Efficient and Optimally Secure Key-Length Extension for Block Ciphers via Randomized Cascading

Peter Gaži[1,2] and Stefano Tessaro[3,4]

[1] Department of Computer Science, Comenius University, Bratislava, Slovakia
[2] Department of Computer Science, ETH Zurich, Switzerland
peter.gazi@inf.ethz.ch
[3] Department of Computer Science and Engineering
University of California San Diego, La Jolla CA, USA
[4] MIT, Cambridge MA, USA
stessaro@cs.ucsd.edu

Abstract. We consider the question of efficiently extending the key length of block ciphers. To date, the approach providing highest security is triple encryption (used e.g. in Triple-DES), which was proved to have roughly $\kappa + \min\{n/2, \kappa/2\}$ bits of security when instantiated with ideal block ciphers with key length κ and block length n, at the cost of three block-cipher calls per message block.

This paper presents a new practical key-length extension scheme exhibiting $\kappa + n/2$ bits of security – hence improving upon the security of triple encryption – solely at the cost of *two* block cipher calls and a key of length $\kappa + n$. We also provide matching generic attacks showing the optimality of the security level achieved by our approach with respect to a general class of two-query constructions.

Keywords: Block ciphers, Cascade encryption, Provable security.

1 Introduction

1.1 Key-Length Extension for Block Ciphers

Several practical block cipher designs have been proposed over the last decades and have been the object of extensive cryptanalytic efforts. Examples include DES [1], IDEA [19], BLOWFISH [28], and the currently in-use AES [4]. Within applications, we typically demand that these block ciphers are a good *pseudorandom permutation* (PRP), i.e., in the eyes of a computationally bounded attacker, they behave as a randomly chosen permutation under a random secret key. For instance, PRP security of the underlying block cipher is necessary to infer security of all modes of operations for message encryption (such as counter-mode and CBC encryption [8]) as well as of message authentication codes like CBC-MAC [9] and PMAC [12].

In practice, we define the PRP security level of a block cipher as the complexity required to distinguish it from a random permutation with non-negligible

D. Pointcheval and T. Johansson (Eds.): EUROCRYPT 2012, LNCS 7237, pp. 63–80, 2012.

advantage. The *key length* κ of a block cipher crucially limits the achievable security level, since the secret key K can be recovered given black-box access to $E(K,\cdot)$ evaluating $E(\cdot,\cdot)$ approximately 2^κ times; obviously, this also yields a PRP distinguishing adversary with equal complexity. Such weakness is *generic*, in the sense that it only depends on κ, and even an *ideal block cipher* suffers from the same attack.[1] In contrast, no real dependency exists between security and the block length n of a block cipher: No generic attack faster than 2^κ exists even if $n = 1$. In the following, let us refer to a block cipher with key and block lengths κ and n, respectively, as a (κ, n)-block cipher.

Key length extension. With a continuous increase of the availability of computing resources, the role of the key length has hence never been more important. Key lengths of say fewer than 64 bits are no longer sufficient to ensure security, making key recovery a matter of a few hours even on modest architectures. This is a serious problem for legacy designs such as DES which have very short keys of length 56 bits, but which otherwise do not seem to present significant non-generic security weaknesses. Constructions based on DES also remain very attractive because of its short block length $n = 64$ which allows enciphering short inputs. This is for example crucial in current applications in the financial industry, such as the EMV standard [6], where the block cipher is applied to PIN numbers, which are very short.

The above described situation motivates the problem of *key-length extension*, which is the main object of this paper: We seek for very efficient constructions provably transforming any (κ, n)-block cipher E into a (κ', n)-block cipher E' with both $\kappa' > \kappa$ and higher PRP security, i.e., the PRP security of E' should be higher than 2^κ whenever E does not exhibit any non-generic weaknesses. We aim both at providing very efficient approaches to key length extension and at understanding the optimal security achievable by such constructions. Our main contribution will be a new and very efficient two-call key-length extension method outperforming the efficiency of existing solutions by a large margin, and achieving security levels which we prove optimal, and which are comparable (and even better) than those of earlier, less efficient, designs.

Ideal cipher model. In our proofs, we model the absence of generic weaknesses of the underlying block cipher by analyzing our constructions when instantiated with an ideal block cipher $\mathbf{E}$. In this model, complexity is measured in terms of the number of queries to $\mathbf{E}$ (so-called *ideal block cipher queries*) and to E' or the given random permutation (we refer to these as *construction* queries). It should be noted that proving security of key-length extension in the ideal cipher model implies absence of generic attacks, treating the underlying cipher as a black-box, and as we will explain in the next section, all attacks on existing schemes are indeed generic.

[1] As usual, an *ideal block cipher* $\mathbf{E} : \{0,1\}^\kappa \times \{0,1\}^n \times \{+,-\} \to \{0,1\}^n$ is the system associating with each key $k \in \{0,1\}^\kappa$ an independent randomly chosen permutation $E(k,\cdot)$ and allowing the adversary to learn $E(k,x)$ and $E^{-1}(k,y)$ for k, x, y of her choice.

1.2 Existing Approaches to Key-Length Extension

The short key length $\kappa = 56$ of DES has constituted the main motivation behind previous work on key-length extension. However, we stress that all previous constructions are generic, and can be applied to *any* block cipher with short keys, hence extending the applicability of these results (as well as the results of this paper) way beyond the specific case of DES.

A first proposal called DESX (due to Rivest) stretches the key length of DES employing a technique called *key whitening* (this approach was later used by Even and Mansour [15]): It is defined such that

$$\mathrm{DESX}_{k_i,k_o,k}(m) = k_o \oplus \mathrm{DES}_k(k_i \oplus m)$$

for all $m, k_i, k_o \in \{0,1\}^{64}$ and $k \in \{0,1\}^{56}$. DESX can be generalized to a generic transformation from a (κ, n)-block cipher to a $(\kappa + 2n, n)$-block cipher whose security was studied by Kilian and Rogaway [18]: They proved that any successful PRP distinguishing attack requires $2^{\frac{\kappa+n}{2}}$ queries.[2] They also observe that the same key may be used in both whitening steps (i.e., $k_i = k_o$) and provide an attack using $2^{\max\{\kappa,n\}}$ queries.

An alternative to whitening is *cascading* (or cascade encryption), i.e., sequentially composing ℓ block-cipher calls with usually different keys. (This is referred to as a cascade of length ℓ.) It is well known that a cascade of length two does not substantially increase security due to the meet-in-the-middle attack [13]. (Even though a security increase in terms of distinguishing advantage is achieved for low attack complexities [7].) The security properties of a cascade of different ciphers was studied by Even and Goldreich [14] showing that a cascade is at least as strong as the strongest of the ciphers used; and by Maurer and Massey [23] proving that it is at least as secure as the *first* cipher of the cascade, however in a more general attack model.

The meet-in-the-middle attack makes triple encryption the shortest cascade with a potential for significant security gain and indeed it has found widespread usage as *Triple-DES* (3DES) [2,3,5], where given keys $k_1, k_2, k_3 \in \{0,1\}^{56}$, a 64-bit message m is mapped to

$$3\mathrm{DES}_{k_1,k_2,k_3}(m) = \mathrm{DES}_{k_1}(\mathrm{DES}_{k_2}(\mathrm{DES}_{k_3}(m))) \ .$$

(A variant with shorter keys $3\mathrm{DES}'_{k_1,k_2}(m) = \mathrm{DES}_{k_1}(\mathrm{DES}^{-1}_{k_2}(\mathrm{DES}_{k_1}(m)))$ is also sometimes used.) For 3DES (and a variant of 3DES$'$ with independent keys), Bellare and Rogaway [11] and subsequently Gaži and Maurer [16] have shown security up to roughly $2^{\kappa+\min\{n,\kappa\}/2}$ queries when DES is replaced by an ideal block cipher. For the case of DES parameters, their result gives concretely security up to 2^{78} queries, whereas the best known attack due to Lucks [21] shows

[2] Their result is in fact more fine-grained, as they show that 2^ρ construction and $2^{\kappa+n-\rho}$ ideal block cipher queries, respectively, are necessary for all integers ρ; while different bounds for both query types are sometimes justified, we adopt a (by now more standard) worst-case approach only bounding the *sum* of both query numbers.

that no security better than 2^{90} can be expected. (It should also be noted that the proof of [16] extends to prove that longer cascades can achieve better security for short keys.)

