

High-Performance Scalar Multiplication using 8-Dimensional GLV/GLS Decomposition

Joppe W. Bos¹, Craig Costello¹, Huseyin Hisil², and Kristin Lauter¹

¹ Microsoft Research, Redmond, USA

² Yasar University, Izmir, Turkey

Abstract. This paper explores the potential for using genus 2 curves over quadratic extension fields in cryptography, motivated by the fact that they allow for an *8-dimensional* scalar decomposition when using a combination of the GLV/GLS algorithms. Besides lowering the number of doublings required in a scalar multiplication, this approach has the advantage of performing arithmetic operations in a 64-bit ground field, making it an attractive candidate for embedded devices. We found cryptographically secure genus 2 curves which, although susceptible to index calculus attacks, aim for the standardized 112-bit security level. Our implementation results on both high-end architectures (Ivy Bridge) and low-end ARM platforms (Cortex-A8) highlight the practical benefits of this approach.

1 Introduction

Elliptic curve cryptography [30, 37] is a popular approach to realize public-key cryptography. One of the main reasons to employ elliptic curves, rather than using more traditional settings like finite fields, is efficiency. According to [44], the performance gain when transferring the Diffie-Hellman protocol [12] from finite fields to elliptic (genus 1) curves at the 128-bit security level is an order of magnitude. There is an active research area dedicated to enhancing the core operation in curve-based protocols: the scalar multiplication. A novel approach that facilitates fast scalar multiplications is the Gallant-Lambert-Vanstone (GLV) method [18]. If an elliptic curve $E(\mathbb{F}_q)$ comes equipped with a non-trivial endomorphism, then a scalar k can be decomposed into two “mini-scalars”, both of which are approximately half the bit-length of k : merging these mini-scalars means that the number of required point doublings in the scalar multiplication can be reduced by a factor of two. The GLV method was extended by Galbraith, Lin and Scott (GLS) [17], who show that regardless of the existence of an endomorphism on $E(\mathbb{F}_q)$, one can achieve a decomposition by considering the points $E(\mathbb{F}_{q^m})$ for $m > 1$. Furthermore, [17] explains that if E already comes equipped with a useful endomorphism over \mathbb{F}_q , then the GLV and GLS endomorphisms can be combined to achieve higher degree decompositions and increased performance. At Asiacrypt 2012, Longa and Sica [36] demonstrated this GLV/GLS combination to achieve a 4-dimensional scalar decomposition on elliptic curves over the quadratic extensions of large prime fields (i.e. $E(\mathbb{F}_{p^2})$), and set the current software speed record for computing scalar multiplications over non-binary fields. The authors of [8] recently showed the practical potential of hyperelliptic (genus 2) curves in cryptography. One attractive aspect of genus 2 curves is that, in general, their Jacobian group $\text{Jac}_C(\mathbb{F}_p)$ has a larger endomorphism ring than that of genus 1 curves. This means that over prime fields or over extension fields of the same degree, the highest possible degree of the GLV/GLS decomposition is twice as large in genus 2 as it is in genus 1.

In this paper we consider 8-dimensional scalar decompositions by exploring the use of genus 2 curves over quadratic extension fields. To the best of our knowledge, this is the first time an 8-dimensional scalar decomposition has been implemented and studied in detail,

addressing two of the open problems posed in the original GLS paper [17, §9]. Using decompositions of this size leads to practical performance issues that do not arise in the 2- and 4-dimensional case; we highlight some pitfalls and present solutions in a variety of scenarios.

In contrast to elliptic curves, “faster-than-generic” attacks are known on genus 2 curves over \mathbb{F}_{p^2} . Namely, one can use the “Weil descent” attack [14] to map the discrete logarithm problem to a higher dimensional abelian variety over \mathbb{F}_p , where index calculus attacks are possible [2, 19]. We assess the current state-of-the-art in index calculus attacks [21, 11] to give conservative security estimates, which present a strong case for the curves we use at the currently standardized 112-bit security level [43].

Since most high-end hardware architectures work with 64-bit words and many embedded platforms work with 32-bit words (like the ARM), using 64-bit primes means that our arithmetic in the ground field is respectively performed using one and two computer words only. We explore different approaches for arithmetic in \mathbb{F}_p , while using lazy reduction techniques from the pairing community [3] to achieve efficient arithmetic in \mathbb{F}_{p^2} . In addition to the 8-dimensional GLV/GLS approach, we consider “generic” genus 2 curves (curves which do not exploit any special properties) and the Kummer surface over \mathbb{F}_{p^2} . Our implementation results on a 64-bit Ivy Bridge processor and a Cortex-A8 ARM CPU show that this approach is competitive with the current state-of-the-art in elliptic curve cryptography, although we reiterate that our work targets the 112-bit security level, while most of the work we (are able to) compare against targets the 128-bit security level. Our implementations targeting 64-bit platforms will be made publicly available through [6].

This paper is organized as follows. Section 2 recalls the preliminaries. Section 3 discusses the security implications of using genus 2 over quadratic extension fields and summarizes our curve choices. Section 4 presents the 8-dimensional GLV/GLS approach and the techniques we employ to overcome various performance issues, while Section 5 explains the various arithmetic approaches we use. Section 6 presents our implementation results and Section 7 concludes the paper.

2 Preliminaries

In this paper we work with “imaginary” hyperelliptic curves of genus 2 over a quadratic extension of large prime fields. Such curves can be written as $C/\mathbb{F}_{p^2} : y^2 = x^5 + f_3x^3 + f_2x^2 + f_1x + f_0$. We use $\text{Jac}_C(\mathbb{F}_{p^2})$ to denote the Jacobian group and we abbreviate the Mumford representation of general (i.e. weight 2) divisors on $\text{Jac}_C(\mathbb{F}_{p^2})$ to write $(x^2 + u_1x + u_0, v_1x + v_0)$ as (u_1, u_0, v_1, v_0) in affine space, or as $(U_1 : U_0 : V_1 : V_0 : Z)$ in homogeneous projective space. We explore three algorithms for computing scalar multiplications on $\text{Jac}_C(\mathbb{F}_{p^2})$: (i) the generic algorithm which computes the scalar multiplication using a sequence of divisor doublings and additions only, (ii) the combination of the GLV [18] and GLS [17] algorithms which both exploit endomorphisms (but in a different way) to accelerate computations, and (iii) Gaudry’s fast formulas [20] for arithmetic on the Kummer surface associated to $\text{Jac}_C(\mathbb{F}_{p^2})$.

2.1 GLV and GLS algorithms.

We point out the difference between the GLV and GLS methods which use endomorphisms to speed up curve-based cryptography. The Gallant, Lambert and Vanstone (GLV) method [18] involves using special curves that come equipped with efficiently computable endomorphisms which are unrelated to the Frobenius endomorphism. For example, when $p \equiv 1 \pmod{\ell}$ for an

odd prime ℓ , the curve $C/\mathbb{F}_p : y^2 = x^\ell + a$ comes equipped with $\phi : (x, y) \mapsto (\xi_\ell x, y)$ for ξ_ℓ a non-trivial ℓ -th root of unity in \mathbb{F}_p . On the other hand, the Galbraith, Lin and Scott (GLS) method [17] does not rely on curves of a special form, but rather exploits the fact that, for any curve defined over \mathbb{F}_p , the p -power Frobenius endomorphism π_p acts non-trivially on points in extension fields of \mathbb{F}_p . The main issue to address in this case is that we can not simply define C over \mathbb{F}_p and consider the points in $C(\mathbb{F}_{p^m})$ for some $m > 1$, because the associated Jacobian group will properly contain the subgroup of points defined over \mathbb{F}_p , meaning that the group order cannot be prime, i.e. it will necessarily contain the large cofactor $\#\text{Jac}_{C/\mathbb{F}_p}(\mathbb{F}_p) \mid \#\text{Jac}_{C/\mathbb{F}_p}(\mathbb{F}_{p^2})$. Alternatively, if we were to define the curve C over \mathbb{F}_{p^m} , the p -power Frobenius would no longer give rise to a map on C . The solution here is to work on a curve C'/\mathbb{F}_{p^m} that is \mathbb{F}_{p^k} -isogenous to C/\mathbb{F}_p , where $m \mid k$. Starting with a point $P \in C'(\mathbb{F}_{p^m})$, we use the (dual) isogeny $\hat{\phi}$ to move P to $\hat{\phi}(P) \in C(\mathbb{F}_{p^k})$ where we can apply π_p (because C is defined over \mathbb{F}_p), before moving back to $C'(\mathbb{F}_{p^m})$ via the isogeny ϕ . The composition of these three maps gives rise to the map $\psi = \phi\pi\hat{\phi}$ which is (technically) defined over \mathbb{F}_{p^k} , but which gives rise to an endomorphism on the large prime order subgroup(s) of $C'(\mathbb{F}_{p^m})$ [17, Th. 1].

Galbraith *et al.* [17, §3] further show how the GLV and GLS ideas can be combined to give more advantageous decompositions. Namely, for curves that are both defined over extension fields *and* have additional (non-trivial) endomorphisms, they show that this is achieved by taking the isogeny ϕ (constituting ψ) to be the twisting isomorphism corresponding to the additional endomorphism(s) on C . For special Buhler-Koblitz curves [9] of the form $C/\mathbb{F}_{p^2} : y^2 = x^5 + a$, we discuss this combined approach in detail in Section 4.

2.2 The Kummer surface.

Gaudry [20] showed that scalar multiplications can be computed more efficiently on the Kummer surface associated to the Jacobian of genus 2 curves than on the Jacobian itself. Recently, the authors of [8] used Gaudry’s fast formulas on genus 2 curves over prime fields to set a new speed record for computing constant-time scalar multiplications. In this work we carry these techniques across to curves defined over quadratic extension fields, and since the method of using the Kummer surface essentially remains unchanged, we refer to [8, §5] for the details.

2.3 The CM Method over Quadratic Extension Fields.

In targeting primes p which are favorable for fast arithmetic as described in Section 5, we use the complex multiplication (CM) method to search for and generate cryptographically strong genus 2 curves over \mathbb{F}_{p^2} , i.e. genus 2 curves C/\mathbb{F}_{p^2} whose Jacobian group $\text{Jac}_C(\mathbb{F}_{p^2})$ has cardinality with a large prime factor. For a summary of the CM method, we refer to [8, §2.2], but note that, to find curves over \mathbb{F}_{p^2} instead of over \mathbb{F}_p , we search for CM fields where p decomposes in a different way than in the case where one is searching for curves defined over prime fields. We explain this below.

The CM method exploits the fact that the number of rational points on the Jacobian of a smooth, projective, irreducible curve of genus g over a finite field, is determined by g pieces of information which are encoded in its zeta function, which is in turn related to the characteristic polynomial of the Frobenius endomorphism, π . This polynomial is of degree four which determines a quartic CM field K , and its four roots come in conjugate pairs $(\pi, \bar{\pi})$ and $(\pi', \bar{\pi}')$. For a genus 2 curve C defined over \mathbb{F}_p , the number of points on the Jacobian over \mathbb{F}_p is $\#\text{Jac}(C)(\mathbb{F}_p) = (1 - \pi)(1 - \bar{\pi})(1 - \pi')(1 - \bar{\pi}')$, while the number of points on the Jacobian

over \mathbb{F}_{p^2} is $\#\text{Jac}(C)(\mathbb{F}_{p^2}) = (1 - \pi^2)(1 - \bar{\pi}^2)(1 - (\pi')^2)(1 - (\bar{\pi}')^2)$. This latter group is not suitable for cryptography because the former group is contained as a proper subgroup and therefore the order cannot be prime. Thus, for this paper, we need to generate curves over \mathbb{F}_{p^2} which are not defined over \mathbb{F}_p .

In order to know the possible group orders for such curves, it suffices to find the prime ideal decomposition of p in \mathcal{O}_K . This determines all possible π 's, because for abelian varieties over \mathbb{F}_q , $q = p^r$, the Frobenius endomorphism satisfies $\pi\bar{\pi} = q$. For primes p which are split completely in K and split completely into principal ideals in the reflex field of K , the abelian varieties with CM by K are defined over \mathbb{F}_p when reduced modulo an ideal lying over p . We therefore need to consider other decompositions so that the varieties are defined over \mathbb{F}_{p^2} when reduced modulo an ideal lying over p .

We do not consider primes that ramify in K , since those primes divide the discriminant of K , and will not be of cryptographic size. Let p be a prime which decomposes in K into a product of two degree 1 primes and one degree 2 prime – this decomposition can only happen in the non-Galois case, and here we assume that K is a primitive quartic CM field so that we avoid the bi-quadratic case anyway. Suppose the degree 1 primes are principal and let α and $\bar{\alpha}$ be their generators. Letting $\beta = \pm\alpha\bar{\alpha}^{-1}p$, we find that β are Weil- p^2 numbers corresponding to abelian varieties defined over \mathbb{F}_{p^2} when reduced modulo a prime lying over p , with group orders $N = \prod_{\sigma}(1 - \beta^{\sigma})$, where σ ranges over the complex embeddings of K (see [24, §3.6.5 Ex 5]). The same can be done for the case where a prime decomposes into the product of two ideals of degree 2, which can also occur in the case that K is Galois cyclic (see [24, §3.6.6 Ex 6]).

For primes p of a particularly favorable form, we then search for CM fields K in which the prime decomposes in one of two ways described above, and for which the group order has a prime factor of a given size. We then find roots in \mathbb{F}_{p^2} of the Igusa class polynomials of K modulo p , before using Mestre's algorithm to reconstruct an equation for the genus 2 curve. For the Igusa class polynomials corresponding to a CM field, we used Kohel's Echidna database [32] and Thomé's Barbecue database [50, 13].

3 Curve Choices and Security

We first discuss security aspects in Section 3.1, before presenting our curve choices in Section 3.2.

