

Leakage-Resilient Cryptography

Stefan Dziembowski*
University of Rome
La Sapienza

Krzysztof Pietrzak
CWI Amsterdam

Abstract

We construct a stream-cipher S whose implementation is secure even if a bounded amount of arbitrary (adversarially chosen) information on the internal state of S is leaked during computation. This captures all possible side-channel attacks on S where the amount of information leaked in a given period is bounded, but overall can be arbitrary large. The only other assumption we make on the implementation of S is that only data that is accessed during computation leaks information.

The stream-cipher S generates its output in chunks K_1, K_2, \dots and arbitrary but bounded information leakage is modeled by allowing the adversary to adaptively choose a function $f_\ell : \{0, 1\}^* \rightarrow \{0, 1\}^\lambda$ before K_ℓ is computed, she then gets $f_\ell(\tau_\ell)$ where τ_ℓ is the internal state of S that is accessed during the computation of K_ℓ . One notion of security we prove for S is that K_ℓ is indistinguishable from random when given $K_1, \dots, K_{\ell-1}, f_1(\tau_1), \dots, f_{\ell-1}(\tau_{\ell-1})$ and also the complete internal state of S after K_ℓ has been computed (i.e. S is forward-secure).

The construction is based on alternating extraction (used in the intrusion-resilient secret-sharing scheme from FOCS'07). We move this concept to the computational setting by proving a lemma that states that the output of any PRG has high HILL pseudoentropy (i.e. is indistinguishable from some distribution with high min-entropy) even if arbitrary information about the seed is leaked. The amount of leakage λ that we can tolerate in each step depends on the strength of the underlying PRG, it is at least logarithmic, but can be as large as a constant fraction of the internal state of S if the PRG is exponentially hard.

1. Introduction

When analyzing the security of a cryptosystem, we can either think of the system as a mathematical object, exactly specifying what kind of access to the functionality a potential adversary has, or try to analyze the security of an actual

implementation. Traditionally, cryptographers have mostly considered the former view and analyzed the security of the mathematical object, and it is generally believed that our current knowledge of cryptography suffices to construct schemes that, when modeled in this way, are extremely secure. On a theoretical side, we know how to construct secure primitives under quite weak complexity-theoretic assumptions, for example secret-key encryption can be based on any one-way function [17]. Also from the practical perspective, the currently used constructions have very strong security properties, e.g. after 30 years of intensive cryptanalytic efforts still the most practical attack on the DES cipher is exhaustive key search.

Side-Channel Attacks. The picture is much more gloomy when the security of *real-life implementations* is considered. The reason is that, when considering the security of an implementation of a cryptosystem, one must take into account the possibility of side-channels, which refers to leakage of any kind of information from the cryptosystem during its execution which cannot be efficiently derived from access to the mathematical object alone. In the last decade many attacks against cryptosystems (still assumed to be sound as mathematical objects) have been found exploiting side-channels like running-time [22], electromagnetic radiation [30, 15], power consumption [23], fault detection [4, 3] and many more (see e.g. [29, 27]).

A typical countermeasure against this type of attacks is to design hardware that minimizes the leakage of secret data (e.g. by shielding any electromagnetic emissions), or to look for an algorithm-specific solution, for example by masking intermediate variables using randomization (see [27] for a list of relevant papers). The problem with hardware-based solutions is that protection against all possible types of leakage is very hard to achieve [1], if not impossible. On the other hand, most algorithm-specific methods proposed so far are only heuristic and do not offer any formal security proof (we mention some exceptions in Sect.1.1). Moreover, they are ad-hoc in the sense that they protect only against some specific attacks that are known at the moment, instead of providing security against a large well-defined class of attacks. This raises the following, natural question: is there

*Part of this work was done when this author was supported by an EU grant MEIF-CT-2006-024300-CRYPTOSENSORS.

a systematic method of designing cryptographic schemes so that already their mathematical description guarantees that they are provably-secure, even if they are implemented on hardware that may be subject to a side-channel attack belonging to a large well-defined class of attacks? Ideally, one would like to develop a theory that (1) provides precise definition of such a class of attacks, and (2) shows how to construct systems that are secure in this model (under the assumptions that are as weak as possible). This should be viewed as moving the task of constructing cryptosystems secure against side-channel attacks from the realm of engineering or security research to cryptography, which over the last 3 decades was extremely successful in defining security models, and constructing cryptosystems that are provably-secure in those models.

General Model for Leakage Resilience. We propose a model for cryptographic computation where the class of possible side-channel attacks is extremely broad, yet simple and natural. Models similar to ours have been proposed before, in particular Micali and Reyzin [25] explicitly stated the “only computation leaks” assumption we will use. The only other assumption on the *implementation* we make is the (trivially necessary) requirement that the amount of leakage in each round is bounded. This approach is inspired by the bounded-storage and bounded-retrieval models and has to best of our knowledge never been used in this context. We stress however, that the main contribution of this paper is not the definition of the model, but the construction of an actual cryptosystem (a stream-cipher) which is provably secure in this model. Details follow.

Consider a cryptosystems CS, let \mathcal{M} denote its memory and \mathcal{M}^0 denote the data initially on \mathcal{M} (i.e. the secret key). Clearly the most general side-channel attack against a cryptosystem CS(\mathcal{M}^0) is one in which the adversary can choose any polynomial-time computable *leakage function* f and retrieve $f(\mathcal{M}^0)$ from the cryptographic machine.¹ Of course no security is achievable in this setting, as defining $f(\mathcal{M}^0) = \mathcal{M}^0$ the adversary learns the complete random key. Thus a necessary restriction we must make on f is that its output range is bounded to $\{0, 1\}^\lambda$ where $\lambda \ll |\mathcal{M}^0|$.

We assume that the adversary can apply this attack many times throughout the lifetime of the device. Technically, this will be done by dividing the execution of the algorithm implementing CS into *rounds*, and allowing the adversary to evaluate a function on the internal state of CS in each of those rounds (let f_j denote the leakage function that she chooses in the j th round, for $j = 1, 2, \dots$). In particular, in this paper we construct a stream cipher which in each round outputs a few bits.

¹Without loss of generality we can assume that the leakage function is applied only to \mathcal{M}^0 since all the other internal variables used in computation are deterministic functions of \mathcal{M}^0 .

Let q be the number of rounds we want our cryptosystem CS to run, and let \mathcal{M}^0 be the secret key that is used in the scheme. At first sight one may think that to hope for any security we would need to assume that $q \cdot \lambda < |\mathcal{M}^0|$, as otherwise the adversary can learn the entire \mathcal{M}^0 , by just retrieving in every round λ different bits of it. This trivial attack does not work any more if we consider cryptosystems which occasionally update their state. For this let \mathcal{M}^j denote the state of CS after round j .

