A Leakage-Resilient Mode of Operation

Krzysztof Pietrzak

CWI Amsterdam, The Netherlands

Abstract. A weak pseudorandom function (wPRF) is a cryptographic primitive similar to – but weaker than – a pseudorandom function: for wPRFs one only requires that the output is pseudorandom when queried on *random* inputs. We show that unlike "normal" PRFs, wPRFs are seed-incompressible, in the sense that the output of a wPRF is pseudorandom even if a bounded amount of information about the key is leaked.

As an application of this result we construct a simple mode of operation which – when instantiated with any wPRF – gives a *leakage-resilient* stream-cipher. The implementation of such a cipher is secure against *every* side-channel attack, as long as the amount of information leaked *per round* is bounded, but overall can be arbitrary large. The construction is simpler than the previous one (Dziembowski-Pietrzak FOCS'08) as it only uses a single primitive (a wPRF) in a straight forward manner.

1 Introduction

Traditionally, cryptographic algorithms are designed to withstand adversaries that can attack the cryptosystem in a black-box fashion. This means that all the adversary can do is to query the system at hand according to the security definition. In many settings this is not a realistic assumption, as real-world adversaries attack concrete *implementations* of cryptosystems, that possibly leak information which cannot be efficiently computed from black-box access alone. Attacks exploiting such leakage are called side-channel attacks. In the last two decades we saw many cryptanalytic attacks exploiting side-channels as running-time [31], electromagnetic radiation [39, 19], power consumption [33] and fault detection [4, 3]. A recent example [18] is the side-channel attack against KeeLoq (which refers to the "KeeLoq block-cipher" and some particular mode in which this cipher is used), which is widely used as e.g. anti-theft mechanisms for cars. Although the KeeLoq block-cipher seems not to be very secure to start with [9, 27], the devastating side-channel attack of [18] exploits a weakness in the mode in which the cipher is used, rather than a weakness in the cipher itself, and it would still be applicable even if the KeeLoq block-cipher was replaced with a strong block-cipher, say AES ([18] Talk of Christof Paar). It is thus an intriguing question whether there exist modes of operation which are provably secure against a wide class of side-channel attacks if instantiated with any block-cipher.

In this paper we answer this question affirmatively, by proposing a mode of operation (cf. Figure 1) which turns any weak PRF into a stream-cipher which is provably secure against *all* side-channel attacks, assuming only that the amount

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: 10.1007/978-3-642-01001-9_35

A. Joux (Ed.): EUROCRYPT 2009, LNCS 5479, pp. 462–482, 2009.

of leakage in each round is bounded, and that only memory which is actually accessed in some round leaks in this round. Such a "leakage-resilient" cipher was recently constructed in [17], the main advantage of our new construction is its simplicity, it can be instantiated with any weak PRF (e.g. with a block-cipher like AES), whereas the construction from [17] additionally required extractors.

The simplicity of the construction (as compared to [17]) comes at the price of more involved security proof. Besides the technical tools we already used in [17], we will need new results concerning the security of weak PRFs when neither the key nor the inputs are uniform. The technique we use to prove this results can also be applied in other settings, e.g. for encryption schemes, and thus could be of independent interest.

Why Leakage-Resilience. Leakage-resilience is an extremely strong security notion considering adversaries who can choose arbitrary leakage functions. To practitioners this may seem like an overkill, after all, why consider unrealistic side-channels which leak some very involved function of the state instead of using some ad-hoc countermeasures against "real" side-channels? A lesson cryptographers have learned in the last decades is that ad-hoc arguments usually result in insecure systems, and this very much applies to the young history of side-channel cryptanalysis. Implementing cryptographic algorithms in a straight forward way, will almost certainly make them very susceptible to side-channel attacks. Often – like in differential power analysis [33, 8] – such attacks extract a little bit of information in each evaluation, and then combine this information to get the secret key. Thus it is crucial that an implementation does not leak even small amounts of (useful) information. In contrast, "leakage-resilient" algorithms as considered in this work guarantee security even if in *each* invocation a *bounded* amount of *arbitrary* information is leaked.

We advocate the following approach to side-channel security: first cryptographers design a leakage-resilient algorithm C, with the guarantee that whenever you implement C such that in each invocation $\leq \lambda$ bits of information leak, the implementation is safe. This still leaves the task of implementing C such that the $\leq \lambda$ leakage bound is met.[1] The rationale here is that this task is clearly much more realistic than having to implement an algorithm in way where nothing leaks at all, as it would be necessary if the algorithm would come with no bound on the leakage that can be tolerated. (cf. Kocher [32] for a similar argument). It is only at this stage that one should consider using ad-hoc measures like masking or blinding, using special circuit designs, and so on. Cryptography seems to be of limited use at this stage, but a background on existing attacks and implementation details is helpful here, thus this task is something that should be left to security researchers and engineers.

Some Related Work. Most papers on side-channel security – like [31, 39, 19, 33, 4, 3] mentioned in the introduction – consider attacks and/or countermeasures

[1] Note that this is unavoidable, as when one cannot keep at least some uncertainty about the internal state, one cannot hope to get a secure implementation.

against a specific side-channel. From the papers considering general models for side-channel attacks, the work of Micali and Reyzin [35] on "physically observable cryptography" is particularly insightful and written in a language accessible to cryptographers. Their model is based on five "axioms", some of which are (more or less explicitly) used in our model.

Ishai et al. [29, 28] consider a model where the adversary can choose some wires in the circuit, and then learns the values carried by those wires during the computation. What makes their work exceptional is that they were the first to *prove* how to implement *any* algorithm secure against an interesting side-channel (i.e. probing attacks).[2] The field of exposure-resilient cryptography [11] considers the more restricting case where the adversary could learn some of the *input* bits.

Very recently [1] showed that some *particular* public-key encryption schemes are surprisingly robust against leakage: the scheme stays secure even if the min-entropy of the key is just a constant fraction of the min-entropy of a random key. We prove a similar result for *any* weak PRFs, but in order prove security even for keys with such low min-entropy, we need the weak PRF to be exponentially hard, whereas [1] can do so with some particular superpolynomial assumptions (learning with error and lattice assumptions).

Papers that consider constructions of *stream-ciphers* which withstand side-channel attacks (as in this work and [17]) include [32, 35, 36]. Kocher [32] considers a very simple construction where one simply iterates a hash function (SHA256 is suggested). This work is kept informal, with no proofs or even formal claims, but contains several interesting conceptual ideas. Micali and Reyzin [35] investigate *reductions* of side-channel resistant primitives, in particular they show that the Blum-Micali construction is secure, assuming the implementation of the underlying permutation already satisfies some strong form of side-channel security. The work which aims at a goal most similar to ours is Petit et al. [36]. They propose and analyze a block-cipher based construction, where security against sides-channels is achieved by making it hard to "combine" leakages from different rounds.[3] Their underlying model [41] is motivated by practical considerations, considering leakage-functions and attacks that have been successfully used to break systems. Compared to [36], we take a much more theoretical approach, our setting is more general and the underlying assumptions are weaker

[2] Formally, Ishai et al. do the following: let $t \geq 0$ be some constant and let $[X]$ denote a $(t+1)$ out of $(t+1)$ secret sharing of the value X. They construct a general compiler, which turns every circuit $G(.)$ into a circuit $G_t(.)$ (of size $t^2|G|$) such that $[G(X)] = G_t([X])$ for all inputs X, and moreover one does not learn any information on $G(X)$ even when given the value carried by any t wires in the circuit $G_t(.)$ while evaluating the input $[X]$. This transformation uses multiparty-computation, which is quite different from all other approaches we discuss here.

[3] By using a forward secure primitive, one can ensure that *past* keys cannot be combined with the current key, as they cannot even be computed. For *future* keys, this is more tricky, as the cipher itself must be able to efficiently derive that keys.

in several aspects.[4] The tools and techniques from [17] and this paper cannot be used to prove security of the constructions from [32, 35, 36] (or any other construction we are aware of), as those constructions are insecure against arbitrary leakage functions as considered in this work, even if the underlying primitives are ideal (e.g. Random oracles in [32] or ideal ciphers in [36]) and only one bit of information leaks per invocation of the underlying primitive. (but this does by no means mean that they are insecure against side-channels that arise in practice.)[5]

Some interesting recent results in settings which are similar or otherwise relevant to general models of side-channel security include [5], who show how to securely realize protocols when perfect deletion is not possible. Goldwasser et al. [22] construct "one-time programs" from simple hardware satisfying some weak form of side-channel security. Dodis and Wichs [12] solve the long standing open problem of two round authenticated key-agreement from non-uniform keys. (See the full version [37] for a more detailed discussion on those papers.)

1.1 Leakage-Resilient Cryptography

In this section we informally introduce and motivate the model of "leakage-resilient cryptography" from [17].

Consider some keyed cryptographic primitive CP. The most general side-channel attack against $\mathsf{CP}(S_0)$ – where S_0 denotes the secret initial state – is to allow an attacker to choose any leakage function f, which then is evaluated on the initial state S_0, and the adversary receives $f(S_0)$.[6] Clearly we cannot hope for any security at all here, as f could simply output the complete state $f(S_0) = S_0$. Thus, it is necessary to somehow restrict the range of the leakage function, we will consider functions with range $\{0,1\}^\lambda$, where $\lambda \ll |S_0|$ is some parameter. The idea to define the set of leakage functions by restricting the output length was inspired by the bounded-retrieval model [10, 14, 13, 6, 16], which in turn was inspired by the bounded-storage model [34, 15, 42].

