The Security of Multiple Encryption in the Ideal Cipher Model

Yuanxi Dai[1], Jooyoung Lee[2], Bart Mennink[3], and John Steinberger[1]

[1] Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing
shustdc@gmail.com, jpsteinb@gmail.com
[2] Faculty of Mathematics and Statistics, Sejong University, Seoul, Korea
jlee05@sejong.ac.kr
[3] Dept. Electrical Engineering, ESAT/COSIC, KU Leuven, and iMinds, Belgium
bart.mennink@esat.kuleuven.be

Abstract. Multiple encryption—the practice of composing a blockcipher several times with itself under independent keys—has received considerable attention of late from the standpoint of provable security. Despite these efforts proving definitive security bounds (i.e., with matching attacks) has remained elusive even for the special case of triple encryption. In this paper we close the gap by improving both the best known attacks and best known provable security, so that both bounds match. Our results apply for arbitrary number of rounds and show that the security of ℓ-round multiple encryption is precisely $\exp(\kappa + \min\{\kappa(\ell' - 2)/2), n(\ell' - 2)/\ell'\})$ where $\exp(t) = 2^t$ and where $\ell' = 2\lceil \ell/2 \rceil$ is the smallest even integer greater than or equal to ℓ, for all $\ell \geq 1$. Our technique is based on Patarin's H-coefficient method and relies on a combinatorial result of Chen and Steinberger originally required in the context of key-alternating ciphers.[1]

1 Introduction

Let $E : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ be a blockcipher with key space $\{0,1\}^\kappa$ and message/ciphertext space $\{0,1\}^n$. The *ℓ-cascade of E*, denoted $E^{(\ell)}$, is the blockcipher of key space $\{0,1\}^{\ell\kappa}$ and of message space $\{0,1\}^n$ obtained by composing E ℓ times with itself under independent keys. Thus

$$E_k^{(\ell)}(x) = E_{k_\ell}(E_{k_{\ell-1}}(\dots(E_{k_1}(x))\dots)) \qquad (1)$$

where $k = k_1 \| \dots \| k_\ell \in \{0,1\}^{\ell\kappa}$. (The inverse of $E^{(\ell)}$ is computed the obvious way.) In particular $E^{(1)} = E$.

Since $E^{(\ell)}$ has longer keys than E for $\ell \geq 2$, the ℓ-cascade can be viewed as a natural mechanism for increasing the key space of a blockcipher and, hence, potentially, enhancing the security level. Security does not necessarily increase linearly with the key length, however. For example there exist meet-in-the-middle

[1] This paper is an independently initiated merge of preprints [9, 23, 30], that were separately submitted to CRYPTO 2014.

(key-recovery) attacks against cascades of length 2 that cost no more[2] than generic (key-recovery) attacks against cascades of length 1 [11]. Indeed, when a variant of DES with longer keys was needed, designers eschewed double encryption (cascades of length 2) in favor of triple encryption [11, 31]. The standard which eventually resulted, so-called Triple DES [2,15,35], is still widely deployed.

Even while generic attacks have guided the considerations of designers since the beginning, finding nontrivial provable security results for multiple encryption in idealized models remained an open problem for a very long time. In the ideal model which we and most previous authors envisage [1, 4, 16, 17, 22] the security of the ℓ-cascade is quantified by the information-theoretic indistinguishability of two worlds, "real" and "ideal". In the "real" world the adversary A is given oracle access to an ideal[3] cipher E, to its inverse E^{-1}, and to a randomly keyed ℓ-cascade instance $E_k^{(\ell)}$ of E (for hidden k) as well as to the inverse $(E_k^{(\ell)})^{-1}$ of the ℓ-cascade; in the "ideal" world the ℓ-cascade instance $E_k^{(\ell)}$ is replaced by a random independent permutation π and its inverse π^{-1}. The adversary knows the value ℓ in question.

The case $\ell = 1$, while quite simple, is already instructive to analyze. In that case the adversary must distinguish between $E_k^{(1)} = E_k$ and a random permutation π, while being given oracle access to E. Since E is ideal, it is easy to argue that the adversary has no advantage as long as it has not queried its oracle E on key k. With k being uniform at random, and with other queries to $E/\pi/E_k$ giving no clue as to the value of k, the adversary's distinguishing advantage is thus upper bounded by—and in fact basically equal to—$q/2^\kappa$, where q is the number of queries made. (We note this bound holds even if n is very small compared to κ, e.g., $n = 1, 2$. For the sake of completeness, we formalize the argument just sketched in Appendix C of our full version [10].) An easy reduction[4] argument, moreover, shows that $E^{(\ell)}$ is at least as secure as $E^{(r)}$ for all $r \leq \ell$. Hence $E^{(\ell)}$ achieves *at least* κ bits of security for all $\ell \geq 1$, and the basic question is to determine how security grows with ℓ.

The first nontrivial results obtained pertaining to this question were by Aiello et al. [1] who show that $E_k^{(2)}$ is *slightly* harder to distinguish from a random π than $E_k^{(1)} = E_k$. More precisely, Aiello et al. show that A's distinguishing advantage for $E^{(2)}$ is upper bounded by an expression of the form $q^2/2^{2\kappa}$, as opposed to $q/2^\kappa$ for $E^{(1)}$, where q is the number of queries made by A. In either event, thus, $E^{(1)}$ and $E^{(2)}$ both essentially offer κ bits of security, given the meet-in-the-middle attack for length two cascades of cost $q = 2^\kappa$ [11]. (See also the full version of this paper [10], which revisits Aiello et al.'s result.)

Subsequently we will write $\exp(\kappa)$ for 2^κ, somewhat in line with the computer science convention of writing $\log(t)$ for $\log_2(t)$. We thus say, e.g., that $E^{(1)}$ and $E^{(2)}$ "achieve security $\exp(\kappa)$", in the sense that it requires about $\exp(\kappa) = 2^\kappa$

[2] This should be qualified: the memory costs are much larger and the query complexity is *slightly* greater [1].

[3] I.e., $E(k, \cdot) : \{0, 1\}^n \to \{0, 1\}^n$ is a random permutation for each key $k \in \{0, 1\}^\kappa$.

[4] Since the adversaries considered are information-theoretic, we note that we don't even have to consider the reduction's running time lossiness.

Table 1. Security lower and upper bounds for cascaded encryption (in log). Here, $\ell' = 2\lceil \ell/2 \rceil$. All results in **bold** are derived in this work.

$E^{(\ell)}$	security	tight
$\ell = 1, 2$	κ [1, 11]	✓
$\ell = 3, 4$	$\kappa + \min\{\kappa/2, n/2\}$ [4, 17]	✗
	$\boldsymbol{\kappa + \min\{\kappa, n/2\}}$	✓
$\ell \geq 5$	$\kappa + \min\{\kappa(\ell' - 2)/\ell', n/2\}$ [17][⋆]	✗
	$\boldsymbol{\kappa + \min\{\kappa(\ell' - 2)/2, n(\ell' - 2)/\ell'\}}$	✓

[⋆] Starting from $\ell \geq 16$, Lee [22] proved an improved security bound of $\exp(\kappa + \min\{\kappa, n\} - 8n/\ell)$.

queries to achieve constant distinguishing advantage between the real and ideal worlds for those cascade lengths.

After Aiello et al. a complicated history of improved security bounds ensues, including work by Bellare and Rogaway [4] for length 3 cascades, by Gaži and Maurer [17] (who corrected some errors in Bellare-Rogaway and who generalized their approach to larger numbers of rounds), and by Lee [22]. For reasons of space, however, we eschew a detailed discussion of these prior results in this proceedings version, and refer the reader to the synopsis in Table 1.