3.1 Weil Descent and Index Calculus.

Attacks which are asymptotically "faster-than-generic" are known to exist on curves over extension fields, using a combination of the ideas of Weil descent and index calculus (see for example [14, 2, 19, 22, 47, 16, 15, 10, 23, 21, 27]). In this work we are concerned with the best-known attacks on the discrete logarithm problem (DLP) in the Jacobian of a genus 2 curve C defined over a quadratic extension field \mathbb{F}_{p^2} . Following [10, 23], one attack transfers the DLP on $\text{Jac}(C)(\mathbb{F}_{p^2})$ to the Jacobian of a higher genus curve \tilde{C} which lies on the abelian variety over \mathbb{F}_p obtained via Weil restriction of scalars from $\text{Jac}(C)(\mathbb{F}_{p^2})$ [15, §7.1 - Ex. 7]. In general it can be hard to find such a curve \tilde{C} , and for the curves we use, the best known technique finds curves \tilde{C} of genus 8 to use in the attack detailed in [10, 23]. Certain cases of genus 2 (imaginary) hyperelliptic curves C over quadratic extension fields \mathbb{F}_{p^2} have been classified as "weak" [49, 38, 26], in that their special form makes it easier than usual to find a

Table 1. An overview of our implementations targeting the 112-bit security level. The security estimate (in bits) resulting from index calculus (i.c.) and Pollard rho (rho) attack are stated. For each instance, we state the prime p and the bit-lengths of the cofactor h and prime r where the group order is $h \cdot r$. For the Kummer instances, we also show the size of the prime (sub)group order r' of the twist. For the 8-GLV/GLS implementations, we report two performance numbers in each scenario: the first is for the constant-time precomputation method and the second is for the faster (but variable time) precomputation method (see Section 4.3). See Section 6 for the specifics about the benchmark platforms.

algorithm	reduction	base field p	$ h _2$	$ r _2$	$ r' _2$	security (bits)		cycles ($\cdot 10^3$)	
						rho	i.c.	IB	ARM
generic	special	$2^{61} - 1$	38	207	-	103	109	204	1492
	Mont.	$(2^{31} - 307656) \cdot 2^{32} - 1$	36	217	-	108	112	-	1808
	NIST	$2^{64} - 189$	36	221	-	110	113	333	-
Kummer	special	$2^{61} - 1$	38	207	228	103	109	108	767
	Mont.	$(2^{31} - 307656) \cdot 2^{32} - 1$	36	217	245	108	112	-	942
	NIST	$2^{64} - 189$	36	221	250	110	113	167	-
8-GLV/GLS	special	$2^{61} - 1$	32	213	-	105	109	100 / 92	617 / 576
	Mont.	$(2^{31} - 201) \cdot 2^{32} - 1$	31	222	-	109	112	-	859 / 810
	NIST	$2^{64} - 2285$	33	224	-	111	113	146 / 136	-

suitable curve \tilde{C} on the Weil restriction of $\text{Jac}(C)$. None of the curves we use fall into these weak classifications: we can essentially rule this out by ensuring that our curves cannot be written as $C : y^2 = (x - \alpha) \cdot h(x)$, with $h(x) \in \mathbb{F}_p[x]$. Thus, to the best of our knowledge, the fastest attack on our curves is due to Gaudry [21] with further improvements provided by Nagao [42]. Gaudry’s attack works directly on the abelian variety obtained as the Weil restriction of scalars, and solves the discrete logarithm problem on genus g hyperelliptic curves over \mathbb{F}_{p^n} , where both n and g are *fixed*, in heuristic asymptotic running time $\tilde{O}(p^{2 - \frac{2}{ng}})$, i.e. not including the “constants” depending on n and g and the logarithmic factors in p . For the sake of obtaining a better comparison with the generic Pollard rho algorithm, we reveal *some* of the factors that are hidden by the \tilde{O} . One of the constants in the \tilde{O} depends exponentially on both g and n as $2^{3n(n-1)g}$ [42]. Hence, a conservative lower bound on the asymptotic running time of this attack, expressed in terms of group operations on the genus g curve, is $\mathcal{O}(p^{2 - \frac{2}{ng}} \cdot 2^{3n(n-1)g} \cdot \log(p)^r)$ for some $r \geq 1$. To give a modest security estimate for our genus 2 curves over quadratic extension fields ($g = n = 2$), we take $r = 1$, ignore other constants involved and keep the \mathcal{O} in terms of group operations on the dimension 4 abelian variety obtained as the Weil restriction of $\text{Jac}(C)$. Hence, we arrive at $p^{3/2} \cdot 2^{12} \cdot \log(p)$ group operations as a conservative estimate of a lower bound on the complexity of Gaudry’s attack for genus 2 curves over \mathbb{F}_{p^2} .

3.2 Generic curves, Buhler-Koblitz curves, and Kummer Surfaces over \mathbb{F}_{p^2} .

For each of the 3 algorithms (generic, Kummer, 8-GLV/GLS) considered in this work, we used the CM method to find curves over quadratic extension fields with characteristic less than 2^{64} that fall into 3 different categories: those which use a Montgomery-friendly prime of the form $(2^{31} - c_1) \cdot 2^{32} - 1$ to target the 32-bit (ARM) environment, those which use a NIST-friendly prime of the form $2^{64} - c_2$ to target 64-bit platforms, and those which use the prime $2^{61} - 1$ that can employ specialized Montgomery- and NIST-like reduction (cf. Section 5). We note that

all our fields³ have $p \equiv 3 \pmod{4}$, so that the quadratic extension can always be constructed as $\mathbb{F}_{p^2} = \mathbb{F}_p[i]/(i^2 + 1)$. Table 1 summarizes the curves that we use in this paper together with the arithmetic approach taken (Montgomery, NIST or special), the security claims and the implementation results. All of the curve parameters are given in Appendix A. The security estimate for the Pollard rho attack [46] is obtained using $\log_2 \left(\sqrt{\frac{\pi r}{2\#\text{Aut}}} \right)$, where $\#\text{Aut}$ is the size of the automorphism group of C . In our case all of the GLV/GLS curves have $\#\text{Aut} = 10$, while all the other curves have $\#\text{Aut} = 2$. As discussed in Section 3.1, the runtime of the index calculus attack depends on p , while the complexity of the Pollard rho attack depends on the (sub)group order r . When searching for curves, we aimed to balance the attack complexity of both approaches in order to enhance performance: relaxing the size of r does not decrease the level of claimed security for index-calculus, but results in smaller scalars (and faster scalar multiplications). This explains why the subgroup orders r in Table 1 are significantly smaller than 256 bits – our target for the Pollard rho security was to aim slightly below our estimate for the index calculus algorithms for the sake of being conservative. Of the 10 isomorphism classes of Buhler-Koblitz curves over $p = 2^{61} - 1$, we chose the one corresponding to the Jacobian group with the largest prime factor of size 213 bits.

Since the index calculus attacks may in practice include large hidden constants, higher powers of $\log p$, or expensive group operations on higher dimensional abelian varieties, we have included the necessary details for curves of prime order (or with very small cofactors) in Table 5 in Appendix B, as well as some timing results on the Ivy Bridge platform to give a rough idea of the difference in performance that arises from different sized scalars.

4 8-dimensional GLV/GLS

In this section we apply the GLV/GLS combination (see Section 2.1) to a particular family of curves, namely Buhler-Koblitz (BK) curves of the form $C/\mathbb{F}_{p^2} : y^2 = x^5 + a$.

4.1 8-GLV/GLS on Buhler-Koblitz Curves over \mathbb{F}_{p^2}

Following the description in [17, §5], we use a BK curve of the form $C/\mathbb{F}_{p^2} : y^2 = x^5 + u^{10}$, with $p \equiv 1 \pmod{10}$ and $u^{10} \in \mathbb{F}_{p^2}$ such that $u \in \mathbb{F}_{p^{20}}$. Let $C'/\mathbb{F}_p : y^2 = x^5 + 1$. The map $\phi^{-1} : C \rightarrow C'$ defined as $\phi^{-1} : (x, y) \mapsto (x/u^2, y/u^5)$ takes points in $C(\mathbb{F}_{p^2})$ to points in $C'(\mathbb{F}_{p^{20}})$, where the p -power Frobenius map $\pi_p : C' \rightarrow C'$ acts non-trivially. Finally, the map $\phi : C' \rightarrow C$ defined as $\phi : (x', y') \mapsto (u^2 x', u^5 y')$ moves the result of π_p back to $C(\mathbb{F}_{p^2})$. Composing these maps into $\psi = \phi \pi_p \phi^{-1}$ gives $\psi : C \rightarrow C$, defined as $\psi : (x, y) \mapsto (x^p \cdot (u^{-2})^{p-1}, y^p \cdot (u^{-5})^{p-1})$; notice that $10 \mid p - 1$ and $u^{10} \in \mathbb{F}_{p^2}$ together imply that this map is defined over \mathbb{F}_{p^2} . Since we use $p \equiv 3 \pmod{4}$ and construct \mathbb{F}_{p^2} as $\mathbb{F}_p[i]/(i^2 + 1)$, we have $z^p = \bar{z}$ for all $z \in \mathbb{F}_{p^2}$, where \bar{z} denotes the *complex conjugate* of z . This ψ map on C/\mathbb{F}_{p^2} extends to give an endomorphism on $\text{Jac}(C)$, given (for general divisors) as

$$\psi : (u_1, u_0, v_1, v_0) \mapsto (\alpha \cdot \bar{u}_1, \beta \cdot \bar{u}_0, \gamma \cdot \bar{v}_1, \delta \cdot \bar{v}_0), \quad (1)$$

where $\alpha = u^{-2(p-1)}$, $\beta = u^{-4(p-1)}$, $\gamma = u^{-3(p-1)}$ and $\delta = u^{-5(p-1)}$ are all precomputed constants in \mathbb{F}_{p^2} . Besides the conjugations which are almost for free, it follows that the cost of computing ψ on general divisors is 4 \mathbb{F}_{p^2} -multiplications, and it is easily verified that the minimal polynomial of ψ on $\text{Jac}(C)$ is $\Phi_{20}(t) = t^8 - t^6 + t^4 - t^2 + 1$ [17, §5].

³ We also considered the prime $p = 2^{64} - 2^{32} + 1 \equiv 1 \pmod{4}$ which looks attractive for 32-bit platforms using NIST-like reduction, however our experiments showed that the Montgomery-friendly primes were faster.

Remark 1 (Higher powers of ψ). Scalar decompositions of dimension greater than 2 require the computation of higher powers of ψ on divisors. In all of our cases, applying ψ^i with $i > 1$ costs no more than applying ψ itself: we simply have a different tuple of 4 precomputed constants $(\alpha_i, \beta_i, \gamma_i, \delta_i) \in \mathbb{F}_{p^2}^4$ that allow us to compute ψ^i as in Eq. (1). In fact, applying even powers of ψ is always cheaper than odd powers, since for ψ^{2j} we always have $(\alpha_{2j}, \beta_{2j}, \gamma_{2j}, \delta_{2j}) \in \mathbb{F}_p^4$, so the multiplications required in (1) are now by base field elements. Additionally, for ψ^{2j} , we also have $\delta = (-1)^j$ which saves one such multiplication, and finally for even powers of ψ the complex conjugations undo themselves, which saves us performing negations. For 8-GLV/GLS, we need to apply powers of ψ up to ψ^7 , so we bear in mind the following order of preference (from cheapest to most expensive): (i) ψ^4 , (ii) $\{\psi^2, \psi^6\}$, and (iii) $\{\psi, \psi^3, \psi^5, \psi^7\}$.

4.2 Decomposing the Scalar

Let r be a large prime factor that divides the Jacobian group order of a BK curve C/\mathbb{F}_{p^2} and let D be a divisor of order r on $\text{Jac}(C)$. Since the minimal polynomial of ψ is $\Phi_{20}(t)$ (see Section 4.1), it follows that $\psi(D) = [\lambda]D$ where $\lambda < r \in \mathbb{Z}$ is a root of $t^8 - t^6 + t^4 - t^2 + 1 \equiv 0 \pmod{r}$. Park, Jeong and Lim [45] gave a simple algorithm that achieves GLV/GLS decompositions through division in the ring $\mathbb{Z}[\psi]$. The first step in this algorithm is to precompute a short vector in the GLV lattice L , which (in our 8-dimensional case) involves finding a short $a = (a_0, \dots, a_7) \in \mathbb{Z}^8$ in the lattice whose basis (matrix) has leading diagonal $(r, 1, \dots, 1) \in \mathbb{Z}^8$ and first column $(r, -\lambda, \dots, \lambda^7) \in \mathbb{Z}^8$, and where all other entries are zero. We then set $\alpha = \sum_{i=0}^7 a_i \cdot \psi^i$ and compute a quotient/remainder pair corresponding to the division k/α in $\mathbb{Z}[\psi]$, namely we find the quotient β and the remainder ρ such that $k = \beta\alpha + \rho$ in $\mathbb{Z}[\psi]$. The first observation here is that since $a \in L$, we have $\alpha D = \mathcal{O}$ for all D of order r , so that $[k]D = \beta\alpha D + \rho D = \rho D$. Since ρ is the remainder in the division by α , its coefficients in $\mathbb{Z}[\psi]$ are also small, so we write $\rho = \sum_{i=0}^7 k_i \psi^i$, from which our 8 mini-scalars are k_0, \dots, k_7 .

Algorithm 1 8-dimensional decomposition of the scalar k on Buhler-Koblitz curves over \mathbb{F}_{p^2} (read the algorithm from left to right and from top to bottom).

Input: The scalar k , the small constants $a_0, \dots, a_7 \in \mathbb{Z}$ and large constants $b_0, \dots, b_7, N \in \mathbb{Z}$.

