The Cokernel Pairing

Krijn Reijnders

COSIC, KU Leuven, Belgium

Abstract. We study a new pairing, beyond the Weil and Tate pairing. The Weil pairing is a non-degenerate pairing $E[m] \times E[m] \to \mu_m$, which operates on the kernel of $[m]$. Similarly, when $\mu_m \subseteq \mathbb{F}_q^*$, the Tate pairing is a non-degenerate pairing $E[m](\mathbb{F}_q) \times E(\mathbb{F}_q)/[m]E(\mathbb{F}_q) \to \mu_m$, which connects the kernel and the rational cokernel of $[m]$. We define a pairing

$$\langle\ ,\ \rangle_m : E(\mathbb{F}_q)/[m]E(\mathbb{F}_q) \times E(\mathbb{F}_q)/[m]E(\mathbb{F}_q) \to \mu_m$$

on the rational cokernels of $[m]$, to complete this triad of pairings. When $E[m] \subseteq E(\mathbb{F}_q)$, this pairing is non-degenerate, and can be computed using three Tate pairings, and two discrete logarithms in μ_m, given a basis for $E[m]$. For $m = \ell$ prime, this pairing allows us to study $E(\mathbb{F}_q)/[\ell]E(\mathbb{F}_q)$ directly and to simplify the computation for a basis of $E[\ell^k]$, and more generally the Sylow ℓ-torsion. This finds natural applications in isogeny-based cryptography when computing ℓ^k-isogenies.

Keywords: pairings · isogenies · elliptic curves

1 Introduction

Pairings are ubiquitous in modern cryptography, from their first uses in the MOV-attacks [MOV91; FR94] to their applications in protocols [HSSI99; Jou04; BLS01; BF01; BKLS02; GHS02; Jou02; BGLS03; BLS04; Gal05; Ver09], exploiting their bilinearity and other unique characteristics. Most commonly, cryptography uses the Weil pairing [Wei40; Mil04] and the Tate pairing [Tat62; Lic69], and their variations [HSV06; Bru11]. For this work, we may think of the Weil pairing of degree m, for $m \in \mathbb{N}$, as the non-degenerate bilinear map

$$e_m : E[m] \times E[m] \to \mu_m,$$

where $E[m]$ is the m-torsion subgroup of an elliptic curve E over a finite field $\mathbb{F}_q$, and μ_m is the group of m-th roots of unity in $\mathbb{F}_q$. Similarly, we may think of the (reduced) Tate pairing of degree m as the bilinear map

$$t_m : E[m](\mathbb{F}_q) \times E(\mathbb{F}_q)/[m]E(\mathbb{F}_q) \to \mu_m,$$

which is non-degenerate when $\mu_m \subseteq \mathbb{F}_q^*$.

More recently, isogeny-based cryptography often uses these pairings, as they find many natural applications in cryptanalysis [CHM+23; MS24] and core algorithm procedures [CJL+17; ZSP+18; KT18; Reij23; LWXZ24; CEMR24; DEF+25]. It is not difficult to see why: Vélu's formulas [Vél71] allow us to compute an isogeny $\phi : E \to E'$ from a description of its kernel $G = \ker \phi$. Hence, given a point $P \in E$ of order n, we can compute a (cyclic) isogeny of degree n with kernel $G = \langle P \rangle$. As the complexity of these formulas is

E-mail: crypto.krijn@gmail.com (Krijn Reijnders)

*This work was supported in part by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement ISOCRYPT - No. 101020788), by the Research Council KU Leuven grant C14/24/099 and by CyberSecurity Research Flanders with reference number VR20192203. Date of this document: 2026-09-01.

$\mathcal{O}(|G|)$, or $\mathcal{O}(\sqrt{|G|})$ using √élu [BDLS20], we improve the performance by factoring ϕ into prime-degree isogenies ϕ_i. As a result, we often want to compute isogenies of prime-power degree ℓ^k, which we may then describe by a point P of order ℓ^k, factored into k isogenies of degree ℓ. To find such points, or to describe them concisely, requires a basis of $E[\ell^k]$, where ℓ^k is such that there are no rational points of order ℓ^{k+1} on E. Equivalently, such points have no rational preimages under $[\ell]$ and we should look for such points in the set $E(\mathbb{F}_q) \setminus [\ell]E(\mathbb{F}_q)$.

This is where the Tate pairing comes in, as it allows us to identify points in $E(\mathbb{F}_q) \setminus [\ell]E(\mathbb{F}_q)$ if we have knowledge of the rational kernel $E[\ell](\mathbb{F}_q)$. It is no surprise that basis generation algorithms in the literature use the Tate pairing, whereas basis change algorithms use either the Tate or Weil pairing. Nevertheless, these pairing techniques only help us indirectly: the Tate pairing allows us to identify points in $E(\mathbb{F}_q) \setminus [\ell]E(\mathbb{F}_q)$, and the Weil pairing allows us to verify a basis for $E[\ell^k]$, but neither operates directly on the cokernel of $[\ell]$.

Contributions. This work introduces the cokernel pairing $\langle\ ,\ \rangle_m$, which operates directly on the cokernel of $[m]$:

$$\langle\ ,\ \rangle_m : E(\mathbb{F}_q)/[m]E(\mathbb{F}_q) \times E(\mathbb{F}_q)/[m]E(\mathbb{F}_q) \to \mu_m.$$

Although the definition of this pairing is elementary when given the Weil pairing, it completes the triad of pairings derived from $[m] : E \to E$ and the properties of the cokernel pairing are interesting for applications in isogeny-based cryptography: When $m = \ell$ is a (small) prime, this pairing allows us to find a basis for $E[\ell^k](\mathbb{F}_q)$ more directly: Our main theorem, Theorem 1, shows that $\langle P, Q\rangle_\ell \neq 1$ if and only if (the classes of) P and Q are a basis for $E(\mathbb{F}_q)/[\ell]E(\mathbb{F}_q)$. Among others, this implies that some multiple of P and Q are a basis for $E[\ell^k](\mathbb{F}_q)$, which is the main application of the cokernel pairing.

From a mathematical point of view, the cokernel pairing fills a symmetry gap: the Weil pairing e_m works with the kernel $E[m]$ for both arguments, the (reduced) Tate pairing t_m connects the rational kernel $E[m](\mathbb{F}_q)$ to the rational cokernel $E(\mathbb{F}_q)/[m]E(\mathbb{F}_q)$, and the cokernel pairing works with this rational cokernel for both arguments.

1. In Section 3, we provide a detailed explanation of the underlying dualities between several key objects, such as the m-torsion $E[m]$ and the Sylow m-torsion $\mathcal{S}_m(E)$, which expands on previous descriptions in the literature [Rob23; CR25; Reij25]. This deepened understanding of the duality between these objects helps us in developing practical applications of pairings.

2. In Section 4, we analyze several ways to compute the cokernel pairing, given different assumptions and situations. Current methods all rely on computation of a certain map Φ_m, and another pairing such as the Tate or Weil pairing. We also demonstrate the cokernel pairing for concrete examples. An overview of these methods and their costs is given in Table 1.

3. In Section 5, we show how the cokernel pairing is related to the Weil and Tate pairing, and how this may be used to find bases of $E[\ell^k]$ on supersingular elliptic curves.

4. In Section 6 we show how the cokernel pairing simplifies (Sylow) torsion basis generation for practical applications, as used in isogeny-based cryptography. This leads to interesting open questions in many directions for the cokernel pairing.

Acknowledgements. We thank Damien Robert for insightful discussions on the cohomological interpretation of the cokernel pairing, and deeper insights into the dualities that appear in this context. We thank the whole ISOCRYPT team from KU Leuven for their support, feedback, and useful ideas. We thank the reviewers from ANTS XVII for their helpful comments and constructive feedback.

2 Preliminaries

In this section, we introduce the necessary background on pairings, profiles, and the Sylow ℓ-torsion.

Notation. We denote the finite field of size q by $\mathbb{F}_q$, and we assume that $q = p^k$ for a prime p, the characteristic of $\mathbb{F}_q$. We denote the multiplicative group of non-zero elements of $\mathbb{F}_q$ by $\mathbb{F}_q^*$, which is a cyclic group of order $q-1$. We denote the algebraic closure of $\mathbb{F}_q$ by $\overline{\mathbb{F}_q}$ and the m-th roots of unity in $\overline{\mathbb{F}_q}$ by μ_m.

For elliptic curves over $\mathbb{F}_q$, we denote the neutral element of E by $\mathbf{0}_E$, and π always denotes the Frobenius endomorphism $(x, y) \mapsto (x^q, y^q)$ with respect to $\mathbb{F}_q$. We use the word *rational* to refer to points, curves, and isogenies defined over the base field $\mathbb{F}_q$, for example, E is rational when it is defined over $\mathbb{F}_q$, and $P \in E$ is rational when $P \in E(\mathbb{F}_q)$.

For an integer $m \in \mathbb{N}$, we denote the Weil pairing by e_m, the Tate pairing by T_m, and the reduced Tate pairing by t_m. We denote the Sylow m-torsion by $\mathcal{S}_m(E)$, which we describe in more detail in Section 2.4. We use the word *cofactor* with respect to some integer m to refer to the smallest integer h such that $[h]P \in \mathcal{S}_m(E)$ for all $P \in E(\mathbb{F}_q)$. When $|E(\mathbb{F}_q)| = N$ and $m = \ell$ prime, we may use[1] $h = N/\ell^{v_\ell(N)}$, where $v_\ell(N)$ is the ℓ-adic valuation of N. Simply put, for $N = \ell^k \cdot h$ with $\gcd(\ell, h) = 1$, we can use the cofactor h. We somewhat misuse the word *basis*, as we sometimes refer to a basis (P, Q) generating a subgroup isomorphic to $\mathbb{Z}_{\ell^f} \times \mathbb{Z}_{\ell^g}$ with $g < f$, even though both P and Q may be of order ℓ^f, i.e., have linear relations.

2.1 Pairings

In general, a pairing $A \times B \to C$ is a bilinear map between abelian groups A, B and C. In this work we are interested in abelian groups A and B that are subgroups or quotient groups of an elliptic curve E over a finite field $\mathbb{F}_q$, and similarly C is a subgroup or quotient group of $\mathbb{F}_q^*$. Central to this work are the subgroups and quotient groups derived from the multiplication-by-m endomorphisms $[m] : P \mapsto [m]P$, namely, the kernel

$$E[m] = \{P \in E(\overline{\mathbb{F}_q}) \mid [m]P = \mathbf{0}_E\},$$

and, viewing $[m]$ as a map $E(\mathbb{F}_q) \to E(\mathbb{F}_q)$, the cokernel

$$E(\mathbb{F}_q)/[m]E(\mathbb{F}_q) = \{P + [m]E(\mathbb{F}_q) : P \in E(\mathbb{F}_q)\}.$$

For the rest of this section, we assume that m is a positive integer coprime to the characteristic of $\mathbb{F}_q$, and we let E be an elliptic curve over $\mathbb{F}_q$. Some results are specialized to the case where $m = \ell$ is a small odd prime, which is our main case of interest.

