Security Amplification for the Cascade of Arbitrarily Weak PRPs: Tight Bounds via the Interactive Hardcore Lemma

Stefano Tessaro

Department of Computer Science and Engineering
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0404, USA
stessaro@cs.ucsd.edu

Abstract. We consider the task of amplifying the security of a weak pseudorandom permutation (PRP), called an ε-PRP, for which the computational distinguishing advantage is only guaranteed to be bounded by some (possibly non-negligible) quantity $\varepsilon < 1$. We prove that the cascade (i.e., sequential composition) of m ε-PRPs (with independent keys) is an $((m-(m-1)\varepsilon)\varepsilon^m + \nu)$-PRP, where ν is a negligible function. In the asymptotic setting, this implies security amplification for all $\varepsilon < 1 - \frac{1}{\mathsf{poly}}$, and the result extends to two-sided PRPs, where the inverse of the given permutation is also queried. Furthermore, we show that this result is essentially tight. This settles a long-standing open problem due to Luby and Rackoff (STOC '86).

Our approach relies on the first hardcore lemma for computational indistinguishability of *interactive* systems: Given two systems whose states do not depend on the interaction, and which no efficient adversary can distinguish with advantage better than ε, we show that there exist events on the choices of the respective states, occurring each with probability at least $1-\varepsilon$, such that the two systems are computationally indistinguishable conditioned on these events.

1 Introduction

1.1 Motivation: Weak PRPs

The security of several cryptographic schemes relies on the assumption that an underlying block cipher is a *pseudorandom permutation* (PRP), a keyed family of permutations $E = \{E_k\}_{k\in\mathcal{K}}$ with the following property: any computationally bounded distinguisher can only decide with negligible advantage over random guessing whether it is given access to E_K (under a random secret key K) or to a uniformly chosen permutation with the same domain.

However, pseudorandomness is a very strong requirement, and continuous progress in cryptanalysis raises some doubts as to whether block-cipher designs such as the *Advanced Encryption Standard* (AES) are indeed secure PRPs. It

Y. Ishai (Ed.): TCC 2011, LNCS 6597, pp. 37–54, 2011.
© International Association for Cryptologic Research 2011

is therefore a prudent approach, as well as a central question in theoretical cryptography, to investigate *weaker* assumptions on a block cipher which are sufficient to efficiently solve a certain cryptographic task at hand.

A natural weakening of a PRP, considered in this paper, is to only require that the best advantage of a computationally restricted distinguisher is bounded by some given quantity $\varepsilon < 1$; we refer to such a primitive as an ε-PRP. In particular, in the asymptotic setting, ε is not required to be a negligible function. Instead, it may be a constant or even moderately converge to one as a function of the security parameter. For instance, common sense dictates that AES is much more likely to be a 0.99-PRP, rather than a fully secure PRP.

1.2 Our Result: Security Amplification of PRPs by Cascading

We investigate the natural and central problem of finding an efficient construction of a fully secure PRP (i.e., a δ-PRP for a negligible δ) from any ε-PRP $E = \{E_k\}_{k\in\mathcal{K}}$. Such constructions should work for arbitrary $\varepsilon < 1$ and call E as few times as possible (ideally, $\log(1/\delta) \cdot (\log(1/\varepsilon))^{-1}$ times to implement a δ-PRP). This is in the same spirit of the long line of research devoted to security amplification initiated by Yao [16] in the context of one-way functions.

The most natural approach is the *m-fold cascade*, a construction which outputs the value

$$(E_{k_1} \circ \cdots \circ E_{k_m})(x)$$

on input x and keys $k_1, \ldots, k_m$ (which are chosen independently). Here, $\circ$ denotes (sequential) composition of permutations.

Despite its simplicity, proving security amplification for the cascade has been a long-standing open problem. On the one hand, Luby and Rackoff [6] and Myers [12] showed that the c-fold cascade is a $((2-\varepsilon)^{c-1}\varepsilon^c + \nu)$-PRP for any *constant* c, where ν is a negligible additive term, but their results fall short of implying that a sufficiently long cascade yields a fully secure PRP for a non-negligible ε. On the other hand, Maurer and Tessaro [10] showed that the cascade of *arbitrary* (polynomial) length m is a $(2^{m-1}\varepsilon^m + \nu)$-PRP, but their bound, which only implies security amplification when $\varepsilon < \frac{1}{2}$, is clearly not tight in view of the superior result for the constant-length case [6,12].

Our Result on Cascades. This paper closes this gap by providing an exact characterization of the security amplification properties of the cascade: We prove that the cascade of m ε-PRPs (with domain $\mathcal{X}$) is a $((m-(m-1)\varepsilon)\varepsilon^m + \nu)$-PRP, i.e., it is security amplifying for essentially any $\varepsilon < 1 - \frac{1}{|\mathcal{X}|}$.[1] The result extends to *two-sided* ε-PRPs, where the inverse can also be queried, and is shown to be nearly tight. Also, this result arises from the application of new *generic* techniques of independent interest, illustrated in the next section.

[1] This restriction is necessary, as an ε-PRP with a fixed point (independent of the key value) can satisfy $\varepsilon = 1 - \frac{1}{|\mathcal{X}|}$, and the cascade obviously preserves such a fixed point.

Further Related Work. Less efficient constructions of fully secure PRPs from ε-PRPs exist: Maurer and Tessaro [11] showed that XORing two independent keys at both ends of the cascade yields an $(\varepsilon^m + \nu)$-PRP. Alternatively, techniques [13,2,10] for strengthening the security of pseudorandom *functions* (PRF) can be used in conjunction with known PRF-to-PRP conversion techniques such as [7]. However, this paper improves on these works in at least two respects: First, we show that similar amplification is achieved with better efficiency by the most natural construction. Second, our approach entails a new set of generic techniques which promise to be applicable in a wider context.

Additionally, let us point out that cascades have been studied in other contexts and models; a comprehensive discussion is deferred to the full version due to space constraints.

1.3 Our General Paradigm: The Interactive Hardcore Lemma and High-Entropy Permutations

The following technique is well known in the study of random processes: One can always define events $\mathcal{A}$ and $\mathcal{B}$ on any two finite random variables X and Y, by means of conditional probability distributions $\mathsf{P}_{\mathcal{A}|X}$ and $\mathsf{P}_{\mathcal{B}|Y}$, such that:

(i) X and Y are equally distributed conditioned on the respective events, i.e., $\mathsf{P}_{X|\mathcal{A}} = \mathsf{P}_{Y|\mathcal{B}}$,
(ii) $\mathsf{P}[\mathcal{A}] = \mathsf{P}[\mathcal{B}] = 1 - d(X, Y)$, where $d(X, Y)$ is the so called *statistical distance*, which equals the best advantage of a computationally *unbounded* distinguisher in distinguishing X and Y.

A computational version of this statement is due to Maurer and Tessaro [11], and was used to prove security amplification results for PRGs. In this paper, we take this approach one step further by presenting a computational version of the above statement for discrete *interactive* systems.

CC-Stateless Systems. We consider the general class of *convex-combination stateless* (or simply *cc-stateless*) interactive systems [10]. Like most cryptographic systems of interest, these systems have the property that the answer of each query *can be seen* as depending solely on the query input and on an *initial state*, but does not depend on previous queries and their answers. A simple example is the cc-stateless system implementing a permutation E_K for a keyed family of permutations $\{E_k\}_{k\in\mathcal{K}}$ and a uniform random key $K \in \mathcal{K}$. A further example is a *uniform random permutation* (URP) $\mathbf{P}$ on a set $\mathcal{X}$, a system choosing a permutation $P : \mathcal{X} \to \mathcal{X}$ uniformly at random, and answering each query x as $P(x)$. Moreover, a randomized encryption scheme where each encryption depends on the random key and some fresh randomness is also cc-stateless.

We stress that being cc-stateless is a property of the input-output behavior of a system, rather than of its actual implementation: Indeed, any implementation using such an initial state may be inefficient (e.g., due to its large size), but at the same time an efficient implementation of a cc-stateless system may be fully stateful. For example, an efficient implementation of a URP keeps an interaction-dependent state (in form of a table of all input-output pairs associated with all

previous queries) and employs lazy sampling, returning for each new query a uniformly distributed value among those not returned yet.

The Hardcore Lemma. Our main technical tool is the *Hardcore Lemma (HCL) for computational indistinguishability* (Theorem 2): Informally, it states that if all computationally bounded distinguishers only achieve advantage at most ε in distinguishing two cc-stateless systems $\mathbf{S}$ and $\mathbf{T}$, then there exist events $\mathcal{A}$ and $\mathcal{B}$, defined on the respective initial states of (the cc-stateless representations of) $\mathbf{S}$ and $\mathbf{T}$, such that the following holds:

(i) The (cc-stateless representations of the) systems $\mathbf{S}$ and $\mathbf{T}$ are computationally indistinguishable conditioned on the respective events $\mathcal{A}$ and $\mathcal{B}$.
(ii) Both events occur with probability at least $1 - \varepsilon$.

In addition, some applications of the HCL require the ability to efficiently simulate $\mathbf{S}$ and $\mathbf{T}$ under the assumption that the associated events occur (or do not occur), possibly with the help of some short (but not necessarily efficiently-samplable[2]) advice. In general, it is unclear whether this is possible given any two events satisfying (i) and (ii), even if both systems are efficiently implementable.

