Practical Leakage-Resilient Pseudorandom Objects with Minimum Public Randomness

Yu Yu[1,2] and François-Xavier Standaert[3]

[1] Tsinghua University, Institute for Interdisciplinary Information Sciences, China
[2] East China Normal University, Department of Computer Science, China
[3] Université catholique de Louvain, ICTEAM/ELEN/Crypto Group, Belgium

Abstract. One of the main challenges in leakage-resilient cryptography is to obtain proofs of security against side-channel attacks, under realistic assumptions and for efficient constructions. In a recent work from CHES 2012, Faust et al. proposed new designs of stream ciphers and pseudorandom functions for this purpose. Yet, a remaining limitation of these constructions is that they require large amounts of public randomness to be proven leakage-resilient. In this paper, we show that tweaked designs with minimum randomness requirements can be proven leakage-resilient in minicrypt. That is, either these constructions are secure, or we are able to construct public-key cryptographic primitives from symmetric-key building blocks and their leakage functions (which is very unlikely). Hence, our results improve the practical relevance of two important leakage-resilient pseudorandom objects.

1 Introduction

Side-channel attacks are an important threat to the security of embedded devices like smart cards and RFID tags. Following the first publications on Differential Power Analysis [19] (DPA) and Electro-Magnetic Analysis [12,29] (EMA), a large body of work has investigated techniques to improve the security of cryptographic implementations. During the first ten years after the publication of these attacks, the solutions proposed were mainly taking advantage of hardware/software modifications. For example, it as been proposed to exploit new circuit technologies or to randomize the time and data in the implementations (see [3,4,36] for early proposals of these ideas, and many improvements and analyzes published at CHES). In general, these countermeasures are successful in the sense that they indeed reduce the amount of information leakage. Yet, security evaluations considering worst-case (profiled) side-channel attacks such as [33] usually reveal that reaching high security levels is expensive and highly dependent of physical assumptions. Taking the example of secret sharing (aka masking), multiple shares are required for this purpose (i.e. so-called higher-order security [34]). However, the implementation cost of higher-order masking schemes is significant [31], and the risk of physical effects leading to exploitable weaknesses (such as glitches [21]) leads to additional design constraints.

E. Dawson (Ed.): RSA 2013, LNCS 7779, pp. 223–238, 2013.

Motivated by the great challenges in physical security, recent works have also considered the possibility to analyze the effectiveness of countermeasures against side-channel attacks in a more formal way, and to design new primitives (aimed to be) inherently more secure against such attacks. Taking the case of symmetric cryptography building blocks (that are important primitives to design as they are usual targets of DPA attacks [20]), a variety of models have been introduced for this purpose, ranging from specialized to general. For example, a PRNG secure against side-channel key recovery attacks was proposed at ASIACCS 2008 by Petit et al. [25], and analyzed in front of a class of (realistic yet specific) leakage functions. Following, a construction of leakage-resilient stream cipher has been presented by Dziembowski and Pietrzak at FOCS 2008, together with a proof of security in the standard model [9]. Quite naturally, such "physical security proofs" raise a number of concerns regarding their relevance to practice, a topic that has been intensively discussed over the last couple of years. In particular, one of the fundamental issues raised by leakage-resilient cryptography is to determine reasonable restrictions of the leakage function, e.g. in terms of informativeness and computational power. As far as computational power is concerned (which will be our main concern in this paper), an appealing solution is to consider the leakage function to be polynomial time computable, as initially proposed by Micali and Reyzin [24], and leading to contrasted observations. On the one hand, polynomial time functions are significantly more powerful than actual leakage functions. For example, they allow so called "precomputation attacks" (aka future computation attacks) that are arguably unrealistic in practice [35]. On the other hand, meaningful alternatives seem quite challenging to specify. Furthermore, given that one obtains proofs of security under such strong leakages without paying too large implementation overheads, polynomial time functions remain a useful abstraction.

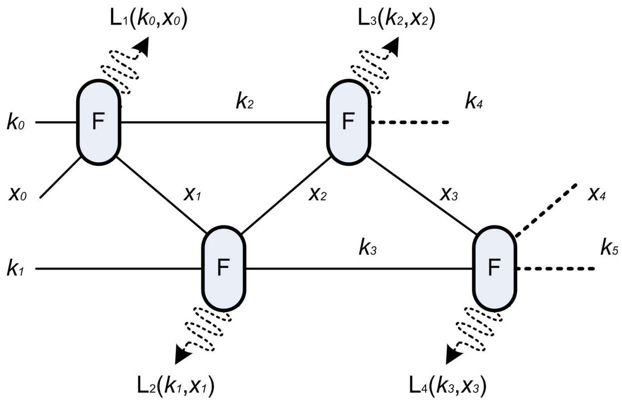

Fig. 1. The Eurocrypt 2009 stream cipher

In this context, one of the design tweaks used by Dziembowsky and Pietrzak is the so-called "alternating structure". Figure 1 depicts such an alternating structure for a simplified stream cipher proposed by Pietrzak at Eurocrypt 2009 [27], that can be instantiated only from (AES-based) weak Pseudo-Random Functions (wPRFs)[1]. If one assumes that the two branches of such an alternating structure leak independently, no leakage occuring in one of the branches can be used to compute bits that will be manipulated in future computations of the other branch, hence ruling out the possibility of precomputation attacks. The main drawback of this proposal is that a key bit-size of $2n$ can only guarantee a security of at most 2^n. Hence, as it appears unrealistic that a circuit actually leaks about something it will only compute during its future iterations, a following work by Yu et al. investigated the possibility to mitigate the need of an alternating structure [37]. In a paper from CCS 2010, they first proposed to design a "natural" (i.e. conform to engineering intuition) leakage-resilient stream cipher, which could only be proven secure under a (non-standard) random oracle based assumption. Next, they proposed a variant of the FOCS 2008 (and Eurocrypt 2009) designs, replacing the alternating structure by alternating public randomness, and under the additional (necessary) assumption that the leakage function is non adaptive. Eventually, in a recent work of CHES 2012, Faust et al. showed that large amounts of public randomness (i.e. linear in the number of stream cipher iterations) were actually required for the proof of Yu et al. to hold [10]. While it remains an open question to determine whether the exact construction proposed in [37] (using only two alternating public values) can be proven secure or attacked in a practical setting, this last result reveals a tension between the proof requirements and how the best known side-channel attacks actually proceed against leakage-resilient constructions [23].

