

Computational Indistinguishability Amplification: Tight Product Theorems for System Composition*

Ueli Maurer and Stefano Tessaro

Department of Computer Science, ETH Zurich, 8092 Zurich, Switzerland
{maurer,tessaros}@inf.ethz.ch

Abstract. Computational indistinguishability amplification is the problem of strengthening cryptographic primitives whose security is defined by bounding the distinguishing advantage of an efficient distinguisher. Examples include pseudorandom generators (PRGs), pseudorandom functions (PRFs), and pseudorandom permutations (PRPs).

The literature on computational indistinguishability amplification consists only of few isolated results. Yao's XOR-lemma implies, by a hybrid argument, that no efficient distinguisher has advantage better than (roughly) $n2^{m-1}\delta^m$ in distinguishing the XOR of m independent n -bit PRG outputs S_1, \dots, S_m from uniform randomness if no efficient distinguisher has advantage more than δ in distinguishing S_i from a uniform n -bit string. The factor 2^{m-1} allows for security amplification only if $\delta < \frac{1}{2}$. For the case of PRFs, a random-offset XOR-construction of Myers was the first result to achieve *strong* security amplification, i.e., also for $\frac{1}{2} \leq \delta < 1$.

This paper proposes a systematic treatment of computational indistinguishability amplification. We generalize and improve the above product theorem for the XOR of PRGs along five axes. First, we prove the *tight* information-theoretic bound $2^{m-1}\delta^m$ (without factor n) also for the computational setting. Second, we prove results for *interactive* systems (e.g. PRFs or PRPs). Third, we consider the general class of *neutralizing combination constructions*, not just XOR. As an application, this yields the first indistinguishability amplification results for the cascade of PRPs (i.e., block ciphers) converting a weak PRP into an arbitrarily strong PRP, both for single-sided and two-sided queries. Fourth, *strong* security amplification is achieved for a subclass of neutralizing constructions which includes as a special case the construction of Myers. As an application we obtain highly practical optimal security amplification for block ciphers, simply by adding random offsets at the input and output of the cascade. Fifth, we show strong security amplification also for *weakened assumptions* like security against random-input (as opposed to chosen-input) attacks.

A key technique is a generalization of Yao's XOR-lemma to (interactive) systems which is of independent interest.

* This research was partially supported by the Swiss National Science Foundation (SNF), project no. 200020-113700/1.

1 Introduction

1.1 Security Amplification

The security of all computationally secure cryptographic systems, even of those called “provably secure” in the literature, relies on unproven assumptions about the underlying cryptographic primitives. Typical assumptions are that a certain construction is a one-way function (OWF), a collision-resistant hash function, a pseudorandom generator (PRG), a pseudorandom function (PRF), a pseudorandom permutation (PRP), etc. To weaken these assumptions is both a fundamental challenge in the theory of cryptography and a major goal for the cautious and prudent design of practical cryptographic systems. Many reductions of strong primitives to weak primitives are known. For example, one of the outstanding results is the construction of a PRG from any OWF [13]. However, this reduction, like many other reductions, is highly inefficient and, while of high theoretical value, not of practical relevance.

A specific way to weaken an assumption is to require only that the security property is mildly true. For instance, a δ -OWF can be efficiently inverted with probability at most δ (rather than a negligible quantity for a regular OWF). Similarly, for a δ -PRG no efficient distinguisher has an advantage more than δ in distinguishing its output from a uniform random string. The corresponding definitions of a δ -PRF, a δ -PRP, etc., are straight-forward. Such a weakened assumption is more likely to be true. For example, it is more conservative to only assume that AES is a 0.99-PRP rather than a fully secure PRP.

The natural question is whether several weak primitives can be efficiently combined to obtain a stronger version of the primitive, ideally one with the full-fledged security property.¹ This is called *security amplification*, in some cases *hardness amplification*. The classical result on security amplification due to Yao [35] is that the parallel composition of m δ -OWFs results in a $(\delta^m + \nu)$ -OWF, where ν is some negligible quantity and for any $\delta < 1$, δ^m can be made negligible for large enough m . Security amplifications of a wide range of cryptographic primitives has subsequently been considered, including for example regular OWFs and OWPs [9,11], two-party protocols [1,29,30,34,12], key-agreement and public-key encryption [7,15,16], collision-resistant hash functions [4], and watermarking schemes [17].²

The term *indistinguishability amplification* refers to security amplification when the relevant security quantity is the *distinguishing advantage* for the best distinguisher from a certain class of distinguishers, typically the class of efficient distinguishers.

¹ Typically one considers several independent instantiations of the *same* weak primitive, but most results actually hold for several different instantiations.

² So-called combiners [14] are another method for relaxing security assumptions: They guarantee that a construction involving several instantiations of a primitive is (fully) secure if at least one (or several, but not all) of them are (fully) secure. However, they do not amplify security of the underlying primitives.

1.2 The XOR-Lemma and Amplification for PRGs

Before we discuss the XOR-lemma, let us compare the prediction advantage and the distinguishing advantage of a biased bit, in an information-theoretic setting, i.e., allowing arbitrary computing power. A bit with bias ϵ takes on the two values with probabilities $\frac{1}{2} - \epsilon$ and $\frac{1}{2} + \epsilon$. When such a bit must be guessed, one would choose the more likely value and be correct with probability $\frac{1}{2} + \epsilon$. To calibrate the guessing advantage, between 0 (when $\epsilon = 0$) and 1 (when the bit is fixed, i.e., $\epsilon = \frac{1}{2}$), one defines the advantage to be 2ϵ . In contrast, the distinguishing advantage is defined as ϵ (with no factor 2) since it is naturally defined for general random variables (not only bits) as the distance of the probability distribution from the uniform one.

As an example, consider two independent bits with biases ϵ_1 and ϵ_2 . It is easy to see that the bias of the XOR is $2\epsilon_1\epsilon_2$. For instance, the XOR of a 0.1-biased bit (40/60) and a 0.2-biased bit (30/70) is a 0.04-biased bit (46/54), where $0.04 = 2 \cdot 0.01 \cdot 0.02$. More generally, the bias of the XOR of m bits is 2^{m-1} times the product of the biases. For the XOR of m bit-strings S_1, \dots, S_m of length n , where S_i has distance δ_i from a uniform n -bit string, the distance from uniform of the XOR of the strings, $S_1 \oplus S_2 \oplus \dots \oplus S_m$, is bounded by $2^{m-1} \prod_{i=1}^m \delta_i$. This bound is tight, as for example the case $n = 1$ discussed above illustrates.

Let us now move to the computational setting, i.e., to Yao's XOR-lemma [35,10], which is much more involved and is another seminal security amplification result. One typically considers a predicate $B(x)$ of the input of a OWF f which is hard to guess when given the output $f(x)$, for uniformly chosen x . But the setting of the XOR-lemma is more general. It states³ that if for bits B_1, \dots, B_m the advantage in guessing B_i given some correlated information X_i is at most α_i for any algorithm with complexity t' , then no algorithm with complexity t has advantage more than $\prod_{i=1}^m \alpha_i + \gamma$ in guessing their XOR-sum, i.e., $B_1 \oplus \dots \oplus B_m$, given X_1, \dots, X_m , where γ can be made arbitrarily small, at the cost of making t smaller with respect to t' .⁴ In terms of distinguishing advantages δ_i , the bound is $2^{m-1} \prod_{i=1}^m \delta_i + \gamma$ (for the reasons described above).

Moreover, a standard hybrid argument, to use the unpredictability of bits to prove the indistinguishability of bit-strings, implies an indistinguishability amplification result for PRGs. Consider m independent PRG outputs, S_1, \dots, S_m , each an n -bit string. If no distinguisher with complexity t' has advantage more than δ_i in distinguishing S_i from a uniform random n -bit string, then no distinguisher with complexity (roughly) t has advantage more than $n(2^{m-1} \prod_{i=1}^m \delta_i + \gamma)$ in

³ In fact, one needs a “tight” version of the XOR-lemma for this statement to hold, such as the one by Levin [20,10], or one obtained from a tight hard-core lemma (e.g. [15]) via the techniques of [18].

⁴ As usual in complexity-theoretic hardness amplification, we experience an *unavoidable* [31] trade-off between the choice of γ (the tightness of the bound) and the complexity of the reduction.

distinguishing $S_1 \oplus S_2 \oplus \dots \oplus S_m$ from a uniform random n -bit string.⁵ The factor n comes from the hybrid argument over the individual bits of the bit-string.

As explained, the factor 2^{m-1} is unavoidable, since it holds even in the information-theoretic setting. Unfortunately, it means that an amplification can be achieved only if the component constructions are better than $\frac{1}{2}$ -secure, i.e., if $\delta_i < \frac{1}{2}$.

1.3 Natural Questions and Previous Results

The above discussion suggests a number of natural questions. (1) Can the factor n in the bound for the XOR of PRGs be eliminated, to obtain a tight bound, namely the equivalent of the information-theoretic counterpart? (2) Can the result be extended to the XOR of PRFs, i.e., primitives for which the security is defined by an *interactive* game, not by the (static) indistinguishability of random variables? (3) If the answer is “yes”, can such a result be extended to other constructions, most importantly the cascade of PRPs? (4) Can the factor 2^{m-1} be eliminated so that security amplification from arbitrarily weak components can be achieved? We will answer all these questions positively.

