

Reduction Modulo $2^{448} - 2^{224} - 1$

Kaushik Nath and Palash Sarkar

Applied Statistics Unit
Indian Statistical Institute
203, B. T. Road
Kolkata - 700108
India
{kaushikn.r,palash}@isical.ac.in

Abstract

An elliptic curve known as Curve448 defined over the finite field \mathbb{F}_p , where $p = 2^{448} - 2^{224} - 1$, has been proposed as part of the Transport Layer Security (TLS) protocol, version 1.3. Elements of \mathbb{F}_p can be represented using 7 limbs where each limb is a 64-bit quantity. This paper describes efficient algorithms for reduction modulo p that are required for performing field arithmetic in \mathbb{F}_p using 7-limb representation. A key feature of our work is that we provide the relevant proofs of correctness of the algorithms. We also report efficient 64-bit assembly implementations for key generation and shared secret computation phases of the Diffie-Hellman key agreement protocol on Curve448. Timings results on the Haswell and Skylake processors demonstrate that the new 64-bit implementations for computing the shared secret are faster than the previously best known 64-bit implementations.

Keywords: Curve448, Goldilocks prime, modulo reduction, elliptic curve cryptography, Diffie-Hellman key agreement.

1 Introduction

As part of the Transport Layer Security (TLS) protocol, version 1.3 [7], RFC 7748 [2] specifies the Montgomery form elliptic curve Curve448 and its birationally equivalent Edwards form elliptic curve Edwards448. The curve Edwards448 was originally proposed in [1] where it was named Ed448-Goldilocks. The underlying field for Curve448 and Edwards448 is \mathbb{F}_p where p is the prime $2^{448} - 2^{224} - 1$.

Implementation of elliptic curve operations require arithmetic over the underlying field \mathbb{F}_p . Specifically, addition, subtraction, multiplication and squaring are required. Additionally, for implementing Montgomery ladder for Curve448, it is required to implement multiplication by a small constant. All of these operations require reduction modulo p .

For 64-bit architecture, an element of \mathbb{F}_p can be represented using 7 limbs where each limb is a 64-bit quantity. Such a representation can be considered to be a packed or, saturated limb representation of the elements of \mathbb{F}_p . Alternatively, elements of \mathbb{F}_p may be represented using 8 limbs where each limb is a 56-bit quantity stored in a 64-bit word. Such a representation can be considered to be a redundant or, unsaturated limb representation. For modern Intel processors such as Skylake and later processors, the implementation of field arithmetic using the saturated limb representation turns out to be faster than that of the unsaturated limb representation.

Our contributions

Algorithms along with their proofs of correctness. In this work, we consider the 7-limb saturated limb representation of elements of \mathbb{F}_p . Our focus is on the reduction algorithms which are required to implement field arithmetic operations in \mathbb{F}_p .

The main contribution of the paper is to present explicit reduction algorithms along with their proofs of correctness for all the field arithmetic operations required to implement Diffie-Hellman key agreement using Curve448. The algorithms proceed over several iterations successively reducing the size of the input. As part of the proof of correctness, it is required to argue that the algorithms terminate without any overflow. The termination argument has a certain amount of subtlety. To the best of our knowledge, no previous work had considered the issue of proof of correctness. Without a formal argument about

termination, a reduction strategy may turn out to be incomplete or may perform redundant operations; we provide a short discussion of these possibilities in the appendix.

Efficient 64-bit assembly implementations of X448. Computing the Diffie-Hellman key agreement over the curve Curve448 demands computation of scalar multiplication over Curve448. This computation has been named as X448 in [2]. Implementation of scalar multiplication requires implementation of field arithmetic over the underlying field. We have implemented field arithmetic over \mathbb{F}_p , and based on it we have developed efficient assembly implementations of the X448 function of Curve448. The performances of our 64-bit implementations for shared secret computation are faster than the previously best known 64-bit implementations. Timing details are provided later. We have made our software publicly available at the following link.

<https://github.com/kn-cs/x448/tree/master/7limb>

2 Arithmetic in \mathbb{F}_p

Let $p = 2^{448} - 2^{224} - 1$ and $\theta = 2^{64}$. For $d \geq 0$, define the polynomial

$$f(\theta) = f_0 + f_1\theta + \cdots + f_d\theta^d \quad (1)$$

where f_0, f_1, \dots, f_d are non-negative integers. Following usual convention, we will call the f_i 's to be limbs of $f(\theta)$.

As mentioned above, we consider the 7-limb representation of the elements of \mathbb{F}_p . So, elements of \mathbb{F}_p can be represented as a polynomial $f(\theta) = f_0 + f_1\theta + \cdots + f_6\theta^6$ where $0 \leq f_0, f_1, \dots, f_6 < \theta$. Note that the set of all such $f(\theta)$ is in one-one correspondence with the set of integers $\{0, 1, \dots, 2^{448} - 1\}$. Since, $p < 2^{448} - 1$, a degree 6 polynomial $f(\theta)$ with $0 \leq f_0, f_1, \dots, f_6 < \theta$ is not necessarily reduced modulo p . So, some elements of \mathbb{F}_p have non-unique representation. This, however, is not a problem for intermediate quantities in an elliptic curve computation. It is only the final result that is reduced to have a unique representation modulo p . Avoiding obtaining unique representations for the intermediate quantities leads to an overall faster algorithm for performing the elliptic curve computation. Consequently, given a polynomial $h(\theta) = h_0 + h_1\theta + \cdots + h_d\theta^d$, with $d > 0$, by reduction modulo p , we will denote the task of obtaining a polynomial $f(\theta) = f_0 + f_1\theta + \cdots + f_6\theta^6$ with $0 \leq f_0, f_1, \dots, f_6 < \theta$ such that $f(\theta) \equiv h(\theta) \pmod{p}$.

For $i \geq 2$, let x and y be two $64i$ -bit integers. Suppose, it is required to compute the integer product $x \cdot y$. If $x = y$, then this corresponds to the squaring operation, while if $x \neq y$, then a general multiplication operation is required. Intel processors from Broadwell (launched in 2014) onwards provide a special set of 64-bit multiplication and addition instructions which allow very fast computation of the product $x \cdot y$. For $i = 4$, the multiplication and squaring algorithms have been illustrated using diagrams in two Intel white papers [5, 6]. Explicit descriptions of the squaring and multiplication algorithms in the general case have been provided in [3].

A field multiplication/squaring in \mathbb{F}_p consists of the following two broad steps. Suppose that $f(\theta)$ and $g(\theta)$ are two 7-limb integers from the set $\{0, 1, \dots, 2^{448} - 1\}$ representing elements of \mathbb{F}_p . In the first step, the integer product of $f(\theta)$ and $g(\theta)$ is obtained in $h(\theta)$. The quantity $h(\theta)$ can be written as a 14-limb quantity $h(\theta) = h_0 + h_1\theta + \cdots + h_{13}\theta^{13}$, where $0 \leq h_0, h_1, \dots, h_{13} < 2^{64}$. The second step consists of reducing $h(\theta)$ to a 7-limb integer which is congruent to $h(\theta)$ modulo p .

The Montgomery ladder algorithm for Curve448 requires multiplying a 7-limb quantity $f(\theta)$ by the constant $c = 39082$ (note, $2^{15} < c < 2^{16}$ and so c is a 16-bit quantity). The integer product $c \cdot f(\theta)$ can be computed much faster than a general integer multiplication of two 7-limb quantities. The result $c \cdot f(\theta)$ can be written as an 8-limb quantity where all the limbs are 64-bit quantities. A reduction algorithm is to be applied to this 8-limb quantity to reduce it to a 7-limb quantity which represents an element of \mathbb{F}_p .