On the attack side Lucks [26] found an attack of cost $\kappa + n/2$ for length 3 cascades (thus matching the Bellare-Rogaway security bound for length 3 cascades in the regime $\kappa \geq n$). Gaži found an attack of cost $\kappa + n(\ell' - 2)/\ell'$ for arbitrary ℓ generalizing Lucks's attack. (Moreover Gaži was the first to give a mathematically rigorous analysis of Lucks's attack.)

Despite this series of results obtaining matching upper and lower bounds on security has remained elusive for all $\ell \geq 3$. In the case $\ell = 3$, for example, all we know is that the security of $E^{(3)}$ lies somewhere in the interval

$$[\exp(\kappa + \min\{\kappa/2, n/2\}), \exp(\kappa + n/2)]$$

which leaves open the question of the true security for $\kappa < n$. For $\ell \geq 5$, moreover, exact security remained open regardless of the ratio between ℓ and κ.

Our results. In this paper we close the remaining gaps between upper and lower bounds for all ℓ, up to customary lower-order terms. More precisely, we show that $E^{(\ell)}$ has security

$$\exp(\kappa + \min\{\kappa(\ell' - 2)/2, n(\ell' - 2)/\ell'\}) \qquad (2)$$

by exhibiting matching attacks and security proofs, for all $\ell \geq 1$. (Note by the form of (2) that new attacks are only needed when $\kappa(\ell' - 2)/2 < n(\ell' - 2)/\ell'$; otherwise the attacks of Gaži suffice.) One can observe from (2) that $\ell = 2r$ rounds buy the same amount of security as $\ell = 2r - 1$ rounds. In fact, we expect the curve describing the adversary's advantage to be slightly more advantageous for $2r-1$ rounds than for $2r$ rounds, as observed by Aiello et al. for $r=1$, but our analysis is not fine-grained enough to verify this.

Techniques. Tightening the security bounds for triple encryption is already an interesting problem in itself. Besides devising a new rather easy attack of cost $\exp(2\kappa)$, it turns out that the bound directly follows from tightening a key combinatorial lemma in Bellare and Rogaway's original proof (Lemma 10 in [5]).

We found the case of larger number of rounds (in particular, $\ell \geq 5$) to be more challenging. While we copied the basic approach of Bellare and Rogaway [4] and of Gaži and Maurer [17] some significant structural changes were required in order to achieve tightness. In particular, we had to rebundle a key two-step game transition from [17] into a single-step transition. Moreover we found that the best way to handle this (now rather delicate) single-step transition was by Patarin's H-coefficient technique [37]. Here we drew inspiration from Chen and Steinberger [8] and, indeed, reused the key combinatorial lemma of that paper. Roughly speaking, this lemma gives an explicit expression for the probability that

$$(P_\ell \circ \cdots \circ P_1)(a) = b$$

where each P_i is a *partially defined* random permutation of $\{0,1\}^n$, where $\circ$ denotes function composition, where $a, b \in \{0,1\}^n$ are two values such that $P_1(a)$ and $P_\ell^{-1}(b)$ are undefined. Here the probability is expressed (in particular, lower-bounded) as a function of the number of edges[5] already defined in the P_i's as well as of the number of "chains" of various lengths[6] formed by those edges in the composition $P_1 \circ \cdots \circ P_\ell$. (In our case $P_i = E_{k_i}$ where $k = k_1 \| \ldots \| k_\ell$ is the secret key.) It is noteworthy that the security proofs for three different classes of composed ciphers (key-alternating ciphers [8], cascade ciphers (this paper), and XOR-cascade ciphers [8,16,18]) now rely on this lemma.

In order to successfully apply the H-coefficient technique and Chen and Steinberger's lemma a crucial step is to upper bound the probability of the adversary obtaining (too many) long chains in $P_\ell \circ \cdots \circ P_1 = E_{k_\ell} \circ \cdots \circ E_{k_1}$. Like Bellare and Rogaway [4] and like Gaži and Maurer [17] before us, we do this by upper bounding the *total* number of query chains of a given length formed by *all* of the adversary's queries to E, regardless of the underlying key, and then by applying a Markov inequality—but in our case we strive for tight bounds on the total number of query chains. At first glance the combinatorial question is nonobvious (especially given the presence of an adaptive adversary) but we observe that on any path of queries at least half the queries are "backwards" (meaning contrary to the path's direction, in this instance) *for at least one of the two possible ways of orienting the path* (as a given path can be traversed right-to-left or left-to-

[5] If $x \in \{0,1\}^n$ is a value such that $y = P_i(x)$ is defined, then the pair (x, y) is also called an *edge* of P_i, equating P_i with a bipartite graph (more precisely, a partial matching) from $\{0,1\}^n$ to $\{0,1\}^n$. The composition $P_\ell \circ \cdots \circ P_1$ is visualized by "gluing" these bipartite graphs sequentially next to one another.

[6] See the previous footnote.

right). Together with some classical balls-in-bins occupancy results, this simple symmetry-breaking observation gives an easy means of upper bounding the total number of query chains formed, and the bounds obtained are also tight. We refer to Proposition 1 for more details.

Other related work. We have already briefly mentioned related work on key-alternating ciphers [7,8,14,21,38] as well as on XOR cascades [16,18,22], to which the beautiful work of Rogaway and Kilian on DESX (a special case of an XOR-cascade) should be added [19].

Coming back to cascade ciphers, Merkle and Hellman [31] show an attack on two-key triple encryption, which attack is revisited by Oorschot and Wiener [34]. (See also [33].) Even and Goldreich [13] present a medley of observations on multiple encryption in various models, including some conclusions which are disputed by Maurer and Massey [27]. The best paper award at CRYPTO 2012, by Dinur et al. [12], concerns, in large part, non-information-theoretic key-recovery attacks on cascade ciphers.

We finally point that similar questions (though using very different techniques) have been pursued in the computational setting, in which one seeks to amplify the *computational* indistinguishability of a PRP by composing it with itself [25,28,29,32]. See in particular [39] which culminates this line of work.

Open questions. As will be seen, our results actually hold even if the adversary is always allowed to make 2^n queries to its permutation oracle (which is $E_k^{(\ell)}$ or π) for free, i.e., to entirely learn its permutation oracle for free. It would be interesting to know if better bounds can be achieved by restricting the number of permutation queries. This is all the more relevant given that many applications will impose limitations on the number of encryptions/decryptions available to the adversary.

2 Definitions

Blockciphers and Cascades. A blockcipher is a function $E : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ such that $E(k, \cdot) : \{0,1\}^n \to \{0,1\}^n$ is a permutation for each key $k \in \{0,1\}^\kappa$. We also write $E_k(x)$ for $E(k,x)$. By the "inverse" E^{-1} of E we mean the blockcipher $E^{-1} : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ such that E_k^{-1} is the inverse permutation of E_k for each $k \in \{0,1\}^\kappa$.

For a blockcipher E and an integer $\ell \geq 1$ we define the *ℓ-cascade* of E, written $E^{(\ell)}$, by equation (1). We note that $E^{(\ell)}$ is a blockcipher of key space $\{0,1\}^{\ell\kappa}$ and of message space $\{0,1\}^n$.

Ideal Ciphers. A blockcipher $E : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$ which is sampled uniformly at random from the space of all blockciphers of key space $\{0,1\}^\kappa$ and of message space $\{0,1\}^n$ is called an *ideal cipher*. In this case E_k is a random independent permutation of $\{0,1\}^n$ for each $k \in \{0,1\}^\kappa$.

