

ARMv8 SIKE: Optimized Supersingular Isogeny Key Encapsulation on ARMv8 Processors

Amir Jalali¹, Reza Azarderakhsh¹, *Member, IEEE*, Mehran Mozaffari Kermani², *Senior Member, IEEE*, Matthew Campagna, and David Jao

Abstract—In this paper, we present highly-optimized constant-time software libraries for supersingular isogeny key encapsulation (SIKE) protocol on ARMv8 processors. Our optimized hand-crafted assembly libraries provide the most efficient timing results on 64-bit ARM-powered devices. Moreover, the presented libraries can be integrated into any other cryptography primitives targeting the same finite field size. We design a new mixed implementation of field arithmetic on 64-bit ARM processors by exploiting the A64 and Advanced SIMD processing units working in parallel. Using these techniques, we are able to improve the performance of the entire protocol by the factor of 5x compared to optimized C implementations on 64-bit ARM high-performance cores, providing 83-, 124-, and 159-bit quantum-security levels. Furthermore, we compare the performance of our proposed library with the previous highly-optimized ARMv8 assembly library available in the literature. The implementation results illustrate the overall 10% performance improvement in comparison with previous work, highlighting the benefit of using mixed implementation over relatively-large finite field size.

Index Terms—ARM assembly, finite field, isogeny-based cryptosystems, key encapsulation mechanism, post-quantum cryptography.

I. INTRODUCTION

IN RECENT years, extensive amount of research has been devoted to quantum computers. These machines are envisioned to be able to solve mathematical problems which are currently unsolvable for conventional computers, because of their exceptional computational power from quantum mechanics. Therefore, if quantum computers are ever built in large scale, they will certainly be able to break many or almost all of the currently in-use public-key cryptosystems, the threat of

which would be catastrophic to the confidentiality and integrity of any secure communication. To counteract this problem, post-quantum cryptography protocols are required to preserve the security in the presence of quantum adversaries. Regardless of whether we can estimate the exact time for the advent of the quantum computing era, we must begin to prepare the security protocols to be resistant against potentially-malicious power of quantum computing. Accordingly, NIST initiated a process to evaluate, and standardize one or more post-quantum public-key cryptography primitives [1]. Recently, the first round of submission of the post-quantum primitives is completed and all the proposals are publicly available¹ to evaluate in terms of the proof of security and efficiency.

The submitted public-key post-quantum cryptography (PQC) proposals are based on five different hard problems and they are categorized as code-based cryptography [2], (ring) lattice-based cryptography [3], [4], hash-based cryptography [5], multivariate cryptography [6], and isogeny-based cryptography [7]. The isogeny-based cryptography is based on the hardness of computing the isogenies between two isomorphic elliptic curves and it provides a complete key encapsulation protocol. The proposed method is denoted as Supersingular Isogeny Key Encapsulation (SIKE) [8], and constructed upon the initial Diffie-Hellman key-exchange scheme proposed by Jao and De Feo [7].

SIKE protocol provides a standard method of key-exchange between two parties, and it has been claimed to be secure against large-scale quantum adversaries running the Shor's quantum algorithm [9]. Compared to other post-quantum candidates, supersingular isogeny problem is a much younger scheme and its security and performance need to be investigated more. In terms of performance, SIKE is not a fast protocol due to the extensive number of point arithmetic which are required for computing large-degree isogenies. However, because of its significant smaller size of secret-key and public-key compared to other PQC candidates, SIKE is a suitable option for the applications where communication bandwidth is critical. Furthermore, since it is the only post-quantum cryptography protocol which is constructed on elliptic curves, hybrid cryptography protocols can be derived from SIKE and classical elliptic curve cryptography (ECC) to make the transition towards post-quantum cryptography more convenient and practical.

Manuscript received March 23, 2019; revised May 17, 2019; accepted May 29, 2019. Date of publication July 22, 2019; date of current version October 30, 2019. This work was supported in part by NSF under Grant CNS-1801341, in part by NIST under Grant 60NANB16D246, in part by NSERC, CryptoWorks21, Public Works and Government Services Canada, in part by the Canada First Research Excellence Fund, and in part by the Royal Bank of Canada. This paper was recommended by Associate Editor G. Maseru. (*Corresponding author: Amir Jalali.*)

A. Jalali and R. Azarderakhsh are with the Department of Computer and Electrical Engineering and Computer Science, Florida Atlantic University, Boca Raton, FL 33431 USA (e-mail: ajalali2016@fau.edu; razarderakhsh@fau.edu).

M. Mozaffari Kermani is with the Department of Computer Science and Engineering, University of South Florida, Tampa, FL 33620 USA (e-mail: mehran2@usf.edu).

M. Campagna is with Amazon Web Services, Inc., Seattle, WA 98108-1207 USA (e-mail: campagna@amazon.com).

D. Jao is with the Department of Mathematics, University of Waterloo, Waterloo, ON N2L 3G1, Canada (e-mail: djao@uwaterloo.ca).

Color versions of one or more of the figures in this article are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TCSI.2019.2920869

¹NIST Standardization Process (Accessed Feb. 2019): <https://csrc.nist.gov/projects/post-quantum-cryptography/round-1-submissions>

The initial idea of constructing cryptography schemes from the isogenies of regular elliptic curves was introduced by Rostovtsev and Stolbunov [10] in 2006. Later, Charles *et al.* [11] presented a set of cryptography hash functions constructed from Ramanujan graphs, i.e., the set of supersingular elliptic curves over \mathbb{F}_{p^2} with ℓ -isogenies. The main breakthrough in constructing a post-quantum cryptography protocol based on the hardness of computing isogenies was proposed by Jao and De Feo [7]. Their proposed scheme presents a set of public-key cryptography schemes such as key-exchange and encryption-decryption protocols with a coherent proof of security. Later, De Feo *et al.* [12] presented the first practical implementation of the Supersingular Isogeny Diffie-Hellman (SIDH) key-exchange protocol using optimized curve arithmetic techniques such as Montgomery arithmetic. Since the introduction of supersingular isogeny public-key protocol, many different schemes and implementations such as digital signature [13], [14], undeniable signature [15], [16], group key agreement [17], and static-static key agreement [18] have been proposed which are all built on the hardness of computing isogenies. The fast hardware architectures for computing isogenies of elliptic curves proposed by Koziel *et al.* [19] demonstrated that isogeny-based cryptography has the potential to be considered as a practical candidates on FPGAs. However, the initial performance evaluations in software were not promising compared to other post-quantum candidates. In particular to those which are constructed over learning with errors problem [20]–[22]. The SIDH projective formulas and implementation by Costello *et al.* [23] smashed the performance bar of the protocol considerably by eliminating field inversions in isogeny computations, and provided an optimized software for key-exchange protocol on Intel processors, taking advantage of hand-written assembly implementation of finite field arithmetic. However, their software still suffered from the large amount of computations due to the arithmetic of elliptic curves. On ARM-powered devices, the performance evaluations are even worse due to the reduced instruction set computing (RISC) which concentrates on power-efficiency rather than performance [24]–[27]. Recently, a new set of optimization techniques in the computation of Montgomery ladder and field arithmetic operations [28] improved the performance of the isogeny-based key-exchange protocol further; however, more investigation on the efficiency of ARM-based implementations is still required.

In this work, we investigate different approaches for implementing highly-optimized arithmetic libraries on 64-bit ARM processors. We introduce a new way of implementing multi-precision multiplication using mixed A64/ASIMD hand-written assembly along with pure A64 assembly implementation. All the previous works on this area concentrated on development of highly-optimized implementation of field arithmetic using pure SIMD instructions [24], [25], [29], or investigated a combination of SIMD and general register implementation on 32-bit ARM platforms, taking advantage of a mixture of 128-bit wide NEON vectors and 32-bit general registers [30], [31]. In this work, we show that such a combination still can be beneficial on 64-bit ARMv8 family of

processors by adopting a novel engineering technique which takes advantage of out-of-order execution pipeline on high-performance ARM cores. We compare the performance of our arithmetic libraries inside the SIKE reference implementation, and conclude the benefits of using mixed implementation over relatively-large finite fields.

