The Multi-User Security of Double Encryption

Viet Tung Hoang[1] and Stefano Tessaro[2]

[1] Dept. of Computer Science, Florida State University
[2] Dept. of Computer Science, University of California Santa Barbara

Abstract. It is widely known that double encryption does not substantially increase the security of a block cipher. Indeed, the classical meet-in-the middle attack recovers the $2k$-bit secret key at the cost of roughly 2^k off-line enciphering operations, in addition to very few known plaintext-ciphertext pairs. Thus, essentially as efficiently as for the underlying cipher with a k-bit key.
This paper revisits double encryption under the lens of multi-user security. We prove that its security degrades only very mildly with an increasing number of users, as opposed to single encryption, where security drops linearly. More concretely, we give a tight bound for the multi-user security of double encryption as a pseudorandom permutation in the ideal-cipher model, and describe matching attacks.
Our contribution is also conceptual: To prove our result, we enhance and generalize the generic technique recently proposed by Hoang and Tessaro for lifting single-user to multi-user security. We believe this technique to be broadly applicable.

Keywords: symmetric security, provable security, multi-user security, double encryption

1 Introduction

A classical problem in cryptography is that of stretching the key length of a block cipher. Namely, from a block cipher E with block length n and key length k, we want to obtain a new one with key length $k' > k$ which is *more* secure than E. The problem was naturally motivated by legacy designs – in particular, DES – with inherently too-short keys (e.g., 56 bits), and the desire to stretch this key length generically without resorting to designing a new cipher.

The common wisdom is that *double encryption* is not useful for key-stretching purposes. Here, by double encryption, we mean the construction that, given an n-bit plaintext M and two k-bit keys K_1, K_2, outputs $E_{K_1}(E_{K_2}(M))$. Indeed, there is a well-known meet-in-the-middle attack recovering the key with only marginally more than 2^k operations given (very few) valid plaintext-ciphertext pairs. This weakness has led to the widespread deployment (which continues to date in some niche areas) of Triple-DES [1], as well as a number of works on analyzing the theory of triple and multiple encryption [7, 13–17, 20], and

alternative constructions with extra whitening steps (and key material) [15,16, 18–20].

In this paper, we revisit double encryption in the context of multi-user security, where we give tight bounds, and show that it constitutes a sound and simple method to mitigate multi-user attacks on block ciphers. However, this problem will also serve as an application for a generic framework to provide good multi-user security bounds, and which we hope to be of wider applicability.

Double encryption in the single user setting. As in previous works, we study the security of double encryption in the ideal-cipher model as a (strong) pseudorandom permutation (PRP). The attacker A is given access to an ideal cipher E to which it can issue p forward or backward queries for any chosen key (these are usually referred to as "offline queries"), and up to q queries (in either direction) to $E_{K_1} \circ E_{K_2}$ (for random secret keys K_1, K_2) or a truly random permutation on the n-bit strings (this being usually called "online queries"). The attacker's goal is to decide which of the two it is accessing. In this model, Aiello *et al.* [2] proved that A's distinguishing advantage satisfies

$$\mathsf{Adv}^{\mathrm{prp}}_{\mathrm{DE}[E]}(A) \leq \left(\frac{p}{2^k}\right)^2 . \tag{1}$$

where DE[E] denotes double encryption. Note that for *single* encryption, the bound is easily shown to be $\mathsf{Adv}^{\mathrm{prp}}_{E}(A) \leq \frac{p}{2^k}$. Both advantages become non-negligible for the same $p \approx 2^k$, although (1) is smaller when $p \ll 2^k$.

The multi-user setting. In the *multi-user (mu) setting*, originally proposed by Bellare, Boldyreva, and Micali [5] for public-key encryption, the attacker can distribute its online queries adaptively across multiple independent key pairs (in the real world) or independent permutations (in the ideal world). A few recent block-cipher analyses [19,24,29] have focused on mu security, and the notion has established itself as a more realistic security target.

One expects security to degrade as the number of users increases, and this loss can be linear in the worst case. For example, for single-encryption, we do have

$$\mathsf{Adv}^{\pm\text{mu-prp}}_{E}(A) \leq \frac{u\,(p+u)}{2^k} \leq \frac{q\,(p+q)}{2^k} , \tag{2}$$

where u is a bound on the number of users A queries, and this bound is tight, i.e., there is a matching attack [10]. Also, we can only guarantee $u \leq q$, as the attacker can decide to only issue one query per user. However, for *double* encryption, we can use a simple hybrid argument to show that

$$\mathsf{Adv}^{\pm\text{mu-prp}}_{\mathrm{DE}[E]}(A) \leq u\left(\frac{p+2q}{2^k}\right)^2 \leq q\left(\frac{p+2q}{2^k}\right)^2 . \tag{3}$$

This bound is already better than the one from (2) – for instance, for roughly $p = q = 2^{k/2}$, this latter bound is still $O(2^{-k/2})$, but (2) gives $\Omega(1)$. However, contrary to the single-encryption case, it is not clear that the bound is tight. We will indeed show a much better bound.

Our bounds. Our main result shows that the security of double encryption does not degrade substantially in the multi-user setting, and that the bound from (3) is overly pessimistic. In particular, we prove that

$$\mathsf{Adv}^{\pm\text{mu-prp}}_{\mathrm{DE}[E]}(A) \leq \frac{1}{2^n} + \frac{5q}{2^{k+n/2}} + \frac{6qB^2 + 222BQ^2}{2^{2k}}$$

where $Q = \max\{p, q\}$ and $B = 5\max\{n + k/2, 2q/2^n\}$. This bound is rather cumbersome, but the key observation is that third-degree monomials in p and q all appear with denominator 2^{2k+n}, whereas any term with denominator 2^{2k} is at most *quadratic* in p, q – very similar to the single-user case.