Security Game. Let ℓ, κ and n be given. Let A be an information-theoretic adversary (or "distinguisher") with oracle access to, among others, an ideal cipher

$E : \{0,1\}^\kappa \times \{0,1\}^n \to \{0,1\}^n$, which we write A^E but by which we mean that A can query *both* E and E^{-1}. (Along the same lines writing A^π indicates that A has access to both π and π^{-1} when π is a permutation.) Then A's *distinguishing advantage* against ℓ-cascades, written $\mathbf{Adv}^{\mathsf{casc}}_{\ell,\kappa,n}(A)$ is defined as

$$\mathbf{Adv}^{\mathsf{casc}}_{\ell,\kappa,n}(A) = \Pr[k^* \xleftarrow{\$} \{0,1\}^{\ell\kappa}; A^{E,E^{(\ell)}_{k^*}} = 1] - \Pr[\pi \xleftarrow{\$} \mathcal{P}; A^{E,\pi} = 1]$$

where the notation

$$k^* \xleftarrow{\$} \{0,1\}^{\ell\kappa}; A^{E,E^{(\ell)}_{k^*}} = 1$$

indicates the event that A outputs 1 after interacting with oracles E/E^{-1} and $E^{(\ell)}_{k^*}/(E^{(\ell)}_{k^*})^{-1}$ where k^* is sampled uniformly at random from the key space of $E^{(\ell)}$, and hidden from A; whereas the notation

$$\pi \xleftarrow{\$} \mathcal{P}; A^{E,\pi} = 1$$

indicates the event that A outputs 1 after interacting with oracles E/E^{-1} and π/π^{-1} where π is a permutation of $\{0,1\}^n$ sampled uniformly at random from the set of all permutations of $\{0,1\}^n$, here denoted $\mathcal{P}$; and where in either case the sampling of the ideal cipher E at the start of the experiment is kept implicit for the sake of succinctness.

We write

$$\mathbf{Adv}^{\mathsf{casc}}_{\ell,\kappa,n}(q)$$

for the supremum of $\mathbf{Adv}^{\mathsf{casc}}_{\ell,\kappa,n}(A)$ taken over all q-query information-theoretic adversaries A. (The notation $\mathbf{Adv}^{\mathsf{casc}}_{\ell,\kappa,n}$ is thus overloaded.)

3 Statement of Results

Lower bounds. Our paper's main result is the following theorem (as always, $\ell' = 2\lceil \ell/2 \rceil$; we also write $(\ell+1)'$ for $2\lceil(\ell+1)/2\rceil$, etc):

Theorem 1. *(a) If $q \geq 2^n$ then, for every real number $C \geq 1$,*

$$\mathbf{Adv}^{\mathsf{casc}}_{\ell,\kappa,n}(q) \leq \frac{\ell^2}{2^{\kappa+1}} + \frac{4}{2^n} + \frac{\alpha}{C} + 2^n \ell C^{(\ell+1)'/2} \left(\frac{8q}{2^{\kappa+n}}\right)^{\ell'/2}$$

where $\alpha = \ell^2 2^\ell (7n)^{\ell'/2}$. Furthermore if $q \geq n2^n$ we can improve α to $\alpha' = \ell^2 2^\ell 14^{\ell'/2} \leq \ell^2 8^{\ell'}$.

(b) If $q \leq 2^n$ then, for every $C \geq 1$ such that $Cq < 2^{\kappa+n-2}$,

$$\mathbf{Adv}^{\mathsf{casc}}_{\ell,\kappa,n}(q) \leq \frac{\ell^2}{2^{\kappa+1}} + \frac{4}{2^n} + \frac{\beta}{C} + \frac{q^2 \ell}{2^n} C^{(\ell+1)'/2} \left(\frac{8}{2^\kappa}\right)^{\ell'/2} + \frac{q\beta}{\ell 2^{\kappa\ell'/2}}$$

where $\beta = \ell^2 2^\ell (3 \log q + 2)^{\ell'/2}$.

Moreover (a) and (b) also hold if the adversary is allowed to ask, for free, all possible 2^n queries to its second oracle.

The presence of the adjustable constant C is typical of security proofs involving a threshold-based "bad event". For given parameters q, n, κ and ℓ there some optimal C that minimizes the bound.

Theorem 1 is, unfortunately and evidently, hard to parse. By analytically optimizing C and making a few other simplifications, however, Theorem 1 yields the following, slightly more digestible corollary:

Corollary 1. (*a*) *If* $q \geq 2^n$ *then*

$$\mathbf{Adv}^{\mathsf{casc}}_{\ell,\kappa,n}(q) \leq \frac{\ell^2}{2^{\kappa+1}} + \frac{4}{2^n} + \alpha(\ell/2+2)\ell^{1/2}\left(\frac{8q}{2^{\kappa+n(\ell'-2)/\ell'}}\right)^{\ell'/(\ell+3)'}$$

where $\alpha = \ell^2 2^\ell (7n)^{\ell'/2}$. *Furthermore if* $q \geq n2^n$ *we can improve* α *to* $\alpha' = \ell^2 2^\ell 14^{\ell'/2} \leq \ell^2 8^{\ell'}$.

(*b*) *If* $q \leq 2^n$ *and* $2^\ell(3n+2)^{\ell'/2} \leq 2^n$ *then*

$$\mathbf{Adv}^{\mathsf{casc}}_{\ell,\kappa,n}(q) \leq \frac{\ell^2}{2^{\kappa+1}} + \frac{4}{2^n} + \beta(\ell/2+2)\left(\frac{\ell 3^{\ell'} q^2}{2^{\kappa\ell'/2+n}}\right)^{2/(\ell+3)'} + \frac{q\beta}{\ell 2^{\kappa\ell'/2}}$$

where $\beta = \ell^2 2^\ell (3\log q + 2)^{\ell'/2}$.

Moreover (a) and (b) also hold if the adversary is allowed to ask, for free, all possible 2^n *queries to its second oracle.*

The proof of Corollary 1 from Theorem 1 can be found in the full version [10].

We note the constraint $2^\ell(3n+2)^{\ell'/2} \leq 2^n$ that appears in the second part of Corollary 1 is almost always satisfied by practical parameters and is always asymptotically verified as $n \to \infty$. (Indeed, we imagine ℓ as fixed whereas $n, \kappa \to \infty$ according to some fixed ratio.)

It directly follows from Corollary 1 that $\mathbf{Adv}^{\mathsf{casc}}_{\ell,\kappa,n}(q)$ is small if

$$q \ll \exp(\kappa + \min\{\kappa(\ell'-2)/2, n(\ell'-2)/\ell'\})$$

(note $\kappa + \kappa(\ell'-2)/2 = \kappa\ell'/2$ and $q^2/2^{\kappa\ell'/2+n} \leq q/2^{\kappa\ell'/2}$ when $q \leq 2^n$) or, a little more precisely, if

$$\begin{aligned} q \ll\ & (2^{-\ell/2}(7n)^{-\ell'/4}\ell^{-2})^{(\ell+3)'} \\ & \cdot \exp(\kappa + \min\{\kappa(\ell'-2)/2 - 2\ell, n(\ell'-2)/\ell' - 3\}). \end{aligned} \qquad (3)$$

We emphasize that the above threshold is a coarse estimate, which takes into account the factors of all three non-negligible expressions in Corollary 1. (Note that $\log q \leq n$ in the second part of Corollary 1, so $\beta \leq \alpha$.) Indeed, if q is a factor r smaller than the expression on the right of (3), then it is easy to see from Corollary 1 that the adversary's advantage is upper bounded by either $r^{\ell'/(\ell+3)'}$ or $r^{4/(\ell+3)'} + r$, disregarding the negligible terms $\ell^2/2^{\kappa+1}$ and $4/2^n$.