2.2 The Weil Pairing

The Weil pairing [Wei40] of degree m is a non-degenerate bilinear pairing

$$e_m : E[m] \times E[m] \to \mu_m. \tag{1}$$

More generally, for abelian varieties A, we can derive the Weil pairing as a pairing

$$A[m] \times \widehat{A}[m] \to \mu_m,$$

where $\widehat{A}$ is the dual abelian variety of A [Sil10]. As elliptic curves are naturally isomorphic to their duals, we get a canonical principal polarization $E \to \widehat{E}$, which allows us to recover

[1] This may not be the smallest integer that clears all but the ℓ-torsion, but works for our needs.

the Weil pairing on E itself. Similarly, Jacobian varieties of non-singular curves with a rational point come equipped with such a principal polarization. In such cases, we may simply write e_m, with no need to specify the polarization $\lambda : A \to \widehat{A}$ against which we define the Weil pairing. More details can be found in the book by Edixhoven, Van der Geer, and Moonen [EVM12, Ch. 11].

The Generalized Weil Pairing. This generalized notion of the Weil pairing on abelian varieties allows us to define the Weil-Cartier pairing [EVM12; Rob23] with respect to an m-isogeny $f : A \to B$. This results in a pairing $\ker f \times \ker \widehat{f} \to \mu_m$, where $\widehat{f} : \widehat{B} \to \widehat{A}$ is the dual isogeny [Sil09, Rem. 8.4, Ex. 3.15].

2.3 The Tate Pairing

Assuming $\mu_m \subseteq \mathbb{F}_q^*$, the Tate-Lichtenbaum pairing [Tat62; Lic69] of degree m, hereafter simply the Tate pairing, is a non-degenerate bilinear pairing

$$T_m : E[m](\mathbb{F}_q) \times E(\mathbb{F}_q)/[m]E(\mathbb{F}_q) \to \mathbb{F}_q^*/\mathbb{F}_q^{*,m}. \tag{2}$$

This pairing was introduced in a cryptographic context by Frey and Rück [FR94]. For cryptographic purposes, working with equivalence classes in $\mathbb{F}_q^*$ is inconvenient. Hence, we may 'reduce' the Tate pairing by applying a final exponentiation by $\frac{q-1}{m}$, which maps to μ_m. This gives the reduced Tate pairing

$$t_m : E[m](\mathbb{F}_q) \times E(\mathbb{F}_q)/[m]E(\mathbb{F}_q) \to \mu_m. \tag{3}$$

The groups $\mathbb{F}_q^*/\mathbb{F}_q^{*,m}$ and μ_m are naturally dual to each other: when we view exponentiation by m as a map $\mathbb{F}_q^* \to \mathbb{F}_q^*$, the final reduction induces an isomorphism

$$\mathbb{F}_q^*/\mathbb{F}_q^{*,m} \xrightarrow{\sim} \mu_m, \quad z \mapsto z^{\frac{q-1}{m}}, \tag{4}$$

between the cokernel $\mathbb{F}_q^*/\mathbb{F}_q^{*,m}$ and the kernel μ_m We stress that, contrary to the Weil pairing, the Tate pairing crucially relies on the field of definition that we are working over, as is clear from its definition.

Link to the Weil Pairing. When $E[m] \subseteq E(\mathbb{F}_q)$, an alternative definition of the reduced Tate pairing can be obtained using a preimage R of Q, e.g., a point such that $[m]R = Q$ [CHM+23]. We get

$$t_m(P, Q) = e_m(P, \pi(R) - R). \tag{5}$$

Written in this way, we clearly see the arithmetic nature of the Tate pairing, as we take Frobenius π with respect to a field $\mathbb{F}_q$. If we compute the Tate pairing for the same points over a different field, we may get a different result. In particular, when we extend to the field of definition of R, we get $\pi(R) = R$, hence $t_m(P, Q) = 1$.

The Generalized Tate Pairing. Similar to the Weil pairing, we may also generalize the Tate pairing with respect to an m-isogeny $f : A \to B$ over $\mathbb{F}_q$. This results in the Tate-Cartier pairing, also known as the generalized f-Tate pairing or the Tate pairing associated to f. Bruin [Bru11] shows that this pairing between the rational kernel and rational cokernel of f is non-degenerate when $\ker f$ is annihilated by $[q-1]$, and gives a description as $\ker \widehat{f}\ (\mathbb{F}_q) \times \operatorname{coker} f\ (\mathbb{F}_q) \to \mathbb{F}_q^*$.

Computation of the Pairings. Miller's algorithm [Mil04; FR94] computes both the Weil and Tate pairing efficiently, which generalize well to Jacobians. In recent work, Robert [Rob24] introduces *cubical arithmetic* to compute pairings on abelian varieties and Kummer varieties, generalizing previous work [LR16; Sta08; Sta11].

2.4 The Sylow ℓ-Torsion

The Sylow m-torsion $\mathcal{S}_m(E)$ is the subgroup of $E(\mathbb{F}_q)$ that contains all points whose orders divide some power of m, sometimes also denoted $E[m^\infty](\mathbb{F}_q)$. We mostly focus on the case where $m = \ell$ is a prime, in which case all points in $\mathcal{S}_\ell(E)$ are in the kernel of $[\ell^k]$ for some large enough k. We define the Sylow torsion generally for abelian varieties.

Definition 1. Let A be a g-dimensional abelian variety over $\mathbb{F}_q$ and ℓ a prime, coprime to p. The *Sylow ℓ-torsion* $\mathcal{S}_\ell(A)$ over $\mathbb{F}_q$ is the subgroup

$$\mathcal{S}_\ell(A) := A[\ell^\infty](\mathbb{F}_q) \cong \mathbb{Z}/\ell^{f_1}\mathbb{Z} \times \mathbb{Z}/\ell^{f_2}\mathbb{Z} \times \ldots \times \mathbb{Z}/\ell^{f_d}\mathbb{Z}$$

with $f_1, \ldots, f_d \in \mathbb{N}$ such that $f_1 \geq f_2 \geq \ldots f_d > 0$, and $d \leq 2g$ the *rank* of $\mathcal{S}_\ell(A)$. We refer to a set of d points as a *basis* of $\mathcal{S}_\ell(A)$ whenever their linear combinations generate all elements of $\mathcal{S}_\ell(A)$.

We say $\mathcal{S}_\ell(A)$ is *symmetric* when the rank $d = 2g$ and all f_i are equal, e.g. $\mathcal{S}_\ell(A) = A[\ell^{f_1}]$. It is clear that $A[\ell](\mathbb{F}_q) \subseteq \mathcal{S}_\ell(A)$, and so d is equal to the rank of $A[\ell](\mathbb{F}_q)$. The exponent equals ℓ^{f_1}. Furthermore, an isogeny $f : A \to B$ gives us a homomorphism $\mathcal{S}_\ell(A) \to \mathcal{S}_\ell(B)$, which is an isomorphism whenever $\deg f$ is coprime to ℓ.

In the case of elliptic curves the rank d of $\mathcal{S}_\ell(E)$ is either 0, 1, or 2. In this paper we assume that $E[\ell]$ is rational, and so $d = 2$ in essentially all cases. Thus, on elliptic curves with $E[\ell]$ rational, a basis for the Sylow ℓ-torsion is simply a pair of points (P, Q) that generate all rational points of order ℓ^k for $0 \leq k \leq f_1$.

The Sylow ℓ-torsion is closely related to the rational cokernel $E(\mathbb{F}_q)/[\ell]E(\mathbb{F}_q)$ and the kernel $E[\ell]$. First, classes in the rational cokernel are in one-to-one correspondence with classes in $\mathcal{S}_\ell(E)/[\ell]\mathcal{S}_\ell(E)$ which comes down to 'ignoring' all other torsion, which we can formalize as multiplication by h, where h is the cofactor with respect to ℓ. Second, as $\mathcal{S}_\ell(E)/[\ell]\mathcal{S}_\ell(E)$ is dual to $E[\ell]$, we may associate a point $P_\ell \in E[\ell]$ to any class $P \in E(\mathbb{F}_q)/[\ell]E(\mathbb{F}_q)$.

Isogeny-based cryptography often works with supersingular curves $E/\mathbb{F}_{p^2}$ of order $(p+1)^2$, where the torsion structure is isomorphic to $\mathbb{Z}/(p+1)\mathbb{Z} \times \mathbb{Z}/(p+1)\mathbb{Z}$. In such cases, the Sylow ℓ-torsion is simply $\mathcal{S}_\ell(E) = E[\ell^f]$, where f is the largest integer such that $\ell^f \mid p+1$. In particular $\mathcal{S}_\ell(E)$ is symmetric, as it is isomorphic to $\mathbb{Z}/(\ell^f)\mathbb{Z} \times \mathbb{Z}/(\ell^f)\mathbb{Z}$.

2.5 The Tate Profile

When the Tate pairing of degree m is non-degenerate, we may study the cokernel $E(\mathbb{F}_q)/[m]E(\mathbb{F}_q)$ more precisely using a rational basis (P, Q) of $E[m](\mathbb{F}_q)$, and the map

$$t_{[m]} : E(\mathbb{F}_q) \to \mu_m^r, \quad R \mapsto (t_m(P, R), t_m(Q, R)).$$

Following [Rob23; CR25; Reij25], we call $t_{[m]}$ the *Tate profile* of R with respect to the basis $(P_1, \ldots, P_r)$. The profile is trivial if and only if $R \in [m]E(\mathbb{F}_q)$, and whenever the profiles of a set (R, R') generate μ_m^r, the set generates the cokernel $E(\mathbb{F}_q)/[m]E(\mathbb{F}_q)$, which holds more generally for principally polarized abelian varieties too [Reij25, Thm. 1]. Under the mild additional assumption $E[m] \subseteq E(\mathbb{F}_q)$, we get $E[m] \xrightarrow{\sim} E(\mathbb{F}_q)/[m]E(\mathbb{F}_q) \xrightarrow{\sim} \mu_m^2$, and we may view the Tate profile as a coordinate system on $E(\mathbb{F}_q)/[m]E(\mathbb{F}_q)$ through the map to μ_m^2. The rational kernel and cokernel are dual, hence isomorphic, but the isomorphism given by the profile crucially depends on a choice of basis.

3 The Cokernel Pairing

Before we introduce the cokernel pairing in Section 3.2, we explore the connection between the kernel $E[m]$ and the cokernel $E(\mathbb{F}_q)/[m]E(\mathbb{F}_q)$ in more detail. We are mostly interested in the case where $m = \ell$ is a small prime and $E[\ell]$ is rational, as studying the Sylow ℓ-torsion in this case is most interesting for our applications. Hence, we will focus on $m = \ell$ prime and curves E with $E[\ell] \subseteq E(\mathbb{F}_q)$, and may sometimes switch to the more general case $m \in \mathbb{N}$ when this does not create extra difficulties or subtleties. In the whole work, m and ℓ are coprime to p.

3.1 Connecting Kernel and Cokernel

The reduced Tate pairing, when defined as in Equation (5), relies on the preimages R of points P under $[\ell]$. We first analyze this map $\Phi_m : P \mapsto \pi(R) - R$, before we analyze the cokernel pairing, to show that the map behaves well and naturally connects the kernel and cokernel of $[\ell]$. The map Φ_m arises naturally from the long exact sequence defining the Tate pairing using Galois cohomology, see [Sil10], but we use a more concrete description.