In the information-theoretic setting, questions 2 and 3 were answered by Maurer, Pietrzak, and Renner [24], whose abstract approach we follow, and the special case of permutations had previously been solved by Vaudenay [32,33]. In contrast, there are only a few isolated results on *computational* indistinguishability amplification, which we now discuss. Myers [27] was the first to consider security amplification for PRFs. Interestingly, he did not solve question 2 above, which remained open, but he actually solved part of question 4. More precisely, he showed for the XOR of PRFs, with the modification that for each PRF a random (secret) offset is XORed to the input, that the stronger bound (without the factor 2^{m-1}) can be achieved. However, his treatment is specific for his construction and does not extend to other settings like the cascade of PRPs. Dodis et al. [6] addressed question 2 and gave a positive answer using techniques originating from the setting of hardness amplification of weakly verifiable puzzles [3,19]. However, their focus is on general interactive cryptographic primitives, including for example message authentication codes (MACs), and the resulting bound for the case of PRFs depends on the number of queries the distinguisher is allowed to ask and is not optimal.

Little is known about the cascade of weak PRPs, which is perhaps the case of highest practical interest as it addresses security amplification for block ciphers.⁶

⁵ It is not clear to us whether this fact has been published, or is unpublished but well-known folklore, or not so well-known (see also [6] for a similar statement about security amplification for the XOR of PRGs).

⁶ Cascades of block ciphers were considered by Even and Goldreich [8] and Maurer and Massey [23], but those results only prove that the cascade is as secure as the strongest component (with no amplification), i.e., that the cascade is a combiner for encryption. Bellare and Rogaway [2] showed a certain security amplification (of a different type) for cascade encryption in the ideal cipher model, which is a purely information-theoretic consideration.

Luby and Rackoff [21] proved an amplification result for the cascade of *two* weak PRPs. This result was extended by Myers [26] to the cascade of a small number of PRPs, but he notes that this result falls short of constructing a (regular) PRP from a weak PRP and states this as an open problem, which we solve.

1.4 Contributions of This Paper

In our attempt at solving the different open questions explained above, we take a very general approach, not targeted at specific constructions. The goal is to develop a deeper and more general understanding and to prove results of a generality that can be useful for other applications.

A first result is a generalization of the XOR-lemma to interactive systems. If a system (as opposed to a random variable for the standard XOR-lemma) of a general type depends on a bit, and no efficient algorithm with access to the system can predict the bit better than with a certain advantage, then the advantage in predicting the XOR of several such bits is bounded by the product of the individual advantages, even if the predictor has complete and arbitrary independent access to all the involved systems.

The XOR of strings or (of the output) of systems, as well as the cascade of systems implementing permutations, are both special cases of a more general concept which was called *neutralizing construction* in [24]. Intuitively, a construction involving several component systems is neutralizing if it is equivalent to an ideal system whenever one component is ideal. For example, the XOR of several PRFs is equivalent to a truly random function if (any) one of the PRFs is replaced by a truly random function.

We prove two tight general product theorems. The first theorem relies on the XOR-lemma and shows that for *all* neutralizing constructions the distinguishing advantage of the combined system is 2^{m-1} times the product of the individual advantages, which is optimal. The second theorem gets rid of the factor 2^{m-1} by considering a special class of *randomized* neutralizing constructions. The applications mentioned in the abstract and the previous sections follow directly from these general theorems.⁷ In particular, one application is a highly practical construction for optimal security amplification for block ciphers, simply by adding random offsets at the input and output of the cascade.

1.5 Notational Preliminaries

Throughout this paper, we use calligraphic letters $\mathcal{X}, \mathcal{Y}, \dots$ to denote sets, upper-case letters X, Y, \dots to denote random variables, and lower-case letters x, y, \dots denote the values they take. Moreover, $P[\mathcal{A}]$ denotes the probability of an event \mathcal{A} , while we use the shorthand $P_X(x) := P[X = x]$, and denote by P_X the probability distribution of X and by $E[X]$ its expected value.

⁷ For each application of the second theorem, one also needs an information-theoretic indistinguishability proof based on the conditional equivalence of two systems, conditioned on an event that must be proved to be unlikely to occur.

We consider *interactive* randomized stateful algorithms in some a-priori fixed (but otherwise unspecified) RAM model of computation. In particular, such an algorithm keeps a state (consisting say of the memory space it uses), and answers each query depending on the input of this query, some coin flips, the current state (which is possibly 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. Additionally, we say that the algorithm A has *time complexity* t_A (where t_A is a function $\mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$) if the sum of the length of the description of A and the total number of steps of A is at most $t_A(q, s)$ for all sequences of q queries, all compatible initial states with size s , and all compatible interactions with an underlying system. We use the shorthand $t_A(q) := t_A(q, 0)$.

This paper adopts a *concrete* approach, i.e. we do not use asymptotics and statements are inherently non-uniform. Still, all results can be extended to the uniform setting by using standard techniques. We comment on the necessary changes in the full version of this paper.

2 Discrete Systems and Constructions

DISCRETE SYSTEMS, CONSTRUCTIONS, AND DISTINGUISHERS. This paper deals with the general notion of a (single-interface) *discrete system* \mathbf{F} taking inputs X_1, X_2, \dots and returning outputs Y_1, Y_2, \dots , where the i -th output Y_i depends (probabilistically) on the first i inputs $X^i = [X_1, \dots, X_i]$ as well as on all previous $i - 1$ outputs $Y^{i-1} = [Y_1, \dots, Y_{i-1}]$. (If all inputs and outputs are in sets \mathcal{X} and \mathcal{Y} , respectively, we call \mathbf{F} an $(\mathcal{X}, \mathcal{Y})$ -system.) Its input-output behavior is minimally described (see e.g. [22]) by the (infinite) sequence of conditional probability distributions $\mathbf{p}_{Y_i|X^i Y^{i-1}}^{\mathbf{F}}$ (for all $i \geq 1$). In general, we use the name “system” (as well as \mathbf{F}) interchangeably to denote both the input-output behavior determined by conditional probability distributions and an actual discrete system realizing this behavior. It thus makes sense to say that two systems \mathbf{F}, \mathbf{G} are *equivalent* (denoted $\mathbf{F} \equiv \mathbf{G}$) if they have the same input-output behavior. A *random variable* X is the simplest type of system, which answers each query with the same value X .

With $\mathbf{C}(\cdot)$ we denote a *construction* invoking one or more underlying compatible *subsystems*, whereas $\mathbf{C}(\mathbf{F})$, $\mathbf{C}(\mathbf{F}, \mathbf{G})$, etc denote the systems obtained when \mathbf{C} is instantiated with \mathbf{F} (and \mathbf{G}). The shorthand $\mathbf{C}(\mathbf{F}, \cdot)$ indicates the construction that behaves as $\mathbf{C}(\mathbf{F}, \mathbf{G})$ given access to the subsystem \mathbf{G} . (All notations extend naturally to constructions with more than two subsystems.) A *distinguisher* \mathbf{D} is a system interacting with another system \mathbf{F} giving inputs X_1, X_2, \dots and obtaining outputs Y_1, Y_2, \dots , outputting a decision bit after a certain number q of queries depending on the transcript (X^q, Y^q) : In particular, we denote as $\mathbf{P}[\mathbf{D}(\mathbf{F}) = 1]$ the probability that it outputs 1.

We say that an interactive algorithm A implements a system \mathbf{F} or a construction $\mathbf{C}(\cdot)$ if it has the same input-output behavior as \mathbf{F} and $\mathbf{C}(\cdot)$, respectively.

In particular, we use A (rather than \mathbf{F}) whenever we want to stress that we use the particular implementation A of \mathbf{F} .

DISTINGUISHING ADVANTAGES. The *distinguishing advantage of a distinguisher \mathbf{D} in distinguishing two systems \mathbf{F} and \mathbf{G}* is the quantity

$$\Delta^{\mathbf{D}}(\mathbf{F}, \mathbf{G}) := |\mathbb{P}[\mathbf{D}(\mathbf{F}) = 1] - \mathbb{P}[\mathbf{D}(\mathbf{G}) = 1]|.$$

We denote as $\Delta_t(\mathbf{F}, \mathbf{G})$, $\Delta_q(\mathbf{F}, \mathbf{G})$, and $\Delta_{t,q}(\mathbf{F}, \mathbf{G})$ the best distinguishing advantages $\Delta^{\mathbf{D}}(\mathbf{F}, \mathbf{G})$ taken over all distinguishers with time complexity at most t , issuing at most q queries, or both, respectively.

SYSTEM COMPOSITION. Given m systems $\mathbf{F}_1, \dots, \mathbf{F}_m$, we use the shorthand $\mathbf{F}_1 \parallel \dots \parallel \mathbf{F}_m$ to denote their *parallel composition*, i.e., the system allowing parallel concurrent interaction with the (independent) m systems.⁸ Moreover, for $(\mathcal{X}, \mathcal{Y})$ -systems \mathbf{F} and \mathbf{G} , and a random bit B (with distribution \mathbb{P}_B), the system $\langle \mathbf{F}, \mathbf{G} \rangle_B$ acts as \mathbf{F} if $B = 0$, and as \mathbf{G} otherwise. Additionally, for any *quasi-group operation*⁹ \star on \mathcal{Y} the $(\mathcal{X}, \mathcal{Y})$ -system $\mathbf{F} \star \mathbf{G}$ on input x invokes both \mathbf{F}, \mathbf{G} with input x , obtaining y, y' , and returns $y \star y'$.¹⁰ Also, for an $(\mathcal{X}, \mathcal{Y})$ -system \mathbf{P} and a $(\mathcal{Y}, \mathcal{Z})$ -system \mathbf{Q} we denote with $\mathbf{P} \triangleright \mathbf{Q}$ the *cascade of \mathbf{P} and \mathbf{Q}* , i.e., the system which on input x first invokes \mathbf{P} on this input, and the resulting output is fed into \mathbf{Q} to obtain the final output.

