How to Fake Auxiliary Input

Dimitar Jetchev and Krzysztof Pietrzak[*]

EPFL Switzerland and IST Austria

Abstract. Consider a joint distribution (X, A) on a set $\mathcal{X} \times \{0,1\}^\ell$. We show that for any family $\mathcal{F}$ of distinguishers $f\colon \mathcal{X} \times \{0,1\}^\ell \to \{0,1\}$, there exists a simulator $h\colon \mathcal{X} \to \{0,1\}^\ell$ such that

1. no function in $\mathcal{F}$ can distinguish (X, A) from $(X, h(X))$ with advantage ε,
2. h is only $O(2^{3\ell}\varepsilon^{-2})$ times less efficient than the functions in $\mathcal{F}$.

For the most interesting settings of the parameters (in particular, the cryptographic case where X has superlogarithmic min-entropy, $\varepsilon > 0$ is negligible and $\mathcal{F}$ consists of circuits of polynomial size), we can make the simulator h *deterministic*.

As an illustrative application of our theorem, we give a new security proof for the leakage-resilient stream-cipher from Eurocrypt'09. Our proof is simpler and quantitatively much better than the original proof using the dense model theorem, giving meaningful security guarantees if instantiated with a standard blockcipher like AES.

Subsequent to this work, Chung, Lui and Pass gave an interactive variant of our main theorem, and used it to investigate weak notions of Zero-Knowledge. Vadhan and Zheng give a more constructive version of our theorem using their new uniform min-max theorem.

1 Introduction

Let $\mathcal{X}$ be a set and let $\ell > 0$ be an integer. We show that for any joint distribution (X, A) over $\mathcal{X} \times \{0,1\}^\ell$ (where we think of A as a short ℓ-bit auxiliary input to X), any family $\mathcal{F}$ of functions $\mathcal{X} \times \{0,1\}^\ell \to \{0,1\}$ (thought of as distinguishers) and any $\varepsilon > 0$, there exists an efficient simulator $h\colon \mathcal{X} \to \{0,1\}^\ell$ for the auxiliary input that fools every distinguisher in $\mathcal{F}$, i.e.,

$$\forall f \in \mathcal{F}\colon |\mathbb{E}[f(X, A)] - \mathbb{E}[f(X, h(X))]| < \varepsilon.$$

Here, "efficient" means that the simulator h is $\tilde{\mathcal{O}}(2^{3\ell}\varepsilon^{-2})$ times more complex than the functions from $\mathcal{F}$ (we will formally define "more complex" in Definition 6). Without loss of generality, we can model the joint distribution (X, A) as $(X, g(X))$, where g is some arbitrarily complex and possibly probabilistic function (where $\mathbb{P}[g(x) = a] = \mathbb{P}[A = a|X = x]$ for all $(x, a) \in \mathcal{X} \times \{0,1\}^\ell$). Let us stress that, as g can be arbitrarily complex, one cannot hope to get an efficient simulator h where $(X, g(X))$ and $(X, h(X))$ are statistically close. Yet, one can still fool all functions in $\mathcal{F}$ in the sense that no function from $\mathcal{F}$ can distinguish the distribution (X, A) from $(X, g(X))$.

[*] Supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC Starting Grant (259668-PSPC)

Relation to [25]. Trevisan, Tulsiani and Vadhan [25, Thm. 3.1] prove a conceptually similar result, stating that if $\mathcal{Z}$ is a set then for any distribution Z over $\mathcal{Z}$, any family $\widetilde{\mathcal{F}}$ of functions $\mathcal{Z} \to [0,1]$ and any function $\widetilde{g} \colon \mathcal{Z} \to [0,1]$, there exists a simulator $\widetilde{h} \colon \mathcal{Z} \to [0,1]$ whose complexity is only $\mathcal{O}(\varepsilon^{-2})$ times larger than the complexity of the functions from $\widetilde{\mathcal{F}}$ such that

$$\forall \widetilde{f} \in \widetilde{\mathcal{F}} \ : \ |\,\mathbb{E}[\widetilde{f}(Z)\widetilde{g}(Z)] - \mathbb{E}[\widetilde{f}(Z)\widetilde{h}(Z)]| < \varepsilon. \tag{1}$$

In [25], this result is used to prove that every high-entropy distribution is indistinguishable from an efficiently samplable distribution of the same entropy. Moreover, it is shown that many fundamental results including the Dense Model Theorem [23,14,21,10,24], Impagliazzo's hardcore lemma [18] and a version of Szémeredi's Regularity Lemma [11] follow from this theorem. The main difference between (1) and our statement

$$\forall f \in \mathcal{F} \colon |\,\mathbb{E}[f(X, g(X))] - \mathbb{E}[f(X, h(X))]| < \varepsilon \tag{2}$$

is that our distinguisher f sees not only X, but also the real or fake auxiliary input $g(X)$ or $h(X)$, whereas in (1), the distinguisher $\widetilde{f}$ only sees X. In particular, the notion of indistinguishability we achieve captures indistinguishability in the standard cryptographic sense. On the other hand, (1) is more general in the sense that the range of $\widetilde{f}, \widetilde{g}, \widetilde{h}$ can be any real number in $[0,1]$, whereas our f has range $\{0,1\}$ and g, h have range $\{0,1\}^{\ell}$.

Nonetheless, it is easy to derive (1) from (2): consider the case of $\ell = 1$ bit of auxiliary input, and only allow families $\mathcal{F}$ of distinguishers where each $f \in \mathcal{F}$ is of the form $f(X, b) = \widehat{f}(X)b$ for some function $\widehat{f} \colon \mathcal{X} \to [0,1]$. For this restricted class, the absolute value in (2) becomes

$$|\,\mathbb{E}[f(X, g(X))] - \mathbb{E}[f(X, h(X))| = |\,\mathbb{E}[\widehat{f}(X)g(X)] - \mathbb{E}[\widehat{f}(X)h(X)]| \tag{3}$$

As $\widehat{f}$ is arbitrary, this restricted class almost captures the distinguishers considered in (1). The only difference is that the function $\widetilde{g}$ has range $[0,1]$ whereas our g has range $\{0,1\}$. Yet, note that in (1), we can replace $\widetilde{g}$ having range $[0,1]$ by a (probabilistic) g with range $\{0,1\}$ defined as $\mathbb{P}[g(x) = 1] = \widetilde{g}(x)$, thus, leaving the expectation $\mathbb{E}[\widetilde{f}(X)\widetilde{g}(X)] = \mathbb{E}[\widetilde{f}(X)g(X)]$ unchanged.[1]

In [25], two different proofs for (1) are given. The first proof uses duality of linear programming in the form of the min-max theorem for two-player zero-sum games. This proof yields a simulator of complexity $\mathcal{O}(\varepsilon^{-4}\log^2(1/\varepsilon))$ times the complexity of the functions in $\mathcal{F}$. The second elegant proof uses boosting and gives a quantitatively much better $\mathcal{O}(\varepsilon^{-2})$ complexity.

[1] The simulator $\widetilde{h}$ from [25] satisfies the additional property $|\,\mathbb{E}[\widetilde{h}(X)] - \mathbb{E}[\widetilde{g}(X)]| = 0$. If this property is needed, we can get it by requiring that the function $f(X, b) = b$ is in $\mathcal{F}$. Then (2) for this f implies $|\,\mathbb{E}[g(X)] - \mathbb{E}[h(X)]| < \varepsilon$. One can make this term exactly zero by slightly biasing h towards 0 if $\mathbb{E}[h(X)] > \mathbb{E}[g(X)]$ or 1 otherwise, slightly increasing the advantage from ε to at most 2ε.

Proof outline. As it was just explained, (1) follows from (2). We do not know if one can directly prove an implication in the other direction, so we prove (2) from scratch. Similarly to [25], the core of our proof uses boosting with the same energy function as the one used in [25].

As a first step, we transform the statement (2) into a "product form" like (1) where $\mathcal{Z} = \mathcal{X} \times \{0,1\}^\ell$ (this results in a loss of a factor of 2^ℓ in the advantage ε; in addition, our distinguishers $\widehat{f}$ will have range $[-1,1]$ instead of $[0,1]$). We then prove that (1) holds for some simulator $\widetilde{h}\colon \mathcal{Z} \to [0,1]$ of complexity ε^{-2} relative to $\mathcal{F}$. Unfortunately, we cannot use the result from [25] in a black-box way at this point as we need the simulator $\widetilde{h}\colon \mathcal{Z} \to [0,1]$ to define a probability distribution in the sense that $\widetilde{h}(x,b) \geq 0$ for all (x,b) and $\sum\limits_{b\in\{0,1\}^\ell} \widetilde{h}(x,b) = 1$ for all x. Ensuring these conditions is the most delicate part of the proof. Finally, we show that the simulator h defined via $\mathbb{P}[h(x) = b] = \widetilde{h}(x,b)$ satisfies (2). Note that for h to be well defined, we need $\widetilde{h}$ to specify a probability distribution as outlined above.

Efficiency of h. Our simulator h is efficient in the sense that it is only $\mathcal{O}(2^{3\ell}\varepsilon^{-2})$ times more complex than the functions in $\mathcal{F}$. We do not know how tight this bounds is, but one can prove a lower bound of $\max\{2^\ell, \varepsilon^{-1}\}$ under plausible assumptions. The dependency on 2^ℓ is necessary under exponential hardness assumptions for one-way functions.[2] A dependency on ε^{-1} is also necessary. Indeed, Trevisan et al. [25, Rem. 1.6] show that such a dependency is necessary for the simulator $\widetilde{h}$ in (1). Since (1) is implied by (2) with h and $\widetilde{h}$ having exactly the same complexity, the ε^{-1} lower bound also applies to our h.

1.1 Subsequent work

The original motivation for this work was to give simpler and quantitatively better proofs for leakage-resilient cryptosystems as we will discuss in Section 4. Our main theorem has subsequently been derived via two different routes.

First, Chung, Lui and Pass [4] investigate weak notions of zero-knowledge. On route, they derive an "interactive" version of our main theorem. In Section 4, we will show how to establish one of their results (with better quantitative bounds), showing that every interactive proof system satisfies a weak notion of zero-knowledge.

Second, Vadhan and Zheng [26, Thm.3.1-3.2] recently proved a version of von Neumann's min-max theorem for two-player zero sum games that does not only guarantee existence of an optimal strategy for the second player, but also constructs a nearly optimal strategy assuming knowledge of several best responses of the second player to strategies of the first player, and provide many applications

[2] More precisely, assume there exists a one-way function where inverting becomes 2^ℓ times easier given ℓ bits of leakage. It is e.g. believed that the AES block-cipher gives such a function as $(K,X) \to (\mathsf{AES}(K,X), X)$.

of this theorem. Their argument is based on relative entropy KL projections and a learning technique known as weight updates and resembles the the proof of the Uniform Hardcore Lemma by Barak, Hardt and Kale [2] (see also [16] for the original application of this method). They derive our main theorem [26, Thm.6.8], but with incomparable bounds. Concretely, to fool circuits of size t, their simulator runs in time $\tilde{O}(t\cdot 2^{\ell}/\varepsilon^2+2^{\ell}/\varepsilon^4)$ compared to ours whose run-time is $\tilde{O}(t\cdot 2^{3\ell}/\varepsilon^2)$. In particular, their bounds are better whenever $1/\varepsilon^2 \leq t\cdot 2^{2\ell}$. The additive $2^{\ell}/\varepsilon^4$ term in their running time appears due to the sophisticated iterative"weight update" procedure, whereas our simulator simply consists of a weighted sum of the evaluation of $\tilde{O}(2^{3\ell}/\varepsilon^2)$ circuits from the family we want to fool (here, circuits of size t).