As an illustrative example, let $\mathbf{S} = E_K$ and $\mathbf{T} = \mathbf{P}$, where $K \in \mathcal{K}$ is uniformly distributed, $E = \{E_k\}_{k\in\mathcal{K}}$ is an efficiently computable family of permutations, and $\mathbf{P}$ is a URP (all permutations are on the n-bit strings). If E is an ε-PRP, the HCL yields an event $\mathcal{A}$ defined on K and an event $\mathcal{B}$ defined on a uniformly chosen permutation table P, both occurring with probability at least $1 - \varepsilon$, such that $E_{K'}$ (for K' sampled from $\mathsf{P}_{K|\mathcal{A}}$) and a system $\mathbf{P}'$ (implementing a permutation table P' sampled from $\mathsf{P}_{P|\mathcal{B}}$) are computationally indistinguishable. While $E_{K'}$ is efficiently implementable given K', a representation of P' requires $2^{\Theta(n)}$ bits, and it is unclear how to define a short advice (i.e., with length $\mathsf{poly}(n)$) that can be used to efficiently simulate $\mathbf{P}'$. However, quite surprisingly, we will show that one can always find events with short advice as long as $\mathbf{S}$ and $\mathbf{T}$ are efficiently implementable. This will be the major challenge in proving the HCL for the interactive setting.

The core of our proof is a tight generalization (Theorem 1) of Impagliazzo's HCL [5] to the setting of guessing a random bit given access to some interactive system whose behavior is correlated with the bit value.

Cascade of Permutations with High Min-Entropy. We briefly illustrate how the HCL is used to prove our bounds for the cascade of ε-PRPs. The main observation is that P' as above has *min-entropy* at least

$$\begin{aligned}\mathsf{H}_\infty(P') &= \log\left(\min_\pi \frac{1}{\mathsf{P}\left[P' = \pi\right]}\right) \\ &= \log\left(\min_\pi \frac{\mathsf{P}\left[\mathcal{B}\right]}{\mathsf{P}\left[P = \pi\right] \cdot \mathsf{P}\left[\mathcal{B} \,\middle|\, P = \pi\right]}\right) \geq \log\left(2^n!\right) - \log\left(\frac{1}{1-\varepsilon}\right),\end{aligned}$$

[2] For now, we only consider the non-uniform setting, thus efficient samplability is not a requirement.

i.e., at most $\log\left((1-\varepsilon)^{-1}\right)$ away from the maximum achievable min-entropy. This gap potentially makes $\mathbf{P}'$ easily distinguishable from a URP. However, we prove (Theorem 3) that the cascade of (at least) two such permutations is indistinguishable from a URP for computationally unbounded distinguishers making at most an exponential number of queries and even when allowing queries to the inverse. (The proof uses techniques from the random systems framework [8], and is of independent interest.)

The main security amplification result (Theorem 4) follows from the observation that by the above at least two (independent) permutations E_{K_i} and E_{K_j} (for $i \neq j$) in the cascade $E_{K_1} \circ \cdots \circ E_{K_m}$ (for independent keys $K_1, \ldots, K_m$) are computationally indistinguishable from $\mathbf{P}'$, except with probability $\varepsilon^m + m(1-\varepsilon)\varepsilon^{m-1}$, and in this case the cascade is computationally indistinguishable from a URP by Theorem 3. The final bound follows from a more fine-grained analysis.

Uniform vs. Non-Uniform Proofs. The results of this paper are formulated in a concrete, non-uniform, computational model. This simplifies the presentation considerably and helps conveying the main ideas. In the full version, we highlight the changes required in order to obtain uniform statements and proofs.

2 Preliminaries

Calligraphic letters $\mathcal{X}, \mathcal{Y}, \ldots$ denote sets and events, upper-case letters $X, Y, \ldots$ random variables (with expected values $\mathsf{E}[X], \mathsf{E}[Y], \ldots$), and lower-case letters $x, y, \ldots$ the values they take. Moreover, $\mathsf{P}[\mathcal{A}]$ is the probability of an event $\mathcal{A}$ (we denote as $\overline{\mathcal{A}}$ its complement) and we use the shorthands $\mathsf{P}_X(x) := \mathsf{P}[X = x]$, $\mathsf{P}_{X|Y}(x,y) := \mathsf{P}[X = x | Y = y]$, $\mathsf{P}_{X\mathcal{A}|Y\mathcal{B}}(x,y) := \mathsf{P}[\mathcal{A} \wedge X = x | \mathcal{B} \wedge Y = y]$, etc. Also, P_X, $\mathsf{P}_{X|Y}$, $\mathsf{P}_{\mathcal{A}X|\mathcal{B}Y}$ denote the corresponding (conditional) probability distributions,[3] and $x \xleftarrow{\$} \mathsf{P}_X$ is the action of sampling a value x with distribution P_X. (We use $x \xleftarrow{\$} \mathcal{S}$ to denote the special case where x is drawn uniformly from a finite set $\mathcal{S}$.) The *statistical distance* $d(X,Y)$ (or $d(\mathsf{P}_X, \mathsf{P}_Y)$) of X and Y (both with range $\mathcal{S}$) is defined as $d(X,Y) := \frac{1}{2}\sum_{x\in\mathcal{S}} |\mathsf{P}_X(x) - \mathsf{P}_Y(x)| = \sum_{x:\mathsf{P}_X(x)\geq\mathsf{P}_Y(x)} (\mathsf{P}_X(x) - \mathsf{P}_Y(x))$. Also, recall that a function is *negligible* if it vanishes faster than the inverse of any polynomial.

Computational Model. We consider *interactive* randomized stateful algorithms in some a-priori fixed RAM model of computation. Such an algorithm keeps a state (consisting, say, of the contents of the memory space it employs), and answers each query depending on the input of this query, some coin flips, the current state (which may be updated), and (possibly) one or more queries to an underlying system. It is also convenient to denote by $A[\sigma]$ the algorithm obtained by *setting* the state of A to σ (provided σ is a compatible state), and then behaving according to A's description. We say that A has *time complexity* t_A (where t_A is a function $\mathbb{N} \times \mathbb{N} \to \mathbb{N}$) if the sum of the length of the description of A, of s, and the total number of steps of A is at most $t_A(q,s)$ for all sequences of

[3] In particular, $\mathsf{P}_{X|Y}$ and $\mathsf{P}_{\mathcal{A}X|\mathcal{B}Y}$ take *two* arguments corresponding to all possible values of X and Y, respectively.

q queries, all compatible initial states with size s, and all compatible interactions with an underlying system. We use the shorthand $t_A(q) := t_A(q, 0)$. Furthermore, $s_A : \mathbb{N} \to \mathbb{N}$ is the *space complexity* of A, where $s_A(q)$ is the worst-case amount of memory used by A when answering any q queries.

Systems and Distinguishers. This paper considers abstract discrete interactive systems [8], denoted by bold-face letters $\mathbf{S}, \mathbf{T}, \ldots$, taking as inputs queries $X_1, X_2, \ldots$ and returning outputs $Y_1, Y_2, \ldots$. Such systems may be implemented by an interactive algorithm A (in which case we sometimes write A as a placeholder for the system it implements to explicit this fact), but may also arise from an arbitrary random process. The *input-output behavior* of the system $\mathbf{S}$ is fully described by the (infinite) family of conditional probability distributions $\mathsf{p}^{\mathbf{S}}_{Y_i|X^iY^{i-1}}$ (for $i \geq 1$) of the i-th output Y_i given the first i queries $X^i = [X_1, \ldots, X_i]$, and the first $i-1$ outputs $Y^{i-1} = [Y_1, \ldots, Y_{i-1}]$. In general, every statement that involves a system $\mathbf{S}$ holds for any realization of the system $\mathbf{S}$, i.e., it only depends on its input-output behavior. In particular, we say that two systems $\mathbf{S}$ and $\mathbf{T}$ are *equivalent*, denoted $\mathbf{S} \equiv \mathbf{T}$, if they have the same input-output behavior, i.e., $\mathsf{p}^{\mathbf{S}}_{Y_i|X^iY^{i-1}} = \mathsf{p}^{\mathbf{T}}_{Y_i|X^iY^{i-1}}$ for all $i \geq 1$. Moreover, we say that an algorithm A *implements* the system $\mathbf{S}$ if $A \equiv \mathbf{S}$.

A *distinguisher* $\mathbf{D}$ is a special type of system which interacts with another system $\mathbf{S}$ by means of q queries and outputs a decision bit $\mathbf{D}(\mathbf{S})$ depending on their outputs: Its *advantage* in distinguishing systems $\mathbf{S}$ and $\mathbf{T}$ is

$$\Delta^{\mathbf{D}}(\mathbf{S}, \mathbf{T}) := \left|\mathsf{P}\left[\mathbf{D}(\mathbf{S}) = 1\right] - \mathsf{P}\left[\mathbf{D}(\mathbf{T}) = 1\right]\right|.$$

Moreover, $\Delta_q(\mathbf{S}, \mathbf{T})$ is the best distinguishing advantage $\Delta^{\mathbf{D}}(\mathbf{S}, \mathbf{T})$ over *all* q-query $\mathbf{D}$, whereas $\Delta_{t,q}(\mathbf{S}, \mathbf{T})$ is used when the maximization is restricted to distinguishers implemented by an algorithm with time complexity t.

Stateless Systems. A system $\mathbf{S}$ is called *stateless* if the i-th answer Y_i only depends on the i-th query X_i, that is, there exists a conditional distribution $\mathsf{p}^{\mathbf{S}}_{Y|X}$ such that $\mathsf{p}^{\mathbf{S}}_{Y_i|X^iY^{i-1}}(y_i, x^i, y^{i-1}) = \mathsf{p}^{\mathbf{S}}_{Y|X}(y_i, x_i)$ for all $i \geq 1$, $x^i = [x_1, \ldots, x_i]$, and $y^i = [y_1, \ldots, y_i]$. Furthermore, $\mathbf{S}$ is *convex-combination-stateless* (or simply *cc-stateless*) [10] if there exists a system $\mathbf{T}(\cdot)$ accessing a random variable S (called the *initial state*) such that $\mathbf{S} \equiv \mathbf{T}(S)$ and $\mathbf{T}(s)$ is stateless for all values s taken by S. To save on notation, we usually write $\mathbf{S}(\cdot)$ instead of $\mathbf{T}(\cdot)$, but we stress that $\mathbf{S}(\cdot)$ and $\mathbf{S}$ are different objects, despite their notational similarity. We refer to $\mathbf{S}(S)$ as the *cc-stateless representation* of $\mathbf{S}$.