[4] In particular 1. We prove security in the standard model, whereas [36] work in the ideal-cipher model 2. The security notion considered in [36] is key-recovery, whereas we use unpredictability (and, in a limited context, indistinguishability). 3. The leakage functions considered in [36] (namely Hamming weight or identity plus noise) are motivated by leakages observed in practice, whereas we bound only the amount, not the type of information leaked 4. Finally, and most importantly, our approach differs in how the observed leakage (cf. point 3.) can be exploited in order to break the security notion (cf. point 2.). [36] show that a so called template attack [7] cannot recover the key, whereas we prove security against every efficient adversary.

[5] A crucial requirement we need from the construction in order to prove leakage-resilience, is that the state can be split in (at least) two parts, and this parts evolve independently, in the sense that any interaction between them is public. Formally, one must be able to express the cipher as a process as in Lemma 5 in this paper.

[6] Here the leakage function is applied only to the state S_0, and not to any internal variables appearing in the computation. This can be done without loss of generality as all the internal variables are simply functions of the state S_0, and thus can be computed by f.

As the implementation of any cryptosystem will leak more information the longer it runs, we want to allow the attacker A to adaptively choose different leakage functions during the lifetime of the system. For this, we assume that CP runs in rounds (where a "round" is just some well defined part of the computation), and denote with S_i the state of CP after round i (to simplify the exposition we assume that the size of the state remains constant).

The attacker A we consider can adaptively choose a leakage function f_i before the ith round, and after round i receives $f_i(S_{i-1})$, i.e. the leakage function evaluated on the state at the beginning of round i. Unfortunately also here no security is possible beyond round t, where $t \cdot \lambda \geq |S_0|$, as A can simply define the f_i's such that $f_i(S_{i-1})$ will be some λ bits of S_t. (note that for $i \leq t$, f_i can compute the future state S_t from its input S_{i-1}.) After round t the attacker A has learned the entire state S_t, and no security is possible beyond this point.

Thus if we want security even after (much) more than $|S_0|$ bits have leaked, we need to further restrict the leakage functions. The restriction we use is one of the "axioms" from [35], and states that "only computation leaks information". This means that f_i does not get the entire state S_{i-1} as input, but only the part of the state that is actually accessed by CP in the ith round.

On Efficient Leakage Functions. As we consider a computational primitive, and the total leakage can be larger than the entire state, we can only allow *efficient* leakage functions.[7] This is not explicitly stated, but naturally comes up in the model, where the main result (Theorem 2) puts an upper bound on the size of a *circuit* computing the entire random experiment in which the cipher is attacked.

On (non)-Uniformity. Throughout, we always consider non-uniform adversaries.[8] In particular, our main reduction is non-uniform, which means we prove that if an adversary *exists* who breaks the stream-cipher, then an adversary (of related complexity) *exists* who breaks the underlying weak PRF. The only step in the proof where we need non-uniformity is a lemma from [2] which relates two types of pseudoentropy notions. As [2] also prove this lemma in a uniform setting (albeit which much worse parameters), it should be possible (though we didn't check the details) to make our reduction uniform, that is to show how to efficiently construct an adversary against the weak PRF from any adversary against the stream-cipher. (we refer to Goldreich's article [20] as to why such a reduction is desirable.)

[7] A computationally unbounded leakage function could simply compute and output the initial state from the output of the stream cipher. If one assumes that the total leakage is smaller than the key [1, 12], considering computationally unbounded leakage functions is meaningful.

[8] Recall 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), which is what we will do.

Relaxing Bounded Leakage. As described above, in each round we allow the adversary to choose any function f with range $\{0,1\}^\lambda$, and she then learns the leakage $f(S)$, where S is the state accessed in this round. Note that this also captures any efficient leakage function g, where there exists another (efficient) leakage function f with range $\{0,1\}^\lambda$ such that $S \rightarrow f(S) \rightarrow g(S)$ is a Markov chain and where one can efficiently sample $g(S)$ given $f(S)$ (as an adversary in our model can ask for $f(S)$, and then compute $g(S)$ himself). This e.g. covers the case where the leakage function outputs a noisy version of S.

We chose to work with bounded leakage as it is a very clean and intuitive model, but for the proof we actually only require that $f(S)$ does not contain more than λ bits of "useful" information on S. Formally, "useful" means that the HILL-pseudoentropy (a notion to be defined in Section 4) of S does not drop by much more than λ bits given $f(S)$. Unfortunately this most general notion is quite unintuitive to work with.[9]

Relaxing the "only computation leaks information" Axiom. The leakage function in round i gets as input only that part of the state which is accessed in that round. This translates to the requirement on the implementation that memory which is not accessed, must not leak at all. In our model and for our particular construction (and also [17]) allowing the adversary to choose a single leakage function f with λ bits output, and then giving her the leakage $f(S^+)$ (where with S^+ we denote the part of the state which is accessed and S^- denotes the remaining state) is equivalent to let her choose two function f' and f'' with $\lambda/2$ bits output respectively, and then output the leakage $f'(S^+)$ and $f''(S^-)$. Thus it is o.k. if the entire state leaks as long the leakage of S^+ and S^- is independent. In particular, we also get security against attacks which seem not to obey the "only computation leaks information" axiom, like the cold boot attack from [23] (see also [1]), who show how measure significant parts of a key that was stored on some memory, even after power is turned off.

1.2 Seed Incompressibility

As main new technical tools we prove bounds on the security of weak PRFs when the key (or the inputs) are not uniformly random as assumed in the security definition for weak PRFs.

Recall that the standard security notion for a pseudorandom function (PRF) $\mathsf{F} : \{0,1\}^\kappa \times \{0,1\}^n \rightarrow \{0,1\}^m$ requires that for a random key $k \in \{0,1\}^\kappa$ no efficient attacker can distinguish $\mathsf{F}(k,.)$ from a uniformly random function. Motivated by the question if random-oracles can (in some settings) be instantiated with efficient functions, Halevi et al. [24] investigate the question whether

[9] A special more intuitive case – which is still more general than bounded leakage – is to consider any (not necessarily) efficient leakage function g where there exists an efficient f with range $\{0,1\}^\lambda$, such that given $f(S)$ one can efficiently sample some "fake" leakage $\tilde{g}(S)$ where $[S, g(S)]$ is computationally indistinguishable from $[S, \tilde{g}(S)]$ (bounded leakage corresponds to $\tilde{g} = g$). Note that here the sampling algorithm only gets $f(S)$, whereas the distinguisher gets S.

"seed-incompressible" functions exist. They consider a setting where an adversary initially gets a "compressed key" $f(k)$ (where $f : \{0,1\}^\kappa \rightarrow \{0,1\}^\lambda$ and $\lambda < \kappa$). A simple observation is that by giving this extra input to an adversary, no function $\mathsf{F}(k,.)$ can possibly be a PRF, as $f(k)$ could e.g. encode the first λ bits of $\mathsf{F}(k,X)$ (for some fixed X), and thus $\mathsf{F}(k,.)$ becomes easily distinguishable from random.

In this paper we revisit the concept of seed incompressibility, but for *weak* pseudorandom functions (wPRF): F is a wPRF, if $\mathsf{F}(k,.)$ cannot be distinguished from random, if queried on *random* inputs. Thus an adversary gets to see $X_1,\ldots,X_q$ and $Z_1,\ldots,Z_q$, and then must guess whether $Z_i = \mathsf{F}(k,X_i)$ or $Z_i = \mathbf{R}(X_i)$ where $\mathbf{R}$ is a uniformly random function. Unlike for normal PRFs, for wPRFs it is not clear if and how a compressed seed $f(k)$ helps the distinguisher, e.g. now simply setting $f(k)$ to denote the λ first bits of $\mathsf{F}(k,X)$ for some fixed input X will not trivially break the security of $\mathsf{F}(k,.)$ as here the adversary cannot choose the inputs X for which she gets to see $\mathsf{F}(k,X)$.

Of course by leaking λ bits of the key, we must tolerate some security loss. In particular, if we use the trivial attack just described (leaking λ bits of $\mathsf{F}(k,X)$), the adversary can get "lucky", and one of the q queries $X_1,\ldots,X_q$ will hit the fixed input X. Because of that, the adversary has some extra advantage of roughly $q/2^n$ (compared to an adversary not getting $f(k)$). Further, if we assume that the best attack against F is brute-force search over the keyspace, then leaking λ bits of the key will degrade the security by a factor of 2^λ. As we prove in Lemma 2, it doesn't get much worse than that: if $\mathsf{F}(k,.)$ cannot be distinguished with advantage more than ϵ, then the advantage (against somewhat smaller adversaries) is still bounded by roughly $2^\lambda(\epsilon + q^2/2^{n+1})$ (here we set t from Lemma 2 to n, and assume that n is large enough so that the last term in (3) can be ignored.)

We actually do not consider the setting where the key k is random, and then $f(k), |f(k)| = \lambda$ is leaked, but the more general case where k is sampled from some distribution with min-entropy at least $|k| - \lambda$. (and we need this more general case later when proving the security of the leakage-resilient stream-cipher), as for any function f and uniformly random k, k has still (expected) min-entropy at least $|k| - \lambda$ given $f(k)$.