Lemma 1. *Let E be an elliptic curve over $\mathbb{F}_q$. For $m \in \mathbb{N}$, assume $E[m] \subseteq E(\mathbb{F}_q)$, and let $P \in E(\mathbb{F}_q)$ with $R \in E(\overline{\mathbb{F}_q})$ such that $[m]R = P$. Then the map*

$$E(\mathbb{F}_q) \to E[m], \qquad P \mapsto \pi(R) - R$$

is a well-defined homomorphism with kernel $[m]E(\mathbb{F}_q)$. The map induces an isomorphism

$$\Phi_m : E(\mathbb{F}_q)/[m]E(\mathbb{F}_q) \xrightarrow{\sim} E[m]. \tag{6}$$

Proof. We have

$$[m](\pi(R) - R) = \pi([m]R) - [m]R = \pi(P) - P = \mathbf{0}_E,$$

and so we get $\Phi_m(P) \in E[m]$. Let $R, R' \in E(\overline{\mathbb{F}_q})$ such that $[m]R = P$ and $[m]R' = P$. Then $R' = R + T$ for some $T \in E[m]$. As $E[m] \subseteq E(\mathbb{F}_q)$, we have $\pi(T) = T$, and so

$$\pi(R') - R' = \pi(R) - R + \pi(T) - T = \pi(R) - R.$$

Thus, the map is well-defined. We get that $\Phi_m(P) = \mathbf{0}_E$ only if $\pi(R) = R$, which implies $R \in E(\mathbb{F}_q)$ and so $P = [m]R \in [m]E(\mathbb{F}_q)$, so $\ker \Phi_m = [m]E(\mathbb{F}_q)$. As both kernel and cokernel have the same size, we get the isomorphism $\Phi_m : E(\mathbb{F}_q)/[m]E(\mathbb{F}_q) \xrightarrow{\sim} E[m]$. □

It is not too difficult to bound the extension degree in which such a preimage R lies, as the following result shows.

Corollary 1. *Let $P \in E(\mathbb{F}_q)$, and let $R \in E(\overline{\mathbb{F}_q})$ such that $[m]R = P$. When $\Phi_m(P) = \mathbf{0}_E$, then $R \in E(\mathbb{F}_q)$. When $\Phi_m(P) \neq \mathbf{0}_E$, then $R \in E(\mathbb{F}_{q^d})$, where d is the order of $\Phi_m(P)$.*

Proof. When $\Phi_m(P) = \mathbf{0}_E$, this is immediate from Lemma 1. When $\Phi_m(P) \neq \mathbf{0}_E$, given that $\pi(R) - R$ is rational, we get $\pi\left(\pi(R) - R\right) = \pi(R) - R$ and so

$$\pi^2(R) - R = \pi\left(\pi(R) - R\right) + \pi(R) - R = [2](\pi(R) - R),$$

and by induction this gives $\pi^k(R) - R = [k](\pi(R) - R)$. Thus, for $k = d$ where d is the order of $\Phi_m(P) = \pi(R) - R$, we find

$$\mathbf{0}_E = [d]\Phi_m(P) = \pi^d(R) - R,$$

that is, $\pi^d(R) = R$, so $R \in E(\mathbb{F}_{q^d})$. □

Thus, when $m = \ell$ a small prime, this implies $R \in E(\mathbb{F}_{q^\ell})$ whenever $\Phi_\ell(P) \neq 0$. Assuming rational m-torsion, the endomorphism $[m]$ divides $\pi - 1$, and so we may also write Φ_m as the endomorphism $\frac{\pi-1}{m}$. As an isomorphism between the rational cokernel and the kernel, Φ_m identifies the 'position' of points in the Sylow m-torsion with respect to their position in $E[m]$. This is equivalent to the positioning given by the Tate profile $t_{[m]}$ with respect to some basis T_1, T_2 of the kernel, encoded as a value in μ_m^2. Furthermore, Φ_m is the curve equivalent to the 'reduction' map $\mathbb{F}_q^*/\mathbb{F}_q^{*,m} \xrightarrow{\sim} \mu_m$ from Equation (4).

3.2 The Cokernel Pairing

The definition of the reduced Tate pairing via Equation (5) motivates us to look at a pairing on the rational cokernels using the map Φ_m for a straightforward definition.

Definition 2. Let E be an elliptic curve over $\mathbb{F}_q$ and $m \in \mathbb{N}$ such that $E[m] \subseteq E(\mathbb{F}_q)$. The *cokernel pairing of degree m* is a pairing

$$\langle\ ,\ \rangle_m : E(\mathbb{F}_q)/[m]E(\mathbb{F}_q) \times E(\mathbb{F}_q)/[m]E(\mathbb{F}_q) \to \mu_m.$$

Given $P, Q \in E(\mathbb{F}_q)$ as representatives of their classes in $E(\mathbb{F}_q)/[m]E(\mathbb{F}_q)$, we define the cokernel pairing in terms of the Weil pairing e_m and the map Φ_m, by

$$\langle P, Q\rangle_m := e_m\left(\Phi_m(P), \Phi_m(Q)\right).$$

The cokernel pairing is naturally connected to the Sylow m-torsion in the same way that the Weil pairing is naturally connected to $E[m]$, with the Tate pairing providing the bridge between these two. Note that the cokernel pairing is formulated independently of a basis, and therefore canonical. This is interesting with regard to later results and applications, such as Theorem 1, as current methods require a (non-canonical) basis of $E[m]$ to achieve similar results.

General properties. We first prove general properties of the cokernel pairing, before we dive deeper into its connection with Sylow torsion.

Proposition 1. *Assuming the m-torsion is rational, the cokernel pairing of degree m is anti-symmetric, bilinear, non-degenerate, and compatible with isogenies $\phi : E \to E'$.*

- *For all $P, Q \in E(\mathbb{F}_q)$, we have $\langle P, Q\rangle_m = \langle Q, P\rangle_m^{-1}$.*
- *For all $P_1, P_2, Q \in E(\mathbb{F}_q)$, we have $\langle P_1 + P_2, Q\rangle_m = \langle P_1, Q\rangle_m \cdot \langle P_2, Q\rangle_m$, and similarly $\langle P, Q_1 + Q_2\rangle_m = \langle P, Q_1\rangle_m \cdot \langle P, Q_2\rangle_m$.*
- *For a given $P \in E(\mathbb{F}_q)$, if $\langle P, Q\rangle_m = 1$ for all $Q \in E(\mathbb{F}_q)$, then $P \in [m]E(\mathbb{F}_q)$, and mutatis mutandis for Q.*
- *For a separable isogeny $\phi : E \to E'$ over $\mathbb{F}_q$ of degree coprime to m and $P, Q \in E(\mathbb{F}_q)$, we have $\langle \phi(P), \phi(Q)\rangle_m = \langle P, Q\rangle_m^{\deg \phi}$.*

Proof. These properties are easily shown using the properties of the Weil pairing, and the isomorphism Φ_m. As the Weil pairing is alternating, we get the same for the cokernel pairing by

$$\langle P, Q\rangle_m = e_m\left(\Phi_m(P), \Phi_m(Q)\right) = e_m\left(\Phi_m(Q), \Phi_m(P)\right)^{-1} = \langle Q, P\rangle_m^{-1}.$$

Similarly, bilinearity follows by bilinearity of the Weil pairing and the fact that Φ_m is a homomorphism. For non-degeneracy, whenever $\langle P, Q\rangle_m = 1$ for all $Q \in E(\mathbb{F}_q)$, we must have by non-degeneracy of the Weil pairing that $\Phi_m(P) = \mathbf{0}_E$, which implies $P \in [m]E(\mathbb{F}_q)$ by Lemma 1. By the first property, the same holds for Q. Furthermore, when $P \in [m]E(\mathbb{F}_q)$

then $R \in E(\mathbb{F}_q)$, and so $\Phi_m(P) = \mathbf{0}_E$ which implies $\langle P, Q\rangle_m = 1$, and similarly for Q. Lastly, compatibility with isogenies follows from the compatibility of the Weil pairing with isogenies, together with the fact that $\Phi_m \circ \phi = \phi \circ \Phi_m$, as $\phi(R)$ is a pre-image under $[m]$ of $\phi(P)$. □

The last property is also interesting when $\deg \phi$ is not coprime to m: the cokernel pairing becomes trivial for any $P, Q \in E(\mathbb{F}_q)$ when mapped through ϕ, which we may intuitively understand as either P or Q getting mapped to $[m]E'(\mathbb{F}_q)$, or that the m-adic torsion of P and Q 'overlaps'. To make this more precise, we discuss the connection to the Sylow ℓ-torsion.

3.3 Connection to the Sylow ℓ-torsion

Recall that we may represent $\mathcal{S}_\ell(E)/[m]\mathcal{S}_\ell(E)$ by classes from $E(\mathbb{F}_q)/[m]E(\mathbb{F}_q)$, where $\mathcal{S}_m(E)$ is the Sylow m-torsion of E, and that $\mathcal{S}_m(E)/[m]\mathcal{S}_m(E)$ is dual to the m-torsion $E[m]$, with the map Φ_m giving us the isomorphism $\mathcal{S}_m(E)/[m]\mathcal{S}_m(E) \xrightarrow{\sim} E[m]$.

When $m = \ell$ a small prime, the Sylow m-torsion is interesting for isogeny-based cryptography, as we usually use rational points of order ℓ^k to compute an ℓ^k-isogeny efficiently, and the Sylow ℓ-torsion describes exactly those elements. Note that a point $P \in E(\mathbb{F}_q) \setminus [\ell]E(\mathbb{F}_q)$ is not per se of order ℓ^k, though one may use the Tate pairing to identify such points [Rob23; CR25]. For our purposes, we mainly need that some multiple of representants of generators $E(\mathbb{F}_q)/[\ell]E(\mathbb{F}_q)$ generate $\mathcal{S}_\ell(E)$ too [Reij25]. The following theorem, for general m, is the crucial connection of $\mathcal{S}_m(E)$ to the cokernel pairing.

Theorem 1. *Let m be coprime to p, and let $P, Q \in E(\mathbb{F}_q)$. Let h be the cofactor of E with respect to m. Then, $[h]P$ and $[h]Q$ generate $\mathcal{S}_m(E)$ if and only if $\langle P, Q\rangle_m = \zeta_m$ is a primitive m-th root of unity.*

Proof. The proof is a combination of two insights: that Φ_m gives us an isomorphism between $\mathcal{S}_m(E)/[m]\mathcal{S}_m(E)$ and $E[m]$, and that the Weil pairing $e_m(P', Q')$ is a primitive m-th root of unity if and only if P', Q' are a basis of $E[m]$.

If we have a basis P, Q of $\mathcal{S}_m(E)$, then the classes of P and Q generate $\mathcal{S}_m(E)/[m]\mathcal{S}_m(E)$. As Φ_m is an isomorphism, this implies that $\Phi_m(P)$ and $\Phi_m(Q)$ are a basis for $E[m]$ and thus the Weil pairing $\zeta = e_m(\Phi_m(P), \Phi_m(Q))$ is a primitive m-th root of unity, which by definition is the cokernel pairing $\langle P, Q\rangle_m$.

Similarly, if $\langle P, Q\rangle_m$ is a primitive m-th root of unity then $\Phi_m(P)$ and $\Phi_m(Q)$ are a basis of $E[m]$. Hence, the classes $[[h]P]$ and $[[h]Q]$ generate $\mathcal{S}_m(E)/[m]\mathcal{S}_m(E)$, which implies $[h]P$ and $[h]Q$ generate $\mathcal{S}_m(E)$. □

When $m = \ell$ is a (small) prime, which is the most interesting case for Sylow torsion, we can simplify the result, as the pairing is a primitive ℓ-th root as soon as it is non-trivial.