1.2 More applications

Apart from reproving one of [4]'s results on weak zero-knowledge mentioned above, we give two more applications of our main theorem in Section 4:

Chain Rules for Computational Entropy. Gentry and Wichs [13] show that black-box reductions cannot be used to prove the security of any succinct non-interactive argument from any falsifiable cryptographic assumption. A key technical lemma used in their proof ([13, Lem. 3.1]) states that if two distributions X and $\widetilde{X}$ over $\mathcal{X}$ are computationally indistinguishable, then for any joint distribution (X,A) over $\mathcal{X}\times\{0,1\}^{\ell}$ (here, A is a short ℓ-bit auxiliary input) there exists a joint distribution $(\widetilde{X},\widetilde{A})$ such that (X,A) and $(\widetilde{X},\widetilde{A})$ are computationally indistinguishable. Our theorem immediately implies the stronger statement that not only such an $(\widetilde{X},\widetilde{A})$ exists, but in fact, it is efficiently samplable, i.e., there exists an efficient simulator $h\colon \mathcal{X}\to\{0,1\}^{\ell}$ such that $(\widetilde{X},h(\widetilde{X}))$ is indistinguishable from $(\widetilde{X},\widetilde{A})$ and thus from (X,A). Reyzin [22, Thm.2] observed that the result of Gentry and Wichs implies a chain rule for conditional "relaxed" HILL entropy. We give a short and simple proof of this chain rule in Proposition 2 of this paper. We then show in Corollary 1 how to deduce a chain rule for (regular) HILL entropy from Proposition 2 using the simple fact (Lemma 1) that short (i.e., logarithmic in the size of the distinguishers) computationally indistinguishable random variables must already be statistically close. Chain rules for HILL entropy have found several applications in cryptography [10,21,7,12]. The chain rule that we get in Corollary 1 is the first one that does not suffer from a significant loss in the distinguishing advantage (we only lose a constant factor of 4). Unlike the case of relaxed HILL-entropy, here we only prove a chain rule for the "non-conditional" case, which is a necessary restriction given a recent counterexample to the (conditional) HILL chain rule by Krenn et al. [19]. We will provide more details on this negative result after the statement of Corallary 1.

Leakage Resilient Cryptography. The original motivation for this work is to simplify the security proofs of leakage-resilient [10,20,7] and other cryptosystems [12] whose security proofs rely on chain rules for computational entropy (as discussed in the previous paragraph). The main idea is to replace the chain rules with simulation-based arguments. In a nutshell, instead of arguing that a variable X must have high (pseudo)entropy in the presence of a short leakage A, one could simply use the fact that the leakage can be efficiently simulated. This not only implies that X has high (pseudo)entropy given the fake leakage $h(X)$, but if X is pseudorandom, it also implies that $(X, h(X))$ is indistinguishable from $(U, h(U))$ for a uniform random variable U on the same set as X. In the security proofs, we would now replace $(X, h(X))$ with $(U, h(U))$ and will continue with a uniformly random intermediate variable U. In contrast, the approach based on chain rules only tells us that we can replace X with some random variable Y that has high min-entropy given A. This is not only much complex to work with, but it often gives weaker quantitative bounds. In particular, in Section 4.3 we revisit the security proof of the leakage-resilient stream-cipher from [20] for which we can now give a conceptually simpler and quantitatively better security proof.

2 Notation and Basic Definitions

2.1 Notation

We use calligraphic letters such as $\mathcal{X}$ to denote sets, the corresponding capital letters X to denote random variables on these sets (equivalently, probability distributions) and lower-case letters (e.g., x) for values of the corresponding random variables. Moreover, $x \leftarrow X$ means that x is sampled according to the distribution X and $x \leftarrow \mathcal{X}$ means that x is sampled uniformly at random from $\mathcal{X}$. Let U_n denote the random variable with uniform distribution on $\{0,1\}^n$. We denote by $\Delta(X;Y) = \frac{1}{2}\sum_{x\in\mathcal{X}} |\,\mathbb{P}[X = x] - \mathbb{P}[Y = x]|$ the statistical distance between X and Y. For $\varepsilon > 0, s \in \mathbb{N}$, we use $X \sim Y$ to denote that X and Y have the same distribution, $X \sim_{\varepsilon} Y$ to denote that their statistical distance is less than ε and $X \sim_{\varepsilon,s} Y$ to denote that no circuit of size s can distinguish X from Y with advantage greater than ε. Note that $X \sim_{\varepsilon,\infty} Y \iff X \sim_{\varepsilon} Y$ and $X \sim_0 Y \iff X \sim Y$.

Finally, if $h \colon \mathcal{X} \to \{0,1\}^{\ell}$ is a probabilistic (randomized) function then we will use $[h]$ to denote the random coins used by h (a notation that will be used in various probabilities and expectations).

2.2 Entropy Measures

A random variable X has min-entropy k, if no (computationally unbounded) adversary can predict the outcome of X with probability greater than 2^{-k}.

Definition 1. (Min-Entropy H_{∞}) *A random variable X has* min-entropy k*, denoted* $\mathsf{H}_{\infty}(X) \geq k$*, if* $\max_x \mathbb{P}[X = x] \leq 2^{-k}$.

Dodis et al. [8] gave a notion of average-case min-entropy defined such that X has average-case min-entropy k conditioned on Z if the probability of the best adversary in predicting X given Z is 2^{-k}.

Definition 2. (Average min-Entropy [8] $\widetilde{\mathsf{H}}_{\infty}$) *Consider a joint distribution (X, Z), then the* average min-entropy *of X conditioned on Z is*

$$\widetilde{\mathsf{H}}_{\infty}(X|Z) = -\log(\mathop{\mathbb{E}}_{z \leftarrow Z}\left[\max_{x} \mathbb{P}[X = x|Z = z]\right]) = -\log(\mathop{\mathbb{E}}_{z \leftarrow Z}\left[2^{-\mathsf{H}_{\infty}(X|Z=z)}\right])$$

HILL-entropy is the computational analogue of min-entropy. A random variable X has HILL-entropy k if there exists a random variable Y having min-entropy k that is indistinguishable from X. HILL-entropy is further quantified by two parameters ε, s specifying this indistinguishability quantitatively.

Definition 3. (HILL-Entropy [15] $\mathsf{H}^{\mathsf{HILL}}$) *$X$ has* HILL entropy *k, denoted by $\mathsf{H}^{\mathsf{HILL}}_{\varepsilon,s}(X) \geq k$, if*

$$\mathsf{H}^{\mathsf{HILL}}_{\varepsilon,s}(X) \geq k \quad \iff \quad \exists Y \ : \ \mathsf{H}_{\infty}(Y) \geq k \quad \text{and} \quad X \sim_{\varepsilon,s} Y$$

Conditional HILL-entropy has been defined by Hsiao, Lu and Reyzin [17] as follows.

Definition 4. (Conditional HILL-Entropy [17]) *X has conditional HILL entropy k (conditioned on Z), denoted $\mathsf{H}^{\mathsf{HILL}}_{\varepsilon,s}(X|Z) \geq k$, if*

$$\mathsf{H}^{\mathsf{HILL}}_{\varepsilon,s}(X|Z) \geq k \quad \iff \quad \exists (Y, Z) \ : \ \widetilde{\mathsf{H}}_{\infty}(Y|Z) \geq k \quad \text{and} \quad (X, Z) \sim_{\varepsilon,s} (Y, Z)$$

Note that in the definition above, the marginal distribution on the conditional part Z is the same in both the real distribution (X, Z) and the indistinguishable distribution (Y, Z). A "relaxed" notion of conditional HILL used implicitly in [13] and made explicit in [22] drops this requirement.

Definition 5. (Relaxed Conditional HILL-Entropy [13,22]) *X has relaxed conditional HILL entropy k, denoted $\mathsf{H}^{\mathsf{rlx\text{-}HILL}}_{\varepsilon,s}(X|Z) \geq k$, if*

$$\mathsf{H}^{\mathsf{rlx\text{-}HILL}}_{\varepsilon,s}(X|Z) \geq k \quad \iff \quad \exists (Y, W) \ : \ \widetilde{\mathsf{H}}_{\infty}(Y|W) \geq k \quad \text{and} \quad (X, Z) \sim_{\varepsilon,s} (Y, W)$$

3 The main theorem

Definition 6. (Complexity of a function) *Let $\mathcal{A}$ and $\mathcal{B}$ be sets and let $\mathcal{G}$ be a family of functions $h \colon \mathcal{A} \to \mathcal{B}$. A function h has* **complexity C relative to** *$\mathcal{G}$ if it can be computed by an oracle-aided circuit of size* $\mathrm{poly}(C \log |\mathcal{A}|)$ *with C oracle gates where each oracle gate is instantiated with a function from $\mathcal{G}$.*

Theorem 1. (Main) *Let $\ell \in \mathbb{N}$ be fixed, let $\varepsilon > 0$ and let $\mathcal{X}$ be any set. Consider a distribution X over $\mathcal{X}$ and any (possibly probabilistic and not necessarily* *efficient) function $g\colon \mathcal{X} \to \{0,1\}^\ell$. Let $\mathcal{F}$ be a family of deterministic (cf. remark below) distinguishers $f\colon \mathcal{X} \times \{0,1\}^\ell \to \{0,1\}$. There exists a (probabilistic) simulator $h : \mathcal{X} \to \{0,1\}^\ell$ with complexity*[3]

$$O(2^{3\ell}\varepsilon^{-2}\log^2(\varepsilon^{-1}))$$

relative to $\mathcal{F}$ which ε-fools every distinguisher in $\mathcal{F}$, i.e.

$$\forall f \in \mathcal{F} \;:\; \left| \mathop{\mathbb{E}}_{x\leftarrow X,[g]}[f(x,g(x))] - \mathop{\mathbb{E}}_{x\leftarrow X,[h]}[f(x,h(x))] \right| < \varepsilon, \tag{4}$$

Moreover, if

$$H_\infty(X) > 2 + \log\log|\mathcal{F}| + 2\log(1/\varepsilon) \tag{5}$$

then there exists a deterministic *h with this property.*

Remark 1 (**Closed and Probabilistic $\mathcal{F}$).** In the proof of Theorem 1 we assume that the class $\mathcal{F}$ of distinguishers is closed under complement, i.e., if $f \in \mathcal{F}$ then also $1-f \in \mathcal{F}$. This is without loss of generality, as even if we are interested in the advantage of a class $\mathcal{F}$ that is not closed, we can simply apply the theorem for $\mathcal{F}' = \mathcal{F} \cup (1-\mathcal{F})$, where $(1-\mathcal{F}) = \{1-f \;:\; f \in \mathcal{F}\}$. Note that if h has complexity t relative to $\mathcal{F}'$, it has the same complexity relative to $\mathcal{F}$. We also assume that all functions $f \in \mathcal{F}$ are deterministic. If we are interested in a class $\mathcal{F}$ of randomized functions, we can simply apply the theorem for the larger class of deterministic functions $\mathcal{F}''$ consisting of all pairs (f,r) where $f \in \mathcal{F}$ and r is a choice of randomness for f. This is almost without loss of generality, except that the requirement in eq.(5) on the min-entropy of X becomes slightly stronger as $\log\log|\mathcal{F}''| = \log\log(|\mathcal{F}|2^\rho)$ where ρ is an upper bound on the number of random coins used by any $f \in \mathcal{F}$.