Recall that the meet-in-the-middle attack on the single user security of double encryption succeeds with advantage $p^2/2^{2k}$, and Biham's key-collision attack [10] achieves advantage $q^2/2^{2k}$. Therefore for the setting that $n \geq k$ (such as DES or AES), our bound is tight. For the setting $n \ll k$ (which occurs in Format-Preserving Encryption [6], and several block-cipher designs), finding matching attacks is difficult, and we leave it as an open problem. However, as an intermediate step, we note that most proofs are in models where the keys are revealed to the distinguisher at the end of the execution. In this model, we can give a matching attack (based on the meet-in-the-middle paradigm) that achieves distinguishing advantage

$$\max\{\lfloor n/8\lg(n)\rfloor, q/2^n\} \cdot \frac{p^2}{3 \cdot 2^{2k}} \ .$$

We discuss attacks below in Section 6.

A disclaimer. We stress that the common wisdom that there is no security increase is obviously still in place. However, the envisioned application is to ciphers whose key length is not an issue in the single-user setting, but becomes too short in a multi-user regime. For instance, a multi-user attack reduces the security of (single) AES128 to 64 bits. Our result shows that iterating AES128 *twice* substantially mitigates the impact of a multi-user attack, and that in fact we obtain almost optimal multi-user security, namely around 115 bits for a total key length of 256 bits. (Also see Figure 2.)

Techniques. Our result is obtained using new techniques we introduce and that we believe to be of broad applicability in lifting existing analyses from the single-user (su) to the mu setting.

Hoang and Tessaro (HT) [19] already proposed a generic approach for this purpose. It is illustrative to briefly review it, and see why it fails for double encryption. HT's idea is to show that the construction (e.g., double encryption) satisfies, in the su case, a property called *point-wise proximity*, a stronger property than indistinguishability, already used in previous works (e.g., in [9]). Concretely, this means that there exists a function $\epsilon = \epsilon(p, q)$ of the query parameters p and q, such that for all transcripts τ containing p offline and q online queries, we have

$$\mathsf{ps}_{\mathrm{ideal}}(\tau) - \mathsf{ps}_{\mathrm{real}}(\tau) \leq \epsilon(p, q) \cdot \mathsf{ps}_{\mathrm{ideal}}(\tau) \ , \tag{4}$$

where $p_{\mathbf{S}_{\text{ideal}}}(\tau)$ and $\mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau)$ are the so-called ideal and real *interpolation probabilities*. Namely, they describe the probability that the real ($\mathsf{p}_{\mathbf{S}_{\text{real}}}$) and the ideal ($\mathsf{p}_{\mathbf{S}_{\text{ideal}}}$) worlds behave consistently with the transcript when the queries the transcript contains are asked in that order.

HT show that then point-wise proximity is achieved in the multi-user experiment, where $\epsilon(p, q)$ is replaced by $\epsilon(p + qt, q)$, where t is the number of calls made by the construction to the underlying primitive (in the case of double encryption, $t = 2$). This implies that the distinguishing advantage is also at most $\epsilon(p + qt, q)$. For this argument to hold, however, ϵ needs to be super-additive, i.e., $\epsilon(x, y) + \epsilon(x, z) \leq \epsilon(x, x + y)$, and moreover, $\epsilon(\cdot, y)$ and $\epsilon(x, \cdot)$ need to be non-decreasing functions for al $x, y \in \mathbb{N}$. For double encryption, no such ϵ can be established. For instance, the natural candidate $\epsilon(p, q) = \left(\frac{p}{2^k}\right)^2$ is not super-additive, as $\epsilon(x, y) + \epsilon(x, z) = 2\epsilon(x, y + z)$.

We take a different approach, by introducing a relaxed notion of *almost proximity*, which in particular akin to the H-coefficient method (cf. e.g. [12, 26]), introduced a partition the set of *single-user* transcripts into good and bad transcripts, and proximity guarantees are shown only on the former. Our main technical insight is the introduction of a precise framework to mitigate the effects of the growth of the probability of a bad transcript when increasing the number of users. We dispense with a formulation here – the conditions are not concise – and refer the reader to Section 3. We note that we also provide simplifications of the framework in Section 4, one of which is in particular sufficient for analyzing double encryption. We finally apply it in Section 5.

Further related work. Multiple encryption has been studied also in the standard computational model, with respect to the question of how it amplifies (weak) PRP security. Luby and Rackoff [21] initially studied double encryption, and bounds for multiple encryption were later provided by Myers [25] and Tessaro [28].

Also, while above we have focused on block cipher analyses, recent works have studied mu security in different contents, in particular for authentication encryption [8] and message-authentication codes [3, 4].

2 Preliminaries

Notation. For a finite set S, we let $x \leftarrow\!\!\$\ S$ denote the uniform sampling from S and assigning the value to x. Let $|x|$ denote the length of the string x, and for $1 \leq i < j \leq |x|$, let $x[i, j]$ denote the substring from the ith bit to the jth bit (inclusive) of x. If A is an algorithm, we let $y \leftarrow A(x_1, \ldots; r)$ denote running A with randomness r on inputs $x_1, \ldots$ and assigning the output to y. We let $y \leftarrow\!\!\$\ A(x_1, \ldots)$ be the resulting of picking r at random and letting $y \leftarrow A(x_1, \ldots; r)$.

Multi-user PRP security of blockciphers. Let $\Pi : \mathcal{K} \times \{0,1\}^n \to \{0,1\}^n$ be a blockcipher, which is built on another blockcipher $E : \{0,1\}^k \times \mathcal{M} \to \mathcal{M}$.

