A Leakage Resilient MAC

D. Martin, E. Oswald, and M. Stam

Department of Computer Science, University of Bristol,
Merchant Venturers Building, Woodland Road,
Bristol, BS8 1UB, United Kingdom.
{dan.martin, elisabeth.oswald, martijn.stam}@bris.ac.uk

Abstract. We put forward the first practical message authentication code (MAC) which is provably secure against continuous leakage under the Only Computation Leaks Information (OCLI) assumption. We introduce a novel, modular proof technique: while most previous schemes are proven secure directly in the face of leakage, we reduce the (leakage) security of our scheme to its non-leakage security. This modularity, while known in other contexts, has two advantages: it makes it clearer which parts of the proof rely on which assumptions (i.e. whether a given assumption is needed for the leakage or the non-leakage security) and it also means that, if the security of the non-leakage version is improved, the security in the face of leakage is improved 'for free'. We feel that this is an advantageous proof technique, providing a better understanding of the scheme's security properties. In practice, we envisage that our scheme would be implemented using pairings on some pairing-friendly elliptic curve, where the 'leakiness' of the group operation can be experimentally estimated. This allows us to compare the resulting instantiation against other leakage resilient MACs (or related schemes), and conclude that ours is the most efficient, as well as being (by far) the most practical.

1 Introduction

Side channel leakage (*e.g.* via timing, power or EM side channels) enables the extraction of secret data out of cryptographic devices, as initially demonstrated by Kocher (*et al.*) in 1996 and 1999 [15, 16]. The engineering community reacted quickly by developing a variety of countermeasures that are commonly described as masking and hiding (see [17]). Such countermeasures intend to *reduce* the overall exploitable leakage via techniques that are cheap to implement.

Initially with hesitance, but more lately with much enthusiasm, the theory community picked up on the fact that schemes are needed which can tolerate some leakage. Complementary to the engineering approach, the aim is to design schemes which do not reduce leakage but cope with it, normally via updating the keys. The most compelling property of this approach is that the security definitions intrinsically incorporate leakage and hence security proofs then hold even in the presence of leakage. The main drawback of having theoretical backing of security seems to be that the resulting schemes are typically considerably less efficient than other schemes. A prime example of such a scheme is the stream cipher by Dziembowski and Pietrzak [6].

Despite the fact that almost all real word cryptographic protocols require some form of authentication, there is a distinct gap in the literature when it comes to leakage resilient message authentication codes (MACs). Hazay *et al.* [12] produce a MAC from minimal assumptions (existence of a one way function). While only relying on minimal assumptions is an advantage from a theoretical perspective, the scheme has a major drawback in that it only allows a bounded amount of leakage (this bound relates to the total leakage of the device). This makes the scheme unsuitable for practice. In his Master's thesis, Schipper [24] discusses a MAC construction in yet another security model. However unfortunately this MAC is also undesirable for practice as the number of AES calls used by verification grows logarithmically in the number of tag queries.

1.1 Our contribution

Inspired by the bilinear ElGamal cryptosystem by Kiltz and Pietrzak [13], we propose a MAC scheme that is secure within the continuous leakage model, using the Only Computation Leaks Information assumption (discussed in Sect. 2). To our knowledge this is the first MAC scheme to be given within this model, which has become one of the more desirable models due to its closer link with practical side channel scenarios.

In Sect. 3 we give our basic MAC construction and prove it secure in the random oracle model without leakage. Unlike previous work (where schemes have to be completely re-proven when considering leakage), we can construct our proof when considering leakage by a reduction to the non-leaky version (see Sect. 4). This is the first proof to achieve such a clean reduction, which has several advantages. Firstly it shows more clearly how much the leakage is impacting on the security of a scheme. This also implies if the security of the basic MAC construction is tightened, the security of the MAC construction with leakage is tightened 'for free'. This manifests itself (as seen in the theorem statement) by having the leakage security bound in terms of the security without leakage. Secondly it becomes clearer which further assumptions are required to prove security when assuming leakage: for example the basic MAC construction requires a Random Oracle assumption, while the Generic Group Model is required when leakage is added.

In Sect. 5 we compare our MAC to the other leakage resilient MACs, as well as other schemes (*i.e.* PRFs, Signatures) which can be converted into MACs. We show that compared to the majority of other provably secure schemes we are considerably more efficient. The only scheme which is comparable with regards to efficiency is a signature scheme [11].

1.2 Related Work

Kiltz and Pietrzak [13] combine two techniques that are commonly used within both communities to build a key encapsulation mechanism on top of a key update scheme. The first technique is masking (or secret sharing as it is known by the theoretical community), which invloves splitting the key into two parts and then working on each share separately. The second technique is frequent rekeying. Unlike other proposals (*e.g.* [14] or [1]), which are stateful (and thus need to be synchronised) or ones which needs to transmit a clue [19] to 'synchronise' parties, [13] can leverage the algebraic properties of the underlying system such that the resulting sytem requires no synchronisation. This is achieved by changing the representation of the shares rather than changing the secret itself. Using the same techniques, Galindo and Vivek [11] create a leakage resilient signature scheme. Both constructions are proven secure in the continuous leakage model using the OCLI assumption [20](see also Sect. 2).

Albeit not related to goal of creating a MAC, there have been several recent papers which design leakage resilient schemes with the balance of provability and useability: schemes that come with some provable guarantee against arbitrary leaks without incurring prohibitively high overhead. When relaxing security notions from completely adaptive inputs (*i.e.* adversaries may choose input messages but also the side channel leakage adaptively) to non-adaptive security, simpler constructions for symmetric key cryptography can be achieved than previously thought [8]. In a differently motivated publication (proving existing schemes secure versus creating provably secure schemes), Balasch *et al.* [2] take a provably secure method, inner product masking, trim it down to implement a masked AES with it, and show this leads to a result which is comparable to other state of the art, yet not formally proven, masking approaches.

Dodis and Pietrzak [5] create a leakage resilient PRF where the leakage functions are chosen non-adaptively before any queries to the PRF are made. Faust *et al.* [8] construct a simpler leakage resilient PRF, which is acheived at the expense of having to make both the input to the PRF and the leakage non-adaptive. All known PRFs in the continual leakage model have the restriction of being non-adaptive (in the leakage), while MACs do not have this restriction. This shows of a seperation between PRFs and MACs which does not exist in the non-leakage model but PRFs will still serve as an interesting comparision.

2 Modelling Leakage

In this section we discuss what assumptions we make when modelling leakage. Clearly some restrictions are required on the leakage, otherwise the adversary will be able to win because he can just ask for the key. One of the first decisions to be made is how to define a bound for the leakage (*i.e.* how many bits about a secret does the adversary get via some side channel). For instance, one could define there to be an overall bound, *i.e.* the adversary gets at most a certain number of bits, irrespective of how often the construction is actually called. Another option would be to impose a per call bound. In this latter case, each call to the construction delivers at most a certain number of bits, while the overall leakage remains unbounded. This type of model is called continuous leakage model and fits best to real world leakage such as power or EM traces.

Whilst some previous works ([6, 9]) make an a priori assumption about the computational complexity of the leakage function, we opted for a concrete security statement. This means that the adversarial advantage is explicitly bounded in the complexity of the leakage function as expressed in the number of queries to the generic group oracles (see Sect. 2.2).

Finally we need to restrict the scope of the leakage function because otherwise (given our choices of assumptions above) no security would be possible (because of the infamous 'future computation attack' [13]), which we discuss in the following.

