

μ Kummer: efficient hyperelliptic signatures and key exchange on microcontrollers

Joost Renes, Peter Schwabe, Benjamin Smith, Lejla Batina

► **To cite this version:**

Joost Renes, Peter Schwabe, Benjamin Smith, Lejla Batina. μ Kummer: efficient hyperelliptic signatures and key exchange on microcontrollers. Cryptographic Hardware and Embedded Systems – CHES 2016, Aug 2016, Santa Barbara, United States. Springer-Verlag, Lecture Notes in Computer Science, 9813, pp.20, 2016, Cryptographic Hardware and Embedded Systems – CHES 2016. <10.1007/978-3-662-53140-2_15>. <hal-01300727v1>

HAL Id: hal-01300727

<https://hal.inria.fr/hal-01300727v1>

Submitted on 11 Apr 2016 (v1), last revised 26 Jan 2017 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright}

μ Kummer: efficient hyperelliptic signatures and key exchange on microcontrollers

Joost Renes¹, Peter Schwabe¹, Benjamin Smith², and Lejla Batina¹ *

¹ Digital Security Group, Radboud University, The Netherlands

`lejla@cs.ru.nl, j.renes@cs.ru.nl, peter@cryptojedi.org`

² INRIA and Laboratoire d'Informatique de l'École polytechnique (LIX), France

`smith@lix.polytechnique.fr`

Abstract. We describe the design and implementation of efficient signature and key-exchange schemes for the AVR ATmega and ARM Cortex M0 microcontrollers, targeting the 128-bit security level. Our algorithms are based on an efficient Montgomery ladder scalar multiplication on the Kummer surface of Gaudry and Schost's genus-2 hyperelliptic curve, combined with the Jacobian point recovery technique of Costello, Chung, and Smith. Our results are the first to show the feasibility of software-only hyperelliptic cryptography on constrained platforms, and represent a significant improvement on the elliptic-curve state-of-the-art for both key exchange and signatures on these architectures. Notably, our key-exchange scalar-multiplication software runs in under 9740k cycles on the ATmega, and under 2650k cycles on the Cortex M0.

Keywords. Hyperelliptic curve cryptography, Kummer surface, AVR ATmega, ARM Cortex M0.

1 Introduction

The current state of the art in asymmetric cryptography, not only on microcontrollers, is elliptic-curve cryptography; the most widely accepted reasonable security is the 128-bit security level. All current speed records for 128-bit secure key exchange and signatures on microcontrollers are held—until now—by elliptic-curve-based schemes. Outside the world of microcontrollers, it is well known that genus-2 hyperelliptic curves and their Kummer surfaces present an attractive alternative to elliptic curves. For example, the current speed record for 128-bit-secure scalar multiplication on a range of architectures is held by Kummer-based software presented at Asiacrypt 2014 by Bernstein, Chuengsatiansup, Lange, and Schwabe [3]. These speed records were achieved by exploiting the computational power of vector units of recent “large” processors such

* This work has been supported by the Netherlands Organisation for Scientific Research (NWO) through Veni 2013 project 13114 and by the Technology Foundation STW (project 13499 - TYPHOON & ASPASIA), from the Dutch government. Permanent ID of this document: `b230ab9b9c664ec4aad0cea0bd6a6732`. Date: 2016-04-07

as Intel Sandy Bridge, Ivy Bridge, and Haswell, or the ARM Cortex-A8. Surprisingly, very little attention has been given to Kummer surfaces on embedded processors. Indeed, this is the first work showing the feasibility of software-only implementations of hyperelliptic-curve based crypto on constrained platforms. There have been some investigations of binary hyperelliptic curves targeting the much lower 80-bit security level, but those are actually examples of software-hardware co-design showing that using hardware acceleration for field operations was necessary to get reasonable performance figures (see eg. [1] and [15]).

In this paper we investigate the potential of genus-2 hyperelliptic curves for both key exchange and signatures on the “classical” 8-bit AVR ATmega architecture, and the more modern 32-bit ARM Cortex-M0 processor. We show that not only are hyperelliptic curves competitive, they clearly outperform state-of-the-art elliptic-curve schemes in terms of speed and size. For example, our variable-basepoint scalar multiplication on a 127-bit Kummer surface is 31% faster on AVR and 26% faster on the M0 than the recently presented speed records for Curve25519 software by Düll, Haase, Hinterwälder, Hutter, Paar, Sánchez, and Schwabe [10]; our implementation is also smaller, and requires less RAM.

We use a recent result by Costello, Chung, and Smith [9] to also set new speed records for 128-bit secure signatures. Specifically, we present a new signature scheme based on fast Kummer surface arithmetic. It is inspired by the EdDSA construction by Bernstein, Duif, Lange, Schwabe, and Yang [4]. On the ATmega, it produces shorter signatures, achieves higher speeds and needs less RAM than the Ed25519 implementation presented in [21].

	ATmega		Cortex M0	
	Cycles	Stack bytes	Cycles	Stack bytes
<code>keygen</code>	10 206 181	812	2 774 087	1 056
<code>sign</code>	10 404 033	926	2 865 351	1 360
<code>verify</code>	16 240 510	992	4 453 978	1 432
<code>dh_exchange</code>	9 739 059	429	2 644 604	584

Table 1. Cycle counts and stack usage in bytes of all functions related to the signature and key exchange schemes, for the AVR ATmega and ARM Cortex M0 microcontrollers.

Our routines handling secret data are constant-time, and are thus naturally resistant to timing attacks. These algorithms are built around the Montgomery ladder, which improves resistance against simple-power-analysis (SPA) attacks. Resistance to DPA attacks can be easily obtained by randomizing the scalar and/or Jacobian points. Re-randomizing the latter after each ladder step also guarantees resistance against horizontal types of attacks.

Source code. We place all of the software described in this paper into the public domain, to maximize reuseability of our results. The software is available at <http://www.cs.ru.nl/~jrenes/>.

2 High-level overview

We begin by describing the details of our signature and Diffie–Hellman schemes, explaining the choices we made in their design. Concrete details on their implementation appear in §3 and §4 below. Experimental results and comparisons follow in §5.

2.1 Signatures

Our signature scheme adheres closely to the proposal of [9, §8], which in turn is a type of Schnorr signature [24]. There are however some differences, and some possible trade-offs, which we discuss below. The full scheme is presented at the end of this section.

Group structure. We build the signature scheme on top of the group structure from the Jacobian $\mathcal{J}_{\mathcal{C}}(\mathbb{F}_q)$ of a genus-2 hyperelliptic curve \mathcal{C} . More specifically, \mathcal{C} is the Gaudry–Schost curve over the prime field \mathbb{F}_q with $q = 2^{127} - 1$ (cf. §3.2). The Jacobian is a group of order $\#\mathcal{J}_{\mathcal{C}}(\mathbb{F}_q) = 2^4N$, where

$$N = 2^{250} - 0x334D69820C75294D2C27FC9F9A154FF47730B4B840C05BD$$

is a 250-bit prime. For more details on the Jacobian and its elements, see §3.3.

Hash function. The hash function H can be any hash function with a 128-bit security level. For our purposes, $H(M) = \text{SHAKE128}(M, 512)$ suffices [11]. While SHAKE128 has variable-length output, we shall only use the 512-bit output implementation.

Encoding. The objects on which we operate on the highest level are points Q in $\mathcal{J}_{\mathcal{C}}(\mathbb{F}_q)$. To minimize communication costs, we compress the common 508-bit representation of Q into 256 bits (see §3.3). To avoid confusion between compressed and uncompressed points, we let \underline{Q} denote the 256-bit encoding of Q . (This notation is the same as in [4].)

Public generator. The public generator can be any element P of $\mathcal{J}_{\mathcal{C}}(\mathbb{F}_q)$ such that $[N]P = 0$. In our implementation we have made the arbitrary choice $P = (X^2 + u_1X + u_0, v_1X + v_0)$, where

$$u_1 = 0x7D5D9C3307E959BF27B8C76211D35E8A,$$

$$u_0 = 0x2703150F9C594E0CA7E8302F93079CE8,$$

$$v_1 = 0x444569AF177A9C1C721736D8F288C942,$$

$$v_0 = 0x7F26CFB225F42417316836CFF8AEFB11.$$

This is the point which we use the most for scalar multiplication. Since it remains fixed, we assume we have its decompressed representation precomputed, so as to avoid having to perform the relatively expensive decompression operation

whenever we need a scalar multiplication; this gives a low-cost speed gain. We further assume we have a “wrapped” representation of the projection of P to the Kummer surface, which is used to speed up the `xDBLADD` function. See §4.1 for more details on the `xWRAP` function.

Public keys. In contrast to the public generator, we assume public keys are compressed: they are communicated much more frequently, and we therefore benefit much more from smaller keys. Moreover, we include the public key in one of the hashes during the `sign` operation [18,20], computing $h = H(\underline{R}||\underline{Q}||M)$ instead of the $h = H(\underline{R}||M)$ originally suggested by Schnorr [24]. This protects against adversaries attacking multiple public keys simultaneously.

Compressed signatures. Schnorr [24] mentions the option of compressing signatures by hashing one of their two components: the hash size only needs to be $b/2$ bits, where b is the key length. Following this suggestion, our signatures are 384-bit values of the form $(h_{128}||s)$, where h_{128} means the lowest 128 bits of $h = H(\underline{R}||\underline{Q}||M)$, and s is a 256-bit scalar. The most obvious upside is that signatures are smaller, reducing communication overhead. Another big advantage is that we can exploit the half-size scalar to speed up signature verification. On the other hand, we lose the possibility of efficient batch verification.