The integer addition of two 7-limb integers $f(\theta)$ and $g(\theta)$ results in an 8-limb integer. In this case, the last limb is a single bit. Nevertheless, the result of the addition has to be reduced to a 7-limb quantity.

Subtraction of two elements $f(\theta)$ and $g(\theta)$ in \mathbb{F}_p is more problematic. The integer operation $f(\theta) - g(\theta)$ can turn out to be negative. To avoid handling negative numbers a suitable multiple of p is added to the result. This creates subtleties in the reduction algorithm.

3 Reduction in \mathbb{F}_p

In Section 3.1 below, we describe the method for reducing a 14-limb quantity to a 7-limb quantity. As part of this algorithm, it is required to reduce an 8-limb quantity to a 7-limb quantity. Correspondingly,

this part can be used to reduce the result obtained either after multiplication by a 64-bit constant or after addition of two 7-limb quantities. This is pointed out in Section 3.2. The case of subtraction in \mathbb{F}_p is described in Section 3.3.

3.1 Reduction from 14-Limb to 7-Limb

Let $h(\theta)$ be the 14-limb polynomial which is to be reduced. The polynomial $h(\theta)$ represents an integer z of $2 \cdot 448 = 896$ bits. A formal description of the algorithm to reduce $h(\theta)$ is given in Function `reduce448` of Algorithm 1. All the operations in `reduce448` can be performed using 64-bit arithmetic instructions available in modern processors. For showing correctness of the algorithm it is required to argue that the output is indeed congruent to the input modulo p . Further, it is also required to argue that the procedure terminates without any overflow.

Let $h^{(0)}(\theta) = h(\theta)$. Function `reduce448` takes the 14-limb polynomial $h^{(0)}(\theta)$ as input and reduces it through the intermediate polynomials $h^{(1)}(\theta)$, $h^{(2)}(\theta)$ finally producing the 7-limb output polynomial $h^{(3)}(\theta)$. A summary of the properties of the polynomials $h^{(1)}(\theta)$, $h^{(2)}(\theta)$ and $h^{(3)}(\theta)$ and the different steps of `reduce448` that produces these polynomials are as follows:

- $h^{(1)}(\theta)$ has 8 limbs. The last limb is at most 2 bits long. The computation of $h^{(1)}(\theta)$ from $h^{(0)}(\theta)$ is achieved by Steps 4-26.
- $h^{(2)}(\theta)$ has 8 limbs. The last limb is at most 1-bit long and further, if $h_7^{(2)} = 1$, then $h_4^{(2)} = h_5^{(2)} = h_6^{(2)} = 0$. The computation of $h^{(2)}(\theta)$ from $h^{(1)}(\theta)$ is achieved by Steps 27-33.
- $h^{(3)}(\theta)$ has 7 limbs where each limb is a 64-bit quantity. The computation of $h^{(3)}(\theta)$ from $h^{(2)}(\theta)$ is achieved by Steps 34-38.

The properties of $h^{(1)}(\theta)$, $h^{(2)}(\theta)$ and $h^{(3)}(\theta)$ stated above are formally proved in Theorem 1. In particular, we note that the second property stated above is required to argue that the procedure terminates without any overflow in the next iteration.

Theorem 1. *Suppose the input $h^{(0)}(\theta) = h_0^{(0)} + h_1^{(0)}\theta + \dots + h_{13}^{(0)}\theta^{13}$ to `reduce448` is such that $0 \leq h_i^{(0)} < 2^{64}$ for $i = 0, 1, \dots, 13$. Then the output $h^{(3)}(\theta)$ of `reduce448` is such that $h^{(3)}(\theta) = h_0^{(3)} + h_1^{(3)}\theta + \dots + h_6^{(3)}\theta^6$ with $0 \leq h_j^{(3)} < 2^{64}$ for $j = 0, 1, \dots, 6$. Further, $h^{(3)}(\theta) \equiv h^{(0)}(\theta) \pmod{p}$.*

Proof. Let $\eta = 64$. We have the prime $p = 2^{448} - 2^{224} - 1$ and since $\theta = 2^{64} = 2^\eta$, we have

$$2^{448} = \theta^7 \equiv 2^{224} + 1 = 2^{\eta/2}\theta^3 + 1 \pmod{p}. \quad (2)$$

Reduction from $h^{(0)}(\theta)$ to $h^{(1)}(\theta)$. The input $h^{(0)}(\theta)$ to `reduce448` can be written as

$$\begin{aligned} h^{(0)}(\theta) &= (h_0^{(0)} + h_1^{(0)}\theta + \dots + h_6^{(0)}\theta^6) + (h_7^{(0)}\theta^7 + h_8^{(0)}\theta^8 + \dots + h_{13}^{(0)}\theta^{13}), \\ &= (h_0^{(0)} + h_1^{(0)}\theta + \dots + h_6^{(0)}\theta^6) + (h_7^{(0)} + h_8^{(0)}\theta + \dots + h_{13}^{(0)}\theta^6)\theta^7, \\ &\equiv (h_0^{(0)} + h_1^{(0)}\theta + \dots + h_6^{(0)}\theta^6) + (h_7^{(0)} + h_8^{(0)}\theta + \dots + h_{13}^{(0)}\theta^6)(2^{\eta/2}\theta^3 + 1) \text{ [using (2)]}, \\ &= (h_0^{(0)} + h_1^{(0)}\theta + \dots + h_6^{(0)}\theta^6) + (h_7^{(0)} + h_8^{(0)}\theta + \dots + h_{13}^{(0)}\theta^6) + \\ &\quad (h_7^{(0)} + h_8^{(0)}\theta + \dots + h_{13}^{(0)}\theta^6)\theta^3 2^{\eta/2}. \end{aligned} \quad (3)$$

Steps 4-8 add the two polynomials $(h_0^{(0)} + h_1^{(0)}\theta + \dots + h_6^{(0)}\theta^6)$ and $(h_7^{(0)} + h_8^{(0)}\theta + \dots + h_{13}^{(0)}\theta^6)$ limb-wise by forwarding the 1-bit carry, producing the polynomial $(r_0^{(0)} + r_1^{(0)}\theta + \dots + r_7^{(0)}\theta^7)$. Hence, from (3) we write

$$h^{(0)}(\theta) \equiv \underbrace{(r_0^{(0)} + r_1^{(0)}\theta + \dots + r_7^{(0)}\theta^7)}_{\text{through Steps 4-8}} + (h_7^{(0)} + h_8^{(0)}\theta + \dots + h_{13}^{(0)}\theta^6)\theta^3 2^{\eta/2}, \quad (4)$$

where $0 \leq r_0^{(0)}, r_1^{(0)}, \dots, r_6^{(0)} < 2^\eta$, and $0 \leq r_7^{(0)} < 2$.