2.1 Only Computation Leaks Information

Micali and Reyzin [20] introduced the Only Computation Leaks Information (OCLI) assumption. It states that data leakage only occurs on data that is currently being computed on and that data at rest will not leak. Whilst this assumption might not strictly hold in practice ([23] shows it to be invalid for some technologies on gate level), it sufficiently captures the behaviour of many state of the art devices.

Application of the OCLI assumption requires splitting a large computation into smaller components that each only operate on a subset of the data available, thus restricting the scope of what can be leaked on. OCLI will be modelled in this paper by splitting a function F into two parts $F^{\text{◓}}$ and $F^{\text{◒}}$. The part of the sensitive/exploitable input S used by $F^{\text{◓}}$ will be denoted $S^{\text{◓}}$ while the parts of the sensitive input used by $F^{\text{◒}}$ will be denoted $S^{\text{◒}}$. Without OCLI, a leakage query could potentially leak on both shares jointly, and thus reveal information about S. However due to OCLI, any leakage query can only ever leak on $S^{\text{◓}}$ and $S^{\text{◒}}$ independently, but never jointly on both.

Concretely, in our model the adversary may adaptively (per function call) choose leakage functions $l^{\text{◓}}, l^{\text{◒}}$ which will leak up to λ bits (this is a security parameter) on $F^{\text{◓}}$ and $F^{\text{◒}}$ respectively. The adversary also gets the output $l^{\text{◓}}(S^{\text{◓}}, x^{\text{◓}}, r^{\text{◓}})$ and $l^{\text{◒}}(S^{\text{◒}}, x^{\text{◒}}, r^{\text{◒}})$ where $x^{\text{◓}}, x^{\text{◒}}$ is the input to the functions and $r^{\text{◓}}, r^{\text{◒}}$ is the randomness that they use.

Note that while the leakage functions $l^{\text{◓}}$ and $l^{\text{◒}}$ can be chosen adaptively from query to query, they do have to be chosen at the same time for a single query. This mild restriction—that the leakage function $l^{\text{◒}}$ is not allowed to depend on the leakage obtained by $l^{\text{◓}}$—is quite common in the literature [11, 13], and reflects the abilities of a real world adversary (they can't change the measurement set-up mid measurement).

If this leakage process is iterated multiple times an index is used to specify which iteration we are on, for example we use $l_i^{\text{◓}}, l_i^{\text{◒}}, S_i^{\text{◓}}, S_i^{\text{◒}}, r_i^{\text{◓}}, r_i^{\text{◒}}$.

2.2 Bilinear Generic Group Model

We briefly recall the definition of bilinear groups and of bilinear maps, where we adhere to asymmetric pairings (see Galbraith *et al.* [10] for an overview). Let

$\mathbb{G}_1, \mathbb{G}_2$, and $\mathbb{G}_3$ be cyclic groups all of prime order p with generators g_1, g_2, and g_3, respectively. A bilinear map is a function $e : \mathbb{G}_1 \times \mathbb{G}_2 \to \mathbb{G}_3$ with the following properties; bilinearity states that $\forall u \in \mathbb{G}_1, v \in \mathbb{G}_2, a, b \in \mathbb{Z}_p : \ e(u^a, v^b) = e(u, v)^{ab}$, while non-degeneracy $e(g_1, g_2) \neq 1$, stops the construction of trivial maps. From this point onwards we define the generator g_3 of $\mathbb{G}_3$ to be $e(g_1, g_2)$.

The generic group model [18, 21, 25] is well established to prove the security of protocols involving elliptic curves. Its goal is to restrict the adversary in such a way that structure of the underlying group cannot be exploited (beyond what follows from the group axioms). This is achieved by representing each element within the group as a random string and providing oracles for the various group operations. As a consequence, given only a representation of a group element, the only ability the adversary has is to check equality (*i.e.* the adversary must use an oracle to perform any required group operations).

In the Generic Bilinear Group (GBG) model each of the three groups (or two when using a symmetric pairing) has its own randomised encoding. Each of these encodings will be represented by an injective encoding function $\xi_1 : \mathbb{Z}_p \to \Xi_1$, $\xi_2 : \mathbb{Z}_p \to \Xi_2$, $\xi_3 : \mathbb{Z}_p \to \Xi_3$ for $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_3$ respectively, where Ξ_1, Ξ_2, Ξ_3 are sets of bitstrings. The adversary has access to the following 4 oracles:

- $\mathcal{O}_1(\xi_1(a), \xi_1(b)) = \xi_1(a + b \mod p)$
- $\mathcal{O}_2(\xi_2(a), \xi_2(b)) = \xi_2(a + b \mod p)$
- $\mathcal{O}_3(\xi_3(a), \xi_3(b)) = \xi_3(a + b \mod p)$
- $\mathcal{O}_e(\xi_1(a), \xi_2(b)) = \xi_3(a \cdot b \mod p)$

for all $a, b \in \mathbb{Z}_p$. Each of the 4 oracles will return $\perp$ if either of the inputs is not a invalid encoding of an underlying group element. $\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_3$ perform the group operations of $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_3$ respectively, while $\mathcal{O}_e$ performs the pairing operation. To work with these groups an adversary only needs to be given $\xi_1(1)$ and $\xi_2(1)$ (corresponding to the generators of $\mathbb{G}_1$ and $\mathbb{G}_2$ respectively) plus access to the four oracles, from which any group element can be computed.

Leaking on generic group elements only reveals information about their representation. In some proofs (without leakage) that use the generic group model, the representation of group elements can be chosen in such a way that even sampling a random group element is hard (for an adversary). This is typically achieved by representing group elements as 'long' random strings. When leakage is included in proofs, such a strategy would not make sense because it would imply that only 'large' amounts of leakage[1] would strengthen the adversary. We instantiate the generic group model using compact representations instead. By setting $\Xi_i = \{0, 1\}^n$ where $n = \lceil \log p \rceil$ we get the unique representations required. This gives the adversary the ability to sample group elements efficiently and directly.

[1] Typically one would need to leak significantly more than $\log p$ bits, where p would be the size of the group.

experiment $\mathbf{Exp}_{\mathcal{M}}^{\mathrm{euf-cma}}(A)$:
$K \xleftarrow{\$} KG()$
$S \leftarrow \{\}$
$(\sigma^*, m^*) \leftarrow A^{Tag(\cdot), Verify(\cdot,\cdot)}$
if $m^* \in S$ **then**
  return 0
end if
Return $VRFY(K, \sigma^*, m^*)$

proc $Tag(m)$:
$S \leftarrow S \cup \{m\}$
$\sigma \leftarrow TAG(K, m)$
Return σ

proc $Verify(\sigma, m)$:
$b \leftarrow VRFY(K, \sigma, m)$
Return b

Fig. 1. EUF-CMA experiment

In contrast, Kiltz and Pietrzak [13] (and similarly, Galindo and Vivek [11]) use indirect sampling by raising some generator to a random exponent. They allow leakage on both the random representations, as well as their discrete logarithms (relative some generator). To model the adversary's ability to leak on the sampling computation itself. Our proof can be seen as more restrictive and our proofs only hold for implementing the sampling directly. We remark that it is possible to sample random elliptic curve points efficiently without performing an exponentiation with an unknown exponent.