STATELESS SYSTEMS. We say that a system \mathbf{F} is *stateless* if there exists a conditional probability distribution $\mathbb{p}_{Y|X}^{\mathbf{F}}$ such that $\mathbb{p}_{Y_i|X^i Y^{i-1}}^{\mathbf{F}}(y_i, x^i, y^{i-1}) = \mathbb{p}_{Y|X}^{\mathbf{F}}(y_i, x_i)$ for all $y_i, x^i = [x_1, \dots, x_i]$, and $y^{i-1} = [y_1, \dots, y_{i-1}]$. Moreover, the system \mathbf{F} is *convex-combination stateless (cc-stateless, for short)* if there exists a random variable S and a construction $\mathbf{F}(\cdot)$ (we abuse notation by recycling the letter \mathbf{F}) such that $\mathbf{F}(S) \equiv \mathbf{F}$, and $\mathbf{F}(s)$ is stateless for *all* values s taken by S . Depending on the context, S may be e.g. a seed, a key, or an internal function table. A non-trivial example of a cc-stateless system is a randomized encryption scheme, which takes a secret key and encrypts each message with independent randomness. Note that $\langle \mathbf{F}, \mathbf{G} \rangle_B$ is cc-stateless if both \mathbf{F}, \mathbf{G} are cc-stateless.

RANDOM FUNCTIONS. A *random function* $\mathbf{F} : \mathcal{X} \rightarrow \mathcal{Y}$ is an $(\mathcal{X}, \mathcal{Y})$ -system which answers consistently, i.e. $X_i = X_j$ implies $Y_i = Y_j$. It is called a *random permutation* if additionally $Y_i = Y_j$ implies $X_i = X_j$. A cc-stateless random function $\mathbf{F} : \mathcal{X} \rightarrow \mathcal{Y}$ is in particular such that $\mathbf{F} \equiv \mathbf{F}(S)$ where $\mathbf{F}(\cdot)$ is deterministic and $\mathbf{F}(s)$ is a function $\mathcal{X} \rightarrow \mathcal{Y}$ for all s . (This is sometimes called a *keyed function family*, but we also consider the case where s is huge and is hence not a key.) Special cases are a *uniform random function (URF)* $\mathbf{R} : \mathcal{X} \rightarrow \mathcal{Y}$ and a *uniform*

⁸ The systems do not interact with each other, and each query to the parallel composition is addressed to one of the systems.

⁹ That is, given $a, c \in \mathcal{Y}$ (or $b, c \in \mathcal{Y}$) there exists a unique b (a) such that $a \star b = c$. An example is bit-wise XOR \oplus for $\mathcal{Y} = \{0, 1\}^n$, but any group operation is a quasi-group operation as well.

¹⁰ We denote as $\mathbf{F}_1 \star \dots \star \mathbf{F}_m$ the system $(\dots((\mathbf{F}_1 \star \mathbf{F}_2) \star \mathbf{F}_3) \dots) \star \mathbf{F}_m$.

random permutation (URP) $\mathbf{P} : \mathcal{X} \rightarrow \mathcal{X}$ that realize a uniformly chosen function $\mathcal{X} \rightarrow \mathcal{Y}$ and permutation $\mathcal{X} \rightarrow \mathcal{X}$, respectively. We denote as $\mathbf{F}(s, x)$ the evaluation of \mathbf{F} with key s and input x .

Informally, in an asymptotic setting, it is convenient to say that an efficient $\mathbf{F}(\cdot)$ is a δ -pseudorandom function (PRF) if $\Delta_{t,q}(\mathbf{F}(S), \mathbf{R}) \leq \delta + \text{negl}$ for a (short) key S , a URF \mathbf{R} , all polynomial t and q , and some negligible¹¹ function negl . Analogously, if an efficient $\mathbf{Q}(\cdot)$ implements a permutation for all keys, it is called a δ -pseudorandom permutation (PRP) if $\Delta_{t,q}(\mathbf{Q}(S), \mathbf{P}) \leq \delta + \text{negl}$ for a URP \mathbf{P} and for all polynomial t and q .

The inverse \mathbf{Q}^{-1} of a cc-stateless random permutation \mathbf{Q} is well-defined, and $\langle \mathbf{Q} \rangle$ is the system accepting *forward queries* $(x, +)$ (answered by $\mathbf{Q}(s, x)$ on key s) and *backward queries* $(y, -)$ (answered as $\mathbf{Q}^{-1}(s, y)$). In particular $\langle \mathbf{Q} \rangle \triangleright \langle \mathbf{Q}' \rangle$ stands for the system $\langle \mathbf{Q} \triangleright \mathbf{Q}' \rangle$. An efficient $\mathbf{Q}(\cdot)$ is called a δ -two-sided PRP¹² if $\Delta_{t,q}(\langle \mathbf{Q}(S) \rangle, \langle \mathbf{P} \rangle) \leq \epsilon + \text{negl}$ for all polynomial q and t . (Of course, one assumes that backward queries can be computed efficiently given s .)

NEUTRALIZING CONSTRUCTIONS. A construction $\mathbf{C}(\cdot)$ is *neutralizing* [24] for systems $\mathbf{F}_1, \dots, \mathbf{F}_m$ and ideal systems $\mathbf{I}_1, \dots, \mathbf{I}_m$, if for $\mathbf{S}_i \in \{\mathbf{F}_i, \mathbf{I}_i\}$ ($i = 1, \dots, m$) we have $\mathbf{C}(\mathbf{S}_1, \dots, \mathbf{S}_m) \equiv \mathbf{C}(\mathbf{I}_1, \dots, \mathbf{I}_m)$ whenever there exists some i with $\mathbf{S}_i = \mathbf{I}_i$.¹³

Every quasi-group operation \star on a set \mathcal{Y} induces a construction $\mathbf{C}(\cdot)$ such that $\mathbf{C}(\mathbf{F}, \mathbf{G}) := \mathbf{F} \star \mathbf{G}$ which is neutralizing for random functions $\mathbf{F}, \mathbf{G} : \mathcal{X} \rightarrow \mathcal{Y}$ and ideal systems \mathbf{I}, \mathbf{J} being independent URFs. In particular, $\mathbf{I} \star \mathbf{J}$ is also a URF. As a special case, this result holds for random variables X, Y over \mathcal{Y} , the ideal systems being uniform random elements of \mathcal{Y} . Moreover, the cascade operator \triangleright induces a construction $\mathbf{C}'(\cdot)$ with $\mathbf{C}'(\mathbf{Q}_1, \mathbf{Q}_2) := \mathbf{Q}_1 \triangleright \mathbf{Q}_2$ which is neutralizing for any two cc-stateless random permutations $\mathbf{Q}_1, \mathbf{Q}_2 : \mathcal{X} \rightarrow \mathcal{X}$ (in fact \mathbf{Q}_1 can possibly be stateful) with ideal systems \mathbf{I}, \mathbf{J} both URPs $\mathcal{X} \rightarrow \mathcal{X}$. In particular, $\mathbf{I} \triangleright \mathbf{J}$ is also a URP. If \mathbf{Q}_1 is cc-stateless, then the same result holds even in the two-sided case for $\langle \mathbf{Q}_1 \rangle$ and $\langle \mathbf{Q}_2 \rangle$ (with ideal system $\langle \mathbf{P} \rangle$ for a URP \mathbf{P}). Both constructions extend naturally to an arbitrary number of subsystems.

3 A General Product Theorem for Neutralizing Constructions

This section presents a very general product theorem showing computational indistinguishability for *every* neutralizing construction. This result relies on a

¹¹ Recall that a function $\nu : \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ is *negligible* if it vanishes faster than the inverse of any polynomial.

¹² In the literature the name *strong* PRP is commonly used, but this term is slightly confusing in the context of this paper.

¹³ Neutralizing constructions capture the notion of a *combiner* [14] for computational indistinguishability properties: Whenever at least one system \mathbf{S}_i is computationally indistinguishable from \mathbf{I}_i , then $\mathbf{C}(\mathbf{S}_1, \dots, \mathbf{S}_m)$ is computationally indistinguishable from $\mathbf{C}(\mathbf{I}_1, \dots, \mathbf{I}_m)$.

generalization of Yao’s XOR-Lemma to discrete interactive systems, which is presented first, and is of independent interest.

3.1 The Generalized XOR-Lemma

SYSTEM-BIT PAIRS. A *system-bit pair* is a system of the form (\mathbf{F}, B) , where $B \in \{0, 1\}$ is a bit value, which is (generally) correlated with the system \mathbf{F} . This can formally be described by the distribution P_B of B and the two systems \mathbf{F}_0 and \mathbf{F}_1 conditioned on the value taken by B , i.e. $(\mathbf{F}, B) = ((\mathbf{F}_0, \mathbf{F}_1)_B, B)$. A possible system-bit pair is a URF $\mathbf{R} : \{0, 1\}^m \rightarrow \{0, 1\}$ and the parity of its function table. The following quantity characterizes the performance of an adversary¹⁴ \mathbf{A} in guessing the bit B when given access to \mathbf{F} only.

