

# Efficient and Secure Algorithms for GLV-Based Scalar Multiplication and their Implementation on GLV-GLS Curves

(*or Keep Calm and Stay with One (and  $p > 3$ )*)

Armando Faz-Hernández<sup>1</sup>, Patrick Longa<sup>2</sup>, and Ana H. Sánchez<sup>1</sup>

<sup>1</sup> Computer Science Department, CINVESTAV-IPN, México

{armfaz, asanchez}@computacion.cs.cinvestav.mx

<sup>2</sup> Microsoft Research,

One Microsoft Way, Redmond, WA 98052, USA

plonga@microsoft.com

**Abstract.** We propose efficient algorithms and formulas that improve the performance of *side-channel protected* elliptic curve computations, with special focus on scalar multiplication exploiting the Gallant-Lambert-Vanstone (CRYPTO 2001) and Galbraith-Lin-Scott (EUROCRYPT 2009) methods. Firstly, by adapting Feng et al.’s recoding to the GLV setting, we derive new regular algorithms for variable-base scalar multiplication that offer protection against simple side-channel and timing attacks. Secondly, we propose an efficient algorithm for fixed-base scalar multiplication that is also protected against side-channel attacks by combining Feng et al.’s recoding with Lim-Lee’s comb method. Thirdly, we propose an efficient technique that interleaves ARM-based and NEON-based multiprecision operations over an extension field, as typically found on GLS curves and pairing computations, to improve performance on modern ARM processors. Finally, we showcase the efficiency of the proposed techniques by implementing a state-of-the-art GLV-GLS curve in twisted Edwards form defined over  $\mathbb{F}_{p^2}$ , which supports a four dimensional decomposition of the scalar and is fully protected against timing attacks. Analysis and performance results are reported for modern x64 and ARM processors. For instance, using a precomputed table of only 512 bytes, we compute a variable-base scalar multiplication in 92,000 and 244,000 cycles on an Intel Ivy Bridge and an ARM Cortex-A15 processor (respect.); using an *off-line* precomputed table of 6KB, we compute a fixed-base scalar multiplication in 53,000 and 116,000 cycles (respect.); and using a precomputed table of 3KB, we compute a double scalar multiplication in 118,000 and 285,000 cycles (respect.). All of these numbers and the proposed techniques represent a significant improvement of the state-of-the-art performance of elliptic curve computations. Most notably, our techniques allow us to reduce the cost of adding protection against timing attacks in the computation of GLV-based variable-base scalar multiplication to below 10%.

**Keywords.** Elliptic curves, scalar multiplication, side-channel protection, GLV method, GLS method, GLV-GLS curve, x64 processor, ARM processor.

## 1 Introduction

Let  $P$  be a point of prime order  $r$  on an elliptic curve over  $\mathbb{F}_p$  containing a degree-2 endomorphism  $\phi$ . The Gallant-Lambert-Vanstone (GLV) method computes the scalar multiplication  $kP$  as  $k_1P + k_2\phi(P)$  [14]. If  $k_1, k_2$  have approximately half the bitlength of the original scalar  $k$ , one should expect an elimination of half the number of doublings by using the Straus-Shamir simultaneous multi-scalar multiplication technique. Thus, the method is especially useful for speeding up the case in which the base point  $P$  is variable, known as variable-base scalar multiplication. Later, Galbraith et al. [13] showed how to exploit the Frobenius endomorphism to enable the use of the GLV approach on a wider set of curves defined over the quadratic extension field  $\mathbb{F}_{p^2}$ . Since then, significant research has been performed to improve the performance [27, 21] and to explore the applicability to other settings [17, 32] or to higher dimensions on genus one curves [21, 28, 15] and genus two curves [7, 8, 15]. Unfortunately, most of the work and comparisons with other approaches

have been carried out with *unprotected* algorithms and implementations. In fact, little effort has been done to investigate methods for protecting GLV-based implementations against side-channel attacks. Just recently, Longa and Sica [28] used the regular windowed recoding by Okeya and Takagi [31] in combination with interleaving [14, 30] to make a four-dimensional implementation constant time. However, the use of this standard approach in the GLV paradigm incurs in a high cost in terms of storage and computing performance because of the high number of required precomputations. This issue worsens for higher dimensions [8].

In this work, we propose a new signed representation, called GLV-based Sign-Aligned Column (GLV-SAC), that gives rise to a new method for scalar multiplication using the GLV method. We depart from the traditional approach based on interleaving or joint sparse form and adapt the recoding by Feng et al. [10], originally intended for standard comb-based fixed-base scalar multiplication, to the computation of GLV-based variable-base scalar multiplication. The method supports a regular execution and thus provides a first layer of protection against some simple side-channel (SSCA) attacks such as simple power analysis (SPA) [24]. Moreover, it does not require dummy operations, making it resilient to safe-error attacks [38, 39], and can be used as a basis for constant-time implementations secure against timing attacks [23, 9, 2, 33]. In addition, we present different variants of the technique that are intended for different scenarios exploiting simple or complex GLV decompositions, and thus provide algorithms that have broad applicability to many settings using GLV, GLS, or a combination of both [13, 17, 27, 21, 28, 32, 7, 8, 15, 35]. In comparison with the best previous approaches, the method improves the computing performance, especially during the potentially expensive precomputation stage, and allows us to save *at least* half of the storage requirement for precomputed values without impacting performance. For instance, the method injects a 17% speedup in the overall computation and a 78% reduction in the memory consumption for a GLV-GLS curve using a 4-GLV decomposition (see §6). The savings in memory without impacting performance are especially relevant for the deployment of GLV-based implementations in constrained devices. Depending on the cost of endomorphisms, the improvement provided by the method is expected to increase for higher-degree decompositions.

Additionally to the standard variable-base scalar multiplication, there are two other core computations that are the basis of most curve-based protocols:  $kP$  with  $P$  known in advance (fixed-base scenario), and  $kP + lQ$  with  $P$  not known in advance and  $Q$  known in advance (variable/fixed-base double scalar scenario), where  $P$  and  $Q$  are points lying on an elliptic curve subgroup of prime order  $r$  and  $k, l$  are integers in  $[1, r - 1]$ . For example, fixed-base scalar multiplication is used during signature generation and variable/fixed-base double scalar multiplication is used during signature verification in the ECDSA protocol.

For the case of the fixed-base scenario, we present an optimized variant of the side-channel protected fixed-base comb method by Feng et al. [10] that exploits the original, multi-table Lim-Lee comb technique [25]. Our approach follows a similar idea proposed by Hamburg [16] using Hedabou et al.’s representation [19]. Our algorithm is generic (i.e., it does not exploit endomorphisms) and is as efficient as Hamburg’s algorithm during the on-line computation with a slightly cheaper off-line precomputation phase.

Processors based on the ARM architecture are widely used in modern smartphones and tablets due to their low power consumption. The ARM architecture comes equipped with 16 32-bit registers and an instruction set including 32-bit operations, which in most cases can be executed in one cycle. To boost performance in certain applications, some ARM processors include a powerful set of vector instructions known as NEON. This consists of a 128-bit Single Instruction Multiple Data (SIMD) engine that includes 16 128-bit registers. Recent research has exploited NEON to accelerate cryptographic operations [6, 16, 34]. On one hand, the interleaving of ARM and NEON instructions is a well-known technique (with increasing potential on modern processors) that can be exploited in cryptography; e.g., see [6]. On the other hand, the vectorized computation using NEON can be advantageously exploited to compute independent multiplications, as found in operations over  $\mathbb{F}_{p^2}$ ; e.g., see [34]. In this work, we take these optimizations further and propose a technique that interleaves ARM-based and NEON-based multiprecision operations, such as multiplication, squaring and modular reduction, in extension field operations in order to maximize the inherent parallelism and hide the execution latency. The technique is especially relevant for implementing the quadratic extension field layer, as found in GLS curves [13] and pairing computations [1]. For instance, it injects a significant speedup in the range 17%-34% in the scalar multiplication execution on a GLV-GLS curve (see §5 and §6).

To demonstrate the efficiency of our techniques, we implement the state-of-the-art twisted Edwards GLV-GLS curve over  $\mathbb{F}_{p^2}$  with  $p = 2^{127} - 5997$  recently proposed by Longa and Sica [28]. This curve, referred to as `Ted127-glv4`, supports a 4-GLV decomposition. Moreover, we also present efficient algorithms for implementing field and quadratic extension field operations targeting our 127-bit prime on x64 and ARM platforms. We combine and exploit incomplete reduction [37] and lazy reduction [36], expanding techniques by [27]. These optimized operations are then applied to state-of-the-art twisted Edwards formulas [3, 20] to speed up computations in the setting of curves over  $\mathbb{F}_{p^2}$ . Our implementations of the *three* core scalar multiplication scenarios, namely, variable-base, fixed-base and double scalar, target modern x64 and ARM processors and include full protection against timing attacks.

Most notably, we show that the proposed algorithms and formulas reduce dramatically the cost of adding protection against timing attacks and the storage requirement for precomputations, and set a new speed record for protected software. For instance, a protected variable-based elliptic curve scalar multiplication on curve `Ted127-glv4` runs in 96,000 cycles on an Intel Sandy Bridge machine, using only 512 bytes of memory for precomputed values. This is 30% faster, using almost 1/5 of the storage, than the state-of-the-art implementation reported by Longa and Sica [28] that computes the same operation in 137,000 cycles using 2.25KB of memory for precomputations. Moreover, this result is only 5% slower, using 1/2 of the storage, than the state-of-the-art *unprotected* computation in [28], which runs in 91,000 cycles using 1KB of memory. This not only represents a new speed record for protected software but also marks the first time that a constant-time variable-base scalar multiplication is performed under 100K cycles on an Intel processor. Similar results are obtained for fixed-base and double scalar multiplication, and for ARM processors exploiting the technique that interleaves NEON and ARM-based operations (see §6 for full benchmark results).

This paper is organized as follows. In §2, we give some preliminaries about the GLV and GLS methods, side-channel attacks and the protected comb methods by Feng et al. [10, 11] and Hedabout et al. [19]. In §3, we present the new GLV-based representation, its variants and the corresponding scalar multiplication method. In §4, we describe the new algorithm for fixed-base scalar multiplication. We describe the implementation of curve `Ted127-glv4` as well as optimized algorithms for field, extension field and point operations targeting x64 and ARM platforms in §5. In this section, we also discuss the interleaving technique for ARM. Finally, in §6, we perform an analysis of the proposed methods and present benchmark results of the core scalar multiplication scenarios on several x64 and ARM processors.

## 2 Preliminaries

### 2.1 The GLV and GLS Methods

In this section, we briefly describe the GLV and GLS methods in a generic,  $m$  dimensional framework. Let  $C$  be a curve defined over a finite field  $\mathbb{F}_p$  equipped with an efficiently computable endomorphism  $\phi$ . The GLV method to compute scalar multiplication [14] consists of first decomposing the scalar  $k$  into sub-scalars  $k_i$  for  $0 \leq i < m$  and then computing  $\sum_{i=0}^{m-1} k_i D_i$  using the Straus-Shamir trick for simultaneous multi-scalar multiplication, where  $D_0$  is the input divisor from the divisor class group of the curve and  $D_i = \phi^i(D_0)$ . If all of the sub-scalars have approximately the same bitlength, the number of required doublings is reduced to approximately  $\log_2 r/m$ , where  $r$  is the prime order of the curve subgroup. Special curves equipped with endomorphisms which are different to the Frobenius endomorphism are known as GLV curves.

The GLS method [13, 12] lifts the restriction to special curves and exploits an endomorphism  $\psi$  arising from the  $p$ -power Frobenius endomorphism on a wider set of curves  $C'$  defined over an extension field  $\mathbb{F}_{p^k}$  that are  $\mathbb{F}_{p^n}$ -isogenous to curves  $C/\mathbb{F}_p$ , where  $k|n$ . Equipped with  $\psi$  to perform the scalar decomposition, one then proceeds to apply the GLV method as above. More complex decompositions arise by applying the GLS paradigm to GLV curves (a.k.a. GLV-GLS curves [13, 28]).

These techniques have received lots of attention recently, given their significant impact in the performance of curve-based systems. Longa and Gebotys [27] report efficient implementations of GLS curves over  $\mathbb{F}_{p^2}$  using two-dimensional decompositions. In [21], Hu, Longa and Xu explore a GLV-GLS curve over  $\mathbb{F}_{p^2}$  supporting a four-dimensional decomposition. In [7], Bos et al. study two and four-dimensional decompositions on genus

2 curves over  $\mathbb{F}_p$ . Bos et al. [8] explore the combined GLV-GLS approach over genus 2 curves defined over  $\mathbb{F}_{p^2}$ , which supports an 8-GLV decomposition. In the case of binary GLS elliptic curves, Oliveira et al. [32] report the implementation of a curve exploiting the 2-GLV method. More recently, Guillevic and Ionica [15] show how to exploit the 4-GLV method on certain genus one curves defined over  $\mathbb{F}_{p^2}$  and genus two curves defined over  $\mathbb{F}_p$ ; and Smith [35] proposes a new family of elliptic curves that support 2-GLV decompositions.

From all the works above, only [28] and [32] include side-channel protection in their GLV-based implementations.

## 2.2 Side-Channel Attacks and Countermeasures

Side-channel attacks [23] exploit leakage information obtained from the physical implementation of a cryptosystem to get access to private key material. Examples of physical information that can be exploited are power, time, electromagnetic emanations, among others. In particular, much attention has been put on timing [23, 9] and simple power attacks (SPA) [24], given their broad applicability and relatively low costs to be realized in practice. Traditionally, the different attacks can also be distinguished by the number of traces that are exploited in the analysis: simple side-channel attacks (SSCA) require only one trace (or very few traces) to observe the leakage that directly reveals the secret bits, whereas differential side-channel attacks (DSCA) require many traces to perform a statistical analysis on the data. The feasibility of these attacks depends on the targeted application, but it is clear that SSCA attacks are feasible in a wider range of scenarios. In this work, we focus on methods that minimize the risk posed by timing attacks and SSCA attacks such as SPA.

In curve-based cryptosystems, the first step to achieve protection against these attacks is to use regular algorithms for performing scalar multiplication (other methods involve the use of unified formulas, but these are generally expensive). One efficient approach in this direction is to recode the scalar to a representation exhibiting a regular pattern. In particular, for the case of variable-base scalar multiplication, the regular windowed recoding proposed by Okeya and Takagi [31] and further analyzed by Joye and Tunstall [22] represents one of the most efficient alternatives. Nevertheless, in comparison with the standard width- $w$  non-adjacent form ( $w$ NAF) [18] used in unprotected implementations, the Okeya-Takagi recoding increases the nonzero density from  $1/(w+1)$  to  $1/(w-1)$ . In contrast, side-channel protected methods for scalar multiplication exploiting the GLV method have not been fully studied. Furthermore, we note that methods typically efficient in the standard case are not necessarily efficient in the GLV paradigm. For example, in [28], Longa and Sica apply the Okeya-Takagi recoding to protect scalar multiplication on a GLV-GLS curve using a four-dimensional GLV decomposition against timing attacks. The resulting protected implementation is about 30% more expensive than the unprotected version. In this work, we aim at reducing that gap, providing efficient methods that can be exploited to improve and protect GLV and GLS-based implementations.

The comb method [25] is an efficient approach for the case of fixed-base scalar multiplication. However, in its original form, the method is unprotected against SSCA and timing attacks. An efficient approach to achieve a regular execution is to recode the scalar using signed nonzero representations such as LSB-set [10], MSB-set [11] or SAB-set [19]. A key observation in this work is that the basic version of the fixed-base comb execution (i.e., without exploiting multiple tables) has several similarities with a GLV-based variable-base execution. So it is therefore natural to adapt these techniques to the GLV setting to achieve side-channel protection. In particular, the LSB-set representation is a good candidate, given that an analogue of this method in the GLV setting minimizes the cost of precomputation.