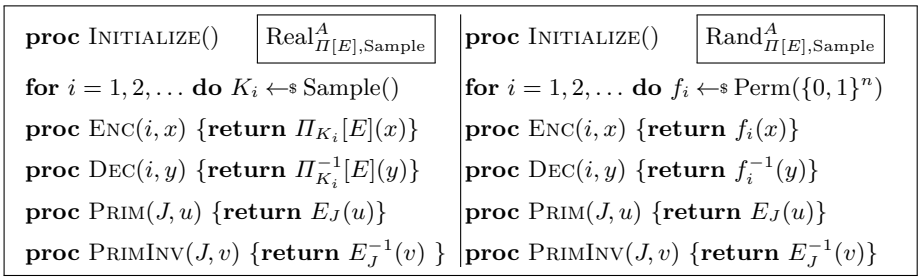

Fig. 1: **Games defining the multi-user security of a blockcipher** $\Pi : \mathcal{K} \times \{0,1\}^n \to \{0,1\}^n$**.** This blockcipher is based on another blockcipher $E : \{0,1\}^k \times \{0,1\}^n \to \{0,1\}^n$. The game is associated with a key-sampling algorithm Sample. Here Perm$(\{0,1\}^n)$ denotes the set of all permutations on $\{0,1\}^n$.

We associate with Π a key-sampling algorithm Sample. Let A be an adversary. Define

$$\mathsf{Adv}^{\pm\text{mu-prp}}_{\Pi[E],\mathrm{Sample}}(A) = \Pr[\mathrm{Real}^A_{\Pi[E],\mathrm{Sample}} \Rightarrow 1] - \Pr[\mathrm{Rand}^A_{\Pi[E],\mathrm{Sample}} \Rightarrow 1]$$

where games Real and Rand are defined in Fig. 1. If Sample is the uniform sampling of $\mathcal{K}$ then we only write $\mathsf{Adv}^{\pm\text{mu-prp}}_{\Pi[E]}(A)$.

In the games above, we first use Sample to sample keys $K_1, K_2, \ldots \in \mathcal{K}$ for Π, and independent, random permutations $f_1, f_2, \ldots$ on $\mathcal{M}$. The adversary is given four oracles Prim, PrimInv, Enc, and Dec. In both games, the oracles Prim and PrimInv always give access to the primitive E and its inverse respectively. The Enc and Dec oracles give access to $f_1(\cdot), f_2(\cdot), \ldots$ and their inverses respectively in game Rand, and access to $\Pi[E](K_1, \cdot), \Pi[E](K_2, \cdot), \ldots$ and their inverses in game Real. The adversary finally needs to output a bit to tell which game it is interacting with.

Single and Double Encryption. Let $k, n \in \mathbb{N}$ and let $E : \{0,1\}^k \times \{0,1\}^n \to \{0,1\}^n$ be a blockcipher. The Single Encryption of E is the blockcipher E itself. The Double Encryption DE$[E]$ of E is a blockcipher with keyspace $(\{0,1\}^k)^2$ and message space $\{0,1\}^n$. On key $K = (J_1, J_2)$ and message $x \in \{0,1\}^n$, $\mathrm{DE}_K[E](x)$ returns $E_{J_2}(E_{J_1}(x))$.

Systems and transcripts. Following up the notation from [19] (which was in turn inspired by Maurer's framework [22]), it is convenient to consider interactions of a distinguisher A with an abstract system $\mathbf{S}$ which answers A's queries. The resulting interaction then generates a transcript $\tau = ((X_1, Y_1), \ldots, (X_q, Y_q))$ of query-answer pairs. It is well known that $\mathbf{S}$ is entirely described by the probabilities $\mathsf{p}_{\mathbf{S}}(\tau)$ that if we make queries in τ to system $\mathbf{S}$, we will receive the answers as indicated in τ. We say in particular that $\mathbf{S}$ is *stateless* if $\mathsf{p}_{\mathbf{S}}(\tau)$ is invariant under permuting the orders of the input-output pairs it contains.

We will generally describe systems informally, or more formally in terms a set of oracles they provide, and only use the fact that they define a corresponding probabilities $\mathsf{p}_{\mathbf{S}}(\tau)$ without explicitly giving these probabilities.

THE EXPECTATION METHOD. In this paper, we shall use the expectation method of Hoang and Tessaro [19]. For a pair of systems $\mathbf{S}_{\text{real}}$ and $\mathbf{S}_{\text{ideal}}$, this method aims to bound the gap $\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau) - \mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau)$, for a fixed (su) transcript τ such that $\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau) > 0$. Under this method, one extends the transcript with a random variable S. In $\mathbf{S}_{\text{real}}$, this S is often a part of the key and suppose that it has marginal distribution μ. In $\mathbf{S}_{\text{ideal}}$, we pick S of the same marginal distribution μ, but independent of τ. Let $\mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau, s)$ denote the probability that $\mathbf{S}_{\text{real}}$ behaves according to τ, and S agrees with s. Let $\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau, s)$ denote the probability that $\mathbf{S}_{\text{ideal}}$ behaves according to τ, and $S \leftarrow\!\!\$\ \mu$ agrees with s. Under the expectation method, one partitions the range of S into two sets, Γ_{good} and Γ_{bad}. For s such that $\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau, s) > 0$, if $s \in \Gamma_{\text{bad}}$ then we say that s is *bad*; otherwise s is *good*. We write $\Pr[S \in \Gamma_{\text{bad}}]$ to denote the probability that $S \leftarrow\!\!\$\ \mu$ independent of τ is bad. Hoang and Tessaro give the following result.

Lemma 1 (The expectation method). *[19] Fix a su transcript τ such that $\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau) > 0$. Assume that there is a partition Γ_{good} and Γ_{bad} of the range $\mathcal{U}$ of S, as well as a function $g : \mathcal{U} \to [0, \infty)$ such that $\Pr[S \in \Gamma_{\text{bad}}] \le \delta$ and for all $s \in \Gamma_{\text{good}}$,*

$$1 - \frac{\mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau, s)}{\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau, s)} \le g(s) \ .$$

Then

$$\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau) - \mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau) \le (\delta + \mathbf{E}[g(S)]) \cdot \mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau) \ . \qquad \square$$

Note that in Lemma 1, the expectation is taken over all possible (good or bad) values of S.

