extension of the base field over which all the endomorphisms of E are defined, so already over $\mathbb{F}_{p^2}$. However, for a supersingular elliptic curve E over a prime field $\mathbb{F}_p$, the number of $\mathbb{F}_p$-isomorphism classes of products $E_1 \times \cdots \times E_n$ with all E_i isogenous to E can be large. For example, this is shown when $n = 2$ in [XYY16, Section 5]. Therefore, one could conceivably obtain a "commutative supersingular" version of the construction above, which would generalize the recent 2-party key exchange protocol CSIDH [CLM$^+$18], assuming that $\mathbb{F}_p$-isomorphism invariants can be computed in that setting. Since those invariants must be arithmetic rather than geometric in nature, however, this seems even more difficult to achieve than in the ordinary case.

5. Some natural candidate cryptographic invariant maps

In order to instantiate a cryptosystem based on the ideas in this paper, it remains to find an efficiently computable map $\operatorname{isom}\colon \operatorname{Ab}(E) \to S$ for some set S, as in the previous section. Below we give evidence that several natural candidates fail, either because efficiently computing them would break the cryptographic security, or because they are not in fact isomorphism invariants.

Our primary roadblock is that while $E_1 \times \cdots \times E_n$ and $E_1' \times \cdots \times E_n'$ can be isomorphic as unpolarized abelian varieties, they are not necessarily isomorphic as polarized abelian varieties with their product polarizations. The first three proposals below for invariants are invariants of the isomorphism class as polarized abelian varieties, but are not invariants of the isomorphism class as unpolarized abelian varieties. We do not know a way for the different parties to choose polarizations on their product varieties in a compatible way, to produce the same invariant, without solving the elliptic curve isogeny problem.

At present, we do not know an invariant of abelian varieties in dimension ≥ 2 that does not require choosing a polarization, with the exception of what we call the "Deligne invariant", described below.

The theta null invariant. One natural candidate is given by Mumford's *theta nulls*, presented in detail in Appendix B. Unfortunately, in order to compute even a single theta null, one must first choose a principal polarization, and the resulting invariant does depend on this choice of polarization in a crucial way. In Proposition B.7 below we show that, as a result, the theta nulls do not in fact provide an isomorphism invariant as unpolarized abelian varieties.

Igusa invariants. Suppose $n = 2$ and $\operatorname{End} E \otimes \mathbb{Q} \cong \mathbb{Q}(\sqrt{-d})$ with $d \in \mathbb{N}$ square-free. If $d \neq 1, 3, 7, 15$, then for E_1 and E_2 in the isogeny class of E, the product $E_1 \times E_2$ is the Jacobian of a genus 2 curve C (see [HN65]). It is possible to compute such a genus 2 curve C, given a suitable principal polarization on $E_1 \times E_2$. For each such C, one could then compute the Igusa invariants [Igu60] of C. The number of genus 2 curves C such that $E_1 \times E_2$ is isomorphic to the Jacobian variety of C is large ([Hay65] and [Lan06, Theorem 5.1]), and unfortunately the Igusa invariants are different for different choices of C. There are many principal polarizations on each element of $\mathrm{Ab}(E)$, and no compatible way for the different parties to choose the same one.

Invariants of Kummer surfaces. When $n = 2$, another approach is to consider the Kummer surface of $A = E_1 \times E_2$, which is the quotient $K = A/\{\pm 1\}$. The surface K itself does not depend on a polarization. But extracting an invariant from K, for example as in [CF96, Chapter 3], *does* depend on having a projective embedding of K.

Deligne invariant. A natural candidate is an isomorphism invariant studied by Deligne [Del69]. Suppose A is an ordinary abelian variety over $k = \mathbb{F}_q$. The Serre-Tate canonical lift of A to characteristic 0 produces an abelian variety over the ring of Witt vectors $W(\bar{k})$. Fixing an embedding of $W(\bar{k})$ into $\mathbb{C}$, we can view this lift as a complex abelian variety $\hat{A}$. Let $T(A)$ denote the first integral homology group of $\hat{A}$. The Frobenius endomorphism F of A also lifts to characteristic 0 and defines an action of F on $T(A)$. The theorem in [Del69, §7] shows that ordinary abelian varieties A and B over $\mathbb{F}_q$ are isomorphic if and only if there is an isomorphism $T(A) \to T(B)$ that respects the action of F.

A natural candidate for a cryptographic invariant map is the map that sends $(E_1, \ldots, E_n)$ to the isomorphism invariant

$$T(E_1 \times \cdots \times E_n) = T(E_1) \oplus \cdots \oplus T(E_n).$$

Specifying the isomorphism class of $T(E_1 \times \cdots \times E_n)$ as a $\mathbb{Z}[F]$-module is equivalent to specifying the action of F as a $2n \times 2n$ integer matrix, unique up to conjugacy over $\mathbb{Z}$. The characteristic polynomial g of F acting on $T(E_1 \times \cdots \times E_n)$ is of the form $g = f^n$, where f is the characteristic polynomial of F acting on $T(E_i)$ for any i. If this characteristic polynomial g had been irreducible over $\mathbb{Q}$, then as in [Tau49] there would be a canonical bijection between the $\mathbb{Z}$-conjugacy classes of matrices with that characteristic polynomial and the class group of $\mathbb{Z}[t]/(g)$, but it is not clear how to give an invariant for a $\mathbb{Z}$-conjugacy class of matrices when the characteristic polynomial is reducible. However, when $n = 1$, the characteristic polynomial *is* irreducible. As a consequence, being able to compute $T(E)$ for an elliptic curve E in polynomial time would yield a polynomial-time algorithm to solve the elliptic curve isogeny problem of recovering $\mathfrak{a}$ given E and $\mathfrak{a} * E$, as follows.

Theorem 5.1. *An efficient algorithm to compute $T(E)$ on an isogeny class of ordinary elliptic curves over a finite field gives an efficient algorithm to solve the elliptic curve isogeny problem in that isogeny class.*

Proof. Suppose that E_1 and E_2 are in the isogeny class, and suppose that for $i = 1, 2$ we have a $\mathbb{Z}$-basis $\{u_i, v_i\}$ for $T(E_i)$ and a 2×2 integer matrix giving the action of F with respect to this basis. Let $f(t)$ be the characteristic polynomial of Frobenius acting on E_1 or E_2; these are the same since E_1 and E_2 are isogenous. Let $R = \mathbb{Z}[t]/(f)$ and $R_{\mathbb{Q}} = R \otimes_{\mathbb{Z}} \mathbb{Q}$. Then $T(E_i)$ is a rank one R-module, with t acting as F.

Compute $a_i, b_i \in \mathbb{Z}$ such that $F(u_i) = a_i u_i + b_i v_i$. Let $\mathfrak{a}_i$ be the fractional R-ideal generated by 1 and $(t - a_i)/b_i$. Compute and output $\mathfrak{a} = \mathfrak{a}_1 \mathfrak{a}_2^{-1}$.