Considering the previous observations, this paper tackles the fundamental question of how much public randomness is actually needed to obtain proofs of leakage-resilience in symmetric cryptography. For this purpose, we investigate (yet another) variant of stream cipher, where only a single public random value is picked up prior to (independent of) the selection of the leakage functions, and then expanded thanks to a PRNG. Quite naturally, a strong requirement for this approach to be interesting is that the seed of the PRNG should *not* be secret (or we would need a leakage-resilient PRNG to process it, i.e. essentially the problem we are trying to solve). Surprisingly, we show that this approach can be proven secure in minicrypt [17] (i.e. the hypothetical world introduced by Impagliazzo, where one-way functions exist, but public-key cryptography does not). More precisely, using the technique of [1] (see also similar ones in earlier literature [7,8,26,28]), we show that either the proposed solution is leakage-resilient,

[1] Besides their possible implementation costs, additional components in leakage-resilient constructions can also become a better target for a side-channel adversary, e.g. as discussed with the case of randomness extractors in the FOCS 2008 stream cipher [22,32]. In this respect, relying only on AES-based primitives (for which the security against side-channel attacks has been carefully analyzed) is an interesting feature of the Eurocrypt 2009 proposal in Figure 1.

or we are able to construct black-box constructions of public-key encryption schemes from symmetric primitives and their leakage functions. When using block ciphers such as the AES to instantiate the stream cipher, the latter is very unlikely due to known separation results between one-way functions and PKE [18]. We then conclude this work by illustrating that this observation also applies to PRFs for which various designs were already proposed [5,10,23,35].

Summarizing, proofs of leakage-resilience require to restrict the leakage function both in terms of informativeness and computing power. As finding useful and realistic restrictions is hard with state-of-the-art techniques, we consider an alternative approach, trying to limit the implementation overheads due to unrealistic models. Admittedly, our analysis is based on the same assumptions as the previously mentioned works (i.e. polynomial time, bounded and non-adaptive leakage functions). The quest for more realistic models remains a very important research direction. Meanwhile, we believe that our intermediate conclusion is important, as it highlights that leakage-resilient (symmetric) cryptography can be obtained with minimum public randomness (i.e. the public seed of a PRNG).

2 Background

2.1 The CCS 2010 Stream Cipher

The CCS 2010 construction, depicted in Figure 2, is based on the observation from the practice of side-channel attacks that leakage functions are more a property of the target device and measurement equipment than something that is adaptively chosen by the adversary. It therefore considers a weaker security model, in which the polynomial time (and bounded) leakage functions are fixed before the stream cipher execution starts. By considering those non-adaptively chosen leakage functions, the construction can be made more efficient and easier to implement in a secure way. This stream cipher is initialized with a secret key k_0 and two values p_0 and p_1 that can be public. Those two values are then used in an alternating way: at round i, one computes k_i and x_i by applying the

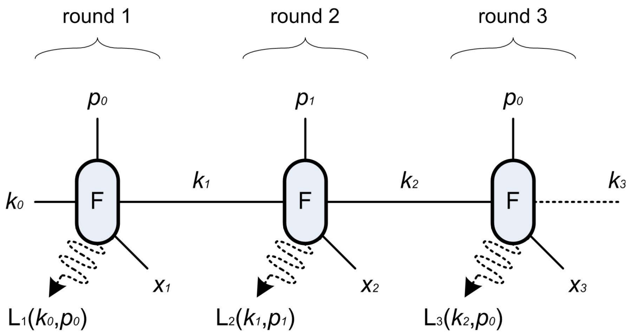

Fig. 2. The CCS 2010 stream cipher

wPRF to inputs k_{i-1} and $p_{i-1 \bmod 2}$. Thanks to the removal of the alternating structure, the complexity of a brute-force attack on this construction becomes directly related to the full length of the key material, which is now exploited in each round.

2.2 The CHES 2012 Stream Cipher

In a paper from CHES 2012, Faust et al. observed that the technical tools used to prove the CCS 2010 construction actually require to use independent public values in all the stream cipher rounds (rather than only two alternating ones). Therefore, only the slightly modified the construction suggested in Figure 3, assuming a common random string $p_0, p_1, p_2, \ldots$, can be proven secure with these tools. The practical advantages of this construction compared to the FOCS 2008 / Eurocrypt 2009 ones naturally become more contrasted. On the positive side, the fact that the values $p_0, p_1, p_2, \ldots$ are public can still make it easier to ensure that rounds leak independently of each other (which is implicitly required by the arguments of the leakage function): for example, a number of public p_i's can be stored in non-volatile memories for this purpose. On the other hand, this amount of public randomness required is linear in the number of stream cipher rounds, which is hardly realistic (hence leading the authors of [10] to pay more attention to leakage-resilient PRFs for which this penalty is less damaging - see Section 4 for a brief discussion).

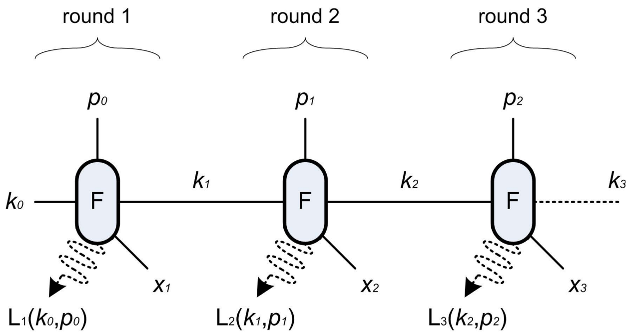

Fig. 3. The CHES 2012 stream cipher

3 Natural PRNG with Minimum Public Randomness

3.1 A New Proposal

As mentioned in introduction, it in unclear whether the need of large public randomness in leakage-resilient stream ciphers is due to proof artifacts or if the lack of such randomness can be exploited in realistic side-channel attacks. This