3 A Generic Method to Bound Multi-user Security

In this section we present a generic method to prove information-theoretic mu security bounds, based (mostly) on upper bounding single-user quantities. The framework is very general, and in fact generalizes the approach by Hoang and Tessaro [19] based on pointwise proximity.

THE GENERIC SETTING. We consider two (stateless) systems $\mathbf{S}_{\text{real}}$ and $\mathbf{S}_{\text{ideal}}$, called the *real* and *ideal* systems, respectively. Each of these two systems can be invoked via two oracles CONS and PRIM, allowing for *construction* and *primitive* queries, respectively. First off, PRIM gives access to an ideal primitive (for example, an ideal cipher, a random function or permutation), whereas CONS's role depends on the context, but always answers queries of the form (i, X), where i is the index of a user and X is the query for that user. More specifically:

1. In $\mathbf{S}_{\text{real}}$, the oracle Cons upon a query (i, X) invokes a construction Π which makes calls to Prim, and additionally depends on some local, initially chosen randomness (or key) K_i. That is, the output is $\Pi^{\text{Prim}}(K_i, X)$.
2. In $\mathbf{S}_{\text{ideal}}$, the oracle Cons samples independent functions $f_1, f_2, \ldots$ from some distribution, and answers a query (i, X) as $f_i(X)$.

For example, the game from Figure 1 can be described as suitable systems $\mathbf{S}_{\text{real}}$ and $\mathbf{S}_{\text{ideal}}$: We would simply handle inversion queries (to Dec and PrimInv) by specifying the direction of the query in the input given to Cons and Prim, i.e., $X = (+, x)$ or $X = (-, y)$. Also, we can model more complex scenarios, like the security of authenticated encryption schemes, as long as we can map the security notion to suitable $\mathbf{S}_{\text{real}}$ and $\mathbf{S}_{\text{ideal}}$.

We generally will assume that there exists a metric of *data complexity* associated with queries made to Cons. For instance, if Cons takes variable-length inputs, σ could be number bits queried to it, whereas if the input length is fixed, this could just be the number of queries. We assume that there exists a parameter t indicating that when answering multiple queries with overall data complexity σ, Π makes at most $t \cdot \sigma$ queries to Prim.

The distinguishing problem. For any adversary A and a system $\mathbf{S}$, we let $\text{Script}(A, \mathbf{S})$ denote the random variable for the transcript of the interaction of A and $\mathbf{S}$. Recall that the advantage of the adversary in distinguishing two systems $\mathbf{S}_{\text{real}}$ and $\mathbf{S}_{\text{ideal}}$ is at most the statistical distance between the distributions of the adversary's transcript in the real and ideal games, which is

$$\mathsf{Adv}^{\text{dist}}_{\mathbf{S}_{\text{real}}, \mathbf{S}_{\text{ideal}}}(A) \leq \sum_{\tau} \max\{0, \mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau) - \mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau)\}, \tag{5}$$

where the sum is taken over all τ such that $\Pr[\text{Script}(A, \mathbf{S}_{\text{ideal}}) = \tau] > 0$.

Note that there might be some context-dependent constraints on the adversary's queries. For example, if part of the inputs to Cons include nonces to a nonce-based authenticated encryption, then one might require that the nonces will not repeat. This is easy to handle, since it will only restrict the set of valid transcripts to be considered. We will usually capture the complexity of A in terms of the number of Prim queries, p, the number of Cons queries q, and the overall data complexity σ for the queries made to Cons. A security bound ϵ is then viewed as a function $\epsilon(p, q, \sigma)$. We say that a function $\epsilon(\cdot, \cdot, \cdot) : \mathbb{N}^3 \to [0, 1]$ is *monotonic* if $\epsilon(\cdot, y, z)$, $\epsilon(x, \cdot, z)$, and $\epsilon(x, y, \cdot)$ are increasing functions, for any $x, y, z \in \mathbb{N}$. Often security bounds are monotonic functions, since increasing the adversary's resources can only help it.

Almost proximity. We now establish a condition on $\mathbf{S}_{\text{real}}$ that we call *almost proximity*, which will allow us to establish mu security from a number of functions, δ_0, δ_1 and δ_2, we define next. In particular, some of these functions (δ_1 and δ_2) are defined with respect to single-user (su) transcript, i.e., transcripts were all queries to Cons are of the form (i, X) *for one single* i.

One begins by defining a context-dependent, undesirable property on su transcripts that we call *bad*, and if a su transcript is not bad then it is *good*. We

partition in particular the set of bad transcripts into two sets, $\mathcal{S}$ and $\mathcal{S}'$. In many cases (such as our Double Encryption application below), one of the two sets $\mathcal{S}$ and $\mathcal{S}'$ is simply the empty set, but we envision more general application scenarios.

Further, we will assume that there exists a function Rate such that for any good su transcript τ,

$$\mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\tau) - \mathsf{ps}_{\mathbf{S}_{\text{real}}}(\tau) \leq \text{Rate}(\tau) \cdot \mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\tau) \ ,$$

where Rate is in particular an *increasing* function mapping a transcript to a number in $[0, 1]$, meaning that for any transcripts τ and τ' such that τ' contains all the query-answer pairs of τ (possibly in a different order), we have $\text{Rate}(\tau') \geq \text{Rate}(\tau)$.

Then, we also assume that there is a monotonic function δ_2 such that for any adversary B attacking a single user via p PRIM queries, q ENC queries with overall data complexity σ, we have

$$\Pr[\text{Script}(B, \mathbf{S}_{\text{ideal}}) \in \mathcal{S}] \leq \delta_2(p, q, \sigma) \ .$$

Note that the bound above is with respect to the ideal system, $\mathbf{S}_{\text{ideal}}$, and thus often easy to compute.