We then prove a similar result (Lemma 3) concerning the security of wPRFs assuming the inputs (as opposed to the key) are not uniformly random.

Proof Sketch. We show that any wPRF is secure even when the secret key is only sampled from some distribution with min-entropy $|k| - \lambda$ by a (uniform) reduction. Assume an adversary A can distinguish $\mathsf{F}(k,.)$ from a random function (when queried on random inputs $X_1,\ldots,X_q$) with advantage ϵ'. Using the Markov bound one can show that this implies that a key k sampled from the above distribution is "weak" with probability at least $\epsilon'/2$, where a key k is said to be weak, if the distinguishing advantage of A, conditioned on the key being k, is at least $\epsilon'/2$. If k is now sampled from the uniform distribution (and not a distribution with min-entropy $|k| - \lambda$), then k will be weak with probability at least $\epsilon'/2^{\lambda+1}$, i.e. we loose at most a factor 2^λ. The crucial point is that when observing the output of a function $g(.)$ on sufficiently many random inputs, then (using the Hoeffding bound) one can almost certainly distinguish the cases where $g(.)$ is $f(k,.)$ for a weak k and the case where $g(.)$ is a random oracle, as by definition of a weak key, the probability of A outputting 1 differs by at least $\epsilon'/2$ for both cases. Thus we can define an adversary which does the above sampling and outputs 0 and 1 respectively in the two above cases. As outlined, this adversary has a distinguishing advantage of at least $\epsilon'/2^{\lambda+1}$.[10] In the above argument it is important that in the case where $g(.)$ is a random oracle, we can sample many *independent* guess bits of A. This is not possible when considering "normal" PRFs, as then the adversary A can simply query $g(.)$ on some fixed inputs, and her guess bits will be completely dependent. This is the point in the proof where we exploit the fact that we consider *weak* PRFs.

1.3 Applications and Reductions

The unpredictability and indistinguishability based notions used in this paper are the strongest possible considering general leakage-functions, and a stream cipher satisfying them is sufficient to realize important primitives like stateful leakage-resilient symmetric authentication and encryption.[11]

It would be very interesting to construct a leakage-resilient pseudorandom function, as then we could implement those symmetric primitives in a *stateless* way. Let us mention here that cryptographic reductions, like the GGM construction of PRFs form PRGs [21], will in general *not* preserve leakage-resilience.

1.4 Notation

For a set $\mathcal{X}$, we denote with $X \overset{*}{\leftarrow} \mathcal{X}$ that X is assigned a value sampled uniformly at random from $\mathcal{X}$. To save on notation, we write X^i to denote a sequence $X_1, \ldots, X_i$. $\mathbf{R}_{n,m}$ denotes a uniformly random function $\{0,1\}^n \rightarrow \{0,1\}^m$, $\mathbf{R}_n$ denotes $\mathbf{R}_{n,n}$.

2 Leakage-Resilient Stream-Cipher from a Weak PRF

Figure 1 illustrates the mode of operation for which we prove that it gives a leakage-resilient stream cipher if instantiated with any weak PRF. Below we first

[10] The expression (3) in Lemma 2 is a bit more complicated than that. The last term in (3) is the error from the Hoeffding bound, and the second to last term is due to the fact that the sampled outputs are not completely independent as required by the Hoeffding bound.

[11] For authentication it is sufficient that the secret X_i used is unpredictable, thus here we can allow the adversary to observe the leakage in the round where X_i is computed. For semantically secure encryption, e.g. when using a one-time pad $C = M \oplus X_i$, we need X_i to be indistinguishable, thus here the adversary cannot get the leakage in round i, but can so for all other rounds $j < i$ (and, as we have forward security, also $j > i$).

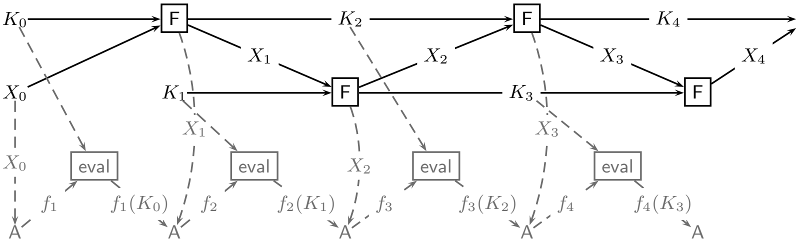

Fig. 1. Leakage resilient stream-cipher S^{F} from a seed-incompressible weak pseudorandom function F. The regular evaluation is shown in black, the attack related part is shown in gray with dashed lines.

define this construction, and state a Theorem which bounds the security of S^{F} as a *normal* stream-cipher. We then define what a *leakage-resilient* stream-cipher is. Then we state our main theorem (Theorem 2) which bounds the security of S^{F} as a leakage-resilient stream-cipher in terms of the security of F as a weak PRF.

The Construction: Let $\mathsf{F} : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^{\kappa+n}$ be a function. Then, with S^{F} we denote the following simple construction of a stream cipher.

Initialization: The initial state is $S_0 = [K_0, K_1, X_0]$, where $K_0, K_1 \overset{*}{\leftarrow} \{0,1\}^\kappa$ and $X_0 \overset{*}{\leftarrow} \{0,1\}^n$. Only K_0, K_1 must be secret, X_0 can be public.
State: The state before the ith round is $S_{i-1} = [K_{i-1}, K_i, X_{i-1}]$.
Computation: In the ith round, $\mathsf{S}^{\mathsf{F}}(S_{i-1})$ computes

$$(K_{i+1}, X_i) := \mathsf{F}(K_{i-1}, X_{i-1})$$

and outputs X_i. Then the state $S_{i-1} = [K_{i-1}, K_i, X_{i-1}]$ is replaced with $S_i = [K_i, K_{i+1}, X_i]$ (note that K_i is not accessed in the ith round).

Security of S without Side-Channels: Theorem 1 below states that the output of S^{F} is pseudorandom (i.e. is a secure stream-cipher in the "classical" sense) if F is a secure weak pseudorandom function. The proof of this theorem is a straight forward hybrid argument and for space reasons only give in the full version of this paper [37]. The security of S^{F} is stated in terms of the security of F as a weak pseudorandom function (wPRF), which is defined like a normal PRF except that the inputs are random and not adversarially chosen.

Definition 1 (wPRF). *$\mathsf{F} : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^m$ is a <u>weak</u> $(\epsilon_{\mathsf{prf}}, s_{\mathsf{prf}}, q_{\mathsf{prf}})$-secure pseudorandom function (wPRF) if for all A of size s_{prf} and for random variables $K \overset{*}{\leftarrow} \{0,1\}^\kappa$ and for $i = 1, \ldots, q_{\mathsf{prf}}$*

$$X_i \overset{*}{\leftarrow} \{0,1\}^m \qquad Y_i = \mathsf{F}(K, X_i) \qquad R_i = \mathbf{R}_{n,m}(X_i)$$

we have $\Pr[\mathsf{A}(X^{q_{\mathsf{prf}}}, Y^{q_{\mathsf{prf}}}) = 1] - \Pr[\mathsf{A}(X^{q_{\mathsf{prf}}}, R^{q_{\mathsf{prf}}}) = 1] \le \epsilon_{\mathsf{prf}}$

Theorem 1 (Security without Leakage). *If* F *is a* $(\epsilon_{\mathsf{prf}}, s_{\mathsf{prf}}, 1)$ *secure wPRF, then for any* $\ell \in \mathbb{N}$*, no adversary of size* $s_{\mathsf{prf}} - \ell \cdot |\mathsf{F}|$ *can distinguish the first* $\ell + 1$ *blocks as output by* S^{F} *from uniformly random with advantage more than* $\ell \cdot \epsilon_{\mathsf{prf}}$*.*

Side-Channel Adversary: As outlined in Section 1.1, we consider an adversary A who can attack S^{F} by choosing any function $f_i : \{0,1\}^{\kappa} \to \{0,1\}^{\lambda}$ before round i, and at the end of the round receives the normal output X_i of S^{F} and also the leakage $\Lambda_i \overset{\mathsf{def}}{=} f_i(K_{i-1})$. In round i, $\mathsf{S}^{\mathsf{F}}(S_{i-1})$ only access K_{i-1} and X_{i-1}, thus giving K_{i-1} as input to f_i means that f_i can use the entire state that S^{F} accesses in round i. Note that we don't have to explicitly give X_{i-1} as input to f_i, as A must only decide on f_i after she got X_{i-1} and thus can hard-code it into f_i. We denote with $\mathcal{A}_\lambda$ the set of adversaries as just described restricted to choose leakage functions with range $\{0,1\}^{\lambda}$.

Leakage-Resilient Security Notion: Let $\mathtt{view}_\ell$ denote the view of the adversary after X_ℓ has been computed, i.e.

$$\mathtt{view}_\ell = [X_0, \ldots, X_\ell, \Lambda_1, \ldots, \Lambda_\ell].$$

With $\mathtt{view}_\ell^- = \mathtt{view}_\ell \setminus X_\ell$ we denote $\mathtt{view}_\ell$ but without the last output X_ℓ. The security notion we consider requires that $X_{\ell+1}$ is *indistinguishable* from random, even when given $\mathtt{view}_\ell$ (which will imply that it is *unpredictable* given $\mathtt{view}_{\ell+1}^-$).