Upper bounds. In Section 4 we present a simple attack of query complexity

$$\ell \cdot \exp(\kappa\ell'/2)$$

that succeeds in distinguishing $(E, E_k^{(\ell)})$ from (E, π) with overwhelming advantage. This complements the previously quoted attack by Gaži, of query complexity

$$\ell \cdot \exp(\kappa + n(\ell' - 2)/\ell')$$

and which also succeeds with overwhelming advantage. Hence the gap left between lower and upper bounds is essentially the gap left between

$$\min\{\ell \cdot \exp(\kappa\ell'/2), \ell \cdot \exp(\kappa + n(\ell' - 2)/\ell')\}$$

and the right-hand side of (3).

4 An Attack of Cost $\exp(\kappa\ell'/2)$

In this section we describe a new "meet-in-the-middle-attack" on $E^{(\ell)}$ of complexity $\exp(\kappa\ell'/2)$, which complements Gaži's attack of query complexity $\exp(\kappa + n(\ell' - 2)/\ell')$. A precise statement is given by the following theorem.

Theorem 2. *For any integer ρ, $1 \leq \rho \leq 2^{n-1}$, there exists an adversary A making at most $\rho\ell 2^{\kappa\ell'/2}$ queries to E and at most ρ queries to $E_k^{(\ell)}/\pi$, such that*

$$\mathbf{Adv}^{\mathsf{casc}}_{\ell,\kappa,n}(A) \geq 1 - 2^{\kappa\ell - \rho(n-1)}.$$

Proof. The adversary A, which implements a meet-in-the-middle attack, is given by the pseudocode of Fig. 1. A starts by querying ρ messages $m_1, \ldots, m_\rho$ to $E_k^{(\ell)}/\pi$, thus obtaining their corresponding ciphertexts c_1, …, c_ρ. Then for each of these message/ciphertext pairs (m_i, c_i) it evaluates the first $\lceil \ell/2 \rceil$ block ciphers for all possible keys starting from m_i and the last $\lfloor \ell/2 \rfloor$ block ciphers in inverse direction starting from c_i. One possible key $k = (k_1 \| \ldots \| k_\ell)$ must "stand out" unless A is in the ideal world. Thus A returns 1 if and only if there is a key compatible with all ρ message-ciphertext pairs (m_i, c_i). It is easy to see that A makes ρ queries to $E_k^{(\ell)}/\pi$ and

$$\rho(\lceil \ell/2 \rceil 2^{\kappa\lceil \ell/2 \rceil} + \lfloor \ell/2 \rfloor 2^{\kappa\lfloor \ell/2 \rfloor}) \leq \rho\ell 2^{\kappa\ell'/2}$$

queries to E, as claimed.

Clearly, in the real world $(E_k^{(\ell)}, E)$, for $(k_L^*, k_R^*) = k$ we have $a_{i,k_L^*} = b_{i,k_R^*}$ for all $i = 1, \ldots, \rho$, so A returns 1. We consider the probability that A returns 1 in the ideal world (π, E). For each key $k = (k_1 \| \ldots \| k_\ell)$, $E_{k_\ell} \circ \cdots \circ E_{k_1}$ becomes a truly random permutation, independent of π. For this key, the probability that $E_{k_\ell} \circ \cdots \circ E_{k_1}(m_i) = c_i$ for every $i = 1, \ldots, \rho$ is upper bounded by

$$\frac{(2^n - \rho)!}{(2^n)!} \leq \left(\frac{1}{2^n - \rho + 1}\right)^\rho \leq \frac{1}{2^{\rho(n-1)}}.$$

The theorem follows by a union bound over all possible keys. □

```
**fix distinct** $m_1, \ldots, m_\rho \in \{0,1\}^n$
**for** $r = 1$ **to** $\rho$ **do**
    **query** $c_i \leftarrow \mathcal{R}(m_i)$
    **forall** $k_L^* = (k_1^* \| \ldots \| k_{\lceil \ell/2 \rceil}^*) \in \{0,1\}^{\kappa \lceil \ell/2 \rceil}$
        **query** $a_{i,k_L^*} \leftarrow E_{k_{\lceil \ell/2 \rceil}^*} \circ \cdots \circ E_{k_1^*}(m_i)$
    **forall** $k_R^* = (k_{\lceil \ell/2 \rceil + 1}^* \| \ldots \| k_\ell^*) \in \{0,1\}^{\kappa \lfloor \ell/2 \rfloor}$
        **query** $b_{i,k_R^*} \leftarrow E_{k_{\lceil \ell/2 \rceil + 1}^*}^{-1} \circ \cdots \circ E_{k_\ell^*}^{-1}(c_i)$
**forall** $(k_L^*, k_R^*) \in \{0,1\}^{\kappa \lceil \ell/2 \rceil} \times \{0,1\}^{\kappa \lfloor \ell/2 \rfloor}$
    **if** $a_{i,k_L^*} = b_{i,k_R^*}$ **for all** $i = 1, \ldots, \rho$
        **return** $1$
**return** $0$
```

Fig. 1. The adversary A for Theorem 2. The oracle to $E_k^{(\ell)}/\pi$ is denoted $\mathcal{R}$.

5 Preliminary Reductions and Proof Overview

Modifications of Bellare and Rogaway [4]. In view of proving Theorem 1, we start by modifying the distinguishability game in the following way. At the very start of the experiment we send a symbol $\star \in \{\bot, \top\}$ to the adversary. In the ideal world we send $\star = \top$, and in the real world we also send $\star = \top$ unless $k_\ell^* = k_i^*$ for some $i < \ell$, where $k^* = k_1^* \| \ldots \| k_\ell^*$ is the secret key, in which case we send $\star = \bot$. Since the adversary is free to disregard $\star$, this modification is without loss of generality.

Next, we make a second modification, namely that if $\star = \bot$ then we forbid the adversary from making any queries. Since $\star$ can only be $\bot$ in the real world this is without loss of generality either (as the adversary already knows which world it is in anyway).

Now we make yet another modification to the real world, by generating a random permutation π like in the ideal world at the beginning of the experiment. If $\star = \top$ we answer queries to $E_{k^*}^{(\ell)}$ by π instead and, to compensate, we define $E_{k_\ell^*} = \pi \circ E_{k_1^*}^{-1} \circ \cdots \circ E_{k_{\ell-1}^*}^{-1}$ (thus "overwriting" $E_{k_\ell^*}$). Since this simply trades the randomness of $E_{k_\ell^*}$ for the randomness in π, it is easy to see that this is an equivalent way of defining the real world.

Note that both worlds now involve an independent[7] random permutation π. For each fixed permutation S one can also consider the distinguishing experiment where π is set to S in each world. A simple averaging argument over π shows, moreover, that there must exist some S for which the adversary's distinguishing advantage is at least as great when π is fixed to S as when π is random. We can thus assume without loss of generality that π is not sampled at random, but set

[7] The real world now has three "random tapes": one for k^*, one for π, and one for the ideal cipher E. Every query made by the adversary is deterministically answered as a function of these three random tapes, and these random tapes are independently sampled. This is the sense in which π is "independent" from other randomness in the real world.

to the same fixed permutation S in both worlds. Since S is fixed, now, and since we are quantifying over all information-theoretic adversaries A, we can assume that A knows S and, hence, makes no queries to its second oracle.

To summarize, modifications so far amount to this: in the real world, we abort the experiment if $k^*_\ell = k^*_i$ for some $i < \ell$, whereas in the contrary (generic) case there is some fixed permutation S, known to the adversary, such that $E_{k^*_\ell} = S \circ E^{-1}_{k^*_1} \circ \cdots \circ E^{-1}_{k^*_{\ell-1}}$. The ideal world never aborts.