Verification efficiency. The most costly operation in signature verification is the two-dimensional scalar multiplication $T = [s]P \oplus [h_{128}]Q$. In [9], the authors propose an algorithm relying on the differential addition chains presented in [2]. However, since we are using compressed signatures, we have a small scalar h_{128} . To abuse this, we simply compute $[s]P$ and $[h_{128}]Q$ separately using the fast scalar multiplication on the Kummer surface and finally add them together on the Jacobian. Not only do we need fewer cycles, but we can also reduce code size by reusing the one-dimensional scalar multiplication routines.

The scheme. We now define our signature scheme, taking the remarks above into account.

Key generation (keygen). Let d be a 256-bit secret key, and P the public generator. Compute $(d' || d'') \leftarrow H(d)$ (with d' and d'' both 256 bits), then $Q \leftarrow [16d']P$. The public key is \underline{Q} .

Signing (sign). Let M be a message, d a 256-bit secret key, P the public generator, and \underline{Q} a compressed public key. Compute $(d' || d'') \leftarrow H(d)$ (with d' and d'' both 256 bits), then $r \leftarrow H(d'' || M)$, then $R \leftarrow [r]P$, then $h \leftarrow H(\underline{R} || \underline{Q} || M)$, and finally $s \leftarrow (r - 16h_{128}d') \bmod N$. The signature is $(h_{128} || s)$.

Verification (verify). Let M be a message with a signature $(h_{128} || s)$ corresponding to a public key \underline{Q} , and let P be the public generator. Compute $T \leftarrow [s]P \oplus [h_{128}]Q$, then $g \leftarrow H(\underline{T} || \underline{Q} || M)$. The signature is correct if $g_{128} = h_{128}$, and incorrect otherwise.

Remark 1. We note that there may be other, faster algorithms to compute this “one-and-a-half-dimensional” scalar multiplication. Since for verification we do

not have to worry about being constant-time, one option might be to alter Montgomery’s PRAC [26, §3.3.1] to make use of the half-size scalar. We have chosen not to pursue this line, preferring the solid benefits of reduced code size instead.

2.2 Diffie-Hellman key exchange.

For key exchange it is not necessary to have a group structure; it is enough to have a pseudo-multiplication. We can therefore carry out our the key exchange directly on the Kummer surface $\mathcal{K}_C = \mathcal{J}_C / \langle \pm \rangle$, gaining efficiency by not projecting from and recovering to the Jacobian \mathcal{J}_C . If Q is a point on \mathcal{J}_C , then its image in \mathcal{K}_C is $\pm Q$. The common representation for points in $\mathcal{K}_C(\mathbb{F}_q)$ is a 512-bit 4-tuple of field elements. For input points (i. e. the generator or public keys), we prefer the 384-bit “wrapped” representation (see §3.5). This not only reduces key size, but it also allows a speed-up in the core `xDBLADD` subroutine. The wrapped representation of a point $\pm Q$ on \mathcal{K}_C is denoted by $\underline{\pm Q}$.

Key exchange (`dh_exchange`). Let d be a 256-bit secret key, and $\underline{\pm P}$ the public generator (respectively public key). Compute $\pm Q \leftarrow \pm[d]P$. The generated public key (respectively shared secret) is $\underline{\pm Q}$.

Remark 2. While it might be possible to reduce the key size even further to 256 bits, we would then have to pay the cost of compressing and decompressing, and also wrapping for `xDBLADD` (see the discussion in [9, App. A]). We therefore choose to keep the 384-bit representation, which is consistent with [3].

3 Building blocks: algorithms and their implementation

We begin by presenting the finite field $\mathbb{F}_{2^{127-1}}$ in §3.1. We then define the curve C in §3.2, before giving basic methods for the elements of \mathcal{J}_C in §3.3. We then present the fast Kummer \mathcal{K}_C and its differential addition operations in §3.4.

3.1 The field \mathbb{F}_q

We work over the prime finite field \mathbb{F}_q , where q is the Mersenne prime

$$q := 2^{127} - 1 .$$

We let \mathbf{M} , \mathbf{S} , \mathbf{a} , \mathbf{s} , and \mathbf{neg} denote the costs of multiplication, squaring, addition, subtraction, and negation in \mathbb{F}_q . Later, we will define a special operation for multiplying by small constants: its cost is denoted by \mathbf{m}_c .

We can represent elements of \mathbb{F}_q as 127-bit values; but since the ATmega and Cortex M0 work with 8- and 32-bit words, respectively, the obvious choice is to represent field elements with 128 bits. That is, an element $g \in \mathbb{F}_q$ is represented as $g = \sum_{i=0}^{15} g_i 2^{8i}$ on the AVR ATmega platform and as $g = \sum_{i=0}^3 g'_i 2^{32i}$ on the Cortex M0, where $g_i \in \{0, \dots, 2^8 - 1\}$, $g'_i \in \{0, \dots, 2^{32} - 1\}$.

For complete field arithmetic we implement modular reduction, addition, subtraction, multiplication, and inversion. We comment on some important aspects here, giving cycle counts in Table 2.

Working with the prime field \mathbb{F}_q , we need integer reduction modulo q ; this is implemented as `bigint_red`. Reduction is very efficient because $2^{128} \equiv 2 \pmod{q}$, which enables us to reduce using only shifts and integer additions. Given this reduction, we implement addition and subtraction operations for \mathbb{F}_q (as `gfe_add` and `gfe_sub`, respectively) in the obvious way.

The most costly operations in \mathbb{F}_q are multiplication (`gfe_mul`) and squaring (`gfe_sqr`), which are implemented as 128×128 -bit bit integer operations (`bigint_mul` and `bigint_sqr`) followed by a call to `bigint_red`. Since we are working on the same platforms as in [10] in which both of these operations are already highly optimized, we chose to take the necessary code for `bigint_mul` and `bigint_sqr` from those implementations:

- On the AVR ATmega: The authors of [17] implement a 3-level Karatsuba multiplication of two 256-bit integers, representing elements f of $\mathbb{F}_{2^{255-19}}$ as $f = \sum_{i=0}^{31} f_i 2^{8i}$ with $f_i \in \{0, \dots, 2^8 - 1\}$. Since the first level of Karatsuba relies on a 128×128 -bit integer multiplication routine named `MUL128`, we simply lift this function out to form a 2-level 128×128 -bit Karatsuba multiplication. Similarly, their 256×256 -bit squaring relies on a 128×128 -bit routine `SQR128`, which we can (almost) directly use. Since the 256×256 -bit squaring is 2-level Karatsuba, the 128×128 -bit squaring is 1-level Karatsuba.
- On the ARM Cortex M0: The authors of [10] make use of optimized Karatsuba multiplication and squaring. In this case their assembly code does not rely on subroutines, but fully inlines 128×128 -bit multiplication and squaring. The 256×256 -bit multiplication and squaring are both 3-level Karatsuba implementations. Hence, using these, we end up with 2-level 128×128 -bit Karatsuba multiplication and squaring.

The function `gfe_invert` computes inversions in \mathbb{F}_q as exponentiations, using the fact that $g^{-1} = g^{q-2}$ for all g in \mathbb{F}_q^\times . To do this efficiently we use an addition chain for $q - 2$, doing the exponentiation in $10\mathbf{M} + 126\mathbf{S}$.

Finally, to speed up our Jacobian point decompression algorithms, we define a function `gfe_powminhalf` which computes $g \mapsto g^{-1/2}$ for g in \mathbb{F}_q (up to a choice of sign). To do this, we note that $g^{-1/2} = \pm g^{-(q+1)/4} = \pm g^{(3q-5)/4}$ in \mathbb{F}_q ; this exponentiation can be done with an addition chain of length 136, using $11\mathbf{M} + 125\mathbf{S}$. We can then define a function `gfe_sqrtinv`, which given (x, y) and a bit b , computes $(\sqrt{x}, 1/y)$ as $(\pm xyz, xyz^2)$ where $z = \text{gfe_powminhalf}(xy^2)$, choosing the sign so that the square root has least significant bit b . Including the `gfe_powminhalf` call, this costs $15\mathbf{M} + 126\mathbf{S} + 1\text{neg}$.

3.2 The curve \mathcal{C} and its theta constants

We define the curve \mathcal{C} “backwards”, starting from its (squared) theta constants

$$a := -11, \quad b := 22, \quad c := 19, \quad \text{and} \quad d := 3 \quad \text{in } \mathbb{F}_q.$$

	AVR ATmega	ARM Cortex M0	Symbolic cost
<code>bigint_mul</code>	1 654	410	
<code>bigint_sqr</code>	1 171	260	
<code>bigint_red</code>	438	71	
<code>gfe_mul</code>	1 952	502	M
<code>gfe_sqr</code>	1 469	353	S
<code>gfe_mulconst</code>	569	83	m_c
<code>gfe_add</code>	400	62	a
<code>gfe_sub</code>	401	66	s
<code>gfe_invert</code>	169 881	46 091	I
<code>gfe_powminhalf</code>	169 881	46 294	11M + 125S
<code>gfe_sqrtinv</code>	178 041	48 593	15M + 126S + 1neg

Table 2. Cycle counts for our field arithmetic implementation (including function-call overhead).

