

Fast FPGA Implementations of Diffie-Hellman on the Kummer Surface of a Genus-2 Curve

Philipp Koppermann¹, Fabrizio De Santis^{2*},
Johann Heyszl¹ and Georg Sigl^{1,3}

¹ Fraunhofer Research Institution AISEC, Munich, Germany
{philipp.koppermann,johann.heyszl,georg.sigl}@aisec.fraunhofer.de

² Siemens AG, Corporate Technology, Munich, Germany
fabrizio.desantis@siemens.com

³ Technische Universität München, EI SEC, Munich, Germany
sigl@tum.de

Abstract. We present the first hardware implementations of Diffie-Hellman key exchange based on the Kummer surface of Gaudry and Schost’s genus-2 curve targeting a 128-bit security level. We describe a single-core architecture for low-latency applications and a multi-core architecture for high-throughput applications. Synthesized on a Xilinx Zynq-7020 FPGA, our architectures perform a key exchange with lower latency and higher throughput than any other reported implementation using prime-field elliptic curves at the same security level. Our single-core architecture performs a scalar multiplication with a latency of 82 microseconds while our multi-core architecture achieves a throughput of 91,226 scalar multiplications per second. When compared to similar implementations of Microsoft’s FourQ on the same FPGA, this translates to an improvement of 48% in latency and 40% in throughput for the single-core and multi-core architecture, respectively. Both our designs exhibit constant-time execution to thwart timing attacks, use the Montgomery ladder for improved resistance against SPA, and support a countermeasure against fault attacks.

Keywords: Diffie-Hellman key exchange · hyperelliptic curve cryptography · Kummer surface · FPGA · Zynq · low-latency · high-throughput · fault countermeasure.

1 Introduction

In 1989, Koblitz [Kob89] first mentioned the application of hyperelliptic curves in cryptography. The Jacobian variety of a genus-2 curve possesses a group structure that can be used to realize cryptographic algorithms such as Diffie-Hellman (DH) key exchange and digital signatures. Unfortunately, group operations on the Jacobian have higher complexity than those on elliptic curves (genus-1 curves). However, using the *pseudo-multiplication* on the Kummer surface of the Jacobian in place of the addition on the Jacobian itself leads to a decrease of the number of field operations per group operation [Gau07]. The Kummer surface is a 2-to-1 point mapping and can be compared to the x -coordinate-only representation of elliptic curves. Table 1 shows the number of field operations for a point addition and a point doubling operation used in DH key exchange for a genus-1 Montgomery curve and a Kummer surface associated to a genus-2 curve. It can be noted that the genus-2 curve requires 1.4-times more multiplications, 3-times more squarings, and

*The work was conducted when the author was with Technische Universität München, EI SEC, Munich, Germany.

Table 1: Comparison of required field operations for point addition and point doubling: multiplication (**M**), squaring (**S**), constant multiplication \mathbf{M}_c , addition (**A**), and subtraction (**Z**).

Genus	Reference	Field size	M	S	\mathbf{M}_c	A	Z
1	Curve25519 [DHH ⁺ 15]	255-bit	5	4	1	4	4
2	Kummer [RSSB16]	127-bit	7	12	12	16	16

4-times more additions and subtractions than the genus-1 curve. However, the Kummer surface based pseudo-multiplication operates on finite fields of half the size than those of elliptic curves while supporting the same security level. This reduced field size can lead to performance benefits in hardware, which is crucial for time critical applications.

In 2006, Bernstein and Lange [BL06] showed in a cost analysis for software that a genus-2 based implementation is potentially 1.5-times faster than a comparable elliptic curve based implementation. At that time, however, a secure Kummer surface of a genus-2 curve was not found yet. Since genus-2 point counting is computationally expensive, it took further six years until Gaudry and Schost [GS12] presented a twist-secure Kummer surface targeting a 128-bit security level. Using this Kummer surface, Bos et al. [BCHL13] were the first to publish a high-speed DH implementation on high-end CPUs proving the earlier cost analysis in [BL06]. A year later, Bernstein et al. [BCLS14] presented another DH implementation using the same Kummer surface and achieved new speed records by heavily vectorizing each ladder step. These speed records were only surpassed by the Four \mathbb{Q} elliptic curve implementation of Costello and Longa [CL15], which exploits a four-dimensional Gallant-Lambert-Vanstone decomposition to minimize the total number of group operations. Renes et al. [RSSB16] published implementation results of a DH key exchange on the Kummer surface of Gaudry and Schost’s genus-2 curve [GS12] for different microcontrollers reporting notable performance results. Their results were recently surpassed by Four \mathbb{Q} implementations for similar architectures by Liu et al. [LLP⁺17]. So far, investigations of the DH key exchange on the Kummer surface of genus-2 curves were confined to software implementations [BCHL13, BCLS14, RSSB16]. While these software implementations already showed the performance advantages of genus-2 curves, the design of efficient hardware is a fundamentally different task. Best performance results are only obtained when each module is carefully optimized with optimally matched timing characteristics to one another. In this work, we show that the Kummer surface of Gaudry and Schost’s genus-2 curve can be used to perform very fast DH key exchanges in hardware.

Contribution. We present the first FPGA implementations of Diffie-Hellman using the Kummer surface of a genus-2 hyperelliptic curve and show its competitiveness compared to elliptic curve based implementations. Following previous high-speed genus-2 implementations in software [BCHL13, BCLS14, RSSB16], we use the Kummer surface [CC86, CCS16, Gau07] of Gaudry and Schost’s genus-2 curve [GS12]. Synthesized on a Zynq-7020, our single-core architecture is about 1.91-times faster than the Four \mathbb{Q} implementation [JMAL16], which has been the fastest prime-field curve scalar multiplication on the same FPGA so far. Our single-core architecture inherently supports a fault countermeasure, but utilizes 1.51-times more slices than the Four \mathbb{Q} implementation. In terms of throughput, our multi-core design shows a factor-1.41 improvement compared to the Four \mathbb{Q} implementation and a factor-2.82 improvement compared to the high-throughput Curve25519 implementation [SG15]. The main design decisions that influenced our results are:

Interleaving two scalar multiplications. Due to the serial nature of the considered ladder, multiple hardware modules operate below full capacity. This allows for a second scalar multiplication to be efficiently interleaved by carefully scheduling the required field operations. The obtained instruction schedule leaves the number of cycles unaltered while effectively doubling the throughput. Note that this interleaved scalar multiplication can also be used as a countermeasure against fault attacks by performing both scalar multiplications on the same input point and check the results for equivalence.

Efficient representation of constant values. For improved performance, we instantiate a dedicated circuit for multiplying field elements with 12-bit constants in each ladder step. Compared to a conventional modular multiplication, the constant modular multiplier requires only 4 clock cycles instead of 7. Some constants, however, are negative; the naive approach would be to convert them to positive elements of the prime field and then use the modular multiplier for multiplication. In order to avoid the increased memory requirements and decreased performance of this naive approach, we neglect the sign when storing the constants and include the conditional negation logic inside the constant modular multiplier.

