Generation of “independent” points on elliptic curves by means of Mordell–Weil lattices

Dmitrii Koshelev[0000−0002−4796−8989]

Computer sciences and networks department, Télécom Paris, France
https://www.telecom-paris.fr
dimitri.koshelev@gmail.com

Abstract. This article develops a novel method of generating “independent” points on an ordinary elliptic curve E over a finite field. Such points are actively used in the Pedersen vector commitment scheme and its modifications. In particular, the new approach is relevant for Pasta curves (of j-invariant 0), which are very popular in the given type of elliptic cryptography. These curves are defined over highly 2-adic fields, hence successive generation of points via a hash function to E is an expensive solution. Our method also satisfies the NUMS (Nothing Up My Sleeve) principle, but it works faster on average. More precisely, instead of finding each point separately in constant time, we suggest to sample several points at once with some probability.

Keywords: elliptic curves · generation of “independent” points · isotrivial elliptic surfaces · Mordell–Weil lattices · Pedersen hash · superelliptic curves · vector commitment schemes

1 Introduction

A *commitment scheme* is a cryptographic primitive that allows one party to commit to a chosen value while keeping it hidden to others, with the ability to reveal the committed value later. Commitment schemes are designed so that the party cannot change the value after they have committed to it. They have important applications in a number of cryptographic protocols including secure coin flipping and zero-knowledge proofs.

There is the classic *Pedersen commitment scheme* [31, Section 3]. It works in any cyclic group with the hard discrete logarithm problem (DLP). However, throughout the article we will deal only with (a large subgroup $\mathbb{G}$ of) the $\mathbb{F}_q$-point group $E(\mathbb{F}_q)$ of an elliptic curve E over a finite field $\mathbb{F}_q$. As is well known, today ordinary (i.e., non-supersingular) curves over fields of large characteristic p are considered the safest. And every cryptographer understands perfectly that the order $\ell := \#\mathbb{G}$ must be prime.

We can use a variant of the original Pedersen commitment (for $n = 1$) to commit to multiple values $(m_1, \cdots, m_n) \in \mathbb{F}_\ell^n$ at once (so-called *vector commitment*). We have to sample a vector of public points $(P_1, \cdots, P_n) \in \mathbb{G}^n$, along with a fixed generator $P_0 \in \mathbb{G}$. Then the commitment is just the sum $m_0P_0 + \sum_{i=1}^n m_iP_i$, where $m_0 \in \mathbb{F}_\ell$ is an auxiliary value to ensure the security of the scheme.

Of course, we can simply commit to each m_i individually, but this solution is much less efficient in terms of memory and computing resources. Indeed, the full multi-scalar multiplication can be performed much more rapidly than each one $r_iP_0 + m_iP_1$ alone. Here $(r_1, \cdots, r_n) \in \mathbb{F}_\ell^n$ is another random vector playing the role of m_0. Besides, vector commitments provide a way to store or transmit only one element of the group $\mathbb{G}$ instead of a vector from $\mathbb{G}^n$. In real-world cryptography it happens that n reaches huge numbers such as $\approx 2^{30}$ as indicated, e.g., in [10].

The aforementioned primitive is also known as the *Pedersen hash* $\mathbb{F}_\ell^n \to \mathbb{G}$ (see, e.g., [5]). It is provably secure, because its resistance is based on the *multi-dimensional DLP*. According to cryptanalysis performed in [18], [19] the given problem does not seem to be simpler in general than the classical DLP. Another advantage of the Pedersen hash is in its additive homomorphic property. All this positively distinguishes it from (Merkle hash tree [30] using) faster standard hash functions such as SHA-3 (Keccak).

Certainly, the Pedersen scheme is resistant only if the points P_i are *"independent"*, that is nobody knows a non-trivial linear relation between them. In other words, it is hard to find values $(k_1, \cdots, k_n) \in \mathbb{F}_\ell^n$ such that $k_0P_0 = \sum_{i=1}^n k_iP_i$ and at least one $k_i \neq 0$. Therefore, every point P_i must be generated in a transparent way. Be careful that, from the mathematical point of view, conversely, any two points depend on each other, since the group $\mathbb{G}$ is prime.

Over time, a malicious user may find some relation between the points P_i through a kind of brute-force attack. We have no guarantee that this cannot happen for concrete points, even though the multi-dimensional DLP is intractable in the general case. The fact is that for the large n there is the huge number of linear relations. At the same time, it is enough to find just one to break the Pedersen scheme. That is why it is desirable for security to periodically change the points.

The author of [16] prefers the word "basis" and he admits that "updating the basis at every round is inefficient". Let's assume the opposite situation when the points P_i remain the same for a long time. Even in this situation, the task of their rapid generation is still important. First, the storage (resp. transmission) of the points requires a lot of memory (resp. bandwidth). And second, there is ground for a potential fault attack, because it is enough for an adversary to replace just one point.

As is known, the points P_i can be obtained by means of a hash function $\mathcal{H} : \{0,1\}^* \to \mathbb{G}$, for example as $P_i = \mathcal{H}(\mathrm{seed}||i)$ (cf. [5, Section 5.1]). This approach forces to evaluate $\mathcal{H}$ exactly n times. The fastest constructed hash functions to elliptic curves extract one radical $\sqrt[m]{\cdot}$ in $\mathbb{F}_q$ for some $m \in \mathbb{N}$. Their actual classification is given in [25, Tables 1-2] (cf. [15]). And the existence of $\mathcal{H}$ without radicals at all is highly unlikely. There is plenty of material devoted to extracting $\sqrt[m]{\cdot}$, starting with the seminal work of *Adleman, Manders, and Miller*

[1]. But despite this, $\sqrt[m]{\cdot}$ continues to be a much more expensive operation than the arithmetic ones in $\mathbb{F}_q$, namely $+$, $-$, $*$, and even $/$.

It is also worth noting the *Kate–Zaverucha–Goldberg (KZG) commitment scheme* (or just the *Kate commitment*) [23] based on pairings of elliptic curves. At the moment, this scheme is recognized by the cryptographic society as one of the best from the computational point of view. However, to deploy it we need a trusted setup. More concretely, the scheme substantially uses the points s^iP_0 (with a secret $s \in \mathbb{F}_\ell$) rather than arbitrary “independent” points. By the way, such points can be equally utilized in the Pedersen scheme.

In comparison with the Pedersen protocol, KZG one is in fact a *polynomial commitment*. By definition, it allows a prover to commit to a polynomial $f = \sum_{i=0}^{n-1} m_{i+1}x^i$, with the property that the prover can later convince a verifier of the equality $f(\alpha) = \beta$, given α, $\beta \in \mathbb{F}_\ell$. In addition, until n points of the form $\big(\alpha, f(\alpha)\big)$ are revealed, the polynomial f remains hidden, as should be clear.

It turns out that the Pedersen vector commitment can be supplemented to give rise to a polynomial commitment without a trusted setup (see, e.g., [9, Section 3], [14, Section 4.5]). Incidentally, those sources are dedicated to a protocol of so-called *recursive proof composition* using an *amicable pair* [37] of prime-order elliptic curves. More precisely, the latter are non-pairing-friendly curves $y^2 = x^3 + 5$ of j-invariant 0 under the name *Pasta curves (Pallas and Vesta)* [20] (cf. [21]).