For $j = 7, 8, \dots, 13$, define

$$h_j^{(0)} = h_{j,0}^{(0)} + h_{j,1}^{(0)} 2^{\eta/2}, \text{ where } h_{j,0}^{(0)} = h_j^{(0)} \pmod{2^{\eta/2}}, \text{ and } h_{j,1}^{(0)} = \lfloor h_j^{(0)} / 2^{\eta/2} \rfloor. \quad (5)$$

Algorithm 1 Reduction from 14-limb to 7-limb in \mathbb{F}_p . In the algorithm, $\eta = 64$.

```

1: function reduce448( $h^{(0)}(\theta)$ )
2: input:  $h^{(0)}(\theta) = h_0^{(0)} + h_1^{(0)}\theta + \dots + h_{13}^{(0)}\theta^{13}$  such that  $0 \leq h_i^{(0)} < 2^\eta$  for  $i = 0, 1, \dots, 13$ .
3: output:  $h^{(3)}(\theta) = h_0^{(3)} + h_1^{(3)}\theta + \dots + h_6^{(3)}\theta^6$  such that  $0 \leq h_i^{(3)} < 2^\eta$  for  $i = 0, 1, \dots, 6$  and
    $h^{(3)}(\theta) \equiv h^{(0)}(\theta) \pmod p$ .

4:    $t \leftarrow h_0^{(0)} + h_7^{(0)}$ ;  $r_0^{(0)} \leftarrow t \bmod 2^\eta$ ; carry  $\leftarrow \lfloor t/2^\eta \rfloor$ 
5:   for  $i \leftarrow 1$  to 6 do
6:      $t \leftarrow h_i^{(0)} + h_{i+7}^{(0)} + \text{carry}$ ;  $r_i^{(0)} \leftarrow t \bmod 2^\eta$ ; carry  $\leftarrow \lfloor t/2^\eta \rfloor$ 
7:   end for
8:    $r_7^{(0)} \leftarrow \text{carry}$ 
9:    $s_0^{(0)} \leftarrow r_0^{(0)}$ ;  $s_1^{(0)} \leftarrow r_1^{(0)}$ ;  $s_2^{(0)} \leftarrow r_2^{(0)}$ 
10:   $t \leftarrow r_3^{(0)} + 2^{\eta/2} \lfloor h_{10}^{(0)} / 2^{\eta/2} \rfloor$ ;  $s_3^{(0)} \leftarrow t \bmod 2^\eta$ ; carry  $\leftarrow \lfloor t/2^\eta \rfloor$ 
11:  for  $i \leftarrow 4$  to 6 do
12:     $t \leftarrow r_i^{(0)} + h_{i+7}^{(0)} + \text{carry}$ ;  $s_i^{(0)} \leftarrow t \bmod 2^\eta$ ; carry  $\leftarrow \lfloor t/2^\eta \rfloor$ 
13:  end for
14:   $s_7^{(0)} \leftarrow r_7^{(0)} + \text{carry}$ 
15:  for  $i \leftarrow 0$  to 2 do
16:     $t_i^{(0)} \leftarrow 2^{\eta/2} (h_{i+11}^{(0)} \bmod 2^{\eta/2}) + \lfloor h_{i+10}^{(0)} / 2^{\eta/2} \rfloor$ 
17:  end for
18:   $t_3^{(0)} \leftarrow 2^{\eta/2} (h_7^{(0)} \bmod 2^\eta) + \lfloor h_{13}^{(0)} / 2^{\eta/2} \rfloor$ 
19:  for  $i \leftarrow 4$  to 6 do
20:     $t_i^{(0)} \leftarrow 2^{\eta/2} (h_{i+4}^{(0)} \bmod 2^{\eta/2}) + \lfloor h_{i+3}^{(0)} / 2^{\eta/2} \rfloor$ 
21:  end for
22:   $t \leftarrow s_0^{(0)} + t_0^{(0)}$ ;  $h_0^{(1)} \leftarrow t \bmod 2^\eta$ ; carry  $\leftarrow \lfloor t/2^\eta \rfloor$ 
23:  for  $i \leftarrow 1$  to 6 do
24:     $t \leftarrow s_i^{(0)} + t_i^{(0)} + \text{carry}$ ;  $h_i^{(1)} \leftarrow t \bmod 2^\eta$ ; carry  $\leftarrow \lfloor t/2^\eta \rfloor$ 
25:  end for
26:   $h_7^{(1)} \leftarrow s_7^{(0)} + \text{carry}$ 

27:   $t \leftarrow h_0^{(1)} + h_7^{(1)}$ ;  $h_0^{(2)} \leftarrow t \bmod 2^\eta$ ; carry  $\leftarrow \lfloor t/2^\eta \rfloor$ 
28:   $t \leftarrow h_1^{(1)} + \text{carry}$ ;  $h_1^{(2)} \leftarrow t \bmod 2^\eta$ ; carry  $\leftarrow \lfloor t/2^\eta \rfloor$ 
29:   $t \leftarrow h_2^{(1)} + \text{carry}$ ;  $h_2^{(2)} \leftarrow t \bmod 2^\eta$ ; carry  $\leftarrow \lfloor t/2^\eta \rfloor$ 
30:   $t \leftarrow h_3^{(1)} + 2^{\eta/2} h_7^{(1)} + \text{carry}$ ;  $h_3^{(2)} \leftarrow t \bmod 2^\eta$ ; carry  $\leftarrow \lfloor t/2^\eta \rfloor$ 
31:   $t \leftarrow h_4^{(1)} + \text{carry}$ ;  $h_4^{(2)} \leftarrow t \bmod 2^\eta$ ; carry  $\leftarrow \lfloor t/2^\eta \rfloor$ 
32:   $t \leftarrow h_5^{(1)} + \text{carry}$ ;  $h_5^{(2)} \leftarrow t \bmod 2^\eta$ ; carry  $\leftarrow \lfloor t/2^\eta \rfloor$ 
33:   $t \leftarrow h_6^{(1)} + \text{carry}$ ;  $h_6^{(2)} \leftarrow t \bmod 2^\eta$ ;  $h_7^{(2)} \leftarrow \lfloor t/2^\eta \rfloor$ 

34:   $t \leftarrow h_0^{(2)} + h_7^{(2)}$ ;  $h_0^{(3)} \leftarrow t \bmod 2^\eta$ ; carry  $\leftarrow \lfloor t/2^\eta \rfloor$ 
35:   $t \leftarrow h_1^{(2)} + \text{carry}$ ;  $h_1^{(3)} \leftarrow t \bmod 2^\eta$ ; carry  $\leftarrow \lfloor t/2^\eta \rfloor$ 
36:   $t \leftarrow h_2^{(2)} + \text{carry}$ ;  $h_2^{(3)} \leftarrow t \bmod 2^\eta$ ; carry  $\leftarrow \lfloor t/2^\eta \rfloor$ 
37:   $h_3^{(3)} \leftarrow h_3^{(2)} + 2^{\eta/2} h_7^{(2)} + \text{carry}$ 
38:   $h_4^{(3)} \leftarrow h_4^{(2)}$ ;  $h_5^{(3)} \leftarrow h_5^{(2)}$ ;  $h_6^{(3)} \leftarrow h_6^{(2)}$ 

39:  return  $h^{(3)}(\theta) = h_0^{(3)} + h_1^{(3)}\theta + \dots + h_6^{(3)}\theta^6$ 

40: end function.