3 A MAC scheme

We define a MAC as a tuple of algorithms $\mathcal{M} = (KG, TAG, VRFY)$ such that:

$$K \xleftarrow{\$} KG()$$
$$\sigma \xleftarrow{\$} TAG(K, m)$$
$$b \leftarrow VRFY(K, \sigma, m).$$

For correctness we require for all valid keys K that $VRFY(K, TAG(K, m), m) = 1$. We use the standard definition of EUF-CMA security for the rest of this section, which is recapped below.

Definition 1 (Existential Unforgability Under Chosen Message Attack (EUF-CMA)).
Let $\mathcal{M} = (KG, TAG, VRFY)$ be a Message Authentication Code. Then Fig. 1 defines the EUF-CMA security game. The advantage of an adversary A winning the game is defined as $\mathbf{Adv}_{\mathcal{M}}^{\mathrm{eufcma}}(A) = \Pr[\mathbf{Exp}_{\mathcal{M}}^{\mathrm{eufcma}}(A) = 1]$.

We now define our basic MAC construction. Using a hash function $H : \{0,1\}^* \times \{0,1\}^k \to \mathbb{G}_2$ our basic MAC scheme $\mathcal{M} = (KG, TAG, VRFY)$ is defined in Fig. 2. It can be shown to provide EUF-CMA security (Thm. 2). The scheme can be understood as follows; key generation consists of generating a random group element of $\mathbb{G}_1$. Tag generation first hashes a message with a random value, then takes the resulting hash as input to a bilinear map, using

proc $KG()$:	proc $TAG(K,m)$:	proc $VRFY(K,(T,w),m)$:
$K \xleftarrow{\$} \mathbb{G}_1$	$w \xleftarrow{\$} \{0,1\}^k$	$W \leftarrow H(m,w)$
Return K	$W \leftarrow H(m,w)$	$T' \leftarrow e(K,W)$
	$T \leftarrow e(K,W)$	Return $T' = T$
	Return (T,w)	

Fig. 2. Our bilinear MAC scheme $\mathcal{M}$

the secret key as other input. The MAC consists of a message, its tag, and the internal randomness used. Verification simply reconstructs the tag T and checks the correctness.

Before we provide the proof of the MAC we introduce a new Bilinear Diffie–Hellman problem, which we will use in the reduction to show the security of the MAC. This new DH problem will have its security 'sandwiched' between two other well known DH problems. We then introduce a vairation of the problem, which makes the proof reduction slightly tidier but will have no effect on the security of the scheme.

3.1 A New Bilinear Diffie–Hellman Problem

In Definition 2 we introduce a bilinear problem, which we coin the target bilinear Diffie–Hellman (TBDH) problem. In Theorem 1 give a reduction to show if Co-Bilinear Diffie–Hellman (CBDH) is assumed to be a hard problem,[2] then so is the TBDH problem. Similarly, it can be shown that if the standard computational Diffie–Hellman (CDH) Problem is easy in $\mathbb{G}_3$ then the TBDH Problem is easy.

Definition 2 (Target Bilinear Diffie–Hellman Problem). *Given $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_3$ with a bilinear map e between them, we say the Target Bilinear Diffie–Hellman (TBDH) Problem is hard if given g_2^x, g_3^y it is hard to compute g_3^{xy}, where x, y are sampled uniformly at random from $\mathbb{Z}_p$. Given an adversary A we define its advantage of winning this game as $\mathbf{Adv}^{tbdh}(A) = \Pr[\mathcal{A} = g_3^{xy} : \mathcal{A} \leftarrow A(g_1, g_2, g_2^x, g_3^y)]$.*

Before relating the TBDH problem to other Diffie–Hellman problems, we recall the CBDH Problem [26]. The CBDH problem states that given g_2^x, g_2^y, where x, y are sampled uniformly at random from $\mathbb{Z}_p$, you must find g_3^{xy}.

Theorem 1. *Let A be an adversary against the TBDH Problem, then there exists an adversary B (with approximately the same runtime as A) against the CBDH Problem, such that:*

$$\mathbf{Adv}^{\mathrm{tbdh}}(A) \leq \mathbf{Adv}^{\mathrm{cbdh}}(B) \ .$$

[2] Where we deliberately leave the notion of hardness informal; of course it is possible to modify our results to the usual notions of negligible advantages against probabilistic polynomial-time adversaries.

Proof. Let adversary A against TBDH be given, then adversary B that breaks CBDH is given in Fig. 3.

adversary $B(g_1, g_2, g_2^a, g_2^b)$:
$g_3^a \leftarrow e(g_1, g_2^a)$
$e(g_1, g_2)^{ab} \leftarrow A(g_1, g_2, g_2^b, g_3^a)$
Return $e(g_1, g_2)^{ab}$

Fig. 3. Constructing a CBDH Adversary from a TBDH Adversary

From this we can see that B will win whenever A does and thus we have $\mathbf{Adv}^{tbdh}(A) \leq \mathbf{Adv}^{cbdh}(B)$. □

We now introduce the TBDHwO problem which will be used in the proof of security for the MAC. While it makes the reduction cleaner it does not weaken the security bound. The TBDHwO problem is the same as the TBDH problem (given g_2^x, g_3^y return g_3^{xy}) with an additional oracle $test(\cdot)$ which given an element checks if it is g_3^{xy}.

Definition 3 (Target Bilinear Diffie-Hellman with Oracle Problem). *We say the Target Bilinear Diffie-Hellman with Oracle (TBDHwO) Problem is hard if given g_2^x, g_3^y it is hard to compute g_3^{xy}, when given access to the $test(\cdot)$ oracle that checks if the given element is g_3^{xy}, where x, y are sampled uniformly at random from $\mathbb{Z}_q$. Given an adversary A we define its advantage of winning this game as $\mathbf{Adv}^{\text{tbdhwo}}(A) = \Pr[\mathcal{A} = g_3^{xy} : \mathcal{A} \leftarrow A^{test(\cdot)}(g_1, g_2, g_2^x, g_3^y)]$*

Lemma 1. *Assuming the TBDH problem is hard then the TBDHwO problem is also hard. Moreover $\mathbf{Adv}^{\text{tbdhwo}}(A) \leq (q_t + 1)\mathbf{Adv}^{\text{tbdh}}(B)$, where q_t is the number of test queries made.*

The proof of Lemma 1 is a standard hybrid argument: since every time A wants a *test* query answered, instead of asking the *test* oracle, he runs a copy of his algorithm and outputs the value he wants testing and based on whether this copy wins or not tells him what value the test function call will return. Since this needs to be done for each test function the bound given holds.