It is crucial to remark that being cc-stateless is a property of the *input-output behavior* of a system: Its (efficient) implementation may well be stateful, and its cc-stateless representation may be completely inefficient (e.g., because the description of the initial state is even too large to be processed by an efficient algorithm).

Random Functions and Permutations. A system $\mathbf{F}$ taking inputs from a set $\mathcal{X}$ and returning outputs in $\mathcal{Y}$ is a *random function* $\mathcal{X} \to \mathcal{Y}$ if for any two equal queries $X_i = X_j$ we have $Y_i = Y_j$ for the respective answers. Furthermore, if $\mathcal{X} = \mathcal{Y}$, it is called a *random permutation* if $X_i \neq X_j$ also implies $Y_i \neq Y_j$. Typical

(cc-stateless) examples are *uniform random function* (URF) $\mathbf{R} : \mathcal{X} \to \mathcal{Y}$, which answers according to a uniformly chosen function $\mathcal{X} \to \mathcal{Y}$, a *uniform random permutation* (URP) $\mathbf{P} : \mathcal{X} \to \mathcal{X}$, implementing a uniformly chosen permutation $\mathcal{X} \to \mathcal{X}$, or E_K for a permutation family $\{E_k\}_{k\in\mathcal{K}}$ and a random $K \in \mathcal{K}$.

The initial state of a cc-stateless random function $\mathbf{F}$ can always be seen without loss of generality as a (randomly chosen) *function table* F according to which $\mathbf{F}$ answers its queries, and usually write $\mathbf{F}(x)$ instead of $F(x)$. In particular, the *inverse* $\mathbf{Q}^{-1}$ of a cc-stateless permutation $\mathbf{Q}$ is well-defined, and $\langle\mathbf{Q}\rangle$ is the *two-sided* random permutation which allows both forward queries $(x, +)$ returning $\mathbf{Q}(x)$ as well as *backward queries* $(y, -)$ returning $\mathbf{Q}^{-1}(y)$. The *cascade* $\mathbf{Q}' \rhd \mathbf{Q}''$ of two random permutations is the system which on input x returns $\mathbf{Q}''(\mathbf{Q}'(x))$, i.e., it implements the composition of the associated permutation tables. (This extends naturally to longer cascades.) Note in particular that for any cascade we have $\mathbf{Q}_1 \rhd \cdots \rhd \mathbf{Q}_m \equiv \mathbf{P}$ whenever there exists i such that $\mathbf{Q}_i \equiv \mathbf{P}$ for a URP $\mathbf{P}$. Moreover, we let $\langle\mathbf{Q}'\rangle \rhd \langle\mathbf{Q}''\rangle := \langle\mathbf{Q}' \rhd \mathbf{Q}''\rangle$.

An efficiently implementable family of permutations $E = \{E_k\}_{k\in\mathcal{K}}$ with domain $\mathcal{X}$ and indexed by keys $k \in \mathcal{K}$ is an *ε-pseudorandom permutation* (ε-PRP) if $\Delta_{t,q}(E_K, \mathbf{P}) \leq \varepsilon$ for all polynomially bounded t and q, a uniform $K \in \mathcal{K}$, and a URP $\mathbf{P}$. It is a *two-sided* ε-PRP if $\langle E_K\rangle$ is efficiently implementable and $\Delta_{t,q}(\langle E_K\rangle, \langle\mathbf{P}\rangle) \leq \varepsilon$ for all polynomially bounded t and q.

3 Hardcore Lemmas for Interactive Systems

3.1 System-Bit Pairs, Measures, and State Samplers

We consider the general setting of *system-bit pairs* [10] $(\mathbf{S}, B)$ consisting of a bit B (with an associated probability distribution P_B), and a system $\mathbf{S} = \mathbf{S}(B)$ whose behavior depends on the outcome of the bit B. A system-bit pair $(\mathbf{S}, B)$ is to be interpreted as a system which parallely composes $\mathbf{S}$ and a correlated bit B (which is *initially* chosen, before any interaction with $\mathbf{S}$ has taken place). The notion of a cc-stateless system-bit pair $(\mathbf{S}(S), B(S))$ is obtained naturally. Also, an implementation $A_{(\mathbf{S},B)}$ of a system-bit pair $(\mathbf{S}, B)$ is without loss of generality an algorithm which outputs the bit B and then simulates the system $\mathbf{S}(B)$.

We associate with every system-bit pair $(\mathbf{S}, B)$ a game where an *adversary* $\mathbf{A}$ interacts with $\mathbf{S}(B)$ and outputs a binary guess $\mathbf{A}(\mathbf{S}(B)) \in \{0, 1\}$ for B: Its *guessing advantage* is defined as the quantity

$$\mathrm{Guess}^{\mathbf{A}}(B \,|\, \mathbf{S}) := 2 \cdot \mathsf{P}[\mathbf{A}(\mathbf{S}(B)) = B] - 1 \in [-1, 1].$$

If $\mathrm{Guess}^{\mathbf{A}}(B \,|\, \mathbf{S}) = 1$, then $\mathbf{A}$ always guesses B correctly, whereas $\mathrm{Guess}^{\mathbf{A}}(B \,|\, \mathbf{S}) = -1$ means that $\mathbf{A}$ is always wrong (though flipping $\mathbf{A}$'s output bit yields an adversary which is always correct.) The shorthand $\mathrm{Guess}_{t,q}(B \,|\, \mathbf{S})$ denotes the best guessing advantage taken over all adversaries with time complexity t and issuing at most q queries to $\mathbf{S}$.

Example 1. An example is the (cc-stateless) system-bit pair $(\mathbf{R}, B)$ for a URF $\mathbf{R} : \mathcal{X} \to \{0,1\}$ and $B := \oplus_{x \in \mathcal{X}} \mathbf{R}(x)$ is the parity of its function table. It is easy to see that $\mathrm{Guess}_q(B \,|\, \mathbf{R}) = 0$ for all $q < |\mathcal{X}|$.

Example 2. If $(\mathbf{F}, B)$ is such that B is uniform, and $\mathbf{F}$ behaves as a system $\mathbf{S}$ if $B = 0$, and as another system $\mathbf{T}$ if $B = 1$, then $\mathrm{Guess}^{\mathbf{D}}(B \,|\, \mathbf{F}) = \Delta^{\mathbf{D}}(\mathbf{S}, \mathbf{T})$ for all $\mathbf{D}$ by a standard argument. Note that if both $\mathbf{S}$ and $\mathbf{T}$ are cc-stateless, then $(\mathbf{F}, B)$ is also cc-stateless.

Measures. A *measure* $\mathcal{M}$ for a cc-stateless system $\mathbf{S} \equiv \mathbf{S}(S)$, where $S \in \mathcal{S}$ is the initial state, is a mapping $\mathcal{M} : \mathcal{S} \to [0,1]$. Its *density* is defined as $\mu(\mathcal{M}) := \mathsf{E}\left[\mathcal{M}(S)\right] = \sum_{s \in \mathcal{S}} \mathsf{P}_S(s) \cdot \mathcal{M}(s)$. The measure $\mathcal{M}$ is naturally associated with a probability distribution $\mathsf{P}_{\mathcal{M}}$ on $\mathcal{S}$ such that $\mathsf{P}_{\mathcal{M}}(s) := \mathsf{P}_S(s) \cdot \mathcal{M}(s) \cdot \mu(\mathcal{M})^{-1}$ for all $s \in \mathcal{S}$. Also, we define the complement of a measure $\mathcal{M}$ as the measure $\overline{\mathcal{M}}$ such that $\overline{\mathcal{M}}(s) := 1 - \mathcal{M}(s)$ for all $s \in \mathcal{S}$. We repeatedly abuse notation writing $S \xleftarrow{\$} \mathcal{M}$ instead of $S \xleftarrow{\$} \mathsf{P}_{\mathcal{M}}$.

Traditionally, measures are seen as "fuzzy" subsets of $\mathcal{S}$. Alternatively, it is convenient to think of $\mathcal{M}$ in terms of a conditional probability distribution $\mathsf{P}_{\mathcal{A}|S}$ with $\mathsf{P}_{\mathcal{A}|S}(s) := \mathcal{M}(s)$ which adjoins the event $\mathcal{A}$ on the choice of S: In particular, $\mu(\mathcal{M}) = \mathsf{P}[\mathcal{A}]$, $\mathsf{P}_{\mathcal{M}} = \mathsf{P}_{S|\mathcal{A}}$, and $\mathsf{P}_{\overline{\mathcal{M}}} = \mathsf{P}_{S|\overline{\mathcal{A}}}$. In the following, we stick to measures for stating and proving our hardcore lemmas, while an event-based view will be useful when exercising these results.

State Samplers. Ideally, the hardcore lemma for a cc-stateless system-bit pair $(\mathbf{S}, B) \equiv (\mathbf{S}(S), B(S))$ (for initial state $S \in \mathcal{S}$) states that if $\mathrm{Guess}_{t,q}(B \,|\, \mathbf{S}) \leq \varepsilon$, then there exists a measure $\mathcal{M}$ on $\mathcal{S}$ such that (i) $\mu(\mathcal{M}) \geq 1 - \varepsilon$ and (ii) $\mathrm{Guess}_{t',q'}(B(S') \,|\, \mathbf{S}(S')) \approx 0$ for $S' \xleftarrow{\$} \mathcal{M}$ and t', q' as close as possible to t, q. Whenever $\mathbf{S}(S)$ is a random variable, this is equivalent to (a tight) version of Impagliazzo's Hardcore Lemma [4]. However, applications of the hardcore lemma (as the one we give later in this paper) require the ability, possibly given some short advice, to efficiently simulate $(\mathbf{S}(S'), B(S'))$ for $S' \xleftarrow{\$} \mathcal{M}$ or $(\mathbf{S}(S''), B(S''))$ for $S'' \xleftarrow{\$} \overline{\mathcal{M}}$.[4] While in the context of random variables the advice is generally a sample of S' itself, this approach fails in the setting of interactive systems: Recall that the representation $(\mathbf{S}(S), B(S))$ is possibly only a thought experiment, and the description of S' may be of exponential size, or no efficient algorithm implementing $(\mathbf{S}, B)$ from S' exists, even if the system-bit pair itself is efficiently implementable.