We claim that $\mathfrak{a} * E_1 \cong E_2$. Define $\lambda_i : T(E_i) \hookrightarrow R_{\mathbb{Q}}$ by sending $w \in T(E_i)$ to the unique $\lambda_i(w) \in R_{\mathbb{Q}}$ such that $\lambda_i(w) \cdot u_i = w$. Then $\lambda_i(u_i) = 1$ and $\lambda_i(v_i) = (t - a_i)/b_i$, so the fractional ideal $\mathfrak{a}_i$ is the image of the map λ_i. Suppose M is a positive integer such that $M\mathfrak{a}$ is an integral ideal of R, and let $h = \lambda_2^{-1} \circ M\lambda_1$. Then $h(T(E_1))$ is an R-submodule of $T(E_2)$. By [Del69, §7], the map $E \mapsto T(E)$ is a fully faithful functor, i.e., it induces a bijection

$$\mathrm{Hom}_k(E_1, E_2) \to \mathrm{Hom}_R(T(E_1), T(E_2)).$$

Thus h arises from a unique isogeny $\phi : E_1 \to E_2$. By [Del69, §4], the kernel of ϕ is isomorphic as an R-module to $T(E_2)/h(T(E_1))$. The latter R-module is isomorphic to $R/M\mathfrak{a}$, and hence is exactly annihilated by $M\mathfrak{a}$. Thus $\ker(\phi) \cong E_1[M\mathfrak{a}]$, so $E_2 \cong E_1/E_1[M\mathfrak{a}] \cong (M\mathfrak{a}) * E_1$. Since $M\mathfrak{a}$ and $\mathfrak{a}$ are in the same ideal class, we have $E_2 \cong \mathfrak{a} * E_1$.

Fractional ideals can be inverted in polynomial time by [Bel04, Algorithm 5.3] or [Coh93, §4.8.4] (see [Bel04, p. 21] for the complexity). □

Acknowledgments

We thank the American Institute of Mathematics (AIM) for supporting a workshop on multilinear maps where the initial seeds for this work were developed, and the Banff International Research Station (BIRS) where our collaboration continued. We also thank Michiel Kosters and Yuri Zarhin. Boneh was partially supported by NSF, DARPA, and ONR. Silverberg was partially supported by a grant from the Alfred P. Sloan Foundation and NSF grant CNS-1703321.

References

[Bel04] Karim Belabas. Topics in computational algebraic number theory. *J. Théor. Nombres Bordeaux*, 16(1):19–63, 2004.

[BGI14] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. In Hugo Krawczyk, editor, *PKC 2014*, volume 8383 of *LNCS*, pages 501–519. Springer, Heidelberg, March 2014.

[BLR$^+$15] Dan Boneh, Kevin Lewi, Mariana Raykova, Amit Sahai, Mark Zhandry, and Joe Zimmerman. Semantically secure order-revealing encryption: Multi-input functional encryption without obfuscation. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 563–594. Springer, Heidelberg, April 2015.

[BLS01] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the Weil pairing. In Colin Boyd, editor, *ASIACRYPT 2001*, volume 2248 of *LNCS*, pages 514–532. Springer, Heidelberg, December 2001.

[BR96] Mihir Bellare and Phillip Rogaway. The exact security of digital signatures: How to sign with RSA and Rabin. In Ueli M. Maurer, editor, *EUROCRYPT'96*, volume 1070 of *LNCS*, pages 399–416. Springer, Heidelberg, May 1996.

[BS03] Dan Boneh and Alice Silverberg. Applications of multilinear forms to cryptography. *Contemporary Mathematics*, 324(1):71–90, 2003.

[BW13] Dan Boneh and Brent Waters. Constrained pseudorandom functions and their applications. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part II*, volume 8270 of *LNCS*, pages 280–300. Springer, Heidelberg, December 2013.

[BZ17] Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. *Algorithmica*, 79(4):1233–1285, 2017. Extended abstract in Crypto 2014.

[CF96] John W. S. Cassels and E. Victor Flynn. *Prolegomena to a middlebrow arithmetic of curves of genus* 2, volume 230 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 1996.

[CLG09] Denis Xavier Charles, Kristin E. Lauter, and Eyal Z. Goren. Cryptographic hash functions from expander graphs. *Journal of Cryptology*, 22(1):93–113, January 2009.

[CLM$^+$18] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: An efficient post-quantum commutative group action. Cryptology ePrint Archive, Report 2018/383, 2018. https://eprint.iacr.org/2018/383.

[Coh93] Henri Cohen. *A course in computational algebraic number theory*, volume 138 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, 1993.

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

[Del69] Pierre Deligne. Variétés abéliennes ordinaires sur un corps fini. *Invent. Math.*, 8:238–243, 1969.

[Deu41] Max Deuring. Die Typen der Multiplikatorenringe elliptischer Funktionenkörper. *Abh. Math. Sem. Hansischen Univ.*, 14:197–272, 1941.

[DY05] Yevgeniy Dodis and Aleksandr Yampolskiy. A verifiable random function with short proofs and keys. In Serge Vaudenay, editor, *PKC 2005*, volume 3386 of *LNCS*, pages 416–431. Springer, Heidelberg, January 2005.

[EHL$^+$18] Kirsten Eisenträger, Sean Hallgren, Kristin E. Lauter, Travis Morrison, and Christophe Petit. Supersingular isogeny graphs and endomorphism rings: Reductions and solutions. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 329–368. Springer, Heidelberg, April / May 2018.

[FHKP13] Eduarda S. V. Freire, Dennis Hofheinz, Eike Kiltz, and Kenneth G. Paterson. Non-interactive key exchange. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 254–271. Springer, Heidelberg, February / March 2013.

[FJP14] Luca De Feo, David Jao, and Jérôme Plût. Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies. *J. Mathematical Cryptology*, 8(3):209–247, 2014.

[FKS18] Luca De Feo, Jean Kieffer, and Benjamin Smith. Towards practical key exchange from ordinary isogeny graphs. Cryptology ePrint Archive, Report 2018/485, 2018. https://eprint.iacr.org/2018/485.

[Fuc14] Georg Fuchsbauer. Constrained verifiable random functions. In Michel Abdalla and Roberto De Prisco, editors, *SCN 14*, volume 8642 of *LNCS*, pages 95–114. Springer, Heidelberg, September 2014.

[Gal18] Steven D. Galbraith. Authenticated key exchange for SIDH. Cryptology ePrint Archive, Report 2018/266, 2018. https://eprint.iacr.org/2018/266.

[GGH13a] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 1–17. Springer, Heidelberg, May 2013.

[GGH$^+$13b] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49. IEEE Computer Society Press, October 2013.

[GGSW13] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 467–476. ACM Press, June 2013.

[GKP$^+$13] Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nickolai Zeldovich. How to run Turing machines on encrypted data. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 536–553. Springer, Heidelberg, August 2013.

[Hay65] Tsuyoshi Hayashida. A class number associated with a product of two elliptic curves. *Natur. Sci. Rep. Ochanomizu Univ.*, 16:9–19, 1965.

[HN65] Tsuyoshi Hayashida and Mieo Nishi. Existence of curves of genus two on a product of two elliptic curves. *J. Math. Soc. Japan*, 17:1–16, 1965.

[Igu60] Jun-ichi Igusa. Arithmetic variety of moduli for genus two. *Ann. of Math. (2)*, 72:612–649, 1960.

[JKP$^+$18] Bruce W. Jordan, Allan G. Keeton, Bjorn Poonen, Eric M. Rains, Nicholas Shepherd-Barron, and John T. Tate. Abelian varieties isogenous to a power of an elliptic curve. *Compositio Math.*, 154(5):934–959, 2018. https://arxiv.org/abs/1602.06237.