Definition 1. *The guessing advantage of an adversary \mathbf{A} in guessing B for a system-bit pair (\mathbf{F}, B) is the quantity $\Gamma^{\mathbf{A}}(\mathbf{F}, B) := 2 \cdot P[\mathbf{A}(\mathbf{F}) = B] - 1$. Additionally, we denote as $\Gamma_{t,q}(\mathbf{F}, B)$ the maximal guessing advantage $\Gamma^{\mathbf{A}}(\mathbf{F}, B)$ taken over all q -query adversaries \mathbf{A} with complexity at most t .*

Note that $\Gamma^{\mathbf{A}}(\mathbf{F}, B) \in [-1, 1]$, where 1 means that \mathbf{A} is able to perfectly predict B by interacting with \mathbf{F} , while -1 means that \mathbf{A} is never correct.¹⁵ The following connection between the guessing and the distinguishing advantages is well known (cf. e.g. [24]).

Lemma 1. *For all \mathbf{F} , \mathbf{G} , and \mathbf{D} , we have $\Delta^{\mathbf{D}}(\mathbf{F}, \mathbf{G}) = |\Gamma^{\mathbf{D}}((\mathbf{F}, \mathbf{G})_B, B)|$ for a uniform bit $B \in \{0, 1\}$.*

THE XOR-LEMMA. Given m system-bit pairs $(\mathbf{G}_1, B_1), \dots, (\mathbf{G}_m, B_m)$, we are interested in the advantage $\Gamma_{t,q_1,\dots,q_m}(\mathbf{G}_1 \parallel \dots \parallel \mathbf{G}_m, B_1 \oplus \dots \oplus B_m)$ of guessing the bit $B_1 \oplus \dots \oplus B_m$ given *parallel* access to the systems $\mathbf{G}_1, \dots, \mathbf{G}_m$, where at most q_i queries to each system \mathbf{G}_i are allowed. That is, we consider the most general attack where the adversary can query each subsystem \mathbf{G}_i individually at most q_i times, *adaptively* depending on the answers of queries to other subsystems. We show that the advantage is upper bounded by the *product* of the individual advantages $\Gamma_{t',q'}(\mathbf{G}_i, B_i)$ for $i = 1, \dots, m$ (for appropriate t', q'), with an extra additive term $\gamma > 0$ which can be made arbitrarily small (but influences the efficiency of the reduction). The result holds provided that all but one of the system-bit pairs are cc-stateless. (Note that the fact that (\mathbf{G}_i, B_i) is cc-stateless implies that \mathbf{G}_i is cc-stateless, but the converse is not always true.) Our result generalizes the original XOR-lemma by Yao [35,10], which considered the special case of system-bit pairs (X_i, B_i) , where X_i is a random variable.

We stress that our result only requires the ability to efficiently implement the cc-stateless system-bit pairs $(\mathbf{G}_i, B_i) = (\mathbf{G}_i(S), B_i(S))$. This may be possible, for instance by using a *stateful* algorithm, even if $\mathbf{G}(\cdot)$ and $B(\cdot)$ themselves are

¹⁴ We stress that distinguishers and adversaries are objects of the same type. The name adversary is used to stress the fact that we are not exclusively considering a distinguishing scenario.

¹⁵ In particular, flipping the output bit of such a \mathbf{A} yields one which is always correct.

not efficiently computable: In fact, S may even be exponentially large. As an example, the aforementioned system-bit pair (\mathbf{R}, B) , where $\mathbf{R} : \{0, 1\}^n \rightarrow \{0, 1\}$ is a URF, and B is the parity of its function table, is clearly cc-stateless, and can efficiently be implemented by first sampling a random B , and then answering queries to \mathbf{R} with independent random bits, with the exception of the last one, which is answered so that the parity equals B .

In the following, we define the quantity $\varphi := 2 \left(\frac{24m}{\gamma} \right)^2 \cdot \ln \left(\frac{7m}{\gamma} \right)$ for understood m and γ . Also, t_{G_i} and s_{G_i} are the time and space¹⁶ complexities of some implementation G_i of the system \mathbf{G}_i , whereas $t_{(G_i, B_i)}$ is the time-complexity of an implementation of the pair (\mathbf{G}_i, B_i) . (Note that an efficient implementation of the latter implies one for the former, but we allow for this distinction.) For all i , we denote $l_i := s_{G_i}(q_i \cdot \varphi)$ and $l_{<i} := \sum_{j=1}^{i-1} l_j$ (for understood q_1, \dots, q_{m-1}).

Theorem 1 (XOR-Lemma). *Let $(\mathbf{G}_1, B_1), \dots, (\mathbf{G}_{m-1}, B_{m-1})$ be cc-stateless system-bit pairs, and let (\mathbf{G}_m, B_m) be an arbitrary system-bit pair. For all $t, q_1, \dots, q_m, \gamma > 0$,*

$$\Gamma_{t, q_1, \dots, q_m}(\mathbf{G}_1 \| \dots \| \mathbf{G}_m, B_1 \oplus \dots \oplus B_m) \leq \prod_{i=1}^m \Gamma_{t'_i, q'_i}(\mathbf{G}_i, B_i) + \gamma,$$

where $t'_i := l_{<i} + \varphi \cdot \left[t + \mathcal{O} \left(\sum_{j=1}^{i-1} t_{G_j}(q_j, l_j) + \sum_{j=i+1}^m t_{(G_j, B_j)}(q_j) \right) \right]$ and $q'_i := \varphi \cdot q_i$ for $i = 1, \dots, m - 1$, whereas $t_m := l_{<m} + t + \mathcal{O} \left(\sum_{j=1}^{m-1} t_{G_j}(q_j, l_j) \right)$ and $q'_m := q_m$.

The asymmetry of our proof technique allows (\mathbf{G}_m, B_m) to be fully stateful.¹⁷ Furthermore, both t'_m and q'_m are much smaller than the corresponding terms t'_i and q'_i for $i = 1, \dots, m - 1$. The following paragraph provides a proof sketch for the case $m = 2$. The full proof is deferred to the full version of this paper.

PROOF IDEA FOR $m = 2$. The proof follows similar lines as Levin’s proof of the XOR-lemma [20,10], but with some major differences due to the peculiarities of reactive systems. For simplicity, we let $(\mathbf{F}, B) = (\mathbf{G}_1, B_1)$ and $(\mathbf{G}, C) = (\mathbf{G}_2, B_2)$. Let \mathbf{A} be an adversary with $\Gamma^{\mathbf{A}}(\mathbf{F} \| \mathbf{G}, B \oplus C) > \delta \cdot \epsilon + \gamma$. We show that either there exists an adversary \mathbf{A}' such that $\Gamma^{\mathbf{A}'}(\mathbf{F}, B) > \delta$ or there exists an adversary \mathbf{A}'' such that $\Gamma^{\mathbf{A}''}(\mathbf{G}, C) > \epsilon$, contradicting the assumed hardness of (\mathbf{F}, B) and/or (\mathbf{G}, C) . The time complexities of \mathbf{A}' and \mathbf{A}'' are strictly related to the one of \mathbf{A} . Recall that the pair $(\mathbf{F}, B) = (\mathbf{F}(S), B(S))$ is cc-stateless, and for all values s taken by the random variable S we define

$$\alpha_1(s) := \Gamma^{\mathbf{A}}(\mathbf{F}(s) \| \mathbf{G}, 1 \oplus C) \text{ and } \alpha(s) := \Gamma^{\mathbf{A}}(\mathbf{F}(s) \| \mathbf{G}, B(s) \oplus C).$$

¹⁶ i.e. the maximal size of the state after the given number of queries.

¹⁷ An orthogonal generalization of the XOR-lemma for stateful interactive systems was proposed by Halevi and Rabin [12]. However, it relies on *sequential* (rather than parallel) access to the systems $\mathbf{G}_1, \dots, \mathbf{G}_m$, which is not sufficient for the applications of this paper.

By definition, $E[\alpha(S)] > \delta \cdot \epsilon + \gamma$. Moreover $\alpha(s) = \alpha_1(s)$ if $B(s) = 1$, and $\alpha(s) = -\alpha_1(s)$ otherwise. This implies that $\alpha_1(S)$ has good correlation with $B(S)$, as an adversary \mathbf{A}' outputting 1 with probability $\frac{1}{2} + \frac{\alpha_1(s)}{2}$ (when given access to $\mathbf{F}(s)$) has advantage at least $\delta \cdot \epsilon + \gamma$. If $|\alpha_1(s)| = |\alpha(s)| \leq \epsilon$ holds for all s , then the advantage can be amplified to be larger than δ by outputting 1 with probability $\frac{1}{2} + \frac{\alpha_1(s)}{2\epsilon}$. Of course, \mathbf{A}' does not know $\alpha_1(s)$, but a statistical estimate can be obtained by repeated interaction with $\mathbf{F}(s)$, as it is stateless: The term γ compensates the possible estimation error.

Note that the existence of a single value s with the property that $|\alpha_1(s)| > \epsilon$ implies that there exists a bit b such that the adversary $\mathbf{A}'' := \mathbf{A}(\mathbf{F}(s)||\cdot) \oplus b$ has advantage larger than ϵ in guessing C from \mathbf{G} , i.e., \mathbf{A}'' is the adversary that simulates the execution of \mathbf{A} with the parallel composition of $\mathbf{F}(s)$ and the given system \mathbf{G} , and outputs \mathbf{A} 's output XORed with b . But such adversary \mathbf{A}'' is not necessarily efficient, because an efficient implementation of $\mathbf{F}(s)$ may not exist. To overcome this issue, we show that for the above adversary \mathbf{A}' to succeed, it is sufficient that the probability over the choice of S that $|\alpha_1(S)| > \epsilon + \gamma/4$ is smaller than $\gamma/4$. Furthermore, if this probability is at least $\gamma/4$, a probabilistic argument yields a (sufficiently) small state σ for the (efficient) implementation F of \mathbf{F} and a (fixed) bit b such that the *efficient* adversary $\mathbf{A}'' := \mathbf{A}(F[\sigma]||\cdot) \oplus b$ achieves advantage at least ϵ .