## 2.3 The Least Significant Bit - Set (LSB-Set) Representation and Variants

Feng, Zhu, Xu and Li [10] proposed a clever signed representation, called LSB-set, that is based on the equivalence  $1 \equiv 1\bar{1} \dots \bar{1}$  (assuming the notation  $-1 \equiv \bar{1}$ ). They used this representation to protect the comb method [25] in the computation of fixed-base scalar multiplication (we refer to this method as LSB-set comb scalar multiplication). Next, we briefly describe the LSB-set recoding and its application to fixed-base scalar multiplication. The reader is referred to [25] and [10] for complete details about the original comb method and the LSB-set comb method, respectively.

Let  $t$  be the bitlength of a given scalar  $k$ . Assume that  $k$  is partitioned in  $w$  consecutive parts of  $d = \lceil t/w \rceil$  bits each, padding  $k$  with  $(dw - t)$  zeros to the left. Let the updated binary representation of  $k$  be  $(k_{l-1}, k_{l-2}, \dots, k_0)$ , where  $l = dw$ . One can visualize the bits of  $k$  in matrix form by considering the  $w$  pieces as the rows and arranging them from top to bottom. The LSB-set recoding consists of first applying the transformation  $1 \mapsto 1\bar{1} \dots \bar{1}$  to the least significant  $d$  bits of the scalar (i.e., the first row in the matrix) and, then, converting every bit  $k_i$  in the remaining rows in such a way that output digits  $b_i$  for  $d \leq i \leq (l-1)$  are in the digit set  $\{0, b_i \bmod d\}$ . That is, digits in the same column are 0 or share the same sign. Then, for computing a comb fixed-base scalar multiplication, one scans the “digit-columns” in the matrix from left to right. Since every digit-column is nonzero by definition, the execution consists of a doubling-addition computation at every iteration, which provides protection against certain SCA attacks such as SPA.

There are other variants in the literature that have also been exploited for implementing protected comb fixed-base scalar multiplication. In [11], Feng, Zhu, Zhao and Li proposed the Most Significant Bit - Set (MSB-set) representation, which reduces slightly the cost of the LSB-set comb method for the case that  $w \mid t$ . The main difference with LSB-set resides in applying the transformation  $1 \mapsto 1\bar{1} \dots \bar{1}$  to the most significant  $d$  bits of the scalar, instead of the least significant portion. In [19], Hedabou, Pinel and Beneteau proposed a *full* signed nonzero representation, referred to as Signed All-Bit-Set (SAB-set), that uses the above transformation to recode the whole scalar. In this case, the cost of precomputation is expected to be slightly higher since no zeros are used in the representation. We remark that the algorithms presented in this work can be modified to use the MSB-set or SAB-set representations.

### 3 The GLV-Based Sign-Aligned Column (GLV-SAC) Representation

In this section, we introduce a variant of the LSB-set recoding that is amenable for the computation of side-channel protected variable-base scalar multiplication in the GLV setting. The new recoding is called GLV-Based Sign-Aligned Column (GLV-SAC). Also, we present a new method for GLV-based scalar multiplication exploiting the proposed representation.

In the following, we first discuss the GLV-SAC representation in a generic setting. In Section 3.2, we discuss variants that are expected to be more efficient when  $m = 2$  and  $m \geq 8$ . To simplify the descriptions, we assume in the remainder that we are working on an elliptic curve. The techniques and algorithms can be easily extended to other settings such as genus 2 curves.

Let  $\{k_0, k_1, \dots, k_j, \dots, k_{m-1}\}$  be a set of positive sub-scalars in the setting of GLV with dimension  $m$ . The basic idea of the new recoding is to have one of the sub-scalars of the  $m$ -GLV decomposition, say  $k_J$ , represented in signed nonzero form and acting as a “sign-aligner”. The latter means that  $k_J$  determines the sign of all the digits of remaining sub-scalars according to their relative position.

The GLV-SAC representation has the following properties

- (i) The length of the digit representation of every sub-scalar  $k_j$  is fixed and given by  $l = \lceil \log_2 r/m \rceil + 1$ , where  $r$  is the prime subgroup order.
- (ii) Exactly one sub-scalar, which should be odd, is expressed by a signed nonzero representation  $k_J = (b_{l-1}^J, \dots, b_0^J)$ , where all digits  $b_i^J \in \{1, -1\}$  for  $0 \leq i < l$ .
- (iii) All the sub-scalars  $k_j$ , with exception of  $k_J$  from (ii), are expressed by signed representations  $(b_{l-1}^j, \dots, b_0^j)$  such that  $b_i^j \in \{0, b_i^J\}$  for  $0 \leq i < l$ .

In the targeted setting, (i) and (ii) guarantee a constant-time execution regardless of the value of the scalar  $k$  and without having to appeal to masking for dealing with the identity element. Item (iii) allows us to reduce the size of the precomputed table by a factor of 2, while minimizing the cost of precomputation.

Note that we do not impose any restriction on which sub-scalar should be designated as  $k_J$ . In some settings, choosing any of the  $k_j$  (with the exception of the one corresponding to the base point, i.e.,  $k_0$ ) could lead to the same performance in the precomputation phase and be slightly faster than  $k_J = k_0$ , if one takes into consideration the use of mixed point additions. The condition that  $k_J$  should be odd enables the conversion of any integer to a full signed nonzero representation using the equivalence  $1 \equiv 1\bar{1} \dots \bar{1}$ . To deal

with this restriction during the scalar multiplication, we first convert the selected sub-scalar  $k_J$  to odd (if even), and then make the corresponding correction in the end (more details can be found in Section 3.1). Finally, the reader should note that the GLV-SAC representation, in the way we describe it above, assumes that the sub-scalars are all positive. This restriction is imposed in order to achieve the minimum length  $l = \lceil \log_2 r/m \rceil + 1$  in the representation. We lift this restriction in Section 3.3.

An efficient algorithm to recode the sub-scalars to GLV-SAC proceeds as follows. Assume that each sub-scalar  $k_j$  is padded with zeros to the left until reaching the fixed length  $l = \lceil \log_2 r/m \rceil + 1$ , where  $r$  is the prime order of the curve subgroup. After choosing a suitable  $k_J$  to act as the “sign-aligner”, the sub-scalar  $k_J$  is recoded to signed nonzero digits  $b_i^J$  using the equivalence  $1 \equiv 1\bar{1} \dots \bar{1}$ . Remaining sub-scalars are then recoded in such a way that output digits at position  $i$  are in the set  $\{0, b_i^J\}$ , i.e., nonzero digits at the same relative position share the same sign. This is shown as Algorithm 1.

---

**Algorithm 1** Protected Recoding Algorithm for the GLV-SAC Representation.

---

**Input:**  $m$   $l$ -bit positive integers  $k_j = (k_{i-1}^j, \dots, k_0^j)_2$  for  $0 \leq j < m$ , an odd “sign-aligner”  $k_J \in \{k_j\}^m$ , where  $l = \lceil \log_2 r/m \rceil + 1$ ,  $m$  is the GLV dimension and  $r$  is the prime group order.

**Output:**  $(b_{l-1}^j, \dots, b_0^j)_{\text{GLV-SAC}}$  for  $0 \leq j < m$ , where  $b_i^j \in \{1, -1\}$ , and  $b_i^j \in \{0, b_i^J\}$  for  $0 \leq j < m$  and  $j \neq J$ .

---

```

1:  $b_{l-1}^J = 1$ 
2: for  $i = 0$  to  $(l - 2)$  do
3:    $b_i^J = 2k_{i+1}^J - 1$ 
4: for  $j = 0$  to  $(m - 1), j \neq J$  do
5:   for  $i = 0$  to  $(l - 1)$  do
6:      $b_i^j = b_i^J \cdot k_0^j$ 
7:      $k_j = \lfloor k_j/2 \rfloor - \lfloor b_i^j/2 \rfloor$ 
8: return  $(b_{l-1}^j, \dots, b_0^j)_{\text{GLV-SAC}}$  for  $0 \leq j < m$ .
```

---

We highlight that, in contrast to [10, Alg. 4] and [11, Alg. 2], our recoding algorithm is simpler and exhibits a regular and constant time execution, making it resilient to timing attacks. Moreover, Algorithm 1 can be implemented very efficiently by exploiting the fact that the only purpose of the recoded digits from the sub-scalar  $k_J$  is, by definition, to determine the sign of their corresponding digit-columns (see details in Alg. 2 below). Since  $k_{i+1}^J = 0$  and  $k_{i+1}^J = 1$  indicate that the corresponding output digit-column  $i$  will be negative and positive, respectively, Step 3 of Algorithm 1 can be reduced to  $b_i^J = k_{i+1}^J$  by assuming the convention  $b_i^J = 0$  to indicate negative and  $b_i^J = 1$  to indicate positive, for  $0 \leq i < l$ . Following this convention, further efficient simplifications are possible for Steps 6 and 7.

### 3.1 GLV-Based Scalar Multiplication using GLV-SAC

We now present a new method for computing variable-base scalar multiplication using the GLV method and the GLV-SAC representation (see Algorithm 2). To simplify the description, we assume that  $k_0$  is fixed as the “sign-aligner”  $k_J$  (it is easy to modify the algorithm to set any other sub-scalar to  $k_J$ ). The basic idea is to arrange the sub-scalars, after being converted to their GLV-SAC representation, in matrix form from top to bottom, with sub-scalar  $k_J = k_0$  at the top, and then run a simultaneous multi-scalar multiplication execution scanning digit-columns from left to right. By using the GLV-SAC recoding, every digit-column  $i$  is expected to be nonzero and have any of the possible combinations  $[b_i^{m-1}, \dots, b_i^2, b_i^1, b_i^0]$ , where  $b_i^0 \in \{1, -1\}$ , and  $b_i^j \in \{0, b_i^0\}$  for  $1 \leq j < m$ . Since nonzero digits in the same column have the same sign, one only needs to precompute all the positive combinations  $P_0 + u_1P_1 + \dots + u_{m-1}P_{m-1}$  with  $u_j \in \{0, 1\}$ , where  $P_j$  are the base points of the sub-scalars. Assuming that negation of group elements is inexpensive in a given curve subgroup, negative values can be computed on-the-fly during the evaluation stage.

Since the GLV-SAC recoding requires that the “sign-aligner”  $k_J$  (in this case,  $k_0$ ) be odd,  $k_0$  is subtracted by one if it is even in Step 3 of Algorithm 2. The correction is then performed at the end of the evaluation stage at Step 9. These computations, as well as the accesses to the precomputed table, should be performed in constant time to guarantee protection against timing attacks. For example, to perform Step 9 in the

---

**Algorithm 2** Protected  $m$ -GLV Variable-Base Scalar Multiplication using the GLV-SAC Representation.

**Input:** Base point  $P_0$  of order  $r$  and  $(m-1)$  points  $P_j$  for  $1 \leq j < m$  corresponding to the endomorphisms,  $m$  scalars  $k_j = (k_{t_j-1}^j, \dots, k_0^j)_2$  for  $0 \leq j < m$ ,  $l = \lceil \frac{\log_2 r}{m} \rceil + 1$  and  $\max(t_j) = \lceil \frac{\log_2 r}{m} \rceil$ .

**Output:**  $kP$ .

---

**Precomputation stage:**

1: Compute  $P[u] = P_0 + u_0P_1 + \dots + u_{m-2}P_{m-1}$  for all  $0 \leq u < 2^{m-1}$ , where  $u = (u_{m-2}, \dots, u_0)_2$ .

**Recoding stage:**

2:  $\text{even} = k_0 \bmod 2$

3: **if**  $\text{even} = 0$  **then**  $k_0 = k_0 - 1$

4: Pad each  $k_j$  with  $(l-t_j)$  zeros to the left for  $0 \leq j < m$  and convert them to the GLV-SAC representation using Algorithm 1 s.t.  $k_j = (b_{l-1}^j, \dots, b_0^j)_{\text{GLV-SAC}}$ . Set digit-columns  $\mathbb{K}_i = [b_i^{m-1}, \dots, b_i^2, b_i^1] \equiv [b_i^{m-1}2^{m-2} + \dots + b_i^22 + b_i^1]$  and digit-column signs  $s_i = b_i^0$  for  $0 \leq i \leq l-1$ .

**Evaluation stage:**

5:  $Q = s_{l-1}P[\mathbb{K}_{l-1}]$

6: **for**  $i = l-2$  **to**  $0$  **do**

7:      $Q = 2Q$

8:      $Q = Q + s_iP[\mathbb{K}_i]$

9:     **if**  $\text{even} = 0$  **then**  $Q = Q + P_0$

10: **return**  $Q$

---

implementation discussed in Section 6 we always carry out the computation  $Q' = Q + P_0$ . Then, we perform a linear pass over the points  $Q$  and  $Q'$  using conditional move instructions to transfer the correct value to the final destination.

Note that Algorithm 2 assumes a decomposed scalar as input. This is sufficient in some settings, in which randomly generated sub-scalars could be provided. However, in others settings, one requires to calculate the sub-scalars in a decomposition phase. We remark that this computation should also be performed in constant time for protecting against timing attacks.

**Example 1.** Let  $m = 4$ ,  $\log_2 r = 16$  and  $kP = 11P_0 + 6P_1 + 14P_2 + 3P_3$ . Using Algorithm 1, the corresponding GLV-SAC representation with fixed length  $l = \lceil 16/4 \rceil + 1 = 5$  is given by (arranged in matrix form from top to bottom as required in Alg. 2)

$$\begin{bmatrix} k_0 \\ k_1 \\ k_2 \\ k_3 \end{bmatrix} \equiv \begin{bmatrix} 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 \end{bmatrix} \equiv \begin{bmatrix} 1 & \bar{1} & 1 & \bar{1} & 1 \\ 1 & \bar{1} & 0 & \bar{1} & 0 \\ 1 & 0 & 0 & \bar{1} & 0 \\ 0 & 0 & 1 & \bar{1} & 1 \end{bmatrix}$$

According to Algorithm 2, digit columns are given by  $\mathbb{K}_0 = [100] = 4$ ,  $\mathbb{K}_1 = [\bar{1}\bar{1}\bar{1}] = 7$ ,  $\mathbb{K}_2 = [100] = 4$ ,  $\mathbb{K}_3 = [00\bar{1}] = 1$  and  $\mathbb{K}_4 = [011] = 3$ , and their corresponding  $s_i$  are  $s_0 = 1$ ,  $s_1 = -1$ ,  $s_2 = 1$ ,  $s_3 = -1$  and  $s_4 = 1$ . Precomputed values  $P[u]$  are given by  $P[0] = P_0$ ,  $P[1] = P_0 + P_1$ ,  $P[2] = P_0 + P_2$ ,  $P[3] = P_0 + P_1 + P_2$ ,  $P[4] = P_0 + P_3$ ,  $P[5] = P_0 + P_1 + P_3$ ,  $P[6] = P_0 + P_2 + P_3$  and  $P[7] = P_0 + P_1 + P_2 + P_3$ . At Step 5 of Alg. 2, we compute  $Q = s_4P[\mathbb{K}_4] = P[3] = P_0 + P_1 + P_2$ . The main loop in the evaluation stage is then executed as follows

$i$	3	2	1	0
$2Q$	$2P_0 + 2P_1 + 2P_2$	$2P_0 + 2P_1 + 4P_2$	$6P_0 + 4P_1 + 8P_2 + 2P_3$	$10P_0 + 6P_1 + 14P_2 + 2P_3$
$Q + s_iP[\mathbb{K}_i]$	$P_0 + P_1 + 2P_2$	$3P_0 + 2P_1 + 4P_2 + P_3$	$5P_0 + 3P_1 + 7P_2 + P_3$	$11P_0 + 6P_1 + 14P_2 + 3P_3$

**Cost Analysis.** To simplify comparisons, we will only consider a setting in which precomputed points are left in some projective system. When converting points to affine is convenient, one should include the cost of this conversion. Also, we do not consider optimizations exploiting cheap endomorphism mappings during

precomputation, since this is dependent on a specific application. The reader is referred to Section 6 for a more precise comparison in a practical implementation using a Twisted Edwards GLV-GLS curve.