```

Using (5) for $j = 10$ we can write (4) as

$$\begin{aligned}
h^{(0)}(\theta) \equiv & (r_0^{(0)} + r_1^{(0)}\theta + \dots + r_7^{(0)}\theta^7) + \\
& (h_7^{(0)} + h_8^{(0)}\theta + h_9^{(0)}\theta^2 + (h_{10,0}^{(0)} + h_{10,1}^{(0)}2^{\eta/2})\theta^3 + h_{11}^{(0)}\theta^4 + h_{12}^{(0)}\theta^5 + h_{13}^{(0)}\theta^6)\theta^{3\cdot 2^{\eta/2}},
\end{aligned}$$

which can be further written as

$$\begin{aligned}
h^{(0)}(\theta) &\equiv (r_0^{(0)} + r_1^{(0)}\theta + \dots + r_7^{(0)}\theta^7) + (h_7^{(0)} + h_8^{(0)}\theta + h_9^{(0)}\theta^2 + h_{10,0}^{(0)}\theta^3)\theta^3 2^{\eta/2} + \\
&\quad (h_{10,1}^{(0)} + h_{11}^{(0)}2^{\eta/2} + h_{12}^{(0)}\theta 2^{\eta/2} + h_{13}^{(0)}\theta^2 2^{\eta/2})\theta^7, \\
&\equiv (r_0^{(0)} + r_1^{(0)}\theta + \dots + r_7^{(0)}\theta^7) + (h_7^{(0)} + h_8^{(0)}\theta + h_9^{(0)}\theta^2 + h_{10,0}^{(0)}\theta^3)\theta^3 2^{\eta/2} + \\
&\quad (h_{10,1}^{(0)} + h_{11}^{(0)}2^{\eta/2} + h_{12}^{(0)}\theta 2^{\eta/2} + h_{13}^{(0)}\theta^2 2^{\eta/2})(\theta^3 2^{\eta/2} + 1) \text{ [using (2)]}, \\
&= (r_0^{(0)} + r_1^{(0)}\theta + \dots + r_7^{(0)}\theta^7) + (h_{10,1}^{(0)} + h_{11}^{(0)}2^{\eta/2} + h_{12}^{(0)}\theta 2^{\eta/2} + h_{13}^{(0)}\theta^2 2^{\eta/2})\theta^3 2^{\eta/2} + \\
&\quad (h_{10,1}^{(0)} + h_{11}^{(0)}2^{\eta/2} + h_{12}^{(0)}\theta 2^{\eta/2} + h_{13}^{(0)}\theta^2 2^{\eta/2}) + (h_7^{(0)} + h_8^{(0)}\theta + h_9^{(0)}\theta^2 + h_{10,0}^{(0)}\theta^3)\theta^3 2^{\eta/2}, \\
&= (r_0^{(0)} + r_1^{(0)}\theta + \dots + r_7^{(0)}\theta^7) + (2^{\eta/2}h_{10,1}^{(0)}\theta^3 + h_{11}^{(0)}\theta^4 + h_{12}^{(0)}\theta^5 + h_{13}^{(0)}\theta^6) + \\
&\quad (h_{10,1}^{(0)} + h_{11}^{(0)}2^{\eta/2} + h_{12}^{(0)}\theta 2^{\eta/2} + h_{13}^{(0)}\theta^2 2^{\eta/2}) + (h_7^{(0)} + h_8^{(0)}\theta + h_9^{(0)}\theta^2 + h_{10,0}^{(0)}\theta^3)\theta^3 2^{\eta/2}. \quad (6)
\end{aligned}$$

Steps 9-14 perform the addition of the polynomials $(r_0^{(0)} + r_1^{(0)}\theta + \dots + r_7^{(0)}\theta^7)$ and $(2^{\eta/2}h_{10,1}^{(0)}\theta^3 + h_{11}^{(0)}\theta^4 + h_{12}^{(0)}\theta^5 + h_{13}^{(0)}\theta^6)$ to produce the polynomial $(s_0^{(0)} + s_1^{(0)}\theta + \dots + s_7^{(0)}\theta^7)$. Hence, from (6) we write

$$\begin{aligned}
h^{(0)}(\theta) &\equiv \underbrace{(s_0^{(0)} + s_1^{(0)}\theta + \dots + s_7^{(0)}\theta^7)}_{\text{through Steps 9-14}} + (h_{10,1}^{(0)} + h_{11}^{(0)}2^{\eta/2} + h_{12}^{(0)}\theta 2^{\eta/2} + h_{13}^{(0)}\theta^2 2^{\eta/2}) + \\
&\quad (h_7^{(0)} + h_8^{(0)}\theta + h_9^{(0)}\theta^2 + h_{10,0}^{(0)}\theta^3)\theta^3 2^{\eta/2}, \quad (7)
\end{aligned}$$

where $0 \leq s_0^{(0)}, s_1^{(0)}, \dots, s_6^{(0)} < 2^\eta$, and $0 \leq s_7^{(0)} \leq 2$. Using the definitions of (5) we can be further write (7) as

$$\begin{aligned}
h^{(0)}(\theta) &\equiv (s_0^{(0)} + s_1^{(0)}\theta + \dots + s_7^{(0)}\theta^7) + h_{10,1}^{(0)} + (h_{11,0}^{(0)} + h_{11,1}^{(0)}2^{\eta/2})2^{\eta/2} + (h_{12,0}^{(0)} + h_{12,1}^{(0)}2^{\eta/2})\theta 2^{\eta/2} + \\
&\quad (h_{13,0}^{(0)} + h_{13,1}^{(0)}2^{\eta/2})\theta^2 2^{\eta/2} + (h_{7,0}^{(0)} + h_{7,1}^{(0)}2^{\eta/2})\theta^3 2^{\eta/2} + (h_{8,0}^{(0)} + h_{8,1}^{(0)}2^{\eta/2})\theta^4 2^{\eta/2} + \\
&\quad (h_{9,0}^{(0)} + h_{9,1}^{(0)}2^{\eta/2})\theta^5 2^{\eta/2} + h_{10,0}^{(0)}\theta^6 2^{\eta/2}, \\
&= (s_0^{(0)} + s_1^{(0)}\theta + \dots + s_7^{(0)}\theta^7) + (h_{10,1}^{(0)} + h_{11,0}^{(0)}2^{\eta/2}) + (h_{11,1}^{(0)} + h_{12,0}^{(0)}2^{\eta/2})\theta + \\
&\quad (h_{12,1}^{(0)} + h_{13,0}^{(0)}2^{\eta/2})\theta^2 + (h_{13,1}^{(0)} + h_{7,0}^{(0)}2^{\eta/2})\theta^3 + (h_{7,1}^{(0)} + (h_{8,0}^{(0)}2^{\eta/2})\theta^4 + \\
&\quad (h_{8,1}^{(0)} + h_{9,0}^{(0)}2^{\eta/2})\theta^5 + (h_{9,1}^{(0)} + h_{10,0}^{(0)}2^{\eta/2})\theta^6, \\
&= (s_0^{(0)} + s_1^{(0)}\theta + \dots + s_7^{(0)}\theta^7) + \underbrace{(t_0^{(0)} + t_1^{(0)}\theta + \dots + t_6^{(0)}\theta^6)}_{\text{through Steps 15-21}}. \quad (8)
\end{aligned}$$

Steps 22-26 add the two polynomials $(s_0^{(0)} + s_1^{(0)}\theta + \dots + s_7^{(0)}\theta^7)$ and $(t_0^{(0)} + t_1^{(0)}\theta + \dots + t_6^{(0)}\theta^6)$ limb-wise by forwarding the 1-bit carry, producing the polynomial $(h_0^{(1)} + h_1^{(1)}\theta + \dots + h_7^{(1)}\theta^7)$. Hence, from (8) we can write

$$\begin{aligned}
h^{(0)}(\theta) &\equiv (s_0^{(0)} + s_1^{(0)}\theta + \dots + s_7^{(0)}\theta^6) + (t_0^{(0)} + t_1^{(0)}\theta + \dots + t_6^{(0)}\theta^6), \\
&= \underbrace{(h_0^{(1)} + h_1^{(1)}\theta + \dots + h_7^{(1)}\theta^7)}_{\text{through Steps 22-26}} = h^{(1)}(\theta), \quad (9)
\end{aligned}$$

where $0 \leq h_0^{(1)}, h_1^{(1)}, \dots, h_6^{(1)} < 2^\eta$, and $0 \leq h_7^{(1)} < 2^2$. In the rest of the proof, we use the looser bound $h_7^{(1)} < 2^{16} = 2^{\eta/4}$. This does not cause any problem. The advantage is that, later we can refer to the subsequent part of the proof to argue about the correctness of the reduction of the quantity obtained after multiplying by the 16-bit curve constant.