4 Applications

4.1 Zero-Knowledge

Chung, Lui and Pass [4] consider the following relaxed notion of zero-knowledge

Definition 7 (distributional (T,t,ε)-zero-knowledge). *Let $(\mathcal{P},\mathcal{V})$ be an interactive proof system for a language L. We say that $(\mathcal{P},\mathcal{V})$ is* distributional (T,t,ε)-zero-knowledge *(where T,t,ε are all functions of n) if for every $n \in \mathbb{N}$, every joint distributions (X_n,Y_n,Z_n) over $(L \cap \{0,1\}^n) \times \{0,1\}^* \times \{0,1\}^*$, and every t-size adversary $\mathcal{V}^*$, there exists a T-size simulator S such that*

$$(X_n, Z_n, \mathsf{out}_{\mathcal{V}^*}[\mathcal{P}(X_n,Y_n) \leftrightarrow \mathcal{V}^*(X_n,Z_n)]) \sim_{\varepsilon,t} (X_n, Z_n, S(X_n,Z_n))$$

where $\mathsf{out}_{\mathcal{V}^*}[\mathcal{P}(X_n,Y_n) \leftrightarrow \mathcal{V}^*(X_n,Z_n)]$ denotes the output of $\mathcal{V}^*(X_n,Z_n)$ after interacting with $\mathcal{P}(X_n,Y_n)$.

[3] If we model h as a Turing machine (and not a circuit) and consider the *expected* complexity of h, then we can get a slightly better $O(2^{3\ell}\varepsilon^{-2})$ bound (i.e. without the $\log^2(\varepsilon^{-1})$ term).

If L in an NP language, then in the definition above, Y would be a witness for $X \in L$. As a corollary of their main theorem, [4] show that every proof system satisfies this relaxed notion of zero-knowledge where the running time T of the simulator is polynomial in t, ε and 2^ℓ. We can derive their Corollary from Theorem 1 with better quantitative bounds for most ranges of parameters than [4]: we get $\tilde{O}(t2^{3\ell}\varepsilon^{-2})$ vs. $\tilde{O}(t^3 2^\ell \varepsilon^{-6})$, which is better whenever $t/\varepsilon^2 \geq 2^\ell$.

Proposition 1. *Let $(\mathcal{P}, \mathcal{V})$ be an interactive proof system for a language L, and suppose that the total length of the messages sent by $\mathcal{P}$ is $\ell = \ell(n)$ (on common inputs X of length n). Then for any $t = t(n) \geq \Omega(n)$ and $\varepsilon = \varepsilon(n)$, $(\mathcal{P}, \mathcal{V})$ is distributional (T, t, ε)-zero-knowledge, where*

$$T = O(t2^{3\ell}\varepsilon^{-2}\log^2(\varepsilon^{-1}))$$

Proof. Let $M \in \{0,1\}^\ell$ denote the messages send by $\mathcal{P}(X_n, Y_n)$ when talking to $\mathcal{V}^*(X_n, Z_n)$. By Theorem 1 (identifying $\mathcal{F}$ from the theorem with circuits of size t) there exists a simulator h of size $O(t \cdot 2^{3\ell}\varepsilon^{-2}\log^2(\varepsilon^{-1}))$ s.t.

$$(X_n, Z_n, M) \sim_{\varepsilon, 2t} (X_n, Z_n, h(X_n, Z_n)) \tag{6}$$

Let $S(X_n, Z_n)$ be defined as follows, first compute $M' = h(X_n, Z_n)$ (with h as above), and then compute $\mathsf{out}^*_{\mathcal{V}}[M' \leftrightarrow \mathcal{V}^*(X_n, Z_n)]$. We claim that

$$(X_n, Z_n, \mathsf{out}_{\mathcal{V}^*}[\mathcal{P}(X_n, Y_n) \leftrightarrow \mathcal{V}^*(X_n, Z_n)]) \sim_{\varepsilon, t} (X_n, Z_n, S(X_n, Z_n)) \tag{7}$$

To see this, note that from any distinguisher D of size t that distinguishes the distributions in (7) with advantage $\delta > \varepsilon$, we get a distinguisher D' of size $2t$ that distinguishes the distributions in (6) with the same advantage by defining D' as $\mathsf{D}'(X_n, Z_n, \tilde{M}) = \mathsf{D}(X_n, Z_n, \mathsf{out}_{\mathcal{V}^*}[\tilde{M} \leftrightarrow \mathcal{V}^*(X_n, Z_n)])$. □

4.2 Chain Rules for (Conditional) Pseudoentropy

The following proposition is a chain rule for relaxed conditional HILL entropy. Such a chain rule for the non-conditional case is implicit in the work of Gentry and Wichs [13], and made explicit and generalized to the conditional case by Reyzin [22].

Proposition 2. ([13,22]) *Any joint distribution $(X, Y, A) \in \mathcal{X} \times \mathcal{Y} \times \{0,1\}^\ell$ satisfies*[4]

$$\mathsf{H}^{\mathsf{rlx\text{-}HILL}}_{\varepsilon,s}(X|Y) \geq k \;\Rightarrow\; \mathsf{H}^{\mathsf{rlx\text{-}HILL}}_{2\varepsilon,\hat{s}}(X|Y,A) \geq k-\ell \textit{ where } \hat{s} = \Omega\left(s \cdot \frac{\varepsilon^2}{2^{3\ell}\log^2(1/\varepsilon)}\right)$$

[4] Using the recent bound from [26] discussed in Section 1.1, we can get $\hat{s} = \Omega\left(s \cdot \frac{\varepsilon^2 \ell}{2^\ell} + \frac{\ell^2 \log^2(1/\varepsilon)}{\varepsilon^4}\right)$

Proof. $\mathsf{H}^{\mathsf{rlx\text{-}HILL}}_{\varepsilon,s}(X|Y) \geq k$ means that there exists a random variable (Z,W) such that $\mathsf{H}_\infty(Z|W) \geq k$ and $(X,Y) \sim_{\varepsilon,s} (Z,W)$. For any $\hat{\varepsilon}, \hat{s}$, by Theorem 1, there exists a simulator h of size $s_h = O\left(\hat{s} \cdot \frac{2^{3\ell}\log^2(1/\hat{\varepsilon})}{\hat{\varepsilon}^2}\right)$ such that (we explain the second step below)

$$(X,Y,A) \sim_{\hat{\varepsilon},\hat{s}} (X,Y,h(X,Y)) \sim_{\varepsilon,s-s_h} (Z,W,h(Z,W))$$

The second step follows from $(X,Y) \sim_{\varepsilon,s} (Z,W)$ and the fact that h has complexity s_h. Using the triangle inequality for computational indistinguishability[5] we get

$$(X,Y,A) \sim_{\hat{\varepsilon}+\varepsilon,\min(\hat{s},s-s_h)} (Z,W,h(Z,W))$$

To simplify this expression, we set $\hat{\varepsilon} := \varepsilon$ and $\hat{s} := \Theta(s\varepsilon^2/2^{3\ell}\log^2(1/\varepsilon))$, then $s_h = O(s)$, and choosing the hidden constant in the Θ such that $s_h \leq s/2$ (and thus $\hat{s} \leq s - s_h = s/2$), the above equation becomes

$$(X,Y,A) \sim_{2\varepsilon,\hat{s}} (Z,W,h(Z,W)) \tag{8}$$

Using the chain rule for average case min-entropy in the first, and $\mathsf{H}_\infty(Z|W) \geq k$ in the second step below we get

$$\widetilde{\mathsf{H}}_\infty(Z|W,h(Z,W)) \geq \widetilde{\mathsf{H}}_\infty(Z|W) - \mathsf{H}_0(h(Z,W)) \geq k - \ell\ . \tag{9}$$

Now equations (8) and (9) imply $\mathsf{H}^{\mathsf{rlx\text{-}HILL}}_{2\varepsilon,\hat{s}}(X|Y,A) = k - \ell$ as claimed. □

By the following lemma, conditional relaxed HILL implies conditional HILL if the conditional part is short (at most logarithmic in the size of the distinguishers considered.)

Lemma 1. *For a joint random variable (X,A) over $\mathcal{X} \times \{0,1\}^\ell$ and $s = \Omega(\ell 2^\ell)$ (more concretely, s should be large enough to implement a lookup table for a function $\{0,1\}^\ell \to \{0,1\}$) conditional relaxed HILL implies standard HILL entropy*

$$\mathsf{H}^{\mathsf{rlx\text{-}HILL}}_{\varepsilon,s}(X|A) \geq k \quad \Rightarrow \quad \mathsf{H}^{\mathsf{HILL}}_{2\varepsilon,s}(X|A) \geq k$$

Proof. $\mathsf{H}^{\mathsf{rlx\text{-}HILL}}_{\varepsilon,s}(X|A) \geq k$ means that there exist (Z,W) where $\widetilde{\mathsf{H}}_\infty(Z|W) \geq k$ and

$$(X,A) \sim_{\varepsilon,s} (Z,W) \tag{10}$$

We claim that if $s = \Omega(\ell 2^\ell)$, then (10) implies that $W \sim_\varepsilon A$. To see this, assume the contrary, i.e., that W and A are not ε-close. There exists then a computationally unbounded distinguisher D where

$$|\,\mathbb{P}[D(W)=1] - \mathbb{P}[D(A)=1]| > \varepsilon.$$

[5] which states that for any random variables α, β, γ we have

$$\alpha \sim_{\varepsilon_1,s_1} \beta \quad \& \quad \beta \sim_{\varepsilon_2,s_2} \gamma \quad \Rightarrow \quad \alpha \sim_{\varepsilon_1+\varepsilon_2,\min(s_1,s_2)} \gamma$$

Without loss of generality, we can assume that D is deterministic and thus, implement D by a circuit of size $\Theta(\ell 2^{\ell})$ via a lookup table with 2^{ℓ} entries (where the ith entry is $D(i)$.) Clearly, D can also distinguish (X,A) from (Z,W) with advantage greater than ε by simply ignoring the first part of the input, thus, contradicting (10). As $A \sim_{\varepsilon} W$, we claim that there exist a distribution (Z,A) such that

$$(Z,W) \sim_{\varepsilon} (Z,A). \tag{11}$$

This distribution (Z,A) can be sampled by first sampling (Z,W) and then outputting $(Z,\alpha(W))$ where α is a function that is the identity with probability at least $1-\varepsilon$ (over the choice of W), i.e., $\alpha(w)=w$ and with probability at most ε, it changes W so that it matches A. The latter is possible since $A \sim_{\varepsilon} W$.

Using the triangle inequality for computational indistinguishability (cf. the proof of Proposition 2) we get with (10) and (11)

$$(X,A) \sim_{2\varepsilon,s} (Z,A) \tag{12}$$

As $\widetilde{\mathsf{H}}_{\infty}(Z|W) \geq k$ (for α as defined above)

$$\widetilde{\mathsf{H}}_{\infty}(Z|W) \geq k \quad \Rightarrow \quad \widetilde{\mathsf{H}}_{\infty}(Z|\alpha(W)) \geq k \quad \Rightarrow \quad \widetilde{\mathsf{H}}_{\infty}(Z|A) \geq k \tag{13}$$

The first implication above holds as applying a function on the conditioned part cannot decrease the min-entropy. The second holds as $(Z,A) \sim (Z,\alpha(W))$. This concludes the proof as (12) and (13) imply that $\mathsf{H}^{\mathsf{HILL}}_{2\varepsilon,s}(X|A) \geq k$. □

As a corollary of Proposition 1 and Lemma 1, we get a chain rule for (non-conditional) HILL entropy. Such a chain rule has been shown by [10] and follows from the more general Dense Model Theorem (published at the same conference) of Reingold et al. [21].