We denote with $\mathsf{S}(S_0) \overset{\ell}{\leadsto} \mathsf{A}$ the random experiment where an adversary $\mathsf{A} \in \mathcal{A}_\lambda$ attacks S (initialized with $S_0 = [K_0, K_1, X_0]$) for ℓ rounds (cf. Fig. 1), and with $\mathtt{view}(\mathsf{S}(S_0) \overset{\ell}{\leadsto} \mathsf{A})$ we denote the view $\mathtt{view}_\ell$ of A at the end of the attack. For any circuit $\mathsf{D} : \{0,1\}^* \to \{0,1\}$ (with one bit output), we denote with $\mathsf{AdvInd}(\mathsf{D}, \mathsf{A}, \mathsf{S}, \ell)$ the advantage of D in distinguishing K_ℓ from a random $U_n \overset{*}{\leftarrow} \{0,1\}^n$ given $\mathtt{view}(\mathsf{S}(S_0) S \overset{\ell-1}{\leadsto} \mathsf{A})$, formally

$$\mathsf{AdvInd}(\mathsf{D}, \mathsf{A}, \mathsf{S}, \ell) = |p_{real} - p_{rand}| \quad \text{where}$$

$$p_{rand} \overset{\mathsf{def}}{=} \Pr_{S_0}[\mathsf{D}(\mathtt{view}(\mathsf{S}(S_0) \overset{\ell-1}{\leadsto} \mathsf{A}), U_n) = 1]$$

$$p_{real} \overset{\mathsf{def}}{=} \Pr_{S_0}[\mathsf{D}(\mathtt{view}(\mathsf{S}(S_0) \overset{\ell-1}{\leadsto} \mathsf{A}), X_\ell) = 1]$$

Security of S against Side-Channel Attacks: The security of S^{F} will depend on the security of F as a weak pseudorandom function. Recall that the complexity a non-uniform adversary is captured by the size of a circuit describing it. For a circuit D, we let $\mathtt{size}(\mathsf{D})$ denote its size. We will also write $\mathtt{size}(\mathsf{S} \overset{\ell-1}{\leadsto} \mathsf{A})$ to denote the size of a circuit needed to implement the entire random experiment $\mathsf{S} \overset{\ell-1}{\leadsto} \mathsf{A}$, as illustrated in Figure 1, where eval denotes a circuit which on input the description of a function $f : \{0,1\}^{\kappa} \to \{0,1\}^{\lambda}$ and $K \in \{0,1\}^{\kappa}$ computes and outputs $f(K)$.

Theorem 2 (Security with Leakage). *Let* $\mathsf{F} : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^{\kappa+n}$ *be a* $(\epsilon_{\mathsf{prf}}, s_{\mathsf{prf}}, n/\epsilon_{\mathsf{prf}})$*-secure wPRF where* $\epsilon_{\mathsf{prf}} \geq n \cdot 2^{-n/3}$ *and* $n \geq 20$*. Let* $\lambda = \log(\epsilon_{\mathsf{prf}}^{-1})/6$ *and* $s' = s_{\mathsf{prf}}\epsilon_{\mathsf{prf}}^2/2^{\lambda+2}(n+\kappa)^3$*. Then for any adversary* $\mathsf{A} \in \mathcal{A}_\lambda$ *and distinguisher* D *where* $\mathtt{size}(\mathsf{S} \overset{\ell-1}{\rightsquigarrow} \mathsf{A}) + \mathtt{size}(\mathsf{D}) \leq s'$ *we have for any* $\ell \in \mathbb{N}$

$$\mathsf{AdvInd}(\mathsf{D}, \mathsf{A}, \mathsf{S}, \ell) \leq 8 \cdot \ell \cdot \epsilon_{\mathsf{prf}}^{1/12} \qquad (1)$$

On λ: Note that the amount of leakage $\lambda = \log(\epsilon_{\mathsf{prf}}^{-1})/6$ we tolerate depends on the hardness of the underlying wPRF. Thus if F is secure against adversaries of super-polynomial size, i.e. $\epsilon_{\mathsf{prf}} = 2^{\omega(\log\kappa)}$, then the amount of leakage is at least super-logarithmic $\lambda = \omega(\log\kappa)$. This already covers many practical attacks like Hamming weight attacks (see e.g. [30]).

If F is exponentially hard, i.e. $\epsilon_{\mathsf{prf}} = 2^{-\Omega(\kappa)}$, then $\lambda = \Omega(\kappa)$, and thus we can even leak a constant fraction of the internal state in each round.

Security loss: The security loss in the theorem is significant: the 12th root of ϵ_{prf} comes up in the distinguishing advantage. In the full version [37] we discuss several approaches which potentially can be used to prove a much better bound.

Unpredictability: Theorem 2 states that when given the view of an adversary A who attacked S for $\ell - 1$ rounds, the next value X_ℓ to be computed is indistinguishable from random. If the adversary is also given $\Lambda_\ell = f_\ell(K_{\ell-1})$ (i.e. the leakage computed in round ℓ), X_ℓ cannot be pseudorandom any more, as Λ_ℓ could e.g. be the λ first bits of X_ℓ. In the case where Λ_ℓ is also leaked, one can still prove (using Lemma 4) that X_ℓ is unpredictable: for any $\delta > 0$, with probability $1-\delta$ the random variable X_ℓ has $n - \lambda - \log(\delta^{-1})$ bits of "HILL-pseudoentropy" (a notion to be defined in Section 4).

Forward Security: Like the construction from [17], also S^F is forward secure: Theorem 2 holds even for a stronger security notion than AdvInd, where the distinguisher D is additionally given entire state of S^F after round $\ell + 1$.

Instantiation with a block-cipher: Our construction requires a wPRF $\mathsf{F} : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^{\kappa+n}$. Such an F can be constructed from any secure block-cipher $\mathsf{BC} : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ like AES. (AES comes with different security parameters $\kappa = n = 128$ and $\kappa = n = 256$). For this we have to do some range expansion, e.g. by setting ($\|$ denotes concatenation)

$$\mathsf{F}(K, X) = \mathsf{BC}(K, X\|0))\|\mathsf{BC}(K, X\|1). \qquad (2)$$

Here $\mathsf{F} : \{0,1\}^\kappa \times \{0,1\}^{n-1} \to \{0,1\}^{2n}$ is a secure PRF (and thus wPRF) assuming that $\mathsf{BC} : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ is a pseudorandom permutation, which is the standard security notion for block-ciphers. [12]

[12] Let us stress, that just assuming that BC is a wPRF is not sufficient as (2) is *not* a secure range expansion of wPRFs (see e.g. [38] for some secure constructions).

3 wPRF with Non-uniform Keys and Inputs

We will need the following classical technical lemma several times.

Lemma 1 (Hoeffding's inequality [26]). *Let $X_1, \dots, X_t$ be independent random variables where for $1 \le i \le t$: $\Pr(X_i \in [a_i, b_i]) = 1$. Then, for the sum of these variables $X = X_1 + \dots + X_t$ we have the inequality:*

$$\Pr[X - \mathrm{E}[X] \ge t\epsilon] \le \exp\left(-\frac{2\,t^2\,\epsilon^2}{\sum_{i=1}^{t}(b_i - a_i)^2}\right)$$

Recall that a random variable Z has min-entropy k, denoted $H_\infty(Z) = k$ if for all z in the range of Z we have $\Pr[Z = z] \le 2^{-k}$.

Definition 2 (wPRF with non-uniform keys and inputs). *We call a function $\mathsf{F} : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^m$ a $(\epsilon_{\mathsf{prf}}, s_{\mathsf{prf}}, q_{\mathsf{prf}})$-secure wPRF with α-low keys, if it's a wPRF as in Definition 1, whenever the key K comes from any distribution with min-entropy $\kappa - \alpha$ (and not uniformly random).*

Similarly, we say F is a $(\epsilon_{\mathsf{prf}}, s_{\mathsf{prf}}, q_{\mathsf{prf}})$-secure wPRF with β-low inputs, if it's a wPRF as in Definition 1, except that the inputs X_i come from any distribution with min-entropy $m - \beta$.

Non-uniform Keys. By the following lemma, every wPRF (using uniform keys) is a wPRF for α-low keys. The loss in security is roughly $2^{\alpha+1}$, which is almost optimal.

