

Fast and simple constant-time hashing to the BLS12-381 elliptic curve

Riad S. Wahby and Dan Boneh

Stanford University

rsw@cs.stanford.edu, dabo@cs.stanford.edu

Abstract. Pairing-friendly elliptic curves in the Barreto-Lynn-Scott family are seeing a resurgence in popularity because of the recent result of Kim and Barbulescu that improves attacks against other pairing-friendly curve families. One particular Barreto-Lynn-Scott curve, called BLS12-381, is the locus of significant development and deployment effort, especially in blockchain applications. This effort has sparked interest in using the BLS12-381 curve for BLS signatures, which requires hashing to one of the groups of the bilinear pairing defined by BLS12-381.

While there is a substantial body of literature on the problem of hashing to elliptic curves, much of this work does not apply to Barreto-Lynn-Scott curves. Moreover, the work that does apply has the unfortunate property that fast implementations are complex, while simple implementations are slow.

In this work, we address these issues. First, we show a straightforward way of adapting the “simplified SWU” map of Brier et al. to BLS12-381. Second, we describe optimizations to this map that both simplify its implementation and improve its performance; these optimizations may be of interest in other contexts. Third, we implement and evaluate. We find that our work yields constant-time hash functions that are simple to implement, yet perform within 9% of the fastest, non-constant-time alternatives, which require much more complex implementations.

Keywords: pairing-friendly · elliptic curves · hashing · Barreto-Lynn-Scott · BLS12-381

1 Introduction

The Barreto-Lynn-Scott family of pairing-friendly elliptic curves [BLS03], and in particular the elliptic curve BLS12-381 [Bow17] (§2.1), has recently seen widespread adoption (e.g., in pairing-based SNARKs [GGPR13, PHGR13, BCTV14, Gro16]), largely because of the recent result of Kim and Barbulescu [KB16] that speeds up attacks on the discrete log problem in finite field extensions (for more information, see [MSS16]).

The availability of high-quality BLS12-381 implementations combined with the desire for aggregatable signatures [BGLS03] has sparked interest [Chi, Eth, BGWZ19, YKS19] in using BLS12-381 for BLS signatures [BLS01] (§2.2). The BLS signature scheme requires a hash function to points in a prime-order subgroup of a pairing-friendly curve. For this purpose, the authors suggest a method based on folklore that they call MapToGroup [BLS01, §3.3] (we call this method “hash-and-check”), which works roughly as follows: pick a random element in the elliptic curve’s base field and check whether it is the x -coordinate of a rational point on the curve. If it is, return that point, otherwise try again.

While hash-and-check is simple to implement and fast in expectation, it is not without downsides. Most importantly, it is not possible to make hash-and-check run in *constant time*—that is, time independent of the hash input—with both good performance and low failure probability. For BLS signatures, a constant-time hash function is not strictly necessary for security. On the other hand, because hash-and-check on a random message

takes k checks with probability $\approx 2^{-k}$, it is relatively easy to (accidentally or adversarially) choose messages that are difficult to hash, which wastes verifiers' and signers' time. Moreover, in practice cryptographic primitives often see “mission creep,” meaning that a constant-time hash function is desirable as a defense against future (mis)use.

Several lines of work in both the number theory and cryptography literature have considered the problem of deterministically mapping to rational points on elliptic curves; we briefly survey in Section 1.1. Unfortunately, most of these constructions do not apply to BLS12-381, because they are restricted to, e.g., elliptic curves of particular shapes or over base fields of specific characteristic.

One exception is the seminal work of Shallue and van de Woestijne [SvdW06] (§2.3), which applies to essentially any elliptic curve. This map can be used for BLS12-381, but fast implementations are complex, while simple ones are slow—especially for constant-time implementations or for embedded applications that rely on special-purpose field arithmetic accelerators (§6). At a high level, this is because evaluating the map requires evaluating Legendre symbols, for which the simple algorithm is an exponentiation (which is expensive) while the fast algorithm involves reductions modulo essentially random integers [Coh93, §1.4.2] (which entails substantial implementation complexity for good performance). Fouque and Tibouchi give an explicit construction tailored to the Barreto-Naehrig curve family, but it is undefined at several points when applied to BLS12-381, further increasing complexity (i.e., to detect and handle the undefined cases) and making implementations with input-independent runtime yet more difficult.

Ulas [Ula07] describes a simpler version of the Shallue–van de Woestijne map; Brier et al. [BCI⁺10, §7] give a further simplification and name it the “simplified SWU” map. This map is attractive because it has somewhat reduced computational cost and is easier to describe and implement compared to the original Shallue–van de Woestijne map. But it only applies to curves with j -invariant $\notin \{0, 1728\}$, almost surgically preventing its application to most pairing-friendly curve families (including Barreto-Lynn-Scott), which have j -invariant either 0 or 1728 for efficiency reasons [BN06, HSV06].

A second issue is that the description by Brier et al. is (somewhat artificially) restricted to curves over fields \mathbb{F} where $\#\mathbb{F} \equiv 3 \pmod{4}$, meaning that it does not apply to curves over extension fields of even degree (since $p^{2k} \equiv 1 \pmod{4}$). This is a concern because, for Barreto-Lynn-Scott and other pairing-friendly curves, one group of the bilinear pairing is a subgroup of an elliptic curve over an even-order extension field (for BLS12-381, a quadratic extension; §2.1). Thus, the simplified SWU map as described does not work for this group.

Our contributions. We show that careful design choices and optimizations yield hash functions that admit fast, simple, and constant-time implementations. Our focus is on the BLS12-381 elliptic curve, but we describe our design methods and optimizations with an eye to straightforward application to other curves. Our specific contributions are:

- In Section 3, we give explicit Shallue and van de Woestijne maps tailored to the BLS12-381 curve. We also describe a simple method for designing exception-free maps of this type, which applies generically to other elliptic curves.
- In Section 4, we give “indirect” maps for BLS12-381 based on the simplified SWU map, which work by mapping to an isogenous curve with nonzero j -invariant, then evaluating the isogeny map. To do so, we extend the simplified SWU map to $\#\mathbb{F} \equiv 1 \pmod{4}$, and thus to even-order extension fields. We also describe several optimizations that make the SWU map simpler to implement and faster to evaluate, including in constant time. Our optimizations apply generically, and can be used to speed up any implementation of the simplified SWU map.
- In Section 5 we describe explicit hash functions built on the above maps, based on known constructions. We briefly discuss security and efficiency in our context.

- In Section 6, we implement and evaluate.

We find that, for implementations built on a rich multi-precision library like GMP [GMP] (in particular, one that supports fast reductions modulo arbitrary integers, and provides fast Legendre symbol and extended Euclidean algorithms), hashes using the map of Section 3 are up to $\approx 9\%$ faster than hashes that use the map of Section 4.

For implementations restricted to using only field operations—typical restrictions for small cryptographic libraries or for compatibility with hardware accelerators (§6)—hashes that use the map of Section 4 are $\approx 1.3\text{--}2\times$ faster than hashes that use the map of Section 3. More surprisingly, our optimizations yield constant-time hashes based on the map of Section 4 that are at worst $\approx 9\%$ slower than the fastest *non-constant-time* implementations of the Section 3 map.

1.1 Related work

Deterministic maps to rational points on elliptic curves. Schinzel and Skałba [SS04] give the first deterministic method of constructing rational points on elliptic curves. For curves $E(\mathbb{F}) : y^2 = x^n + k$ with $x \in \{3, 4\}$ and \mathbb{F} of characteristic greater than 3, their construction yields at most four points, which are parameterized by k . Skałba [Ska05] subsequently gives a more general construction for points on curves of the form $E(\mathbb{F}) : y^2 = x^3 + ax + b$ where $a \neq 0$ and \mathbb{F} has characteristic greater than 3.

Shallue and van de Woestijne [SvdW06] generalize and simplify Skałba’s construction, giving concretely efficient rational maps to essentially any elliptic curve. Ulas [Ula07] further simplifies the construction, at the cost of restricting it to curves of the form $E(\mathbb{F}_p) : y^2 = x^3 + ax + b, ab \neq 0$, i.e., curves with j -invariant $\notin \{0, 1728\}$; the author also generalizes the results to some hyperelliptic curves. Brier et al. [BCI⁺10] give yet a further simplification, and Fouque and Tibouchi [FT10] tweak this version to simplify their analysis of the size of its image. In Section 4, we further optimize this map. Fouque and Tibouchi [FT12] also analyze the original construction of Shallue and van de Woestijne specifically for the case of Barreto-Naehrig curves [BN06]; in Section 3, we give a very similar construction tailored to the BLS12-381 curve.

Icart [Ica09] describes a different approach for elliptic curves over fields of characteristic $p \equiv 2 \pmod 3$, for which Farashahi et al. [FSV09] and Fouque and Tibouchi [FT10] give improved analyses. Kammerer et al. [KLR10] generalize this approach to Hessian curves and certain hyperelliptic curves. Farashahi [Far11] also independently gives a map to Hessian curves based on Icart’s approach. Couveignes and Kammerer [CK11] give a further generalization to an infinite family of such maps. Because of their restriction on field characteristic, none of these maps apply to BLS12-381.

Following the work of Farashahi [Far11], Fouque et al. [FJT13] study injective encodings to elliptic curves. Their results apply to curves $E(\mathbb{F}_p), p \equiv 3 \pmod 4$ with $4 \mid \#E(\mathbb{F}_p)$. Bernstein et al. [BHKL13] extend this work, giving a related encoding to all curves over fields of odd characteristic having a point of order 2. All of these restrictions rule out application to BLS12-381.

Hash functions to curves as random oracles. Brier et al. [BCI⁺10] study hash functions to elliptic curves in the indifferentiability framework of Maurer et al. [MRH04].¹ The authors build two indifferentiable hash functions by composing random oracles $H : \{0, 1\}^* \rightarrow \mathbb{F}$ with deterministic maps $\mathbb{F} \rightarrow E(\mathbb{F})$. Their first construction is tailored to Icart’s map; the second applies to essentially all known deterministic maps, but is roughly five times more costly to evaluate. We describe and evaluate both constructions in Sections 5 and 6.

¹Informally, an implementation is indifferentiable from an ideal primitive if no PPT Turing machine can distinguish between the two with non-negligible probability, even with access to internals of the implementation. This generalization of indistinguishability sidesteps certain impossibility results.

More recently, Farashahi et al. [FFS⁺13] give a new analysis showing that the more efficient hash construction of Brier et al. is indifferentiable from a random oracle for essentially any deterministic map. Fouque and Tibouchi [FT12] give a different analysis for a particular version of the map of Shallue and van de Woestijne tailored to Barreto-Naehrig curves. Both of the above analyses apply to the maps of Sections 3 and 4.

Kim and Tibouchi [KT15] further improve the analyses of Brier et al. and Farashahi et al. Notably, they refine the analysis of the costlier of the two maps of Brier et al., adding a degree of freedom that allows implementers to trade off cost and security. We describe and evaluate hashes based on this construction in Sections 5 and 6. They also show that replacing the random oracle H in the constructions of Brier et al. with $\hat{H} : \{0, 1\}^* \rightarrow \{0, 1\}^{\lceil \log \# \mathbb{F} \rceil}$ preserves indifferentiability with little security loss; this improves efficiency.

Fast cofactor multiplication when hashing to G_2 . Barreto-Lynn-Scott curves define two groups G_1 and G_2 , where $\#G_1 = \#G_2$, G_1 is a subgroup of $E_1(\mathbb{F}_p)$, and G_2 is a subgroup of $E_2(\mathbb{F}_{p^{2k}})$ ($k = 1$ for BLS12-381). Hasse’s theorem [Has] dictates that the cofactor is much larger for $E_2(\mathbb{F}_{p^{2k}})$ than for $E_1(\mathbb{F}_p)$, so exponentiating by this cofactor to obtain an element of G_2 is much costlier than the corresponding computation for G_1 .

Scott et al. [SBC⁺09] show how to reduce this cost by exploiting efficiently-computable endomorphisms, building on prior methods [GLV01, IMCT02, GS08]. Fuentes-Castañeda et al. [FKR12] describe another approach exploiting the same endomorphism, reducing costs for some curves. Budroni and Pintore [BP17] study both methods for Barreto-Lynn-Scott curves and give explicit constructions, which all of our hashes to G_2 (§5) use.