The cost of the proposed  $m$ -GLV variable-base scalar multiplication using the GLV-SAC representation (Alg. 2) is given by  $(l - 1)$  doublings and  $l$  additions during the evaluation stage using  $2^{m-1}$  points, where  $l = \lceil \frac{\log_2 r}{m} \rceil + 1$ . Naively, precomputation costs  $2^{m-1} - 1$  additions (in practice, several of these additions might be performed using cheaper mixed additions). So the total cost is given by  $(l - 1)$  doublings and  $(l + 2^{m-1} - 1)$  additions.

In contrast, the method based on the regular windowed recoding [31] used in [28] requires  $(l - 1)$  doublings and  $m \cdot (l - 1)/(w - 1) + 2m - 1$  additions during the evaluation stage and  $m$  doublings with  $m \cdot (2^{w-2} - 1)$  additions during the precomputation stage, using  $m \cdot (2^{w-2} + 1)$  points (naive approach without exploiting endomorphisms). If, for example,  $r = 256$ ,  $m = 4$  and  $w = 5$  (typical parameters to achieve 128-bit security on a curve similar to `Ted127-g1v4`), the new method costs 64 doublings and 72 additions using 8 points, whereas the regular windowed method costs 68 doublings and 99 additions using 36 points. Thus, the new method improves performance while reduces dramatically the number of precomputations (in this case, to almost 1/5 of the storage). Assuming that one addition costs 1.3 doublings, the expected speedup is 20%.

Certainly, one can reduce the number of precomputations when using the regular windowed recoding by only precomputing multiples corresponding to one or some of the sub-scalars. However, these savings in memory come at the expense of computing endomorphisms during the evaluation stage, which can cost from several multiplications [7] to approximately one full point addition each (see Appendix B). The proposed method always requires the minimum storage without impacting performance.

The basic GLV-SAC representation and its scalar multiplication are particularly efficient for four-dimensional GLV. In the following section, we discuss variants that are efficient for  $m = 2$  and  $m \geq 8$ .

### 3.2 Windowed and Partitioned GLV-SAC: Case of Dimension 2 and 8 (or Higher)

In some cases, the performance of the proposed scalar multiplication can be improved further by combining windowed techniques with the GLV-SAC recoding. Given a window width  $w$ , assume that a set of sub-scalars  $k_j$  has been padded with enough zeros to the left to guarantee that  $w|l$ , where  $l$  is the expected length of an extended GLV-SAC representation that we refer to as  $w$ GLV-SAC. The basic idea is to join every  $w$  consecutive digits in the  $w$ GLV-SAC representation, and precompute all possible values  $P[u] = u'P_0 + u_0P_1 + \dots + u_{m-2}P_{m-1}$  for  $0 \leq u < 2^{wm-1}$  and  $u' \in \{1, 3, \dots, 2^w - 1\}$ , where  $u_j = \{0, 1, 2, \dots, 2^w - 1\}$  (again, points corresponding to negative values of  $u'$  can be computed on-the-fly). Scalar multiplication then proceeds by scanning  $w$ -digit columns from left to right.

Conveniently, Algorithm 1 can also be used to obtain  $w$ GLV-SAC( $k_j$ ), with the only change in the fixed length to  $l = (\lceil \log_2 r/w \rceil + 1) \bmod w + (\lceil \log_2 r/w \rceil + 1)$ .

**Example 2.** Let  $m = 2$ ,  $\log_2 r = 8$ ,  $w = 2$  and  $kP = 11P_0 + 14P_1$ . Using Algorithm 1, the corresponding  $w$ GLV-SAC representation with fixed length  $l = (\lceil 8/2 \rceil + 1) \bmod 2 + \lceil 8/2 \rceil + 1 = 6$ , arranged in matrix form from top to bottom, is given by

$$\begin{bmatrix} k_0 \\ k_1 \end{bmatrix} \equiv \begin{bmatrix} 0 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 0 \end{bmatrix} \equiv \begin{bmatrix} 1 & \bar{1} & | & \bar{1} & 1 & | & \bar{1} & 1 \\ 1 & \bar{1} & | & 0 & 0 & | & \bar{1} & 0 \end{bmatrix}$$

In this case, one can consider 2-digit columns  $\mathbb{K}_i = [b_{2i+1}^1 b_{2i}^1, b_{2i+1}^0 b_{2i}^0] \equiv (3 \cdot |2b_{2i+1}^0 + b_{2i}^0| + 1)/2 + |2b_{2i+1}^1 + b_{2i}^1| - 1$  with signs  $s_i = b_{2i+1}^0$ , for  $0 \leq i < l/2$  (using the notation  $k_j = (b_{l-1}^j, \dots, b_0^j)_{w\text{GLV-SAC}}$ ). Precomputed values are given by  $P[u] = u'P_0 + u_0P_1$  for  $0 \leq u < 8$  and  $u' \in \{1, 3\}$ , where  $u_0 \in \{0, 1, 2, 3\}$ . Thus, in the example, the 2-digit columns are given by  $\mathbb{K}_0 = [\bar{2}\bar{1}] = 3$ ,  $\mathbb{K}_1 = [0\bar{1}] = 1$  and  $\mathbb{K}_2 = [11] = 2$ , and their corresponding  $s_i$  are  $s_0 = -1$ ,  $s_1 = -1$  and  $s_2 = 1$ . The precomputed values  $P[u]$  are given by  $P[0] = P_0 - P_1$ ,  $P[1] = P_0$ ,  $P[2] = P_0 + P_1$ ,  $P[3] = P_0 + 2P_1$ ,  $P[4] = 3P_0$ ,  $P[5] = 3P_0 + P_1$ ,  $P[6] = 3P_0 + 2P_1$  and  $P[7] = 3P_0 + 3P_1$ . In the evaluation stage we first compute  $Q = s_2P[\mathbb{K}_2] = P[2] = P_0 + P_1$  and then execute



$i$	1	0
$2^w Q$	$4P_0 + 4P_1$	$12P_0 + 16P_1$
$Q + s_i P[\mathbb{K}_i]$	$3P_0 + 4P_1$	$11P_0 + 14P_1$

Since the requirement of precomputations, given by  $2^{wm-1}$ , increases rapidly as  $w$  and  $m$  grow, windowed GLV-SAC is especially attractive for two-dimensional GLV implementations. In this case, by fixing  $w = 2$  the number of precomputed points is only 8. At the same performance level (in terms of number of additions and doublings in the evaluation stage), this is approximately half the memory requirement of a method based on the regular windowed recoding [31]. Although Oliveira et al. [32] manage to reduce the number of precomputations from 16 to 8, these savings in memory come at the expense of computing endomorphisms on-the-fly during the evaluation stage.

Whereas joining columns in the representation matrix is amenable for small  $m$  using windowing, for large  $m$  it is recommended to join rows instead. We illustrate the approach with  $m = 8$ . Given a set of sub-scalars  $k_j$  for  $0 \leq j < 8$ , we first partition it in  $c$  consecutive sub-sets  $k'_i$  such that  $c|8$ , and then convert every sub-set to the GLV-SAC representation (using Algorithm 1). In this case, every column in the matrix consists of  $c$  sub-columns, each one corresponding to a sub-set  $k'_i$ . Scalar multiplication then proceeds by scanning  $c$  sub-columns per iteration from right to left. Thus, with this “partitioned” GLV-SAC approach, one increases the number of point additions per iteration in the main loop of Algorithm 2 from one to  $c$ . However, the number of required precomputations is reduced significantly from  $2^{m-1}$  to  $c \cdot 2^{\frac{m}{c}-1}$ . For example, for  $m = 8$ , this introduces a reduction in the number of points from 128 to only 16 if  $c$  is fixed to 2 (each sub-table corresponding to a sub-set of scalars contains 8 points). At the same performance level (in terms of number of additions and doublings in the evaluation stage), this is approximately half the memory requirement of a method based on the regular windowed recoding [31], as discussed by the recent work by Bos et al. [8]. Performance is also expected to be improved since the number of point operations in the precomputation stage is significantly reduced. Note that, if one only considers positive sub-scalars and the endomorphism mapping is inexpensive in comparison to point addition, then sub-tables can be computed by simply applying the endomorphism to the first sub-table arising from the base point  $P_0$ . In some instances, such as the 8-GLV in [8], this approach is expected to reduce further the cost of precomputation. However, an issue arises when sub-scalars can also be negative. We solve this problem in the next subsection.

### 3.3 Extended GLV-SAC Representation: Dealing with Negative Sub-Scalars

Typically, sub-scalars obtained from decomposition methods can be positive as well as negative. However, for efficiency reasons (specifically, to get the minimum length in the representation) the basic version of GLV-SAC only works for positive integers (see at the beginning of Section 3). In general, a straightforward solution during the scalar multiplication is to simply convert negative sub-scalars to positive and then negate the corresponding base  $P_i$ . Recoding to GLV-SAC can then be performed with Algorithm 1 as before. However, an issue with this approach arises when the precomputed table needs to be partitioned in more than one sub-table (see, for example, the case of 8-GLV in the previous subsection). If the bases  $P_i$  are negated depending on the sign of their corresponding sub-scalar then cheap applications of the GLV endomorphism cannot be conveniently applied to compute instances of the original sub-table.

To solve this problem, we present an extended version of the GLV-SAC representation, exhibiting the following properties

- (i) The length of the digit representation of every sub-scalar  $k_j$  is fixed and given by  $l = \lceil \log_2 r/m \rceil + 2$ , where  $r$  is the prime group order.
- (ii) and (iii) are the same as the original GLV-SAC representation (see at the beginning of Section 3).

Thus, the only difference with the original GLV-SAC representation is the addition of one digit to the length, in order to deal with the sign. We present Algorithm 3 for the efficient recoding to this representation.

---

**Algorithm 3** Protected Recoding Algorithm for the Extended GLV-SAC Representation.
 

---

**Input:**  $m$   $l$ -bit integers  $k_j = (k_{l-1}^j, \dots, k_0^j)_2$  and their corresponding signs  $s_j \in \{-1, 1\}$  for  $0 \leq j < m$ , an odd “sign-aligner”  $k_J \in \{k_j\}^m$ , where  $l = \lceil \log_2 r/m \rceil + 2$ ,  $m$  is the GLV dimension and  $r$  is the prime group order.

**Output:**  $(b_{l-1}^j, \dots, b_0^j)_{\text{Ext-GLV-SAC}}$  for  $0 \leq j < m$ , where  $b_i^J \in \{1, -1\}$ , and  $b_i^j \in \{0, b_i^J\}$  for  $0 \leq j < m$  and  $j \neq J$ .

---

```

1:  $b_{l-1}^J = 1$ 
2: for  $i = 0$  to  $(l - 2)$  do
3:    $b_i^J = s_J \cdot (2k_{i+1}^J - 1)$ 
4: for  $j = 0$  to  $(m - 1), j \neq J$  do
5:   for  $i = 0$  to  $(l - 1)$  do
6:      $b_i^j = b_i^J \cdot k_0^j$ 
7:      $k_j = \lfloor k_j/2 \rfloor - s_j \cdot \lfloor b_i^J/2 \rfloor$ 
8: return  $(b_{l-1}^j, \dots, b_0^j)_{\text{Ext-GLV-SAC}}$  for  $0 \leq j < m$ .

```

---

## 4 Modified LSB-Set Comb Method for Fixed-Base Scalar Multiplication

The side-channel protected comb methods by Feng et al. [10, 11] and Hedabou et al. [19] used a simple version of the Lim-Lee’s comb method that is restricted to only one table (see [25] for more details). Recently, Hamburg [16] precisely proposed to combine the original multi-table Lim-Lee’s comb approach with Hedabou et al.’s SAB-set representation for improved performance. In this section, we follow a similar observation to extend Feng et al.’s LSB-set comb method with the use of multiple tables.

Let  $k$  be a scalar with bitlength  $t$ . Assume that the binary representation of  $k$  is updated to  $(k_{l-1}, k_{l-2}, \dots, k_0)$  by padding with  $(dw - t)$  zeros to the left for some window width  $w$ , where  $l = dw$ ,  $d = \lceil t/w \rceil$  and  $w \geq 2$  (see Section 2.3). First, we note that the basic LSB-set comb method [10] requires additional operations in comparison with the original comb method when  $w \nmid t$ . In [11], Feng et al. fixed this deficiency by proposing an MSB-set representation that recodes the most significant  $d$  bits to a nonzero representation, instead of the least significant  $d$  bits. In addition, MSB-set lifts the restriction to odd integers, inherent to the LSB-set representation. We show here that it is possible to modify the original LSB-set representation to fix the problem with the extra operations. Moreover, we easily solve the restriction to odd integers by exploiting the group order during scalar multiplication, and thus benefit from a somewhat simpler recoding algorithm.

As in [25], we add a parameter, referred to as  $v$ , that represents the number of tables to use. The modified LSB-set representation (denoted by  $m\text{LSB-set}$ ) has the following properties

- (i) Given a window width  $w \geq 2$ , a table parameter  $v \geq 1$  and an odd integer  $k$  with bitlength  $t$ , the length of  $m\text{LSB-set}(k)$  is given by  $l = dw$ , where  $d = ev$  and  $e = \lceil t/wv \rceil$ .
- (ii) The least significant  $d$  digits  $(b_{d-1}, \dots, b_0)$  are in signed nonzero form, i.e.,  $b_i \in \{1, -1\}$  for  $0 \leq i < d$ .
- (iii) Remaining digits  $b_i$  for  $d \leq i < l$  are expressed by a signed representation  $(c, b_{l-1}, \dots, b_d)$  s.t.  $b_i \in \{0, b_{i \bmod d}\}$  and  $c \in \{0, 1\}$ . If  $wv \nmid t$ ,  $c$  is discarded.

Similarly to GLV-SAC, (i) and (ii) will enable a constant-time execution regardless of the value of  $k$  without having to appeal to masking for dealing with the identity element. Item (iii) will allow us to reduce the size of the precomputed table by a factor of 2, while minimizing the cost of precomputation. In our optimized setting using multiple tables, there is an additional carry bit  $c$  whenever  $wv \mid t$ . We remark that the latter reveals nothing since the appearance of the carry bit depends on public parameters. To deal with this extra bit during scalar multiplication we perform a correction with a precomputed value (see next section). Again, the condition that  $k$  should be odd enables the conversion of the least significant  $d$  bits to a signed nonzero representation using the equivalence  $1 \equiv \bar{1} \dots \bar{1}$ . To deal with this restriction during scalar multiplication, we take advantage that our setting consists of a subgroup of prime order  $r$ . Hence,  $-k \pmod{r} = r - k$  gives an odd result if  $k$  is even.

An efficient algorithm to recode a given scalar to its  $m$ LSB-set representation is given in Algorithm 4. Conveniently, the recoding algorithm is simple and regular, exhibiting resistance to timing attacks. Note that, depending on the targeted platform, other implementation alternatives are possible. For example, Step 3 can be computed as  $b_i = (k_{i+1} - 1) | 1$ , where  $|$  represents a logical OR operation.

---

**Algorithm 4** Protected Odd-Only Recoding Algorithm for the Modified LSB-Set Representation.

---

**Input:** An odd  $l$ -bit integer  $k = (k_{l-1}, \dots, k_0)_2$ , window width  $w \geq 2$  and table parameter  $v \geq 1$ , where  $l = dw$ ,  $d = ev$  and  $e = \lceil t/wv \rceil$ .