[JMV09] David Jao, Stephen D. Miller, and Ramarathnam Venkatesan. Expander graphs based on GRH with an application to elliptic curve cryptography. *J. Number Theory*, 129(6):1491–1504, 2009. https://arxiv.org/abs/0811.0647.

[Jou04] Antoine Joux. A one round protocol for tripartite Diffie-Hellman. *Journal of Cryptology*, 17(4):263–276, September 2004.

[Kan11] Ernst Kani. Products of CM elliptic curves. *Collectanea Math.*, 62(3):297–339, 2011. http://www.mast.queensu.ca/~kani/papers/CMprod3.pdf.

[KPTZ13] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 13*, pages 669–684. ACM Press, November 2013.

[Lan06] Herbert Lange. Principal polarizations on products of elliptic curves. In *The geometry of Riemann surfaces and abelian varieties*, volume 397 of *Contemp. Math.*, pages 153–162. Amer. Math. Soc., Providence, RI, 2006.

[Lau02] Kristin Lauter. The maximum or minimum number of rational points on genus three curves over finite fields. *Compositio Math.*, 134(1):87–111, 2002. https://arxiv.org/abs/math/0104086.

[Lit28] John E. Littlewood. On the class number of the corpus $P(\sqrt{-k})$. *Proc. London Math. Soc.*, 27(1):358–372, 1928. https://doi.org/10.1112/plms/s2-27.1.358.

[Lys02] Anna Lysyanskaya. Unique signatures and verifiable random functions from the DH-DDH separation. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 597–612. Springer, Heidelberg, August 2002.

[MRV99] Silvio Micali, Michael O. Rabin, and Salil P. Vadhan. Verifiable random functions. In *40th FOCS*, pages 120–130. IEEE Computer Society Press, October 1999.

[Mum66] David Mumford. On the equations defining abelian varieties. I. *Invent. Math.*, 1:287–354, 1966.

[Rao14] Vanishree Rao. Adaptive multiparty non-interactive key exchange without setup in the standard model. Cryptology ePrint Archive, Report 2014/910, 2014. http://eprint.iacr.org/2014/910.

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

[Ser02] Jean-Pierre Serre. Modules hermitiens et courbes algébriques. Appendix to [Lau02], 2002.

[Shi77] Tetsuji Shioda. Some remarks on abelian varieties. *J. Fac. Sci. Univ. Tokyo, Sec. 1A*, 24(1):11–21, 1977. http://repository-old.dl.itc.u-tokyo.ac.jp/dspace/handle/2261/6164.

[Shi78] Tetsuji Shioda. Supersingular K3 surfaces. In Knud Lønsted, editor, *Algebraic Geometry*, volume 732 of *Lecture Notes in Mathematics*, pages 564–591. Springer, 1978.

[Ste11] Ernst Steinitz. Rechteckige Systeme und Moduln in algebraischen Zahlköppern. I. *Math. Ann.*, 71(3):328–354, 1911.

[Tau49] Olga Taussky. On a theorem of Latimer and MacDuffee. *Canadian J. Math.*, 1:300–302, 1949.

[XYY16] Jiangwei Xue, Tse-Chung Yang, and Chia-Fu Yu. On superspecial abelian surfaces over finite fields. *Documenta Mathematica*, 21:1607–1643, 2016.

[Zha16] Mark Zhandry. How to avoid obfuscation using witness PRFs. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A, Part II*, volume 9563 of *LNCS*, pages 421–448. Springer, Heidelberg, January 2016.

Appendix A. Background on isogenies

This appendix provides background on isogenies of ordinary elliptic curves over finite fields and their products, which expands upon the short discussion at the beginning of Section 4 and makes it explicit. Throughout this appendix, E is a fixed ordinary elliptic curve over the finite field $\mathbb{F}_q$.

A.1. **Endomorphism rings.**

Let $\mathscr{O} = \mathrm{End}(E)$ be the endomorphism ring of the elliptic curve E. Note that since, in the ordinary case, all endomorphisms over the algebraic closure are already defined over $\mathbb{F}_q$, it is not necessary to distinguish between the endomorphisms of E defined over $\mathbb{F}_q$ or over $\overline{\mathbb{F}_q}$.

Denote by $\pi \in \mathscr{O}$ the Frobenius endomorphism of E, and by $\mathscr{O}_K$ the integral closure of $\mathscr{O}$, i.e., the ring of integers of $\mathrm{End}(E) \otimes_{\mathbb{Z}} \mathbb{Q}$. We have an inclusion of orders $\mathbb{Z}[\pi] \subset \mathscr{O} \subset \mathscr{O}_K$ and both inclusions are strict in general.

However, the conductor $m = \big[\mathscr{O}_K : \mathbb{Z}[\pi]\big]$ is typically quite small. Let t be the trace of the Frobenius π (computed using Schoof's algorithm) and $D_\pi = t^2 - 4p$ the discriminant of $\mathbb{Z}[\pi]$; then m is given by $D_\pi = m^2 D$ where D is the (fundamental) discriminant of $\mathscr{O}_K$. Hence, m^2 is the quotient of D_π by its squarefree part (or 4 times its squarefree part when D_π is even). This is usually a small number, and equal to 1 with constant probability for a randomly chosen curve E.

Therefore, we will assume for simplicity that $m = 1$, and thus $\mathbb{Z}[\pi] = \mathscr{O} = \mathscr{O}_K$. Note that strictly speaking, it is not known how to test for that condition in polynomial time (squarefreeness is not known to be easier than factoring), but we can easily test for stronger conditions, such as D_π being prime or the product of a large prime by a small squarefree integer, and since restricting attention to such special discriminants is desirable anyway (to avoid smooth class numbers), this is a reasonable test to carry out.

In principle, it should be possible to extend what follows to the general case of an arbitrary endomorphism ring $\mathscr{O}$ as long as we restrict attention to, on the one hand, the subset $\mathrm{Ell}(\mathscr{O})$ of the isogeny class of E over $\mathbb{F}_q$ consisting of curves with the same endomorphism ring $\mathscr{O}$ as E, and on the other hand, isogenies of degree prime to m. In particular, the latter are guaranteed to preserve the endomorphism ring, and suffice to apply the expander graph argument of Jao–Miller–Venkatesan [JMV09]; see below. There are a few complications in considering the general case, however, so the extension is not considered in the present discussion.

A.2. **Classification of abelian varieties isogenous to a power of E.** Let $\mathrm{Ab}(E)$ denote the category of abelian varieties over $\mathbb{F}_q$ isogenous[2] to a product of copies of E. In the case when $\mathbb{Z}[\pi] = \mathscr{O}_K = \mathscr{O}$, the functor $\mathrm{Hom}(E, -)$ is an equivalence between the category $\mathrm{Ab}(E)$ and the category $\mathrm{Mod}(\mathscr{O})$ of torsion-free $\mathscr{O}$-modules of finite type. The inverse functor associates to any module $M \in \mathrm{Mod}(\mathscr{O})$ with finite presentation $\mathscr{O}^n \to \mathscr{O}^m \to M \to 0$ the abelian variety $M \otimes_{\mathscr{O}} E$ defined by taking the corresponding cokernel $E^n \to E^m \to M \otimes_{\mathscr{O}} E \to 0$. This is established in Serre's appendix [Ser02] to [Lau02]. See also [Kan11] and [JKP+18] for more general related results.

