Bounded Tamper Resilience: How to go beyond the Algebraic Barrier

Ivan Damgård[1], Sebastian Faust[2], Pratyay Mukherjee[1], and Daniele Venturi[1]

[1] Department of Computer Science, Aarhus University
[2] Security and Cryptography Laboratory, EPFL

Abstract. Related key attacks (RKAs) are powerful cryptanalytic attacks where an adversary can change the secret key and observe the effect of such changes at the output. The state of the art in RKA security protects against an a-priori unbounded number of certain algebraic induced key relations, e.g., affine functions or polynomials of bounded degree. In this work, we show that it is possible to go beyond the algebraic barrier and achieve security against *arbitrary* key relations, by restricting the number of tampering queries the adversary is allowed to ask for. The latter restriction is necessary in case of arbitrary key relations, as otherwise a generic attack of Gennaro *et al.* (TCC 2004) shows how to recover the key of almost any cryptographic primitive. We describe our contributions in more detail below.

1. We show that standard ID and signature schemes constructed from a large class of Σ-protocols (including the Okamoto scheme, for instance) are secure even if the adversary can *arbitrarily* tamper with the prover's state a *bounded* number of times and obtain some bounded amount of leakage. Interestingly, for the Okamoto scheme we can allow also independent tampering with the public parameters.
2. We show a *bounded* tamper and leakage resilient CCA secure public key cryptosystem based on the DDH assumption. We first define a weaker CPA-like security notion that we can instantiate based on DDH, and then we give a general compiler that yields CCA-security with tamper and leakage resilience. This requires a public tamper-proof common reference string.
3. Finally, we explain how to boost bounded tampering and leakage resilience (as in 1. and 2. above) to *continuous* tampering and leakage resilience, in the so-called *floppy model* where each user has a personal hardware token (containing leak- and tamper-free information) which can be used to refresh the secret key.

We believe that bounded tampering is a meaningful and interesting alternative to avoid known impossibility results and can provide important insights into the security of existing standard cryptographic schemes.

Keywords: related key security, bounded tamper resilience, public key encryption, identification schemes

1 Introduction

Related key attacks (RKAs) are powerful cryptanalytic attacks against a cryptographic implementation that allow an adversary to change the key, and subsequently observe the effect of such modification on the output. In practice, such attacks can be carried out, e.g., by heating up the device or altering the internal power supply or clock [4,11], and may have severe consequences for the security of a cryptographic implementation. To illustrate such key tampering, consider a digital signature scheme Sign with public/secret key pair (pk, sk). The tampering adversary obtains pk and can replace sk with $T(sk)$ where T is some arbitrary tampering function. Then, the adversary gets access to an oracle $\mathsf{Sign}(T(sk), \cdot)$, i.e., to a signing oracle running with the tampered key $T(sk)$. As usual the adversary wins the game by outputting a valid forgery with respect to the original public key pk. Notice that T may be the identity function, in which case we get the standard security notion of digital signature schemes.

Bellare and Kohno [8] pioneered the formal security analysis of cryptographic schemes in the presence of related key attacks. In their setting an adversary tampers *continuously* with the key by applying functions T chosen from a set of *admissible* tampering functions $\mathcal{T}$. In the signature example from above, each signing query for message m would be accompanied with a tampering function $T \in \mathcal{T}$ and the adversary obtains $\mathsf{Sign}(T(sk), m)$. Clearly, a result in the RKA setting is stronger if the class of admissible functions $\mathcal{T}$ is larger, and hence several recent works have focussed on further broadening $\mathcal{T}$. The current state of the art (see discussion in Section 1.2) considers certain algebraic relations of the key, e.g., $\mathcal{T}$ is the set of all affine functions or all polynomials of bounded degree. A natural question that arises from these works is if we can further broaden the class of tampering functions — possibly showing security for *arbitrary* relations. In this work, we study this question and show that under certain assumptions security against arbitrary key relations can be achieved.

Is arbitrary key tampering possible? Unfortunately, the answer to the above question in its most general form is negative. As shown by Gennaro *et al.* [25], it is *impossible* to protect any cryptographic scheme against arbitrary key relations. In particular, there is an attack that allows to recover the secret key of most stateless cryptographic primitives after only a few number of tampering queries.[3] To prevent this attack the authors propose to use a *self-destruct* mechanism. That is, before each execution of the cryptographic scheme the key is checked for its validity. In case the key was changed the device self-destructs. In practice, such self-destruct can for instance be implemented by overwriting the secret key with the all-zero string, or by switching to a special mode in which the device outputs $\perp$.[4] In this work, we consider an alternative setting to avoid the

[3] The impossibility result of [25] leaves certain loopholes, which however seem very hard to exploit.

[4] We notice that the self-destruct has to be permanent as otherwise the attack of [25] may still apply.

impossibility results of [25], and assume that an adversary can only carry out a bounded number of (say t) tampering queries. To explain our setting consider again the example of a digital signature scheme. In our model, we give the adversary access to t tampered signing oracles $\mathsf{Sign}(T_i(sk), \cdot)$, where T_i can be an arbitrary adaptively chosen tampering function. Notice that of course each of these oracles can be queried a polynomial number of times, while t is typically linear in the security parameter.

Is security against bounded tampering useful? Besides from being a natural and non-trivial security notion, we believe that our adversarial model of *arbitrary*, *bounded* tampering is useful for a number of reasons:

1. It is a natural alternative to continuous restricted tampering: our security notion of *bounded*, *arbitrary* tampering is orthogonal to the traditional setting of RKA security where the adversary can tamper *continuously* but is *restricted* to certain classes of attacks. Most previous work in the RKA setting considers algebraic key relations that are tied to the scheme's algebra and may not reflect attacks in practice. For instance, it is not clear that heating up the device or shooting with a laser on the memory can be described by, e.g., an affine function — a class that is usually considered in the literature. We also notice that physical tampering may completely destroy the device, or may be detected by hardware countermeasures, and hence our model of bounded but arbitrary tampering may be sufficient in such settings.
2. It allows to analyze the security of *standard* cryptoschemes: as outlined above a common countermeasure to protect against arbitrary tampering is to implement a key validity check and self-destruct (or output a special failure symbol) in case such check fails. Unfortunately, most standard cryptographic implementations do not come with such a built-in procedure to check the validity of the key. Our notion of bounded tamper resilience allows to make formal security guarantees of *standard* cryptographic schemes where neither the construction, nor the implementation needs to be specially engineered.
3. It can be a useful as a building-block: even if the restriction of bounded tamper resilience may be too strong in some settings, it can be useful to achieve results in the stronger continuous tampering setting (we provide some first preliminary results on this in the full version [17]). Notice that this is similar to the setting of leakage resilient cryptography which also started mainly with "bounded leakage" that later turned out to be very useful to get results in the continuous leakage setting.

We believe that due to the above points the bounded tampering model is an interesting alternative to avoid known impossibility results for arbitrary tampering attacks.

1.1 Our Contribution

We initiate a general study of schemes resilient to both *bounded* tamper and leakage attacks. We call this model the *bounded leakage and tampering model*

Tampering Model	ID Schemes		IND-CCA PKE
	Σ-Protocols	Okamoto	BHHO
Secret Key	✓	✓	✓
Public Parameters	n.a.	✓	n.a.
Continuous Tampering iFloppy	✓	✓	✓
Key Length	$\log \lvert\mathcal{X}\rvert$	$\ell \log p$	$\ell \log p$
Tampering Queries	$\lfloor \log \lvert\mathcal{X}\rvert / \log \lvert\mathcal{Y}\rvert \rfloor - 2$	$\ell - 2$	$\ell - 3$

Table 1. An overview of our results for bounded leakage and tamper resilience. All parameters $|\mathcal{X}|$, $|\mathcal{Y}|$ ℓ, p and n are a function of the security parameter k. For the case of Σ-protocol, the set $\mathcal{X}$ is the set of all possible witnesses and the set $\mathcal{Y}$ is the set of all possible statements for the language; we can achieve a better bound depending on the conditional average min-entropy of the witness given the statement (cf. Section 3).