3.2 A Product Theorem from the XOR-Lemma

Throughout this section, let $\mathbf{C}(\cdot)$ be a neutralizing construction for systems $\mathbf{F}_1, \dots, \mathbf{F}_m, \mathbf{I}_1, \dots, \mathbf{I}_m$ (of which all but \mathbf{F}_m and \mathbf{I}_m have to be cc-stateless). We provide a very general product theorem upper bounding the distinguishing advantage $\Delta_{t,q}(\mathbf{C}(\mathbf{F}_1, \dots, \mathbf{F}_m), \mathbf{C}(\mathbf{I}_1, \dots, \mathbf{I}_m))$ in terms of the individual advantages $\Delta_{t'_i, q'_i}(\mathbf{F}_i, \mathbf{I}_i)$ (for some related t'_i, q'_i). The theorem is a computational version of the information-theoretic product theorem from [24]: In particular, we inherit the same bounds, with an unavoidable additive term.

The theorem relies on the canonical implementation $\langle F_i, I_i \rangle_{B_i}$ of $\langle \mathbf{F}_i, \mathbf{I}_i \rangle_{B_i}$ which chooses a random bit $B_i \in \{0, 1\}$ and answers each query using the implementations F_i and I_i (with respective complexities t_{F_i} and t_{I_i}) of \mathbf{F}_i or of \mathbf{I}_i , respectively, depending on the value of B_i . (B_i is in particular part of the state.) It can be implemented with complexity $t_{\langle F_i, I_i \rangle_{B_i}}(q, s) = \max\{t_{F_i}(q, s), t_{I_i}(q, s)\} + \mathcal{O}(1)$. This also yields an implementation of $\langle \langle \mathbf{F}_i, \mathbf{I}_i \rangle_{B_i}, B_i \rangle$ with the same complexity (by additionally outputting the bit B_i). Finally, we let l_i and $l_{<i}$ as above be defined with respect to $\langle F_i, I_i \rangle_{B_i}$, and let t_C be the time complexity of an efficient implementation of $\mathbf{C}(\cdot)$.

Theorem 2 (Product Theorem). *Let $\mathbf{C}(\cdot)$ be as above, and let $q > 0$ be such that $\mathbf{C}(\cdot)$ makes q_i queries to its i -th subsystem when invoked q times. Then, for all $t, \gamma > 0$, if $\Delta_{t'_i, q'_i}(\mathbf{F}_i, \mathbf{I}_i) \leq \frac{1}{2}$ for all $i = 1, \dots, m - 1$,*

$$\Delta_{t,q}(\mathbf{C}(\mathbf{F}_1, \dots, \mathbf{F}_m), \mathbf{C}(\mathbf{I}_1, \dots, \mathbf{I}_m)) \leq 2^{m-1} \cdot \prod_{i=1}^m \Delta_{t'_i, q'_i}(\mathbf{F}_i, \mathbf{I}_i) + 2\gamma,$$

where $t'_i := l_{<i} + \varphi \cdot [t + t_C(q) + \mathcal{O}(\sum_{j=1}^{i-1} t_{\langle F_j, I_j \rangle_{B_j}}(q_j, l_j) + \sum_{j=i+1}^m t_{\langle F_j, I_j \rangle_{B_j}}(q_j))]$ and $q'_i := \varphi \cdot q_i$ for all $i = 1, \dots, m - 1$, whereas $t'_m := l_{<m} + t + t_C(q) + \mathcal{O}(\sum_{j=1}^{m-1} t_{\langle F_j, I_j \rangle_{B_j}}(q_j, l_j))$ and $q'_m := q_m$.

PROOF SKETCH. We present a proof sketch of the above theorem for the case $m = 2$. For simplicity, let $\mathbf{F}_1 = \mathbf{F}$, $\mathbf{F}_2 = \mathbf{G}$, $\mathbf{I}_1 = \mathbf{I}$, and $\mathbf{I}_2 = \mathbf{J}$. The core of the proof is a generic argument (i.e. it holds for all distinguishers, regardless of their computing power) reducing the task of upper bounding the distinguishing advantage for a neutralizing construction to the setting of the XOR-lemma.¹⁸ It is easy to verify that (also cf. [24])

$$\begin{aligned} \Delta^{\mathbf{D}}(\mathbf{C}(\mathbf{F}, \mathbf{G}), \mathbf{C}(\mathbf{I}, \mathbf{J})) &= 2 \cdot \Delta^{\mathbf{D}}(\langle \mathbf{C}(\mathbf{F}, \mathbf{G}), \mathbf{C}(\mathbf{I}, \mathbf{J}) \rangle_B, \mathbf{C}(\mathbf{I}, \mathbf{J})) \\ &= |\Gamma^{\mathbf{D}}(\langle \langle \mathbf{C}(\mathbf{F}, \mathbf{G}), \mathbf{C}(\mathbf{I}, \mathbf{J}) \rangle_B, \mathbf{C}(\mathbf{I}, \mathbf{J}) \rangle_{B'}, B')|, \end{aligned}$$

where B and B' are independent uniformly distributed random bits. Note that conditioned on $B' = 0$, the system $\langle \langle \mathbf{C}(\mathbf{F}, \mathbf{G}), \mathbf{C}(\mathbf{I}, \mathbf{J}) \rangle_B, \mathbf{C}(\mathbf{I}, \mathbf{J}) \rangle_{B'}$ behaves as $\mathbf{C}(\mathbf{F}, \mathbf{G})$ with probability $\frac{1}{2}$, and as $\mathbf{C}(\mathbf{I}, \mathbf{J})$ otherwise. On the other hand, conditioned on $B' = 1$ it always behaves as $\mathbf{C}(\mathbf{I}, \mathbf{J})$. In particular, this implies that (for independent uniform random bits B_1, B_2)

$$(\langle \langle \mathbf{C}(\mathbf{F}, \mathbf{G}), \mathbf{C}(\mathbf{I}, \mathbf{J}) \rangle_B, \mathbf{C}(\mathbf{I}, \mathbf{J}) \rangle_{B'}, B') \equiv (\mathbf{C}(\langle \mathbf{F}, \mathbf{I} \rangle_{B_1}, \langle \mathbf{G}, \mathbf{J} \rangle_{B_2}), B_1 \oplus B_2),$$

because of the neutralizing property. We thus obtain

$$\Gamma^{\mathbf{D}}(\langle \langle \mathbf{C}(\mathbf{F}, \mathbf{G}), \mathbf{C}(\mathbf{I}, \mathbf{J}) \rangle_B, \mathbf{C}(\mathbf{I}, \mathbf{J}) \rangle_{B'}, B') = \Gamma^{\mathbf{D}}(\mathbf{C}(\langle \mathbf{F}, \mathbf{I} \rangle_{B_1}, \langle \mathbf{G}, \mathbf{J} \rangle_{B_2}), B_1 \oplus B_2)$$

and we conclude the proof by “absorbing” the computation of $\mathbf{C}(\cdot)$ into \mathbf{D} , clearly without modifying the advantage. Using the XOR-lemma (Theorem 1) for $m = 2$ we obtain

$$\begin{aligned} \Delta_{t,q}(\mathbf{C}(\mathbf{F}, \mathbf{G}), \mathbf{C}(\mathbf{I}, \mathbf{J})) &\leq 2 \cdot \Gamma_{t+t_C(q), q_1, q_2}(\langle \mathbf{F}, \mathbf{I} \rangle_{B_1}, \langle \mathbf{G}, \mathbf{J} \rangle_{B_2}, B_1 \oplus B_2) \\ &\leq 2 \cdot \Gamma_{t'_1, q'_1}(\langle \mathbf{F}, \mathbf{I} \rangle_{B_1}, B_1) \cdot \Gamma_{t'_2, q'_2}(\langle \mathbf{G}, \mathbf{J} \rangle_{B_2}, B_2) + 2\gamma. \end{aligned}$$

for appropriate t'_1, q'_1 and t'_2, q'_2 . Extending the proof to *arbitrary* neutralizing constructions for $m > 2$ requires some extra care. The details can be found in the full version of this paper.