2 Background

Notation. We write $E(\mathbb{F})$ for the group (in multiplicative notation) of rational points on elliptic curve E over field \mathbb{F} of order $\#E(\mathbb{F})$. $a || b$ is the concatenation of a and b .

$\chi(\cdot)$ is a quadratic character. For $\alpha \in \mathbb{F}_p$, this is the Legendre symbol, which can be computed as $\alpha^{(p-1)/2}$. For $\delta \in \mathbb{F}_{p^2}$, this can be computed as $\delta^{(p^2-1)/2} = (\delta^p \delta)^{(p-1)/2}$, i.e., the Legendre symbol of the norm $\|\delta\|$. For any field \mathbb{F} , $\forall \rho, \tau \in \mathbb{F}$, $\chi(\rho \cdot \tau) = \chi(\rho) \cdot \chi(\tau)$.

$\text{Sgn}_0(\beta)$ is a function that returns the “sign” of β . For $\beta \in \mathbb{F}_p$, let $\text{Sgn}_0(\beta) = -1$ if $\beta > (p-1)/2$, and 1 otherwise. For $\gamma = \gamma_0 + \gamma_1 \sqrt{d} \in \mathbb{F}_{p^2} \triangleq \mathbb{F}_p[\sqrt{d}]/(x^2 - d)$, let $\text{Sgn}_0(\gamma) = \text{Sgn}_0(\gamma_1)$ if $\gamma_1 \neq 0$, and $\text{Sgn}_0(\gamma_0)$ otherwise.

We regard the square root in \mathbb{F} as a function, so we fix a canonical representation. For $\mathbb{F}_p, p \equiv 3 \pmod{4}$, $\sqrt{\alpha} \triangleq \alpha^{(p+1)/4} \in \mathbb{F}_p$. Otherwise, $\beta \triangleq \sqrt{\alpha} \in \mathbb{F}$ such that $\text{Sgn}_0(\beta) = 1$.

Jacobian coordinates. It is often convenient to represent points (x, y) on $E(\mathbb{F})$ in Jacobian projective coordinates $(X : Y : Z)$, which are defined as follows:

$$(x, y) \mapsto (x : y : 1) \quad (X : Y : Z) \mapsto (X/Z^2, Y/Z^3)$$

Projective coordinates are generally useful to avoid computing inversions, and Jacobian coordinates give some of the fastest group operations for this curve shape [EFD]. Moreover, the point addition law in this representation is independent of the coefficients of the curve equation. Looking ahead, this fact will be useful for SWU-based hash functions (§4, §5).

2.1 The BLS12-381 elliptic curve

Barreto-Lynn-Scott curves [BLS03] are a pairing-friendly family with j -invariant = 0, parameterized by a value $z \equiv 1 \pmod{3}$ (see also [FST10, AFK⁺13]). BLS12-381 [Bow17] is a member of this family with parameter $z = -0\text{xd}20100000010000$ that defines the following bilinear group pair G_1, G_2 [BLS01, Definition 2.2]:

G_1 is the order- q subgroup of $E_1(\mathbb{F}_p) : y^2 = x^3 + 4$, $\#E_1(\mathbb{F}_p) = h_1q$, where

$$\begin{aligned} p &= 0x1a0111ea397fe69a4b1ba7b6434bacd764774b84f38512bf6730d2a0f6b0f624 \\ &\quad 1eabfffeb153ffffb9feffffffffffaaab = z + (z^4 - z^2 + 1)(z - 1)^2/3 \\ q &= 0x73eda753299d7d483339d80809a1d80553bda402fffe5bfeffffffffff00000001 \\ &= z^4 - z^2 + 1 \\ h_1 &= 0x396c8c005555e1568c00aaab0000aaab = (z - 1)^2/3 \end{aligned}$$

G_2 is the order- q subgroup of $E_2(\mathbb{F}_{p^2}) : y^2 = x^3 + 4(1 + \sqrt{-1})$, where $\mathbb{F}_{p^2} \triangleq \mathbb{F}_p[\sqrt{-1}]/(x^2 + 1)$. $\#E_2(\mathbb{F}_{p^2}) = h_2q$, where

$$\begin{aligned} h_2 &= 0x5d543a95414e7f1091d50792876a202cd91de4547085abaa68a205b2e5a7ddf a \\ &\quad 628f1cb4d9e82ef21537e293a6691ae1616ec6e786f0c70cf1c38e31c7238e5 \end{aligned}$$

2.2 BLS signatures

This description follows the one due to Boneh et al. [BLS01]; it uses the following primitives:

- G_1 and G_2 are cyclic groups of prime order q with generators g_1 and g_2 , respectively.
- $\psi_t : G_2 \rightarrow G_1$ is an efficiently computable isomorphism with $\psi_t(g_2) = g_1$. If such a ψ_t does not exist, a stronger assumption is necessary for security [BDN18].
- $e : G_1 \times G_2 \rightarrow G_T$ is a bilinear map.
- The groups G_1, G_2 and the maps ψ_t, e comprise a bilinear group pair.
- $H : \{0, 1\}^* \rightarrow G_1$ is a hash function modeled as a random oracle.

The BLS signature scheme is the triple of algorithms ($\text{gen}^{\text{BLS}}, \text{sign}^{\text{BLS}}, \text{verify}^{\text{BLS}}$) defined as

$\text{gen}^{\text{BLS}}() \rightarrow (pk, sk) : \text{Sample } x \xleftarrow{\mathbb{R}} \{0, \dots, q - 1\}$, compute $g_2^x \in G_2$, and output (g_2^x, x) .

$\text{sign}^{\text{BLS}}(sk, msg) \rightarrow sig : \text{Output } H(msg)^{sk} \in G_1$.

$\text{verify}^{\text{BLS}}(pk, msg, sig) \rightarrow \{\text{OK}, \perp\} : \text{OK if } e(H(msg), pk) = e(sig, g_2)$, else \perp .

It is also possible to swap the functions of G_1 and G_2 in the above. This reduces the size of public keys at the cost of longer signatures; in practice, when using aggregatable signatures this tradeoff may be desirable. In this case, signing and verifying requires a hash function $H : \{0, 1\}^* \rightarrow G_2$. This work considers hashing to both G_1 and G_2 .

2.3 The Shallue–van de Woestijne map

For any elliptic curve $E(\mathbb{F}) : y^2 = f(x) = x^3 + ax + b$, $\#\mathbb{F} > 5$, Shallue and van de Woestijne give a map from $L \subseteq \mathbb{F}$ to the curve $E(\mathbb{F})$ [SvdW06]. They observe, generalizing and simplifying the result of Skałba [Ska05], that for any rational point on the threefold

$$V(\mathbb{F}) : f(x_1)f(x_2)f(x_3) = x_4^2$$

such that $x_4 \neq 0$, at least one of $f(x_j), j \in \{1, 2, 3\}$ must be a square. This implies that one of the x_j is the x -coordinate of a rational point on $E(\mathbb{F})$.

To construct a rational point on $V(\mathbb{F})$, the authors define the surface $S(\mathbb{F})$ and the rational map $\phi_1 : S(\mathbb{F}) \mapsto V(\mathbb{F})$, which is invertible on its image [SvdW06, Lemma 6]:

$$S(\mathbb{F}) : y^2 (u^2 + uv + v^2 + a) = -f(u)$$

$$\phi_1 : (u, v, y) \mapsto \left(v, -u - v, u + y^2, f(u + y^2) \cdot \frac{y^2 + uv + v^2 + a}{y} \right).$$

Next, the authors observe [SvdW06, Lemma 7] that fixing $u = u_0$ satisfying $f(u_0) \neq 0$ and $3u_0^2 + 4a \neq 0$ specializes $S(\mathbb{F})$ to a curve that is birational to a conic with a rational parameterization. This gives a rational map $\phi_2 : \mathbb{A}^1 \mapsto S(\mathbb{F})$ that is invertible on its image.

Putting it all together, define $L = \{t \in \mathbb{F} : \phi_1(\phi_2(t)) \text{ is defined}\}$. Then, to map $t \in L$ to $E(\mathbb{F})$, first compute $\phi_1(\phi_2(t))$, which is a rational point (x_1, x_2, x_3, x_4) on $V(\mathbb{F})$, so at least one $f(x_j), j \in \{1, 2, 3\}$ is square. Choose the smallest j where this is the case, compute the corresponding y -coordinate, and return (x_j, y) .

2.4 The simplified Shallue–van de Woestijne–Ulas map of Brier et al.

Brier et al. [BCI⁺10] define the simplified SWU map (with a small modification due to Fouque and Tibouchi [FT10]) as follows. Let $E(\mathbb{F}_p) : y^2 = g(x) = x^3 + ax + b, ab \neq 0, p \equiv 3 \pmod 4$, and

$$X_0(t) = -\frac{b}{a} \left(1 + \frac{1}{t^4 - t^2} \right) \quad X_1(t) = -t^2 X_0(t)$$

Then the simplified SWU map is given by

$$t \mapsto \begin{cases} \infty & \text{if } t \in \{-1, 0, 1\} \\ \left(X_0(t), \sqrt{g(X_0(t))} \right) & \text{if } \chi(g(X_0(t))) = 1 \\ \left(X_1(t), -\sqrt{g(X_1(t))} \right) & \text{otherwise} \end{cases}$$

To see why this map works, assume u is non-square and assume we have x such that $g(ux) = u^3 g(x)$. Since u is non-square, either $g(x)$ or $g(ux)$ is square, and thus either x or ux is the x -coordinate of a rational point on $E(\mathbb{F}_p)$. Expanding and solving for x ,

$$\begin{aligned} u^3 x^3 + aux + b &= u^3(x^3 + ax + b) \\ aux + b &= u^3(ax + b) \\ x &= -\frac{b}{a} \cdot \frac{u^3 - 1}{u^3 - u} = -\frac{b}{a} \left(\frac{u^3 - u - 1}{u^3 - u} + \frac{u}{u^3 - u} \right) = -\frac{b}{a} \left(1 + \frac{u - 1}{u(u^2 - 1)} \right) \\ &= -\frac{b}{a} \left(1 + \frac{1}{u^2 + u} \right) \end{aligned}$$

Since $p \equiv 3 \pmod 4$, -1 is non-square, so $u = -t^2$ is, too. Substituting yields X_0 and X_1 .

3 A Shallue–van de Woestijne map for BLS12-381

We construct an explicit Shallue–van de Woestijne map (§2.3) for the BLS12-381 curves $E_1(\mathbb{F}_p)$ and $E_2(\mathbb{F}_{p^2})$ (§2.1). Our description follows the one by Fouque and Tibouchi [FT12].

For both of the BLS12-381 curves, we have $E(\mathbb{F}) : y^2 = f_i(x) = x^3 + b_i$ (§2.1), so we restrict our analysis to the case $a = 0$. For now, we work with $S(\mathbb{F})$ generically in terms of $u = u_0$; we discuss convenient choices of u_0 below. Rewriting [SvdW06, Lemma 7]:

$$\begin{aligned} y^2 \left(\frac{3}{4}u_0^2 + \left(v + \frac{u_0}{2} \right)^2 \right) &= -f_i(u_0) \\ z^2 + f_i(u_0)w^2 &= -\frac{3}{4}u_0^2 \quad \text{where } z = v + \frac{u_0}{2}, \quad w = \frac{1}{y} \end{aligned}$$

-3 is square in \mathbb{F}_p (and thus in \mathbb{F}_{p^2}), so $(z_0, w_0) \triangleq (\sqrt{-3u_0^2/4}, 0)$ is a rational point on this conic. Parameterizing in t by setting $z = z_0 + tw$ and substituting gives

$$t\sqrt{-3u_0^2} + (t^2 + f_i(u_0))w = 0 \quad w \neq 0$$

Solving for y and v ,

$$y = \frac{1}{w} = -\frac{t^2 + f_i(u_0)}{t\sqrt{-3u_0^2}}$$

$$v = z_0 + tw - \frac{u_0}{2} = \frac{\sqrt{-3u_0^2} - u_0}{2} - \frac{t^2\sqrt{-3u_0^2}}{t^2 + f_i(u_0)}$$

Finally, from the map ϕ_1 (§2.3), we have

$$x_1 = v = \frac{\sqrt{-3u_0^2} - u_0}{2} - \frac{t^2\sqrt{-3u_0^2}}{t^2 + f_i(u_0)}$$

$$x_2 = -u_0 - v = \frac{t^2\sqrt{-3u_0^2}}{t^2 + f_i(u_0)} - \frac{\sqrt{-3u_0^2} + u_0}{2}$$

$$x_3 = u_0 + y^2 = u_0 - \frac{(t^2 + f_i(u_0))^2}{3u_0^2 t^2}$$

This map is undefined when $t = 0$ or $t^2 + f_i(u_0) = 0$. Fouque and Tibouchi [FT12] choose u_0 such that $-f_i(u_0)$ is non-square (ensuring $t^2 + f_i(u_0) \neq 0$) and add a special case for $t = 0$. We take a slightly different approach in order to avoid an explicit special case.