**Output:**  $(c, b_{l-1}, \dots, b_0)_{m\text{LSB-set}}$ , where  $b_i \in \{1, -1\}$  for  $0 \leq i < d$  and  $b_i \in \{0, b_{i \bmod d}\}$  for  $d \leq i \leq l - 1$ . If  $wv \nmid t$  then  $c \in \{0, 1\}$ , otherwise,  $c$  is discarded.

---

```

1:  $b_{d-1} = 1$ 
2: for  $i = 0$  to  $(d - 2)$  do
3:    $b_i = 2k_{i+1} - 1$ 
4:  $c = \lfloor k/2^d \rfloor$ 
5: for  $i = d$  to  $(l - 1)$  do
6:    $b_i = b_{i \bmod d} \cdot c_0$ 
7:    $c = \lfloor c/2 \rfloor - \lfloor b_i/2 \rfloor$ 
8: return  $(c, b_{l-1}, \dots, b_0)_{m\text{LSB-set}}$ . Discard  $c$  if  $wv \nmid t$ .

```

---

#### 4.1 Fixed-Based Scalar Multiplication using $m$ LSB-Set

The new algorithm for computing fixed-base scalar multiplication using the modified LSB-set representation is shown as Algorithm 5. Disregarding  $c$ , the basic idea is to split  $m\text{LSB-set}(k) = (c, b_{l-1}, \dots, b_0)$  in consecutive  $d$ -digit parts  $K^{w'}$ , for  $0 \leq w' < w$ , and arrange them in matrix form from top to bottom, with the least significant  $d$  digits (i.e.,  $K^0$ ) at the top. Then, partition each  $K^{w'}$  in  $v$  strings of  $e$  digits each, where  $v \geq 1$  represents a suitably chosen number of tables (each table consists of  $v$  consecutive digit-columns). Thus, we have that  $K_{v',e'}^{w'} = b_{dw'+ev'+e'}$ , where  $K_{v',e'}^{w'}$  denotes the  $e'$ -th digit in the  $v'$ -th string of a given  $K^{w'}$ . We then run a simultaneous multi-scalar multiplication scanning digit-columns from left to right, taking an entry from each table. By using the  $m$ LSB-set recoding, every digit-column is nonzero by definition and has any of the possible combinations  $[K_{v',e'}^{w'-1}, \dots, K_{v',e'}^1, K_{v',e'}^0]$ , where  $K_{v',e'}^0 \in \{1, -1\}$ , and  $K_{v',e'}^{w'} \in \{0, K_{v',e'}^0\}$  for  $1 \leq w' < w$ . Since nonzero digits in the same column have the same sign, one only needs to precompute all the positive combinations  $P[u][v'] = 2^{ev'}(1 + u_02^d + \dots + u_{w-2}2^{(w-1)d})P$  for all  $0 \leq u < 2^{w-1}$  and  $0 \leq v' < v$ , where  $u = (u_{w-2}, \dots, u_0)_2$  and  $P$  is the base point. Assuming that negation of group elements is inexpensive in a given curve subgroup, negative values can be computed on-the-fly during the evaluation stage. Note that, in the fixed-base scenario,  $P$  is assumed to be known in advance and, hence, the precomputation can be computed off-line. Finally, if  $wv \mid t$  for the chosen values for  $w$  and  $v$ , we need to precompute the extra point  $2^{wd}P$  and apply a correction at the end of the evaluation stage (Step 10 of Alg. 5). A detailed proof is shown in Appendix A.

As can be seen, the main loop of Algorithm 5 computes  $kP$  using the regular pattern one doubling,  $v$  additions,  $\dots$ , one doubling,  $v$  additions and, hence, it provides the first step to protect against timing attacks and SSCA attacks such as SPA. Following previous recommendations, table accesses and conditional statements (e.g., Steps 3, 9 and 10) should be performed in constant-time to guarantee protection against timing attacks. We also note that for the GLV-setting Algorithm 5 does not exploit endomorphisms. It is an open problem to find a scalar multiplication method that exploits endomorphisms efficiently in the fixed-base case.

In the on-line computation, the method requires  $e - 1 = \lceil \frac{t}{w \cdot v} \rceil - 1$  doublings and  $ev - 1 = v \lceil \frac{t}{w \cdot v} \rceil - 1$  additions, using  $v \cdot 2^{w-1}$  precomputed points. This cost should be increased in one addition using an extra precomputed point for the case in which  $wv \nmid t$ . This cost is similar to Hamburg's method using the SAB-set representation [16]. In comparison, the SAB-set [19] and MSB-set [11] comb methods require approximately  $(d - 1)$  doublings and  $d$  additions using  $2^{w-1}$  points, where  $d = \lceil t/w \rceil$ . For example, if we assume that one addition costs 1.3 doublings and  $t = 256$  bits using 256 precomputed points, previous comb methods cost

approximately 66 doublings (with  $w = 9$ ), whereas the proposed method costs approximately 57 doublings (with  $w = 8, v = 2$  optimal), injecting a 14% speedup.

**Example 3.** Let  $w = 2, v = 2$  and  $kP = 395P$ . Using Algorithm 4, the corresponding  $m$ LSB-set representation with fixed length  $l = 6 \cdot 2 = 12$ , where  $e = 3$  and  $d = 6$ , arranged in matrix form from top to bottom, is given by

$$\begin{bmatrix} K^0 \\ K^1 \end{bmatrix} \equiv \begin{bmatrix} 0 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 0 \end{bmatrix} \equiv \begin{bmatrix} 1 & \bar{1} & \bar{1} & 1 & \bar{1} & 1 \\ 1 & \bar{1} & \bar{1} & 0 & \bar{1} & 0 \end{bmatrix}$$

The digit columns are given by  $\mathbb{K}_{0,2} = [0] = 0, \mathbb{K}_{0,1} = [\bar{1}] = 1, \mathbb{K}_{0,0} = [0] = 0, \mathbb{K}_{1,2} = [1] = 1, \mathbb{K}_{1,1} = [\bar{1}] = 1$  and  $\mathbb{K}_{1,0} = [\bar{1}] = 1$ , and their corresponding  $s_{v',e'}$  are  $s_{0,2} = 1, s_{0,1} = -1, s_{0,0} = 1, s_{1,2} = 1, s_{1,1} = -1$  and  $s_{1,0} = -1$ . Precomputed values  $P[u][v']$  are given by  $P[0][0] = P, P[1][0] = (1+64)P = 65P, P[0][1] = 8P$  and  $P[1][1] = 8 \cdot (1+64)P = 520P$ . In the evaluation stage we first compute  $Q = s_{0,2}P[\mathbb{K}_{0,2}][0] + s_{1,2}P[\mathbb{K}_{1,2}][1] = P + 520P = 521P$  and then execute

$i$	1	0
$2Q$	$1042P$	$914P$
$Q + s_{v',i}P[\mathbb{K}_{v',i}][v']$	$(1024 - 65 - 520)P = 457P$	$(914 + 1 - 520)P = 395P$

---

**Algorithm 5** Protected Fixed-Base Scalar Multiplication using the Modified LSB-Set Comb Method.

---

**Input:** A point  $P \in E(\mathbb{F}_q)$  of prime order  $r$ , a scalar  $k = (k_{t-1}, \dots, k_0)_2 \in [1, r-1]$ , window width  $w \geq 2$ , table parameter  $v \geq 1$ , where  $d = ev$  and  $e = \lceil t/wv \rceil$ .

**Output:**  $kP$ .

---

**Precomputation stage:**

1: Compute  $P[u][v'] = 2^{ev'}(1 + u_02^d + \dots + u_{w-2}2^{(w-1)d})P$  for all  $0 \leq u < 2^{w-1}$  and  $0 \leq v' < v$ , where  $u = (u_{w-2}, \dots, u_0)_2$ . If  $wv \mid t$  then compute  $2^{wd}P$ .

**Recoding stage:**

2:  $\text{even} = k \bmod 2$

3: **if**  $\text{even} = 0$  **then**  $k = r - k$

4: Pad  $k$  with  $(dw - t)$  zeros to the left and convert it to the  $m$ LSB-set representation using Algorithm 4 s.t.  $k = (c, b_{l-1}, \dots, b_0)_{m\text{LSB-set}}$ , where  $l = dw$ . Set  $k = K^{w-1} \parallel \dots \parallel K^1 \parallel K^0$ , where each  $K^{w'}$  consists of  $v$  strings of  $e$  digits each. Let the  $v'$ -th string in a given  $K^{w'}$  be denoted by  $K_{v'}^{w'}$ , and the  $e'$ -th digit in a given  $K_{v'}^{w'}$  be denoted by  $K_{v',e'}^{w'}$ , s.t.  $K_{v',e'}^{w'} = b_{dw'+ev'+e'}$ . Set digit-columns  $\mathbb{K}_{v',e'} = [K_{v',e'}^{w-1}, \dots, K_{v',e'}^2, K_{v',e'}^1] \equiv |K_{v',e'}^{w-1}2^{w-2} + \dots + K_{v',e'}^2 + K_{v',e'}^1|$  and digit-column signs  $s_{v',e'} = K_{v',e'}^0$ .

**Evaluation stage:**

5:  $Q = s_{0,e-1}P[\mathbb{K}_{0,e-1}][0] + s_{1,e-1}P[\mathbb{K}_{1,e-1}][1] + \dots + s_{v-1,e-1}P[\mathbb{K}_{v-1,e-1}][v-1]$

6: **for**  $i = e - 2$  **to** 0 **do**

7:      $Q = 2Q$

8:      $Q = Q + s_{0,i}P[\mathbb{K}_{0,i}][0] + s_{1,i}P[\mathbb{K}_{1,i}][1] + \dots + s_{v-1,i}P[\mathbb{K}_{v-1,i}][v-1]$

9:     **if**  $\text{even} = 0$  **then**  $Q = -Q$

10: **if**  $(wv \mid t \wedge c = 1)$  **then**  $Q = Q + 2^{wd}P$

11: **return**  $Q$

---

## 5 High-Speed Implementation on GLV-GLS Curves

In this section, we describe implementation aspects of the GLV-GLS curve `Ted127-glv4`. We present optimized algorithms for prime field, extension field and point arithmetic. We also present the technique of interleaving NEON and ARM-based multiprecision operations over  $\mathbb{F}_{p^2}$ . Although our techniques are especially tuned for the targeted curve, we remark that they can be adapted and exploited in other scenarios.

## 5.1 The Curve

For complete details about the four-dimensional method using GLV-GLS curves, the reader is referred to [13] and [29]. We use the following GLV-GLS curve in twisted Edwards form [28], referred to as `Ted127-glv4`:

$$E'_{TE}/\mathbb{F}_{p^2} : -x^2 + y^2 = 1 + dx^2y^2, \quad (1)$$

where  $\mathbb{F}_{p^2}$  is defined as  $\mathbb{F}_p[i]/(i^2 - \beta)$ ,  $\beta = -1$  is a quadratic non-residue in  $\mathbb{F}_p$  and  $u = 1 + i$  is a quadratic non-residue in  $\mathbb{F}_{p^2}$ . Also,  $p = 2^{127} - 5997$ ,  $d = 170141183460469231731687303715884099728 + 116829086847165810221872975542241037773i$  and  $\#E'_{TE}(\mathbb{F}_{p^2}) = 8r$ , where  $r$  is the 251-bit prime  $2^{251} - 255108063403607336678531921577909824432295$ .  $E'_{TE}$  is isomorphic to the Weierstrass curve  $E'_W/\mathbb{F}_{p^2} : y^2 = x^3 - 15/2 u^2x - 7u^3$ , which is the quadratic twist of a curve isomorphic to the GLV curve  $E_W/\mathbb{F}_p : y^2 = 4x^3 - 30x - 28$  (see [28, Section 5]).  $E'_{TE}/\mathbb{F}_{p^2}$  is equipped with two efficiently computable endomorphisms  $\Phi$  and  $\Psi$  defined over  $\mathbb{F}_{p^2}$ , which enable a four-dimensional decomposition for any scalar  $k \in [1, r-1]$  in the subgroup generated by a point  $P$  of order  $r$  and, consequently, enable a four-dimensional scalar multiplication given by

$$kP = k_1P + k_2\Phi(P) + k_3\Psi(P) + k_4\Psi\Phi(P), \quad \text{with } \max_i(|k_i|) < Cr^{1/4}$$

for some explicit  $C > 0$  [28]. Let  $\zeta_8 = u/\sqrt{-2a'}$ , where  $a' = 27u^3(\frac{\sqrt{2}}{2} - 1)$ , be a primitive 8th root of unity. The affine formulas for  $\Phi(x, y)$  and  $\Psi(x, y)$  are given by

$$\Phi(x, y) = \left( -\frac{(\zeta_8^3 + 2\zeta_8^2 + \zeta_8)xy^2 + (\zeta_8^3 - 2\zeta_8^2 + \zeta_8)x}{2y}, \frac{(\zeta_8^2 - 1)y^2 + 2\zeta_8^3 - \zeta_8^2 + 1}{(2\zeta_8^3 + \zeta_8^2 - 1)y^2 - \zeta_8^2 + 1} \right) \quad \text{and} \quad \Psi(x, y) = \left( \zeta_8x^p, \frac{1}{y^p} \right),$$

respectively. It can be verified that  $\Phi^2 + 2 = 0$  and  $\Psi^2 + 1 = 0$ . The formulas in homogeneous projective coordinates can be found in Appendix B.

Note that `Ted127-glv4` has  $a = -1$  (in the twisted Edwards equation; see [3]), corresponding to the most efficient set of formulas proposed by Hisil et al. [20]. Although GLV-GLS curves with suitably chosen parameters when transformed to twisted Edwards form offer roughly the same performance, as discussed in [28], there are certain differences in the cost of formulas for computing the endomorphisms  $\Phi$  and  $\Psi$ . Curve `Ted127-glv4` exhibits relatively efficient formulas for computing the endomorphisms in comparison with other GLV-GLS curves from [28]. On the other hand, our selection of the pseudo-Mersenne prime  $p = 2^{127} - 5997$  enables efficient field arithmetic by exploiting lazy and incomplete reduction techniques (see the next section for details). Also, since  $p \equiv 3 \pmod{4}$ ,  $-1$  is a quadratic non-residue in  $\mathbb{F}_p$ , which minimizes the cost of multiplication over  $\mathbb{F}_{p^2}$  by transforming multiplications by  $\beta$  to inexpensive subtractions.

## 5.2 Field Arithmetic

For field inversion, we use the modular exponentiation  $a^{p-2} \pmod{p} \equiv a^{-1}$  using a fixed and short addition chain. This method is simple to implement and inherently protected against timing attacks.

In the case of a pseudo-Mersenne prime of the form  $p = 2^m - c$ , with  $c$  small, field multiplication can be efficiently performed by computing an integer multiplication followed by a modular reduction exploiting the special form of the prime. This separation of operations also enables the use of lazy reduction in the extension field arithmetic. For x64, integer multiplication is implemented in product scanning form (a.k.a Comba's method), mainly exploiting the powerful 64-bit unsigned multiplier instruction. Let  $0 \leq a, b < 2^{m+1}$ . To exploit the extra room of one bit in our targeted prime  $2^{127} - 5997$ , we first compute  $M = a \cdot b = 2^{m+1}M_H + M_L$  followed by the reduction step  $R = M_L + 2cM_H \leq 2^{m+1}(2c+1) - 2$ . Then, given  $R = 2^mR_H + R_L$ , we compute  $R_L + cR_H \pmod{p}$ , where  $R_L, cR_H < 2^m$ . This final operation can be efficiently carried out by employing the modular addition proposed by Bos et al. [7] to get the final result in the range  $[0, p-1]$ . Note that the computation of field multiplication above naturally accepts inputs in unreduced form without incurring in

extra costs, enabling the use of additions without correction or operations with incomplete reduction (see below for more details). We follow a similar procedure for computing field squaring. For ARM, we implement the integer multiplication using the schoolbook method. In this case, and also for modular reduction, we extensively exploit the parallelism of ARM and NEON instructions. The details are discussed in Section 5.4.