In dimension 1, this implies that the set $\mathrm{Ell}(\mathscr{O})$ of isomorphism classes of elliptic curves over $\mathbb{F}_q$ isogenous to E is in bijection with the set of isomorphism classes of torsion-free modules of rank 1 over $\mathscr{O}$. Since any such module is isomorphic to an ideal of $\mathscr{O}$, this yields a bijection between $\mathrm{Ell}(\mathscr{O})$ and the ideal class group $\mathrm{Cl}(\mathscr{O})$ of $\mathscr{O}$.

Under that correspondence, an ideal $\mathfrak{a}$ maps to the elliptic curve $\mathfrak{a} \otimes_{\mathscr{O}} E$ given by $E/E[\mathfrak{a}]$, where $E[\mathfrak{a}]$ is the intersection of the kernels of all endomorphisms in $\mathfrak{a}$ (that notation is consistent with the usual one for torsion subgroups). Denote that elliptic curve by $\mathfrak{a} * E$. Its endomorphism ring is clearly still $\mathscr{O}$, and the identification is canonical (e.g., because the Frobenius endomorphism on each curve is well-defined). We can therefore consider the elliptic curve $\mathfrak{b} * \big(\mathfrak{a} * E\big)$ for some other ideal $\mathfrak{b}$ of $\mathscr{O}$, and it turns out to be isomorphic to $\mathfrak{ab} * E$. Thus, $*$ endows $\mathrm{Ell}(\mathscr{O})$ with a group action by $\mathrm{Cl}(\mathscr{O})$, under which it is a principal homogenous space.

By Steinitz's classification theorem of torsion-free modules of finite type over Dedekind domains [Ste11] (which says that any such module is isomorphic to a sum of ideals, with equality when the product ideals are in the same class) and the equivalence of categories, we see that any abelian variety $A \in \mathrm{Ab}(E)$ is of the form $E_1 \times \cdots \times E_g$ for some elliptic curves E_i of the form $\mathfrak{a}_i * E$. Moreover, two such abelian varieties,

$$\begin{aligned} A &= E_1 \times \cdots \times E_g \qquad (E_i = \mathfrak{a}_i * E) \quad \text{and} \\ A' &= E_1' \times \cdots \times E_g' \qquad (E_i' = \mathfrak{a}_i' * E), \end{aligned}$$

are isomorphic if and only if $[\mathfrak{a}_1 \cdots \mathfrak{a}_g] = [\mathfrak{a}_1' \cdots \mathfrak{a}_g']$ in $\mathrm{Cl}(\mathscr{O})$. In particular, for any fixed $g \geq 1$, the isomorphism classes of abelian varieties of dimension g are in bijection with $\mathrm{Cl}(\mathscr{O})$, given by mapping $[\mathfrak{a}] \in \mathrm{Cl}(\mathscr{O})$ to $(\mathfrak{a} * E) \times E^{g-1}$.

A.3. **Random sampling and isogeny walks.** The group $\mathrm{Cl}(\mathscr{O})$ is usually large; by Dirichlet's class number formula combined with Littlewood's estimate for the values at 1 of the

[2] The isogeny can be defined over $\mathbb{F}_q$ or the algebraic closure; the two notions coincide.

L-functions of quadratic characters [Lit28], its order $h(D)$ satisfies

$$h(D) \gg \frac{\sqrt{|D|}}{\log\log|D|}$$

if one assumes that the Generalized Riemann Hypothesis holds for quadratic fields (where $D = t^2 - 4q$ can be chosen of the same order of magnitude as q in absolute value). Moreover, the best known algorithms to compute its structure or even its order are subexponential.

We would therefore like to be able to sample uniformly from $\mathrm{Cl}(\mathscr{O})$ without having to compute the group structure or the order, and do so in such a way that if the sampling algorithm returns the class of an ideal $\mathfrak{a}$, it is feasible to compute $\mathfrak{a} * E$ (which cannot be done directly with an ideal of large norm, as it involves computing a large degree isogeny).

To solve this problem, the same approach as in the supersingular case basically works: simply obtain an element of the class group as a product of sufficiently many ideals of small prime norm (and the corresponding isogeny as the composition of the corresponding isogenies of small prime degrees). Contrary to the supersingular case, it is not possible to use a fixed degree, but using all small primes up to a polynomial bound is enough, as we will discuss below.

Prime ideals and ℓ-isogenies. Before we discuss the sampling of random classes in $\mathrm{Cl}(\mathscr{O})$, we must recall a few facts about the ideals of the quadratic ring $\mathscr{O}$ and the corresponding isogenies. All of this is described in particular in [Cou06], for example. Any ideal $\mathfrak{a}$ of $\mathscr{O}$ can be factored uniquely $\mathfrak{a} = \mathfrak{l}_1^{e_1} \cdots \mathfrak{l}_r^{e_r}$ into a product of powers of prime ideals $\mathfrak{l}_i$. To any prime ideal $\mathfrak{l}$, one can associate the unique rational prime $\ell \in \mathbb{Z}^+$ above which it sits (ℓ is such that $\mathfrak{l} \cap \mathbb{Z} = \ell\mathbb{Z}$), and prime ideals can thus be obtained from the ideal decomposition in $\mathscr{O}$ of principal ideals generated by rational primes ℓ.

Since $\mathscr{O}$ is quadratic, that decomposition is easy to describe in terms of the Legendre symbol $\left(\frac{D}{\ell}\right)$. A rational prime ℓ such that $\left(\frac{D}{\ell}\right) = -1$ is inert in $\mathscr{O}$, i.e., $\ell\mathscr{O}$ is itself prime. As a result, these prime ideals are trivial in the class group, and not interesting for our purposes. On the other hand, a rational prime ℓ such that $\left(\frac{D}{\ell}\right) = 1$ is completely split: we have $\ell\mathscr{O} = \mathfrak{l} \cdot \bar{\mathfrak{l}}$ for distinct prime ideals $\mathfrak{l} \neq \bar{\mathfrak{l}}$ of norm ℓ. This case is the most relevant for us; it gives rise to ℓ-isogenies.

Finally, there are finitely many rational primes ℓ such that $\left(\frac{D}{\ell}\right) = 0$, i.e., primes dividing the discriminant. Those primes are ramified: $\ell\mathscr{O} = \mathfrak{l}^2$ for some prime ideal of norm ℓ. As a result, a ramified prime $\mathfrak{l}$ is of order at most 2 in the class group; moreover, it is a consequence of the Chinese remainder theorem that one can always find an ideal $\mathfrak{a}$ of $\mathscr{O}$ coprime to D such that $[\mathfrak{l}] = [\mathfrak{a}]$ in $\mathrm{Cl}(\mathscr{O})$. Thus, we can always ignore these primes in order to make the presentation more consistent.