We emphasize that despite the availability of modern designs with built-in larger keys (e.g., $\kappa \in \{128, 192, 256\}$ for AES), Triple-DES remains nowadays popular, not only because of backwards compatibility, but also because its short block size ($n = 64$ vs. $n \geq 128$ for AES) is well suited to applications enciphering short inputs such as personal-identification numbers (PINs). For example, it is the basis of the EMV standard for PIN-based authentication of debit and credit card transactions [6]. However, the use of three calls per processed message block is widely considered a drawback within applications which we address and solve in this paper.

Other related work. It is worth mentioning that several works have studied cascading-based security amplification of block ciphers only assumed to satisfy weaker forms of PRP security, both in the information-theoretic [32,24,25,17] as well as in the computational settings [20,26,31]. These results however consider an orthogonal model to ours and are hence incomparable.

1.3 Our Results

None of the above efficient constructions provably achieves security beyond $2^{\kappa+\min\{\kappa,n\}/2}$, and such security is achieved only at the fairly expensive cost of at least three block-cipher calls per message block. This paper aims at improving the efficiency-security trade-off in key-length extension. We ask the following question: Suppose that we only consider constructions making at most *two* calls to the underlying cipher. *What is the best security level we are expected to achieve?*

Better security and better efficiency. Quite surprisingly, our main result (presented in Section 4) exposes a "win-win" situation: We devise a *two*-call construction of a $(\kappa + n, n)$-block cipher from any (κ, n)-block cipher with security $2^{\kappa+n/2}$ in the ideal block cipher model, i.e., the obtained security is *higher* than that of existing three-call designs studied in [11,16].[3] Our construction – which we refer to as the *double XOR-cascade* (2XOR) – is obtained by careful insertion of two randomization steps (with the *same* key value) to a related-key version of double encryption. Concretely, we map each n-bit message m to

$$2\mathrm{XOR}_{k,z}(m) = E_{\widetilde{k}}(E_k(m \oplus z) \oplus z)$$

for all key values $k \in \{0,1\}^\kappa$ and $z \in \{0,1\}^n$, and where $\widetilde{k}$ is, for example, obtained from k by flipping one individual bit.

We note that the key length is comparable to the one of the two-key variant of 3DES (assuming $\kappa \approx n$). Intuitively, our construction requires some mild

[3] In fact, our construction tolerates arbitrarily many construction queries (i.e., up to 2^n) *and* $2^{\kappa+n/2}$ ideal block cipher queries. However, we stress that in all practically relevant cases $\kappa \geq n/2$, hence we tacitly assume this property throughout the paper.

Table 1. Required number of block-cipher queries, key lengths, security lower bounds and best known attacks for various key-length extension schemes. The bounds are parameterized by the key length of the underlying block cipher (denoted by κ) and its block size (denoted by n), and are for the usual case where $\kappa \geq n/2$.

	# of	key	log of the number of queries	
construction	queries	length	security lower bound	best known attack
(κ, n)-block cipher	1	κ	κ	κ
DESX [15,18]	1	$\kappa + n$	$(\kappa + n)/2$	$\max\{\kappa, n\}$
double encryption [13]	2	2κ	κ	κ
triple encryption [11,16,21]	3	3κ	$\kappa + \min\{\kappa, n\}/2$	90 (for 3DES)
double XOR-cascade [here]	2	$\kappa + n$	$\kappa + n/2$ (Thm. 3)	$\kappa + n/2$ (Thm. 2)

form of related-key security [10] which we obtain for free when the underlying block cipher is ideal, but may be a concern in practice. However, it should be noted that an alternative version of the construction where $\widetilde{k}$ is replaced by an independent and unrelated key value k' achieves the same security level at the cost of a longer $(2\kappa + n)$-bit key, which is for instance still shorter than in DESX with independent whitening keys (for DES parameters).

The core of our security proof (cf. Theorem 3) is a technical argument of independent interest: Namely, we prove that it is hard to distinguish two random, independent, permutations π_1, π_2 on the n-bit strings from two randomly chosen permutations π_1, π_2 with the property that $\pi_2(\pi_1(x \oplus Z) \oplus Z) = x$ for all x and a random secret value Z even if we are allowed arbitrary queries to each of π_1, π_2, π_1^{-1}, and π_2^{-1}. This fact yields our main theorem by a careful adaptation of the techniques from [11,16] to take into account both randomization and the use of related keys.

Generic attacks and optimality. With the above result at hand, it is legitimate to ask whether we should expect two-call constructions with even better security: In Section 3, we answer this in the negative, at least for a class of natural constructions.

As a warm up of independent interest, we confirm that only much weaker security can be achieved by a one-call construction: Regardless of the amount of key material employed in the construction, an attack with query complexity $2^{\max\{\kappa,n\}}$ always exists (using memory $2^{\max\{\kappa,n\}}$),[4] showing the optimality of DESX-like constructions in the case $\kappa = n$. We then turn to two-call constructions, which are necessary to achieve higher security: Here, we prove that any construction for which distinct inputs map to distinct first queries and distinct answers from the first call imply distinct inputs to the second call admits a distinguishing attack making $2^{\kappa+n/2}$ ideal block cipher queries and 2^n construction queries. This class contains as a special case all constructions obtained by randomizing the cascade of length two using arbitrarily many key bits, including ours.

[4] More precisely, our attack requires roughly 2^κ ideal block cipher queries and 2^n construction queries.

In addition, we also show that simpler randomization methods for length-two cascades admit distinguishing attacks with even lower complexity. For example, randomizing the cascade of length two as $E_{k_2}(E_{k_1}(m \oplus z_1)) \oplus z_2$ instead of using our approach yields a simple $2^{\max\{\kappa,n\}}$ meet-in-the middle attack. This shows an interesting feature of our constructions, namely that while targeting CCA security (i.e., we allow for forward and backward queries to the construction), our design requires *asymmetry*, a fact which seems to contradict common wisdom.

Finally, note that all generic attacks presented in this paper (both against one-query and two-query constructions) can be mounted even if the distinguisher is only allowed to ask forward construction queries (i.e., in the CPA setting). In contrast, the construction we propose is proven to be secure even with respect to an adversary allowed to ask inverse construction queries (CCA adversary).

Final remarks. Table 1 summarizes the results of this paper in the context of previously known results. To serve as an overview, some bounds are presented in a simplified form. Note that the security of *any* key-length extension construction in our model can be upper-bounded by $2^{\kappa+n}$ which corresponds to the trivial attack asking all possible block cipher and construction queries.

Our results and proofs are presented using Maurer's random systems framework [22], which we review in Section 2 in a self-contained way sufficient to follow the contents of the paper.

2 Preliminaries

2.1 Basic Notation

We denote sets by calligraphic letters $\mathcal{X}, \mathcal{Y}, \ldots$, and by $|\mathcal{X}|, |\mathcal{Y}|, \ldots$ their cardinalities. We also let $\mathcal{X}^k$ be the set of k-tuples $x^k = (x_1, \ldots, x_k)$ of elements of $\mathcal{X}$. Strings are elements of $\{0,1\}^k$ and are usually denoted as $s = s_1 s_2 \ldots s_k$, with $\|$ denoting the usual string concatenation. Additionally, we let $\mathrm{Func}(m, \ell)$ be the set of all functions from $\{0,1\}^m$ to $\{0,1\}^\ell$ and $\mathrm{Perm}(n)$ be the set of all permutations of $\{0,1\}^n$. In particular, $id \in \mathrm{Perm}(n)$ represents the identity mapping when n is understood from the context. Throughout this paper logarithms will always be to the base 2.

We denote random variables and concrete values they can take by upper-case letters $X, Y, \ldots$ and lower-case letters $x, y, \ldots$, respectively. For events A and B and random variables U and V with ranges $\mathcal{U}$ and $\mathcal{V}$, respectively, we let $\mathsf{P}_{UA|VB}$ be the corresponding conditional probability distribution, seen as a (partial) function $\mathcal{U} \times \mathcal{V} \to [0,1]$. Here the value $\mathsf{P}_{UA|VB}(u,v) = \mathsf{P}[U = u \wedge A | V = v \wedge B]$ is well defined for all $u \in \mathcal{U}$ and $v \in \mathcal{V}$ such that $\mathsf{P}_{VB}(v) > 0$ and undefined otherwise. Two probability distributions P_U and $\mathsf{P}_{U'}$ on the same set $\mathcal{U}$ are equal, denoted $\mathsf{P}_U = \mathsf{P}_{U'}$, if $\mathsf{P}_U(u) = \mathsf{P}_{U'}(u)$ for all $u \in \mathcal{U}$. Conditional probability distributions are equal if the equality holds for all arguments for which both of them are defined. To emphasize the random experiment $\mathcal{E}$ in consideration, we sometimes write it in the superscript, e.g. $\mathsf{P}^{\mathcal{E}}_{U|V}(u,v)$. Finally, the complement of an event A is denoted by $\overline{A}$.