Let  $0 \leq a, b < 2^m - c$ . Field subtraction is computed as  $(a - b) + borrow \cdot 2^m - borrow \cdot c$ , where  $borrow = 0$  if  $a \geq b$ , otherwise  $borrow = 1$ . Notice that in practice the addition with  $borrow \cdot 2^m$  can be efficiently implemented by clearing the  $(m + 1)$ -th bit of  $a - b$ .

**Incomplete Reduction.** Similar to [27], we exploit the form of the pseudo-Mersenne prime in combination with the incomplete reduction technique to speedup computations. We also mix incompletely reduced and completely reduced operands in novel ways.

Let  $0 \leq a < 2^m - c$  and  $0 \leq b < 2^m$ . Field addition with incomplete reduction is computed as  $(a + b) - carry \cdot 2^m + carry \cdot c$ , where  $carry = 0$  if  $a + b < 2^m$ , otherwise  $carry = 1$ . Again, in practice the subtraction with  $carry \cdot 2^m$  can be efficiently implemented by clearing the  $(m + 1)$ -th bit of  $a + b$ . The result is correct modulo  $p$ , but falls in the range  $[0, 2^m - 1]$ . Thus, this addition operation with incomplete reduction works with both operands in completely reduced form or with one operand in completely reduced form and one in incompletely reduced form. A similar observation applies to subtraction. Consider two operands  $a$  and  $b$ , such that  $0 \leq a < 2^m$  and  $0 \leq b < 2^m - c$ . The standard field subtraction  $(a - b) \bmod (2^m - c)$  described above will then produce an incompletely reduced result in the range  $[0, 2^m - 1]$ , since  $a - b$  with  $borrow = 0$  produces a result in the range  $[0, 2^m - 1]$  and  $a - b$  with  $borrow = 1$  produces a result in the range  $[-2^m + c + 1, -1]$ , which is then fully reduced by adding  $2^m - c$ . Thus, performance can be improved by using incomplete reduction for an addition preceding a subtraction. For example, this technique is exploited in the point doubling computation (see Steps 7-8 of Algorithm 11). Note that, in contrast to addition, only the first operand is allowed to be in incompletely reduced form for subtraction.

To guarantee correctness in our software, and following the previous description, incompletely reduced results are always fed to one of the following: one of the operands of an incompletely reduced addition, the first operand of a field subtraction, a field multiplication or squaring (which ultimately produces a completely reduced output), or a field addition without correction preceding a field multiplication or squaring.

In the targeted setting, there are only a limited number of spots in the curve arithmetic in which incompletely reduced numbers cannot be efficiently exploited. For these few cases, we require a standard field addition. We use the efficient implementation proposed by Bos et al. [7]. Again, let  $0 \leq a, b < 2^m - c$ . Field addition is then computed as  $((a + c) + b) - carry \cdot 2^m - (1 - carry) \cdot c$ , where  $carry = 0$  if  $a + b + c < 2^m$ , otherwise  $carry = 1$ . Similar to previous cases, the subtraction with  $carry \cdot 2^m$  can be efficiently carried out by clearing the  $(m + 1)$ -th bit in  $(a + c) + b$ . As discussed above, this efficient computation is also advantageously exploited in the modular reduction for multiplication and squaring.

### 5.3 Quadratic Extension Field Arithmetic

For the remainder, we use the following notation: (i)  $I, M, S, A$  and  $R$  represent field inversion, field multiplication, field squaring, field addition and modular reduction, respectively, (ii)  $M_i$  and  $A_i$  represent integer multiplication and integer addition, respectively, and (iii)  $i, m, s, a$  and  $r$  represent analogous operations over  $\mathbb{F}_{p^2}$ . When representing registers in algorithms, capital letters are used to allocate operands with “double precision” (in our case, 256 bits). For simplification purposes, in the operation counting an integer operation with double-precision is considered equivalent to two integer operations with single precision. We assume that addition, subtraction, multiplication by two and negation have roughly the same cost.

Let  $a = a_0 + a_1 i \in \mathbb{F}_{p^2}$  and  $b = b_0 + b_1 i \in \mathbb{F}_{p^2}$ . Inversion over  $\mathbb{F}_{p^2}$  is computed as  $a^{-1} = (a_0 - a_1 i) / (a_0^2 + a_1^2)$ . Addition and subtraction over  $\mathbb{F}_{p^2}$  consist in computing  $(a_0 + b_0) + (a_1 + b_1)i$  and  $(a_0 - b_0) + (a_1 - b_1)i$ , respectively. We compute multiplication over  $\mathbb{F}_{p^2}$  using the Karatsuba method. In this case, we fully exploit lazy reduction and the room of one bit that is gained by using a prime of 127 bits. The details for the x64 implementation are shown in Algorithm 6. Remarkably, note that only the subtraction in Step 3 requires a correction to produce a positive result. No other addition or subtraction requires correction to positive or to modulo  $p$ . That is,  $\times, +$  and  $-$  represent operations over the integers. In addition, the algorithm accepts

inputs in completely or incompletely reduced form and always produces a result in completely reduced form. Optionally, one may “delay” the computation of the final modular reductions (by setting  $reduction = FALSE$  in Alg. 6) if lazy reduction could be exploited in the curve arithmetic. This has been proven to be useful to formulas for the Weierstrass form [1], but unfortunately the technique cannot be advantageously exploited in the most efficient formulas for Twisted Edwards (in this case, one should set  $reduction = TRUE$ ). Squaring over  $\mathbb{F}_{p^2}$  is computed using the complex method. The details for the x64 implementation are shown in Algorithm 7. In this case, all the additions are computed as integer operations since, again, results can be let to grow up to 128 bits, letting subsequent multiplications take care of the reduction step.

---

**Algorithm 6** Multiplication in  $\mathbb{F}_{p^2}$  with reduction ( $m = 3M_i + 9A_i + 2R$ ) and without reduction ( $m_u = 3M_i + 9A_i$ ), using completely or incompletely reduced inputs (x64 platform).

---

**Input:**  $a = (a_0 + a_1i)$  and  $b = (b_0 + b_1i) \in \mathbb{F}_{p^2}$ , where  $0 \leq a_0, a_1, b_0, b_1 \leq 2^{127} - 1, p = 2^{127} - c, c$  small.

**Output:**  $a \cdot b \in \mathbb{F}_{p^2}$ .

---

1: $T_0 \leftarrow a_0 \times b_0$	[0, $2^{254}$ >
2: $T_1 \leftarrow a_1 \times b_1$	[0, $2^{254}$ >
3: $C_0 \leftarrow T_0 - T_1$	< $-2^{254}, 2^{254}$ >
4: <b>if</b> $C_0 < 0$ , <b>then</b> $C_0 \leftarrow C_0 + 2^{128} \cdot p$	[0, $2^{255}$ >
5: <b>if</b> $reduction = TRUE$ , <b>then</b> $c_0 \leftarrow C_0 \bmod p$	[0, $p$ >
6: $t_0 \leftarrow a_0 + a_1$	[0, $2^{128}$ >
7: $t_1 \leftarrow b_0 + b_1$	[0, $2^{128}$ >
8: $T_2 \leftarrow t_0 \times t_1$	[0, $2^{256}$ >
9: $T_2 \leftarrow T_2 - T_0$	[0, $2^{256}$ >
10: $C_1 \leftarrow T_2 - T_1$	[0, $2^{256}$ >
11: <b>if</b> $reduction = TRUE$ , <b>then</b> $c_1 \leftarrow C_1 \bmod p$	[0, $p$ >
12: <b>return</b> <b>if</b> $reduction = TRUE$ <b>then</b> $a \cdot b = (c_0 + c_1i)$ , <b>else</b> $a \cdot b = (C_0 + C_1i)$ .	

---

**Algorithm 7** Squaring in  $\mathbb{F}_{p^2}(s = 2M + 1A + 2A_i)$ , using completely reduced inputs (x64 platform).

---

**Input:**  $a = (a_0 + a_1i) \in \mathbb{F}_{p^2}$ , where  $0 \leq a_0, a_1 \leq p - 1, p = 2^{127} - c, c$  small.

**Output:**  $a^2 \in \mathbb{F}_{p^2}$ .

---

1: $t_0 \leftarrow a_0 + a_1$	[0, $2^{128}$ >
2: $t_1 \leftarrow a_0 - a_1 \bmod p$	[0, $p$ >
3: $c_0 \leftarrow t_0 \times t_1 \bmod p$	[0, $p$ >
4: $t_0 \leftarrow a_0 + a_0$	[0, $2^{128}$ >
5: $c_1 \leftarrow t_0 \times a_1 \bmod p$	[0, $p$ >
6: <b>return</b> $a^2 = (c_0 + c_1i)$ .	

---

#### 5.4 Extension Field Arithmetic on ARM: Efficient Interleaving of ARM-Based and NEON-Based Multiprecision Operations

The potential performance gain when interleaving ARM and NEON operations is well-known. This feature was exploited in [6] to speed up the Salsa20 stream cipher. On the other hand, Sánchez and Rodríguez-Henríquez [34] showed how to take advantage of NEON instructions to perform independent multiplications in operations over  $\mathbb{F}_{p^2}$ . In the following, we go a step further and show how to exploit the increasingly efficient capacity of modern ARM processors for executing ARM and NEON instructions “simultaneously” to implement multiprecision operations, such as multiplication, squaring and modular reduction, over  $\mathbb{F}_{p^2}$ . In other words, we exploit the fact that when ARM code produces a data hazard in the pipeline, the NEON unit may be ready to execute vector instructions, and vice versa. Note that loading/storing values from ARM to NEON registers still remains relatively expensive, so in order to achieve an effective performance improvement, one should carefully interleave *independent* operations while minimizing the loads and stores from one unit to the other. Hence, operations such as multiplication and squaring over  $\mathbb{F}_{p^2}$  are particularly friendly to this technique, given the availability of internal independent multiplications in their formulas.

---

**Algorithm 8** Double 128-bit integer product with ARM and NEON interleaved (`double_mul_neonarm`).

---

**Input:**  $a = \{a_i\}, b = \{b_i\}, c = \{c_i\}, d = \{d_i\}, i \in \{0, \dots, 3\}$ .

**Output:**  $(F, G) \leftarrow (a \times b, c \times d)$ .

---

```

1:  $(F, G) \leftarrow (0, 0)$ 
2: for  $i = 0$  to 1 do
3:    $(C_0, C_1, C_2) \leftarrow (0, 0, 0)$ 
4:   for  $j = 0$  to 3 do
5:      $(C_0, F_{i+j}, C_1, F_{i+j+2}) \leftarrow (F_{i+j} + a_i b_j + C_0, F_{i+j+2} + a_{i+2} b_j + C_1)$            {done by NEON}
6:   for  $j = 0$  to 3 do
7:      $(C_2, G_{i+j}) \leftarrow G_{i+j} + c_j d_i + C_2$                                            {done by ARM}
8:    $(F_{i+4}, F_{i+6}, G_{i+4}) \leftarrow (F_{i+4} + C_0, C_1, C_2)$ 
9: for  $i = 2$  to 3 do
10:  for  $j = 0$  to 3 do
11:     $(C_2, G_{i+j}) \leftarrow G_{i+j} + c_j d_i + C_2$                                            {done by ARM}
12:   $G_{i+4} \leftarrow C_2$ 
13: return  $(F, G)$ 

```

---



---

**Algorithm 9** Triple 128-bit integer product with ARM and NEON interleaved (`triple_mul_neonarm`).

---

**Input:**  $a = \{a_i\}, b = \{b_i\}, c = \{c_i\}, d = \{d_i\}, e = \{e_i\}, f = \{f_i\}, i \in \{0, \dots, 3\}$ .

**Output:**  $(F, G, H) \leftarrow (a \times b, c \times d, e \times f)$ .

---

```

1:  $(F, G, H) \leftarrow (0, 0, 0)$ 
2: for  $i = 0$  to 3 do
3:    $(C_0, C_1, C_2) \leftarrow (0, 0, 0)$ 
4:   for  $j = 0$  to 3 do
5:      $(C_0, F_{i+j}, C_1, G_{i+j}) \leftarrow (F_{i+j} + a_j b_i + C_0, G_{i+j} + c_j d_i + C_1)$            {done by NEON}
6:   for  $j = 0$  to 3 do
7:      $(C_2, H_{i+j}) \leftarrow H_{i+j} + e_j f_i + C_2$                                            {done by ARM}
8:    $(F_{i+4}, G_{i+4}, H_{i+4}) \leftarrow (C_0, C_1, C_2)$ 
9: return  $(F, G, H)$ 

```

---



---

**Algorithm 10** Double modular reduction with ARM and NEON interleaved (`double_red_neonarm`).

---

**Input:** A prime  $p = 2^{127} - c$ ,  $a = \{a_i\}, b = \{b_i\}, i \in \{0, \dots, 7\}$ .

**Output:**  $(F, G) \leftarrow (a \bmod p, b \bmod p)$ .

---

```

1:  $(F_i, G_i) \leftarrow (a_i, b_i)_{i \in \{0, \dots, 3\}}$ 
2:  $(C_0, C_1, C_2) \leftarrow (0, 0, 0)$ 
3: for  $j = 0$  to 1 do
4:    $(C_0, F_j, C_1, F_{j+2}) \leftarrow (F_j + a_{j+4} c + C_0, F_{j+2} + a_{j+6} c + C_1)$            {done by NEON}
5: for  $j = 0$  to 3 do
6:    $(C_2, G_j) \leftarrow G_j + b_{j+4} c + C_2$                                            {done by ARM}
7:  $(F_2, F_4, G_4) \leftarrow (F_2 + C_0, C_1, C_2)$ 
8:  $(F_0, G_0) \leftarrow (F_4 c + F_0, G_4 c + G_0)$ 
9: return  $(F, G)$ 

```

---



Thus, using this approach, we implemented:

- a double integer multiplier (`double_mul_neonarm`) detailed in Algorithm 8, which interleaves a single 128-bit multiplication using NEON and a single 128-bit multiplication using ARM,
- a triple integer multiplier (`triple_mul_neonarm`) detailed in Algorithm 9, which interleaves two single 128-bit multiplication using NEON and one single 128-bit multiplication using ARM, and
- a double reduction algorithm (`double_red_neonarm`) detailed in Algorithm 10, that interleaves a single modular reduction using NEON and a single modular reduction using ARM.

Note that integer multiplication is implemented using the schoolbook method, which requires one multiplication, two additions, one shift and one bit-wise AND operation per iteration. These operations were implemented using efficient fused instructions such as UMLAL, UMAAL, VMLAL and VSRA [26], which add the result of a multiplication or shift operation to the destination register in one single operation, reducing code size.

To validate the efficiency of our approach, we compared the interleaved algorithms above with standard implementations using only NEON or ARM. In all the cases, we observed a reduction of costs in favor of our novel interleaved ARM/NEON implementations (see Section 6 for benchmark results).

`Triple_mul_neonarm` is nicely adapted to the computation of multiplication over  $\mathbb{F}_{p^2}$ , since this operation requires three integer multiplications of 128 bits (Steps 1, 2 and 8 of Algorithm 6). For the case of squaring over  $\mathbb{F}_{p^2}$ , we use `double_mul_neonarm` to compute the two independent integer multiplications (Steps 3 and 5 of Algorithm 7). Finally, for each case we can efficiently use a `double_red_neonarm`. The final algorithms for ARM are shown as Algorithms 13 and 14 in Appendix C.