To formalize the concept of an advice distribution, we introduce the notion of a state sampler for a cc-stateless system (such as e.g. a system-bit pair).

Definition 1 (State Samplers). *Let $\mathbf{S} \equiv \mathbf{S}(S)$ be a cc-stateless system with implementation $A_{\mathbf{S}}$ and $S \in \mathcal{S}$, let $\zeta_1, \zeta_2 \in [0,1]$, and let $\mathcal{M} : \mathcal{S} \to [0,1]$ be a*

[4] Formally, one actually needs to prove that $\mathrm{Guess}_{t',q'}(B(S') \,|\, \mathbf{S}(S')) \approx 0$ holds even given access to the advice: While this is implicit in the non-uniform setting (every adversary with advice can be turned in an equally good one without advice), the proof is more challenging in the uniform setting, cf. the full version.

measure for $\mathbf{S}$. *A* (ζ_1, ζ_2)-(state) sampler O *for* $\mathcal{M}$ *and* $A_{\mathbf{S}}$ *with length* ℓ *is a random process* O *such that:*

(i) O *always returns a pair* (σ, z) *with* σ *being a valid state for* $A_{\mathbf{S}}$ *with* $|\sigma| \leq \ell$ *and* $z \in [0,1]$*;*

(ii) For $(\Sigma, Z) \xleftarrow{\$} \mathsf{O}$*, we have*[5]

$$(A_{\mathbf{S}}[\Sigma], Z) \equiv (\mathbf{S}(S), Z'(S)),$$

where $Z'(S) \in [0,1]$ *is a random variable (which only depends on* S*) that differs from* $\mathcal{M}(S)$ *by at most* ζ_1*, except with probability* ζ_2*, for any value taken by* S*.*

Example 3. For all implementations $A_{\mathbf{S}}$ of $\mathbf{S}$, the all-one measure (i.e., $\mathsf{P}_{\mathcal{M}} = \mathsf{P}_S$) admits an error-less sampler O which returns the initial (void) state for $A_{\mathbf{S}}$ and $z = 1$.

Note that O is *not* required to be efficiently implementable, but black-box access to state samplers allow for efficient simulation of $\mathbf{S}(S')$ for $S' \xleftarrow{\$} \mathcal{M}$ using reject-sampling (provided $\mathbf{S}$ admits an efficient implementation): Given the output (Σ, Z) sampled from a (ζ_1, ζ_2)-sampler O, we flip a coin B with $\mathsf{P}_B(1) = Z$: Consider the distribution $\mathsf{P}_{\Sigma|B=1}$ of Σ conditioned on the outcome $B = 1$. If $\zeta_1 = \zeta_2 = 0$, it is not hard to verify that $A_{\mathbf{S}}[\Sigma'] \equiv \mathbf{S}(S')$ for $\Sigma' \xleftarrow{\$} \mathsf{P}_{\Sigma|B=1}$. This is because, by definition, we have $(A_{\mathbf{S}}[\Sigma], Z, B) \equiv (\mathbf{S}(S), \mathcal{M}(S), B')$, where B' is a bit which is 1 with probability $\mathcal{M}(S)$, and thus in particular $A_{\mathbf{S}}[\Sigma'] \equiv \mathbf{S}(S')$ where $S' \xleftarrow{\$} \mathsf{P}_{S|B'=1}$. In addition, since $\mathsf{P}_{B'|S}(1, s) := \mathcal{M}(s)$ and $\mathsf{P}_{B'}(1) := \sum_{s \in \mathcal{S}} \mathsf{P}_S(s) \cdot \mathcal{M}(s) = \mu(\mathcal{M})$,

$$\mathsf{P}_{S|B'}(s, 1) = \mathcal{M}(s) \cdot \mathsf{P}_S(s) \cdot \mu(\mathcal{M})^{-1} = \mathsf{P}_{\mathcal{M}}(s).$$

Of course, one can similarly simulate $\mathbf{S}(S'')$ for $S'' \xleftarrow{\$} \mathsf{P}_{\overline{\mathcal{M}}}$, as we obtain a corresponding sampler by just replacing z by $1 - z$ in the output (σ, z). This approach can be extended to non-zero errors ζ_1 and ζ_2 with some care.

3.2 The Hardcore Lemma for System-Bit Pairs

In the following, for understood parameters γ, ε, ζ_1, and ζ_2, we define

$$\varphi_{\mathsf{hc}} := \frac{6400}{\gamma^2(1-\varepsilon)^4} \cdot \ln\left(\frac{160}{\gamma(1-\varepsilon)^3}\right) \text{ and } \psi_{\mathsf{hc}} := \frac{200}{\gamma^2(1-\varepsilon)^4\zeta_1^2} \cdot \ln\left(\frac{2}{\zeta_2}\right).$$

We now state the HCL for cc-stateless system-bit pairs. Even though we apply the result only in a more restricted setting, we prove a more general statement for arbitrary cc-stateless system-bit pairs.

[5] That is, we consider the parallel composition of a system (either $A_{\mathbf{S}}[\Sigma]$ or $\mathbf{S}(S)$) and a correlated $[0,1]$-valued random variable.

Theorem 1 (HCL for System-Bit Pairs). *Let* $(\mathbf{S}, B) \equiv (\mathbf{S}(S), B(S))$ *be a cc-stateless system-bit pair admitting an implementation* $A_{(\mathbf{S},B)}$ *with space complexity* $s_{A_{(\mathbf{S},B)}}$*. Furthermore, for some integers* $t, q > 0$ *and some* $\varepsilon \in [0, 1)$*,*

$$\mathrm{Guess}_{t,q}(B \,|\, \mathbf{S}) \leq \varepsilon.$$

Then, for all $0 < \zeta_1, \zeta_2 < 1$ *and all* $0 < \gamma \leq \frac{1}{2}$*, there exists a measure* $\mathcal{M}$ *for* $(\mathbf{S}, B)$ *with* $\mu(\mathcal{M}) \geq 1 - \varepsilon$ *such that the following two properties are satisfied:*

(i) For $S' \xleftarrow{\$} \mathcal{M}$*,* $t' := t/\varphi_{\mathsf{hc}}$*, and* $q' := q/\varphi_{\mathsf{hc}}$*,*

$$\mathrm{Guess}_{t',q'}(B(S') \,|\, \mathbf{S}(S')) \leq \gamma.$$

(ii) There exists a (ζ_1, ζ_2)*-sampler for* $\mathcal{M}$ *and* $A_{(\mathbf{S},B)}$ *with length* $s_{A_{(\mathbf{S},B)}}(\psi_{\mathsf{hc}} \cdot q')$*. Moreover, if* $(\mathbf{S}(s), B(s))$ *is deterministic for all* s*, then there also exists a* $(0, 0)$*-sampler for* $\mathcal{M}$ *and* $A_{(\mathbf{S},B)}$ *with length* $s_{A_{(\mathbf{S},B)}}((7 \cdot \gamma^{-2} \cdot (1-\varepsilon)^{-3} + 1) \cdot q')$*.*

In the remainder of this section, we outline the main ideas behind the proof. The complete proof is found in the full version.

Proof Outline. The proof is by contradiction: We assume that for all measures $\mathcal{M}$ with $\mu(\mathcal{M}) \geq 1 - \varepsilon$ admitting a (ζ_1, ζ_2)-sampler as in (ii), there exists an adversary $\mathbf{A}$ with time complexity t' and query complexity q' such that $\mathrm{Guess}^{\mathbf{A}}(B(S') \,|\, \mathbf{S}(S')) > \gamma$ for $S' \xleftarrow{\$} \mathcal{M}$. The core of the proof consists of proving that, under this assumption, there exists a sufficiently small family of adversaries $\mathcal{A}$ (more specifically, $|\mathcal{A}| = 7 \cdot \gamma^{-2} \cdot (1-\varepsilon)^{-3} + 1$) such that either (A) $\alpha(S) > \gamma$ holds with probability higher than $1 - \frac{1-\varepsilon}{4}$ over the choice of S, where $\alpha(s) := \mathsf{E}\left[\mathrm{Guess}^{\mathbf{A}'}(B(s) \,|\, \mathbf{S}(s))\right]$ for all s, where $\mathbf{A}' \xleftarrow{\$} \mathcal{A}$, or (B) $\mathsf{E}\left[\alpha(S')\right] > \Theta\left((1-\varepsilon)^2\gamma\right)$ *for all* measures $\mathcal{M}$ with density $1 - \varepsilon$ and $S' \xleftarrow{\$} \mathcal{M}$.

In Case (A), a simple majority-voting based strategy yields a good adversary breaking the assumed hardness of $(\mathbf{S}, B)$, whereas in Case (B) such an adversary can be built from $\mathcal{A}$ using techniques similar to the case of random variables [5,3]. Both adversaries heavily rely on the cc-stateless property of $(\mathbf{S}, B)$.