We also define another, context-dependent, desired property on mu transcripts that we call *nice* — we let $\mathcal{N}$ be the set of all nice mu transcripts. (We stress that niceness is with respect to mu transcripts, whereas being good/bad is only with respect to su ones.) The notion of niceness involves only the CONS query-answer pairs: for any two transcripts τ and τ' that have the same CONS query-answer pairs (possibly in different orders), if $\tau \in \mathcal{N}$ then so is τ'. Also, for a mu transcript τ involving queries to exactly r users, and for each $i \in \{1, \ldots, r\}$, let $\text{Map}(i, \tau)$ denote the su transcript obtained by deleting the $\text{CONS}(j, \cdot)$ queries and answers for any $j \neq i$. We require the following conditions:

- For any transcript $\tau \in \mathcal{N}$ and all i, $\text{Map}(i, \tau) \notin \mathcal{S}'$.
- There is a monotonic function δ_0 such that for any mu adversary A making p PRIM queries, q CONS queries, and data complexity σ,

$$\Pr[\text{Script}(A, \mathbf{S}_{\text{ideal}}) \notin \mathcal{N}] \leq \delta_0(p, q, \sigma) \ .$$

- There is a monotonic function δ_1 such that for any $\tau \in \mathcal{N}$ of r users that contains p PRIM, q CONS queries of total data complexity at most σ,

$$\sum_{i=1}^{r} \text{Rate}(\text{Map}(i, \tau)) \leq \delta_1(p, q, \sigma) \ . \tag{6}$$

  We refer to this last property as mu-boundedness.

We refer to the existence of suitable functions $\delta_0, \delta_1, \delta_2$ for corresponding Rate, Map, $\mathcal{S}$, $\mathcal{S}'$ and $\mathcal{N}$ as meeting the *almost-proximity* conditions.

MU SECURITY VIA ALMOST PROXIMITY. The following result bounds the mu advantage in distinguishing $\mathbf{S}_{\text{real}}$ and $\mathbf{S}_{\text{ideal}}$, granted the almost-proximity conditions defined above are met.

Lemma 2 (Mu-security via almost proximity). *Assume that the almost-proximity conditions above are met, for some δ_2, δ_0 and δ_1. Then for any adversary A that makes at most q* CONS *queries of total data complexity σ, and p* PRIM *queries, we have*

$$\mathsf{Adv}^{\text{dist}}_{\mathbf{S}_{\text{real}},\mathbf{S}_{\text{ideal}}}(A) \le \delta_0(p,q,\sigma) + 2\delta_1(p+t\sigma,q,\sigma) + 2q\cdot\delta_2(p+t\sigma,q,\sigma) \ .$$

DISCUSSION. A meaningful question is why we need to separate the set of bad su transcripts into $\mathcal{S}$ and $\mathcal{S}'$. The reason is that, when we move from su to mu setting, under our method, the term δ_2 will blow up to $q\delta_2$, which is similar to the hybrid argument. To avoid an inferior mu bound, we would like to minimize the term δ_2 as much as possible, by carving out $\mathcal{S}'$ from the set of bad su transcripts. Due to the requirement that $\text{Map}(i,\tau) \notin \mathcal{S}'$ for every nice mu transcript τ and every i, the set $\mathcal{S}'$ and the notion of niceness needs to be chosen in tandem to minimize $q\delta_2 + \delta_0(p,q,\sigma)$. Bounding $\Pr[\text{Script}(A,\mathbf{S}_{\text{ideal}}) \notin \mathcal{N}]$ requires working directly in the mu setting, but recall that we are in the ideal game, which is often simple to deal with.

Proof (of Lemma 2). Since we consider a computationally unbounded adversary, without loss of generality, assume that the adversary is deterministic. For simplicity, from this point, we will write δ_2 and δ_1 instead of $\delta_2(p+t\sigma,q,\sigma)$ and $\delta_1(p+t\sigma,q,\sigma)$. Without loss of generality, assume that $\delta_1 < 1/2$; otherwise the the claimed bound in the statement of this lemma is moot. We also assume that the adversary's transcript involves at most r users.

RESTRICTING TO NICE TRANSCRIPTS. Recall that in the ideal system, the probability that the adversary A can produce a mu transcript that is not nice is at most $\delta_0(p,q,\sigma)$. From Equation (5), what is left is to show that

$$\sum_{\tau} \mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\tau) - \mathsf{ps}_{\mathbf{S}_{\text{real}}}(\tau) \le 2\delta_1 + 2q\delta_2, \tag{7}$$

where the sum in the left hand side is taken over all nice transcripts τ in the support $\text{supp}(\text{Script}(A,\mathbf{S}_{\text{ideal}}))$ of $\text{Script}(A,\mathbf{S}_{\text{ideal}})$ such that $\mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\tau) > \mathsf{ps}_{\mathbf{S}_{\text{real}}}(\tau)$. Below, when we talk about a *valid* transcript τ, this means that τ meets the constraint above.

BUILDING HYBRIDS. For each $i \in \{0,\ldots,r\}$, consider the hybrid system $\mathbf{S}_i$ that provides the interface compatible with the real and ideal systems, but queries for user u_j are answered via the actual construction $\Pi^{\text{PRIM}}(K_j,\cdot)$ for $j > i$, and via an independent, perfect simulation of the $\text{CONS}(j,\cdot)$ oracle of the ideal game

if $j \le i$. Then $\mathbf{S}_0 = \mathbf{S}_{\text{real}}$ and $\mathbf{S}_r = \mathbf{S}_{\text{ideal}}$ and thus for any valid transcript τ,

$$\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau) - \mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau) = \sum_{i=1}^{r} \mathsf{p}_{\mathbf{S}_i}(\tau) - \mathsf{p}_{\mathbf{S}_{i-1}}(\tau) \ . \tag{8}$$

Let B_i be the following hybrid su adversary. It samples key K_j for Π^{Prim} for every $i < j \le r$, and then runs A. Queries for user u_j are answered via $\Pi^{\text{Prim}}(K_j, \cdot)$ if $j > i$, and via the $\text{Cons}(1, \cdot)$ oracle of B_i if $j = i$, and via an independent, perfect simulation of the $\text{Cons}(j, \cdot)$ oracle of the ideal game if $j < i$. In other words, adversary B_i simulates system $\mathbf{S}_{i-1}$ in its su real game, and simulates system $\mathbf{S}_i$ in its su ideal game. It makes at most q Cons queries of total data complexity σ and at most $p + t\sigma$ Prim queries.