High-speed modular multiplier. The performance of the scalar multiplication is strongly correlated with the performance of the modular multiplier. We take the multiplier presented in [KSHS17, Fra17], which is explicitly optimized for Mersenne prime fields, and modify it by applying the non-standard tiling technique [RMIT14] to further improve its performance. In this way, we also reduce the number of required DSP blocks by 10%.

Organization. In Sect. 2, we describe the basics of DH key exchange using a genus-2 curve, describe Gaudry and Schost’s hyperelliptic curve and its Kummer surface, and summarize the scalar multiplication on this Kummer surface using the Montgomery ladder. In Sect. 3, a description of the single-core and multi-core hardware architectures is provided. In Sect. 4 we present the performance analysis and compare our results with related work. Finally, we conclude and discuss the results in Sect. 5.

2 Diffie-Hellman key exchange using Kummer surfaces

For elliptic and hyperelliptic curve based cryptosystems, the main operation is the scalar multiplication $Q = sP$, where Q, P are two points on a curve and $s \in \mathbb{Z}$ is a scalar value. In case of elliptic curve cryptography, the two points are located on an elliptic curve defined over a finite field. An abelian group is formed by all points on the elliptic curve together with the point at infinity under the addition law, which is obtained by the chord-and-tangent rule. A point can be multiplied with a scalar by using an algorithm such as the Montgomery ladder [Mon87], which repetitively performs point doublings and differential additions.

For a genus-2 hyperelliptic curve \mathcal{C} , a group structure can be formed with the corresponding Jacobian $\mathcal{J}_{\mathcal{C}}(\mathbb{F}_q)$ where \mathbb{F}_q is a finite field. The Jacobian forms an abelian group and is defined as the quotient space $\mathcal{J}_{\mathcal{C}} := \text{Div}_{\mathcal{C}}^0 / \text{Princ}_{\mathcal{C}}$, where $\text{Div}_{\mathcal{C}}^0$ denotes the group of divisors of degree zero of curve \mathcal{C} and $\text{Princ}_{\mathcal{C}}$ the group of principal divisors of curve \mathcal{C} . A detailed description of Jacobians of hyperelliptic curves can be found in [CFA⁺05]. Although a point can be multiplied on the Jacobian, it is computationally expensive making it non-competitive compared to elliptic curve based implementations. Instead, we make use of the Kummer surface $\mathcal{K}_{\mathcal{C}}$ that is associated with $\mathcal{J}_{\mathcal{C}}$ of the hyperelliptic curve \mathcal{C} . The Kummer surface is defined as the quotient space of the Jacobian by its involution, which we denote by $\mathcal{K}_{\mathcal{C}} := \mathcal{J}_{\mathcal{C}} / \langle \pm 1 \rangle$. Gaudry [Gau07] showed that scalar multiplication on the Kummer surface can be computed faster than on the corresponding Jacobian. Even though the group structure is lost when points on $\mathcal{J}_{\mathcal{C}}$ are mapped to $\mathcal{K}_{\mathcal{C}}$, a pseudo-multiplication

Algorithm 1 `scalar_mult`: unwrap input point to Montgomery ladder on \mathcal{K}_C followed by point wrapping. It is assumed that the public key (respectively public generator) is in 381-bit wrapped representation.

Input: $(s = \sum_{i=0}^{250} s_i 2^i) \in [0, 2^{251}), \underline{\pm P}$ for $\pm P$ in \mathcal{K}_C .

Output: $\underline{\pm Q}$ for $\pm Q \leftarrow \pm[s]P$ in \mathcal{K}_C .

- 1: $\pm P \leftarrow \text{unwrap}(\underline{\pm P})$ ▷ compute 4-tuple representation of $\pm P$
 - 2: $\pm Q \leftarrow \text{mont_ladder}(s, \pm P, \underline{\pm P})$
 - 3: $\underline{\pm Q} \leftarrow \text{wrap}(\pm Q)$ ▷ compute wrapped 381-bit representation of $\pm Q$
 - 4: **return** $\underline{\pm Q}$
-

[CCS16] can be defined. For DH key exchange pseudo-multiplication is sufficient, and thus we perform all computations on \mathcal{K}_C .

Our implemented DH key exchange works the same as the one described by Renes et al. [RSSB16]; we also follow their notation throughout this paper. If a point P is on the Jacobian \mathcal{J}_C , we denote its image on the Kummer surface by $\pm P$. Each point $\pm P$ is represented by a 4-tuple where each element is 127-bit wide which sums up to 508 bit in total. As described in [CCS16, RSSB16], we assume that the public key (respectively public generator) is represented by a 3-tuple in its wrapped 381-bit representation denoted by $\underline{\pm P}$. Renes et al. [RSSB16] showed that keeping the input points in their wrapped representation offers two advantages: first, it reduces the required amount of data that needs to be transmitted and second, it results in a speed-up for the ladder computation.

For an ephemeral key exchange, the scalar multiplication is performed twice: once for computing an entity's public key, where the public generator is the input point, and once for computing a shared secret, where the other entity's public-key is the input point.

Key exchange. *Let $\underline{\pm P}$ be the public generator (respectively public key) in its wrapped representation and s be the 251-bit secret key. We then compute $Q \leftarrow \pm[s]P$ and derive the generated public key (respectively the shared secret) as $\underline{\pm Q}$.*

The scalar multiplication is implemented by Algorithm 1 (`scalar_mult`) and uses three functions: `unwrap` computes the 4-tuple representation of the input point, `mont_ladder` multiplies the unwrapped input point by a scalar value using the Montgomery ladder [Mon87], and `wrap` finally computes the 381-bit wrapped representation of the output point; all these functions are described in detail in Sect. 2.3. As these functions depend on various parameters of the Kummer surface of Gaudry and Schost's genus-2 hyperelliptic curve [GS12], we first summarize the definition of this curve in Sect. 2.1 and then describe the associated Kummer surface in Sect. 2.2. More details can be found in [BL06, RSSB16].