Corollary 2. *Let ℓ be coprime to p, and let $P, Q \in E(\mathbb{F}_q)$. Let h be the cofactor of E with respect to ℓ. Then, $[h]P$ and $[h]Q$ generate $\mathcal{S}_\ell(E)$ if and only if $\langle P, Q\rangle_\ell \neq 1$.*

It is this version that we mainly use when applying the cokernel pairing in later applications. Thus, the cokernel pairing plays a dual role to the Weil pairing: whereas a non-trivial Weil pairing implies a basis for $E[\ell]$, a non-trivial cokernel pairing implies a basis for $\mathcal{S}_\ell(E)$. More generally, for composite m, we want the cokernel pairing value to be a primitive m-th root of unity, similar to how the Weil pairing indicates a basis for $E[m]$ when the pairing is of order m. The isomorphism Φ_ℓ connects the dual objects $E[\ell]$ and $\mathcal{S}_\ell(E)/[\ell]\mathcal{S}_\ell(E)$, and the (reduced) Tate pairing allows us to transfer knowledge from one to the other. We have visualised this in Figure 1.

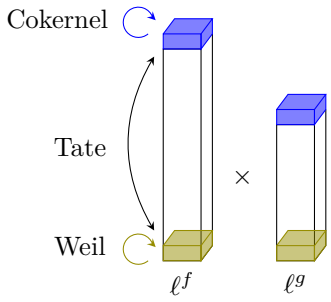

Figure 1: A visualisation of the Sylow ℓ-torsion $\mathcal{S}_\ell(E)$, indicating in blue the cokernel $\mathcal{S}_\ell(E)/[\ell]\mathcal{S}_\ell(E)$ where the cokernel pairing operates and in olive the kernel $E[\ell]$ where the Weil pairing operates. The Tate pairing transforms information from the kernel to the cokernel and back, as an elevator from the first to the top floor.

Remark 1. The cokernel pairing requires the assumption that $E(\mathbb{F}_q)$ has rational ℓ-torsion, which restricts its usage to specific situations: When $E[\ell](\mathbb{F}_q)$ has rank 1, then $\mathcal{S}_(\ell)(E)$ has rank 1 too, and so there is no need for a cokernel pairing, as we may use the Tate pairing to find a generator of $\mathcal{S}_\ell(E)$ in this case. In higher dimensions, partially-rational cokernels are more interesting to study, in particular through the isomorphism with the partially-rational kernel.

3.4 The Generalized Cokernel Pairing

Similar to the generalized Weil and Tate pairing described in Sections 2.2 and 2.3, there seems to be no obstruction to generalizing the cokernel pairing to any separable isogeny $f : A \to B$ between abelian varieties with rational kernels $\ker f$ and $\ker \hat{f}$. Let m be the exponent of $\ker f$. By restricting the map Φ_m, which we denote Φ_f, we get an isomorphism $(\operatorname{coker} f)\ (\mathbb{F}_q) \xrightarrow{\sim} (\ker f)\ (\mathbb{F}_q)$, and similarly $(\operatorname{coker} \widehat{f})\ (\mathbb{F}_q) \xrightarrow{\sim} (\ker \widehat{f})\ (\mathbb{F}_q)$. Thus, we may define the following generalization.

Definition 3. Let $f : A \to B$ be a separable isogeny between abelian varieties A and B over a finite field $\mathbb{F}_q$. Let m be the exponent of $(\ker f)\ (\mathbb{F}_q)$. Then, the *generalized f-cokernel pairing* is a pairing

$$\langle\ ,\ \rangle_f : (\operatorname{coker} f)\ (\mathbb{F}_q) \times (\operatorname{coker} \widehat{f})\ (\mathbb{F}_q) \to \mu_m.$$

Given $P \in B(\mathbb{F}_q)$ and $Q \in \hat{A}(\mathbb{F}_q)$, we define

$$\langle P, Q\rangle_f := e_f\left(\Phi_f(P), \Phi_{\widehat{f}}(Q)\right),$$

where e_f is the generalized Weil pairing $\ker f \times \ker \widehat{f} \to \mu_m$ with respect to f.

4 Computation of the Cokernel Pairing

We describe several methods to compute the cokernel pairing. All of these methods rely on the computation of Φ_m and an application of a pairing. First, in Section 4.1, we describe three methods to compute Φ_m, assuming either no knowledge of $E[m]$ or the knowledge of a basis of $E[m]$. Then, in Section 4.2, we analyze two methods to compute the cokernel pairing, using either the Weil or Tate pairing, assuming we can compute Φ_m. We summarize the costs and limitations of each approach in Table 1. Lastly, we give concrete examples in Section 4.3.

4.1 Computing Φ_m

We describe three methods to compute Φ_m: First, naively, by computing the point R such that $[m]R = P$ in $E(\mathbb{F}_{q^m})$. Second, by working over the fiber $[m]^{-1}(P)$. Third, assuming knowledge of a basis of $E[m]$, using a generic endomorphism computation algorithm.

4.1.1 Naive computation

The most straightforward computation of Φ_m computes the point R with $[m]R = P$ directly in $E(\mathbb{F}_{q^m})$. Writing $\psi_m(x)$ for the m-th division polynomial, and ϕ_m and ω_m as in [Sil09, III, Ex. 3.7], we may write the map $[m] : E \to E$ as

$$[m](x, y) = \left(\frac{\phi_m(x)}{\psi_m^2(x)}, \frac{\omega_m(x, y)}{\psi_m^3(x)}\right). \tag{7}$$

Thus, to find a preimage R of P, we compute a root of

$$G(x) = \phi_m(x) - x_P \cdot \psi_m^2(x), \tag{8}$$

with $\deg G = m^2$ to find the x-coordinate x_R of R. We find an associated y-coordinate y_R by solving the curve equation. Given R, we may then compute $\Phi_m(P) = \pi(R) - R$.

Cost. The dominant cost in this approach is the computation of the root x_R in the extension field $\mathbb{F}_{q^m}$. In the usual procedure this requires a subroutine computing $x^{q^m} \bmod G(x)$ at a cost of $\mathcal{O}(M(m^2) \cdot m \log q)$, where $M(n)$ denotes the cost of multiplying two polynomials of degree n over $\mathbb{F}_q$. With schoolbook multiplication, we get $M(m^2) = m^4$, with Karatsuba $M(m^2) = m^{2 \cdot \log_2(3)} \approx m^{3.17}$.

4.1.2 Computing symbolically over the fiber

By Lemma 1, the value of $\Phi_m(P)$ does not depend on the choice of $R \in [m]^{-1}(P)$. Therefore, instead of computing an actual root x_R as in the naive computation, we may compute $\pi(R) - R$ symbolically, on the generic point of the fiber $[m]^{-1}(P)$, by working in the algebra

$$A := \mathbb{F}_q[X]/(G(X)),$$

with G as in Equation (8). This allows us to symbolically compute 'over the fiber' with the x-coordinate of R as X. Thus, to compute the x-coordinate $x_{\Phi_m(P)}$ of $\Phi_m(P)$, we first compute $w = f(X)^{\frac{q-1}{2}}$, where $E : y^2 = f(x)$, to get $\pi(R)$ as (X^q, y^q), as $y^q = y \cdot y^{q-1} = y \cdot w$. We then obtain the x-coordinate of $\pi(R) - R$ using the standard chord-tangent construction with $\pi(R)$ and $-R$: For a short Weierstrass curve $y^2 = x^3 + ax + b$, addition of (x_1, y_1) and (x_2, y_2) gives $x_3 = \lambda^2 - x_1 - x_2$ with $\lambda = \frac{y_2 - y_1}{x_2 - x_1}$. For $\pi(R) = (X^q, y \cdot w)$ and $-R = (X, -y)$, this gives $\lambda = y \cdot \mu$ where $\mu = \frac{w+1}{X^q - X}$. Note that $y^2 = f(x)$, so that the final x-coordinate is simply $f(X) \cdot \mu^2 - X^q - X \in \mathbb{F}_q \subseteq A$, now a constant.

Remark 2. The above computation fails when $R \in E(\mathbb{F}_q)$ as $X^q - X$ vanishes and so μ is undefined. In this situation, we may recover the rational x_R as a root of $G(x)$ similar to the previous section.

The symbolic computation is summarized in Algorithm 1 for a short Weierstrass curve. On a Montgomery curve, where $x_3 = \lambda^2 - x_1 - x_2 - A$, one should adjust the final result by subtracting A.

Remark 3. Algorithm 1 returns only the x-coordinate of $\Phi_m(P)$. For the applications in this work, this is enough information, as we can compute pairings with only the x-coordinate using cubical arithmetic [Rob24; PRR+25]. If the y-coordinate is required, one can extend the above approach to include y, but this requires the y-argument of Equation (7).

Algorithm 1 Symbolic computation of Φ_m

Input: A point $P \in E(\mathbb{F}_q)$, where $E : y^2 = f(x)$ with $f(x) = x^3 + ax + b$.
Output: The x-coordinate of the point $\Phi_m(P) \in E[m]$
1: $G(x) \leftarrow \phi_m(x) - x_P \cdot \psi_m^2(x)$
2: Define A as $\mathbb{F}_q[X]/(G(X))$
3: $\delta \leftarrow X^q - X$
4: $w \leftarrow f(X)^{\frac{q-1}{2}}$
5: $\mu \leftarrow \frac{w+1}{\delta}$
6: **return** $f(X) \cdot \mu^2 - X^q - X$

Cost. The dominant cost of Algorithm 1 are the two exponentiations X^q and $f(X)^{\frac{q-1}{2}}$, giving an overall cost of $\mathcal{O}(M(m^2)\log q)$, where $M(n)$ denotes the cost of multiplying two polynomials of degree n over $\mathbb{F}_q$. This saves a factor m compared to the naive approach, essentially by decreasing the degree of the required exponentiations in A from q^m to q.

4.1.3 When a basis of $E[m]$ is given

We may write Φ_m as the endomorphism $\frac{\pi-1}{m} \in \mathrm{End}(E)$. In [DEF+25, Appendix D], the authors describe an algorithm that computes endomorphisms of the form $\frac{a+b\alpha}{m}$ over $\mathbb{F}_q$, assuming a basis for the m-torsion. For completeness, we repeat their algorithm for our situation $a = -1$, $b = 1$, $\alpha = \pi$ in Algorithm 2. Let T_1, T_2 be a basis for $E[m]$, and set $\zeta_0 = e_m(T_1, T_2)$ as a fixed m-th root of unity.

Algorithm 2 Rational computation of Φ_m

Input: A point $P \in E(\mathbb{F}_q)$, a basis T_1, T_2 of $E[m]$, and $\zeta_0 = e_m(T_1, T_2) \in \mu_m$.
Output: The point $\Phi_m(P) \in E[m]$
1: $(\zeta_1, \zeta_2) \leftarrow (t_m(T_1, P), t_m(T_2, -P))$
2: $(r, s) \leftarrow (\log_{\zeta_0}(\zeta_2), \log_{\zeta_0}(\zeta_1))$
3: **return** $rT_1 + sT_2$

Cost. This algorithm costs two Tate pairings of degree m, which takes $\mathcal{O}(\log q + \log m)$ using Miller's algorithm or cubical arithmetic, and two discrete logarithms in μ_m, which takes $\mathcal{O}(\sqrt{m})$ assuming generic Pollard's rho. However, Algorithm 2 assumes the knowledge of a basis for $E[m]$, which is a rather strong assumption in practical applications, and may be the dominating cost.

4.2 Computing cokernel pairing using Weil and Tate pairings

In this section, we describe two algorithms to compute the cokernel pairing, assuming the methods to compute Φ_m from Section 4.1. We give an algorithm using the Weil pairing based on Algorithm 2 and a generic algorithm based on the Tate pairing using either Algorithm 1 or Algorithm 2.