Reducing to transcript-wise gap. Fix a valid transcript τ. Let $\mathcal{T}(i, \tau)$ denote the set of extended transcripts of B_i in its su ideal game that are enhanced with the simulated Cons queries and answers as well as the simulated keys K_j, such that the corresponding simulated transcript for A is τ. For each $\tau_i \in \mathcal{T}(i, \tau)$, let $\text{Tr}(\tau_i)$ be the transcript of B_i derived from τ_i. For $\mathbf{S} \in \{\mathbf{S}_{\text{real}}, \mathbf{S}_{\text{ideal}}\}$, let $\mathsf{p}_{\mathbf{S}}(\tau_i)$ denote the probability that, when B_i interacts with $\mathbf{S}$, its enhanced transcript is τ_i. Note that compared to $\text{Tr}(\tau_i)$, the additional information τ_i contains is the keys K_j, and the queries/answers on the simulated oracle $\text{Cons}(j, \cdot)$ of the ideal game for users $j < i$. Since this information is independent of $\mathbf{S}_{\text{real}}$ and $\mathbf{S}_{\text{ideal}}$,

$$\frac{\mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau_i)}{\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau_i)} = \frac{\mathsf{p}_{\mathbf{S}_{\text{real}}}(\text{Tr}(\tau_i))}{\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\text{Tr}(\tau_i))} \ . \tag{9}$$

Let $\mathcal{S}_i$ be the set of extended transcripts τ_i of B_i such that $\text{Tr}(\tau_i) \in \mathcal{S}$. We claim that

$$\begin{aligned}
\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau) - \mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau) &\le 2\Big(\sum_{i=1}^{r} \sum_{\tau_i \in \mathcal{T}(i,\tau) \cap \mathcal{S}_i} \mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau_i)\Big) + 2\delta_1 \sum_{\tau_1 \in \mathcal{T}(1,\tau)} \mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau_1) \\
&= 2\Big(\sum_{i=1}^{r} \sum_{\tau_i \in \mathcal{T}(i,\tau) \cap \mathcal{S}_i} \mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau_i)\Big) + 2\delta_1 \cdot \mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau) \\
&\le 2\Big(\sum_{i=1}^{r} \sum_{\tau_i \in \mathcal{T}(i,\tau) \cap \mathcal{S}_i} \mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau_i)\Big) + 2\delta_1 \cdot \mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau),
\end{aligned} \tag{10}$$

where the last inequality is due to the assumption that τ is valid. This claim will be justified later. By summing both sides of Equation (10) over all valid τ, we can bound the left-hand side of Equation (7) by

$$2\Big(\sum_{i=1}^{r} \Pr[\text{Script}(B_i, \mathbf{S}_{\text{ideal}}) \in \mathcal{S}]\Big) + 2\delta_1 \le 2q \cdot \delta_2 + 2\delta_1$$

which is the right-hand side of Equation (7). To justify Equation (10), note that

$$\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau) - \mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau) = \sum_{i=1}^{r} \mathsf{p}_{\mathbf{S}_i}(\tau) - \mathsf{p}_{\mathbf{S}_{i-1}}(\tau) \ .$$

Moreover, for each $i \leq r$,

$$\mathsf{ps}_{\mathbf{S}_i}(\tau) = \sum_{\tau_i \in \mathcal{T}(i,\tau)} \mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\tau_i),$$

whereas

$$\mathsf{ps}_{\mathbf{S}_{i-1}}(\tau) \geq \sum_{\tau_i \in \mathcal{T}(i,\tau)} \mathsf{ps}_{\mathbf{S}_{\text{real}}}(\tau_i),$$

because (a) the left-hand side is the chance that adversary B_i in its real world (recall that the real world of B_i is the ideal world of B_{i-1}) can generate τ, which is $\sum_{\tau'} \mathsf{ps}_{\mathbf{S}_{\text{real}}}(\tau')$ over all enhanced transcripts τ' that B_i can witness such that the corresponding transcript for A is τ, and (b) the right-hand side is $\sum_{\tau'} \mathsf{ps}_{\mathbf{S}_{\text{real}}}(\tau')$ over *some* (but probably not all) such τ'. Hence

$$\begin{aligned}
&\mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\tau) - \mathsf{ps}_{\mathbf{S}_{\text{real}}}(\tau) \leq \sum_{i=1}^{r} \sum_{\tau_i \in \mathcal{T}(i,\tau)} \mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\tau_i) - \mathsf{ps}_{\mathbf{S}_{\text{real}}}(\tau_i) \\
&= \Big(\sum_{i=1}^{r} \sum_{\tau_i \in \mathcal{T}(i,\tau) \cap \mathcal{S}_i} \mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\tau_i) - \mathsf{ps}_{\mathbf{S}_{\text{real}}}(\tau_i)\Big) + \sum_{i=1}^{r} \sum_{\tau_i \in \mathcal{T}(i,\tau) \setminus \mathcal{S}_i} \mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\tau_i) - \mathsf{ps}_{\mathbf{S}_{\text{real}}}(\tau_i) \\
&\leq \Big(\sum_{i=1}^{r} \sum_{\tau_i \in \mathcal{T}(i,\tau) \cap \mathcal{S}_i} \mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\tau_i)\Big) + \sum_{i=1}^{r} \sum_{\tau_i \in \mathcal{T}(i,\tau) \setminus \mathcal{S}_i} \mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\tau_i) - \mathsf{ps}_{\mathbf{S}_{\text{real}}}(\tau_i) \ .
\end{aligned}$$