Theorem 2. *Let $H : \{0,1\}^* \times \{0,1\}^k \rightarrow \mathbb{G}_2$ be modelled as a random oracle and A be an EUF-CMA adversary against $\mathcal{M}$ who makes q_h queries to the hash function and q_v verification queries, then there exists an adversary B (of similar computational complexity) against the TBDH problem such that:*

$$\mathbf{Adv}_{\mathcal{M}}^{\text{eufcma}}(A) \leq (q_h + 1)(q_v + 1)\mathbf{Adv}^{\text{tbdh}}(B).$$

Proof. The proof works by reducing the problem of forging the MAC to the problem of solving the TBDH problem; let A be an adversary against the EUF-CMA security of $\mathcal{M}$, Fig. 4 shows the reduction on how an adversary B can solve the TBDHwO problem using the adversary A.

adversary $B(\mathcal{H} = g_2^x, \mathcal{F} = g_3^y)$:
$i \leftarrow 0$
$j \xleftarrow{\$} [q_h]$
$((T, w), m) \leftarrow A^{H(\cdot),Tag(\cdot),VRFY(\cdot)}()$
Return T

simulator $H(m, w)$:
$i \leftarrow i + 1$
if $W[m, w] = \perp$ **then**
  if $i = j$ **then**
    $W[m, w] \leftarrow \times$
    Return $\mathcal{H}$
  else
    $W[m, w] \xleftarrow{\$} \mathbb{Z}_q$
  end if
end if
Return $h^{W[m,w]}$

simulator $Tag(m)$:
$w \xleftarrow{\$} \{0, 1\}^k$
if $W[m, w] = \perp$ **then**
  $W[m, w] \xleftarrow{\$} \mathbb{Z}_q$
else if $W[m, w] = \times$ **then**
  ABORT
end if
$T \leftarrow \mathcal{F}^{W[m,w]}$
Return T

simulator $VRFY((T, w), m)$:
if $W[m, w] = \perp$ **then**
  $W[m, w] \xleftarrow{\$} \mathbb{Z}_q$
else if $W[m, w] = \times$ **then**
  Return $test(T)$
end if
Return $(T = \mathcal{F}^{W[m,w]})$

Fig. 4. Constructing a TBDHwO Adversary from a EUF-CMA Adversary

Without loss of generality we assume that, if A creates a forgery (T, w) on m, it has queried $H(m, w)$ (if it hasn't we can simply create an equivalent adversary that performs the same operations but hashes $H(m, w)$ before outputting the forgery, at the cost of one extra hash query).

From this we can see that when the message–randomness pair forged on was the j^{th} query to the RO, B has won the TBDH game. The reduction constructs tags in such a way that it simulates having the key as g_1^x. If the adversary subsequently forges on a point whose hash is g_2^y, the resulting tag will be the answer to the TBDH problem (g_3^{xy}).

The probability of the message–randomness pair being forged on being the j^{th} RO call is $(q_h+1)^{-1}$. The *ABORT* does not affect this probability, since the only time B will abort is if A tries to tag on the value B wants it to make a forgery on and hence a forgery on this value is no longer possible. Thus $\mathbf{Adv}_{\mathcal{M}}^{\text{eufcma}}(A) \leq (q_h + 1)\mathbf{Adv}^{\text{tbdhwo}}(B)$ and by Lemma 1 the theorem holds (where q_t from Lemma 1 equals q_v). □

4 A Leakage Resilient MAC

We start this section by introducing the definition of a key update mechanism. Kiltz and Pietrzak [13] implicitly constructed and used a key update mechanism within their KEM. This key update mechanism was then used again in the signature scheme by Galindo and Vivek [11]. After showing that our definition aligns with the KP key update mechanism, we define what it means for a scheme to be compatible with a key update mechanism. We showi this is the case for our MAC given in the previous section and then go on to prove our MAC secure in the face of leakage.

4.1 Key Update Mechanism

We define a key update mechanism as a set of tuples $\mathcal{KU} = (Share, Recombine, U^{◓}, U^{◒})$ such that:

$$\begin{aligned}
(S_0^{◓}, S_0^{◒}) &\xleftarrow{\$} Share(K) \\
(S_{i+1}^{◓}, r_u) &\xleftarrow{\$} U^{◓}(S_i^{◓}) \\
S_{i+1}^{◒} &\xleftarrow{\$} U^{◒}(S_i^{◒}, r_u) \\
K_i &\leftarrow Recombine(S_i^{◓}, S_i^{◒})
\end{aligned}$$

For correctness we require that $Recombine(Share(K)) = K$.

We define an equivalence class as follows; we say $(S_i^{◓}, S_i^{◒}) \equiv (S_j^{◓}, S_j^{◒})$ if $Recombine(S_i^{◓}, S_i^{◒}) = Recombine(S_j^{◓}, S_j^{◒})$. Then the final requirement is that the algorithms $U^{◓}, U^{◒}$ preserve the equivalence class of the shares (and thus $\forall i : K_i = K$). Formally we require $(S_i^{◓}, S_i^{◒}) \equiv (S_{i+1}^{◓}, S_{i+1}^{◒})$ where $(S_{i+1}^{◓}, O_i) \xleftarrow{\$} U^{◓}(S_i^{◓}), S_i^{◒} \xleftarrow{\$} U^{◒}(S_i^{◒}, O_i)$.

The KP key update mechanism used within the KEM [13] can be seen to fit within this framework. This is due to the fact that the key is initially split into two shares which multiply together to give back the original key. The first share is updated by multiplying it by a random value, while the second share is updated by multiplying it by the inverse of the random value. This forms our equivalence class and thus when the two shares are multiplied together we will recover the orignal key, regardlesss of how many times the shares have been updated. The KP key update mechanism will be used for the remainder of this paper (and denoted $\mathcal{KU}$).

Definition 4 (Key Update Splittable). *We say that a tuple of functions $(F^{◓}, F^{◒})$ is a split of F conforming to key update mechanism $\mathcal{KU}$ if the following two properties hold. Firstly:*

$$\{F(K, x)\}_{\mathcal{R}} = \{F^*(Share(K), x)\}_{\mathcal{R}^*}$$

where F^* is defined in Fig. 5, the equivalence is over the randomness from sets $\mathcal{R}, \mathcal{R}^*$ used by F, F^* respectively. Secondly, that for all sharings $(S_0^{◓}, S_0^{◒})$ the joint distribution on $(S_1^{◓}, S_1^{◒})$ after F^* has been called once is the same as if $(S_0^{◓}, S_0^{◒})$ had been updated using $(U^{◓}, U^{◒})$.

Claim. The MAC $\mathcal{M}$ given in Sec. 3 is Key Update Splittable conforming to the KP Key Update Mechanism $\mathcal{KU}$.

proc $F^*(S_i^{◓}, S_i^{◒}, x)$:
$(S_{i+1}^{◓}, O) \xleftarrow{\$} F^{◓}(S_i^{◓}, x)$
$(S_{i+1}^{◒}, y) \xleftarrow{\$} F^{◒}(S_i^{◒}, O)$
Return y

Fig. 5. The algorithm F^*

proc $KG()$:
$K \xleftarrow{\$} \mathbb{G}_1$
$S_0^{◓} \xleftarrow{\$} \mathbb{G}_1$
$S_0^{◒} \leftarrow K \cdot (S_0^{◓})^{-1}$
Return $(S_0^{◓}, S_0^{◒})$

proc $TAG^{◓}(S_i^{◓}, m)$:
$w \xleftarrow{\$} R$
$W \leftarrow H(m, w)$
$t_i^{◓} \leftarrow e(S_i^{◓}, W)$
$r_{i+1} \xleftarrow{\$} \mathbb{G}_1$
$s_{i+1}^{◓} \leftarrow S_i^{◓} \cdot r_{i+1}$
Return $(S_{i+1}^{◓}, r_{i+1}, t_i^{◓}, W, w)$

proc $VRFY(K, (T, w), m)$:
$W \leftarrow H(m, w)$
$T' \leftarrow e(K, W)$
Return $(T' = T)$

proc $TAG^{◒}(S_i^{◒}, t_i^{◓}, W, w, r_{i+1})$:
$t_i^{◒} \leftarrow e(S_i^{◒}, W)$
$S_{i+1}^{◒} \leftarrow S_i^{◒} \cdot r_{i+1}^{-1}$
$T \leftarrow t_i^{◓} \cdot t_i^{◒}$
Return $(S_{i+1}^{◒}, (T, w))$

Fig. 6. Leakage Resilient MAC $\mathcal{M}^*$

Proof. $\mathcal{M}$ can be converted into $\mathcal{M}^*$ which is given in Fig. 6.