First, notice that, once we have fixed u_0 , we can evaluate the map using only one inversion by applying Montgomery’s trick [Mon87], i.e., inverting the product $t^2(t^2 + f_i(u_0))$. Evaluating the x_j then entails a few inexpensive arithmetic operations involving t^2 and precomputed constants. Computing Legendre symbols and a square root yields y .

Now we can handle the exceptional cases. Notice that when applying Montgomery’s trick, the map is undefined just when the value being inverted is zero. If we use an inversion algorithm that returns zero on input zero (which is true, e.g., for inversion by Fermat’s little theorem), the resulting value of x_1 will be $x_0 \triangleq (\sqrt{-3u_0^2} - u_0)/2$. Choosing u_0 such that $f_i(x_0)$ is square then yields an exception-free map. For $E_1(\mathbb{F}_p)$, the smallest u_0 in absolute value for which this holds is $u_0 = -3$; for $E_2(\mathbb{F}_{p^2})$, it is $u_0 = -1$.

Finally, Fouque and Tibouchi observe that the x_j values depend only on t^2 , i.e., t and $-t$ map to the same x -coordinate. They suggest choosing the sign of y to be the sign of $\chi(t)$, but this costs an extra Legendre symbol evaluation and does not work in \mathbb{F}_{p^2} (because -1 is square). A more easily computed choice that works for both \mathbb{F}_p and \mathbb{F}_{p^2} is $y = \text{Sgn}_0(t) \cdot \sqrt{f_i(x_j)}$. This fully specifies both maps.

Putting it all together. Precompute $f_i(u_0)$, $(\sqrt{-3u_0^2} \pm u_0)/2$, $\sqrt{-3u_0^2}$, and $1/3u_0^2$. On input t , compute $\alpha = 1/(t^2(t^2 + f_i(u_0)))$ setting $\alpha = 0$ if $t^2(t^2 + f_i(u_0)) = 0$. Use α to compute x_1 , x_2 , and x_3 (e.g., $x_1 = (\sqrt{-3u_0^2} - u_0)/2 - \alpha t^4 \sqrt{-3u_0^2}$). Choose the smallest $j \in \{1, 2, 3\}$ such that $\chi(f_i(x_j)) = 1$, compute $y = \text{Sgn}_0(t) \cdot \sqrt{f_i(x_j)}$, and return (x_j, y) .

Computing this map in constant time requires evaluating all x_j and all $\chi(f_i(x_j))$, all in constant time. Note that fast Legendre symbol and inversion algorithms [Coh93, §1.3.2, §1.4.2] are *not* constant time. Fouque and Tibouchi suggest blinding the Legendre symbol by choosing random r_j and computing $\chi(r_j^2 f_i(x_j))$ [FT12, §6]; a similar trick for inversion is standard (to invert β , choose random r , invert $r\beta$ and then multiply by r). Of course, computing Legendre symbols or inversions by exponentiation is easily made constant time, but is also far more costly. Our constant-time implementations using only field operations take an approach that we describe in the Section 4; we discuss specifics in Section 6.

4 An optimized SWU map for BLS12-381

The simplified SWU map of Brier et al. [BCI⁺10] (§2.4) applies only to curves of the form $E(\mathbb{F}) : y^2 = g(x) = x^3 + ax + b$ where $ab \neq 0$ and $\#\mathbb{F} \equiv 3 \pmod{4}$. For BLS12-381 (§2.1), $E_1(\mathbb{F}_p)$ meets the second requirement but not the first; $E_2(\mathbb{F}_{p^2})$ meets neither requirement. As a result, this map cannot be applied directly to either curve.

In this section we show how to solve these problems. To avoid the requirement that $\#\mathbb{F} \equiv 3 \pmod{4}$, we tweak the map’s definition. To sidestep the issue that $a = 0$ for both curves, we construct an “indirect” map, with two steps: first, map to some $E'(\mathbb{F})$ isogenous to $E(\mathbb{F})$ with $ab \neq 0$, and second, apply the isogeny map to obtain a point on $E(\mathbb{F})$.

One potential issue with the indirect approach is that, for an isogeny of degree d , the resulting map is to the points $\{P^d : P \in E(\mathbb{F})\}$ —that is, it maps only to a *subset* of $E(\mathbb{F})$. Our concern is twofold: first, recall from Section 2.2 that our goal is to hash to a subgroup of $E(\mathbb{F})$, in particular the order- q subgroup whose elements are $\{P^h : P \in E(\mathbb{F})\}$ for $\#\mathbb{F} = hq$, so we must ensure that this subgroup is in the map’s image. Second, in Section 5 we will construct hash functions that rely on the statistical properties of the SWU map, so we must ensure that the indirect approach does not alter those properties.

Fortunately, both concerns are easily dispensed with. For the first, we choose an isogeny of degree d coprime to q ; exponentiation by h yields $\{P^{hd} : P \in E(\mathbb{F})\} \simeq \{P^h : P \in E(\mathbb{F})\}$. For the second, since choosing d in this way induces the above isomorphism, it suffices to show that exponentiation by h preserves the relevant statistical properties. Boneh and Franklin [BF03, §5.2] (and, in the same vein, Brier et al. [BCI⁺10, §6.1]) show that this is the case as long as $q \nmid h$, which is true for BLS12-381.

In the remainder of this section, we generalize the SWU map to $\#\mathbb{F} \not\equiv 3 \pmod{4}$ (§4.1); describe two optimizations that make the map simpler to implement and faster to evaluate (§4.2); and give explicit curves $E'_1(\mathbb{F}_p)$ 11-isogenous to $E_1(\mathbb{F}_p)$ and $E'_2(\mathbb{F}_{p^2})$ 3-isogenous to $E_2(\mathbb{F}_{p^2})$, plus a small hint for evaluating the isogeny maps quickly (§4.3). Finally, we put all of the above together, yielding the SWU maps to $E_1(\mathbb{F}_p)$ and $E_2(\mathbb{F}_{p^2})$ (§4.4).

4.1 Generalizing the map to $E(\mathbb{F})$, $\#\mathbb{F} \not\equiv 3 \pmod{4}$

Recall from Section 2.4 that for an elliptic curve $E'(\mathbb{F}) : y^2 = g(x) = x^3 + ax + b$, the SWU map uses a parameterization of x in terms of u such that $g(ux) = u^3g(x)$, for u a non-square. When $\#\mathbb{F} \equiv 3 \pmod{4}$ (as in §2.4), -1 is non-square, so $u = -t^2$ is a convenient choice. More generally, for some non-square $\xi \in \mathbb{F}$, let $u = \xi t^2$. Then we have

$$g(X_0(t)) \cdot g(X_1(t)) = (g(X_0(t)))^2 \xi^3 t^6 = (t^3 g(X_0(t)))^2 \xi^3$$

Since ξ^3 —and thus the right-hand side—is non-square, exactly one of $g(X_0(t))$ and $g(X_1(t))$ must be square, and thus either $X_0(t)$ or $X_1(t) \triangleq \xi t^2 X_0(t)$ is the x -coordinate of a rational point on $E'(\mathbb{F})$. In the next section, it will be convenient for $g(b/\xi a)$ to be square. For $E'_1(\mathbb{F}_p)$, $\xi_1 \triangleq -1 \in \mathbb{F}_p$ satisfies this requirement; for $E'_2(\mathbb{F}_{p^2})$, $\xi_2 \triangleq 1 + \sqrt{-1} \in \mathbb{F}_{p^2}$ does.

4.2 Optimizing the map

As described in Section 2.4, the SWU map requires computing a field inversion, a quadratic character, and a square root. We now describe how to avoid computing both the quadratic character and the inversion, in a way that is amenable to constant-time implementation. We describe each optimization for $E'_1(\mathbb{F}_p)$, then show how they apply to $E'_2(\mathbb{F}_{p^2})$.

Notation. In this section we use the generic $g(\cdot)$, a , b , and ξ in expressions that apply to both $E'_1(\mathbb{F}_p)$ and $E'_2(\mathbb{F}_{p^2})$. These correspond to the curve equation $y^2 = g(x) = x^3 + ax + b$ and the map’s distinguished non-square. We give the ξ_i , a_i , and b_i in Sections 4.1 and 4.3.

Eliminating the quadratic character computation. For $E'_1(\mathbb{F}_p)$, $p \equiv 3 \pmod{4}$, so one can evaluate the y -coordinate as follows: compute $\alpha = g_1(X_0(t))^{(p+1)/4}$. If $g_1(X_0(t))$ is square, α is its square root; this is easily checked by comparing α^2 with $g_1(X_0(t))$. If they are equal, α is the y -coordinate (up to sign; §4.4). Otherwise, compute the y -coordinate as $\sqrt{g_1(X_1(t))} = t^3\alpha$. To see why this works, assume $g_1(X_0(t))$ is non-square in \mathbb{F}_p ; then

$$\left(t^3 \cdot g_1(X_0(t))^{(p+1)/4}\right)^2 = t^6 \cdot g_1(X_0(t)) \cdot g_1(X_0(t))^{(p-1)/2} = -t^6 \cdot g_1(X_0(t))$$

The final equality comes from the fact that $g_1(X_0(t))^{(p-1)/2} = \chi(g_1(X_0(t))) = -1$. Recall (§2.4) that $g(ux) = u^3g(x)$, $X_1(t) = uX_0(t)$, and $u = -t^2$, establishing the claim.

Avoiding inversions. We borrow and adapt a trick from Bernstein et al. [BDL⁺12, §5] that merges an inversion and a square root computation into a single exponentiation:

$$(U/V)^{(p+1)/4} = U^{(p+1)/4} \cdot V^{p-1-(p+1)/4} = U^{(p+1)/4} \cdot V^{(3p-5)/4} = UV (UV^3)^{(p-3)/4}$$

We can rewrite $g(X_0(t))$ into the required form as follows:

$$X_0(t) \triangleq \frac{N}{D} = \frac{b(\xi^2 t^4 + \xi t^2 + 1)}{-a(\xi^2 t^4 + \xi t^2)} \quad g(X_0(t)) \triangleq \frac{U}{V} = \frac{N^3 + aND^2 + bD^3}{D^3}$$

The above is undefined just when $D = 0$. To facilitate constant-time evaluation in this case, set N and D to predefined “good” values and continue. Recall (§4.1) that we chose ξ such that $g(b/\xi a)$ is square, and notice that if $D = 0$, then $N = b$ (since $\xi^2 t^4 + \xi t^2 = 0$). Thus, all that is required to recover from the exception is to set $D = \xi a$.

Finally, the x -coordinate is either N/D or $\xi t^2 N/D$. To avoid inverting D , return a point in Jacobian projective coordinates (§2); we give specifics in Section 4.4.