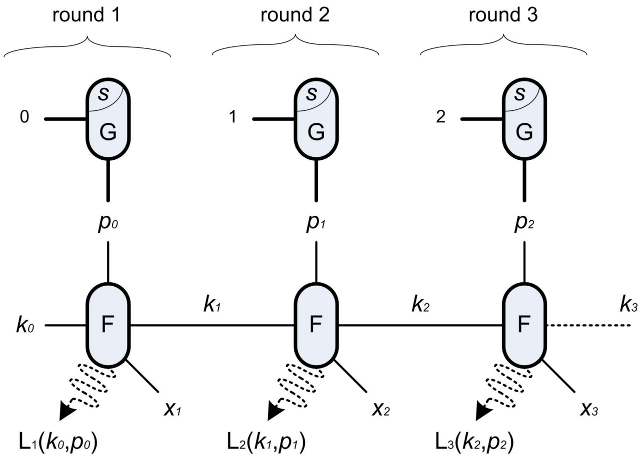

Fig. 4. Leakage-resilient stream cipher with minimum randomness

question is important as such attacks would most likely reveal an issue in the most natural construction of [37], where no public randomness is used at all and the proof is based on a random oracle assumption. In order to answer it, we propose an alternative stream cipher depicted in Figure 4.

The Proposed Stream Cipher. We denote our stream cipher with SC, let n be the security parameter, and (k_0,s) be the initial state of SC, where $k_0 \in \{0,1\}^n$ is a secret key and $s \in \{0,1\}^n$ a public seed, both randomly chosen. SC expands s into $p_0, p_1, p_2, \ldots$ on-the-fly by running a PRF $\mathsf{G} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$ in counter mode[2], i.e., $p_i := \mathsf{G}(s,i)$. Then, SC uses the generated public strings $p_0, p_1, p_2, \ldots$ to randomize another PRF $\mathsf{F} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^{2n}$, which updates the secret state k_i and produces the output x_i, i.e. $(k_i, x_i) := \mathsf{F}(k_{i-1}, p_{i-1})$. That is, the stream cipher SC in Figure 4 is essentially similar to the previous ones, excepted that any public string p_i is obtained by running a PRF on a counter value, using the public seed s.

Instantiation and Efficiency. Following [27], we instantiate F and G with a block cipher $\mathsf{BC} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^n$, e.g. the AES. As will be shown in Lemma 4, it is sufficient to produce $\log(1/\varepsilon)$ bits of fresh pseudo-randomness for every p_i (and pad the rest with zero's), with ε a security parameter of the PRF F (see Definition 1). This further improves efficiency, as we only need to run G once every $\lfloor n/\log(1/\varepsilon) \rfloor$ iterations of F.

Leakage Models of the CCS 2010/CHES 2012 stream ciphers. For every i^{th} iteration, let $\mathsf{L}_i : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^\lambda$ be a function (on k_{i-1}

[2] Alternatively, we can also expand s by iterating a length-doubling PRNG in a forward-secure way, but this would lead to less efficient designs and is not needed (since s is public).

and p_{i-1}) that outputs the leakage incurred during the computation of F on (k_{i-1},p_{i-1}). The CCS 2010/CHES 2012 constructions model the leakages as follows [10,37]:

1. (*Efficient computability*). L_i can be computed by polynomial-size circuits.
2. (*Bounded leakage per iteration*). The leakage function has bounded range given by $\lambda \in O(\log(1/\varepsilon))$, where ε is a security parameter of the PRF F (see Definition 1).
3. (*Non-adaptivity*). The selection of the leakage functions L_i is made prior to (or independent of) s, and thus only depends on k_{i-1} and p_{i-1}.

Note that strictly speaking, the leakage models needed to prove the security of the CCS 2012 and CHES 2012 stream ciphers are not exactly equivalent. Namely, the CHES 2012 stream cipher can further tolerate that each L_i not only depends on the current state (k_{i-1},p_{i-1}), but also on the past transcript $\mathsf{T}_{i-1} \stackrel{\mathsf{def}}{=} (x_1, \cdots, x_{i-1}, p_0, \cdots, p_{i-2}, \mathsf{L}_1(k_0, p_0), \cdots, \mathsf{L}_{i-1}(k_{i-2}, p_{i-2}))$. This is naturally impossible if only two p_i's are used.

Leakage models of FOCS 2008/Eurocrypt 2009 stream ciphers. The FOCS 2008/Eurocrypt 2009 constructions consider a model similar to the above one, but they do not require condition #3 and allow the adaptive selection of the leakage functions. That is, at the beginning of each round, the adversary adaptively chooses a function L_i based on his current view. As previously mentioned, this leads to unrealistic attacks as the adversary can simply recover a future secret state, say k_{100}, by letting each L_i leak some different λ bits about it. The authors of [9,27] deal with this issue by tweaking their stream cipher design with an alternating structure (as in Figure 1).

In the next sections, we will prove the leakage-resilient security of our stream cipher in the (non-adaptive) model from CCS 2010/CHES 2012. More precisely, we will also consider its less restrictive version where the leakage functions can depend on the past transcript. Yet, for brevity, we will not explicitly put T_{i-1} as an input of each L_i, as an adversary can hardwire them into L_i. Note also that we do not need to model leakages on G since the seed s (from which all p_0, $\cdots$, p_i can be efficiently computed) is public.

3.2 Security Analysis

Notations and Definitions. For security parameter n, a function $\mathsf{negl} : \mathbb{N} \to [0,1]$ is *negligible* if for any $c > 0$ there is a n_0 such that $\mathsf{negl}(n) \leq 1/n^c$ for all $n \geq n_0$. We use uppercase letters (e.g. X) to denote a random variable and lowercase letters (e.g. x) to denote a specific value, with exceptions being n, t and q which are reserved for security parameter, circuit-size (or running time) and query complexity, respectively. We write $x \leftarrow X$ to denote the sampling of a random x according to X. We use U_n to denote the uniform distribution over $\{0,1\}^n$. For function f, we denote its circuit-size complexity by $\mathsf{size}(f)$ or t_f. We denote with $\Delta_D(X,Y)$ the advantage of a circuit D in distinguishing the random variables X, Y: $\Delta_{\mathsf{D}}(X,Y) \overset{\text{def}}{=} |\ \Pr[\mathsf{D}(X)=1] - \Pr[\mathsf{D}(Y)=1]\ |$. The *computational distance* between two random variables X, Y is defined with $\mathsf{CD}_t(X,Y) \overset{\text{def}}{=} \max_{\mathsf{size(D)} \le t} \Delta_{\mathsf{D}}(X,Y)$, which takes the maximum over all distinguishers D of size t. For convenience, we use $\mathsf{CD}_t(X,Y|Z)$ as shorthand for $\mathsf{CD}_t((X,Z),(Y,Z))$. The min-entropy of X is defined as $\mathbf{H}_\infty(X) \overset{\text{def}}{=} -\log(\max_x \Pr[X=x])$. We finally define average (aka conditional) min-entropy of a random variable X conditioned on Z as:

$$\widetilde{\mathbf{H}}_\infty(X|Z) \overset{\text{def}}{=} -\log\left(\ \mathbb{E}_{z\leftarrow Z}\left[\ \max_x \Pr[X=x|Z=z]\ \right]\ \right),$$

where $\mathbb{E}_{z\leftarrow Z}$ denotes the expected value computed over all $z \leftarrow Z$.

Standard Security Notions. Indistinguishability requires that no efficient adversary is able to distinguish a real distribution from an idealized one (e.g. uniform randomness) with non-negligible advantage. In this paper, we will work in the concrete non-uniform setting[3]. Yet, we note that the proof can be made uniform by adapting the technique from [2,38] (see [9] for a discussion). Given this precision, a standard PRF is defined as:

Definition 1 (PRF). $\mathsf{F} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^m$ *is a pseudorandom function (PRF) if for all polynomial-size distinguisher* D *making up to any polynomial number of queries, we have:*

$$|\Pr[\mathsf{D}^{\mathsf{F}(k,\cdot)} = 1 \mid k \leftarrow U_n] - \Pr[\mathsf{D}^{\mathsf{R}(\cdot)} = 1]\ | \ \le\ \mathsf{negl}(n),$$

where R *is a random function uniformly drawn from function family* $\{\{0,1\}^n \to \{0,1\}^m\}$. *Furthermore, we say that* F *is a (t,q,ε)-secure PRF if for all distinguishers* D *of size t making q queries, the above advantage is bounded by ε.*

Security without Leakages. Without considering side-channel adversaries, the security of SC is easily proven using a standard hybrid argument, by considering F (on any fixed input) as a PRG, and without any security requirement about G (which could just output any constant). This is formalized by the following theorem:

Theorem 1 (Security without Leakages). *If* F *is a* $(t,1,\varepsilon)$*-secure PRF, then* SC *is* (t',ℓ,ε')*-secure, i.e.* $\mathsf{CD}_{t'}((X_1,X_2,\cdots,X_\ell),U_{n\ell}|S) \le \varepsilon'$, *with* $t' \approx t - \ell{\cdot}t_{\mathsf{F}}$ *and* $\varepsilon' \le \ell \cdot \varepsilon$.

Leakage-Resilient Security. We first observe that as soon as some leakage is given to the adversary, he can easily exploit it to distinguish x_i from

[3] An efficient uniform adversary can be considered as a Turing-machine which on input 1^n (security parameter in unary) terminates in time polynomial in n, whereas its non-uniform counterpart will, for each n, additionally get some polynomial-length advice.

uniform randomness (e.g. $\mathsf{L}_i(k_{i-1}, p_{i-1})$ leaking the first bit of x_i is enough for this purpose). Thus, all previous approaches in leakage-resilient cryptography require that any (computationally bounded) adversary observing the leakages for as many rounds as he wishes should not be able to distinguish the next x_ℓ without seeing $\mathsf{L}_\ell(k_{\ell-1}, p_{\ell-1})$ [9,10,27,37]. Formally, let:

$$\mathsf{view}_\ell(\mathsf{A}, \mathsf{SC}, K_0, S) \stackrel{\text{def}}{=} (S, X_1, \cdots, X_{\ell-1}, \mathsf{L}_1(K_0, P_0), \cdots, \mathsf{L}_{\ell-1}(K_{\ell-2}, P_{\ell-2})) \quad (1)$$

denote the view of adversary A after attacking SC (initialized with K_0 and S) for ℓ rounds, for which we use shorthand view_ℓ in the rest of the paper. Given a distinguisher D, we then define its indistinguishability advantage (on uniform K_0 and S) as:

$$\mathsf{AdvInd}(\mathsf{SC}, \mathsf{A}, \mathsf{D}, \ell) \stackrel{\text{def}}{=} |\Pr_{K_0,S}[\mathsf{D}(\mathsf{view}_\ell, X_\ell) = 1] - \Pr_{K_0,S}[\mathsf{D}(\mathsf{view}_\ell, U_n) = 1]|.$$

We will use $\mathsf{size}(\mathsf{A}) \stackrel{\text{def}}{=} \ell(t_\mathsf{G} + t_\mathsf{F}) + \sum_{i=1}^{\ell-1} t_{\mathsf{L}_i}$ to denote the circuit-size complexity of the physical implementation of SC and $\mathsf{size}(\mathsf{D})$ to denote the circuit-size complexity of D.

Using these notations, our main result can be stated as follows.

Theorem 2 (Leakage-Resilient Security). *If* F *is (t,2,ε)-secure PRF, and* G *is a (t,q,ε)-secure PRF, then for any* $\ell \leq q$*, adversary* A*, distinguisher* D *with* $(\mathsf{size}(\mathsf{A}) + \mathsf{size}(\mathsf{D})) \in \Omega(2^{3\lambda}\varepsilon \cdot t/n)$ *and for any leakage size (per round)* λ*, we have that either:*

$$\mathsf{AdvInd}(\mathsf{SC}, \mathsf{A}, \mathsf{D}, \ell) \in O(\ell\sqrt{2^{3\lambda} \cdot \varepsilon}),$$

or otherwise there exist efficient black-box constructions of public key encryption (PKE) from the PRFs F *and* G *and the leakage functions* $\mathsf{L}_1, \cdots, \mathsf{L}_{\ell-1}$*.*

How to Interpret the Result? The above theorem is a typical "win-win" situation, similar to those given in [1,7,8,26,28], where a contradiction to one task gives rise to an efficient protocol for another seemingly unrelated (and sometimes more useful) task. As mentioned in introduction, we know from [18] that black box constructions of PKE from PRFs are very unlikely to exist. Thus, if the building primitives F and G are one-way function equivalents (i.e. they are not PKE primitives), for example using practical block ciphers such as the AES, and the leakage functions are intrinsic to hardware implementation (i.e. not artificially chosen) then the stream cipher SC will be leakage-resilient as stated above. Before giving the proof, we recall the notion of HILL pseudo-entropy:

Definition 2 (HILL Pseudo-entropy [14,16]). *X has at least k bits of HILL pseudo-entropy, denoted by* $\mathbf{H}^{\mathsf{HILL}}_{\varepsilon,t}(X) \geq k$*, if there exists Y so that* $\mathbf{H}_\infty(Y) \geq k$ *and* $\mathsf{CD}_t(X, Y) \leq \varepsilon$*. X has at least k bits of HILL pseudo-entropy conditioned on Z , denoted by* $\mathbf{H}^{\mathsf{HILL}}_{\varepsilon,t}(X|Z) \geq k$*, if there exists* (Y, Z') *such that* $\widetilde{\mathbf{H}}_\infty(Y|Z') \geq k$ *and* $\mathsf{CD}_t((X, Z), (Y, Z')) \leq \varepsilon$*.*

Outline of the Proof. We will present the proof in two main steps. First, we will show the security of our stream cipher when the seed is kept secret. This part of the proof essentially borrows techniques from previously published papers. Next, we will show our main result, i.e. that either leakage-resilience is maintained when S is public, or we have efficient black box constructions of PKE from PRFs as stated in Theorem 2.

Lemma 1 (Security of SC with Secret S). *Let* $P_{[0\cdots\ell-1]} \stackrel{\text{def}}{=} (P_0, \cdots, P_{\ell-1})$. *For the same* F, G, ℓ, A, D *as given in Theorem 2, we have that:*

$$|\Pr_{K_0,S}[\mathsf{D}(\mathsf{view}_\ell \setminus S, P_{[0\cdots\ell-1]}, X_\ell) = 1] - \mathsf{D}(\mathsf{view}_\ell \setminus S, P_{[0\cdots\ell-1]}, U_n) = 1]| \in O(\ell\sqrt{2^{3\lambda}\cdot\varepsilon}).$$

Proof sketch. Since G is a secure PRF and S is leak-free, it suffices to prove the security by replacing every P_i by true randomness P'_i. The conclusion follows from Lemma 2 below, by letting $i = \ell$ and applying computational extractor[4] F on $K_{\ell-1}$ and $P'_{\ell-1}$. It essentially holds because $P'_{\ell-1}$ is independent of all preceding random variables. □

Lemma 2 (The i^{th} round HILL Pseudo-entropy). *Assume that we use uniform randomness* $P'_0, \cdots, P'_{\ell-1}$ *and define the view accordingly as below:*

$$\mathsf{view}'_\ell \stackrel{\text{def}}{=} (P'_0, \cdots, P'_{\ell-1}, X_1, \cdots, X_{\ell-1}, \mathsf{L}_1(K_0, P'_0), \cdots, \mathsf{L}_{\ell-1}(K_{\ell-2}, P'_{\ell-2})). \quad (2)$$

Then we have:

$$\mathbf{H}^{\mathsf{HILL}}_{\varepsilon_i,t_i}(K_{i-1}|\mathsf{view}'_i \setminus P_{i-1}) \geq n - \lambda, \quad (3)$$

where $\varepsilon_i = 2(i-1)\sqrt{2^{3\lambda}\cdot\varepsilon}$ *and* $(t_i + (i-1)t_{\mathsf{F}} + \sum_{j=1}^{i-1} t_{\mathsf{L}_i}) \in \Omega(2^{3\lambda}\varepsilon \cdot t/n)$.

A proof of this Lemma can be found in [10] (and implicitly in [9,27,37]). We will provide an alternative proof with improved bounds in Section 3.3, by utilizing recent technical lemmata from [11] (slightly improving the dense model theorem [9,30]) and Lemma 4 from [6], which explicitly states that a PRF used as computational exactor only needs $\log(1/\varepsilon)$ bits of randomness (which, as mentioned in Section 3.1, is desirable for efficiency).

The only difference between Lemma 1 and our final goal (i.e. Theorem 2) is that the security guarantee of the former one forbids adversary to see S (it only makes P_0, $\cdots$, $P_{\ell-1}$ public). We now argue why this security guarantee remains when additionally conditioned on S. Beforehand, we introduce preliminaries about key-agreement and PKE.

---

[4] As shown in Lemma 4, PRFs are computational extractors in the sense that when applied to min-entropy sources (or their computational analogue HILL pseudo-entropy sources), one obtains pseudo-random outputs provided that independent P'_is are used.

KEY-AGREEMENT AND PKE. PKE is equivalent to a 2-pass key-agreement protocol [18], which in turn can be obtained from a 2-pass bit-agreement protocol with noticeable correlation and overwhelming security [15]. *Bit-agreement* refers to a protocol in which two efficient parties Alice and Bob (without any pre-shared secrets) communicate over an authenticated channel. At the end of the protocol, Alice and Bob output a bit b_{A} and b_{B}, respectively. The protocol has *correlation* ϵ, if it holds that $\Pr[b_{\mathsf{A}} = b_{\mathsf{B}}] \geq \frac{1+\epsilon}{2}$. Furthermore, the protocol has security δ, if for every efficient adversary Eve, which can observe the whole communication C, it holds that $\Pr[\mathsf{Eve}(1^k, C) = b_{\mathsf{B}}] \leq 1 - \frac{\delta}{2}$.

The following Lemma completes the proof of Theorem 2.

Lemma 3 (Secret vs. Public S). *For the same* F, G, ℓ, A, D *as given in Theorem 2 such that by keeping S secret, the stream cipher* SC *is secure as stated in Lemma 1, i.e.*

$$\left| \Pr_{K_0,S}[\mathsf{D}(\mathsf{view}_\ell \setminus S, P_{[0\cdots\ell-1]}, X_\ell) = 1] - \mathsf{D}(\mathsf{view}_\ell \setminus S, P_{[0\cdots\ell-1]}, U_n) = 1] \right| = \mathsf{negl}(n), \tag{4}$$

we have that either the above is still negligible when additionally conditioned on S, or otherwise there exists efficient black-box constructions of public key encryption from the PRFs F *and* G *and the leakage functions* $\mathsf{L}_1,\cdots,\mathsf{L}_{\ell-1}$.