Output: The mini-scalars k_0, \dots, k_7 .

$$\begin{aligned}
& y_0 \leftarrow \lfloor \frac{k \cdot b_0}{N} \rfloor, \quad y_1 \leftarrow \lfloor \frac{k \cdot b_1}{N} \rfloor, \quad y_2 \leftarrow \lfloor \frac{k \cdot b_2}{N} \rfloor, \quad y_3 \leftarrow \lfloor \frac{k \cdot b_3}{N} \rfloor, \quad y_4 \leftarrow \lfloor \frac{k \cdot b_4}{N} \rfloor, \quad y_5 \leftarrow \lfloor \frac{k \cdot b_5}{N} \rfloor, \quad y_6 \leftarrow \lfloor \frac{k \cdot b_6}{N} \rfloor, \quad y_7 \leftarrow \lfloor \frac{k \cdot b_7}{N} \rfloor, \\
& k_0 \leftarrow k, \quad u \leftarrow a_0 \cdot y_0, \quad k_0 \leftarrow k - u, \quad u \leftarrow a_0 \cdot y_1, \quad v \leftarrow a_1 \cdot y_0, \quad u \leftarrow u + v, \quad k_1 \leftarrow -u, \quad u \leftarrow a_2 \cdot y_0, \\
& v \leftarrow a_0 \cdot y_2, \quad u \leftarrow u + v, \quad v \leftarrow a_1 \cdot y_1, \quad u \leftarrow u + v, \quad k_2 \leftarrow -u, \quad u \leftarrow a_3 \cdot y_0, \quad v \leftarrow a_0 \cdot y_3, \quad u \leftarrow u + v, \\
& v \leftarrow a_1 \cdot y_2, \quad u \leftarrow u + v, \quad v \leftarrow a_2 \cdot y_1, \quad u \leftarrow u + v, \quad k_3 \leftarrow -u, \quad u \leftarrow a_0 \cdot y_4, \quad v \leftarrow a_4 \cdot y_0, \quad u \leftarrow u + v, \\
& v \leftarrow a_1 \cdot y_3, \quad u \leftarrow u + v, \quad v \leftarrow a_3 \cdot y_1, \quad u \leftarrow u + v, \quad v \leftarrow a_2 \cdot y_2, \quad u \leftarrow u + v, \quad k_4 \leftarrow -u, \quad u \leftarrow a_0 \cdot y_5, \\
& v \leftarrow a_5 \cdot y_0, \quad u \leftarrow u + v, \quad v \leftarrow a_1 \cdot y_4, \quad u \leftarrow u + v, \quad v \leftarrow a_4 \cdot y_1, \quad u \leftarrow u + v, \quad v \leftarrow a_2 \cdot y_3, \quad u \leftarrow u + v, \\
& v \leftarrow a_3 \cdot y_2, \quad u \leftarrow u + v, \quad k_5 \leftarrow -u, \quad u \leftarrow a_0 \cdot y_6, \quad v \leftarrow a_6 \cdot y_0, \quad u \leftarrow u + v, \quad v \leftarrow a_1 \cdot y_5, \quad u \leftarrow u + v, \\
& v \leftarrow a_5 \cdot y_1, \quad u \leftarrow u + v, \quad v \leftarrow a_2 \cdot y_4, \quad u \leftarrow u + v, \quad v \leftarrow a_4 \cdot y_2, \quad u \leftarrow u + v, \quad v \leftarrow a_3 \cdot y_3, \quad u \leftarrow u + v, \\
& k_6 \leftarrow -u, \quad u \leftarrow a_0 \cdot y_7, \quad v \leftarrow a_7 \cdot y_0, \quad u \leftarrow u + v, \quad v \leftarrow a_1 \cdot y_6, \quad u \leftarrow u + v, \quad v \leftarrow a_6 \cdot y_1, \quad u \leftarrow u + v, \\
& v \leftarrow a_2 \cdot y_5, \quad u \leftarrow u + v, \quad v \leftarrow a_5 \cdot y_2, \quad u \leftarrow u + v, \quad v \leftarrow a_3 \cdot y_4, \quad u \leftarrow u + v, \quad v \leftarrow a_4 \cdot y_3, \quad u \leftarrow u + v, \\
& k_7 \leftarrow -u, \quad u \leftarrow a_1 \cdot y_7, \quad v \leftarrow a_7 \cdot y_1, \quad u \leftarrow u + v, \quad v \leftarrow a_2 \cdot y_6, \quad u \leftarrow u + v, \quad v \leftarrow a_6 \cdot y_2, \quad u \leftarrow u + v, \\
& v \leftarrow a_3 \cdot y_5, \quad u \leftarrow u + v, \quad v \leftarrow a_4 \cdot y_4, \quad u \leftarrow u + v, \quad v \leftarrow a_5 \cdot y_3, \quad u \leftarrow u + v, \quad k_0 \leftarrow k_0 + u, \quad k_2 \leftarrow k_2 - u, \\
& k_4 \leftarrow k_4 + u, \quad k_6 \leftarrow k_6 - u, \quad u \leftarrow a_6 \cdot y_3, \quad v \leftarrow a_7 \cdot y_2, \quad u \leftarrow u + v, \quad v \leftarrow a_3 \cdot y_6, \quad u \leftarrow u + v, \quad v \leftarrow a_4 \cdot y_5, \\
& u \leftarrow u + v, \quad v \leftarrow a_5 \cdot y_4, \quad u \leftarrow u + v, \quad v \leftarrow a_2 \cdot y_7, \quad u \leftarrow u + v, \quad k_1 \leftarrow k_1 + u, \quad k_3 \leftarrow k_3 - u, \quad k_5 \leftarrow k_5 + u, \\
& k_7 \leftarrow k_7 - u, \quad u \leftarrow a_4 \cdot y_6, \quad v \leftarrow a_5 \cdot y_5, \quad u \leftarrow u + v, \quad v \leftarrow a_6 \cdot y_4, \quad u \leftarrow u + v, \quad v \leftarrow a_7 \cdot y_3, \quad u \leftarrow u + v, \\
& v \leftarrow a_3 \cdot y_7, \quad u \leftarrow u + v, \quad k_0 \leftarrow k_0 + u, \quad u \leftarrow a_7 \cdot y_4, \quad v \leftarrow a_6 \cdot y_5, \quad u \leftarrow u + v, \quad v \leftarrow a_4 \cdot y_7, \quad u \leftarrow u + v, \\
& v \leftarrow a_5 \cdot y_6, \quad u \leftarrow u + v, \quad k_1 \leftarrow k_1 + u, \quad u \leftarrow a_7 \cdot y_5, \quad v \leftarrow a_6 \cdot y_6, \quad u \leftarrow u + v, \quad v \leftarrow a_5 \cdot y_7, \quad u \leftarrow u + v, \\
& k_2 \leftarrow k_2 + u, \quad u \leftarrow a_7 \cdot y_6, \quad v \leftarrow a_6 \cdot y_7, \quad u \leftarrow u + v, \quad k_3 \leftarrow k_3 + u, \quad u \leftarrow a_7 \cdot y_7, \quad k_4 \leftarrow k_4 + u.
\end{aligned}$$

Besides the 8 precomputed “short” constants a_0, \dots, a_7 that must be input into the decomposition routine, there are 9 additional precomputed constants that aid a faster division [45]. Let $g(t) \in \mathbb{Z}[t]$ be the minimal polynomial of $\alpha \in \mathbb{Z}[\psi]$ with constant term N , so that we can write it as $g(t) = t \cdot h(t) + N$. We precompute $\hat{\alpha} = -h(\alpha) = \sum_{i=0}^7 b_i \psi^i$, which is N/α in $\mathbb{Z}[\psi]$. Along with the scalar k , we input the 8 values a_0, \dots, a_7 , the 8 values b_0, \dots, b_7 , and N into the decomposition algorithm from [45, §5.2], which we present in three-operand form in Algorithm 1. For each of the BK curves that are used in this work, we provide the shortest vector (a_0, \dots, a_7) in the GLV lattice L in Appendix A.2. The first line of Algorithm 1 shows the most non-trivial part of decomposing k on the fly, while the rest of the algorithm is straightforward. For $i = 0, \dots, 7$, we compute the division $y_i = \lfloor \frac{k \cdot b_i}{N} \rfloor$ using only integer operations as follows. We find the smallest b' such that $N < 2^{bb'}$, where b is the width of the machine word-size (32 or 64 in practice). We then precompute $\ell_i = \lfloor \frac{2^{bb'} \cdot b_i}{N} \rfloor \geq 0$, so that the division can now be computed as $y_i = \lfloor \frac{\ell_i \cdot k}{2^{bb'}} \rfloor$. The division by $2^{bb'}$ comes for free: it can be implemented by a shift of the machine words of the results. Depending on the sign of k , the result can be off by one due to the rounding, but in practice this does not influence the size of the mini-scalars.

4.3 Constructing the Lookup Table

After the scalar k is decomposed into 8 mini-scalars $k_i < 2^m$, each corresponding to the divisor $D_i = \text{sign}(k_i) \cdot \psi^i(D)$, following the standard approach [17, 36, 8] (for 2- and 4-dimensional decompositions) would mean computing the scalar multiplication by first precomputing a lookup table $L[i] = \sum_{\ell=0}^7 (\lfloor \frac{i}{2^\ell} \rfloor \bmod 2) \cdot D_\ell$, for $0 \leq i < 2^8$. When simultaneously processing the j^{th} bit of each of the mini-scalars, the precomputed multiple $L[i]$ is added to the accumulator of the main loop, for $i = \sum_{\ell=0}^7 2^\ell \left(\lfloor \frac{k_\ell}{2^j} \rfloor \bmod 2 \right)$. The advantage here is that only one doubling and one addition are used for each of the m bits in the mini-scalar. The precomputation phase, computing the entries of the $L[i]$, is relatively inexpensive for 2- and 4-dimensional GLV/GLS. In the setting of 8-GLV/GLS however, computing these $2^8 = 256$ entries is computationally significant: roughly speaking, constructing this full-sized lookup table would be as expensive as computing the scalar multiplication in the generic way (i.e. not using endomorphisms). An observation is that in practice m is usually small ($m < 34$), so that we do not need to precompute the entire table and we can compute the required entries on-the-fly. Unfortunately, computing a random table element might require multiple additions (in the worst case) and no performance gain can be expected when using this approach.

We present two different approaches which solve this problem (which can both be seen as an extension to the approach described in [34], but in the special case of two tables). Both approaches generate *two* lookup tables consisting of $2^4 = 16$ elements each. So instead of computing the single large table L , one can compute two significantly smaller tables T_1 and T_2 such that $T_1[i] = \sum_{\ell=0}^3 (\lfloor \frac{i}{2^\ell} \rfloor \bmod 2) \cdot D_\ell$ and $T_2[i] = \sum_{\ell=0}^3 (\lfloor \frac{i}{2^\ell} \rfloor \bmod 2) \cdot D_{\ell+4}$, for $0 \leq i < 2^4$. This has the advantage of significantly lowering the precomputation cost of the tables, but increases the number of “per bit” curve additions from one to two when processing the mini-scalars. The two methods we present differ in how the tables are generated: the first approach is slightly slower than the second, but has the advantage that it runs in constant time.

The constant-time approach. The straight-forward approach to generate the two lookup tables T_1 and T_2 is to first compute $T_i[j]$ for $i \in \{1, 2\}$ and $j \in \{1, 2, 4, 8\}$ using (at most) the

Table 2. Generating the lookup table in constant time for 8-dimensional GLV/GLS, where the divisors D_i are computed efficiently from D and k_i sequentially as follows (the cost is stated in the table): $D_0 = D$, $D_1 = \phi(D_0)$, $D_i = \phi^2(D_{i-2})$ for $i \in \{2, 3\}$, $D_i = \phi^4(D_{i-4})$ for $i \in \{4, 5, 6, 7\}$, $D_i = \text{sign}(k_i) \cdot D_i$ for $0 \leq i < 8$. The second argument in the mixed sums is the affine divisor.

operation	D_3	D_2	D_1	D_0	method	operation	D_7	D_6	D_5	D_4	method
$T_1[0] \leftarrow \mathcal{O}$	0	0	0	0	-	$T_2[0] \leftarrow \mathcal{O}$	0	0	0	0	-
$T_1[1] \leftarrow D_0$	0	0	0	1	-	$T_2[1] \leftarrow D_4$	0	0	0	1	ψ^4
$T_1[2] \leftarrow D_1$	0	0	1	0	ψ	$T_2[2] \leftarrow D_5$	0	0	1	0	ψ^4
$T_1[3] \leftarrow T_1[1] + T_1[2]$	0	0	1	1	AFF	$T_2[3] \leftarrow T_2[1] + T_2[2]$	0	0	1	1	AFF
$T_1[4] \leftarrow D_2$	0	1	0	0	ψ^2	$T_2[4] \leftarrow D_6$	0	1	0	0	ψ^4
$T_1[5] \leftarrow T_1[1] + T_1[4]$	0	1	0	1	AFF	$T_2[5] \leftarrow T_2[1] + T_2[4]$	0	1	0	1	AFF
$T_1[6] \leftarrow T_1[2] + T_1[4]$	0	1	1	0	AFF	$T_2[6] \leftarrow T_2[2] + T_2[4]$	0	1	1	0	AFF
$T_1[7] \leftarrow T_1[6] + T_1[1]$	0	1	1	1	MIX	$T_2[7] \leftarrow T_2[6] + T_2[1]$	0	1	1	1	MIX
$T_1[8] \leftarrow D_3$	1	0	0	0	ψ^2	$T_2[8] \leftarrow D_7$	1	0	0	0	ψ^4
$T_1[9] \leftarrow T_1[1] + T_1[8]$	1	0	0	1	AFF	$T_2[9] \leftarrow T_2[1] + T_2[8]$	1	0	0	1	AFF
$T_1[10] \leftarrow T_1[2] + T_1[8]$	1	0	1	0	AFF	$T_2[10] \leftarrow T_2[2] + T_2[8]$	1	0	1	0	AFF
$T_1[11] \leftarrow T_1[10] + T_1[1]$	1	0	1	1	MIX	$T_2[11] \leftarrow T_2[10] + T_2[1]$	1	0	1	1	MIX
$T_1[12] \leftarrow T_1[8] + T_1[4]$	1	1	0	0	AFF	$T_2[12] \leftarrow T_2[8] + T_2[4]$	1	1	0	0	AFF
$T_1[13] \leftarrow T_1[12] + T_1[1]$	1	1	0	1	MIX	$T_2[13] \leftarrow T_2[12] + T_2[1]$	1	1	0	1	MIX
$T_1[14] \leftarrow T_1[12] + T_1[2]$	1	1	1	0	MIX	$T_2[14] \leftarrow T_2[12] + T_2[2]$	1	1	1	0	MIX
$T_1[15] \leftarrow T_1[14] + T_1[1]$	1	1	1	1	MIX	$T_2[15] \leftarrow T_2[14] + T_2[1]$	1	1	1	1	MIX

ψ map for each computation – we prioritize higher even powers of ψ following Remark 1. Next, the other elements are computed as $T_i[j] = T_i[k] + T_i[j - k]$ for $j > 1$ and $k < j$, and where k is chosen so that the fastest possible formulas can be applied each time. Namely, elements that are obtained by using an addition become projective divisors, whilst elements that are in $T_{1,2}[j]$ for $j \in \{1, 2, 4, 8\}$ (which are computed using the ψ map) are affine. Adding two affine divisors together to give a projective divisor is faster than performing a mixed addition between an affine and projective divisor, so we prioritize this affine-only addition where possible. We modified the formulas for the mixed-addition operation to formulas for an affine-affine addition operation, which are given in Algorithm 2. Compared to mixed-additions, this lowers the required number of multiplications in \mathbb{F}_{p^2} from 37 to 29. We denote the operations of projective doubling, projective addition, mixed addition and addition between two affine divisors by **DBL**, **ADD**, **MIX** and **AFF** respectively. Table 2 outlines our approach to compute both lookup tables in constant time. Table 3 summarizes the total cost for both the precomputation of the lookup tables and the computation of the scalar of the 8-GLV/GLS routine as a function of the maximum bit-length m of the mini-scalars k_i . We use **m**, **s** and **a** to denote the costs of computing multiplications, squarings and additions in \mathbb{F}_{p^2} respectively.

