

Implementation of a leakage-resilient ElGamal key encapsulation mechanism

David Galindo¹ · Johann Großschädl² · Zhe Liu³ · Praveen Kumar Vadnala² · Srinivas Vivek⁴

Received: 8 February 2015 / Accepted: 6 February 2016 / Published online: 29 February 2016
© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract Leakage-resilient cryptography aims to extend the rigorous guarantees achieved through the provable security paradigm to physical implementations. The constructions designed on basis of this new approach inevitably suffer from an Achilles heel: a bounded leakage assumption is needed. Currently, a huge gap exists between the theory of such designs and their implementation to confirm the leakage resilience in practice. The present work tries to narrow this gap for the leakage-resilient bilinear ElGamal key encapsulation mechanism (BEG-KEM) proposed by Kiltz and Pietrzak in 2010. Our first contribution is a variant of the bounded leakage and the only-computation-leaks model that is closer to practice. We weaken the restriction on the image size of the leakage functions in these models and only

insist that the inputs to the leakage functions have sufficient min-entropy left, in spite of the leakage, with no limitation on the quantity of this leakage. We provide a novel security reduction for BEG-KEM in this relaxed leakage model using the generic bilinear group axiom. Secondly, we show that a naive implementation of the exponentiation in BEG-KEM makes it impossible to meet the leakage bound. Instead of trying to find an exponentiation algorithm that meets the leakage axiom (which is a non-trivial problem in practice), we propose an advanced scheme, BEG-KEM+, that avoids exponentiation by a secret value, but rather uses an encoding into the base group due to Fouque and Tibouchi. Thirdly, we present a software implementation of BEG-KEM+ based on the Miracl library and provide detailed experimental results. We also assess its (theoretical) resistance against power analysis attacks from a practical perspective, taking into account the state-of-the-art in side-channel cryptanalysis.

A preliminary version of this paper has appeared at PROOFS 2014.

✉ David Galindo
D.Galindo@cs.bham.ac.uk
Johann Großschädl
johann.groszschaedl@uni.lu
Zhe Liu
zhelu.liu@uwaterloo.ca
Praveen Kumar Vadnala
praveen.vadnala@uni.lu
Srinivas Vivek
sv.venkatesh@bristol.ac.uk

Keywords Secure implementation · Side-channel cryptanalysis · Leakage-resilient cryptography · Security proof · Public-key encryption · Pairings

1 Introduction

How to secure cryptographic algorithms embedded in devices that can eventually “fall in the hands” of an adversary? Answering this question is probably the holy grail in cryptography nowadays. Two paths are taken to explore the possible solutions, a destructive one and a constructive one. In the first path, we find the rich contributions of the practice and theory of side-channel attacks. In the second path, we find the no less precious body of countermeasures against the attacks unveiled in the first path. Lately, a novel approach called leakage-resilient cryptography is being studied, which

¹ School of Computer Science, University of Birmingham, Birmingham B15 2TT, UK
² University of Luxembourg, FSTC, 6 rue Richard Coudenhove-Kalergi, 1359 Luxembourg, Luxembourg
³ Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, Canada
⁴ University of Bristol, Merchant Venturers Building, Woodland Road, Clifton, Bristol, UK

aims at extending the guarantees delivered by the provable security paradigm to the physical world. Despite the clever discoveries and constructions provided by this new approach, it persistently presents an Achilles heel: a bounded leakage assumption is needed. Ensuring this is unfortunately a challenging endeavor on its own and, admittedly, the leakage-resilient cryptography body of work has not significantly helped to argue why this could be a reasonable assumption.

In this work, we consider the *only computation leaks information* (OCL) leakage model by Micali and Reyzin [22]. In this model only actual computations are supposed to leak sensitive information. This captures the usual situation in side-channel attacks, where leakage data only depend on the current state of the target device and some independent randomness [32]. The internal data of the device are divided into two parts, an *active* and a *passive* part, the active part being the input data used in the current computation. Therefore, at a given time frame, only the active data is leaking. The main non-invasive attacks against embedded devices, like the attacks based on power consumption [21], electromagnetic radiations [15] or running-time [20] measurements, belong to this category.

It is currently agreed upon that not only the OCL model but also the bounded retrieval/memory leakage models [1,2] or the auxiliary input model [8], rely on a strange combination of both strong and weak assumptions. On the one side, the information leakage is supposed to be bounded in a somewhat artificial manner; on the other side, the leakage considered is overly general, for instance it might come from any polynomial time function. However, these assumptions are actually far from the reality that practitioners experience in their daily work in a side-channel analysis lab.

Several contemporary works [4,27,32] have put forward ways to redefine the above models and bring them closer to practice, for symmetric cryptography primitives. This comes at the cost of algorithmic-level specialization, providing models that are indeed more realistic, but which apply to a more restrained class of primitives (i.e., pseudorandom generators, block ciphers).

We aim at contributing to the challenge of bringing leakage-resilient cryptography closer to the practice. In this work, we do so by analyzing, modifying, implementing and evaluating a previous leakage-resilient key encapsulation mechanism proposed by Kiltz and Pietrzak [18]. This is one of the very few schemes admitting continual leakage (maybe the only one?) that one could dare to implement in an embedded processor, for instance in a smartphone. It is a pairing-based stateful variant of the ElGamal encryption scheme (called BEG-KEM), where the secret key is an element of the pairing base group (essentially a point in the group of points of an elliptic curve). The secret key is divided into two shares, which are re-shared at each new decryption call using multiplicative blinding. To decrypt, one takes the

first half of the secret key, refreshes it, and uses it as the input to a pairing calculation. In the second step, the second half of the secret key is updated with the blinding used for refreshing; it is then used as the input to a new pairing calculation; and finally the two pairing values are multiplied to obtain a decapsulated symmetric key (for the details see Sect. 2).

The result proven in [18], which holds under a variant of the generic group model tailored to pairing groups uses a bounded leakage assumption. Roughly speaking, it is required that the *data leaked* against side-channel attacks that satisfy the OCL axiom, shall be significantly smaller than κ for a single measurement, where κ is the security parameter (e.g., $\kappa = 128$). These leakages are modeled as an oracle that answers values $f(\cdot)$ for adaptively chosen arbitrary (but efficiently computable) functions f on input the secret data being used in the calculation. This kind of requirement that may look reasonable for a theoretician used to study cryptographic primitives in the so-called *black-box model* might seem completely unrealistic to the practitioner. As an example, let us recall the figure gathered in [32], where it is pointed out that the leaking of a block cipher recently reported in [24], consisted of 200,000 traces leading to more than 1.5 Gb of data storage.

We start our investigation by proposing and testing a relaxation on the requirement of ‘bounded leakage size’ in the OCL model. We weaken the restriction on the image size of the leakage functions in these models to asking that the random variables used to refresh the secret key shall have enough min-entropy left given the leakage, with no limitation on the ‘size’ of this leakage. This is an altogether more reasonable leakage bound assumption, which could eventually be met by clever implementations (in fact we provide an implementation candidate). We give a new security reduction using the generic bilinear group axiom for BEG-KEM in this relaxed leakage model, which turns out to be tighter than the original reduction in [18] in the OCL model. Due to space limitations, we only include here a short description of the proof. The complete proof can be found in the full version [11].