The $E'_2(\mathbb{F}_{p^2})$ case is only slightly more complicated. In particular, we need to show both how to take one square root in \mathbb{F}_{p^2} in constant time, and how to avoid taking a second. For $p^2 \equiv 9 \pmod{16}$, we recall a standard trick for computing $\sqrt{\delta} \in \mathbb{F}_{p^2}$. For square δ , define

$$\gamma = \delta^{(p^2+7)/16} = \left(\delta \cdot \delta^{(p^2-1)/8}\right)^{1/2} = \delta^{1/2} \left(\delta^{(p^2-1)/2}\right)^{1/8} = \delta^{1/2} \cdot 1^{1/8}$$

The final equality comes from the fact that $\delta^{(p^2-1)/2} = \chi(\delta) = 1$. This implies that $\sqrt{\delta}$ must be one of $\pm\gamma$, $\pm\gamma\sqrt{-1}$, or $\pm\gamma\sqrt{\pm\sqrt{-1}}$ (all exist). Exponentiating and checking in constant time is straightforward if the three constants are precomputed.²

This gives us an analogous approach to the $E'_1(\mathbb{F}_p)$ case: compute $\gamma = g_2(X_0(t))^{(p^2+7)/16}$ and check the four possible square roots. If none is found, $g_2(X_0(t))$ is non-square, and one of the four possible values $t^3\gamma\eta = \sqrt{g_2(X_1(t))}$, where $\eta^2 = \xi_2^3(-1)^{-1/4}$. This is because

$$\left(\eta t^3 \cdot g_2(X_0(t))^{(p^2+7)/16}\right)^2 = \eta^2 t^6 \cdot g_2(X_0(t)) \cdot g_2(X_0(t))^{(p-1)/8} = \eta^2 t^6 \cdot g_2(X_0(t)) \cdot (-1)^{1/4}$$

Recalling (§4.1) that $g(X_1(t)) = \xi^3 t^6 g(X_0(t))$, we have $\eta^2 = \xi_2^3(-1)^{-1/4}$ as claimed. Note that $(-1)^{-1/4}$ and ξ_2^3 are non-square in \mathbb{F}_{p^2} , so four unique values of η exist (along with their negations). Once again, for efficiency these values should be precomputed.

Finally, this requires computing $\gamma = (U/V)^{(p^2+7)/16}$ for U and V defined above. In this case, the trick of Bernstein et al. yields $\gamma = UV^7 (UV^{15})^{(p^2-9)/16}$, avoiding an inversion.

²Adj and Rodríguez-Henríquez describe a different algorithm for computing square roots in \mathbb{F}_{p^2} , $p \equiv 3 \pmod{4}$ [AR12], which essentially always requires two exponentiations in \mathbb{F}_{p^2} with exponents of size $\approx \log p$. The method we describe uses one exponentiation in \mathbb{F}_{p^2} with an exponent of size $\approx 2\log p$. The costs are almost the same; the method we describe is slightly easier to implement in constant time.

4.3 The isogeny maps

Recall from above that we want an isogeny of degree d coprime to q . Small d is best for efficiency. For $E_1(\mathbb{F}_p)$, the smallest prime d giving a curve in the base field for which $a \neq 0$ is 11; thus, we use the 11-isogenous curve³ $E'_1(\mathbb{F}_p) : y^2 = x^3 + a_1x + b_1$ where

$$a_1 = 0x144698a3b8e9433d693a02c96d4982b0ea985383ee66a8d8e8981aefd881ac98936f8da0e0f97f5cf428082d584c1d$$

$$b_1 = 0x12e2908d11688030018b12e8753eee3b2016c1f0f24f4070a0b9c14fcef35ef55a23215a316ceaa5d1cc48e98e172be0$$

The 11-isogeny from $E'_1(\mathbb{F}_p)$ to $E_1(\mathbb{F}_p)$ is given by the following map due to Vélu [Vél71]:

$$(x, y) \mapsto \left(\frac{\sum_{i=0}^{11} k_{1,i} x^i}{\sum_{i=0}^{10} k_{2,i} x^i}, y \frac{\sum_{i=0}^{15} k_{3,i} x^i}{\sum_{i=0}^{15} k_{4,i} x^i} \right)$$

The $E_2(\mathbb{F}_{p^2})$ case is much nicer: there is a 3-isogenous curve $E'_2(\mathbb{F}_{p^2}) : y^2 = x^3 + a_2x + b_2$ where $a_2 = 240\sqrt{-1}$ and $b_2 = 1012 + 1012\sqrt{-1}$. The 3-isogeny from $E'_2(\mathbb{F}_{p^2})$ to $E_2(\mathbb{F}_{p^2})$ is given by a map similar to the map for $E'_1(\mathbb{F}_p)$, but with polynomials of degree at most 3.

Appendix A gives Sage [SM] scripts for constructing the isogeny maps.

Avoiding inversions. Above, we mentioned that returning a point in Jacobian projective coordinates saves an inversion when computing the SWU map. Fortunately, it is easy to evaluate the isogeny maps on points in Jacobian coordinates without an inversion. For example, given $(X : Y : Z)$ on $E'_1(\mathbb{F}_p)$, where $x = X/Z^2$ and $y = Y/Z^3$, we evaluate the isogeny map to give a point $(X_o : Y_o : Z_o)$ in Jacobian coordinates on $E_1(\mathbb{F}_p)$, as follows.

Rewriting the x -coordinate map given above in terms of the projective coordinates:

$$\frac{X}{Z^2} \mapsto \frac{N_x}{D_x} \triangleq \frac{\sum_{i=0}^{11} k_{1,i} \left(\frac{X}{Z^2}\right)^i}{\sum_{i=0}^{10} k_{2,i} \left(\frac{X}{Z^2}\right)^i} = \frac{\sum_{i=0}^{11} k_{1,i} (Z^2)^{11-i} X^i}{Z^2 \sum_{i=0}^{10} k_{2,i} (Z^2)^{10-i} X^i}$$

Similarly, for the y -coordinate map:

$$\frac{Y}{Z^3} \mapsto \frac{N_y}{D_y} \triangleq \frac{Y \sum_{i=0}^{15} k_{3,i} \left(\frac{X}{Z^2}\right)^i}{\sum_{i=0}^{15} k_{4,i} \left(\frac{X}{Z^2}\right)^i} = \frac{Y \sum_{i=0}^{15} k_{3,i} (Z^2)^{15-i} X^i}{Z^3 \sum_{i=0}^{15} k_{4,i} (Z^2)^{15-i} X^i}$$

To evaluate the above maps, first compute $Z^{2i}, i \in \{1, \dots, 15\}$. Then, to evaluate, for example, the numerator of the X map, compute the products $k_{1,i} (Z^2)^{11-i}, i \in \{1, \dots, 11\}$ and then evaluate the polynomial using Horner's method [Knu97].

Finally, after evaluating the numerator and denominator of each map, compute

$$Z_o = D_x D_y \quad X_o = N_x D_y Z_o \quad Y_o = N_y D_x Z_o^2$$

Then $y_o = Y_o/Z_o^3 = N_y/D_y$ and $x_o = X_o/Z_o^2 = N_x/D_x$, as required.

³Alternatively, we might map to a curve isogenous to E_1 over an extension field in which there are suitable isogenies of lower degree, then map to $E_1(\mathbb{F}_p)$ via the trace map (see, e.g., [GR04, §3]). But this entails costly arithmetic in the extension—and in particular, a much costlier square-root computation—so it is almost certainly much more expensive than using the 11-isogenous curve (e.g., see Tables 2 and 3, §6).

4.4 Putting it all together

We slightly modify the SWU map of Section 2.4. In particular, since $X_0(t)$ and $X_1(t)$ only depend on t^2 , we can arbitrarily choose the sign of the resulting y -coordinate; we set $y = \text{Sgn}_0(t) \cdot \sqrt{g(X_j(t))}$, which we justified in Section 3.

Thus, the map to a point on $E_1(\mathbb{F}_p)$ is computed as follows on input $t \in \mathbb{F}_p$:

- (1) Compute N, D, U, V (§4.2) and $\alpha \triangleq UV (UV^3)^{(p-3)/4}$.
- (2) If $\alpha^2 V - U = 0$, then $g_1(X_0(t))$ is square in \mathbb{F}_p , so $(x, y) = (N/D, \text{Sgn}_0(t) \cdot \alpha) \in E'_1(\mathbb{F}_p)$.
Set $(X : Y : Z) = (ND : \text{Sgn}_0(t) \cdot \alpha D^3 : D)$.
- (3) Otherwise, $g_1(X_1(t))$ is square in \mathbb{F}_p , so $(x, y) = (\xi t^2 N/D, t^3 \alpha) \in E'_1(\mathbb{F}_p)$. (No $\text{Sgn}_0(\cdot)$ is necessary in this case, because multiplication by t^3 preserves the sign of t .)
Set $(X : Y : Z) = (\xi t^2 ND : t^3 \alpha D^3 : D)$.
- (4) Evaluate the 11-isogeny map on $(X : Y : Z)$ (§4.3) and return the resulting point.

The map to $E_2(\mathbb{F}_{p^2})$ is analogous.

As a final optimization, if one is computing this map multiple times and summing the results (§5), one can avoid repeatedly evaluating the isogeny map by summing the points on $E'_1(\mathbb{F}_p)$ or $E'_2(\mathbb{F}_{p^2})$ and then applying the isogeny map to the sum. This optimization comes essentially for free: the point addition law in Jacobian coordinates is independent of the coefficients in the curve equation (§2), meaning that one can use the same point addition routine for $E_1(\mathbb{F}_p)$ and $E'_1(\mathbb{F}_p)$, and similarly for $E_2(\mathbb{F}_{p^2})$ and $E'_2(\mathbb{F}_{p^2})$ [EFD].

5 Hashing to the groups G_1 and G_2 of BLS12-381

We describe six hash functions that use the maps of Sections 3 and 4 to output elements of G_1 or G_2 (§2.1). All but the first and fourth are based on the work of Brier et al. [BCI⁺10], Farashahi et al. [FFS⁺13], and Kim and Tibouchi [KT15], who show that they are indifferentiable from a random oracle (§1.1). The hashes use the following primitives:⁴

- The functions $H_p : \{0, 1\}^* \rightarrow \mathbb{F}_p$, $H_{p^2} : \{0, 1\}^* \rightarrow \mathbb{F}_{p^2}$, and $H_{128} : \{0, 1\}^* \rightarrow \{0, 1\}^{128}$ are hashes that we model as random oracles.
- $\text{Map}_1 : \mathbb{F}_p \rightarrow E_1(\mathbb{F}_p)$ and $\text{Map}_2 : \mathbb{F}_{p^2} \rightarrow E_2(\mathbb{F}_{p^2})$ can be either of the maps given in Sections 3 and 4.
- z is the Barreto-Lynn-Scott parameter for the BLS12-381 curve (§2.1).
- $\Psi : E_2(\mathbb{F}_{p^2}) \rightarrow E_2(\mathbb{F}_{p^2})$ is the endomorphism of Budroni and Pintore [BP17, §4.1], based on the work of Scott et al. [SBC⁺09] and Fuentes-Castañeda et al. [FKR12]. This is effectively a fast method of exponentiating by the cofactor h_2 of $E_2(\mathbb{F}_{p^2})$.

Clearing cofactors. The maps of Sections 3 and 4 output elements of $E_1(\mathbb{F}_p)$ or $E_2(\mathbb{F}_{p^2})$, but we want elements of G_1 or G_2 . For G_2 , we use the endomorphism Ψ (see above). For G_1 , Scott observes [Sco19] that there is a faster way than exponentiating by h_1 . Recall (§2.1) that $h_1 = (z - 1)^2/3$; for BLS12-381, $1 - z = 3 \prod_k p_k$ for primes p_k , i.e., $h_1 = 3 \prod_k p_k^2$. It is easy to check that $E_1(\mathbb{F}_p)$ has no points of order p_k^2 for any k . As a result, exponentiating by $1 - z \approx \sqrt{h_1}$ gives an element of G_1 . This speeds up cofactor clearing by more than $2\times$ versus exponentiating by h_1 , in part because $1 - z$ has low Hamming weight.

⁴For constant-time hash functions, these primitives must be constant time, too. Reducing a $2 \log p$ -bit integer modulo p gives an element of \mathbb{F}_p with negligible bias. Alternatively, Kim and Tibouchi [KT15] show that H_p can be replaced with $H_{\hat{p}} : \{0, 1\}^* \rightarrow \{0, 1\}^{\lceil \log p \rceil}$. H_{p^2} can be implemented via two evaluations of H_p or $H_{\hat{p}}$. All of these are easily built from a PRG (e.g., AES-CTR or ChaCha20 [Ber08]) seeded with a hash of the input. Ψ is a straight-line computation, so it is constant time if field operations are.