3.3 Applications of Theorem 2

SUMS OF PRFS. Let $\mathbf{F}_1, \dots, \mathbf{F}_m : \mathcal{X} \rightarrow \mathcal{Y}$ be cc-stateless random functions (in fact, \mathbf{F}_m can possibly be stateful), and let \star be a quasi-group operation on \mathcal{Y} . The operator \star is neutralizing, as discussed in Section 2, for $\mathbf{F}_1, \dots, \mathbf{F}_m$ and ideal systems $\mathbf{I}_1 = \dots = \mathbf{I}_m = \mathbf{R}$, where $\mathbf{R} : \mathcal{X} \rightarrow \mathcal{Y}$ is a URF. In order to simplify the time complexity statements, we assume that there exist efficient algorithms

¹⁸ A similar argument was implicitly used in the information-theoretic product theorem of [24].

implementing $\mathbf{F}_i(\cdot)$ such that $\mathbf{F}_i(s, x)$ is computed in time t_{F_i} given s and x (this holds in the interesting case where we apply the result to PRFs) and elements of \mathcal{Y} can be encoded using $\ell \approx \log |\mathcal{Y}|$ bits. Note that the canonical implementation of \mathbf{R} keeps a linearly-growing state of size $s = \mathcal{O}(q \cdot \ell)$ after q queries, and answers each query in time $\mathcal{O}(\log(s))$. Therefore, with $t_{\langle F_i, R \rangle_{B_i}}(q, s) = \mathcal{O}(q \cdot \max\{t_{F_i}, \log(s + q\ell)\})$ and $l_{<i} = \mathcal{O}((i - 1)\varphi q\ell)$, we apply Theorem 2 to obtain the following result (we tacitly assume that all advantages are bounded by $\frac{1}{2}$):

Corollary 1. *For all $t, q, \gamma > 0$,*

$$\Delta_{t,q}(\mathbf{F}_1 \star \cdots \star \mathbf{F}_m, \mathbf{R}) \leq 2^{m-1} \cdot \prod_{i=1}^m \Delta_{t'_i, q'_i}(\mathbf{F}_i, \mathbf{R}) + 2\gamma.$$

A version of this result with weaker bounds was shown by Dodis et al. [6] for $\star = \oplus$. (Their bounds depend in particular on the number of queries.) We remark that the analogous result for PRGs follows as a special case, since a PRG can be seen as a one-input PRF.

In the asymptotic setting, if $\mathbf{F}(\cdot)$ is a δ -PRF (for some $\delta < \frac{1}{2}$), it follows that $\mathbf{F}(S_1) \star \cdots \star \mathbf{F}(S_m)$, for independent keys S_1, \dots, S_m , is a $2^{m-1} \cdot \delta^m$ -PRF: For t, q polynomial (in n), we have $\Delta_{t,q}(\mathbf{F}(S_1) \star \cdots \star \mathbf{F}(S_m), \mathbf{R}) \leq 2^{m-1} \cdot \delta^m + \text{negl} + 1/p(n)$ for all polynomials p , as both t'_i and q'_i are polynomial as well.

CASCADE OF PRPs. Let $\mathbf{P} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a URP and let $\mathbf{Q}_1, \dots, \mathbf{Q}_m : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be cc-stateless random permutations. Recall that the \triangleright operator is neutralizing for $\mathbf{Q}_1, \dots, \mathbf{Q}_m$ (all with ideal system \mathbf{P}), as well as for $\langle \mathbf{Q}_1 \rangle, \dots, \langle \mathbf{Q}_m \rangle$ (all with ideal system $\langle \mathbf{P} \rangle$). As above, we assume that both $\mathbf{Q}_i(s, x)$ and $\mathbf{Q}_i^{-1}(s, y)$ are computable in time t_{Q_i} . Furthermore, simulating the URP \mathbf{P} (as well as the two-sided URP $\langle \mathbf{P} \rangle$) requires the same complexity as implementing a URF. Therefore, with $t_{\langle Q_i, P \rangle_{B_i}}(q, s) = t_{\langle \langle Q_i \rangle, \langle P \rangle \rangle_{B_i}}(q, s) = \mathcal{O}(q \cdot \max\{t_{Q_i}, \log(s + qn)\})$ and $l_{<i} = \mathcal{O}((i - 1)\varphi qn)$, Theorem 2 yields the following corollary:

Corollary 2. *For all $t, q, \gamma > 0$,*

$$\Delta_{t,q}(\mathbf{Q}_1 \triangleright \cdots \triangleright \mathbf{Q}_m, \mathbf{P}) \leq 2^{m-1} \cdot \prod_{i=1}^m \Delta_{t'_i, q'_i}(\mathbf{Q}_i, \mathbf{P}) + 2\gamma,$$

and

$$\Delta_{t,q}(\langle \mathbf{Q}_1 \rangle \triangleright \cdots \triangleright \langle \mathbf{Q}_m \rangle, \langle \mathbf{P} \rangle) \leq 2^{m-1} \cdot \prod_{i=1}^m \Delta_{t'_i, q'_i}(\langle \mathbf{Q}_i \rangle, \langle \mathbf{P} \rangle) + 2\gamma.$$

We remark that this is the first result considering *two-sided* PRPs, and even in the one-sided setting only the case $m = 2$ was considered by Luby and Rackoff [21], and subsequently extended to $m = \mathcal{O}(\log n)$ by Myers [26].

Furthermore, we note that \mathbf{Q}_1 is allowed to be stateful in the one-sided case, as Theorem 2 allows one system to be stateful: In fact, \triangleright is not necessarily neutralizing whenever at least two permutations are stateful.

4 A Strong Product Theorem for Randomized Neutralizing Constructions

4.1 A Product Theorem from Self-independence

Since Theorem 2 holds for arbitrary neutralizing constructions, one cannot avoid the factor 2^{m-1} in the bound. This section shows that a subclass of neutralizing constructions satisfying a simple information-theoretic property yield a *strong* product theorem, i.e., the obtained upper bound is roughly the product of the individual advantages.

SELF-INDEPENDENCE. The notion of self-independence of an ideal system \mathbf{I} under a construction $\mathbf{C}(\cdot)$ captures the fact that a computationally unbounded distinguisher cannot tell apart the scenario where the *same* instance of \mathbf{I} is accessed through independent instances of $\mathbf{C}(\cdot)$ from the setting where each instance of $\mathbf{C}(\cdot)$ accesses an independent instance of \mathbf{I} .

Definition 2. *The system \mathbf{I} is η -self-independent under $\mathbf{C}(\cdot)$ for a function $\eta : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$, if for all $q, \lambda > 0$, the best (information-theoretic) distinguishing advantage when allowing q queries to each subsystem satisfies*

$$\Delta_{q, \dots, q}(\mathbf{C}_1(\mathbf{I}) \| \dots \| \mathbf{C}_\lambda(\mathbf{I}), \mathbf{C}_1(\mathbf{I}_1) \| \dots \| \mathbf{C}_\lambda(\mathbf{I}_\lambda)) \leq \eta(q, \lambda),$$

where $\mathbf{C}_1(\cdot), \dots, \mathbf{C}_\lambda(\cdot)$ and $\mathbf{I}_1, \dots, \mathbf{I}_\lambda$ are independent copies of $\mathbf{C}(\cdot)$ and \mathbf{I} , respectively.

As an example, consider the construction $\mathbf{C}(\cdot)$ which generates a (secret) random n -bit offset Z , and given access to a random function $\mathbf{F} : \{0, 1\}^n \rightarrow \{0, 1\}^n$, $\mathbf{C}(\mathbf{F})$ returns $\mathbf{F}(x \oplus Z)$ upon each query x . It is not hard to show, e.g. using the tools from [22], that a URF $\mathbf{R} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is η -self-independent under $\mathbf{C}(\cdot)$ for $\eta(q, \lambda) \leq \frac{q^2 \lambda^2}{2} \cdot 2^{-n}$, i.e., the probability that for some distinct $i \neq j$ the instances $\mathbf{C}_i(\cdot)$ and $\mathbf{C}_j(\cdot)$ invoke \mathbf{R} with the same input.

RESTRICTED ATTACKS ON CRYPTOGRAPHIC FUNCTIONS. Indistinguishability-based security definitions can also be weakened by restricting the distinguisher's access to the given system. For instance, the standard PRF notion considering an (adaptive) *chosen-input attack* can be weakened to non-adaptive chosen-input attacks or even (*known*) *random-input attacks*. (Keyed functions which are secure under the latter notion are usually called *weak PRFs* [28] in the literature.¹⁹) This is conveniently modeled by letting the distinguisher access either of $\mathbf{E}(\mathbf{F})$ and $\mathbf{E}(\mathbf{G})$, where the construction $\mathbf{E}(\cdot)$ enforces a particular type of access, and \mathbf{F} and \mathbf{G} are the systems to be distinguished. For a chosen-input attack, \mathbf{E} would just give full access to the underlying system (i.e. $\mathbf{E}(\cdot)$ is the *identity*), and the following are two additional examples:

¹⁹ The name is slightly misleading within the context of this paper, as it can be used [27] to describe an ϵ -PRF for a non-negligible $\epsilon < 1$.

- *Random-input attacks* against an $(\mathcal{X}, \mathcal{Y})$ -system are modeled by $\mathbf{K}(\cdot)$ that, upon each invocation (with some dummy input), generates a fresh uniformly-chosen element $r \in \mathcal{X}$, makes a query with input r to the given subsystem, obtaining $y \in \mathcal{Y}$, and returns (r, y) .
- For a quasi-group operation $*$ on \mathcal{X} (usually \oplus), a *random-offset attack* is modeled by a construction $\mathbf{Z}(\cdot)$ which initially generates a random offset $Z \in \mathcal{X}$, and upon each invocation with input $x \in \mathcal{X}$, makes a query to the given subsystem with input $x \star Z$, and outputs the returned value y . (To our knowledge, this notion was not previously considered in the literature.)

A feature of the product theorem of this section is that it is easily applicable also to the restricted-access case.