(BLT) model. While our general techniques use ideas from the leakage realm, we emphasize that bounded leakage resilience does *not* imply bounded tamper resilience. In fact, it is easy to find contrived schemes that are leakage resilient but completely break for a single tampering query. At a more technical level, we observe that a trivial strategy using leakage to simulate, e.g., faulty signatures, has to fail as the adversary can get any polynomial number of faulty signatures — which clearly cannot be simulated with bounded leakage only. Nevertheless, as we show in this work, we are able to identify certain classes of cryptoschemes for which a small amount of leakage is sufficient to simulate faulty outputs. We discuss this in more detail below.

Our concrete schemes are proven secure under standard assumptions (DL, factoring or DDH) and are efficient and simple. Moreover, we show that our schemes can easily be extended to the continual setting by putting an additional simple assumption on the hardware. We elaborate more on our main contributions in the following paragraphs (see also Table 1.1 for an overview of our results). Importantly, all our results allow arbitrary key tampering and do not need any kind of tamper detection mechanism.

Identification schemes. It is well known that the Generalized Okamoto identification scheme [34] provides security against bounded leakage from the secret key [3,30]. In Section 3, we show that additionally it provides strong security against tampering attacks. While in general the tampered view may contain a polynomial number of faulty transcripts that may potentially reveal a large amount of information about the secret key, we can show that fortunately this is not the case for the Generalized Okamaoto scheme. More concretely, we are able to identify a short amount of information that for each tampering query allows us to simulate any number of corresponding faulty transcripts. Hence, BLT security of the Generalized Okamoto scheme is implied by its leakage resilience.

Our results on the Okamoto identification can be further generalized to a large class of identification schemes (and signature schemes based on the Fiat-Shamir heuristic). More concretely, we show that Σ-protocols where the secret

key is significantly longer than the public key are BLT secure. We can instantiate our result with the generalized Guillou-Quisquater ID scheme [27], and its variant based on factoring [24] yielding tamper resilient identification based on factoring. We give more details in Section 3.

Interestingly, for Okamoto identification security still holds in a stronger model where the adversary is allowed to tamper not only with the secret key of the prover, but also with the description of the public parameters (i.e., the generator g of a group $\mathbb{G}$ of prime order p). The only restrictions are: (i) tampering with the public parameters is independent from tampering with the secret key and (ii) the tampering with public parameters must map to its domain. We also show that the latter restrictions are necessary, by presenting explicit attacks when the adversary can tamper jointly with the secret key and the public parameters or he can tamper the public parameters to some particular range.

Public key encryption. We show how to construct IND-CCA secure public key encryption (PKE) in the BLT model. To this end, we first introduce a weaker CPA-like security notion, where an adversary is given access to a restricted (faulty) decryption oracle. Instead of decrypting adversarial chosen ciphertexts such an oracle accepts inputs (m, r), encrypts the message m using randomness r under the original public key, and returns the decryption using the faulty secret key. This notion already provides a basic level of tamper resilience for public key encryption schemes. Consider for instance a setting where the adversary can tamper with the decryption key, but has no control over the ciphertexts that are sent to the decryption oracle, e.g., the ciphertexts are sent over a secure authenticated channel.

Our notion allows the adversary to tamper adaptively with the secret key; intuitively this allows him to learn faulty decryptions of ciphertexts for which he already knows the corresponding plaintext (under the original public key) and the randomness. We show how to instantiate our basic tamper security notion under DDH. More concretely, we prove that the BHHO cryptosystem [12] is BLT and CPA secure. The proof uses similar ideas as in the proof of the Okamoto identification scheme.

We then show how to transform our extended CPA-like notion to CCA security in the BLT model. To this end, we follow the classical paradigm to transform IND-CPA security into IND-CCA security by adding an argument of "plaintext knowledge" π to the ciphertext. Our transformation requires a public tamper-proof common reference string similar to earlier work [29]. Intuitively, this works because the argument π enforces the adversary to submit to the faulty decryption oracle only ciphertexts for which he knows the corresponding plaintext (and the randomness used to encrypt it). The pairs (m, r) can then be extracted from the argument π, allowing to reduce IND-CCA BLT security to our extended IND-CPA security notion.

Updating the key in the i*Floppy model.* As mentioned earlier, if the key is not updated BLT security is the best we can hope for when we consider arbitrary tampering. To go beyond the bound of $|sk|$ tampering queries we may regularly

update the secret key with fresh randomness, which renders information that the adversary has learned about earlier keys useless. The effectiveness of key updates in the context of tampering attacks has first been used in the important work of Kalai *et al.* [29]. We follow this idea but add an additional hardware assumption that allows for much simpler and more efficient key updates. More concretely, we propose the i*Floppy model* which is a variant of the floppy model proposed by Alwen *et al.* [3] and recently studied in depth by Agrawal *et al.* [2]. In the floppy model a user of a cryptodevice possesses a so-called *floppy* – a secure hardware token – that stores an update key.[5] The floppy is leakage and tamper proof and the update key that it holds is solely used to refresh the actual secret key kept on the cryptodevice. One may think of the floppy as a particularly secure device that the user keeps at home, while the cryptodevice, e.g., a smart-card, runs the actual cryptographic task and is used out in the wild prone to leakage and tampering attacks. We consider a variant called the *i*Floppy model (here "i" stands for individual). While in the floppy model of [2,3] all users can potentially possess an identical hardware token, in the *i*Floppy model we require that each user has an individual floppy storing some secret key related data. We note that from a practical point of view the *i*Floppy model is incomparable to the original floppy model. It may be more cumbersome to produce personalized hardware tokens, but on the other hand, in practice one would not want to distribute hardware tokens that all contain the same global update key as this constitutes a single point of failure.

We show in the *i*Floppy model a simple compiler that "boosts" any ID scheme with BLT security into a scheme with *continuous* leakage and tamper resilience (CLT security). Similarly, we show how to extend IND-CCA BLT security to the CLT setting for the BHHO cryptosystem (borrowing ideas from [2]). We emphasize that while the *i*Floppy model puts additional requirements on the way users must behave in order to guarantee security, it greatly simplifies cryptographic schemes, and allows us to base security on standard assumptions. Our results in the *i*Floppy model are mainly deferred to the full version [17].

Tampering with the computation via the BRM. Finally, we make a simple observation showing that if we instantiate the above ID compiler with an ID scheme that is secure in the bounded retrieval model [15,20,3] we can provide security in the *i*Floppy model even when the adversary can replace the original cryptoscheme with an arbitrary adversarial chosen functionality, i.e., we can allow arbitrary tampering with the computation (see the full version [17]). While easy to prove, we believe this is nevertheless noteworthy: it seems to us that results in the BRM naturally provide some form of tamper resilience and leave it as an open question for future research to explore this direction further.

[5] Notice that "floppy" is just terminology and we use it for consistency with earlier works.

1.2 Previous Work

Related key security. We already discussed the relation between BLT security and the traditional notion of RKA security above. Below we give further details on some important results in the RKA area. Bellare and Kohno [8] initiated the theoretical study of related-key-attacks. Their result mainly focused on symmetric key primitives (e.g. PRP, PRF). They proposed various block-cipher based constructions which are RKA-secure against certain restricted classes of tampering functions. Their constructions were further improved by [32,6]. Following these works other cryptographic primitives were constructed that are provably secure against certain classes of related key attacks. Most of these works consider rather restricted tampering functions that, e.g., can be described by a linear or affine function [8,32,6,5,35,37,10]. A few important exceptions are described below.

In [9] the authors show how to go beyond the linear barrier by extending the class of allowed tampering functions to the class of polynomials of bounded degree for a number of public-key primitives. Also, the work of Goyal, O'Neill and Rao [26] considers polynomial relations that are induced to the inputs of a hash function. Finally Bellare, Cash and Miller [7] develop a framework to transfer RKA security from a pseudorandom function to other primitives (including many public key primitives).