Construction #1. The hash function $H_1 : \{0, 1\}^* \rightarrow G_1$ is given by

$$H_1(msg) \triangleq \text{Map}_1(H_p(msg))^{1-z}$$

Since the maps of Sections 3 and 4 are to a subset of $E_1(\mathbb{F}_p)$, and since both maps are invertible on their image, this function is easily distinguished from a random oracle. We note that a random oracle is not necessary for BLS signatures—hashing to a constant fraction of G_1 suffices (see also [BCI⁺10, Section 5.2]). Still, we recommend using one of the other hash functions in this section, for three reasons. First, a random oracle simplifies the BLS security proof. Second, an indifferentiable hash function is suitable for uses beyond signatures (another defense against “mission creep”; §1). And third, as discussed above, an exponentiation is anyway necessary to clear the cofactor; this high fixed cost helps moderate the (relative) overheads of the other hash constructions (§6).

Construction #2. The hash function $H_2 : \{0, 1\}^* \rightarrow G_1$ is given by

$$H_2(msg) \triangleq (\text{Map}_1(H_p(msg || 0)) \cdot \text{Map}_1(H_p(msg || 1)))^{1-z}$$

As mentioned in Section 4.4, when using that section’s map it is most efficient to evaluate the map to $E'_1(\mathbb{F}_p)$ twice, sum the resulting points, and then apply the isogeny map.

Construction #3. Let g_3 be a generator of G_1 with unknown discrete log relation to the base point g_1 of the BLS signature scheme (§2.2). Then $H_3 : \{0, 1\}^* \rightarrow G_1$ is given by

$$H_3(msg) \triangleq \text{Map}_1(H_p(msg))^{1-z} \cdot g_3^{H_{128}(msg)}$$

This is an aggressive choice of parameters: we estimate that this function’s outputs have statistical distance from uniform only $\approx 2^{-56}$ for either choice of Map_1 [KT15, Cor. 1]. Increasing the range of H_{128} improves uniformity at the cost of performance.

For BLS12-381, $1 - z$ is a 64-bit integer, so directly using multi-exponentiation [Möl01] to compute H_3 gives little savings (because computing $g_3^{H_{128}(msg)}$ requires about twice as many squarings as computing $\text{Map}_1(H_p(msg))^{1-z}$). This can be addressed straightforwardly as follows. First, compute r_1 and r_2 at most 64 bits such that $H_{128}(msg) = r_2 \cdot 2^{64} + r_1$. Then, letting $g_{3,64} = g_3^{2^{64}}$ (which can be precomputed, because g_3 is fixed),

$$H_3(msg) \triangleq \text{Map}_1(H_p(msg))^{1-z} \cdot g_3^{r_1} \cdot g_{3,64}^{r_2}$$

This saves squarings because $1 - z$, r_1 , and r_2 are all the same size in bits.

Finally, it is very important that the discrete log relation between the generator g_3 and the BLS base point g_1 is unknown. To see why, assume the discrete log of g_3 base g_1 is ℓ , and recall (§2.2) that for key $(pk, sk) = (g_2^x, x)$, a signature on msg is $H_3(msg)^x$. Then

$$\begin{aligned} H_3(msg)^x &= \text{Map}_1(H_p(msg))^{(1-z) \cdot x} \cdot g_3^{H_{128}(msg) \cdot x} \\ &= \text{Map}_1(H_p(msg))^{(1-z) \cdot x} \cdot ((g_1^\ell)^x)^{H_{128}(msg)} \\ &= \text{Map}_1(H_p(msg))^{(1-z) \cdot x} \cdot (\psi_t(pk)^\ell)^{H_{128}(msg)} \quad (\psi_t \text{ is defined in §2.2}) \end{aligned}$$

Knowing ℓ renders moot the multiplication by $g_3^{H_{128}(msg)}$ in H_3 : since $(\psi_t(pk)^\ell)^{H_{128}(msg)}$ is public, forging a signature for $H_3(msg)$ just requires forging one for $\text{Map}_1(H_p(msg))^{1-z}$.

Construction #4. The hash function $H_4 : \{0, 1\}^* \rightarrow G_2$ is given by

$$H_4(msg) \triangleq \Psi(\text{Map}_2(H_{p^2}(msg)))$$

This is the equivalent of construction #1 for G_2 , and the same caveats apply.

Construction #5. The hash function $H_5 : \{0, 1\}^* \rightarrow G_2$ is given by

$$H_5(msg) \triangleq \Psi \left(\text{Map}_2(H_{p^2}(msg || 0)) \cdot \text{Map}_2(H_{p^2}(msg || 1)) \right)$$

This is the equivalent of construction #2 for G_2 . As in that case, when using the map of Section 4, it is faster to add points on $E'_2(\mathbb{F}_{p^2})$ and then apply the isogeny.

Construction #6 Let g_4 be a generator of G_2 with unknown discrete log relation to the base point g_2 of the BLS signature scheme (§2.2). Then $H_6 : \{0, 1\}^* \rightarrow G_2$ is given by

$$H_6(msg) \triangleq \Psi \left(\text{Map}_2(H_{p^2}(msg)) \right) \cdot g_4^{H_{128}(msg)}$$

This is like construction #3, but for G_2 ; as there, parameter choices are aggressive for performance, and the discrete log relation between g_4 and g_2 must be unknown.

To evaluate H_6 quickly, we integrate the exponentiation of g_4 into the evaluation of Ψ . For $\psi : E_2(\mathbb{F}_{p^2}) \rightarrow E_2(\mathbb{F}_{p^2})$ the “untwist-Frobenius-twist” endomorphism of Galbraith and Scott [GS08, §5], Budroni and Pintore define Ψ [BP17, §4.1] as

$$\Psi(P) \triangleq P^{z^2 - z - 1} \cdot \psi(P)^{z-1} \cdot \psi(\psi(P^2))$$

Letting $P = \text{Map}_2(H_{p^2}(msg))$, $H_{128}(msg) = r_2 \cdot 2^{64} + r_1$, and $g_{4,64} = g_4^{2^{64}}$,

$$H_6(msg) \triangleq (P^{z-1} \cdot \psi(P))^z \cdot g_4^{r_1} \cdot g_{4,64}^{r_2} \cdot P^{-1} \cdot \psi(P)^{-1} \cdot \psi(\psi(P^2))$$

This lets us use a multi-exponentiation to evaluate the product of the leftmost three terms.

6 Implementation and evaluation

We implement and evaluate hash-and-check, plus the hashes of Section 5 using the maps of Sections 3 and 4. We compare performance for three implementation styles of varying complexity. The most complex style uses rich functionality from a full-featured multi-precision library, and is not constant time. The other two styles are simpler: they are restricted to using only field operations, i.e., fixed-modulus arithmetic. They differ in that one is constant time and one is not. We justify our interest in simplicity below.

In sum, we find that our optimizations to the map of Section 4 yield hash functions that are at worst only $\approx 9\%$ slower than the fastest alternatives, yet are considerably simpler to implement. Moreover, when comparing implementations restricted to field operations and constant-time execution, the map of Section 4 is faster by ≈ 1.3 – $2\times$ than the map of Section 3. Our experiments show that this speed advantage is due to our optimizations.

Implementation complexity: why restrict to field operations? We are interested in implementations restricted to field operations because this is the bare minimum functionality required for elliptic curve operations—so these primitives are guaranteed to be available and likely to be highly optimized. This is especially germane in hardware implementations and embedded cryptographic co-processors (e.g., [Gui10, CDF⁺11, SLA]), which are usually restricted to field arithmetic because general multi-precision arithmetic is too expensive in terms of area or energy. Such restrictions are also typical of small software libraries (e.g., [NaC, Pai]); reasons include ease of optimization and constant-time implementation, and a simpler codebase, which generally leads to easier maintenance and fewer bugs.

Special-purpose arithmetic libraries like GMP [GMP] and large cryptographic libraries like Botan [Bota], Crypto++ [Cry], and OpenSSL [Opea] *do* implement full-featured multi-precision arithmetic—but at best only the very basic operations are constant time [Opeb, Botb]. This means that implementations aiming for input-independent runtime that use

rich functionality from these libraries must resort to techniques like blinding [Koc96, FT12, Bos14]. But this is no silver bullet: such techniques increase complexity, require a good source of randomness, and must be carefully analyzed to ensure that all leaks are plugged.

Still, it is reasonable to wonder about the performance of implementations with input-independent runtime built on full-featured multi-precision libraries. Since the main cost of blinding is in complexity rather than execution time, a good first-order estimate is that such implementations would be similar in speed to their unblinded counterparts. A performance comparison with a blinded variant of the map of Section 3 is future work.

As a rough comparison of complexity, our routines for constant-time arithmetic in \mathbb{F}_p and \mathbb{F}_{p^2} (briefly described below) comprise 342 lines of C code plus 614 lines of headers, constants, and automatically-generated addition chains. In contrast, the bare-bones mini-GMP library [GMP] (a subset of GMP’s integer arithmetic that, e.g., does not include the Legendre symbol function) comprises about 3600 lines of C.⁵ The multi-precision arithmetic implementations of Botan, Crypto++, and OpenSSL weigh in at about 10800 lines of C++, 5500 lines of C++, and 11700 lines of C, respectively (not including necessary support code, e.g., for memory management). Obviously, these libraries provide *much* richer functionality than ours! But our interest is in hash functions that are fast *without* rich functionality, so that they can be implemented in simple, compact code or hardware.

Implementation. We implement four hash-and-check variants and 30 variants of the hash functions of Section 5 (detailed below) in 3520 lines of C, including constant-time field arithmetic and curve operations for BLS12-381 (§2.1). Our field operations use reduced-radix Montgomery arithmetic geared to 64-bit processors. Our curve operations use the “slothful reduction” approach due to Scott [Sco17]. For exponentiations by fixed exponents (e.g., for square roots, etc.) we automatically generate addition chains by trying the methods of Bos and Coster [BC90], Bergeron et al., [BBBD89, BBB94], and Yacobi [Yac91, Yac98], and selecting the best result [Add]. For hashing and PRG we use the OpenSSL version 1.1.1.c implementations of SHA-256 and AES-CTR [Opea]. We use GMP version 6.1.2 [GMP] as our full-featured multi-precision library.

We have released our implementation under an open-source license [BH].

Benchmarks and baselines. For each of the hash functions of Section 5 plus the hash-and-check method, we evaluate up to three variants. The first uses GMP for modular arithmetic, Legendre symbols, and field inversions. The second also uses GMP, but is restricted to using only field operations. The third uses our constant-time field arithmetic library, and executes with input-independent runtime and memory access patterns. We discuss implementation specifics immediately below.

- For hash-and-check using full GMP functionality, we use `mpz_legendre` for $\chi(\cdot)$; in microbenchmarks, this is orders of magnitude faster than an exponentiation in \mathbb{F}_p . When using only field operations, we avoid computing Legendre symbols by computing $\beta = \alpha^{(p+1)/4}$ and checking whether $\beta^2 = \alpha$ (and analogously for \mathbb{F}_{p^2}).
- For the Shallue–van de Woestijne maps (§3) using full GMP functionality, we implement as suggested in Section 3 (i.e., using Montgomery’s trick). We use `mpz_legendre` for $\chi(\cdot)$, and compute inversions with the `mpz_invert` function; in our microbenchmarks, this costs an order of magnitude less than an exponentiation in \mathbb{F}_p . Since it is not constant time, this implementation computes each x_j only if necessary.

For the Shallue–van de Woestijne implementations that do not use rich functionality, we avoid inversions and Legendre symbol computations using the trick of Bernstein

⁵The full GMP is faster because it includes optimized assembly, but it is also almost two orders of magnitude more code and includes unneeded functionality, e.g., rational and floating-point support.

et al. [BDL⁺12] described in Section 4: to check x_j we compute the numerator U_j and denominator V_j of $f(x_j)$, compute $U_j V_j (U_j V_j^3)^{(p-3)/4}$, and check whether we have found a square root (and analogously for \mathbb{F}_{p^2}). The constant-time Shallue–van de Woestijne map thus requires three exponentiations.