Unfortunately, no security is possible even if we allow CS to update its state (i.e. when \mathcal{M}^j is not necessarily equal to \mathcal{M}^{j+1}) if we allow *any* (poly-time computable) f_j , to see this let $t = \lceil |\mathcal{M}|/\lambda \rceil$ and consider $f_j, j \leq t$ where each f_j outputs different λ bits of \mathcal{M}^t (note that the function $f_j, j \leq t$ can compute the future state \mathcal{M}^t from the current state \mathcal{M}^j). After the t th round the adversary has learned the complete state \mathcal{M}^t , and no security is possible beyond this point. We call this the *key-precomputation attack*.

Hence, we have to somehow restrict the leakage function if we want security even when the total amount of leaked information is (much) larger than the internal state. The restriction that we will use is that in each round, the leakage function f_j only gets as input the part of the state \mathcal{M}^j that is actually accessed in the j th round by CS. This translates into a requirement on the implementation: we assume that only computation leaks information, and the “untouched memory cells” are completely secure. As illustrated in Fig. 1, in our construction of a stream-cipher, \mathcal{M} will consist of just three parts $\mathcal{M}_0, \mathcal{M}_1$ and \mathcal{O} (where \mathcal{O} is the output tape), and in the j th round CS (and thus the leakage function f_j) will access only $\mathcal{M}_{j \bmod 2}$ and \mathcal{O} . We give the leakage function (in the j th round) access to the complete $\mathcal{M}_{j \bmod 2}, \mathcal{O}$, even if the computation of CS only access a small part of it. Thus in an actual implementation, one only must ensure that in the j th round $\mathcal{M}_{j+1 \bmod 2}$ does not leak. This requirement should easily be realizable by an actual implementation having \mathcal{M}_0 and \mathcal{M}_1 use different static memory cells (here “static” refers to the fact that this memory needs not to be refreshed, and thus should not leak any kind of radiation when not used).²

Let us mention that the above restriction is not the only natural restriction that one could make on the leakage functions to avoid the key-precomputation attack. One other op-

²Let us mention that this model also covers the case where (the not accessed) $\mathcal{M}_{j+1 \bmod 2}$ does leak in round j , as long as this leakage is independent of the leakage of (the accessed) $\mathcal{M}_{j \bmod 2}$ (i.e. when we consider an adversary Q' who can in round j choose two functions f'_j and f''_j and then gets $f'_j(\mathcal{M}_{j \bmod 2})$ and also $f''_j(\mathcal{M}_{j+1 \bmod 2})$). The reason is that we can simulate Q' by an adversary Q who just chooses one function f_j which outputs $f'_j(\mathcal{M}_{j \bmod 2})$ and also $f''_{j+1}(\mathcal{M}_{j \bmod 2})$ (thus Q in round j simply precomputes the information that Q' will learn in round $j+1$ on the non-leaking part). Note that it's not a problem that Q' might compute f''_{j+1} adaptively as a function of the information leaked in round j , as the leakage function f_j has this information too, and thus can compute the f''_{j+1} that Q' would have chosen.

tion might be to allow the state to be refreshed using external randomness. This option might be difficult to handle for many cryptosystems – including ciphers – for several reasons. For example one must make sure that all legitimate parties get the randomness in each refresh cycle, which means that parties have to be often “online” to keep their key valid, even if they almost never actually use it. Another option is to require that the leakage function is in some very weak complexity class not including the function used for key evolution.³

Leakage Resilient Stream-Cipher. The main contribution of this paper is the construction of a stream cipher S which is provably secure in the model described above. Let τ_ℓ denote the data on S 's memory which is accessed in the ℓ th round, and let K_ℓ denote the output written by S on its output tape \mathcal{O} in the ℓ th round.

The classical security notion for stream ciphers implies that one cannot distinguish K_ℓ from a random string given $K_1, \dots, K_{\ell-1}$, of course our construction satisfies this notion. But we prove much more, namely that K_ℓ is indistinguishable from random even when not only given $K_0, \dots, K_{\ell-1}$, but additionally $\Lambda_1, \dots, \Lambda_{\ell-1}$ where $\Lambda_j = f_j(\tau_j)$ and each f_j is a function with range $\{0, 1\}^\lambda$ chosen adaptively (as a function of $K_1, \dots, K_{j-1}, \Lambda_1, \dots, \Lambda_{j-1}$) by an adversary. If the adversary also gets Λ_ℓ , we cannot hope that K_ℓ is indistinguishable from random any more, as f_ℓ could for example simply output the λ first bits of K_ℓ . The best we can hope for in this case, is that K_ℓ is unpredictable (or equivalently, has high HILL-pseudoentropy), in the full version of this paper [14] we will show that for our construction this indeed is the case.

Forward Security. In many settings, it is not enough that K_ℓ is indistinguishable (or unpredictable) given the view of the adversary after round $\ell - 1$ as just described, but it should stay indistinguishable even if S leaks some information *in the future*. In our construction such “forward-security” comes up naturally, as the key K_ℓ is almost independent (in a computational sense) from the state of S after K_ℓ was output. Precise security definitions are given in Sect. 2.

Our Construction. The starting point of our construction is the concept of alternating extraction previously used in the intrusion-resilient secret-sharing scheme from [13]. We move this concept to the computational setting by proving a lemma that states that the output of any PRG has high HILL

³Interestingly, that would probably be the first case of a real-life cryptographic application where it makes sense to assume that the computational power of the adversary (in some parts of the attack scenario) is smaller than the computational power needed to execute the scheme.

pseudoentropy (i.e. is indistinguishable from some distribution with high min-entropy) even if arbitrary information about the seed is leaked. Our construction can be instantiated with any pseudorandom-generator, and the amount of leakage λ that we can tolerate in each step depends on the strength of the underlying PRG, it is at least logarithmic, but can be as large as a constant fraction of the internal state of S if the PRG is exponentially secure. The impatient reader might want to skip ahead to Section 2.2 and have a look at the actual the definition.

On (Non-)Uniformity. Throughout, we always consider non-uniform adversaries.⁴ In particular, our stream-cipher is secure against non-uniform adversaries, and we require the PRG used in the construction to be secure against non-uniform adversaries. The only step in the security proof where it matters that we are in a non-uniform setting, is in Section 5, where we use a theorem due to Barak et al. [2] which shows that two notions of pseudoentropy (called HILL and metric-type) are equivalent for circuits. In [2] this equivalence is also proved in a uniform setting, and one could use this to get a stream-cipher secure against uniform adversaries from any PRG secure against uniform adversaries. We will not do so, as for one thing the non-uniform setting is the more interesting one in our context, and moreover the exact security we could get in the uniform setting is much worse (due to the security loss in the reduction from [2] in the uniform setting).

1.1. Related work

A general theory of side-channel attacks was put forward by Micali and Reyzin [25], who propose a number of “axioms” on which such a theory should be based. In particular they formulate and motivate the assumption that “only computation leaks information”, used subsequently in e.g. [16, 28] and also in this work. As mentioned in the introduction, most published work on securing cryptosystems against side-channel attacks are ad-hoc solutions trying to prevent some particular attack or heuristics coming without security proofs, we mention some notable exceptions below.