Secondly, we observe that the blinding mechanism originally proposed is susceptible to invalidate the leakage bound assumption. This is because to perform blinding, one computes an exponentiation G^{r_i} for a random integer r_i , which if implemented in a naive way, can almost completely leak r_i , even with a simple power analysis attack (i.e., with a single power trace), as we discuss in Sect. 5. The authors in [18] did not discuss how exponentiation shall be implemented to meet the leakage bound, nor we can currently find an exponentiation algorithm with these guarantees. Thus, their positive result risks to be void.

This is why we propose an advanced BEG-KEM+, where we avoid blinding by an exponentiation G^{r_i} for a random

integer r_i . Our modification is based on the observation that knowledge of the exponent r_i is not needed to perform a successful decryption, but it suffices to build a *random element* in a suitable pairing base group. We propose instead to use a random encoding into asymmetric pairing groups by Fouque and Tibouchi [10]. It turns out that this encoding produces a random element in the base group, and can naturally be implemented in such a way that the leakage expected against a single measurement is arguably minimal (see Sect. 5).

Fourthly, we stress that the idea of leakage-resilient cryptography—like any other theoretical concept—can only be brought into practice by actual implementation. For this reason, we implemented BEG-KEM+ in ANSI C on an ARM-based microcontroller. BEG-KEM+ is, to our knowledge, the first implementation and evaluation of a public-key scheme from the leakage-resilient literature.

2 Stateful bilinear ElGamal KEM

In this section we present the stateful bilinear ElGamal key encapsulation mechanism (BEG-KEM) from [18]. First, we recall the basics of the notion of min-entropy. Then we introduce the concept of stateful KEM and security under non-adaptive chosen-ciphertext attacks in the presence of continual min-entropy leakage (CCmLA1). We note again that the class of leakage functions allowed in our model (based on lowering min-entropy) is broader than the bounded length model (CCLA1) used in [18].¹

2.1 Min-entropy

Let X be a finite random variable with probability distribution \Pr . The *min-entropy* of X , denoted $\mathbf{H}_\infty(X)$, is defined as

$$\mathbf{H}_\infty(X) := -\log_2 \left(\max_x \Pr[X = x] \right).$$

Min-entropy is a standard measure of the worst-case predictability of a random variable. Let Z be a random variable. The *average conditional min-entropy* of X given Z , denoted $\tilde{\mathbf{H}}_\infty(X | Z)$, is defined as

$$\tilde{\mathbf{H}}_\infty(X | Z) := -\log_2 \left(\mathbb{E}_{z \leftarrow Z} \left[\max_x \Pr[X = x | Z = z] \right] \right).$$

Average conditional min-entropy is a measure of the worst-case predictability of a random variable given a correlated random variable.

¹ We point out the authors of [18] mention that their results also carry over to a relaxed leakage model, close in spirit to ours. However, this model is not fully detailed, and additionally no justification of this fact is given in [18] nor in [19].

Lemma 1 [9] *Let $f : X \rightarrow \{0, 1\}^{\lambda'}$ be a function on X . Then $\tilde{\mathbf{H}}_\infty(X | f(X)) \geq \mathbf{H}_\infty(X) - \lambda'$.*

The following result is a variant of the Schwartz–Zippel Lemma [13, 28, 38].

Lemma 2 (Schwartz–Zippel; min-entropy version) *Let $F \in \mathbb{Z}_q[X_1, \dots, X_n]$ be a non-zero polynomial of (total) degree at most d . Let P_i ($i = 1, \dots, n$) be probability distributions on \mathbb{Z}_q such that $\mathbf{H}_\infty(P_i) \geq \log q - \lambda'$, where $0 \leq \lambda' \leq \log q$. If $x_i \xleftarrow{P_i} \mathbb{Z}_q$ ($i = 1, \dots, n$) are independent, then $\Pr[F(x_1, \dots, x_n) = 0] \leq 2^{\lambda'} \frac{d}{q}$.*

Corollary 1 *If $\lambda' < \log q - \omega(\log \log q)$ in Lemma 2, then $\Pr[F(x_1, \dots, x_n) = 0]$ is negligible (in $\log q$).*

2.2 Stateful key encapsulation mechanism

Formally, a split-state key encapsulation mechanism $\text{KEM} = (\text{KG}, \text{Enc}, \text{Dec1}, \text{Dec2})$ consists of four polynomial time algorithms. Let κ denote the security parameter and λ denote the *leakage parameter*. The key generation procedure $\text{KG}(\kappa, \lambda)$ takes as input κ and λ , and outputs the public key pk , a pair of initial (stateful) secret states (σ_0, σ'_0) , and the public parameters \mathbb{PP} . The encapsulation procedure $\text{Enc}(pk)$ takes as input pk , and outputs a secret symmetric key K and the corresponding ciphertext C . The stateful decapsulation procedure takes C as an input and outputs $K \in \mathcal{K}$. This procedure is split into two consecutive steps Dec1 and Dec2 , where each step accesses distinct parts of the two secret states. The procedures Dec1 and Dec2 may also update the secret key using locally generated fresh randomness:

$$(\sigma_i, w_i) \xleftarrow{r_i} \text{Dec1}(\sigma_{i-1}, C) ; (\sigma'_i, K) \xleftarrow{r'_i} \text{Dec2}(\sigma'_{i-1}, w_i).$$

The scheme KEM is required to satisfy the following correctness property:

$$\Pr \left[\text{Dec2}(\text{Dec1}(\text{Enc}(pk), \sigma_{i-1}) \setminus \sigma_i, \sigma'_{i-1}) = K : (pk, (\sigma_{i-1}, \sigma'_{i-1})) \leftarrow (\text{KG}, \text{Dec1}, \text{Dec2}), K \leftarrow \text{Enc}(pk) \right] = 1.$$

The security of the scheme KEM is defined by the experiment in Table 1.