From these, we define the dual theta constants

$$\begin{aligned} A &:= a + b + c + d = 33, & B &:= a + b - c - d = -11, \\ C &:= a - b + c - d = -17, & D &:= a - b - c + d = -49. \end{aligned}$$

Observe that projectively,

$$\left(\frac{1}{a} : \frac{1}{b} : \frac{1}{c} : \frac{1}{d}\right) = (114 : -57 : -66 : -418)$$

and

$$\left(\frac{1}{A} : \frac{1}{B} : \frac{1}{C} : \frac{1}{D}\right) = (-833 : 2499 : 1617 : 561).$$

Crucially, all of these constants can be represented using just 16 bits. Since Kummer arithmetic involves many multiplications by these constants, we implement a separate 16×128 -bit multiplication function `gfe_mulconst`. For the AVR ATmega, we store the constants in two 8-bit registers. For the Cortex M0, the values fit into a halfword; this works well with the 16×16 -bit multiplication. Multiplication by any of these 16-bit constants costs **m_c**.

Continuing, we define $e/f := (1 + \alpha)/(1 - \alpha)$, where $\alpha^2 = CD/AB$ (we take the square root with least significant bit 0), and thus

$$\begin{aligned} \lambda &:= ac/bd = 0x15555555555555555555555555555552, \\ \mu &:= ce/df = 0x73E334FBB315130E05A505C31919A746, \\ \nu &:= ae/bf = 0x552AB1B63BF799716B5806482D2D21F3. \end{aligned}$$

These are the *Rosenhain invariants* of the curve \mathcal{C} , found by Gaudry and Schost [13], which we are (finally!) ready to define as

$$\mathcal{C} : Y^2 = f_{\mathcal{C}}(X) := X(X - 1)(X - \lambda)(X - \mu)(X - \nu).$$

The curve constants are the coefficients of $f_{\mathcal{C}}(X) = \sum_{i=0}^5 f_i X^i$: so $f_0 = 0$, $f_5 = 1$,

$$\begin{aligned} f_1 &= \text{0x1EDD6EE48E0C2F16F537CD791E4A8D6E} , \\ f_2 &= \text{0x73E799E36D9FCC210C9CD1B164C39A35} , \\ f_3 &= \text{0x4B9E333F48B6069CC47DC236188DF6E8} , \\ f_4 &= \text{0x219CC3F8BB9DFE2B39AD9E9F6463E172} . \end{aligned}$$

We store the squared theta constants $(a : b : c : d)$, along with $(1/a : 1/b : 1/c : 1/d)$, and $(1/A : 1/B : 1/C : 1/D)$; the Rosenhain invariants λ , μ , and ν , together with $\lambda\mu$ and $\lambda\nu$; and the curve constants f_1 , f_2 , f_3 , and f_4 , for use in our Kummer and Jacobian arithmetic functions. Obviously, none of the Rosenhain or curve constants are small; multiplying by these costs a full **M**.

3.3 Elements of $\mathcal{J}_{\mathcal{C}}$, compressed and decompressed.

Our algorithms use the usual Mumford representation for elements of $\mathcal{J}_{\mathcal{C}}(\mathbb{F}_q)$: they correspond to pairs $\langle u(X), v(X) \rangle$, where u and v are polynomials over \mathbb{F}_q with u monic, $\deg v < \deg u \leq 2$, and $v(X)^2 \equiv f_{\mathcal{C}}(X) \pmod{u(X)}$. We compute the group operation \oplus in $\mathcal{J}_{\mathcal{C}}(\mathbb{F}_q)$ using a function **ADD**, which implements the algorithm found in [14] (after a change of coordinates to meet their Assumption 1)³ at a cost of $28\mathbf{M} + 2\mathbf{S} + 11\mathbf{a} + 24\mathbf{s} + 1\mathbf{I}$.

For transmission, we compress the 508-bit Mumford representation to a 256-bit form. Our functions **compress** (Algorithm 1) and **decompress** (Algorithm 2) implement Stahlke’s compression technique (see [25] and [9, App. A] for details).

Algorithm 1: compress: compresses points on $\mathcal{J}_{\mathcal{C}}$ to 256-bit strings. Symbolic cost: $3\mathbf{M} + 1\mathbf{S} + 2\mathbf{a} + 2\mathbf{s}$. ATmega: 8 016 cycles. Cortex M0: 2 186 cycles.

Input: $\langle X^2 + u_1X + u_0, v_1X + v_0 \rangle = P \in \mathcal{J}_{\mathcal{C}}$.
Output: A string $b_0 \dots b_{255}$ of 256 bits.

- 1 $w \leftarrow 4((u_1 \cdot v_0 - u_0 \cdot v_1) \cdot v_1 - v_0^2)$ // $3\mathbf{M} + 1\mathbf{S} + 2\mathbf{a} + 2\mathbf{s}$
- 2 $b_0 \leftarrow \text{LeastSignificantBit}(v_1)$
- 3 $b_{128} \leftarrow \text{LeastSignificantBit}(w)$
- 4 **return** $b_0 || u_0 || b_{128} || u_1$

3.4 The Kummer surface $\mathcal{K}_{\mathcal{C}}$

The Kummer surface of \mathcal{C} is the quotient $\mathcal{K}_{\mathcal{C}} := \mathcal{J}_{\mathcal{C}} / \langle \pm 1 \rangle$; points on $\mathcal{K}_{\mathcal{C}}$ correspond to points on $\mathcal{J}_{\mathcal{C}}$ taken up to sign. If P is a point in $\mathcal{J}_{\mathcal{C}}$, then we write

$$(x_P : y_P : z_P : t_P) = \pm P$$

³ We only call **ADD** once in our algorithms, so for lack of space we omit its description.

Algorithm 2: decompress: decompresses 256-bit string to a point on \mathcal{J}_C .
Symbolic cost: $46\mathbf{M} + 255\mathbf{S} + 17\mathbf{a} + 12\mathbf{s} + 6\mathbf{neg}$. ATmega: 386 524 cycles
Cortex M0: 106 013 cycles

Input: A string $b_0 \cdots b_{255}$ of 256 bits.
Output: $\langle X^2 + u_1X + u_0, v_1X + v_0 \rangle = P \in \mathcal{J}_C$.

- 1 $\mathbf{U}_1 = b_{129} \cdots b_{256}$ as an element of \mathbb{F}_q
- 2 $\mathbf{U}_0 = b_1 \cdots b_{127}$ as an element of \mathbb{F}_q
- 3 $\mathbf{T}_1 \leftarrow \mathbf{U}_1^2$ // 1S
- 4 $\mathbf{T}_2 \leftarrow \mathbf{U}_0 - \mathbf{T}_1$ // 1s
- 5 $\mathbf{T}_3 \leftarrow \mathbf{U}_0 + \mathbf{T}_2$ // 1a
- 6 $\mathbf{T}_4 \leftarrow \mathbf{U}_0 \cdot (\mathbf{T}_3 \cdot f_4 + (\mathbf{U}_1 \cdot f_3 - 2f_2))$ // $3\mathbf{M} + 1\mathbf{a} + 2\mathbf{s}$
- 7 $\mathbf{T}_3 \leftarrow -\mathbf{T}_3$ // 1neg
- 8 $\mathbf{T}_1 \leftarrow \mathbf{T}_3 - \mathbf{U}_0$ // 1s
- 9 $\mathbf{T}_4 \leftarrow 2(\mathbf{T}_4 + (\mathbf{T}_1 \cdot \mathbf{U}_0 + f_1) \cdot \mathbf{U}_1)$ // $2\mathbf{M} + 3\mathbf{a}$
- 10 $\mathbf{T}_1 \leftarrow 2(\mathbf{T}_1 - \mathbf{U}_0)$ // $1\mathbf{a} + 1\mathbf{s}$
- 11 $\mathbf{T}_5 \leftarrow ((\mathbf{U}_0 - (f_3 + \mathbf{U}_1 \cdot (\mathbf{U}_1 - f_4))) \cdot \mathbf{U}_0 + f_1)^2$ // $2\mathbf{M} + 1\mathbf{S} + 2\mathbf{a} + 2\mathbf{s}$
- 12 $\mathbf{T}_5 \leftarrow \mathbf{T}_4^2 - 2\mathbf{T}_5 \cdot \mathbf{T}_1$ // $1\mathbf{M} + 1\mathbf{S} + 1\mathbf{a} + 1\mathbf{s}$
- 13 $(\mathbf{T}_6, \mathbf{T}_5) \leftarrow \mathbf{gfe_sqrtinv}(\mathbf{T}_5, \mathbf{T}_1, b_1)$ // $19\mathbf{M} + 127\mathbf{S} + 2\mathbf{neg}$
- 14 $\mathbf{T}_4 \leftarrow (\mathbf{T}_5 - \mathbf{T}_4) \cdot \mathbf{T}_6$ // $1\mathbf{M} + 1\mathbf{s}$
- 15 $\mathbf{T}_5 \leftarrow -f_4 \cdot \mathbf{T}_2 - ((\mathbf{T}_3 - f_3) \cdot \mathbf{U}_1) + f_2 + \mathbf{T}_4$ // $2\mathbf{M} + 2\mathbf{s} + 2\mathbf{a} + 1\mathbf{neg}$
- 16 $\mathbf{T}_6 = \mathbf{gfe_powminhalf}(4\mathbf{T}_6)$ // $= 1/(2v_1)$. $11\mathbf{M} + 125\mathbf{S} + 2\mathbf{a}$
- 17 $\mathbf{V}_1 \leftarrow 2\mathbf{T}_5 \cdot \mathbf{T}_6$ // $1\mathbf{M} + 1\mathbf{a}$
- 18 **if** $b_0 \neq \mathbf{LeastSignificantBit}(\mathbf{V}_1)$ **then** $(\mathbf{V}_1, \mathbf{T}_6) \leftarrow (-\mathbf{V}_1, -\mathbf{T}_6)$ // 2neg
- 19 $\mathbf{T}_5 \leftarrow (\mathbf{U}_1 \cdot f_4 + (\mathbf{T}_2 - f_3)) \cdot \mathbf{U}_0$ // $2\mathbf{M} + 1\mathbf{a} + 1\mathbf{s}$
- 20 $\mathbf{V}_0 \leftarrow (\mathbf{U}_1 \cdot \mathbf{T}_4 + \mathbf{T}_5 + f_1) \cdot \mathbf{T}_6$ // $2\mathbf{M} + 2\mathbf{a}$
- 21 **return** $\langle X^2 + \mathbf{U}_1X + \mathbf{U}_0, \mathbf{V}_1X + \mathbf{V}_0 \rangle$

for its image in $\mathcal{K}_{\mathcal{C}}$. To avoid subscript explosion, we make the following convention: when points P and Q on $\mathcal{J}_{\mathcal{C}}$ are clear from the context, we write

$$(x_{\oplus} : y_{\oplus} : z_{\oplus} : t_{\oplus}) = \pm(P \oplus Q) \quad \text{and} \quad (x_{\ominus} : y_{\ominus} : z_{\ominus} : t_{\ominus}) = \pm(P \ominus Q) .$$