A. Contributions

In this work, we study different approaches of implementing SIKE on 64-bit ARM. We engineer the finite field arithmetic implementation accurately to provide the fastest timing records of the protocol on our target platforms. Our contributions can be categorized as follows:

- We propose a new approach for implementing finite field arithmetic on 64-bit ARM processors. We combine general register and vector limbs in an efficient way to reduce the pipeline stalls and improve the overall performance. To the best of our knowledge, this work is the first implementation of such a technique on ARMv8 processors.
- We implement different optimized versions of finite field multiplication using Karatsuba multiplication method which outperforms the previous implementation of the field multiplication with the same size on ARMv8 target platform. The proposed implementations are constant-time and resistant to timing attacks.
- Our optimized software provides a constant-time implementation of the post-quantum SIKE protocol over three different quantum security levels. We state that, this work is the first implementation of SIKE_{p964} which provides 159-bit quantum security level.
- We provide a comprehensive timing results of the SIKE protocol, using different hand-written assembly techniques along with optimized C benchmarks on ARMv8 processors. We are able to improve the overall performance of the SIKE library significantly on target processors.

In Section 2 we recall the essential concepts from [7], [8] which are required in SIKE. In Section 3 we describe our implementation techniques and methodology to develop a highly-optimized field multiplier on ARMv8 processors. In Section 4 we present the efficient implementation of key components inside SHA-3 inside the SIKE protocol. We present performance results of our implementation and the comparison with previous work in Section 5. We conclude the paper in Section 6.

B. Code Availability

For reproducibility, we provide our optimized implementation publicly available. The source codes are available at: <https://github.com/amirjalali65/armv8-sike>.

II. BACKGROUND

This section provides a brief presentation of the SIKE protocol. We refer readers to [7], [8] for more detailed explanation of the supersingular isogeny problem and the base key-exchange protocol which the SIKE is constructed upon.

A. Isogenies of Elliptic Curves

Let p be a prime of the form $p = 2^{e_A}3^{e_3} - 1$, and let E be a supersingular elliptic curve defined over a field of characteristic p . E can be also defined over \mathbb{F}_{p^2} up to its isomorphism. An isogeny $\phi : E \rightarrow E'$ is a non-constant map from E to E' which translates the identity into the identity. An isogeny map is defined by its degree and kernel. The degree of an isogeny is its degree as morphism. An isogeny with degree ℓ map is called ℓ -isogeny. Let G be a subgroup of points on E which contains $\ell + 1$ cyclic subgroups of order ℓ . This subgroup is the torsion group $E[\ell]$ and each element of this group is corresponding to an isogeny of degree ℓ ; accordingly, an isogeny also can be identified by G , i.e., the kernel of isogeny, and it can be computed using Vélu's formula [32]. We denote this map as $\phi : E \rightarrow E'/\langle G \rangle$. Vélu's formula can only compute the isogeny of small degrees, while isogeny-based cryptography requires the evaluation of large-degree isogenies on curves. An efficient recursive strategy of computing large-degree isogeny is described in [7], by representing full binary trees and dynamic algorithms. The proposed strategy is adopted inside the SIKE software to compute large-degree isogeny.

Isogenies of elliptic curves divide isomorphic curves into isomorphism classes over \mathbb{F}_{p^2} . These classes of curves are categorized with their j -invariants [33]. This value is unique for isomorphic curves of the same class.

SIKE implementation is constructed on Montgomery curves, taking advantage of their fast and compact arithmetic. Let E be a Montgomery curve which defined by $E : By^2 = x^3 + Ax^2 + x$ equation. The j -invariant of E can be derived by the following equation:

$$j(E) = \frac{256(A^2 - 3)^3}{A^2 - 4}. \quad (1)$$

Therefore, in order to verify whether two elliptic curves are on the same class of isomorphisms, we can evaluate their j -invariant values. This feature is exploited to construct the supersingular isogeny key encapsulation mechanism [8], where two parties compute two isomorphic curves of the same class, and the shared secret is computed as the shared j -invariant values. Moreover, from (1), in order to compute the j -invariant of a Montgomery curve E , we only need to push the curve coefficient A into the formula. Further, we can compute the curve coefficient using the x -abscissas of two points x_P and x_Q on the curve as follows:

$$A = \frac{(1 - x_P x_Q - x_P x_R - x_Q x_R)^2}{4x_P x_Q x_R} - x_P - x_Q - x_R, \quad (2)$$

where $R = P - Q$ is also a point on E . Therefore, the j -invariant of a Montgomery curve can be evaluated using the x -abscissas of two points and their difference. The above abstraction is used inside the supersingular isogeny encryption procedure which is explained in the next section.

B. Supersingular Isogeny Key Encapsulation (SIKE) Mechanism

1) *Public Parameters*: SIKE protocol [8], like any other supersingular isogeny-based scheme, is defined over a set of

public parameters. These parameters need to be agreed by the parties prior to the key encapsulation mechanism and they are listed as follows:

- A prime p of the form $p = 2^{e_A}3^{e_3} - 1$, where e_A, e_B are two positive integers. The corresponding finite field is defined over \mathbb{F}_{p^2} . This explicit form is required for two main reasons. First, the isogeny computations using Vélu's formula need to be constructed over two different torsion subgroups, i.e., $E[2^{e_2}]$ and $E[3^{e_3}]$ of points on a starting curve for each party. Second, for efficiency reasons, primes of this form are Montgomery-friendly primes and they provide faster arithmetic [23].
- A starting supersingular Montgomery curve E_0 defined over \mathbb{F}_{p^2} .
- Two sets of generators, i.e., 3-tuple x -coordinates from $E_0[2^{e_2}]$ and $E_0[3^{e_3}]$. Note that, as it is discussed in detail in [8] and [23], the entire protocol can be defined and implemented only using two x -coordinates bases of each torsion subgroup, i.e., $\{x_{P_2}, x_{Q_2}\} \in E[2^{e_2}]$ and $\{x_{P_3}, x_{Q_3}\} \in E[3^{e_3}]$; however, for efficiency reasons, two auxiliary x -coordinates are used to encode these bases, i.e., $x_{R_2} = x_{P_2} - x_{Q_2}$ and $x_{R_3} = x_{P_3} - x_{Q_3}$.

In order to describe the SIKE protocol, first, we need to understand the supersingular isogeny public-key encryption scheme [7], because key encapsulation and decapsulation algorithms are defined based on supersingular isogeny encryption and decryption operations.

2) *Supersingular Isogeny Public-Key Encryption*: In this section, we recall the public-key encryption method from the isogenies of supersingular elliptic curves which was first introduced in [7], and it is recently described with some changes for achieving better performance in [8].

Similar to any public-key encryption schemes, supersingular isogeny PKE contains three main operations, i.e., key-generation (KeyGen), encryption (Enc), and decryption (Dec). Fig. 1 illustrates these operations from the isogeny map perspective.

Key Generation: The secret-key sk , is a random positive integer which is randomly generated from a key-space corresponding to each torsion subgroup's order. We denote them as $\mathcal{K}_2 = \{0, \dots, 2^{e_2} - 1\}$ and $\mathcal{K}_3 = \{0, \dots, 3^{e_3} - 1\}$, accordingly. Next, the secret key and the x -coordinates of generators construct the kernel of the first isogeny ϕ_ℓ such that $x_{S_\ell} = \langle x_{P_\ell + [\text{sk}_\ell]Q_\ell} \rangle$, where $\ell \in \{2, 3\}$ is the degree of isogeny. Subsequently, public-key pk is computed by evaluating the isogeny of (2^{e_2} or 3^{e_3}) iteratively from the small isogeny evaluation and point multiplication as it is discussed previously. Eventually, public-key contains a 3-tuple of x -coordinates, i.e., $\text{pk}_m = (x_{P_m}, x_{Q_m}, x_{R_m})$, where $x_{R_m} = x_{P_m} - x_{Q_m}$ and m is the degree of the isogeny. In Fig. 1 we choose sk from \mathcal{K}_3 ; therefore, the key generation algorithm maps E_0 to E_3 .