Tamper resilient encodings. A generic method for tamper protection has been put forward by Gennaro *et al.* [25]. The authors propose a general "compiler" that transforms any cryptographic device CS with secret state st, e.g., a block cipher, into a "transformed" cryptoscheme CS' running with state st' that is resilient to arbitrary tampering with st'. In their construction the original state is signed and the signature is checked before each usage. While the above works for any tampering function, it is limited to settings where CS does not change its state as it would need access to the secret signing key to authenticate the new state. This drawback is resolved by the concept of non-malleable codes pioneered by Dziembowski, Pietrzak and Wichs [21]. The original construction of [21] considers an adversary that can tamper independently with bits. This has been extended to small size blocks in [13], and recently to so-called split-state tampering [31,1]. While the above schemes provide surprisingly strong security guarantees, they all require certain assumptions on the hardware (e.g., the memory has to be split into two parts that cannot tampered with jointly), and require significant changes to the implementation for decoding, tamper detection and self-destruct.

Continuous tamper resilience via key updates. Kalai *et al.* [29] provide first feasibility results in the so-called *continuous leakage and tampering* model (CLT). Their constructions achieve strong security requirements where the adversary can arbitrarily tamper continuously with the state. This is achieved by updating the secret key after each usage. While the tampering adversary considered in [29] is clearly stronger (continuous as opposed to bounded tampering), the

proposed schemes are non-standard, rather inefficient and rely on non-standard assumptions. Moreover, the approach of key updates requires a stateful device and large amounts of randomness which is costly in practice. The main focus of this work, are simple standard cryptosystems that neither require randomness for key updates nor need to keep state.

Tampering with computation. In all the above works (including ours) it is assumed that the circuitry that computes the cryptographic algorithm using the potentially tampered key runs correctly and is not subject to tampering attacks. An important line of works analyze to what extend we can guarantee security when the complete circuitry is prone to tampering attacks [28,22,16]. These works typically consider a restricted class of tampering attacks (e.g., individual bit tampering) and assume that large parts of the circuit (and memory) remain un-tampered.

2 Preliminaries

For space reasons, we defer some of the basic definitions to the full version [17].

Basic notation. We review the basic terminology used throughout the paper. For $n \in \mathbb{N}$, we write $[n] := \{1, \ldots, n\}$. Given a set $\mathcal{S}$, we write $s \leftarrow \mathcal{S}$ to denote that element s is sampled uniformly from $\mathcal{S}$. If A is an algorithm, $y \leftarrow \mathsf{A}(x)$ denotes an execution of A with input x and output y; if A is randomized, then y is a random variable. Vectors are denoted in bold. Given a vector $\mathbf{x} = (x_1, \ldots, x_\ell)$ and some integer a, we write $a^{\mathbf{x}}$ for the vector $(a^{x_1}, \ldots, a^{x_\ell})$.

We denote with k the security parameter. A function $\delta(k)$ is called *negligible* in k (or simply negligible) if it vanishes faster than the inverse of any polynomial in k. A machine A is called *probabilistic polynomial time* (PPT) if for any input $x \in \{0,1\}^*$ the computation of $\mathsf{A}(x)$ terminates in at most $poly(|x|)$ steps and A is probabilistic (i.e., it uses randomness as part of its logic). Random variables are usually denoted by capital letters. We sometimes abuse notation and denote a distribution and the corresponding random variable with the same capital letter, say X.

Languages and relations. A *decision problem* related to a language $\mathfrak{L} \subseteq \{0,1\}^*$ requires to determine if a given string y is in $\mathfrak{L}$ or not. We can associate to any $\mathcal{NP}$-language $\mathfrak{L}$ a polynomial-time recognizable relation $\mathfrak{R} \subseteq \{0,1\}^* \times \{0,1\}^*$ defining $\mathfrak{L}$ itself, i.e. $\mathfrak{L} = \{y : \exists x \text{ s.t. } (y,x) \in \mathfrak{R}\}$ for $|x| \leq poly(|y|)$. The string x is called a *witness* for membership of $y \in \mathfrak{L}$.

Information theory. The min-entropy of a random variable X over a set $\mathcal{X}$ is defined as $\mathbf{H}_\infty(X) := -\log \max_x \Pr[X = x]$, and measures how X can be predicted by the best (unbounded) predictor. The conditional average min-entropy [19] of X given a random variable Z (over a set $\mathcal{Z}$) possibly dependent on X, is defined as $\widetilde{\mathbf{H}}_\infty(X|Z) := -\log \mathbb{E}_{z \leftarrow Z}[2^{-\mathbf{H}_\infty(X|Z=z)}]$. Following [3],

we sometimes rephrase the notion of conditional min-entropy in terms of predictors A that are given some information Z (presumably correlated with X), so $\widetilde{\mathbf{H}}_\infty(X|Z) = -\log(\max_{\mathsf{A}} \Pr[\mathsf{A}(Z) = X])$. The above notion of conditional min-entropy can be generalized to the case of interactive predictors A, which participate in some randomized experiment $\mathcal{E}$. An experiment is modeled as interaction between A and a challenger oracle $\mathcal{E}(\cdot)$ which can be randomized, stateful and interactive.

Definition 1 ([3]). *The conditional min-entropy of a random variable X, conditioned on the experiment $\mathcal{E}$ is $\widetilde{\mathbf{H}}_\infty(X|\mathcal{E}) = -\log(\max_{\mathsf{A}} \Pr\left[A^{\mathcal{E}(\cdot)}() = X\right])$. In the special case that $\mathcal{E}$ is a non-interactive experiment which simply outputs a random variable Z, then $\widetilde{\mathbf{H}}_\infty(X|Z)$ can be written to denote $\widetilde{\mathbf{H}}_\infty(X|\mathcal{E})$ abusing the notion.*

We will rely on the following basic properties (see [19, Lemma 2.2]).

Lemma 1. *For all random variables X, Z and Λ over sets $\mathcal{X}$, $\mathcal{Z}$ and $\{0,1\}^\lambda$ such that $\widetilde{\mathbf{H}}_\infty(X|Z) \geq \alpha$, we have*

$$\widetilde{\mathbf{H}}_\infty(X|Z, \Lambda) \geq \widetilde{\mathbf{H}}_\infty(X|Z) - \lambda \geq \alpha - \lambda.$$

3 ID Schemes with BLT Security

In an identification scheme a prover tries to convince a verifier of its identity (corresponding to its public key pk). Formally, an identification scheme is a tuple of algorithms $\mathcal{ID} = (\mathsf{Setup}, \mathsf{Gen}, \mathsf{P}, \mathsf{V})$ defined as follows:

- $pp \leftarrow \mathsf{Setup}(1^k)$: Algorithm Setup takes the security parameter as input and outputs public parameters pp. The set of all public parameter is denoted by $\mathcal{PP}$.
- $(pk, sk) \leftarrow \mathsf{Gen}(1^k)$: Algorithm Gen outputs the public key and the secret key corresponding to the prover's identity. The set of all possible secret keys is denoted by $\mathcal{SK}$.
- (P, V): We let $(\mathsf{P}(pp, sk) \rightleftarrows \mathsf{V}(pp))(pk)$ denote the interaction between prover P (holding sk and using public parameters pp) and verifier V on common input pk. Such interaction outputs a result in $\{accept, reject\}$, where *accept* means P's identity is considered as valid.