The Kummer surface of this \mathcal{C} has a “fast” model in \mathbb{P}^3 defined by

$$\mathcal{K}_{\mathcal{C}} : E \cdot xyzt = \left(\begin{array}{c} (x^2 + y^2 + z^2 + t^2) \\ -F \cdot (xt + yz) - G \cdot (xz + yt) - H \cdot (xy + zt) \end{array} \right)^2$$

where

$$F = \frac{a^2 - b^2 - c^2 + d^2}{ad - bc}, \quad G = \frac{a^2 - b^2 + c^2 - d^2}{ac - bd}, \quad H = \frac{a^2 + b^2 - c^2 - d^2}{ab - cd},$$

and $E = 4abcd(ABCD/((ad - bc)(ac - bd)(ab - cd)))^2$ (see eg. [7], [8], and [12]). The identity point $\langle 1, 0 \rangle$ of $\mathcal{J}_{\mathcal{C}}$ maps to

$$\pm 0_{\mathcal{J}_{\mathcal{C}}} = (a : b : c : d) .$$

Algorithm 3 (Project) projects general points from $\mathcal{J}_{\mathcal{C}}(\mathbb{F}_q)$ into $\mathcal{K}_{\mathcal{C}}$. The “special” case, where u is linear, is treated in [9, §7.2].

Algorithm 3: Project: $\mathcal{J}_{\mathcal{C}} \rightarrow \mathcal{K}_{\mathcal{C}}$. Symbolic cost: $8\mathbf{M} + 1\mathbf{S} + 4\mathbf{m}_{\mathbf{c}} + 7\mathbf{a} + 4\mathbf{s}$. ATmega: 20 205 cycles. Cortex M0: 5 667 cycles.

Input: $\langle X^2 + u_1X + u_0, v_1X + v_0 \rangle = P \in \mathcal{J}_{\mathcal{C}}$.

Output: $(x_P : y_P : z_P : t_P) = \pm P \in \mathcal{K}_{\mathcal{C}}$.

1	$(\mathbb{T}_1, \mathbb{T}_2, \mathbb{T}_3, \mathbb{T}_4) \leftarrow (\mu - u_0, \lambda\nu - u_0, \nu - u_0, \lambda\mu - u_0)$	// 4s
2	$\mathbb{T}_5 \leftarrow \lambda + u_1$	// 1a
3	$\mathbb{T}_7 \leftarrow u_0 \cdot ((\mathbb{T}_5 + \mu) \cdot \mathbb{T}_3)$	// 2M + 1a
4	$\mathbb{T}_5 \leftarrow u_0 \cdot ((\mathbb{T}_5 + \nu) \cdot \mathbb{T}_1)$	// 2M + 1a
5	$(\mathbb{T}_6, \mathbb{T}_8) \leftarrow (u_0 \cdot ((\mu + u_1) \cdot \mathbb{T}_2 + \mathbb{T}_2), u_0 \cdot ((\nu + u_1) \cdot \mathbb{T}_4 + \mathbb{T}_4))$	// 4M + 4a
6	$\mathbb{T}_1 \leftarrow v_0^2$	// 1S
7	$(\mathbb{T}_5, \mathbb{T}_6, \mathbb{T}_7, \mathbb{T}_8) \leftarrow (\mathbb{T}_5 - \mathbb{T}_1, \mathbb{T}_6 - \mathbb{T}_1, \mathbb{T}_7 - \mathbb{T}_1, \mathbb{T}_8 - \mathbb{T}_1)$	// 4s
8	return $(a \cdot \mathbb{T}_5 : b \cdot \mathbb{T}_6 : c \cdot \mathbb{T}_7 : d \cdot \mathbb{T}_8)$	// 4m _c

3.5 Pseudo-addition on $\mathcal{K}_{\mathcal{C}}$.

While the points of $\mathcal{K}_{\mathcal{C}}$ do not form a group, we have a pseudo-addition operation (differential addition), which computes $\pm(P \oplus Q)$ from $\pm P$, $\pm Q$, and $\pm(P \ominus Q)$. The function `xADD` (Algorithm 4) implements the standard differential addition. The special case where $P = Q$ yields a pseudo-doubling operation.

To simplify the presentation of our algorithms, we define three operations on points in \mathbb{P}^3 . First, $\mathcal{M} : \mathbb{P}^3 \times \mathbb{P}^3 \rightarrow \mathbb{P}^3$ multiplies the corresponding coordinates of a pair of points:

$$\mathcal{M} : ((x_1 : y_1 : z_1 : t_1), (x_2 : y_2 : z_2 : t_2)) \mapsto (x_1 x_2 : y_1 y_2 : z_1 z_2 : t_1 t_2) .$$

The special case $(x_1 : y_1 : z_1 : t_1) = (x_2 : y_2 : z_2 : t_2)$ is denoted by

$$\mathcal{S} : (x : y : z : t) \mapsto (x^2 : y^2 : z^2 : t^2) .$$

Finally, the Hadamard transform⁴ is defined by

$$\mathcal{H} : (x : y : z : t) \mapsto (x' : y' : z' : t') \quad \text{where} \quad \begin{cases} x' = x + y + z + t , \\ y' = x + y - z - t , \\ z' = x - y + z - t , \\ t' = x - y - z + t . \end{cases}$$

Clearly \mathcal{M} and \mathcal{S} , cost $4\mathbf{M}$ and $4\mathbf{S}$, respectively. The Hadamard transform can easily be implemented with $4\mathbf{a} + 4\mathbf{s}$. However, the additions and subtractions are relatively cheap, making function call overhead a large factor. To minimize this we inline the Hadamard transform, trading a bit of code size for efficiency.

Algorithm 4: xADD: Differential addition on $\mathcal{K}_{\mathcal{C}}$. Symbolic cost: $14\mathbf{M} + 4\mathbf{S} + 4\mathbf{m}_{\mathcal{C}} + 12\mathbf{a} + 12\mathbf{s}$. ATmega: 34 774 cycles. Cortex M0: 9 598 cycles.

Input: $(\pm P, \pm Q, \pm(P \ominus Q)) \in \mathcal{K}_{\mathcal{C}}^3$ for some P and Q on $\mathcal{J}_{\mathcal{C}}$.
Output: $\pm(P \oplus Q) \in \mathcal{K}_{\mathcal{C}}$.

1	$(V_1, V_2) \leftarrow (\mathcal{H}(\pm P), \mathcal{H}(\pm Q))$	// 8a + 8s
2	$V_1 \leftarrow \mathcal{M}(V_1, V_2)$	// 4M
3	$V_1 \leftarrow \mathcal{M}(V_1, (1/A : 1/B : 1/C : 1/D))$	// 4m _C
4	$V_1 \leftarrow \mathcal{H}(V_1)$	// 4a + 4s
5	$V_1 \leftarrow \mathcal{S}(V_1)$	// 4S
6	$(C_1, C_2) \leftarrow (z_{\ominus} \cdot t_{\ominus}, x_{\ominus} \cdot y_{\ominus})$	// 2M
7	$V_2 \leftarrow \mathcal{M}((C_1 : C_1 : C_2 : C_2), (y_{\ominus} : x_{\ominus} : t_{\ominus} : z_{\ominus}))$	// 4M
8	return $\mathcal{M}(V_1, V_2)$	// 4M

Lines 5 and 6 of Algorithm 4 only involve the third argument, $\pm(P \ominus Q)$; essentially, they compute the point $(y_{\ominus} z_{\ominus} t_{\ominus} : x_{\ominus} z_{\ominus} t_{\ominus} : x_{\ominus} y_{\ominus} t_{\ominus} : x_{\ominus} y_{\ominus} z_{\ominus})$ (which is projectively equivalent to $(1/x_{\ominus} : 1/y_{\ominus} : 1/z_{\ominus} : 1/t_{\ominus})$, but requires no inversions; note that this is generally *not* a point on $\mathcal{K}_{\mathcal{C}}$). In practice, the pseudoadditions used in our scalar multiplication all use a fixed third argument, so it makes sense to precompute this “inverted” point and to scale it by x_{\ominus} so that the first coordinate is 1, thus saving $7\mathbf{M}$ in each subsequent differential

⁴ Observe that $(A : B : C : D) = \mathcal{H}((a : b : c : d))$ and, dually, $(a : b : c : d) = \mathcal{H}((A : B : C : D))$.

addition for a one-off cost of **1I**. The resulting data can be stored as the 3-tuple $(x_\ominus/y_\ominus, x_\ominus/z_\ominus, x_\ominus/t_\ominus)$, ignoring the trivial first coordinate: this is the *wrapped* form of $\pm(P \ominus Q)$. The function `xWRAP` (Algorithm 5) applies this transformation.