Proof. By contradiction, let us assume that for some $c > 0$ and for infinitely many n's, there exists efficient $\tilde{\mathsf{D}}$ such that: $\Pr_{K_0,S}[\tilde{\mathsf{D}}(\mathsf{view}_\ell, X_\ell) = 1] - \Pr_{K_0,S}[\tilde{\mathsf{D}}(\mathsf{view}_\ell, U_n) = 1] \geq \frac{1}{n^c}$. We construct a 2-pass bit-agreement protocol as in Figure 5.

It follows from Equation (4) that no efficient passive adversary Eve (observing the communication) will be able to guess b_B (i.e. whether r is x_ℓ or uniform randomness) with more than negligible advantage. Furthermore, the bit-agreement also achieves correlation:

Alice		Bob
$s \leftarrow U_n$	$\xrightarrow{p_0,\cdots,p_{\ell-1}}$	$k_0 \leftarrow U_n$
$p_0,\cdots,p_{\ell-1} \leftarrow \mathsf{G}(s,0),\cdots,\mathsf{G}(s,\ell-1)$		Evaluate SC on $k_0, p_0,\cdots,p_{\ell-1}$
		to get $\mathsf{view}_\ell \setminus s$ and x_ℓ
		$b_{\mathsf{B}} \leftarrow U_1$
	$\xleftarrow{r, \mathsf{view}_\ell \setminus s}$	if $b_{\mathsf{B}} = 0$ then $r := x_\ell$
$b_{\mathsf{A}} \leftarrow \tilde{\mathsf{D}}(r, \mathsf{view}_\ell)$		else if $b_{\mathsf{B}} = 1$ then $r \leftarrow U_n$

Fig. 5. A bit agreement protocol from any PRFs F, G and leakage functions $\mathsf{L}_1, \cdots, \mathsf{L}_{\ell-1}$

$$\begin{aligned}\Pr[b_{\mathsf{A}}=b_{\mathsf{B}}] &= \underbrace{\Pr[b_{\mathsf{B}}=1]}_{=1/2}\Pr[b_{\mathsf{A}}=1|b_{\mathsf{B}}=1]+\underbrace{\Pr[b_{\mathsf{B}}=0]}_{=1/2}\underbrace{\Pr[b_{\mathsf{A}}=0|b_{\mathsf{B}}=0]}_{=1-\Pr[b_{\mathsf{A}}=1|b_{\mathsf{B}}=0]}\\ &= \frac{1}{2}\left(\Pr[b_{\mathsf{A}}=1|b_{\mathsf{B}}=1]+1-\Pr[b_{\mathsf{A}}=1|b_{\mathsf{B}}=0]\right)\\ &= \frac{1}{2}\left(1+\Pr_{K_0,S}[\tilde{\mathsf{D}}(\mathsf{view}_\ell,X_\ell)=1]-\Pr_{K_0,S}[\tilde{\mathsf{D}}(\mathsf{view}_\ell,U_n)=1]\right)\geq\frac{1+\frac{1}{n^c}}{2},\end{aligned}$$

which implies 2-pass key agreement and PKE (by privacy amplification and parallel repetition [15]). Intuitively, the protocol can be seen as a bit-PKE. That is, Alice generates secret and public key pair $sk = s$ and $pk = (p_0, \cdots, p_{\ell-1})$ respectively, and sends her public key to Bob for him to encrypt his message b_B such that only Alice (with secret key sk) can decrypt (with non-negligible correlation). This completes the proof. □

As observed in [1], we can further extend this type of bit-PKE to a 1-out-of-2 Oblivious Transfer (OT) against curious-but-honest adversaries[5] as follows. For choice bit b, Alice first samples $pk_b := (p_0, \cdots, p_{\ell-1})$ and $pk_{1-b} \leftarrow U_{n\ell}$ and then sends pk_0, pk_1 to Bob. Bob, who holds two bits σ_0 and σ_1, uses the bit-PKE to encrypt σ_0 and σ_1 under pk_0 and pk_1, respectively. Finally, Alice recover σ_b and learns no information about σ_{1-b} (since it is computationally hidden by uniform randomness pk_{1-b}).

Additional remark about the protocol in Figure 5. In the non-uniform setting, any insecurity already implies efficient protocols for PKE and OT (using the hypothetical non-uniform $\tilde{\mathsf{D}}$), whereas in the uniform setting we will get practical and useful protocols, uniformly generated given the security parameter. See more discussion in [1].

3.3 Alternative Proof of Lemma 2

We will need the two following technical lemmata for the proof.

Theorem 3 (Dense Model Theorem [9,11]). *Let $(X, Z) \in \{0,1\}^n \times \{0,1\}^\lambda$ be random variables such that $\mathsf{CD}_t(X, U_n) < \varepsilon$ and let $\varepsilon_{\mathsf{HILL}} > 0$. Then we have:*

$$\mathbf{H}^{\mathsf{HILL}}_{2^\lambda\varepsilon+\varepsilon_{\mathsf{HILL}},t_{\mathsf{HILL}}}(X|Z) \geq n-\lambda, \quad \text{where } t_{\mathsf{HILL}} \in \Omega(\varepsilon^2_{\mathsf{HILL}}\cdot t/n).$$

Lemma 4 (PRFs on Weak Keys and Inputs [6,27]). *If $\mathsf{F} : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}^m$ is a $(2t, 2, \varepsilon)$-secure PRF, then for (K,Z) with $\widetilde{\mathbf{H}}_\infty(K|Z) \geq n-\lambda$, and independent P with $\mathbf{H}_\infty(P) \geq \log(1/\varepsilon)$, we have $\mathsf{CD}_t(\mathsf{F}(K,P), U_m \mid P, Z) \leq \sqrt{2^\lambda\cdot\varepsilon}$.*

[5] A 1-out-of-2 oblivious transfer refers to a protocol, where Alice has a bit b and Bob has two messages m_0 and m_1 such that Alice wishes to receive m_b without Bob learning b, while Bob wants to be assured that the Alice receives only one of the two messages.