To show the existence of an appropriate family, we associate with each family $\mathcal{A}$ and $\tau \in \mathbb{N}$ a measure $\mathcal{M}_{\mathcal{A},\tau}$ such that elements for which $\mathcal{A}$ is worse, i.e., $|\mathcal{A}| \cdot \alpha(s) \leq \tau$, are given high weight (i.e. $\mathcal{M}_{\mathcal{A},\tau}(s) = 1$), whereas elements for which $\mathcal{A}$ performs well, i.e., $|\mathcal{A}| \cdot \alpha(s) \geq \tau + \frac{1}{\gamma(1-\varepsilon)}$, are not chosen ($\mathcal{M}_{\mathcal{A},\tau}(s) = 0$). An intermediate measure value is assigned to states not falling into one of these two categories. In particular, $\mathcal{M}_{\emptyset,0}$ is the all-one measure (i.e., $\mathsf{P}_{\mathcal{M}}$ equals the state distribution P_S), which has density $1 \geq 1 - \varepsilon$. A crucial property is that $\mathcal{M}_{\mathcal{A},\tau}$ admits an (ζ_1, ζ_2)-state sampler for all $\mathcal{A}$ and τ, which is shown by using the observation that $\mathcal{M}_{\mathcal{A},\tau}(S)$ can always be estimated given *black-box* access to $(\mathbf{S}(S), B(S))$. We then consider the following iterative process: It starts with $\mathcal{A} := \emptyset$ and then, at each round, it possibly increases τ to ensure that $\mu(\mathcal{M}_{\mathcal{A},\tau}) \geq 1 - \varepsilon$ and then uses the assumption of the HCL being wrong to find an adversary achieving advantage larger than γ for $\mathcal{M}_{\mathcal{A},\tau}$, and adds it to $\mathcal{A}$. We prove that within $7 \cdot \gamma^{-2} \cdot (1-\varepsilon)^{-3} + 1$ iterations, $\mathcal{A}$ satisfies (A) or (B).

Remark 1. A natural question is whether the HCL can be extended to arbitrary system-bit pairs, where the measure is defined on the randomness of the implementation of the system-bit pair, regardless of the system having a cc-stateless representation. Yet, techniques similar to the ones used in counter-examples to soundness amplification for interactive arguments via parallel repetition [1,14] yield (non cc-stateless) efficiently implementable system-bit pairs for which, given multiple independent instances of the system-bit pair, the probability of guessing all of the bits given access to all of the associated systems in parallel does not decrease with the number of instances. If such a a general HCL were true, then it is not hard to prove that the guessing probability *would* decrease exponentially in the number of instances.

3.3 The Hardcore Lemma for Computational Indistinguishability

This section presents the hardcore lemma for computational indistinguishability of *interactive* systems, which generalizes the statement for random variables previously shown in [11].

Theorem 2 (HCL for Computational Indistinguishability). *Let* $\mathbf{S} \equiv \mathbf{S}(S)$ *and* $\mathbf{T} \equiv \mathbf{T}(T)$ *be cc-stateless systems, with respective implementations* $A_{\mathbf{S}}$ *(with space complexity* $s_{A_{\mathbf{S}}}$*) and* $A_{\mathbf{T}}$ *(with space complexity* $s_{A_{\mathbf{T}}}$*). Furthermore, for some integers* $t, q > 0$ *and some* $\varepsilon \in [0, 1)$*,*

$$\Delta_{t,q}(\mathbf{S}, \mathbf{T}) \leq \varepsilon.$$

Then, for all $0 < \zeta_1, \zeta_2 < 1$ *and all* $0 < \gamma \leq \frac{1}{2}$*, there exist measures* $\mathcal{M}_{\mathbf{S}}$ *and* $\mathcal{M}_{\mathbf{T}}$ *such that* $\mu(\mathcal{M}_{\mathbf{S}}) \geq 1 - \varepsilon$ *and* $\mu(\mathcal{M}_{\mathbf{T}}) \geq 1 - \varepsilon$ *and the following properties hold:*

(i) For $S' \xleftarrow{\$} \mathcal{M}_{\mathbf{S}}$*,* $T' \xleftarrow{\$} \mathcal{M}_{\mathbf{T}}$*,* $t' := t/\varphi_{\mathsf{hc}}$*, and* $q' := q/\varphi_{\mathsf{hc}}$*, we have*

$$\Delta_{t',q'}(\mathbf{S}(S'), \mathbf{T}(T')) \leq 2\gamma;$$

(ii) There exist a (ζ_1, ζ_2)*-sampler* $\mathsf{O}_{\mathbf{S}}$ *for* $\mathcal{M}_{\mathbf{S}}$ *and* $A_{\mathbf{S}}$ *with length* $s_{A_{\mathbf{S}}}(\psi_{\mathsf{hc}} \cdot q')$ *and a* (ζ_1, ζ_2)*-sampler* $\mathsf{O}_{\mathbf{T}}$ *for* $\mathcal{M}_{\mathbf{T}}$ *and* $A_{\mathbf{T}}$ *with length* $s_{A_{\mathbf{T}}}(\psi_{\mathsf{hc}} \cdot q')$*. Furthermore, if both* $\mathbf{S}$ *and* $\mathbf{T}$ *are random functions, then both samplers can be made error-less with lengths* $s_{A_{\mathbf{S}}}(\psi \cdot q')$ *and* $s_{A_{\mathbf{T}}}(\psi \cdot q')$*, where* $\psi := 7 \cdot \gamma^{-2} \cdot (1 - \varepsilon)^{-3} + 1$*.*

We postpone the proof to the full version, which relies on Theorem 1, and only present the main ideas in the following.

Proof Sketch. We define $(\mathbf{F}, B) \equiv (\mathbf{F}(X, B), B)$ to be the cc-stateless system-bit pair with a uniform random bit B and where $\mathbf{F}$ behaves as $\mathbf{S}$ if $B = 0$ and as $\mathbf{T}$ if $B = 1$. In particular, the initial state (X, B) of $(\mathbf{F}, B)$ is sampled by first letting $B \xleftarrow{\$} \{0, 1\}$, and then choosing $X \xleftarrow{\$} \mathsf{P}_S$ if $B = 0$ and $X \xleftarrow{\$} \mathsf{P}_T$ otherwise, and

$$(\mathbf{F}(x, b), B(x, b)) = \begin{cases} (\mathbf{S}(x), 0) & \text{if } b = 0, \\ (\mathbf{T}(x), 1) & \text{if } b = 1. \end{cases}$$

By a standard argument $\Delta_{t,q}(\mathbf{S},\mathbf{T}) = \mathrm{Guess}_{t,q}(B \,|\, \mathbf{F}) \leq \varepsilon$ holds (also cf. Example 2), and Theorem 1 thus implies that there exists a measure $\mathcal{M}$ for $(\mathbf{F}, B)$ such that $\mu(\mathcal{M}) \geq 1-\varepsilon$, and $\mathrm{Guess}_{t',q'}(B' \,|\, \mathbf{F}(X')) \leq \gamma$, where $(X', B') \xleftarrow{\$} \mathcal{M}$, $t' = t/\varphi_{\mathsf{hc}}$, and $q' = q/\varphi_{\mathsf{hc}}$. Define $\mathcal{M}_{\mathbf{S}}(s) := \mathcal{M}(s,0)$ and $\mathcal{M}_{\mathbf{T}}(t) := \mathcal{M}(t,1)$, and note that

$$\begin{aligned} \mathsf{P}_{X'B'}(s,0) &= \frac{1}{2\mu(\mathcal{M})} \cdot \mathsf{P}_S(s) \cdot \mathcal{M}_{\mathbf{S}}(s), \\ \mathsf{P}_{X'B'}(t,1) &= \frac{1}{2\mu(\mathcal{M})} \cdot \mathsf{P}_T(t) \cdot \mathcal{M}_{\mathbf{T}}(t). \end{aligned} \tag{1}$$

If B' were uniformly distributed (i.e., $\sum_s \mathsf{P}_{X'B'}(s,0) = \sum_t \mathsf{P}_{X'B'}(t,1) = \frac{1}{2}$), we then would have $\mu(\mathcal{M}_{\mathbf{S}}) = \mu(\mathcal{M}_{\mathbf{T}}) = \mu(\mathcal{M}) \geq 1-\varepsilon$ by (1), and (X', B') could be sampled by choosing B' uniformly, and letting $X' = S' \xleftarrow{\$} \mathcal{M}_{\mathbf{S}}$ if $B' = 0$, and $X' = T' \xleftarrow{\$} \mathcal{M}_{\mathbf{T}}$ if $B' = 1$. This would also yield

$$\Delta_{t',q'}(\mathbf{S}(S'), \mathbf{T}(T')) = \mathrm{Guess}_{t',q'}(B' \,|\, \mathbf{F}(X')) \leq \gamma,$$

concluding the proof. The main challenge in the full proof is dealing with the fact that B' is generally only $\Theta(\gamma)$-close to uniform.

Remark 2. Theorem 2 can be seen as a computational analogue of Lemma 5 in [9], which shows a similar property for information-theoretic indistinguishability (i.e., with respect to computationally unbounded distinguishers). Theorem 2 can of course also be used in the IT setting, and it is somewhat stronger in that it yields events defined on the initial state of the system, instead of interaction-dependent sequences of events as in [9]. However, Lemma 5 in [9] holds for *arbitrary* systems and presents a tight reduction with $q' = q$ and no additive term γ, which we do not know how to achieve in the computational setting.