Exposure-resilient functions [5, 9, 20] are functions whose output remains secure, even if an adversary can learn the value of some *input* bits, this model has been extensively investigated and very strong results have been obtained.

Ishai et al. [19, 18] consider the more general case of making circuits provably secure [19] and even tamper resistant [18] against adversaries who can read/tamper the value

⁴Recall that a uniform adversary can be modelled as a Turing-machine which as input gets a security parameter, whereas (more powerful) non-uniform adversaries will, for each security parameter, additionally get a different polynomial-length advice string. Equivalently, we can model non-uniform adversaries as a sequence of circuits (indexed by the security parameter).

of a bounded number of arbitrary wires in the circuit (and not just the input bits). It is interesting to compare the result from this paper with the approach of Ishai et al. On one hand, their results are generic, in the sense that they provide a method to transform any cryptosystem given as a circuit C into another circuit C_t that is secure against an adversary that can read-off up to t wires, whereas we only construct a particular primitive (a stream-cipher). On the other hand, we prove security against any side-channel attack, whereas Ishai et al. consider the particular case where the adversary can read-off the values of a few individual wires. Moreover Ishai et al. require special gates that can generate random bits, we do not assume any special hardware.

Canetti et al. [6] consider the possibility of secure computation in a setting where perfect deletion of most of the memory is not possible. Although the goal is different, their model is conceptually very similar to ours: non-perfect deletion of X is modelled by giving an adversary $f(X)$ for a sufficiently compressing function f of its choice. In their setting, the assumption that parts of the state can be perfectly erased is well motivated, unfortunately in our context this would translate to the very unrealistic requirement that some computations can be done perfectly leakage free.

The idea to define the set of leakage functions by restricting the length of function's output is taken from the bounded-retrieval model [8, 11, 10, 7, 13], which in turn was inspired by the bounded-storage model [24].⁵ Finally let us mention that some constructions of ciphers secure against general leakages were also proposed in the literature, however, their security proofs rely on very strong assumptions like the ideal-cipher model [28], or one-way permutations which do not leak any information at all [25].

1.2. Probability-theoretic preliminaries

We denote with U_n the random variable with distribution uniform over $\{0, 1\}^n$. With $X \sim Y$ we denote that X and Y have the same distribution. Let random variables X_0, X_1 be distributed over some set \mathcal{X} and let Y be a random variable distributed over \mathcal{Y} . Define the *statistical distance between X_0 and X_1* as $\delta(X_0; X_1) = 1/2 \sum_{x \in \mathcal{X}} |P_{X_0}(x) - P_{X_1}(x)|$. Moreover let $\delta(X_0; X_1|Y) := d(X_0, Y; X_1, Y)$ be the *statistical dis-*

⁵The bounded-storage model is limited in its usability by the fact that the secret key must be larger than the memory of a potential adversary, which means in the range of terabytes. In the bounded-retrieval model, the key must only be larger than the amount of data adversary can retrieve without being detected (say, by having a computer-virus send the data from an infected machine), which means in the range of Mega- or Gigabytes. Whereas in our setting the key length depends on the amount of side-channel information that leaks (in one round) from the cryptosystem considered, which (given a reasonable construction) we can assume to be as small as a few (or a few hundred) bits. In particular, unlike the bounded-storage and bounded-retrieval models, our keys need not to be made artificially huge.

tance between X_0 and X_1 conditioned on Y . If X is distributed over $\{0, 1\}^n$ then let $d(X) := \delta(X; U_n)$ denote the *statistical distance of X from a uniform distribution (over $\{0, 1\}^n$)*, and let $d(X|Y) := \delta(X; U_n|Y)$ denote the *statistical distance of X from a uniform distribution, given Y .* If $d(X) \leq \epsilon$ then we will say that X is ϵ -close to uniform. We will say that a variable X has *min-entropy k* , denoted $\mathbf{H}_\infty(X) = k$, if $\max_x \Pr[X = x] = 2^{-k}$.

Definition 1 (Extractor) A function $\text{ext} : \{0, 1\}^{k_{\text{ext}}} \times \{0, 1\}^r \rightarrow \{0, 1\}^{m_{\text{ext}}}$ is an $(\epsilon_{\text{ext}}, n_{\text{ext}})$ extractor if for any X with $\mathbf{H}_\infty(X) \geq n_{\text{ext}}$ and $K \sim U_{k_{\text{ext}}}$ we have that $d(\text{ext}(K, X), K) \leq \epsilon_{\text{ext}}$.

2. A Leakage-Resilient Stream-Cipher

We will now formally define our security notions which we already informally discussed and motivated in Sect. 1.

Initialization. The secret key of our stream cipher S consists of the three variables $A, B \in \{0, 1\}^r$ and $K_0 \in \{0, 1\}^k$. The values A, B, K_0 should be sampled uniformly at random, but only A, B must be secret, K_0 must not, one can think of K_0 as the first k bits of output of S . In an implementation, the memory of S is assumed to be split in three parts, $\mathcal{M}_0, \mathcal{M}_1, \mathcal{O}$, and for $j > 0$ we denote with $\mathcal{M}_0^{j-1}, \mathcal{M}_1^{j-1}, \mathcal{O}^{j-1}$ the contents of $\mathcal{M}_0, \mathcal{M}_1, \mathcal{O}$ at the beginning of the j th round, in particular the initial state is $\mathcal{M}_0^0 = A, \mathcal{M}_1^0 = B$ and $\mathcal{O}^0 = K_0$.

Computation. As illustrated in Fig. 1, in the ℓ th round S does only access (which means reads and rewrites) $\mathcal{M}_{\ell \bmod 2}$ and the output tape \mathcal{O} . Let τ_ℓ denote the value (on either \mathcal{M}_0 or \mathcal{M}_1) that is accessed in the ℓ th round, and $\bar{\tau}_\ell$ the value which is not accessed, i.e.

$$\tau_\ell \stackrel{\text{def}}{=} \mathcal{M}_{\ell \bmod 2}^{\ell-1} \quad \bar{\tau}_\ell \stackrel{\text{def}}{=} \mathcal{M}_{\ell+1 \bmod 2}^{\ell-1} \quad (1)$$

We will refer to the output of the ℓ th round (i.e. the value \mathcal{O}^ℓ on the output tape \mathcal{O} at the end of this round) as K_ℓ .