Corollary 1. *For any distribution* $(X,A) \in \mathcal{X} \times \{0,1\}^{\ell}$ *and* $\hat{s} = \Omega\left(\frac{s\cdot\varepsilon^2}{2^{3\ell}\log^2(\ell)}\right)$

$$\mathsf{H}^{\mathsf{HILL}}_{\varepsilon,s}(X) \geq k \;\Rightarrow \mathsf{H}^{\mathsf{rlx\text{-}HILL}}_{2\varepsilon,\hat{s}}(X|A) \geq k-\ell \;\Rightarrow \mathsf{H}^{\mathsf{HILL}}_{4\varepsilon,\hat{s}}(X|A) \geq k-\ell$$

Note that unlike the chain rule for relaxed HILL given in Proposition 2, the chain rule for (standard) HILL given by the corollary above requires that we start with some non-conditional variable X. It would be preferable to have a chain rule for the conditional case, i.e., and expression of the form $\mathsf{H}^{\mathsf{HILL}}_{\varepsilon,s}(X|Y)=k \;\Rightarrow \mathsf{H}^{\mathsf{HILL}}_{\varepsilon',s'}(X|Y,A)=k-\ell$ for some $\varepsilon'=\varepsilon\cdot p(2^{\ell}), s'=s/q(2^{\ell},\varepsilon^{-1})$ (for polynomial functions $p(.),q(.)$), but as recently shown by Krenn et al. [19], such a chain rule does not hold (all we know is that such a rule holds if we also allow the security to degrade exponentially in the length $|Y|$ of the conditional part.)

4.3 Leakage-Resilient Cryptography

We now discuss how Theorem 1 can be used to simplify and quantitatively improve the security proofs for leakage-resilient cryptosystems. These proofs currently rely on chain rules for HILL entropy given in Corollary 1. As an illustrative example, we will reprove the security of the leakage-resilient stream-cipher based on any weak pseudorandom function from Eurocrypt'09 [20], but with much better bounds than the original proof.

For brevity, in this section we often write B^i to denote a sequence $B_1, \ldots, B_i$ of values. Moreover, $A \| B \in \{0,1\}^{a+b}$ denotes the concatenation of the strings $A \in \{0,1\}^a$ and $B \in \{0,1\}^b$.

A function $\mathsf{F}\colon \{0,1\}^k \times \{0,1\}^n \to \{0,1\}^m$ is an (ε, s, q)-**secure weak PRF** if its outputs on q random inputs look random to any size s distinguisher, i.e., for all D of size s

$$\left| \underset{K, X^q}{\mathbb{P}}[\mathsf{D}(X^q, \mathsf{F}(K, X_1), \ldots, \mathsf{F}(K, X_q) = 1] - \underset{X^q, R^q}{\mathbb{P}}[\mathsf{D}(X^q, R^q) = 1 \right| \leq \varepsilon,$$

where the probability is over the choice of the random $X_i \leftarrow \{0,1\}^n$, the choice of a random key $K \leftarrow \{0,1\}^k$ and random $R_i \leftarrow \{0,1\}^m$ conditioned on $R_i = R_j$ if $X_i = X_j$ for some $j < i$.

A **stream-cipher** $\mathsf{SC}\colon \{0,1\}^k \to \{0,1\}^k \times \{0,1\}^n$ is a function that, when initialized with a secret initial state $S_0 \in \{0,1\}^k$, produces a sequence of output blocks $X_1, X_2, \ldots$ recursively computed by

$$(S_i, X_i) := \mathsf{SC}(S_{i-1})$$

We say that SC is (ε, s, q)-secure if for all $1 \leq i \leq q$, no distinguisher of size s can distinguish X_i from a uniformly random $U_n \leftarrow \{0,1\}^n$ with advantage greater than ε given $X_1, \ldots, X_{i-1}$ (here, the probability is over the choice of the initial random key S_0)[6], i.e.,

$$\left| \underset{S_0}{\mathbb{P}}[\mathsf{D}(X^{i-1}, X_i) = 1] - \underset{S_0, U_n}{\mathbb{P}}[\mathsf{D}(X^{i-1}, U_n)] \right| \leq \varepsilon$$

A **leakage-resilient** stream-cipher is $(\varepsilon, s, q, \ell)$-secure if it is (ε, s, q)-secure as just defined, but where the distinguisher in the jth round not only gets X_j, but also ℓ bits of arbitrary adaptively chosen leakage about the secret state accessed during this round. More precisely, before $(S_j, X_j) := \mathsf{SC}(S_{j-1})$ is computed, the distinguisher can choose any leakage function f_j with range $\{0,1\}^\ell$, and then not only get X_j, but also $\Lambda_j := f_j(\hat{S}_{j-1})$, where $\hat{S}_{j-1}$ denotes the part of the secret state that was modified (i.e., read and/or overwritten) in the computation $\mathsf{SC}(S_{j-1})$.

Figure 1 illustrates the construction of a leakage-resilient stream cipher SC^{F} from any weak PRF $\mathsf{F}\colon \{0,1\}^k \times \{0,1\}^n \to \{0,1\}^{k+n}$ from [20]. The initial state is $S_0 = \{K_0, K_1, X_0\}$. Moreover, in the ith round (starting with round 0), one computes $K_{i+2} \| X_{i+1} := \mathsf{F}(K_i, X_i)$ and outputs X_{i+1}. The state after

[6] A more standard notion would require $X_1, \ldots, X_q$ to be indistinguishable from random; this notion is implied by the notion we use by a standard hybrid argument losing a multiplicative factor of q in the distinguishing advantage.

this round is $(K_{i+1}, K_{i+2}, X_{i+1})$.[7] In this section we will sketch a proof of the following security bound on SC^{F} as a leakage-resilient stream cipher in terms of the security of F as a weak PRF.

Lemma 2. *If* F *is a* $(\varepsilon_{\mathsf{F}}, s_{\mathsf{F}}, 2)$*-secure weak PRF then* SC^{F} *is a* $(\varepsilon', s', q, \ell)$*-secure leakage resilient stream cipher where*

$$\varepsilon' = 4q\sqrt{\varepsilon_{\mathsf{F}} 2^{\ell}} \qquad s' = \Theta(1) \cdot \frac{s_{\mathsf{F}}\varepsilon'^2}{2^{3\ell}}$$

The bound above is quantitatively much better than the one in [20]. Setting the leakage bound $\ell = \log \varepsilon_{\mathsf{F}}^{-1}/6$ as in [20], we get (for small q) $\varepsilon' \approx \varepsilon_{\mathsf{F}}^{5/12}$, which is by over a power of 5 better than the $\varepsilon_{\mathsf{F}}^{1/13}$ from [20], and the bound on $s' \approx s_{\mathsf{F}}\varepsilon_{\mathsf{F}}^{4/3}$ improves by a factor of $\varepsilon_{\mathsf{F}}^{5/6}$ (from $s_{\mathsf{F}}\varepsilon_{\mathsf{F}}^{13/6}$ in [20] to $s_{\mathsf{F}}\varepsilon_{\mathsf{F}}^{8/6}$ here). This improvement makes the bound meaningful if instantiated with a standard block-cipher like AES which has a keyspace of 256 bits, making the assumption that it provides $s_{\mathsf{F}}/\varepsilon_{\mathsf{F}} \approx 2^{256}$ security.[8]

Besides our main Theorem 1, we need another technical result which states that if F is a weak PRF secure against two queries, then its output on a single random query is pseudorandom, even if one is given some short auxiliary information about the uniform key K. The security of weak PRFs with non-uniform keys has first been proven in [20], but we will use a more recent and elegant bound from [1]. As a corollary of [1, Thm.3.7 in eprint version], we get that for any $(\varepsilon_{\mathsf{F}}, s_{\mathsf{F}}, 2)$-secure weak PRF $\mathsf{F}\colon \{0,1\}^k \times \{0,1\}^n \to \{0,1\}^m$, uniform and independent key and input $K \sim U_k, X \sim U_n$ and any (arbitrarily complex) function $g\colon \{0,1\}^k \to \{0,1\}^{\ell}$, one has[9]

$$(X, \mathsf{F}(K,X), g(K)) \sim_{\hat{\varepsilon}, s_{\mathsf{F}}/2} (X, U_m, g(K)) \text{ where } \hat{\varepsilon} = \varepsilon_{\mathsf{F}} + \sqrt{\varepsilon_{\mathsf{F}} 2^{\ell}} + 2^{-n} \approx \sqrt{\varepsilon_{\mathsf{F}} 2^{\ell}} \tag{14}$$

Generalizing the notation of $\sim_{\varepsilon,s}$ from variables to interactive distinguishers, given two (potentially stateful) oracles G, G', we write $G \sim_{\varepsilon,s} G'$ to denote that no oracle-aided adversary A of size s can distinguish G from G', i.e.,

$$G \sim_{\varepsilon,s} G' \iff \forall \mathsf{A}, |\mathsf{A}| \le s \ : \ |\mathbb{P}[\mathsf{A}^{G} \to 1] - \mathbb{P}[\mathsf{A}^{G'} \to 1]| \le \varepsilon.$$

Proof (of Lemma 2 (Sketch)). We define an oracle G_0^{real} that models the standard attack on the leakage-resilient stream cipher. That is, G_0^{real} samples a random initial state S_0. When interacting with an adversary $\mathsf{A}^{G_0^{real}}$, the oracle

[7] Note that X_i is not explicitly given as input to f_i even though the computation depends on X_i. The reason is that the adversary can choose f_i adaptively after seeing X_i, so X_i can be hard-coded it into f_i.

[8] We just need security against two random queries, so the well known non-uniform upper bounds on the security of block-ciphers of De, Trevisan and Tulsiani [6,5] do not seem to contradict such an assumption even in the non-uniform setting.

[9] The theorem implies a stronger statement where one only requires that K has $k - \ell$ bits average-case min-entropy (which is implied by having K uniform and leaking ℓ bits), we state this weaker statement as it is sufficient for our application.

G_0^{real} expects as input adaptively chosen leakage functions $f_1, f_2, \ldots, f_{q-1}$. On input f_i, it computes the next output block $(X_i, K_{i+1}) := \mathsf{SC}(K_{i-1}, X_{i-1})$ and the leakage $\Lambda_i = f_i(K_{i-1})$. It forwards X_i, Λ_i to A and deletes everything except the state $S_i = \{X_i, K_i, K_{i+1}\}$. After round $q-1$, G_0^{real} computes and forwards X_q (i.e., the next output block to be computed) to A. The game G_0^{rand} is defined in the same way, but the final block X_q is replaced with a uniformly random U_n.

To prove that SC^{F} is an $(\varepsilon', s', \ell, q)$-secure leakage-resilient stream cipher, we need to show that

$$G_0^{real} \sim_{\varepsilon', s'} G_0^{rand}, \tag{15}$$

for ε', s' as in the statement of the lemma.

Defining games G_i^{real} and G_i^{rand} for $1 \leq i \leq q-1$. We define a series of games $G_1^{real}, \ldots, G_{q-1}^{real}$ where G_{i+1}^{real} is derived from G_i^{real} by replacing X_i, K_{i+1} with uniformly random values $\tilde{X}_i, \tilde{K}_{i+1}$ and the leakage Λ_i with simulated fake leakage $\tilde{\Lambda}_i$ (the details are provided below). Games G_i^{rand} will be defined exactly as G_i^{real} except that (similarly to the case $i = 0$), the last block X_q is replaced with a uniformly random value.

For every i, $1 \leq i \leq q-1$, the variables $\tilde{K}_i, \tilde{X}_i$ as defined by the oracles realizing the games G_j^{rand} and G_j^{real} where $j \geq i$ will satisfy the following properties (as the initial values (X_0, K_0, K_1) never get replaced, for notational convenience we define $(\tilde{X}_0, \tilde{K}_0, \tilde{K}_1) \stackrel{\mathsf{def}}{=} (X_0, K_0, K_1)$)

i. $\tilde{K}_i, \tilde{X}_i$ are uniformly random.
ii. Right before the $(i-1)$th round (i.e. the round where the oracle computes $X_i \| K_{i+1} := \mathsf{F}(\tilde{X}_{i-1}, \tilde{K}_{i-1})$), the oracle has leaked no information about $\tilde{K}_{i-1}$ except for the ℓ bits fake leakage $\tilde{\Lambda}_i$.
iii. Right before the $(i-1)$th round $\tilde{K}_{i-1}$ and $\tilde{X}_{i-1}$ are independent given everything the oracle did output so far.