These curves (and many others [2]) are defined over *highly* 2*-adic fields*, i.e., $2^e \mid q - 1$ for a fairly large $e \in \mathbb{N}$. Such fields allow to utilize the *fast Fourier transform (FFT)* [14, Section 4.2] to speed up the polynomial arithmetic in numerous modern protocols. The downside is that one cannot express a square root in $\mathbb{F}_q$ via one exponentiation in $\mathbb{F}_q$. We can always resort to the *Tonelli–Shanks algorithm* [13, Algorithm 5.14], but it is substantially slower than the exponentiation operation. That is why we should avoid square roots as far as possible.

2 Underlying mathematical preliminaries

Consider a finite field $\mathbb{F}_q$ of characteristic $p > 3$. The notion of an *elliptic surface* [32, Section 5], [35, Chapter III] over $\mathbb{F}_q$ is key for us. Without loss of generality, we can confine to a short Weierstrass form

$$\mathcal{E}\colon y^2 = x^3 + a(t)x + b(t) \quad \subset \quad \mathbb{A}^3_{(x,y,t)}$$

with polynomial coefficients $a(t)$, $b(t) \in \mathbb{F}_q[t]$. As usual, $\mathcal{E}$ is interpreted as an elliptic curve over the function field $F := \mathbb{F}_q(t)$ in one variable. From time to time, we will equally need the field $F' := \overline{\mathbb{F}_q}(t)$ over the algebraic closure $\overline{\mathbb{F}_q}$.

Recall that the *Mordell–Weil group* of $\mathcal{E}$ is the abelian group $\mathcal{E}(F)$ of all F-points on $\mathcal{E}$. Due to a special case of the *Mordell–Weil theorem* [32, Section 3.3] the group $\mathcal{E}(F)$ is finitely generated. Its rank r is called the *Mordell–Weil rank* of $\mathcal{E}$. As always, we denote by $\mathcal{E}(F)_{tor}$ the (finite) torsion subgroup of $\mathcal{E}(F)$. The quotient $\mathcal{E}(F)/\mathcal{E}(F)_{tor} \simeq \mathbb{Z}^r$ enjoys a positive-definite quadratic form $\widehat{h}$ under the name the *canonical height* or the *Néron–Tate height* [35, Section III.4]. The corresponding symmetric bilinear form $\langle \cdot, \cdot \rangle$ and r-dimensional lattice are said to be the *height pairing* and the *Mordell–Weil lattice* respectively (see [32, Section 6.5]).

As is customary, we are given an ordinary elliptic $\mathbb{F}_q$-curve $E\colon y^2 = x^3+ax+b$. Throughout the article we assume the coincidence of the j-invariants: $j(E) = j(\mathcal{E})$. Such a surface $\mathcal{E}$ is said to be *isotrivial*. Note that $r = 0$ for *trivial (constant) elliptic surfaces* (s.t. $E \simeq_{F'} \mathcal{E}$), because elliptic curves are not rational. Hence trivial surfaces are excluded from our consideration. By definition, $\mathcal{E}$ is a non-trivial twist of E. Since $p > 3$, the results of [36, Section X.5] about twisting elliptic curves are still relevant even though F is not a perfect field.

Let's define the function

$$c(t) := \begin{cases} \dfrac{a(t)b(t)}{ab} & \text{if} \qquad ab \neq 0, \text{ i.e., } j(E) \notin \{0, 1728\}, \\ \dfrac{a(t)}{a} & \text{if} \qquad b = 0, \text{ i.e., } j(E) = 1728, \\ \dfrac{b(t)}{b} & \text{if} \qquad a = 0, \text{ i.e., } j(E) = 0. \end{cases}$$

Let $d \in \{2, 4, 6\}$ be the order of the cyclic group $\mathrm{Aut}(E)$ and $s := \sqrt[d]{c(t)}$. The curves E, $\mathcal{E}$ are isomorphic precisely over the Kummer extension $F(s)/F$ of degree d. It is the function field of the superelliptic curve $C\colon s^d = c(t) \subset \mathbb{A}^2_{(t,s)}$. The corresponding isomorphism (from the proof of [36, Proposition III.1.4.(b)]) has the form

$$\varphi\colon \mathcal{E} \to E \qquad (x, y) \mapsto \left(\frac{x}{z^2}, \frac{y}{z^3}\right),$$

where

$$z := \begin{cases} \dfrac{as}{a(t)} = \dfrac{b(t)}{bs} & \text{if} \qquad ab \neq 0, \text{ i.e., } j(E) \notin \{0, 1728\}, \\ s & \text{otherwise.} \end{cases}$$

It is worth saying that points from $\mathcal{E}(F)$ are nothing but $\mathbb{F}_q$-sections $\mathbb{A}^1_t \to \mathcal{E}$ of the projection pr_t to the variable t. Below $\mathcal{E}_t$ stands for its fiber over $t \in \overline{\mathbb{F}_q}$. Similarly, $\varphi_t : \mathcal{E}_t \to E$ denotes the specialization of φ. There is only a finite number of degenerate fibers, namely those for which the discriminant $\Delta(\mathcal{E}_t) = -16\big(4a^3(t) + 27b^2(t)\big)$ vanishes. Clearly, this happens exactly when φ_t is meaningless. In this situation, $\varphi_t(x, y) = \infty$ is a convenient notation (for any map). Finally, given $t \in \mathbb{F}_q$, the condition $\mathcal{E}_t \simeq_{\mathbb{F}_q} E$ occurs iff φ_t is defined over $\mathbb{F}_q$ iff $s \in \mathbb{F}_q^*$.

We need certain results about the endomorphism rings of elliptic curves, which can be found in any classical source like [36, Sections III.9, V.3]. Since E is an ordinary $\mathbb{F}_q$-curve, new endomorphisms on it are not added when extending $\mathbb{F}_q$. It is also readily shown by exploiting φ that the coefficients of $F(s)$-endomorphisms on $\mathcal{E}$ in fact belong to F. Eventually, we have:

$$\mathrm{End}(E/\mathbb{F}_q) = \mathrm{End}(E/F) = \mathrm{End}\big(E/F(s)\big) \simeq \mathrm{End}\big(\mathcal{E}/F(s)\big) = \mathrm{End}(\mathcal{E}/F).$$

By abuse of notation, we will identify all these rings by means of the single symbol $\mathcal{O}$.

As is well known, $\mathcal{O}$ is an order in the imaginary quadratic field $\mathbb{Q}(\sqrt{D})$, where $D := t_q^2 - 4q$ and t_q is the trace (of the Frobenius) of $E/\mathbb{F}_q$. Furthermore, $\mathcal{O} = \mathbb{Z} \oplus \mathbb{Z}\phi$ for some endomorphism ϕ (with the dual one $\widehat{\phi}$). Recall that its characteristic (and at the same time minimal) polynomial equals

$$\chi_\phi = x^2 - \mathrm{tr}(\phi)x + \deg(\phi), \qquad \text{where} \qquad \mathrm{tr}(\phi) := \phi + \widehat{\phi}, \qquad \deg(\phi) := \phi\cdot\widehat{\phi}.$$

As explained in [4, Appendix A], original *Schoof's algorithm* (see, e.g., [13, Algorithm 2.4]) for computing t_q is easily modified to compute $\mathrm{tr}(\phi)$ whenever ϕ is the composition of a bounded number of small-degree isogenies.