Lemma 2. *For any $\alpha > 0$ and $t \in \mathbb{N}$: If $\mathsf{F} : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^m$ is a $(\epsilon_{\mathsf{prf}}, s_{\mathsf{prf}}, q_{\mathsf{prf}})$-secure wPRF (for uniform keys), then it is a $(\epsilon'_{\mathsf{prf}}, s'_{\mathsf{prf}}, q'_{\mathsf{prf}})$-secure wPRF with α-low keys if the following holds*[13]

$$\begin{aligned} q_{\mathsf{prf}} &\ge q'_{\mathsf{prf}} \cdot t \\ \epsilon_{\mathsf{prf}} &\le \epsilon'_{\mathsf{prf}}/2^{\alpha+1} - \frac{q_{\mathsf{prf}}^2}{2^{n+1}} - 2 \cdot \exp\left(-\frac{t^2 \cdot \epsilon'^2_{\mathsf{prf}}}{8}\right) \\ s_{\mathsf{prf}} &\ge s'_{\mathsf{prf}} \cdot t \end{aligned} \tag{3}$$

Proof. Assume there exists a random variable K_α where $H_\infty(K_\alpha) = \kappa - \alpha$, but where F is not a $(\epsilon'_{\mathsf{prf}}, s'_{\mathsf{prf}}, q'_{\mathsf{prf}})$-secure wPRF if the key is K_α. To prove the Lemma, we must show that then F is not $(\epsilon_{\mathsf{prf}}, s_{\mathsf{prf}}, q_{\mathsf{prf}})$-secure wPRF for uniformly random keys. By assumption, there exists an adversary $\mathsf{A}, |\mathsf{A}| \le s'_{\mathsf{prf}}$ such that

$$\sum_{k \in \{0,1\}^\kappa} \Pr[k = K_\alpha] \cdot \xi_k > \epsilon'_{\mathsf{prf}} \tag{4}$$

[13] As ϵ'_{prf} appears twice in eq.(3), we cannot easily express ϵ'_{prf} as a function of ϵ_{prf}. One can get a closed expression at the price of a worse bound by e.g. replacing ϵ'_{prf} in (3) with ϵ_{prf}, one then gets (for $t \in \mathbb{N}$ of our choice): $q'_{\mathsf{prf}} := q_{\mathsf{prf}}/t, \quad s'_{\mathsf{prf}} := s_{\mathsf{prf}}/t, \quad \epsilon'_{\mathsf{prf}} := 2^{\alpha+1}\left(\epsilon_{\mathsf{prf}} + q_{\mathsf{prf}}^2/2^{n+1} + 2 \cdot \exp\left(-t^2 \cdot \epsilon_{\mathsf{prf}}^2/8\right)\right)$.

where ξ_k denotes A's advantage conditioned on the key being k, i.e. with $X_i \overset{*}{\leftarrow} \{0,1\}^n$, $Y_i = \mathsf{F}(k, X_i)$, $R_i \leftarrow \mathbf{R}_{n,m}(X_i)$ (for $i = 1, \ldots, q'_{\mathsf{prf}}$)

$$\xi_k \overset{\mathsf{def}}{=} \Pr[\mathsf{A}(X^{q'_{\mathsf{prf}}}, Y^{q'_{\mathsf{prf}}}) = 1] - \Pr[\mathsf{A}(X^{q'_{\mathsf{prf}}}, R^{q'_{\mathsf{prf}}}) = 1]$$

We say $k \in \{0,1\}^\kappa$ is weak if $\xi_k \geq \epsilon'_{\mathsf{prf}}/2$, and let $\mathcal{K} \subset \{0,1\}^\kappa$ denote the set of all weak keys. From (4) we get by Markov

$$\Pr[K_\alpha \in \mathcal{K}] \geq \epsilon'_{\mathsf{prf}}/2.$$

Let K be uniform over $\{0,1\}^\kappa$. If we define an event $\mathcal{E}$ depending on K by $\Pr[\mathcal{E}|K = k] = \Pr[K_\alpha = k]/2^{\alpha-\kappa}$ it satisfies (see [37] for the proof)

$$\Pr[\mathcal{E}] = 2^{-\alpha} \quad \text{and} \quad \Pr[K = k|\mathcal{E}] = \Pr[K_\alpha = k]$$

With this we can lower bound the probability that the uniformly random key K is weak as

$$\Pr[K \in \mathcal{K}] \geq \Pr[\mathcal{E}]\Pr[K \in \mathcal{K}|\mathcal{E}] = \Pr[\mathcal{E}]\Pr[K_\alpha \in \mathcal{K}] = \frac{\Pr[K_\alpha \in \mathcal{K}]}{2^\alpha} \geq \frac{\epsilon'_{\mathsf{prf}}}{2^{\alpha+1}} \tag{5}$$

We will construct an adversary $\tilde{\mathsf{A}}$, where for $X_i \overset{*}{\leftarrow} \{0,1\}^n, Y_i = \mathsf{F}(k, X_i), R_i \leftarrow \mathbf{R}_{n,m}(X_i)$ the adversary $\tilde{\mathsf{A}}(X^{q_{\mathsf{prf}}}, R^{q_{\mathsf{prf}}})$ (where $q_{\mathsf{prf}} = q'_{\mathsf{prf}} \cdot t$) will almost always output 0, whereas $\tilde{\mathsf{A}}(X^{q_{\mathsf{prf}}}, Y^{q_{\mathsf{prf}}})$ will almost always output 1 if $k \in \mathcal{K}$. So $\tilde{\mathsf{A}}$ breaks the security of F as a weak PRF with advantage at least $\epsilon_{\mathsf{prf}} \approx \Pr[k \in \mathcal{K}] \geq \epsilon'_{\mathsf{prf}}/2^{\alpha+1}$. Let

$$\phi = \Pr[\mathsf{A}(X^{q_{\mathsf{prf}}}, R^{q_{\mathsf{prf}}}) = 1] \tag{6}$$

where the probability if over the choice of the $X_i \overset{*}{\leftarrow} \{0,1\}^n$, the random function $\mathbf{R}_{n,m}$ used to compute $R_i = \mathbf{R}_{n,m}(X_i)$ and also A (if it's not deterministic). Our adversary $\tilde{\mathsf{A}}$ on input $X^{q_{\mathsf{prf}}}, Z^{q_{\mathsf{prf}}}$, does the following.

- Split the input in t equal parts which we denote $(\hat{X}_1, \hat{Z}_1), \ldots, (\hat{X}_t, \hat{Z}_t)$ (so e.g. $\hat{X}_i = X_{(i-1)q'_{\mathsf{prf}}+1}, \ldots, X_{i \cdot q'_{\mathsf{prf}}}$).
- For $i = 1, \ldots, t$ compute $T_i \leftarrow \mathsf{A}(\hat{X}_i, \hat{Z}_i)$ and let

$$T := \sum_{i=1}^t T_i$$

If $(T - t \cdot \phi) \leq t \cdot \epsilon'_{\mathsf{prf}}/4$ then $\tilde{\mathsf{A}}$ outputs 0, otherwise she outputs 1.

By the following two claims, $\tilde{\mathsf{A}}$ will almost never output 1 if the Z_i are random, but will output 1 with probability almost $\epsilon_{\mathsf{prf}}/2^{\alpha+1}$ if the the Z_i were computed by F.

Claim 1. *Let $X_i \overset{*}{\leftarrow} \{0,1\}^n$ and $R_i = \mathbf{R}_{n,m}(X_i)$, then*

$$\Pr[\tilde{\mathsf{A}}(X^{q_{\mathsf{prf}}}, R^{q_{\mathsf{prf}}}) = 1] \leq \exp\left(-\frac{t^2 \cdot \epsilon'^2_{\mathsf{prf}}}{8}\right) + \frac{q^2_{\mathsf{prf}}}{2^{n+1}}$$

Proof. By definition $\tilde{\mathsf{A}}$ will output 1 iff $(T - t \cdot \phi) > t \cdot \epsilon'_{\mathsf{prf}}/4$. In the case where the Z_i are computed as $\mathbf{R}_{n,m}(X_i)$ (as it is the case for the R_i in this claim) we have by eq.(6) $t \cdot \phi = \mathrm{E}[T]$, thus

$$\Pr[\tilde{\mathsf{A}}(X^{q_{\mathsf{prf}}}, R^{q_{\mathsf{prf}}}) = 1] = \Pr\left[T - \mathrm{E}[T] > t \cdot \frac{\epsilon'_{\mathsf{prf}}}{4}\right] \tag{7}$$

Let $T'_1, \ldots, T'_t$ be independent binary random variables, where for $j = 1, \ldots, t$ the T_j is sampled by choosing a uniformly random function $\mathbf{R}^j : \{0,1\}^n \to \{0,1\}^m$ and (for $i = 1, \ldots, q'_{\mathsf{prf}}$) $X_{j,i} \xleftarrow{*} \{0,1\}^n$, $R_{j,i} = \mathbf{R}^j(X_i)$ and setting $T'_j = \mathsf{A}(X_{j,1}, \ldots, X_{j,q'_{\mathsf{prf}}}, R_{j,1}, \ldots, R_{j,q'_{\mathsf{prf}}})$. Further let $T' := \sum_{j=1}^{t} T'_j$. As the T'_j's are independent, we can use Hoeffding's inequality (Lemma 1) to upper bound

$$\Pr\left[T' - \mathrm{E}[T'] > t \cdot \frac{\epsilon'_{\mathsf{prf}}}{4}\right] \leq \exp\left(-\frac{t^2 \cdot \epsilon'^2_{\mathsf{prf}}}{8}\right) \tag{8}$$

This bound does not apply to (7), as unlike the T'_j, the T_j are not completely independent, as we use *the same* random function $\mathbf{R}_{n,m}$ for each T_j. We will show that this is not a big problem if the domain is large enough, as conditioned on all the X_i's being different, the R_i's will have the same distribution in the computation of the T_j and T'_j; Let $\mathcal{E}$ denote the event, which holds if all the $q_{\mathsf{prf}} = q'_{\mathsf{prf}} \cdot t$ values $X_{j,i}$ (sampled to compute T or T') are pairwise distinct. As those values are all sampled independently and uniformly from $\{0,1\}^n$, by the birthday bound

$$\Pr[\neg\mathcal{E}] \leq \frac{q^2_{\mathsf{prf}}}{2^{n+1}} \tag{9}$$

Conditioned on $\mathcal{E}$, the distribution of the T_i's and T'_i (and thus of T and T') is identical, in particular

$$\Pr\left[T' - \mathrm{E}[T'] > t \cdot \frac{\epsilon'_{\mathsf{prf}}}{4} \middle| \mathcal{E}\right] = \Pr\left[T - \mathrm{E}[T] > t \cdot \frac{\epsilon'_{\mathsf{prf}}}{4} \middle| \mathcal{E}\right] \tag{10}$$