Adversary. As illustrated in Fig. 1, we consider adversaries Q which in the ℓ th round can adaptively choose a function f_ℓ with range $\{0, 1\}^\lambda$, and at the end of the round gets the output K_ℓ and

$$\Lambda_\ell \stackrel{\text{def}}{=} f_\ell(\tau_\ell)$$

i.e. the output of f_ℓ on input the data accessed by S in this round. We denote with \mathcal{A}_λ adversaries as just described restricted to choose leakage functions with range $\{0, 1\}^\lambda$. Let view_ℓ denote the view of the adversary after K_ℓ has been computed, i.e.

$$\text{view}_\ell = [K_0, \dots, K_\ell, \Lambda_1, \dots, \Lambda_\ell].$$

K_1, K_2, \dots be computed as $K_i = \text{ext}(K_{i-1}^{\text{next}}, C_i)$ (where K^{next} denotes the k_{ext} first bits of K and $C_i = B$ if i is odd and $C_i = A$ otherwise). So the K_i 's are computed by alternately extracting from A and B . It is not hard to show that $K_i = \text{ext}(K_{i-1}^{\text{next}}, C_i)$ is $i\epsilon_{\text{ext}}$ close to uniformly random given K_0, \dots, K_{i-1} while C_i has still enough min-entropy for our extractor (i.e. $\mathbf{H}_\infty(C_i | K_1, \dots, K_{i-1}) \geq n_{\text{ext}}$).

As shown in [13], the key K_i is even close to uniformly random when not only given K_1, \dots, K_{i-1} but also some values $f_1(C_1), \dots, f_{i-1}(C_{i-1})$ for arbitrary functions f_i as long as C_i has min-entropy at least n_{ext} (conditioned on K_0, \dots, K_{i-1} , and $f_1(C_1), \dots, f_{i-1}(C_{i-1})$).

Consider a ‘‘stream cipher’’ $S^*(A, B, K_0)$ which outputs K_1, K_2, \dots computed as described above, and an adversary Q which, before K_i is computed, can adaptively choose a function f_i and then gets $K_i, f_i(C_i)$.⁶ As explained in the previous paragraph, we can give the following security guarantee for S^* : as long as the min-entropy of C_i is at least n_{ext} (given the adversary’s view), the next output K_i is close to uniformly random (given the view of the adversary so far).

Pseudorandomity. The stream cipher S^* just described is not very useful, as it only provides security (in the sense that the next output looks random given the current view as required by our AdvInd security notion) as long as the output (i.e. the K_i 's plus the leaked information) is shorter (by at least n_{ext} bits) than the initial key.

To get security beyond that bound, we will ‘‘refresh’’ the values A, B after we extracted from them. Let $A_i = \mathcal{M}_0^i$ and $B_i = \mathcal{M}_1^i$ denote the values on \mathcal{M}_0 and \mathcal{M}_1 after round i respectively. In round i (we assume i is odd, otherwise replace the role of A and B) we extract $(K_i, X_i) = \text{ext}(K_{i-1}^{\text{next}}, B_{i-1})$, and use X_i to compute the fresh $B_i := \text{prg}(X_i)$ using a pseudorandom generator prg as illustrated in Fig. 2. If at the beginning of the i th round B_{i-1} has min-entropy at least n_{ext} (given the adversary’s view), K_{i-1}^{next} is pseudorandom (given B_i) and we assume that during this i th round no information is leaked, then X_i , and thus also $B_i = \text{prg}(X_i)$ is pseudorandom given the view of the adversary.

Of course assuming that the refreshing phase does not leak any information is completely unjustified, and we do not want to make such an assumption. As we give $\Lambda_i = f_i(B_i)$ to the adversary, we cannot hope for B_i to be pseudorandom (just consider the case where $f_i(B_i)$ are the λ first bits of B_i). Fortunately, B_i needs not to be (pseudo)random to apply alternating extraction, all we need is that B_i has high min-entropy. Of course $B_i = \text{prg}(X_i)$ cannot have more min-entropy than X_i , but as we consider computationally bounded adversaries, it is enough if B_i

⁶As K_{i-1} can be hard-coded into f_i , this function has access to all the data accessed during the computation of $K_i = \text{ext}(K_{i-1}^{\text{next}}, C_i)$

is indistinguishable from some distribution with high min-entropy. A random variable which is computationally indistinguishable from some variable with min-entropy k is said to have HILL-pseudorandomness k . It is not hard to see that a pseudorandom value B_i has high HILL-pseudorandomness when given $f_i(B_i)$ for some efficient function f_i , but this is not enough for our application, as the leakage function f_i is given access to B_{i-1} (and not just B_i), from which it can compute the seed X_i used to compute $B_i = \text{prg}(X_i)$. We will prove (Lemma 3) that for any pseudorandom generator prg , the output of $\text{prg}(X)$ on a random seed X has high HILL-pseudorandomness even if some function (with sufficiently short output) of X (and not only $\text{prg}(X)$) is leaked.

Using this lemma, we can prove that refreshing using a PRG as just described actually works, and will result in a ‘‘fresh’’ value B_i (or A_i for even i) having high HILL-pseudorandomness.

2.2. The construction

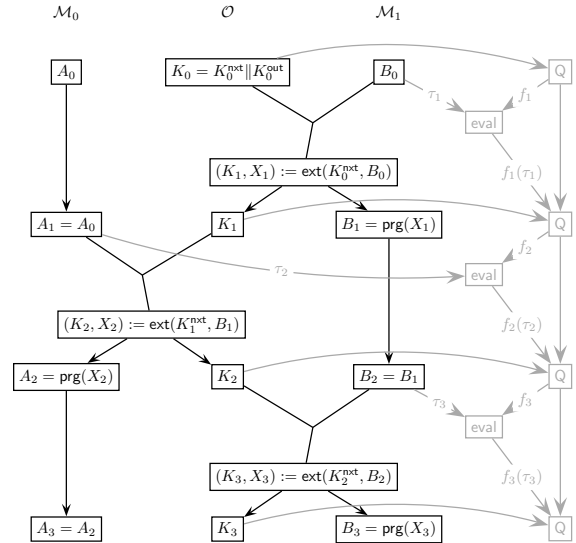


Figure 2. Illustration of the random experiment $S(A_0, B_0, K_0) \stackrel{3}{\rightsquigarrow} Q$ for the stream cipher S as described in Section 2.2.

We will now formally define the construction just outlined, which is based on an extractor $\text{ext} : \{0, 1\}^{k_{\text{ext}}} \times \{0, 1\}^r \rightarrow \{0, 1\}^{m_{\text{ext}}}$ and a pseudorandom generator $\text{prg} : \{0, 1\}^{k_{\text{prg}}} \rightarrow \{0, 1\}^r$.

State: The state of S at the beginning of round ℓ is $\mathcal{M}_0^{\ell-1}, \mathcal{M}_1^{\ell-1}, \mathcal{O}^{\ell-1}$. The computation done in round ℓ is defined below.

Get Key: Read $K_{\ell-1} = \mathcal{O}^{\ell-1}$ and parse it as $K_{\ell-1} = (K_{\ell-1}^{\text{next}}, K_{\ell-1}^{\text{out}}) \in \{0, 1\}^{k_{\text{ext}}} \times \{0, 1\}^{k_{\text{out}}}$.

Extract: Compute $\text{ext}(K_{\ell-1}^{\text{next}}, \mathcal{M}_{\ell \bmod 2}^{\ell-1})$ and parse it as $(K_\ell, X_\ell) \in \{0, 1\}^k \times \{0, 1\}^{k_{\text{prg}}}$.

Write output: Write K_ℓ on \mathcal{O} .