Reduction from $h^{(1)}(\theta)$ to $h^{(2)}(\theta)$. Polynomial $h^{(1)}(\theta)$ can further be written as

$$\begin{aligned}
h^{(1)}(\theta) &\equiv h_0^{(1)} + h_1^{(1)}\theta + \dots + h_6^{(1)}\theta^6 + h_7^{(1)}(2^{\eta/2}\theta^3 + 1) \text{ [using (2)]}, \\
&= (h_0^{(1)} + h_1^{(1)}\theta + \dots + h_6^{(1)}\theta^6) + (h_7^{(1)} + 2^{\eta/2}h_7^{(1)}\theta^3). \quad (10)
\end{aligned}$$

Steps 27-33 add the polynomial $(2^{\eta/2}\theta^3 + 1)h_7^{(1)} = (h_7^{(1)} + 2^{\eta/2}h_7^{(1)}\theta^3)$ to the polynomial $(h_0^{(1)} + h_1^{(1)}\theta + \dots + h_6^{(1)}\theta^6)$, which produces $(h_0^{(2)} + h_1^{(2)}\theta + \dots + h_7^{(2)}\theta^7)$, where $0 \leq h_0^{(2)}, h_1^{(2)}, \dots, h_6^{(2)} < 2^\eta$, and $0 \leq h_7^{(2)} < 2$.

Hence, from (10) we write

$$\begin{aligned} h^{(1)}(\theta) &\equiv (h_0^{(1)} + h_1^{(1)}\theta + \cdots + h_6^{(1)}\theta^6) + (h_7^{(1)} + 2^{\eta/2}h_7^{(1)}\theta^3), \\ &= \underbrace{(h_0^{(2)} + h_1^{(2)}\theta + \cdots + h_7^{(2)}\theta^7)}_{\text{through Steps 27-33}} = h^{(2)}(\theta), \end{aligned} \quad (11)$$

where $0 \leq h_0^{(2)}, h_1^{(2)}, \dots, h_6^{(2)} < 2^\eta$, and $0 \leq h_7^{(2)} < 2$. Note that in Steps 27-33, the value of carry is at most 1. In Step 33, $h_7^{(2)} = 1$ if and only if $h_6^{(1)} = 2^\eta - 1$ and carry = 1 which implies $h_6^{(2)} = t \bmod 2^\eta = 2^\eta \bmod 2^\eta = 0$. Moving one step backward, in Step 32 the output carry is 1 if and only if the conditions $h_5^{(1)} = 2^\eta - 1$ and the input carry = 1 hold, which results in setting $h_5^{(2)}$ to 0. Moving another step backward, in Step 31, the output carry is 1 if and only if the conditions $h_4^{(1)} = 2^\eta - 1$ and the input carry = 1 hold, which results in setting $h_4^{(2)}$ to 0. Moving one more step backward, in Step 30, the output carry is 1 if and only if the conditions $h_3^{(1)} = 2^\eta - 1$ and the input carry = 1 hold, and so the value of $h_3^{(2)}$ is bounded above by $(2^\eta - 1 + 2^{\eta/4} \cdot 2^{\eta/2} + 1) \bmod 2^\eta = 2^{3\eta/4}$. Hence, if $h_7^{(2)} = 1$, the conditions

$$h_3^{(2)} < 2^{3\eta/4}, \quad h_4^{(2)} = h_5^{(2)} = h_6^{(2)} = 0. \quad (12)$$

have to hold.

Reduction from $h^{(2)}(\theta)$ to $h^{(3)}(\theta)$. Polynomial $h^{(3)}(\theta)$ can further be written as

$$\begin{aligned} h^{(2)}(\theta) &\equiv h_0^{(2)} + h_1^{(2)}\theta + \cdots + h_6^{(2)}\theta^6 + h_7^{(2)}(2^{\eta/2}\theta^3 + 1) \text{ [using (2)],} \\ &= (h_0^{(2)} + h_1^{(2)}\theta + \cdots + h_6^{(2)}\theta^6) + (h_7^{(2)} + 2^{\eta/2}h_7^{(2)}\theta^3). \end{aligned} \quad (13)$$

If $h_7^{(2)} = 0$, then after Steps 34-38 we get $h_j^{(3)} = h_j^{(2)}$, $j = 0, 1, \dots, 6$; else, if $h_7^{(2)} = 1$, then using (12) we can say that the reduction surely terminates by the addition in Step 37. Using the bound of $h_3^{(2)} < 2^{3\eta/4}$ from (12) the maximum possible value of $h_3^{(3)}$ through Step 37 is $2^{3\eta/4} + 2^{\eta/2} + 1 < 2^\eta$. This implies after Steps 34-38 $0 \leq h_j^{(3)} < 2^\eta$, $j = 0, 1, 2, 3$, and $h_4^{(3)} = h_5^{(3)} = h_6^{(3)} = 0$. Hence, in any case from (13) it follows that

$$\begin{aligned} h^{(2)}(\theta) &\equiv (h_0^{(2)} + h_1^{(2)}\theta + \cdots + h_6^{(2)}\theta^6) + (h_7^{(2)} + 2^{\eta/2}h_7^{(2)}\theta^3), \\ &= \underbrace{(h_0^{(3)} + h_1^{(3)}\theta + \cdots + h_6^{(3)}\theta^6)}_{\text{through Steps 34-38}} = h^{(3)}(\theta), \end{aligned} \quad (14)$$

where $0 \leq h_0^{(3)}, h_1^{(3)}, \dots, h_6^{(3)} < 2^\eta$. Also, by combining (9), (11) and (14) we have $h^{(3)}(\theta) \equiv h^{(0)}(\theta) \bmod p$, which proves the theorem. \square

3.2 Reduction from 8-Limb to 7-Limb

Integer addition of two field elements in \mathbb{F}_p will produce an 8-limb quantity, the eighth limb of which has a size of at most 1 bit. Multiplying a field element by a field constant will also produce an 8-limb quantity. Considering Curve448, the field constant with which a multiplication of a field element arises in the Montgomery ladder is $(A + 2)/4 = (156326 + 2)/4 = 39082 < 2^{16} = 2^{\eta/4}$. Hence, given an 8-limb quantity, the reduction to 7-limb can be performed as follows. Consider the 8-limb quantity to be $h^{(1)}(\theta)$ and apply the part of reduce448 which reduces $h^{(1)}(\theta)$ to $h^{(3)}(\theta)$. The correctness of the reduction is guaranteed by the part of the proof of Theorem 1 which argues the correctness of the reduction from $h^{(1)}(\theta)$ to $h^{(2)}(\theta)$ and from $h^{(2)}(\theta)$ to $h^{(3)}(\theta)$.