Algorithm 5: `xWRAP`: $(x : y : z : t) \mapsto (x/y, x/z, x/t)$. Symbolic cost: **7M** + **1I** ATmega: 182 251 cycles. Cortex M0: 49 609 cycles.

Input: $(x : y : z : t) \in \mathbb{F}^3$
Output: $(x/y, x/z, x/t) \in \mathbb{F}_q^3$.

```

1  $V_1 \leftarrow y \cdot z$  // 1M
2  $V_2 \leftarrow x/(V_1 \cdot t)$  // 2M + 1I
3  $V_3 \leftarrow V_2 \cdot t$  // 1M
4 return  $(V_3 \cdot z, V_3 \cdot y, V_1 \cdot V_2)$  // 3M
```

Algorithm 6 combines the pseudo-doubling with the differential addition, sharing intermediate operands, to define a differential double-and-add `xDBLADD`. This is the fundamental building block of the Montgomery ladder.

Algorithm 6: `xDBLADD`: Combined differential double-and-add. The difference point is wrapped. Symbolic cost: **7M** + **12S** + **12m_c** + **16a** + **16s**. ATmega: 36 706 cycles. Cortex M0: 9 861 cycles.

Input: $(\pm P, \pm Q, (x_\ominus/y_\ominus, x_\ominus/z_\ominus, x_\ominus/t_\ominus)) \in \mathcal{K}_C^2 \times \mathbb{F}_q$.
Output: $(\pm[2]P, \pm(P \oplus Q)) \in \mathcal{K}_C^2$.

```

1  $(V_1, V_2) \leftarrow (\mathcal{S}(\pm P), \mathcal{S}(\pm Q))$  // 8S
2  $(V_1, V_2) \leftarrow (\mathcal{H}(V_1), \mathcal{H}(V_2))$  // 8a + 8s
3  $(V_1, V_2) \leftarrow (\mathcal{S}(V_1), \mathcal{M}(V_1, V_2))$  // 4M+4S
4  $(V_1, V_2) \leftarrow (\mathcal{M}(V_1, (\frac{1}{A} : \frac{1}{B} : \frac{1}{C} : \frac{1}{D})), \mathcal{M}(V_2, (\frac{1}{A} : \frac{1}{B} : \frac{1}{C} : \frac{1}{D})))$  // 8mc
5  $(V_1, V_2) \leftarrow (\mathcal{H}(V_1), \mathcal{H}(V_2))$  // 8a + 8s
6 return  $(\mathcal{M}(V_1, (\frac{1}{a} : \frac{1}{b} : \frac{1}{c} : \frac{1}{d})), \mathcal{M}(V_2, (1 : \frac{x_\ominus}{y_\ominus} : \frac{x_\ominus}{y_\ominus} : \frac{x_\ominus}{t_\ominus})))$  // 3M + 4mc
```

4 Scalar multiplication

All of our cryptographic routines are built around scalar multiplication in \mathcal{J}_C and pseudo-scalar multiplication in \mathcal{K}_C . We implement pseudo-scalar multiplication using the classic Montgomery ladder in §4.1. In §4.2, we extend this to full scalar multiplication on \mathcal{J}_C using the point recovery technique proposed in [9].

4.1 Pseudomultiplication on \mathcal{K}_C

Since $[m](\ominus P) = \ominus[m]P$ for all m and P , we have a pseudo-scalar multiplication operation $(m, \pm P) \mapsto \pm[m]P$ on \mathcal{K}_C , which we compute using Algorithm 7

	M	S	m _c	a	s	neg	I	ATmega	Cortex M0
ADD	28	2	0	11	24	0	1	228 552	62 886
Project	8	1	4	7	8	0	0	20 205	5 667
xWRAP	7	0	0	0	0	0	1	182 251	49 609
xUNWRAP	4	0	0	0	0	0	0	7 297	2 027
xADD	14	4	4	12	12	0	0	34 774	9 598
xDBLADD	7	12	12	16	16	0	0	36 706	9 861
recoverGeneral	77	8	0	19	10	3	1	318 910	88 414
fast2genPartial	11	0	0	9	0	0	0	21 339	6 110
fast2genFull	15	0	0	12	0	0	0	29 011	8 333
recoverFast	139	12	4	70	22	5	1	447 176	124 936
compress	3	1	0	2	2	0	0	8 016	2 186
decompress	46	255	0	17	12	6	0	386 524	106 013

Table 3. Operation and cycle counts of basic functions on the Kummer and Jacobian.

(the Montgomery ladder), implemented as `crypto_scalarmult`. The loop of Algorithm 7 maintains the following invariant: at the end of iteration i we have

$$(V_1, V_2) = (\pm[k]P, \pm[k+1]P) \quad \text{where} \quad k = \sum_{j=i}^{\beta-1} m_j 2^{\beta-1-i}.$$

Hence, at the end we return $\pm[m]P$, and also $\pm[m+1]P$ as a (free) byproduct. We assume that we have a constant-time conditional swap routine `CSWAP`($b, (V_1, V_2)$), which returns (V_1, V_2) if $b = 0$ and (V_2, V_1) if $b = 1$. This makes the execution of Algorithm 7 uniform and constant-time, which means it is suitable for use with secret values of m .

Algorithm 7: `crypto_scalarmult`: Montgomery ladder on \mathcal{K}_C . Uniform and constant-time: may be used for secret scalars. The point is wrapped. Symbolic cost: $(4 + 7\beta)\mathbf{M} + 12\beta\mathbf{S} + 12\beta\mathbf{m}_c + 16\beta\mathbf{a} + 16\beta\mathbf{s}$, where $\beta =$ scalar bitlength. ATmega: 9 513 536 cycles. Cortex: 2 633 662 cycles.

Input: $(m = \sum_{i=0}^{\beta-1} m_i 2^i, (x_P/y_P, x_P/z_P, x_P/t_P)) \in [0, 2^\beta) \times \mathbb{F}_q^3$ for $\pm P$ in \mathcal{K}_C .
Output: $(\pm[m]P, \pm[m+1]P) \in \mathcal{K}_C^2$.

- 1 $V_1 \leftarrow (a : b : c : d)$
- 2 $V_2 \leftarrow \text{xUNWRAP}(x_P/y_P, x_P/z_P, x_P/t_P)$ // = $\pm P$. 4M
- 3 **for** $i = 250$ **down to** 0 **do** // $7\beta\mathbf{M} + 12\beta\mathbf{S} + 12\beta\mathbf{m}_c + 16\beta\mathbf{a} + 16\beta\mathbf{s}$
- 4 $(V_1, V_2) \leftarrow \text{CSWAP}(m_i, (V_1, V_2))$
- 5 $(V_1, V_2) \leftarrow \text{xDBLADD}(V_1, V_2, (x_P/y_P, x_P/z_P, x_P/t_P))$
- 6 $(V_1, V_2) \leftarrow \text{CSWAP}(m_i, (V_1, V_2))$
- 7 **return** (V_1, V_2)

Our implementation of `crypto_scalarmult` assumes that its input Kummer point $\pm P$ is wrapped. This follows the approach of [3]. Indeed, many calls

to `crypto_scalarmult` involve Kummer points that are stored or transmitted in wrapped form. However, `crypto_scalarmult` does require the unwrapped point internally—if only to initialize one variable. We therefore define a function `xUNWRAP` (Algorithm 8) to invert the `xWRAP` transformation at a cost of only 4M.

Algorithm 8: `xUNWRAP`: $(x/y, x/z, x/t) \mapsto (x : y : z : t)$. Symbolic cost: 4M. ATmega: 7 297 cycles. Cortex: 2 027 cycles.

Input: $(u, v, w) \in \mathbb{F}_q^3$ s.t. $u = x_P/y_P, v = x_P/z_P, w = x_P/t_P$ for $\pm P \in \mathcal{K}_C$
Output: $(x_P : y_P : z_P : t_P) \in \mathbb{P}^3$

```

1  $(T_1, T_2, T_3) \leftarrow (v \cdot w, u \cdot w, u \cdot v)$  // 3M
2 return  $(T_3 \cdot w : T_1 : T_2 : T_3)$  // 1M
```

4.2 Point recovery from \mathcal{K}_C to \mathcal{J}_C

Point recovery means efficiently computing $[m]P$ on \mathcal{J}_C given $\pm[m]P$ on \mathcal{K}_C and some additional information. In our case, the additional information is the base point P and the second output of the Montgomery ladder, $\pm[m+1]P$.

Algorithm 9 (`Recover`) implements the point recovery algorithm described in [9]. This is the genus-2 analogue of the point recovery methods defined for elliptic curves in [19], [22], and [5].

Algorithm 9: `Recover`: From \mathcal{K}_C to \mathcal{J}_C . Symbolic cost: 139M + 12S + 4m_c + 70a + 22s + 3neg + 1I. ATmega: 447 176 cycles. Cortex: 124 936 cycles.