It results from the above that any ideal class $[\mathfrak{a}] \in \mathrm{Cl}(\mathscr{O})$ can be written as a product $[\mathfrak{l}_1]^{e_1} \cdots [\mathfrak{l}_r]^{e_r}$ of powers of classes of prime ideals over split rational primes ℓ_i (such that $\left(\frac{D}{\ell_i}\right) = 1$). Computing $\mathfrak{a} * E$ therefore reduces to computing $\mathfrak{l} * C$ for an arbitrary prime ideal $\mathfrak{l}$ over a split rational prime ℓ, and C an arbitrary elliptic curve isogenous to E. This still cannot be done when the norm ℓ is large, but if it is small (say $\mathrm{polylog}(q)$), it can be done efficiently.

Indeed, for such a split ℓ, there are exactly two ℓ-isogenies from C: if we let $\ell\mathscr{O} = \mathfrak{l} \cdot \bar{\mathfrak{l}}$, these are $C \to C/C[\mathfrak{l}]$ and $C \to C/C[\bar{\mathfrak{l}}]$. As usual for low degree isogenies, they can be computed efficiently by evaluating modular polynomials. Moreover, we can easily distinguish between the two by looking at the action of the Frobenius endomorphism π on the ℓ-torsion $C[\ell]$ (viewed as an $\mathbb{F}_\ell$-vector space of dimension 2).

More precisely, since $\left(\frac{D}{\ell}\right) = 1$, the characteristic polynomial $X^2 - tX + q$ of the Frobenius π has two distinct roots $\mu \neq \bar{\mu}$ modulo ℓ, and we can then write the ideals $\mathfrak{l}, \bar{\mathfrak{l}}$ as $(\ell, \pi - \mu)$ and $(\ell, \pi - \bar{\mu})$ respectively (up to reordering). If we view $C[\ell]$ as an $\mathbb{F}_\ell$-vector space of dimension 2, the two kernels $C[\mathfrak{l}]$ and $C[\bar{\mathfrak{l}}]$ will be vector subspaces given by

$$C[\mathfrak{l}] = C[\ell] \cap \operatorname{Ker}(\pi - \mu) \quad \text{and} \quad C[\bar{\mathfrak{l}}] = C[\ell] \cap \operatorname{Ker}(\pi - \bar{\mu})$$

respectively, according to their generators above. In other words, these are the eigenspaces of the Frobenius on $C[\ell]$, and π acts by multiplication by μ on $C[\mathfrak{l}]$ (resp. $\bar{\mu}$ on $C[\bar{\mathfrak{l}}]$).

Random sampling in $\mathrm{Cl}(\mathscr{O})$. We have just seen that we can efficiently compute $\mathfrak{l} * C$ for any prime ideal $\mathfrak{l}$ of $\mathscr{O}$ over a split rational prime $\ell = \mathrm{polylog}(q)$. We now argue that those prime ideals are in fact sufficient to sample from a distribution close to the uniform distribution in $\mathrm{Cl}(\mathscr{O})$.

This results from a theorem of Jao, Miller and Venkatesan [JMV09, Th. 1.1] essentially saying that the Cayley graph of $\mathrm{Cl}(\mathscr{O})$ obtained using the classes of the ideals $\mathfrak{l}$ associated to all split rational primes $\ell \leq x$ with $x = \log^B |D|$ is an expander graph for any $B > 2$. Moreover, there is an explicit bound on the spectral gap; namely, if we denote by $\lambda_{\text{triv}} \sim 2x \log x$ the largest eigenvalue of the adjacency matrix, then all other eigenvalues λ satisfy

$$|\lambda| = O\big((\lambda_{\text{triv}} \log \lambda_{\text{triv}})^{1/2+1/B}\big).$$

If we take $B = 2 + \varepsilon$, this implies that, for a suitable absolute constant C, a random walk of length

$$r \geq C \cdot \frac{\delta + \log h(D)}{\varepsilon \cdot \log\log |D|} \tag{3}$$

in the graph yields a distribution that is $\exp(-\delta)$-statistically close to uniform.

Therefore, if we denote by S the set of ideals $\mathfrak{l}$ associated to all split rational primes $\ell \leq x$ as above, and we choose r elements $\mathfrak{l}_1, \dots, \mathfrak{l}_r$ in S at random with r as in (3), then the class of the ideal $\mathfrak{a} = \mathfrak{l}_1 \cdots \mathfrak{l}_r$ is close to uniformly distributed in $\mathrm{Cl}(\mathscr{O})$.

We can then compute the elliptic curve $\mathfrak{a} * E$ as $\mathfrak{l}_r * \Big(\mathfrak{l}_{r-1} * \big(\cdots * (\mathfrak{l}_1 * E)\big)\Big)$, and its isomorphism class is close to uniformly distributed in the isogeny class $\mathrm{Ell}(\mathscr{O})$ of E.

A.4. **A low-brow approach to the classification result.** In Appendix A.2 above, we recalled results of Serre and Steinitz that are sufficient to establish the following properties, used in the main construction of Section 4. Suppose that E is such that $\mathscr{O} = \mathbb{Z}[\pi]$ is integrally closed, and let $\mathfrak{a}_1, \dots, \mathfrak{a}_n, \mathfrak{a}'_1, \dots, \mathfrak{a}'_n$ be ideals of $\mathscr{O}$. Then, the following abelian varieties are isomorphic:

$$(\mathfrak{a}_1 * E) \times \cdots \times (\mathfrak{a}_n * E) \cong (\mathfrak{a}_1 \cdots \mathfrak{a}_n) * E \times E^{n-1}. \tag{4}$$

and more generally, we have

$$\begin{aligned} (\mathfrak{a}_1 * E) \times \cdots \times (\mathfrak{a}_n * E) \cong (\mathfrak{a}'_1 * E) \times \cdots \times (\mathfrak{a}'_n * E) \quad & \text{if and only if} \\ \mathfrak{a}_1 \cdots \mathfrak{a}_n = \mathfrak{a}'_1 \cdots \mathfrak{a}'_n \quad & \text{as ideal classes in } \mathrm{Cl}(\mathscr{O}). \end{aligned} \tag{5}$$

As a side note, we now mention that those properties can *in part* be established using elementary techniques. More precisely, (4) is a consequence of the following elementary result.

Theorem A.1. *Let E be an elliptic curve over a finite field $\mathbb{F}_q$, and K a finite étale subgroup of E (i.e., the map $E \to E/K$ is separable) defined over $\mathbb{F}_q$. Suppose that K contains* *subgroups K_i defined over $\mathbb{F}_q$, for $1 \leq i \leq n$, whose orders are pairwise coprime, and suppose $K = K_1 + \cdots + K_n$. Then:*

$$(E/K_1) \times \cdots \times (E/K_n) \cong (E/K) \times E^{n-1}.$$

Proof. The result is immediate for $n = 1$. We next prove the result for $n = 2$ by constructing an explicit isomorphism. Consider the commutative diagram:

$$\begin{tikzcd}
E \arrow[r, "\varphi_1"] \arrow[d, "\varphi_2"'] \arrow[dr, "\theta"] & E/K_1 \arrow[d, "\psi_1"] \\
E/K_2 \arrow[r, "\psi_2"'] & E/K
\end{tikzcd}$$

where all maps are the natural quotient isogenies. If we denote by m_1 and m_2 the orders of K_1 and K_2, we have $\deg \varphi_1 = \deg \psi_2 = m_1$ and $\deg \varphi_2 = \deg \psi_1 = m_2$. Now choose integers $a, b \in \mathbb{Z}$ such that $am_1 + bm_2 = 1$. We define morphisms

$$f \colon E \times (E/K) \to (E/K_1) \times (E/K_2) \qquad \text{and} \qquad g \colon (E/K_1) \times (E/K_2) \to E \times (E/K)$$

by the following matrices:

$$\mathrm{Mat}(f) = \begin{pmatrix} \varphi_1 & \widehat{\psi_1} \\ -b\varphi_2 & a\widehat{\psi_2} \end{pmatrix} \qquad \text{and} \qquad \mathrm{Mat}(g) = \begin{pmatrix} a\widehat{\varphi_1} & -\widehat{\varphi_2} \\ b\psi_1 & \psi_2 \end{pmatrix}.$$