The claim now follows from (7)-(10). □

Claim 2. *Let* $K \xleftarrow{*} \{0,1\}^\kappa$ *and for* $i = 1, \ldots, q_{\mathsf{prf}}$: $X_i \xleftarrow{*} \{0,1\}^n$ *and* $Y_i = \mathsf{F}(K, X_i)$*, then*

$$\Pr[\tilde{\mathsf{A}}(X^{q_{\mathsf{prf}}}, Y^{q_{\mathsf{prf}}}) = 1] \geq \frac{\epsilon'_{\mathsf{prf}}}{2^{\alpha+1}}\left(1 - \exp\left(-\frac{t^2 \cdot \epsilon'^2_{\mathsf{prf}}}{8}\right)\right)$$

Proof. We have

$$\Pr[\tilde{\mathsf{A}}(X^{q_{\mathsf{prf}}}, Y^{q_{\mathsf{prf}}}) = 1] \geq \Pr[K \in \mathcal{K}] \cdot \Pr[\tilde{\mathsf{A}}(X^{q_{\mathsf{prf}}}, Y^{q_{\mathsf{prf}}}) = 1 | K \in \mathcal{K}] \tag{11}$$

By (5) we can lower bound the first term on the right side in (11) as

$$\Pr[K \in \mathcal{K}] \geq \epsilon'_{\mathsf{prf}}/2^{\alpha+1} \tag{12}$$

It remains to upper bound the second term. For this recall that $\tilde{\mathsf{A}}$ outputs 0 if $|T - t \cdot \phi| > t \cdot \epsilon'_{\mathsf{prf}}/4$, where $T = \sum_{j=1}^{t} T_j$ and each T_j is the output of $\mathsf{A}(X^{q'_{\mathsf{prf}}}, Y^{q'_{\mathsf{prf}}})$ where $Y_i = \mathsf{F}(K, X_i)$ (here the $X^{q'_{\mathsf{prf}}}$ are independent for each j but K is fixed). If $K \in \mathcal{K}$, then by definition of $\mathcal{K}$ we have $|\mathrm{E}[T_j] - \phi| \geq \epsilon'_{\mathsf{prf}}/2$, and thus $\tilde{\mathsf{A}}$ will only output 0, if the value of T is bounded away by at least $t \cdot \epsilon'_{\mathsf{prf}}/4$ from its expectation, again using the Hoeffding bound

$$\Pr[\tilde{\mathsf{A}}(X^{q_{\mathsf{prf}}}, Y^{q_{\mathsf{prf}}}) = 0 | K \in \mathcal{K}] = \Pr\left[T - \phi > t \cdot \frac{\epsilon'_{\mathsf{prf}}}{4}\right] \leq \exp\left(-\frac{t^2 \cdot \epsilon'^2_{\mathsf{prf}}}{8}\right)$$

The claim follows from this equation and (11),(12). □

The bound on $\tilde{\mathsf{A}}$'s advantage ϵ_{prf} as claimed in the lemma follows from the two claims above. The bound on the size s_{prf} and number of queries q_{prf} made by $\tilde{\mathsf{A}}$ follows directly from the definition of $\tilde{\mathsf{A}}$. ■

Non-uniform Inputs. We just showed that a wPRF stays secure even if the key is not uniform. In the full version of the paper we prove a similar result for the case where the *inputs* are not uniformly random. We only consider the case where the adversary gets a single input/output pair.

Lemma 3. *Let $\beta > 0$, then if $\mathsf{F} : \{0,1\}^{\kappa} \times \{0,1\}^n \to \{0,1\}^m$ is a $(\epsilon_{\mathsf{prf}}, s_{\mathsf{prf}}, 1)$-secure wPRF (for uniform inputs), it's also a $(\epsilon'_{\mathsf{prf}}, s'_{\mathsf{prf}}, 1)$-secure wPRF for β-low entropy input, where for any $t \in \mathbb{N}$*

$$\epsilon_{\mathsf{prf}} \leq \epsilon'_{\mathsf{prf}}/2^{\beta+1} - 2 \cdot \exp\left(-\frac{2 \cdot t^2 \cdot \epsilon'^2_{\mathsf{prf}}}{64}\right)$$

$$s_{\mathsf{prf}} \geq s'_{\mathsf{prf}} \cdot 2t$$

4 Proof of Theorem 2

Proof Sketch. We will prove the security of S^{F} (cf. Figure 1) by proving that if the state X_{i-1}, K_{i-1} accessed in round i is independent and has HILL-pseudoentropy $n - 2\lambda$ and $\kappa - 2\lambda$, respectively, then also the output X_i, K_{i+1} has such a HILL-pseudoentropy given the leakage $\Lambda_i = f(X_{i-1}, K_{i-1})$ (Lemma 7). Though we unavoidably get some degradation in the "quality" of the pseudoentropy (in terms of ϵ, s in Definition 3 below), this degradation is only additive, and thus we can sum it up over all rounds.[14]

[14] This summation to bound the degradation in security is quite tedious. It might seem that one could get a much simpler proof using a hybrid argument, where for the jth hybrid one would simply replace the output in the first j rounds (having high HILL-pseudoentropy) with some (indistinguishable) output having high min-entropy. Unfortunately we can't make this intuition work, the reason is that high HILL-pseudoentropy only implies existence of an indistinguishable random variable with high min-entropy, but gives no means as to how to sample it. Thus it is not clear how to efficiently sample the hybrids just described.

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

Definition 3 (HILL-pseudoentropy[25, 2]). *We say X has* HILL *pseudoentropy k, denoted by $\mathbf{H}^{\mathsf{HILL}}_{\epsilon,s}(X) \geq k$, if there exists a distribution Y with min-entropy $\mathbf{H}_\infty(Y) = k$ where $\delta_s(X;Y) \leq \epsilon$.*

Definition 4 (PRG). *A function* $\mathsf{prg} : \{0,1\}^n \to \{0,1\}^m$ *is a (δ, s)-secure pseudorandom generator (PRG) if $\delta_s(\mathsf{prg}(U_n)\ ;\ U_m) \leq \delta$.*

Thus $\mathsf{prg}(Z)$ is indistinguishable from random if $Z \stackrel{*}{\leftarrow} \{0,1\}^n$. If some function $f(Z)$ of the seed is leaked, then $\mathsf{prg}(Z)$ will not look random any more, as e.g. $f(Z)$ could just output some bits of $\mathsf{prg}(Z)$. The following lemma states that if the range of f is not too big, then $\mathsf{prg}(Z)$ will still have high HILL-pseudoentropy.

Lemma 4 (Pseudoentropy of a PRG, [17]). *Let* $\mathsf{prg} : \{0,1\}^n \to \{0,1\}^m$ *and $f : \{0,1\}^n \to \{0,1\}^\lambda$ (where $1 \leq \lambda < n < m$) be any functions. If* prg *is a $(\epsilon_{\mathsf{prg}}, s_{\mathsf{prg}})$-secure pseudorandom-generator, then for any $\epsilon, \Delta > 0$ satisfying $\epsilon_{\mathsf{prg}} \leq \frac{\epsilon^2}{2^\lambda} - 2^{-\Delta}$, we have with $Z \stackrel{*}{\leftarrow} \{0,1\}^n$ and for any $\epsilon_{\mathsf{HILL}} > 0$*

$$\Pr_{Z \stackrel{*}{\leftarrow} \{0,1\}^n}[\mathbf{H}^{\mathsf{HILL}}_{\epsilon+\epsilon_{\mathsf{HILL}},\hat{s}}(\mathsf{prg}(Z)|f(Z)) \geq m - \Delta] \geq 1 - \epsilon \tag{13}$$

where $\hat{s} \approx \epsilon^2_{\mathsf{HILL}} s_{\mathsf{prg}}/8m$.

We will use the following technical lemma about some general random processes to show that the inputs X_i and keys K_i in the computation of S^{F} are independent.

Lemma 5 ([16]). *Let A_0, B_0 be independent random variables, and $\phi_1, \phi_2, \ldots$ be any sequence of functions. Let $A_1, A_2, \ldots,\ B_1, B_2, \ldots$ and $V_1, V_2, \ldots$ be defined as*

$$\begin{aligned}((A_{i+1}, V_{i+1}), B_{i+1}) &:= (\phi_{i+1}(A_i, V_1, \ldots, V_i), B_i) \\ &\qquad \textit{if } i \textit{ is even} \\ (A_{i+1}, (V_{i+1}, B_{i+1})) &:= (A_i, \phi_{i+1}(B_i, V_1, \ldots, V_i)) \\ &\qquad \textit{otherwise}\end{aligned}$$

Then $B_i \to \{V_1, \ldots, V_i\} \to A_i$ (and $A_i \to \{V_1, \ldots, V_i\} \to B_i$) is a Markov chain (or equivalently, A_i and B_i are independent given the $V_1, \ldots, V_i$)

Combining Lemmata 2, 3 and 4, we can prove Lemma 6 below, which states that the output $\mathsf{F}(K,X)$ of a wPRF has high HILL-pseudoentropy, even if K and X have high min-entropy (but are independent) and given some leakage $f(K,X)$. We set $t = n/\epsilon_{\mathsf{prf}}$ in Lemma 2 and 3, moreover we need the domain $\{0,1\}^n$ of F to be large enough, in particular, we will assume that (with ϵ_{prf} as in the lemma below)

$$\epsilon_{\mathsf{prf}} \geq \frac{n^2}{2^{n+1} \cdot \epsilon^2_{\mathsf{prf}}} + 2\exp(-n^2/32) \tag{14}$$

Note that the term on the right side drops exponentially in n, thus this restriction is a very weak one, and is e.g. satisfied for any $\epsilon_{\mathsf{prf}} \geq n \cdot 2^{-n/3}$ and $n \geq 20$.