There is a natural action of the group $\mu_d \simeq \mathbb{Z}/d$ on the curve C and hence on its Jacobian J_C. Let's introduce the number

$$k := \max\left\{k' \in \mathbb{N} \mid \text{exists a surjective } \mu_d\text{-equivariant } \mathbb{F}_q\text{-morphism } J_C \to E^{k'}\right\},$$

where μ_d acts diagonally on $E^{k'}$. Evidently, k does not exceed the geometric genus of C. And when k attains the genus, J_C is said to be a *ρ-maximal* (or *singular*) *abelian variety* [6, Proposition 3].

By virtue of [27] (cf. [24, Sections 6, 7]) we have the sequence of homomorphisms of $\mathcal{O}$-modules

$$\mathcal{E}(F) \simeq \mathrm{Mor}_{\mu_d}(C, E) \to \mathrm{Hom}_{\mu_d}(J_C, E) \simeq \mathrm{Hom}(E^k, E) \simeq \mathrm{End}(E)^k. \qquad (1)$$

The first homomorphism maps $P \mapsto \varphi_P$ through φ in a clear way. The kernel of the second one consists of constant morphisms, which implies the equality $r = 2k$. At last, the third one is not in any way a canonical isomorphism.

Likewise, we possess the sequence starting from

$$\mathcal{E}\big(F(s)\big) \simeq E\big(F(s)\big) \simeq \mathrm{Mor}(C, E).$$

Looking ahead, this $\mathcal{O}$-module gives an advantage over $\mathcal{E}(F)$ whenever the rank of the former is greater than that of the latter. Indeed, in Algorithm 1 one can evaluate any (not necessarily μ_d-equivariant) covers $\psi_1, \cdots, \psi_n : C \to E$ that are independent over $\mathcal{O}$. However, $\mathcal{E}\big(F(s)\big)$ is an awkward object that is more difficult to analyze than $\mathcal{E}(F)$. In Section 5.1 we carry out such an analysis in a simple example.

3 New generation method and its running time

Let us keep the notation of the previous section. At the same time, consider an arbitrary cyclic $\mathbb{F}_q$-cover $\chi\colon C \to \mathbb{P}^1$ of degree $m \mid q-1$. In other words, the curve can be represented in the form $C\colon v^m = f(u) \subset \mathbb{A}^2_{(u,v)}$ for some $f \in \mathbb{F}_q[u]$ without roots of multiplicity $\geqslant m$. In particular, the earlier coordinates t, s are expressed via rational $\mathbb{F}_q$-functions in u, v and vice versa: $(t, s) = \tau(u, v)$ and

$(u, v) = \tau^{-1}(t, s)$. When $m = d$, for our purposes, it will be sufficient to take $f = c(t)$ (or, equivalently, $\chi = pr_t$) and $\tau = \mathrm{id}$.

Pick any points $P_1, \cdots, P_n \in \mathcal{E}(F) \setminus \mathcal{E}(F)_{tor}$ linearly independent over $\mathcal{O}$. Given $u \in \mathbb{F}_q$, the condition $\mathcal{E}_t \simeq_{\mathbb{F}_q} E$ evidently amounts to the fact that $v = \sqrt[m]{f(u)} \in \mathbb{F}_q$, unless $\tau(u, v)$ is meaningless or $\mathcal{E}_t$ is singular. If actually $v \in \mathbb{F}_q$, we obtain the n points $P_i(t) \in \mathcal{E}_t(\mathbb{F}_q) \simeq E(\mathbb{F}_q)$ at least for integral P_i. Since in discrete logarithm cryptography the group $E(\mathbb{F}_q)$ is (almost) prime, the specialized points (very often) become dependent for $n > 1$. However, we do not see a non-trivial relation between them. Finally, whenever $\mathbb{G} \subsetneq E(\mathbb{F}_q)$, it remains to clear the cofactor to definitely fall into $\mathbb{G}$, but the resulting points are still "independent".

It is worth avoiding torsion points P_i, because they and hence $P_i(t)$ have tiny orders with respect to ℓ. We also emphasize that the points must be independent precisely over $\mathcal{O}$, and not just over $\mathbb{Z}$. Although 1, ϕ are linearly independent endomorphisms, their restrictions on $\mathbb{G}$ are not. Indeed, from $\mathbb{G} \simeq \mathbb{Z}/\ell$ it follows that $\mathrm{End}(\mathbb{G}) \simeq \mathbb{F}_\ell$. On the other hand, in practice $\mathbb{G} = E(\mathbb{F}_q)[\ell]$. As a consequence, there exists $\lambda \in \mathbb{F}_\ell$ such that $\phi(P) = \lambda P$ for all $P \in \mathbb{G}$. In other terms, λ is a root of the characteristic polynomial $\chi_\phi \in \mathbb{F}_\ell[x]$, i.e., λ is an eigenvalue of $\phi|_{E[\ell]}$. Eventually, knowing χ_ϕ, we can determine λ with 50-percent confidence (100-percent one when ϕ is easy to evaluate).

Realizing ϕ as an abstract element of $\mathbb{Q}(\sqrt{D})$, we immediately get χ_ϕ, because $\widehat{\phi}$ is the complex conjugate of ϕ. In turn, the latter can be found via randomized *Bisson–Sutherland's algorithm* [7] (resp. deterministic *Kohel's* one [17, Section 25.4.2]). While in the worst case its running time is sub-exponential (resp. exponential), the curve E is usually generated once and for all by a certain regulator. It is not ruled out that ϕ is in its sleeve. That is why we should not rely on the hardness of finding ϕ. In addition, implementors of elliptic cryptosystems often choose E for which, conversely, ϕ is a small-degree endomorphism known to all. This is done in order to enjoy the *GLV (Galbraith–Lambert–Vanstone) scalar multiplication method* [17, Section 11.3.3].

Fix a hash function $\eta : \{0, 1\}^* \to \mathbb{F}_q$. We need to change $u = \eta(\mathrm{seed}||i)$, where $i \in \mathbb{N}$, while the desired requirement $\sqrt[m]{f(u)} \in \mathbb{F}_q$ is not met. So the new generation method (formalized in Algorithm 1) is a priori non-constant-time. Nevertheless, this is not dangerous as regards timing attacks, because $\mathrm{seed}||i$ is public information. Frankly speaking, we have to continue sampling u when we encounter one of the degenerate situations $\tau(u, v) = \infty$, $\varphi_t(x, y) = \infty$, or $P_i(t) = \infty$. They arise with negligible probability, so we do not pay attention to them anymore, with the permission of the reader.

We have the *power residue symbol* $\left(\frac{x}{q}\right)_m := x^{(q-1)/m}$ generalizing the Legendre symbol (for $m = 2$). It is obviously a surjective homomorphism $\mathbb{F}_q^* \to \mu_m$ to the group of all m-th roots of unity. As is well known, to determine whether $f(u)$ is an m-th residue in $\mathbb{F}_q$ it is sufficient to check the equality $\left(\frac{f(u)}{q}\right)_m = 1$. It turns out that for $m \leqslant 11$ computing the residue symbol is a much cheaper operation due to [22] than extracting any root in $\mathbb{F}_q$. Thus, unlike the generation method with a hash function $\mathcal{H} : \{0,1\}^* \to E(\mathbb{F}_q)$, we obtain a set of “independent” $\mathbb{F}_q$-points on E by extracting only one root in $\mathbb{F}_q$ (of degree m).