The first two properties above will follow from the definition of the games. The third point follows using Lemma 4 from [9], we will not discuss this here in detail, but only mention that the reason for the alternating structure of the cipher as illustrated in Figure 1, with an upper layer computing $K_0, K_2, \ldots$ and the lower layer computing $K_1, K_3, \ldots$, is to achieve this independence.

We now describe how the oracle G_{i+1}^{real} is derived from G_i^{real}. For concreteness, we set $i = 2$. In the third step, G_2^{real} computes $(X_3, K_4) := \mathsf{F}(\tilde{K}_2, \tilde{X}_2), \Lambda_3 = f_3(\tilde{K}_2)$ and forwards X_3, Λ_3 to A. The state stored after this step is $S_3 = \{X_3, \tilde{K}_3, K_4\}$. Let $V_2 \stackrel{\mathsf{def}}{=} \{\tilde{X}^2, \tilde{\Lambda}^2\}$ be the view (i.e. all the outputs she got from her oracle) of the adversary A after the second round.

Defining an intermediate oracle. We now define an oracle $G_{2/3}^{real}$ (which will be in-between G_2^{real} and G_3^{real}) derived from G_2^{real} by replacing $\Lambda_3 = f_3(\tilde{K}_2)$ with fake leakage $\tilde{\Lambda}_3$ computed as follows: let $h(\cdot)$ be a simulator for the leakage $\tilde{\Lambda}_3 := f_3(\tilde{K}_2)$ such that (for $\hat{\varepsilon}, \hat{s}$ to be defined)

$$(Z, h(Z)) \sim_{\hat{\varepsilon}, \hat{s}} (Z, \tilde{\Lambda}_3) \quad \text{where} \quad Z = \{V_2, X_3, K_4\} \tag{16}$$

By Theorem 1, there exists such a simulator of size $s_h \overset{\mathsf{def}}{=} O(\hat{s}2^{3\ell}/\hat{\varepsilon}^2)$. Note that h not only gets the pseudorandom output X_3, K_4 whose computation has leaked bits, but also the view V_2. The reason for the latter is that we need to fool an adversary who learned V_2. Equation (16) then yields

$$G_2^{real} \sim_{\hat{\varepsilon},\hat{s}-s_0} G_{2/3}^{real}, \tag{17}$$

where s_0 is the size of a circuit required to implement the real game G_0^{real}. The reason we loose s_0 in the circuit size here is that in a reduction where we use a distinguisher for G_2^{real} and $G_{2/3}^{real}$ to distinguish $(Z, h(Z))$ and $(Z, \tilde{\Lambda}_3)$ we must still compute the remaining $q-4$ rounds, and s_0 is an upper bound on the size of this computation.

The game G_3^{real} is derived from $G_{2/3}^{real}$ by replacing the values $X_3 \| K_4 := \mathsf{F}(\tilde{K}_2, \tilde{X}_2)$ with uniformly random $\tilde{X}_3 \| \tilde{K}_4$ right after they have been computed (let us stress that also the fake leakage that is computed as in (16) now uses these random values, i.e., $Z = \{V_2, \tilde{X}_3, \tilde{K}_4\}$).

Proving indistinguishability. We claim that the games are indistinguishable with parameters

$$G_{2/3}^{real} \sim_{\sqrt{\varepsilon_{\mathsf{F}} 2^{\ell}}, s_{\mathsf{F}}/2 - s_h - s_0} G_3^{real} \tag{18}$$

Recall that in $G_{2/3}^{real}$, we compute $X_3 \| K_4 := \mathsf{F}(\tilde{K}_2, \tilde{X}_2)$ where by i. $\tilde{X}_2, \tilde{K}_2$ are uniformly random, by ii. only ℓ bits of $\tilde{K}_2$ have leaked and iii. $\tilde{X}_2$ and $\tilde{K}_2$ are independent. Using these properties, equation (14) implies that the outputs are roughly $(\sqrt{\varepsilon_{\mathsf{F}} 2^{\ell}}, s_{\mathsf{F}}/2)$ pseudorandom, i.e.,

$$(\tilde{X}_2, X_3 \| K_4, \tilde{\Lambda}_1) \sim_{\sqrt{\varepsilon_{\mathsf{F}} 2^{\ell}}, s_{\mathsf{F}}/2} (\tilde{X}_2, \tilde{X}_3 \| \tilde{K}_4, \tilde{\Lambda}_1), \tag{19}$$

from which we derive (18). Note the loss of s_h in circuit size in equation (18) due to the fact that given a distinguisher for $G_{2/3}^{real}$ and G_3^{real}, we must recompute the fake leakage given only distributions as in (19).

We will assume that $s_0 \leq \hat{s}/2$, i.e., the real experiment is at most half as complex as the size of the adversaries we will consider (the setting where this is not the case is not very interesting anyway.) Then $\hat{s} - s_0 \geq \hat{s}/2$.

Up to this point, we have not yet defined what $\hat{\varepsilon}$ and $\hat{s}$ are, so we set them to

$$\hat{\varepsilon} \overset{\mathsf{def}}{=} \sqrt{\varepsilon_{\mathsf{F}} 2^{\ell}} \quad \text{and} \quad \hat{s} \overset{\mathsf{def}}{=} \Theta(1)\frac{s_{\mathsf{F}}\hat{\varepsilon}^2}{2^{3\ell}} \quad \text{then} \quad s_{\mathsf{F}} = 8 \cdot s_h = \Theta(1)\frac{\hat{s}2^{3\ell}}{\hat{\varepsilon}^2}.$$

With (17) and (18), we then get $G_2^{real} \sim_{2\hat{\varepsilon},\hat{s}/2} G_3^{real}$. The same proof works for any $1 \leq i \leq q-1$, i.e., we have

$$G_i^{real} \sim_{2\hat{\varepsilon},\hat{s}/2} G_{i+1}^{real}\ ,\ G_i^{rand} \sim_{2\hat{\varepsilon},\hat{s}/2} G_{i+1}^{rand}.$$

Moreover, using i.-iii. with (14),

$$G_{q-1}^{real} \sim_{2\hat{\varepsilon},\hat{s}/2} G_{q-1}^{rand}.$$

Using the triangle inequality $2q$ times, the two equations above yield

$$G_0^{real} \sim_{4q\hat{\varepsilon},\hat{s}/2} G_0^{rand},$$

which which completes the proof of the lemma.

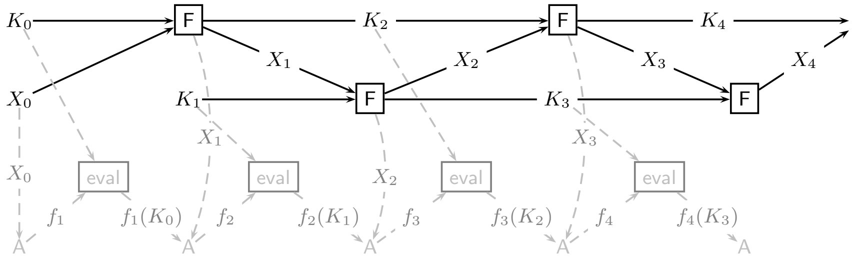

Fig. 1. Leakage resilient stream-cipher SC^{F} from a any weak pseudorandom function F. The regular evaluation is shown in black, the attack related part is shown in gray with dashed lines. The output of the cipher is $X_0, X_1, \ldots$.

References

1. Barak, B., Dodis, Y., Krawczyk, H., Pereira, O., Pietrzak, K., Standaert, F.X., Yu, Y.: Leftover hash lemma, revisited. In: CRYPTO 2011. pp. 1–20. LNCS, Springer (Aug 2011)
2. Barak, B., Hardt, M., Kale, S.: The uniform hardcore lemma via approximate bregman projections. In: Mathieu, C. (ed.) SODA. pp. 1193–1200. SIAM (2009)
3. Bellare, M., Rompel, J.: Randomness-efficient oblivious sampling. In: FOCS. pp. 276–287 (1994)
4. Chung, K.M., Lui, E., Pass, R.: From weak to strong zero-knowledge and applications. Cryptology ePrint Archive, Report 2013/260 (2013), `http://eprint.iacr.org/`
5. De, A., Trevisan, L., Tulsiani, M.: Non-uniform attacks against one-way functions and prgs. Electronic Colloquium on Computational Complexity (ECCC) 16, 113 (2009)
6. De, A., Trevisan, L., Tulsiani, M.: Time space tradeoffs for attacks against one-way functions and PRGs. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 649–665. Springer (Aug 2010)
7. 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 (Aug 2010)
8. Dodis, Y., Reyzin, L., Smith, A.: Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In: Cachin, C., Camenisch, J. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 523–540. Springer (May 2004)
9. Dziembowski, S., Pietrzak, K.: Intrusion-resilient secret sharing. In: 48th FOCS. pp. 227–237. IEEE Computer Society Press (Oct 2007)

10. Dziembowski, S., Pietrzak, K.: Leakage-resilient cryptography. In: 49th FOCS. pp. 293–302. IEEE Computer Society Press (Oct 2008)
11. Frieze, A.M., Kannan, R.: Quick approximation to matrices and applications. Combinatorica 19(2), 175–220 (1999)
12. Fuller, B., O'Neill, A., Reyzin, L.: A unified approach to deterministic encryption: New constructions and a connection to computational entropy. In: TCC 2012. pp. 582–599. LNCS, Springer (2012)
13. Gentry, C., Wichs, D.: Separating succinct non-interactive arguments from all falsifiable assumptions. In: Fortnow, L., Vadhan, S.P. (eds.) 43rd ACM STOC. pp. 99–108. ACM Press (Jun 2011)
14. Gowers, T.: Decompositions, approximate structure, transference, and the Hahn–Banach theorem. Bull. London Math. Soc. 42(4), 573–606 (2010)
15. 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)
16. Herbster, M., Warmuth, M.K.: Tracking the best linear predictor. Journal of Machine Learning Research 1, 281–309 (2001)
17. 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 (May 2007)
18. Impagliazzo, R.: Hard-core distributions for somewhat hard problems. In: FOCS. pp. 538–545 (1995)
19. Krenn, S., Pietrzak, K., Wadia, A.: A counterexample to the chain rule for conditional hill entropy - and what deniable encryption has to do with it. In: TCC. pp. 23–39 (2013)
20. Pietrzak, K.: A leakage-resilient mode of operation. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 462–482. Springer (Apr 2009)
21. Reingold, O., Trevisan, L., Tulsiani, M., Vadhan, S.P.: Dense subsets of pseudorandom sets. In: 49th FOCS. pp. 76–85. IEEE Computer Society Press (Oct 2008)
22. Reyzin, L.: Some notions of entropy for cryptography - (invited talk). In: ICITS. pp. 138–142 (2011), `http://www.cs.bu.edu/~reyzin/papers/entropy-survey.pdf`
23. Tao, T., Ziegler, T.: The primes contain arbitrarily long polynomial progressions. Acta Math. 201, 213–305 (2008)
24. Trevisan, L.: Guest column: additive combinatorics and theoretical computer science. SIGACT News 40(2), 50–66 (2009)
25. Trevisan, L., Tulsiani, M., Vadhan, S.P.: Regularity, boosting, and efficiently simulating every high-entropy distribution. In: IEEE Conference on Computational Complexity. pp. 126–136 (2009)
26. Vadhan, S.P., Zheng, C.J.: A uniform min-max theorem with applications in cryptography. In: CRYPTO (1). pp. 93–110 (2013)

A Proof of Theorem 1

We will prove Theorem 1 not for the family $\mathcal{F}$ directly, but for a family $\widehat{\mathcal{F}}$ which for every $f \in \mathcal{F}$ contains the function $\widehat{f} \colon \mathcal{X} \times \{0,1\}^\ell \to [-1,1]$ defined as

$$\widehat{f}(x,b) = f(x,b) - w_f(x) \quad \text{where} \quad w_f(x) = \mathop{\mathbb{E}}_{b \leftarrow \{0,1\}^\ell}[f(x,b)] = 2^{-\ell} \sum_{b \in \{0,1\}^\ell} f(x,b)$$