Since we have that:

$$\begin{aligned} Tag^*(S_i^{◓}, S_i^{◒}, m) &= (T, w) \\ &= (t^{◓} \cdot t^{◒}, w) \\ &= (e(S_i^{◓}, H(m, w)) \cdot e(S_i^{◓}, H(m, w)), w) \\ &= (e(S_i^{◓} \cdot S_i^{◒}, H(m, w)), w) \\ &= (e(K, H(m, w)), w) \\ &= Tag(K, m). \end{aligned}$$

Hence $\mathcal{M}$ is Key Update Splittable □

The security notion used for including leakage is called EUF-CMLA security and is defined below. There are three algorithms in our leakage resilient MAC: Key Generation, Tag and Verify. Our security definition only allows to leak on Tag, and we now explain why this is necessary. Clearly the Key Generation must not leak because it would leak on the original key. This constraint actually matches practice because typical (security) devices would be shipped with their

experiment $\mathbf{Exp}^{\mathrm{eufcmla}}_{\mathcal{M}}(A)$:
$K \xleftarrow{\$} KG()$
$(S_0^{◓}, S_0^{◒}) \xleftarrow{\$} SHARE(K)$
$S \leftarrow \{\}$
$(\sigma^*, m^*) \leftarrow A^{Tag(\cdot), Verify(\cdot,\cdot)}$
if $m^* \in S$ **then**
  return 0
end if
Return $VRFY(K, \sigma^*, m^*)$

proc $Verify(\sigma, m)$:
$b \leftarrow VRFYK, \sigma, m$
Return b

proc $Tag(m, l_i^{◓}, l_i^{◒})$:
$S \leftarrow S \cup \{m\}$
$(S_{i+1}^{◓}, O_i) \xleftarrow{r_i^{◓}} TAG^{◓}(S_i^{◓}, m)$
$\Lambda_i^{◓} \leftarrow l_i^{◓}(S_i^{◓}, r_i^{◓})$
$(S_{i+1}^{◒}, \sigma) \xleftarrow{r_i^{◒}} TAG^{◒}(S_i^{◒}, O_i)$
$\Lambda_i^{◒} \leftarrow l_i^{◒}(S_i^{◒}, r_i^{◒}, O_i)$
Return $(\sigma, \Lambda_i^{◓}, \Lambda_i^{◒})$

Fig. 7. EUF-CMLA experiment

keys preinstalled. This leaves to consider whether Tag or Verify (or both) are allowed to leak. This question has not been considered before in the continual poly-time leakage model in the case of symmetric schemes. This is because all previous schemes in this model were public-key in which the question simply does not arise.

Making Verify leaky is problematic, because we are allowing adaptive leakage, which this following example shows. Assume the adversary takes a random group element and a message and sends both to Verify. In the majority of schemes, Verify has to calculate the correct tag to compare the submitted tag against. Hence the adversary can keep submitting the same message and random elements until he has completely leaked the tag created for comparison. This tag can then be submitted as a forgery since it was never requested from the Tag oracle. There are currently no known methods to do this form of equality check in a secure manner when leakage is involved.

Thankfully only leaking on Tag seems to align with practice because in small devices such as smart cards, it tends to be the card that must authenticate itself to the reader and not the other way around. This means that the card will be using the Tag algorithm, which will leak since it is a small device, while it is the reader who will be verifying the tags and since this is a more powerful device, it also implement side channel protections to reduce leakage.

Definition 5 (Existential Unforgability Under Chosen Message Leakage Attack (EUF-CMLA)).

Let $\mathcal{M}^* = (\mathcal{KU}, TAG^{◓}, TAG^{◒}, VRFY)$ *be a Message Authentication Code. Then Fig. 7 defines the EUF-CMLA security game. The advantage of an adversary A winning the game is defined as* $\mathbf{Adv}^{\mathrm{eufcmla}}_{\mathcal{M}}(A) = \Pr[\mathbf{Exp}^{\mathrm{eufcma}}_{\mathcal{M}}(A) = 1]$.

Theorem 3. *The MAC* $\mathcal{M}^*$ *is EUF-CMLA secure in the Generic Group Model. The advantage of a q-query (to the generic group oracles) adversary who is* *allowed λ bits of leakage is given by:*

$$\mathbf{Adv}_{\mathcal{M}^*}^{\text{eufcmla}}(A) \leq 2^{2\cdot\lambda} \cdot \mathbf{Adv}_{\mathcal{M}}^{\text{eufcma}}(B) + \frac{q^2}{p}.$$

Proof. This proof is given in the Generic Group Model and shows that even with the use of leakage the adversary cannot get any elements that they could not get when no leakage was involved. After this has been shown it is reasonably straightforward to argue that without learning any new elements from the leakage then the leakage can increase the adversary's advantage by at most the number of bits that is leaked on for a single element. By showing that each element is only leaked on twice we get that the advantage can only be increased by at most $2^{2\lambda}$ over the advantage in the game where no leakage is involved.

We will represent group elements with polynomials, which will be instantiated at the end of the computation. The polynomials allow the game to keep track of which elements the adversary has asked for in a straightforward manner and because they are instantiated at the end, the adversary's decisions clearly can not be dependant on the actual values of the elements. Instatiation of the polynomials at the end is a common trick used within the litrature but means that if two (non equal) polynomials, when instantiated, collide the simulation fails as a single group element now has multiple representations. Thus we must also show that the chance of an adversary forcing this collision is also small.

Let $K, \{R_i\}_{i=0}^{q_T}, \{H_i\}_{i=1}^{q_H}, \{U_i\}_{i=1}^{2q_O}, \{V_i\}_{i=1}^{2q_O}, \{W_i\}_{i=1}^{2q_O}$ be indeterminants where q_H is the number of hash queries, q_T is the number of Tag calls and q_O is the number of group oracle calls (let $q = q_H + q_T + 6 \cdot q_O$). The indeterminants represent the following; K is the secret key, $\{R_i\}_{i=0}^{q_T}$ are the randomness used to update the key, $\{H_i\}_{i=1}^{q_H}$ represent any hash function queries and $\{U_i\}_{i=1}^{2q_O}, \{V_i\}_{i=1}^{2q_O}, \{W_i\}_{i=1}^{q_O}$ represent any elements that are guessed in $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_3$ respectively. The lists $\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3$ are used to keep track of polynomials and their representations in $\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_3$ respectively. They are initialised as follows:

$$\mathcal{L}_1 = \{(1, \xi_1^1\} \cup \{(R_i, \xi_{i+2}^1)\}_{i=0}^{q_T}$$
$$\mathcal{L}_2 = \{(1, \xi_1^2)\}$$
$$\mathcal{L}_3 = \{(1, \xi_1^3)\}$$

All three lists are initially instantiated with the identity, note that it is not strictly necessary to instantiate the identity in $\mathbb{G}_3$ since it can be calculated. We precompute the representations of the randomness but since the adversary does not have access to this list of elements (and instantiation, defining the element happens at the end) this does not change the game but makes it notationally simpler. All of the representations in this step are chosen randomly with the requirement that they are unique.

The oracles used by the Generic Group model are given in Fig. 8.