The same thought occurs in [41, Section 3] to speed up (de)compression in the post-quantum *protocol SIDH (Supersingular Isogeny Diffie–Hellman).* Instead of applying a constant-time encoding (essentially $\mathcal{H}$) to an elliptic curve, the authors of that article prefer to “subvert” it to produce at once two independent (in the strict sense) points with high probability. They agree with us that a randomized algorithm with one square root (and several Legendre symbols) is faster on average than a deterministic one with two square roots in the same field. This is especially relevant for SIDH, since the given protocol is deployed over a highly 2-adic field.

Algorithm 1: New generation method

Data: a seed $\in \{0,1\}^*$ and a hash function $\eta\colon \{0,1\}^* \to \mathbb{F}_q$,
an elliptic $\mathbb{F}_q$-curve E and an elliptic $\mathbb{F}_q$-surface $\mathcal{E}$ of the same j-invariant,
points $P_1, \cdots, P_n \in \mathcal{E}(F) \setminus \mathcal{E}(F)_{tor}$ independent over $\mathcal{O}$,
a superelliptic curve $C\colon v^m = f(u)$ (where $m \mid q-1$ and $f \in \mathbb{F}_q[u]$) such that $E \simeq \mathcal{E}$ over the function field $\mathbb{F}_q(C)$.
Result: n “independent” points in $E(\mathbb{F}_q)$.

```
begin
    i := 0;
    u := η(seed||i);
    while (f(u)/q)_m ≠ 1 do
        i := i + 1;
        u := η(seed||i);
    end
    v := ᵐ√f(u);
    (t, s) := τ(u, v);
    return φ_{P_1}(t, s), ···, φ_{P_n}(t, s).
end
```

We do not claim the authorship of the following lemma, but we prove it for the sake of completeness.

Lemma 1. *Let $m \mid q-1$ and $f \in \mathbb{F}_q[u]$ be a polynomial without roots of multiplicity $\geqslant m$. Given a random $u \in \mathbb{F}_q$, the probability that $\sqrt[m]{f(u)} \in \mathbb{F}_q$ equals*

$$\rho := \frac{N}{q} = \frac{1}{m} + O\Big(\frac{1}{\sqrt{q}}\Big), \qquad \textit{where} \qquad N := \#\big\{u \in \mathbb{F}_q \mid \sqrt[m]{f(u)} \in \mathbb{F}_q\big\},$$

m, f are fixed, but $q \to +\infty$.

Proof. We are going to extend the reasoning of [13, Section 8.2.1] from the case $m = 2$. As is customary, $p_a \in \mathbb{N}$ stands for the arithmetic genus of $C\colon v^m = f(u)$. Let n_0 be the number of $\mathbb{F}_q$-points on C of the form $(u, 0)$ and n_∞ be the number of those at infinity. Trivially, $p_a,\ n_0,\ n_\infty = O(1)$. Since C is known to be an absolutely irreducible curve, we have the *Weil–Aubry–Perret inequality*

$$|\#C(\mathbb{F}_q) - (q+1)| \leqslant 2p_a\sqrt{q} \qquad [3, \text{Corollary } 2.4],$$

where n_∞ is taken into account in $\#C(\mathbb{F}_q)$. Therefore, $\#C(\mathbb{F}_q) - q = O(\sqrt{q})$.

For compactness, we use the auxiliary notation

$$\alpha(u) := \sum_{i=0}^{m-1} \Big(\frac{f^i(u)}{q}\Big)_m, \qquad A := \sum_{u \in \mathbb{F}_q} \alpha(u).$$

From the equality $(x^m - 1)/(x-1) = \sum_{i=0}^{m-1} x^i$ it follows that

$$\alpha(u) = \begin{cases} m & \text{if} \quad \sqrt[m]{f(u)} \in \mathbb{F}_q^*, \\ 1 & \text{if} \quad f(u) = 0, \\ 0 & \text{otherwise.} \end{cases}$$

Consequently,

$$\#C(\mathbb{F}_q) = A + n_\infty, \qquad N = \frac{A + n_0(m-1)}{m}.$$

Eventually,

$$\rho - \frac{1}{m} = \frac{A - q + n_0(m-1)}{mq} = \frac{\#C(\mathbb{F}_q) - n_\infty - q + n_0(m-1)}{mq} =$$
$$= \frac{O(\sqrt{q}) - n_\infty + n_0(m-1)}{mq} = O\Big(\frac{1}{\sqrt{q}}\Big).$$

Lemma 2. *The average-case complexity of Algorithm* 1 *is that of computing* m *symbols* $\big(\frac{\cdot}{q}\big)_m$ *and one radical* $\sqrt[m]{\cdot}$ *in* $\mathbb{F}_q$.

Proof. We suggest to consider the probability ρ_k that for $k \in \mathbb{N}$ random independent elements $u_i \in \mathbb{F}_q$ the root $\sqrt[m]{f(u_i)} \notin \mathbb{F}_q$ and for the $(k+1)$-th one, conversely, $\sqrt[m]{f(u_{k+1})} \in \mathbb{F}_q$. By virtue of Lemma 1 we get:

$$\rho_k = x^k \cdot \frac{1}{m} = \frac{(m-1)^k}{m^{k+1}}, \qquad \text{where} \qquad x := \frac{m-1}{m}.$$

Denote by X the random variable returning $k+1$ with the probability ρ_k. It corresponds to the number of symbols $\big(\frac{\cdot}{q}\big)_m$ arising during the work of our algorithm.

By definition of average-case complexity, we need to compute the expected value

$$\mathbb{E}[X] = \sum_{k=0}^{\infty} (k+1)\rho_k = \frac{1}{m}\sum_{k=0}^{\infty}(k+1)x^k$$

It is a classical fact that under the condition $|x| < 1$ (fulfilled for $m \in \mathbb{N}$) the geometric series $\sum_{k=0}^{\infty} x^k = 1/(1-x)$, hence

$$\mathbb{E}[X] = \frac{1}{m}\Big(\sum_{k=0}^{\infty} x^{k+1}\Big)' = \frac{1}{m}\Big(\frac{x}{1-x}\Big)' = \frac{1}{m(1-x)^2} = m.$$

Bearing in mind the final m-th root extraction, the lemma is proved.

Due to (1) the number $n = r/2 \in \mathbb{N}$ is the most optimal in Algorithm 1. Besides, the smaller number m, simpler methods exist (over a general field $\mathbb{F}_q$) for finding $\big(\frac{\cdot}{q}\big)_m$ and $\sqrt[m]{\cdot}$, not to mention Lemma 2. The minimal possible m is the cyclic analogue

$$\gamma_c := \min\big\{ \deg(\chi) \mid \chi\colon C \to \mathbb{P}^1 \text{ is a cyclic (i.e., Kummer) } \mathbb{F}_q\text{-cover}\big\}$$

of the *gonality* γ [28, Section 6.5.3] of the curve C. Trivially, $2 \leqslant \gamma \leqslant \gamma_c \leqslant d \leqslant 6$.

Thus, we see that our generation method works more productively if the fraction $\delta := r/\gamma_c$ is greater. It is natural to call it the *relative Mordell–Weil rank* of $\mathcal{E}$. Of course, this notion is useless when $j(\mathcal{E}) \notin \{0, 1728\}$, that is $d = 2$. In the opposite case, it seems quite difficult to determine the exact value γ_c, so it is reasonable to also define $\delta(\chi) := r/\deg(\chi)$. Then $\delta = \max_\chi\{\delta(\chi)\}$.