## 5.5 Point Arithmetic

In this section, we describe implementation details and our optimized formulas for the point arithmetic. We use as basis the most efficient set of formulas proposed by Hisil et al. [20], corresponding to the case  $a = -1$ , that uses a combination of homogeneous projective coordinates  $(X : Y : Z)$  and extended homogeneous coordinates of the form  $(X : Y : Z : T)$ , where  $T = XY/Z$ .

The basic algorithms for computing point doubling and addition are shown in Algorithms 11 and 12, respectively. In these algorithms, we extensively exploit incomplete reduction, following the details given in Section 5.2. To ease coupling of doubling and addition in the main loop of the scalar multiplication computation, we make use of Hamburg’s “extensible” strategy and output values  $\{T_a, T_b\}$ , where  $T = T_a \cdot T_b$ , at every point operation, so that a subsequent operation may compute coordinate  $T$  if required. Note that the cost of doubling is given by  $4m + 3s + 5a$ . We do not apply the usual transformation  $2XY = (X + Y)^2 - (X^2 + Y^2)$  because in our case it is faster to compute one multiplication and one incomplete addition than one squaring, one subtraction and one addition. In the setting of variable-base scalar multiplication (see Alg. 2), the main loop of the evaluation stage consists of a doubling-addition computation, which corresponds to the successive execution of Algorithms 11 and 12. For this case, precomputed points are more efficiently represented as  $(X + Y, Y - X, 2Z, 2T)$  (corresponding to setting `EXT_COORD = TRUE` in Alg. 12), so the cost of addition is given by  $8m + 6a$ . In the fixed-base scenario, the main loop of the evaluation stage consists of one doubling and  $v$  mixed additions. In this case, we consider two possibilities: representing precomputations as  $(x, y)$  or as  $(x + y, y - x, 1, 2t)$ , where  $t = xy$ . The latter case (also corresponding to setting `EXT_COORD = TRUE`, but with  $Z = 1$ ) allows saving four additions and one multiplication per iteration (Steps 2, 8 and 11 of Alg. 12), but increases the memory requirements to store the additional coordinate. Hence, each option ends up being optimal for certain storage values. We evaluate these options in Section 6. For the case  $(x, y)$ , mixed addition costs  $8m + 10a$  and, for the case  $(x + y, y - x, 1, 2t)$ , mixed addition costs  $7m + 7a$ . In the variable/fixed-base double scalar scenario, precomputed points corresponding to the variable base are stored as  $(X + Y, Y - X, 2Z, 2T)$  and, thus, addition with these points costs  $8m + 6a$ ; whereas points corresponding to the fixed base can again be represented as  $(x, y)$  or as  $(x + y, y - x, 1, 2t)$ , following the same trade-offs and costs for mixed addition discussed above. These options are also evaluated in Section 6.

---

**Algorithm 11** Twisted Edwards point doubling over  $\mathbb{F}_{p^2}$  (DBL =  $4m + 3s + 5a$ ).

---

**Input:**  $P = (X_1, Y_1, Z_1)$ .

**Output:**  $2P = (X_2, Y_2, Z_2)$  and  $\{T_a, T_b\}$  such that  $T_2 = T_a \cdot T_b$ .

---

1: $T_a \leftarrow X_1^2$	$(X_1^2)$
2: $t_1 \leftarrow Y_1^2$	$(Y_1^2)$
3: $T_b \leftarrow T_a \oplus t_1$	$(X_1^2 + Y_1^2)$
4: $T_a \leftarrow t_1 - T_a$	$(Y_1^2 - X_1^2)$
5: $Y_2 \leftarrow T_b \times T_a$	$(Y_2 = (X_1^2 + Y_1^2)(Y_1^2 - X_1^2))$
6: $t_1 \leftarrow Z_1^2$	$(Z_1^2)$
7: $t_1 \leftarrow t_1 \oplus t_1$	$(2Z_1^2)$
8: $t_1 \leftarrow t_1 \ominus T_a$	$(2Z_1^2 - (Y_1^2 - X_1^2))$
9: $Z_2 \leftarrow T_a \times t_1$	$(Z_2 = (Y_1^2 - X_1^2)[2Z_1^2 - (Y_1^2 - X_1^2)])$
10: $T_a \leftarrow X_1 \oplus X_1$	$(2X_1)$
11: $T_a \leftarrow T_a \times Y_1$	$(2X_1 Y_1)$
12: $X_2 \leftarrow T_a \times t_1$	$(X_2 = 2X_1 Y_1 [2Z_1^2 - (Y_1^2 - X_1^2)])$
13: <b>return</b> $2P = (X_2, Y_2, Z_2)$ and $\{T_a, T_b\}$ such that $T_2 = T_a \cdot T_b$ .	

---



---

**Algorithm 12** Twisted Edwards point addition over  $\mathbb{F}_{p^2}$  (ADD =  $8m + 6a$ , mADD =  $7m + 7a$  or  $8m + 10a$ ).

---

**Input:**  $P = (X_1, Y_1, Z_1)$  and  $\{T_a, T_b\}$  such that  $T_1 = T_a \cdot T_b$ . If  $EXT\_COORD = FALSE$  then  $Q = (x_2, y_2)$ , else  $Q = (X_2 + Y_2, Y_2 - X_2, 2Z_2, 2T_2)$ .

**Output:**  $P + Q = (X_3, Y_3, Z_3)$  and  $\{T_a, T_b\}$  such that  $T_3 = T_a \cdot T_b$ .

---

1: $T_1 \leftarrow T_a \times T_b$	$(T_1)$
2: <b>if</b> $EXT\_COORD = FALSE$ <b>then</b> $T_2 = x_2 \oplus x_2, T_2 = T_2 \times y_2$	$(2T_2)$
3: $t_1 \leftarrow T_2 \times Z_1$	$(2T_2 Z_1)$
4: <b>if</b> $Z_2 = 1$ <b>then</b> $t_2 \leftarrow T_1 \oplus T_1$ <b>else</b> $t_2 \leftarrow T_1 \times Z_2$	$(2T_1 Z_2)$
5: $T_a \leftarrow t_2 - t_1$	$(T_a = \alpha = 2T_1 Z_2 - 2T_2 Z_1)$
6: $T_b \leftarrow t_1 \oplus t_2$	$(T_b = \theta = 2T_1 Z_2 + 2T_2 Z_1)$
7: $t_2 \leftarrow X_1 \oplus Y_1$	$(X_1 + Y_1)$
8: <b>if</b> $EXT\_COORD = TRUE$ <b>then</b> $Y_3 = Y_2$ , <b>else</b> $Y_3 = y_2 - x_2$	$(Y_2 - X_2)$
9: $t_2 \leftarrow Y_3 \times t_2$	$(X_1 + Y_1)(Y_2 - X_2)$
10: $t_1 \leftarrow Y_1 - X_1$	$(Y_1 - X_1)$
11: <b>if</b> $EXT\_COORD = TRUE$ <b>then</b> $X_3 = X_2$ , <b>else</b> $X_3 = x_2 \oplus y_2$	$(X_2 + Y_2)$
12: $t_1 \leftarrow X_3 \times t_1$	$(X_2 + Y_2)(Y_1 - X_1)$
13: $Z_3 \leftarrow t_2 - t_1$	$\beta = (X_1 + Y_1)(Y_2 - X_2) - (X_2 + Y_2)(Y_1 - X_1)$
14: $t_1 \leftarrow t_1 \oplus t_2$	$\omega = (X_1 + Y_1)(Y_2 - X_2) + (X_2 + Y_2)(Y_1 - X_1)$
15: $X_3 \leftarrow T_b \times Z_3$	$(X_3 = \beta\theta)$
16: $Z_3 \leftarrow t_1 \times Z_3$	$(Z_3 = \beta\omega)$
17: $Y_3 \leftarrow T_a \times t_1$	$(Y_3 = \alpha\omega)$
18: <b>return</b> $P + Q = (X_3, Y_3, Z_3)$ and $\{T_a, T_b\}$ such that $T_3 = T_a \cdot T_b$ .	

---

## 6 Performance Analysis and Experimental Results

In this section, we carry out the performance analysis of the proposed GLV-based scalar multiplication method using the GLV-SAC representation, and present benchmark results of our constant-time implementations of curve `Ted127-glv4` on x64 and ARM platforms. We also assess the performance improvement obtained with the proposed ARM/NEON interleaving technique. For our experiments, we targeted a 3.4GHz Intel Core i7-2600 Sandy Bridge processor and a 3.4GHz Intel Core i7-3770 Ivy Bridge processor, from the Intel family, and a Samsung Galaxy Note with a 1.4GHz Exynos 4 Cortex-A9 processor and an Arndale

Board with a 1.7GHz Exynos 5 Cortex-A15 processor, from the ARM family, both equipped with the NEON vector unit. The x64 implementation was compiled with Microsoft Visual Studio 2012 and ran on 64-bit Windows (Microsoft Windows 8 OS). In our experiments, we turned off hyperthreading and Intel’s Turbo Boost; we averaged the cost of  $10^4$  operations which were measured with the timestamp counter instruction `rdtsc`. The ARM implementation was developed and compiled with the Android NDK (ndk8d) toolkit. In this case, we averaged the cost of  $10^4$  operations which were measured with the `clock_gettime()` function and scaled to clock cycles using the processor frequency.

First, we present timings for all the fundamental operations of scalar multiplication in Table 1. Implementation details for quadratic extension field operations and point operations over  $\mathbb{F}_{p^2}$  can be found in Section 5. “IR” stands for incomplete reduction and “extended” represents the use of the extended coordinates  $(X + Y, Y - X, 2Z, 2T)$  to represent precomputed points.

Next, we analyze the different scalar multiplication scenarios on curve `Ted127-glv4`.

**Variable-Base Scenario.** Based on Algorithm 2, scalar multiplication on curve `Ted127-glv4` involves the computation of one  $\Phi$  endomorphism, 2  $\Psi$  endomorphisms, 3 additions and 4 mixed additions in the precomputation stage; 63 doublings, 63 additions, one mixed addition and 64 protected table lookups in the evaluation stage; and one inversion and 2 multiplications over  $\mathbb{F}_{p^2}$  for converting the final result to affine:

$$\text{COST}_{\text{variable\_kP}} = 1i + 833m + 191s + 769a + 64LUT8 + 4M + 9A.$$

This operation count does not take into account other additional computations, such as the recoding to the GLV-SAC representation or the decomposition to 4-GLV, which are relatively inexpensive (see Table 1).

Compared to [28], which uses a method based on the regular windowed recoding [31], the optimized GLV-SAC method for variable-base scalar multiplication introduces a reduction in about 181 multiplications, 26 squarings and 228 additions over  $\mathbb{F}_{p^2}$ . Additionally, it only requires 8 precomputed points, which involve 64 protected table lookups over 8 points (denoted by `LUT8`) during scalar multiplication, whereas the method in [28] requires 36 precomputed points, which involve 68 protected table lookups over 9 points. For example, this represents in practice a 17% speedup in the computation and a 78% reduction in the memory consumption of precomputation on curve `Ted127-glv4`.

**Fixed-Base Scenario.** In this case, we analyze costs when using the  $m$ LSB-set comb method (Algorithm 5) with 32, 64, 128, 256 and 512 precomputed points. Recalling Section 4, and since  $t = 251$ , the method costs  $\lceil \frac{251}{w \cdot v} \rceil - 1$  doublings and  $v \lceil \frac{251}{w \cdot v} \rceil - 1$  additions using  $v \cdot 2^{w-1}$  points. Again, there are two options: storing points as  $(x, y)$  (affine coordinates) or as  $(x + y, y - x, 2t)$  (“extended” affine coordinates). We show in Table 4, Appendix D, the costs in terms of multiplications over  $\mathbb{F}_{p^2}$  per bit for curve `Ted127-glv4`. Best results for a given memory requirement are highlighted. And, as can be seen, on the target platforms each precomputation representation is optimal for determined storage values. The latter is true in most cases; however, for high memory values, the results are mixed and depend on the platform.

**Variable/Fixed-Base Double Scalar Scenario.** This computation can be efficiently computed using  $w$ NAF with interleaving [14, 30]. Since precomputation for the fixed base is performed offline, one may arbitrarily increase the window size for this case, only taking into consideration any memory restriction. When using the  $m$ -dimensional GLV, the two scalars can be split in two sets of  $m$  sub-scalars with maximum bitlength  $l = \lceil \log_2 r/m \rceil$  each. The cost is given by  $m \cdot (\frac{l}{w_1+1}) - 1$  additions,  $m \cdot (\frac{l}{w_2+1})$  mixed additions and  $(l - 1)$  doublings using  $m \cdot (2^{w_1-2} + 2^{w_2-2})$  precomputed points, where  $w_1$  is the window size for the variable base and  $w_2$  is the window size for the fixed base. Note that it is possible to choose different window sizes for each sub-scalar, giving more flexibility in the selection of the optimal number of precomputations. In this direction, we analyze the cost of `Ted127-glv4` for different values of  $w_{2,j}$  (corresponding to each of the  $j$  sub-scalars of the fixed base, where  $0 \leq j < 4$ ), with the optimal value  $w_1 = 4$  for the sub-scalars of the variable base. This value for  $w_1$  was determined during experimentation on the targeted platforms. Let  $l = 63$  be the maximum bitlength of the eight sub-scalars. The computation approximately involves 62 doublings, 48 additions and  $(\sum_{j=0}^3 \frac{63}{w_{2,j}+1}) + 3$  mixed additions in the evaluation stage using 16 “ephemeral”

precomputed points and  $\sum_{j=0}^3 2^{w_{2,j}-2}$  “permanent” precomputed points; 2 doublings, 6 additions, 8  $\Psi$  endomorphisms and one  $\Phi$  endomorphism in the online precomputation stage; and one inversion with 2 multiplications over  $\mathbb{F}_{p^2}$  to convert the final result to affine. Again, we examine storing points as  $(x, y)$  or as  $(x + y, y - x, 2t)$ . We show in Table 5, Appendix E, the costs in terms of multiplications over  $\mathbb{F}_{p^2}$  per bit for curve **Ted127-g1v4**. In this case, extended coordinates offer a higher performance in all cases. This is mainly due to the reduced cost for extracting points from the precomputed table, which is not required to be performed in constant time in this scenario.

**Table 1.** Cost (in cycles) of basic operations on curve **Ted127-g1v4**.

Operation		ARM Cortex-A9	ARM Cortex-A15	Intel Sandy Bridge	Intel Ivy Bridge
$\mathbb{F}_{p^2}$	ADD with IR	20	19	12	12
	ADD	39	37	15	15
	SUB	39	37	12	12
	SQR	223	141	59	56
	MUL	339	185	78	75
	INV	13,390	9,675	6,060	5,890
ECC	DBL	2,202	1,295	545	525
	ADD	3,098	1,831	690	665
	mADD ( $Z_1 = 1$ )	2,943	1,687	622	606
	$\phi$ endomorphism ( $Z_1 = 1$ )	3,118	1,724	745	712
	$\psi$ endomorphism ( $Z_1 = 1$ )	1,644	983	125	119
Misc	8-point LUT (extended)	291	179	83	79
	GLV-SAC recoding	1,236	873	482	482
	4-GLV decomposition	756	430	305	290

In Table 2 we summarize our benchmark results for all the core scalar multiplication operations: variable-base, fixed-base and variable/fixed-base double scalar scenario. In Table 3 we compare them with other implementations in the literature for the variable-base scenario. The reader should note that we only compare timings for constant-time implementations. Hence, results are not directly comparable with an analysis based on *unprotected* implementations [13, 12].