THE PRODUCT THEOREM. In the following, let $\mathbf{C}(\cdot)$ be a neutralizing construction for systems $\mathbf{F}_1, \dots, \mathbf{F}_m$ and ideal system $\mathbf{I}_1, \dots, \mathbf{I}_m$, all of which (with the possible exception of \mathbf{F}_m and \mathbf{I}_m) are cc-stateless. Furthermore, we assume that $\mathbf{F}_i(\cdot)$ is efficiently implementable for all $i = 1, \dots, m - 1$,²⁰ and the corresponding (short) random variable S_i is drawn from the set \mathcal{S}_i . Also, we let $\mathbf{E}(\cdot)$ be construction restricting access to \mathbf{F}_i and \mathbf{I}_i . Finally, for $i = 1, \dots, m$, and for $s_1 \in \mathcal{S}_1, \dots, s_{i-1} \in \mathcal{S}_{i-1}$ we define

$$\mathbf{C}_{s_1, \dots, s_{i-1}}^{(i)}(\cdot) := \mathbf{C}(\mathbf{F}_1(s_1), \dots, \mathbf{F}_{i-1}(s_{i-1}), \cdot, \mathbf{F}_{i+1}, \dots, \mathbf{F}_m)$$

and consider the following two properties:

- (i) For all $i = 1, \dots, m - 1$ (the property is not necessary for $i = m$) and all $s_1 \in \mathcal{S}_1, \dots, s_{i-1} \in \mathcal{S}_{i-1}$, the ideal system \mathbf{I}_i is η -self-independent under the construction $\mathbf{C}_{s_1, \dots, s_{i-1}}^{(i)}(\cdot)$ for some small function η .
- (ii) For all $i = 1, \dots, m$ and $s_1 \in \mathcal{S}_1, \dots, s_{i-1} \in \mathcal{S}_{i-1}$, there exists a construction $\mathbf{T}_{s_1, \dots, s_{i-1}}^{(i)}(\cdot)$ with the property that for independent instances $\mathbf{T}_1(\cdot), \dots, \mathbf{T}_\lambda(\cdot)$ and $\mathbf{C}_1(\cdot), \dots, \mathbf{C}_\lambda(\cdot)$ of $\mathbf{T}_{s_1, \dots, s_{i-1}}^{(i)}(\cdot)$ and $\mathbf{C}_{s_1, \dots, s_{i-1}}^{(i)}(\cdot)$, respectively, and all compatible systems \mathbf{S} ,

$$\mathbf{T}_1(\mathbf{E}(\mathbf{S})) \parallel \dots \parallel \mathbf{T}_\lambda(\mathbf{E}(\mathbf{S})) \equiv \mathbf{C}_1(\mathbf{S}) \parallel \dots \parallel \mathbf{C}_\lambda(\mathbf{S}).$$

We define t_{T_i} as the maximal complexity (taken over all s_1, \dots, s_{i-1}) for implementing the construction $\mathbf{T}_{s_1, \dots, s_{i-1}}^{(i)}(\cdot)$.

In the following, we define $\lambda := \left(\frac{4m}{\gamma}\right)^2 \cdot \ln\left(\frac{4m}{\gamma}\right)$, for understood m and γ .

Theorem 3 (Strong Product Theorem). *Let $q > 0$ and $\mathbf{C}(\cdot)$ be as above satisfying conditions (i) and (ii), and assume that upon q queries, $\mathbf{C}(\cdot)$ makes at most q_i queries to the i -th subsystem. Then, for all $t, \gamma > 0$,*

²⁰ While the same techniques as in the proof of Theorem 1 could be used to address general cc-stateless systems where $\mathbf{F}(\cdot)$ is not necessarily efficient, this will not be necessary for our applications.

$$\Delta_{t,q}(\mathbf{C}(\mathbf{F}_1, \dots, \mathbf{F}_m), \mathbf{C}(\mathbf{I}_1, \dots, \mathbf{I}_m)) \leq \prod_{i=1}^m \Delta_{t'_i, q'_i}(\mathbf{E}(\mathbf{F}_i), \mathbf{E}(\mathbf{I}_i)) + \sum_{i=1}^{m-1} \eta(q_i, \lambda) + \gamma,$$

where $t'_i := \lambda \cdot (t + \mathcal{O}(t_{T_i}(q_i)))$ and $q'_i := \lambda \cdot q_i$ for all $i = 1, \dots, m - 1$, whereas $t'_m := t + \mathcal{O}(t_{T_m}(q))$ and $q'_m := q_m$.

The proof of Theorem 3 is deferred to the full version of this paper. It abstracts and generalizes the proof technique used by Myers [27] (which was in turn based on Levin’s proof of the XOR-lemma [20,10]).

4.2 Applications of the Strong Product Theorem

We present a number of new results which follow as simple applications of Theorem 3. Let $\mathbf{Q}_1, \dots, \mathbf{Q}_m : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be cc-stateless random permutations, and let $\mathbf{F}_1, \dots, \mathbf{F}_m : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ be cc-stateless random functions. Furthermore, let $\mathbf{P} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ and $\mathbf{R} : \{0, 1\}^n \rightarrow \{0, 1\}^\ell$ be a URF and URP, respectively. Assume that $\mathbf{Q}_i(s, x)$ (and $\mathbf{Q}_i^{-1}(s, y)$) and $\mathbf{F}_i(s, x)$ can be computed in time t_{Q_i} and t_{F_i} , respectively, for all s, x , and y .

RANDOMIZED CASCADE OF PRPs. The perhaps most surprising application is a *strong* product theorem for (two-sided) PRPs. We modify the (two-sided) cascade $\langle \mathbf{Q}_1 \rangle \triangleright \dots \triangleright \langle \mathbf{Q}_m \rangle$ by choosing two independent random offsets that are added to the inputs and the outputs, i.e., we consider $\langle \oplus_{Z_1} \rangle \triangleright \langle \mathbf{Q}_1 \rangle \triangleright \dots \triangleright \langle \mathbf{Q}_m \rangle \triangleright \langle \oplus_{Z_2} \rangle$ for two independent uniform n -bit strings Z_1, Z_2 , where for some $z \in \{0, 1\}^n$ the system $\langle \oplus_z \rangle$ is the bi-directional mapping which answers a forward query $(x, +)$ with $x \oplus z$ and a backward query $(y, -)$ with $y \oplus z$. The computational overhead is minimal compared to the regular cascade, and requires only additional storage for two n -bit strings (which are to be seen as part of the secret key).

Clearly the neutralizing property of the original cascade is preserved. Furthermore, using techniques from [22], we show in the full version that the construction satisfies condition (i) above with $\eta(q, \lambda) \leq q^2 \lambda^2 2^{-n}$. Therefore, Theorem 3 (with $\mathbf{E}(\cdot)$ being the identity) yields the following result.

Corollary 3. *For all $t, q, \gamma > 0$, and independent uniform n -bit strings Z_1, Z_2 ,*

$$\Delta_{t,q}(\langle \oplus_{Z_1} \rangle \triangleright \langle \mathbf{Q}_1 \rangle \triangleright \dots \triangleright \langle \mathbf{Q}_m \rangle \triangleright \langle \oplus_{Z_2} \rangle, \langle \mathbf{P} \rangle) \leq \prod_{i=1}^m \Delta_{t'_i, q'_i}(\langle \mathbf{Q}_i \rangle, \langle \mathbf{P} \rangle) + \frac{mq^2 \lambda^2}{2^n} + \gamma,$$

where $t'_i := \lambda \cdot (t + \mathcal{O}(q \cdot \sum_{j \neq i} t_{Q_j}))$ and $q'_i := \lambda \cdot q$ for all $i = 1, \dots, m - 1$, whereas $t'_m := t + \mathcal{O}(q \cdot \sum_{j=1}^{m-1} t_{Q_j})$ and $q'_m := q$.

The result can be used to obtain a δ^m -two-sided PRP from any δ -two-sided PRP. (Note that the η -dependent term is negligible for polynomial t, q and any γ which is the inverse of a polynomial.) It can be shown that the second random offset Z_2 is superfluous in the one-sided case.

SUM OF RANDOM-INPUT PRFS. The construction $\mathbf{K}(\mathbf{F}_1 \oplus \cdots \oplus \mathbf{F}_m)$ (i.e. the XOR of the functions accessed in a random-input attack) is clearly neutralizing (the ideal system being $\mathbf{K}(\mathbf{R})$). In the full version, we show that it also satisfies condition (i) with $\eta(q, \lambda) \leq \frac{q^2 \lambda^2}{2} \cdot 2^{-n}$. Moreover, for all i and keys $s_1 \in \mathcal{S}_1, \dots, s_{i-1} \in \mathcal{S}_{i-1}$, the appropriate construction $\mathbf{T}_{s_1, \dots, s_{i-1}}^{(i)}(\cdot)$ generates random keys S_{i+1}, \dots, S_m and whenever invoked, it issues a query to $\mathbf{K}(\mathbf{S})$, obtaining (r, y) , and outputs the pair

$$\left(r, \bigoplus_{j=1}^{i-1} \mathbf{F}_j(s_j, r) \oplus y \oplus \bigoplus_{j=i+1}^m \mathbf{F}_j(S_j, r) \right).$$

It is easy to see that these constructions satisfy property (ii), since $\mathbf{K}(\cdot)$ evaluates the given function at a fresh random input upon each invocation. Theorem 3 yields the following result.

Corollary 4. *For all $t, q, \gamma > 0$,*

$$\Delta_{t,q}(\mathbf{K}(\mathbf{F}_1 \oplus \cdots \oplus \mathbf{F}_m), \mathbf{K}(\mathbf{R})) \leq \prod_{i=1}^m \Delta_{t'_i, q'_i}(\mathbf{K}(\mathbf{F}_i), \mathbf{K}(\mathbf{R})) + \frac{(m-1)q^2 \lambda^2}{2^{n+1}} + \gamma,$$

where $t'_i := \lambda(t + \mathcal{O}(q \cdot \sum_{j \neq i} t_{F_j}))$ and $q'_i := \lambda \cdot q$ for all $i = 1, \dots, m-1$, whereas $t'_m := t + \mathcal{O}(q \cdot \sum_{j=1}^{m-1} t_{F_j})$ and $q'_m := q$.