4.2.1 Using the Weil pairing

The most straightforward computation computes $\Phi_m(P)$ and $\Phi_m(Q)$ directly and then computes $\langle P, Q\rangle_m$ as the Weil pairing of $\Phi_m(P)$ and $\Phi_m(Q)$. Such a Weil pairing takes roughly two Miller loops at a cost of $\mathcal{O}(\log m)$ and an inversion, for a total cost of two computations of Φ, two Miller loops, and one final inversion.

When using Algorithm 2 to compute Φ, with an assumed basis (T_1, T_2) of $E[m]$, we may improve significantly: this algorithm already describes $\Phi(P)$ and $\Phi(Q)$ as linear

combinations $rT_1 + sT_2$ with $r, s \in \mathbb{Z}/m\mathbb{Z}$ so that the Weil pairing between $\Phi(P)$ and $\Phi(Q)$ reduces to

$$\langle P, Q\rangle_m = e_m\left(r_P T_1 + s_P T_2, r_Q T_1 + s_Q T_2\right) = \zeta_0^{r_P s_Q - r_Q s_P}.$$

This is summarized in Algorithm 3.

Algorithm 3 Cokernel Pairing Computation using Weil Pairing

Input: Points $P, Q \in E(\mathbb{F}_q)$, a basis T_1, T_2 of $E[m]$, and $\zeta_0 = e_m(T_1, T_2) \in \mu_m$.
Output: The cokernel pairing $\langle P, Q\rangle_m \in \mu_m$

1: $(\zeta_1, \zeta_2) \leftarrow (t_m(T_1, P), t_m(T_2, -P))$
2: $(\zeta_3, \zeta_4) \leftarrow (t_m(T_1, Q), t_m(T_2, -Q))$
3: $(r_P, s_P) \leftarrow \left(\log_{\zeta_0}(\zeta_2), \log_{\zeta_0}(\zeta_1)\right)$
4: $(r_Q, s_Q) \leftarrow \left(\log_{\zeta_0}(\zeta_4), \log_{\zeta_0}(\zeta_3)\right)$
5: **return** $\zeta_0^{r_P s_Q - r_Q s_P}$

Cost. Assuming (T_1, T_2) and ζ_0 are known, the above algorithm computes four Tate pairings of degree m in $\mathcal{O}(\log q + \log m)$, and four discrete logarithms in μ_m in $\mathcal{O}(\sqrt{m})$.

Remark 4. We stress that, at its core, the above computation of $\langle P, Q\rangle_m$ requires us to first compute the positions of P and Q in $\mathcal{S}_m(E)/[m]\mathcal{S}_m(E)$. This happens in Algorithm 2 by computing the Tate profile $(t_m(T_1, P), t_m(T_2, P))$ with respect to a basis of $E[m]$. The resulting Weil pairing essentially verifies the independence of these Tate profiles. Thus, if we are only interested in non-triviality of the cokernel pairing, we may forgo the final Weil pairing and the discrete logarithm computations, and verify the independence of the Tate profiles themselves in μ_m^2.

Inverse of Φ_m. Given Algorithm 2, we may similarly wonder if we can compute the inverse $\Phi_m^{-1}(T)$ for $T \in E[m]$ as a map $E[m] \to E(\mathbb{F}_q)/[m]E(\mathbb{F}_q)$, given a basis (P, Q) of $E(\mathbb{F}_q)/[m]E(\mathbb{F}_q)$. This can be done with an algorithm very similar to Algorithm 2: given $\zeta = \langle P, Q\rangle_m$, we compute $\zeta_1 = t_m(T, -P)$ and $\zeta_2 = t_m(T, Q)$. We then find $a = \log_\zeta(\zeta_2)$ and similarly $b = \log_\zeta(\zeta_1)$ to get $\Phi_m^{-1}(T) = aP + bQ$.

Remark 5. Given the above description of an algorithm to invert the map Φ_m, we may complete the duality between the Weil pairing and the cokernel pairing. That is, given rational points $P, Q \in E[m]$, we may also compute $e_m(P, Q)$ as $\langle \Phi_m^{-1}(P), \Phi_m^{-1}(Q)\rangle_m$. This serves only for completeness, as it does not seem to give any benefit in computing $e_m(P, Q)$.

4.2.2 Using the Tate pairing

We may improve on the previous approach by the observation that we defined the (reduced) Tate pairing $t_m(P, Q)$ of degree m by $e_m\left(P, \pi(R) - R\right)$, where $[m]R = Q$. Thus, to compute $\langle P, Q\rangle_m$, we do not require both $\Phi_m(P)$ and $\Phi_m(Q)$: it is enough to compute $\Phi_m(P)$ and use $\langle P, Q\rangle_m = t_m(\Phi_m(P), Q)$. This gives Algorithm 4, which may be computed using either Algorithm 1 or Algorithm 2, depending on the knowledge of $E[m]$.

Algorithm 4 Cokernel Pairing Computation using Tate Pairing

Input: Points $P, Q \in E(\mathbb{F}_q)$.
Output: The cokernel pairing $\zeta = \langle P, Q\rangle_m \in \mu_m$

1: $P_m \leftarrow \Phi_m(P)$
2: $\zeta \leftarrow t_m(P_m, Q)$
3: **return** ζ

Cost. The above algorithm requires an additional Tate pairing of degree m in cost $\mathcal{O}(\log q + \log m)$ beyond the computation of Φ_m using either Algorithm 1 or Algorithm 2. In the former case, the cost is dominated by the computation of Φ_m, which takes $\mathcal{O}(m^{3.17} \log q)$ whereas in the latter, the total cost becomes three Tate pairings of degree m in $\mathcal{O}(\log q + \log m)$ together with two discrete logarithms in μ_m in $\mathcal{O}(\sqrt{m})$.

Table 1: Overview of the computational methods in this work for Φ_m and $\langle P, Q\rangle_m$ on an elliptic curve E over $\mathbb{F}_q$, in number of $\mathbb{F}_q$-operations. We denote the cost of polynomial multiplication of degree n over $\mathbb{F}_q$ by $M(n)$, so that $M(m^2) \approx m^{3.17}$ with Karatsuba. Discrete logarithms in $\mu_m \subseteq \mathbb{F}_q^*$ are computed with generic Pollard rho. Additional requirements are denoted with superscripts: Methods with † require knowledge of a torsion basis for $E[m]$, whereas ‡ assumes the curve is supersingular.

Method	Reference	Cost breakdown	Complexity
Computing Φ_m			
Naive (root of G)	Sec. 4.1	1 root of G ($\deg m^2$)	$\mathcal{O}(M(m^2)\, m \log q)$
Symbolic over the fiber	Alg. 1	2 exp. mod G ($\deg m^2$)	$\mathcal{O}(M(m^2) \log q)$
Rational, with basis†	Alg. 2	2 pairings, 2 DLogs	$\mathcal{O}(\log q + \sqrt{m})$
Scalar multiplication‡	Sec. 5.1	1 scalar mult.	$\mathcal{O}(\log p - \log m)$
Computing the cokernel pairing $\langle P, Q\rangle_m$			
Via Weil pairing†	Alg. 3	4 pairings, 4 DLogs	$\mathcal{O}(\log q + \sqrt{m})$
Via Tate pairing (symbolic)	Alg. 4	1 symbolic Φ_m, 1 pairing	$\mathcal{O}(M(m^2) \log q)$
Via Tate pairing (basis)†	Alg. 4	3 pairings, 2 DLogs	$\mathcal{O}(\log q + \sqrt{m})$
Via Tate pairing (scalar mult.)‡	Alg. 4	1 scalar mult., 1 pairing	$\mathcal{O}(\log p)$

4.3 Two concrete examples of cokernel pairings

We describe two concrete examples of cokernel pairings, one for degree $m = 2$ and one for degree $m = 5$, using the above methods to compute the pairing. We first discuss a supersingular example with symmetric Sylow m-torsion, for $m = 5$.

Example 1. Let $p = 4 \cdot 5^3 - 1$, and let $\mathbb{F}_q = \mathbb{F}_p(i)$ with $i^2 = -1$. Let $A = 439 + 245 \cdot i$ and let $E_A : y^2 = x^3 + Ax^2 + x$. The curve E_A is supersingular, and $E_A(\mathbb{F}_q) \cong \mathbb{Z}/(p+1)\mathbb{Z} \times \mathbb{Z}/(p+1)\mathbb{Z}$. Thus, for $m = 5$, we find that the Sylow m-torsion is symmetric:

$$\mathcal{S}_5(E_A) \cong \mathbb{Z}/(5^3)\mathbb{Z} \times \mathbb{Z}/(5^3)\mathbb{Z}.$$

A basis T_1, T_2 for $E_A[5]$ is given by $T_1 = (269 + 210 \cdot i, 319 + 35 \cdot i)$ and $T_2 = (498 + 486 \cdot i, 271 + 66 \cdot i)$, with $\zeta_0 = e_5(T_1, T_2) = 137 + 72 \cdot i$. We will compute the cokernel pairing for $P = (72 + 448 \cdot i, 433 + 172 \cdot i)$ and $Q = (467 + 169 \cdot i, 438 + 298 \cdot i)$. We first describe the naive approach: Solving Equation (7) over $\mathbb{F}_q(\alpha)$, where $\alpha^5 + i + 2 = 0$ gives the point $R = (x_R, y_R)$ such that $[5]R = P$, where

$$x_R = (439 + 185i)\alpha^4 + (321 + 112i)\alpha^3 + (461 + 412i)\alpha^2 + (178 + 222i)\alpha + 357 + 67i,$$
$$y_R = (295 + 234i)\alpha^4 + (335 + 159i)\alpha^3 + (141 + 409i)\alpha^2 + (317 + 20i)\alpha + 37 + 322i.$$

Then, we compute $\Phi_5(P) = \pi(R) - R = (492 + 177i, 399 + 442i) \in E[5]$. A similar computation gives $\Phi_5(Q) = (269 + 210i, 180 + 464i) \in E[5]$. The Weil pairing gives us

$$\langle P, Q\rangle_5 = e_5\left(\Phi_5(P), \Phi_5(Q)\right) = 137 + 427i.$$

By Theorem 1 for $h = 4$, we get that $[4]P$ and $[4]Q$ generate $\mathcal{S}_5(E)$, and, by its symmetry, both $[4]P$ and $[4]Q$ have order 5^3. For the approach using Algorithm 4, we first compute $\Phi_5(P) = [4]T_1 + [4]T_2$, and use $t_5(\Phi_5(P), Q)$ as another way to compute $\zeta = 137 + 427i$.

We get a different behavior in the following example of an ordinary curve with asymmetric Sylow m-torsion for $m = 2$.