The problem of maximizing δ has much in common with a classic one of pure mathematics about how big the conventional Mordell–Weil rank r can theoretically be for elliptic surfaces. Over an algebraically closed field of zero characteristic (or just $\mathbb{C}$) the current record equals 68 for the surfaces $\mathcal{E}_m\colon y^2 = x^3 + t^m + 1$ such that $360 \mid m$ (see [32, Section 13.2]). Be careful, there is a discrepancy with our previous notation $\mathcal{E}_t$ of a fiber.

Circumstances are drastically different in a prime characteristic p. There is no upper bound on r in the class of non-isotrivial surfaces [39], whose j-invariants are always ordinary. The same is true for (isotrivial) surfaces of supersingular j-invariants [38]. In fact, among those it is enough to confine to $\mathcal{E}_{p^e+1}$ (where $e \in \mathbb{N}$) as shown in [33]. Surprisingly, in accordance with the articles [8], [12] the rank r can be made arbitrarily large even in the class of isotrivial surfaces of ordinary j-invariants. However, it is not clear how constructive the results established in those articles.

4 The case of j-invariant 0

Hereafter we focus on elliptic curves of j-invariant 0, that is $a = a(t) = 0$, because they are popular in practice. Since we deal only with ordinary curves, $3 \mid q-1$ or, equivalently, a primitive cubic root $\omega = (-1+\sqrt{-3})/2$ of unity lies in $\mathbb{F}_q$ (see [36, Example V.4.4]). There is on E, $\mathcal{E}$ the automorphism $[\omega](x,y) = (\omega x, y)$ of order 3 and moreover $\mathcal{O} = \mathbb{Z}[\omega]$.

For any $m \mid q-1$ and $c \in \mathbb{F}_q^*$ consider the twist $\mathcal{E}_m : y^2 = x^3 + t^m + c$ of the aforementioned elliptic surface. Remarkably, the group $\mathcal{E}_m(F)$ is torsion-free regardless of m. Further, for $\mathcal{E}_m$ to be a *rational surface* it is necessary and sufficient that $m \leqslant 6$. These and other details about the surfaces $\mathcal{E}_m$ can be found in [33]. And the general theory of rational elliptic surfaces is discussed, e.g., in [32, Chapter 7].

It is also natural to denote the curve C from the previous sections by $C_m : bs^6 = t^m + c$. Its geometric genus $g(C_m)$ can be computed via a formula from [28, Section 5.1]. In this article we decided to focus only on the case $m \leqslant 6$, because it is the simplest and investigated in the literature. For instance, C_6 is a twist of the *Fermat sextic curve* [6, Proposition 7], [27, Example 4.3]. We hope to study the opposite case $m > 6$ in the future articles. So from now on, we represent the curve in the form $C_m : t^m = bs^6 - c$. In terms of Section 3 this means that $f = bs^6 - c$, i.e., $\chi = pr_s$ and $(t, s) = (v, u)$.

Table 1 exhibits main information about the rational surfaces $\mathcal{E}_m$ over $\overline{\mathbb{F}_q}$. We provide it for the convenience of the reader, no more no less. First, (up to an isomorphism) the Mordell–Weil lattices $\mathcal{E}_m(F')$ are dual to some *root lattices* (E_8 is self-dual). By the way, a good survey of root lattices and their dual ones is given in [32, Section 2.3]. And second, the column $d_{\min}$ (resp. disc) contains the minimum norm (resp. discriminant) of the lattices.

<table>
<tr><th>m</th><th>$\mathcal{E}_m(F')$</th><th>$\delta(\chi)$</th><th>$d_{\min}$</th><th>disc</th><th>$g(C_m)$</th></tr>
<tr><td>1</td><td>0</td><td>0</td><td></td><td></td><td>0</td></tr>
<tr><td>2</td><td>A_2^*</td><td>1</td><td>2/3</td><td>1/3</td><td>2</td></tr>
<tr><td>3</td><td>D_4^*</td><td>4/3</td><td>1</td><td>1/4</td><td>4</td></tr>
<tr><td>4</td><td>E_6^*</td><td>3/2</td><td>4/3</td><td>1/3</td><td>7</td></tr>
<tr><td>5</td><td rowspan="2">E_8</td><td>8/5</td><td rowspan="2">2</td><td rowspan="2">1</td><td rowspan="2">10</td></tr>
<tr><td>6</td><td>4/3</td></tr>
</table>

Table 1. The rational surfaces $\mathcal{E}_m/\overline{\mathbb{F}_q}$

Note that $0 < 1 < 4/3 < 3/2 < 8/5$ for the values from the column $\delta(\chi)$. For Algorithm 1 the surface $\mathcal{E}_6$ does not provide any advantage with respect to $\mathcal{E}_3$. That is why we do not consider it in detail. In turn, the surface $\mathcal{E}_5$ is the best. Unfortunately, $5 \nmid q - 1$ for Pasta curves, hence for them we have to be content with $\mathcal{E}_4$. So we are able to generate 3 "independent" $\mathbb{F}_q$-points on E in such a way that the average running time coincides with that of computing 4 symbols $\left(\frac{\cdot}{q}\right)_4$ and one quartic root in $\mathbb{F}_q$. The latter can be obviously represented as 2 successive square roots. Alternatively, one can apply (a variation of) the Adleman–Manders–Miller algorithm in order to directly find $\sqrt[4]{\cdot}$.

It is time to remind that Pasta curves were designed, taking into account the existence of $\mathbb{F}_q$-isogenies of small degree (namely 3) from auxiliary elliptic curves of j-invariants different from 0. As a result, the state-of-the-art hash function for Pasta curves is the *Wahby–Boneh hash function* $\mathcal{H}_{WB}$ [40] based on the *simplified SWU (Shallue–van de Woestijne–Ulas)* one [15, Section 6.6.2]. It requires to compute one square root in $\mathbb{F}_q$ during the execution.

As we said before, $\sqrt{\cdot}$ (as well as $\sqrt[4]{\cdot}$) is a laborious operation over highly 2-adic fields and Pasta curves are defined over such fields. Fortunately, their fields $\mathbb{F}_q$ are not highly 3-adic (more concretely, $27 \nmid q-1$). Therefore, the cubic root extraction in $\mathbb{F}_q$ can be performed by one exponentiation by virtue of [11, Proposition 1]. Thus, instead of $\mathcal{E}_4$, it might be wise to use the surface $\mathcal{E}_3$ to obtain 2 “independent” points at the price of 3 symbols $\left(\frac{\cdot}{q}\right)_3$ and one cubic root in $\mathbb{F}_q$.

Finally, the surface $\mathcal{E}_2$ is useless, because we always have the opportunity to exploit the more advantageous surface $\mathcal{E}_3$. Indeed, we are not aware of practical situations in which $\mathbb{F}_q$ is a highly 3-adic field and at the same time $4, 5 \nmid q-1$. Even if this situation occurs (and $\mathcal{H}_{WB}$ is not applicable), it is enough to use the universal *SW hash function* [13, Sections 8.3.4, 8.4.2] with the same running time as for the method built on $\mathcal{E}_2$.