Then:

$$\begin{aligned}
\mathrm{Mat}(g \circ f) &= \begin{pmatrix} a\widehat{\varphi_1} & -\widehat{\varphi_2} \\ b\psi_1 & \psi_2 \end{pmatrix} \begin{pmatrix} \varphi_1 & \widehat{\psi_1} \\ -b\varphi_2 & a\widehat{\psi_2} \end{pmatrix} \\
&= \begin{pmatrix} a\widehat{\varphi_1}\varphi_1 + b\widehat{\varphi_2}\varphi_2 & a\widehat{\varphi_1}\widehat{\psi_1} - a\widehat{\varphi_2}\widehat{\psi_2} \\ b\psi_1\varphi_1 - b\psi_2\varphi_2 & b\psi_1\widehat{\psi_1} + a\psi_2\widehat{\psi_2} \end{pmatrix} \\
&= \begin{pmatrix} a[m_1] + b[m_2] & (a-a)\widehat{\theta} \\ (b-b)\theta & b[m_2] + a[m_1] \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.
\end{aligned}$$

Thus, $g \circ f$ is the identity on $E \times (E/K)$, and similarly $f \circ g$ is the identity on $(E/K_1) \times (E/K_2)$. So f and g are mutually inverse isomorphisms, proving the case $n = 2$.

We now proceed by induction. Suppose the desired result holds for some n, and suppose K can be written as $K = K_1 + \cdots + K_{n+1}$ with the K_i's of pairwise coprime order. Let $H = K_1 + \cdots + K_n$. The induction hypothesis implies

$$(E/K_1) \times \cdots \times (E/K_n) \cong (E/H) \times E^{n-1}. \tag{6}$$

On the other hand, $K = H + K_{n+1}$ and the subgroups H and K_{n+1} have coprime order, so by the $n = 2$ case above:

$$(E/H) \times (E/K_{n+1}) \cong (E/K) \times E. \tag{7}$$

Combining (6) and (7) gives:

$$\begin{aligned}
(E/K_1) \times \cdots \times (E/K_n) \times (E/K_{n+1}) &\cong ((E/H) \times E^{n-1}) \times (E/K_{n+1}) \\
&\cong ((E/H) \times (E/K_{n+1})) \times E^{n-1} \cong (E/K) \times E^n
\end{aligned}$$

which completes the induction. □

Applying Theorem A.1 to $K_i = E[\mathfrak{a}_i]$, we obtain (4) in the special case when the K_i's have pairwise coprime orders, i.e., when the ideals $\mathfrak{a}_i$ have pairwise coprime norms. The general case follows by observing that for any choice of the $\mathfrak{a}_i$'s, one can find ideals $\mathfrak{b}_i$ with pairwise coprime norms such that $\mathfrak{b}_i$ is in the same ideal class as $\mathfrak{a}_i$ (just let $\mathfrak{b}_1 = \mathfrak{a}_1$, and construct $\mathfrak{b}_{i+1}$ as an ideal in the class of $\mathfrak{a}_{i+1}$ coprime to $N(\mathfrak{a}_1 \cdots \mathfrak{a}_i) \cdot \mathscr{O}$).

As a result, the "if" part of property (5) also follows. Indeed, by (4),

$$(\mathfrak{a}_1 * E) \times \cdots \times (\mathfrak{a}_n * E) \cong (\mathfrak{a}_1 \cdots \mathfrak{a}_n) * E \times E^{n-1}$$

and

$$(\mathfrak{a}'_1 * E) \times \cdots \times (\mathfrak{a}'_n * E) \cong (\mathfrak{a}'_1 \cdots \mathfrak{a}'_n) * E \times E^{n-1}.$$

If the ideal classes of $\mathfrak{a}_1 \cdots \mathfrak{a}_n$ and $\mathfrak{a}'_1 \cdots \mathfrak{a}'_n$ are equal in $\mathrm{Cl}(\mathscr{O})$, then

$$(\mathfrak{a}_1 * E) \times \cdots \times (\mathfrak{a}_n * E) \cong (\mathfrak{a}'_1 * E) \times \cdots \times (\mathfrak{a}'_n * E)$$

as desired.

However, the converse (the "only if" part of property (5)) is more difficult, and essentially states that for E' and E'' isogenous to E, if $E' \times E^{n-1} \cong E'' \times E^{n-1}$ then $E' \cong E''$. Such a cancellation result does not hold for a supersingular E, so the proof in the ordinary case has to rely on the structure of the endomorphism ring in a crucial way. See [Shi77] for an illuminating discussion, particularly Section 4 for the $n = 2$ case.

Appendix B. Mumford's invariant

In this appendix we give details of the result on Mumford's theta null invariants that was discussed in Section 5.

B.1. **Theta functions.** We review what we need from Mumford's theory of theta functions. Let A be a g-dimensional abelian variety over an algebraically closed field k. In practice, we need only work over a low degree extension of our base field. We will assume A comes equipped with a principal polarization given by an invertible sheaf L_0; the choice of such a polarization will later be crucial for our analysis. Let L be an ample symmetric invertible sheaf on A with $\mathrm{char}(k) \nmid \deg L$. Going forwards, all sheaves will be of this form. Typically we let $L = L_0^m$ for some m not divisible by $\mathrm{char}(k)$.

For $P \in A$, write τ_P for the translation-by-P map. Let $H(L) \subset A$ be the set of P for which $\tau_P^* L \cong L$. For convenience, we will assume $H(L) = A[m]$ for some choice of m. Let $G(L)$ be the *theta group* of L, which consists of pairs (P, ψ_P) where $P \in H(L)$ and $\psi_P : L \to L$ is an isomorphism that factors through $\tau_P^* L$; that is, there is some isomorphism $\tilde{\psi}_P : \tau_P^* L \to L$ such that $\psi_P = \tilde{\psi}_P \circ \tau_P^*$. The group operation is

$$(P, \psi_P)(Q, \psi_Q) = (P + Q, \psi_P \circ \psi_Q),$$

which one checks is well-defined. The theta group lies in an exact sequence

$$0 \longrightarrow k^* \longrightarrow G(L) \longrightarrow H(L) \longrightarrow 0 \qquad (8)$$

where the image of $\alpha \in k^*$ is (O, α), and the map $G(L) \to H(L)$ is $(P, \psi_P) \mapsto P$. Let $V = \Gamma(A, L)$. Then $G(L)$ has a natural action on V.