Example 2. Let $p = 62723$, and let $\mathbb{F}_q = \mathbb{F}_p(i)$ with $i^2 = -1$. Let $a = 29939 + 47523 \cdot i$ and $b = 10859 + 6507 \cdot i$, and take $E : y^2 = x^3 + ax + b$. The curve E is ordinary, and $E(\mathbb{F}_q) \cong \mathbb{Z}/n_1\mathbb{Z} \times \mathbb{Z}/n_2\mathbb{Z}$, where $n_1 = 2^6 \cdot 3 \cdot 7 \cdot 365903$ and $n_2 = 2^3$. Thus, for $m = 2$, the Sylow m-torsion is asymmetric, with structure

$$\mathcal{S}_2(E) \cong \mathbb{Z}/(2^6)\mathbb{Z} \times \mathbb{Z}/(2^3)\mathbb{Z}.$$

We use the basis $T_1 = (54664 + 59102 \cdot i, 0)$ and $T_2 = (18942 + 2030 \cdot i, 0)$ for $E[2]$, with $\zeta_0 = e_2(T_1, T_2) = -1$. We compute the cokernel pairing for $P = (29237 + 15619 \cdot i, 1514 + 12755 \cdot i)$ and $Q = (51627 + 57123 \cdot i, 17021 + 6724 \cdot i)$ using Algorithm 4. From $t_2(T_1, P) = 1$ and $t_2(T_2, -P) = -1$, we get $r = 1$ and $s = 0$, so $\Phi_2(P) = T_1$. Then, by $t_2(\Phi_2(P), Q) = t_2(T_1, Q) = -1$ we find that $\langle P, Q\rangle_2 = -1$, hence P and Q generate $E(\mathbb{F}_q)/[2]E(\mathbb{F}_q)$. Thus, some multiple of P and Q must generate $\mathcal{S}_2(E)$, and this multiple is given by the cofactor $h = 3 \cdot 7 \cdot 365903$, so that we get the Sylow-2 basis $P' = [h]P$ and $Q' = [h]Q$. In this case, both P' and Q' are of order 2^6. For elegance, we may take $R' = [5]P' - Q'$ instead, which has order 2^3, and use the basis (P', R') for $\mathcal{S}_2(E)$.

5 Connections to the Weil and Tate Pairing

We derive results that connect the cokernel pairing to the Weil pairing and Tate profiles. Let $m = \ell$ be prime. When the Sylow ℓ-torsion is symmetric, say $\mathcal{S}_\ell(E) \cong \left(\mathbb{Z}/\ell^f\mathbb{Z}\right)^2$, we find an easy connection to the Weil pairing of degree ℓ^f. Namely, we get $\mathcal{S}_\ell(E) = E[\ell^f]$, and both the Weil pairing of degree ℓ^f as well as the cokernel pairing of degree ℓ describe conditions for a basis P, Q of this object: the Weil pairing $e_{\ell^f}(P, Q)$ has order ℓ^f for $P, Q \in E[\ell^f]$ if and only if P and Q are a basis for $E[\ell^f]$.

Lemma 2. *Let E be an elliptic curve over $\mathbb{F}_q$ with symmetric Sylow ℓ-torsion of exponent ℓ^f. Let $P, Q \in E(\mathbb{F}_q)$, and let h denote the cofactor with respect to ℓ. Then*

$$\langle P, Q\rangle_\ell \neq 1 \quad \Rightarrow \quad e_{\ell^f}([h]P, [h]Q) \text{ is a primitive } \ell^f\text{-th root of unity.}$$

Furthermore, if $P, Q \in E[\ell^f]$, then we get the other direction

$$e_{\ell^f}(P, Q) \text{ is a primitive } \ell^f\text{-th root of unity} \quad \Rightarrow \quad \langle P, Q\rangle_\ell \neq 1.$$

Proof. When $\langle P, Q\rangle_\ell \neq 1$, we get that $[h]P$ and $[h]Q$ generate $\mathcal{S}_\ell(E) = E[\ell^f]$, and so the Weil pairing $e_{\ell^f}([h]P, [h]Q)$ is a primitive ℓ^f-th root of unity. Conversely, if $e_{\ell^f}(P, Q)$ is a primitive ℓ^f-th root of unity, then P and Q generate $E[\ell^f] = \mathcal{S}_\ell(E)$, and so $\langle P, Q\rangle_\ell \neq 1$. □

In the above case, we get that both the cokernel pairing and the Weil pairing can be seen as pairings on $E[\ell^f] \times E[\ell^f]$, however, the cokernel pairing maps to μ_ℓ, whereas the Weil pairing maps to μ_{ℓ^f}. When Frobenius acts as a scalar, we can make this connection more explicit, as we see in Section 5.1.

5.1 The cokernel pairing on maximal supersingular curves

For supersingular elliptic curves E over $\mathbb{F}_{p^2}$ with trace $t = \pm 2p$, Frobenius acts as a scalar. We take $t = -2p$ for this example, i.e., E is a maximal supersingular curve of order $(p + 1)^2$, although the same reasoning holds for $t = 2p$ for E of order $(p - 1)^2$. In this case, the characteristic polynomial of π factors as $(x + p)^2$, hence π acts as $[-p]$, and $E(\mathbb{F}_{p^2}) = E[\pi - 1] = E[p + 1]$. We therefore must have symmetric Sylow ℓ-torsion, given by $\ell^f \parallel p + 1$. This simplifies the computation of the endomorphism Φ_ℓ, as the endomorphism $\frac{\pi-1}{\ell} : E(\mathbb{F}_q) \to E[\ell]$ simplifies to the scalar multiplication by $-\left[\frac{p+1}{\ell}\right]$, i.e., we clear everything except the ℓ-torsion.

Computing the cokernel pairing.

This alternative approach to computing Φ_ℓ requires a scalar multiplication by $[\frac{p+1}{\ell}]$, in comparison to the methods of Section 4.1, at a cost of $\mathcal{O}(\log p - \log \ell)$, with no need for a basis of $E[\ell]$. We may then compute the cokernel pairing using Algorithm 4 as

$$\langle P, Q\rangle_\ell = t_\ell\left(-[\tfrac{p+1}{\ell}]P, Q\right). \tag{9}$$

Writing $p+1 = \ell^k \cdot h$ with $\gcd(h,\ell) = 1$, this approach reinterprets the basis computation approach described in [CJL+17] using descent:

Corollary 3. *Let E be a supersingular elliptic curve over $\mathbb{F}_{p^2}$ with trace $t = -2p$ and $\ell^f \parallel p+1$. Let $P, Q \in E(\mathbb{F}_q)$ and let $h = \frac{p+1}{\ell^f}$. Then, if $t_\ell([h\cdot\ell^{f-1}]P, Q) \neq 1$, the points $([h]P, [h]Q)$ form a basis for $E[\ell^f]$.*

Furthermore, Equation (9) allows us to express the cokernel pairing of degree ℓ in terms of the Weil pairing of degree ℓ^f on such supersingular curves.

Lemma 3. *Let E be a supersingular elliptic curve over $\mathbb{F}_{p^2}$ with trace $t = -2p$ and with symmetric Sylow ℓ-torsion of exponent ℓ^f. Let $P, Q \in E[\ell^f]$. Let $\alpha \in \mathbb{Z}/\ell\mathbb{Z}$ such that $[\alpha] = -\left[\frac{p+1}{\ell^f}\right]$ on $E[\ell]$, so that $\Phi_\ell(P) = -\left[\frac{p+1}{\ell}\right]P = [\alpha\cdot\ell^{f-1}]P$. Then*

$$\langle P, Q\rangle_\ell = e_{\ell^f}(P, Q)^{\alpha^2\cdot\ell^{f-1}}.$$

Proof. Using $e_{nm}(P,Q) = e_n([m]P, Q)$ when $P \in E[nm]$ and $Q \in E[n]$, we get

$$e_{\ell^f}(P,Q)^{\ell^{f-1}} = e_{\ell^f}(P, [\ell^{f-1}]Q) = e_\ell([\ell^{f-1}]P, [\ell^{f-1}]Q).$$

Then, as $\Phi_\ell(P) = [\alpha\cdot\ell^{f-1}]P$, we get

$$\begin{aligned}\langle P, Q\rangle_\ell &= e_\ell(\Phi_\ell(P), \Phi_\ell(Q))\\ &= e_\ell([\alpha\cdot\ell^{f-1}]P, [\alpha\cdot\ell^{f-1}]Q)\\ &= e_\ell([\ell^{f-1}]P, [\ell^{f-1}]Q)^{\alpha^2}\\ &= e_{\ell^f}(P,Q)^{\alpha^2\cdot\ell^{f-1}}.\end{aligned}$$

□

Remark 6. Although this section describes an alternative way to compute the cokernel pairing on minimal/maximal supersingular curves which does not require a basis for $E[\ell]$, this approach requires either a large scalar multiplication, using Corollary 3 or a Weil pairing of degree ℓ^f, using Lemma 3. Both computations are expensive when f is large, and will quickly exceed the cost of Algorithm 4 using Algorithm 2 if $E[\ell]$ is known.

5.2 When Frobenius is not a scalar

The simplified computation from Section 5.1 relies crucially on the fact that we can rewrite the action of $\frac{\pi-1}{\ell}$ as a scalar multiplication. For other curves, even if the Sylow ℓ-torsion is symmetric, which implies that a scalar multiplication $E(\mathbb{F}_q) \to E[\ell]$ exists, the action of $\frac{\pi-1}{\ell}$ is different. For example, let $E/\mathbb{F}_q$ be an ordinary curve with $\mathcal{S}_\ell(E) = E[\ell^f]$. Write $E(\mathbb{F}_q) \cong \mathbb{Z}/n_1\mathbb{Z} \times \mathbb{Z}/n_2\mathbb{Z}$ with $n_2 \mid n_1$, so that we may write $n_1 = e\cdot r\cdot\ell^f$ and $n_2 = r\cdot\ell^f$ for some integers e and r. Let (P,Q) generate $E(\mathbb{F}_q)$ with P of order n_1 and Q of order n_2. Then one can similarly define a map $E(\mathbb{F}_q) \to E[\ell]$ by $P \mapsto [e\cdot r\cdot\ell^{f-1}]P$ and $Q \mapsto [r\cdot\ell^{f-1}]Q$, that is, we are clearing the cofactor. However, this map is crucially different from $\Phi_\ell = \frac{\pi-1}{\ell}$.

Example 3. We take $p = 535303$, $\ell = 3$ and $f = 3$. Then the curve

$$E/\mathbb{F}_p : y^2 = x^3 + 262034x^2 + x,$$

satisfies $\mathcal{S}_\ell(E) = E[\ell^f]$. We take $P = (533658, 176488)$ and $Q = (402889, 457605)$ as generators of $E(\mathbb{F}_q)$ and as an implicit basis of $\mathcal{S}_\ell(E)$. By going to $\mathbb{F}_{p^3}$, we are able to compute $\Phi_\ell(P) = (503161, 476634)$ and $\Phi_\ell(Q) = (525015, 190104)$, whereas clearing the cofactor gives us points $P' = (434272, 111323)$ and $Q' = (525015, 345199)$. These are connected by $\Phi_\ell(P) = [2](P' - Q')$ and $\Phi_\ell(Q) = [2]Q'$.

In general, with no assumption on the Sylow ℓ-torsion or the trace of Frobenius, the more natural connection of the cokernel pairing is to Tate profiles, as the following result shows.

Lemma 4. *The cokernel pairing $\langle P, Q\rangle_\ell$ is non-trivial if and only if the Tate profiles $t_{[\ell]}(P)$ and $t_{[\ell]}(Q)$ are non-trivial and independent.*

Proof. This is straightforward from Theorem 1 and the fact that non-trivial and independent Tate profiles imply a basis for the Sylow ℓ-torsion $\mathcal{S}_\ell(E)$ [Reij25]. □

We stress the difference between these two objects: the Tate profile explicitly requires a basis for $E[\ell]$ to give coordinates to $\mathcal{S}_\ell(E)/[\ell]\mathcal{S}_\ell(E)$, from which we derive that two independent profiles generate the Sylow ℓ-torsion, after computing the position of these points with respect to the given basis. The cokernel pairing, however, is formulated independently of a basis, and does not give us the position of these points. On the one hand, this implies that we have less information, however, we still have enough information to obtain a basis for the Sylow ℓ-torsion. On the other hand, this implies that we may hope to compute $\langle P, Q\rangle_\ell$ without a basis for $E[\ell]$, which is impossible for the Tate profile.