Encryption: The encryption procedure encrypts an n -bit message m from a message space $\mathcal{M} = \{0, 1\}^n$, and generates two ciphertexts c_0 and c_1 . As it will be discussed later, in the case of key encapsulation mechanism, we need to generate a random string along with the encryption process for security reason. The secret-key which is generated in the key generation procedure is randomly generated from either

Public Parameters

$$E_0/F_{p^2}$$

$$p = 2^{e_2} 3^{e_3} - 1$$

$$(x_{P_2}, x_{Q_2}, x_{R_2}) \in E_0[2^{e_2}]$$

$$(x_{P_3}, x_{Q_3}, x_{R_3}) \in E_0[3^{e_3}]$$

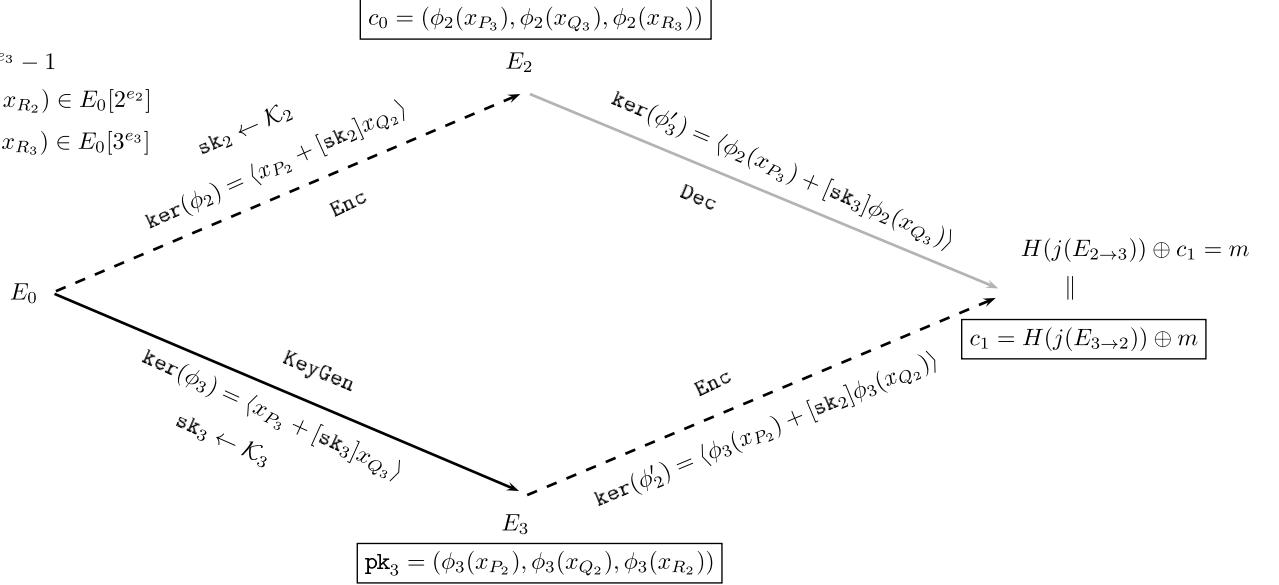


Fig. 1. Supersingular isogeny PKE protocol.

\mathcal{K}_2 or \mathcal{K}_3 . Let sk_3 be the secret-key from key generation algorithm from \mathcal{K}_3 , and pk_3 be its corresponding public key. We choose the other keyspace to generate the randomness inside the encryption algorithm, i.e., \mathcal{K}_2 . The first ciphertext c_0 is generated as a 3-tuple of x -coordinates from evaluating the large degree 2^{e_2} -isogeny, starting from E_0 using the kernel ϕ_2 which is computed as $x_{S_2} = \langle x_{P_2} + [\text{sk}_2]x_{Q_2} \rangle$.

The second ciphertext is generated after 3 steps:

- 1) First, the j -invariant of curve $E_{3 \rightarrow 2}$ is evaluated by computing the large-degree 2_3^e -isogeny, initializing from the kernel $\phi'_2 = \langle \phi_3(x_{P_2}) + [\text{sk}_2]\phi_3(x_{Q_2}) \rangle$, and a starting curve E_3 which is retrieved from pk_3 using equation (2).
- 2) Second, the hash of the retrieved j -invariant from Step 1 is computed.
- 3) Computing the ciphertext as $c_1 = H(j) \oplus m$.

The computed pair of ciphertexts is the output of encryption procedure.

Decryption: Decryption algorithm computes m from (c_0, c_1) using sk_3 . First, we compute the j -invariant of curve $E_{2 \rightarrow 3}$ which is computed by evaluating 3_3^e -isogeny of a starting curve E_2 and the kernel $\phi'_3 = \langle \phi_2(x_{P_3}) + [\text{sk}_3]\phi_2(x_{Q_3}) \rangle$. Note that E_2 is the evaluated curve from 3-tuple x -coordinate c_0 . Second, we compute hash of the computed j -invariant using the same hash function used in encryption algorithm. Finally, we retrieve the message using $m = H(j) \oplus c_1$.

3) **Key Encapsulation Mechanism:** A key encapsulation mechanism contains 3 main functions: key-generation, encapsulation, and decapsulation. In this section, we describe the SIKE protocol briefly based on the supersingular isogeny public-key encryption algorithm. We refer readers to [16] for more details.

Key Generation: Similar to PKE protocol, the key generation algorithm generates a secret-key from keyspace \mathcal{K}_3 and computes the corresponding 3-tuple x -coordinates pk_3 by evaluating 3^{e_2} -degree isogeny from starting curve E_0 .

Moreover, an n -bit secret random message $s \in \{0, 1\}^n$ is generated and concatenated to sk_3 and pk_3 to construct the SIKE secret-key sk_3 . The generated pk_3 and sk_3 are the output of this operation.

$$E_0 \rightarrow E_3 / \langle x_{P_3} + [\text{sk}_3]x_{Q_3} \rangle \rightarrow \text{sk}_3 : (s, \text{sk}_3, \text{pk}_3). \quad (3)$$

Key Encapsulation: Key encapsulation defines on top of the supersingular isogeny encryption method. The input of this operation is the public-key pk_3 which is generated in key-generation procedure. First, an n -bit random string $m \in \{0, 1\}^n$ is generated and concatenated with the public-key pk_3 . Next, the result is hashed using a custom-SHAKE256 (cSHAKE256) hash function G . This hash value is the *ephemeral secret-key* r , and it is pushed along with pk_3 into encryption function to construct the SIKE ciphertext. The hash function H inside the encryptor is also a cSHAKE256 function. The generated ciphertexts are further concatenated with m and hashed to generate the secret shared-key K :

$$\begin{aligned} \text{Enc}(\text{pk}_3, m, G(m \parallel \text{pk}_3)) &\rightarrow (c_0, c_1) \\ H(m \parallel (c_0, c_1)) &\rightarrow K. \end{aligned} \quad (4)$$

Key Decapsulation: The decapsulation algorithm computes the shared-key K from the outputs of equations (3) and (4). First, 2-tuple ciphertext is decrypted using secret-key sk_3 and hashed to retrieve m' . Further, m' is concatenated with public-key pk_3 and hashed using the G function to retrieve ephemeral secret-key r' .