Refresh: Compute $\text{prg}(X_\ell)$ and write it on $\mathcal{M}_{\ell \bmod 2}$.

3. Security of S

Total Size. We denote with $\text{size}(D)$ the size of a circuit D . For an adversary $Q \in \mathcal{A}_\lambda$, $\text{size}(S \stackrel{\ell}{\rightsquigarrow} Q)$ denotes the size of a circuit needed to implement the random experiment $S \stackrel{\ell}{\rightsquigarrow} Q$.

Theorem 1 (Security of S) Let $\text{ext} : \{0, 1\}^{k_{\text{ext}}} \times \{0, 1\}^r \rightarrow \{0, 1\}^{m_{\text{ext}}}$ be an $(\epsilon_{\text{ext}}, n_{\text{ext}})$ -extractor, and let $\text{prg} : \{0, 1\}^{k_{\text{prg}}} \rightarrow \{0, 1\}^r$ be an $(\epsilon_{\text{prg}}, s_{\text{prg}})$ pseudorandom generator. Consider any $\epsilon_{\text{HILL}} > 0$ and let $\hat{s} \approx \epsilon_{\text{HILL}}^2 s_{\text{prg}} / 8r$.⁷ Consider any $\epsilon_{\text{gap}} > 0, \Delta > 0$ where

$$\epsilon_{\text{prg}} \leq \frac{\epsilon_{\text{gap}}^2}{2^\lambda} - 2^{-\Delta}, \quad n_{\text{ext}} \leq r - \Delta - (\lambda + m_{\text{ext}}) - 2 \log(1/\epsilon_{\text{gap}}) \quad (2)$$

Then for all adversaries $Q \in \mathcal{A}_\lambda$ and D where $\text{size}(S \stackrel{\ell}{\rightsquigarrow} Q) + \text{size}(D) \leq \hat{s}$ with $\delta_\ell \stackrel{\text{def}}{=} \ell^2(3\epsilon_{\text{gap}} + \epsilon_{\text{HILL}} + \epsilon_{\text{ext}})$

$$\text{AdvInd}(D, Q, S, \ell) \leq \delta_\ell \quad \text{and} \quad \text{AdvIndFwd}(D, Q, S, \ell) \leq \delta_\ell \quad (3)$$

We actually do not even need the initial key to S to be uniformly random, but only require a weaker condition as give by equations (10) and (11).

The proof of Theorem 1 is split in three parts. The first part in Section 4 on alternating extraction is information theoretic and uses ideas from the intrusion-resilient secret-sharing scheme from [13]. In the second part (Section 5) we revisit some notions and results on computational pseudoentropy. We then prove that the output of any pseudorandom generator has high HILL pseudoentropy even if information about the seed is leaked. In Section 6 we prove Theorem 1 by using the result from Section 5 to get a computational version of alternating extraction from Section 4.

How Much Leakage can we Tolerate? The amount of leakage λ we can tolerate is bounded by (2) as $\epsilon_{\text{prg}} \leq \epsilon_{\text{gap}}^2 / 2^\lambda - 2^{-\Delta}$. For concreteness, assume we set Δ such that $2^{-\Delta} \leq \epsilon_{\text{prg}} / 2$ and $\epsilon_{\text{gap}} \geq \sqrt[4]{\epsilon_{\text{prg}} / 4}$, then we can set

$$\lambda = \left\lfloor \frac{\log \epsilon_{\text{prg}}^{-1}}{2} \right\rfloor$$

To see what this means it is convenient to take an asymptotic viewpoint and think of S as a *family* of stream ciphers

⁷See Lemma 2 as to what \hat{s} exactly is.

indexed by a security parameter which we identify with k_{prg} , i.e. the input length to prg . If prg is secure against polynomial-size circuits, then $\epsilon_{\text{prg}} = 2^{-\omega(\log k_{\text{prg}})}$ (and thus $\lambda \in \omega(\log k_{\text{prg}})$), and if prg is secure against exponential size circuits, then $\epsilon_{\text{prg}} = 2^{-\Theta(k_{\text{prg}})}$ (and $\lambda \in \Theta(k_{\text{prg}})$).

Already the $\lambda \in \omega(\log k_{\text{prg}})$ case covers quite a large class of real-life attacks. In particular many attacks based on measuring the power consumption result in logarithmic-size leakages, e.g. in a so-called *Hamming weight attack* (see e.g. [21]) the adversary just learns the number of wires carrying the bit 1. Of course this value is of logarithmic length in the size of the circuit, and hence also in k_{prg} .

In the case $\lambda \in \Theta(k_{\text{prg}})$ (i.e. if prg is exponentially hard) one can leak even a constant fraction of the entire state of S .

4. Random Keys by Alternating Extraction




Figure 3. The “alternating extraction” random experiment $S^*(A, B, K_0) \stackrel{3}{\rightsquigarrow} Q$ as considered in Lemma 1.

In this section we state an information theoretic result which is very similar to the main main technical lemma used in the security proof of the intrusion-resilient secret-sharing scheme from [13], a proof appears in [14].

Basically, we consider the random experiment $S \stackrel{\ell}{\rightsquigarrow} Q$ but without the refreshing. For this let S^* denoted the construction S but where A and B are never replaced: thus in the random experiment $S^*(A, B, K_0) \stackrel{\ell}{\rightsquigarrow} Q$ where $Q \in \mathcal{A}_\lambda$, in the j th round Q chooses a function $f_j : \{0, 1\}^r \rightarrow \{0, 1\}^\lambda$ and as output gets $K_j = \text{ext}(K_{j-1}^{\text{next}}, \tau_j)$ and $\Lambda_j = f_j(\tau_j)$ where $\tau_j = B$ if j is odd and $\tau_j = A$ otherwise.

As Q attacks S^* , she learns information on A and B , and thus the min-entropy of A and B degrades. We show that as long as the min-entropy of A and B is high enough (which means more than n_{ext} as required by the extractor ext), the

next key K_j to be output is close to uniformly random when given the view after K_{j-1} has been computed.

Lemma 1 belows similar to Lemma 8 from [13] (for the special case of two players).