Recycling inputs to MIX, AFF. Additional (but minor) savings are possible by recycling some overlapping operations in the computation of T_1 and T_2 . The same divisors are often used as input for the mixed addition (**MIX**) and affine addition (**AFF**), which each contain a few operations that only depend on one of the inputs. For example, focusing only on the second input into Algorithm 2, the operations $t_5 \leftarrow u'_1 \cdot u'_0$, $t_6 \leftarrow u_1'^2$ which are required in the computation of $T_1[3] = T_1[2] + T_1[1]$ can be reused 6 more times in $T_1[2i + 1]$ for $i = 2, \dots, 7$, and similarly for T_2 . This accounts for all of the odd indexes, while $T_1[2]$ and/or $T_1[4]$ can be recycled into all of the other entries with even indexes, and similarly for T_2 .

Using ψ to speed up precomputations. If we are not concerned with implementations which need to run in constant time and aim to optimize for performance only, then the endomorphism ψ can be used to accelerate the computation of T_1 and T_2 . The reason we can

Algorithm 2 Addition between two general affine divisors to give a projective divisor on the Jacobian of $C : y^2 = x^5 + f_3x^3 + f_2x^2 + f_1x + f_0$ (read the algorithm from left to right and from top to bottom).

Input: $P = (u_1, u_0, v_1, v_0)$ and $Q = (u'_1, u'_0, v'_1, v'_0)$.

Output: $P + Q = (U''_1, U''_0, V''_1, V''_0, Z'')$.

$$\begin{aligned}
V''_0 &\leftarrow v'_0 - v_0, & t_1 &\leftarrow v_1 + v'_1, & t_2 &\leftarrow v_1 - v'_1, & t_3 &\leftarrow u_1 + u'_1, & t_4 &\leftarrow u_1 \cdot u_0, & t_5 &\leftarrow u'_1 \cdot u'_0, & t_5 &\leftarrow t_5 - t_4, \\
U''_0 &\leftarrow u'_0 - u_0, & t_4 &\leftarrow u_1^2, & t_6 &\leftarrow u_1'^2, & U''_1 &\leftarrow u_1 - u'_1, & t_7 &\leftarrow t_4 - t_6, & t_4 &\leftarrow t_4 + t_6, & t_7 &\leftarrow U''_0 + t_7, \\
t_6 &\leftarrow t_5 \cdot t_2, & t_2 &\leftarrow U''_0 \cdot t_2, & U''_0 &\leftarrow U''_0 \cdot t_7, & t_5 &\leftarrow t_5 \cdot U''_1, & U''_1 &\leftarrow V''_0 \cdot U''_1, & t_7 &\leftarrow V''_0 \cdot t_7, & t_6 &\leftarrow t_6 - t_7, \\
U''_0 &\leftarrow U''_0 - t_5, & t_5 &\leftarrow t_2 - U''_1, & t_7 &\leftarrow U''_0'^2, & t_2 &\leftarrow t_5, & U''_0 &\leftarrow U''_0 \cdot t_5, & t_2 &\leftarrow t_2^2, & V''_1 &\leftarrow t_2 \cdot u_1, \\
t_4 &\leftarrow t_2 \cdot t_4, & V''_0 &\leftarrow t_6 \cdot t_5, & t_5 &\leftarrow t_5^2, & t_6 &\leftarrow t_6^2, & t_8 &\leftarrow u_0 \cdot t_5, & Z'' &\leftarrow U''_0 \cdot t_5, & U''_1 &\leftarrow 2 \cdot V''_0, \\
U''_1 &\leftarrow U''_1 - t_7, & t_1 &\leftarrow U''_0 \cdot t_1, & t_6 &\leftarrow t_6 + t_1, & t_1 &\leftarrow t_2 \cdot t_3, & U''_1 &\leftarrow U''_1 - t_1, & t_7 &\leftarrow U''_1 - t_7, & t_3 &\leftarrow t_3 \cdot t_7, \\
t_3 &\leftarrow t_3 + t_4, & t_3 &\leftarrow t_3/2, & t_6 &\leftarrow t_6 - t_3, & t_7 &\leftarrow V''_1 - V''_0, & V''_0 &\leftarrow V''_0 - U''_1, & t_4 &\leftarrow t_6 - t_8, & V''_1 &\leftarrow V''_1 \cdot t_7, \\
t_2 &\leftarrow t_2 \cdot t_4, & t_2 &\leftarrow t_2 + V''_1, & V''_1 &\leftarrow U''_1 \cdot V''_0, & V''_1 &\leftarrow V''_1 + t_2, & t_8 &\leftarrow t_8 \cdot t_7, & V''_0 &\leftarrow V''_0 \cdot t_6, & V''_0 &\leftarrow t_8 + V''_0, \\
t_8 &\leftarrow Z'' \cdot v_1, & V''_1 &\leftarrow V''_1 - t_8, & U''_1 &\leftarrow U''_1 \cdot U''_0, & U''_0 &\leftarrow t_6 \cdot U''_0, & t_6 &\leftarrow Z'' \cdot v_0, & V''_0 &\leftarrow V''_0 - t_6,
\end{aligned}$$

not use ψ in the same way for each scalar is that its usefulness and applicability depends on the signs of the k_i , which change each time. We use an example to illustrate: define $s_i = \mathbf{sign}(k_i) \in \{-1, +1\}$, and suppose that after computing D_0, \dots, D_7 (which are negated according to the signs of k_0, \dots, k_7), we compute $T_1[3] \leftarrow T_1[1] + T_1[2] = D_0 + D_1$. When computing $T_1[6]$, which is usually computed using an affine addition as $T_1[6] = T_1[2] + T_1[4] = D_1 + D_2$, we can possibly use ψ to compute $D_1 + D_2$. If the signs s_0, s_1, s_2 are equal then $T_1[6] = D_1 + D_2 = \psi(D_0 + D_1)$, while if s_0 and s_2 are equal and $s_1 = -s_0$, then $T_1[6] = D_1 + D_2 = -\psi(D_0 + D_1)$. Alternatively, if $s_0 \neq s_2$, then we still need (at least) one addition on top of $\psi(D_0 + D_1)$ to compute $D_1 + D_2$ and so using the original addition between $T_1[2]$ and $T_1[4]$ is preferred.

In Algorithm 3 we outline the complete strategy which exhausts each possibility of using ψ to recycle prior computations before resorting to a divisor addition. As in the above example, the usefulness of previous values is completely dependent on the combinations of the associated signs. As we proceed further into the algorithm, the chances of reusing previous computations generally increases. For example, $T_2[12]$ would ordinarily require the addition $T_2[12] = T_2[8] + T_2[4] = D_7 + D_6$, but it could also possibly be computed as any of $\psi(D_6 + D_5)$, $\psi^2(D_5 + D_4)$, $\psi^4(D_3 + D_2)$, $\psi^5(D_2 + D_1)$ or $\psi^6(D_1 + D_0)$, depending on

Table 3. Summary of costs for a single 8-GLV/GLS scalar multiplication, where $\max\{|k_i|\} < 2^m$. The left side of the table gives the cost and number of occurrences of the 5 divisor operations used, which are combined to give a total cost of 8-GLV/GLS in terms of m . For each of the implementations in this work, the right side of the table uses the average value of m to give the average number of multiplications, squarings and additions required in 8-GLV/GLS. While the costs reported correspond to the simple, constant-time precomputation strategy, the final column on the right side gives the number of additions (both mixed and affine) that are replaced with ψ 's if the faster precomputation strategy is employed. All averages were taken over 10 million scalar decompositions.

divisor	formulas found in	cost per. operation	# ops table	# ops scalar	curve	av. m	average cost [m, s, a]	av. ψ 's
DBL	[8, Alg. 1]	$36\mathbf{m} + 6\mathbf{s} + 34\mathbf{a}$	-	m	p_{611}	26.43	[4039, 490, 3070]	9.15
ADD	[8, Alg. 2]	$44\mathbf{m} + 4\mathbf{s} + 29\mathbf{a}$	-	$2m$	Mont. (i)	27.53	[4176, 505, 3171]	8.90
MIX	[8, Alg. 3]	$37\mathbf{m} + 5\mathbf{s} + 29\mathbf{a}$	12	-	Mont. (ii)	31.10	[4618, 555, 3499]	8.41
AFF	Alg. 2	$29\mathbf{m} + 6\mathbf{s} + 29\mathbf{a}$	10	-	NIST (i)	27.84	[4214, 509, 3199]	9.25
ψ	Eq. (1)	$4\mathbf{m}$	7	-	NIST (ii)	31.71	[4693, 563, 3554]	8.41

total: $(124m + 762)\mathbf{m} + (14m + 120)\mathbf{s} + (92m + 638)\mathbf{a}$

Algorithm 3 The fast (non-constant time) approach to generate the lookup table for 8-dimensional GLV/GLS.

Input: The eight divisors $D_i = \text{sign}(k_i) \cdot \psi^i(D)$ corresponding to the eight mini-scalars k_i .