Input: $(P, \pm P, \pm Q, \pm(P \oplus Q)) \in \mathcal{J}_C \times \mathcal{K}_C^3$ for some P, Q in \mathcal{J}_C .
Output: $Q \in \mathcal{J}_C$.

```

1  $gP \leftarrow \text{fast2genPartial}(\pm P)$  // 11M + 9a
2  $gQ \leftarrow \text{fast2genFull}(\pm Q)$  // 15M + 12a
3  $gS \leftarrow \text{fast2genPartial}(\pm(P \oplus Q))$  // 11M + 9a
4  $\times D \leftarrow \text{xADD}(\pm P, \pm Q, \pm(P \oplus Q))$  // 14M + 4S + 4m_c + 12a + 12s
5  $gD \leftarrow \text{fast2genPartial}(\times D)$  // 11M + 9a
6 return recoverGeneral( $P, gP, gQ, gS, gD$ ) // 77M+8S+19a+10s+3neg+1I
```

While we refer the reader to [9] for technical details on this method, and proof of its correctness, there is one important mathematical detail that we should mention, since it is reflected in the structure of our code. Namely, point recovery is more naturally computed starting from the general Flynn model $\tilde{\mathcal{K}}_C$ of the Kummer, because it is more closely related to the Mumford model for \mathcal{J}_C . Algorithm 9 therefore proceeds in two steps: first we map the problem onto $\tilde{\mathcal{K}}_C$

using Algorithms 10 and 11 (`fast2genFull` and `fast2genPartial`), and then we recover from $\widetilde{\mathcal{K}}_C$ to \mathcal{J}_C using Algorithm 12 (`recoverGeneral`).

Since the general Kummer \mathcal{K}_C only appears briefly in our recovery procedure (we never use its relatively slow arithmetic operations), we will not investigate it any further here—but the curious reader may refer to [6] for the general theory. For our purposes, it suffices to recall that $\widetilde{\mathcal{K}}_C$ is, like \mathcal{K}_C , embedded in \mathbb{P}^3 ; and the isomorphism $\mathcal{K}_C \rightarrow \widetilde{\mathcal{K}}_C$ is defined (in eg. [9, §7.4]) by the linear transformation

$$(x_P : y_P : z_P : t_P) \mapsto (\tilde{x}_P : \tilde{y}_P : \tilde{z}_P : \tilde{t}_P) := (x_P : y_P : z_P : t_P)L,$$

where L is (any scalar multiple of) the matrix

$$\begin{pmatrix} a^{-1}(\nu - \lambda) & a^{-1}(\mu\nu - \lambda) & a^{-1}\lambda\nu(\mu - 1) & a^{-1}\lambda\nu(\mu\nu - \lambda) \\ b^{-1}(\mu - 1) & b^{-1}(\mu\nu - \lambda) & b^{-1}\mu(\nu - \lambda) & b^{-1}\mu(\mu\nu - \lambda) \\ c^{-1}(\lambda - \mu) & c^{-1}(\lambda - \mu\nu) & c^{-1}\lambda\mu(1 - \nu) & c^{-1}\lambda\mu(\lambda - \mu\nu) \\ d^{-1}(1 - \nu) & d^{-1}(\lambda - \mu\nu) & d^{-1}\nu(\lambda - \mu) & d^{-1}\nu(\lambda - \mu\nu) \end{pmatrix},$$

which we precompute and store. If $\pm P$ is a point on \mathcal{K}_C , then $\widetilde{\pm P}$ denotes its image on $\widetilde{\mathcal{K}}_C$; we compute $\widetilde{\pm P}$ using Algorithm 10 (`fast2genFull`).

Algorithm 10: `fast2genFull`: The map $\mathcal{K}_C \rightarrow \widetilde{\mathcal{K}}_C$. Symbolic cost: 15M + 12a. ATmega: 29 011 cycles. Cortex: 8 333 cycles.

Input: $\pm P \in \mathcal{K}_C$
Output: $\widetilde{\pm P} \in \widetilde{\mathcal{K}}_C$.

```

1  $\tilde{x}_P \leftarrow x_P + (L_{12}/L_{11})y_P + (L_{13}/L_{11})z_P + (L_{14}/L_{11})t_P$  // 3M + 3a
2  $\tilde{y}_P \leftarrow (L_{21}/L_{11})x_P + (L_{22}/L_{11})y_P + (L_{23}/L_{11})z_P + (L_{24}/L_{11})t_P$  // 4M + 3a
3  $\tilde{z}_P \leftarrow (L_{31}/L_{11})x_P + (L_{32}/L_{11})y_P + (L_{33}/L_{11})z_P + (L_{34}/L_{11})t_P$  // 4M + 3a
4  $\tilde{t}_P \leftarrow (L_{41}/L_{11})x_P + (L_{42}/L_{11})y_P + (L_{43}/L_{11})z_P + (L_{44}/L_{11})t_P$  // 4M + 3a
5 return  $(\tilde{x}_P : \tilde{y}_P : \tilde{z}_P : \tilde{t}_P)$ 

```

Sometimes we only require the first three coordinates of $\widetilde{\pm P}$. Algorithm 11 (`fast2genPartial`) saves 4M + 3a per point by not computing \tilde{t}_P .

4.3 Full scalar multiplication on \mathcal{J}_C

We now combine our pseudo-scalar multiplication function `crypto_scalarmult` with the point-recovery function `Recover` to define a full scalar multiplication function `jacobian_scalarmult` (Algorithm 13) on \mathcal{J}_C .

Remark 3. We have designed `jacobian_scalarmult` to take not only a scalar m and a Jacobian point P in its Mumford representation, but also the wrapped form of $\pm P$ as an auxiliary argument: that is, we assume that `xP` \leftarrow `Project(P)`

Algorithm 11: fast2genPartial: The map $\mathcal{K}_C \rightarrow \mathbb{P}^2$. Symbolic cost: 11M + 9a. ATmega: 21 339 cycles. Cortex: 8 333 cycles.

Input: $\pm P \in \mathcal{K}_C$.
Output: $(\tilde{x}_P : \tilde{y}_P : \tilde{z}_P) \in \mathbb{P}^2$

- 1 $\tilde{x}_P \leftarrow x_P + (L_{12}/L_{11})y_P + (L_{13}/L_{11})z_P + (L_{14}/L_{11})t_P$ // 3M + 3a
- 2 $\tilde{y}_P \leftarrow (L_{21}/L_{11})x_P + (L_{22}/L_{11})y_P + (L_{23}/L_{11})z_P + (L_{24}/L_{11})t_P$ // 4M + 3a
- 3 $\tilde{z}_P \leftarrow (L_{31}/L_{11})x_P + (L_{32}/L_{11})y_P + (L_{33}/L_{11})z_P + (L_{34}/L_{11})t_P$ // 4M + 3a
- 4 **return** $(\tilde{x}_P : \tilde{y}_P : \tilde{z}_P)$

Algorithm 12: recoverGeneral: From $\tilde{\mathcal{K}}_C$ to \mathcal{J}_C . Symbolic cost: 77M + 8S + 19a + 10s + 3neg + 1I. ATmega: 318 910 cycles. Cortex: 88 414 cycles.

Input: $(P, \pm \tilde{P}, \pm \tilde{Q}, \pm(\tilde{P} \oplus \tilde{Q}), \pm(\tilde{P} \ominus \tilde{Q})) \in \mathcal{J}_C \times \tilde{\mathcal{K}}_C^4$ for some P and Q in \mathcal{J}_C .
The values of \tilde{t}_P , \tilde{t}_\oplus , and \tilde{t}_\ominus are not required.

Output: $Q \in \mathcal{J}_C$.