To sum up, methods of generating n “independent” $\mathbb{F}_q$-points on elliptic curves of j-invariant 0 are exhibited in Table 2. To justify its bottom row, in the next section we will demonstrate that for all $m \leqslant 5$ there is $c \in \mathbb{F}_q^*$ for which the Mordell–Weil lattices of $\mathcal{E}_m/\mathbb{F}_q$ and $\mathcal{E}_m/\overline{\mathbb{F}_q}$ coincide. To be precise, we will explicitly construct r (resp. $r/2$) minimal points $P_i \in \mathcal{E}_m(F)$ independent over $\mathbb{Z}$ (resp. $\mathbb{Z}[\omega]$). We are guided by the fact that the size of point formulas is proportional to their canonical height. Furthermore, P_i form a basis (cf. [29]), although this property is not applied anywhere by us.

method	n	average complexity	conditions on q
classical with $\mathcal{H}\colon \{0,1\}^* \to E(\mathbb{F}_q)$	1	[25, Tables 1-2]	($\sqrt{\cdot}$ for $\mathcal{H}_{WB}$)
new with $\mathcal{E}_m$, where $2 \leqslant m \leqslant 5$	$m-1$	$m\left(\frac{\cdot}{q}\right)_m + \sqrt[m]{\cdot}$	$m \mid q-1$

Table 2. Generation methods for elliptic $\mathbb{F}_q$-curves E of j-invariant 0

5 Linearly independent points in $\mathcal{E}_m(F)$

In the current section we tacitly resort to the computer algebra system Magma. The corresponding code is loaded on the web page [26]. Besides, we will regularly use a folklore result that, given a pair of lattices $L' \subset L$ of the same rank, the squared index $[L : L']^2 = \mathrm{disc}(L')/\mathrm{disc}(L)$. In particular, $L = L'$ if and only if $\mathrm{disc}(L) = \mathrm{disc}(L')$.

5.1 The case $m = 2$

Assume that $\sqrt[3]{c} \in \mathbb{F}_q$. It is readily seen that the points $P_i := (-\omega^{i-1}\sqrt[3]{c}, t)$ belong to $\mathcal{E}_2(F)$. Any two of them are clearly independent over $\mathbb{Z}$ and dependent over $\mathbb{Z}[\omega]$. The height pairing on the sublattice $\langle P_1, P_2 \rangle \subset \mathcal{E}_2(F)$ is given by the Gram matrix

$$M = \begin{pmatrix} \dfrac{2}{3} & -\dfrac{1}{3} \\ -\dfrac{1}{3} & \dfrac{2}{3} \end{pmatrix},$$

where the i-th row and column correspond to P_i. Since $\det(M) = 1/3 = \mathrm{disc}(\mathrm{A}_2^*)$, the minimal points P_1, P_2 in fact constitute a $\mathbb{Z}$-basis of $\mathcal{E}_2(F) = \mathcal{E}_2(F') \simeq \mathrm{A}_2^*$. Consequently, $P := P_1$ is a generator over $\mathbb{Z}[\omega]$.

In comparison with $m > 2$, the case under consideration is easier, hence let's dwell on it in more detail. The curve $C_2\colon t^2 = bs^6 - c$ is a famous hyperelliptic curve of geometric genus 2 (see, e.g., [6, Example 1]). There are two quadratic $\mathbb{F}_q$-covers

$$\begin{array}{lll} \varphi_P\colon C_2 \to E & (s,t) \mapsto \left(\dfrac{-\sqrt[3]{c}}{s^2}, \dfrac{t}{s^3} \right), & \quad E\colon y^2 = x^3 + b, \\ & & \text{where} \\ \varphi'\colon C_2 \to E' & (s,t) \mapsto (bs^2, bt), & \quad E'\colon y^2 = x^3 - b^2 c. \end{array}$$

Notice that $\varphi_P \in \mathrm{Mor}_{\mu_6}(C_2, E)$ as always, but $\varphi' \notin \mathrm{Mor}_{\mu_6}(C_2, E')$. This immediately implies independency of φ_P, φ' over $\mathbb{Z}[\omega]$ after identifying the curves E, E' by an isomorphism, which exists at most over $\mathbb{F}_{q^6}$. In other words, $J_{C_2} \sim_{\mathbb{F}_q} E \times E' \sim_{\mathbb{F}_{q^6}} E^2$.

By analogy with the sequence (1), there are two $\mathbb{F}_q$-covers $C_2 \to E$ independent over $\mathbb{Z}[\omega]$ if and only if

$$J_{C_2} \sim_{\mathbb{F}_q} E^2 \ \Leftrightarrow\ E \sim_{\mathbb{F}_q} E' \ \Leftrightarrow\ E \simeq_{\mathbb{F}_q} E' \ \Leftrightarrow\ \sqrt[6]{-bc} \in \mathbb{F}_q \ \Leftrightarrow\ \sqrt{-bc}, \sqrt[3]{b} \in \mathbb{F}_q.$$

The first $\Leftrightarrow$ takes place according to uniqueness (up to an $\mathbb{F}_q$-isogeny) of the Jacobian decomposition into simple components. The second one follows from the fact that E, E' are ordinary twists of each other. The remaining ones are evident. The restriction $\sqrt{-bc} \in \mathbb{F}_q$ is surmountable by picking $c = -b^3$. But despite this, for many curves E (including Pasta curves) $\sqrt[3]{b} \notin \mathbb{F}_q$.

5.2 The case $m = 3$

From the proof of [33, Proposition 5.2] we know that the sought points in $\mathcal{E}_3(F)$ have the form $(x, y) = (a_1 t + a_0, b_1 t + b_0)$ for some a_i, $b_i \in \mathbb{F}_q$. Substituting it into the equation of $\mathcal{E}_3$, we get the points

$$P_1 := \left(-t, \frac{\sqrt{-3}\cdot u^3}{18} \right), \qquad P_2 := \left(-t + \frac{u^2}{3}, ut - \frac{u^3}{6} \right),$$

where $u := \sqrt[6]{-108c}$, as well as $P_3 := [\omega]P_1$ and $P_4 := [\omega]P_2$. Note that $u \in \mathbb{F}_q$ if and only if $\sqrt{c}$, $\sqrt[3]{4c} \in \mathbb{F}_q$. Of course, it is sufficient to just take $c = -1/108$ or, equivalently, $u = \sqrt[6]{1}$.

The height pairing on the sublattice $\langle P_i \rangle_{i=1}^4 \subset \mathcal{E}_3(F)$ is given by the Gram matrix

$$M = \begin{pmatrix} 1 & 0 & -\dfrac{1}{2} & \dfrac{1}{2} \\ 0 & 1 & -\dfrac{1}{2} & -\dfrac{1}{2} \\ -\dfrac{1}{2} & -\dfrac{1}{2} & 1 & 0 \\ \dfrac{1}{2} & -\dfrac{1}{2} & 0 & 1 \end{pmatrix},$$

where the i-th row and column correspond to P_i. Since $\det(M) = 1/4 = \mathrm{disc}(\mathrm{D}_4^*)$, the minimal points P_i constitute a $\mathbb{Z}$-basis of $\mathcal{E}_3(F) = \mathcal{E}_3(F') \simeq \mathrm{D}_4^*$. As a result, P_1, P_2 do a $\mathbb{Z}[\omega]$-basis.