Output: $\begin{cases} T_1[i] = \sum_{\ell=0}^3 (\lfloor \frac{i}{2^\ell} \rfloor \bmod 2) D_\ell, \\ T_2[i] = \sum_{\ell=0}^3 (\lfloor \frac{i}{2^\ell} \rfloor \bmod 2) D_{\ell+4}, \text{ for } 0 \leq i < 2^4 \end{cases}$

```

1:  $T_1[0] \leftarrow \mathcal{O}$ 
2:  $T_1[1] \leftarrow D_0$ 
3:  $T_1[2] \leftarrow D_1$ 
4:  $T_1[3] \leftarrow D_0 + D_1$ 
5:  $T_1[4] \leftarrow D_2$ 
6:  $T_1[5] \leftarrow D_2 + D_0$ 
7: if  $s_0 = s_2$  then
8:    $T_1[6] \leftarrow \psi^1(T_1[3])$ 
9:   if  $s_0 \neq s_1$  then
10:      $T_1[6] \leftarrow -T_1[6]$ 
11: else
12:    $T_1[6] \leftarrow D_1 + D_2$ 
13:  $T_1[7] \leftarrow T_1[6] + T_1[1]$ 
14:  $T_1[8] \leftarrow D_3$ 
15:  $T_1[9] \leftarrow D_0 + D_3$ 
16: if  $s_0 \cdot s_2 = s_1 \cdot s_3$  then
17:    $T_1[10] \leftarrow \psi^1(T_1[5])$ 
18:   if  $s_0 \neq s_1$  then
19:      $T_1[10] \leftarrow -T_1[10]$ 
20: else
21:    $T_1[10] \leftarrow D_1 + D_3$ 
22:  $T_1[11] \leftarrow T_1[10] + T_1[1]$ 
23: if  $s_0 \cdot s_1 = s_2 \cdot s_3$  then
24:    $T_1[12] \leftarrow \psi^2(T_1[3])$ 
25:   if  $s_0 \neq s_2$  then
26:      $T_1[12] \leftarrow -T_1[12]$ 
27: else if  $s_1 = s_3$  then
28:    $T_1[12] \leftarrow \psi^1(T_1[6])$ 
29:   if  $s_1 \neq s_2$  then
30:      $T_1[12] \leftarrow -T_1[12]$ 
31: else
32:    $T_1[12] \leftarrow D_2 + D_3$ 
33:  $T_1[13] \leftarrow T_1[12] + T_1[1]$ 
34: if  $s_0 = s_1 = s_2 = s_3$  then
35:    $T_1[14] \leftarrow \psi^1(T_1[7])$ 
36: else
37:    $T_1[14] \leftarrow T_1[12] + T_1[2]$ 
38:  $T_1[15] \leftarrow T_1[14] + T_1[1]$ 
39:  $T_2[0] \leftarrow \mathcal{O}$ 
40:  $T_2[1] \leftarrow D_4$ 
41:  $T_2[2] \leftarrow D_5$ 
42: if  $s_0 \cdot s_1 = s_4 \cdot s_5$  then
43:    $T_2[3] \leftarrow \psi^4(T_1[3])$ 
44:   if  $s_0 \neq s_4$  then
45:      $T_2[3] \leftarrow -T_2[3]$ 
46: else if  $s_2 \cdot s_3 = s_4 \cdot s_5$  then
47:    $T_2[3] \leftarrow \psi^2(T_1[12])$ 
48:   if  $s_2 \neq s_4$  then
49:      $T_2[3] \leftarrow -T_2[3]$ 
50: else if  $s_1 \cdot s_2 = s_4 \cdot s_5$  then
51:    $T_2[3] \leftarrow \psi^3(T_1[6])$ 
52:   if  $s_1 = s_4$  then
53:      $T_2[3] \leftarrow -T_2[3]$ 
54: else
55:    $T_2[3] \leftarrow D_4 + D_5$ 
56:  $T_2[4] \leftarrow D_6$ 
57: if  $s_0 \cdot s_2 = s_4 \cdot s_6$  then
58:    $T_2[5] \leftarrow \psi^4(T_1[5])$ 
59:   if  $s_0 \neq s_4$  then
60:      $T_2[5] \leftarrow -T_2[5]$ 
61: else if  $s_1 \cdot s_3 = s_4 \cdot s_6$  then
62:    $T_2[5] \leftarrow \psi^3(T_1[10])$ 
63:   if  $s_1 \neq s_4$  then
64:      $T_2[5] \leftarrow -T_2[5]$ 
65: else
66:    $T_2[5] \leftarrow D_4 + D_6$ 
67: if  $s_1 \cdot s_2 = s_5 \cdot s_6$  then
68:    $T_2[6] \leftarrow \psi^4(T_1[6])$ 
69:   if  $s_1 \neq s_5$  then
70:      $T_2[6] \leftarrow -T_2[6]$ 
71: else if  $s_4 = s_6$  then
72:    $T_2[6] \leftarrow \psi^1(T_2[3])$ 
73:   if  $s_4 \neq s_5$  then
74:      $T_2[6] \leftarrow -T_2[6]$ 
75: else if  $s_2 \cdot s_3 = s_5 \cdot s_6$  then
76:    $T_2[6] \leftarrow \psi^3(T_1[12])$ 
77:   if  $s_2 \neq s_5$  then
78:      $T_2[6] \leftarrow -T_2[6]$ 
79: else if  $s_0 \cdot s_1 = s_5 \cdot s_6$  then
80:    $T_2[6] \leftarrow \psi^5(T_1[3])$ 
81:   if  $s_0 \neq s_5$  then
82:      $T_2[6] \leftarrow -T_2[6]$ 
83: else
84:    $T_2[6] \leftarrow D_5 + D_6$ 
85: if  $s_0 \cdot s_4 = s_1 \cdot s_5 = s_2 \cdot s_6$  then
86:    $T_2[7] \leftarrow \psi^4(T_1[7])$ 
87:   if  $s_0 \neq s_4$  then
88:      $T_2[7] \leftarrow -T_2[7]$ 
89: else if  $s_1 \cdot s_4 = s_2 \cdot s_5 = s_3 \cdot s_6$  then
90:    $T_2[7] \leftarrow \psi^3(T_1[14])$ 
91:   if  $s_1 \neq s_4$  then
92:      $T_2[7] \leftarrow -T_2[7]$ 
93: else
94:    $T_2[7] \leftarrow T_2[6] + T_2[1]$ 
95:  $T_2[8] \leftarrow D_7$ 
96: if  $s_0 \cdot s_4 = s_3 \cdot s_7$  then
97:    $T_2[9] \leftarrow \psi^4(T_1[9])$ 
98:   if  $s_0 \neq s_4$  then
99:      $T_2[9] \leftarrow -T_2[9]$ 
100: else
101:    $T_2[9] \leftarrow T_2[1] + T_2[8]$ 
102: if  $s_1 \cdot s_5 = s_3 \cdot s_7$  then
103:    $T_2[10] \leftarrow \psi^4(T_1[10])$ 
104:   if  $s_1 \neq s_5$  then
105:      $T_2[10] \leftarrow -T_2[10]$ 
106: else if  $s_4 \cdot s_5 = s_6 \cdot s_7$  then
107:    $T_2[10] \leftarrow \psi^1(T_2[5])$ 
108:   if  $s_4 \neq s_5$  then
109:      $T_2[10] \leftarrow -T_2[10]$ 
110: else if  $s_0 \cdot s_5 = s_2 \cdot s_7$  then
111:    $T_2[10] \leftarrow \psi^5(T_1[5])$ 
112:   if  $s_0 \neq s_5$  then
113:      $T_2[10] \leftarrow -T_2[10]$ 
114: else
115:    $T_2[10] \leftarrow D_5 + D_7$ 
116: if  $s_0 \cdot s_4 = s_1 \cdot s_5 = s_3 \cdot s_7$  then
117:    $T_2[11] \leftarrow \psi^4(T_1[11])$ 
118:   if  $s_0 \neq s_4$  then
119:      $T_2[11] \leftarrow -T_2[11]$ 
120: else
121:    $T_2[11] \leftarrow T_2[10] + T_2[1]$ 
122: if  $s_2 \cdot s_6 = s_3 \cdot s_7$  then
123:    $T_2[12] \leftarrow \psi^4(T_1[12])$ 
124:   if  $s_2 \neq s_6$  then
125:      $T_2[12] \leftarrow -T_2[12]$ 
126: else if  $s_4 \cdot s_6 = s_5 \cdot s_7$  then
127:    $T_2[12] \leftarrow \psi^2(T_2[3])$ 
128:   if  $s_4 \neq s_6$  then
129:      $T_2[12] \leftarrow -T_2[12]$ 
130: else if  $s_0 \cdot s_6 = s_1 \cdot s_7$  then
131:    $T_2[12] \leftarrow \psi^6(T_1[3])$ 
132:   if  $s_0 \neq s_6$  then
133:      $T_2[12] \leftarrow -T_2[12]$ 
134: else if  $s_5 = s_7$  then
135:    $T_2[12] \leftarrow \psi^1(T_2[6])$ 
136:   if  $s_5 \neq s_6$  then
137:      $T_2[12] \leftarrow -T_2[12]$ 
138: else if  $s_1 \cdot s_6 = s_2 \cdot s_7$  then
139:    $T_2[12] \leftarrow \psi^5(T_1[6])$ 
140:   if  $s_1 \neq s_6$  then
141:      $T_2[12] \leftarrow -T_2[12]$ 
142: else
143:    $T_2[12] \leftarrow D_6 + D_7$ 
144: if  $s_0 \cdot s_4 = s_2 \cdot s_6 = s_3 \cdot s_7$  then
145:    $T_2[13] \leftarrow \psi^4(T_1[13])$ 
146:   if  $s_0 \neq s_4$  then
147:      $T_2[13] \leftarrow -T_2[13]$ 
148: else
149:    $T_2[13] \leftarrow T_2[12] + T_2[1]$ 
150: if  $s_1 \cdot s_5 = s_2 \cdot s_6 = s_3 \cdot s_7$  then
151:    $T_2[14] \leftarrow \psi^4(T_1[14])$ 
152:   if  $s_1 \neq s_5$  then
153:      $T_2[14] \leftarrow -T_2[14]$ 
154: else if  $s_4 = s_5 = s_6 = s_7$  then
155:    $T_2[14] \leftarrow \psi^1(T_2[7])$ 
156: else if  $s_0 \cdot s_5 = s_1 \cdot s_6 = s_2 \cdot s_7$  then
157:    $T_2[14] \leftarrow \psi^5(T_1[7])$ 
158:   if  $s_0 \neq s_5$  then
159:      $T_2[14] \leftarrow -T_2[14]$ 
160: else
161:    $T_2[14] \leftarrow T_2[12] + T_2[2]$ 
162: if  $s_0 \cdot s_4 = s_1 \cdot s_5 = s_2 \cdot s_6 = s_3 \cdot s_7$  then
163:    $T_2[15] \leftarrow \psi^4(T_1[15])$ 
164:   if  $s_0 \neq s_4$  then
165:      $T_2[15] \leftarrow -T_2[15]$ 
166: else
167:    $T_2[15] \leftarrow T_2[14] + T_2[1]$ 

```

whether the associated s_i align favorably. Again, we prioritize the possible application of even powers of ψ according to the hierarchy discussed in Remark 1. We note that any-time ψ is used to recycle previously computed sums, they are now acting on projective (instead of affine) divisors. This requires an updated description of ψ , which is given as $\psi : (U_1 : U_0 : V_1 : V_0 : Z) \mapsto (\alpha \cdot \bar{U}_1 : \beta \cdot \bar{U}_0 : \gamma \cdot \bar{V}_1 : \delta \cdot \bar{V}_0 : \bar{Z})$, for which the only difference from the affine version in Eq. (1) is that the Z coordinate must also be conjugated. We point out that Remark 1 applies identically to the projective case. Of the 22 additions that would otherwise be required, the final column in the right part of Table 3 gives the average number of additions that are replaced by ψ 's in the six different 8-GLV/GLS scenarios we implemented⁴. In all cases this gives over a 30% speedup when constructing the lookup table.

5 Arithmetic

In this paper we are concerned with arithmetic modulo primes $p < 2^{64}$ to realize scalar multiplications in $\text{Jac}_C(\mathbb{F}_{p^2})$. We optimize this arithmetic on two different levels: on the one hand the extension field arithmetic in \mathbb{F}_{p^2} is optimized in terms of multiplications in \mathbb{F}_p , and on the other hand we aim to optimize the multiplications in \mathbb{F}_p by choosing p such that modular reduction is particularly efficient. On architectures where the 64-bit modulus p fits in a single machine word, the modular multiplication can be computed by doing the multiplication first, followed by a NIST-like reduction [48, 51]. Other popular embedded platforms, like the ARM, have a smaller machine word size of 32 bits. Since representing the prime p requires two such words, other techniques (besides the NIST-like reduction) might be attractive to explore. Following the observations from [8], we choose the primes p to be Montgomery-friendly to accelerate the implementation of the modular arithmetic on such 32-bit platforms. Since the use of Montgomery-friendly primes only makes sense when the prime can be represented by two or more machine words, our approach for 64-bit architectures follows the more conventional NIST-like reduction.

5.1 Modular Arithmetic

NIST-like reduction. It is well-known that modular reduction can be computed efficiently, without using any multiplications, when the modulus has a special form. Typically, the mod-

⁴ The differences in these averages is due to the fact that the scalars in each case are all chosen modulo different primes, meaning that the signs of the eight k_i behave slightly differently in each case.

Algorithm 4 This algorithm, including Line 1 and Line 3, computes the radix- 2^b interleaved Montgomery multiplication [39] ($\text{MontMul}(A, B, p) = A \cdot B \cdot 2^{-bn} \bmod p$) for an n -word modulus p . Excluding Line 1 and Line 3 gives the algorithm for computing the radix- 2^b Montgomery reduction only ($\text{MontRed}(C, p) = C \cdot 2^{-bn} \bmod p$).

Input: $\left\{ \begin{array}{l} (A = \sum_{i=0}^{n-1} a_i 2^{bi}, B) \text{ or } C \text{ and } p, \mu \text{ such that } 0 \leq a_i < 2^b, 0 \leq A \leq S_0 < 2^{bn}, 0 \leq B \leq S_1 < 2^{bn}, \\ 0 \leq C \leq S_1 S_2 < 2^{2bn}, 2^{b(n-1)} \leq p < 2^{bn}, 2 \nmid p, \mu = -p^{-1} \bmod 2^b, \end{array} \right.$

Output: $(C' \equiv A \cdot B \cdot 2^{-bn} \bmod p)$ or $(C' \equiv C \cdot 2^{-bn} \bmod p)$ such that $0 \leq C' < r_{(b,n)}(S_0 S_1, p)$

- 1: $[C \leftarrow 0]$
 - 2: **for** $i = 0$ to $n - 1$ **do**
 - 3: $[C \leftarrow C + a_i \cdot B]$
 - 4: $q \leftarrow \mu \cdot C \bmod 2^b, C \leftarrow (C + q \cdot p) / 2^b$
 - 5: **return** $C' \leftarrow C$
-

ular multiplication and the modular reduction are computed sequentially. An example of a family of such primes are generalized Mersenne primes, whose adoption usually results in significant performance gains; this is why NIST has standardized multiple instances of such primes [51]. Let us illustrate the basic idea with $p_{611} = 2^{61} - 1$, which belongs to this class of primes. Computing the modular multiplication $c \equiv a \cdot b \pmod{p_{611}}$ with $0 \leq a, b < p_{611}$ can be done by first computing the multiplication and shifting this value (note that the result still fits in a 128-bit data-type) as $t = (2^3 \cdot a) \cdot b$. Due to the special form of p_{611} , we have $t = t_1 \cdot 2^{64} + t_0 \equiv t_1 \cdot 2^{64} + t_0 - t_1 \cdot 2^3 \cdot p_{611} \equiv t_0 + 2^3 \cdot t_1 \pmod{p_{611}}$, for $0 \leq t_0, t_1 < 2^{64}$, and hence we can compute the reduction as $\left(\lfloor t/2^{64} \rfloor + (t \bmod 2^{64})/2^3\right) \pmod{p_{611}}$. Since $0 \leq \lfloor t/2^{64} \rfloor, (t \bmod 2^{64})/2^3 < 2^{61}$, we can use these integers as input to a modular addition to reduce the result properly to the range $[0, p_{611})$. For numbers of the form $2^x - 1$, modular addition is especially efficient, since if $c = a + b$, where $0 \leq a, b < 2^x - 1$, then $c' = \lfloor \frac{c}{2^x} \rfloor + c - 2^x \equiv c \pmod{2^x - 1}$, where c' is properly reduced and can be computed using only a **shift**, an **add** and a **bit-reset** instruction (and possibly data movements).