Proof sketch. Similar to [9,27], we show the above by induction on ε_i and t_i. For $i = 1$, Equation (3) is trivially satisfied ($t_1 = \infty$ and $\varepsilon_1 = 0$). It remains to show that if Equation (3) holds for case i with parameter ε_i and t_i, then it must hold for case $i+1$ with $\varepsilon_{i+1} \leq \varepsilon_i + 2\sqrt{2^{3\lambda} \cdot \varepsilon}$ and $t_{i+1}{=}\min\{t_i - (t_{\mathsf{F}} + t_{\mathsf{L}_i}), \Theta(2^{3\lambda}\varepsilon \cdot t/n)\}$. By Definition 2, Equation (3) with (ε_i,t_i) refers to the fact that conditioned on $\mathsf{view}'_i \setminus P'_{i-1}$, there exists $\tilde{K}_{i-1}$ with $n - \lambda$ bits of average min-entropy such that K_{i-1} is (t_i, ε_i)-close to $\tilde{K}_{i-1}$. By our leakage assumptions, P'_{i-1} is independent of (K_{i-1},$\mathsf{view}'_i \setminus P'_{i-1}$), so if we apply F to $\tilde{K}_{i-1}$ and P'_{i-1}, Lemma 4 directly implies that:

$$\mathsf{CD}_{t/2}(\ (\tilde{K}_i, \tilde{X}_i) := \mathsf{F}(\tilde{K}_{i-1}, P'_{i-1})\ , U_{2n} \mid \mathsf{view}'_i\ ) \leq \sqrt{2^{\lambda} \cdot \varepsilon}.$$

Taking into account $\mathsf{L}_i(\tilde{K}_{i-1}, P'_{i-1})$, we know by Theorem 3 that:

$$\mathbf{H}^{\mathsf{HILL}}_{2\sqrt{2^{3\lambda}\cdot\varepsilon},\Theta(2^{3\lambda}\varepsilon\cdot t/n)}(\ \tilde{K}_i, \tilde{X}_i \mid \mathsf{view}'_i, \mathsf{L}_i(\tilde{K}_{i-1}, P'_{i-1})\ )\ \geq\ 2n - \lambda,$$

which implies (using the chain rule for min-entropy) that $\tilde{K}_i$ has $n - \lambda$ bits of HILL pseudo-entropy (for the same parameters) conditioned on $\tilde{X}_i$. Note that this is almost what we want except that F is applied to $\tilde{K}_{i-1}$ rather than K_{i-1}. Hence, we need to pay $2\sqrt{2^{3\lambda} \cdot \varepsilon}$ for $\varepsilon_{i+1} - \varepsilon_i$, and lose $t_{\mathsf{F}}{+}t_{\mathsf{L}_i}$ in complexity (to simulate the experiment). □

4 Leakage-Resilient PRFs

By minimizing their randomness requirements, the previous results improve the relevance of leakage-resilient stream ciphers. Besides, they also increases our confidence that simple constructions such as the first proposal in [37] are indeed secure against side-channel attacks. Hence, a natural question is to investigate whether a similar situation is observed for PRFs. In this context, three proposals have been analyzed in the literature. Standaert et al. first observed in [35] that a tree-based construction such as the one of Goldreich, Goldwasser and Micali [13] inherently brings improved resistance against side-channel attacks. They proved its leakage-resilience under a (non-standard) random oracle based assumption. Next, Dodis and Pietrzak proposed a similar tree-based design using an alternating structure, and proved its leakage-resilience in the standard model. Finally, Faust et al. replaced the alternating structure by public randomness (following the approach they used for the stream cipher in Figure 3) [10]. In this last case, a fresh p_i is required in each step of the PRF tree. The techniques described in the previous section can be directly applied to mitigate this requirement. That is, one can run a PRF on a counter and public seed to generate the p_i's. As in Lemma 3, either this construction is secure, or we can build a bit agreement protocol using the PRFs and leakage functions of the figure. While the randomness saving may be not substantial for a regular PRF (with input size linear in n), it will be desirable for variants that handle long (polynomial-size) inputs, e.g.

for Message Authentication Codes (MACs). Finally, we note that as in [10], the constructed leakage-resilient PRF is only secure against non-adaptive inputs.

Acknowledgements. Yu Yu was supported by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the National Natural Science Foundation of China Grant 61033001, 61172085, 61061130540, 61073174, 61103221, 11061130539, 61021004 and 61133014. François-Xavier Standaert is an associate researcher of the Belgian Fund for Scientific Research (FNRS-F.R.S.). This work has been funded in part by the ERC project 280141 on CRyptographic Algorithms and Secure Hardware (CRASH).

References

1. Barak, B., Dodis, Y., Krawczyk, H., Pereira, O., Pietrzak, K., Standaert, F.-X., Yu, Y.: Leftover Hash Lemma, Revisited. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 1–20. Springer, Heidelberg (2011)
2. Barak, B., Shaltiel, R., Wigderson, A.: Computational Analogues of Entropy. In: Arora, S., Jansen, K., Rolim, J.D.P., Sahai, A. (eds.) RANDOM 2003 and APPROX 2003. LNCS, vol. 2764, pp. 200–215. Springer, Heidelberg (2003)
3. Chari, S., Jutla, C.S., Rao, J.R., Rohatgi, P.: Towards Sound Approaches to Counteract Power-Analysis Attacks. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 398–412. Springer, Heidelberg (1999)
4. Clavier, C., Coron, J.-S., Dabbous, N.: Differential Power Analysis in the Presence of Hardware Countermeasures. In: Paar, C., Koç, Ç.K. (eds.) CHES 2000. LNCS, vol. 1965, pp. 252–263. Springer, Heidelberg (2000)
5. Dodis, Y., Pietrzak, K.: Leakage-Resilient Pseudorandom Functions and Side-Channel Attacks on Feistel Networks. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 21–40. Springer, Heidelberg (2010)
6. Dodis, Y., Yu, Y.: Overcoming weak expectations. Short version appears in Information Theory Workshop ITW 2012 (2012), `http://www.cs.nyu.edu/~dodis/ps/weak-expe.pdf`
7. Dubrov, B., Ishai, Y.: On the randomness complexity of efficient sampling. In: Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC 2006), pp. 711–720 (2006)
8. Dziembowski, S.: On Forward-Secure Storage. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 251–270. Springer, Heidelberg (2006)
9. Dziembowski, S., Pietrzak, K.: Leakage-resilient cryptography. In: Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2008), pp. 293–302 (2008)
10. Faust, S., Pietrzak, K., Schipper, J.: Practical Leakage-Resilient Symmetric Cryptography. In: Prouff, E., Schaumont, P. (eds.) CHES 2012. LNCS, vol. 7428, pp. 213–232. Springer, Heidelberg (2012)
11. Fuller, B., O'Neill, A., Reyzin, L.: A Unified Approach to Deterministic Encryption: New Constructions and a Connection to Computational Entropy. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 582–599. Springer, Heidelberg (2012)
12. Gandolfi, K., Mourtel, C., Olivier, F.: Electromagnetic Analysis: Concrete Results. In: Koç, Ç.K., Naccache, D., Paar, C. (eds.) CHES 2001. LNCS, vol. 2162, pp. 251–261. Springer, Heidelberg (2001)

13. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions. In: Proceedings of the 25th Annual Symposium on Foundations of Computer Science (FOCS 1984), pp. 464–479 (1984)
14. Håstad, J., Impagliazzo, R., Levin, L.A., Luby, M.: A pseudorandom generator from any one-way function. SIAM J. Comput. 28(4), 1364–1396 (1999)
15. Holenstein, T.: Key agreement from weak bit agreement. In: Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005), pp. 664–673 (2005)
16. Hsiao, C.-Y., Lu, C.-J., Reyzin, L.: Conditional Computational Entropy, or Toward Separating Pseudoentropy from Compressibility. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 169–186. Springer, Heidelberg (2007)
17. Impagliazzo, R.: A personal view of average-case complexity. In: Proceedings of Structure in Complexity Theory Conference, pp. 134–147 (1995)
18. Impagliazzo, R., Rudich, S.: Limits on the provable consequences of one-way permutations. In: The 21st Annual ACM Symposium on Theory of Computing (STOC 1989), pp. 44–61 (1989)
19. 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)
20. Mangard, S., Oswald, E., Popp, T.: Power analysis attacks - revealing the secrets of smart cards. Springer (2007)
21. Mangard, S., Popp, T., Gammel, B.M.: Side-Channel Leakage of Masked CMOS Gates. In: Menezes, A. (ed.) CT-RSA 2005. LNCS, vol. 3376, pp. 351–365. Springer, Heidelberg (2005)
22. Medwed, M., Standaert, F.-X.: Extractors against Side-Channel Attacks: Weak or Strong? In: Preneel, B., Takagi, T. (eds.) CHES 2011. LNCS, vol. 6917, pp. 256–272. Springer, Heidelberg (2011)
23. Medwed, M., Standaert, F.-X., Joux, A.: Towards Super-Exponential Side-Channel Security with Efficient Leakage-Resilient PRFs. In: Prouff, E., Schaumont, P. (eds.) CHES 2012. LNCS, vol. 7428, pp. 193–212. Springer, Heidelberg (2012)
24. Micali, S., Reyzin, L.: Physically Observable Cryptography (Extended Abstract). In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 278–296. Springer, Heidelberg (2004)
25. 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: Abe, M., Gligor, V.D. (eds.) ASIACCS, pp. 56–65. ACM (2008)
26. Pietrzak, K.: Composition Implies Adaptive Security in Minicrypt. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 328–338. Springer, Heidelberg (2006)
27. Pietrzak, K.: A Leakage-Resilient Mode of Operation. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 462–482. Springer, Heidelberg (2009)
28. Pietrzak, K., Sjödin, J.: Weak Pseudorandom Functions in Minicrypt. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfsdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part II. LNCS, vol. 5126, pp. 423–436. Springer, Heidelberg (2008)
29. Quisquater, J.-J., Samyde, D.: ElectroMagnetic Analysis (EMA): Measures and Counter-Measures for Smart Cards. In: Attali, S., Jensen, T. (eds.) E-smart 2001. LNCS, vol. 2140, pp. 200–210. Springer, Heidelberg (2001)
30. Reingold, O., Trevisan, L., Tulsiani, M., Vadhan, S.P.: Dense subsets of pseudorandom sets. In: Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2008), pp. 76–85 (2008)

31. Rivain, M., Prouff, E.: Provably Secure Higher-Order Masking of AES. In: Mangard, S., Standaert, F.-X. (eds.) CHES 2010. LNCS, vol. 6225, pp. 413–427. Springer, Heidelberg (2010)
32. Standaert, F.-X.: How Leaky Is an Extractor? In: Abdalla, M., Barreto, P.S.L.M. (eds.) LATINCRYPT 2010. LNCS, vol. 6212, pp. 294–304. Springer, Heidelberg (2010)
33. Standaert, F.-X., Malkin, T.G., 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)
34. Standaert, F.-X., Veyrat-Charvillon, N., Oswald, E., Gierlichs, B., Medwed, M., Kasper, M., Mangard, S.: The World Is Not Enough: Another Look on Second-Order DPA. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 112–129. Springer, Heidelberg (2010)
35. Standaert, F.-X., Pereira, O., Yu, Y., Quisquater, J.-J., Yung, M., Oswald, E.: Leakage resilient cryptography in practice. In: Towards Hardware Intrinsic Security: Foundation and Practice, pp. 105–139. Springer, Heidelberg (2010); Cryptology ePrint Archive, Report 2009/341 (2009), http://eprint.iacr.org/
36. Tiri, K., Verbauwhede, I.: Securing Encryption Algorithms against DPA at the Logic Level: Next Generation Smart Card Technology. In: Walter, C.D., Koç, Ç.K., Paar, C. (eds.) CHES 2003. LNCS, vol. 2779, pp. 125–136. Springer, Heidelberg (2003)
37. Yu, Y., Standaert, F.-X., Pereira, O., Yung, M.: Practical leakage-resilient pseudorandom generators. In: Al-Shaer, E., Keromytis, A.D., Shmatikov, V. (eds.) ACM Conference on Computer and Communications Security, pp. 141–151. ACM (2010)
38. Zheng, C.J.: A uniform min-max theorem and its applications. STOC 2012, Poster (2012), http://cs.nyu.edu/~stoc2012/acceptedposters.pdf