In this experiment, $f_i(\sigma_{i-1}, r_i)$ and $h_i(\sigma'_{i-1}, r'_i, w_i)$ are (efficiently computable) leakage functions that the adversary can choose adaptively between the rounds. The functions $f_i(\cdot)$ and $h_i(\cdot)$ are such that the min-entropy of the individual inputs of the leakage functions is decreased by at most λ bits, given the corresponding leakages. More precisely, the requirement on the leakage functions is that

$$\tilde{\mathbf{H}}_\infty(t | f_i(\sigma_{i-1}, r_i)) \geq \mathbf{H}_\infty(t) - \lambda \quad \forall t \in \sigma_{i-1} \cup r_i,$$

Table 1 CCmLA1 security experiment for KEM

KEM-CCmLA1 _{KEM} ($\mathcal{A}, \kappa, \lambda$)	KEM-Leak-Oracle $O^{\text{CCmLA1}}(C, f_i, h_i)$
$(pk, (\sigma_0, \sigma'_0)) \leftarrow \text{KG}(\kappa, \lambda)$	
$i := 1, w \leftarrow \mathcal{A}^{O^{\text{CCmLA1}}(\cdot)}(pk)$	$(\sigma_i, w_i) \xleftarrow{r_i} \text{Dec1}(\sigma_{i-1}, C)$
$b \xleftarrow{\$} \{0, 1\}$	$(\sigma'_i, K) \xleftarrow{r'_i} \text{Dec2}(\sigma'_{i-1}, w_i)$
$(C, K_0) \leftarrow \text{Enc}(pk)$	$\Lambda_i := f_i(\sigma_{i-1}, r_i)$
$K_1 \xleftarrow{\$} \mathcal{K}$	$\Lambda'_i := h_i(\sigma'_{i-1}, r'_i, w_i)$
$b' \leftarrow \mathcal{A}(w, CK_b)$	$i := i + 1$
	Return $(K, \Lambda_i, \Lambda'_i)$

and

$$\tilde{\mathbf{H}}_\infty(t | h_i(\sigma'_{i-1}, r'_i, w_i)) \geq \mathbf{H}_\infty(t) - \lambda \quad \forall t \in \sigma'_{i-1} \cup r'_i \cup w_i.$$

Essentially, the above equations restrict the class of allowed leakage functions to those that do not decrease the min-entropy of each atomic parameter of the secret state by more than λ bits. For instance, if $w_i = \{w_{i,1}, w_{i,2}\}$, then we require that individually $w_{i,1}$ and $w_{i,2}$ have their min-entropy reduced by at most λ bits given the leakages.

Definition 1 (CCmLA1 security for KEM) A key encapsulation mechanism KEM is secure under non-adaptive chosen-ciphertext attacks in the presence of continual split-state leakage (CCmLA1), with min-entropy leakage bound λ , if $\Pr[b' = b]$ is at most negligibly greater than $\frac{1}{2}$ in the Experiment $\text{KEM-CCmLA1}_{\text{KEM}}(\mathcal{A}, \kappa, \lambda)$ for any efficient adversary \mathcal{A} .

Note that if in the above definition we would force the leakage functions to have output length of at most λ bits, then we would obtain the CCLA1 security for KEM as defined in [18]. From Lemma 1, we have that the conditional min-entropy of a random variable, given the leakage output of at most λ bits, cannot decrease by more than λ bits. Hence, if a KEM is CCLA1 secure, then it is also CCmLA1 secure.

2.2.1 Bilinear groups

Let $\text{BGen}'(\kappa, \lambda)$ be a probabilistic bilinear group generator that outputs $(\mathbb{G}, \mathbb{G}_T, q, e', g)$ such that:

1. $\mathbb{G} = \langle g \rangle$ and \mathbb{G}_T are (multiplicatively written) cyclic groups of prime order q with binary operations \cdot and \star , respectively. The size of q is κ bits.
2. $e' : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ is a map that is:
 - (a) bilinear: $\forall u, v \in \mathbb{G}$ and $\forall a, b \in \mathbb{Z}, e'(u^a, v^b) = e'(u, v)^{ab}$.
 - (b) non-degenerate: $e'(g, g) \neq 1$.

Such a group \mathbb{G} is said to be a bilinear group if the above properties hold and the group operations in \mathbb{G} and \mathbb{G}_T , and the map e' are efficiently computable. The group \mathbb{G} is called as *base group* and \mathbb{G}_T as *target group*.

2.2.2 Generic bilinear group model

The generic bilinear group (GBG) model [6] is an extension of the generic group model [31]. The encodings of the elements of \mathbb{G} and \mathbb{G}_T are given by random bijective maps $\xi : \mathbb{Z}_q \rightarrow \mathcal{E}$ and $\xi_T : \mathbb{Z}_q \rightarrow \mathcal{E}_T$, respectively, where \mathcal{E} and \mathcal{E}_T are sets of bit-strings. The group operations in \mathbb{G} and \mathbb{G}_T , and evaluation of the bilinear map e are performed by three public oracles $\mathcal{O}, \mathcal{O}_T$ and \mathcal{O}_e , respectively, defined as follows. For all $a, b \in \mathbb{Z}_q$

- $\mathcal{O}(\xi(a), \xi(b)) := \xi(a + b \bmod q)$
- $\mathcal{O}_T(\xi_T(a), \xi_T(b)) := \xi_T(a + b \bmod q)$
- $\mathcal{O}_e(\xi(a), \xi(b)) := \xi_T(ab \bmod q)$

We assume that the (fixed) generator g of \mathbb{G} satisfies $g = \xi(1)$, and also the (fixed) generator g_T of \mathbb{G}_T satisfies $g_T = e(g, g) = \xi_T(1)$. The encoding of g is provided to all users of the group oracles. The users can thus efficiently sample random elements in both \mathbb{G} and \mathbb{G}_T .

We further assume that $\mathcal{E} \cap \mathcal{E}_T = \phi, |\mathcal{E}| = |\mathcal{E}_T| = q$, and that the elements of \mathcal{E} and \mathcal{E}_T are efficiently recognizable. For instance, the encodings in \mathcal{E} can comprise of the binary representation of the set $\{0, 1, \dots, q - 1\}$, where every string begins with ‘0’ and all are of uniform length. The encodings in \mathcal{E}_T are similarly defined but instead begin with ‘1’. Since the encodings are efficiently recognizable, the queries to a group oracle with an invalid encoding can be detected and an error can be raised. For simplicity, we assume that the users’ queries to the oracles are all valid.

2.3 Bilinear ElGamal KEM

The scheme $\text{BEG} = (\text{KG}_{\text{BEG}}, \text{Enc}_{\text{BEG}}, \text{Dec1}_{\text{BEG}}, \text{Dec2}_{\text{BEG}})$ is as follows:

1. $\text{KG}_{\text{BEG}}(\kappa)$: Compute $\mathbb{PP} = (\mathbb{G}, \mathbb{G}_T, e', q, g) \leftarrow \text{BGen}'(\kappa, \lambda)$ and randomly choose $x, t_0 \xleftarrow{\$} \mathbb{F}_q$. Set $X = g^x, \sigma_0 = g^{t_0}, \sigma'_0 = g^{x-t_0}$, and $X_T = e'(g, g)^x$. Return (pk, sk_0) , where
 - (a) the public key is $pk = (\mathbb{PP}, X_T)$.
 - (b) the secret state is $sk_0 = (\sigma_0, \sigma'_0) \in \mathbb{G} \times \mathbb{G}$.
2. $\text{Enc}_{\text{BEG}}(pk)$: Choose a random $r \xleftarrow{\$} \mathbb{F}_q$. Compute the ciphertext $C = g^r$, and the derived key $K = X_T^r$. Return (C, K) .