2.1 Gaudry and Schost's genus-2 hyperelliptic curve

The genus-2 hyperelliptic curve \mathcal{C} of Gaudry and Schost [GS12] is defined over the prime field \mathbb{F}_p with $p = 2^{127} - 1$. The curve \mathcal{C} can be mapped to a curve of equation

$$\mathcal{C}_{[\lambda, \mu, \nu]} : Y^2 := X(X-1)(X-\lambda)(X-\mu)(X-\nu) ,$$

where the so-called Rosenhain invariants [Gau07] are

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

The squared theta constants are

$$a = -11, \quad b = 22, \quad c = 19, \quad \text{and} \quad d = 3, \\ e/f = (1 + \sqrt{CD/AB}) / (1 - \sqrt{CD/AB}),$$

where A, B, C and D are the dual theta constants

$$A := a + b + c + d = 33, \quad B := a + b - c - d = -11, \\ C := a - b + c - d = -17, \quad D := a - b - c + d = -49.$$

2.2 Kummer surface

Similar to previous works, we use the *fast* Kummer surface $\mathcal{K}_{\mathcal{C}} \in \mathbb{P}^3$ of [CC86, CCS16, Gau07], which is defined as:

$$\mathcal{K}_{\mathcal{C}} : Exyzt = ((x^2 + y^2 + z^2 + t^2) - F(xt + yz) - G(xz + yt) - H(xy + zt))^2,$$

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 $E = 4abcd(ABCD / ((ad - bc)(ac - bd)(ab - cd)))^2$. For a point P in $\mathcal{J}_{\mathcal{C}}$, its image $\mathcal{K}_{\mathcal{C}}$ is denoted by

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

The identity point $\langle 1, 0 \rangle$ of $\mathcal{J}_{\mathcal{C}}$ maps to

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

2.3 Scalar multiplication on the Kummer surface

As described in Algorithm 1 (`scalar_mult`), we assume that the input and output points are in their wrapped representation. The wrapped representation of the point $\pm P = (x : y : z : t)$ in $\mathcal{K}_{\mathcal{C}}$ is a 3-tuple which is denoted by $\underline{\pm P} = (x/y, x/z, x/t)$. Algorithm 2

Algorithm 2 `unwrap`: $(x/y, x/z, x/t) \mapsto (x : y : z : t)$ `unwrap` point to its 508-bit representation.

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

Output: $(x : y : z : t) \in \mathbb{P}^3$.

- 1: $(V_1, V_2, V_3) \leftarrow ((x/z)(x/t), (x/y)(x/t), (x/y)(x/z))$
 - 2: $V_4 \leftarrow V_3(x/t)$
 - 3: **return** $(V_4 : V_1 : V_2 : V_3)$
-

(`unwrap`) implements the point unwrapping, which consists of 4 multiplications in \mathbb{F}_p . The wrapping function is described in Algorithm 3 (`wrap`); it consists of a finite field inversion and 7 multiplications. As in [RSSB16], we define three operations in the projective space \mathbb{P}^3 to improve the readability of the Montgomery ladder. First, the multiplication \mathcal{M} that multiplies the corresponding pairs of coordinates from two distinct points in \mathbb{F}_p :

$$\mathcal{M} : ((x_1 : y_1 : z_1 : t_1), (x_2 : y_2 : z_2 : t_2)) \mapsto (x_1x_2 : y_1y_2 : z_1z_2 : t_1t_2).$$

Algorithm 3 wrap: $(x : y : z : t) \mapsto (x/y, x/z, x/t)$ compute wrapped 381-bit representation.

Input: $(x : y : z : t) \in \mathbb{P}^3$.

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

- 1: $V_1 \leftarrow yz$
 - 2: $V_2 \leftarrow x / (V_1 t)$ ▷ inversion
 - 3: $V_3 \leftarrow V_2 t$
 - 4: **return** $(V_3 z, V_3 y, V_1 V_2)$
-

Algorithm 4 mont_ladder: Montgomery ladder using combined differential double-and-add.

Input: $(s = \sum_{i=0}^{250} s_i 2^i) \in [0, 2^{251}), (\pm P, \pm P) \in \mathcal{K}_C^2$.

Output: $\pm Q = (x_Q : y_Q : z_Q : t_Q) \in \mathbb{P}^3$ for $\pm Q \leftarrow \pm[s]P$ in \mathcal{K}_C .

- 1: $V_5 \leftarrow (a : b : c : d)$
 - 2: $V_6 \leftarrow (x_P : y_P : z_P : t_P)$ ▷ representation of $\pm P$
 - 3: $V_7 \leftarrow (\frac{1}{A} : \frac{1}{B} : \frac{1}{C} : \frac{1}{D})$
 - 4: $V_8 \leftarrow (\frac{1}{a} : \frac{1}{b} : \frac{1}{c} : \frac{1}{d})$
 - 5: $V_9 \leftarrow (1 : \frac{x_P}{y_P} : \frac{x_P}{z_P} : \frac{x_P}{t_P})$ ▷ representation of $\pm P$
 - 6: **for** $i = 250$ down to 0 **do**
 - 7: $(V_1, V_2) \leftarrow \text{cswap}(s_i \oplus s_{i+1}, (V_5, V_6))$ ▷ $s_{251} = 0$
 - 8: $(V_1, V_2) \leftarrow (\mathcal{H}(V_1), \mathcal{H}(V_2))$
 - 9: $(V_3, V_4) \leftarrow (\mathcal{S}(V_1), \mathcal{M}(V_1, V_2))$
 - 10: $(V_5, V_6) \leftarrow (\mathcal{M}(V_3, V_7), \mathcal{M}(V_4, V_7))$
 - 11: $(V_1, V_2) \leftarrow (\mathcal{H}(V_5), \mathcal{H}(V_6))$
 - 12: $(V_3, V_4) \leftarrow (\mathcal{S}(V_1), \mathcal{S}(V_2))$
 - 13: $(V_5, V_6) \leftarrow (\mathcal{M}(V_3, V_8), \mathcal{M}(V_4, V_9))$
 - 14: **end for**
 - 15: $(V_1, V_2) \leftarrow \text{cswap}(s_0, (V_5, V_6))$
 - 16: **return** $\pm Q = V_2$
-

Second, the special case where the two points are equal, i.e. squaring in \mathbb{F}_p the corresponding pairs of coordinates:

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

Third, the Hadamard transform $\mathcal{H} : (x : y : z : t) \mapsto (x_{\mathcal{H}} : y_{\mathcal{H}} : z_{\mathcal{H}} : t_{\mathcal{H}})$ with

$$x_{\mathcal{H}} = \overbrace{(x+y)}^u + \overbrace{(z+t)}^v, \quad z_{\mathcal{H}} = \overbrace{(x-y)}^r + \overbrace{(z-t)}^s, \quad (1)$$

$$y_{\mathcal{H}} = (x+y) - (z+t), \quad t_{\mathcal{H}} = (x-y) - (z-t). \quad (2)$$

Finally, Algorithm 4 (`mont_ladder`) describes the Montgomery ladder for the scalar multiplication on the Kummer surface of Gaudry and Schost's genus-2 curve. The constants that are stored in V_7 and V_8 are projectively derived from the squared theta constants (a, b, c, d) and the dual theta constants (A, B, C, D) respectively (see Sect. 2.1):

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

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

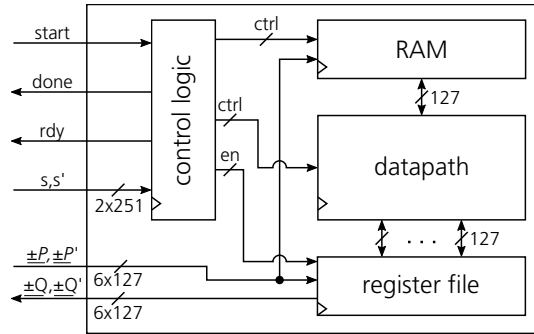


Figure 1: Single-core architecture, which contains all control and datapath logic for computing Algorithm 1 (`scalar_mult`).