Lemma 1 (Alternating Extraction) *Let $\text{ext} : \{0, 1\}^{k_{\text{ext}}} \times \{0, 1\}^r \rightarrow \{0, 1\}^{m_{\text{ext}}}$ be an $(\epsilon_{\text{ext}}, n_{\text{ext}})$ -extractor. Let $A, B \in \{0, 1\}^r$ and $K_0 \in \{0, 1\}^k$ be random variables where A and B are independent and*

$$d(K_0|B) \leq \epsilon_0 \quad \mathbf{H}_{\infty}(A) \geq r - \Delta \quad \mathbf{H}_{\infty}(B) \geq r - \Delta,$$

Consider any $\lambda, \Delta, r \geq 0$ and $1 \geq \epsilon_{\text{gap}} > 0$ which satisfy

$$n_{\text{ext}} \leq r - \Delta - \lceil \ell/2 \rceil (\lambda + m_{\text{ext}}) - \log(1/\epsilon_{\text{gap}}).$$

Consider any adversary $Q \in \mathcal{A}_{\lambda}$ and the random experiment $S^*(A, B, K_0) \stackrel{\ell}{\rightsquigarrow} Q$. Recall that $\text{view}_{\ell} = [K_0, \dots, K_{\ell}, \Lambda_1, \dots, \Lambda_{\ell}]$ and $\tau_{\ell} = B$ if ℓ is odd and $\tau_{\ell} = A$ otherwise. We have

$$d(K_{\ell+1}|\text{view}_{\ell}, \tau_{\ell}) \leq (\ell + 1)\epsilon_{\text{ext}} + 2\epsilon_{\text{gap}} + \epsilon_0,$$

i.e. given τ_{ℓ} and the view of Q after the computation of K_{ℓ} , the next key $K_{\ell+1} = \text{ext}(K_{\ell}, \tau_{\ell})$ to be output by S is $(\ell\epsilon_{\text{ext}} + 2\epsilon_{\text{gap}} + \epsilon_0)$ -close to uniformly random.

5. Pseudoentropy

In this section we will prove that the output of a PRG has high HILL-pseudoentropy even if some function of the seed is leaked. We first prove this result for a weaker notion of pseudoentropy called ‘‘metric-type’’, and then use the equivalence of metric-type and HILL-pseudoentropy (Lemma 2) to get our lower bound for HILL-pseudoentropy.

Basic Definitions We denote with $\delta^D(X; Y)$ the advantage of a circuit D in distinguishing the random variables X, Y , i.e.: $\delta^D(X; Y) \stackrel{\text{def}}{=} |\Pr[D(X) = 1] - \Pr[D(Y) = 1]|$. With $\delta_s(X; Y)$ we denote $\max_D \delta^D(X; Y)$ where the maximum is over all circuits D of size s . For a random variable X over $\{0, 1\}^z$, $d_s(X) \stackrel{\text{def}}{=} \delta_s(X; U_z)$.

Definition 2 (Pseudorandom Generator) *A function $\text{prg} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ is a (δ, s) -secure pseudorandom generator (PRG) if $d_s(\text{prg}(U_n)) \leq \delta$.*

Definition 3 (HILL pseudoentropy [17, 2]) *We say X has HILL pseudoentropy k , denoted by $\mathbf{H}_{\epsilon, s}^{\text{HILL}}(X) \geq k$, if there exists a distribution Y where $\mathbf{H}_{\infty}(Y) \geq k$ and $\delta_s(X, Y) \leq \epsilon$.*

The above definition requires that there exists a distribution Y with high min-entropy that is indistinguishable from X by all distinguishers. One can also consider a notion where the quantifiers are exchanged, i.e. to allow the distribution to depend on the distinguisher.

Definition 4 (Metric-type pseudoentropy [2]) *We say X has metric-type pseudoentropy k , denoted $\mathbf{H}_{\epsilon, s}^{\text{Metric}}(X) \geq k$, if for every circuit D of size s there exists a distribution Y with $\mathbf{H}_{\infty}(Y) \geq k$ and $\delta^D(X, Y) \leq \epsilon$.*

Barak et al. [2] use the von Neumann’s min-max theorem [26] to prove the equivalence of \mathbf{H}^{HILL} and $\mathbf{H}^{\text{Metric}}$.

Lemma 2 *Let X be a distribution over $\{0, 1\}^n$. For every $\epsilon, \epsilon_{\text{HILL}} > 0$ and k , if $\mathbf{H}_{\epsilon, s}^{\text{Metric}}(X) \geq k$ then $\mathbf{H}_{\epsilon + \epsilon_{\text{HILL}}, \hat{s}}^{\text{HILL}}(X) \geq k$ where $s \in O(n\hat{s}/\epsilon_{\text{HILL}}^2)$ or equivalently $\hat{s} \in \Omega(\epsilon_{\text{HILL}}^2 s/n)$. More precisely (by inspection of the proof of Thm.5.2 in [2]) $s \leq 8n\hat{s}/\epsilon_{\text{HILL}}^2 - \zeta$ where ζ is the size of a circuit needed to compute the majority of $8n/\epsilon_{\text{HILL}}^2$ bits.*

5.1. Pseudoentropy of a PRG

By the following lemma, the output of a PRG has high metric-type pseudoentropy (and thus by Lemma 2 also high HILL-pseudoentropy) even if some function of its input is leaked.

Lemma 3 (Metric/HILL Pseudoentropy of a PRG) *Let $\text{prg} : \{0, 1\}^n \rightarrow \{0, 1\}^m$ and $f : \{0, 1\}^n \rightarrow \{0, 1\}^{\lambda}$ (where $1 \leq \lambda < n < m$) be any functions. If prg is a $(\epsilon_{\text{prg}}, s_{\text{prg}})$ -secure pseudorandom-generator, then for any $\epsilon, \Delta > 0$ satisfying $\epsilon_{\text{prg}} \leq \frac{\epsilon^2}{2^{\lambda}} - 2^{-\Delta}$, we have with $X \sim U_n$*

$$\Pr_{y:=f(X)} [\mathbf{H}_{\epsilon, s_{\text{prg}}}^{\text{Metric}}(\text{prg}(X)|f(X) = y) \geq m - \Delta] \geq 1 - \epsilon \quad (4)$$

and for any $\epsilon_{\text{HILL}} > 0$

$$\Pr_{y:=f(X)} [\mathbf{H}_{\epsilon + \epsilon_{\text{HILL}}, \hat{s}}^{\text{HILL}}(\text{prg}(X)|f(X) = y) \geq m - \Delta] \geq 1 - \epsilon \quad (5)$$

where $\hat{s} \approx \epsilon_{\text{HILL}}^2 s_{\text{prg}}/8m$.

Proof: Eq. (5) follows from (4) by Lemma 2. To prove (4) assume for contradiction that it does not hold. Hence, by Def. 4, there exists a subset

$$\mathcal{S} \subseteq \{0, 1\}^{\lambda} \quad \text{where} \quad \Pr[f(U_n) \in \mathcal{S}] > \epsilon \quad (6)$$

such that for each $a \in \mathcal{S}$ there exists a distinguisher D_a of size at most s_{prg} such that for every random variable Z with $\mathbf{H}_{\infty}(Z) \geq m - \Delta$ we have (again $X \sim U_n$)

$$|\Pr[D_a(Z) = 1] - \Pr[D_a(\text{prg}(X)) = 1|f(X) = a]| \geq \epsilon \quad (7)$$

Consider some $a \in \mathcal{S}$ for which

$$\Pr[f(U_n) = a] > 2^{-\lambda} \cdot \epsilon \quad (8)$$

Such an a exists by (6) and as $|\mathcal{S}| = 2^{\lambda}$. Our distinguisher for the PRG prg will be the distinguisher D_a satisfying (7) and (8). It remains to prove that D_a breaks prg with advantage higher than ϵ_{prg} . For $\beta \in \{0, 1\}$ let $\mathcal{I}_{\beta} := \{x \in \{0, 1\}^m : D_a(x) = \beta\}$