Definition 2. *Let $\lambda = \lambda(k)$, $t = t(k)$ and $\delta = \delta(k)$ be parameters and let $\mathcal{T}$ be some set of functions such that $T \in \mathcal{T}$ has a type $T : \mathcal{SK} \times \mathcal{PP} \to \mathcal{SK} \times \mathcal{PP}$. We say that $\mathcal{ID}$ is (λ, t, δ)-bounded leakage and tamper secure (in short BLT-secure) against impersonation attacks with respect to $\mathcal{T}$ if the following properties are satisfied.*

(i) Correctness. *For all $pp \leftarrow \mathsf{Setup}(1^k)$ and $(pk, sk) \leftarrow \mathsf{Gen}(1^k)$ we have that $(\mathsf{P}(pp, sk) \rightleftarrows \mathsf{V}(pp))(pk)$ outputs accept.*

(ii) Security. *For all PPT adversaries A we have that $\Pr[\mathsf{A}$ wins$] \leq \delta(k)$ in the following game:*

1. *The challenger runs* $pp \leftarrow \mathsf{Setup}(1^k)$ *and* $(pk, sk) \leftarrow \mathsf{Gen}(1^k)$*, and gives* (pp, pk) *to* A*.*
2. *The adversary is given oracle access to* $\mathsf{P}(pp, sk)$ *that outputs polynomially many proof transcripts with respect to secret key* sk*.*
3. *The adversary may adaptively ask* t *tampering queries. During the* i*th query,* A *chooses a function* $T_i \in \mathcal{T}$ *and gets oracle access to* $\mathsf{P}(\widetilde{pp}_i, \widetilde{sk}_i)$*, where* $(\widetilde{sk}_i, \widetilde{pp}_i) = T_i(sk, pp)$*. The adversary can interact with the oracle* $\mathsf{P}(\widetilde{pp}_i, \widetilde{sk}_i)$ *a polynomially number of times, where it uses the tampered secret key* $\widetilde{sk}_i$ *and the public parameter* $\widetilde{pp}_i$*.*
4. *The adversary may adaptively ask leakage queries. In the* j*th query,* A *chooses a function* $L_j : \{0,1\}^* \to \{0,1\}^{\lambda_j}$ *and receives back the output of the function applied to* sk*.*
5. *The adversary loses access to all other oracles and interacts with an honest verifier* V *(holding* pk*). We say that* A wins *if* $(\mathsf{A} \rightleftarrows \mathsf{V}(pp))(pk)$ *outputs accept and* $\sum_j \lambda_j \leq \lambda$*.*

Notice that in the above definition the leakage is from the original secret key sk. This is without loss of generality as our tampering functions are modeled as deterministic circuits.

In a slightly more general setting, one could allow A to leak on the original secret key also in the last phase where it has to convince the verifier. In the terminology of [3] this is reminiscent of so-called *anytime leakage* attacks. Our results can be generalized to this setting, however we stick to Definition 2 for simplicity.

The rest of this section is organized as follows. In Section 3.1 we prove that a large class of Σ-protocols are secure in the BLT model, where the tampering function is allowed to modify the secret state of the prover but not the public parameters. In Section 3.2 we look at a concrete instantiation based on the Okamoto ID scheme, and prove that this construction is secure in a stronger model where the tampering function can modify both the secret state of the prover and the public parameters (but independently). Finally, in Section 3.3 we illustrate that the latter assumption is necessary, as otherwise the Okamoto ID scheme can be broken by (albeit contrived) attacks.

3.1 Σ-protocols are Tamper Resilient

We start by considering ID schemes based on Σ-protocols [14]. Σ-protocols are a special class of interactive proof systems for membership in a language $\mathfrak{L}$, where a prover $\mathsf{P} = (\mathsf{P}_0, \mathsf{P}_1)$ wants to convince a verifier $\mathsf{V} = (\mathsf{V}_0, \mathsf{V}_1)$ (both modelled as PPT algorithms) that a shared string y belongs to $\mathfrak{L}$. Denote with x the witness corresponding to y and let pp be public parameters. The protocol proceeds as follows: (1) The prover computes $a \leftarrow \mathsf{P}_0(pp)$ and sends it to the verifier; (2) The verifier chooses $c \leftarrow \mathsf{V}_0(pp, y)$ uniformly at random and sends it to the prover; (3) The prover answers with $z \leftarrow \mathsf{P}_1(pp, (a, c, x))$; (4) The verifier outputs a result $\mathsf{V}_1(pp, y, (a, c, z)) \in \{accept, reject\}$. We call this a *public-coin*

ID Scheme from Σ-Protocol

Let $((\mathsf{P}_0, \mathsf{P}_1), (\mathsf{V}_0, \mathsf{V}_1))$ be a Σ-protocol for a relation $\mathfrak{R}$.

$\mathsf{Setup}(1^k)$**:** Sample public parameters $pp \leftarrow \mathcal{PP}$ for the underlying relation $\mathfrak{R}$.

$\mathsf{Gen}(1^k)$**:** Output a pair $(y, x) \in \mathfrak{R}$, where $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ and $|x|$ is polynomially bounded by $|y|$.

$(\mathsf{P}(pp, x) \rightleftarrows \mathsf{V}(pp))(y)$**:** The protocol works as follows.

1. The prover sends $a \leftarrow \mathsf{P}_0(pp)$ to the verifier.
2. The verifier chooses a random challenge $c \leftarrow \mathsf{V}_0(pp, y)$ and sends it to the prover.
3. The prover computes the answer $z \leftarrow \mathsf{P}_1(pp, (a, c, x))$.
4. The verifier accepts iff $\mathsf{V}_1(pp, y, (a, c, z))$ outputs *accept*.

Fig. 1. ID scheme based on Σ-protocol for relation $\mathfrak{R}$

three round interactive proof system. A formal definition of Σ-protocols can be found in the full version [17].

It is well known that Σ-protocols are a natural tool to design ID schemes. The construction is depicted in Figure 1. Consider now the class of tampering functions $\mathcal{T}_{\mathsf{sk}} \subset \mathcal{T}$ such that $T \in \mathcal{T}_{\mathsf{sk}}$ has the following form: $T = (T^{sk}, ID^{pp})$ where $T^{sk} : \mathcal{SK} \to \mathcal{SK}$ is an arbitrary polynomial time computable function and $ID^{pp} : \mathcal{PP} \to \mathcal{PP}$ is the identity function. This models tampering with the secret state of P without changing the public parameters (these must be hard-wired into the prover's code). The proof of the following theorem (which appears in the full version [17]) uses ideas of [3], but is carefully adjusted to incorporate tampering attacks.

Theorem 1. *Let $k \in \mathbb{N}$ be the security parameter and let (P, V) be a Σ-protocol for relation $\mathfrak{R}$ with $|\mathcal{Y}| = O(k^{\log k})$, such that the representation problem is hard for $\mathfrak{R}$. Assume that conditioned on the distribution of the public input $y \in \mathcal{Y}$, the witness $x \in \mathcal{X}$ has high average min entropy β, i.e., $\widetilde{\mathbf{H}}_\infty(X|Y) \geq \beta$. Then, the ID scheme of Figure 1 is $(\lambda(k), t(k), negl(k))$-BLT secure against impersonation attacks with respect to $\mathcal{T}_{\mathsf{sk}}$, where*

$$\lambda \leq \beta - t \log |\mathcal{Y}| - k \qquad \textit{and} \qquad t \leq \left\lfloor \frac{\beta}{\log |\mathcal{Y}|} \right\rfloor - 1.$$

3.2 Concrete Instantiation with more Tampering

We extend the power of the adversary by allowing him to tamper not only with the witness, but also with the public parameters (used by the prover to generate the transcripts). However the tampering has to be independent on the two components. This is reminiscent of the so-called split-state model (considered

Generalized Okamoto ID Scheme

Let $\ell = \ell(k)$ be some function of the security parameter. Consider the following identification scheme.

Setup: Choose a group $\mathbb{G}$ of prime order p with generator g and a vector $\boldsymbol{\alpha} \leftarrow \mathbb{Z}_p^\ell$, and output $pp = (\mathbb{G}, g, g^{\boldsymbol{\alpha}})$ where $g^{\boldsymbol{\alpha}} = (g_1, \ldots, g_\ell)$.

$\mathsf{Gen}(1^k)$: Select a vector $\mathbf{x} \leftarrow \mathbb{Z}_p^\ell$ and set $y = pk = \prod_{i=1}^{\ell} g_i^{x_i}$ and $sk = \mathbf{x}$.

$(\mathsf{P}(pp, sk) \rightleftarrows \mathsf{V}(pp))(pk)$: The protocol works as follows.