- 1 $(Z1, Z2) \leftarrow (\tilde{y}_P \cdot \tilde{x}_Q - \tilde{x}_Q \cdot \tilde{y}_P, \tilde{x}_P \cdot \tilde{z}_Q - \tilde{z}_P \cdot \tilde{x}_Q)$ // 4M+2s
- 2 $T1 \leftarrow Z1 \cdot \tilde{z}_P$ // 1M
- 3 $mZ3 \leftarrow Z2 \cdot \tilde{y}_P + T1$ // 1M + 1a
- 4 $D \leftarrow Z2^2 \cdot \tilde{x}_P + mZ3 \cdot Z1$ // 2M + 1S + 1a
- 5 $T2 \leftarrow Z1 \cdot Z2$ // 1M
- 6 $T3 \leftarrow \tilde{x}_P \cdot \tilde{x}_Q$ // 1M
- 7 $E \leftarrow T3 \cdot (f_2 \cdot Z2^2 - f_1 \cdot T2) + \tilde{t}_Q \cdot D$ // 5M + 1S + 1a + 1s
- 8 $E \leftarrow E + mZ3 \cdot \tilde{x}_Q^2 \cdot (f_3 \cdot Z2 \cdot \tilde{x}_P + f_4 \cdot mZ3)$ // 5M + 1S + 2a
- 9 $E \leftarrow E + mZ3 \cdot \tilde{x}_Q \cdot (mZ3 \cdot \tilde{y}_Q - Z2 \cdot \tilde{x}_P \cdot \tilde{z}_Q)$ // 5M + 1a + 1s
- 10 $X1 \leftarrow \tilde{x}_P \cdot (Z2 \cdot v_1(P) - Z1 \cdot v_0(P))$ // 3M + 1s
- 11 $T4 \leftarrow Z1 \cdot \tilde{y}_P + Z2 \cdot \tilde{x}_P$ // 2M + 1a
- 12 $X2 \leftarrow T1 \cdot v_1(P) + T4 \cdot v_0(P)$ // 2M + 1a
- 13 $C5 \leftarrow Z1^2 - T4 \cdot \tilde{x}_Q$ // 1M + 1S + 1s
- 14 $C6 \leftarrow T1 \cdot \tilde{x}_Q + T2$ // 1M + 1a
- 15 $T5 \leftarrow \tilde{z}_\oplus \cdot \tilde{x}_\ominus - \tilde{x}_\oplus \cdot \tilde{z}_\ominus$ // 2M + 1s
- 16 $X3 \leftarrow X1 \cdot T5 - X2 \cdot (\tilde{x}_\oplus \cdot \tilde{y}_\ominus - \tilde{y}_\oplus \cdot \tilde{x}_\ominus)$ // 4M + 2s
- 17 $(X5, X6) \leftarrow (X3 \cdot C5, X3 \cdot C6)$ // 2M
- 18 $X4 \leftarrow T3 \cdot (X1 \cdot (\tilde{z}_\oplus \cdot \tilde{y}_\ominus - \tilde{y}_\oplus \cdot \tilde{z}_\ominus) + T5 \cdot X2)$ // 5M + 1a + 1s
- 19 $(X7, X8) \leftarrow (X5 + Z1 \cdot X4, X6 + Z2 \cdot Z4)$ // 2M + 2a
- 20 $T6 \leftarrow \tilde{x}_\oplus \cdot \tilde{x}_\ominus$ // 1M
- 21 $E \leftarrow -T6 \cdot T3 \cdot (E \cdot \tilde{x}_P^2 + (X1 \cdot T3)^2)$ // 5M + 2S + 1a + 1neg
- 22 $(X9, X10) \leftarrow (E \cdot X7, E \cdot X8)$ // 2M
- 23 $F \leftarrow X2 \cdot (\tilde{x}_\oplus \cdot \tilde{y}_\ominus + \tilde{y}_\oplus \cdot \tilde{x}_\ominus) + X1 \cdot (\tilde{z}_\oplus \cdot \tilde{x}_\ominus + \tilde{x}_\oplus \cdot \tilde{z}_\ominus)$ // 6M + 3a
- 24 $F \leftarrow X1 \cdot F + 2(X2^2 \cdot T6)$ // 2M + 1S + 2a
- 25 $F \leftarrow -2(F \cdot D \cdot T6 \cdot T3 \cdot T3^2 \cdot \tilde{x}_P)$ // 5M + 1S + 1a + 1neg
- 26 $(U1, U0) \leftarrow (-F \cdot \tilde{y}_Q, F \cdot \tilde{z}_Q)$ // 2M + 1neg
- 27 $Fi \leftarrow 1/(F \cdot \tilde{x}_Q)$ // 1M + 1I
- 28 $(u'_1, u'_0, v'_1, v'_0) \leftarrow (Fi \cdot U1, Fi \cdot U0, Fi \cdot X9, Fi \cdot X10)$ // 4M
- 29 **return** $\langle X^2 + u'_1 X + u'_0, v'_1 X + v'_0 \rangle$

Algorithm 13: `jacobian_scalarmult`: Scalar multiplication on \mathcal{J}_C , using the Montgomery ladder on \mathcal{K}_C and recovery to \mathcal{J}_C . Assumes wrapped projected point as auxiliary input. Symbolic cost: $(7\beta + 143)\mathbf{M} + (12\beta + 12)\mathbf{S} + (12\beta + 4)\mathbf{m}_c + (70 + 16\beta)\mathbf{a} + (22 + 16\beta)\mathbf{s} + 3\mathbf{neg} + \mathbf{I}$. ATmega: 9 968 127 cycles. Cortex: 2 709 401 cycles.

Input: $(m, P, (x_P/y_P, x_P/z_P, x_P/t_P)) \in [0, 2^\beta) \times \mathcal{J}_C$
Output: $[m]P \in \mathcal{J}_C$

```

1 (X0, X1) ← crypto_scalarmult( $m, (x_P/y_P, x_P/z_P, x_P/t_P)$ )
                                     // (7β + 4)M+12βS+12βmc+16βa+16βs
2 xP ← xUNWRAP( $(x_P/y_P, x_P/z_P, x_P/t_P)$ ) // 4M
3 return Recover( $P, xP, X_0, X_1$ ) // 139M+12S+4mc+70a+22s+3neg+1I

```

and `xWRAP(xP)` have already been carried out as a precomputation. This saves redundant `Projecting` and `xWRAPPING` when we are operating on fixed base points, as is often the case in our protocols. Nevertheless, `jacobian_scalarmult` could easily be converted to a “pure” Jacobian scalar multiplication function (with no auxiliary input) by inserting appropriate `Project` and `xWRAP` calls at the start, and removing the `xUNWRAP` call at Line 2. These modifications would increase the cost of `jacobian_scalarmult` by $11\mathbf{M} + 1\mathbf{S} + 4\mathbf{m}_c + 7\mathbf{a} + 8\mathbf{s} + 1\mathbf{I}$.

5 Results and comparison

The high-level cryptographic functions for our signature scheme are named `keygen`, `sign` and `verify`. Their implementations contain no surprises: they do exactly what was specified in §2.1, calling the lower-level functions described in §3 and §4 as required. Our key exchange uses only the function `dh_exchange`, for both Diffie-Hellman key generation and key exchange. It implements exactly what we specified in §2.2: it is a call to `crypto_scalarmult` plus a call to `xWRAP` to convert to the correct 384-bit representation. Table 1 (in the introduction) presents the cycle counts and stack usage for all of our high-level functions.

5.1 Code and compilation

In our experiments, the code for the AVR ATmega was compiled with `avr-gcc` at optimization level `-O2`. The ARM Cortex M0 code uses the `clang` compiler, also with optimization level `-O2`. We experimented with different optimization levels (`-O3`, `-O1`, and `-Os`), but the results were fairly similar. The total size of the program is 20 242 bytes for the AVR ATmega, and 19 606 bytes for the ARM Cortex M0. This consists of the full signature and key-exchange code, including the hash function `SHAKE128` implemented with fixed 512-bit output, with the code taken from the reference implementation.⁵

⁵ We used the reference C implementation for the Cortex M0, and the assembly implementation for AVR; both are available on <http://keccak.noekeon.org/>. The only

5.2 Comparison

As we believe ours to be the first genus-2 hyperelliptic curve implementation on both the AVR ATmega and the ARM Cortex M0 architectures, it is difficult to make a comparison. On the other hand, we may compare with elliptic curve-based alternatives at the same 128-bit security level: notably [28], [16], [27], and [10].

If one is only interested in a key exchange scheme, it is enough to have an object which merely has a differential addition. A well-known example is the Montgomery model for elliptic curves, with very efficient x -coordinate-only arithmetic, which was used in eg. [28], [16], and [10] to obtain efficient Diffie–Hellman key exchange. It is also possible to use x -only arithmetic on Weierstrass curves, as we see in [27] (which is based on the pseudo-addition presented in [5]). In genus 2, Kummer surfaces have similar properties. Although they only have a pseudo-addition, it is very efficient and therefore highly suitable for key exchange.

If one also wants to implement signatures, differential addition is no longer sufficient. In this case, one has to work with points in an elliptic curve group (or the Jacobian of a genus 2 curve). To still make use of the efficient x -only arithmetic, one must project a point (x, y) to the x -only representation, do a pseudo-scalar multiplication, and then recover the correct y -coordinate of the result. (This is done in [27] using a Weierstrass curve.) As recovery is generally quite slow, this imposes non-negligible overhead.

In genus 2 we can use a completely analogous technique: we project points from the Jacobian to the Kummer surface, use its efficient arithmetic for the scalar multiplication, and recover the correct element of the Jacobian. As in the elliptic case, this does impose some overhead. We therefore only do this when really necessary: that is, for signatures. When computing shared secrets in key exchange, we remain on the Kummer surface.

	Implementation	Object	Clock cycles	Code size	Stack
S,DH	Wenger et al. [27]	NIST P-256	$\approx 10\,730\,000$	7 168 bytes	540 bytes
DH	Düll et al. [10]	Curve25519	3 589 850	7 900 bytes	548 bytes
DH	This work	\mathcal{K}_C	2 633 662	$\approx 4\,328$ bytes	248 bytes
S	This work	\mathcal{J}_C	2 709 401	$\approx 9\,874$ bytes	968 bytes

Table 4. Comparison of scalar multiplication routines on the ARM Cortex M0 architecture at the 128-bit security level. **S** denotes signature-compatible full scalar multiplication; **DH** denotes Diffie–Hellman pseudo-scalar multiplication.

As we see in Table 4, genus-2 techniques give great results for Diffie–Hellman key exchange on the ARM Cortex M0 architecture. Comparing with the current fastest implementation [10], we reduce the number of clock cycles by about 27%,

change required is to the padding, which must take domain separation into account according to [11, p.28].

while about halving code size and stack usage. For signatures, the state-of-the-art is [27]: here we reduce the cycle count for the underlying scalar multiplications by a very impressive 75%, at the cost of a moderate increase in code size and stack usage.

	Implementation	Object	Cycles	Code size	Stack
DH	Liu et al. [28]	256-bit curve	$\approx 21\,078\,200$	14 700 bytes*	556 bytes
S,DH	Wenger et al. [27]	NIST P-256	$\approx 34\,930\,000$	16 112 bytes	590 bytes
DH	Hutter, Schwabe [16]	Curve25519	22 791 579	n/a [†]	677 bytes
DH	Düll et al. [10]	Curve25519	13 900 397	17 710 bytes	494 bytes
DH	This work	\mathcal{K}_c	9 513 536	$\approx 9\,490$ bytes	99 bytes
S	This work	\mathcal{J}_c	9 968 127	$\approx 16\,516$ bytes	735 bytes

Table 5. Comparison of scalar multiplication routines on the AVR ATmega architecture at the 128-bit security level. **S** denotes signature-compatible full scalar multiplication; **DH** denotes Diffie–Hellman pseudo-scalar multiplication. The implementation marked * also contains a fixed-basepoint scalar multiplication routine, whereas the implementation marked [†] does not report code size for the separated scalar multiplication.