5.3 The case $m = 4$

In this section, $i := \sqrt{-1} \in \mathbb{F}_q$. Also, we need the values $v := 26\sqrt{3} - 45$ and $u := \sqrt[12]{2^6 3vc}$. The surface $\mathcal{E}_4$ over a non-closed field is discussed in the article [34]. From there we know that one can search for the desired points from $\mathcal{E}_4(F)$ in the form $(x, y) = (a_1 t + a_0, t^2 + b_1 t + b_0)$, substituting it into the equation of $\mathcal{E}_4$. In addition to $P_1 := (-\sqrt[3]{c}, t^2)$, we find a point P_2 with the coordinates

$$x_2 := ut + \frac{\sqrt{3}+3}{12}u^4, \qquad y_2 := t^2 + \frac{u^3}{2}t + \frac{\sqrt{3}+2}{8}u^6$$

and $P_3(t) := -P_2(it)$. Obviously, $u \in \mathbb{F}_q$ if and only if $\sqrt[4]{2^2 3vc}$, $\sqrt[3]{3vc} \in \mathbb{F}_q$. Inter alia, $\sqrt[3]{c} \in \mathbb{F}_q$, because $\sqrt[3]{3v} = 2\sqrt{3} - 3$. Of course, it is enough to just pick $c = 1/(2^6 3v)$ or, equivalently, $u = \sqrt[12]{1}$.

As usual, there are equally the counterparts $P_{3+j} := [\omega]P_j$, where $1 \leqslant j \leqslant 3$. The height pairing on the sublattice $\langle P_k \rangle_{k=1}^6 \subset \mathcal{E}_4(F)$ is given by the Gram matrix M such that

$$3M = \begin{pmatrix} 4 & -2 & -2 & -2 & 1 & 1 \\ -2 & 4 & 1 & 1 & -2 & 1 \\ -2 & 1 & 4 & 1 & -2 & -2 \\ -2 & 1 & 1 & 4 & -2 & -2 \\ 1 & -2 & -2 & -2 & 4 & 1 \\ 1 & 1 & -2 & -2 & 1 & 4 \end{pmatrix},$$

where the k-th row and column correspond to P_k. Since $\det(M) = 1/3 = \mathrm{disc}(\mathrm{E}_6^*)$, the minimal points P_k constitute a $\mathbb{Z}$-basis of $\mathcal{E}_4(F) = \mathcal{E}_4(F') \simeq \mathrm{E}_6^*$. As a result, P_j do a $\mathbb{Z}[\omega]$-basis.

5.4 The case $m = 5$

As well as $\mathcal{E}_4$, the surface $\mathcal{E}_5$ over a non-closed field is studied in article [34] on which we rely. First of all, possessing $\zeta := \sqrt[5]{1} \in \mathbb{F}_q$, $\zeta \neq 1$, we besides have the root $\sqrt{5} = 2\zeta^3 + 2\zeta^2 + 1$. Also, we need the values

$$v := \sqrt{\frac{3(\sqrt{5}+5)}{2}} = \zeta^2(\zeta - 1)\sqrt{-3},$$

$$\theta := 564300 + 252495\sqrt{5} + 170252{\cdot}v + 76074\sqrt{5}{\cdot}v, \qquad u := \sqrt[30]{60{\cdot}\theta c}.$$

Without further ado, one can just take $c = 1/(60{\cdot}\theta)$, that is $u = \sqrt[30]{1}$.

It turns out to be enough to confine to points of the form

$$Q_u = \Big(\frac{1}{u^2}t^2 + a_1 t + a_0,\ \frac{1}{u^3}t^3 + b_2 t^2 + b_1 t + b_0\Big).$$

As earlier, the substitution of Q_u into the equation of $\mathcal{E}_5$ gives rise to a polynomial system. After that, we find its solution

$$a_0 := -\frac{(8289\zeta^3 + 35113\zeta^2 + 43402\zeta + 21701)\omega + (26238\zeta^3 + 39650\zeta^2 + 21701\zeta - 2804)}{15}u^{10},$$

$$a_1 := -\frac{(58\zeta^3 + 246\zeta^2 + 304\zeta + 152)\omega + (184\zeta^3 + 278\zeta^2 + 152\zeta - 19)}{5}u^4,$$

$$b_0 := \frac{12a_0a_1 - a_1^3u^2 - 12a_0u^4 + 15a_1^2u^6 + 9a_1u^{10} + u^{14}}{16}u,$$

$$b_1 := \frac{12a_0 + 3a_1^2u^2 - 6a_1u^6 - u^{10}}{8u}, \qquad b_2 := \frac{3a_1 + u^4}{2u}.$$

Consider the points $P_i := Q_{\zeta^{i-1}u}$ and $P_{4+i} := [\omega]P_i = Q_{\omega\zeta^{i-1}u}$, where $1 \leqslant i \leqslant 4$. The height pairing on the sublattice $\langle P_k\rangle_{k=1}^{8} \subset \mathcal{E}_5(F)$ is given by the Gram matrix

$$M = \begin{pmatrix} 2 & -1 & 0 & 0 & -1 & 1 & 0 & 0 \\ -1 & 2 & -1 & 0 & 0 & -1 & 1 & 0 \\ 0 & -1 & 2 & -1 & 0 & 0 & -1 & 1 \\ 0 & 0 & -1 & 2 & 0 & 0 & 0 & -1 \\ -1 & 0 & 0 & 0 & 2 & -1 & 0 & 0 \\ 1 & -1 & 0 & 0 & -1 & 2 & -1 & 0 \\ 0 & 1 & -1 & 0 & 0 & -1 & 2 & -1 \\ 0 & 0 & 1 & -1 & 0 & 0 & -1 & 2 \end{pmatrix},$$

where the k-th row and column correspond to P_k. Since $\det(M) = 1 = \mathrm{disc}(\mathrm{E}_8)$, the minimal points P_k constitute a $\mathbb{Z}$-basis of $\mathcal{E}_5(F) = \mathcal{E}_5(F') \simeq \mathrm{E}_8$. As a result, P_i do a $\mathbb{Z}[\omega]$-basis.

Acknowledgements. The author expresses his gratitude to Alistair Stewart, Antonio Sanso, Geoffroy Couteau, Jeffrey Burdges, and Sergey Vasilyev for their useful comments on the role of generating "independent" points on elliptic curves in real-world vector commitment schemes.