The Montgomery ladder consists of 251 ladder steps, each one performing a differential-addition and a differential-doubling operation. Each ladder step includes a conditional swap of two pairs of coordinates.

3 Hardware architectures

The implementation of Algorithm 1 (`scalar_mult`) is the essential task of our hardware design. We present a single-core architecture for low-latency applications and a multi-core architecture for high-throughput applications. Our single-core architecture performs two scalar multiplications on the Kummer surface at a time by scheduling the field operations for point addition and point doubling such that it is possible to interleave a second scalar multiplication with no cycle penalty. The top-view architecture is illustrated in Fig. 1. It takes two points in their wrapped representation as input, processes them, and returns two points in their wrapped representation as output. We logically divide our single-core design into three parts that are described in the next subsections: memory, datapath, and control logic. Further we describe a multi-core architecture that instantiates 4 independently operating cores and can perform up to 8 scalar multiplications with different keys and input points.

A note on fault attacks. The two interleaved scalar multiplications can inherently be used as a redundancy countermeasure to thwart fault attacks in our designs, i.e. by performing two interleaved scalar multiplications on the same input data and subsequently compare the result bit-wise for equivalence. This can prevent a large number of fault attacks such as the powerful Biehl-Meyer-Müller DFA [BMM00]. The countermeasure can be applied to both our single- and multi-core architectures without applying any changes to the presented hardware designs.

3.1 Memory

The memory consists of a 16×127 -bit register file and a 6×127 -bit simple dual-port RAM. The register file is divided in four larger blocks, where each block is 4×127 -bit wide. We follow the logical structure of Algorithm 4 (`mont_ladder`) in which operations are performed on two points at a time (e.g. V_1, V_2 on line 8). We also use a simple dual-port RAM for storing the wrapped input point $\frac{x_p}{y_p}, \frac{x_p}{z_p}$, and $\frac{x_p}{t_p}$, which is accessed in read-only mode. Note that when no design constraints are set, the used synthesis tool instantiates distributed

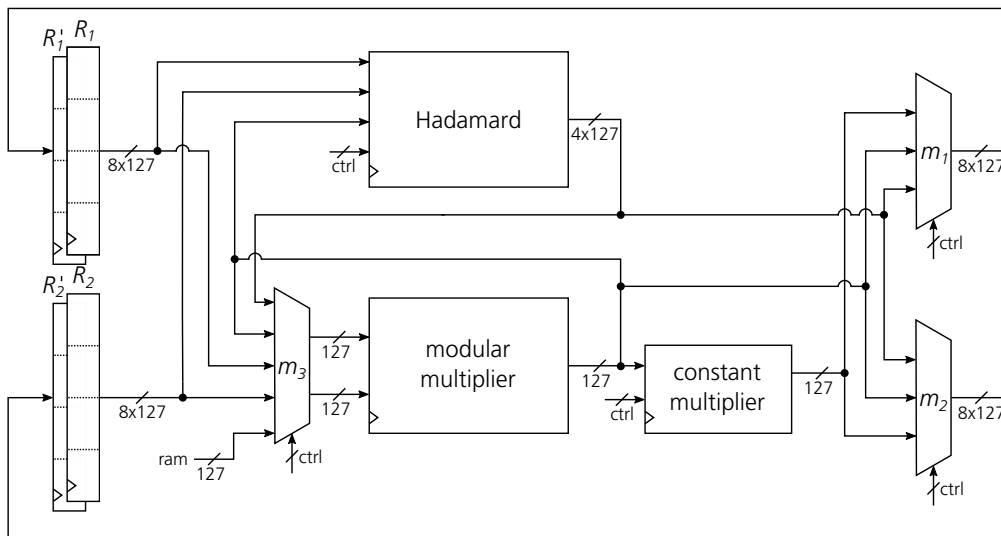


Figure 2: Datapath including register file.

RAM instead of block RAM for storing this point. We found out that forcing the synthesis tool to use block RAM resulted in a 10% decrease of the maximum clock frequency.

3.2 Datapath

The datapath including the register file is shown in Fig. 2. It implements the required field operations in \mathbb{F}_p . The register blocks R_i and R'_i for $i \in [1, 2]$ are required for storing intermediate values of the first and the second scalar multiplication, respectively. The register blocks R_1 and R'_1 are initialized with the constants $V_5 = (a : b : c : d)$ whenever Algorithm 1 (scalar_mult) is started. The modular multiplier is preceded by the multiplexer m_3 that allows to perform field operations using various input sources. The output of the constant modular multiplier and the Hadamard module serve as fast forward input paths for the modular multiplier. These fast forward paths are required when data needs to be processed immediately without any further delay. Moreover, the modular multiplier can process 127-bit inputs that originate from the RAM and are required in each ladder step (e.g. multiplication by $\frac{x_p}{y_p}$). We can store each field operation output in the register blocks, i.e. R_i and R'_i , by accordingly selecting the signals with the multiplexers m_1 and m_2 . Although large multiplexers result in an increased area utilization, they allow greater flexibility in scheduling instructions which leads to higher overall performance. All select and enable signals in Fig. 2 are driven by the control logic (see Sect. 3.3).

Modular multiplier. We have implemented a modular multiplier that computes and accumulates its digit-products in full parallel. Combined with carefully placed pipeline stages, this parallel approach enables us to continuously fetch new input operands and return the result after 7 cycles including the reduction step. This property is not only beneficial for the performance, but also required in order to interleave a second scalar multiplication. Our implemented modular multiplier is used for both squaring and multiplication in \mathbb{F}_p . Fig. 3 shows the hardware architecture of our modular multiplier. After all digit-products have been computed by the DSP blocks, they need to be summed up by an adder-tree.

This adder-tree is commonly implemented in slower standard LUT logic and poses the bottleneck in most multiplier designs. A major problem is constituted by the large adder sizes that scale up to twice the operand width, i.e. 254-bit. Deriving a high-speed

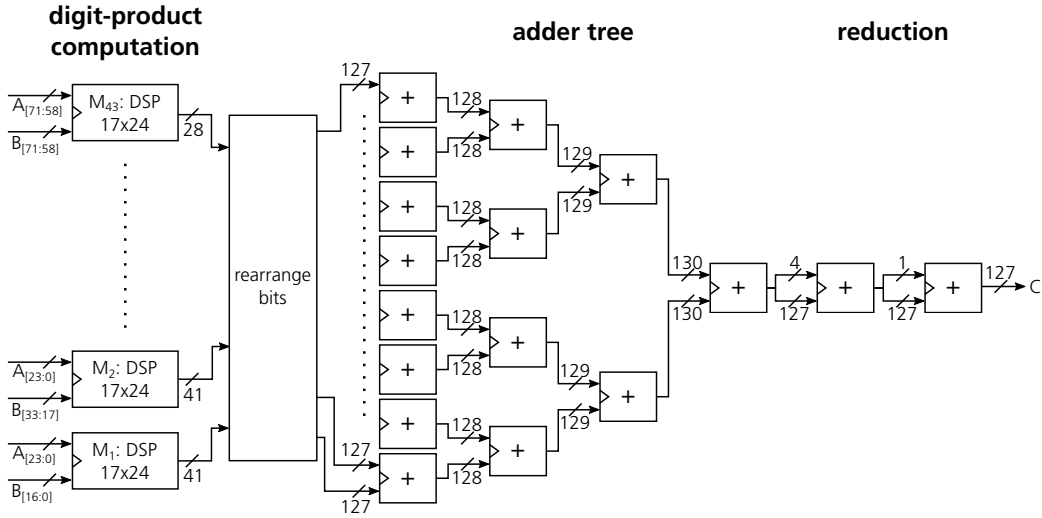


Figure 3: Architecture of the 127×127 -bit multiplier modulo $2^{127} - 1$.