Further Normalizations. Since A is information-theoretic we can assume without loss of generality that A is *deterministic*.

As in [8] we will also modify the experiment by *giving the secret key to A after it has finished making all its queries*. More precisely, in the real world we give the "real" key k^* used to key the second oracle $E^{(\ell)}_{k^*}$ whereas in the ideal world (where no such key exists) we sample a "dummy" key $k^* \in \{0,1\}^{\kappa\ell}$ uniformly at random and give this dummy key to A. Since A is free to disregard this extra information this is also without loss of generality.

Transcripts. The interaction of A with its oracles is encoded by a *transcript* which, basically, is a list of questions asked and answers received, together also with the key value received at the end of the experiment.

More precisely, a transcript can be encoded by a triple of the form $(\star, Q_E, k^*)$ where $\star \in \{\bot, \top\}$, where $k^* \in \{0,1\}^{\kappa\ell}$ is the final key value received, and where Q_E is an *unordered* set of triples of the form $(k, x, y) \in \{0,1\}^\kappa \times \{0,1\}^n \times \{0,1\}^n$ with each such tuple indicating that either $E(k, x)$ was queried with answer y or that $E^{-1}(k, y)$ was queried with answer x. Indeed, A's interaction with its oracles can be unambiguously reconstructed from such an "unordered and undirected" set Q_E by using the fact that A is deterministic, cf. [8].

We write $\mathcal{T}$ for the set of all possible transcripts.

Probability space of Oracles. Let $\mathcal{P}$ be the set of all permutations from $\{0,1\}^n$ to $\{0,1\}^n$. Then a blockcipher of key space $\{0,1\}^\kappa$ and message space $\{0,1\}^n$ can be viewed as an element of $\mathcal{P}^{\exp(\kappa)}$ (2^κ-fold direct product). Thus, an ordered pair

$$(E', k^*) \in \mathcal{P}^{\exp(\kappa)} \times \{0,1\}^{\kappa\ell}$$

uniquely determines a real-world environment for A. More precisely, unless $\star = \bot$ in which case A receives no further information except for k^*, A's ideal cipher oracle E is defined by

$$E_k = \begin{cases} E'_k & \text{if } k \neq k^*_\ell \\ S \circ {E'}^{-1}_{k^*_1} \circ \cdots \circ {E'}^{-1}_{k^*_{\ell-1}} & \text{if } k = k^*_\ell \end{cases}$$

where $k^* = k^*_1 \| \ldots \| k^*_\ell$. We thus identify elements of

$$\Omega_X := \mathcal{P}^{\exp(\kappa)} \times \{0,1\}^{\kappa\ell}$$

with real-world oracles. We view Ω_X as a probability space with uniform measure (indeed, the definition of the real-world experiment induces uniform measure on Ω_X).

We similarly define

$$\Omega_Y := \mathcal{P}^{\exp(\kappa)} \times \{0,1\}^{\kappa\ell}$$

to be identified with the set of all ideal-world oracles, and which we also view as a probability space with uniform measure. Here the last coordinate corresponds to the "dummy key" given to the adversary at the end of the experiment. We emphasize that, for $(E, k^*) \in \Omega_Y$, the ideal cipher oracle to which A has access is precisely E, i.e., with no key being overwritten as a function of k^* and S; this is precisely the difference between the real and ideal worlds in the (generic) case when $k^*_\ell \notin \{k^*_1, \ldots, k^*_{\ell-1}\}$.

We can view the transcript produced by A in the real world as a random variable defined over Ω_X. Formally, let $X : \Omega_X \to \mathcal{T}$ be the function defined by letting $X(\omega)$ be the transcript obtained by running A on oracle ω. Thus X is a random variable of range $\mathcal{T}$, and the distribution of X is exactly the distribution of transcripts in the real world. We similarly define $Y : \Omega_Y \to \mathcal{T}$, so that Y is the transcript distribution in the ideal world.

The H-coefficient technique [36,37], in its simplest form, states that if we can divide $\mathcal{T}$ into a set of (so-called) "good" transcripts $\mathcal{T}_1$ and (so-called) "bad" transcripts $\mathcal{T}_2$, such that[8]

$$\frac{\Pr[X = \tau]}{\Pr[Y = \tau]} \geq 1 - \varepsilon_1 \tag{4}$$

for some $\varepsilon_1 > 0$ and for all $\tau \in \mathcal{T}_1$, then the adversary's distinguishing advantage is upper bounded by

$$\Pr[Y \in \mathcal{T}_2] + \varepsilon_1.$$

We refer to [8] for more details.

Computing transcript probabilities. Another key insight of the H-coefficient technique is that the probability of obtaining a transcript in either world can be computed via the formulas

$$\Pr[X = \tau] = \frac{|\mathsf{comp}_X(\tau)|}{|\Omega_X|}, \qquad \Pr[Y = \tau] = \frac{|\mathsf{comp}_Y(\tau)|}{|\Omega_Y|} \tag{5}$$

as long as $\Pr[Y = \tau] > 0$, and where $\mathsf{comp}_X(\tau) \subseteq \Omega_X$ (resp. $\mathsf{comp}_Y(\tau) \subseteq \Omega_Y$) is the set of real-world (resp. ideal-world) oracles that are compatible with a transcript τ, where "compatibility" is defined the obvious[9] way: an oracle ω is compatible with a transcript τ if each individual query in τ is compatible with ω (in particular, τ's key value should match ω's key value). See [8] and Appendix D of our full version [10] for further discussion of these identities.

[8] By convention, the ratio $\Pr[X = \tau]/\Pr[Y = \tau]$ is considered to be ∞ if $\Pr[Y = \tau] = 0$.

[9] Slightly more formally—but less intuitively—an oracle (or "environment") ω is compatible with a transcript τ if there exists *some* (wlog, deterministic) adversary A' that produces τ as transcript when given ω as oracle.

Terminology: Chains. Let $\tau = (\star, Q_E, k^*)$ be a transcript, where $k^* = k_1^* \| \ldots \| k_\ell^*$. Loosely following [17], a tuple $(h, x_h, k_{h+1}, x_{h+1}, k_{h+2}, \ldots, k_{h+r}, x_{h+r})$ where $0 \leq h \leq \ell - 1$ is an integer, where $1 \leq r \leq \ell$, and where

$$\begin{cases} (k_i, x_{i-1}, x_i) \in Q_E & \text{if } i-1 \neq \ell \\ (k_i, S^{-1}(x_{i-1}), x_i) \in Q_E & \text{if } i-1 = \ell \end{cases}$$

for $h+1 \leq i \leq h+r$ (in particular, $x_i \in \{0,1\}^n$ and $k_i \in \{0,1\}^\kappa$ for each x_i, k_i) is called an *r-chain of τ starting at index h* or simply an *r-chain of τ*. Moreover, an r-chain is said to *fit* τ if $k_{h+i} = k_{h+i}^*$ for $1 \leq i \leq r$, indices taken mod ℓ and in the range $\{1, \ldots, \ell\}$. We sometimes commit a slight abuse of language by saying that a chain "fits k^*" instead of "fits τ" when it is clear which transcript τ is intended.

By means of emphasis, a chain which doesn't (necessarily) fit the key of τ is said to be *generic*; thus all r-chains of τ are by definition generic.

The rest of the proof in a nutshell. Broadly, our "bad transcripts" are transcripts that either have a bad key (i.e., $k_i^* = k_j^*$ for some $i \neq j$) or transcripts with too many (long) fitting chains, where "too many" depends geometrically on the chain length r, as might be expected. When there are not too many long chains that fit the transcript's key, indeed, we are in a position to apply the lemma of Chen and Steinberger [8] to show that the probability of obtaining the given transcript in the real world is not far off from the probability of obtaining the same transcript in the ideal world, as required by (4).