Lemma 6. *Let* $\mathsf{F} : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^m$ *be a* $(\epsilon_{\mathsf{prf}}, s_{\mathsf{prf}}, n/\epsilon_{\mathsf{prf}})$*-secure wPRF. Let* $K \in \{0,1\}^\kappa$ *and* $X \in \{0,1\}^n$ *be independent where* $H_\infty(K) = \kappa - 2\lambda$ *and* $H_\infty(X) = n - 2\lambda$ *and let* $f : \{0,1\}^{\kappa+n} \to \{0,1\}^\lambda$ *be any leakage function, then for large enough* n *(as just described) and* $\lambda \le \log(\epsilon_{\mathsf{prf}}^{-1})/6$

$$\Pr_{X,Y}[\mathbf{H}^{\mathsf{HILL}}_{\epsilon',s'}(\mathsf{F}(K,X)|X, f(K,X)) \ge m - 2\lambda] \ge 1 - 2^{-\lambda/2+1}$$

with $\epsilon' = 2^{-\lambda/2+2}$ *and* $s' = s_{\mathsf{prf}}/2^{\lambda+3}(n+\kappa)^3$.

Proof. We set $\varDelta := 2\lambda$ and $\epsilon = \epsilon_{\mathsf{HILL}} := 2^{-\lambda/2+1}$, and require that

$$\lambda \le 2 + \log(\epsilon_{\mathsf{prg}}^{-1})/2 \tag{15}$$

so that it satisfies the condition $\epsilon_{\mathsf{prg}} \le \frac{\epsilon^2}{2^\lambda} - 2^{-\varDelta}$ from Lemma 4, where now we can write (13) as

$$\Pr_{Z \xleftarrow{*} \{0,1\}^n}[\mathbf{H}^{\mathsf{HILL}}_{2^{-\lambda/2+2},\hat{s}}(\mathsf{prg}(Z)|f(Z)) \ge m - 2\lambda] \ge 1 - 2^{-\lambda/2+1} \tag{16}$$

where $\hat{s} = s_{\mathsf{prg}}/2^{\lambda+1}(n+\kappa)$. Now consider the wPRF F from the statement of the lemma, first we apply Lemma 2 with $t = n/\epsilon_{\mathsf{prf}}$ and $q_{\mathsf{prf}} = t$ to get for a uniformly random X' (in the second step below we use eq.(14)).

$$\delta_{s_{\mathsf{prf}}\epsilon_{\mathsf{prf}}/n}(\mathsf{F}(K,X')\|X' \ ; \ U_m\|X') \le$$
$$\epsilon_{\mathsf{prf}} \cdot 2^{\varDelta+1} + 2^{\varDelta+1}\left(n^2/2^{n+1} \cdot \epsilon_{\mathsf{prf}}^2 + 2\exp\left(-n^2/8\right)\right) \le \epsilon_{\mathsf{prf}} \cdot 2^{\varDelta+2}$$

Thus F is a $(s_{\mathsf{prf}}\epsilon_{\mathsf{prf}}/n, \epsilon_{\mathsf{prf}} \cdot 2^{\varDelta+1}, 1)$ secure wPRF even if we use a non-uniform key K. Now we apply Lemma 3 (again with $t = n/\epsilon_{\mathsf{prf}}$ and using eq.(14) in the second step)

$$\delta_{s_{\mathsf{prf}}\epsilon_{\mathsf{prf}}^2/2n^2}(\mathsf{F}(K,X)\|X \ ; \ U_m\|X) \le$$
$$\epsilon_{\mathsf{prf}} \cdot 2^{2\varDelta+3} + 2^{\varDelta+1} \cdot 2 \cdot \exp(-n^2/32) \le \epsilon_{\mathsf{prf}} \cdot 2^{2\varDelta+4}$$

Thus we can see F on input K, X as an $(\epsilon_{\mathsf{prg}}, s_{\mathsf{prg}})$-secure pseudorandom generator where $s_{\mathsf{prg}} = s_{\mathsf{prf}}\epsilon_{\mathsf{prf}}^2/2n^2$ and $\epsilon_{\mathsf{prg}} = \epsilon_{\mathsf{prf}} \cdot 2^{2\varDelta+4}$ (note that eq.(15) is still satisfied as in the statement of the lemma we require $\lambda \le \log(\epsilon_{\mathsf{prf}}^{-1})/6$).

Now consider any function $f : \{0,1\}^{\kappa+n} \to \{0,1\}^\lambda$, by (16)

$$\Pr_{K,X}[\mathbf{H}^{\mathsf{HILL}}_{\epsilon',s'}(\mathsf{F}(K,X)|f(K,X),X) \ge m - 2\lambda] \ge 1 - 2^{-\lambda/2+1}$$

with $\epsilon' = 2^{-\lambda/2+2}$ and $s' = s_{\mathsf{prg}}/2^{\lambda+1}(n+\kappa) > s_{\mathsf{prf}}\epsilon_{\mathsf{prf}}^2/2^{\lambda+2}(n+\kappa)^3$. ■

The following lemma quantifies the security loss in one round of our stream cipher. Let $\mathtt{size}_i$ denote the size of the circuit realizing the ith round of the experiment $\mathsf{S}^{\mathsf{F}} \overset{\ell}{\rightsquigarrow} \mathsf{A}$, then $\sum_{i=1}^{\ell} \mathtt{size}_i = \mathtt{size}(\mathsf{S} \overset{\ell}{\rightsquigarrow} \mathsf{A})$.

Lemma 7 (The ith round). *Consider the random experiment* $\mathsf{S}^\mathsf{F} \overset{\ell}{\rightsquigarrow} \mathsf{A}$. *Then if before round* $i \leq \ell$ *for some* $s_{i-1} \leq s'$ *(with* s', ϵ', λ *are as in the previous lemma)*

$$\mathbf{H}^{\mathsf{HILL}}_{\epsilon_{i-1},s_{i-1}}(K_{i-1}|\mathtt{view}_{i-1}) \geq \kappa - 2\lambda$$
$$\mathbf{H}^{\mathsf{HILL}}_{\epsilon_{i-1},s_{i-1}}(X_{i-1}|\mathtt{view}^-_{i-1}) \geq n - 2\lambda$$

then with probability $1 - 2^{-\lambda/2+1}$ *the output* $(K_{i+1}, X_i) = \mathsf{F}(K_{i-1}, X_{i-1})$ *satisfies*

$$\mathbf{H}^{\mathsf{HILL}}_{\epsilon_i,s_i}(\mathsf{F}(K_{i-1}, X_{i-1})|\mathtt{view}^-_i) \geq \kappa + n - 2\lambda$$

where $\epsilon_i = \epsilon_{i-1} + \epsilon'$, $s_i = s_{i-1} + \mathtt{size}_i$

Proof. Consider random variables K'_{i-1}, X'_{i-1} which have high *min-entropy*

$$H_\infty(K'_{i-1}|\mathtt{view}_{i-1}) \geq \kappa - \lambda \quad \text{and} \quad H_\infty(X'_{i-1}|\mathtt{view}^-_{i-1}) \geq n - \lambda$$

By Lemma 6 with probability at least $1 - 2^{-\lambda/2+1}$

$$\mathbf{H}^{\mathsf{HILL}}_{\epsilon',s'}(\mathsf{F}(K'_{i-1}, X'_{i-1})|\mathtt{view}^-_i) \geq \kappa + n - 2\lambda$$

holds with $\epsilon' = 2^{-\lambda/2+2}$ and $s' = \frac{s_{\mathsf{prf}}}{2^{\lambda+3}\cdot(n+\kappa)^3}$. If we now use the random variables K_{i-1}, X_{i-1} (only having high HILL-pseudoentropy) instead of K'_{i-1}, X'_{i-1}, we get (recall that $s_{i-1} < s'$)

$$\mathbf{H}^{\mathsf{HILL}}_{\epsilon'+\epsilon_{i-1},s_{i-1}-\mathtt{size}_i}(\mathsf{F}(K_{i-1}, X_{i-1})|\mathtt{view}^-_i) \geq \kappa + n - 2\lambda$$

Let us stress that here the new error ϵ_i is $\epsilon' + \epsilon_{i-1}$, and not $\epsilon' + 2\epsilon_{i-1}$, as one would think because we must add an error term of ϵ_{i-1} for K_{i-1} and X_{i-1} respectively. Such a weaker bound would render the lemma useless, as then ϵ_i would grow exponentially in i. The reason we only have to add ϵ_{i-1}, is that in round $i-1$, F outputs (X_{i-1}, K_i), and it's this tuple that cannot be distinguished with advantage more than ϵ_{i-1}. Thus by adding an error ϵ_{i-1} for X_{i-1} in round i, we also account for K_i to be used in the next round, and we won't have to add an extra error term there. ■

The bound on the security of S^F as stated in Theorem 2 now follows by summing up the security decrease in each round as stated in the previous lemma. To apply the lemma, one must show that for each i, the K_i and X_i are independent given the view of the adversary, this follows from Lemma 5 by identifying A_i (from the Lemma) with $K_{2(i-1)}$ (as computed by S^F), identifying B_i with $K_{2(i-1)+1}$ and V_i with $\mathtt{view}_i$. In particular, after ℓ round, the error adds up to

$$\mathsf{AdvInd}(\mathsf{D}, \mathsf{A}, \mathsf{S}, \ell) \leq \ell \cdot 2^{-\lambda/2+3}.$$

Note that this is a bit strange, as the advantage *decreases* by increasing the leakage λ, but this is only due to the fact that we explicitly set the error parameters ϵ and ϵ_{HILL} as functions of λ in the proof of Lemma 6 in order to keep the number of parameters down. Setting $\lambda = \log(\epsilon^{-1}_{\mathsf{prf}})/6$ (note that this is the largest value allowed in the statement of Lemma 6), we get the bound as claimed in the theorem.