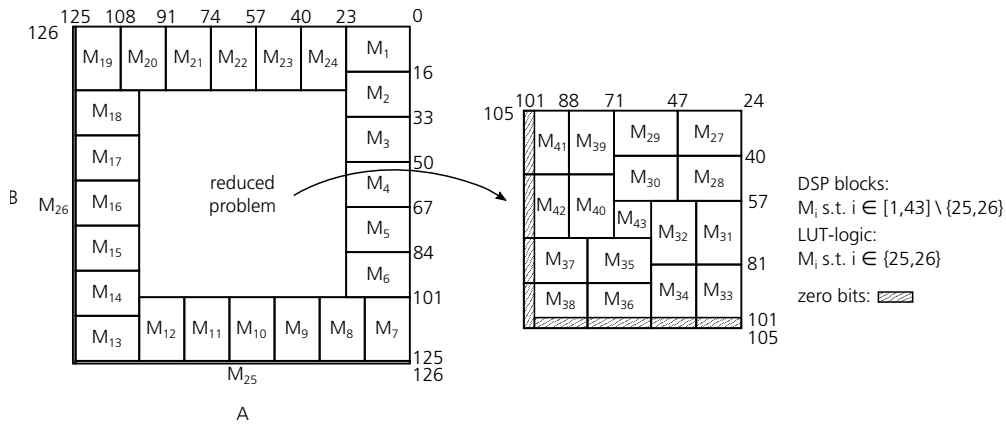


Figure 4: Left: Non-standard tiling [RMIT14] for a 127×127 -bit multiplier. Right: Non-standard tiling for smaller a 78×78 -bit multiplier.

design is also complicated by the varying sizes of the adders in the adder tree, which can lead to inefficient pipelining. To overcome this problem, Koppermann et al. [KSHS17] presented a technique for high-speed multiplication in Mersenne prime fields that reduces and equalizes the adder sizes. The main idea is to rearrange the digit-products on the bit-level while combining the multiplication with the fast reduction method proposed by Crandall in [Cra92].

In modern FPGAs, DSP blocks typically contain asymmetric multipliers, e.g. in case of the Zynq-7020 FPGA a 17×24 -bit multiplier is contained in each DSP block. In order to exploit these asymmetries to reduce the amount of DSP blocks used to perform large multiplications, different optimization strategies were proposed [dDP09, GACL12, SC10]. In particular, the authors of [dDP09] showed that operand decomposition boils down to a tiling problem, where each tile represents the result of a smaller digit-product computation. Roy et al. [RMIT14] proposed the non-standard tiling algorithm as a solution to this tiling problem. They presented a formal procedure to compute a non-standard tiling for large multipliers with arbitrary operand sizes. For a 127×127 -bit multiplier, Fig. 4 presents the

implemented non-standard tiling [RMIT14]. The left side of Fig. 4 illustrates the initial tiling for the 127×127-bit multiplier. With this initial tiling, the problem of finding an efficient placement for a 127×127-bit multiplier is reduced to a 78×78-bit multiplier. Again, we perform non-standard tiling for the reduced problem which results in a smaller 14×14-bit multiplier M_{43} . The size of the tiles M_i where $i \in [1, 43] \setminus \{25, 26\}$ corresponds to the asymmetric multiplier widths and can consequently be implemented in a single DSP block. The two tiles M_{25} and M_{26} , however, correspond to a 126×1-bit multiplier and a 1×127-bit multiplier, respectively, both implemented in LUT logic. The horizontal side represents operand A and the vertical side represents operand B . Comparing non-standard-tiling with standard-tiling, only 41 DSP blocks are required instead of 64 [SC10].

Constant modular multiplier. In order to speed up the Montgomery ladder, we instantiate a constant modular multiplier that multiplies one of the constants in $\{\frac{1}{a}, \frac{1}{b}, \frac{1}{c}, \frac{1}{d}, \frac{1}{A}, \frac{1}{B}, \frac{1}{C}, \frac{1}{D}\}$ with a variable 127-bit operand. The constant modular multiplier returns with a latency of 4 cycles, which is 3 cycles less than the generic modular multiplier, and is implemented using 6 DSP blocks only. The multiplication itself is pipelined and followed by two reduction steps including a conditional negation. The conditional negation is required for the multiplication with projectively negative constants, i.e. $\frac{1}{b}, \frac{1}{c}, \frac{1}{d}$, and $\frac{1}{A}$. For all other constants, i.e. $\frac{1}{a}, \frac{1}{B}, \frac{1}{C}$, and $\frac{1}{D}$, the negation output is ignored. All constants are hard-decoded and then selected for multiplication via a select signal. Overall, 12 modular multiplications in each ladder step can be replaced by constant multiplications.

Hadamard transform. A core computation step in Algorithm 4 (`mont_ladder`) is the Hadamard transform. It is essentially composed of 4 modular additions and 4 modular subtractions, which we implemented using 2 modular adders and 2 modular subtractors. In order to parallelize the execution of independent operations, a modular adder is implemented using two addition circuits that are connected in series, each one having a clocked register output. The first adds two 127-bit wide operands and the second reduces the sum again by using Crandall’s fast reduction [Cra92]. Because a register is placed after each addition circuit, a result is obtained each cycle after an initial delay of 2 cycles. The modular subtraction circuit is implemented similarly; modular addition and modular subtraction are both implemented in LUT logic.

Two successive Hadamard transforms, i.e. $\mathcal{H}(V_1), \mathcal{H}(V_2)$, are computed at the beginning of each ladder step before any other computation can take place. Therefore, the modular adder and the modular subtractor circuits are connected with a multiplexer in a way that two Hadamard transforms are finished in successive clock cycles. Table 2 shows the scheduling for a Hadamard transform of two points, i.e. $V_1 = (x_1 : y_1 : z_1 : t_1)$ and $V_2 = (x_2 : y_2 : z_2 : t_2)$, plotted over cycles to compute Equation (1) and Equation (2) (see Sect. 2.3). The cycles plotted under the corresponding component (e.g. modular adder \mathbf{A}_1) represent the processing stage. To give an example, u_1 in cycle 1 means that $u_1 = x_1 + y_1$ is in the first processing stage in the modular adder. In cycle 3, the computation of u_1 is finished and can be further processed by other modules. The transformed points $\mathcal{H}(V_1)$ and $\mathcal{H}(V_2)$ are returned in the 5th cycle and in the 6th cycle, respectively.