Let δ be the g-tuple $(m, m, \ldots, m)$; we call δ the *level* or *type* of L. Let $K_1 = (\mathbb{Z}/m\mathbb{Z})^g$, let $K_2 = \mathrm{Hom}(K_1, k^*)$, and let $H(\delta) = K_1 \times K_2$. Write $\langle x, y \rangle$ for the evaluation pairing on $H(\delta)$, which we extend to a symplectic form by letting $K_1 \times \{0\}$ and $\{0\} \times K_2$ be isotropic subspaces. Let $G(\delta)$, the *abstract theta group of level* δ, be the set $k^* \times H(\delta)$ with group operation

$$(a, x, y) \cdot (b, x', y') = (\langle x, y' \rangle\, ab, x + x', y + y').$$

Then $G(\delta)$ lies in the exact sequence

$$0 \longrightarrow k^* \longrightarrow G(\delta) \longrightarrow H(\delta) \longrightarrow 0.$$

Let $V(\delta)$ be the k-vector space of maps from K_1 to k. Then $G(\delta)$ acts on $V(\delta)$ via

$$(a, x, y)(f)(z) = a\,\langle z, y\rangle\, f(x + z).$$

A **level δ theta structure** for A is a choice of L as above as well as a choice of isomorphism $\alpha : G(\delta) \to G(L)$ that is the identity on k^*. Mumford [Mum66, p. 298] shows that α induces an isomorphism $\beta : V(\delta) \to V$, unique up to a scalar, that respects α. That is, for all $g \in G(\delta)$ and $f \in V(\delta)$, we have $\beta(gf) = \alpha(g)\beta(f)$. Let $x_0, \ldots, x_N$ be a fixed ordering of the elements of K_1, and let f_i be the delta function at x_i; that is, $f_i(x_i) = 1$ and $f_i(x_j) = 0$ for $i \neq j$. The f_i then furnish a canonical basis for $V(\delta)$. Let $\theta_i = \beta(f_i)$. Define $\phi : A \longrightarrow \mathbb{P}^N$ by $\phi(P) = [\theta_0(P) : \cdots : \theta_N(P)]$. Then $\phi(O)$ is the **theta null associated to** α. Thus, a choice of theta structure gives rise to a point $\phi(O)$ in projective space.

B.2. **Properties of theta structures.**

Proposition B.1. *Given a theta structure $\alpha : G(\delta) \to G(L)$, if $4 \mid \delta$, then the associated map $\phi : A \to \mathbb{P}^N$ is an embedding. Furthermore, $\phi(A)$ is determined by the theta null $\phi(O)$ associated to α.*

Since for us, $\delta = (m, m, \ldots, m)$, $4 \mid \delta$ means $4 \mid m$. The above proposition is the Corollary on p. 349 of [Mum66]. The equations defining $\phi(A)$ are usually called the *Riemann equations.* As a consequence, we see that the theta null associated to a theta structure of type δ, with $4 \mid \delta$, is a nondegenerate invariant; that is, with δ chosen beforehand, nonisomorphic abelian varieties cannot have the same theta null.

Lemma B.2. *Let $\alpha : G(\delta) \to G(L)$ be a theta structure. Then the induced map $H(\delta) \to H(L)$ is a symplectic isomorphism, where the pairing on $H(L)$ is the Weil pairing.*

Proof. See p. 317 of [Mum66]. □

Lemma B.3. *Suppose $H(L) = A[m]$. Given a symplectic isomorphism $h : H(\delta) \to H(L)$, there are exactly m^{2g} theta structures α that induce h.*

Proof. Let $x_1, \ldots, x_g, y_1, \ldots, y_g$ be a symplectic basis for $H(\delta)$. Let $P_i = h(x_i)$ and $Q_i = h(y_i)$. We have that $G(\delta)$ is generated by k^* and the elements $(1, x_i, 0)$ and $(1, 0, y_i)$. Each of these elements has order m in $G(\delta)$. Let $\alpha : G(\delta) \to G(L)$ be an isomorphism that induces h. Then $\alpha((1, x_i, 0)) = (P_i, \psi_{P_i})$ and $\alpha((1, 0, y_i)) = (Q_i, \psi_{Q_i})$ for appropriate choices of ψ_{P_i} and ψ_{Q_i}; in particular, each of these maps $V \to V$ must have order dividing m. From the exact sequence (8), we see that for fixed i, any two choices of ψ_{P_i} (resp. ψ_{Q_i}) differ by a nonzero scalar. But such a scalar must be an mth root of unity so that ψ_{P_i} (resp. ψ_{Q_i}) has order m. Thus, for each P_i, there are at most m choices of ψ_{P_i}, and similarly for the Q_i. Since the choices of $\alpha((1, x_i, 0))$ and $\alpha((1, 0, y_i))$ completely determine α, the claim follows. □

Lemma B.4. *Let A_1, A_2 be abelian varieties and L_1, L_2 ample invertible sheaves on A_1, A_2, respectively. For $i = 1, 2$ let $p_i : A_1 \times A_2 \to A_i$ be the projection maps. Then*

$$G(p_1^* L_1 \otimes p_2^* L_2) = (G(L_1) \times G(L_2))/\{(a, a^{-1}) | a \in k^*\}.$$

Proof. This is Lemma 1 on p. 323 of [Mum66]. □

We will write the right-hand group above as $G(L_1) \times_{k^*} G(L_2)$; if we think of the relevant groups as group schemes, then we do in fact get the fiber product.

Observe that if δ_1, δ_2 are the types of L_1, L_2, respectively, then $\delta = (\delta_1, \delta_2)$ is the type of $p_1^* L_1 \otimes p_2^* L_2$ and $G(\delta) = G(\delta_1) \times_{k^*} G(\delta_2)$. It follows that if α_i is a theta structure for the pair (A_i, L_i), then we have a product theta structure $\alpha : G(\delta) \to G(L)$.

Proposition B.5. *Let A_1 and A_2 be abelian varieties, let L_1 and L_2 be ample invertible sheaves on A_1 and A_2 of types δ_1 and δ_2, respectively, and let $L = p_1^* L_1 \otimes p_2^* L_2$, an invertible sheaf on $A = A_1 \times A_2$. Let α_1 and α_2 be theta structures for (A_1, L_1) and (A_2, L_2), respectively, and let α be the product theta structure for (A, L). For the theta structure α_i, let $\phi_i : A_i \to \mathbb{P}^{N_i}$ be the associated morphism, and similarly define $\phi : A \to \mathbb{P}^N$. Then, up to possible reordering of coordinates, $\phi = s \circ (\phi_1, \phi_2)$, where $(\phi_1, \phi_2) : A_1 \times A_2 \to \mathbb{P}^{N_1} \times \mathbb{P}^{N_2}$ and s is the Segre map $\mathbb{P}^{N_1} \times \mathbb{P}^{N_2} \to \mathbb{P}^N$.*

Proof. We first check that the Segre map makes sense; that is, that

$$N = (N_1 + 1)(N_2 + 1) - 1.$$

Setting $V_i = \Gamma(A_i, L_i)$ and $V = \Gamma(A, L)$, we have $\dim V_i = N_i + 1$, $\dim V = N + 1$, but also $V = V_1 \otimes V_2$. Observe that similarly, $V(\delta) = V(\delta_1) \otimes V(\delta_2)$. Therefore the dimension counts are compatible.