3. $\text{Dec1}_{\text{BEG}}(\sigma_{i-1}, C)$: Choose a random $t_i \xleftarrow{\$} \mathbb{F}_q$, set $\sigma_i = \sigma_{i-1} \cdot g^{t_i}$, $Y_i = e'(\sigma_i, C)$. Return (t_i, Y_i) .
4. $\text{Dec2}_{\text{BEG}}(\sigma'_{i-1}, (t_i, Y_i), C)$: Set $\sigma'_i = \sigma'_{i-1} \cdot g^{-t_i}$, and $Y'_i = e'(\sigma'_i, C)$. Compute the derived key $K = Y_i \cdot Y'_i \in \mathbb{G}_T$. Return K .

The correctness of the scheme follows from the fact that $\sigma_i \cdot \sigma'_i = X \forall i \geq 0$ and using the bilinearity of $e'()$.

Theorem 1 [18, Theorem 1] *The scheme BEG (also called BEG-KEM) is CCLA1 secure in the generic bilinear group model. The advantage of an s -query adversary who gets at most λ bits of leakage per each invocation of Dec1_{BEG} or Dec2_{BEG} is at most $\frac{s^3}{q} 2^{2\lambda+1}$.*

3 A CCmLA1 security reduction in the generic bilinear group Model

We show that BEG-KEM is also leakage resilient in the min-entropy leakage model introduced above, where leakage functions are not necessarily size bounded. The only restriction is that the inputs to the leakage functions shall have enough min-entropy left, as a function of a leakage parameter λ , given the corresponding outputs. Interestingly, using a different proof technique than [19], we obtain a tighter bound on the adversarial CCLmA1 advantage than the bound claimed in [18] for the adversarial CCLA1 advantage, w.r.t. the number of oracle queries s . In other words, with respect to the previous work, we provide here a new security reduction under a more realistic leakage model, and surprisingly we achieve better tightness.

Theorem 2 *The scheme BEG-KEM is CCmLA1 secure in the GBG model. The advantage of an s -query adversary with min-entropy leakage bound λ is $\left(\frac{9s^2+3s}{q}\right) 2^{2\lambda}$.*

At a high level, the proof of this theorem proceeds in two steps as in [12, 13]. First we show in Theorem 3 that the scheme is secure if there is no leakage, i.e., CCA1 security. Note that the adversary is transparent to the internal details of secret state updates. Then, we complete the proof of CCmLA1 security by analyzing the effect of leakage on the CCA1 security.

The main idea to prove the CCA1 security is that the adversary will not be able to compute the derived symmetric key K_0 even after seeing the challenge ciphertext. To show this we just need to prove that K_0 cannot be written as a “linear combination” of the elements of \mathbb{G}_T that it has got as input or can compute itself using the pairing oracle along with the input elements of \mathbb{G} . Hence, in the GBG model it will not be able to distinguish the actual derived key or a randomly chosen key in \mathbb{G}_T . The challenger simulates the security game \mathcal{G}

to the adversary in the naive way. In addition, the challenger simulates the generic bilinear group oracles in the usual way by maintaining lists of pairs of encodings and polynomials that represent the relation amongst group elements.

We then argue that the proof for the non-leakage setting (i.e. proof of Theorem 3) and that for the leakage setting would be the same conditioned on the fact that the adversary is unable to derive useful relation amongst the elements it has seen or guessed, and that it will not be able to compute and hence leak the full secret key X through the leakage functions, if λ is sufficiently small. Finally, we show that the probability of this event is increased by a factor of at most $2^{2\lambda}$ compared to the non-leakage setting. The formal proof of the next theorem can be found in [11].

Theorem 3 *The scheme BEG is CCA1 secure in the generic bilinear group model, i.e., it is secure against non-adaptive chosen-ciphertext attacks if there is no leakage of the secret states. The advantage of an s -query adversary is at most $\frac{1}{2} + \frac{9s^2}{q}$.*

3.1 Leakage setting: completing proof of Theorem 2

Let us first briefly sketch the main ideas of the proof. Working on the lines of the proof of the previous theorem, the advantage of \mathcal{A} is bounded by its success probabilities conditioned on the event whether or not a collision has occurred in the lists consisting of elements of \mathbb{G} and \mathbb{G}_T . It is important to note that the proof for the non-leakage setting (i.e., proof of Theorem 3) and the leakage setting would be the same conditioned on the fact that a collision has not occurred, and that the leakage functions will not be able to compute the “polynomial X ” corresponding to the secret key nor guess the correct representations of the group elements for which it only partially obtains information through the leakage functions. The reason is that in the event of no collision, the adversary gets to see only distinct group elements and hence it will not have enough information on the relation amongst the group elements it can compute. The fact that the leakage functions cannot compute the full secret key shows that the adversary will never be able to continually leak the whole of the secret key. Hence, leakage on the secret state will not be useful in this case. Hence, the success probability of \mathcal{A} is the same in the event of no collision (that includes the event of guessing the representations of group elements using partial information about them).

However, the probability that a collision occurs in the leakage setting is increased by a factor of at most $2^{2\lambda}$. This is because when \mathcal{A} has access to leakage output $f_i(\sigma_{i-1}, t_i)$ and $h_i(\sigma'_{i-1}, (t_i, Y_i))$ during i th decryption query, then in adversary’s view the parameters t_i ($i \geq 1$) are no longer uniformly distributed even though they are still independent. Hence, \mathcal{A} can now cause collisions among polynomials with

increased probability. Since t_i appears in both $f_i()$ and $h_i()$, its (average conditional) min-entropy will be reduced by at most 2λ bits.

The only useful information that the leakage functions can provide to \mathcal{A} is about the secret key X . This is because the values t_i are independent of the derived shared secret key. However, \mathcal{A} can use the leakages of t_i to eventually leak X . If \mathcal{A} is able to compute X , then it can trivially compute the symmetric key corresponding to the challenge ciphertext. The event of no collision, and the fact that X is not a “linear combination” of the inputs to the leakage functions, guarantees that \mathcal{A} is unable to compute X . Note that because the representations of group elements in the GBG model are randomized, the probability of guessing the complete representations of each of σ_{i-1} , σ'_{i-1} and Y_i , given the leakages, is increased by a factor of at most $2^{2\lambda}$. For a formal proof see [11].

4 BEG-KEM+: a leakage-resilient KEM closer to practice

Our choice of BEG-KEM for this investigation is entirely motivated by the fact that a similar leakage resilience result as that proven in [18] cannot be expected for a pairing-less group, as shown in [14]. This motivates using pairing groups to implement ElGamal.

On the other hand, while Theorem ensures a protection against side-channel attacks that combine traces of different computations (e.g., differential power analysis attacks), we still need protection against single trace attacks, i.e. Simple Power Analysis (SPA). The use of pairing groups can help on this respect, as pointed out by Scott in [29]:

“[...] it is of interest to consider the resistance of pairing-based protocols to so-called SPA attacks [...] one might with reasonable confidence expect that the power consumption profile of (and execution time for) such protocols will be constant and independent of any secret values.”

We continue by proposing a tweak to BEG-KEM with the aim to make the most, from a minimizing leakage perspective, out of our choice of using pairing groups to realize leakage-resilient public-key cryptographic primitives.