A note on lazy reduction. To reduce the number of modular reductions and hence the number of required cycles, lazy reduction is a popular technique. In software, lazy reduction comes typically for free because field elements are often smaller than a multiple of the word size which results in unused bits at higher positions. In hardware, however, lazy reduction leads to increased memory requirements, larger multipliers, and a more complex control logic to distinguish between reduced and unreduced field elements when initiating a modular multiplication. As a consequence, we neglected the lazy reduction technique.

Table 2: Instruction scheduling for two successive Hadarmard computations as in line 8 of Algorithm 4 (`mont_ladder`) using modular addition (**A**) and subtraction (**Z**).

Cycle	A ₁		A ₂		Z ₂		Z ₂	
	1	3	1	3	1	3	1	3
1	u_1	-	v_1	-	r_1	-	s_1	-
2	u_2	-	v_2	-	r_2	-	s_2	-
3	$x_{\mathcal{H}_1}$	u_1	$z_{\mathcal{H}_1}$	v_1	$y_{\mathcal{H}_1}$	r_1	$t_{\mathcal{H}_1}$	s_1
4	$x_{\mathcal{H}_2}$	u_2	$z_{\mathcal{H}_2}$	v_2	$y_{\mathcal{H}_2}$	r_2	$t_{\mathcal{H}_2}$	s_2
5	-	$\mathbf{x}_{\mathcal{H}_1}$	-	$\mathbf{z}_{\mathcal{H}_1}$	-	$\mathbf{y}_{\mathcal{H}_1}$	-	$\mathbf{t}_{\mathcal{H}_1}$
6	-	$\mathbf{x}_{\mathcal{H}_2}$	-	$\mathbf{z}_{\mathcal{H}_2}$	-	$\mathbf{y}_{\mathcal{H}_2}$	-	$\mathbf{t}_{\mathcal{H}_2}$

Table 3: Latency and throughput of field operations.

Operation	Latency (cycles)	Throughput (op/cycles)
Addition/subtraction in \mathbb{F}_p	2	1
Multiplication/squaring in \mathbb{F}_p	7	1
Constant multiplication in \mathbb{F}_p	4	1
Inversion in \mathbb{F}_p	952	1/476
Hadamard transform	4	1/2

3.3 Control logic

The control logic takes care of performing the necessary memory operations in the register file and RAM, and schedules the instructions required by Algorithm 1 (`scalar_mult`). The unwrapping and wrapping function, and the Montgomery ladder logically divide the control logic into separate control blocks. The control logic is implemented using a Finite State Machine (FSM). Inside the FSM multiple counters are used to track the processing status of arithmetic modules such as the modular multiplier. For an efficient instruction scheduling, the latency and throughput characteristics of the underlying functions such as modular multiplication and Hadamard transform are required. Table 3 shows the performance of the field operations in \mathbb{F}_p and the Hadamard transform. The throughput denotes how often an instruction can be scheduled, e.g. a throughput of 1/2 (op/cycles) means 1 instruction can be scheduled in 2 cycles. Table 4 reports the latency of all high-level operations.

Montgomery ladder. Over 90 percent of all cycles are spent for the Montgomery ladder, and hence it is crucial to efficiently schedule field-level instructions. Table 5 shows the instruction scheduling for a Montgomery ladder step for two scalar multiplications. Instructions of the second scalar multiplication are complemented by a prime symbol, e.g. y'_1 . Montgomery ladder calls 251 Montgomery ladder steps, each implementing a combined differential double-and-add which takes 41 cycles to run. All scheduled instructions denote the expected output, e.g. in cycle 5 the squaring y_3 is an abbreviation and stands for the computation of $y_3 = V_{3,y} = V_{1,y}V_{1,y}$ as described in line 9 of Algorithm 4 (`mont_ladder`). The conditional-swap function is implemented with no timing-penalty by simply swapping the arguments of the first two Hadamard transforms. Our control logic schedules modular multiplications and multiplications by constants in parallel for best performance results. Note that the constant multiplier uses the direct output of the modular multiplier.

Table 4: Latency of high-level functions.

Operation	Latency (cycles)
Unwrap	30
Combined differential double-and-add	41
Montgomery ladder	10,302
Wrap	998
Scalar multiplication	11,330

Table 5: Instruction scheduling for single ladder step as described in Algorithm 4 (mont_ladder) for the modular multiplier (**M**), the constant modular multiplier (**M_c**), and the Hadamard transform module (**H**).

Cycle	M		H		M_c		Cycle	M		H		M_c	
	1	8	1	5	1	5		1	8	1	5	1	4
1	-	-	\mathcal{H}_1	-	-	-	28	z_3	z'_3	-	-	z'_5	y'_6
2	-	-	\mathcal{H}_2	-	-	-	29	t_3	t'_3	-	-	t'_5	z'_6
...	-	-	-	-	-	-	30	x_3	x'_3	-	-	x'_5	t'_6
5	y_3	-	-	\mathcal{H}_1	-	-	31	y_6	y_4	\mathcal{H}'_2	-	-	x'_6
6	y_4	-	-	\mathcal{H}_2	-	-	32	z_6	z_4	-	-	-	z'_5
7	z_4	-	-	-	-	-	33	t_6	t_4	-	-	-	t'_5
8	t_4	-	-	-	-	-	34	x_4	y_3	\mathcal{H}'_1	-	y_5	x'_5
9	x_4	-	-	-	-	-	35	y'_4	z_3	-	\mathcal{H}'_2	z_5	-
10	z_3	-	-	-	-	-	36	z'_4	t_3	-	-	t_5	-
11	t_3	-	-	-	-	-	37	t'_4	x_3	-	-	x_5	-
12	x_3	y_3	\mathcal{H}'_1	-	y_5	-	38	y'_3	y6	-	\mathcal{H}'_1	-	y5
13	-	y_4	\mathcal{H}'_2	-	y_6	-	39	z'_3	z6	-	-	-	z5
14	-	z_4	-	-	z_6	-	40	t'_3	t6	-	-	-	t5
15	-	t_4	-	-	t_6	-	41	x'_3	x4	-	-	-	x5
16	y'_3	x_4	-	\mathcal{H}'_1	x_6	y_5	1	y'_6	y'_4	-	-	-	-
17	y'_4	z_3	-	\mathcal{H}'_2	z_5	y_6	2	z'_6	z'_4	-	-	-	-
18	z'_4	t_3	-	-	t_5	z_6	3	t'_6	t'_4	-	-	-	-
19	t'_4	x_3	-	-	x_5	t_6	4	x'_4	y'_3	-	-	y'_5	-
20	x'_4	-	\mathcal{H}_2	-	-	x_6	5	-	z'_3	-	-	z'_5	-
21	z'_3	-	-	-	-	z_5	6	-	t'_3	-	-	t'_5	-
22	t'_3	-	-	-	-	t_5	7	-	x'_3	-	-	x'_5	-
23	x'_3	y'_3	\mathcal{H}_1	-	y'_5	x_5	8	-	y6'	-	-	-	y5'
24	y_4	y'_4	-	\mathcal{H}_2	y'_6	-	9	-	z6'	-	-	-	z5'
25	z_4	z'_4	-	-	z'_6	-	10	-	t6'	-	-	-	t5'
26	t_4	t'_4	-	-	t'_6	-	11	-	x4'	-	-	-	x5'
27	y_3	x'_4	-	\mathcal{H}_1	x'_6	y'_5	-	-	-	-	-	-	-

Table 6: Device utilization and maximum clock frequency on Xilinx Zynq-7020 FPGA.