$$\begin{aligned} \text{Dec}(\text{sk}_3, (c_0, c_1)) &\rightarrow m' \\ G(m' \parallel \text{pk}_3) &\rightarrow r'. \end{aligned}$$

Next, c'_0 is computed by evaluating 2_2^e -isogeny of starting curve E_0 using the kernel $\langle x_{P_2} + [r']x_{Q_2} \rangle$:

$$E_0 \rightarrow E_2 / \langle x_{P_2} + [r']x_{Q_2} \rangle \rightarrow c'_0.$$

TABLE I
SIKE FINITE FIELD PARAMETERS [8]

$p = 2^{e_2}3^{e_3} - 1$	Length (bits)	$\min(\sqrt[3]{2^{e_2}}, \sqrt[3]{3^{e_3}})$	Quantum Security
$2^{250}3^{159} - 1$	503	1.26×2^{83}	83
$2^{372}3^{239} - 1$	751	1.00×2^{124}	124
$2^{486}3^{301} - 1$	964	1.02×2^{159}	159

The final correction is performed by comparing the c_0 value with c'_0 and if they are equal, the shared-key K is computed as $K = H(m' \parallel (c_0, c_1))$, otherwise $K = H(s \parallel (c_0, c_1))$.

C. SIKE Implementation Parameters

The reference implementation of SIKE [8] contains three sets of parameters which correspond to three finite fields, denoted as \mathbb{F}_{p503} , \mathbb{F}_{p751} , and \mathbb{F}_{p964} because of their bit-length. Our optimized software is forked from the reference implementation and it offers the same security levels using the same parameters. Table I presents the information regarding each prime and its provided quantum security level. The starting elliptic curve of the reference implementation is defined as $E_0/\mathbb{F}_{p_2} : y^2 = x^3 + x$, with the cardinality equal to $\#E_0 = (2^{e_2}3^{e_3})^2$ and $j(E_0) = 1728$. This special instance of the Montgomery curve is chosen for efficiency reason, while any other supersingular curve can be used inside the protocol. We remark that the best known quantum attack on isogeny-based cryptography is based on claw-finding algorithm using quantum walks [34], which theoretically can find the isogeny between two curves in $O(\sqrt[3]{\ell^{e_i}})$, where ℓ^{e_i} is the size of the isogeny kernel; accordingly, the provided quantum security level for each prime in SIKE is determined by the minimum bit-length of each isogeny kernel, i.e., $\min(\sqrt[3]{2^{e_2}}, \sqrt[3]{3^{e_3}})$. Detailed discussion about the security of the SIKE protocol is out of scope of this work and we refer the readers to [8] for further details.

III. OPTIMIZED FIELD ARITHMETIC IMPLEMENTATION ON ARMV8

In this section, we explain our design approach and techniques. Our state-of-the-art implementation technique concentrates on the base field multiplication, since it is the most expensive operation inside the SIKE protocol. However, the proposed library takes advantage of other optimization in different layers of arithmetic.

A. Quadratic Extension Field Arithmetic

The main arithmetic operations inside the SIKE protocol are elliptic curve point arithmetic over quadratic extension field \mathbb{F}_{p_2} . Since the bit-length of the SIKE primes are smaller than the multiple of our target platform word size (64-bit), in the reference implementation, optimization techniques such as lazy reduction are exploited to improve the overall performance of \mathbb{F}_{p_2} arithmetic. As a result, finite field multiplication and reduction are implemented separately, and the reduction operation is delayed as much as possible to minimize the number of field arithmetic.

Since the quadratic extension arithmetic over \mathbb{F}_{p_2} is implemented efficiently in the SIKE reference implementation [8]

TABLE II
SIKE BASE FIELD OPERATION COUNTS OVER DIFFERENT PARAMETER SETS

Scheme	mult.	red.	add.	sub.
SIKE _{p503}	195,889	149,138	56,978	83,142
SIKE _{p751}	307,946	234,253	88,764	131,618
SIKE _{p964}	408,786	310,707	117,666	172,910

using above optimization techniques, we adapt the same strategy for computing quadratic extension field arithmetic. However, we focus on the underlying base field arithmetic implementation to improve the efficiency of the protocol further on the target platform.

B. SIKE Performance Profiling

In order to optimize the SIKE reference implementation, first we need to profile the number of operation counts for underlying field arithmetic. Since the isogeny computations for different degrees require different number of operations in corresponding with the isogeny graph, the number of field operations is different for each security level. Table II presents the number of field arithmetic operations in the SIKE protocol over different finite fields. This helps us to concentrate on the performance-bottleneck operations and try to improve them further compared to previous optimized implementations. As mentioned above, using lazy reduction technique reduces the number of field reduction to field multiplication significantly. Moreover, projective coordinates arithmetic replaces all the field inversions with extra multiplications and the performance of multiplier directly affects the overall performance of the protocol. Therefore, in this work, we focus on the efficiency of the field multiplication and engineer a highly-optimized multiplier for each quantum-security level. We find different levels of Karatsuba multiplication very effective to improve the performance compared to operand-scanning method, while we can exploit a parallel strategy to compute the multiplication of each half using mixed A64 and ASIMD assembly instructions. We elaborate more on this technique in this paper, but first we describe the most relevant architectural capabilities of the target platform which are considered in the design of our optimized library.

C. Target Architecture

We present a new design of arithmetic implementation on ARMv8 platforms by combining A64 and ASIMD instructions. To have these units working in parallel, the target platform should support *out-of-order* super-scalar pipeline. This feature is provided by ARMv8 high-performance cores such as Cortex-A57, Cortex-A72, and Cortex-A73. Therefore, we expect to get the best performance results of our *mixed* implementation on these family of processors, while our A64 implementation is anticipated to outperform the mixed version on power-efficient cores such as Cortex-A53.

1) *ARMv8 Cortex-A57 and Cortex-A72*: This family of processors are designed for the high-performance applications where they combine with power-efficient cores to achieve power-performance efficiency in the ARM big.LITTLE technology. Instructions are fetched and decoded in order

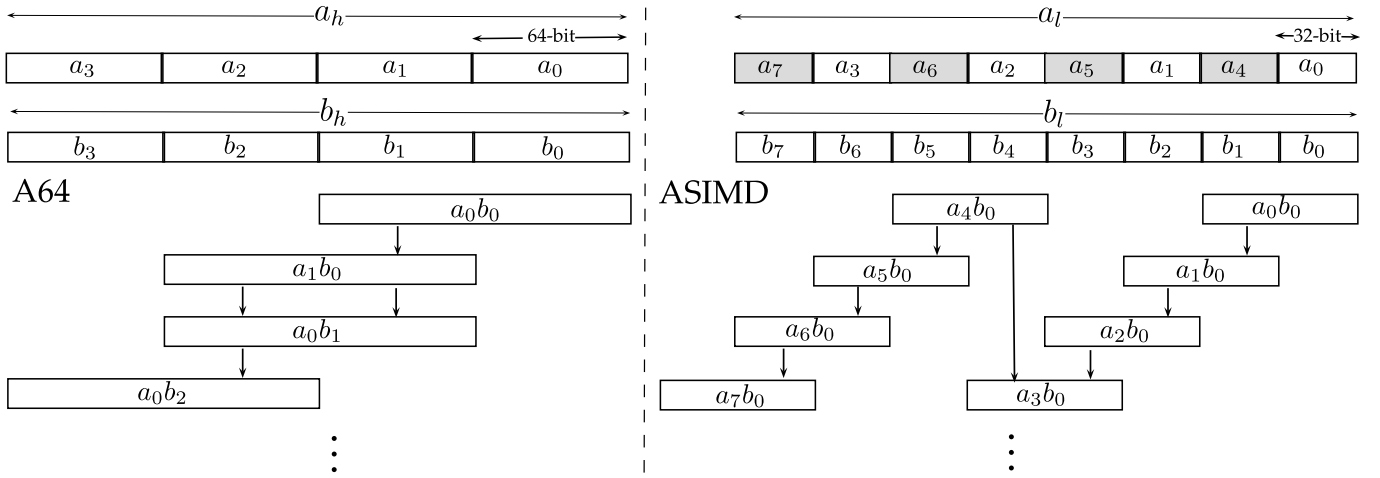


Fig. 2. 512-bit mixed-multiplication implementation overview.

into internal micro-operations (μops), and issued out-of-order to one of eight execution pipelines [35]; accordingly, two separate pipelines are dedicated for A64 and ASIMD μops . Once A64 instructions are dispatched through the integer pipeline, they are queued to be executed. Meanwhile, ASIMD pipeline can execute vector operations in the background when the A64 pipeline is stalled. This approach removes pipeline stalls to some extent and it is expected to boost the performance of arithmetic. We exploited this capability to implement optimized \mathbb{F}_p multiplication.