What is left is to prove that

$$\begin{aligned}
&\sum_{i=1}^{r} \sum_{\tau_i \in \mathcal{T}(i,\tau) \setminus \mathcal{S}_i} \mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\tau_i) - \mathsf{ps}_{\mathbf{S}_{\text{real}}}(\tau_i) \\
&\leq \Big(\sum_{i=1}^{r} \sum_{\tau_i \in \mathcal{T}(i,\tau) \cap \mathcal{S}_i} \mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\tau_i)\Big) + 2\delta_1 \sum_{\tau_1 \in \mathcal{T}(1,\tau)} \mathsf{ps}_{\mathbf{S}_{\text{real}}}(\tau_1) \ . \qquad (11)
\end{aligned}$$

Now, recall that for each $\tau_i \in \mathcal{T}(i,\tau) \setminus \mathcal{S}_i$, the su transcript $\mathrm{Tr}(\tau_i)$ is good. Since the two systems satisfy the almost proximity condition,

$$\mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\mathrm{Tr}(\tau_i)) - \mathsf{ps}_{\mathbf{S}_{\text{real}}}(\mathrm{Tr}(\tau_i)) \leq \mathrm{Rate}(\mathrm{Tr}(\tau_i)) \cdot \mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\mathrm{Tr}(\tau_i)) \ .$$

Recall that from Equation (9), the ratio between $\mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\mathrm{Tr}(\tau_i))$ and $\mathsf{ps}_{\mathbf{S}_{\text{real}}}(\mathrm{Tr}(\tau_i))$ is exactly that between $\mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\tau_i)$ and $\mathsf{ps}_{\mathbf{S}_{\text{real}}}(\tau_i)$. Then

$$\mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\tau_i) - \mathsf{ps}_{\mathbf{S}_{\text{real}}}(\tau_i) \leq \mathrm{Rate}(\mathrm{Tr}(\tau_i)) \cdot \mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\tau_i) \ . \qquad (12)$$

This in turn implies that

$$\mathsf{ps}_{\mathbf{S}_{\text{ideal}}}(\tau_i) \leq \frac{\mathsf{ps}_{\mathbf{S}_{\text{real}}}(\tau_i)}{1 - \mathrm{Rate}(\mathrm{Tr}(\tau_i))} \ . \qquad (13)$$

To justify that the denominator of the right-hand side is nonzero so that Equation (13) above is well-defined, let τ' be the mu transcript that has the same

Cons queries/answers as τ, and the same Prim queries/answers as τ_i. Since τ is nice, so is τ'. Thus, $1 - \mathrm{Rate}(\mathrm{Tr}(\tau_i)) = 1 - \mathrm{Rate}(\mathrm{Map}(i, \tau')) \geq 1 - \delta_1 > 0$. From Equation (12), to justify Equation (11), we need to bound each sum

$$\sum_{\tau_i \in \mathcal{T}(i,\tau) \setminus \mathcal{S}_i} \mathrm{Rate}(\mathrm{Tr}(\tau_i)) \cdot \mathsf{p}_{\mathbf{S}_{\mathrm{ideal}}}(\tau_i),$$

for every $i \in \{1, \ldots, r\}$. For $\ell \leq i$, define $\mathrm{Rate}(i, \tau_\ell)$ as follows. Let τ' be the su transcript induced by τ_ℓ in which we only keep Cons queries/answers for user u_i, and all Prim queries/answers. Let $\mathrm{Rate}(i, \tau_\ell) = \mathrm{Rate}(\tau')$. The special case $\mathrm{Rate}(i, \tau_i)$ coincides with $\mathrm{Rate}(\mathrm{Tr}(\tau_i))$. We claim that for each i, the sum above is at most

$$\sum_{\tau_1 \in \mathcal{T}(1,\tau)} 2\mathrm{Rate}(i, \tau_1) \cdot \mathsf{p}_{\mathbf{S}_{\mathrm{real}}}(\tau_1) + \sum_{s=1}^{i} \sum_{\tau_s \in \mathcal{T}(s,\tau) \cap \mathcal{S}_s} 2\mathrm{Rate}(i, \tau_s) \cdot \mathsf{p}_{\mathbf{S}_{\mathrm{ideal}}}(\tau_s) \ . \quad (14)$$

Note that for any $s \geq 1$ and any $\tau_s \in \mathcal{T}(s, \tau)$, if we let τ' be the mu transcript that has the same Cons queries/answers as τ, and the same Prim queries/answers as τ_s, then τ' is also nice, because τ is nice. Then

$$\sum_{i=s}^{r} \mathrm{Rate}(i, \tau_s) = \sum_{i=s}^{r} \mathrm{Rate}(\mathrm{Map}(i, \tau')) \leq \delta_1 \ . \quad (15)$$