3.3 Subtraction

Let $f(\theta)$ and $g(\theta)$ be 7-limb quantities representing elements of \mathbb{F}_p . The requirement is to compute $(f(\theta) - g(\theta)) \bmod p$. Function sub448 of Algorithm 2 performs this computation. The description of sub448 uses the instruction sub which is defined as follows. Let x and y be 64-bit quantities and \mathbf{b}_0 be a bit. The instruction $\text{sub}(x, y, \mathbf{b}_0)$ produces as output the pair (z, \mathbf{b}_1) where z is a 64-bit quantity and \mathbf{b}_1 is a bit. The definitions of z and \mathbf{b}_1 are as follows.

$$z = \begin{cases} x - (y + \mathbf{b}_0) & \text{if } x \geq y + \mathbf{b}_0, \\ 2^{64} + x - (y + \mathbf{b}_0) & \text{if } x < y + \mathbf{b}_0; \end{cases} \quad (15)$$

$$\mathbf{b}_1 = \begin{cases} 0 & \text{if } x \geq y + \mathbf{b}_0, \\ 1 & \text{if } x < y + \mathbf{b}_0. \end{cases} \quad (16)$$

The assembly instruction **sub** can be used to implement $\text{sub}(x, y, 0)$ while the assembly instruction **sbb** can be used to implement the more general $\text{sub}(x, y, \mathbf{b}_0)$.

Algorithm 2 Subtraction in \mathbb{F}_p .

```

1: function sub448( $(f(\theta), g(\theta))$ )
2: input: 7-limb quantities  $f(\theta)$  and  $g(\theta)$  such that  $0 \leq f_i, g_j < 2^{64}$  for  $i, j = 0, 1, \dots, 6$ .
3: output:  $h^{(2)}(\theta) = h_0^{(2)} + h_1^{(2)}\theta + \dots + h_6^{(2)}\theta^6$  such that  $0 \leq h_i^{(2)} < 2^{64}$  for  $i = 0, 1, \dots, 6$  and
    $h^{(2)}(\theta) \equiv (f(\theta) - g(\theta)) \pmod{p}$ .

4:  $\mathbf{b} \leftarrow 0$ 
5: for  $i \leftarrow 0$  to 6 do
6:    $(h_i^{(0)}, \mathbf{b}) \leftarrow \text{sub}(f_i, g_i, \mathbf{b})$ 
7: end for

8:  $\mathfrak{d} \leftarrow \mathbf{b}; \mathfrak{d}' \leftarrow \mathfrak{d} \ll 32$ 
9:  $\mathbf{b} \leftarrow 0$ 
10:  $(h_0^{(1)}, \mathbf{b}) \leftarrow \text{sub}(h_0^{(0)}, \mathfrak{d}, \mathbf{b})$ 
11:  $(h_1^{(1)}, \mathbf{b}) \leftarrow \text{sub}(h_1^{(0)}, 0, \mathbf{b})$ 
12:  $(h_2^{(1)}, \mathbf{b}) \leftarrow \text{sub}(h_2^{(0)}, 0, \mathbf{b})$ 
13:  $(h_3^{(1)}, \mathbf{b}) \leftarrow \text{sub}(h_3^{(0)}, \mathfrak{d}', \mathbf{b})$ 
14:  $(h_4^{(1)}, \mathbf{b}) \leftarrow \text{sub}(h_4^{(0)}, 0, \mathbf{b})$ 
15:  $(h_5^{(1)}, \mathbf{b}) \leftarrow \text{sub}(h_5^{(0)}, 0, \mathbf{b})$ 
16:  $(h_6^{(1)}, \mathbf{b}) \leftarrow \text{sub}(h_6^{(0)}, 0, \mathbf{b})$ 

17:  $\mathfrak{d} \leftarrow \mathbf{b}; \mathfrak{d}' \leftarrow \mathfrak{d} \ll 32$ 
18:  $\mathbf{b} \leftarrow 0$ 
19:  $(h_0^{(2)}, \mathbf{b}) \leftarrow \text{sub}(h_0^{(1)}, \mathfrak{d}, \mathbf{b})$ 
20:  $(h_1^{(2)}, \mathbf{b}) \leftarrow \text{sub}(h_1^{(1)}, 0, \mathbf{b})$ 
21:  $(h_2^{(2)}, \mathbf{b}) \leftarrow \text{sub}(h_2^{(1)}, 0, \mathbf{b})$ 
22:  $(h_3^{(2)}, \mathbf{b}) \leftarrow \text{sub}(h_3^{(1)}, \mathfrak{d}', \mathbf{b})$ 
23:  $h_4^{(2)} \leftarrow h_4^{(1)}; h_5^{(2)} \leftarrow h_5^{(1)}; h_6^{(2)} \leftarrow h_6^{(1)}$ 

24: return  $h^{(2)}(\theta) = h_0^{(2)} + h_1^{(2)}\theta + \dots + h_6^{(2)}\theta^6$ 

25: end function.

```

The correctness of **sub448** is stated in the following theorem.

Theorem 2. *The output $h^{(2)}(\theta) = h_0^{(2)} + h_1^{(2)}\theta + \dots + h_6^{(2)}\theta^6$ of **sub448** satisfies $0 \leq h_i^{(2)} < 2^{64}$ for $i = 0, 1, \dots, 6$ and $h^{(2)}(\theta) \equiv (f(\theta) - g(\theta)) \pmod{p}$.*

Proof. The limbs $h_i^{(2)}$, $i = 0, 1, \dots, 6$ are obtained as the first components of the outputs of some invocations of the **sub** instruction. Consequently, it follows that all of these are 64-bit quantities. This settles the point about the bounds on these limbs. So, we have to argue two things. First, $h^{(2)}(\theta) = (f(\theta) - g(\theta)) \pmod{p}$ and second that the procedure terminates without any overflow. The congruency argument is obtained from the following observations.

1. Let $\delta = 2^{224} + 1$. Steps 8-16 of **sub448** correspond to the subtraction of δ from the integer represented by $h^{(0)}(\theta)$. Similarly, Steps 17-22 correspond to the subtraction of δ from the integer represented by $h^{(1)}(\theta)$.
2. Suppose $f(\theta) \geq g(\theta)$ (as integers). Then, after Step 7, we have $h^{(0)}(\theta) = f(\theta) - g(\theta)$ and $\mathbf{b} = 0$. As a consequence of $\mathbf{b} = 0$ at Step 7, it follows that $h^{(0)}(\theta) = h^{(1)}(\theta) = h^{(2)}(\theta)$ establishing the result for this particular case.
3. In view of the previous point, assume $f(\theta) < g(\theta)$. In this case, after Step 7, we have that $h^{(0)}$ represents the integer $2^{448} + f(\theta) - g(\theta)$ and $\mathbf{b} = 1$. Steps 10-16 subtract δ from $h^{(0)}(\theta) = 2^{448} + f(\theta) - g(\theta)$.
 - (a) If $h^{(0)}(\theta) \geq \delta$, then after Step 16, $h^{(1)}(\theta)$ represents the integer $h^{(0)}(\theta) - \delta = 2^{448} + f(\theta) - g(\theta) - \delta = p + f(\theta) - g(\theta) \equiv (f(\theta) - g(\theta)) \pmod{p}$ and $\mathbf{b} = 0$. As a consequence of $\mathbf{b} = 0$ at Step 16, it follows that $h^{(2)}(\theta) = h^{(1)}(\theta)$ establishing the result for this case.