Claim 1 For some $\beta \in \{0, 1\}$ we have $|\mathcal{I}_\beta| < 2^{m-\Delta}$

Proof of Claim: Assume for contradiction that $|\mathcal{I}_\beta| \geq 2^{m-\Delta}$ for $\beta = 0$ and $\beta = 1$. For $\beta \in \{0, 1\}$ and $X \sim U_n$ let $p_\beta = \Pr[\text{D}_a(\text{prg}(X)) = \beta | f(X) = a]$. Let Z' be a random variable distributed uniformly over $S'_0 \cup S'_1$ where S'_β is an arbitrary subset of \mathcal{I}_β of size $p_\beta 2^{m-\Delta}$ (here we use the fact that $|\mathcal{I}_\beta| \geq 2^{m-\Delta}$). Clearly since $|S'_0 \cup S'_1| = 2^{m-\Delta}$ we have that $\mathbf{H}_\infty(Z') = m - \Delta$ and by construction (with $X \sim U_n$)

$$\underbrace{\Pr[\text{D}_a(Z') = 1]}_{=\Pr[Z' \in S'_1]=p_1} - \Pr[\text{D}_a(\text{prg}(X)) = 1 | f(X) = a] = 0$$

contradicting (7). This finishes the proof of the claim. \triangle
For β as guaranteed by the above claim, we have

$$\Pr[\text{D}_a(U_m) = \beta] = |\mathcal{I}_\beta|/2^m < 2^{-\Delta}. \quad (9)$$

By equation (7), using that $\mathbf{H}_\infty(U_m) = m > m - \Delta$ we get:

$$\underbrace{|\Pr[\text{D}_a(U_m) = \beta] - \Pr[\text{D}_a(\text{prg}(X)) = \beta | f(X) = a]|}_{< 2^{-\Delta}} \geq \epsilon$$

By assumption, we have $\epsilon \geq \epsilon^2/2^\lambda > 2^{-\Delta}$. As for any $x, y, \epsilon \geq 0$ we have that $|x-y| \geq \epsilon$ and $\epsilon > x$ implies $y \geq \epsilon$, the above equation implies $\Pr[\text{D}_a(\text{prg}(X)) = \beta | f(X) = a] \geq \epsilon$, and further with $X \sim U_n$

$$\begin{aligned} & \Pr[\text{D}_a(\text{prg}(X)) = \beta] \\ & \geq \underbrace{\Pr[\text{D}_a(\text{prg}(X)) = \beta | f(X) = a]}_{\geq \epsilon} \cdot \underbrace{\Pr[f(X) = a]}_{> 2^{-\lambda} \cdot \epsilon \text{ by (8)}} > \frac{\epsilon^2}{2^\lambda}. \end{aligned}$$

By (9) and the above equation, the advantage of D_a for U_m and $\text{prg}(U_n)$ is at least

$$\Pr[\text{D}_a(\text{prg}(U_n)) = \beta] - \Pr[\text{D}_a(U_m) = \beta] > \frac{\epsilon^2}{2^\lambda} - 2^{-\Delta} \geq \epsilon_{\text{prg}}$$

which contradicts the $(\epsilon_{\text{prg}}, s_{\text{prg}})$ -security of prg . \square

6 Putting Things Together

In this section we show how the security of S as stated in Theorem 1 follows from the results in the two previous sections. S is based on ext and prg where

- $\text{ext} : \{0, 1\}^{k_{\text{ext}}} \times \{0, 1\}^r \rightarrow \{0, 1\}^{m_{\text{ext}}}$ is an $(\epsilon_{\text{ext}}, n_{\text{ext}})$ -extractor.
- $\text{prg} : \{0, 1\}^{k_{\text{prg}}} \rightarrow \{0, 1\}^r$ is an $(\epsilon_{\text{prg}}, s_{\text{prg}})$ pseudorandom generator.

Further we set $k_{\text{out}} := m_{\text{ext}} - k_{\text{ext}} + k_{\text{prg}}$ (thus $m_{\text{ext}} = k_{\text{ext}} + k_{\text{out}} + k_{\text{prg}}$) and fix parameters $\Delta, \epsilon_{\text{gap}}, \lambda$ satisfying

$$\epsilon_{\text{prg}} \leq \frac{\epsilon_{\text{gap}}^2}{2^\lambda} - 2^{-\Delta} \text{ and } n_{\text{ext}} \leq r - \Delta - (\lambda + m_{\text{ext}}) - \log \epsilon_{\text{gap}}^{-1}$$

We also fix some $\epsilon_{\text{HILL}} > 0$ and set $\hat{s} := \epsilon_{\text{HILL}}^2 s_{\text{prg}}/8r$.

The following lemma quantifies how much security is “lost” by one round of our stream cipher. Let size_i denote the size of the circuit realizing the i th round of the experiment $S \rightsquigarrow Q$, then $\sum_{i=1}^\ell \text{size}_i = \text{size}(S \rightsquigarrow Q)$.

Lemma 4 (The i th round) Consider the random experiment $S \rightsquigarrow Q$. Then if before round $i \leq \ell$ (recall that $\bar{\tau}_{i-1} = \tau_i = B_i$ if i is odd and $\bar{\tau}_{i-1} = \tau_i = A_i$ otherwise) for some $s_{i-1} \leq \hat{s}$ and $\epsilon' \stackrel{\text{def}}{=} \epsilon_{\text{HILL}} + \epsilon_{\text{gap}}$

$$\mathbf{H}_{\epsilon', s_{i-1}}^{\text{HILL}}(A_{i-1} | \text{view}_{i-1}, B_{i-1}) \geq r - \Delta$$

$$\mathbf{H}_{\epsilon', s_{i-1}}^{\text{HILL}}(B_{i-1} | \text{view}_{i-1}, A_{i-1}) \geq r - \Delta$$

$$d_{s_{i-1}}(K_{i-1} | \text{view}_{i-2}, \bar{\tau}_{i-1}) \leq \epsilon_{i-1}$$

then with $s_i \stackrel{\text{def}}{=} s_{i-1} - \text{size}_i, s'_i \stackrel{\text{def}}{=} s_{i-1} - \text{size}(\text{ext})$ and $\epsilon_i \stackrel{\text{def}}{=} \epsilon_{i-1} + \epsilon_{\text{ext}} + \epsilon_{\text{gap}} + \epsilon'$.

$$d_{s'_i}(K_i, X_i | \text{view}_{i-1}, \bar{\tau}_i) \leq \epsilon_i$$

with probability $1 - \epsilon_{\text{gap}} - \epsilon_i$

$$\mathbf{H}_{\epsilon', s_i}^{\text{HILL}}(A_i | \text{view}_i, B_i) \geq r - \Delta$$

$$\mathbf{H}_{\epsilon', s_i}^{\text{HILL}}(B_i | \text{view}_i, A_i) \geq r - \Delta$$