1. The prover chooses a random vector $\mathbf{r} \leftarrow \mathbb{Z}_p^\ell$ and sends $a = \prod_{i=1}^{\ell} g_i^{r_i}$ to the verifier.
2. The verifier chooses a random challenge $c \leftarrow \mathbb{Z}_p$ and sends it to the prover.
3. The prover computes the answer $\mathbf{z} = (r_1 + cx_1, \ldots, r_\ell + cx_\ell)$.
4. The verifier accepts if and only if $\prod_{i=1}^{\ell} g_i^{z_i} = a \cdot y^c$.

Fig. 2. Generalized Okamoto Identification Scheme

for instance in [31]), with the key difference that in our case the secret state does not need to be split into two parts.

We model this through the following class of tampering functions $\mathcal{T}_{\mathsf{split}}$: We say that $T \in \mathcal{T}_{\mathsf{split}}$ if we can write $T = (T^{sk}, T^{pp})$ where $T^{sk} : \mathcal{SK} \to \mathcal{SK}$ and $T^{pp} : \mathcal{PP} \to \mathcal{PP}$ are arbitrary polynomial time computable functions. Recall that the input/output domains of T^{sk}, T^{pp} are identical, hence the size of the witness and the public parameters cannot be changed. As we show in the next section, this restriction is necessary. Note also that $\mathcal{T}_{\mathsf{sk}} \subseteq \mathcal{T}_{\mathsf{split}} \subseteq \mathcal{T}$.

Generalized Okamoto. Consider the generalized version of the Okamoto ID scheme [34], depicted in Figure 2. The underlying hard relation here is the relation $\mathfrak{R}_{\mathsf{DL}}$ and the representation problem for $\mathfrak{R}_{\mathsf{DL}}$ is the ℓ-representation problem in a group $\mathbb{G}$. As proven in [3], this problem is equivalent to the Discrete Log problem in $\mathbb{G}$. The proof of the following corollary appears in the full version [17].

Corollary 1. *Let $k \in \mathbb{N}$ be the security parameter and assume the Discrete Log problem is hard in $\mathbb{G}$. Then, the generalized Okamoto ID scheme is $(\lambda(k), t(k), negl(k))$-BLT secure against impersonation attacks with respect to $\mathcal{T}_{\mathsf{split}}$, where*

$$\lambda \leq (\ell - 1 - t)\log(p) - k \qquad \text{and} \qquad t \leq \ell - 2.$$

3.3 Some Attacks

We show that for the Okamoto scheme it is hard to hope for BLT security beyond the class of tampering functions $\mathcal{T}_{\mathsf{split}}$. We illustrate this by concrete attacks which work in case one tries to extend the power of the adversary in two different ways: (1) Allowing A to tamper jointly with the witness and the public parameters; (2) Allowing A to tamper independently with the witness and with the public parameters but increase their size.

Tampering jointly with the public parameters. Consider the class of functions $\mathcal{T}$ introduced in Definition 2.

Claim. The generalized Okamoto ID scheme is *not* BLT-secure against impersonation attacks with respect to $\mathcal{T}$.

Proof. The attack uses a single tampering query. Define the tampering function $T(\mathbf{x}, pp) = (\widetilde{\mathbf{x}}, \widetilde{pp})$ to be as follows:

- The witness is unchanged, i.e., $\mathbf{x} = \widetilde{\mathbf{x}}$.
- The value $\widetilde{p}$ is some prime of size $|\widetilde{p}| \approx |p|$ such that the Discrete Log problem is easy in the corresponding group $\widetilde{\mathbb{G}}$. (This can be done efficiently by choosing $\widetilde{p} - 1$ to be the product of small prime (power) factors [36].)
- Let $\widetilde{g}$ be a generator of $\widetilde{\mathbb{G}}$ (which exists since $\widetilde{p}$ is a prime) and define the new generators as $\widetilde{g}_i = \widetilde{g}^{x_i} \bmod \widetilde{p}$.

Consider now a transcript $(a, c, \mathbf{z})$ produced by a run of $\mathsf{P}(\widetilde{pp}, \mathbf{x})$. We have $a = \widetilde{g}^{\sum_{i=1}^{\ell} x_i r_i} \bmod \widetilde{p}$ for random $r_i \in \mathbb{Z}_{\widetilde{p}}$. By computing the Discrete Log of a in base $\widetilde{g}$ (which is easy by our choice of $\widetilde{\mathbb{G}}$), we get one equation $\sum_{i=1}^{\ell} x_i r_i = \log_{\widetilde{g}}(a) \bmod \widetilde{p}$. Asking for polynomially many transcripts, yields ℓ linearly independent equations (with overwhelming probability) and thus allows to solve for $(x_1, \ldots, x_\ell)$. (Note here that with high probability $x_i \bmod p = x_i \bmod \widetilde{p}$ since $|p| \approx |\widetilde{p}|$.)

Tampering by "inflating" the prime p. Consider the following class of tampering functions $\mathcal{T}_{\mathsf{split}} \subseteq \mathcal{T}^*_{\mathsf{split}}$: We say that $T \in \mathcal{T}^*_{\mathsf{split}}$ if $T = (T^{sk}, T^{pp})$, where $T^{sk} : \mathcal{SK} \to \{0,1\}^*$ and $T^{pp} : \mathcal{PP} \to \{0,1\}^*$.

Claim. The generalized Okamoto ID scheme is *not* BLT-secure against impersonation attacks with respect to $\mathcal{T}^*_{\mathsf{split}}$.

Proof. The attack uses a single tampering query. Consider the following tampering function $T = (T^{sk}, T^{pp}) \in \mathcal{T}^*_{\mathsf{split}}$:

- Choose $\widetilde{p}$ to be a prime of size $|\widetilde{p}| = \Omega(\ell|p|)$, such that the Discrete Log problem is easy in $\widetilde{\mathbb{G}}$. (This can be done as in the proof of Claim 3.3.)
- Choose a generator $\widetilde{g}$ of $\widetilde{\mathbb{G}}$; define $\widetilde{g}_1 = \widetilde{g}$ and $\widetilde{g}_j = 1$ for all $j = 2, \ldots, \ell$.
- Define the witness to be $\widetilde{\mathbf{x}}$ such that $\widetilde{x}_1 = x_1 || \ldots || x_\ell$ and $\widetilde{x}_j = 0$ for all $j = 2, \ldots, \ell$.

Given a single transcript $(a, c, \mathbf{z})$ the adversary learns $a = \widetilde{g}^{r_1}$ for some $r_1 \in \mathbb{Z}_{\widetilde{p}}$. Since the Discrete Log is easy in this group, A can find r_1. Now the knowledge of c and $z_1 = r_1 + c\widetilde{x}_1$, allows to recover $\widetilde{x}_1 = (x_1, \ldots, x_\ell)$.

3.4 BLT-Secure Signatures

It is well known that every Σ-protocol can be turned into a signature scheme via the Fiat-Shamir heuristic [23]. By applying the Fiat-Shamir transformation to the protocol of Figure 1, we get efficient BLT-secure signatures in the random oracle model.

4 IND-CCA PKE with BLT Security

We start by defining IND-CCA public key encryption (PKE) with BLT security. A PKE scheme is a tuple of algorithms $\mathcal{PKE} = (\mathsf{Setup}, \mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ defined as follows. (1) Algorithm Setup takes as input the security parameter and outputs the description of public parameters pp; the set of all public parameters is denoted by $\mathcal{PP}$. (2) Algorithm KGen takes as input the security parameter and outputs a public/secret key pair (pk, sk); the set of all secret keys is denoted by $\mathcal{SK}$ and the set of all public keys by $\mathcal{PK}$. (3) The randomized algorithm Enc takes as input the public key pk, a message $m \in \mathcal{M}$ and randomness $r \in \mathcal{R}$ and outputs a ciphertext $c = \mathsf{Enc}(pk, m; r)$; the set of all ciphertexts is denoted by $\mathcal{C}$. (4) The deterministic algorithm Dec takes as input the secret key sk and a ciphertext $c \in \mathcal{C}$ and outputs $m = \mathsf{Dec}(sk, c)$ which is either equal to some message $m \in \mathcal{M}$ or to an error symbol $\perp$.