6 Applications of the Cokernel Pairing

We derive the main application of the cokernel pairing directly from Theorem 1, namely, finding a basis for the Sylow ℓ-torsion of an elliptic curve E over a finite field $\mathbb{F}_q$. As isogeny-based cryptography often works with supersingular curves over $\mathbb{F}_{p^2}$, we are specifically interested in finding a basis of $E[\ell^f]$. From a practical point of view, the cokernel allows us to compute an *implicit basis* for $E[\ell^f]$, which improves the efficiency of computing kernel points of order ℓ^f in practice.

6.1 Computing a Sylow torsion basis

Computing a basis for $E[\ell^f]$ is highly optimized for $\ell = 2$ on Montgomery curves using entangled basis generation [ZSP+18], and can be generalized to other curve models when $E[2]$ is known [Reij25]. Nevertheless, these versions of entangled basis generation are unenlightening for $\ell > 2$, and in this case, basis generation algorithms are more ad-hoc, especially when the Sylow torsion is asymmetric.

Using the cokernel pairing, we find an intuitive and straightforward approach, given a basis for the kernel $E[m]$, even for composite m. If we do not care for efficiency and simply want an easy method to find such a basis, we may sample random points in $E(\mathbb{F}_q)$, until we find a pair (P, Q) where $\langle P, Q\rangle_m$ is a primitive m-th root of unity.

A more efficient approach uses a combination of Tate pairings and cokernel pairings. Let T_1, T_2 be a basis for $E[m]$. First, we use the Tate pairings of degree m with kernel points T_1 and T_2 with random cokernel points $P \in E(\mathbb{F}_q)/[m]E(\mathbb{F}_q)$ until we get a non-trivial pair $\zeta_1 = t_m(T_1, P)$ and $\zeta_2 = t_m(T_2, -P)$, i.e., they generate an order-m subgroup of μ_m^2. Given ζ_1 and ζ_2, we compute $\Phi_m(P)$, and sample $Q \in E(\mathbb{F}_q)$ until $\langle P, Q\rangle_m$ is an m-th root of unity, which requires the $\Phi_m(P)$ computed before. After multiplication by $[h]$, we then have a basis for $\mathcal{S}_m(E)$. This is summarized in Algorithm 5

Algorithm 5 Sylow Torsion Basis generation

Input: A basis T_1, T_2 of $E[m]$, and $\zeta_0 = e_m(T_1, T_2) \in \mu_m$.
Output: A basis (P, Q) for $\mathcal{S}_m(E)$

repeat
  $P \xleftarrow{\$} E(\mathbb{F}_q)$
  $\zeta_1 \leftarrow t_m(T_1, P)$
  $\zeta_2 \leftarrow t_m(T_2, -P)$
until (ζ_1, ζ_2) has order m in μ_m^2
$(r, s) \leftarrow \left(\log_{\zeta_0}(\zeta_2), \log_{\zeta_0}(\zeta_1)\right)$
$P_m \leftarrow rT_1 + sT_2$
repeat
  $Q \xleftarrow{\$} E(\mathbb{F}_q)$
  $\zeta \leftarrow \langle P, Q\rangle_m = t_m(P_m, Q)$
until ζ has order m
return $([h]P, [h]Q)$

Cost analysis. To compare against the performance of the approach using Tate profiles, we compute the probability of success of both approaches, and the expected number of degree-ℓ Tate pairings we need to compute.

First, both approaches sample a point P at random until the Tate pairings $\zeta_1 = t_\ell(T_1, P)$ and $\zeta_2 = t_\ell(T_2, -P)$ together are non-trivial.[2] The probability of failure is $\frac{1}{\ell^2}$, as this only happens when we sample $P \in [\ell]E(\mathbb{F}_q)$. Hence, on average, this requires $2 \cdot \frac{\ell^2}{\ell^2-1}$ Tate pairings for both approaches.

Both approaches then need to sample a random Q that completes the basis (P, Q). As P generates a subgroup of order ℓ in $E(\mathbb{F}_q)/[\ell]E(\mathbb{F}_q)$, which itself has order ℓ^2, we have a success probability $1 - \frac{\ell}{\ell^2} = \frac{\ell-1}{\ell}$, and so, we expect to require $\ell/(\ell-1)$ samples of Q on average. Per Q, the approach using the cokernel pairing needs a single Tate pairing to confirm Q is correct. The approach using Tate profiles requires two Tate pairings to confirm Q is independent of P, except when one of ζ_1 or ζ_2 is trivial, in which case it only requires one. Such trivial ζ_i happen with probability $\frac{2(\ell-1)}{\ell^2-1}$, and so on average, we need $\frac{2\cdot(\ell-1)^2+2(\ell-1)}{\ell^2-1} \approx 2 - \frac{2}{\ell+1}$ Tate pairings for Q.

Additionally, the cokernel approaches requires discrete logarithms to compute the coefficients of P_ℓ, whereas the Tate profile approach requires discrete logarithms to verify the correctness of Q. The cost of these is equal for both approaches. In practical applications, ℓ is almost always a small prime, so that discrete logarithms in μ_ℓ are efficient. Overall, we find that the cokernel approach requires an expected $2 \cdot \frac{\ell^2}{\ell^2-1} + \frac{\ell}{\ell-1} = 3 + \frac{\ell+3}{\ell^2-1}$ Tate pairings, and therefore saves on average $\frac{\ell}{\ell+1}$ Tate pairings in comparison to the approach using Tate profiles, by directly computing the 'correct' Tate pairing $t_\ell(P_\ell, Q)$. Table 2 lists the expected number of degree-ℓ Tate pairings for both approaches for small primes ℓ.

Table 2: Expected number of degree-ℓ Tate pairings for the two approaches.

ℓ	Average	3	5	7	11	13	17	19
Tate profiles	$\frac{4\ell^2}{\ell^2-1}$	4.50	4.17	4.08	4.03	4.02	4.01	4.01
Cokernel	$3 + \frac{\ell+3}{\ell^2-1}$	3.75	3.33	3.21	3.12	3.10	3.07	3.06

[2] One may also consider the approach of randomly sampling P until ζ_1 is non-trivial, and then computing ζ_2. The expected value is slightly worse for this approach.

6.2 Describing a kernel point of an ℓ^f-isogeny

Most isogeny-based schemes require us to compress a point $K \in E[\ell^f]$ using a deterministic basis (P', Q') of $E[\ell^f]$ as $K = [a]P' + [b]Q'$ with $a, b \in \mathbb{Z}/(\ell^f)\mathbb{Z}$, so that we may communicate K, and therefore the isogeny $E \to E/\langle K \rangle$, using only the values[3] a and b, instead of, e.g., the x-coordinate of K. This is a common technique in many isogeny-based schemes [JAC+17; DKL+20; BDF+24; AAA+25].

We have seen already that we may generalize the requirement on P and Q to a basis of $\mathcal{S}_m(E)$ instead, and even more generally to a basis of $E(\mathbb{F}_q)/[m]E(\mathbb{F}_q)$. The definition of an *implicit basis* [CEMR24] captures the difference between a basis for $E(\mathbb{F}_q)/[m]E(\mathbb{F}_q)$ and $\mathcal{S}_m(E)/[m]\mathcal{S}_m(E)$.

Definition 4. Let $P, Q \in E(\mathbb{F}_q)$. We say that (P, Q) is an *implicit basis* for $\mathcal{S}_m(E)$ if there is an $h \in \mathbb{Z}_{>0}$, co-prime to m, such that $([h]P, [h]Q)$ is a basis for $\mathcal{S}_m(E)$.

Given a deterministically-sampled implicit basis (P, Q), we can similarly write $K = [a]([h]P) + [b]([h]Q)$. The gain is that we may now compute K as $K = [h]([a]P + [b]Q)$, which saves one application of the map $[h]$. For large cofactors, saving such a scalar multiplication can be significant [CEMR24]. That is, we may consider the implicit basis (P, Q) as an actual basis for compression, but the points themselves are not in $\mathcal{S}_m(E)$, only after applying the map $[h] : E(\mathbb{F}_q) \to \mathcal{S}_m(E)$.

The cokernel pairing allows us to generate a basis (P, Q) for $E(\mathbb{F}_q)/[m]E(\mathbb{F}_q)$, which is by definition an implicit basis for $\mathcal{S}_m(E)$. For example, in Algorithm 5, we may simply ignore the last multiplication by $[h]$ and return the implicit basis P, Q. Thus, the cokernel pairing gives us a very simple and natural tool to compute implicit bases (P, Q), summarized in the following result.

Corollary 4. *Let $P, Q \in E(\mathbb{F}_q)$. When $\langle P, Q \rangle_m$ is a primitive m-th root of unity, the pair (P, Q) is an implicit basis of $\mathcal{S}_m(E)$. The cofactor $h \in \mathbb{Z}_{>0}$ maps the implicit basis (P, Q) to an explicit basis $([h]P, [h]Q)$ of $\mathcal{S}_m(E)$.*

We conclude that we only need an implicit basis (P, Q) of $\mathcal{S}_m(E)$ to describe an isogeny effectively, and the cokernel pairing gives us the exact criterion for when P and Q are an implicit basis. Although torsion basis generation and point compression are elementary building blocks of isogeny-based cryptography, they are a significant part of the computational cost of isogeny-based primitives. Hence, the simplified approach using the cokernel pairing may improve several isogeny-based schemes.

7 Future Work

We have introduced the cokernel pairing and explored initial computations, applications, and connections to the Weil and Tate pairing. This opens up many questions for future work.

The main question is on an improved computation of the cokernel pairing. Both Algorithm 3 and Algorithm 4 explicily compute one or both of $\Phi_m(P)$ and $\Phi_m(Q)$, either symbolically using Algorithm 1 or by using a basis of $E[m]$ using Algorithm 2, and then apply a Weil or Tate pairing. Remarkably, the Tate pairing $e_m(P, \Phi_m(Q))$ may instead be computed without explicitly computing $\Phi_m(Q)$. We may hope that we can compute $\langle P, Q \rangle_m$ similarly, without a direct computation of Φ_m or using the Tate pairing. On the other hand, we were unable to show that a cokernel pairing compuation *needs* to compute Φ_m, and it seems difficult to show that such a computation is required. The peculiar case of maximal supersingular elliptic curves shows a different but expensive way to compute Φ_m.

[3] Often we only need one of these two: as we need K of order ℓ^f, either a or b is invertible, and we may choose P and Q in such a way that we can always express a generator of the same kernel by $P + a^{-1}bQ$.

Future work may explore the usage of cubical arithmetic [Rob24] for efficient computations of cokernel pairings.

Another direction is in generalizations of the cokernel pairing. We may similarly explore the cokernel pairing defined on more general abelian varieties.

Another direction for generalization is inspired by [CHM+23], which explores Tate pairings $T_f^\alpha(P,Q)$ defined as $e_f(P,\alpha(R))$ for α a suitable endomorphism and $f(R)=Q$, which coincides with the f-Tate pairing for $\alpha=\pi-1$. For the cokernel pairing, we could similarly replace the role of $\pi-1$ either in one argument, or in both arguments, given suitable endomorphisms α,β.

A more concrete direction of research is related to entangled basis generation [ZSP+18], which results in a basis (P,Q) for $E[2^f]$ on specific supersingular elliptic curves $E/\mathbb{F}_{p^2}$ by a rather arbitrary choice of x_P and x_Q. Understanding why this choice of x_P and x_Q ensures $\langle P,Q\rangle_2 \neq 1$ may help in generalizing entangled basis generation to primes $\ell>2$ or genus $g>1$.