We implement the field multiplication using two different strategies for each level of security. First, we implemented the straightforward A64 implementation using one-level Karatsuba technique. Next, we designed a mixed version of Karatsuba multiplication to be utilized in our target architecture.

D. Karatsuba Multiplication in A64

Previous implementations of field multiplication over $\mathbb{F}_{p^{751}}$ [23], [25] and $\mathbb{F}_{p^{964}}$ [25] are based on Comba-based and two-level Karatsuba multiplication, respectively. In this work, we propose one-level Karatsuba multiplication method for both $\mathbb{F}_{p^{503}}$ and $\mathbb{F}_{p^{751}}$, while we follow two-level implementation over $\mathbb{F}_{p^{964}}$ for achieving better performance as discussed in [25]. In particular, we achieved better performance results over $\mathbb{F}_{p^{751}}$ using Karatsuba multiplication compared to previous Comba-method implementations.

Using A64 assembly instructions, one-level Karatsuba implementation is straightforward; field operands are represented in radix- 2^{64} and two n -bit field elements, e.g., a and b are divided into two $\frac{n}{2}$ -bit halves, i.e., $a_h 2^{\frac{n}{2}} + a_l$ and $b_h 2^{\frac{n}{2}} + b_l$. Consequently, one n -bit long multiplication is replaced with three $\frac{n}{2}$ -bit multiplications, at the cost of extra additions/subtractions: $c = (a_h + a_l).(b_h + b_l) - a_h b_h - a_l b_l$.

Multiplication implementation in A64 starts with loading input operands into general 64-bit registers using `ldp` instruction which can load a pair of 64-bit data in each instruction. Next, the multiplication of 64-bit registers is implemented using `mul` and `umulh` instructions, computing the low and high halves of the 128-bit result, respectively. The intermediate

additions, subtractions, and carry propagation operations are performed using `adcs` and `sbcscs`. Eventually, the result is stored back to memory iteratively in pairs of 64-bit wide using `stp` instruction.

E. Mixed A64-ASIMD Karatsuba Multiplication

One of the benefits of using divide-and-conquer method such as Karatsuba multiplication is that the parts which build the final result, e.g., $(a_h + a_l).(b_h + b_l)$, $a_h b_h$, and $a_l b_l$ can be computed independently. This fact is in contrast with other multiplication methods such as Comba-method which requires the carry to be propagated step-by-step between each slot of the final result; as a result, the divided parts can be computed in parallel and combined at the end. In our mixed version of implementation, we exploit this advantage and compute $a_h b_h$ and $a_l b_l$ concurrently. We choose to implement $a_h b_h$ using A64 and $a_l b_l$ using ASIMD assembly instructions. Fig. 2 illustrates the overall overview of our mixed A64-ASIMD multiplication for 512-bit input operands.

As it is already discussed in details [16], [25], it is rather inefficient to attempt ASIMD implementation for field arithmetic on 64-bit ARM due to the smaller radix representation of operands, i.e., radix- 2^{32} and therefore extended number of arithmetic operations. However, in our engineered mixed implementation, many pipeline stalls are removed; this may result in improvement in overall performance the multiplication.

The first step to design our mixed multiplier is to choose what algorithm to use for each half of multiplication, i.e., $a_h b_h$ and $a_l b_l$. For A64 implementation, we found the Comba-multiplication method fast and optimum. In particular, since the multiplication size is cut in half ($\frac{n}{2}$ -bit), we have access to redundant number of 64-bit general registers that can be used to implement the operand-scanning method without any need to load/store data back and forth in the middle of process. On the other hand, for the ASIMD implementation of $a_l b_l$, Comba-multiplication is not a suitable option due to the lack of carry-propagation instruction. Previous works on the efficient implementation of multi-precision multiplication on ARMv7 NEON [24], [36] present an efficient field

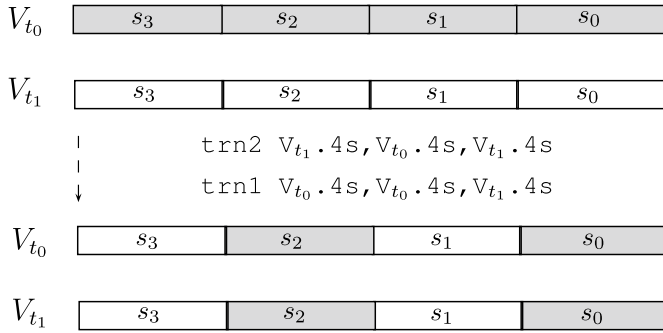


Fig. 3. Transposition of data into two vectors to handle carry-overflow.

multiplication using parallel school-book method. We adopted the same strategy; however, since the vector manipulation in ARMv8 ASIMD is totally different from ARMv7 NEON, we customized our version of school-book multiplication on ARMv8 platforms.

A64-based product-scanning method is implemented by loading the most significant halves of input operands, and computing each slot of the result by multiplication and addition of corresponding input words. ASIMD-based school-book multiplication, however, is more complicated to implement. We explain the implementation steps in detail in the following:

- 1) After loading the least-significant halves of input operands into vectors using `ld1` instruction, one of the input vectors is rearranged in correspond with the intermediate additions positions (highlighted words in Fig. 2) using transpose instructions `trn1` and `trn2`. Note that, ASIMD multiplication instructions compute two 32×32 -bit multiplications concurrently, i.e., a_0b_0 and a_4b_0 are computed using one multiplication instruction in parallel.
- 2) Each 32-bit limb of “non-shuffled” operand is multiplied to the entire shuffled operand in each step using `umul1` and `umul12` instructions. These instruction computes the multiplication of a 32-bit limb to the first and second halves of a 128-bit ASIMD vector, respectively. We use these instructions only in the first step of the multiplication process since we need no addition. For the next steps, we exploit `umlal` and `umlal2` instructions which compute the multiplication and addition at the same time.
- 3) Vectorized additions are performed in parallel and can add two 128-bit vector using only one `add` instruction. However, vectorized addition does not handle carry-propagation between vector slots; accordingly, intermediate multiplication results need to be transposed from 4×32 -bit values into 2×32 -bit ones, while the remaining two slots are preserved for carry overflows. Fig. 3, shows this process for a 128-bit wide ASIMD vector. This transposition is performed by compounding the multiplication result vector (V_{t_0}) with a zero vector (V_{t_1}) using `trn1` and `trn2` instructions.
- 4) The ASIMD-based multiplication performs Steps 2 and 3 iteratively for each 32-bit word of the second operand. In the final step we compute the last additions and carry propagation operations and move the result vectors

into general registers using `umov` instruction to perform subtraction ($a_h b_h - a_l b_l$).

We implement the rest of Karatsuba multiplication operations, i.e., $(a_h + a_l) \cdot (b_h + b_l)$ and subtractions, using A64 instruction, taking advantage of its 64-bit wide arithmetic. Furthermore, we design 512-, 768-, and 1024-bit mixed Karatsuba multiplication using the above technique and integrate them into the SIKE software, providing three different quantum security levels.