Definition 3. *Let $\lambda = \lambda(k)$, $t = t(k)$ and $\delta = \delta(k)$ be parameters and let $\mathcal{T}_{\mathsf{sk}}$ be some set of functions such that $T \in \mathcal{T}_{\mathsf{sk}}$ has a type $T : \mathcal{SK} \to \mathcal{SK}$. We say that $\mathcal{PKE}$ is IND-CCA $(\lambda(k), t(k), \delta(k))$-BLT secure with respect to $\mathcal{T}_{\mathsf{sk}}$ if the following properties are satisfied.*

(i) Correctness. *For all $pp \leftarrow \mathsf{Setup}(1^k)$, $(pk, sk) \leftarrow \mathsf{KGen}(1^k)$ we have that $\Pr[\mathsf{Dec}(sk, \mathsf{Enc}(pk, m)) = m] = 1$ (where the randomness is taken over the internal coin tosses of algorithm Enc).*

(ii) Security. *For all PPT adversaries A we have that $\Pr[\mathsf{A} \text{ wins}] \leq \frac{1}{2} + \delta(k)$ in the following game:*

1. *The challenger runs $pp \leftarrow \mathsf{Setup}(1^k)$, $(pk, sk) \leftarrow \mathsf{KGen}(1^k)$ and gives (pp, pk) to A.*
2. *The adversary is given oracle access to $\mathsf{Dec}(sk, \cdot)$. This oracle outputs polynomially many decryptions of ciphertexts using secret key sk.*
3. *The adversary may adaptively ask t tampering queries. During the ith query, A chooses a function $T_i \in \mathcal{T}_{\mathsf{sk}}$ and gets oracle access to $\mathsf{Dec}(\widetilde{sk}_i, \cdot)$, where $\widetilde{sk}_i = T_i(sk)$. This oracle outputs polynomially many decryptions of ciphertexts using secret key $\widetilde{sk}_i$.*
4. *The adversary may adaptively ask polynomially many leakage queries. In the jth query, A chooses a function $L_j : \{0,1\}^* \to \{0,1\}^{\lambda_j}$ and receives back the output of the function applied to sk.*
5. *The adversary outputs two messages of the same length $m_0, m_1 \in \mathcal{M}$ and the challenger computes $c_b \leftarrow \mathsf{Enc}(pk, m_b)$ where b is a uniformly random bit.*
6. *The adversary keeps access to $\mathsf{Dec}(sk, \cdot)$ and outputs a bit b'. We say A wins if $b = b'$, $\sum_j \lambda_j \leq \lambda$ and c_b has not been queried for.*

In case $t = 0$ we get the notion of leakage resilient IND-CCA from [33] as a special case. Notice that A is not allowed to tamper with the secret key after seeing the challenge ciphertext. As we show in the full version [17], this restriction is necessary because otherwise A could overwrite the secret key depending on the

plaintext encrypted in c_b, and thus gain some advantage in guessing the value of b by asking additional decryption queries.

We build an IND-CCA BLT-secure PKE scheme in two steps. In Section 4.1 we define a weaker notion which we call IND-CPA BLT security. In Section 4.2 we show a general transformation from IND-CPA BLT security to IND-CCA BLT security relying on tSE NIZK proofs [18] in the common reference string (CRS) model. The CRS is supposed to be tamper-free and must be hard-wired into the code of the encryption algorithm; however tampering and leakage can depend adaptively on the CRS and the public parameters. Finally, in Section 4.3, we prove that a variant of the BHHO encryption scheme [33] satisfies our notion of IND-CPA BLT security.

4.1 IND-CPA BLT Security

The main idea of our new security notion is as follows. Instead of giving A full access to a tampering oracle (as in Definition 3) we restrict his power by allowing him to see the output of the (tampered) decryption oracle only for ciphertexts c for which A already knows both the corresponding plaintext m and the randomness r used to generate c (via the real public key). Essentially this restricts A to submit to the tampering oracle only "well-formed" ciphertexts.

Definition 4. *Let $\lambda = \lambda(k)$, $t = t(k)$ and $\delta = \delta(k)$ be parameters and let $\mathcal{T}_{\mathsf{sk}}$ be some set of functions such that $T \in \mathcal{T}_{\mathsf{sk}}$ has a type $T : \mathcal{SK} \to \mathcal{SK}$. We say that $\mathcal{PKE}$ is IND-CPA $(\lambda(k), t(k), \delta(k))$-BLT secure with respect to $\mathcal{T}_{\mathsf{sk}}$ if it satisfies property (i) of Definition 3 and property (ii) is modified as follows:*

(ii) Security. *For all PPT adversaries A we have that $\Pr[\mathsf{A}\ \textit{wins}] \leq \frac{1}{2} + \delta(k)$ in the following game:*

1. *The challenger runs $pp \leftarrow \mathsf{Setup}(1^k)$, $(pk, sk) \leftarrow \mathsf{KGen}(1^k)$ and gives (pp, pk) to A.*
2. *The adversary may adaptively ask t tampering queries. During the ith query, A chooses a function $T_i \in \mathcal{T}_{\mathsf{sk}}$ and gets oracle access to $\mathsf{Dec}^*(\widetilde{sk}_i, \cdot, \cdot)$, where $\widetilde{sk}_i = T_i(sk)$. This oracle answers polynomially many queries of the following form: Upon input a pair $(m, r) \in \mathcal{M} \times \mathcal{R}$, compute $c \leftarrow \mathsf{Enc}(pk, m; r)$ and output a plaintext $\widetilde{m} = \mathsf{Dec}(\widetilde{sk}_i, c)$ using the current tampered key.*
3. *The adversary may adaptively ask leakage queries. In the jth query, A chooses a function $L_j : \{0,1\}^* \to \{0,1\}^{\lambda_j}$ and receives back the output of the function applied to sk.*
4. *The adversary outputs two messages of the same length $m_0, m_1 \in \mathcal{M}$ and the challenger computes $c_b \leftarrow \mathsf{Enc}(pk, m_b)$ where b is a uniformly random bit.*
5. *The adversary loses access to all oracles and outputs a bit b'. We say that A wins if $b = b'$ and $\sum_j \lambda_j \leq \lambda$.*

From IND-CPA BLT Security to IND-CCA BLT Security

Let $\mathcal{PKE} = (\mathsf{Setup}, \mathsf{KGen}, \mathsf{Enc}, \mathsf{Dec})$ be a PKE scheme and $(\mathsf{Gen}, \mathsf{Prove}, \mathsf{Verify})$ be a tSE NIZK argument system for the relation:

$$\mathfrak{R}_{\mathsf{PKE}} = \{(pk, c), (m, r) : \ c = \mathsf{Enc}(pk, m; r)\}\,.$$

Define the following PKE scheme $\mathcal{PKE}' = (\mathsf{Setup}', \mathsf{KGen}', \mathsf{Enc}', \mathsf{Dec}')$.

Setup': Sample $pp \leftarrow \mathsf{Setup}(1^k)$ and $(\omega, \mathsf{tk}, \mathsf{ek}) \leftarrow \mathsf{Gen}(1^k)$ and let $pp' = (pp, \omega)$.
KGen': Run $(pk, sk) \leftarrow \mathsf{KGen}(1^k)$ and set $pk' = pk$ and $sk' = sk$.
Enc': Sample $r \leftarrow \mathcal{R}$ and compute $c \leftarrow \mathsf{Enc}(pk, m; r)$. Output (c, π), where $\pi \leftarrow \mathsf{Prove}^{\omega}((pk, c), (m, r))$.
Dec': Check that $\mathsf{Verify}^{\omega}((pk, c), \pi) = 1$. If not output $\perp$; otherwise, output $m = \mathsf{Dec}(sk, c)$.

Fig. 3. How to transform IND-CPA BLT-secure PKE into IND-CCA BLT-secure PKE

4.2 A General Transformation

We compile an arbitrary IND-CPA BLT-secure encryption scheme into an IND-CCA BLT-secure one by appending to the ciphertext c an argument of "plaintext knowledge" π computed through a tSE NIZK argument system. The same construction has been already used by Dodis *et al.* [18] to go from IND-CPA security to IND-CCA security in the context of memory leakage.