The Adversary A outputs $m, (T, w)$ and is said to have won if:

proc $\mathcal{O}_1(\xi_1, \xi_2)$:
if $\xi_1 \notin \mathcal{L}_1$ **then**
  $F_1 \leftarrow Guess_1(\xi_1)$
end if
if $\xi_2 \notin \mathcal{L}_1$ **then**
  $F_2 \leftarrow Guess_1(\xi_2)$
end if
get F_1 and F_2 from $\mathcal{L}_1$
$F_3 \leftarrow F_1 + F_2$
if $F_3 \in \mathcal{L}_1$ **then**
  get ξ_3 from $\mathcal{L}_1$
else
  $\xi_3 \leftarrow Sample_1(F_3)$
end if
Return ξ_3

proc $\mathcal{O}_e(\xi_1, \xi_2)$:
if $\xi_1 \notin \mathcal{L}_1$ **then**
  $F_1 \leftarrow Guess_1(\xi_1)$
end if
if $\xi_2 \notin \mathcal{L}_2$ **then**
  $F_2 \leftarrow Guess_2(\xi_2)$
end if
get F_1 from $\mathcal{L}_1$
get F_2 from $\mathcal{L}_2$
$F_3 \leftarrow F_1 \cdot F_2$
if $F_3 \in \mathcal{L}_3$ **then**
  get ξ_3 from $\mathcal{L}_3$
else
  $\xi_3 \leftarrow Sample_3(F_3)$
end if
Return ξ_3

proc $Sample_1(F)$:
$\xi \stackrel{\$}{\leftarrow} \Xi_1 \backslash \mathcal{L}_1$
add (F, ξ) to $\mathcal{L}_1$
Return ξ

proc $Guess_1(\xi)$:
$d \stackrel{\$}{\leftarrow} \mathbb{Z}_p$
add (d, ξ) to $\mathcal{L}_1$
Return d

Fig. 8. GGM group oracles used within the proof ($\mathcal{O}_2, \mathcal{O}_3, Sample_2, Sample_3, Guess_2, Guess_3$ are not included due to their similarity to the oracles for $\mathbb{G}_1$.)

1. $F_i^l = F_j^l$ for $l \in \{1, 2, 3\}$ and $i \neq j$
2. $K \cdot H^* - T = 0$ where H^* is the indeterminant corresponding to the hash of m and w

The first case corresponds to the adversary being able to create two polynomials which evaluate to the same value. The second case corresponds to the adversary being able to create a forgery on the MAC.

Since all polynomials in the orginal lists are degree one and the only operation that increases the degree is the pairing operation which can only be called on elements in $\mathbb{G}_1, \mathbb{G}_2$. This means that we can have degree two polynomials in $\mathbb{G}_3$ but not the other two groups. Hence by the Swartz-Zippel lemma we have that the probaility of two (non-zero) polynomials evaluate to the same value is $\frac{2}{p}$. Since there are at most q polynomials we get that there are $\binom{q}{2} \leq \frac{q^2}{2}$ pairs of polynomials that could collide and thus the probability of any two polynomials colliding is $\frac{q^2}{p}$.

Without loss of generality, we will now only look at leakage in the target group $\mathbb{G}_3$ since any element from $\mathbb{G}_1, \mathbb{G}_2$ calculated by the leakage can be transfered over to $\mathbb{G}_3$ using a pairing and any elements known to the adversary can easily be recomputed within the leakage.

While we now only need to look at leakage in the target group since each leakage function has access to different secret elements, the set of elements that can be calculated by each will be different and thus must be considered individually. The adversary will win if he can repeatedly get a (new) Tag into the leakage function to reveal it, or he can cause collisions within the leakage. If he can only create a forgery once he will not be able to leak the complete tag and thus he is required to get in into multiple leakage sets

Let $L_i^{\text{◓}}$ be the set of elements that could be computed by the leakage function $l_i^{\text{◓}}$, then:

$$L_i^{\text{◓}} = \{A \cdot S_i^{\text{◓}} + B \cdot R_{i+1} + C\}$$

Where $A, B \in \mathbb{F}_p[\{H_i\}_{i=0}^{q_H}, \{V_i\}_{i=0}^{2q_O}]$ and $C \in \mathbb{F}_p[K\{H_i\}_{i=0}^{q_H}, \{U_i\}_{i=0}^{2q_O}, \{V_i\}_{i=0}^{2q_O}, \{W_i\}_{i=0}^{q_O}]$ and we use $S_i^{\text{◓}}$ to denote $\sum_{j=0}^{i} R_i$.

Let $L_i^{\text{◒}}$ be the set of elements that could be computed by the leakage function $l_i^{\text{◒}}$, then:

$$L_i^{\text{◒}} = \{A \cdot S_i^{\text{◒}} + B \cdot R_{i+1} + C + d \cdot S_i^{\text{◓}} \cdot H_i\}$$

Without loss of generality we will assume that i$^{\text{th}}$ tag call maps to H_i. Also we have that $d \in \mathbb{F}_p$ and$S_i^{\text{◒}}$ denotes $K - \sum_{j=0}^{i} R_i$.

The only Tags that can be contained in $L_i^{\text{◓}}$ are tags of the form $H_j \cdot K$ for $j < i$ in which the adversary has included the tag via F and thus there is no advantage in leaking upon $H_j \cdot K$. Ignoring this trivial case, there are no linear combinations possible that will reveal an unknown tag or the key itself. Similarly for $L_i^{\text{◓}}$ the only tags that can be included are of the form $H_j \cdot K$ for $j \leq i$ this is because he can again ask to leak on tags he has already seen by embedding them in F but this time can also leak on $H_i \cdot K$ but this is also not of any use because the leakage on this tag will be received at the same time the adversary is given the Tag and thus no extra information is gained.

From this point on we will only consider leakage on elements that contain an unknown component, since if all components are completely known by the adversary it is not worth learning leakage on.

Claim. An element x which has been leaked on can be in at most 2 of the $L_i^{\text{◓}}$ or $L_i^{\text{◒}}$. More formally $\forall i, j, k : i \neq j \neq k$ $L_i^{\text{◓}} \cap L_j^{\text{◓}} \cap L_k^{\text{◓}} = \emptyset$ and $L_i^{\text{◒}} \cap L_j^{\text{◒}} \cap L_k^{\text{◒}} = \emptyset$ and if $x \in L_i^{\text{◓}} \cap L_j^{\text{◒}}$ then $x \notin L_k^{\text{◓}}$ or $x \notin L_k^{\text{◒}}$.

Since R_i is only used within $L_i^{\text{◓}}, L_i^{\text{◒}}$ any element involving R_i as a component can only be leaked on from one of these sets (and thus with leakage functions $l_i^{\text{◓}}, l_i^{\text{◒}}$). Looking at $S_i^{\text{◓}}$ ($S_i^{\text{◒}}$ behaves similarly) it can only appear in $L_i^{\text{◓}}$ and $L_{i-1}^{\text{◓}}$ (since $S_i^{\text{◓}} = S_{i-1}^{\text{◓}} + R_i$) and thus again can only be leaked on by two leakage functions. The element $S_i^{\text{◓}} \cdot H_i$ which is passed between $Tag^{\text{◓}}$ and $Tag^{\text{◒}}$ can only be leaked on by $L_i^{\text{◓}}, L_i^{\text{◒}}$.

It can not be leaked on from $L_{i-1}^{\text{◓}}$ because the randomness w will not be chosen until the next invocation and thus the leakage function can only guess the correct w with negligible probability. Since each element can only be leaked on twice the adversary can only learn up to $2 \cdot \lambda$ bits of information per secret. Thus we get the bound as stated in the theorem, since after collisions the adversary's advantage can be at most $2^{2\cdot\lambda}$ times the advantage of playing the standard non-leakage game.