AI disclosure. The main contribution of LLMs in this work lies in the symbolic computation of Section 4.1.2, which was derived in a discussion with Claude Opus 5. The author takes full responsibility for this section. Other usage of Opus 5 was editorial in nature, with feedback on storyline, grammar, and readability.

References

[AAA+25] Marius A. Aardal et al. SQIsign 2.0: Algorithm specifications and supporting documentation. Technical report, 2025.

[BDF+24] Andrea Basso, Pierrick Dartois, Luca De Feo, Antonin Leroux, Luciano Maino, Giacomo Pope, Damien Robert, and Benjamin Wesolowski. Sqisign2d-west - the fast, the small, and the safer. In Kai-Min Chung and Yu Sasaki, editors, *Advances in Cryptology - ASIACRYPT 2024 - 30th International Conference on the Theory and Application of Cryptology and Information Security, Kolkata, India, December 9-13, 2024, Proceedings, Part III*, volume 15486 of *Lecture Notes in Computer Science*, pages 339–370. Springer, 2024. ISBN: 978-981-96-0890-4. DOI: 10.1007/978-981-96-0891-1_11.

[BDLS20] Daniel J. Bernstein, Luca De Feo, Antonin Leroux, and Benjamin Smith. Faster computation of isogenies of large prime degree. In *ANTS-XIV-14th Algorithmic Number Theory Symposium*, volume 4, pages 39–55. Mathematical Sciences Publishers, 2020.

[BF01] Dan Boneh and Matt Franklin. Identity-based encryption from the weil pairing. In *Annual international cryptology conference*, pages 213–229. Springer, 2001.

[BGLS03] Dan Boneh, Craig Gentry, Ben Lynn, and Hovav Shacham. Aggregate and verifiably encrypted signatures from bilinear maps. In *Advances in Cryptology—EUROCRYPT 2003: International Conference on the Theory and Applications of Cryptographic Techniques, Warsaw, Poland, May 4–8, 2003 Proceedings 22*, pages 416–432. Springer, 2003.

[BKLS02] Paulo S. L. M. Barreto, Hae Y. Kim, Ben Lynn, and Michael Scott. Efficient algorithms for pairing-based cryptosystems. In *Advances in Cryptology—CRYPTO 2002: 22nd Annual International Cryptology Conference Santa Barbara, California, USA, August 18–22, 2002 Proceedings 22*, pages 354–369. Springer, 2002.

[BLS01] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the weil pairing. In *International conference on the theory and application of cryptology and information security*, pages 514–532. Springer, 2001.

[BLS04] Paulo S. L. M. Barreto, Ben Lynn, and Michael Scott. Efficient implementation of pairing-based cryptosystems. *Journal of Cryptology*, 17:321–334, 2004.

[Bru11] Peter Bruin. The tate pairing for abelian varieties over finite fields. *Journal de theorie des nombres de Bordeaux*, 23(2):323–328, 2011.

[CEMR24] Maria Corte-Real Santos, Jonathan Komada Eriksen, Michael Meyer, and Krijn Reijnders. Aprèssqi: extra fast verification for sqisign using extension-field signing. In *Advances in Cryptology - EUROCRYPT 2024 - 43nd Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2024.

[CHM+23] Wouter Castryck, Marc Houben, Simon-Philipp Merz, Marzio Mula, Sam van Buuren, and Frederik Vercauteren. Weak instances of class group action based cryptography via self-pairings. In *Annual International Cryptology Conference*, pages 762–792. Springer, 2023.

[CJL+17] Craig Costello, David Jao, Patrick Longa, Michael Naehrig, Joost Renes, and David Urbanik. Efficient compression of sidh public keys. In *Advances in Cryptology–EUROCRYPT 2017: 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30–May 4, 2017, Proceedings, Part I 36*, pages 679–706. Springer, 2017.

[CR25] Maria Corte-Real Santos and Krijn Reijnders. Return of the kummer: a toolbox for genus-2 cryptography. *IACR Communications in Cryptology*, 2(3), 2025.

[DEF+25] Pierrick Dartois, Jonathan Komada Eriksen, Tako Boris Fouotsa, Arthur Herlédan Le Merdy, Riccardo Invernizzi, Damien Robert, Ryan Rueger, Frederik Vercauteren, and Benjamin Wesolowski. PEGASIS: practical effective class group action using 4-dimensional isogenies. In Yael Tauman Kalai and Seny F. Kamara, editors, *Advances in Cryptology - CRYPTO 2025 - 45th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 17-21, 2025, Proceedings, Part I*, volume 16000 of *Lecture Notes in Computer Science*, pages 67–99. Springer, 2025. ISBN: 978-3-032-01854-0. DOI: 10.1007/978-3-032-01855-7_3.

[DKL+20] Luca De Feo, David Kohel, Antonin Leroux, Christophe Petit, and Benjamin Wesolowski. SQISign: Compact Post-quantum Signatures from Quaternions and Isogenies. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2020*, pages 64–93, Cham. Springer International Publishing, 2020. ISBN: 978-3-030-64837-4.

[EVM12] Bas Edixhoven, Gerard Van der Geer, and Ben Moonen. *Abelian varieties*. 2012. URL: https://www.math.ru.nl/~bmoonen/research.html.

[FR94] Gerhard Frey and Hans-Georg Rück. A remark concerning m-divisibility and the discrete logarithm in the divisor class group of curves. *Mathematics of computation*, 62(206):865–874, 1994.

[Gal05] Steven D. Galbraith. Pairings. london mathematics society lecture note series, vol. 317, 2005.

[GHS02] Steven D. Galbraith, Keith Harrison, and David Soldera. Implementing the tate pairing. In *International Algorithmic Number Theory Symposium*, pages 324–337. Springer, 2002.

[HSSI99] Ryuichi Harasawa, Junji Shikata, Joe Suzuki, and Hideki Imai. Comparing the mov and fr reductions in elliptic curve cryptography. In *Advances in Cryptology—EUROCRYPT'99: International Conference on the Theory and Application of Cryptographic Techniques Prague, Czech Republic, May 2–6, 1999 Proceedings 18*, pages 190–205. Springer, 1999.

[HSV06] Florian Hess, Nigel P. Smart, and Frederik Vercauteren. The eta pairing revisited. *IEEE Transactions on Information Theory*, 52(10):4595–4602, 2006. DOI: 10.1109/TIT.2006.881709.

[JAC+17] David Jao, Reza Azarderakhsh, Matt Campagna, Craig Costello, Luca de Feo, Basil Hess, Amir Jalili, Brian Koziel, Brian Lamacchia, Patrick Longa, Michael Naehrig, Joost Renes, Vladimir Soukharev, and David Urbanik. SIKE: Supersingular Isogeny Key Encapsulation. Technical report, 2017.

[Jou02] Antoine Joux. The weil and tate pairings as building blocks for public key cryptosystems: survey. In *International Algorithmic Number Theory Symposium*, pages 20–32. Springer, 2002.

[Jou04] Antoine Joux. A one round protocol for tripartite diffie–hellman. *Journal of cryptology*, 17(4):263–276, 2004.

[KT18] Takeshi Koshiba and Katsuyuki Takashima. New assumptions on isogenous pairing groups with applications to attribute-based encryption. In *International Conference on Information Security and Cryptology*, pages 3–19. Springer, 2018.

[Lic69] Stephen Lichtenbaum. Duality theorems for curves over p-adic fields. *Inventiones mathematicae*, 7(2):120–136, 1969.

[LR16] David Lubicz and Damien Robert. Arithmetic on abelian and Kummer varieties. *Finite Fields and Their Applications*, 39:130–158, May 2016. DOI: 10.1016/j.ffa.2016.01.009. eprint: 2014/493, HAL: hal-01057467.

[LWXZ24] Kaizhan Lin, Weize Wang, Zheng Xu, and Chang-An Zhao. A faster software implementation of sqisign. *IEEE Transactions on Information Theory*, 2024.

[Mil04] Victor S Miller. The Weil pairing, and its efficient calculation. *Journal of cryptology*, 17(4):235–261, 2004.

[MOV91] Alfred Menezes, Tatsuaki Okamoto, and Scott Vanstone. Reducing elliptic curve logarithms to logarithms in a finite field. In *Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing*, STOC '91, pages 80–89, New Orleans, Louisiana, USA. Association for Computing Machinery, 1991. ISBN: 0897913973. DOI: 10.1145/103418.103434.

[MS24] Joseph Macula and Katherine E Stange. Extending class group action attacks via sesquilinear pairings. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 371–395. Springer, 2024.

[PRR+25] Giacomo Pope, Krijn Reijnders, Damien Robert, Alessandro Sferlazza, and Benjamin Smith. Simpler and faster pairings from the montgomery ladder. Cryptology ePrint Archive, Paper 2025/672, 2025. URL: https://eprint.iacr.org/2025/672.

[Reij23] Krijn Reijnders. Effective Pairings in Isogeny-based Cryptography. In *Progress in Cryptology - LATINCRYPT 2023 - 8th International Conference on Cryptology and Information Security in Latin America*, pages 109–128. Springer, 2023.

[Reij25] Krijn Reijnders. The Tate Profile. In *Progress in Cryptology - LATINCRYPT 2025 - 9th International Conference on Cryptology and Information Security in Latin America*. Springer, 2025.

[Rob23] Damien Robert. The geometric interpretation of the tate pairing and its applications. Cryptology ePrint Archive, Paper 2023/177, 2023. URL: https://eprint.iacr.org/2023/177.

[Rob24] Damien Robert. Fast pairings via biextensions and cubical arithmetic. April 2024. eprint: 2024/517, HAL: hal-04848028.

[Sil09] Joseph H Silverman. *The arithmetic of elliptic curves*, volume 106. Springer, 2009.

[Sil10] Joseph H Silverman. A survey of local and global pairings on elliptic curves and abelian varieties. In *Pairing-Based Cryptography-Pairing 2010: 4th International Conference, Yamanaka Hot Spring, Japan, December 2010. Proceedings 4*, pages 377–396. Springer, 2010.

[Sta08] Katherine Stange. *Elliptic nets and elliptic curves*. PhD thesis, Brown University, 2008. URL: https://repository.library.brown.edu/studio/item/bdr:309/PDF/.

[Sta11] Katherine Stange. Elliptic nets and elliptic curves. *Algebra & Number Theory*, 5(2):197–229, 2011.

[Tat62] John Tate. Duality theorems in Galois cohomology over number fields. In *Proc. Internat. Congr. Mathematicians (Stockholm, 1962)*, pages 288–295, 1962.

[Vél71] Jacques Vélu. Isogénies entre courbes elliptiques. *Comptes-Rendus de l'Académie des Sciences*, 273:238–241, 1971. URL: https://gallica.bnf.fr/ark:/12148/cb34416987n/date. https://gallica.bnf.fr/ark:/12148/cb34416987n/date.

[Ver09] Frederik Vercauteren. Optimal pairings. *IEEE transactions on information theory*, 56(1):455–461, 2009.

[Wei40] André Weil. Sur les fonctions algébriquesa corps de constantes fini. *CR Acad. Sci. Paris*, 210(1940):592–594, 1940.

[ZSP+18] Gustavo HM Zanon, Marcos A Simplicio, Geovandro CCF Pereira, Javad Doliskani, and Paulo SLM Barreto. Faster key compression for isogeny-based cryptosystems. *IEEE Transactions on Computers*, 68(5):688–701, 2018.