Montgomery arithmetic. Montgomery proposed a new way of computing modular multiplication in the mid 1980s [39]. The idea behind Montgomery multiplication is to replace the relatively expensive divisions by computationally inexpensive logical shifts on computers, lowering the computational complexity by a constant factor compared to the classical method. We present the algorithm for a computer platform which works on b -bit ($b > 2$) words: i.e. we use a 2^b -radix system. Montgomery multiplication modulo an odd n -word moduli p , $2^{(n-1)b} \leq p < 2^{nb}$, is computed by transforming each of the operands to its Montgomery residue $\tilde{A} = A \cdot 2^{bn} \pmod{p}$. Montgomery multiplication is defined as $\tilde{C} \equiv \tilde{A} \cdot \tilde{B} \cdot 2^{-bn} \equiv C \cdot 2^{bn} \pmod{p}$. Algorithm 5, including the lines in brackets, outlines interleaved Montgomery multiplication, while if the bracketed lines are excluded this computes the Montgomery reduction only. Note that modular addition and subtraction can be done in the usual way when working with Montgomery residues since $\tilde{A} \pm \tilde{B} \equiv (A \pm B) \cdot 2^{bn} \equiv \widetilde{A \pm B} \pmod{p}$. The result of the Montgomery multiplication of two positive integers $\tilde{A} \leq S_0$ and $\tilde{B} \leq S_1$ can be bounded by $r_{(b,n)}(S_0 \cdot S_1, p) = \frac{S_0 S_1}{2^{bn}} + p$. Hence, if both inputs are bounded by 2^{bn} , then the result is at most $r_{(b,n)}(2^{2bn}, p) = 2^{2bn} + p$: a conditional subtraction with p is required when the output is required to be less than 2^{2bn} . It follows that if both inputs are bounded by 2^{bn-1} and $2^{b(n-1)} \leq p < 2^{2bn-2}$, then this conditional subtraction can be omitted since $r_{(b,n)}(2^{2(bn-1)}, p) = 2^{2bn-2} + p < 2^{2bn-1}$, and the output of Montgomery multiplication can be reused as input directly (this is the idea behind subtraction-less Montgomery multiplication [52]).

Montgomery-friendly arithmetic. The idea behind Montgomery-friendly primes [33, 29, 1, 25, 8] is to reduce the number of multiplications and registers used by taking $\mu = -p^{-1} \pmod{2^b} = \pm 1$; this is achieved when $p \equiv \mp 1 \pmod{2^b}$. Note that all NIST primes, as standardized in [51], have this property for $b \leq 32$. The number of multiplications can be reduced further when the $(n-1)$ most significant words of p have a special form, such that multiplication by p can be transformed into a sequence of shifts and additions or subtractions (just as in the NIST-like reduction). For $b = 32$, examples of Montgomery-friendly primes are those primes of the form $(2^{31} - c) \cdot 2^{32} - 1$, with $0 \leq c < 2^{31}$, as mentioned in Section 3.2. Here we intentionally use 63-bit primes (instead of the full double-word length of 64 bits) to allow accumulation in the Montgomery reduction without using an additional word. Note that the prime used in our NIST-like reduction example is Montgomery-friendly as well, since $p_{611} = 2^{61} - 1 = 2^{29} \cdot 2^{32} - 1$.

Algorithm 5 \mathbb{F}_{p^2} multiplication using Karatsuba and lazy reduction following [3].

Input: $\begin{cases} (a_0 + a_1i), (b_0 + b_1i) \in \mathbb{F}_{p^2}, \text{ with} \\ 0 \leq a_0, a_1 < S_0, 0 \leq b_0, b_1 < S_1. \\ \tilde{m} = m \times p \text{ such that } \tilde{m} \geq S_0S_1 \end{cases}$

Output: $(c_0 + c_1i) = (a_0 + a_1i)(b_0 + b_1i)$

- 1: $T_0 \leftarrow a_0 \times b_0$ ($< S_0S_1$)
- 2: $T_1 \leftarrow a_1 \times b_1$ ($< S_0S_1$)
- 3: $t_0 \leftarrow a_0 + a_1$ ($< 2S_0$)
- 4: $t_1 \leftarrow b_0 + b_1$ ($< 2S_1$)
- 5: $T_2 \leftarrow t_0 \times t_1$ ($< 4S_0S_1$)
- 6: $T_3 \leftarrow T_2 - (T_0 + T_1)$ ($< 2S_0S_1$)
- 7: $c_1 \leftarrow \text{MontRed}(T_3)$ ($< r_{(b,n)}(2S_0S_1, p)$)
- 8: $T_4 \leftarrow T_0 + \tilde{m} - T_1$ ($< S_0S_1 + \tilde{m}$)
- 9: $c_0 \leftarrow \text{MontRed}(T_4)$ ($< r_{(b,n)}(S_0S_1 + \tilde{m}, p)$)

Algorithm 6 \mathbb{F}_{p^2} squaring.

Input: $(a_0 + a_1i) \in \mathbb{F}_{p^2}$, with $0 \leq a_0, a_1 < S_0$.

Output: $(c_0 + c_1i) = (a_0 + a_1i)^2$

- 1: $T_0 \leftarrow a_0 + a_1$ ($< 2S_0$)
- 2: $T_1 \leftarrow a_0 + p - a_1$ ($< 2S_0$)
- 3: $c_0 \leftarrow \text{MontMul}(T_0, T_1)$ ($< r_{(b,n)}(4S_0^2, p)$)
- 4: $T_3 \leftarrow 2a_0$ ($< 2S_0$)
- 5: $c_1 \leftarrow \text{MontMul}(T_3, a_1)$ ($< r_{(b,n)}(2S_0^2, p)$)

5.2 Extension Field Arithmetic

Arithmetic in $\mathbb{F}_{p^2} = \mathbb{F}_p[i]/(i^2 + 1)$ is realized using arithmetic operations from \mathbb{F}_p . For instance, the result of multiplying two elements $a_0 + a_1i, b_0 + b_1i \in \mathbb{F}_{p^2}$ is $(a_0b_0 - a_1b_1) + (a_0b_1 + a_1b_0)i \in \mathbb{F}_{p^2}$. This can be achieved using four \mathbb{F}_p -multiplications, one \mathbb{F}_p -subtraction and one \mathbb{F}_p -addition or, when using a single level of Karatsuba, using three \mathbb{F}_p -multiplications, two \mathbb{F}_p -subtractions and three \mathbb{F}_p -additions. To optimize this further, we follow the lazy-reduction techniques described in [3], where the idea is to delay the modular reductions until the end of the computation. This has the advantage of reducing the number of reductions at the cost of performing the intermediate additions and subtractions on numbers of twice the bit-length. When using Karatsuba, this approach is outlined in Algorithm 5, together with the bounds on all intermediate values (given the bounds S_0 and S_1 on the inputs). In order to avoid working with negative numbers, we also require an additional precomputed input value \tilde{m} , which is a multiple of p such that $\tilde{m} \geq S_0S_1$. In practice the bounds on the input are chosen such that both $2S_0$ and $2S_1$ are less 2^{bn} , to avoid making the multiplication $t_0 \times t_1$ in Line 5 of Algorithm 5 work on more computer words. We found that the approach outlined Algorithm 5 (using Karatsuba and postponing the reductions) to be preferable on the 32-bit ARM Cortex-A8 platform. However, on our 64-bit Ivy Bridge platform, calculating the \mathbb{F}_{p^2} multiplication is more efficient using the “naive” schoolbook multiplication (but still using the lazy-reduction techniques to postpone the modular reductions). This requires one additional modular multiplication compared to Karatsuba, but lowers the modular additions/subtractions to only two. Due to the relatively low cost ratio between 64-bit modular multiplications and 64-bit additions it is more efficient to use schoolbook on such 64-bit platforms. Note that due to our representation of \mathbb{F}_{p^2} , squaring can be computed using only two \mathbb{F}_p multiplications, since $(a_0 + a_1i)^2 = (a_0 + a_1)(a_0 - a_1) + 2a_0a_1i$. This approach (including the bounds on the output) is given in Algorithm 6.

For computations modulo $p_{611} = 2^{61} - 1$ on the ARM, we choose to use Montgomery multiplication in combination with a conditional final subtraction, since such a subtraction is particularly efficient (see Section 5.1). This has the advantage of allowing us to add (or subtract) numbers *without* reducing them and using them as input, since if $S_0 = 2(2^{61} - 1)$ and $S_1 = (2^{61} - 1)$, then the first Montgomery reduction in Algorithm 5 is bounded by

$r_{(4,32)}(8p_{611}^2, p_{611}) - p_{611} < p_{611}$, so that the result is automatically properly reduced. For the second reduction, we could choose $\tilde{m} = (2^{63} + 1) \cdot p_{611}$ such that $r_{(4,32)}(2p_{611}^2, p_{611}) + (2^{63} + 1) \cdot p_{611} - p_{611} < p_{611}$ is also properly reduced. Another possibility is to choose $\tilde{m} = 2^{64} \cdot p_{611}$ to avoid adding the least significant 64 bits of \tilde{m} , which reduces the number of required addition instructions. However, in this case we would need one more conditional subtraction since $r_{(4,32)}(2p_{611}^2, p_{611}) + 2^{64} \cdot p_{611} - k \cdot p_{611} < p_{611}$ holds for $k \geq 2$. For the other Montgomery-friendly primes, we performed a similar analysis to minimize the number of reductions after additions and subtractions.

5.3 Using Mixed Additions

As outlined in Table 3, mixed divisor additions are significantly faster than using regular (projective) divisor additions. It is a common approach to convert the projective divisors in the lookup table to affine divisors in order to use these faster formulas when computing the scalar multiplication. This can be done efficiently using Montgomery’s *simultaneous inversion* method [40]. Supposing there are w such projective divisors in our lookup table(s), the simultaneous inversion method finds the w independent inverses using a single inversion and $3(w - 1)$ multiplications. For each of the w projective divisors of the form $(U_1 : U_0 : V_1 : V_0 : Z)$, normalization (given Z^{-1}) costs four additional multiplications. Hence, the total cost of converting the entire lookup table to affine coordinates is $(7w - 3)\mathbf{m} + \mathbf{I}$, where \mathbf{I} is the cost of an inversion in \mathbb{F}_{p^2} . To compute the inverse in $\mathbb{F}_{p^2} = \mathbb{F}_p[i]/(i^2 + 1)$, we use $(a_0 + a_1i)^{-1} = a_0/(a_0^2 + a_1^2) + (-a_1/(a_0^2 + a_1^2))i$, which costs, besides the \mathbb{F}_p -inversion, two \mathbb{F}_p -squarings, two \mathbb{F}_p -multiplications, a single \mathbb{F}_p -addition and a single \mathbb{F}_p -negation. Our implementations on both platforms revealed that it was always preferable to perform this normalization, i.e. that the cost of normalizing the lookup table is outweighed by the savings achieved when processing the scalar.

6 Results and Discussion

We implemented the generic, Kummer and 8-dimensional GLV/GLS (see Section 4) techniques using the different arithmetic approaches (as outlined in Section 5). The performance numbers of *all* our curves are presented in our overview in Table 1 (refer back to Section 3). In this section we use our fastest curves (for comparisons with other work) in two settings: one aims solely for performance (non-constant time) while the other provides some side-channel resistance [31] (i.e. runs in constant time). In Table 4 we summarize all the fastest software scalar multiplication results for genus g curves over both \mathbb{F}_p and \mathbb{F}_{p^2} for both 64-bit processors and 32-bit ARM architectures.

High-end 64-bit architecture. The 64-bit implementations cover the fastest overall constant time performance numbers [35] the fastest constant time performance numbers for elliptic curves over prime fields by Bernstein [5], the fastest (non-constant time) implementation for elliptic curves by Longa and Sica [36], the fastest constant time (Kummer) and non-constant time (4-GLV) performance numbers on genus 2 curves over prime fields [8] by Bos et al., and the fastest implementation of the NIST curve NIST-p224 by Käsper [28]. Note that all of these curves aim to provide 128-bit security, except the NIST curve which is designed to provide 112-bit security. We ran all these implementations on the same CPU: an Intel Core i7-3520M (Ivy Bridge) processor at 2893.484 MHz with hyperthreading turned off and over-clocking

Table 4. Performance comparison of scalar multiplication on an Intel Core i7-3520M Ivy Bridge (IB) and various ARM processors. We display the genus g of the curve, if the implementation runs in constant time (CT) or not, the field K over which the curve is defined, the security in bits (sec) provided by the curves and finally the performance number in 10^3 cycles.

	reference	g	CT	K	sec	10^3 cycles
Ivy-Bridge	Longa [35]	1	✓	\mathbb{F}_{p^2}	125	94
	Bernstein [5]	1	✓	\mathbb{F}_p	126	182
	Bos et al. [8] (Kummer)	2	✓	\mathbb{F}_p	125	117
	Bos et al. [8] (4-GLV)	2	✗	\mathbb{F}_p	125	156
	Käsper [28] (NISTp-224)	1	✓	\mathbb{F}_p	112	302
	Longa, Sica [36] (2-GLV)	1	✗	\mathbb{F}_p	127	145
	this work (Kummer)	2	✓	\mathbb{F}_{p^2}	103	108
	this work (Kummer)	2	✓	\mathbb{F}_{p^2}	110	167
	this work (8-GLV/GLS)	2	✗	\mathbb{F}_{p^2}	105	92
	this work (8-GLV/GLS)	2	✗	\mathbb{F}_{p^2}	111	136
ARM	Bernstein, Schwabe [7] (Cortex-A8 with NEON)	1	✓	\mathbb{F}_p	126	527
	Hamburg [25] (Cortex-A9)	1	✓	\mathbb{F}_p	125	616
	Morozov et al. [41] (Cortex-A8) (NISTp-224)	1	?	\mathbb{F}_p	112	7805
	this work (Kummer) (Cortex-A8)	2	✓	\mathbb{F}_{p^2}	103	767
	this work (Kummer) (Cortex-A8)	2	✓	\mathbb{F}_{p^2}	108	942
	this work (8-GLV/GLS) (Cortex-A8)	2	✗	\mathbb{F}_{p^2}	105	576
this work (8-GLV/GLS) (Cortex-A8)	2	✗	\mathbb{F}_{p^2}	109	810	

(“turbo boost”) disabled. We either compiled the code on our machine (for [8, 28, 5]) or used a precompiled binary provided to us by the authors (for [36, 35]).

Table 4 includes our fastest constant time implementation (Kummer) and our fastest non-constant time one (8-dimensional GLV/GLS) which will be made publicly available through [6]. A direct comparison to the state-of-the-art performance numbers is difficult; different curves of varying genus defined over different fields are used and most of the curves in Table 4 aim to provide 128-bit security while our curves aim for the 112-bit security level. Nevertheless, it is clear from our performance numbers that genus 2 curves over quadratic extension fields are competitive (and often faster) in terms of performance, even when taking the security into account. For instance, when compared to the fast implementation of curve NIST-p224 by Käsper [28], also aiming to provide 112-bit security, we are able to reduce the throughput by roughly a factor three. Interestingly, while implementations on the Kummer surface proved to be faster than 4-GLV/GLS implementations on genus 2 curves over 128-bit prime fields [8], our work over quadratic extension fields of 64-bit primes shows that 8-GLV/GLS overtakes \mathbb{F}_{p^2} Kummer implementations in terms of speed.