2.2 Random Systems

The presentation of this paper relies on Maurer's random systems framework [22]. However, we stress that most of the paper remains understandable at a very high level, even without the need of a deeper understanding of the techniques behind the framework; we provide a self-contained introduction.

The starting point of the random-system framework is the basic observation that the input-output behavior of any kind of discrete system with inputs in $\mathcal{X}$ and outputs in $\mathcal{Y}$ can be described by an infinite family of functions describing, for each $i \geq 1$, the probability distribution of the i-th output $Y_i \in \mathcal{Y}$ given the values of the first i inputs $X^i \in \mathcal{X}^i$ and the previous $i-1$ outputs $Y^{i-1} \in \mathcal{Y}^{i-1}$. Formally, hence, an $(\mathcal{X}, \mathcal{Y})$-*(random) system* $\mathbf{F}$ is an infinite sequence of functions $\mathsf{p}^{\mathbf{F}}_{Y_i|X^iY^{i-1}} : \mathcal{Y} \times \mathcal{X}^i \times \mathcal{Y}^{i-1} \to [0,1]$ such that, $\sum_{y_i} \mathsf{p}^{\mathbf{F}}_{Y_i|X^iY^{i-1}}(y_i, x^i, y^{i-1}) = 1$ for all $i \geq 1$, $x^i \in \mathcal{X}^i$ and $y^{i-1} \in \mathcal{Y}^{i-1}$. We stress that the notation $\mathsf{p}^{\mathbf{F}}_{Y_i|X^iY^{i-1}}$, by itself, involves some abuse, as we are not considering any particular random experiment with well-defined random variables Y_i, X^i, Y^{i-1} until the system will be interacting with a distinguisher (see below), in which case the random variables will exist and take the role of the transcript. In general, we shall also typically define discrete systems by a high level description, as long as the resulting conditional probability distributions can be derived uniquely from this description.

We say that a system $\mathbf{F}$ is *deterministic* if the range of $\mathsf{p}^{\mathbf{F}}_{Y_i|X^iY^{i-1}}$ is $\{0,1\}$ for all $i \geq 1$. Moreover, it is *stateless* if the probability distribution of each output depends only on the current input, i.e., if there exists a distribution $\mathsf{p}_{Y|X} : \mathcal{Y} \times \mathcal{X} \to [0,1]$ such that $\mathsf{p}^{\mathbf{F}}_{Y_i|X^iY^{i-1}}(y_i, x^i, y^{i-1}) = \mathsf{p}_{Y|X}(y_i, x_i)$ for all y_i, x^i and y^{i-1}.

We also consider systems $\mathbf{C}^{\mathbf{F}}$ that arise from *constructions* $\mathbf{C}^{(\cdot)}$ accessing a sub-system $\mathbf{F}$. Note that while a construction $\mathbf{C}^{(\cdot)}$ does not define a random system by itself, $\mathbf{C}^{\mathbf{F}}$ does define a random system. The notions of being deterministic and of being stateless naturally extend to constructions.[5] We also consider the *parallel composition* of two (possibly dependent) discrete systems $\mathbf{F}$ and $\mathbf{G}$, denoted $(\mathbf{F}, \mathbf{G})$, which is the system that allows queries to both systems $\mathbf{F}$ and $\mathbf{G}$.

Examples. A *random function* $\mathbf{F} : \{0,1\}^m \to \{0,1\}^n$ is a system which implements a function f initially chosen according to some distribution on $\mathrm{Func}(m,n)$.[6] In particular, the *uniform random function (URF)* $\mathbf{R} : \{0,1\}^m \to \{0,1\}^\ell$ realizes a uniformly chosen function $f \in \mathrm{Func}(m,\ell)$, whereas the *uniform random permutation (URP) on* $\{0,1\}^n$, denoted $\mathbf{P} : \{0,1\}^n \times \{+,-\} \to \{0,1\}^n$, realizes a uniformly chosen permutation $P \in \mathrm{Perm}(n)$ allowing both forward queries of the form $(x,+)$ returning $P(x)$ as well as backward queries $(y,-)$ returning $P^{-1}(y)$. More generally, we meet the convention (for the purpose of this paper) that any

[5] We dispense with a formal definition. However, we point out that we allow a stateless construction to keep a state during invocations of its subsystem.

[6] As is the case with the notion of a random variable, the word "random" does not imply uniformity of the distribution.

system realizing a random function (possibly by means of a construction) which is a permutation will *always* allow both forward and backward queries.

Another important example of a random function is the *ideal block cipher* $\mathbf{E} : \{0,1\}^\kappa \times \{0,1\}^n \times \{+,-\} \to \{0,1\}^n$ which realizes an independent uniform random permutation $\mathbf{E}_k \in \mathrm{Perm}(n)$ for each key $k \in \{0,1\}^\kappa$; in particular, the system allows both forward and backward queries to each $\mathbf{E}_k$.

Finally, note that with some abuse of notation, we often write $\mathbf{E}_k$ or $\mathbf{P}$ to refer to the randomly chosen permutation P implemented by the system $\mathbf{E}_k$ or $\mathbf{P}$, respectively.

Distinguishers and indistinguishability. A *distinguisher* $\mathbf{D}$ for an $(\mathcal{X},\mathcal{Y})$-random system asking q queries is a $(\mathcal{Y},\mathcal{X})$-random system which is "one query ahead:" its input-output behavior is defined by the conditional probability distributions of its queries $\mathsf{p}^{\mathbf{D}}_{X_i|X^{i-1}Y^{i-1}}$ for all $1 \leq i \leq q$. (The first query of $\mathbf{D}$ is determined by $\mathsf{p}^{\mathbf{D}}_{X_1}$.) After the distinguisher asks all q queries, it outputs a bit W_q depending on the transcript (X^q, Y^q). For a random system $\mathbf{F}$ and a distinguisher $\mathbf{D}$, let $\mathbf{DF}$ be the random experiment where $\mathbf{D}$ interacts with $\mathbf{F}$, with the distributions of the transcript (X^q, Y^q) and of the bit W_q being uniquely defined by their conditional probability distributions. Then, for two $(\mathcal{X},\mathcal{Y})$-random systems $\mathbf{F}$ and $\mathbf{G}$, the *distinguishing advantage* of $\mathbf{D}$ in distinguishing systems $\mathbf{F}$ and $\mathbf{G}$ by q queries is the quantity $\Delta^{\mathbf{D}}(\mathbf{F},\mathbf{G}) = \left|\mathsf{P}^{\mathbf{DF}}(W_q = 1) - \mathsf{P}^{\mathbf{DG}}(W_q = 1)\right|$. We are usually interested in the maximal distinguishing advantage over all distinguishers asking q queries, which we denote by $\Delta_q(\mathbf{F},\mathbf{G}) = \max_{\mathbf{D}} \Delta^{\mathbf{D}}(\mathbf{F},\mathbf{G})$ (with $\mathbf{D}$ ranging over all such distinguishers).

For a random system $\mathbf{F}$, we often consider an internal *monotone condition* defined on it. Such a condition is initially satisfied (true), but once it gets violated, it cannot become true again (hence the name monotone). Typically, the condition captures whether the behavior of the system meets some additional requirement (e.g. distinct outputs, consistent outputs) or this was already violated during the interaction. We formalize such a condition by a sequence of events $\mathcal{A} = A_0, A_1, \ldots$ such that A_0 always holds, and A_i holds if the condition holds after query i. The probability that a distinguisher $\mathbf{D}$ issuing q queries to $\mathbf{F}$ makes a monotone condition $\mathcal{A}$ fail in the random experiment $\mathbf{DF}$ is denoted by $\nu^{\mathbf{D}}(\mathbf{F}, \overline{A_q}) = \mathsf{P}^{\mathbf{DF}}(\overline{A_q})$ and we are again interested in the maximum over all such distinguishers, denoted by $\nu(\mathbf{F}, \overline{A_q}) = \max_{\mathbf{D}} \nu^{\mathbf{D}}(\mathbf{F}, \overline{A_q})$. For a random system $\mathbf{F}$ with a monotone condition $\mathcal{A} = A_0, A_1, \ldots$ and a random system $\mathbf{G}$, we say that $\mathbf{F}$ *conditioned on* $\mathcal{A}$ *is equivalent to* $\mathbf{G}$, denoted $\mathbf{F}|\mathcal{A} \equiv \mathbf{G}$, if $\mathsf{p}^{\mathbf{F}}_{Y_i|X^iY^{i-1}A_i} = \mathsf{p}^{\mathbf{G}}_{Y_i|X^iY^{i-1}}$ for $i \geq 1$, for all arguments for which $\mathsf{p}^{\mathbf{F}}_{Y_i|X^iY^{i-1}A_i}$ is defined. Intuitively, this captures the fact that as long as the condition $\mathcal{A}$ holds in $\mathbf{F}$, it behaves the same as $\mathbf{G}$.