The results for the representative variable-base scenario set a new speed record for protected curve-based scalar multiplication on x64 and ARM processors. In comparison with the previously fastest x64 implementation by Longa and Sica [28], which runs in 137,000 cycles, the presented result injects a cost reduction of 30% on a Sandy Bridge machine. Likewise, in comparison with the state-of-the-art genus 2 implementation by Bos et al. [7], our results are between 21%-24% faster on x64 processors. It is also between 17%-19% faster than the very recent implementation by Oliveira et al. [32] based on a binary GLS curve using the 2-GLV method, and about 2 times faster than Bernstein et al.’s implementation using a Montgomery curve over  $\mathbb{F}_p$  [4]. Moreover, our results also demonstrate that the proposed techniques bring a dramatic reduction in the overhead for protecting against timing attacks. An unprotected version of our implementation computes a scalar multiplication in 87,000 cycles on the Sandy Bridge processor, which is only 9% faster than our protected version. In the case of ARM, our implementation of variable-base scalar multiplication on curve **Ted127-g1v4** is 27% and 32% faster than Bernstein and Schwabe’s [5] and Hamburg’s [16] implementation (respect.) of **curve25519** on a Cortex-A9 processor. Note, however, that comparisons on ARM are particularly difficult. The implementation of [6] was originally optimized for Cortex-A8, and the implementation of [16] does not exploit NEON.

We achieve similar results in the fixed-base and double scalar scenarios. For instance:

- Hamburg [16] computes key generation (dominated by a fixed-base scalar multiplication) in 60K cycles on a Sandy Bridge and 254K cycles on a Cortex-A9 (without NEON) using a table of size 7.5KB. Using only 6KB, our software runs a fixed-base scalar multiplication in 54K cycles and 204K cycles, respectively.

- Hamburg [16] computes signature verification (dominated by a double scalar multiplication) in 169K cycles on a Sandy Bridge and 618K cycles on a Cortex-A9 (without NEON) using a table of size 3KB. Our software runs a double scalar multiplication in only 123K cycles and 495K cycles, respectively, using the same table size.

Finally, in our experiments to assess the improvement obtained with the ARM/NEON interleaving technique on the Cortex-A9 processor, we observed speedups close to 17% and 24% in comparison with implementations exploiting only ARM or NEON instructions, respectively. Remarkably, for the same figures on the Cortex-A15, we observed speedups in the order of 34% and 35%, respectively. These experimental results confirm the significant performance improvement enabled by the proposed technique, which exploits the increasing capacity of the latest ARM processors for parallelizing ARM and NEON instructions.

**Table 2.** Cost (in  $10^3$  cycles) of core scalar multiplication operations on curve **Ted127-glv4** with full protection against timing-type side-channel attacks at approximately 128-bit security level. Results are approximated to the nearest  $10^3$  cycles.

Scalar Multiplication		ARM	ARM	Intel	Intel
Type	Parameters	Cortex-A9	Cortex-A15	Sandy Bridge	Ivy Bridge
$kP$ , variable base	8 points, 512 bytes, extended	417	244	96	92
$kP$ , fixed base	$v = 4, w = 5$ , 64 points, 6KB, extended	204	116	54	52
	$v = 4, w = 6$ , 128 points, 12KB, extended	181	106	50	49
	$v = 8, w = 6$ , 256 points, 24KB, extended	172	100	48	46
$kP + lQ$	$w_2 = 3$ , 8 points, 768 bytes, extended	560	321	136	130
	$w_2 = 5$ , 32 points, 3KB, extended	495	285	123	118
	$w_2 = 7$ , 128 points, 12KB, extended	463	266	116	111

**Table 3.** Cost (in  $10^3$  cycles) of implementations of variable-base scalar multiplication with protection against timing-type side-channel attacks at approximately 128-bit security level. Results are approximated to the nearest  $10^3$  cycles.

Work		ARM	ARM	Intel	Intel
Curve	Precomputations	Cortex-A9	Cortex-A15	Sandy Bridge	Ivy Bridge
Ted127-glv4 (this work)	512 bytes (8 points)	417	244	96	92
Ted127-glv4, Longa-Sica [28]	2.25 KB (36 points)	-	-	137	-
Binary GLS $E/\mathbb{F}_{2^{254}}$ , Oliveira et al. [32]	512 bytes (8 points)	-	-	115	113
Genus 2 Kummer $C/\mathbb{F}_p$ , Bos et al. [8]	0	-	-	126 (*)	117
Curve25519, Bernstein et al. [4]	0	-	-	194 (*)	183 (*)
Curve25519, Bernstein et al. [6]	0	568 (*)	-	-	-
Curve25519, Hamburg [16]	0	616	-	153	-

(\*) Source: eBACS [5].

## References

1. D.F. Aranha, K. Karabina, P. Longa, C. Gebotys, and J. López. Faster explicit formulas for computing pairings over ordinary curves. In *Advances in Cryptology - EUROCRYPT*, volume 6632, pages 48–68. Springer, 2011.
2. D. Bernstein. Cache-timing attacks on AES. 2005. <http://cr.yyp.to/antiforgery/cachetiming-20050414.pdf>.
3. D. Bernstein, P. Birkner, M. Joye, T. Lange, and C. Peters. Twisted Edwards curves. In S. Vaudenay, editor, *Proceedings of Africacrypt 2008*, volume 5023 of *LNCS*, pages 389–405. Springer, 2008.

4. D. Bernstein, N. Duif, T. Lange, P. Schwabe, and B.-Y. Yang. High-speed high-security signatures. In B. Preneel and T. Takagi, editors, *Proceedings of CHES 2011*, volume 6917 of *LNCS*, pages 124–142. Springer, 2011.
5. D. Bernstein and T. Lange. eBACS: ECRYPT Benchmarking of Cryptographic Systems, accessed on May 21, 2013. <http://bench.cr.yp.to/results-dh.html>.
6. D. Bernstein and P. Schwabe. NEON crypto. In *Cryptographic Hardware and Embedded Systems - CHES 2012*, volume 7428 of *Lecture Notes in Computer Science*, pages 320–339. Springer, 2012.
7. J.W. Bos, C. Costello, H. Hisil, and K. Lauter. Fast cryptography in genus 2. In T. Johansson and P.Q. Nguyen, editors, *Advances in Cryptology - EUROCRYPT*, volume 7881 of *LNCS*, pages 194–210. Springer, 2013. Also in Cryptology ePrint Archive, Report 2012/670, <http://eprint.iacr.org/2012/670>.
8. J.W. Bos, C. Costello, H. Hisil, and K. Lauter. High-performance scalar multiplication using 8-dimensional GLV/GLS decomposition. In *Cryptographic Hardware and Embedded Systems - CHES 2013 (to appear)*, 2013. Also in Cryptology ePrint Archive, Report 2013/146. Available at: <http://eprint.iacr.org/2013/146>.
9. D. Brumley and D. Boneh. Remote timing attacks are practical. In S. Mangard and F.-X. Standaert, editors, *Proceedings of the 12th USENIX Security Symposium*, volume 6225 of *LNCS*, pages 80–94. Springer, 2003.
10. M. Feng, B.B. Zhu, M. Xu, and S. Li. Efficient comb elliptic curve multiplication methods resistant to power analysis. In *Cryptology ePrint Archive, Report 2005/222*, 2005. Available at: <http://eprint.iacr.org/2005/222>.
11. M. Feng, B.B. Zhu, C. Zhao, and S. Li. Signed MSB-set comb method for elliptic curve point multiplication. In K. Chen, R. Deng, X. Lai, and J. Zhou, editors, *Proceedings of Information Security Practice and Experience (ISPEC 2006)*, volume 3903 of *LNCS*, pages 13–24. Springer, 2006.
12. S. D. Galbraith, X. Lin, and M. Scott. Endomorphisms for faster elliptic curve cryptography on a large class of curves. In *J. Cryptology*, volume 24(3), pages 446–469, 2011.
13. S.D. Galbraith, X. Lin, and M. Scott. Endomorphisms for faster elliptic curve cryptography on a large class of curves. In A. Joux, editor, *Advances in Cryptology - EUROCRYPT*, volume 5479 of *LNCS*, pages 518–535. Springer, 2009.
14. R.P. Gallant, J.L. Lambert, and S.A. Vanstone. Faster Point Multiplication on Elliptic Curves with Efficient Endomorphisms. In J. Kilian, editor, *Advances in Cryptology - CRYPTO*, volume 2139 of *LNCS*, pages 190–200. Springer, 2001.
15. A. Guillevic and S. Ionica. Four dimensional GLV via the Weil restriction. In *Cryptology ePrint Archive, Report 2013/311*, 2013. Available at: <http://eprint.iacr.org/2013/311>.
16. M. Hamburg. Fast and compact elliptic-curve cryptography. In *Cryptology ePrint Archive, Report 2012/309*, 2012. Available at: <http://eprint.iacr.org/2012/309>.
17. D. Hankerson, K. Karabina, and A. Menezes. Analyzing the Galbraith-Lin-Scott point multiplication method for elliptic curves over binary fields. *IEEE Trans. Computers*, 58(10):1411–1420, 2009.
18. D. Hankerson, A. Menezes, and S. Vanstone. *Guide to elliptic curve cryptography*. Springer Verlag, 2004.
19. M. Hedabou, P. Pinel, and L. Beneteau. Countermeasures for preventing comb method against SCA attacks. In R. Deng, F. Bao, H. Pang, and J. Zhou, editors, *Proceedings of Information Security Practice and Experience (ISPEC 2005)*, volume 3439 of *LNCS*, pages 85–96. Springer, 2005.
20. H. Hisil, K. Wong, G. Carter, and E. Dawson. Twisted Edwards curves revisited. In J. Pieprzyk, editor, *Advances in Cryptology - ASIACRYPT*, volume 5350 of *LNCS*, pages 326–343. Springer, 2008.
21. Z. Hu, P. Longa, and M. Xu. Implementing 4-dimensional GLV method on GLS elliptic curves with j-invariant 0. *Designs, Codes and Cryptography*, 63(3):331–343, 2012. Also in Cryptology ePrint Archive, Report 2011/315, <http://eprint.iacr.org/2011/315>.
22. M. Joye and M. Tunstall. Exponent recoding and regular exponentiation algorithms. In M. Joye, editor, *Proceedings of Africacrypt 2003*, volume 5580 of *LNCS*, pages 334–349. Springer, 2009.
23. P. C. Kocher. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In N. Kobitz, editor, *Advances in Cryptology - CRYPTO*, volume 1109 of *LNCS*, pages 104–113. Springer, 1996.
24. P. C. Kocher, J. Jaffe, and B. Jun. Differential Power Analysis. In M. Wiener, editor, *Advances in Cryptology - CRYPTO*, volume 1666 of *LNCS*, pages 388–397. Springer, 1999.
25. C.H. Lim and P.J. Lee. More flexible exponentiation with precomputation. In Y. Desmedt, editor, *Advances in Cryptology - CRYPTO*, volume 839 of *LNCS*, pages 95–107. Springer, 1994.
26. ARM Limited. ARM Architecture Reference Manual: ARMv7-A and ARMv7-R edition, 2012.
27. P. Longa and C. Gebotys. Efficient techniques for high-speed elliptic curve cryptography. In S. Mangard and F.-X. Standaert, editors, *Proceedings of CHES 2010*, volume 6225 of *LNCS*, pages 80–94. Springer, 2010.
28. P. Longa and F. Sica. Four-dimensional Gallant-Lambert-Vanstone scalar multiplication. In X. Wang and K. Sako, editors, *Advances in Cryptology - ASIACRYPT*, volume 7658 of *LNCS*, pages 718–739. Springer, 2012.
29. P. Longa and F. Sica. Four-dimensional Gallant-Lambert-Vanstone scalar multiplication. In *Journal of Cryptology (to appear)*, LNCS. Springer, 2013.

30. B. Möller. Algorithms for multi-exponentiation. In S. Vaudenay and A.M. Youssef, editors, *Proceedings of SAC 2001*, volume 2259 of *LNCS*, pages 165–180. Springer, 2001.
31. K. Okeya and T. Takagi. The width- $w$  NAF method provides small memory and fast elliptic curve scalars multiplications against side-channel attacks. In M. Joye, editor, *Proceedings of CT-RSA 2003*, volume 2612 of *LNCS*, pages 328–342. Springer, 2003.
32. T. Oliveira, J. López, D.F. Aranha, and F. Rodríguez-Henríquez. Two is the fastest prime. In *Cryptographic Hardware and Embedded Systems - CHES 2013 (to appear)*, 2013. Also in Cryptology ePrint Archive, Report 2013/131. Available at: <http://eprint.iacr.org/2013/131>.
33. D.A. Osvik, A. Shamir, and E. Tromer. Cache attacks and countermeasures: the case of AES. In D. Pointcheval, editor, *Topics in Cryptology - CT-RSA 2006*, volume 3860, pages 1–20. Springer, 2006.
34. A.H. Sánchez and F. Rodríguez-Henríquez. NEON implementation of an attribute-based encryption scheme. In M. Jacobson, M. Locasto, P. Mohassel, and R. Safavi-Naini, editors, *International Conference on Applied Cryptography and Network Security - ACNS 2013 (to appear)*, 2013. Available at: <http://cacr.uwaterloo.ca/techreports/2013/cacr2013-07.pdf>.
35. B. Smith. Families of fast elliptic curves from  $\mathbb{Q}$ -curves. In *Cryptology ePrint Archive, Report 2013/312*, 2013. Available at: <http://eprint.iacr.org/2013/312>.
36. D. Weber and T.F. Denny. The solution of McCurley’s discrete log challenge. In H. Krawczyk, editor, *Advances in Cryptology - CRYPTO*, volume 1462 of *LNCS*, pages 458–471. Springer, 1998.
37. T. Yanik, E. Savaş, and Ç.K. Koç. Incomplete reduction in modular arithmetic. In *IEE Proc. of Computers and Digital Techniques*, volume 149(2), pages 46–52, 2002.
38. S.-M. Yen and M. Joye. Checking before output may not be enough against fault- based cryptanalysis. *IEEE Trans. Computers*, 49(9):967–970, 2000.
39. S.-M. Yen, S. Kim, S. Lim, and S.-J. Moon. A countermeasure against one physical cryptanalysis may benefit another attack. In *Information Security and Cryptology - ICISC 2001*, volume 2288 of *Lecture Notes in Computer Science*, pages 414–427. Springer, 2002.