Low-end 32-bit architecture. For our low-end platform we consider the 32-bit ARM platform. More specifically we run our experiments on the BeagleBoard-xM [4], a low-power open-source hardware single-board computer, which contains an TI DM3730 processor (1 GHz Cortex-A8 ARM core). Unlike the setting of the 64-bit platforms, we were unable to run implementations from Table 4 on our platform since not all implementations were made available; hence, we copied the performance numbers directly from the papers and mention which ARM processor is used. We point out that the fast performance result by Bernstein and Schwabe [7] was obtained using ARM’s NEON instruction set (a combined 64- and 128-bit single instruction, multiple data instruction set), a possibility which has not been studied in this nor the other ARM papers mentioned in Table 4. Just as the implementation from [7], Hamburg studies fast curve arithmetic using Edwards (genus 1) curves [25].

A direct comparison is again difficult in this case because our curves in Table 1 provide a lower level of security. However, compared to the work by Morozov et al. [41] which also targets the 112-bit security level using the standard NIST curves, our numbers are an order of magnitude faster.

7 Conclusions

In this paper we have explored the possibility of using genus 2 curves over quadratic extension fields in cryptography, where the size of ground field fits into a *single* 64-bit word. This setting allows one to use 8-dimensional GLV/GLS scalar decompositions, which we explored in a variety of scenarios. The downside of using primes of this size for genus 2 based cryptography is that there exist faster-than-generic index calculus attacks which affect the security. Nevertheless, we show how to obtain 112-bit security and present performance numbers for both high-end 64-bit architectures and low-end 32-bit ARM platforms.

References

1. T. Acar and D. Shumow. Modular reduction without pre-computation for special moduli. Technical report, Microsoft Research, 2010.
2. L. Adleman, J. DeMarrais, and M. Huang. A subexponential algorithm for discrete logarithms over hyperelliptic curves of large genus over $\text{GF}(q)$. *Theoretical Computer Science*, 226(1-2):7–18, 1999.
3. D. F. Aranha, K. Karabina, P. Longa, C. H. Gebotys, and J. López. Faster explicit formulas for computing pairings over ordinary curves. In K. G. Paterson, editor, *Advances in Cryptology – EUROCRYPT 2011*, volume 6632 of *Lecture Notes in Computer Science*, pages 48–68. Springer Berlin Heidelberg, 2011.
4. Beagle Board. BeagleBoard-xM System Reference Manual. http://beagleboard.org/static/BBxMSRM_latest.pdf, 2013.
5. D. J. Bernstein. Curve25519: New Diffie-Hellman speed records. In M. Yung, Y. Dodis, A. Kiayias, and T. Malkin, editors, *Public Key Cryptography – PKC 2006*, volume 3958 of *Lecture Notes in Computer Science*, pages 207–228. Springer, Heidelberg, 2006.
6. D. J. Bernstein and T. Lange (editors). eBACS: ECRYPT Benchmarking of Cryptographic Systems. <http://bench.cr.yp.to>, accessed 1 March 2013.
7. D. J. Bernstein and P. Schwabe. NEON crypto. In E. Prouff and P. Schaumont, editors, *CHES 2012*, volume 7428 of *Lecture Notes in Computer Science*, pages 320–339. Springer, 2012.
8. J. W. Bos, C. Costello, H. Hisil, and K. Lauter. Two is greater than one. Cryptology ePrint Archive, Report 2012/670, 2012. <http://eprint.iacr.org/>.
9. J. Buhler and N. Koblitz. Lattice basis reduction, Jacobi sums and hyperelliptic cryptosystems. *Bulletin of the Australian Mathematical Society*, 58(1):147–154, 1998.
10. C. Diem. The GHS attack in odd characteristic. *J. Ramanujan Math. Soc.*, 18(1):1–32, 2003.
11. C. Diem. On the discrete logarithm problem in elliptic curves. *Compositio Mathematica*, 147(01):75–104, 2011.
12. W. Diffie and M. E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.
13. A. Enge and E. Thomé. Work in progress. Private communication, February 2013.
14. G. Frey. How to disguise an elliptic curve (Weil descent). Talk at ECC: slides available at <http://cacr.uwaterloo.ca/conferences/1998/ecc98/frey.ps>, September 1998.
15. S. D. Galbraith. Weil descent of Jacobians. *Discrete Applied Mathematics*, 128(1):165–180, 2003.
16. S. D. Galbraith, F. Hess, and N. P. Smart. Extending the GHS Weil descent attack. volume 2332 of *Lecture Notes in Computer Science*, pages 29–44. Springer, 2002.
17. S. D. Galbraith, X. Lin, and M. Scott. Endomorphisms for faster elliptic curve cryptography on a large class of curves. *J. Cryptology*, 24(3):446–469, 2011.
18. R. P. Gallant, R. J. Lambert, and S. A. Vanstone. Faster point multiplication on elliptic curves with efficient endomorphisms. In J. Kilian, editor, *CRYPTO*, volume 2139 of *Lecture Notes in Computer Science*, pages 190–200. Springer, 2001.

19. P. Gaudry. An algorithm for solving the discrete log problem on hyperelliptic curves. *EUROCRYPT*, 1807:19–34, 2000.
20. P. Gaudry. Fast genus 2 arithmetic based on theta functions. *Journal of Mathematical Cryptology JMC*, 1(3):243–265, 2007.
21. P. Gaudry. Index calculus for abelian varieties of small dimension and the elliptic curve discrete logarithm problem. *J. Symb. Comput.*, 44(12):1690–1702, 2009.
22. P. Gaudry, F. Hess, and N. P. Smart. Constructive and destructive facets of Weil descent on elliptic curves. *J. Cryptology*, 15(1):19–46, 2002.
23. P. Gaudry, E. Thomé, N. Thériault, and C. Diem. A double large prime variation for small genus hyperelliptic index calculus. *Math. Comput.*, 76(257):475–492, 2007.
24. E. Z. Goren and K. E. Lauter. Genus 2 curves with complex multiplication. *International Mathematics Research Notices*, 2012(5):1068–1142, 2012.
25. M. Hamburg. Fast and compact elliptic-curve cryptography. Cryptology ePrint Archive, Report 2012/309, 2012. <http://eprint.iacr.org/>.
26. T. Iijima, F. Momose, and J. Chao. Classification of elliptic/hyperelliptic curves with weak coverings against GHS attack without isogeny condition. Cryptology ePrint Archive, Report 2009/613, 2009. <http://eprint.iacr.org/>.
27. A. Joux and V. Vitse. Cover and decomposition index calculus on elliptic curves made practical - application to a previously unreachable curve over \mathbb{F}_{p^6} . In D. Pointcheval and T. Johansson, editors, *EUROCRYPT*, volume 7237 of *Lecture Notes in Computer Science*, pages 9–26. Springer, 2012.
28. E. Käsper. Fast elliptic curve cryptography in OpenSSL. In G. Danezis, S. Dietrich, and K. Sako, editors, *Financial Cryptography Workshops*, volume 7126 of *Lecture Notes in Computer Science*, pages 27–39. Springer, 2012.
29. M. Knežević, F. Vercauteren, and I. Verbauwhede. Speeding up bipartite modular multiplication. In M. Hasan and T. Hellesest, editors, *Arithmetic of Finite Fields – WAIFI 2010*, volume 6087 of *Lecture Notes in Computer Science*, pages 166–179. Springer Berlin / Heidelberg, 2010.
30. N. Koblitz. Elliptic curve cryptosystems. *Mathematics of Computation*, 48(177):203–209, 1987.
31. P. C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In N. Koblitz, editor, *Crypto 1996*, volume 1109 of *Lecture Notes in Computer Science*, pages 104–113. Springer, Heidelberg, 1996.
32. D. R. Kohel. Databases for Elliptic Curves and Higher Dimensional Analogues (Echidna). <http://echidna.maths.usyd.edu.au/kohel/dbs/>.
33. A. K. Lenstra. Generating RSA moduli with a predetermined portion. In K. Ohta and D. Pei, editors, *Asiacrypt’98*, volume 1514 of *Lecture Notes in Computer Science*, pages 1–10. Springer Berlin / Heidelberg, 1998.
34. C. H. Lim and P. J. Lee. More flexible exponentiation with precomputation. In *CRYPTO ’94*, volume 839 of *Lecture Notes in Computer Science*, pages 95–107. Springer, 1994.
35. P. Longa. Personal communication, February 2013.
36. P. Longa and F. Sica. Four-dimensional Gallant-Lambert-Vanstone scalar multiplication. In X. Wang and K. Sako, editors, *Asiacrypt 2012*, volume 7658 of *Lecture Notes in Computer Science*, pages 718–739. Springer, 2012.
37. V. S. Miller. Use of elliptic curves in cryptography. In H. C. Williams, editor, *Crypto 1985*, volume 218 of *Lecture Notes in Computer Science*, pages 417–426. Springer, Heidelberg, 1986.
38. F. Momose and J. Chao. Scholten forms and elliptic/hyperelliptic curves with weak Weil restrictions. Cryptology ePrint Archive, Report 2005/277, 2005. <http://eprint.iacr.org/>.
39. P. L. Montgomery. Modular multiplication without trial division. *Mathematics of Computation*, 44(170):519–521, April 1985.
40. P. L. Montgomery. Speeding the Pollard and elliptic curve methods of factorization. *Mathematics of Computation*, 48(177):243–264, 1987.
41. S. Morozov, C. Tergino, and P. Schaumont. System integration of elliptic curve cryptography on an OMAP platform. In *IEEE 9th Symposium on Application Specific Processors – SASP*, pages 52–57. IEEE Computer Society, 2011.
42. K. Nagao. Decomposition attack for the Jacobian of a hyperelliptic curve over an extension field. In G. Hanrot, F. Morain, and E. Thomé, editors, *ANTS*, volume 6197 of *Lecture Notes in Computer Science*, pages 285–300. Springer, 2010.
43. National Institute of Standards and Technology. Special publication 800-57: Recommendation for key management part 1: General (revised). http://csrc.nist.gov/publications/nistpubs/800-57/sp800-57-Part1-revised2_Mar08-2007.pdf.

44. National Security Agency. The case for elliptic curve cryptography. http://www.nsa.gov/business/programs/elliptic_curve.shtml, 2009.
45. Y.-H. Park, S. Jeong, and J. Lim. Speeding up point multiplication on hyperelliptic curves with efficiently-computable endomorphisms. In L. R. Knudsen, editor, *EUROCRYPT*, volume 2332 of *Lecture Notes in Computer Science*, pages 197–208. Springer, 2002.
46. J. M. Pollard. Monte Carlo methods for index computation (mod p). *Mathematics of Computation*, 32(143):918–924, 1978.
47. N. P. Smart. How secure are elliptic curves over composite extension fields? In B. Pfitzmann, editor, *EUROCRYPT*, volume 2045 of *Lecture Notes in Computer Science*, pages 30–39. Springer, 2001.
48. J. A. Solinas. Generalized Mersenne numbers. Technical Report CORR 99–39, Centre for Applied Cryptographic Research, University of Waterloo, 1999.
49. N. Thériault. Weil descent attack for Kummer extensions. *J. Ramanujan Math. Soc.*, 18(3):218–312, 2003.
50. E. Thomé. Barbeque – a genus 2 CM database: barbeque.loria.fr:46015/cm/cm.html.
51. U.S. Department of Commerce/National Institute of Standards and Technology. Digital Signature Standard (DSS). FIPS-186-3, 2009. http://csrc.nist.gov/publications/fips/fips186-3/fips_186-3.pdf.
52. C. D. Walter. Montgomery exponentiation needs no final subtractions. *Electronics Letters*, 35(21):1831–1832, 1999.

A Curves

A.1 Generic Curves over \mathbb{F}_{p^2} .