Let $\mathbf{F}$ be a random system with a monotone condition $\mathcal{A}$. Following [25], we define $\mathbf{F}$ *blocked by* $\mathcal{A}$ to be a new random system that behaves exactly like $\mathbf{F}$ while the condition $\mathcal{A}$ is satisfied. Once $\mathcal{A}$ is violated, it only outputs a special blocking symbol $\perp$ not contained in the output alphabet of $\mathbf{F}$.

We make use of the following helpful claims proven in previous papers. Below, we also present an informal explanation of their merits.

Lemma 1. *Let $\mathbf{C}^{(\cdot)}$ and $\mathbf{C}'^{(\cdot)}$ be two constructions invoking a subsystem, and let $\mathbf{F}$ and $\mathbf{G}$ be random systems. Let $\mathcal{A}$ and $\mathcal{B}$ be two monotone conditions defined on $\mathbf{F}$ and $\mathbf{G}$, respectively.*

(i) *[22, Theorem 1] If $\mathbf{F}|\mathcal{A} \equiv \mathbf{G}$ then $\Delta_q(\mathbf{F}, \mathbf{G}) \leq \nu(\mathbf{F}, \overline{A_q})$.*

(ii) *[16, Lemma 2] Let $\mathbf{F}^{\perp}$ denote the random system $\mathbf{F}$ blocked by $\mathcal{A}$ and let $\mathbf{G}^{\perp}$ denote $\mathbf{G}$ blocked by $\mathcal{B}$. Then for every distinguisher $\mathbf{D}$ asking q queries we have $\Delta^{\mathbf{D}}(\mathbf{F}, \mathbf{G}) \leq \Delta_q(\mathbf{F}^{\perp}, \mathbf{G}^{\perp}) + \nu^{\mathbf{D}}(\mathbf{F}, \overline{A_q})$.*

(iii) *[22, Lemma 5] $\Delta_q(\mathbf{C}^{\mathbf{F}}, \mathbf{C}^{\mathbf{G}}) \leq \Delta_{q'}(\mathbf{F}, \mathbf{G})$, where q' is the maximum number of invocations of any internal system $\mathbf{H}$ for any sequence of q queries to $\mathbf{C}^{\mathbf{H}}$, if such a value is defined.*

(iv) *[16, Lemma 3] There exists a fixed permutation $S \in \mathrm{Perm}(n)$ (represented by a deterministic stateless system) such that $\Delta_q(\mathbf{C}^{\mathbf{P}}, \mathbf{C}'^{\mathbf{P}}) \leq \Delta_q(\mathbf{C}^{S}, \mathbf{C}'^{S})$.*

The first claim can be seen as a generalized version of the Fundamental Lemma of Game-Playing for the context of random systems, stating that if two systems are equivalent as long as some condition is satisfied, then the advantage in distinguishing these systems can be upper-bounded by the probability of violating this condition. The second claim is even more general, analyzing the situation where the systems are not equivalent even if the conditions defined on them are satisfied, but their behavior is similar (which is captured by the term $\Delta_q(\mathbf{F}^{\perp}, \mathbf{G}^{\perp})$). The third claim states the intuitive fact that interacting with the distinguished systems through an additional enveloping construction $\mathbf{C}$ cannot improve the distinguishing advantage and the last claim is just an averaging argument over all the possible values taken by $\mathbf{P}$.

3 Generic Attacks against Efficient Key-Length Extension Schemes

We start by addressing the following question: *What is the maximum achievable security level for very efficient key-length extension schemes?* To this end, this section presents generic distinguishing attacks against one- and two-call block-cipher constructions in Sections 3.1 and 3.2, respectively. These attacks are in the same spirit as the recent line of work on generic attacks on hash functions (cf. e.g. [27,29,30]). Along the same lines, here attack complexity will be measured in terms of query- rather than time-complexity. This allows us to consider arbitrary constructions, while being fully sufficient to assess security in the ideal cipher model, where distinguishers are computationally unrestricted.

More formally, we consider stateless and deterministic (keyed) constructions $\mathbf{C}^{(\cdot)}$ invoking an ideal cipher $\mathbf{E} : \{0,1\}^{\kappa} \times \{0,1\}^{n} \times \{+,-\} \to \{0,1\}^{n}$ to implement a function $\mathbf{C}^{\mathbf{E}} : \{0,1\}^{\kappa'} \times \{0,1\}^{n} \times \{+,-\} \to \{0,1\}^{n}$ to serve as a block cipher with key length κ'. We assume that the construction $\mathbf{C}^{\mathbf{E}}$ realizes a permutation for each $k' \in \{0,1\}^{\kappa'}$ and hence it also provides the interface for inverse queries as indicated. Consequently, for a random (secret) κ'-bit string K', we let $\mathbf{C}^{\mathbf{E}}_{K'}$ denote the system which only gives access to the permutation

$\mathbf{C}^{\mathbf{E}}(K', \cdot)$ and its inverse (i.e., takes inputs from $\{0,1\}^n \times \{+,-\}$). (In fact, none of the attacks in this section will require backward queries.)

Since the goal of this section is mainly to serve as a supporting argument for the optimality of our construction presented in Section 4, due to space restrictions we omit the proofs of our claims and only provide some intuition. All statements are proved in a more general setting in the full version of this paper.

3.1 One-Query Constructions

Throughout this section, we assume that $\mathbf{C}^{(\cdot)}$, to evaluate input $(x, +)$ for $x \in \{0,1\}^n$ under a key $k' \in \{0,1\}^{\kappa'}$, makes exactly one query to the underlying subsystem, and we denote this query as $q(k', x)$. We consider two different cases, depending on the structure of $q(\cdot, \cdot)$, before deriving the final attack.

The injective case. We first consider the case where the mapping $x \mapsto q(k', x)$ is injective for each k'. We shall denote this as a *one-injective-query construction*. In this case, distinct queries to $\mathbf{C}_{k'}^{\mathbf{E}}$ lead to distinct internal queries to $\mathbf{E}$ and hence if the distinguisher queries both $\mathbf{C}_{K'}^{\mathbf{E}}$ and $\mathbf{E}$ at sufficiently many random positions, one can expect that during the evaluation of the outer queries, $\mathbf{C}_{K'}^{(\cdot)}$ asks $\mathbf{E}$ for a value that was also asked by the distinguisher. If this occurs, the distinguisher can, while trying all possible keys k', evaluate $\mathbf{C}_{k'}^{(\cdot)}$ on its own by simulating $\mathbf{C}^{(\cdot)}$ and using the response from $\mathbf{E}$; and by comparing the outcomes it can distinguish the construction from a truly random permutation. This is the main idea behind the following lemma.

Lemma 2. *Let* $\mathbf{E}\colon \{0,1\}^{\kappa} \times \{0,1\}^n \times \{+,-\} \to \{0,1\}^n$ *be an ideal block cipher, let* $\mathbf{C}^{(\cdot)}\colon \{0,1\}^{\kappa'} \times \{0,1\}^n \times \{+,-\} \to \{0,1\}^n$ *be a one-injective-query construction and let* $\mathbf{P}$ *be a URP on* $\{0,1\}^n$*. Then, for a random key* $K' \in \{0,1\}^{\kappa'}$ *and every parameter* $0 < t < 2^{\min\{n,k\}-1}$*,*[7] *there exists a distinguisher* $\mathbf{D}$ *such that*

$$\Delta^{\mathbf{D}}((\mathbf{E}, \mathbf{C}_{K'}^{\mathbf{E}}), (\mathbf{E}, \mathbf{P})) \geq 1 - 2/t - 2^{\kappa' - t\cdot(n-1)} ,$$

and which makes at most $4t \cdot 2^{\max\{(\kappa+n)/2, \kappa\}}$ *queries to the block cipher* $\mathbf{E}$*, as well as at most* $2 \cdot 2^{\min\{(\kappa+n)/2, n\}}$ *forward queries to either of* $\mathbf{C}_{K'}^{\mathbf{E}}$ *and* $\mathbf{P}$*.*

The above lemma covers most of the natural one-query constructions, since these typically satisfy the injectivity requirement (e.g. the DESX construction). In the following we see that constructions asking non-injective queries do not achieve any improvement in security.