By our isomorphisms $G(\delta) \cong G(\delta_1) \times_{k^*} G(\delta_2)$ and $G(L) \cong G(L_1) \times_{k^*} G(L_2)$ and the fact that α is obtained as

$$(\alpha_1, \alpha_2) : G(\delta_1) \times_{k^*} G(\delta_2) \longrightarrow G(L_1) \times_{k^*} G(L_2),$$

one sees that the induced isomorphism $\beta : V(\delta) \to V$ can instead be written

$$(\beta_1, \beta_2) : V(\delta_1) \otimes V(\delta_2) \longrightarrow V_1 \otimes V_2$$

where the β_i are the isomorphisms induced by the α_i. Let $f_0, \ldots, f_{N_1}$ be the standard basis for $V(\delta_1)$, and $g_0, \ldots, g_{N_2}$ the standard basis for $V(\delta_2)$. Then $f_i \otimes g_j$ is a basis for $V(\delta) = V(\delta_1) \otimes V(\delta_2)$, and one sees that it will be a permutation of the standard basis. Let us assume that the standard ordering is given by

$$f_0 \otimes g_0, f_1 \otimes g_0, \ldots, f_0 \otimes g_1, f_1 \otimes g_1, \ldots.$$

Let $r_i = \beta_1(f_i)$ and $s_j = \beta_2(g_j)$. Then for $P = (P_1, P_2) \in A_1 \times A_2$,

$$\begin{aligned}
\phi_1(P_1) &= [r_0(P_1) : \cdots : r_{N_1}(P_1)], \\
\phi_2(P_2) &= [s_0(P_2) : \cdots : s_{N_2}(P_2)], \text{ and} \\
\phi(P) &= [(r_0 s_0)(P) : (r_1 s_0)(P) : \cdots : (r_{N_1} s_{N_2})(P)] \\
&= [r_0(P_1) s_0(P_2) : \cdots : r_{N_1}(P_1) s_{N_2}(P_2)].
\end{aligned}$$

Then visibly $s(\phi_1(P_1), \phi_2(P_2)) = \phi(P)$. The claim follows. □

Observe that under the hypotheses of the proposition, if ϕ is an embedding, then $\phi(A)$ is simply the Segre embedding applied to the product embedding

$$\phi_1(A_1) \times \phi_2(A_2) \subset \mathbb{P}^{N_1} \times \mathbb{P}^{N_2}.$$

B.3. **Mumford's theta null invariant.** Fix $n \in \mathbb{Z}_{>0}$ and mutually isogenous elliptic curves $E_1, \dots, E_n$. Let

$$A = E_1 \times E_2 \times \cdots \times E_n$$

with its product polarization, i.e., the principal polarization with divisor

$$D = (\{O\} \times E_2 \times \cdots \times E_n) + (E_1 \times \{O\} \times E_3 \times \cdots \times E_n) + \cdots + (E_1 \times \cdots \times \{O\}). \tag{9}$$

Fix $m \geq 2$ and let L be the invertible sheaf associated to the divisor mD, so that $H(L) = A[m]$ and $\delta = (m, m, \dots, m)$. Loop over every theta structure for L and, for each one, compute the associated theta null. We define

$$\mathrm{Mumf}_m(E_1, \dots, E_n)$$

to be the resulting set of theta nulls.

Proposition B.6. *For fixed m and n, we have*

$$\#\,\mathrm{Mumf}_m(E_1, \dots, E_n) \leq m^{2n} \cdot \#\mathrm{Sp}_{2n}(\mathbb{Z}/m\mathbb{Z}).$$

Proof. To choose a theta structure, we first choose a symplectic isomorphism $h : H(\delta) \to H(L)$. The number of such isomorphisms is the size of $\mathrm{Sp}_{2n}(\mathbb{Z}/m\mathbb{Z})$, the group of $2n \times 2n$ symplectic matrices with entries in $\mathbb{Z}/m\mathbb{Z}$. For each symplectic isomorphism h, by Lemma B.3 there are m^{2n} theta structures lying above it. The claim now follows. □

Proposition B.7. *Let E_1 and E_2 be isogenous ordinary elliptic curves over $\mathbb{F}_q$ with complex multiplication by $\mathbb{Q}(\sqrt{-d})$ with d square-free. Let $h(d)$ be the class number of $\mathbb{Q}(\sqrt{-d})$. Suppose $4 \mid m \in \mathbb{Z}_{>0}$ and $h(d) > m^4\#\mathrm{Sp}_{2n}(\mathbb{Z}/m\mathbb{Z})$. Then there exist E_3 and E_4 isogenous to E_1 such that $E_1 \times E_2 \cong E_3 \times E_4$ but $\mathrm{Mumf}_m(E_1, E_2) \neq \mathrm{Mumf}_m(E_3, E_4)$.*

Proof. Let $A = E_1 \times E_2$. Let P be the set of pairs of elliptic curves (E_3, E_4) isogenous to E_1 for which $A \cong E_3 \times E_4$. Suppose for the sake of contradiction that $M := \mathrm{Mumf}_m(E_1, E_2) = \mathrm{Mumf}_m(E_3, E_4)$ for all $(E_3, E_4) \in P$. Construct a map $F : P \to M$ as follows. Given $(E_3, E_4) \in P$, define the invertible sheaves $L_3 = L(mO)$ on E_3, $L_4 = L(mO)$ on E_4, and $L = p_3^* L_3 \otimes p_4^* L_4$ on $E_3 \times E_4$, where the p_i are the obvious projections. Let $\delta = (m, m)$. Then our theta nulls for $E_3 \times E_4$ are computed using theta structures of the form $\alpha : G(\delta) \to G(L)$. Choose a product theta structure α_0. Then define $F(E_3, E_4)$ to be the theta null associated to α_0.

We claim that F is injective. For given a theta null in the image of F, it comes from a product theta structure for some $(E', E'') \in P$. Since $4 \mid m$, by the second part of Proposition B.1 the associated embedding $\phi(A)$ is determined by the theta null $\phi(O)$. By Proposition B.5, we have $\phi(A) = s(\phi_1(E'), \phi_2(E''))$, where s is the Segre map. As ϕ is an embedding (by Proposition B.1), the ϕ_i must also be embeddings. Therefore we can recover the pair (E', E'') from the theta null. This gives a map $\mathrm{im}(F) \to P$ which is a left inverse to F, which shows that F is injective.

Let $\mathscr{O} = \mathrm{End}(E_1)$. The map $\mathrm{Cl}(\mathscr{O}) \to P$ given by $[\mathfrak{a}] \mapsto (\mathfrak{a} * E_1, \mathfrak{a}^{-1} * E_2)$ is injective, so $\#P \geq h(d)$. By Proposition B.6 and our assumption on m we have

$$\#M \leq m^4\#\mathrm{Sp}_4(\mathbb{Z}/m\mathbb{Z}) < h(d) \leq \#P,$$

contradicting the injectivity of F. The claim follows. □

As stated in Appendix A.3, assuming the Generalized Riemann Hypothesis, the class number $h(d)$ is roughly of size $\sqrt{q}$ as $q \to \infty$. Thus the inequality $h(d) > m^4 \# \mathrm{Sp}_{2n}(\mathbb{Z}/m\mathbb{Z})$ in the statement of Proposition B.7 holds for typical use-cases.