Component	Single-core	Multi-core	Available
	@138.7 MHz	@129.2 MHz	
LUTs	8,764 (16%)	35,015 (66%)	53,200
Registers	6,852 (6%)	27,300 (26%)	106,400
DSP48E1	49 (22%)	196 (89%)	220
Block RAM	0 (0%)	0 (0%)	140
Occupied slices	2,657 (20%)	10,554 (79%)	13,300

Modular inversion. We use Fermat’s little theorem to compute the multiplicative inverse x^{-1} of an integer $x \in \mathbb{F}_p \setminus \{0\}$. The finite field inversion is given by $x^{-1} \equiv x^{2^{127}-3}$. This exponentiation is computed with a sequence of 126 modular squarings and 10 modular multiplications as described by Renes et al. [RSSB16]. We implemented the modular inversion such that two elements of the prime field are inverted simultaneously by interleaving the field multiplication and squaring operations.

3.4 Multi-core architecture

For multi-core architectures, the amount of cores which can be instantiated in parallel is strongly limited by the number of DSP blocks available on the target FPGA device. Our multi-core architecture implements 4 independently operating single-cores each featuring its own control logic. As a result, up to 8 scalar multiplications with different keys and input points can be computed. Instantiating multiple single-cores is a common concept and was similarly applied by Sasdrich and Güneysu [SG15] for Curve25519 and Järvinen et al. [JMAL16] for FourQ. Sasdrich and Güneysu used a shared inversion module and Järvinen et al. used a shared control logic component. We also implemented a multi-core architecture with a shared control logic using a single shared key to reduce the area utilization. However, the LUT logic was only reduced by approximately 10% which is a rather small improvement compared to its limitations. In fact, this shared control logic architecture requires all scalar multiplications to be started in parallel as there is only one control logic for all cores.

4 Results and analysis

We synthesized our single-core and multi-core architectures with Xilinx Vivado 2017.2 on a Xilinx Zynq-7020 FPGA (XC7Z020CLG484-3). All our results are obtained after place-and-route. Table 6 presents the area utilization including the maximum clock frequency for the single-core and multi-core architecture. Our single-core architecture requires 20% of the available slices and 22% of the available DSP blocks. Through according design methods and proper constraining we achieve a maximum clock frequency of 138.7 MHz, which corresponds to a clock period of 7.21 ns. Two interleaved scalar multiplications require 11,330 cycles, and thus a session-key can be computed with a latency of 82 μ s. The interleaving of two scalar multiplications can then be either used to effectively double the throughput to 24,482 scalar multiplications per second or provide resistance against fault attacks. To enable a fair comparison with other works, we assume that only a single input point and secret scalar is available at a time for the single-core architecture i.e. the interleaving of two scalar multiplications is used as an additional fault countermeasure. For our multi-core design we instantiate the maximum amount of 4 single-cores on the

Table 7: Comparison of single- and multi-core architectures of variable-base scalar multiplications featuring a 128-bit security level on a Zynq-7020.

Reference	Curve	Cores	Resources			Latency (μ s)	T-put (op/s)
			Slices	DSP	BRAM		
[SG15]	Curve25519	1	1,029	20	2	397	2,519
[JMAL16]	FourQ (Mont.)	1	565	16	7	310	3,222
[JMAL16]	FourQ (End.)	1	1,691	27	10	157	6,389
This work	Kummer	1	2,657	49	0	82	12,224
[SG15]	Curve25519	11	11,277	220	22	397	32,304
[JMAL16]	FourQ (End.)	11	5,697	187	110	170	64,730
This work	Kummer	4	10,554	196	0	88	91,226

Zynq-7020 FPGA. Compared to our single-core design, we see a decrease in the maximum clock frequency; using Vivado tools, we can place-and-route our design with a clock frequency of 129.2 MHz which corresponds to a clock period of 7.74 ns. For the multi-core architecture with independently operating single-cores we report a throughput of 91,226 scalar multiplications per second.

Table 7 provides a comparison of our results with state-of-the-art scalar multiplication implementations on the same Zynq-7020 FPGA device all featuring a 128-bit security level. Namely, we compare our work with the Curve25519 implementation by Sasdrich and Güneysu [SG15] and the FourQ implementation by Järvinen et al. [JMAL16]. Comparing the latency of the single-core designs, our proposed implementation is 1.91-times faster than FourQ using endomorphisms, 3.78-times faster than FourQ using the Montgomery ladder, and 4.84-times faster than Curve25519. The improvement in latency is related to the increased area utilization i.e. our design demands 1.57-times and 2.58-times more slices than FourQ using endomorphisms and Curve25519, respectively. Yet, our implementation performs better than the fastest implementation so far (FourQ with End.) in both the LUT-latency product (217,787 against 265,487) as well as the DSP-latency product (4,018 against 4,239).

Our multi-core architecture with independently operating single-cores offers a throughput that is 1.41-times higher than FourQ and 2.82-times higher than the Curve25519 implementation. In terms of latency, we also report the fastest scalar multiplication, i.e. our architecture is 1.93-times faster than FourQ and 4.51-times faster than Curve25519. Note that all reported multi-core designs use the maximum number of cores that can be successfully placed on the target device. However, only our multi-core design features fully independent single-cores, i.e. neither the inversion unit, such as the Curve25519 implementation [SG15], nor the scalar multiplication unit, such as FourQ implementation [JMAL16], are shared. Also note that we make use of distributed RAM implemented by LUT logic for memory, which leaves a notable amount of BRAM available for other applications. We emphasize that Curve25519 and FourQ could also benefit from interleaved scalar multiplication. However, this was not included in the corresponding implementations and thus no results can be compared.

5 Conclusions

We presented the first hardware implementation results for a key exchange on the Kummer surface of Gaudry and Schost’s genus-2 curve. Although a Kummer surface based key

exchange has an increased number of field operations per ladder step when compared to elliptic curves, our presented architectures perform a scalar multiplication with lower latency and higher throughput than any other reported prime-field elliptic curve key exchange featuring a 128-bit security level on a Zynq-7020 FPGA. These results set new records for latency and throughput among state-of-the-art 128-bit secure key exchange implementations known so far, such as Curve25519 [SG15] and FourQ [JMAL16].