We remark that since all the curve arithmetic in SIKE implementation is performed on Montgomery space, the most optimized algorithm for reduction is Montgomery reduction which is already used in the SIKE implementation submission. Moreover, SIKE primes, i.e., \mathbb{F}_{p503} , \mathbb{F}_{p751} , and \mathbb{F}_{p964} , have a special form which is utilized to improve the reduction algorithm by eliminating several single-precision multiplications; we refer the reader to [23], [25] for more details. Therefore, we believe the most efficient implementation of reduction algorithm on ARMv8 is based on product-scanning method using A64 assembly instructions.

IV. OPTIMIZED SHA-3 IMPLEMENTATION

SIKE reference implementation requires three hash functions F , G , and H which are used inside encapsulation and decapsulation algorithms. These hash functions are all instances of `cSHAKE256` with different custom input strings which is based on SHA-3 library and specified by NIST. Compared to isogeny computations, the computational cost of these functions is negligible and they barely affect the overall performance of the SIKE protocol. However, in this work, we focus on using the most efficient implementation of this protocol on ARMv8 platforms; therefore, we replace the generic implementation of SHA-3 with the hand-written assembly version developed by OpenSSL² project authors in our software. The essential function inside the SHA-3 is the `Keccak-1600` function which performs the core permutations. Note that, the optimized implementation of this function is developed using A64 assembly instructions without taking advantage of ASIMD vectors for the following reason. Addressing 64-bit lanes of ASIMD vectors is not as trivial as 32-bit NEON on ARMv7 processors. In particular, we found it is rather complicated to perform `rotate` operation using available ASIMD instructions on 64-bit lanes which adds up significant number of clock cycles to the overall timing results. Therefore, 64-bit ASIMD implementation of `Keccak-1600` is slower than A64 general register implementation on ARMv8 processors.

V. PERFORMANCE RESULTS AND DISCUSSION

In this section, we present our benchmark procedure and provide SIKE performance evaluation results on two famous ARMv8 family of processors, i.e., Cortex-A57 and Cortex-A72.

A. Implementation Details

As it is mentioned before, our optimized implementation targets high-performance ARMv8 processors which

²Available in: <https://github.com/openssl/openssl>

TABLE III

PERFORMANCE RESULTS (PRESENTED IN MILLIONS OF CLOCK CYCLES) OF THE PROPOSED SOFTWARE LIBRARIES IN COMPARISON WITH REFERENCE IMPLEMENTATION ON ARMV8 PLATFORMS. (BENCHMARKS WERE OBTAINED ON 1.95 GHz CORTEX-A57 AND 2.4 GHz CORTEX-A72 CORES RUNNING ANDROID 7.1.1 AND 8.1.0, RESPECTIVELY)

Scheme	Operation	A64 assembly		(ASIMD + A64) assembly		Optimized C	
		Cortex-A57	Cortex-A72	Cortex-A57	Cortex-A72	Cortex-A57	Cortex-A72
SIKEp503	KeyGen.	31.2	31.2	31.2	33.6	150.1	147.5
	Encap.	50.7	50.4	52.6	55.2	245.7	245
	Decap.	54.6	54.1	56.5	57.6	261.3	260
	Total	136.5	135.7	140.3	146.4	657.1	652.5
SIKEp751	KeyGen.	101.4	100.8	99.4	98.4	528.4	512.5
	Encap.	163.8	162.6	161.8	161.2	858	835
	Decap.	175.5	174.8	173.5	172.4	922.3	892.5
	Total	440.7	438.2	434.7	432	2,308	2,255
SIKEp964	KeyGen.	222.3	220.8	216.4	214.2	1,201.2	1,167.6
	Encap.	374.4	372	362.7	361.1	2,020.2	1,918.8
	Decap.	393.9	393.6	382.2	380.7	2,131.3	2,090.9
	Total	990.6	986.4	961.3	956	5,352.7	5,177.3

support out-of-order pipeline. We use a Huawei Nexus 6P cellphone, running Android Nougat 7.1.1 and a Google Pixel 2, running Android Oreo v8.1.0, for benchmarking our software on Cortex-A57 and Cortex-A72, respectively. The binaries are cross-compiled with clang 5.0.3 using `-O3 -fomit-frame-pointer -pie` flags and run via `adb-shell`. The `-pie` flag is used to generate position independent executables which can be run on cellphones. We benchmarked the executables on one high-performance core using `taskset` command to ensure power-efficient cores were not involved.

Our optimized software libraries using both A64 and mixed-assembly are publicly available³ for different levels of security.

B. Performance Results

Table III presents benchmark results of our SIKE implementation on target platforms for different security levels. The results highlight the performance of our optimized ARMv8 assembly implementation using A64 assembly and mixed assembly in comparison with the portable optimized C reference implementation [8]. Results are averaged over 10^3 iterations and converted to clock cycle counts corresponding to the target processor frequency. Based on the benchmark results, the optimized arithmetic libraries provide roughly $5\times$ faster results than generic implementation on the target platforms. Moreover, we notice that for larger finite fields, our mixed ASIMD/A64 version outperforms the A64 optimized library by a very small factor, while in relatively smaller fields, the A64 implementation is more efficient. This highlights the fully pipeline utilization that we obtain using mixed assembly implementation. However, as it is discussed before, since ASIMD arithmetic operations are performed in radix-2³², the overall performance improvement is not remarkable. We also observe that the clock cycle count on Cortex-A72 core is slightly less than Cortex-A57. This performance enhancement is related to architecture design improvements in Cortex-A72 family of processors which contain many micro-architecture updates for performance improvement compared to its prior generation.

³<https://github.com/amirjalali65/armv8-sike>

TABLE IV

PERFORMANCE COMPARISON OF THIS WORK WITH THE PREVIOUS ARMV8 OPTIMIZED IMPLEMENTATION OF SIKEp751 ON A CORTEX-A57 CORE PRESENTED IN MILLIONS OF CLOCK CYCLES

Implementation	KeyGen	Encap.	Decap.	Total
This Work	99.4	161.8	173.5	434.7
SIDHv3.0*	109.2	179.4	193.2	481.8

* Supports optimized ARMv8 implementation of SIKEp751

In order to evaluate the efficiency of Karatsuba multiplication over \mathbb{F}_{p751} , we compared our mixed optimized library with the only available⁴ highly-optimized ARMv8 SIKE implementation. We compared the performance of the SIDH v3.0 which supports optimized ARMv8 SIKE implementation with our software for the only available security level, i.e., SIKEp751. We have noticed the available ARMv8 implementation of SIKEp503 inside the SIDHv3.0 library is implemented in C and therefore the performance comparison with this work is not fair. Table IV presents the performance evaluation of the SIKE implementation inside SIDHv3.0 in comparison with this work. The presented benchmarks show roughly 10% performance improvement of our library over highly-optimized previous library for the same field size. We state that the performance improvement of this results corresponds to different multiplication algorithm and implementation technique that we use in this work. Moreover, it proves the fact that ARMv8 vector instructions can still be useful inside field arithmetic operations. In particular, using mixed implementation and taking advantage of 256-bit wide ASIMD vectors can result in performance improvement over larger fields. That is because the scarce number of available 64-bit general registers enforces an excessive number of memory-register load and stores instructions, and using vectors will reduce these costly operations, and improve the overall performance.