Acknowledgements

I'd like to thank Stefan Dziembowski and Sebastian Faust for some interesting discussions and the authors of [36, 41] for clarifying their work.

References

1. Akavia, A., Goldwasser, S., Vaikuntanathan, V.: Simultaneous hardcore bits and cryptography against memory attacks. In: TCC (2009)
2. Barak, B., Shaltiel, R., Wigderson, A.: Computational analogues of entropy. In: RANDOM-APPROX, pp. 200–215 (2003)
3. Biham, E., Shamir, A.: Differential fault analysis of secret key cryptosystems. In: Kaliski Jr., B.S. (ed.) CRYPTO 1997. LNCS, vol. 1294, pp. 513–525. Springer, Heidelberg (1997)
4. Boneh, D., DeMillo, R.A., Lipton, R.J.: On the importance of checking cryptographic protocols for faults. In: Fumy, W. (ed.) EUROCRYPT 1997. LNCS, vol. 1233, pp. 37–51. Springer, Heidelberg (1997)
5. Canetti, R., Eiger, D., Goldwasser, S., Lim, D.-Y.: How to protect yourself without perfect shredding. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfsdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part II. LNCS, vol. 5126, pp. 511–523. Springer, Heidelberg (2008)
6. Cash, D.M., Ding, Y.Z., Dodis, Y., Lee, W., Lipton, R.J., Walfish, S.: Intrusion-resilient key exchange in the bounded retrieval model. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 479–498. Springer, Heidelberg (2007)
7. Chari, S., Rao, J.R., Rohatgi, P.: Template attacks. In: Kaliski Jr., B.S., Koç, Ç.K., Paar, C. (eds.) CHES 2002. LNCS, vol. 2523, pp. 13–28. Springer, Heidelberg (2003)
8. Coron, J.-S.: Resistance against differential power analysis for elliptic curve cryptosystems. In: Koç, Ç.K., Paar, C. (eds.) CHES 1999. LNCS, vol. 1717, p. 292. Springer, Heidelberg (1999)
9. Courtois, N.T., Bard, G.V., Wagner, D.: Algebraic and slide attacks on keeLoq. In: Nyberg, K. (ed.) FSE 2008. LNCS, vol. 5086, pp. 97–115. Springer, Heidelberg (2008)
10. Di Crescenzo, G., Lipton, R.J., Walfish, S.: Perfectly secure password protocols in the bounded retrieval model. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 225–244. Springer, Heidelberg (2006)
11. Dodis, Y., Sahai, A., Smith, A.: On perfect and adaptive security in exposure-resilient cryptography. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 301–324. Springer, Heidelberg (2001)
12. Dodis, Y., Wichs, D.: One-round authenticated key agreement from weak secrets. Cryptology ePrint Archive, Report 2008/503 (2008), http://eprint.iacr.org/
13. Dziembowski, S.: Intrusion-resilience via the bounded-storage model. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 207–224. Springer, Heidelberg (2006)
14. Dziembowski, S.: On forward-secure storage (extended abstract). In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 251–270. Springer, Heidelberg (2006)
15. Dziembowski, S., Maurer, U.M.: On generating the initial key in the bounded-storage model. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 126–137. Springer, Heidelberg (2004)

16. Dziembowski, S., Pietrzak, K.: Intrusion-resilient secret sharing. In: FOCS, pp. 227–237 (2007)
17. Dziembowski, S., Pietrzak, K.: Leakage-resilient cryptography. In: FOCS (2008)
18. Eisenbarth, T., Kasper, T., Moradi, A., Paar, C., Salmasizadeh, M., Shalmani, M.T.M.: On the power of power analysis in the real world: A complete break of the KEELOQ code hopping scheme. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 203–220. Springer, Heidelberg (2008)
19. Gandolfi, K., Mourtel, C., Olivier, F.: Electromagnetic analysis: Concrete results. In: CHES, pp. 251–261 (2001)
20. Goldreich, O.: A uniform-complexity treatment of encryption and zero-knowledge. Journal of Cryptology 6(1), 21–53 (1993)
21. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions. In: FOCS, pp. 464–479 (1984)
22. Goldwasser, S., Kalai, Y.T., Rothblum, G.N.: One-time programs. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 39–56. Springer, Heidelberg (2008)
23. Halderman, J.A., Schoen, S.D., Heninger, N., Clarkson, W., Paul, W., Calandrino, J.A., Feldman, A.J., Appelbaum, J., Felten, E.W.: Lest we remember: Cold boot attacks on encryption keys. In: USENIX Security Symposium, pp. 45–60 (2008)
24. Halevi, S., Myers, S., Rackoff, C.: On seed-incompressible functions. In: Canetti, R. (ed.) TCC 2008. LNCS, vol. 4948, pp. 19–36. Springer, Heidelberg (2008)
25. Håstad, J., Impagliazzo, R., Levin, L.A., Luby, M.: A pseudorandom generator from any one-way function. SIAM Journal on Computing 28(4), 1364–1396 (1999)
26. Hoeffding, W.: Probability inequalities for sums of bounded random variables. Journal of the American Statistical Association 58(301), 13–30 (1963)
27. Indesteege, S., Keller, N., Dunkelman, O., Biham, E., Preneel, B.: A practical attack on KeeLoq. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 1–18. Springer, Heidelberg (2001)
28. Ishai, Y., Prabhakaran, M., Sahai, A., Wagner, D.: Private circuits II: Keeping secrets in tamperable circuits. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 308–327. Springer, Heidelberg (2006)
29. Ishai, Y., Sahai, A., Wagner, D.: Private circuits: Securing hardware against probing attacks. In: Boneh, D. (ed.) CRYPTO 2003. LNCS, vol. 2729, pp. 463–481. Springer, Heidelberg (2003)
30. Kelsey, J., Schneier, B., Wagner, D., Hall, C.: Side channel cryptanalysis of product ciphers. In: Quisquater, J.-J., Deswarte, Y., Meadows, C., Gollmann, D. (eds.) ESORICS 1998. LNCS, vol. 1485, pp. 97–110. Springer, Heidelberg (1998)
31. Kocher, P.C.: Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In: Koblitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 104–113. Springer, Heidelberg (1996)
32. Kocher, P.C.: Design and validation strategies for obtaining assurance in countermeasures to power analysis and related attacks. In: Proceedings of the NIST Physical Security Workshop (2005)
33. Kocher, P.C., Jaffe, J., Jun, B.: Differential Power Analysis. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 388–397. Springer, Heidelberg (1999)
34. Maurer, U.M.: A provably-secure strongly-randomized cipher. In: Damgård, I.B. (ed.) EUROCRYPT 1990. LNCS, vol. 473, pp. 361–373. Springer, Heidelberg (1991)
35. Micali, S., Reyzin, L.: Physically observable cryptography (extended abstract). In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 278–296. Springer, Heidelberg (2004)

36. Petit, C., Standaert, F.-X., Pereira, O., Malkin, T., Yung, M.: A block cipher based pseudo random number generator secure against side-channel key recovery. In: ASIACCS, pp. 56–65 (2008)
37. Pietrzak, K.: Full version of this paper, http://homepages.cwi.nl/ pietrzak/publications.html
38. Pietrzak, K., Sjödin, J.: Range extension for weak pRFs; the good, the bad, and the ugly. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 517–533. Springer, Heidelberg (2007)
39. Quisquater, J.-J., Samyde, D.: Electromagnetic analysis (ema): Measures and counter-measures for smart cards. In: E-smart, pp. 200–210 (2001)
40. Reingold, O., Trevisan, L., Tulsiani, M., Vadhan, S.P.: Dense subsets of pseudorandom sets. In: FOCS, pp. 76–85 (2008)
41. Standaert, F.-X., Malkin, T., Yung, M.: A unified framework for the analysis of side-channel key recovery attacks. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 443–461. Springer, Heidelberg (2009)
42. Vadhan, S.P.: Constructing locally computable extractors and cryptosystems in the bounded-storage model. Journal of Cryptology 17(1), 43–77 (2004)