Acknowledgements

The authors would like to thank the anonymous reviewers for their useful comments that improved the quality of the paper. The authors also acknowledge Abhijith Chikrapla Danappa and Zohaib Khan for their valuable effort in developing the initial prototypes of the presented implementations. The work presented in this contribution was supported by the German Federal Ministry of Education and Research in the project *secUnity* through grant number 16KIS0394.

References

- [BCHL13] Joppe W. Bos, Craig Costello, Hüseyin Hisil, and Kristin E. Lauter. Fast cryptography in genus 2. In *Advances in Cryptology - EUROCRYPT 2013, 32nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Athens, Greece, May 26-30, 2013. Proceedings*, pages 194–210, 2013.
- [BCLS14] Daniel J. Bernstein, Chitchanok Chuengsatiansup, Tanja Lange, and Peter Schwabe. Kummer strikes back: New DH speed records. In *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*, pages 317–337, 2014.
- [BL06] Daniel J. Bernstein and Tanja Lange. Elliptic vs. hyperelliptic, part 1, 2006. <http://cr.y.p.to/talks.html#2006.09.20>.
- [BMM00] Ingrid Biehl, Bernd Meyer, and Volker Müller. Differential fault attacks on elliptic curve cryptosystems. In *Advances in Cryptology - CRYPTO 2000, 20th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 2000, Proceedings*, pages 131–146, 2000.
- [CC86] David V. Chudnovsky and Gregory V. Chudnovsky. Sequences of numbers generated by addition in formal groups and new primality and factorization tests. *Advances in Applied Mathematics*, 7(4):385–434, December 1986.
- [CCS16] Ping Ngai Chung, Craig Costello, and Benjamin Smith. Fast, uniform scalar multiplication for genus 2 jacobians with fast kummers. In *Selected Areas in Cryptography - SAC 2016 - 23rd International Conference, St. John's, NL, Canada, August 10-12, 2016, Revised Selected Papers*, pages 465–481, 2016.
- [CFA⁺05] Henri Cohen, Gerhard Frey, Roberto Avanzi, Christophe Doche, Tanja Lange, Kim Nguyen, and Frederik Vercauteren, editors. *Handbook of Elliptic and Hyperelliptic Curve Cryptography*. Chapman and Hall/CRC, 2005.
- [CL15] Craig Costello and Patrick Longa. FourQ: Four-dimensional decompositions on a \mathbb{Q} -curve over the mersenne prime. In *Advances in Cryptology - ASIACRYPT 2015 - 21st International Conference on the Theory and Application*

- of Cryptology and Information Security, Auckland, New Zealand, November 29 - December 3, 2015, Proceedings, Part I*, pages 214–235, 2015.
- [Cra92] Richard E Crandall. Method and apparatus for public key exchange in a cryptographic system, October 27 1992. US Patent 5,159,632.
- [dDP09] Florent de Dinechin and Bogdan Pasca. Large multipliers with fewer DSP blocks. In *19th International Conference on Field Programmable Logic and Applications, FPL 2009, August 31 - September 2, 2009, Prague, Czech Republic*, pages 250–255, 2009.
- [DHH⁺15] 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. *Des. Codes Cryptography*, 77(2-3):493–514, 2015.
- [Fra17] Fraunhofer Institute for Applied and Integrated Security. GitHub repository: mod-mul-merenne. <https://github.com/Fraunhofer-AISEC>, 2017.
- [GACL12] Shuli Gao, Dhamin Al-Khalili, Nouredine Chabini, and J. M. Pierre Langlois. Asymmetric large size multipliers with optimised FPGA resource utilisation. *IET Computers & Digital Techniques*, 6(6):372–383, 2012.
- [Gau07] Pierrick Gaudry. Fast genus 2 arithmetic based on theta functions. *J. Mathematical Cryptology*, 1(3):243–265, 2007.
- [GS12] Pierrick Gaudry and Éric Schost. Genus 2 point counting over prime fields. *J. Symb. Comput.*, 47(4):368–400, 2012.
- [JMAL16] Kimmo Järvinen, Andrea Miele, Reza Azarderakhsh, and Patrick Longa. FourQ on FPGA: New hardware speed records for elliptic curve cryptography over large prime characteristic fields. In *Cryptographic Hardware and Embedded Systems - CHES 2016 - 18th International Conference, Santa Barbara, CA, USA, August 17-19, 2016, Proceedings*, pages 517–537, 2016.
- [Kob89] Neal Koblitz. Hyperelliptic cryptosystems. *J. Cryptology*, 1(3):139–150, 1989.
- [KSHS17] Philipp Koppermann, Fabrizio De Santis, Johann Heyszl, and Georg Sigl. Automatic generation of high-performance modular multipliers for arbitrary mersenne primes on fpgas. In *2017 IEEE International Symposium on Hardware Oriented Security and Trust, HOST 2017, McLean, VA, USA, May 1-5, 2017*, pages 35–40, 2017.
- [LLP⁺17] Zhe Liu, Patrick Longa, Geovandro C. C. F. Pereira, Oscar Reparaz, and Hwajeong Seo. FourQ on embedded devices with strong countermeasures against side-channel attacks. In *Cryptographic Hardware and Embedded Systems - CHES 2017 - 19th International Conference, Taipei, Taiwan, September 25-28, 2017, Proceedings*, pages 665–686, 2017.
- [Mon87] Peter L. Montgomery. Speeding the pollard and elliptic curve methods of factorization. *Mathematics of computation*, 48(177):243–264, 1987.
- [RMIT14] Debapriya Basu Roy, Debdeep Mukhopadhyay, Masami Izumi, and Junko Takahashi. Tile before multiplication: An efficient strategy to optimize DSP multiplier for accelerating prime field ECC for NIST curves. In *The 51st Annual Design Automation Conference 2014, DAC '14, San Francisco, CA, USA, June 1-5, 2014*, pages 177:1–177:6, 2014.

-
- [RSSB16] Joost Renes, Peter Schwabe, Benjamin Smith, and Lejla Batina. μ Kummer: Efficient hyperelliptic signatures and key exchange on microcontrollers. In *Cryptographic Hardware and Embedded Systems - CHES 2016 - 18th International Conference, Santa Barbara, CA, USA, August 17-19, 2016, Proceedings*, pages 301–320, 2016.
- [SC10] Shreesha Srinath and Katherine Compton. Automatic generation of high-performance multipliers for fpgas with asymmetric multiplier blocks. In *Proceedings of the ACM/SIGDA 18th International Symposium on Field Programmable Gate Arrays, FPGA 2010, Monterey, California, USA, February 21-23, 2010*, pages 51–58, 2010.
- [SG15] Pascal Sasdrich and Tim Güneysu. Implementing curve25519 for side-channel-protected elliptic curve cryptography. *TRETS*, 9(1):3:1–3:15, 2015.