Non-injective queries. We now permit that the construction $\mathbf{C}^{(\cdot)}$ might, for some key k', invoke the underlying ideal cipher in a *non-injective* way, i.e., $q(k', \cdot)$ is not an injective map. We prove that, roughly speaking, such a construction $\mathbf{C}_{K'}^{\mathbf{E}}$ might be distinguishable from a URP $\mathbf{P}$ based solely on an entropy argument. The intuitive reasoning is that if $\mathbf{C}^{(\cdot)}$ allows on average (over the choice of

[7] Roughly speaking, higher t increases the advantage but also the required number of queries; we obtain the desired bound using a constant t. For a first impression, consider e.g. $t = 4$ and $\kappa' \approx 2n$.

the key k') that too many queries x map to the same $q(k', x)$, then it also does not manage to obtain sufficient amount of randomness from the underlying random function to simulate $\mathbf{P}$ convincingly, opening the door to a distinguishing attack. In the following, let $q(k') = |\{q(k', x) : x \in \{0,1\}^n\}|$ for all $k' \in \{0,1\}^\kappa$.

Lemma 3. *Let* $\mathbf{C}^{(\cdot)} : \{0,1\}^{\kappa'} \times \{0,1\}^n \times \{+,-\} \to \{0,1\}^n$ *be a one-query construction, let* $\mathbf{P}$ *be a URP on* $\{0,1\}^n$ *and let* $\mathbf{E} : \{0,1\}^\kappa \times \{0,1\}^n \times \{+,-\} \to \{0,1\}^n$ *be an ideal block cipher. Also, let* $K' \in \{0,1\}^{\kappa'}$ *be a random key, and assume that there exists* q^* *such that* $q(K') \leq q^*$ *with probability at least* $\frac{1}{2}$*. Then, there exists a distinguisher* $\mathbf{D}$ *asking* 2^n *forward queries such that*

$$\Delta^{\mathbf{D}}\left(\mathbf{C}^{\mathbf{E}}_{K'}, \mathbf{P}\right) \geq \tfrac{1}{2} - 2^{\kappa' + n \cdot q^* - \log(2^n!)} .$$

Putting the pieces together. We can combine the techniques used to prove Lemma 2 (somewhat relaxing the injectivity requirement) and Lemma 3, to obtain the following final theorem yielding an attack for arbitrary one-query block-cipher constructions.

Theorem 1. *Let* $n \geq 6$ *and* $\kappa' \leq 2^n - 1$*, let* $\mathbf{E} : \{0,1\}^\kappa \times \{0,1\}^n \times \{+,-\} \to \{0,1\}^n$ *be an ideal block cipher, let* $\mathbf{C}^{(\cdot)} : \{0,1\}^{\kappa'} \times \{0,1\}^n \times \{+,-\} \to \{0,1\}^n$ *be a one-query construction, and let* $\mathbf{P}$ *be a URP on* $\{0,1\}^n$*. Then, for a random key* $K' \in \{0,1\}^{\kappa'}$ *and for all parameters* $0 < t < 2^{n-2}$*, there exists a distinguisher* $\mathbf{D}$ *such that*

$$\Delta^{\mathbf{D}}\left((\mathbf{E}, \mathbf{C}^{\mathbf{E}}_{K'}), (\mathbf{E}, \mathbf{P})\right) \geq \min\left\{\tfrac{1}{4}, \tfrac{1}{2} - \tfrac{2}{t} - 2^{\kappa' - t\cdot(n-1)}\right\} ,$$

and which asks at most $8t \cdot 2^\kappa$ *queries to* $\mathbf{E}$ *and* 2^n *forward queries to either of* $\mathbf{C}^{\mathbf{E}}_{K'}$ *and* $\mathbf{P}$*.*

Theorem 1 shows that no one-query construction can achieve security beyond $2^{\max\{\kappa, n\}}$ queries, hence in our search for efficient key-length extension schemes we have to we turn our attention towards constructions issuing at least two queries.

3.2 Two-Query Constructions

We now consider an arbitrary deterministic stateless construction $\mathbf{C}^{(\cdot)} : \{0,1\}^{\kappa'} \times \{0,1\}^n \times \{+,-\} \to \{0,1\}^n$ that makes exactly two queries to an ideal block cipher $\mathbf{E} : \{0,1\}^k \times \{0,1\}^n \times \{+,-\} \to \{0,1\}^n$ to evaluate each query. In the following, these constructions shall be referred to as *two-query constructions.* We denote by $q_1(k', x) \in \{0,1\}^\kappa \times \{0,1\}^n \times \{+,-\}$ the first query $\mathbf{C}^{(\cdot)}$ asks its subsystem when it is itself being asked a forward query $(k', x, +)$. Moreover, we denote by $q_2(k', x, s) \in \{0,1\}^\kappa \times \{0,1\}^n \times \{+,-\}$ the second query it asks when it is itself being asked a forward query $(k', x, +)$ and the answer to the first query $q_1(k', x)$ was $s \in \{0,1\}^n$. Since $\mathbf{C}^{(\cdot)}$ is deterministic and stateless, both q_1 and q_2 are well-defined mappings.

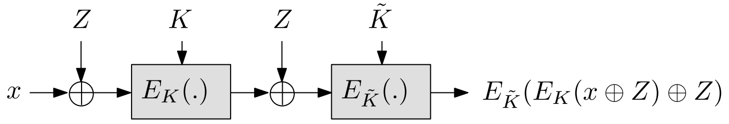

Fig. 1. The double XOR-cascade construction analyzed in Theorem 3

Theorem 2. *Let* $\mathbf{C}^{(\cdot)}\colon \{0,1\}^{\kappa'} \times \{0,1\}^n \times \{+,-\} \to \{0,1\}^n$ *be a two-query construction satisfying the following two conditions:*

1. *for every* $k' \in \{0,1\}^{\kappa'}$ *the mapping* $q_1(k',\cdot)$ *is injective,*
2. *distinct answers to the first query imply distinct second queries, i.e., for every* $k' \in \{0,1\}^{\kappa'}$ *and every* $x, x' \in \{0,1\}^n$ *if* $s \neq s'$ *then* $q_2(k',x,s) \neq q_2(k',x',s')$.

Then for a random key $K' \in \{0,1\}^{\kappa'}$*, for a URP* $\mathbf{P}$ *on* $\{0,1\}^n$ *and for every parameter* $0 < t < 2^{n/2-1}$*, there exists a distinguisher* $\mathbf{D}$ *such that*

$$\Delta^{\mathbf{D}}((\mathbf{E},\mathbf{C}^{\mathbf{E}}_{K'}),(\mathbf{E},\mathbf{P})) \geq 1 - 2/t - 13 \cdot 2^{-\frac{n}{2}} - 2^{\kappa'-t\cdot(n-1)},$$

where $\mathbf{D}$ *makes at most* $2(t+4)\cdot 2^{\kappa+n/2}$ *queries to* $\mathbf{E}$ *as well as* 2^n *forward queries to either of* $\mathbf{C}^{\mathbf{E}}_{K'}$ *and* $\mathbf{P}$.

Hence, no two-query construction from a large class described in the above theorem can achieve security beyond $2^{\kappa+n/2}$ queries. In the following section we present a simple and efficient construction from this class that achieves the above limit.