References

1. Adleman, L., Manders, K., Miller, G.: On taking roots in finite fields. In: Symposium on Foundations of Computer Science (SFCS 1977). pp. 175–178 (1977)
2. Aranha, D.F., El Housni, Y., Guillevic, A.: A survey of elliptic curves for proof systems (2022), `https://eprint.iacr.org/2022/586`
3. Aubry, Y., Perret, M.: A Weil theorem for singular curves. In: Pellikaan, R., Perret, M., Vlăduţ, S.G. (eds.) Arithmetic, Geometry, and Coding Theory. pp. 1–7. Proceedings in Mathematics, De Gruyter, Berlin (1996)
4. Bank, E., Camacho-Navarro, C., Eisenträger, K., Morrison, T., Park, J.: Cycles in the supersingular ℓ-isogeny graph and corresponding endomorphisms. In: Balakrishnan, J., Folsom, A., Lalín, M., Manes, M. (eds.) Research Directions in Number Theory. Association for Women in Mathematics Series, vol. 19, pp. 41–66. Springer, Cham (2019)
5. Baylina, J., Bellés, M.: 4-bit window Pedersen hash on the Baby Jubjub elliptic curve (2019), `https://iden3-docs.readthedocs.io/en/latest/_downloads/4b929e0f96aef77b75bb5cfc0f832151/Pedersen-Hash.pdf`
6. Beauville, A.: Some surfaces with maximal Picard number. Journal de l'École polytechnique – Mathématiques **1**, 101–116 (2014)
7. Bisson, G., Sutherland, A.V.: Computing the endomorphism ring of an ordinary elliptic curve over a finite field. Journal of Number Theory **131**(5), 815–831 (2011)
8. Bouw, I.I., Diem, C., Scholten, J.: Ordinary elliptic curves of high rank over $\overline{\mathbb{F}_p}(x)$ with constant j-invariant. Manuscripta Mathematica **114**, 487–501 (2004)
9. Bowe, S., Grigg, J., Hopwood, D.: Recursive proof composition without a trusted setup (2019), `https://eprint.iacr.org/2019/1021`
10. Buterin, V.: Open problem: ideal vector commitment (2020), `https://ethresear.ch/t/open-problem-ideal-vector-commitment/7421`
11. Cho, G.H., Koo, N., Ha, E., Kwon, S.: New cube root algorithm based on the third order linear recurrence relations in finite fields. Designs, Codes and Cryptography **75**(3), 483–495 (2015)
12. Diem, C., Scholten, J.: Ordinary elliptic curves of high rank over $\overline{\mathbb{F}_p}(x)$ with constant j-invariant II. Journal of Number Theory **124**(1), 31–41 (2007)
13. El Mrabet, N., Joye, M.: Guide to pairing-based cryptography. Cryptography and Network Security Series, Chapman and Hall/CRC, New York (2017)
14. Electric Coin Company: The halo2 book, `https://zcash.github.io/halo2`
15. Faz-Hernandez, A., Scott, S., Sullivan, N., Wahby, R.S., Wood, C.A.: Hashing to elliptic curves (2022), `https://datatracker.ietf.org/doc/draft-irtf-cfrg-hash-to-curve`
16. Feist, D.: Inner product arguments (2021), `https://dankradfeist.de/ethereum/2021/07/27/inner-product-arguments.html`
17. Galbraith, S.D.: Mathematics of public key cryptography. Cambridge University Press, New York (2012)

18. Galbraith, S.D., Ruprai, R.S.: An improvement to the Gaudry–Schost algorithm for multidimensional discrete logarithm problems. In: Parker, M.G. (ed.) Cryptography and Coding. IMACC 2009. LNCS, vol. 5921, pp. 368–382. Springer, Berlin, Heidelberg (2009)
19. Gaudry, P., Schost, É.: A low-memory parallel version of Matsuo, Chao, and Tsujii's algorithm. In: Buell, D. (ed.) Algorithmic Number Theory. ANTS 2004. LNCS, vol. 3076, pp. 208–222. Springer, Berlin, Heidelberg (2004)
20. Hopwood, D.: The Pasta curves for Halo 2 and beyond (2020), `https://electriccoin.co/blog/the-pasta-curves-for-halo-2-and-beyond`
21. Hopwood, D.: Pluto/Eris supporting evidence (2021), `https://github.com/daira/pluto-eris`
22. Joye, M., Lapiha, O., Nguyen, K., Naccache, D.: The eleventh power residue symbol. Journal of Mathematical Cryptology **15**(1), 111–122 (2021)
23. Kate, A., Zaverucha, G.M., Goldberg, I.: Constant-size commitments to polynomials and their applications. In: Abe, M. (ed.) Advances in Cryptology – ASIACRYPT 2010. LNCS, vol. 6477, pp. 177–194. Springer, Berlin, Heidelberg (2010)
24. Kloosterman, R.: Higher Noether–Lefschetz loci of elliptic surfaces. Journal of Differential Geometry **76**(2), 293–316 (2007)
25. Koshelev, D.: Some remarks on how to hash faster onto elliptic curves (2021), `https://eprint.iacr.org/2021/1082`
26. Koshelev, D.: Magma code (2022), `https://github.com/dishport/Generation-of-independent-points-on-elliptic-curves-by-means-of-Mordell-Weil-lattices`
27. Libgober, A.: Factors of Jacobians and isotrivial elliptic surfaces. In: Maxim, L. (ed.) International Conference on Singularity Theory and Applications, 2011. Journal of Singularities, vol. 5, pp. 115–123. Worldwide Center of Mathematics, Cambridge, Massachusetts (2012)
28. Malmendier, A., Shaska, T.: From hyperelliptic to superelliptic curves. Albanian Journal of Mathematics **13**(1), 107–200 (2019)
29. Martinet, J., Schürmann, A.: Bases of minimal vectors in lattices, III. International Journal of Number Theory **8**(2), 551–567 (2012)
30. Merkle, R.C.: A digital signature based on a conventional encryption function. In: Pomerance, C. (ed.) Advances in Cryptology – CRYPTO 1987. LNCS, vol. 293, pp. 369–378. Springer, Berlin, Heidelberg (1988)
31. Pedersen, T.P.: Non-interactive and information-theoretic secure verifiable secret sharing. In: Feigenbaum, J. (ed.) Advances in Cryptology – CRYPTO 1991. LNCS, vol. 576, pp. 129–140. Springer, Berlin, Heidelberg (1992)
32. Schütt, M., Shioda, T.: Mordell–Weil lattices, A Series of Modern Surveys in Mathematics, vol. 70. Springer, Singapore (2019)
33. Shioda, T.: Mordell–Weil lattices and sphere packings. American Journal of Mathematics **113**(5), 931–948 (1991)
34. Shioda, T.: Cyclotomic analogue in the theory of algebraic equations of type E_6, E_7, E_8. In: Kim, M.H., Hsia, J.S., Kitaoka, Y., Schulze-Pillot, R. (eds.) Integral Quadratic Forms and Lattices. Contemporary Mathematics, vol. 249, pp. 87–96. American Mathematical Society, Providence, Rhode Island (1999)
35. Silverman, J.H.: Advanced topics in the arithmetic of elliptic curves, Graduate Texts in Mathematics, vol. 151. Springer, New York (1994)
36. Silverman, J.H.: The arithmetic of elliptic curves, Graduate Texts in Mathematics, vol. 106. Springer, New York (2009)
37. Silverman, J.H., Stange, K.E.: Amicable pairs and aliquot cycles for elliptic curves. Experimental Mathematics **20**(3), 329–357 (2011)

38. Tate, J.T., Shafarevich, I.R.: The rank of elliptic curves. Doklady Akademii Nauk SSSR **175**(4), 770–773 (1967)
39. Ulmer, D.: Elliptic curves with large rank over function fields. Annals of Mathematics **155**(1), 295–315 (2002)
40. Wahby, R.S., Boneh, D.: Fast and simple constant-time hashing to the BLS12-381 elliptic curve. IACR Transactions on Cryptographic Hardware and Embedded Systems **2019**(4), 154–179 (2019)
41. Zanon, G.H.M., Simplicio, M.A., Pereira, G.C.C.F., Doliskani, J., Barreto, P.S.L.M.: Faster isogeny-based compressed key agreement. In: Lange, T., Steinwandt, R. (eds.) Post-Quantum Cryptography. PQCrypto 2018. LNCS, vol. 10786, pp. 248–268. Springer, Cham (2018)