- (b) If $h^{(0)}(\theta) < \delta$, then after Step 16, $h^{(1)}(\theta)$ represents the integer $2^{448} + h^{(0)}(\theta) - \delta = 2^{448} + 2^{448} + f(\theta) - g(\theta) - \delta = 2^{448} + p + f(\theta) - g(\theta)$ and $\mathbf{b} = 1$. Steps 19-22 subtract δ from $h^{(1)}(\theta) = 2^{448} + p + f(\theta) - g(\theta)$ to obtain $h^{(2)}(\theta) = h^{(1)}(\theta) - \delta = 2^{448} + p + f(\theta) - g(\theta) - \delta = 2p + f(\theta) - g(\theta) \equiv (f(\theta) - g(\theta)) \pmod{p}$.

It only remains to argue that \mathbf{b} produced by the `sub` instruction in Step 22 is necessarily 0. If the value of \mathbf{b} in the input of `sub` in Step 22 is 0, then of course, the value of \mathbf{b} produced by this `sub` call is also 0. So, suppose that the value of \mathbf{b} in the input of `sub` in Step 22 (which is the output value of \mathbf{b} in Step 21) is 1. This implies that the value of \mathfrak{v}' in Step 22 is 2^{32} , since otherwise if, $\mathfrak{v}' = 0$ then output of \mathbf{b} in Step 21 has to be 0, which contradicts the assumption that the output value of \mathbf{b} in Step 21 is 1. Then the value of \mathbf{b} produced by the `sub` call in Step 22 is 0, if and only if $h_3^{(1)} \geq 2^{32} + 1$. The value of \mathbf{b} in the input of `sub` in Step 22 is 1, only if the value of \mathbf{b} produced by the `sub` call in Step 16 is 1. Arguing backwards, the value of \mathbf{b} produced by the `sub` call in Step 13 must be 1. The input to the `sub` call in Step 13 is $(h_3^{(0)}, \mathfrak{v}', \mathbf{b})$ and so the value of \mathbf{b} produced by this `sub` call is 1 if and only if the value of \mathbf{b} in the input to this `sub` call is 1 and $0 \leq h_3^{(0)} < 2^{32} + 1$. This implies $2^{64} - 2^{32} - 1 \leq 2^{64} + h_3^{(0)} - 2^{32} - 1 < 2^{64}$, or, $2^{64} - 1 - 2^{32} \leq h_3^{(1)} < 2^{64}$, from which it follows that $h_3^{(1)} \geq 2^{32} + 1$, as required. \square

4 Implementation and Timings

We present two 7-limb 64-bit implementations for shared secret computation phase of the X448 function. We term these implementations as `mxaa`- and `maax`-type implementations. The implementations based on the instructions `mulx`, `add`, `adc` are collectively termed as `mxaa`, and the implementations based on the instructions `mulx`, `adcx`, `adox` are collectively termed as `maax`. All the implementations are based on 64-bit assembly instructions targeting the Intel architectures. The `mxaa` type implementations are supported across a wide range of Intel processors. The `maax` type implementations are supported on modern Intel processors such as Skylake, but are not supported on previous generation processors such as Haswell.

Implementation of X448 requires implementation of field arithmetic over \mathbb{F}_p . Field multiplication and squaring are done in two steps. The first step multiplies two 7-limb field elements (considered as integers) to obtain a 14-limb integer. The second step reduces the 14-limb integer to a 7-limb integer. For the reduction, we have used Function `reduce448` while for the integer multiplication we have used the algorithms given in [3]. Implementations of field addition, subtraction and multiplication by the curve constant are as described in Sections 3.2 and 3.3. Overall, the implementation of X448 requires an implementation of the Montgomery ladder. We have made a careful constant-time assembly implementation of the Montgomery ladder. A major goal of the implementation has been to make efficient use of the available registers so that a minimal number of load/store instructions are required. Below, we provide timing results for the new implementations. The timing experiments were carried out on single cores of the Haswell and Skylake processors. The turbo-boost and hyper-threading features were turned off while measuring the cpu-cycles.

Platform specifications. The specifications of the hardware and software tools used in our software implementations are given below.

Haswell: Intel[®]Core[™] i7-4790 4-core CPU 3.60 Ghz. The OS was 64-bit Ubuntu 14.04 LTS and the source code was compiled using GCC version 7.3.0.

Skylake: Intel[®]Core[™] i7-6500U 2-core CPU @ 2.50GHz. The OS was 64-bit Ubuntu 14.04 LTS and the source code was compiled using GCC version 7.3.0.

Operation	Haswell	Skylake	Implementation	Implementation Type
Shared secret	732013	587389	[4]	<code>mxaa</code> , inline assembly
	-	530984	[4]	<code>maax</code> , inline assembly
	719217	461379	this work	<code>mxaa</code> , assembly
	-	434831	this work	<code>maax</code> , assembly

Table 1: CPU-cycle counts on Haswell and Skylake processors for shared secret computation on Curve448.

Timings in the form of cpu-cycles are provided in Table 1. For comparison we have considered the timings of the most efficient (to the best of our knowledge) publicly available 64-bit implementation of Curve448, which is the software implementation corresponding to the work [4]¹. We downloaded the mentioned software and measured the cpu-cycles on the same platforms on which we have measured the cpu-cycles of our implementations. This has been done to keep the comparisons consistent. We summarize the following observations from the timings of Table 1.

- On Skylake, the new implementations are substantially better than the the previous implementations. For `maax`-type implementation a speed-up of about 18% is obtained, while for `mxaa` type implementations, a speed-up of about 22% is obtained.
- For Haswell, the new `mxaa`-type implementation is better than the previous implementation by about 13K cpu-cycles. While this is an improvement, it is not as substantial an improvement as has been achieved in Skylake.

While the reduction algorithms that we have described avoid certain redundant operations performed by the code corresponding to [4], and consequently, do contribute to the speed improvement, it is not the only reason for the speed-up. A major reason for the speed improvement is a very careful assembly implementation making judicious use of the available registers so that the number of load/store operations is minimal.

In addition to the shared secret computation, we have also developed assembly code for the key generation phase of Curve448. The corresponding timings are reported in Table 2.

Operation	Haswell	Skylake	Implementation	Implementation Type
Key generation	653035	427058	this work	<code>mxaa</code> , assembly
	-	396583	this work	<code>maax</code> , assembly

Table 2: CPU-cycle counts on Haswell and Skylake processors for key generation on Curve448.

5 Conclusion

In this work we have presented reduction algorithms and their proofs of correctness required for computation in the field \mathbb{F}_p where $p = 2^{448} - 2^{224} - 1$. Based on these algorithms and other previously known techniques, we have made efficient 64-bit assembly implementations of the X448 function of Curve448 leading to new speed records for 64-bit implementations. While our work has concentrated entirely on the prime $2^{448} - 2^{224} - 1$, we note that the ideas involved can be applied to other primes having a similar form such as the prime $2^{480} - 2^{240} - 1$.