- Our SWU map (§4) needs neither inversions nor Legendre symbol computations, so there is no implementation that uses full GMP functionality. Also, the constant-time versions require only one exponentiation because of our optimizations (§4).
- We do not implement constant-time hash-and-check, because it would either be very slow or have high failure probability.
- All hash functions return points in Jacobian projective coordinates (§2), meaning they do not need to compute an inversion after clearing the cofactor. Our implementation of Ψ (§5) also works in projective coordinates to avoid inversions. This is reasonable for BLS signatures because another exponentiation would immediately follow, and its input would be in Jacobian projective form in any case.⁶

Setup and method. We measure each hash function’s cost by executing it 10^6 times. We give all hash functions the same sequence of inputs, which we generate by hashing a seed. Each execution invokes the SHA-256 compression function once, uses the result to seed an AES-CTR PRG, extracts one or more field elements as required by the construction, and executes the hash function on the extracted field element(s).

We report cost on an Intel Xeon E3-1535M v6 with hyperthreading and dynamic frequency scaling disabled. Our testbench machine runs Arch Linux [Arc] (current as of July 10, 2019) with kernel version 5.2 and GCC version 9.1.0 [GCC].

Results. Table 1 shows the results.

- *Hash-and-check*, for both G_1 and G_2 , is reasonably fast in the average case, even when restricted to field operations. As discussed in Section 1, however, it is relatively easy to find messages that take many iterations to hash. To illustrate this, we report the mean of the worst 10% of runtimes for each hash-and-check experiment. When hash-and-check is implemented using a full-featured multi-precision library, the worst decile is within ≈ 3 –9% of the average case, because additional iterations (i.e., Legendre symbol computations) are inexpensive.

When hash-and-check is restricted to field operations, worst-decile performance is considerably worse than the average case: $\approx 83\%$ worse for G_1 , $\approx 87\%$ for G_2 . In other words, about 10% of the time on a random message hash-and-check performs almost as badly as the *slowest* of the alternatives that we consider. For adversarially-chosen messages, the performance could easily be even worse.

- *Construction #1*, when built on either of the maps of Sections 3 and 4 (hereafter “SW” and “SWU”, respectively), gives roughly similar performance to average-case hash-and-check, whether implemented either using a full-featured multi-precision library or restricted to field operations.

For constant-time implementations of construction #1, SW is $\approx 60\%$ slower than SWU; this is because the SW map’s cost is dominated by three exponentiations, whereas

⁶In principle, exponentiating a point P is slightly faster when $Z = 1$, because it allows using a mixed addition law [EFD], which is less expensive. In practice there is little difference, because fast exponentiation routines usually precompute small powers of P , and the resulting points have $Z \neq 1$ whether or not P does. Such an exponentiation routine using mixed addition would anyway require one inversion (to clear denominators of the precomputed points via Montgomery’s trick), meaning that P ’s denominator can be cleared for free. The difference in cost is thus only in the precomputation, which is negligible.

Table 1: Cost in thousands of CPU cycles for hash-and-check and the constructions of Section 5 applied to the maps of Sections 3 (“SW”) and 4 (“SWU”). For hash-and-check, “worst 10%” is the mean of the slowest 10% of runtimes. “Full MP lib” implementations use complex functionality from GMP (§6), “Field ops only” implementations are restricted to field operations, and “Constant time” implementations are further restricted to executing in time independent of the value (but not the length) of the hash input.

Group	Hash function	Map	Full MP lib	Field ops only	Constant time
G_1	Hash-and-check (worst 10%)	—	319	389	—
		—	348	712	—
	Construction #1	SW	330	376	577
		SWU	—	341	361
	Construction #2	SW	459	564	965
		SWU	—	456	496
	Construction #3	SW	675	721	944
		SWU	—	694	735
	G_2	Hash-and-check (worst 10%)	—	2327	3123
—			2390	5833	—
Construction #4		SW	2345	2954	3990
		SWU	—	2372	2364
Construction #5		SW	3259	4474	6555
		SWU	—	3280	3264
Construction #6		SW	3759	4363	5482
		SWU	—	3767	3822

the SWU map requires only one because of our optimizations (§4). Without our optimizations, the constant-time SWU map also requires three exponentiations (one inversion and two square roots), so we expect its performance to be almost identical to the constant-time SW map’s. We confirmed this hypothesis by implementing and measuring an unoptimized constant-time SWU map.

Importantly, the gap in performance between the constant-time SWU-based hash and the fastest non-constant-time SW-based hash is only $\approx 9\%$. In other words, even if blinding had no cost, a constant-time SW-based hash built on a full-featured multi-precision library would be only slightly faster than our SWU-based hash, even though a blinded SW-based hash entails much more implementation complexity.

- *Construction #2* is slower than construction #1 by $\approx 34\text{--}67\%$, depending on the map and implementation style. Ignoring the constant-time SW map, which is an outlier, the range is $\approx 34\text{--}50\%$; roughly speaking, this is the cost of making the hash function indifferentiable from a random oracle without the downsides of hash-and-check.

The SW and SWU maps are about the same speed when the SW map’s implementation uses rich multi-precision functionality. When restricted to field operations, however, the SW map is $\approx 24\%$ slower than the SWU map. The constant-time implementation of the SW map is dramatically worse: $\approx 95\%$ slower than the constant-time SWU map. The gap is bigger than in construction #1 because this construction entails two map evaluations rather than one.

Similarly to construction #1, the constant-time SWU-based hash is only $\approx 8\%$ slower than the fastest non-constant-time SW-based hash.

- *Construction #3* is the slowest hash for G_1 , $\approx 28\text{--}52\%$ slower than #2, in spite of aggressive parameter selection (§5). Improving uniformity to match construction #2 would make the situation even worse.

The exception is the constant-time SW-based hash: constructions #2 and #3 are similar in (lack of) speed, because the cost of evaluating the constant-time SW map twice (in #2) is similar to the cost of the constant-time multi-exponentiation (in #3). For the other variants, the multi-exponentiation dominates execution time, so the differences between the SW and SWU map and between full-featured and restricted multi-precision operations are small.

- *Construction #4*, like #1, is roughly competitive with hash-and-check in G_2 when both are implemented with the full-featured multi-precision library.

For implementations using only field arithmetic, however, the much higher cost of square root computations ($\approx 7\times$: compared to \mathbb{F}_p , a square-root computation in \mathbb{F}_{p^2} is an exponentiation of twice the length, where each step is $3\text{--}4\times$ more expensive) means that SW and hash-and-check are $\approx 25\text{--}32\%$ slower than SWU. This difference is even more pronounced in the constant-time case: SW is $\approx 69\%$ slower.

In this and the next construction, the constant-time SWU hash is—somewhat counterintuitively—slightly faster than the non-constant-time one. Microbenchmarks show that the constant-time and non-constant-time field operations have similar costs, so all else equal we should expect the constant-time hash to be slightly slower. We traced this anomaly to function call overhead: whereas the non-constant-time field operations entail calls into a shared library, the constant-time field operations are defined locally, so they are aggressively inlined by the compiler. This effect is especially pronounced for G_2 because each operation over \mathbb{F}_{p^2} entails three or four operations over \mathbb{F}_p and thus three or four times higher function call overhead.

- *Construction #5* continues the trend: the constant-time SWU map is about as fast as the fastest SW map, $\approx 36\%$ faster than the SW map restricted to field operations, and $\approx 2\times$ faster than the constant-time SW map.
- *Construction #6*, like construction #3, has generally poor performance in spite of the aggressive parameter choice. The constant-time SW-based hash is the exception: construction #6 is faster than #5, because in \mathbb{F}_{p^2} the cost of the extra SW map evaluation (in #5) is much higher than the cost of the multi-exponentiation (in #6).

Costs in more detail. Tables 2 and 3 give detailed cost breakdowns for the constant-time versions of constructions #1, #2, #4, and #5. Each table lists costs in terms of field operations, plus estimates of total costs expressed in terms of \mathbb{F}_p multiplications under the assumption that a squaring costs 80% as much as a multiplication in \mathbb{F}_p . (Costs of \mathbb{F}_{p^2} operations in terms of \mathbb{F}_p operations are given in the caption of Table 3.)

As expected, clearing the cofactor is a major contributor to overall cost. For SWU-based hashes it accounts for $\approx 38\text{--}53\%$ of the cost for G_1 and $\approx 44\text{--}61\%$ for G_2 . (Relative figures are lower for SW-based hashes because of the greater cost of the SW map.)

For SWU-based hashes to G_1 (Table 2), the cost of evaluating the 11-isogeny map is almost a third of the cost of a square root in \mathbb{F}_p , in part because evaluating the isogeny on a point in projective coordinates (§4.3) nearly doubles the number of multiplications required (in exchange for saving an inversion, a worthwhile tradeoff). Nevertheless, the isogeny map is only about 11% of the total cost in the worst case (construction #1). The isogeny for hashes to G_2 (Table 3) costs only about 5% as much as a square root in \mathbb{F}_{p^2} (primarily because the square root in \mathbb{F}_{p^2} is much more costly than in \mathbb{F}_p) and comprises less than 2% of the total cost in the worst case (construction #4).

Table 2: Costs for constant-time hash constructions #1 and #2 to G_1 (Table 1). mul is one multiplication in \mathbb{F}_p ; sqr is one squaring. The rightmost column estimates the total cost assuming that a sqr costs 80% of a mul.

	mul	sqr	total
Basic operations			
$\sqrt{U/V}$ (§4.2, [BCI ⁺ 10])	85	379	388
11-isogeny map (§4.3)	120	5	124
point addition	12	5	16
point doubling	4	5	8
cofactor clearing: P^{1-z} (§5)	320	340	592
point addition $\times 6$ + point doubling $\times 62$			
SW-Map ₁ (§3)	268	1146	1185
$\sqrt{U/V} \times 3$ + mul $\times 13$ + sqr $\times 9$			
SWU-Map ₁ without 11-isogeny map (§4)	99	383	405
$\sqrt{U/V}$ + mul $\times 14$ + sqr $\times 4$			
Hash construction #1			
§3 map	588	1486	1777
SW-Map ₁ + cofactor clearing			
§4 map	539	728	1121
SWU-Map ₁ + 11-isogeny map + cofactor clearing			
Hash construction #2			
§3 map	868	2637	2978
SW-Map ₁ $\times 2$ + point addition + cofactor clearing			
§4 map	650	1116	1543
SWU-Map ₁ $\times 2$ + point addition + 11-isogeny map + cofactor clearing			
256-bit scalar multiplication	1908	1645	3224
point addition $\times 74$ + point doubling $\times 255$			

Finally, both tables list the cost of a 256-bit scalar multiplication on the respective curves, as would be used to generate a BLS signature (§2.2). (The listed figures are for a 4-bit windowed exponentiation, including the cost of precomputation.) For signatures on G_1 (i.e., hash evaluation plus scalar multiplication), the SWU-based maps result in a savings of ≈ 15 – 30% compared to the SW-based maps. For G_2 , the savings is ≈ 27 – 47% .

Discussion. In sum, our optimizations yield a constant-time SWU map that is at worst $\approx 9\%$ slower than the *fastest* SW map implementation, even though our optimized SWU map’s implementation is very simple—it uses only field operations—while the SW map’s implementation requires a much richer multi-precision arithmetic implementation.

We also find that our optimizations are effective: in their absence, a constant-time SWU implementation using only field operations requires three exponentiations (an inversion and two square roots). Our experiments show that this puts its cost on par with the cost of our constant-time SW map implementation, whose execution time is also dominated by three exponentiations. In other words, our optimizations give roughly a 1.3 – $2\times$ speed-up.

Finally, we argue that a few percent performance overhead is a worthwhile trade for the simplicity of our optimized SWU map. In fact, even implementations that already have a full-featured multi-precision integer library might prefer the constant-time SWU map to a blinded SW map: as previously discussed, blinded implementations require a good source of randomness, or the blinding may be rendered ineffective; and they are intrinsically more complex to maintain and debug, e.g., because blinded execution is nondeterministic.