This lemma can be proven by using Lemma 3 to get a computational version of the alternating extraction Lemma 1 (see the full version [14] for details). We'll now see how Theorem 1 is implied by this lemma. Let $\epsilon_0 = 0$ and $s_0 = \hat{s}$, then $\epsilon_\ell = \ell(2\epsilon_{\text{gap}} + \epsilon_{\text{ext}} + \epsilon_{\text{HILL}})$ and $s_\ell = \hat{s} - \text{size}(S \rightsquigarrow Q)$. If the initial key A_0, B_0, K_0 satisfies

$$\mathbf{H}_{\epsilon', s_0}^{\text{HILL}}(A_0 | B_0) \geq r - \Delta \quad \mathbf{H}_{\epsilon', s_0}^{\text{HILL}}(B_0 | A_0) \geq r - \Delta \quad (10)$$

$$d(K_0 | B_0) = \epsilon_0 \quad (11)$$

as it is the case in Theorem 1, by Lemma 4 with probability $1 - \sum_{i=1}^\ell (\epsilon_{\text{gap}} + \epsilon_i) \geq 1 - \ell(\epsilon_{\text{gap}} + \epsilon_\ell)$ we have

$$d_{s'_\ell}(K_\ell, X_\ell | \text{view}_{\ell-1}, \bar{\tau}_\ell) \leq \epsilon_\ell$$

This proves (note that $s'_\ell < s_\ell$) the bound for the AdvInD as stated in Theorem 1. To prove the bound for AdvInDFwd we move $X_\ell, K_\ell^{\text{next}}$ to the conditioned part (recall that $K_\ell = K_\ell^{\text{next}} \| K_\ell^{\text{out}}$)

$$d_{s'_\ell}(K_\ell^{\text{out}} | K_\ell^{\text{next}}, X_\ell, \text{view}_{\ell-1}, \bar{\tau}_\ell) \leq \epsilon_\ell$$

Then we apply the prg to X_ℓ

$$\underbrace{d_{s'_\ell - |\text{prg}|}}_{> s_\ell} (K_\ell^{\text{next}} | K_\ell^{\text{out}}, \text{view}_{\ell-1}, \underbrace{\overbrace{(\bar{\tau}_{\ell+1}, \bar{\tau}_\ell)}^{A_\ell, B_\ell}}_{\text{prg}(X_\ell)}) \leq \epsilon_\ell$$

Which proves the bound on AdvInDFwd from Theorem 1.

References

- [1] R. Anderson and M. Kuhn. Tamper resistance: a cautionary note. In *WOEC'96*
- [2] B. Barak, R. Shaltiel, and A. Wigderson. Computational analogues of entropy. In *RANDOM-APPROX*, 2003.
- [3] E. Biham and A. Shamir. Differential fault analysis of secret key cryptosystems. In *CRYPTO* 1997.
- [4] D. Boneh, R. A. DeMillo, and R. J. Lipton. On the importance of checking cryptographic protocols for faults (extended abstract). In *EUROCRYPT* 1997.
- [5] R. Canetti, Y. Dodis, S. Halevi, E. Kushilevitz, and A. Sahai. Exposure-resilient functions and all-or-nothing transforms. In *EUROCRYPT* 2000.
- [6] R. Canetti, D. Eiger, S. Goldwasser, and D.-Y. Lim. How to protect yourself without perfect shredding. In *ICALP (2)* 2008.
- [7] D. Cash, Y. Z. Ding, Y. Dodis, W. Lee, R. J. Lipton, and S. Walfish. Intrusion-resilient key exchange in the bounded retrieval model. In *TCC* 2007.
- [8] G. D. Crescenzo, R. J. Lipton, and S. Walfish. Perfectly secure password protocols in the bounded retrieval model. In *TCC* 2006.
- [9] Y. Dodis, A. Sahai, and A. Smith. On perfect and adaptive security in exposure-resilient cryptography. In *EUROCRYPT* 2001.
- [10] S. Dziembowski. Intrusion-resilience via the bounded-storage model. In *TCC* 2006.
- [11] S. Dziembowski. On forward-secure storage. In *CRYPTO* 2006.
- [12] S. Dziembowski and U. M. Maurer. On generating the initial key in the bounded-storage model. In *EUROCRYPT* 2004.
- [13] S. Dziembowski and K. Pietrzak. Intrusion-resilient secret sharing. In *FOCS* 2007.
- [14] S. Dziembowski and K. Pietrzak. Full version of this paper. Cryptology ePrint Archive, Report 2008/240, 2008. <http://eprint.iacr.org>.
- [15] K. Gandolfi, C. Mourtel, and F. Olivier. Electromagnetic analysis: Concrete results. In *CHES* 2001.
- [16] S. Goldwasser, Y. T. Kalai, and G. Rothblum. One-time programs. In *CRYPTO* 2008.
- [17] J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999.
- [18] Y. Ishai, M. Prabhakaran, A. Sahai, and D. Wagner. Private Circuits II: Keeping Secrets in Tamperable Circuits. In *EUROCRYPT* 2006.
- [19] Y. Ishai, A. Sahai, and D. Wagner. Private Circuits: Securing Hardware against Probing Attacks. In *CRYPTO* 2003.
- [20] J. Kamp and D. Zuckerman. Deterministic extractors for bit-fixing sources and exposure-resilient cryptography. *SIAM J. Comput.*, 36(5):1231–1247, 2007.
- [21] J. Kelsey, B. Schneier, D. Wagner, and C. Hall. Side channel cryptanalysis of product ciphers. In J.-J. Quisquater, Y. Deswarte, C. Meadows, and D. Gollmann, editors, *ESORICS* 1998.
- [22] P. C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In *CRYPTO* 1996.
- [23] P. C. Kocher, J. Jaffe, and B. Jun. Differential power analysis. In *CRYPTO* 1999.
- [24] U. M. Maurer. A provably-secure strongly-randomized cipher. In *EUROCRYPT* 1990.
- [25] S. Micali and L. Reyzin. Physically observable cryptography (extended abstract). In *TCC* 2004.
- [26] J. V. Neumann. Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320, 1928.
- [27] E. N. of Excellence (ECRYPT). The side channel cryptanalysis lounge. http://www.crypto.ruhr-uni-bochum.de/en_sclounge.html.
- [28] C. Petit, F.-X. Standaert, O. Pereira, T. G. Malkin, and M. Yung. A block cipher based prng secure against side-channel key recovery. Cryptology ePrint Archive, Report 2007/356, 2007. <http://eprint.iacr.org/>.
- [29] J.-J. Quisquater and F. Koene. Side channel attacks: State of the art, October 2002. [27].
- [30] J.-J. Quisquater and D. Samyde. Electromagnetic analysis (ema): Measures and counter-measures for smart cards. In *E-smart* 2001.
- [31] S. P. Vadhan. Constructing locally computable extractors and cryptosystems in the bounded-storage model. *J. Cryptology*, 17(1):43–77, 2004.