4.1 An advanced BEG-KEM+ more resistant to side-channel attacks

Let us first make the observation that $\text{Dec1}_{\text{BEG}}^*$ is picking a random point in the pairing-based group \mathbb{G} by computing an exponentiation g^r for a random r . As is well known, a naïve implementation of exponentiation can leak the entire

exponent r , which would, of course, invalidate the required bound of maximum leakage in our new (as well as in the old) model. This leads us to the question whether it is possible, given the large body of side-channel-resistant exponentiation techniques, to find an algorithm that would likely meet the leakage bound for single measurements. In other words, we have to answer the question of whether the exponentiation can be made resistant against SPA attacks.

Exponentiation in a multiplicative group (or scalar multiplication in an elliptic curve group) of large order involves hundreds or even thousands of low-level arithmetic operations such as modular multiplication. Unfortunately, all these low-level operations are (either directly or indirectly) controlled by the secret exponent, which means that each of them can potentially leak sensitive information (see e.g., [33, 35, 36] for further details). Consequently, we need both an SPA-resistant exponentiation algorithm and an SPA-resistant implementation of the underlying multiple-precision operations. The latter is difficult to achieve in software due to side-channel leakage induced by certain micro-architectural features such as the early-termination mechanism of integer multipliers in ARM processors [16]. For example, it was shown in [16] that highly regular exponentiation (resp. scalar multiplication) techniques, which are (in theory) perfectly SPA-resistant, succumb to an SPA attack when exploiting the early-termination mechanism. Therefore, we avoid exponentiation with a secret exponent in our modified scheme.²

A careful analysis of BEG-KEM reveals that $\text{Dec1}_{\text{BEG}}^*$ only needs to sample uniformly at random an element u of \mathbb{G} , and that knowledge of $\log_g u$ is *not necessary*. It suffices then to use a method that computes a random point in the base group.

One possibility is to use a variant of the so-called *try-and-increment* approach [7, 34], where a random coordinate x for an elliptic curve point is chosen; next if a point in the curve exists with that x -coordinate, its y -coordinate is computed and the procedure is stopped. Otherwise, the procedure is iterated until a point in the curve is found. We have chosen not to follow this approach, in particular because its running time depends on the consecutive seeds x_1, x_2, \dots used, which could eventually lead to timing leakages or attacks. It should be noted that the original try-and-increment approach has been found to be vulnerable to timing attacks in some contexts (when used to build a Password Authenticated Key-Exchange protocol, see [34, Section 3.2] for details). We prefer to use instead a method that will run in (almost) con-

² As mentioned previously, the secret exponent controls a large number of multiple-precision arithmetic operations, which execute an even larger number of `mul` instructions. Each of these `mul` instructions can potentially trigger the early-termination mechanism and, hence, leak information about the secret exponent. In our modified scheme, the secret value is only used as input of a multiple-precision operation and does not control any other operations.

stant time, which is a common approach when thwarting timing attacks.

For this reason, we decided to build a random element in the pairing base group using a so-called *encoding* to the base group [10, 17, 30]. Roughly speaking, an encoding is a deterministic function mapping an arbitrary string to a point in an elliptic curve. Recently, Fouque and Tibouchi [10] proposed a modification of the Shallue and van de Woestijne encoding into arbitrary elliptic curves [30] that maps arbitrary strings to Barreto-Naehrig asymmetric pairing groups [3]. Let $f : \mathbb{F}_p^* \rightarrow E(\mathbb{F}_p)$ be the Fouque-Tibouchi encoding. Then, $(t_1, t_2) \mapsto u = u_1 \cdot_E u_2$ builds a point $u \in E(\mathbb{F}_p)$ distributed uniformly at random if $t_1, t_2 \xleftarrow{\$} \mathbb{F}_p^*$, where \cdot_E is the addition operation in $E(\mathbb{F}_p)$. Additionally, [10] points out that f can be naturally implemented so that its computation is completely independent of the inputs, which clearly helps us towards meeting our desired min-entropy leakage bound.

4.2 BEG-KEM+

Let **ABGen** be an asymmetric bilinear group generator that outputs $(\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, q, g_1, g_2)$ with $|\mathbb{G}_1| = |\mathbb{G}_2| = |\mathbb{G}_T| = q$, where q is a prime, κ be the security parameter, and λ be the leakage parameter. We will again use the multiplicative notation for group operations in $\mathbb{G}_1, \mathbb{G}_2$, and \mathbb{G}_T . Let $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ be a type 3 pairing map, i.e., e is a non-degenerate bilinear map with no known efficiently computable isomorphism $\psi : \mathbb{G}_2 \rightarrow \mathbb{G}_1$. These groups are instantiated using the BN curves, denoted $E(\mathbb{F}_p)$, of the form $y^2 = x^3 + b$, where $b \in \mathbb{F}_p$ [3]. In addition, let G_1 and G_2 be generators of \mathbb{G}_1 and \mathbb{G}_2 , respectively, and $f : \mathbb{F}_p^* \rightarrow \mathbb{G}_1$ be the Fouque-Tibouchi encoding of the elements of \mathbb{G}_1 .

The advanced **BEG – KEM+** = $(\text{KG}_{\text{BEG}}^+, \text{Enc}_{\text{BEG}}^+, \text{Dec1}_{\text{BEG}}^+, \text{Dec2}_{\text{BEG}}^+)$ is defined as follows:

1. $\text{KG}_{\text{BEG}}^+(\kappa)$: Compute $\mathbb{PP} = (\mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, q, G_1, G_2) \leftarrow \text{ABGen}(\kappa)$ and randomly choose $x, t_0 \xleftarrow{\$} \mathbb{F}_q$. Set $X = G_1^x, \sigma_0 = G_1^{t_0}, \sigma'_0 = G_1^{(x-t_0)}$, and $X_T = e(G_1, G_2)^x$. Return (pk, sk_0) , where
 - (a) the public key is $pk = (\mathbb{PP}, X_T)$.
 - (b) the secret state is $sk_0 = (\sigma_0, \sigma'_0)$.
2. $\text{Enc}_{\text{BEG}}^+(pk)$: Choose a random $r \xleftarrow{\$} \mathbb{F}_p$. Compute the ciphertext $C = G_2^r$, and the derived key $K = X_T^r$. Return (C, K) .
3. $\text{Dec1}_{\text{BEG}}^+(\sigma_{i-1}, C)$: Choose random $t_i, z_i \xleftarrow{\$} \mathbb{F}_p^*$, set $u_i = f(t_i) \cdot f(z_i)$, and compute $\sigma_i = \sigma_{i-1} \cdot u_i, Y_i = e(\sigma_i, C)$. Return (u_i, Y_i) .
4. $\text{Dec2}_{\text{BEG}}^+(\sigma'_{i-1}, (u_i, Y_i), C)$: Set $\sigma'_i = \sigma'_{i-1} \cdot (u_i)^{-1}$, and $Y'_i = e(\sigma'_i, C)$. Compute the derived key $K = Y_i \cdot Y'_i \in \mathbb{G}_T$. Return K .