The main technical challenge that arises is that of upper bounding the probability of obtaining too many length r chains that fit the key. Here one must emphasize that this probability (which is the probability of obtaining a "bad" transcript) is being computed in the ideal world. In the ideal world, the key value $k^* \in \{0,1\}^{\kappa\ell}$ is chosen at random *after* all queries are completed. Hence, by a Markov bound, it suffices to show that, with high probability, not too many *generic* r-chains are created by the adversary's queries. We deliver a tight bound on the number of generic chains by using a fairly simple argument, as already discussed in the paper's introduction (see in particular Proposition 1 in Section 6). See further details in Section 6.

6 Proof of Theorem 1

For the remainder of the proof of Theorem 1 we will assume that $n \geq 2$ and also, if $q \geq 2^n$, that

$$4Cq \leq 2^{\kappa+n} \qquad \text{and} \qquad C2^n \left(\frac{q}{2^{\kappa+n}}\right)^{\lceil \ell/2 \rceil} < 1. \tag{6}$$

These assumptions are without loss of generality because the first part of Theorem 1 is void otherwise, as can easily be checked. We also let $N = 2^n$.

We start by making a few more definitions that will be useful for the definition of bad transcripts and thereafter. Firstly, for a transcript $\tau = (\star, Q_E, k^*)$

we let Q_E^+, Q_E^- be the sets of queries in Q_E obtained respectively by *forward* and *backward* queries to E by the adversary. (To wit, a query to E is forward, a query to E^{-1} is backward.) We note that while Q_E does not explicitly encode forward/backward information by design, such information can be uniquely reconstructed from Q_E given the fact that A is deterministic; hence, this information is implicitly contained in Q_E.

The *maximum forward query occupancy* of τ, denoted $\mathsf{fwd}(\tau)$, is given by

$$\mathsf{fwd}(\tau) := \max_{y_0 \in \{0,1\}^n} |\{(k, x, y) \in Q_E^+ : y = y_0\}| \tag{7}$$

and $\mathsf{bwd}(\tau)$, the *maximum backward query occupancy*, is similarly given by

$$\mathsf{bwd}(\tau) = \max_{x_0 \in \{0,1\}^n} |\{(k, x, y) \in Q_E^- : x = x_0\}|.$$

We also define

$$\mathsf{fitkey}(\tau, r, h)$$

as the number of r-chains in τ that fit k^* and that start at position h.

Note that back-of-the-envelope computations suggest that $\mathsf{fwd}(\tau)$ and $\mathsf{bwd}(\tau)$ should be around q/N for $q \geq N = 2^n$ and should be around $\log(q) \leq n$ for $q \leq N$. This motivates the definition of the following threshold $\zeta(q)$:

$$\zeta(q) := \begin{cases} 3\log(q) + 2 & \text{if } q \leq N, \\ 7nq/N & \text{if } N \leq q \leq nN, \\ 14q/N & \text{if } nN \leq q. \end{cases}$$

For now, the factors $3\log(q) + 2$, $7n$ and 14 that appear in the definition of $\zeta(q)$ should be more or less ignored; these coefficients are necessary to make bad transcripts, as defined next, unlikely. (We distinguish between the cases $N \leq q \leq nN$ and $nN \leq q$ only so that we can give a slightly sharper bound in the latter case. Also we allow cases to overlap for the sake of typographical and conceptual convenience.) In fact, we find it convenient to factor $\zeta(q)$ into "essential" an "non-essential" parts $\zeta'(q)$ and $\zeta''(q)$:

$$\zeta''(q) = \begin{cases} 3\log(q) + 2 & \text{if } q \leq N, \\ 7n & \text{if } N \leq q \leq nN, \\ 14 & \text{if } nN \leq q. \end{cases} \qquad \zeta'(q) = \begin{cases} 1 & \text{if } q \leq N, \\ q/N & \text{if } q \geq N. \end{cases} \tag{8}$$

Thus $\zeta(q) = \zeta''(q)\zeta'(q)$. Note also that $\zeta(q) \leq 2^\kappa$ by the wlog assumptions made in (6).

Bad transcripts. We say that a transcript $\tau = (\star, Q_E, k^*)$ is *bad* if either (i) $k_i^* = k_j^*$ for some $i \neq j$, or (ii) $\mathsf{fwd}(\tau) \geq \zeta(q)$ or $\mathsf{bwd}(\tau) \geq \zeta(q)$, or (iii) there exists some h, $0 \leq h \leq \ell - 1$ such that

$$\mathsf{fitkey}(\tau, \ell, h) \geq 1,$$

or (iv) there exists some r, $1 \le r \le \ell$ and some h, $0 \le h \le \ell - 1$ such that

$$\mathsf{fitkey}(\tau, r, h) \ge C z_r.$$

where

$$z_r := \min\{q, N\} \cdot \left(\frac{\zeta'(q)}{2^\kappa}\right)^{\lceil r/2 \rceil}. \tag{9}$$

We let $\mathcal{T}_2$ be the set of bad transcripts, and let $\mathcal{T}_1 = \mathcal{T} \backslash \mathcal{T}_2$. One can note that every transcript with $\star = \bot$ is a bad transcript, since in that case $k^*_\ell = k^*_i$ for some $i \ne \ell$.

Bounding the probability of bad transcripts. Here we attach ourselves to upper bounding $\Pr[Y \in \mathcal{T}_2]$, as required by the H-coefficient technique. This is the probability of obtaining a bad transcript in the ideal world.

The probability that two subkeys of k^* are equal is obviously at most $\binom{\ell}{2} 2^{-\kappa} \le \ell^2 / 2^{\kappa+1}$. For the other two events we need the help of the following lemmas:

Lemma 1. *One has*

$$\Pr_{\tau \sim Y}[\mathsf{fwd}(\tau) \ge \zeta(q)] \le \frac{2}{N} \qquad \textit{and} \qquad \Pr_{\tau \sim Y}[\mathsf{bwd}(\tau) \ge \zeta(q)] \le \frac{2}{N}$$

for all q, n.

(Here $\Pr_{\tau \sim Y}$ indicates that τ is sampled according to the ideal world distribution on transcripts. The same probabilities could equivalently be written $\Pr[\mathsf{fwd}(Y) \ge \zeta(q)]$, $\Pr[\mathsf{bwd}(Y) \ge \zeta(q)]$.)

Lemma 2. *One has*

$$\Pr_{\tau \sim Y}[\mathsf{fitkey}(\tau, \ell, h) \ge 1 \wedge \mathsf{fwd}(\tau) \le \zeta(q) \wedge \mathsf{bwd}(\tau) \le \zeta(q)] \le 2^\ell \zeta''(q)^{\lceil \ell/2 \rceil} z_\ell$$

for each $0 \le h \le \ell - 1$, and

$$\Pr_{\tau \sim Y}[\mathsf{fitkey}(\tau, r, h) \ge C z_r \wedge \mathsf{fwd}(\tau) \le \zeta(q) \wedge \mathsf{bwd}(\tau) \le \zeta(q)] \le \frac{2^r \zeta''(q)^{\lceil r/2 \rceil}}{C}$$

for each $1 \le r \le \ell$, $0 \le h \le \ell - 1$ with z_r as defined in (9).