$\square$

5 Practical Aspects of our Scheme

In this section we focus on some practical considerations: how efficient is it in comparison to other leakage resilient MAC constructions, and what would a practical implementation need to guarantee to meet our leakage bound/assumptions?

Before giving a comparision we need to make some choices for parameters of the various schemes. In case of schemes which have as underlying primitive a pseudo random function (PRF), we chose to instantiate this PRF with AES-128. This is motivated by the fact that this reflects the current state of the art. For our own scheme, and a somewhat comparable signature scheme, we use as instantiation of the bilinear map a pairing defined over a suitable pairing friendly elliptic curve. In this case we choose as group size parameter 2^{160}, again with the motivation to reflect a state-of-the art security bound (i.e. 2^{80} is regarded as minimum for bound which is implied by a group of double this size).

5.1 Security considerations

The leakage bound we have for our scheme is that we can tolerate up to approx. 50 bits of leakage assuming a group size of about 2^{160} per invocation of the scheme. For a practical implementation it is important to acknowledge that this will not include the initial sharing out of the key. Consequently in a strict sense this would need to be done in a secure environment.

When considering what '50 bits of leakage' means given our constructions, it helps to think about the tagging and verification algorithms in a concrete instantiation: i.e. we are working with pairings that are defined over some pairing friendly curve. Consequently, $S_i^{◓}$, $S_i^{◒}$, r_i and ${r_i}^{-1}$ are elliptic curve points, $S_i^{◓} \cdot r_i$ and $S_i^{◒} \cdot {r_i}^{-1}$ are elliptic curve point additions, and e is a pairing. Now that we have a concrete instantiation of our scheme, we can argue more concretely about implementation options. In order to provide some resistance against simple side channel attacks such as SPA we would hence essentially need to blind the EC points which correspond to secrets. A recent contribution [3] gives a good overview on various side channel attacks on pairings from which one can again conclude that also the pairing operation should best be implemented on blinded EC points.

The obvious question arising now might be what we have gained in practice, as we yet again need to apply the basic techniques that would prevent attacks such as SPA and DPA? The answer is for any non leakage resilient version, even partial leakage on the random values will probably allow lattice based attacks such as [22]. The leakage resilient scheme however guarantees that in the presence of partial leakage no such attack can succeed.

5.2 Comparison with other leakage-resilient schemes

In our comparison we essentially look at the number of elliptic curve or AES operations that tagging and verification require. We also report on the tolerated leakage, the key size and the tag size. Table 1 provides an overview and the following text explains and discusses the provided numbers. As can be seen from this table, our scheme is highly competitive when compared with other provably secure schemes.

The only leakage resilient MAC scheme in the literature that isn't built on top of another leakage resilient scheme is by Hazay *et al.* [12] to the best of our knowledge. The advantage of their work is that it only relies on very minimal assumptions (the existence of one-way functions). However the leakage bound is a total bound, *i.e.* it holds regardless of the number of times the tag and verify algorithms are called. Assuming that we instantiate the PRF required for this scheme with AES-128 (*i.e.* setting $\lambda = s = 128$ to use the notation in the paper), and using the equations they give in Theorem 5.6, it turns out that AES will be called 512 times per tag and verify query. While this already makes the scheme computationally expensive, the larger problem is the overall key size: the key must be of size approximately 2^{18} bits, which is impractical for many applications. Even worse, under these parameters the MAC can leak 512 bits over the lifetime of the system (i.e. not per MAC invocation). This also means that with these choices AES and the elliptic curve schemes will give approximately the same level of security.

Due to the small number of leakage resilient MACs, we also provide a comparison against PRFs and signatures because although there is not an immediate strategy to convert them into leakage resilient MACs available, there is potential that one might use PRFs or signatures to instantiate a MAC with some leakage resilient properties.

The PRF by Dodis and Pietrzak [5] requires that the leakage functions are fixed prior to the attack and are not adaptively chosen. They define the PRF $\Gamma^F : \Sigma^{3k+n} \times \Sigma^m \rightarrow \Sigma^{4k+2n}$ created from a wPRF $F : \Sigma^k \times \Sigma^n \rightarrow \Sigma^{4k+2n}$. If we instantiate F with AES-128 with something like $F(x) = Enc_K(x||000)$ $||Enc_K(x||001)||Enc_K(x||010)||Enc_K(x||011)||Enc_K(x||100)||Enc_K(x||101)$ and then take the desired number of output bits, this gives us $k = m = 128$, $n = 125$ and $\Gamma^F : \Sigma^{509} \times \Sigma^{128} \rightarrow \Sigma^{768}$. As stated in the paper each time the PRF is called, the function F is called $m + 1$ times and thus AES will be called 774 times. Hence even if this can be converted into a MAC reasonably inexpensively, the PRF itself is very expensive.

Schipper's construction [24] requires the use of a EUF-CMA MAC and a leakage resilient PRF. Hence, the timings will be very similar to the numbers from [5] and thus we do not include it explicitly in the table.

The PRF by Faust *et al.* [8] requires that neither the leakage or the PRF inputs are queried adaptively but are fixed prior to the start of the game. They construct a scheme $\Gamma^{F,m} : \{0,1\}^{k+(m+1)l} \times \{0,1\}^m \rightarrow \{0,1\}^n$ which uses the wPRF $F : \{0,1\}^k \times \{0,1\}^l \rightarrow \{0,1\}^{2k}$ and $m + 1$ public random values of length l. If we again instantiated the wPRF with AES-128 to get

$F(x) = Enc_K(x||0)||Enc_K(x||1)$ meaning we get $m = k = 128$, $l = 127$,$n = 256$ and $\Gamma^{F,128} : \{0,1\}^{16511} \times \{0,1\}^{128} \to \{0,1\}^{256}$. This will call AES 258 times per invocation, which while an improvement still seems prohibitive for a practical implementation.

The signature scheme by Faust *et al.* [7] uses $2l$ exponentiations, $4(l-1)$ multiplications, $l-1$ additions and 2 hash function calls in the signing algorithm and tl exponentiations, tl multiplications and t hash function calls in the verification algorithm, where l is related to the underlying l-representation problem [4] (assumed to be hard) and t is the depth of the signature chain. The downside of this scheme is that even if $l = 2$, while signing is efficient, verification takes longer depending how deep the signature chain is. This could mean that verification quickly becomes too expensive for an embedded device to perform.

The signature scheme by Galindo and Vivek [11] is the only one comes close to our construction in terms of performance, which is unsurprising given that it is also based on the same key update mechanism. For signing, the algorithm uses 2 Elliptic Curve scalar multiplications, 5 Elliptic Curve additions and also generates a random curve point (and its' inverse), while verification uses 2 Elliptic Curve point additions, 1 Elliptic Curve scalar multiplication and 2 pairings. Since a pairing is currently only slightly more expensive than an exponentiation, our tagging algorithm will be almost equivalent in timing to their signing algorithm. Their verification is faster than ours, however, their keys are larger. Furthermore, while there is no hash function explicit in their scheme, it is assumed that the message comes from $\mathbb{Z}_p$ and thus in practice to sign arbitrary messages, the message would have to be hashed onto $\mathbb{Z}_p$.