## A Comb Method using the Modified LSB-Set Representation

Let  $t$  be the bitlength of a given scalar  $k$ . Assume that  $k$  is partitioned in  $w$  consecutive parts of  $d$  digits each, and each part is partitioned in  $v$  strings of  $e$  digits each, padding  $k$  with  $(dw - t)$  zeros to the left, where  $l = dw$ ,  $d = ev$  and  $e = \lceil t/wv \rceil$ . The modified LSB-set representation of  $k$  is given by

$$k = c \cdot 2^l + \sum_{i=0}^{l-1} b_i 2^i \equiv (c, b_{l-1}, \dots, b_0)_{m\text{LSB-set}}, \quad (2)$$

where  $b_i \in \{1, -1\}$  for  $0 \leq i < d$ , and  $b_i \in \{0, b_{i \bmod d}\}$  for  $d \leq i \leq l-1$ . If  $wv \mid t$  then the carry bit  $c \in \{0, 1\}$ . Otherwise,  $c$  is always zero. Disregarding the carry bit, rewrite the representation (2) in matrix form as follows

$$k' = k - c \cdot 2^l \equiv \begin{bmatrix} K^0 \\ \vdots \\ K^{w'} \\ \vdots \\ K^{w-1} \end{bmatrix} = \begin{bmatrix} K_{v-1}^0 & \cdots & K_{v'}^0 & \cdots & K_0^0 \\ \vdots & & \vdots & & \vdots \\ K_{v-1}^{w'} & \cdots & K_{v'}^{w'} & \cdots & K_0^{w'} \\ \vdots & & \vdots & & \vdots \\ K_{v-1}^{w-1} & \cdots & K_{v'}^{w-1} & \cdots & K_0^{w-1} \end{bmatrix} \equiv \sum_{w'=0}^{w-1} K^{w'} 2^{dw'} \quad (3)$$

where each  $K^{w'}$  consists of  $v$  strings of  $e$  digits each. Let the  $v'$ -th string in a given  $K^{w'}$  be denoted by  $K_{v'}^{w'}$ , and the  $e'$ -th digit in a given  $K_{v'}^{w'}$  be denoted by  $K_{v',e'}^{w'}$ , such that  $K_{v',e'}^{w'} = b_{dw'+ev'+e'}$ . Then, to compute the scalar multiplication we have

$$k'P = \sum_{w'=0}^{w-1} K^{w'} 2^{dw'} P = \sum_{w'=0}^{w-1} \sum_{v'=0}^{v-1} K_{v'}^{w'} 2^{ev'} 2^{dw'} P = \sum_{w'=0}^{w-1} \sum_{v'=0}^{v-1} \sum_{e'=0}^{e-1} K_{v',e'}^{w'} 2^{e'} 2^{ev'} 2^{dw'} P. \quad (4)$$

Assuming that  $P[w'] = 2^{dw'}P$  for  $0 \leq w' \leq w-1$ , then

$$k'P = \sum_{e'=0}^{e-1} 2^{e'} \left( \sum_{v'=0}^{v-1} \sum_{w'=0}^{w-1} K_{v',e'}^{w'} 2^{ev'} P[w'] \right) = \sum_{e'=0}^{e-1} 2^{e'} \left( \sum_{v'=0}^{v-1} 2^{ev'} \left( K_{v',e'}^0 P[0] + \sum_{w'=1}^{w-1} K_{v',e'}^{w'} P[w'] \right) \right). \quad (5)$$

Recall that by definition of the  $m$ LSB-set representation  $K_{v',e'}^0 \in \{1, -1\}$  and  $K_{v',e'}^{w'} \in \{0, K_{v',e'}^0\}$  for  $1 \leq w' \leq w-1$ , for a given pair of indices  $(v', e')$ . Assume that the following values are precomputed for all  $0 \leq u < 2^{w-1}$  and  $0 \leq v' < v$

$$P[u][v'] = 2^{ev'} (P[0] + u_0 P[1] + \dots + u_{w-2} P[w-1]) = 2^{ev'} (1 + u_0 2^d + \dots + u_{w-2} 2^{(w-1)d}) P, \quad (6)$$

where  $u = (u_{w-2}, \dots, u_0)_2$ . Then,  $k'P$  can be rewritten as

$$k'P = \sum_{e'=0}^{e-1} 2^{e'} \left( \sum_{v'=0}^{v-1} s_{v',e'} P[\mathbb{K}_{v',e'}][v'] \right), \quad (7)$$

where digit-columns  $\mathbb{K}_{v',e'} = [K_{v',e'}^{w-1}, \dots, K_{v',e'}^2, K_{v',e'}^1, K_{v',e'}^0] \equiv |K_{v',e'}^{w-1} 2^{w-2} + \dots + K_{v',e'}^2 + K_{v',e'}^1|$ , and the sign  $s_{v',e'} = K_{v',e'}^0 \in \{1, -1\}$ .

Based on equation (7),  $k'P$  can be computed from left-to-right using precomputed points (6) together with a variant of the double-and-add algorithm (see Algorithm 5). The final result is obtained after a final correction computing  $kP = k'P + c \cdot 2^{wd}P$  using the precomputed value  $2^{wd}P$ .

## B Formulas for Endomorphisms $\Phi$ and $\Psi$ on Curve Ted127-g1v4

Let  $P = (X_1, Y_1, Z_1)$  be a point in homogeneous projective coordinates on a Twisted Edwards curve with eq. (1),  $u = 1 + i$  be a quadratic non-residue in  $\mathbb{F}_{p^2}$ , and  $\zeta_8 = u/\sqrt{-2a'}$  be a primitive 8th root of unity, where  $a' = 27u^3(\frac{\sqrt{2}}{2} - 1)$ . Then, we can compute  $\Phi(P) = (X_2, Y_2, Z_2, T_2)$  as follows

$$\begin{aligned} X_2 &= -X_1 (\alpha Y_1^2 + \theta Z_1^2) [\mu Y_1^2 - \phi Z_1^2], & Y_2 &= 2Y_1 Z_1^2 [\phi Y_1^2 + \gamma Z_1^2], \\ Z_2 &= 2Y_1 Z_1^2 [\mu Y_1^2 - \phi Z_1^2], & T_2 &= -X_1 (\alpha Y_1^2 + \theta Z_1^2) [\phi Y_1^2 + \gamma Z_1^2], \end{aligned}$$

where  $\alpha = \zeta_8^3 + 2\zeta_8^2 + \zeta_8$ ,  $\theta = \zeta_8^3 - 2\zeta_8^2 + \zeta_8$ ,  $\mu = 2\zeta_8^3 + \zeta_8^2 - 1$ ,  $\gamma = 2\zeta_8^3 - \zeta_8^2 + 1$  and  $\phi = \zeta_8^2 - 1$ .

For curve **Ted127-g1v4**, we have the fixed values

$$\begin{aligned} \zeta_8 &= 1 + Ai, & \alpha &= A + 2i, & \theta &= A + Bi, \\ \mu &= (A-1) + (A+1)i, & \gamma &= (A+1) + (A-1)i, & \phi &= (B+1) + i, \end{aligned}$$

where  $A = 143485135153817520976780139629062568752$ ,  $B = 170141183460469231731687303715884099729$ .

Computing an endomorphism  $\Phi$  with the formula above costs  $12m + 2s + 5a$  or only  $8m + 1s + 5a$  if  $Z_1 = 1$ . Similarly, we can compute  $\Psi(P) = (X_2, Y_2, Z_2, T_2)$  as follows

$$X_2 = \zeta_8 X_1^p Y_1^p, \quad Y_2 = Z_1^{p^2}, \quad Z_2 = Y_1^p Z_1^p, \quad T_2 = \zeta_8 X_1^p Z_1^p.$$

Given the value for  $\zeta_8$  on curve **Ted127-g1v4** computing an endomorphism  $\Psi$  with the formula above costs approximately  $3m + 1s + 2M + 5A$  or only  $1m + 2M + 4A$  if  $Z_1 = 1$ .

## C Algorithms for Quadratic Extension Field Operations exploiting Interleaved ARM/NEON Operations

Below are the algorithms for multiplication and squaring over  $\mathbb{F}_{p^2}$ , with  $p = 2^{127} - c$ , for ARM platforms. They exploit functions interleaving ARM/NEON-based operations, namely `triple_mul_neonarm`, `double_mul_neonarm` and `double_red_neonarm`, detailed in Algorithms 9, 8 and 10, respectively.



---

**Algorithm 13** Multiplication in  $\mathbb{F}_{p^2}$  using completely or incompletely reduced inputs,  $m = 3M_i + 9A_i + 2R$  (ARM platform).

---

**Input:**  $a = (a_0 + a_1i)$  and  $b = (b_0 + b_1i) \in \mathbb{F}_{p^2}$ , where  $0 \leq a_0, a_1, b_0, b_1 \leq 2^{127} - 1, p = 2^{127} - c, c$  small.

**Output:**  $a \cdot b \in \mathbb{F}_{p^2}$ .

---

1: $t_0 \leftarrow a_0 + a_1$	[0, $2^{128}$ >
2: $t_1 \leftarrow b_0 + b_1$	[0, $2^{128}$ >
3: $(T_0, T_1, T_2) \leftarrow \text{triple\_mul\_neonarm}(a_0, b_0, a_1, b_1, t_0, t_1)$	[0, $2^{256}$ >
4: $C_0 \leftarrow T_0 - T_1$	< $-2^{254}, 2^{254}$ >
5: <b>if</b> $C_0 < 0$ , <b>then</b> $C_0 \leftarrow C_0 + 2^{128} \cdot p$	[0, $2^{255}$ >
6: $T_2 \leftarrow T_2 - T_0$	[0, $2^{256}$ >
7: $C_1 \leftarrow T_2 - T_1$	[0, $2^{256}$ >
8: <b>return</b> $(c_0, c_1) \leftarrow \text{double\_red\_neonarm}(C_0, C_1)$	[0, $p$ >

---



---

**Algorithm 14** Squaring in  $\mathbb{F}_{p^2}$  using completely reduced inputs,  $s = 2M + 1A + 2A_i$  (ARM platform).

---

**Input:**  $a = (a_0 + a_1i) \in \mathbb{F}_{p^2}$ , where  $0 \leq a_0, a_1 \leq p - 1, p = 2^{127} - c, c$  small.

**Output:**  $a^2 \in \mathbb{F}_{p^2}$ .

---

1: $t_0 \leftarrow a_0 + a_1$	[0, $2^{128}$ >
2: $t_1 \leftarrow a_0 - a_1 \bmod p$	[0, $p$ >
3: $t_2 \leftarrow a_0 + a_0$	[0, $2^{128}$ >
4: $(C_0, C_1) \leftarrow \text{double\_mul\_neonarm}(t_0, t_1, t_2, a_1)$	[0, $p^2$ >
5: <b>return</b> $a^2 = \text{double\_red\_neonarm}(C_0, C_1)$	[0, $p$ >

---

## D Cost of Fixed-Base Scalar Multiplication using the $m$ LSB-Set Comb Method

Below, we present estimated costs in terms of multiplications over  $\mathbb{F}_{p^2}$  per bit for fixed-based scalar multiplication on curve **Ted127-glv4** using the  $m$ LSB-set comb method (Algorithm 5). Precomputed points are stored as  $(x, y)$  coordinates (“affine”) or as  $(x + y, y - x, 2t)$  coordinates (“extended”).

**Table 4.** Cost (in  $\mathbb{F}_{p^2}$  multiplications per bit) of fixed-based scalar multiplication using the  $m$ LSB-set representation on curve **Ted127-glv4**.

$v, w, \#$ of points, memory	precomp coordinates	ARM Cortex-A9	ARM Cortex-A15	Intel Sandy Bridge	Intel Ivy Bridge
1, 6, 32 points, 2KB	affine	2.88	3.24	3.05	3.08
2, 5, 32 points, 2KB		<b>2.66</b>	<b>2.97</b>	<b>2.81</b>	<b>2.83</b>
4, 4, 32 points, 2KB		2.77	3.06	2.91	2.93
1, 6, 32 points, 3KB	extended	2.76	3.11	2.92	2.94
2, 5, 32 points, 3KB		<b>2.46</b>	2.73	<b>2.58</b>	<b>2.60</b>
4, 4, 32 points, 3KB		2.47	<b>2.71</b>	2.58	2.60
2, 6, 64 points, 4KB	affine	2.33	2.63	2.46	2.49
4, 5, 64 points, 4KB		<b>2.32</b>	<b>2.59</b>	<b>2.45</b>	<b>2.47</b>
8, 4, 64 points, 4KB		2.56	2.83	2.68	2.70
2, 6, 64 points, 6KB	extended	2.21	2.50	2.33	2.35
4, 5, 64 points, 6KB		<b>2.12</b>	<b>2.35</b>	<b>2.22</b>	<b>2.24</b>
8, 4, 64 points, 6KB		2.26	2.48	2.36	2.38
2, 7, 128 points, 8KB	affine	2.26	2.60	2.40	2.43
4, 6, 128 points, 8KB		<b>2.07</b>	<b>2.34</b>	<b>2.18</b>	<b>2.21</b>
8, 5, 128 points, 8KB		2.17	2.42	2.28	2.30
2, 7, 128 points, 12KB	extended	2.25	2.60	2.37	2.40
4, 6, 128 points, 12KB		<b>1.95</b>	2.20	<b>2.05</b>	<b>2.07</b>
8, 5, 128 points, 12KB		1.96	<b>2.17</b>	2.05	2.07
4, 7, 256 points, 16KB	affine	2.02	2.34	2.14	2.18
8, 6, 256 points, 16KB		<b>1.94</b>	<b>2.19</b>	<b>2.04</b>	<b>2.07</b>
16, 5, 256 points, 16KB		2.09	2.33	2.19	2.21
4, 7, 256 points, 24KB	extended	2.02	2.34	2.12	2.15
8, 6, 256 points, 24KB		<b>1.82</b>	<b>2.06</b>	<b>1.91</b>	<b>1.93</b>
16, 5, 256 points, 24KB		1.88	2.09	1.96	1.98
8, 7, 512 points, 32KB	affine	1.92	2.22	2.03	2.06
16, 6, 512 points, 32KB		<b>1.86</b>	<b>2.10</b>	<b>1.96</b>	<b>1.98</b>
32, 5, 512 points, 32KB		2.04	2.27	2.14	2.16
8, 7, 512 points, 48KB	extended	1.91	2.22	2.00	2.04
16, 6, 512 points, 48KB		<b>1.75</b>	<b>1.97</b>	<b>1.82</b>	<b>1.85</b>
32, 5, 512 points, 48KB		1.83	2.03	1.91	1.93

## E Cost of Variable/Fixed-Base Double Scalar Multiplication on Curve Ted127-glv4 using $w$ NAF with Interleaving

Below, we present estimated costs in terms of multiplications over  $\mathbb{F}_{p^2}$  per bit for variable/fixed-base double scalar multiplication on curve Ted127-glv4 using  $w$ NAF with interleaving. Precomputations for the fixed base are stored as  $(x, y)$  coordinates (“affine”) or as  $(x+y, y-x, 2t)$  coordinates (“extended”). The window size  $w_{2,j}$  for each sub-scalar  $j$ , number of points and memory listed in the first column correspond to requirements for the fixed base. For the variable base, we fix  $w_1 = 4$ , corresponding to the use of 16 precomputed points (see §6).

**Table 5.** Cost (in  $\mathbb{F}_{p^2}$  multiplications per bit) of variable/fixed-base double scalar multiplication on curve Ted127-glv4 using  $w$ NAF with interleaving,  $w_1 = 4$ , 16 “ephemeral” precomputed points.

$w_{2,j}$ , # of points, memory	precomp coordinates	ARM Cortex-A9	ARM Cortex-A15	Intel Sandy Bridge	Intel Ivy Bridge
2/2/3/3, 6 points, 384 bytes	affine	6.66	7.28	7.11	7.14
2/2/2/2, 4 points, 384 bytes	extended	<b>6.57</b>	<b>7.14</b>	<b>7.00</b>	<b>7.03</b>
3/3/4/4, 12 points, 768 bytes	affine	6.07	6.64	6.51	6.53
3/3/3/3, 8 points, 768 bytes	extended	<b>5.95</b>	<b>6.47</b>	<b>6.36</b>	<b>6.38</b>
4/4/5/5, 24 points, 1.5KB	affine	5.71	6.24	6.13	6.15
4/4/4/4, 16 points, 1.5KB	extended	<b>5.58</b>	<b>6.07</b>	<b>5.98</b>	<b>6.00</b>
5/5/6/6, 48 points, 3KB	affine	5.42	5.92	5.82	5.84
5/5/5/5, 32 points, 3KB	extended	<b>5.33</b>	<b>5.80</b>	<b>5.72</b>	<b>5.74</b>
6/6/7/7, 96 points, 6KB	affine	5.20	5.68	5.59	5.61
6/6/6/6, 64 points, 6KB	extended	<b>5.08</b>	<b>5.53</b>	<b>5.46</b>	<b>5.48</b>
7/7/8/8, 192 points, 12KB	affine	5.05	5.52	5.44	5.46
7/7/7/7, 128 points, 12KB	extended	<b>4.95</b>	<b>5.40</b>	<b>5.33</b>	<b>5.35</b>
8/8/9/9, 384 points, 24KB	affine	4.98	5.44	5.37	5.39
8/8/8/8, 256 points, 24KB	extended	<b>4.83</b>	<b>5.26</b>	<b>5.20</b>	<b>5.22</b>