The intuition why the transformation works is fairly simple: The argument π enforces the adversary to submit to the tampered decryption oracle only ciphertexts for which he knows the corresponding plaintext (and the randomness used to encrypt it). In the security proof the pair (m, r) can indeed be extracted from such argument, allowing to reduce IND-CCA BLT security to IND-CPA BLT security.

Theorem 2. *Let $k \in \mathbb{N}$ be the security parameter. Assume that $\mathcal{PKE}$ is an IND-CPA $(\lambda(k), t(k), \delta(k))$-BLT secure encryption scheme and that $(\mathsf{Gen}, \mathsf{Prove}, \mathsf{Verify})$ is a strong tSE NIZK argument system for relation $\mathfrak{R}_{\mathsf{PKE}}$. Then, the encryption scheme $\mathcal{PKE}'$ of Figure 3 is IND-CCA $(\lambda(k), t(k), \delta'(k))$-BLT secure for $\delta' \leq \delta + negl(k)$.*

Proof. We prove the theorem by a series of games. All games are a variant of the IND-CCA BLT game and in all games the adversary gets correctly generated public parameters (pp, ω, pk). Leakage and tampering queries are answered using the corresponding secret key sk. The games will differ only in the way the challenge ciphertext is computed or in the way the decryption oracles work.

Game G_1. This is the IND-CCA BLT game of Definition 3 for the scheme $\mathcal{PKE}'$. Note in particular that all decryption oracles expect to receive as

input a ciphertext of the form (c, π) and proceed to verify the proof π before decrypting the ciphertext (and output $\perp$ if such verification fails). The challenge ciphertext is a pair (c_b, π_b) such that $c_b = \mathsf{Enc}(pk, m_b; r)$ and $\pi_b \leftarrow \mathsf{Prove}^{\omega}((pk, c_b), (m_b, r))$, where $m_b \in \{m_0, m_1\}$ for a uniformly random bit b. By assumption we have that

$$\Pr[\mathsf{A} \text{ wins in } \mathsf{G}_1] \leq \frac{1}{2} + \delta'(k).$$

Game G_2. In this game we change the way the challenge ciphertext is computed by replacing the argument π_b with a simulated argument $\pi_b \leftarrow \mathsf{S}((pk, c_b), \mathsf{tk})$. It follows from the composable NIZK property of the argument system that G_1 and G_2 are computationally close. In particular there exists a negligible function $\delta_1(k)$ such that $|\Pr[\mathsf{A} \text{ wins in } \mathsf{G}_1] - \Pr[\mathsf{A} \text{ wins in } \mathsf{G}_2]| \leq \delta_1(k)$.

Game G_3. We change the way decryption queries are handled. Queries (c, π) to $\mathsf{Dec}(sk, \cdot)$ (such that π accepts) are answered by running the extractor Ext on π, yielding $(m, r) \leftarrow \mathsf{Ext}((pk, c), \pi, \mathsf{ek})$, and returning m.

Queries (c, π) to $\mathsf{Dec}(\widetilde{sk}_i, \cdot)$ (such that π accepts) are answered as follows. We first extract $(m, r) \leftarrow \mathsf{Ext}((pk, c), \pi, \mathsf{ek})$ as above. Then, instead of returning m, we recompute $c = \mathsf{Enc}(pk, m; r)$ and return $\widetilde{m} = \mathsf{Dec}(\widetilde{sk}_i, c)$.

It follows from true simulation extractability that G_2 and G_3 are computationally close. The reason for this is that A gets to see only a single simulated proof for a true statement (i.e., the pair (pk, c_b)) and thus cannot produce a pair $(c, \pi) \neq (c_b, \pi_b)$ such that the proof π accepts and Ext fails to extract the corresponding plaintext m. In particular there exists a negligible function $\delta_2(k)$ such that $|\Pr[\mathsf{A} \text{ wins in } \mathsf{G}_2] - \Pr[\mathsf{A} \text{ wins in } \mathsf{G}_3]| \leq \delta_2(k)$.

Game G_4. In the last game we replace the ciphertext c_b in the challenge with an encryption of $0^{|m_b|}$, whereas we still compute the proof as $\pi_b \leftarrow \mathsf{S}((pk, c_b), \mathsf{tk})$. We claim that G_3 and G_4 are computationally close. This follows from IND-CPA BLT-security of $\mathcal{PKE}$. Assume there exists a distinguisher D between G_3 and G_4. We build an adversary B breaking IND-CPA BLT security for $\mathcal{PKE}$. The adversary B uses D as a black-box as follows.

Reduction B^{D}:

1. Receive (pp, pk) from the challenger, sample $(\omega, \mathsf{tk}, \mathsf{ek}) \leftarrow \mathsf{Gen}(1^k)$ and give $pp' = (pp, \omega)$ and $pk' = pk$ to A.
2. Upon input a normal decryption query (c, π) from A, run the extractor to compute $(m, r) \leftarrow \mathsf{Ext}((pk, c), \pi, \mathsf{ek})$ and return m.
3. Upon input a tampering query $T_i \in \mathcal{T}_{\mathsf{sk}}$, forward T_i to the tampering oracle for $\mathcal{PKE}$. To answer a query (c, π), run the extractor to compute $(m, r) \leftarrow \mathsf{Ext}((pk, c), \pi, \mathsf{ek})$. Submit (m, r) to oracle $\mathsf{Dec}^*(\widetilde{sk}_i, \cdot, \cdot)$ and receive the answer $\widetilde{m}$. Return $\widetilde{m}$ to A.
4. Upon input a leakage query L_j, forward L_j to the leakage oracle for $\mathcal{PKE}$.
5. When A outputs $m_0, m_1 \in \mathcal{M}$, sample a random bit b' and output $(m_{b'}, 0^{|m_{b'}|})$. Let c_b be the corresponding challenge ciphertext. Compute $\pi_b \leftarrow \mathsf{S}((pk, c_b), \mathsf{tk})$ and forward (c_b, π_b) to A. Continue to answer normal decryption queries (c, π) from A as above.

6. Output whatever D does.

Notice that the reduction perfectly simulates the environment for A; in particular c_b is either the encryption of randomly chosen message among (m_0, m_1) (as in G_3) or an encryption of zero (as in G_4). Since $\mathcal{PKE}$ is (λ, t, δ)-BLT secure, it must be $|\Pr[\mathsf{A} \text{ wins in } \mathsf{G}_3] - \Pr[\mathsf{A} \text{ wins in } \mathsf{G}_4]| \leq \delta(k)$.

As clearly $\Pr[\mathsf{A} \text{ wins in } \mathsf{G}_4] = 0$, we have obtained

$$\begin{aligned}
\delta' &= |\Pr[\mathsf{A} \text{ wins in } \mathsf{G}_1] - \Pr[\mathsf{A} \text{ wins in } \mathsf{G}_4]| \\
&\leq |\Pr[\mathsf{A} \text{ wins in } \mathsf{G}_1] - \Pr[\mathsf{A} \text{ wins in } \mathsf{G}_2]| + |\Pr[\mathsf{A} \text{ wins in } \mathsf{G}_2] \\
&\qquad - \Pr[\mathsf{A} \text{ wins in } \mathsf{G}_3]| + |\Pr[\mathsf{A} \text{ wins in } \mathsf{G}_3] - \Pr[\mathsf{A} \text{ wins in } \mathsf{G}_4]| \\
&\leq \delta_1(k) + \delta_2(k) + \delta(k) = \delta(k) + negl(k).
\end{aligned}$$

This concludes the proof.

4.3 Instantiation from BHHO

We show that the variant of the encryption scheme introduced by Boneh *et al.* [12] used in [33] is IND-CPA BLT-secure. The proof relies on the observation that one can simulate polynomially many decryption queries for a given tampered key by only leaking a bounded amount of information from the secret key. Hence, security follows from leakage resilience. The formal description of the scheme and the proof can be found in the full version [17].