References

1. Abdalla, M., Belaïd, S., Fouque, P.A.: Leakage-resilient symmetric encryption via re-keying. In: Bertoni, G., Coron, J.S. (eds.) CHES 2013. LNCS, vol. 8086, pp. 471–488. Springer, Berlin, Germany, Santa Barbara, California, US (Aug 20–23, 2013)
2. Balasch, J., Faust, S., Gierlichs, B., Verbauwhede, I.: Theory and practice of a leakage resilient masking scheme. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 758–775. Springer, Berlin, Germany, Beijing, China (Dec 2–6, 2012)
3. Blömer, J., Günther, P., Liske, G.: Improved side channel attacks on pairing based cryptography. Cryptology ePrint Archive, Report 2011/706 (2011), http://eprint.iacr.org/2011/706
4. Brands, S.: An efficient off-line electronic cash system based on the representation problem. Technical Report CS–R9323
5. Dodis, Y., Pietrzak, K.: Leakage-resilient pseudorandom functions and side-channel attacks on Feistel networks. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 21–40. Springer, Berlin, Germany, Santa Barbara, CA, USA (Aug 15–19, 2010)
6. Dziembowski, S., Pietrzak, K.: Leakage-resilient cryptography. In: 49th FOCS. pp. 293–302. IEEE Computer Society Press, Philadelphia, Pennsylvania, USA (Oct 25–28, 2008)

Table 1. A comparison of possible EUF-CMLA schemes

Scheme	Leakage	Key Size	Tag Size	Tag Time	Verification Time
Our scheme	$\frac{1}{2}(\log p - 2\log q_u - n)$ p group size q queries n security bound (approx 50 bits)	2 EC points (approx 320 bits)	1 EC point 1 random string (approx 228 bits)	1 hash 1 random point (and inverse) generated 3 EC additions 2 pairings	1 hash 1 random point (and inverse) generated 3 EC additions 2 pairings
HLWW:[12]	512 bits total	approx 2^{18} bits	approx 2^{19} bits	512 AES calls	512 AES calls
DP:[5]	$\frac{\log(\epsilon^{-1})}{6}$ for $\epsilon = \mathbf{Adv}_F^{wprf}(A)$ (approx 22 bits)	509 bits	762 bits	774 AES calls	774 AES calls
FPS:[8]	$\frac{\log(\epsilon^{-1})}{4}$ for $\epsilon = \mathbf{Adv}_F^{wprf}(A)$ (approx 32 bits)	128 bits 16,383 public bits	256 bits	258 AES calls	258 AES calls
FKPR:[7]	$2\text{kl}(\frac{1}{2} - \frac{1}{2l} - \delta)$ (approx 40 bits)	Varies	Varies	$2l$ exponentiations $4(l-1)$ multiplications $l-1$ additions 2 hashes	Varies
GV:[11]	$<< \frac{\log p}{2}$ p is group size (approx 60 bits)	sk: 2 EC points pk: 3 EC points (approx 800 bits)	2 EC points (approx 320 bits)	1 random point (and inverse) generated 2 scalar multiplications 5 EC additions 1 hash	1 scalar multiplication 2 EC additions 2 pairings 1 hash

7. Faust, S., Kiltz, E., Pietrzak, K., Rothblum, G.N.: Leakage-resilient signatures. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 343–360. Springer, Berlin, Germany, Zurich, Switzerland (Feb 9–11, 2010)
8. Faust, S., Pietrzak, K., Schipper, J.: Practical leakage-resilient symmetric cryptography. In: Prouff, E., Schaumont, P. (eds.) CHES 2012. LNCS, vol. 7428, pp. 213–232. Springer, Berlin, Germany, Leuven, Belgium (Sep 9–12, 2012)
9. Faust, S., Rabin, T., Reyzin, L., Tromer, E., Vaikuntanathan, V.: Protecting circuits from leakage: the computationally-bounded and noisy cases. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 135–156. Springer, Berlin, Germany, French Riviera (May 30 – Jun 3, 2010)
10. Galbraith, S.D., Paterson, K.G., Smart, N.P.: Pairings for cryptographers. Discrete Applied Mathematics, Vol 156 #14 pages 3113–3121 (2008)
11. Galindo, D., Vivek, S.: A practical leakage-resilient signature scheme in the generic group model. In: Knudsen, L.R., Wu, H. (eds.) SAC 2012. LNCS, vol. 7707, pp. 50–65. Springer, Berlin, Germany, Windsor, Ontario, Canada (Aug 15–16, 2012)
12. Hazay, C., López-Alt, A., Wee, H., Wichs, D.: Leakage-resilient cryptography from minimal assumptions. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 160–176. Springer, Berlin, Germany, Athens, Greece (May 26–30, 2013)

13. Kiltz, E., Pietrzak, K.: Leakage resilient ElGamal encryption. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 595–612. Springer, Berlin, Germany, Singapore (Dec 5–9, 2010)
14. Kocher, P.: Blind signature systems. U.S. Patent #4,759,063
15. Kocher, P.C.: Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In: Koblitz, N. (ed.) CRYPTO'96. LNCS, vol. 1109, pp. 104–113. Springer, Berlin, Germany, Santa Barbara, CA, USA (Aug 18–22, 1996)
16. Kocher, P.C., Jaffe, J., Jun, B.: Differential power analysis. In: Wiener, M.J. (ed.) CRYPTO'99. LNCS, vol. 1666, pp. 388–397. Springer, Berlin, Germany, Santa Barbara, CA, USA (Aug 15–19, 1999)
17. Mangard, S., Oswald, E., Popp, T.: Power analysis attacks: Revealing the secrets of smart cards. Springer (2008)
18. Maurer, U.M.: Abstract models of computation in cryptography (invited paper). In: Smart, N.P. (ed.) 10th IMA International Conference on Cryptography and Coding. LNCS, vol. 3796, pp. 1–12. Springer, Berlin, Germany, Cirencester, UK (Dec 19–21, 2005)
19. Medwed, M., Standaert, F.X., Großschädl, J., Regazzoni, F.: Fresh re-keying: Security against side-channel and fault attacks for low-cost devices. In: Bernstein, D.J., Lange, T. (eds.) AFRICACRYPT 10. LNCS, vol. 6055, pp. 279–296. Springer, Berlin, Germany, Stellenbosch, South Africa (May 3–6, 2010)
20. Micali, S., Reyzin, L.: Physically observable cryptography (extended abstract). In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 278–296. Springer, Berlin, Germany, Cambridge, MA, USA (Feb 19–21, 2004)
21. Nechaev, V.I.: Complexity of a determinate algorithm for the discrete logarithm. Mathematical Notes 55(2), 165–172 (1994)
22. Nguyen, P.Q., Shparlinski, I.: The insecurity of the digital signature algorithm with partially known nonces. Journal of Cryptology 15(3), 151–176 (2002)
23. Renauld, M., Standaert, F.X., Veyrat-Charvillon, N., Kamel, D., Flandre, D.: A formal study of power variability issues and side-channel attacks for nanoscale devices. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 109–128. Springer, Berlin, Germany, Tallinn, Estonia (May 15–19, 2011)
24. Schipper, J.: Leakage-Resilient Authentication. Ph.D. thesis, Utrecht University (2010)
25. Shoup, V.: Lower bounds for discrete logarithms and related problems. In: Fumy, W. (ed.) EUROCRYPT'97. LNCS, vol. 1233, pp. 256–266. Springer, Berlin, Germany, Konstanz, Germany (May 11–15, 1997)
26. Yacobi, Y.: A note on the bilinear Diffie-Hellman assumption. Cryptology ePrint Archive, Report 2002/113 (2002), http://eprint.iacr.org/2002/113