Any simulator that fools $\widehat{\mathcal{F}}$ also fools $\mathcal{F}$ with the same advantage since $\forall \widehat{f} \in \widehat{\mathcal{F}}$,

$$\begin{aligned} &\left| \mathop{\mathbb{E}}_{x \leftarrow X, [g]} [\widehat{f}(x, g(x))] - \mathop{\mathbb{E}}_{x \leftarrow X, [h]} [\widehat{f}(x, h(x))] \right| \\ = &\left| \mathop{\mathbb{E}}_{x \leftarrow X, [g]} [f(x, g(x)) - w_f(x)] - \mathop{\mathbb{E}}_{x \leftarrow X, [h]} [f(x, h(x)) - w_f(x)] \right| \\ = &\left| \mathop{\mathbb{E}}_{x \leftarrow X, [g]} [f(x, g(x))] - \mathop{\mathbb{E}}_{x \leftarrow X, [h]} [f(x, h(x))] \right| \end{aligned}$$

Evaluating $\widehat{f}$ requires 2^ℓ evaluations of f as we need to compute $w_f(x)$. We thus lose a factor of 2^ℓ in efficiency by considering $\widehat{\mathcal{F}}$ instead of $\mathcal{F}$. The reason that we prove the theorem for $\widehat{\mathcal{F}}$ instead of for $\mathcal{F}$ is because in what follows, we will need that for any x, the expectation over a uniformly random $b \in \{0,1\}^\ell$ is 0, i.e.,

$$\forall \widehat{f} \in \widehat{\mathcal{F}}, x \in \mathcal{X} \colon \mathop{\mathbb{E}}_{b \leftarrow \{0,1\}^\ell} [\widehat{f}(x, b)] = 0. \tag{20}$$

To prove the theorem, we must show that for any joint distribution $(X, g(X))$ over $\mathcal{X} \times \{0,1\}^\ell$, there exists an efficient simulator $h \colon \mathcal{X} \to \{0,1\}^\ell$ such that

$$\forall \widehat{f} \in \widehat{\mathcal{F}} \colon \left| \mathop{\mathbb{E}}_{x \leftarrow X} [\widehat{f}(x, g(x)) - \widehat{f}(x, h(x))] \right| < \varepsilon. \tag{21}$$

Moving to product form. We define the function $\widetilde{g} \colon \mathcal{X} \times \{0,1\}^\ell \to [0,1]$ as $\widetilde{g}(x, a) := \mathop{\mathbb{P}}_{[g]}[g(x) = a]$. Note that for every $x \in \mathcal{X}$, we have

$$\sum_{a \in \{0,1\}^\ell} \widetilde{g}(x, a) = 1. \tag{22}$$

We can write the expected value of $\widehat{f}(X, g(X))$ as follows:

$$\begin{aligned} \mathop{\mathbb{E}}_{x \leftarrow X, [g]} [\widehat{f}(x, g(x))] &= \sum_{a \in \{0,1\}^\ell} \mathop{\mathbb{E}}_{x \leftarrow X} \left[\widehat{f}(x, a) \mathop{\mathbb{P}}_{[g]}[g(x) = a] \right] = \\ &= \sum_{a \in \{0,1\}^\ell} \mathop{\mathbb{E}}_{x \leftarrow X} \left[\widehat{f}(x, a) \widetilde{g}(x, a) \right] = \\ &= 2^\ell \mathop{\mathbb{E}}_{x \leftarrow X, u \leftarrow \{0,1\}^\ell} [\widehat{f}(x, u) \widetilde{g}(x, u)]. \end{aligned} \tag{23}$$

We will construct a simulator $\widetilde{h} \colon \mathcal{X} \times \{0,1\}^\ell \to [0,1]$ such that for $\gamma > 0$ (to be defined later),

$$\forall \widehat{f} \in \widehat{\mathcal{F}} \colon \mathop{\mathbb{E}}_{x \leftarrow X, b \leftarrow \{0,1\}^\ell} [\widehat{f}(x, b)(\widetilde{g}(x, b) - \widetilde{h}(x, b))] < \gamma. \tag{24}$$

From this $\widetilde{h}$, we can then get a simulator $h(\cdot)$ like in (21) assuming that $\widetilde{h}(x,\cdot)$ is a probability distribution for all x, i.e., $\forall x \in \mathcal{X}$,

$$\sum_{b\in\{0,1\}^\ell} \widetilde{h}(x,b) = 1, \tag{25}$$

$$\forall b \in \{0,1\}^\ell \ : \ \widetilde{h}(x,b) \geq 0. \tag{26}$$

We will define a sequence $h_0, h_1, \ldots$ of functions where $h_0(x,b) = 2^{-\ell}$ for all x, b.[10] Define the energy function

$$\Delta_t = \mathop{\mathbb{E}}_{x\leftarrow X, b\leftarrow\{0,1\}^\ell} [(\widetilde{g}(x,b) - h_t(x,b))^2].$$

Assume that after the first t steps, there exists a function $\widehat{f}_{t+1} \colon \mathcal{X} \times \{0,1\}^\ell \to [-1,1]$ such that

$$\mathop{\mathbb{E}}_{x\leftarrow X, b\leftarrow\{0,1\}^\ell} [\widehat{f}_{t+1}(x,b)(g(x,b) - h_t(x,b))] \geq \gamma,$$

and define

$$h_{t+1}(x,b) = h_t(x,b) + \gamma \widehat{f}_{t+1}(x,b) \tag{27}$$

The energy function then decreases by γ^2, i.e.,

$$\begin{aligned}
&\Delta_{t+1} \\
&= \mathop{\mathbb{E}}_{x\leftarrow X, b\leftarrow\{0,1\}^\ell} [(\widetilde{g}(x,b) - h_t(x,b) - \gamma \widehat{f}_{t+1}(x,b))^2] = \\
&= \Delta_t + \underbrace{\mathop{\mathbb{E}}_{x\leftarrow X, b\leftarrow\{0,1\}\ell} [\gamma^2 \widehat{f}_{t+1}(x,b)]}_{\leq\gamma^2} - \underbrace{\mathop{\mathbb{E}}_{x\leftarrow X, b\leftarrow\{0,1\}^\ell} [2\gamma f_{t+1}(x,b)(\widetilde{g}(x,b) - h_t(x,b))]}_{\geq 2\gamma^2} \\
&\leq \Delta_t - \gamma^2.
\end{aligned}$$

Since $\Delta_0 \leq 1$, $\Delta_t \geq 0$ for any t (as it is a square) and $\Delta_i - \Delta_{i+1} \geq \gamma^2$, this process must terminate after at most $1/\gamma^2$ steps meaning that we have constructed $\widetilde{h} = h_t$ that satisfies (24). Note that the complexity of the constructed $\widetilde{h}$ is bounded by $2^\ell \gamma^{-2}$ times the complexity of the functions from $\mathcal{F}$ since, as mentioned earlier, computing $\widetilde{f}$ requires 2^ℓ evaluations of f. In other words, $\widetilde{h}$ has complexity $\mathcal{O}(2^\ell \gamma^{-2})$ relative to $\mathcal{F}$.

Moreover, since for all $x \in \mathcal{X}$ and $\widehat{f} \in \widehat{\mathcal{F}}$, we have $\sum\limits_{b\in\{0,1\}^\ell} h_0(x,b) = 1$ and $\sum\limits_{b\in\{0,1\}^\ell} \widehat{f}(x,b) = 0$, condition (25) holds as well. Unfortunately, (26) does not hold since it might be the case that $h_{t+1}(x,b) < 0$. We will explain later how to fix this problem by replacing $\widehat{f}_{t+1}$ in (27) with a similar function $\widehat{f}^*_{t+1}$ that

[10] It is not relevant how exactly h_0 is defined, but we need $\sum_{b\leftarrow\{0,1\}^\ell}[h_0(x,b)] = 1$ for all $x \in \mathcal{X}$.

satisfies $h_{t+1}(x,b) = h_t + \gamma \widehat{f}^*_{t+1} \geq 0$ for all x and b in addition to all of the properties just discussed. Assume for now that $\widetilde{h}$ satisfies (24)-(26).

Let $h\colon \mathcal{X} \to \{0,1\}^\ell$ be a probabilistic function defined as follows: we set $h(x) = b$ with probability $\widetilde{h}(x,b)$. Equivalently, imagine that we have a biased dice with 2^ℓ faces labeled by $b \in \{0,1\}^\ell$ such that the probability of getting the face with label b is $\widetilde{h}(x,b)$. We then define $h(x)$ by simply throwing this dice and reading off the label. It follows that $\underset{[h]}{\mathbb{P}}[h(x) = b] = \widetilde{h}(x,b)$. This probabilistic function satisfies

$$\begin{aligned}\underset{[h],x\leftarrow X}{\mathbb{E}}[\widehat{f}(x,h(x))] &= \underset{x\leftarrow X}{\mathbb{E}} \sum_{a\in\{0,1\}^\ell} \widehat{f}(x,a) \underset{[h]}{\mathbb{P}}[h(x) = a] \\ &= \underset{x\leftarrow X}{\mathbb{E}} \sum_{a\in\{0,1\}^\ell} \widehat{f}(x,a) h_t(x,a) \\ &= \underset{x\leftarrow X, u\leftarrow\{0,1\}^\ell}{\mathbb{E}} 2^\ell \widehat{f}(x,u) h_t(x,u). \qquad (28)\end{aligned}$$

Plugging (28) and (23) into (24), we obtain

$$\forall \widehat{f} \in \mathcal{F} \ : \ \underset{x\leftarrow X,[h]}{\mathbb{E}} \left[\frac{\widehat{f}(x,g(x))}{2^\ell} - \frac{\widehat{f}(x,h(x)]}{2^\ell} \right] < \gamma.$$

Equivalently,

$$\forall \widehat{f} \in \mathcal{F} \ : \ \underset{x\leftarrow X,[h]}{\mathbb{E}} \left[\widehat{f}(x,g(x)) - \widehat{f}(x,h(x)) \right] < \gamma 2^\ell \qquad (29)$$

We get (4) from the statement of the theorem by setting $\gamma := \varepsilon/2^\ell$. The simulator $\widetilde{h}$ is thus of complexity $\mathcal{O}(2^{3\ell}(1/\varepsilon)^2)$ relative to $\mathcal{F}$.

Enforcing $h_t(x,b) \geq 0$ *for* $\ell = 1$. We now fix the problem with the positivity of $h_t(x,b)$. Consider the case $\ell = 1$. Consider the following properties:

i. $\sum_{b\in\{0,1\}} h_t(x,b) = 1$ for $x \in \mathcal{X}$,

ii. $\forall b \in \{0,1\}, h_t(x,b) \geq 0$ for $x \in \mathcal{X}$,

iii. $\underset{x\leftarrow X,b\leftarrow\{0,1\}}{\mathbb{E}} [\widehat{f}_{t+1}(x,b)(g(x,b) - h_t(x,b))] \geq \gamma$ for $\gamma > 0$.