Connection to Computational Entropy. Let $\mathbf{Q}$ be a cc-stateless random permutation on $\mathcal{X}$ (with $N := |\mathcal{X}|$) with function table Q and such that $\Delta_{t,q}(\mathbf{Q},\mathbf{P}) \leq \varepsilon$ for a URP $\mathbf{P}$. Theorem 2 yields events $\mathcal{A}$ on Q and $\mathcal{B}$ on a uniform permutation table P such that $\mathsf{P}[\mathcal{A}] \geq 1-\varepsilon$, $\mathsf{P}[\mathcal{B}] \geq 1-\varepsilon$, and $\Delta_{t',q'}(\mathbf{Q}',\mathbf{P}') \leq \gamma$, where $\mathbf{Q}'$ and $\mathbf{P}'$ are cc-stateless random functions with function tables $Q' \xleftarrow{\$} \mathsf{P}_{Q|\mathcal{A}}$ and $P' \xleftarrow{\$} \mathsf{P}_{P|\mathcal{B}}$, respectively. In particular, $\mathsf{P}_{P'}(\pi) = \frac{\mathsf{P}_P(\pi)\cdot\mathsf{P}_{\mathcal{B}|P}(\pi)}{\mathsf{P}[\mathcal{B}]} \leq \frac{1}{(1-\varepsilon)\cdot(N!)}$ for all permutations π, and the *min-entropy* $\mathsf{H}_\infty(P') := -\log\max_\pi \mathsf{P}_{P'}(\pi)$ is at least $\log(N!) - \log\left((1-\varepsilon)^{-1}\right)$. Informally, this can be interpreted as Q having *"computational"* min-entropy at most $\log\left((1-\varepsilon)^{-1}\right)$ away from the maximum achievable entropy $\log(N!)$ with probability $1-\varepsilon$.[6] Clearly, the statement also extends to the two-sided case as well as to other types of systems.

[6] We stress, however, that the distribution P' depends on t, q, as well as on γ.

Remark 3. Another useful fact is that P' has statistical distance ε from P. This follows from the observation that the distribution of P' is a convex combination of flat distributions over subsets of size at least $(1-\varepsilon)\cdot(N!)$: As each such distribution is ε-away from uniform, the bound follows from the convexity of the statistical distance. Therefore, $\Delta_{t,q}(\mathbf{P}', \mathbf{P}) \leq \Delta_{t,q}(\langle \mathbf{P}' \rangle, \langle \mathbf{P} \rangle) \leq d(P', P) \leq \varepsilon$ for all t, q.

4 Cascade of Weak Permutations

4.1 Cascade of Permutations with Large Entropy

Let $\mathbf{Q}_1$ and $\mathbf{Q}_2$ be two independent cc-stateless random permutations on the set $\mathcal{X}$ (with $N := |\mathcal{X}|$) with the property that the min-entropies of their respective function tables Q_1 and Q_2 satisfy $\mathsf{H}_\infty(Q_1) \geq \log(N!) - \log\left((1-\varepsilon)^{-1}\right)$ and $\mathsf{H}_\infty(Q_2) \geq \log(N!) - \log\left((1-\varepsilon)^{-1}\right)$ for some $\varepsilon \in \left[0, 1 - \frac{1}{N}\right)$. We prove that the cascade $\mathbf{Q}_1 \rhd \mathbf{Q}_2$ is indistinguishable from a URP $\mathbf{P}$ for computationally *unbounded* distinguishers, both in the one- and in the two-sided cases.

Theorem 3 (Cascade of Large-Entropy Permutations). *For all $q, \Lambda \geq 1$,*

$$\Delta_q(\langle \mathbf{Q}_1 \rhd \mathbf{Q}_2 \rangle, \langle \mathbf{P} \rangle) \leq \tfrac{4q\Lambda}{N} + \tfrac{2\Lambda(q+\Lambda)}{(1-\varepsilon)N} + 2\left(\frac{q\log\left((1-\varepsilon)^{-1}\right)}{\Lambda}\right)^{\frac{1}{2}}.$$

The same bound applies to any cascade $\mathbf{Q}'_1 \rhd \cdots \rhd \mathbf{Q}'_m$ of m independent cc-stateless random permutations such that $\mathbf{Q}'_i \equiv \mathbf{Q}_1$ and $\mathbf{Q}'_j \equiv \mathbf{Q}_2$ for some $i < j$, as such a cascade can be seen as the cascade of two permutations $\overline{\mathbf{Q}}_1 := \mathbf{Q}'_1 \rhd \cdots \rhd \mathbf{Q}'_i$ and $\overline{\mathbf{Q}}_2 := \mathbf{Q}'_{i+1} \rhd \cdots \rhd \mathbf{Q}'_m$ with the same min-entropy guarantees on their function tables. The theorem allows free choice of Λ: For our purposes, it suffices to set $\Lambda := (\log N)^\zeta$ (for a slowly growing $\zeta = \omega(1)$ in the security parameter $\log N$) to achieve indistinguishability for $q = \mathsf{poly}(\log N)$ queries and any $\varepsilon \leq 1 - \frac{(\log N)^{3\zeta}}{N}$.

The core of the proof, which is omitted for lack of space, is a lemma stating that $\langle \mathbf{Q}_i \rangle$ (for $i = 1, 2$) is indistinguishable from a random permutation $\langle \mathbf{Q}_i \rangle_{\overline{\mathbf{D}}_i}$ which is initialized by letting a carefully chosen distinguisher $\overline{\mathbf{D}}_i$ (making Λ queries) interact with $\langle \mathbf{Q}_i \rangle$, and then answering queries according to a randomly chosen permutation consistent with $\overline{\mathbf{D}}_i$'s interaction. (This extends a previous result by Unruh [15] to random permutations.) We employ tools from the random systems framework [8] (including a new lemma) to prove that the cascade of two independent such permutations is indistinguishable from a URP.

4.2 Security Amplification of Weak PRPs

Let $\mathbf{Q}$ be a cc-stateless random permutation with domain $\mathcal{X}$ (for $N := |\mathcal{X}| = 2^n$, where n is the security parameter) such that $\langle \mathbf{Q} \rangle$ is implemented by the algorithm $A_{\langle \mathbf{Q} \rangle}$ with time complexity $t_{A_{\langle \mathbf{Q} \rangle}}$ and space complexity $s_{A_{\langle \mathbf{Q} \rangle}}$. We also consider the canonical (efficient) implementation of a two-sided URP $\langle \mathbf{P} \rangle$ that maintains

a table consisting of all input-output pairs (x_i, y_i) of previous queries as its state, and, upon a new query $(x, +)$ or $(y, -)$, it chooses uniformly at random a y' (or x') not appearing as the second (first) element in a previous input-output pair, and adds (x, y') (or (x', y)) to the table. (If a corresponding pair is in the table, it answers accordingly.) Thus each query is answered in time $\mathcal{O}(\log(s))$, where s is the size of the table, and $s = \mathcal{O}(q \cdot n)$ after q queries.

The following is the main security amplification result of this paper.

Theorem 4. *Let $\mathbf{Q}_1, \ldots, \mathbf{Q}_m$ be independent instances of $\mathbf{Q}$ and let $\mathbf{P}$ be a URP, and assume that for some t, q we have $\Delta_{t,q}(\langle \mathbf{Q} \rangle, \langle \mathbf{P} \rangle) \leq \varepsilon$. For all $\gamma > 1$ and $\Lambda > 0$,*

$$\Delta_{t'',q''}(\langle \mathbf{Q}_1 \rhd \cdots \rhd \mathbf{Q}_m \rangle, \langle \mathbf{P} \rangle) \leq (m - (m-1)\varepsilon) \cdot \varepsilon^m + \tfrac{4q''\Lambda}{N} + \tfrac{2\Lambda(q''+\Lambda)}{(1-\varepsilon)N} + 2\left(\tfrac{q'' \log\left((1-\varepsilon)^{-1}\right)}{\Lambda}\right)^{\frac{1}{2}} + (2m+2)\gamma,$$

where $t'' := t/\varphi_{\mathsf{hc}} - (m-1)\max\left\{t_{A_{\langle \mathbf{Q} \rangle}}(q'', s_{A_{\langle \mathbf{Q} \rangle}}(q'' \cdot \psi)), \mathcal{O}\left(q'' \log(q'' \cdot (\psi+1)n)\right)\right\}$ and $q'' := q/\varphi_{\mathsf{hc}}$, for $\psi := 7 \cdot \gamma^{-2} \cdot (1-\varepsilon)^{-3} + 1$ and φ_{hc} as in Theorem 2.

Essentially the same result can be proven for the single-sided case. The proof of Theorem 4 follows from the observation that, with very high probability, at least two permutations in the cascade are computational indistinguishable from random permutations with large entropy, allowing application of Theorem 3. Extra work is required to prove a non-trivial bound for the case where *at most* one permutation is guaranteed to have high-entropy. The tightness of these bounds is discussed in Section 4.3.

Proof. Theorem 2 implies that we can define (two-sided) random permutations $\langle \mathbf{Q}' \rangle, \langle \mathbf{Q}'' \rangle$, and $\langle \mathbf{P}' \rangle$ such that the following three properties hold for some $p \leq \varepsilon$: (i) The function table of $\langle \mathbf{P}' \rangle$ has min-entropy at least $\log(N!) - \log\left((1-\varepsilon)^{-1}\right)$, (ii) $\langle \mathbf{Q} \rangle$ behaves as $\langle \mathbf{Q}' \rangle$ with probability $1-p$ and as $\langle \mathbf{Q}'' \rangle$ with probability p, and (iii) $\Delta_{t',q''}(\langle \mathbf{Q}' \rangle, \langle \mathbf{P}' \rangle) \leq 2\gamma$ for $t' := t/\varphi_{\mathsf{hc}}$. Furthermore, $\langle \mathbf{Q}' \rangle$ and $\langle \mathbf{Q}'' \rangle$ can both be perfectly implemented using $A_{\langle \mathbf{Q} \rangle}$ initialized with some appropriately distributed state of length at most $s_{A_{\langle \mathbf{Q} \rangle}}(q'' \cdot \psi)$ given as advice. Similarly, $\langle \mathbf{P}' \rangle$ can be simulated by running the above canonical algorithm initialized with an appropriate state of length $\mathcal{O}(q'' \cdot \psi \cdot n)$. (See the discussion in Section 3.1.)