Algorithm 1 Shallue-van de Woestijne encoding to BN curves $y^2 = x^3 + b$ [10]

Input: A random number $t \in \mathbb{F}_p^*$.
Output: Point $P \in E(\mathbb{F}_p)$

- 1: $w \leftarrow \sqrt{-3} \cdot t / (1 + b + t^2)$
- 2: $x_1 \leftarrow (-1 + \sqrt{-3}) / 2 - tw$
- 3: $x_2 \leftarrow -1 - x_1$
- 4: $x_3 \leftarrow 1 + 1/w^2$
- 5: $r_1, r_2, r_3 \xleftarrow{\$} \mathbb{F}_p^*$
- 6: $\alpha \leftarrow \chi_q(r_1^2 \cdot (x_1^3 + b))$
- 7: $\beta \leftarrow \chi_q(r_2^2 \cdot (x_2^3 + b))$
- 8: $i \leftarrow [(\alpha - 1) \cdot \beta \bmod 3] + 1$
- 9: **return** $P[x_i, \chi_q(r_3^2 \cdot t) \cdot \sqrt{(x_i^3 + b)}]$

Algorithm 1 describes the constant-time hashing function to BN curves from [10]. As described in the original paper, implementing this algorithm against timing and simple power analysis (SPA) attacks is not difficult to be achieved. In step 6 and 7, instead of computing the values $\chi_q(x_1^3 + b)$ and $\chi_q(x_2^3 + b)$ in a straightforward way, which can leak secret data, the authors suggested to use blinding. Namely, to get α and β , we actually evaluate $\chi_q(r_1^2 \cdot (x_1^3 + b))$ and $\chi_q(r_2^2 \cdot (x_2^3 + b))$, where r_1 and r_2 are random field elements generated in Step 5. On the other hand, to prevent the leakage while computing the index i , they employ a specific algebraic function $\phi(\alpha, \beta) = [(\alpha - 1) \cdot \beta \bmod 3] + 1$, which runs in constant time.

5 Secure implementation and performance analysis

In this section, we first describe a software implementation of BEG-KEM+ (along with the instantiation of the underlying pairing groups) and present the execution times we measured on an ARM Cortex M-3 processor. The second part of this section is devoted to a “practical” security evaluation of BEG-KEM+ by analyzing potential sources of information leakage in the underlying arithmetic operations that could be exploited to mount a side-channel attack.

5.1 Implementation details and performance analysis

We implemented both BEG-KEM and BEG-KEM+ in Magma and ANSI C, whereby the former implementation served as a reference for the latter. The C implementation is based on the MIRACL library to ensure an efficient execution of the pairing evaluation and all other arithmetic operations in the diverse groups and fields. We instantiated both BEG-KEM and our improved scheme using the Ate pairing over a 254-bit Barreto-Naehrig (BN) curve. More specifically, our implementations adopts the curve BN254 from [26], which provides a security level roughly comparable to that of 128-

Table 2 Running times for field exponentiation, square root, inversion, group exponentiation and pairing operations (in 10^6 clock cycles)

Operation	Running time
Square root \mathbb{F}_q	0.7
Inversion \mathbb{F}_q	0.087
Encoding to \mathbb{G}_2	3.7
Exponentiation \mathbb{G}_1	4.5
Exponentiation \mathbb{G}_2	10.0
Exponentiation \mathbb{G}_T	27.1
Pairing	65.0

Table 3 Comparison of running times for key generation, encapsulation and decapsulation for BEG-KEM and BEG-KEM+ (in 10^6 clock cycles)

Operation	BEG-KEM	BEG-KEM+
KeyGen	108	108
Encryption	34	34
Decryption	131	140

bit AES. BN curves are defined by a Weierstrass equation of the form $y^2 = x^3 + b$ over a prime field \mathbb{F}_q , whereby q can be written as polynomial $p(u) = 36u^4 + 36u^3 + 24u^2 + 6u + 1$ for some parameter u [3]. In our case, $u = -(2^{62} + 2^{55} + 1) = -0x4080000000000001$ and, hence, q has a length of 254 bits. The curve BN254 is given by the equation $y^2 = x^3 + 2$ (i.e. $b = 2$) and has prime order with embedding degree $k = 12$.

The execution times for various arithmetic operations in the different fields and groups are summarized in Table 2, whereby all timings are specified in millions of clock cycles. Our prototype platform for performance evaluation is an Arduino Due microcontroller board equipped with an ARM Cortex-M3 CPU. Even though the three groups \mathbb{G}_1 , \mathbb{G}_2 , and \mathbb{G}_T have the same order, the underlying multiple-precision arithmetic operations are performed with operands of different size. \mathbb{G}_1 and \mathbb{G}_2 are elliptic curve groups over \mathbb{F}_q and \mathbb{F}_{q^2} , the elements of which have, in our case, a bitlength of 254 and 508 bits, respectively. The group \mathbb{G}_T is a subgroup of the multiplicative group of the extension field $\mathbb{F}_{q^{12}}$, i.e. the modular multiplications for exponentiation in \mathbb{G}_T are carried out on 3048-bit operands.

The execution times for key generation, encapsulation as well as decapsulation for both BEG-KEM and BEG-KEM+ are given in Table 3. Our results show that an encapsulation can be carried out in 34 million clock cycles, while the decapsulation takes about 140 million cycles. We observe that our modified decapsulation algorithm is roughly 6 % slower than the original one.

5.2 Side-channel resistance from a practical point of view

One of the fundamental principles of leakage-resilient cryptography is to use a critical secret only once (or a few times), which ensures that an attacker is not able to retrieve the secret key if the per-invocation leakage is in some way “limited” or “bounded.” In every invocation of the scheme or function, the secret is either “refreshed” or a completely new secret is generated randomly. The original BEG-KEM scheme from [18], and also our variant BEG-KEM+, follow this principle. As a consequence, all forms of side-channel attack that require several executions of a cryptographic function with one and the same secret key, e.g., differential power analysis (DPA), are obviously not applicable to BEG-KEM+ (and in fact the latter is guaranteed by Theorem 2). However, attacks that aim to recover the secret key from information leaked from a single invocation of a cryptographic function (i.e. SPA attacks) may succeed under certain conditions. The group exponentiation computed in the BEG-KEM scheme to derive a random group element $\sigma_0 = g^{t_0}$ serves as a good example. If this exponentiation is implemented in completely straightforward way (e.g., using the square-and-multiply method) an attacker can obtain t_0 if he is able to distinguish group squarings from group products in the power consumption profile. Such SPA attacks on unprotected or insufficiently protected ECC implementations are fairly easy and have been reported extensively in the literature, see, e.g., [5, Chapter IV] and the references therein. Therefore, we advocated to replace the aforementioned group exponentiation by a deterministic encoding into an elliptic curve group [10].