Table 3: Costs for constant-time hash constructions #4 and #5 to G_2 (Table 1). mul2 is one multiplication in \mathbb{F}_{p^2} , sqr2 is one squaring, and mulh is a multiplication by an element of \mathbb{F}_p . mul2 costs 4 \mathbb{F}_p multiplications, mulh costs 2 \mathbb{F}_p multiplications, and sqr2 costs one multiplication and two squarings in \mathbb{F}_p . The rightmost column estimates the total cost in \mathbb{F}_p multiplications, assuming that squaring costs 80% of multiplication in \mathbb{F}_p .

	mul2	mulh	sqr2	total
Basic operations				
$\sqrt{U/V}$ (§4.2, [BCI ⁺ 10])	156	0	761	2603
3-isogeny map (§4.3)	30	0	3	128
point addition	11	3	5	63
point doubling	2	3	5	27
ψ , “untwist-Frobenius-twist” (§5, [GS08, §5])	11	5	2	59
Ψ , cofactor clearing (§5, [BP17, §4.1])	448	435	706	4498
$\psi \times 3 + \text{point addition} \times 15 + \text{point doubling} \times 125$				
SW-Map ₂ (§3)	477	6	2292	7879
$\sqrt{U/V} \times 3 + \text{mul2} \times 9 + \text{mulh} \times 6 + \text{sqr2} \times 9$				
SWU-Map ₂ without 3-isogeny map (§4)	176	10	769	2723
$\sqrt{U/V} + \text{mul2} \times 20 + \text{mulh} \times 10 + \text{sqr2} \times 8$				
Hash construction #4				
§3 map	925	441	2998	12377
SW-Map ₂ + Ψ				
§4 map	654	445	1478	7349
SWU-Map ₂ + 3-isogeny map + Ψ				
Hash construction #5				
§3 map	1413	450	5295	20319
SW-Map ₂ $\times 2 + \text{point addition} + \Psi$				
§4 map	841	458	2252	10135
SWU-Map ₂ $\times 2 + \text{point addition} + 3\text{-isogeny map} + \Psi$				
256-bit scalar multiplication	1324	987	1645	11547
point addition $\times 74 + \text{point doubling} \times 255$				

7 Conclusion

We tackled the problem of hashing to Barreto-Lynn-Scott pairing-friendly elliptic curves [BLS03], focusing on BLS12-381 [Bow17]. To do so, we revisited the Shallue-van de Woestijne [SvdW06] and “simplified” Shallue-van de Woestijne-Ulas [Ula07, BCI⁺10] maps.

We proposed an “indirect” SWU map for Barreto-Lynn-Scott curves. Specifically, we showed a simple way of extending the SWU map to curves with j -invariant $\in \{0, 1728\}$, by mapping to an isogenous curve and then evaluating the isogeny map. We also proposed a small change that extends the SWU map to curves over fields where $\#\mathbb{F} \not\equiv 3 \pmod{4}$.

We then described several optimizations that make the SWU map simpler to implement and faster to evaluate, including in constant time. Specifically, our optimizations eliminate field inversions and quadratic character computations in the SWU map, and make it possible to evaluate the map in constant time by computing just one modular square root.

Finally, we implemented and evaluated 34 hash function variants built on the Shallue-van de Woestijne and optimized SWU maps. All told, we found that our optimizations to the SWU map yield hash functions that are fast, simple to implement, and constant time. Specifically, constant-time hash functions based on this map implemented using only field arithmetic are within 9% of the best-performing non-constant-time hash functions, which require significantly more complex implementations.

Acknowledgments

This work was supported in part by the NSF, the ONR, the Simons Foundation, the Stanford Center for Blockchain Research, and the Ripple Foundation. We thank Michael Scott for describing the trick in Section 5 for quickly clearing the cofactor on $E_1(\mathbb{F}_p)$, Fraser Brown for detailed comments, and the anonymous reviewers for their valuable feedback.

References

- [Add] kwantam/addchain. <https://github.com/kwantam/addchain>.
- [AFK⁺13] Diego F. Aranha, Laura Fuentes-Castañeda, Edward Knapp, Alfred Menezes, and Francisco Rodríguez-Henríquez. Implementing pairings at the 192-bit security level. In Michel Abdalla and Tanja Lange, editors, *PAIRING 2012*, volume 7708 of *LNCS*, pages 177–195. Springer, Heidelberg, May 2013.
- [AR12] Gora Adj and Francisco Rodríguez-Henríquez. Square root computation over even extension fields. *Cryptology ePrint Archive*, Report 2012/685, 2012. <http://eprint.iacr.org/2012/685>.
- [Arc] Arch Linux. <https://www.archlinux.org>.
- [BBB94] F. Bergeron, J. Berstel, and S. Brlek. Efficient computation of addition chains. *Journal de Théorie des Nombres de Bordeaux*, 6(1):21–38, 1994.
- [BBBD89] F. Bergeron, J. Berstel, S. Brlek, and C. Duboc. Addition chains using continued fractions. *J. Algorithms*, 10(3):403–412, September 1989.
- [BC90] Jurjen N. Bos and Matthijs J. Coster. Addition chain heuristics. In Gilles Brassard, editor, *CRYPTO’89*, volume 435 of *LNCS*, pages 400–407. Springer, Heidelberg, August 1990.
- [BCI⁺10] Eric Brier, Jean-Sébastien Coron, Thomas Icart, David Madore, Hugues Randriam, and Mehdi Tibouchi. Efficient indifferentiable hashing into ordinary elliptic curves. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 237–254. Springer, Heidelberg, August 2010.
- [BCTV14] Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Scalable zero knowledge via cycles of elliptic curves. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 276–294. Springer, Heidelberg, August 2014.
- [BDL⁺12] Daniel J. Bernstein, Niels Duif, Tanja Lange, Peter Schwabe, and Bo-Yin Yang. High-speed high-security signatures. *Journal of Cryptographic Engineering*, 2(2):77–89, September 2012.
- [BDN18] Dan Boneh, Manu Drijvers, and Gregory Neven. Compact multi-signatures for smaller blockchains. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part II*, volume 11273 of *LNCS*, pages 435–464. Springer, Heidelberg, December 2018.
- [Ber08] Daniel J. Bernstein. ChaCha, a variant of Salsa20. <https://cr.yp.to/chacha/chacha-20080128.pdf>, January 2008.
- [BF03] Dan Boneh and Matthew K. Franklin. Identity based encryption from the Weil pairing. *SIAM Journal on Computing*, 32(3):586–615, 2003.

- [BGLS03] Dan Boneh, Craig Gentry, Ben Lynn, and Hovav Shacham. Aggregate and verifiably encrypted signatures from bilinear maps. In Eli Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 416–432. Springer, Heidelberg, May 2003.
- [BGWZ19] Dan Boneh, Sergey Gorbunov, Hoeteck Wee, and Zhenfei Zhang. Bls signature scheme. Technical Report draft-boneh-bls-signature-00, Internet Engineering Task Force, February 2019.
- [BH] bls12-381_hash: Fast and constant-time hashing to the BLS12-381 elliptic curve. https://github.com/kwantam/bls12-381_hash.
- [BHKL13] Daniel J. Bernstein, Mike Hamburg, Anna Krasnova, and Tanja Lange. Elligator: elliptic-curve points indistinguishable from uniform random strings. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 2013*, pages 967–980. ACM Press, November 2013.
- [BLS01] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the Weil pairing. In Colin Boyd, editor, *ASIACRYPT 2001*, volume 2248 of *LNCS*, pages 514–532. Springer, Heidelberg, December 2001.
- [BLS03] Paulo S. L. M. Barreto, Ben Lynn, and Michael Scott. Constructing elliptic curves with prescribed embedding degrees. In Stelvio Cimato, Clemente Galdi, and Giuseppe Persiano, editors, *SCN 02*, volume 2576 of *LNCS*, pages 257–267. Springer, Heidelberg, September 2003.
- [BN06] Paulo S. L. M. Barreto and Michael Naehrig. Pairing-friendly elliptic curves of prime order. In Bart Preneel and Stafford Tavares, editors, *SAC 2005*, volume 3897 of *LNCS*, pages 319–331. Springer, Heidelberg, August 2006.
- [Bos14] Joppe W. Bos. Constant time modular inversion. *Journal of Cryptographic Engineering*, 4(4):275–281, November 2014.
- [Bota] Botan: Crypto and TLS for modern C++. <https://botan.randombit.net>.
- [Botb] Botan: Side channels. https://botan.randombit.net/manual/side_channels.html.
- [Bow17] Sean Rowe. BLS12-381: New zk-SNARK elliptic curve construction. <https://electriccoin.co/blog/new-snark-curve/>, March 2017.
- [BP17] Alessandro Budroni and Federico Pintore. Efficient hash maps to \mathbb{G}_2 on BLS curves. Cryptology ePrint Archive, Report 2017/419, 2017. <http://eprint.iacr.org/2017/419>.
- [CDF⁺11] Ray C. C. Cheung, Sylvain Duquesne, Junfeng Fan, Nicolas Guillermine, Ingrid Verbauwhede, and Gavin Xiaoxu Yao. FPGA implementation of pairings using residue number system and lazy reduction. In Bart Preneel and Tsuyoshi Takagi, editors, *CHES 2011*, volume 6917 of *LNCS*, pages 421–441. Springer, Heidelberg, September / October 2011.
- [Chi] Chia-network/bls-signatures. <https://github.com/Chia-Network/bls-signatures>.
- [CK11] Jean-Marc Couveignes and Jean-Gabriel Kammerer. The geometry of flex tangents to a cubic curve and its parameterizations. Cryptology ePrint Archive, Report 2011/033, 2011. <http://eprint.iacr.org/2011/033>.

- [Coh93] Henri Cohen. *A Course in Computational Algebraic Number Theory*. Springer-Verlag, 1993.
- [Cry] Crypto++ Library. <https://www.cryptopp.com/>.
- [EFD] Explicit formulas database: Genus-1 large-characteristic short weierstrass curves. <http://www.hyperelliptic.org/EFD/g1p/auto-shortw.html>.
- [Eth] Ethereum 2.0 spec: Bls signatures. https://github.com/ethereum/eth2.0-specs/blob/master/specs/bls_signature.md.
- [Far11] Reza Rezaeian Farashahi. Hashing into hessian curves. In Abderrahmane Nitaj and David Pointcheval, editors, *AFRICACRYPT 11*, volume 6737 of *LNCS*, pages 278–289. Springer, Heidelberg, July 2011.
- [FFS⁺13] Reza R. Farashahi, Pierre-Alain Fouque, Igor E. Shparlinski, Mehdi Tibouchi, and J. Felipe Voloch. Indifferentiable deterministic hashing to elliptic and hyperelliptic curves. *AMS Mathematics of Computation*, 82(281):491–512, 2013.
- [FJT13] Pierre-Alain Fouque, Antoine Joux, and Mehdi Tibouchi. Injective encodings to elliptic curves. In Colin Boyd and Leonie Simpson, editors, *ACISP 13*, volume 7959 of *LNCS*, pages 203–218. Springer, Heidelberg, July 2013.
- [FKR12] Laura Fuentes-Castañeda, Edward Knapp, and Francisco Rodríguez-Henríquez. Faster hashing to G_2 . In Ali Miri and Serge Vaudenay, editors, *SAC 2011*, volume 7118 of *LNCS*, pages 412–430. Springer, Heidelberg, August 2012.
- [FST10] David Freeman, Michael Scott, and Edlyn Teske. A taxonomy of pairing-friendly elliptic curves. *Journal of Cryptology*, 23(2):224–280, April 2010.
- [FSV09] Reza R. Farashahi, Igor E. Shparlinski, and Jos’e Felipe Voloch. On hashing into elliptic curves. *Journal of Mathematical Cryptology*, 3(4):353–360, 2009.
- [FT10] Pierre-Alain Fouque and Mehdi Tibouchi. Estimating the size of the image of deterministic hash functions to elliptic curves. In Michel Abdalla and Paulo S. L. M. Barreto, editors, *LATINCRYPT 2010*, volume 6212 of *LNCS*, pages 81–91. Springer, Heidelberg, August 2010.
- [FT12] Pierre-Alain Fouque and Mehdi Tibouchi. Indifferentiable hashing to Barreto-Naehrig curves. In Alejandro Hevia and Gregory Neven, editors, *LATINCRYPT 2012*, volume 7533 of *LNCS*, pages 1–17. Springer, Heidelberg, October 2012.
- [GCC] GCC, the GNU compiler collection. <https://gcc.gnu.org/>.
- [GGPR13] Rosario Gennaro, Craig Gentry, Bryan Parno, and Mariana Raykova. Quadratic span programs and succinct NIZKs without PCPs. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 626–645. Springer, Heidelberg, May 2013.
- [GLV01] Robert P. Gallant, Robert J. Lambert, and Scott A. Vanstone. Faster point multiplication on elliptic curves with efficient endomorphisms. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 190–200. Springer, Heidelberg, August 2001.
- [GMP] The GNU Multi-Precision arithmetic library. <https://gmplib.org/>.