4 The Double XOR-Cascade Construction

We present a two-query construction matching the upper bound $2^{\kappa+n/2}$ on security proved in the previous section. The construction, which we call the *double XOR-cascade construction* (2XOR), consists of two applications of the block-cipher interleaved with two XOR operations: Given a (κ,n)-block cipher E, we define the $(\kappa+n,n)$-block cipher 2XOR^E such that

$$2\text{XOR}^E_{k,z}(m) = E_{\widetilde{k}}(E_k(m \oplus z) \oplus z)$$

for all $k \in \{0,1\}^\kappa$, $z, m \in \{0,1\}^n$, and where $\widetilde{k} = \pi(k)$ for some understood fixpoint-free permutation $\pi \in \text{Perm}(\kappa)$ (e.g., $\pi(k) = k \oplus 0^{\kappa-1}1$, i.e., π flips the last bit). The construction is depicted in Figure 1. Note that both XOR transformations use the same value z and the two block-cipher calls use two distinct keys such that one can be deterministically derived from the other one. We also consider a construction $2\text{XOR}'$ of a $(2\kappa+n)$-block cipher where $\widetilde{k}$ is replaced by an (independent) κ-bit key.

Security of 2XOR. We now discuss the security of the double XOR-cascade construction in the ideal cipher model. To this end, let $\mathbf{X}^{(\cdot)}\colon \{0,1\}^\kappa \times \{0,1\}^n \times \{0,1\}^n \times \{+,-\} \to \{0,1\}^n$ denote a (deterministic stateless) construction which expects a subsystem $\mathbf{E}\colon \{0,1\}^\kappa \times \{0,1\}^n \times \{+,-\} \to \{0,1\}^n$ realizing a block cipher. $\mathbf{X}^{\mathbf{E}}$ then answers each query $(k,z,x,+)$ by $\mathbf{E}_{\tilde{k}}\left(\mathbf{E}_k\left(x \oplus z\right) \oplus z\right)$ and each query $(k,z,y,-)$ by $\mathbf{E}_k^{-1}(\mathbf{E}_{\tilde{k}}^{-1}(y) \oplus z) \oplus z$. As before, for randomly chosen (secret) keys $(K,Z) \in \{0,1\}^\kappa \times \{0,1\}^n$, we let $\mathbf{X}^{\mathbf{E}}_{K,Z}$ be the system which gives access to the permutation $\mathbf{X}^{\mathbf{E}}(K,Z,\cdot)$ in both directions (i.e., takes inputs from $\{0,1\}^n \times \{+,-\}$).

Theorem 3. *Let* $\mathbf{P}$ *and* $\mathbf{E}$ *denote a URP on* $\{0,1\}^n$ *and an ideal* (κ,n)*-block cipher, respectively; let* $(K,Z) \in \{0,1\}^\kappa \times \{0,1\}^n$ *be uniformly chosen keys. For the construction* $\mathbf{X}^{(\cdot)}_{K,Z}$ *defined as above, and for every distinguisher* $\mathbf{D}$ *making* q *queries to* $\mathbf{E}$,

$$\Delta^{\mathbf{D}}\left(\left(\mathbf{E}, \mathbf{X}^{\mathbf{E}}_{K,Z}\right), (\mathbf{E}, \mathbf{P})\right) \leq 4 \cdot \left(\frac{q}{2^{\kappa+n/2}}\right)^{2/3} .$$

In particular, $\mathbf{D}$ *can make arbitrarily many queries to either of* $\mathbf{X}^{\mathbf{E}}_{K,Z}$ *and* $\mathbf{P}$.

We also note that an analogous statement for the construction 2XOR′ can be easily derived from the presented claim.

Proof intuition. The proof, given below, follows a two-step approach. In the first part, we prove that for any parameter $h \leq 2^{n/2}$, the above advantage is upper bounded by $\varepsilon(h) + \frac{q}{h2^{\kappa-1}}$, where $\varepsilon(h)$ is an upper bound on the advantage of a h-query distinguisher in telling apart the following two settings, in both of which it issues both forward and backward queries to two permutations $\pi_1, \pi_2 \in \mathrm{Perm}(n)$:

1. In the first case, π_1, π_2 are chosen uniformly and independently.
2. In the second setting, a uniform n-bit string Z is chosen, and π_1 and π_2 are chosen uniformly at random such that $\pi_2(\pi_1(\cdot \oplus Z) \oplus Z) = id$.

This step follows a pattern similar to the one used in [11,16] to analyze the security of plain cascades, but with obvious modifications and extra care to take into account randomization as well as key dependency.

Then, the main technical part of the proof consists of proving a bound $3h^2/2^{n+1}$ on $\varepsilon(h)$, which is a new result of independent interest. The intuition here is that without knowing Z, it is hard to come up with two queries, one to π_1 and one to π_2, which result in input-output pairs $\pi_1(x) = y$ and $\pi_2(x') = y'$ satisfying $x = y' \oplus Z$ and $x' = y \oplus Z$ simultaneously. However, as long as this does not happen, both permutations appear independent and random.

We stress that our double-randomization is crucial here: omitting one of the randomization steps, as well as adding a third randomization step for the same Z, would all result in invalidating the argument. The full version of this paper also provides some useful extra intuition as for why other simpler randomization methods for the cascade fail to provide the required security level.

Proof. We start by noting that the system $(\mathbf{E}, \mathbf{X}^{\mathbf{E}}_{K,Z})$ simply can be seen as providing an interface to query $2^\kappa + 1$ (dependent) permutations

$$\mathbf{E}_{k_1}, \mathbf{E}_{k_2}, \ldots, \mathbf{E}_{k_{2^\kappa}}, \mathbf{E}_{\tilde{K}} \left(\mathbf{E}_K \left(\cdot \oplus Z\right) \oplus Z\right) \ ,$$

each both in forward and backward direction, where $k_1, k_2, \ldots, k_{2^\kappa}$ is an enumeration of the κ-bit strings. By the group structure of $\mathrm{Perm}(n)$ under composition, the joint distribution of these permutations does not change if we start by choosing the last permutation uniformly at random, i.e., we replace it by $\mathbf{P}$, then choose K and Z and finally choose the permutations of the block cipher independently and randomly except for the one corresponding to the key $\tilde{K}$, which we set to $x \mapsto \mathbf{P}\left(\mathbf{E}_K^{-1}\left(x \oplus Z\right) \oplus Z\right)$. Hence, let $\mathbf{G}^{(\cdot)}$ be a system that expects a single permutation as its subsystem (let us denote it P) and itself provides an interface to a block cipher (let us denote it G). It answers queries to G in the following way: in advance, it chooses random keys (K, Z) and then generates random independent permutations for G used with any key except $\tilde{K}$. For $\tilde{K}$, $\mathbf{G}$ realizes the permutation $x \mapsto P\left(G_K^{-1}\left(x \oplus Z\right) \oplus Z\right)$, querying P for any necessary values. Then the above argument shows that $(\mathbf{E}, \mathbf{X}^{\mathbf{E}}_{K,Z}) = (\mathbf{G}^{\mathbf{P}}, \mathbf{P})$ and hence we obtain

$$\Delta_q\left(\left(\mathbf{E}, \mathbf{X}^{\mathbf{E}}_{K,Z}\right), (\mathbf{E}, \mathbf{P})\right) = \Delta_q\left(\left(\mathbf{G}^{\mathbf{P}}, \mathbf{P}\right), (\mathbf{E}, \mathbf{P})\right) \leq \Delta_q\left(\left(\mathbf{G}^{S}, S\right), (\mathbf{E}, S)\right),$$

where the last inequality follows from claim (iv) in Lemma 1 and S denotes the fixed permutation whose existence is guaranteed by this claim. Since S is fixed and hence known to the distinguisher, it makes no sense to query it and thus it remains to bound $\Delta_q\left(\mathbf{G}^S, \mathbf{E}\right)$ for any permutation S. From now on, we denote the system $\mathbf{G}^S$ by $\mathbf{G}$ to simplify the notation.

We shall refer to all forward or backwards queries to G involving the permutations indexed by K or $\tilde{K}$ as *relevant.* Similarly, the system $\mathbf{E}$ can be seen as also choosing some random key K (and hence $\tilde{K}$), this just does not affect its behavior, and we can hence define relevant queries for $\mathbf{E}$ in an analogous way. Let $\mathcal{A}^h$ and $\mathcal{B}^h$ denote monotone conditions defined on systems $\mathbf{E}$ and $\mathbf{G}$ respectively, such that each of these conditions remains satisfied as long as at most h of the queries asked so far were relevant. The parameter h will be chosen optimally at the end of the proof. We require $h < 2^{n/2}$.