5.2.1 SPA resistance of pairing evaluation

Section 4.1 quotes a statement of Scott [29, Section 3.1] saying that one can expect the power consumption profile of a pairing-based protocol to be independent of any secret values. An intuitive explanation why pairings are fairly “robust” against SPA leakage is also given in [29]: the target of the attack is a secret point, which is generally much harder to reveal than, e.g., a secret scalar or a secret exponent. As mentioned before, our implementation uses the Ate pairing instantiated on a BN curve over a 254-bit prime field \mathbb{F}_p . Consequently, the secret is the x and y coordinate of an elliptic curve point, which are in our case simply elements of \mathbb{F}_p . The only way in which an attacker can hope to gain information about x and y is by inspecting the power consumption and execution time of the \mathbb{F}_p -arithmetic operations (e.g., addition, multiplication) performed on them. However, the operand-related SPA leakage from field-arithmetic operations is generally very small. To explain this in detail, let us use the addition in \mathbb{F}_p as example, which is nothing else than a modular addition $r = a + b \bmod p$. We assume that a is a

secret value and that b is known to the attacker. A modular addition consists of an ordinary addition $s = a + b$, followed by a subtraction if the sum s is equal to or bigger than p . Conventional wisdom from the side-channel community says that such a conditional subtraction causes differences in the power consumption profile (and also execution time), which is observable by an attacker. However, the information content is very small; in fact, when the subtraction is executed the attacker just knows that $a + b \geq p$, i.e. he has learned that $a \geq p - b$.

The situation is similar for multiplication in \mathbb{F}_p , which is nothing else than a modular multiplication $r = a \cdot b \bmod p$. Again, we assume that a is the secret value and that b is known to the attacker. A modular multiplication involves a conventional multiplication $t = a \cdot b$, followed by a modular reduction $r = t \bmod p$, which is in pairing-based cryptography typically implemented using Montgomery’s algorithm [23]. Both the multiplication and Montgomery reduction are highly regular (i.e. do not have to execute any conditional statements), except for the so-called “final subtraction.” Montgomery’s reduction technique does not directly compute $t \bmod p$ but produces the following output

$$x = (t + (t \cdot p' \bmod 2^n) \cdot p) / 2^n \tag{1}$$

where $p' = -p^{-1} \bmod 2^n$ and n is the bitlength of p . Note that x may be not fully reduced, which means a final subtraction of p is necessary to get the least non-negative residue as result. An attacker able to observe whether or not this final subtraction is executed learns only whether $x \geq p$ or not, which does not reveal much information about a . The same also holds for subtraction and squaring in \mathbb{F}_p . However, a noteworthy exception is the inversion operation, which we will further discuss below. In summary, a straightforward implementation of the arithmetic operations (bar inversion) in \mathbb{F}_p leaks only very little information about the operands, which confirms that pairing evaluation is, in general, not susceptible to SPA attacks. To our knowledge, the recent literature contains only two papers in which SPA attacks on pairings are discussed [25,37], but both of them are only relevant for pairings over binary fields where the multiplication is implemented in a highly irregular way. The attack from [35] is only applicable to scalar multiplication with a secret scalar, but not to pairings with a secret point.

5.2.2 SPA resistance of encoding function

The encoding function shown in Algorithm 1 consists of a number of basic arithmetic operations (e.g., addition, multiplication) in the field \mathbb{F}_p . Furthermore, two inversions are executed, one in step 1 and the other in step 4. The straightforward approach to invert an element of a finite field is the

Extended Euclidean Algorithm (EEA). Conventional wisdom from the side-channel community says that the EEA is a highly irregular algorithm, executing many conditional operations, which is likely to leak SPA-relevant information about the operand to be inverted. To prevent an SPA attack on the inversion operation, we apply a simple multiplicative masking; that is, instead of inverting a field element v directly, we first multiply it by a random number r , which yields the product $t = v \cdot r$. Then, we invert this product using the EEA to obtain $1/t = 1/(v \cdot r)$, which we finally multiply again by r to get $1/v$ as result.

The function χ in step 6 and 7 of Algorithm 1 is essentially an evaluation of the Legendre Symbol, which, in turn, consists of an exponentiation using a constant public exponent (i.e., $(p+1)/4$). The input to the χ function is “blinded” by the random value r_1^2 and r_2^2 , which means the underlying exponentiation cannot leak any SPA-relevant information. As mentioned in Sect. 4.1, a constant-time algebraic function is adopted for the calculation of the index i in step 8, which also cannot leak.