Additionally, for $\mathcal{I} \subseteq \{1, \ldots, m\}$, let $\mathcal{A}_I$ be the event that $\langle \mathbf{Q}_i \rangle$ behaves as $\langle \mathbf{Q}' \rangle$ for all $i \in \mathcal{I}$ whereas $\langle \mathbf{Q}_i \rangle$ behaves as $\langle \mathbf{Q}'' \rangle$ for all $i \notin \mathcal{I}$. Likewise, for independent instances $\langle \mathbf{Q}'_i \rangle$ and $\langle \mathbf{Q}''_i \rangle$ (for $i = 1, \ldots, m$) of $\langle \mathbf{Q}' \rangle$ and $\langle \mathbf{Q}'' \rangle$, respectively, let $\mathbf{Q}_{\mathcal{I}} := \mathbf{S}_1 \rhd \cdots \rhd \mathbf{S}_m$, where $\mathbf{S}_i := \mathbf{Q}'_i$ for all $i \in \mathcal{I}$ and $\mathbf{S}_i := \mathbf{Q}''_i$ for all $i \notin \mathcal{I}$.

We now fix some distinguisher $\mathbf{D}$ with time complexity t'' and making q'' queries, and we first observe that

$$\delta^{\mathbf{D}}(\langle \mathbf{Q}_1 \rhd \cdots \rhd \mathbf{Q}_m \rangle, \langle \mathbf{P} \rangle) = \sum_{\mathcal{I} \subseteq \{1,\ldots,m\}} q_{\mathcal{I}} \cdot \delta^{\mathbf{D}}(\langle \mathbf{Q}_{\mathcal{I}} \rangle, \langle \mathbf{P} \rangle), \tag{2}$$

where $\delta^{\mathbf{D}}(\mathbf{F},\mathbf{G}) := \mathsf{P}\left[\mathbf{D}(\mathbf{F})=1\right] - \mathsf{P}\left[\mathbf{D}(\mathbf{G})=1\right]$ and $q_{\mathcal{I}} := \mathsf{P}\left[\mathcal{A}_{\mathcal{I}}\right] = (1-p)^{|\mathcal{I}|} \cdot p^{m-|\mathcal{I}|}$. Note that the maximum of $\delta^{\mathbf{D}}(\mathbf{F},\mathbf{G})$ over all distinguishers $\mathbf{D}$ with time complexity t'' and space complexity q'' is $\Delta_{t'',q''}(\mathbf{F},\mathbf{G})$.

We first upper bound the summands corresponding to sets $\mathcal{I}$ with at most one element. To this end, for all $i=1,\ldots,m$, we define the distinguisher $\mathbf{D}_i$ which, given access to a two-sided random permutation $\langle\mathbf{S}\rangle$, outputs

$$\mathbf{D}(\langle\mathbf{Q}''_1 \rhd \cdots \rhd \mathbf{Q}''_{i-1} \rhd \mathbf{S} \rhd \mathbf{Q}''_{i+1} \rhd \cdots \rhd \mathbf{Q}''_m\rangle),$$

and is implemented with time complexity $t''+(m-1)t_{A_{\langle\mathbf{Q}\rangle}}(q', s_{A_{\langle\mathbf{Q}\rangle}}(\psi\cdot q')) \leq t'$ given the appropriate advice.

We have $\delta'_i := \delta^{\mathbf{D}_i}(\langle\mathbf{Q}'\rangle, \langle\mathbf{P}\rangle) = \delta^{\mathbf{D}_i}(\langle\mathbf{Q}'\rangle, \langle\mathbf{P}'\rangle) + \delta^{\mathbf{D}_i}(\langle\mathbf{P}'\rangle, \langle\mathbf{P}\rangle) \leq 2\gamma+\varepsilon$, where the bound on the first term follows from the hardcore lemma (for every fixed value of the advice), whereas the bound on the second term follows from Remark 3. Additionally, $\delta^{\mathbf{D}_i}(\langle\mathbf{Q}\rangle, \langle\mathbf{P}\rangle) = (1-p)\cdot\delta'_i + p\cdot\delta''_i \leq \varepsilon$ with $\delta''_i := \delta^{\mathbf{D}_i}(\langle\mathbf{Q}''\rangle, \langle\mathbf{P}\rangle)$ by the indistinguishability assumption on $\langle\mathbf{Q}\rangle$ and the fact that $t' < t$. Since

$$\left\langle\mathbf{Q}''_1 \rhd \cdots \rhd \mathbf{Q}''_{i-1} \rhd \mathbf{P} \rhd \mathbf{Q}''_{i+1} \rhd \cdots \rhd \mathbf{Q}''_m\right\rangle \equiv \langle\mathbf{P}\rangle,$$

we obtain $\delta^{\mathbf{D}}(\langle\mathbf{Q}_\emptyset\rangle, \langle\mathbf{P}\rangle) = \delta''_i$ and $\delta^{\mathbf{D}}(\langle\mathbf{Q}_{\{i\}}\rangle, \langle\mathbf{P}\rangle) = \delta'_i$ for all $i \in \{1,\ldots,m\}$, and thus

$$\begin{aligned}\sum_{|\mathcal{I}|\leq 1} q_{\mathcal{I}}\cdot\delta^{\mathbf{D}}(\langle\mathbf{Q}_{\mathcal{I}}\rangle, \langle\mathbf{P}\rangle) &= \sum_{i=1}^{m}\frac{1}{m}\cdot p^m\cdot\delta''_i + p^{m-1}(1-p)\cdot\delta'_i \\ &\leq \max_{i\in\{1,\ldots,m\}}\left\{p^m\cdot\delta''_i + m\cdot p^{m-1}\cdot(1-p)\cdot\delta'_i\right\}.\end{aligned}$$

However, for all $i \in \{1,\ldots,m\}$, we combine all of the above observations to obtain

$$\begin{aligned}p^m\delta''_i + mp^{m-1}(1-p)\delta'_i &= p^{m-1}(p\delta''_i + (1-p)\delta'_i) + (m-1)p^{m-1}(1-p)\delta'_i \\ &\leq p^{m-1}\varepsilon + (m-1)p^{m-1}(1-p)\varepsilon + 2\gamma \\ &\leq \varepsilon^m + (m-1)\varepsilon^m(1-\varepsilon) + 2\gamma \\ &= \varepsilon^m(m-(m-1)\varepsilon) + 2\gamma,\end{aligned}$$

where we also have used $p \leq \varepsilon$ and the fact that $p^m + (m-1)p^{m-1}(1-p)$ grows monotonically for $p \in [0,1]$.

To bound the remaining summands of Equation (2) with $|\mathcal{I}| \geq 2$, we use a standard hybrid argument and Theorem 3 to obtain

$$\delta^{\mathbf{D}}(\langle\mathbf{Q}_{\mathcal{I}}\rangle, \langle\mathbf{P}\rangle) \leq m\cdot\gamma + \frac{4q''\Lambda}{N} + \frac{2\Lambda(q''+\Lambda)}{(1-\varepsilon)N}.$$

This concludes the proof. □

The following corollary follows by applying the theorem to all $\gamma = 1/p$ (for some polynomial p in n) and to all polynomially bounded t, q, and by choosing an appropriate $\Lambda := n^{\omega(1)}$:

Corollary 1. *Let $E = \{E_k\}_{k\in\mathcal{K}}$ be a (two-sided) ε-PRP for $\varepsilon \leq 1 - \frac{1}{\mathsf{poly}(n)}$, where n is the security parameter. Then, for any $m = \mathsf{poly}(n)$, the cascade $\{E_{k_1} \circ \cdots \circ E_{k_m}\}_{k_1,\ldots,k_m\in\mathcal{K}}$ is a (two-sided) $(\varepsilon^m(m-(m-1)\varepsilon)+\nu)$-PRP for some negligible function ν, where $\circ$ denotes permutation composition.*

4.3 Tightness

Let $\varepsilon < 1 - 2^{-n}$ be such that $\log\left((1-\varepsilon)^{-1}\right) \in \{1,\ldots,n\}$. Let $\mathbf{Q} : \{0,1\}^n \to \{0,1\}^n$ be the cc-stateless random permutation which initially chooses $B \in \{0,1\}$ with $\mathsf{P}_B(0) = \varepsilon$. If $B = 0$, then $\mathbf{Q}$ behaves as the identity permutation id, whereas if $B = 1$ it behaves as a uniformly chosen permutation Q' with the constraint that the first $\log\left((1-\varepsilon)^{-1}\right)$ bits of $Q'(0^n)$ are all equal to 0. Clearly, it is possible to give an efficient *stateful* algorithm implementing $\mathbf{Q}$ (or $\langle\mathbf{Q}\rangle$) by using lazy sampling.[7] Also, let $\mathbf{Q}_1,\ldots,\mathbf{Q}_m$ be independent instances of $\mathbf{Q}$. We prove the following two statements:

(i) For all distinguishers $\mathbf{D}$ and an n-bit URP $\mathbf{P}$, we have $\Delta^{\mathbf{D}}(\langle\mathbf{Q}\rangle, \langle\mathbf{P}\rangle) \leq \varepsilon$, regardless of their computing power.
(ii) There exists a constant-time distinguisher $\mathbf{D}^*$ making *one* single (forward) query such that

$$\Delta^{\mathbf{D}^*}(\mathbf{Q}_1 \rhd \cdots \rhd \mathbf{Q}_m, \mathbf{P}) \geq (m-(m-1)\varepsilon)\varepsilon^m - \frac{1}{2^n}.$$

Hence, the bound of Theorem 4 cannot be substantially improved, even if allowing a huge security loss (i.e., $t'' << t$ and $q'' << q$). This extends to arbitrary m a previous tightness result given by Myers [12] for the special case $m = 2$.