It is easy to upper-bound the probability of asking more than h relevant queries in $\mathbf{E}$: since the key K does not affect the responses of the system (and therefore the behavior is also independent of the associated monotone condition), we only have to consider non-adaptive strategies. Hence, for any distinguisher $\mathbf{D}$ asking q queries, the expected number of relevant queries among them is $q \cdot 2^{1-\kappa}$ and using Markov inequality, we obtain $\nu(\mathbf{E}, \overline{\mathcal{A}}^h_q) \leq q/h2^{\kappa-1}$. Let $\mathbf{E}^\perp$ and $\mathbf{G}^\perp$ denote the systems $\mathbf{E}$ and $\mathbf{G}$ blocked by $\mathcal{A}^h$ and $\mathcal{B}^h$, respectively. Then we can apply claim (ii) of Lemma 1 to obtain

$$\Delta_q(\mathbf{G}, \mathbf{E}) \leq \Delta_q(\mathbf{G}^\perp, \mathbf{E}^\perp) + \nu(\mathbf{E}, \overline{\mathcal{A}}^h_q) \leq \Delta_q(\mathbf{G}^\perp, \mathbf{E}^\perp) + q/h2^{\kappa-1} \ .$$

Now, one can observe that the systems $\mathbf{G}^\perp$ and $\mathbf{E}^\perp$ only differ in a small part. More specifically, we have $\mathbf{G}^\perp = \mathbf{C}^{\mathbf{S}}$ and $\mathbf{E}^\perp = \mathbf{C}^{\mathbf{T}}$, where:

- **S** is a system that chooses $Z \in \{0,1\}^n$ at random and provides access (by means of both forward and backward queries) to two randomly chosen permutations π_1, π_2 on $\{0,1\}^n$ such that they satisfy the equation $\pi_2(\pi_1(\cdot \oplus Z) \oplus Z) = id$;
- **T** is a system providing access (by means of both forward and backward queries) to two independent uniformly random permutations $\pi_1, \pi_2 \in \mathrm{Perm}(n)$;
- $\mathbf{C}^{(\cdot)}$ is a (randomized) construction that expects a subsystem which provides two permutations π_1 and π_2. $\mathbf{C}^{(\cdot)}$ itself provides access to a block cipher C as follows: it chooses a uniformly random key K and sets $C_K := \pi_1$ and $C_{\tilde{K}} := \pi_2 \circ S$. (**C** only queries its subsystem once it is asked a relevant query). The permutations for all other keys are chosen independently at random. Moreover, **C** only allows h relevant queries, after that it returns $\perp$.

By Lemma 1(iii), the above observation gives us $\Delta_q(\mathbf{G}^\perp, \mathbf{E}^\perp) \leq \Delta_h(\mathbf{S}, \mathbf{T})$ and hence it remains to bound $\Delta_h(\mathbf{S}, \mathbf{T})$. We start by taking a different view of the internal workings of the system **S**. Once the values Z, π_1, π_2 are chosen, the internal state of **S** can be represented by a set $\mathcal{T}$ of 2^n 4-tuples (x_1, y_1, x_2, y_2) such that $\pi_1(x_1) = y_1$ and $\pi_2(x_2) = y_2$, and $x_2 = y_1 \oplus Z$ and $x_1 = y_2 \oplus Z$. For any $\mathcal{I} \subseteq \{1, \ldots, 4\}$, let $\mathcal{T}_{\mathcal{I}}$ be the projection of $\mathcal{T}$ on the components in $\mathcal{I}$. Then note that for any two distinct tuples $(x_1, y_1, x_2, y_2), (x'_1, y'_1, x'_2, y'_2) \in \mathcal{T}$ we have $x_1 \neq x'_1$, $y_1 \neq y'_1$, $x_2 \neq x'_2$, and $y_2 \neq y'_2$, in other words $\mathcal{T}_{\{i\}} = \{0,1\}^n$ for every $i \in \{1, \ldots, 4\}$.

Equivalently, it is not hard to verify that **S** can be implemented using lazy-sampling to set up $\mathcal{T}$: Initially, $\mathcal{T} = \emptyset$ and Z is a uniform n-bit string. Then, **S** answers queries as follows:

- Upon a query $\pi_1(x)$, it returns y if $(x, y) \in \mathcal{T}_{\{1,2\}}$ for some y. Otherwise, it returns a random $y \in \{0,1\}^n \setminus \mathcal{T}_{\{2\}}$ and adds $(x, y, y \oplus Z, x \oplus Z)$ to $\mathcal{T}$.
- Upon a query $\pi_1^{-1}(y)$, it returns x if $(x, y) \in \mathcal{T}_{\{1,2\}}$ for some x. Otherwise, it returns a random $x \in \{0,1\}^n \setminus \mathcal{T}_{\{1\}}$ and adds $(x, y, y \oplus Z, x \oplus Z)$ to $\mathcal{T}$.
- Upon a query $\pi_2(x)$, it returns y if $(x, y) \in \mathcal{T}_{\{3,4\}}$ for some y. Otherwise, it returns a random $y \in \{0,1\}^n \setminus \mathcal{T}_{\{4\}}$ and adds $(y \oplus Z, x \oplus Z, x, y)$ to $\mathcal{T}$.
- Upon a query $\pi_2^{-1}(y)$, it returns x if $(x, y) \in \mathcal{T}_{\{3,4\}}$ for some x. Otherwise, it returns a random $x \in \{0,1\}^n \setminus \mathcal{T}_{\{3\}}$ and adds $(y \oplus Z, x \oplus Z, x, y)$ to $\mathcal{T}$.

We consider an intermediate system $\mathbf{S}'$ obtained from **S**: In addition to $\mathcal{T}$, it also keeps track of sets $\mathcal{P}_1$ and $\mathcal{P}_2$, both consisting of ordered pairs of n-bit strings. (Again $\mathcal{P}_{i,1}$ and $\mathcal{P}_{i,2}$ denote the strings appearing as first and second component in $\mathcal{P}_i$, respectively.) Initially each $\mathcal{P}_i$ is empty and during the experiment, $\mathcal{P}_i$ keeps track of input-output pairs for π_i which were already defined by directly answering a π_i query in either direction (as opposed to those that were defined internally by $\mathbf{S}'$ when answering a π_{3-i} query). Concretely, $\mathbf{S}'$ answers a query $\pi_1(x)$ by y if $(x, y) \in \mathcal{T}_{\{1,2\}} \cup \mathcal{P}_1$ for some y. Otherwise, it returns a uniformly chosen $y \in \{0,1\}^n \setminus \mathcal{P}_{1,2}$ and adds (x, y) to $\mathcal{P}_1$. Moreover, if $y \notin \mathcal{T}_{\{2\}}$, it also adds the tuple $(x, y, y \oplus Z, x \oplus Z)$ to $\mathcal{T}$. Queries $\pi_1^{-1}(y)$, $\pi_2(x)$, and $\pi_2^{-1}(y)$ are answered in a symmetric fashion. Having this description of $\mathbf{S}'$, note that we obtain the system $\mathbf{T}$ if a query $\pi_1(x)$ is answered by some given y only if $(x, y) \in \mathcal{P}_1$, and otherwise a fresh random output is generated (but the 4-tuples are still added to $\mathcal{T}$ as above).

We now define two monotone conditions $\mathcal{A}$ and $\mathcal{B}$ on $\mathbf{S}'$:

- $\mathcal{A} = A_0, A_1, \ldots$ fails at the first query $\pi_i(x)$ answered by a random y which satisfies $y \in \mathcal{T}_{\{2(i-1)+2\}}$, or $\pi_i^{-1}(y)$ answered by a random x such that $x \in \mathcal{T}_{\{2(i-1)+1\}}$.
- $\mathcal{B} = B_0, B_1, \ldots$ fails at the first query $\pi_i(x)$ such that there exists y satisfying $(x, y) \in \mathcal{T}_{\{2(i-1)+1,2(i-1)+2\}} \setminus \mathcal{P}_i$, or $\pi_i^{-1}(y)$ such that there exists x satisfying $(x, y) \in \mathcal{T}_{\{2(i-1)+1,2(i-1)+2\}} \setminus \mathcal{P}_i$.

By the above representations of $\mathbf{S}$ and $\mathbf{T}$, one can easily verify that $\mathbf{S}'|\mathcal{A} \equiv \mathbf{S}$ and $\mathbf{S}'|\mathcal{B} \equiv \mathbf{T}$. Therefore, by the triangle inequality and by claim (i) from Lemma 1,

$$\Delta_h(\mathbf{S}, \mathbf{T}) \leq \Delta_h(\mathbf{S}, \mathbf{S}') + \Delta_h(\mathbf{S}', \mathbf{T}) \leq \nu(\mathbf{S}', \overline{A}_h) + \nu(\mathbf{S}', \overline{B}_h).$$

To upper bound $\nu(\mathbf{S}', \overline{A}_h)$, note that each time a fresh random value is chosen from $\{0,1\}^n \setminus \mathcal{P}_{i,j}$ when answering the i^{th} query, it is in $\mathcal{T}_{2(i-1)+j}$ with probability at most $\frac{i-1}{2^n-i} \leq 2\frac{i-1}{2^n}$, hence the union bound gives us $\nu(\mathbf{S}', \overline{A}_h) \leq \frac{h^2}{2^n}$.