- [GR04] Steven D. Galbraith and Victor Rotger. Easy decision Diffie-Hellman groups. *London Mathematical Society Journal of Computation and Mathematics*, 7:201–218, 2004.
- [Gro16] Jens Groth. On the size of pairing-based non-interactive arguments. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 305–326. Springer, Heidelberg, May 2016.
- [GS08] Steven D. Galbraith and Michael Scott. Exponentiation in pairing-friendly groups using homomorphisms. In Steven D. Galbraith and Kenneth G. Paterson, editors, *PAIRING 2008*, volume 5209 of *LNCS*, pages 211–224. Springer, Heidelberg, September 2008.
- [Gui10] Nicolas Guillermine. A high speed coprocessor for elliptic curve scalar multiplications over \mathbb{F}_p . In Stefan Mangard and François-Xavier Standaert, editors, *CHES 2010*, volume 6225 of *LNCS*, pages 48–64. Springer, Heidelberg, August 2010.
- [Has] Helmut Hasse. Zur Theorie der abstrakten elliptischen Funktionenkörper, I–III. *Journal für die reine und angewandte Mathematik*, 1936(175).
- [HSV06] F. Hess, N.P. Smart, and F. Vercauteren. The eta pairing revisited. Cryptology ePrint Archive, Report 2006/110, 2006. <http://eprint.iacr.org/2006/110>.
- [Ica09] Thomas Icart. How to hash into elliptic curves. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 303–316. Springer, Heidelberg, August 2009.
- [IMCT02] Tsutomu Iijima, Kazuto Matsuo, Jinhui Chao, and Shigeo Tsujii. Construction of Frobenius maps of twists [of] elliptic curves and its application to elliptic scalar multiplication. In *Proc. SCIS*, January 2002.
- [KB16] Taechan Kim and Razvan Barbulescu. Extended tower number field sieve: A new complexity for the medium prime case. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 543–571. Springer, Heidelberg, August 2016.
- [KLR10] Jean-Gabriel Kammerer, Reynald Lercier, and Guénaél Renault. Encoding points on hyperelliptic curves over finite fields in deterministic polynomial time. In Marc Joye, Atsuko Miyaji, and Akira Otsuka, editors, *PAIRING 2010*, volume 6487 of *LNCS*, pages 278–297. Springer, Heidelberg, December 2010.
- [Knu97] Donald E. Knuth. *Seminumerical Algorithms*, volume 2 of *The Art of Computer Programming*, chapter 4.6.4. Addison-Wesley, 3rd edition, 1997.
- [Koc96] Paul C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In Neal Koblitz, editor, *CRYPTO'96*, volume 1109 of *LNCS*, pages 104–113. Springer, Heidelberg, August 1996.
- [KT15] Taechan Kim and Mehdi Tibouchi. Improved elliptic curve hashing and point representation. In *Proc. Workshop on Coding and Cryptography*, April 2015.
- [Möl01] Bodo Möller. Algorithms for multi-exponentiation. In Serge Vaudenay and Amr M. Youssef, editors, *SAC 2001*, volume 2259 of *LNCS*, pages 165–180. Springer, Heidelberg, August 2001.
- [Mon87] Peter L. Montgomery. Speeding the Pollard and elliptic curve methods of factorization. *Math. Comp.*, 48(177):243–264, 1987.

- [MRH04] Ueli M. Maurer, Renato Renner, and Clemens Holenstein. Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology. In Moni Naor, editor, *TCC 2004*, volume 2951 of *LNCS*, pages 21–39. Springer, Heidelberg, February 2004.
- [MSS16] Alfred Menezes, Palash Sarkar, and Shashank Singh. Challenges with assessing the impact of NFS advances on the security of pairing-based cryptography. In *Mycrypt*, December 2016.
- [NaC] NaCL. <https://nacl.cr.yp.to>.
- [Opea] OpenSSL: Cryptography and SSL/TLS toolkit. <https://www.openssl.org/>.
- [Opeb] BIGNUM code is not constant-time due to bn_correct_top. <https://github.com/openssl/openssl/issues/6640>.
- [Pai] GitHub: zkcrypto/pairing BLS12-381 implementation. https://github.com/zkcrypto/pairing/tree/master/src/bls12_381.
- [PHGR13] Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. Pinocchio: Nearly practical verifiable computation. In *2013 IEEE Symposium on Security and Privacy*, pages 238–252. IEEE Computer Society Press, May 2013.
- [SBC⁺09] Michael Scott, Naomi Benger, Manuel Charlemagne, Luis J. Dominguez Perez, and Ezekiel J. Kachisa. Fast hashing to G_2 on pairing-friendly curves. In Hovav Shacham and Brent Waters, editors, *PAIRING 2009*, volume 5671 of *LNCS*, pages 102–113. Springer, Heidelberg, August 2009.
- [Sco17] Michael Scott. Slothful reduction. Cryptology ePrint Archive, Report 2017/437, 2017. <http://eprint.iacr.org/2017/437>.
- [Sco19] Michael Scott. Personal communication, April 2019.
- [Ska05] Mariusz Skalba. Points on elliptic curves over finite fields. *Acta Arithmetica*, 117(3):293–301, 2005.
- [SLA] Silicon Labs EFM32PG12 Pearl Gecko Microcontroller Reference Manual. <https://www.silabs.com/documents/public/reference-manuals/efm32pg12-rm.pdf>.
- [SM] SageMath. <http://www.sagemath.org/>.
- [SS04] Andrzej Schinzel and Mariusz Skalba. On equations $y^2 = x^n + k$ in a finite field. *Bulletin of the Polish Academy of Sciences Mathematics*, 52(3):223–226, 2004.
- [SvdW06] Andrew Shallue and Christiaan E. van de Woestijne. Construction of rational points on elliptic curves over finite fields. In *Algorithmic Number Theory Symposium*, July 2006.
- [Ula07] Maciej Ulas. Rational points on certain hyperelliptic curves over finite fields. *Bulletin of the Polish Academy of Sciences Mathematics*, 55(2):97–104, 2007.
- [Vél71] Jacques Vélu. Isogénies entre courbes elliptiques. *Comptes rendus de l’Académie des Sciences de Paris, Série A*, 273:238–241, 1971.
- [Yac91] Yacov Yacobi. Exponentiating faster with addition chains. In Ivan Damgård, editor, *EUROCRYPT’90*, volume 473 of *LNCS*, pages 222–229. Springer, Heidelberg, May 1991.

- [Yac98] Yacov Yacobi. Fast exponentiation using data compression. *SIAM J. Comput.*, 28(2):700–703, 1998.
- [YKS19] Shoko Yonezawa, Tetsutaro Kobayashi, and Tsunekazu Saito. Pairing-friendly curves. Technical Report draft-yonezawa-pairing-friendly-curves-02, Internet Engineering Task Force, July 2019.

A The isogeny maps

The following Sage [SM] script shows one way to find a suitable isogeny for use with the map of Section 4.

```
#!/usr/bin/env sage

import sage.schemes.elliptic_curves.isogeny_small_degree as isd

# look for isogenous curves having j-invariant not in {0, 1728}
# Caution: this can take a while!
def find_iso(E):
    for p_test in primes(30):
        isos = [ i for i in isd.isogenies_prime_degree(E, p_test)
                 if i.codomain().j_invariant() not in (0, 1728) ]
        if len(isos) > 0:
            return isos[0].dual()
    return None

# BLS12-381 parameters
z = -0xd201000000010000
h = (z - 1) ** 2 // 3
q = z ** 4 - z ** 2 + 1
p = z + h * q
assert is_prime(p)
assert is_prime(q)

# E1
F = GF(p)
E11 = EllipticCurve(F, [0, 4])
assert E11.order() == h * q

# E2
F2.<X> = GF(p^2, modulus=[1,0,1])
E112 = EllipticCurve(F2, [0, 4 * (1 + X)])
assert E112.order() % q == 0

iso_G1 = find_iso(E11)          # an isogeny from E' to E,
E11_prime = iso_G1.domain()    # where this is E'
assert iso_G1(E11_prime.random_point()).curve() == E11

iso_G2 = find_iso(E112)        # an isogeny from E2' to E2,
E112_prime = iso_G2.domain()  # where this is E2'
assert iso_G2(E112_prime.random_point()).curve() == E112
```

The following Sage script prints the rational maps $E'_1(\mathbb{F}_p) \mapsto E_1(\mathbb{F}_p)$ and $E'_2(\mathbb{F}_{p^2}) \mapsto E_2(\mathbb{F}_{p^2})$ used in the evaluation of Section 6.

```
#!/usr/bin/env sage

z = -0xd201000000010000
h = (z - 1) ** 2 // 3
q = z ** 4 - z ** 2 + 1
p = z + h * q
F = GF(p)
F2.<X> = GF(p^2, modulus=[1,0,1])

##
## Ell1' -> Ell1
##
# the following integers are base36-encoded to compress horizontally
vv = [ '12rutfybsn0bkhh9qg5qlfrpugmxeczd757bmomkcfj8qj94vfcibfpfz6778pphed5apz4t'
      , '74t8nxvndcq4fsu3byqf2ubacd8zpfj5htxn621zauzk8jeti9eg3iakb1qzc44dbo59g94ue8'
      , '793jv4j16d1k90qpfb51iyn9gbkoakins8196hjOrcw750ya3xz7bklkkyi2zsf1alzxka0q2v'
      , 'wjf4r8fn0t5ud2l8mj6qtxj6vwqkfV403p6t8rrlalpyli69k7yrbkyfqv3h3k4bup3ef0vqo'
      , '798csdr2a2qwd6hz4bjk1l53bbxbr7ndhgwI663wezfpv82hqp2iae52dm4atf4z2xgzaar09'
      , '3i0ielgn7to06ad2bvkiqmaxy3k9yn6fgqkh2hks4qt2xrlsgy6s45tc32fwk2ar4e8yajsjo'
      , '1rju5l2fizu0w476g37nhlfgbsfyylt153doespo31onzy21q72irl9s7s4p051r2160gw3p'
      , '1']
a_prime = int(vv[0], base=36)
b_prime = int(vv[1], base=36)
kernel_poly = [ int(v, base=36) for v in vv[2:] ]
Ell = EllipticCurve(F, [0, 4])
EllP = EllipticCurve(F, [a_prime, b_prime])
iso11 = EllipticCurveIsogeny(EllP, kernel_poly, codomain=Ell, degree=11)
# choose isogeny with opposite sign for y (this is arbitrary)
iso11.switch_sign()

##
## Ell2' -> Ell2
##
Ell2 = EllipticCurve(F2, [0, 4 * (X + 1)])
Ell2p = EllipticCurve(F2, [240 * X, 1012 * (X + 1)])
iso3 = EllipticCurveIsogeny(Ell2p, [6 * (1 - X), 1], codomain=Ell2)

# print rational maps
print iso11.rational_maps()
print iso3.rational_maps()

## you can also access the rational maps separately
# xmap = iso.rational_maps[0]
# ymap = iso.rational_maps[1]

## and also their numerators and denominators
# xmap_num = xmap.numerator()
# xmap_den = xmap.denominator()
# ymap_num = ymap.numerator()
# ymap_den = ymap.denominator()
```