$\mathbf{Q}$ is a two-sided ε-PRP. In the following, let Q and P be random variables representing the distributions of the permutation tables of $\mathbf{Q}$ and $\mathbf{P}$, respectively. There are $(1-\varepsilon)(2^n!)$ permutations π for which the last $\log\left((1-\varepsilon)^{-1}\right)$ bits of $\pi(0^n)$ all equal to 0, and the identity id is one such permutation. Hence,

$$\mathsf{P}_Q(\mathrm{id}) = \varepsilon + (1-\varepsilon)\cdot\frac{1}{(1-\varepsilon)(2^n!)} = \varepsilon + \frac{1}{2^n!} \geq \frac{1}{2^n!} = \mathsf{P}_P(\mathrm{id}).$$

For all $\pi \neq \mathrm{id}$, we have $\mathsf{P}_Q(\pi) \leq (1-\varepsilon)\cdot\frac{1}{(1-\varepsilon)(2^n!)} = \frac{1}{2^n!} = \mathsf{P}_P(\pi)$. This yields $\Delta^{\mathbf{D}}(\langle\mathbf{P}\rangle, \langle\mathbf{Q}\rangle) \leq d(P,Q) = \mathsf{P}_Q(\mathrm{id}) - \mathsf{P}_P(\mathrm{id}) = \varepsilon$ for *all* distinguishers $\mathbf{D}$.

[7] Also, from any PRP $E = \{E_k\}_{k\in\{0,1\}^n}$ with n-bit string domain, we can define a permutation family $E' = \{E'_{k'}\}_{k'\in\{0,1\}^{\log(1/\varepsilon)+n}}$ which is computationally indistinguishable from $\mathbf{Q}$ under a uniform $(\log(1/\varepsilon)+n)$-bit random key: For all $k' \in \{0,1\}^{\log(1/\varepsilon)}$ and $k \in \{0,1\}^n$, let $E'_{k'\|k}(x) := x$ if $k' = 0^{\log(1/\varepsilon)}$, and $E'_{k'\|k}(x) := E_k(x) \oplus E_k(0^n)|_{\log\left((1-\varepsilon)^{-1}\right)}$ otherwise, where $z|_r$ sets the last $n-r$ bits of $z \in \{0,1\}^n$ to be 0 (and leaves the first r unchanged) and $\|$ denotes string concatenation.

Lower Bound for Distinguishing the Cascade. We define $\mathbf{D}^*$ as the distinguisher querying 0^n and outputting 1 if and only if the first $\log\left((1-\varepsilon)^{-1}\right)$ bits of the resulting output are all 0, and outputting 0 otherwise. In particular, $\mathsf{P}[\mathbf{D}^*(\mathbf{P})=1] = 2^{-\log\left((1-\varepsilon)^{-1}\right)} = 1-\varepsilon$, as the output of $\mathbf{P}$ on input 0^n is uniform.

Denote as B_i the bit B associated with the i-th instance $\mathbf{Q}_i$, and let $\mathcal{A}_{\mathcal{I}}$ for $\mathcal{I} \subseteq \{1,\ldots,m\}$ be the event that $B_i = 1$ for all $i \in \mathcal{I}$ and $B_i = 0$ for all $i \notin \mathcal{I}$. Furthermore, let $\mathcal{E}$ be the event that $\mathcal{A}_{\mathcal{I}}$ occurs for some $\mathcal{I}$ with $|\mathcal{I}| \leq 1$. Clearly, $\mathsf{P}[\mathcal{E}] = \varepsilon^m + m(1-\varepsilon)\varepsilon^{m-1}$ and $\mathsf{P}\left[\mathbf{D}^*(\mathbf{Q}_1 \rhd \cdots \rhd \mathbf{Q}_m) = 1 \,\middle|\, \mathcal{E}\right] = 1$, since $\mathbf{Q}_1 \rhd \cdots \rhd \mathbf{Q}_m$ under $\mathcal{E}$ behaves either as the identity or as Q', and in both cases the first $\log\left((1-\varepsilon)^{-1}\right)$ output bits are all 0.

Let us fix $\mathcal{I}$ with $k := |\mathcal{I}| \geq 2$, and let $\mathbf{Q}'_1,\ldots,\mathbf{Q}'_k$ be independent random permutations answering according to Q'. Then,

$$\mathsf{P}\left[\mathbf{D}^*(\mathbf{Q}_1 \rhd \cdots \rhd \mathbf{Q}_m) = 1 \,\middle|\, \mathcal{A}_{\mathcal{I}}\right] = \mathsf{P}\left[\mathbf{D}^*(\mathbf{Q}'_1 \rhd \cdots \rhd \mathbf{Q}'_k) = 1\right].$$

For any input $x \neq 0^n$ the probability that the first $\log\left((1-\varepsilon)^{-1}\right)$ output bits of $\mathbf{Q}'_k(x)$ are all 0 is exactly $1-\varepsilon$, whereas the probability that $\mathbf{Q}'_k$ is invoked on 0^n is at most $\frac{1}{(1-\varepsilon)2^n}$ (as regardless of the input, the output $\mathbf{Q}'_{k-1}$ is uniformly distributed on a set of at least size $(1-\varepsilon)2^n$), and therefore

$$\mathsf{P}\left[\mathbf{D}^*(\mathbf{Q}'_1 \rhd \cdots \rhd \mathbf{Q}'_k) = 1\right] \geq \left(1 - \frac{1}{(1-\varepsilon)2^n}\right) \cdot (1-\varepsilon) = 1 - \varepsilon - \frac{1}{2^n},$$

which in turn implies $\mathsf{P}\left[\mathbf{D}^*(\mathbf{Q}_1 \rhd \cdots \rhd \mathbf{Q}_m) = 1 \,\middle|\, \overline{\mathcal{E}}\right] \geq 1 - \varepsilon - \frac{1}{2^n}$. From this, we conclude $\Delta^{\mathbf{D}^*}(\mathbf{Q}_1 \rhd \cdots \rhd \mathbf{Q}_m, \mathbf{P}) \geq (m-(m-1)\varepsilon)\varepsilon^m - \frac{1}{2^n}$.

5 Conclusions and Open Problems

This paper has presented the first tight analysis of the security amplification properties of the cascade of weak PRPs, both in the one- and two-sided cases. Our main tool is a hardcore lemma (Theorem 2) for computational indistinguishability of discrete interactive cc-stateless systems. It is our belief that the generality of this result makes it suitable to the solution of a number of other problems. For instance, an interesting problem is whether *parallel* and *deterministic* security-amplifying constructions for *arbitrarily* weak pseudorandom *functions* exist. To date, the best known constructions are either randomized [13,10], or only work for moderately weak PRFs [2,10]. Also, quantitative improvements of our results should also be of interest. One may try to minimize the length of the state output by the state sampler or to improve the bound of Theorem 3.

Acknowledgments. The author is grateful to Peter Gaži, Thomas Holenstein, Russell Impagliazzo, Ueli Maurer, and Salil Vadhan for helpful discussions and insightful feedback. This work was done while the author was a graduate student at ETH Zurich, supported in part by the Swiss National Science Foundation (SNF), project no. 200020-113700/1. He is currently at UCSD, partially supported by NSF grant CNS-0716790.

References

1. Bellare, M., Impagliazzo, R., Naor, M.: Does parallel repetition lower the error in computationally sound protocols? In: FOCS 1997: Proceedings of the 38th IEEE Annual Symposium on Foundations of Computer Science, pp. 374–383 (1997)
2. Dodis, Y., Impagliazzo, R., Jaiswal, R., Kabanets, V.: Security amplification for *interactive* cryptographic primitives. In: Reingold, O. (ed.) TCC 2009. LNCS, vol. 5444, pp. 128–145. Springer, Heidelberg (2009)
3. Holenstein, T.: Key agreement from weak bit agreement. In: STOC 2005: Proceedings of the 37th Annual ACM Symposium on Theory of Computing, pp. 664–673 (2005)
4. Holenstein, T.: Pseudorandom generators from one-way functions: A simple construction for any hardness. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 443–461. Springer, Heidelberg (2006)
5. Impagliazzo, R.: Hard-core distributions for somewhat hard problems. In: FOCS 1995: Proceedings of the 36th IEEE Annual Symposium on Foundations of Computer Science, pp. 538–545 (1995)
6. Luby, M., Rackoff, C.: Pseudo-random permutation generators and cryptographic composition. In: STOC 1986: Proceedings of the 18th Annual ACM Symposium on Theory of Computing, pp. 356–363 (1986)
7. Luby, M., Rackoff, C.: How to construct pseudorandom permutations from pseudorandom functions. SIAM Journal on Computing 17(2), 373–386 (1988)
8. Maurer, U.: Indistinguishability of random systems. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 110–132. Springer, Heidelberg (2002)
9. Maurer, U., Pietrzak, K., Renner, R.: Indistinguishability amplification. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 130–149. Springer, Heidelberg (2007)
10. Maurer, U., Tessaro, S.: Computational indistinguishability amplification: Tight product theorems for system composition. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 355–373. Springer, Heidelberg (2009)
11. Maurer, U., Tessaro, S.: A hardcore lemma for computational indistinguishability: Security amplification for arbitrarily weak prgs with optimal stretch. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 237–254. Springer, Heidelberg (2010)
12. Myers, S.: On the development of block-ciphers and pseudo-random function generators using the composition and XOR operators. Master's thesis, University of Toronto (1999)
13. Myers, S.: Efficient amplification of the security of weak pseudo-random function generators. Journal of Cryptology 16, 1–24 (2003)
14. Pietrzak, K., Wikström, D.: Parallel repetition of computationally sound protocols revisited. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 86–102. Springer, Heidelberg (2007)
15. Unruh, D.: Random oracles and auxiliary input. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 205–223. Springer, Heidelberg (2007)
16. Yao, A.C.: Theory and applications of trapdoor functions. In: FOCS 1982: Proceedings of the 23rd IEEE Annual Symposium on Foundations of Computer Science, pp. 80–91 (1982)