We can combine Lemmas 1 and 2 by a union bound. When $q \ge N$ condition (iii) is implied by condition (iv) since

$$C z_\ell = CN \cdot \left(\frac{q}{2^{\kappa+n}}\right)^{\lceil \ell/2 \rceil}$$

is less than 1 by (6). In this case, therefore, we don't need to incorporate the first part of Lemma 2 into the union bound. Using the fact that $r \le \ell$ and that $\zeta''(q) \ge 1$ we can upper bound $\zeta''(q)^{\lceil r/2 \rceil}$ by $\zeta''(q)^{\lceil \ell/2 \rceil}$, thus obtaining

$$\Pr[Y \in \mathcal{T}_2] \le \begin{cases} \frac{\ell^2}{2^{\kappa+1}} + \frac{4}{N} + 2^\ell \zeta''(q)^{\lceil \ell/2 \rceil} \cdot (\ell^2/C) & \text{if } q \ge N, \\ \frac{\ell^2}{2^{\kappa+1}} + \frac{4}{N} + 2^\ell \zeta''(q)^{\lceil \ell/2 \rceil} \cdot (\ell z_\ell + \ell^2/C) & \text{if } q \le N \end{cases} \tag{10}$$

since there are ℓ choices for h and ℓ^2 choices for the pair (r, h).

The proof of Lemma 1 (which involves a few subtleties because permutations "lose randomness" after $\approx 2^n$ queries) can be found in the paper's full version [10].

For Lemma 2, a key component is given by the following proposition, which happens to be a key part of our proof and which sharpens similar bounds found in [4, 17]:

Proposition 1. *Assume $\tau = (\star, Q_E, k^*)$ is a q-query transcript such that $\mathsf{fwd}(\tau) \leq \zeta(q)$, $\mathsf{bwd}(\tau) \leq \zeta(q)$. Then the total number of r-chains of τ starting at position h is at most*

$$2^r \cdot \min\{q, N\} \cdot \zeta(q)^{\lceil r/2 \rceil} 2^{\kappa \lfloor r/2 \rfloor}.$$

Proof. Let $\nu = (h, x_h, k_{h+1}, x_{h+1}, \ldots, k_{h+r}, x_{h+r})$ be an r-chain of τ. Thus either $(k_i, x_{i-1}, x_i) \in Q_E^+$ or $(k_i, x_{i-1}, x_i) \in Q_E^-$ for $h + 1 \leq i \leq h + r$. Let ν's *signature* be the string $sig^\nu \in \{+, -\}^r$ such that $(k_i, x_{i-1}, x_i) \in Q_E^{sig_i^\nu}$ for $h+1 \leq i \leq h+r$.

We start by fixing a signature $sig^0 \in \{+, -\}^r$ and by upper bounding the number of r-chains ν of τ starting at position h such that $sig^\nu = sig^0$. We can assume without loss of generality that sig^0 contains at least as many $-$'s as $+$'s, i.e., that the number of $-$'s is at least $\lceil r/2 \rceil$.

If $\nu = (h, x_h, k_{h+1}, x_{h+1}, \ldots, k_{h+r}, x_{h+r})$ is a ν-chain with signature sig^0 then there are, firstly, at most

$$\min\{q, N\}$$

choices for x_h given that $(k_{h+1}, x_h, x_{h+1}) \in Q_E$. Then, presuming x_h fixed, there are at most 2^κ choices for x_{h+1} if $sig_1^0 = +$ and at most $\zeta(q)$ choices for x_{h+1} if $sig_1^0 = -$, given that τ is a transcript such that $\mathsf{bwd}(\tau) \leq \zeta(q)$. Similarly, each subsequent step introduces a factor of either 2^κ or $\zeta(q)$ depending on the sign of that step in sig^0. Hence (and since $2^\kappa \geq \zeta(q)$) the total number of choices for $x_h, k_{h+1}, \ldots, x_{h+r}$ is at most

$$\min\{q, N\} \cdot \zeta(q)^{\lceil r/2 \rceil} 2^{\kappa \lfloor r/2 \rfloor}.$$

Multiplying by 2^r to account for all possible signatures concludes the proof. □

Proof of Lemma 2. Since $\Pr[A \wedge B] \leq \Pr[A|B]$ we have

$$\begin{aligned} &\Pr_{\tau \sim Y}[\mathsf{fitkey}(\tau, r, h) \geq T \wedge \mathsf{fwd}(\tau) \leq \zeta(q) \wedge \mathsf{bwd}(\tau) \leq \zeta(q)] \\ \leq &\Pr_{\tau \sim Y}[\mathsf{fitkey}(\tau, r, h) \geq T \mid \mathsf{fwd}(\tau) \leq \zeta(q) \wedge \mathsf{bwd}(\tau) \leq \zeta(q)] \end{aligned}$$

where $T \in \{C z_r, 1\}$ is the bound we want to prove. When we condition on $\mathsf{fwd}(\tau) \leq \zeta(q) \wedge \mathsf{bwd}(\tau) \leq \zeta(q)$, however, k^* is still independent uniformly at random (being entirely independent from Q_E in the ideal world), and so the expected number of r-chains that fit τ at position h is upper bounded by

$$2^r \cdot \min\{q, N\} \cdot \zeta(q)^{\lceil r/2 \rceil} 2^{\kappa \lfloor r/2 \rfloor} \frac{1}{2^{\kappa r}} \tag{11}$$

by Proposition 1. (Each r-chain of Q_E, indeed, has probability of exactly $1/2^{\kappa r}$ of being "hit" by k^*.) Since $r - \lfloor r/2 \rfloor = \lceil r/2 \rceil$, (11) can be written

$$2^r \zeta''(q)^{\lceil r/2 \rceil} \min\{q, N\} \left(\frac{\zeta'(q)}{2^\kappa}\right)^{\lceil r/2 \rceil} = 2^r \zeta''(q)^{\lceil r/2 \rceil} z_r$$

with z_r as defined in (9). It thus follows by Markov's inequality that

$$\Pr_{\tau \sim Y}[\mathsf{fitkey}(\tau, \ell, h) \geq 1 \wedge \mathsf{fwd}(\tau) \leq \zeta(q) \wedge \mathsf{bwd}(\tau) \leq \zeta(q)] \leq 2^\ell \zeta''(q)^{\lceil \ell/2 \rceil} z_\ell$$

and

$$\Pr_{\tau \sim Y}[\mathsf{fitkey}(\tau, r, h) \geq C z_r \mid \mathsf{fwd}(\tau) \leq \zeta(q) \wedge \mathsf{bwd}(\tau) \leq \zeta(q)] \leq \frac{2^r \zeta''(q)^{\lceil r/2 \rceil}}{C}$$

which proves Lemma 2 and inequality (10). □

Remaining Steps. Having upper bounded the probability of bad transcripts, the rest of the proof concerns itself with lower bounding the ratio

$$\frac{\Pr[X = \tau]}{\Pr[Y = \tau]}$$

for good transcripts τ, and more precisely of showing this ratio is at least $1 - \varepsilon$ for

$$\varepsilon = \begin{cases} \ell C^{\lceil (\ell+1)/2 \rceil} N \left(\frac{8q}{2^{\kappa+n}}\right)^{\lceil \ell/2 \rceil} & \text{if } q \leq N, \\ \frac{q^2 \ell}{N} C^{\lceil (\ell+1)/2 \rceil} \left(\frac{8}{2^\kappa}\right)^{\lceil \ell/2 \rceil} & \text{if } q \geq N. \end{cases}$$

For reasons of space we leave this part of the proof to the full version [10]. The overall approach, however, is quite similar to that espoused by Chen and Steinberger [8], and the main technical tool required for this part of the proof is indeed their own "path-completion lemma".

Acknowledgments. Yuanxi Dai was supported by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301 and the National Natural Science Foundation of China Grant 61033001, 61361136003. Jooyoung Lee was supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Education (NRF-2013R1A1A2007488). Bart Mennink was supported by the Research Fund KU Leuven, OT/13/071 and the Research Council KU Leuven: GOA TENSE (GOA/11/007). John Steinberger was supported by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the National Natural Science Foundation of China Grant 61033001, 61361136003, and by the China Ministry of Education grant number 20121088050.