Generic curve for a 63-bit Montgomery-friendly prime. Let $p = (2^{31} - 307656) \cdot 2^{32} - 1$. The CM field $K = \mathbb{Q}[x]/(x^4 + 18x + 60)$ has class number 4 and gives rise to a curve C over $\mathbb{F}_{p^2} = \mathbb{F}_p[i]/(i^2 + 1)$ whose Jacobian has order $\#\text{Jac}_C(\mathbb{F}_{p^2}) = h \cdot r$, where $h = 2^6 \cdot 293 \cdot 3407003$ and

$$r = 113211333261982024217022596120436785263778756288209705887768360919,$$

which is a 217-bit prime. A possible degree 5 model is $C : y^2 = x^5 + f_3x^3 + f_2x^2 + f_1x + f_0$, where

$$\begin{aligned} f_3 &= 3481766375391715547i + 5725587270663764102, & f_2 &= 7780325096854741248i + 1284140721825785735, \\ f_1 &= 6300910221402262603i + 8769731505802494108, & f_0 &= 1369518879755098155i + 7792045068676593973. \end{aligned}$$

Generic curve for a 64-bit NIST-friendly prime. Let $p = 2^{64} - 189$. The CM field $K = \mathbb{Q}[x]/(x^4 + 110x + 2497)$ has class number 4 and gives rise to a curve C over $\mathbb{F}_{p^2} = \mathbb{F}_p[i]/(i^2 + 1)$ whose Jacobian has order $\#\text{Jac}_C(\mathbb{F}_{p^2}) = h \cdot r$, where $h = 2^6 \cdot 16879 \cdot 62683$ and

$$r = 115792089237316190691681685414194564710330229202610583744872853364031268362653,$$

which is a 221-bit prime. A possible degree 5 model is $C : y^2 = x^5 + f_3x^3 + f_2x^2 + f_1x + f_0$, where

$$\begin{aligned} f_3 &= 12544884702866578275i + 6148275953345011740, & f_2 &= 4257865936154958333i + 1215557756379749393, \\ f_1 &= 16519109982018842029i + 3600247297987423813, & f_0 &= 6416264942088602071i + 7008878166144108916. \end{aligned}$$

Generic curve for the 61-bit special prime. Let $p = 2^{61} - 1$. The CM field $K = \mathbb{Q}[x]/(x^4 + 34x + 29)$ has class number 4 and gives rise to a curve C over $\mathbb{F}_{p^2} = \mathbb{F}_p[i]/(i^2 + 1)$ whose Jacobian has order $\#\text{Jac}_C(\mathbb{F}_{p^2}) = h \cdot r$, where $h = 2^6 \cdot 2681508737$ and

$$r = 164725089312508207602342831124551763828685547226980673422949427,$$

which is 207 bits. A possible degree 5 model is $C : y^2 = x^5 + f_3x^3 + f_2x^2 + f_1x + f_0$, where

$$\begin{aligned} f_3 &= 1092171533470960661i + 200282817898000997, & f_2 &= 1010904730175095615i + 222880891256166859, \\ f_1 &= 2188897028088906113i + 2121806688923765783, & f_0 &= 2023072108862458234i + 1545665884212463459. \end{aligned}$$

A.2 Buhler-Koblitz Curves for 8-Dimensional GLV/GLS

BK curve for a 63-bit Montgomery-friendly prime. Let $p = (2^{31} - 201) \cdot 2^{32} - 1$ and construct $\mathbb{F}_{p^2} = \mathbb{F}_p[i]/(i^2 + 1)$. The Jacobian of the curve $C/\mathbb{F}_{p^2} : y^2 = x^5 + a$ with $a = 4923475266616025179i + 5793761506378637713$ has order $\#\text{Jac}_C(\mathbb{F}_{p^2}) = h \cdot r$, where $h = 821 \cdot 2101481$ and

$$r = 4194595820778911700452727695673246028118799060661534138832805061921,$$

which is a 222-bit prime. The endomorphism ψ in (1) is defined by

$$\begin{aligned} \alpha &= 5513956703956065859i + 2946903236728575314, & \beta &= 8439477163420991741i + 7773202624005877183, \\ \gamma &= 3782408741013213751i + 2442117051565934285, & \delta &= 7951546392112689792i + 7345844203826764325. \end{aligned}$$

The shortest vector in the GLV lattice is given by

$$(a_0, \dots, a_7) = (13549510, -112063962, 50889612, 53347019, 108477091, -14113633, 92683994, -71784510).$$

BK curve for a 64-bit NIST-friendly prime. Let $p = 2^{64} - 2285$ and construct $\mathbb{F}_{p^2} = \mathbb{F}_p[i]/(i^2 + 1)$. The Jacobian of the curve $C/\mathbb{F}_{p^2} : y^2 = x^5 + a$ with $a = 17275736705428938862i + 9597225893510653377$ has order $\#\text{Jac}_C(\mathbb{F}_{p^2}) = h \cdot r$, where $h = 421 \cdot 1181 \cdot 10601$ and

$$r = 21968482676544214814720042345948911463859528450701256799597520244461,$$

which is a 224-bit prime. The endomorphism ψ in (1) is defined by

$$\begin{aligned} \alpha &= 16936375061285296480i + 1975550208383908832, & \beta &= 5062688676656924599i + 9526106207035922484, \\ \gamma &= 17154700957512463544i + 6621962486378311879, & \delta &= 9429501515427699879i + 15905244445742514264. \end{aligned}$$

The shortest vector in the GLV lattice is given by

$$(a_0, \dots, a_7) = (87166103, 121861322, -104402854, 130055674, -121298598, -27534036, 19420196, -72633646).$$

BK curve for the 61-bit special prime. Let $p = 2^{61} - 1$ and construct $\mathbb{F}_{p^2} = \mathbb{F}_p[i]/(i^2 + 1)$. The Jacobian of the curve $C/\mathbb{F}_{p^2} : y^2 = x^5 + a$ with $a = 389685887377831749i + 1917550166465917276$ has order $\#\text{Jac} = h \cdot r$, where $h = 4196633321$ and

$$r = 6736245669830859461503996919463094254829566207348665703723122861,$$

which is a 213-bit prime. The endomorphism ψ in (1) is defined by

$$\begin{aligned} \alpha &= 1043675927557798232i + 711577460084601543, & \beta &= 1298223254182651606i + 1876811897552284385, \\ \gamma &= 1508434156926104763i + 2150923980995833762, & \delta &= 1646230938444105124i + 205317303568012635. \end{aligned}$$

The shortest vector in the GLV lattice is given by

$$(a_0, \dots, a_7) = (8638881, -34043662, 43566183, 22144585, 53872798, 1112704, 40233559, -29304313).$$

A.3 Kummer Surfaces over \mathbb{F}_{p^2}

The three Kummer surfaces we use all come from the generic curves presented in A.1.

Kummer surface for a 63-bit Montgomery-friendly prime. Let $p = (2^{31} - 307656) \cdot 2^{32} - 1$ and construct $\mathbb{F}_{p^2} = \mathbb{F}_p[i]/(i^2 + 1)$. One choice of the Rosenhain model corresponding to the generic curve over this prime (from A.1) is given by

$$\begin{aligned} \lambda &= 2535046569008734284i + 3027622460227942962, & \mu &= 6168887938445916986i + 3306366075534497189, \\ \nu &= 1342856250081051940i + 556271600778613627. \end{aligned}$$

The group orders of $\text{Jac}(C) \cong \text{Jac}(C_{\text{Ros}})$ and $\text{Jac}(C') \cong \text{Jac}(C'_{\text{Ros}})$ are given by $h \cdot r$ and $h' \cdot r'$ respectively, where $h = 2^6 \cdot 293 \cdot 3407003$, $h' = 2^4 \cdot 3 \cdot 5$, and

$$\begin{aligned} r &= 113211333261982024217022596120436785263778756288209705887768360919, \\ r' &= 30136913640765001310593303136401614891746213772217580804569131889303082401, \end{aligned}$$

which are 217 and 245-bit primes respectively. The corresponding Kummer surface \mathcal{K} is parameterized by

$$\begin{aligned} E' &= 6885924737355512695i + 5731565823886889093, & F &= 2971432930196930029i + 7041132631915596671, \\ G &= 5802556419299219191i + 406612783433840227, & H &= 1976054157537666133i + 5209763529069961553. \end{aligned}$$

For a generator of the smaller subgroup, we take

$$\begin{aligned} (P_x : P_y : P_z : P_t) &= (1 : 6780948843060404695i + 2778810341654550438 : \\ &8114001334824362978i + 177626812522210373 : 8746905464522843317i + 2846346116089482067), \end{aligned}$$

which has order r on \mathcal{K} .

Kummer surface for a 64-bit NIST-friendly prime. Let $p = 2^{64} - 189$ and construct $\mathbb{F}_{p^2} = \mathbb{F}_p[i]/(i^2 + 1)$. One choice of the Rosenhain model corresponding to the generic curve over this prime (from A.1) is given by

$$\begin{aligned} \lambda &= 2952413239496295680i + 18323491867119280217, & \mu &= 17185170913358942570i + 17939661102861094359, \\ \nu &= 2480178925665143594i + 11935307278311786128. \end{aligned}$$

The group orders of $\text{Jac}(C) \cong \text{Jac}(C_{\text{Ros}})$ and $\text{Jac}(C') \cong \text{Jac}(C'_{\text{Ros}})$ are given by $h \cdot r$ and $h' \cdot r'$ respectively, where $h = 2^6 \cdot 16879 \cdot 62683$, $h' = 2^6$, and

$$\begin{aligned} r &= 1710024880158127742499387005384745007574557462363960828332983222797, \\ r' &= 1809251394333065479352804274970805210144581724733222509461653232757833765789, \end{aligned}$$

which are 221 and 250-bit primes respectively. The corresponding Kummer surface \mathcal{K} is parameterized by

$$\begin{aligned} E' &= 16043557314664054699i + 1643743012492832694, & F &= 5130901660039278381i + 6378763037882534459, \\ G &= 1608091444136491268i + 17326398380347202807, & H &= 9399469748702509058i + 14187747531033270771. \end{aligned}$$

For a generator of the smaller subgroup, we take

$$\begin{aligned} (P_x : P_y : P_z : P_t) &= (1 : 12020423121669216257i + 14412747545880565171 : \\ &17907431373007144635i + 9160419332789117271 : 5435335636436947021i + 179642190672187846), \end{aligned}$$

which has order r on \mathcal{K} .

Kummer surface for the 61-bit special prime. Let $p = 2^{61} - 1$ and construct $\mathbb{F}_{p^2} = \mathbb{F}_p[i]/(i^2 + 1)$. One choice of the Rosenhain model corresponding to the generic curve over this prime (from A.1) is given by

$$\begin{aligned} \lambda &= 819188565154922720i + 895962277282143337, & \mu &= 1814120660984670986i + 1949052274256106659, \\ & & \nu &= 366557254231698774i + 1066444538207623350. \end{aligned}$$

The group orders of $\text{Jac}(C) \cong \text{Jac}(C_{\text{Ros}})$ and $\text{Jac}(C') \cong \text{Jac}(C'_{\text{Ros}})$ are given by $2^6 \cdot 2681508737 \cdot r$ and $2^6 \cdot 1567 \cdot r'$ respectively, where

$$\begin{aligned} r &= 164725089312508207602342831124551763828685547226980673422949427, \\ r' &= 281883705293296797185233622320210153913603412961528577419551223300313, \end{aligned}$$

which are 207- and 228-bit primes respectively. The corresponding Kummer surface \mathcal{K} is parameterized by

$$\begin{aligned} E' &= 1376950493669180641i + 392025351069600249, & F &= 2182231190395397901i + 1046607026783845849, \\ G &= 1949345643105151894i + 420871672382413506, & H &= 949840083031512991i + 1462105493516569052. \end{aligned}$$

The point

$$\begin{aligned} (P_x : P_y : P_z : P_t) &= (1 : 980711647388742169i + 1130772398590331836 : \\ &1498799964986104091i + 2182769176191920281 : 520921417271535836i + 1317163629809380394) \end{aligned}$$

has order r on \mathcal{K} .

B Prime order curves

In Table 5 we provide details for curves with prime group orders (or minimal cofactors) over the three types of primes (NIST-friendly, Montgomery-friendly and $p = 2^{61} - 1$) that were explored in this work. Both the generic and Buhler-Koblitz (BK) curves for 8-GLV/GLS have prime order, whilst the curves C corresponding to the Kummer surfaces have optimal cofactors of 16 (see [8, §5.2]). We note that in two of the Kummer scenarios, the twists also have optimal cofactors of 16, whilst over $p = 2^{61} - 1$ the best Kummer we could find had a cofactor of $16 \cdot 3^4$ on the twist – we thank Thomé for sending us the class polynomials corresponding to the CM field in Table 5 which has class number 154. We provide the values of A and B which define the quartic CM field $x^4 + Ax^2 + B$ for the generic curves and Kummer surfaces, and note that the CM field for the BK curves is defined by the polynomial $\Phi_5(x) = x^4 + x^3 + x^2 + x + 1$. For any particular prime p , there are at most 10 different group orders (isomorphism classes)

Table 5. Details for curves with Jacobians of prime order or with optimal cofactors of $h = 16$ in the case of the Kummer surfaces. The difference in performance between the prime order curves ($r \approx p^4$) and the curves (given in Table 1) which takes smaller subgroups ($r \ll p^4$) is illustrated in the final two columns. These numbers were obtained using the 64-bit environment described in Section 6, which is why there are no numbers reported for the Montgomery-friendly primes (which target 32-bit platforms).

algorithm	reduction	CM field [A, B , class no.]	base field p	$ h _2$	$ r _2$	cycles ($\cdot 10^3$)	
						$r \approx p^4$	$r \ll p^4$
generic	special	[50, 541, 2]	$2^{61} - 1$	1	244	240	204
	Mont.	[37, 5, 2]	$(2^{31} - 701) \cdot 2^{32} - 1$	1	252	-	-
	NIST	[62, 689, 2]	$2^{64} - 189$	1	256	388	333
Kummer	special	[369, 59, 154]	$2^{61} - 1$	4	244	125	108
	Mont.	[97, 1901, 3]	$(2^{31} - 1608) \cdot 2^{32} - 1$	4	248	-	-
	NIST	[32, 67, 4]	$2^{64} - 2289$	4	252	190	167
8-GLV/GLS	Mont.	$[\Psi_5(x), 1]$	$(2^{31} - 307656) \cdot 2^{32} - 1$	1	252	-	-
	NIST	$[\Psi_5(x), 1]$	$2^{64} - 131465$	1	256	148	136

of BK curves defined over \mathbb{F}_{p^2} , which are obtained as $C : y^2 = x^5 + \tilde{a}^e$ for $e = 0, \dots, 9$, where \tilde{a} is a primitive element in \mathbb{F}_{p^2} ; both of the BK curves in Table 5 have $e = 7$.

Table 5 shows that transferring to prime order curves gives a slight performance penalty, as can be expected. Note that we could not give timings for a prime order 8-GLV/GLS implementation over $p = 2^{61} - 1$, because there is no such curve (recall that we presented the group order containing the largest prime factor in Table 1). In the generic and Kummer surfaces implementations, where scalars are of length $\approx r$ and there are only minor precomputations, the performance ratios match closely with the ratios of the subgroup sizes. On the other hand, the 8-GLV/GLS implementations do not benefit as much from taking smaller prime subgroup sizes, since the significant cost of the precomputation remains fixed regardless of the subgroup size.

Finally, we note that for prime order BK curves, substantial savings can be made to the 8-GLV/GLS decomposition algorithm (see Algorithm 1) depending on the scenario: in the case of the prime order BK curve over the Montgomery-friendly field, the shortest vector in the GLV lattice has $a_1 = a_3 = a_7 = 0$ and $a_5 = 1$, whilst in the case of the prime order BK curve over the NIST-friendly field, the shortest vector has $a_3 = a_5 = a_7 = 0$ and $a_1 = -1$, both of which give solid speedups to the decomposition phase. The absence of any advantageous terms in the other cases is due to the Jacobian group orders having large cofactors.