5 Updating the Key in the *i*Floppy Model

We complement the results from the previous two sections by showing how to obtain security against an unbounded number of tampering queries in the floppy model of [3,2]. Recall that in this model we assume the existence of an external tamper-free and leakage-free storage (the floppy), which is needed to refresh the secret key on the tamperable device. An important difference between the floppy model considered in this paper and the model of [2] is that in our case the floppy can contain "user-specific" information, whereas in [2] it contains a *unique* master key which in principle could be equal for all users. To stress this difference, we refer to our model as the *i*Floppy model.

Clearly, the assumption of a unique master key makes production easier but it is also a single point of failure in the system since in case the content of the floppy is published (e.g., by a malicious user) the entire system needs to be re-initialized.[6] A solution for this is to assume that each floppy contains a different master key as is the case in the *i*Floppy model, resulting in a trade-off between security and production cost. Due to space restrictions, we defer the formal definitions and proofs to the full version [17].

[6] We note that in the schemes of [2] making the content of the floppy public does not constitute a total breach of security; however the security proof completely breaks down, leaving no security guarantee for the schemes at hand.

Acknowledgments

Ivan Damgård and Daniele Venturi acknowledge support from the Danish National Research Foundation, the National Science Foundation of China (under the grant 61061130540), the Danish Council for Independent Research (under the DFF Starting Grant 10-081612) and also from the CFEM research center within which part of this work was performed. Sebastian Faust was partially funded by the above grants. Pratyay Mukherjee's work at Aarhus University was supported by the above grants and a European Research Commission Starting Grant (no. 279447). Part of this work was done while this author was at the University of Warsaw and was supported by the WELCOME/2010-4/2 grant founded within the framework of the EU Innovative Economy Operational Programme.

References

1. Divesh Aggarwal, Yevgeniy Dodis, and Shachar Lovett. Non-malleable codes from additive combinatorics. *IACR Cryptology ePrint Archive*, 2013:201, 2013.
2. Shweta Agrawal, Yevgeniy Dodis, Vinod Vaikuntanathan, and Daniel Wichs. On continual leakage of discrete log representations. *IACR Cryptology ePrint Archive*, 2012:367, 2012.
3. Joël Alwen, Yevgeniy Dodis, and Daniel Wichs. Leakage-resilient public-key cryptography in the bounded-retrieval model. In *CRYPTO*, pages 36–54, 2009.
4. Ross Anderson and Markus Kuhn. Tamper resistance: a cautionary note. In *WOEC'96: Proceedings of the 2nd conference on Proceedings of the Second USENIX Workshop on Electronic Commerce*, pages 1–1, Berkeley, CA, USA, 1996. USENIX Association.
5. Benny Applebaum, Danny Harnik, and Yuval Ishai. Semantic security under related-key attacks and applications. In *ICS*, pages 45–60, 2011.
6. Mihir Bellare and David Cash. Pseudorandom functions and permutations provably secure against related-key attacks. In *CRYPTO*, pages 666–684, 2010.
7. Mihir Bellare, David Cash, and Rachel Miller. Cryptography secure against related-key attacks and tampering. In *ASIACRYPT*, pages 486–503, 2011.
8. Mihir Bellare and Tadayoshi Kohno. A theoretical treatment of related-key attacks: RKA-PRPs, RKA-PRFs, and applications. In *EUROCRYPT*, pages 491–506, 2003.
9. Mihir Bellare, Kenneth G. Paterson, and Susan Thomson. RKA security beyond the linear barrier: IBE, encryption and signatures. In *ASIACRYPT*, pages 331–348, 2012.
10. Rishiraj Bhattacharyya and Arnab Roy. Secure message authentication against related key attack. In *FSE*, 2013.
11. Dan Boneh, Richard A. DeMillo, and Richard J. Lipton. On the importance of eliminating errors in cryptographic computations. *J. Cryptology*, 14(2):101–119, 2001.
12. Dan Boneh, Shai Halevi, Michael Hamburg, and Rafail Ostrovsky. Circular-secure encryption from decision diffie-hellman. In *CRYPTO*, pages 108–125, 2008.
13. Seung Geol Choi, Aggelos Kiayias, and Tal Malkin. Bitr: Built-in tamper resilience. In *ASIACRYPT*, pages 740–758, 2011.
14. Ronald Cramer. *Modular Design of Secure yet Practical Cryptographic Protocols*. PhD thesis, University of Amsterdam, November 1996.

15. Giovanni Di Crescenzo, Richard J. Lipton, and Shabsi Walfish. Perfectly secure password protocols in the bounded retrieval model. In *TCC*, pages 225–244, 2006.
16. Dana Dachman-Soled and Yael Tauman Kalai. Securing circuits against constant-rate tampering. In *CRYPTO*, pages 533–551, 2012.
17. Ivan Damgård, Sebastian Faust, Pratyay Mukherjee, and Daniele Venturi. Bounded tamper resilience: How to go beyond the algebraic barrier. *IACR Cryptology ePrint Archive*, 2013.
18. Yevgeniy Dodis, Kristiyan Haralambiev, Adriana López-Alt, and Daniel Wichs. Cryptography against continuous memory attacks. In *FOCS*, pages 511–520, 2010.
19. Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.*, 38(1):97–139, 2008.
20. Stefan Dziembowski. Intrusion-resilience via the bounded-storage model. In *TCC*, pages 207–224, 2006.
21. Stefan Dziembowski, Krzysztof Pietrzak, and Daniel Wichs. Non-malleable codes. In *ICS*, pages 434–452, 2010.
22. Sebastian Faust, Krzysztof Pietrzak, and Daniele Venturi. Tamper-proof circuits: How to trade leakage for tamper-resilience. In *ICALP (1)*, pages 391–402, 2011.
23. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO*, pages 186–194, 1986.
24. Marc Fischlin and Roger Fischlin. The representation problem based on factoring. In *CT-RSA*, pages 96–113, 2002.
25. Rosario Gennaro, Anna Lysyanskaya, Tal Malkin, Silvio Micali, and Tal Rabin. Algorithmic tamper-proof (atp) security: Theoretical foundations for security against hardware tampering. In *TCC*, pages 258–277, 2004.
26. Vipul Goyal, Adam O'Neill, and Vanishree Rao. Correlated-input secure hash functions. In *TCC*, pages 182–200, 2011.
27. Louis C. Guillou and Jean-Jacques Quisquater. A "paradoxical" indentity-based signature scheme resulting from zero-knowledge. In *CRYPTO*, pages 216–231, 1988.
28. Yuval Ishai, Manoj Prabhakaran, Amit Sahai, and David Wagner. Private circuits II: Keeping secrets in tamperable circuits. In *EUROCRYPT*, pages 308–327, 2006.
29. Yael Tauman Kalai, Bhavana Kanukurthi, and Amit Sahai. Cryptography with tamperable and leaky memory. In *CRYPTO*, pages 373–390, 2011.
30. Jonathan Katz and Vinod Vaikuntanathan. Signature schemes with bounded leakage resilience. In *ASIACRYPT*, pages 703–720, 2009.
31. Feng-Hao Liu and Anna Lysyanskaya. Tamper and leakage resilience in the split-state model. In *CRYPTO*, pages 517–532, 2012.
32. Stefan Lucks. Ciphers secure against related-key attacks. In *FSE*, pages 359–370, 2004.
33. Moni Naor and Gil Segev. Public-key cryptosystems resilient to key leakage. In *CRYPTO*, pages 18–35, 2009.
34. Tatsuaki Okamoto. Provably secure and practical identification schemes and corresponding signature schemes. In *CRYPTO*, pages 31–53, 1992.
35. Krzysztof Pietrzak. Subspace LWE. In *TCC*, pages 548–563, 2012.
36. Stephen Pohlig and Martin Hellman. An improved algorithm for computing logarithms over and its cryptographic significance. *IEEE Transactions on Information Theory*, 24(1):106–110, 1978.
37. Hoeteck Wee. Public key encryption against related key attacks. In *Public Key Cryptography*, pages 262–279, 2012.