Looking at Table 5, on the AVR ATmega architecture we reduce the cycle count for Diffie–Hellman by about 32% compared with the current record [10], again roughly halving the code size, and reducing stack usage by about 80%. The Jacobian scalar multiplication needed for signatures, reduces the cycle count by 71% compared to [27], while increasing the stack usage by 25%.

Finally we can compare to the currently fastest full signature implementation [21], shown in Table 6.

Implementation	Object	Function	Cycles	Stack
Nascimento et al. [21]	Ed25519	sig. gen.	19 047 706	1 473 bytes
Nascimento et al. [21]	Ed25519	sig. ver.	30 776 942	1 226 bytes
This work	\mathcal{J}_c	sign	10 404 033	926 bytes
This work	\mathcal{J}_c	verify	16 240 510	992 bytes

Table 6. Comparison of a full signature scheme on the AVR ATmega architecture at the 128-bit security level.

We see that we almost half the number of cycles, while also reducing the stack usage by a decent margin. We do not compare code size, as this is not reported in [21].

References

1. Lejla Batina, David Hwang, Alireza Hodjat, Bart Preneel, and Ingrid Verbauwhede. Hardware/software co-design for hyperelliptic curve cryptography (HECC) on the 8051 μ P. In Josyula R. Rao and Berk Sunar, editors, *Cryptographic Hardware and Embedded Systems - CHES 2005, 7th International Workshop, Edinburgh, UK, August 29 - September 1, 2005, Proceedings*, volume 3659 of *Lecture Notes in Computer Science*, pages 106–118. Springer, 2005. 2
2. Daniel J. Bernstein. Differential addition chains. preprint, 2006. 4
3. Daniel J. Bernstein, Chitchanok Chuengsatiansup, Tanja Lange, and Peter Schwabe. Kummer strikes back: New DH speed records. In Sarkar and Iwata [23], pages 317–337. 1, 5, 13
4. Daniel J. Bernstein, Niels Duif, Tanja Lange, Peter Schwabe, and Bo-Yin Yang. High-speed high-security signatures. *J. Cryptographic Engineering*, 2(2):77–89, 2012. 2, 3
5. Eric Brier and Marc Joye. Weierstraß elliptic curves and side-channel attacks. In David Naccache and Pascal Paillier, editors, *Public Key Cryptography, 5th International Workshop on Practice and Theory in Public Key Cryptosystems, PKC 2002, Paris, France, February 12-14, 2002, Proceedings*, volume 2274 of *Lecture Notes in Computer Science*, pages 335–345. Springer, 2002. 14, 18
6. J. W. S. Cassels and E. V. Flynn. *Prolegomena to a middlebrow arithmetic of curves of genus 2*, volume 230. Cambridge University Press, 1996. 15
7. David V. Chudnovsky and Gregory V. Chudnovsky. Sequences of numbers generated by addition in formal groups and new primality and factorization tests. *Adv. in Appl. Math.*, 7:385–434, 1986. 10
8. Romain Cosset. *Applications of theta functions for hyperelliptic curve cryptography*. Ph.D Thesis, Université Henri Poincaré - Nancy I, November 2011. 10
9. Craig Costello, Ping-Ngai Chung, and Benjamin Smith. Fast, uniform, and compact scalar multiplication for elliptic curves and genus 2 jacobians with applications to signature schemes. Cryptology ePrint Archive, Report 2015/983, 2015. <https://eprint.iacr.org/2015/983>. 2, 3, 4, 5, 8, 10, 12, 14, 15
10. Michael Düll, Björn Haase, Gesine Hinterwälder, Michael Hutter, Christof Paar, Ana Helena Sánchez, and Peter Schwabe. High-speed curve25519 on 8-bit, 16-bit and 32-bit microcontrollers. *Design, Codes and Cryptography*, 77(2), 2015. <http://cryptojedi.org/papers/#mu25519>. 2, 6, 18, 19
11. Morris J. Dworkin. SHA-3 standard: Permutation-based hash and extendable-output functions. Technical report, jul 2015. 3, 18
12. Pierrick Gaudry. Fast genus 2 arithmetic based on Theta functions. *J. Mathematical Cryptology*, 1(3):243–265, 2007. 10
13. Pierrick Gaudry and Eric Schost. Genus 2 point counting over prime fields. *J. Symb. Comput.*, 47(4):368–400, 2012. 7
14. Huseyin Hisil and Craig Costello. Jacobian coordinates on genus 2 curves. In Sarkar and Iwata [23], pages 338–357. 8
15. Alireza Hodjat, Lejla Batina, David Hwang, and Ingrid Verbauwhede. HW/SW co-design of a hyperelliptic curve cryptosystem using a microcode instruction set coprocessor. *Integration*, 40(1):45–51, 2007. 2
16. Michael Hutter and Peter Schwabe. NaCl on 8-bit AVR microcontrollers. In Amr Youssef and Abderrahmane Nitaj, editors, *Progress in Cryptology – AFRICACRYPT 2013*, volume 7918 of *Lecture Notes in Computer Science*, pages 156–172. Springer-Verlag Berlin Heidelberg, 2013. <http://cryptojedi.org/papers/#avrnacl>. 18, 19

17. Michael Hutter and Peter Schwabe. Multiprecision multiplication on AVR revisited, 2014. <http://cryptojedi.org/papers/#avrmul>. 6
18. Jonathan Katz and Nan Wang. Efficiency improvements for signature schemes with tight security reductions. In Sushil Jajodia, Vijayalakshmi Atluri, and Trent Jaeger, editors, *Proceedings of the 10th ACM Conference on Computer and Communications Security, CCS 2003, Washington, DC, USA, October 27-30, 2003*, pages 155–164. ACM, 2003. 4
19. Julio López and Ricardo Dahab. Fast multiplication on elliptic curves over $\text{GF}(2^m)$ without precomputation. In Çetin Kaya Koç and Christof Paar, editors, *Cryptographic Hardware and Embedded Systems, First International Workshop, CHES'99, Worcester, MA, USA, August 12-13, 1999, Proceedings*, volume 1717 of *Lecture Notes in Computer Science*, pages 316–327. Springer, 1999. 14
20. David M'Raihi, David Naccache, David Pointcheval, and Serge Vaudenay. Computational alternatives to random number generators. In Stafford E. Tavares and Henk Meijer, editors, *Selected Areas in Cryptography '98, SAC'98, Kingston, Ontario, Canada, August 17-18, 1998, Proceedings*, volume 1556 of *Lecture Notes in Computer Science*, pages 72–80. Springer, 1998. 4
21. Erick Nascimento, Julio López, and Ricardo Dahab. Efficient and secure elliptic curve cryptography for 8-bit avr microcontrollers. In Rajat Subhra Chakraborty, Peter Schwabe, and Jon Solworth, editors, *Security, Privacy, and Applied Cryptography Engineering*, volume 9354 of *LNCS*. Springer, 2015. 2, 19
22. Katsuyuki Okeya and Kouichi Sakurai. Efficient elliptic curve cryptosystems from a scalar multiplication algorithm with recovery of the y-coordinate on a Montgomery-form elliptic curve. In Çetin Kaya Koç, David Naccache, and Christof Paar, editors, *Cryptographic Hardware and Embedded Systems - CHES 2001, Third International Workshop, Paris, France, May 14-16, 2001, Proceedings*, volume 2162 of *Lecture Notes in Computer Science*, pages 126–141. Springer, 2001. 14
23. Palash Sarkar and Tetsu Iwata, editors. *Advances in Cryptology - ASIACRYPT 2014 - 20th International Conference on the Theory and Application of Cryptology and Information Security, Kaoshiung, Taiwan, R.O.C., December 7-11, 2014. Proceedings, Part I*, volume 8873 of *Lecture Notes in Computer Science*. Springer, 2014. 20
24. Claus-Peter Schnorr. Efficient identification and signatures for smart cards. In Gilles Brassard, editor, *Advances in Cryptology - CRYPTO '89, 9th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 1989, Proceedings*, volume 435 of *Lecture Notes in Computer Science*, pages 239–252. Springer, 1989. 3, 4
25. Colin Stahlke. Point compression on jacobians of hyperelliptic curves over \mathbb{F}_q . Cryptology ePrint Archive, Report 2004/030, 2004. <https://eprint.iacr.org/2004/030>. 8
26. Martijn Stam. *Speeding up subgroup cryptosystems*. PhD thesis, Technische Universiteit Eindhoven, 2003. 5
27. Erich Wenger, Thomas Unterluggauer, and Mario Werner. 8/16/32 shades of elliptic curve cryptography on embedded processors. In Goutam Paul and Serge Vaudenay, editors, *Progress in Cryptology - INDOCRYPT 2013*, volume 8250 of *Lecture Notes in Computer Science*, pages 244–261. Springer-Verlag Berlin Heidelberg, 2013. https://online.tugraz.at/tug_online/voe_main2.getvolltext?pCurrPk=72486. 18, 19
28. Johann Großschädl Zhe Liu, Erich Wenger. MoTE-ECC: Energy-scalable elliptic curve cryptography for wireless sensor networks. In *Applied Cryptography*

and Network Security, volume 8479 of *Lecture Notes in Computer Science*, pages 361–379. Springer-Verlag Berlin Heidelberg, 2014. https://online.tugraz.at/tug_online/voe_main2.getvolltext?pCurrPk=77985. 18, 19