References

- [1] Mike Hamburg. Ed448-goldilocks, a new elliptic curve. *IACR Cryptology ePrint Archive*, 2015:625, 2015.
- [2] Adam Langley and Mike Hamburg. Elliptic curves for security. Internet Research Task Force (IRTF), Request for Comments: 7748, <https://tools.ietf.org/html/rfc7748>, 2016. Accessed on 16 September, 2019.
- [3] Kaushik Nath and Palash Sarkar. Efficient Arithmetic in (Pseudo-)Mersenne Prime Order Fields. *IACR Cryptology ePrint Archive*, 2018:985, 2018.
- [4] Thomaz Oliveira, Julio López Hernandez, Hüseyin Hisil, Armando Faz-Hernández, and Francisco Rodríguez-Henríquez. How to (pre-)compute a ladder - improving the performance of X25519 and X448. In *Selected Areas in Cryptography - SAC 2017 - 24th International Conference, Ottawa, ON, Canada, August 16-18, 2017, Revised Selected Papers*, pages 172–191, 2017.
- [5] E. Ozturk, J. Guilford, and V. Gopal. Large integer squaring on Intel architecture processors, intel white paper. <https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/large-integer-squaring-ia-paper.pdf>, 2013.
- [6] E. Ozturk, J. Guilford, V. Gopal, and W. Feghali. New instructions supporting large integer arithmetic on Intel architecture processors, intel white paper. <https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/ia-large-integer-arithmetic-paper.pdf>, 2012.
- [7] Version 1.3 TLS Protocol. RFC 8446. https://datatracker.ietf.org/doc/rfc8446/?include_text=1, 2018. Accessed on 16 September, 2019.

¹Program code from https://github.com/armfazh/rfc7748_precomputed was accessed on June 25, 2020.

A Comparison to [4]

Page 17 of [4], provides an abstraction of the reduction strategy used to convert the integer z represented by the 14-limb polynomial $h^{(0)}(\theta)$ to a reduced integer in \mathbb{F}_p , which are given as below.

$$z \leftarrow (z \bmod 2^{672}) + (2^{448} + 2^{224})\lfloor z/2^{672} \rfloor, \quad (17)$$

$$z \leftarrow (z \bmod 2^{448}) + (2^{224} + 1)\lfloor z/2^{448} \rfloor, \quad (18)$$

$$z \leftarrow (z \bmod 2^{448}) + (2^{224} + 1)\lfloor z/2^{448} \rfloor. \quad (19)$$

The first two steps (17) and (18) convert the 14-limb input quantity $h^{(0)}(\theta)$ to the 8-limb quantity $h^{(1)}(\theta)$, such that the size of the eighth limb of $h^{(1)}(\theta)$ is at most 2 bits long. Step (17) reduces $h^{(0)}(\theta)$ to an 11-limb polynomial, say $f(\theta) = f_0 + f_1\theta + \dots + f_{10}\theta^{10}$ and Step (18) reduces $f(\theta)$ to $h^{(1)}(\theta)$. Step (19) further reduces $h^{(1)}(\theta)$ to $h^{(2)}(\theta)$ which is also an 8-limb quantity whose final limb is at most 1 bit long. The final reduction round that converts $h^{(2)}(\theta)$ to $h^{(3)}(\theta)$, which is a 7-limb quantity is missing. As a result, the reduction strategy suggested in [4] is incomplete.

Reduction Algorithms Used in the Code Accompanying [4]. We have studied the latest version² of the implementation corresponding to [4]. The reduction algorithm used in this code is different from the strategy outlined in the paper. While the strategy suggested in the paper is incomplete, the algorithms implemented in the code are indeed complete. They, however, perform some redundant operations. Recall that in the final round which reduces $h^{(2)}(\theta)$ to $h^{(3)}(\theta)$, Algorithm 1 proceeds only up to the fourth limb. Theorem 1 shows that this is sufficient. The present version of the code corresponding to [4] performs the additions till the last limb. The three extra additions in the final round are redundant. Similar redundancies are also present in addition, subtraction and multiplication by the field constant.

An efficiency issue while reducing $h^{(0)}(\theta)$ to $h^{(1)}(\theta)$. Define $\phi = \theta^3 2^{n/2} = 2^{224}$, which implies $\phi^2 \equiv \phi + 1$ using (2). We can also view $h^{(0)}(\theta)$ as an equivalent polynomial $\mathfrak{h}^{(0)}$ in base ϕ defined as $\mathfrak{h}^{(0)}(\phi) = a + b\phi + c\phi^2 + d\phi^3$, where $a, b, c, d < \phi$. Under such a consideration, the reduction from $h^{(0)}(\theta)$ to $h^{(1)}(\theta)$ of Algorithm 1 can be described as a reduction from $\mathfrak{h}^{(0)}(\phi)$ to $\mathfrak{h}^{(1)}(\phi)$ through the following steps.

$$\begin{aligned} \mathfrak{h}^{(0)}(\phi) &= a + b\phi + c\phi^2 + d\phi^3 \\ &= (a + b\phi) + (c + d\phi)\phi^2 \\ &\equiv (a + b\phi) + (c + d\phi)(\phi + 1) && \text{mod } p \\ &= (a + b\phi) + (c + d\phi) + c\phi + d\phi^2 && \text{mod } p \\ &\equiv (a + b\phi) + (c + d\phi) + c\phi + d\phi + d && \text{mod } p \\ &= (a + b\phi) + (c + d\phi) + d\phi + (d + c\phi) && \text{mod } p \\ &= \mathfrak{h}^{(1)}(\phi) && \text{mod } p. \end{aligned}$$

Algorithm 1 computes $\mathfrak{h}^{(1)}(\phi)$ from $\mathfrak{h}^{(0)}(\phi)$ by first adding $(c + d\phi)$ to $(a + b\phi)$ through Steps 4-8. Then it adds $d\phi$ to the result through Steps 9-14. Finally, $(d + c\phi)$ is computed through Steps 15-21 and added to the previous result through Steps 22-26 to produce $\mathfrak{h}^{(1)}(\phi)$. The x86 architecture has 15 64-bit registers (keeping aside the stack pointer register `rsp`) to work with. To store the value of the product $\mathfrak{h}^{(0)}(\phi) = (a + b\phi + c\phi^2 + d\phi^3)$ we need 14 64-bit registers. So, the polynomial $(a + b\phi)$ is stored in 7 registers and $(c + d\phi)$ is stored in another 7. We need d to compute $(d + c\phi)$, so it is better to keep d undisturbed until we compute the value of $(d + c\phi)$. We first add the register values of $(c + d\phi)$ to the registers of $(a + b\phi)$. The register values of $(a + b\phi)$ gets updated to produce the temporary sum and the register values of $(c + d\phi)$ remain unchanged. After that, we only copy the middle limb of $(c + d\phi)$ to a temporary register and mask off its lower 32 bits to achieve the first 32 bits of d . The remaining 192 bits are easily obtained from the last three register values of $(c + d\phi)$ without any extra operations. Now, we have d and add it to the previous sum to get a modified sum. Finally, we compute $(d + c\phi)$ from $(c + d\phi)$ through `shrd` instructions which works in a circular manner and add the obtained value to the previous sum to get the final value.

An alternative way to compute $\mathfrak{h}^{(1)}(\phi)$ would be to first add $(c + 2d\phi)$ to $(a + b\phi)$ and then add $(d + c\phi)$ to the result. The code corresponding to [4] uses this method. However, depending on the number of available 64-bit registers in the x86 architectures, this is going to be less efficient. This is

²Program code from https://github.com/armfazh/rfc7748_precomputed/blob/master/src/fp448_x64.c was accessed on June 25, 2020.

because, computing $2d$ from d will necessitate extra operations to back up d for computing $(d + c\phi)$ later on. As a result, the number of load/stores will increase.