C. Broader Impact of the Proposed Work

The proposed method of multiplication in this work is independent of the shape of the SIKE primes and thus it can be used in other cryptographic schemes. In particular, we observe

⁴<https://github.com/microsoft/pqcrypto-sidh> (accessed in Feb. 2018)

TABLE V

PERFORMANCE COMPARISON OF CLASSICAL AND POST-QUANTUM ELLIPTIC CURVE BASED CRYPTOGRAPHY ON CORTEX-A57 AND CORTEX-A72 PRESENTED IN 10^6 CLOCK CYCLES

Scheme	Post-Quantum	Cortex-A57	Cortex-A72
Curve25519	No	3.04	3.01
FourQ	No	1.82	1.79
SIKEp503	Yes	657.1	652.5

larger bit-length multiplications benefit more from the proposed mixed implementation, since it offers less number of memory-register transfers and pipeline stalls. As a result, we expect to obtain better performance of different classical and post-quantum cryptographic schemes on high-performance ARMv8 processors by adopting the proposed implementation method. For instance, classical and post-quantum RSA [37] and ECC implementations over large finite fields are the most promising candidates which can benefit significantly from the proposed method on ARMv8 platforms. Moreover, non-RNS variants of homomorphic encryption schemes such as original Fan-Vercauteren method [38] and its inherited schemes can also profit from the proposed method inside their underlying large integer multiplications of polynomial coefficients.

In all the above applications, we expect to achieve better performance improvement of our proposed method over relatively larger parameter sets.

D. Comparison With Classical ECC

To provide a clear performance comparison of classical and post-quantum elliptic curve cryptography, we benchmarked two popular classical elliptic curve key-exchange cryptography schemes, Curve25519 [39] and FourQ [40] on the target platforms, and provided the obtained results in Table V. Since the optimized assembly implementation of these schemes are not publicly available, we only present the portable implementation results in comparison with SIKEp503 portable implementation. As we expected, the quantum resistant SIKE shows slower results by orders of magnitude compared to classical ECC.

VI. CONCLUSION

The main motivation behind this work was to push the performance bar of the post-quantum SIKE protocol further on 64-bit ARM-powered embedded devices; we presented an optimized implementation of SIKE protocol on 64-bit high-performance ARM processors. We investigated different design of implementing multi-precision multiplication on ARMv8 processors, with and without using ASIMD vectorization. We showed that incorporating ASIMD implementation, mixed with A64 design, over relatively large finite fields, can improve the overall performance of projective coordinates cryptography protocols such as SIKE.

Our proposed approach can be used inside any other cryptography primitives which require significant amount of field multiplication over large primes. Moreover, we benchmarked our implementations on two popular Android cellphones to evaluate the practical efficiency of the SIKE protocol. We compared our results with the previous SIKE optimized implementation to justify the performance benefits we gained using our proposed implementation techniques. We believe the proposed

method can be used on other platforms which support out-of-order execution pipeline, improving the overall performance of the protocol by reducing pipeline stalls. We plan to investigate this possibility in the future works.

ACKNOWLEDGMENT

The authors would like to thank the reviewers for their comments.

REFERENCES

- [1] T. Lange. (2015). *PQCRYPTO Project in the EU, NIST Workshop on Cybersecurity in a Post-Quantum World*. Accessed: Feb. 2018. [Online]. Available: <http://csrc.nist.gov/groups/ST/post-quantum-2015/presentations/session7-lange-tanja.pdf>
- [2] R. J. McEliece, "A public-key cryptosystem based on algebraic," *Coding Thv.*, vol. 4244, pp. 114–116, Apr. 1978.
- [3] J. Hoffstein, J. Pipher, and J. H. Silverman, "NTRU: A ring-based public key cryptosystem," in *Proc. Int. Algorithmic Number Theory Symp.* Berlin, Germany: Springer, 1998, pp. 267–288.
- [4] M. Ajtai, "Generating hard instances of lattice problems," in *Proc. ACM Symp. Theory Comput.*, 1996, pp. 99–108.
- [5] R. C. Merkle, "Secrecy, authentication, and public key systems," Ph.D. dissertation, Dept. Elect. Eng., Stanford Univ., Stanford, CA, USA, 1979.
- [6] T. Matsumoto and H. Imai, "Public quadratic polynomial-tuples for efficient signature-verification and message-encryption," in *Proc. Workshop Theory Appl. Cryptogr. Techn.* Berlin, Germany: Springer, 1988, pp. 419–453.
- [7] D. Jao and L. De Feo, "Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies," in *Proc. Int. Workshop Post-Quantum Cryptogr.* Berlin, Germany: Springer, 2011, pp. 19–34.
- [8] D. Jao *et al.*, "Supersingular isogeny key encapsulation," NIST, Gaithersburg, MD, USA, Tech. Rep., 2017.
- [9] P. W. Shor, "Algorithms for quantum computation: Discrete logarithms and factoring," in *Proc. Annu. Symp. Found. Comput. Sci.*, 1994, pp. 124–134.
- [10] A. Rostovtsev and A. Stolbunov, "Public-key cryptosystem based on isogenies," *IACR Cryptol. ePrint Arch.*, Tech. Rep. 2006/145, 2006, p. 145.
- [11] D. X. Charles, K. E. Lauter, and E. Z. Goren, "Cryptographic hash functions from expander graphs," *J. Cryptol.*, vol. 22, no. 1, pp. 93–113, 2009.
- [12] L. De Feo, D. Jao, and J. Plût, "Towards quantum-resistant cryptosystems from supersingular elliptic curve isogenies," *J. Math. Cryptol.*, vol. 8, no. 3, pp. 209–247, 2014.
- [13] S. D. Galbraith, C. Petit, and J. Silva, "Signature schemes based on supersingular isogeny problems," *IACR Cryptol. ePrint Arch.*, Tech. Rep. 2016/1154, 2016. Accessed: Feb. 2018.
- [14] Y. Yoo, R. Azarderakhsh, A. Jalali, D. Jao, and V. Soukharev, "A post-quantum digital signature scheme based on supersingular isogenies," in *Proc. Int. Conf. Financial Cryptogr. Data Secur.* Cham, Switzerland: Springer, 2017, pp. 163–181.
- [15] D. Jao and V. Soukharev, "Isogeny-based quantum-resistant undeniable signatures," in *Proc. Int. Workshop Post-Quantum Cryptogr.* Cham, Switzerland: Springer, 2014, pp. 160–179.
- [16] A. Jalali, R. Azarderakhsh, and M. Mozaffari-Kermani, "Efficient post-quantum undeniable signature on 64-bit arm," in *Proc. Int. Conf. Sel. Areas Cryptogr.* Cham, Switzerland: Springer, 2017, pp. 281–298.
- [17] R. Azarderakhsh, A. Jalali, D. Jao, and V. Soukharev, "Practical supersingular isogeny group key agreement," *IACR Cryptol. ePrint Arch.*, Tech. Rep. 2019/330, 2019. [Online]. Available: <https://eprint.iacr.org/2019/330>
- [18] R. Azarderakhsh, D. Jao, and C. Leonardi, "Post-quantum static-static key agreement using multiple protocol instances," in *Proc. Int. Conf. Sel. Areas Cryptogr.* Cham, Switzerland: Springer, 2017, pp. 45–63.
- [19] B. Kozziel, R. Azarderakhsh, M. M. Kermani, and D. Jao, "Post-quantum cryptography on FPGA based on isogenies on elliptic curves," *IEEE Trans. Circuits Syst.*, vol. 64, no. 1, pp. 86–99, Jan. 2017.
- [20] J. W. Bos, C. Costello, M. Naehrig, and D. Stebila, "Post-quantum key exchange for the TLS protocol from the ring learning with errors problem," in *Proc. IEEE Symp. Secur. Privacy*, May 2015, pp. 553–570.
- [21] E. Alkim, L. Ducas, T. Pöppelmann, and P. Schwabe, "Post-quantum key exchange—A new hope," *IACR Cryptol. ePrint Arch.*, Tech. Rep. 2015/1092, 2015. Accessed: Feb. 2018. [Online]. Available: <http://eprint.iacr.org>