Assume that $h_t\colon \mathcal{X} \to \{0,1\}$ and $\widehat{f}_{t+1}\colon \mathcal{X} \times \{0,1\} \to [-1,1]$ satisfy *i*) and *ii*) for all $x \in \mathcal{X}$ and *iii*) for some $\gamma > 0$. Recall that $\Delta_t = \underset{x\leftarrow X,b\leftarrow\{0,1\}}{\mathbb{E}} [(\widetilde{g}(x,b) - h_t(x,b))^2]$. We have shown that $h_{t+1} = h_t + \gamma \widehat{f}_{t+1}$ satisfies

$$\Delta_{t+1} \leq \Delta_t - \gamma^2. \qquad (30)$$

Moreover, for all $x \in \mathcal{X}$, h_{t+1} will still satisfy *i*) but not necessarily *ii*). We define a function $\widehat{f}^*_{t+1}$ such that setting $h_{t+1} = h_t + \gamma \widehat{f}^*_{t+1}$ will satisfy *i*) and *ii*) for all $x \in \mathcal{X}$ and an inequality similar to (30).

First, for any $x \in \mathcal{X}$ for which condition *ii)* is satisfied, let $f^*_{t+1} = \widehat{f}_{t+1}$. Consider now $x \in \mathcal{X}$ for which *ii)* fails for some $b \in \{0,1\}$, i.e., for which $h_t(x,b) + \gamma \widehat{f}_{t+1}(x,b) < 0$. Let $\gamma' = -h_t(x,b)/f_{t+1}(x,b)$. Note that $0 \leq \gamma' \leq \gamma$ and $h_t(x,b) + \gamma' \widehat{f}_{t+1}(x,b) = 0$. Let

$$\widehat{f}^*_{t+1}(x,b) = \frac{\gamma'}{\gamma} \widehat{f}_{t+1}(x,b) \qquad \widehat{f}^*_{t+1}(x,1-b) = \widehat{f}_{t+1}(x,1-b) + \frac{1-\gamma'}{\gamma} \widehat{f}_{t+1}(x,b).$$

Let $h_{t+1}(x,\cdot) = h_t(x,\cdot) + \gamma \widehat{f}^*_{t+1}(x,\cdot)$ and note that

$$\sum_{b \in \{0,1\}} \widehat{f}^*_{t+1}(x,b) = \sum_{b \in \{0,1\}} \widehat{f}_{t+1}(x,b) = 0.$$

Condition *i)* is then satisfied for h_{t+1} for any $x \in \mathcal{X}$. By the definition of γ', condition *ii)* is satisfied for any $x \in \mathcal{X}$ as well. Condition *iii)* is more delicate and in fact need not hold. Yet, we will prove the following:

Lemma 3. *If $\widehat{f}_{t+1}$ and h_t satisfy i) and ii) for every $x \in \mathcal{X}$, and iii) then*

$$\begin{aligned} &\mathop{\mathbb{E}}_{x \leftarrow X, b \leftarrow \{0,1\}} [\widehat{f}_{t+1}(x,b)(g(x,b) - h_t(x,b))] \\ &- \mathop{\mathbb{E}}_{x \leftarrow X, b \leftarrow \{0,1\}} [\widehat{f}^*_{t+1}(x,b)(g(x,b) - h_t(x,b))] \leq \frac{\gamma}{4}. \end{aligned} \tag{31}$$

Proof. To prove (31), it suffices to show that for every $x \in \mathcal{X}$,

$$\sum_{b \in \{0,1\}} \widehat{f}_{t+1}(x,b)(g(x,b) - h_t(x,b)) - \sum_{b \in \{0,1\}} \widehat{f}^*_{t+1}(x,b)(g(x,b) - h_t(x,b)) \leq \frac{\gamma}{2}. \tag{32}$$

If $x \in \mathcal{X}$ is such that *ii)* is satisfied for h_{t+1} then there is nothing to prove. Suppose that *ii)* fails for some $x \in \mathcal{X}$ and $b \in \{0,1\}$. For brevity, let $f := \widehat{f}_{t+1}(x,b)$, $g := g(x,b)$, $h = h_t(x,b)$. We have $-1 \leq f < 0$, $h + \gamma f < 0$, $0 \leq g \leq 1$ and $h = -\gamma f^*$. Using $g - h \geq -h$, the left-hand side of (32) then satisfies

$$\begin{aligned} &2(f + h/\gamma)(g - h) \leq 2(f + h/\gamma)(-h) \\ &= \frac{2}{\gamma}(-f\gamma - h)h \leq \frac{2}{\gamma}\left(\frac{-f\gamma - h + h}{2}\right)^2 = \frac{\gamma f^2}{2} \leq \frac{\gamma}{2}, \end{aligned} \tag{33}$$

where we have used the inequality $uv \leq \left(\dfrac{u+v}{2}\right)^2$.

If *iii)* holds then Lemma 3 implies $\gamma - \mathop{\mathbb{E}}_{x \leftarrow X, b \leftarrow \{0,1\}} [\widehat{f}^*_{t+1}(x,b)(g(x,b) - h_t(x,b))] \leq \dfrac{\gamma}{4}$. Equivalently,

$$\mathop{\mathbb{E}}_{x \leftarrow X, b \leftarrow \{0,1\}} [\widehat{f}^*_{t+1}(x,b)(g(x,b) - h_t(x,b))] \geq \frac{3\gamma}{4}. \tag{34}$$

Defining $h_{t+1} = h_t + \gamma \widehat{f}^*_{t+1}$, we still get

$$\varDelta_{t+1} \leq \varDelta_t - \left(\frac{3\gamma}{4}\right)^2 = \varDelta_t - \frac{9\gamma^2}{16}. \tag{35}$$

Remark 2. In this case, the slightly worse inequality (35) will increase the complexity of $\widetilde{h}$, but only by a constant factor of 16/9, i.e., $\widetilde{h}$ will still have complexity $\mathcal{O}(2^\ell \gamma^{-2})$ relative to $\mathcal{F}$.

Enforcing $h_t(x, b) \geq 0$ for general ℓ. Let $\widehat{f}_{t+1}(x, b)$ be as before and suppose that there exists $x \in \mathcal{X}$ such that $h_t(x, b) + \gamma \widehat{f}_{t+1}(x, b) < 0$ for at least one $b \in \{0, 1\}^\ell$. We will show how to replace $\widehat{f}_{t+1}$ with another function $\widehat{f}^*_{t+1}$ such that it satisfies an inequality of type (34) and such that $h_{t+1}(x, b) = h_t(x, b) + \gamma \widehat{f}^*_{t+1}(x, b) \geq 0$. Let S be the set of all elements $b \in \{0, 1\}^\ell$ for which $h_t(x, b) + \gamma \widehat{f}_{t+1}(x, b) < 0$. For $b \in S$, it follows that $\widehat{f}_{t+1}(x, b) < 0$. As before, for $b \in S$, define $\widehat{f}^*_{t+1}(x, b) = -\dfrac{h_t(x, b)}{\gamma}$. Note that for each such b, we have added a positive mass $-\dfrac{h_t(x, b) + \gamma \widehat{f}_{t+1}(x, b)}{\gamma}$ to modify each $\widehat{f}_{t+1}(x, b)$. Let

$$M = \sum_{b \in S} - \left(\widehat{f}_{t+1}(x, b) + \frac{h_t(x, b)}{\gamma}\right) \tag{36}$$

be the total mass. For $b \notin S$, define $\widehat{f}^*_{t+1}(x, b) = \widehat{f}_{t+1}(x, b) - \dfrac{M}{2^\ell - s}$. Clearly, $\mathop{\mathbb{E}}_{b \leftarrow \{0,1\}^\ell} \widehat{f}^*_{t+1}(x, b) = 0$. We will now show the following

Lemma 4. *For every $x \in \mathcal{X}$, the function $\widehat{f}^*_{t+1}$ satisfies*

$$\sum_{b \in \{0,1\}^\ell} (\widehat{f}_{t+1}(x, b) - \widehat{f}^*_{t+1}(x, b))(g(x, b) - h_t(x, b)) < 2^{\ell-1}\gamma.$$

Proof. Let $s = |S|$ and $h_S = \sum_{i=1}^{s} h_t(x, b_i)$. First, note that (as in the case $\ell = 1$)

$$\begin{aligned} \forall b \in S\colon\ & \left(\widehat{f}_{t+1}(x, b) + \frac{h_t(x, b)}{\gamma}\right)(g(x, b) - h_t(x, b)) \\ & \leq - \left(\widehat{f}_{t+1}(x, b) + \frac{h_t(x, b)}{\gamma}\right) h_t(x, b). \end{aligned} \tag{37}$$

Moreover,

$$\sum_{b \notin S} g(x, b) \leq \sum_{b \in \{0,1\}^\ell} g(x, b) = 1. \tag{38}$$

The difference that we want to estimate is then

$$\begin{aligned}
\varDelta &= \sum_{b\in\{0,1\}^\ell} (\widehat{f}_{t+1}(x,b) - \widehat{f}^*_{t+1}(x,b))(g(x,b) - h_t(x,b)) \\
&= \sum_{b\in S} \left(\widehat{f}_{t+1}(x,b) + \frac{h_t(x,b)}{\gamma}\right)(g(x,b) - h_t(x,b)) + \frac{M}{2^\ell - s}\sum_{b\notin S}(g(x,b) - h_t(x,b)) \\
&\overset{(37),(38)}{\le} \sum_{b\in S} - \left(\widehat{f}_{t+1}(x,b) + \frac{h_t(x,b)}{\gamma}\right) h_t(x,b) + \frac{M}{2^\ell - s}\underbrace{\left(1 - \sum_{b\notin S} h_t(x,b)\right)}_{=h_S} \\
&\overset{(36)}{=} \sum_{b\in S}\underbrace{- \left(\widehat{f}_{t+1}(x,b) + \frac{h_t(x,b)}{\gamma}\right) h_t(x,b)}_{\le\gamma/4} + \frac{h_S}{2^\ell - s}\sum_{b\in S} - \left(\widehat{f}_{t+1}(x,b) + \frac{h_t(x,b)}{\gamma}\right) \\
&\overset{(33)}{\le} \frac{s\gamma}{4} - \frac{h_S}{2^\ell - s}\sum_{b\in S}\widehat{f}_{t+1}(x,b) - \frac{h_S^2}{\gamma(2^\ell - s)} = \frac{s\gamma}{4} + \frac{h_S f_S}{2^\ell - s} - \frac{h_S^2}{\gamma(2^\ell - s)},
\end{aligned}$$

where $f_S = -\sum_{b\in S}\widehat{f}_{t+1}(x,b)$. Note that $\sum_{b\in S} -\widehat{f}_{t+1}(x,b) \le s$ and (using (20)) $\sum_{b\in S} -\widehat{f}_{t+1}(x,b) = \sum_{b\notin S}\widehat{f}_{t+1}(x,b) \le 2^\ell - s$, i.e., $f_S \le \min\{s, 2^\ell - s\}$. Since

$$\begin{aligned}
\frac{h_S f_S}{2^\ell - s} - \frac{h_S^2}{\gamma(2^\ell - s)} &= \frac{1}{\gamma(2^\ell - s)} h_S(\gamma f_S - h_S) \\
&\le \frac{1}{\gamma(2^\ell - s)}\left(\frac{h_S + (\gamma f_S - h_S)}{2}\right)^2 \le \frac{s\gamma}{4},
\end{aligned}$$

where we have used that $f_S^2 \le s(2^\ell - s)$. Since $s < 2^\ell$, we obtain $\varDelta \le \frac{s\gamma}{2} < 2^{\ell-1}\gamma$ which proves the lemma.

To complete the proof, note that the above lemma implies that

$$\begin{aligned}
&\mathop{\mathbb{E}}_{x\leftarrow X, b\leftarrow\{0,1\}^\ell}[\widehat{f}_{t+1}(x,b)(g(x,b) - h_t(x,b))] \\
&\quad - \mathop{\mathbb{E}}_{x\leftarrow X, b\leftarrow\{0,1\}^\ell}[\widehat{f}^*_{t+1}(x,b)(g(x,b) - h_t(x,b))] < \frac{\gamma}{2},
\end{aligned}$$

and hence,

$$\mathop{\mathbb{E}}_{x\leftarrow X, b\leftarrow\{0,1\}^\ell}[\widehat{f}^*_{t+1}(x,b)(g(x,b) - h_t(x,b))] > \frac{\gamma}{2}. \tag{39}$$

Remark 3. Similarly, the slightly worse inequality (39) will increase the complexity of $\widetilde{h}$ by a constant factor of 4, i.e., $\widetilde{h}$ will still have complexity $\mathcal{O}(2^\ell\gamma^{-2})$ relative to $\mathcal{F}$.