The result holds for any other quasi-group operation. It is remarkable that XOR satisfies much stronger indistinguishability amplification properties under random-input attacks than under chosen-input attacks. This is particularly interesting, as a wide number of applications, such as secure symmetric message encryption, can efficiently be based on this weaker PRF notion (cf. [5,25]).

RANDOMIZED XOR OF PRFS. The first product theorem for PRFs, due to Myers [27], considered the neutralizing composition $\mathbf{Z}_1(\mathbf{F}_1) \oplus \cdots \oplus \mathbf{Z}_m(\mathbf{F}_m)$ for independent instances of $\mathbf{Z}(\cdot)$. This result is directly implied by Theorem 3, which in fact also implies the same result for the construction $\mathbf{Z}(\mathbf{F}_1 \oplus \cdots \oplus \mathbf{F}_m)$ using the *same* offset for all invocations: As we show in the full version, both compositions satisfy property (i) with $\eta(q, \lambda) \leq \frac{q^2 \lambda^2}{2} 2^{-n}$.

However, a major advantage of Myers' original construction (which was unobserved so far) is that independent instances of the construction can be simulated even when only given access to $\mathbf{Z}(\mathbf{S})$ (with $\mathbf{S} \in \{\mathbf{F}_i, \mathbf{R}\}$). The corresponding construction $\mathbf{T}_{s_1, \dots, s_{i-1}}^{(i)}(\cdot)$ chooses independent instances $\mathbf{F}_{i+1}, \dots, \mathbf{F}_m, \mathbf{Z}_1(\cdot), \dots, \mathbf{Z}_{i-1}(\cdot), \mathbf{Z}_{i+1}(\cdot), \dots, \mathbf{Z}_m(\cdot)$, and a random n -bit string Z , and on input x queries $x \oplus Z$ to $\mathbf{Z}(\mathbf{S})$, obtaining $y \in \{0, 1\}^\ell$, and outputs

$$y \oplus \bigoplus_{j=1}^{i-1} \mathbf{Z}_j(\mathbf{F}_j(s_j))(x) \oplus \bigoplus_{j=i+1}^m \mathbf{Z}_j(\mathbf{F}_j)(x),$$

where $\mathbf{Z}_j(\mathbf{F}_j)(x)$ is the result of invoking the system $\mathbf{Z}_j(\mathbf{F}_j)$ on input x .

Once again, condition (ii) is easily verified by the fact that access through $\mathbf{Z}(\cdot)$ can be re-randomized by simply adding a fresh random offset to all inputs. Thus, Theorem 3 yields the following strengthened version of the main result of [27].

Corollary 5. *For all $t, q, \gamma > 0$, and for independent instances $\mathbf{Z}_1(\cdot), \dots, \mathbf{Z}_m(\cdot)$ of $\mathbf{Z}(\cdot)$,*

$$\Delta_{t,q}(\mathbf{Z}_1(\mathbf{F}_1) \oplus \dots \oplus \mathbf{Z}_m(\mathbf{F}_m), \mathbf{R}) \leq \prod_{i=1}^m \Delta_{t'_i, q'_i}(\mathbf{Z}(\mathbf{F}_i), \mathbf{Z}(\mathbf{R})) + \frac{(m-1)q^2\lambda^2}{2^{n+1}} + \gamma,$$

where $t'_i := \lambda(t + \mathcal{O}(q \cdot \sum_{j \neq i} t_{F_j}))$ and $q'_i := \lambda \cdot q$ for all $i = 1, \dots, m-1$, whereas $t'_m := t + \mathcal{O}(q \cdot \sum_{j=1}^{m-1} t_{F_j})$ and $q'_m := q$.

The best advantage under $\mathbf{Z}(\cdot)$ can be significantly smaller than under direct access: Consider e.g. a good PRF with the additional property of outputting the zero string when evaluated at some fixed known input, regardless of the key.

References

1. Bellare, M., Impagliazzo, R., Naor, M.: Does parallel repetition lower the error in computationally sound protocols? In: FOCS 1997, pp. 374–383 (1997)
2. Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 409–426. Springer, Heidelberg (2006)
3. Canetti, R., Halevi, S., Steiner, M.: Hardness amplification of weakly verifiable puzzles. In: Kilian, J. (ed.) TCC 2005. LNCS, vol. 3378, pp. 17–33. Springer, Heidelberg (2005)
4. Canetti, R., Rivest, R.L., Sudan, M., Trevisan, L., Vadhan, S.P., Wee, H.: Amplifying collision resistance: A complexity-theoretic treatment. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 264–283. Springer, Heidelberg (2007)
5. Damgård, I.B., Nielsen, J.B.: Expanding pseudorandom functions; or: From known-plaintext security to chosen-plaintext security. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 449–464. Springer, Heidelberg (2002)
6. Dodis, Y., Impagliazzo, R., Jaiswal, R., Kabanets, V.: Security amplification for interactive cryptographic primitives. In: TCC 2009. LNCS, vol. 5444, pp. 128–145 (2009)
7. Dwork, C., Naor, M., Reingold, O.: Immunizing encryption schemes from decryption errors. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 342–360. Springer, Heidelberg (2004)
8. Even, S., Goldreich, O.: On the power of cascade ciphers. ACM Trans. Comput. Syst. 3(2), 108–116 (1985)
9. Goldreich, O., Impagliazzo, R., Levin, L.A., Venkatesan, R., Zuckerman, D.: Security preserving amplification of hardness. In: FOCS 1990, pp. 318–326 (1990)
10. Goldreich, O., Nisan, N., Wigderson, A.: On Yao’s XOR-lemma. Electronic Colloquium on Computational Complexity (ECCC) 2(50) (1995)
11. Haitner, I., Harnik, D., Reingold, O.: On the power of the randomized iterate. In: Dwork, C. (ed.) CRYPTO 2006. LNCS, vol. 4117, pp. 22–40. Springer, Heidelberg (2006)
12. Halevi, S., Rabin, T.: Degradation and amplification of computational hardness. In: Canetti, R. (ed.) TCC 2008. LNCS, vol. 4948, pp. 626–643. Springer, Heidelberg (2008)
13. 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)
14. Herzberg, A.: On tolerant cryptographic constructions. In: Menezes, A. (ed.) CT-RSA 2005. LNCS, vol. 3376, pp. 172–190. Springer, Heidelberg (2005)

15. Holenstein, T.: Key agreement from weak bit agreement. In: STOC 2005, pp. 664–673 (2005)
16. Holenstein, T., Renner, R.: One-way secret-key agreement and applications to circuit polarization and immunization of public-key encryption. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 478–493. Springer, Heidelberg (2005)
17. Hopper, N., Molnar, D., Wagner, D.: From weak to strong watermarking. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 362–382. Springer, Heidelberg (2007)
18. Impagliazzo, R.: Hard-core distributions for somewhat hard problems. In: FOCS 1995, pp. 538–545 (1995)
19. Impagliazzo, R., Jaiswal, R., Kabanets, V.: Chernoff-type direct product theorems. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 500–516. Springer, Heidelberg (2007)
20. Levin, L.A.: One way functions and pseudorandom generators. *Combinatorica* 7(4), 357–363 (1987)
21. Luby, M., Rackoff, C.: Pseudo-random permutation generators and cryptographic composition. In: STOC 1986, pp. 356–363 (1986)
22. Maurer, U.: Indistinguishability of random systems. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 110–132. Springer, Heidelberg (2002)
23. Maurer, U., Massey, J.L.: Cascade ciphers: The importance of being first. *Journal of Cryptology* 6(1), 55–61 (1993)
24. Maurer, U., Pietrzak, K., Renner, R.: Indistinguishability amplification. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 130–149. Springer, Heidelberg (2007)
25. Maurer, U., Sjödin, J.: A fast and key-efficient reduction of chosen-ciphertext to known-plaintext security. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 498–516. Springer, Heidelberg (2007)
26. 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)
27. Myers, S.: Efficient amplification of the security of weak pseudo-random function generators. *Journal of Cryptology* 16, 1–24 (2003)
28. Naor, M., Reingold, O.: Synthesizers and their application to the parallel construction of pseudo-random functions. *Journal of Computer and System Sciences* 58(2), 336–375 (1999)
29. Pass, R., Venkatasubramanian, M.: An efficient parallel repetition theorem for Arthur-Merlin games. In: STOC 2007, pp. 420–429 (2007)
30. 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)
31. Shaltiel, R., Viola, E.: Hardness amplification proofs require majority. In: STOC 2008, pp. 589–598 (2008)
32. Vaudenay, S.: Provable security for block ciphers by decorrelation. In: Meinel, C., Morvan, M. (eds.) STACS 1998. LNCS, vol. 1373, pp. 249–275. Springer, Heidelberg (1998)
33. Vaudenay, S.: Adaptive-attack norm for decorrelation and super-pseudorandomness. In: Heys, H.M., Adams, C.M. (eds.) SAC 1999. LNCS, vol. 1758, pp. 49–61. Springer, Heidelberg (2000)
34. Wullschlegel, J.: Oblivious-transfer amplification. In: Naor, M. (ed.) EUROCRYPT 2007. LNCS, vol. 4515, pp. 555–572. Springer, Heidelberg (2007)
35. Yao, A.C.: Theory and applications of trapdoor functions. In: FOCS 1982, pp. 80–91 (1982)