6 Conclusion

In this paper, we aimed to bring the concept of leakage-resilient cryptography closer to practice. Most of the leakage-resilient public-key cryptography schemes proposed until now are too inefficient for real-world applications. Even though they provide provable security against a large class of side-channel attacks, they do so under certain leakage models and leakage bound requirements that are far from what we can ensure in practice. On the other hand, the side-channel countermeasures are often ad hoc and do not provide enough security guarantees. We addressed this problem by bringing best practices from both worlds together. First, we argued that a naive implementation of the pairing group exponentiation in the leakage-resilient ElGamal key encapsulation mechanism proposed by Kiltz and Pietrzak makes it impossible to reach the required leakage bound. To overcome this problem, we have made two additional contributions. On the one hand, we have proposed a relaxed leakage model, that we call min-entropy leakage, that lifts the restriction on the image size of leakage functions, and proposes instead to require that the inputs to the leakage functions have sufficient min-entropy left, in spite of the leakage. On the other hand, we adopted a different mechanism for finding a random point in an elliptic curve group, namely the encoding of Fouque and Tibouchi. We assessed the security of our implementation from both a theoretical and a practical perspective and argued that it is indeed secure in both the worlds. BEG-KEM+ is, to our knowledge, the first leakage-resilient public-key scheme that has been successfully implemented and evaluated on an embedded 32-bit processor.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Akavia, A., Goldwasser, S., Vaikuntanathan, V.: Simultaneous hardcore bits and cryptography against memory attacks. In: Reingold, O. (ed.) TCC, LNCS, vol. 5444, pp. 474–495. Springer (2009)
- Alwen, J., Dodis, Y., Wichs, D.: Leakage-resilient public-key cryptography in the bounded-retrieval model. In: CRYPTO, pp. 36–54 (2009)
- Barreto, P.S.L.M., Naehrig, M.: Pairing-friendly elliptic curves of prime order. In: Preneel, B., Tavares, S.E. (eds.) Selected Areas in Cryptography, LNCS, vol. 3897, pp. 319–331. Springer (2005)
- Belaïd, S., Grosso, V., Standaert, F.X.: Masking and leakage-resilient primitives: one, the other(s) or both? Cryptology ePrint archive, report 2014/053 (2014)
- Blake, I.F., Seroussi, G., Smart, N.P.: Advances in Elliptic Curve Cryptography, London Mathematical Society Lecture Notes Series, vol. 317. Cambridge University Press, Cambridge (2005)
- Boneh, D., Boyen, X., Goh, E.J.: Hierarchical Identity Based Encryption with constant size ciphertext. In: Cramer, R. (ed.) EUROCRYPT, LNCS, vol. 3494, pp. 440–456. Springer (2005)
- Boneh, D., Lynn, B., Shacham, H.: Short signatures from the weil pairing. In: Boyd, C. (ed.) Advances in Cryptology—ASIACRYPT 2001, 7th International Conference on the Theory and Application of Cryptology and Information Security, Gold Coast, Australia, 9–13 December 2001, Proceedings, Lecture Notes in Computer Science, vol. 2248, pp. 514–532. Springer (2001). doi:10.1007/3-540-45682-1_30
- Dodis, Y., Goldwasser, S., Kalai, Y.T., Peikert, C., Vaikuntanathan, V.: Public-key encryption schemes with auxiliary inputs. In: Micciancio, D. (ed.) TCC, LNCS, vol. 5978, pp. 361–381. Springer (2010)
- Dodis, Y., Ostrovsky, R., Reyzin, L., Smith, A.: Fuzzy extractors: how to generate strong keys from biometrics and other noisy data. SIAM J. Comput. **38**(1), 97–139 (2008)
- Fouque, P.A., Tibouchi, M.: Indifferentiable hashing to Barreto–Naehrig curves. In: LATINCRYPT, pp. 1–17 (2012)
- Galindo, D., Großschädl, J., Liu, Z., Vadnala, P.K., Vivek, S.: Implementation and evaluation of a leakage-resilient elgamal key encapsulation mechanism. Cryptology ePrint archive, report 2014/835 (2014)
- Galindo, D., Vivek, S.: A leakage-resilient pairing-based variant of the Schnorr signature scheme. In: Stam, M. (ed.) IMA International Conference, LNCS, vol. 8308, pp. 173–192. Springer (2013)
- Galindo, D., Vivek, S.: A practical leakage-resilient signature scheme in the generic group model. In: SAC 2012, LNCS, vol. 7707, pp. 50–65. Springer (2013)
- Galindo, D., Vivek, S.: Limits of a conjecture on a leakage-resilient cryptosystem. Inf. Process. Lett. **114**(4), 192–196 (2014)
- Gandolfi, K., Mourtél, C., Olivier, F.: Electromagnetic analysis: Concrete results. In: Çetin Kaya Koç, D. Naccache, C. Paar (eds.) CHES, LNCS, vol. 2162, pp. 251–261. Springer (2001)
- Großschädl, J., Oswald, E., Page, D., Tunstall, M.: Side-channel analysis of cryptographic software via early-terminating multiplications. In: Lee, D., Hong, S. (eds.) Information Security and Cryptology—ICISC 2009, LNCS, vol. 5984, pp. 176–192. Springer (2010)
- Icart, T.: How to hash into elliptic curves. In: CRYPTO, pp. 303–316 (2009)
- Kiltz, E., Pietrzak, K.: Leakage resilient ElGamal encryption. In: Abe, M. (ed.) ASIACRYPT, LNCS, vol. 6477, pp. 595–612. Springer (2010)
- Kiltz, E., Pietrzak, K.: Leakage resilient ElGamal encryption. Full version of [18]. http://homepage.ruhr-uni-bochum.de/Eike.Kiltz/papers/elgamal_leak.pdf. Accessed 4 June 2014 (2010)
- Kocher, P.C.: Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In: Kobitz, N. (ed.) CRYPTO, LNCS, vol. 1109, pp. 104–113. Springer (1996)
- Kocher, P.C., Jaffe, J., Jun, B.: Differential power analysis. In: Wiener, M.J. (ed.) CRYPTO, LNCS, vol. 1666, pp. 388–397. Springer (1999)
- Micali, S., Reyzin, L.: Physically observable cryptography (extended abstract). In: Naor, M. (ed.) TCC, LNCS, vol. 2951, pp. 278–296. Springer (2004)
- Montgomery, P.L.: Modular multiplication without trial division. Math. Comput. **44**(170), 519–521 (1985)
- Moradi, A.: Statistical tools flavor side-channel collision attacks. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT, LNCS, vol. 7237, pp. 428–445. Springer (2012)
- Page, D., Vercauteren, F.: Fault and side-channel attacks on pairing based cryptography. IACR Cryptol. ePrint Arch. **2004**, 283 (2004)
- Pereira, G.C., Simplicio, M.A., Naehrig, M., Barreto, P.S.: A family of implementation-friendly BN elliptic curves. J. Syst. Softw. **84**(8), 1319–1326 (2011)
- Prouff, E., Rivain, M.: Masking against side-channel attacks: a formal security proof. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT, LNCS, vol. 7881, pp. 142–159. Springer (2013)
- Schwartz, J.T.: Fast probabilistic algorithms for verification of polynomial identities. J. ACM **27**(4), 701–717 (1980)
- Scott, M.: Computing the Tate pairing. In: Menezes, A. (ed.) CT-RSA, LNCS, vol. 3376, pp. 293–304. Springer (2005)
- Shallue, A., van de Woestijne, C.: Construction of rational points on elliptic curves over finite fields. In: Hess, F., Pauli, S., Pohst, M.E. (eds.) ANTS, LNCS, vol. 4076, pp. 510–524. Springer (2006)
- Shoup, V.: Lower bounds for discrete logarithms and related problems. In: Fumy, W. (ed.) EUROCRYPT, LNCS, vol. 1233, pp. 256–266. Springer (1997)
- Standaert, F.X., Pereira, O., Yu, Y.: Leakage-resilient symmetric cryptography under empirically verifiable assumptions. In: Canetti, R., Garay, J.A. (eds.) CRYPTO (1), LNCS, vol. 8042, pp. 335–352. Springer (2013)
- Stebila, D., Thériault, N.: Unified point addition formulæ and side-channel attacks. In: Goubin, L., Matsui, M. (eds.) Cryptographic Hardware and Embedded Systems—CHES 2006, LNCS, vol. 4249, pp. 354–368. Springer (2006)
- Tibouchi, M.: A note on hashing to bn curves. In: 29th Japanese Symposium on Cryptography and Information Security—SCIS 2012. <http://www.normalesup.org/~tibouchi/papers/bnhash-scis.pdf> (2012). Accessed 27 Feb 2016
- Walter, C.D.: Simple power analysis of unified code for ECC double and add. In: Joye, M., Quisquater, J.J. (eds.) Cryptographic Hardware and Embedded Systems—CHES 2004, LNCS, vol. 3156, pp. 191–204. Springer (2004)
- Walter, C.D., Thompson, S.: Distinguishing exponent digits by observing modular subtractions. In: Naccache, D. (ed.) Topics in Cryptology—CT-RSA 2001, LNCS, vol. 2020, pp. 192–207. Springer (2001)
- Whelan, C., Scott, M.: Side channel analysis of practical pairing implementations: Which path is more secure? In: Nguyen, P.Q. (ed.) VIETCRYPT, LNCS, vol. 4341, pp. 99–114. Springer (2006)
- Zippel, R.: Probabilistic algorithms for sparse polynomials. In: Ng, E.W. (ed.) EUROSAM, LNCS, vol. 72, pp. 216–226. Springer (1979)