In order to bound $\nu(\mathbf{S}', \overline{B}_h)$, let us introduce a monotone condition $\mathcal{C} = C_0, C_1, \ldots$ on $\mathbf{T}$ which fails under the same circumstances as $\mathcal{B}$ in $\mathbf{S}'$ (note that this can be done since $\mathbf{T}$ also keeps track of the sets $\mathcal{T}$ and $\mathcal{P}_i$). As a consequence of these equivalent definitions and the fact that the behaviors of $\mathbf{S}'$ and $\mathbf{T}$ are the same as long as the respective associated conditions are satisfied, we have $\nu(\mathbf{S}', \overline{B}_h) = \nu(\mathbf{T}, \overline{C}_h)$. However, the input-output behavior of $\mathbf{T}$ is independent of Z (and $\mathcal{C}$ failing), and hence we can equivalently postpone the sampling of Z to the end of the interaction, go through the generated transcript to construct $\mathcal{T}$, and upper bound the probability that $\mathcal{C}$ has failed at some query. This implies that for the choice of Z, one query must have been *bad* in the following sense:

- query $\pi_1(x)$ is preceded by a π_2-query resulting in an input-output pair (x', y') such that $y' \oplus Z = x$;
- query $\pi_1^{-1}(y)$ preceded by a π_2-query resulting in pair (x', y') s.t. $x' \oplus Z = y$;
- query $\pi_2(x')$ preceded by a π_1-query resulting in pair (x, y) s.t. $y \oplus Z = x'$;
- query $\pi_2^{-1}(y')$ is preceded by a π_1-query resulting in pair (x, y) s.t. $x \oplus Z = y'$.

Given the transcript, and for randomly chosen Z, the i^{th} query is bad with probability at most $(i-1)/2^n$, and the probability that at least one query is bad is thus at most $\frac{h^2}{2^{n+1}}$ by the union bound.

Putting all the obtained terms together, the part of the distinguisher's advantage that depends on h is $f(h) = q/h2^{\kappa-1} + 3h^2/2^{n+1}$. This term is minimal for $h^* = (\frac{1}{3}q2^{n-\kappa+1})^{1/3}$ which gives us $f(h^*) < 4 \cdot \left(\frac{q}{2^{\kappa+n/2}}\right)^{2/3}$ as desired. □

Acknowledgements. We thank Mihir Bellare and Ueli Maurer for insightful feedback. Peter Gaži was partially supported by grants SNF 200020-132794, VEGA 1/0266/09 and UK/95/2011. Stefano Tessaro's work was done while at UC San Diego, partially supported by NSF grant CNS-0716790 and by Calit2. He is currently supported by DARPA, contract number FA8750-11-2-0225.

References

1. FIPS PUB 46: Data Encryption Standard (DES). National Institute of Standards and Technology (1977)
2. ANSI X9.52: Triple Data Encryption Algorithm Modes of Operation (1998)
3. FIPS PUB 46-3: Data Encryption Standard (DES). National Institute of Standards and Technology (1999)
4. FIPS PUB 197: Advanced Encryption Standard (AES). National Institute of Standards and Technology (2001)
5. NIST SP 800-67: Recommendation for the Triple Data Encryption Algorithm (TDEA) Block Cipher. National Institute of Standards and Technology (2004)
6. EMV Integrated Circuit Card Specifications for Payment Systems. Book 2: Security and Key Management, v.4.2. EMVCo (June 2008)
7. Aiello, W., Bellare, M., Di Crescenzo, G., Venkatesan, R.: Security Amplification by Composition: The Case of Doubly-Iterated, Ideal Ciphers. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, pp. 390–407. Springer, Heidelberg (1998)
8. Bellare, M., Desai, A., Jokipii, E., Rogaway, P.: A concrete security treatment of symmetric encryption. In: FOCS 1997: Proceedings of the 38th IEEE Annual Symposium on Foundations of Computer Science, pp. 394–403 (1997)
9. Bellare, M., Kilian, J., Rogaway, P.: The Security of Cipher Block Chaining Message Authentication Code. In: Desmedt, Y.G. (ed.) CRYPTO 1994. LNCS, vol. 839, pp. 341–358. Springer, Heidelberg (1994)
10. Bellare, M., Kohno, T.: A Theoretical Treatment of Related-key Attacks: RKA-PRPs, RKA-PRFs, and Applications. In: Biham, E. (ed.) EUROCRYPT 2003. LNCS, vol. 2656, pp. 491–506. Springer, Heidelberg (2003)
11. 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), http://eprint.iacr.org/2004/331
12. Black, J., Rogaway, P.: A Block-Cipher Mode of Operation for Parallelizable Message Authentication. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 384–397. Springer, Heidelberg (2002)
13. Diffie, W., Hellman, M.E.: Exhaustive Cryptanalysis of the NBS Data Encryption Standard. Computer 10(6), 74–84 (1977)
14. Even, S., Goldreich, O.: On the power of cascade ciphers. ACM Trans. Comput. Syst. 3(2), 108–116 (1985)
15. Even, S., Mansour, Y.: A construction of a cipher from a single pseudorandom permutation. In: Journal of Cryptology, pp. 151–161. Springer, Heidelberg (1991)
16. Gaži, P., Maurer, U.: Cascade Encryption Revisited. In: Matsui, M. (ed.) ASIACRYPT 2009. LNCS, vol. 5912, pp. 37–51. Springer, Heidelberg (2009)
17. Gaži, P., Maurer, U.: Free-Start Distinguishing: Combining Two Types of Indistinguishability Amplification. In: Kurosawa, K. (ed.) ICITS 2009. LNCS, vol. 5973, pp. 28–44. Springer, Heidelberg (2010)
18. Kilian, J., Rogaway, P.: How to Protect DES Against Exhaustive Key Search (an Analysis of DESX). Journal of Cryptology 14, 17–35 (2001)
19. Lai, X., Massey, J.L.: A Proposal for a New Block Encryption Standard. In: Damgård, I.B. (ed.) EUROCRYPT 1990. LNCS, vol. 473, pp. 389–404. Springer, Heidelberg (1991)
20. Luby, M., Rackoff, C.: Pseudo-random permutation generators and cryptographic composition. In: STOC 1986: Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing, pp. 356–363 (1986)

21. Lucks, S.: Attacking Triple Encryption. In: Vaudenay, S. (ed.) FSE 1998. LNCS, vol. 1372, pp. 239–253. Springer, Heidelberg (1998)
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.: Composition of Random Systems: When Two Weak Make One Strong. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 410–427. Springer, Heidelberg (2004)
25. Maurer, U., Pietrzak, K., Renner, R.: Indistinguishability Amplification. In: Menezes, A. (ed.) CRYPTO 2007. LNCS, vol. 4622, pp. 130–149. Springer, Heidelberg (2007)
26. 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)
27. Rogaway, P., Steinberger, J.P.: Security/Efficiency Tradeoffs for Permutation-Based Hashing. In: Smart, N.P. (ed.) EUROCRYPT 2008. LNCS, vol. 4965, pp. 220–236. Springer, Heidelberg (2008)
28. Schneier, B.: Description of a New Variable-Length Key, 64-bit Block Cipher (Blowfish). In: Anderson, R. (ed.) FSE 1993. LNCS, vol. 809, pp. 191–204. Springer, Heidelberg (1994)
29. Stam, M.: Beyond Uniformity: Better Security/Efficiency Tradeoffs for Compression Functions. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 397–412. Springer, Heidelberg (2008)
30. Steinberger, J.P.: Stam's Collision Resistance Conjecture. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 597–615. Springer, Heidelberg (2010)
31. Tessaro, S.: Security Amplification for the Cascade of Arbitrarily Weak PRPs: Tight Bounds via the Interactive Hardcore Lemma. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 37–54. Springer, Heidelberg (2011)
32. Vaudenay, S.: Decorrelation: a theory for block cipher security. Journal of Cryptology 16(4), 249–286 (2003)