References

1. William Aiello, Mihir Bellare, Giovanni Di Crescenzo, and Ramarathnam Venkatesan. Security amplification by composition: the case of doubly-iterated, ideal ciphers, CRYPTO 1998, LNCS 1462, pp. 390–407.
2. ANSI X9.52: Triple Data Encryption Algorithm Modes of Operation (withdrawn), 1998.
3. Frederik Armknecht, Ewan Fleischmann, Matthias Krause, Jooyoung Lee, Martijn Stam and John Steinberger, The preimage security of double-block length compression functions. Asiacrypt 2011, LNCS 7073, Springer, 233–251.
4. Mihir Bellare and Phillip Rogaway, The security of triple encryption and a framework for code-based game-playing proofs. Eurocrypt 2006, LNCS 4004 pp. 409–426.
5. Mihir Bellare and Phillip Rogaway, Code-based game-playing proofs and the security of triple encryption. IACR eprint report. `eprint.iacr.org/2004/331`
6. John Black, Phillip Rogaway, Thomas Shrimpton, Black-Box Analysis of the Block Cipher-Based Hash-Function Constructions from PGV. CRYPTO 2002, LNCS 2442, pp. 320–335.
7. Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, Francois-Xavier Standaert, John Steinberger and Elmar Tischhauser, Key-Alternating Ciphers in a Provable Setting: Encryption Using a Small Number of Public Permutations. EUROCRYPT 2012, LNCS 7237, pp. 45–62.
8. Shan Chen and John Steinberger, Tight security bounds for key-alternating ciphers. EUROCRYPT 2014, LNCS 8441, pp. 327–350.
9. Yuanxi Dai and John Steinberger, Tight security bounds for multiple encryption. IACR Cryptology ePrint Archive, 2014/096, `http://eprint.iacr.org/2014/096.pdf`.
10. Yuanxi Dai, Jooyoung Lee, Bart Mennink and John Steinberger, The security of multiple encryption in the ideal cipher model. (Full version of this paper.) IACR Cryptology ePrint Archive.
11. Whitfield Diffie and Martin Hellman, Exhaustive cryptanalysis of the NBS data encryption standard. Computer 10 (6), 74–84, 1997.
12. Itai Dinur, Orr Dunkelman, Nathan Keller and Adi Shamir, Efficient Dissection of Composite Problems, with Applications to Cryptanalysis, Knapsacks, and Combinatorial Search Problems. CRYPTO 2012, LNCS 7417, pp. 719–740.
13. Shimon Even and Oded Goldreich, On the power of cascade ciphers. ACM Transactions on Computer Systems, vol. 3, no. 2, pp. 108–116, 1985.
14. Shimon Even and Yishay Mansour, A Construction of a Cipher From a Single Pseudorandom Permutation. ASIACRYPT 1991, LNCS 739, pp. 210–224, Springer-Verlag, 1993.
15. FIPS46-3: Data Encryption Standard. National Institute of Standards and Technology (withdrawn), 1999.
16. Peter Gaži, Plain versus Randomized Cascading-Based Key-Length Extension for Block Ciphers, CRYPTO 2013, LNCS 8042, pp551–570.
17. Peter Gaži and Ueli Maurer, Cascade encryption revisited, Asiacrypt 2009, LNCS 5912, pp37–51.
18. Peter Gaži and Stefano Tessaro, Efficient and Optimally Secure Key-Length Extension for Block Ciphers via Randomized Cascading. Eurocrypt 2012, LNCS 7237, pp. 63–80, Springer, Heidelberg (2012).
19. Joe Kilian and Phillip Rogaway, How to protect DES against exhaustive key search (an analysis of DESX). Journal of Cryptology 14 (1), 17-35 (2001).

20. Matthias Krause, Frederik Armknecht and Ewan Fleischmann, Preimage resistance beyond the birthday bound: Double-length hashing revisited. IACR eprint report, `http://eprint.iacr.org/2010/519.pdf`.
21. Rudolphe Lampe, Jacques Patarin and Yannick Seurin, An Asymptotically Tight Security Analysis of the Iterated Even-Mansour Cipher, Asiacrypt 2012, Lecture Notes in Computer Science Volume 7658, pp 278-295, 2012.
22. Jooyoung Lee, Towards Key-Length Extension with Optimal Security: Cascade Encryption and Xor-cascade Encryption, Eurocrypt 2013, LNCS 7881, pp405–425.
23. Jooyoung Lee, Tight Security for Triple Encryption. IACR Cryptology ePrint Archive, 2014/015, `http://eprint.iacr.org/2014/015.pdf`.
24. Jooyoung Lee, John Steinberger and Martijn Stam, The preimage security of double-block-length compression functions. IACR eprint report, `http://eprint.iacr.org/2011/210.pdf`.
25. Michael Luby and Charles Rackoff, Pseudo-random permutation generators and cryptographic composition. In: STOC 1986: Proceedings of the 18th Annual ACM Symposium on Theory of Computing, pp. 356–363.
26. Stefan Lucks, Attacking triple encryption. *Fast Software Encryption* 1998, LNCS 1372, pp. 239–253.
27. Ueli Maurer and James L. Massey, Cascade ciphers: The importance of being first. Journal of Cryptology 6(1), pp. 55–61, 1993.
28. Ueli Maurer, Kryzstof Pietrzak and Renato Renner, Indistinguishability amplification. CRYPTO 2007, LNCS 4622, pp. 130–149.
29. Ueli Maurer and Stefano Tessaro, Computational indistinguishability amplification: Tight product theorems for system composition. CRYPTO 2009, LNCS 5677, pp. 355–373.
30. Bart Mennink and Bart Preneel, Triple and Quadruple Encryption: Bridging the Gap. IACR Cryptology ePrint Archive, 2014/016, `http://eprint.iacr.org/2014/016.pdf`.
31. Ralph Merkle and Martin Hellman, On the Security of Multiple Encryption, Communications of the ACM, vol. 24, no. 7, pp. 465–467, ,July 1981. See also: Communicutionr of the ACM, vol. 24, no. 11, p. 776, November 1981.
32. Myers, S., On the development of block-ciphers and pseudo-random function generators using the composition and XOR operators. Masters thesis, University of Toronto (1999).
33. Paul C. van Oorschot and Michael Wiener, Improving implementable meet-in-the-middle attacks by orders of magnitude, CRYPTO 1996, LNCS 1109 pp. 229–236.
34. Paul C. van Oorschot and Michael Wiener, A Known-Plaintext Attack on Two-Key Triple Encryption, Eurocrypt 1990, LNCS 473 pp. 318–325.
35. NIST SP 800-67, Revision 1: Recommendation for the Triple Data Encryption Algorithm (TDEA) Block Cipher. National Institute of Standards and Technology, 2012.
36. Jacques Patarin, Etude de Génerateurs de Permutations Bases sur les Schemas du DES. In Ph.D. Thesis. Inria, Domaine de Voluceau, France, 1991.
37. Jacques Patarin, The "Coefficients H" Technique, Selected Areas in Cryptography, LNCS 5381, 2009, pp. 328-345.
38. John Steinberger, Improved Security Bounds for Key-Alternating Ciphers via Hellinger Distance, `http://eprint.iacr.org/2012/481.pdf.`
39. Stefano Tessaro, Security Amplification for the Cascade of Arbitrarily Weak PRPs: Tight Bounds via the Interactive Hardcore Lemma. Theory of Cryptography Conference 2011, LNCS 6597, pp. 37-54.