A.1 Derandomizing $\widetilde{h}$

Next, we discuss how to derandomize $\widetilde{h}$. We can think of the probabilistic function $\widetilde{h}$ as a deterministic function $\widetilde{h}'$ taking two inputs where the second input represents the random coins used by $\widetilde{h}$. More precisely, for $R \leftarrow \{0,1\}^\rho$ (ρ is an upper bound on the number of random bits used by $\widetilde{h}$) and for any x in the support of X, we have $\widetilde{h}'(x,R) \sim \widetilde{h}(x)$.

To get our derandomized $\widehat{h}$, we replace the randomness R with the output of a function ϕ chosen from a family of t-wise independent functions for some large t, i.e., we set $\widehat{h}(x) = \widetilde{h}'(x,\phi(x))$. Recall that a family Φ of functions $\mathcal{A} \to \mathcal{B}$ is t-wise independent if for any t distinct inputs $a_1,\ldots,a_t \in \mathcal{A}$ and a randomly chosen $\phi \leftarrow \Phi$, the outputs $\phi(a_1),\ldots,\phi(a_t)$ are uniformly random in $\mathcal{B}^t$. In the proof, we use the following tail inequality for variables with bounded independence:

Lemma 5 (Lemma 2.2 from [3]). *Let $t \geq 6$ be an even integer and let $Z_1,\ldots,Z_n$ be t-wise independent variables taking values in $[0,1]$.*
Let $Z = \sum_{i=1}^n Z_i$, then for any $A > 0$

$$\mathbb{P}[|Z - \mathbb{E}[Z]| \geq A] \leq \left(\frac{nt}{A^2}\right)^{t/2}$$

Recall that the min-entropy of X is $H_\infty(X) = -\log\left(\max_x \mathbb{P}[X=x]\right)$, or equivalently, X has min-entropy k if $\mathbb{P}[X=x] \leq 2^{-k}$ for all $x \in \mathcal{X}$.

Lemma 6. (Deterministic Simulation) *Let $\varepsilon > 0$ and assume that*

$$H_\infty(X) > 2 + \log\log|\mathcal{F}| + 2\log(1/\varepsilon). \tag{40}$$

For any (probabilistic) $\widetilde{h}\colon \mathcal{X} \to \{0,1\}^\ell$, there exists a deterministic *$\widehat{h}$ of the same complexity relative to $\mathcal{F}$ as $\widetilde{h}$ such that*

$$\forall f \in \mathcal{F}\colon \left| \mathop{\mathbb{E}}_{x\leftarrow X,[\widetilde{h}]}[f(x,\widetilde{h}(x))] - \mathop{\mathbb{E}}_{x\leftarrow X}[f(x,\widehat{h}(x))] \right| < \varepsilon \tag{41}$$

Remark 4. **About the condition (40).** A lower bound on the min-entropy of X in terms of $\log\log|\mathcal{F}|$ and $\log(1/\varepsilon)$ as in (40) is necessary. For example one can show that for $\varepsilon < 1/2$, (41) implies $H_\infty(X) \geq \log\log|\mathcal{F}|$. To see this, consider the case when X is uniform over $\{0,1\}^m$ (so $H_\infty(X) = m$), $\mathcal{F}$ contains all 2^{2^m} functions $f\colon \{0,1\}^m \times \{0,1\} \to \{0,1\}$ satisfying $f(x,1-b) = 1 - f(x,b)$ for all $x,b \in \{0,1\}^{m+1}$, and $\widetilde{h}(x) \sim U_1$ is uniformly random for all x (so it ignores its input). Now, given any deterministic $\widehat{h}$, we can choose $f \in \mathcal{F}$ where $f(x,\widehat{h}(x)) = 1$ for all $x \in \{0,1\}^m$ (such an f exists by definition of $\mathcal{F}$). For this f,

$$\left| \underbrace{\mathop{\mathbb{E}}_{x\leftarrow X,[\widetilde{h}]}[f(x,\widetilde{h}(x))]}_{=1/2} - \underbrace{\mathop{\mathbb{E}}_{x\leftarrow X}[f(x,\widehat{h}(x))]}_{=1} \right| = 1/2.$$

In terms of $\log(1/\varepsilon)$, one can show that (41) implies $H_\infty(X) \geq \log(1/\varepsilon) - 1$ (even if $|\mathcal{F}| = 1$). For this, let $\widetilde{h}$ and X be as above, $\mathcal{F} = \{f\}$ is defined as $f(x, b) = b$ if $x = 0^m$ and $f(x, b) = 0$ otherwise. For any deterministic $\widehat{h}$, we get

$$|\underbrace{\mathop{\mathbb{E}}_{x \leftarrow X, [\widetilde{h}]}[f(x, \widetilde{h}(x))]}_{1/2^{m+1}} - \underbrace{\mathop{\mathbb{E}}_{x \leftarrow X}[f(x, \widehat{h}(x))]}_{1/2^m \text{ or } 0}| = 1/2^{m+1}$$

and thus, $\varepsilon = 1/2^{m+1}$. Equivalently $H_\infty(X) = m = \log(1/\varepsilon) - 1$. The condition (40) is mild and in particular, it covers the cryptographically interesting case where $\mathcal{F}$ is the family of polynomial-size circuits (i.e., for a security parameter n and a constant c, $|\mathcal{F}| \leq 2^{n^c}$), X has superlogarithmic min-entropy $H_\infty(X) = \omega(\log n)$ and $\varepsilon > 0$ is negligible in n. Here, (40) becomes

$$\omega(\log n) > 2 + c \log n + 2 \log \varepsilon^{-1}$$

which holds for a negligible $\varepsilon = 2^{-\omega(\log n)}$.

Proof (Proof of Lemma 5). Let $m = H_\infty(X)$. We will only prove the lemma for the restricted case where X is flat, i.e., it is uniform on a subset $\mathcal{X}' \subseteq \mathcal{X}$ of size 2^m. [11] Consider any fixed $f \in \mathcal{F}$ and the 2^m random variables $Z_x \in \{0, 1\}$ indexed by $x \in \mathcal{X}'$ sampled as follows: first, sample $\phi \leftarrow \Phi$ from a family of t-wise independent functions $\mathcal{X} \to \{0, 1\}^\rho$ (recall that ρ is a upper bound on the number of random bits used by $\widetilde{h}$). Now, Z_x is defined as

$$Z_x = f(x, \widetilde{h}'(x, \phi(x))) = f(x, \widehat{h}(x))$$

and $Z = \sum_{x \in \mathcal{X}'} Z_x$. Note that the same ϕ is used for all Z_x.

1. The variables Z_x for $x \in \mathcal{X}'$ are t-wise independent, i.e., for any t distinct $x_1, \ldots, x_t$, the variables $Z_{x_1}, \ldots, Z_{x_t}$ have the same distribution as $Z'_{x_1}, \ldots, Z'_{x_t}$ sampled as $Z'_{x_i} \leftarrow f(x_i, \widetilde{h}'(x_i, R))$. The reason is that the randomness $\phi(x_1), \ldots, \phi(x_t)$ used to sample the $Z_{x_1}, \ldots, Z_{x_t}$ is uniform in $\{0, 1\}^\rho$ as ϕ is t-wise independent.
2. $\mathbb{E}[Z_x] = \mathop{\mathbb{E}}_{\phi \leftarrow \Phi}[f(x, \widetilde{h}'(x, \phi(x))] = \mathop{\mathbb{E}}_{[\widetilde{h}]}[f(x, \widetilde{h}(x))]$.
3. $\mathop{\mathbb{P}}_{x \leftarrow X, [\widetilde{h}]}[f(x, \widetilde{h}(x)) = 1] = \mathop{\mathbb{E}}_{\phi \leftarrow \Phi}[Z/2^m]$.

[11] Any distribution satisfying $H_\infty(X) = m$ can be written as a convex combination of flat distributions with min-entropy m. Often, this fact is sufficient to conclude that a result proven for flat distributions with min-entropy m implies the result for any distribution with the same min-entropy. Here, this is not quite the case, because we might end up using a different ϕ for every flat distribution. But as the only property we actually require from X is $\mathbb{P}[X = x] \leq 2^{-m}$, the proof goes through for general X, but becomes somewhat more technical.

Let $\mu = \underset{\phi \leftarrow \Phi}{\mathbb{E}}[Z] = \underset{\phi \leftarrow \Phi}{\mathbb{E}}\left[\sum_{x \in \mathcal{X}'} Z_x\right]$. By Lemma 5, we have

$$\mathbb{P}[|Z - \mu| \geq \varepsilon 2^m] \leq \left(\frac{t}{\varepsilon^2 2^m}\right)^{t/2}.$$

Let us call ϕ bad for f if $|Z - \mu| \geq \varepsilon 2^m$ (or equivalently, using iii)),

$$\left| \underset{x \leftarrow X, [\widetilde{h}]}{\mathbb{E}}[f(x, \widetilde{h}(x))] - \underset{x \leftarrow X}{\mathbb{E}}[f(x, \phi(x))] \right| \geq \varepsilon$$

We want to choose t such that the probability of ϕ being bad for any particular $f \in \mathcal{F}$ is less than $1/|\mathcal{F}|$, i.e.

$$(t/\varepsilon^2 2^m)^{t/2} < |\mathcal{F}|^{-1}. \tag{42}$$

We postpone for a second how to choose t and discussing when this is even possible. Assuming (42),

$$\underset{\phi \leftarrow \Phi}{\mathbb{P}}[|Z - \mu| \geq \varepsilon 2^m] \leq \left(\frac{t}{\varepsilon^2 2^m}\right)^{t/2} < |\mathcal{F}|^{-1},$$

and by taking a union bound over all $f \in \mathcal{F}$, we get

$$\underset{\phi \leftarrow \Phi}{\mathbb{P}}[\exists f \in \mathcal{F} \ : \ |Z - \mu| \geq \varepsilon 2^m] < 1,$$

which implies that there exits $\phi \in \Phi$ such that

$$\forall f \in \mathcal{F} \ : \ |Z - \mu| < \varepsilon 2^m.$$

Equivalently, using how Z and μ were defined,

$$\forall f \in \mathcal{F} \ : \ \left| \sum_{x \in \mathcal{X}'} f(x, \widehat{h}(x)) - \sum_{x \in \mathcal{X}'} f(x, \widetilde{h}(x)) \right| < \varepsilon 2^m.$$

Finally, using that X is uniform over $\mathcal{X}'$, we get (for the above choice of ϕ) the statement of the lemma

$$\forall f \in \mathcal{F} \ : \ \left| \underset{x \leftarrow X}{\mathbb{E}}[f(x, \widehat{h}(x))] - \underset{x \leftarrow X, [\widetilde{h}]}{\mathbb{E}}[f(x, \widetilde{h}(x))] \right| < \varepsilon.$$

We still have to determine when t can be chosen so that (42) holds. By taking logarithm and rearranging the terms, (42) becomes

$$mt/2 > \log|\mathcal{F}| + (t/2)\log(t) + t\log(1/\varepsilon),$$

i.e.,

$$m > 2\log|\mathcal{F}|/t + \log(t) + 2\log(1/\varepsilon).$$

Setting $t = \log|\mathcal{F}|$, we get

$$m > 2 + \log\log|\mathcal{F}| + 2\log(1/\varepsilon).$$

which holds as it is the condition (5) we made on the min-entropy $m = H_\infty(X)$.