- [22] J. Bos *et al.*, “Frodo: Take off the ring! Practical, quantum-secure key exchange from lwe,” in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2016, pp. 1006–1018.
- [23] C. Costello, P. Longa, and M. Naehrig, “Efficient algorithms for supersingular isogeny Diffie-Hellman,” in *Advances in Cryptology*. Berlin, Germany: Springer, 2016.
- [24] B. Koziel, A. Jalali, R. Azarderakhsh, D. Jao, and M. Mozaffari-Kermani, “NEON-SIDH: Efficient implementation of supersingular isogeny Diffie-Hellman key exchange protocol on ARM,” in *Proc. Int. Conf. Cryptol. Netw. Secur.* Cham, Switzerland: Springer, 2016, pp. 88–103.
- [25] A. Jalali, R. Azarderakhsh, M. M. Kermani, and D. Jao, “Supersingular isogeny Diffie-Hellman key exchange on 64-bit ARM,” *IEEE Trans. Depend. Sec. Comput.*, to be published.
- [26] A. Jalali, R. Azarderakhsh, and M. M. Kermani, “NEON SIKE: Supersingular isogeny key encapsulation on ARMv7,” in *Proc. Int. Conf. Secur., Privacy, Appl. Cryptogr. Eng.*, 2018, pp. 37–51.
- [27] A. Jalali, R. Azarderakhsh, M. M. Kermani, and D. Jao, “Towards optimized and constant-time CSIDH on embedded devices,” in *Proc. 10th Int. Workshop Constructive Side-Channel Anal. Secure Design (COSADE)*, 2019, pp. 215–231.
- [28] A. Faz-Hernández, J. López, E. Ochoa-Jiménez, and F. Rodríguez-Henríquez, “A faster software implementation of the supersingular isogeny diffie-hellman key exchange protocol,” *IEEE Trans. Comput.*, vol. 67, no. 11, pp. 1622–1636, Nov. 2018.
- [29] S. Ali and M. Cenk, “Faster residue multiplication modulo 521-bit Mersenne prime and an application to ECC,” *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 65, no. 8, pp. 2477–2490, Aug. 2018.
- [30] P. Longa, “Four@NEON: Faster elliptic curve scalar multiplications on ARM processors,” in *Proc. 23rd Int. Conf. Sel. Areas Cryptogr. (SAC)*, 2016, pp. 501–519.
- [31] H. Seo, Z. Liu, P. Longa, and Z. Hu, “SIDH on ARM: Faster modular multiplications for faster post-quantum supersingular isogeny key exchange,” *IACR Trans. Cryptograph. Hardw. Embedded Syst.*, vol. 2018, no. 3, pp. 1–20, 2018.
- [32] J. Vélú, “Isogénies entre courbes elliptiques,” *CR Acad. Sci. Paris A*, vol. 273, pp. A238–A241, Jan. 1971.
- [33] J. H. Silverman, *The Arithmetic of Elliptic Curves*, vol. 106. New York, NY, USA: Springer, 2009.
- [34] S. Tani, “Claw finding algorithms using quantum walk,” *Theor. Comput. Sci.*, vol. 410, no. 50, pp. 5285–5297, 2009.
- [35] ARM Limited. (2016). *Cortex-A57 Software Optimization Guide*. Accessed: Feb. 2018. [Online]. Available: http://infocenter.arm.com/help/topic/com.arm.doc.uan0015b/cortex_a57_software_optimization_guide_external.pdf
- [36] H. Seo, Z. Liu, J. Großschädl, J. Choi, and H. Kim, “Montgomery modular multiplication on ARM-NEON revisited,” in *Proc. Int. Conf. Inf. Secur. Cryptol.* Cham, Switzerland: Springer, 2014, pp. 328–342.
- [37] D. J. Bernstein, N. Heninger, P. Lou, and L. Valenta, “Post-quantum RSA,” in *Proc. 8th Int. Workshop Post-Quantum Cryptogr. (PQCrypto)*, 2017, pp. 311–329.
- [38] J. Fan and F. Vercauteren, “Somewhat practical fully homomorphic encryption,” *IACR Cryptol. ePrint Arch.*, Tech. Rep. 2012/144, 2012, p. 144.
- [39] D. J. Bernstein, “Curve25519: New Diffie-Hellman speed records,” in *Proc. 9th Int. Workshop Public Key Cryptogr. (PKC)*, 2006, pp. 207–228.
- [40] C. Costello and P. Longa, “four Q: four-dimensional decompositions on a Q-curve over the mersenne prime,” *IACR Cryptol. ePrint Arch.*, Tech. Rep. 2015/565, 2015, p. 565.



Amir Jalali received the B.Sc. degree in computer engineering from Shahid Beheshti University, Tehran, Iran, in 2009, the M.Sc. degree in computer engineering from the Department of Computer Engineering and Information Technology, Amirkabir University of Technology, Tehran, Iran, in 2012, and the Ph.D. degree in computer engineering from the Department of Computer, Electrical Engineering and Computer Science, Florida Atlantic University, USA, in 2018. His current research interests include efficient software implementation of elliptic curve cryptography, post-quantum cryptography, and homomorphic encryption.



Reza Azarderakhsh (M’12) received the Ph.D. degree in electrical and computer engineering from Western University in 2011. He was a recipient of the NSERC Post-Doctoral Research Fellowship with the Center for Applied Cryptographic Research, Department of Combinatorics and Optimization, University of Waterloo. He is currently an Assistant Professor with the Department of Electrical and Computer Engineering, Florida Atlantic University. His current research interests include finite field and its application, elliptic curve cryptography, pairing based cryptography, and post-quantum cryptography. He is serving as an Associate Editor for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS (TCAS-I).



Mehran Mozaffari Kermani (S’00–M’11–SM’16) received the B.Sc. degree from the University of Tehran, Iran, and the M.E.Sc. and Ph.D. degrees from the University of Western Ontario, London, Canada, in 2007 and 2011, respectively. In 2012, he joined the Electrical Engineering Department, Princeton University, NJ, USA, as an NSERC Post-Doctoral Research Fellow. From 2013 to 2017, he was an Assistant Professor with the Rochester Institute of Technology. In 2017, he has joined the Computer Science and Engineering Department, University of South Florida. He has been the TPC member for a number of conferences including HOST (Publications Chair), DAC, DATE, RFIDSec, LightSec, WAIFI, FDTC, and DFT. He is currently serving as an Associate Editor for the IEEE TVLSI, ACM TECS, and IEEE TCAS I. He has been the Guest Editor for the IEEE TDSC, IEEE/ACM TCBB, and the IEEE TETC for special issues on security.



Matthew Campagna received the bachelor’s degree in mathematics from Fordham University and the Ph.D. degree in mathematics (group theory) from Wesleyan University. He managed the Secure Systems Research Group, Pitney Bowes. He was with the U.S. Department of Defense. He led the Certicom Research Group, BlackBerry, managing research in cryptography and standards participation in ANSI, ZigBee, SECG, ETSI’s SAGE, and the 3GPP-SA3 working group. He is currently the Principal Security Engineer with the Cryptography Group, Amazon Web Services, Inc., where he currently leads the cryptographic design of AWS cryptographic services such as AWS KMS, CloudHSM, and AWS Certificate Manager. He is also an affiliate of the Institute for Quantum Computing, University of Waterloo, and a member of the ETSI Security Algorithms Group Experts (SAGE).



David Jao received the Ph.D. degree in mathematics from Harvard University, Cambridge, MA, USA, in 2003. From 2003 to 2006, he was with the Cryptography and Anti-Piracy Group, Microsoft Research, contributing cryptographic software modules for several Microsoft products. He is currently an Associate Professor with the Mathematics Faculty, University of Waterloo, Canada, and the Director of the Centre for Applied Cryptographic Research. His research interests include elliptic curve cryptography, protocol design and implementation, and post-quantum cryptography.