From Equation (15),

$$\sum_{i=1}^{r} \sum_{\tau_1 \in \mathcal{T}(1,\tau)} 2\mathrm{Rate}(i, \tau_1) \cdot \mathsf{p}_{\mathbf{S}_{\mathrm{real}}}(\tau_1) \leq \sum_{\tau_1 \in \mathcal{T}(1,\tau)} 2\delta_1 \cdot \mathsf{p}_{\mathbf{S}_{\mathrm{real}}}(\tau_1), \quad (16)$$

and

$$\begin{aligned} &\sum_{i=1}^{r} \sum_{s=1}^{i} \sum_{\tau_s \in \mathcal{T}(s,\tau) \cap \mathcal{S}_s} 2\mathrm{Rate}(i, \tau_s) \cdot \mathsf{p}_{\mathbf{S}_{\mathrm{ideal}}}(\tau_s) \quad (17) \\ &= \sum_{s=1}^{r} \sum_{\tau_s \in \mathcal{T}(s,\tau) \cap \mathcal{S}_s} \sum_{i=s}^{r} 2\mathrm{Rate}(i, \tau_s) \cdot \mathsf{p}_{\mathbf{S}_{\mathrm{ideal}}}(\tau_s) \\ &\leq \sum_{s=1}^{r} \sum_{\tau_s \in \mathcal{T}(s,\tau) \cap \mathcal{S}_s} 2\delta_1 \cdot \mathsf{p}_{\mathbf{S}_{\mathrm{ideal}}}(\tau_s) \leq \sum_{s=1}^{r} \sum_{\tau_s \in \mathcal{T}(s,\tau) \cap \mathcal{S}_s} \mathsf{p}_{\mathbf{S}_{\mathrm{ideal}}}(\tau_s) \ . \quad (18) \end{aligned}$$

Combining Equations (12), (14), (16), and (18) gives us Equation (11).

To justify Equation (14), fix $i \in \{1, \ldots, r\}$. We create a binary tree whose weight at the root is exactly the sum above for i. In this tree, for any two children of a node, the left one must be a leaf node. Moreover, we will put weights on the nodes so that the weight of a parent node is bounded by the sum of the weights of its children. Hence the weight at the root is bounded by the total weight of the leaves.

Starting at the root, from Equation (13), we can bound the weight at the root by a linear combination of $\mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau_i)$, where $\tau_i \in \mathcal{T}(i,\tau)\backslash\mathcal{S}_i$. For each such τ_i, if we enhance it with the key of user u_i and the internal PRIM queries/answers due to the CONS queries of user u_i then we will get an extended transcript τ_{i-1} for adversary B_{i-1}. (Recall that the real world of B_i is the ideal world of B_{i-1}.) Hence the linear combination of $\mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau_i)$ becomes a linear combination of $\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau_{i-1})$, for $\tau_{i-1} \in \mathcal{T}(i-1,\tau)$. We divide this into two parts, one for $\tau_{i-1} \in \mathcal{S}_{i-1}$, and another for $\tau_{i-1} \notin \mathcal{S}_{i-1}$. The first partial sum will be the weight of the left child of the root, and the second the weight of the right child. So far, we have placed the weights up to the second level of the tree. We will repeat the process above, starting at the right child of the root, until we reach the i-th level. At that point, the weight of the right-most leaf is a linear combination of $\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau_1)$, for $\tau_1 \in \mathcal{T}(1,\tau)$.

Recall that the weight of each node of the binary tree above is a linear combination. We now specify the coefficients. At the root, each coefficient for $\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau_i)$ is $\text{Rate}(i,\tau_i)$. We will have to bound $\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau_i)$ via $\mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau_i)$ by Equation (13), so the coefficients for the left and right children of the root are at most

$$\frac{\text{Rate}(i,\tau_i)}{1-\text{Rate}(i,\tau_i)} \leq \frac{\text{Rate}(i,\tau_{i-1})}{1-\text{Rate}(\tau_{i-1})},$$

where the inequality is due to the fact that Rate is increasing and τ_{i-1} contains all queries/answers of τ_i, and thus $\text{Rate}(i-1,\tau_{i-1}) \geq \text{Rate}(i,\tau_i)$. By repeating this process, for nodes at the $(i+1-s)$-th level, the coefficients are at most

$$\frac{\text{Rate}(i,\tau_s)}{\prod_{\ell=s+1}^{i}\big(1-\text{Rate}(\ell,\tau_s)\big)} \ .$$

Now, for the right most leaf, its weight is currently a linear combination of $\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau_1)$, but we want to have its weight as a linear combination of $\mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau_1)$ instead. To achieve this, we will again use Equation (13) (but i replaced by 1), and the new coefficients for this leaf are at most

$$\frac{\text{Rate}(i,\tau_1)}{\prod_{\ell=1}^{i}\big(1-\text{Rate}(\ell,\tau_1)\big)} \ .$$

Hence the coefficients for a leaf at the $(i+1-s)$-th level of the tree are at most

$$\frac{\text{Rate}(i,\tau_s)}{\prod_{\ell=s}^{i}\big(1-\text{Rate}(\ell,\tau_s)\big)} \leq \frac{\text{Rate}(i,\tau_s)}{1-\sum_{\ell=s}^{i}\text{Rate}(\ell,\tau_s)} \leq \frac{\text{Rate}(i,\tau_s)}{1-\delta_1} \leq 2\text{Rate}(i,\tau_s),$$

where the first inequality is due to the fact that $(1-x)(1-y) \geq 1-x-y$ for any $0 \leq x,y < 1$, and the second inequality is due to Equation (15). The total weight of the leaves therefore is at most

$$\sum_{\tau_1 \in \mathcal{T}(1,\tau)} 2\text{Rate}(i,\tau_1)\cdot \mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau_1) + \sum_{s=1}^{i}\sum_{\tau_s \in \mathcal{T}(s,\tau)\cap\mathcal{S}_s} 2\text{Rate}(i,\tau_s)\cdot \mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau_s) \ .$$

This concludes the proof. □

4 Simplification of the Framework for Specific Settings

Since the framework in Section 3 aims to provide an umbrella for *all* settings, it appears unnecessarily complex in many important settings. To improve the usability of our framework, in this section, we consider some simplified treatments of our general framework for specific settings. Each such specialized result is somewhat more limited in scope, but simpler to use.

4.1 A simple specialization of the framework

We now describe a specialization of the framework that is very simple, but might be powerful enough for typical real-world cryptographic schemes, such as the authenticated encryption scheme GCM [23]. This simple treatment however is not enough for Double Encryption, and thus in the next subsection, we will consider another specialized result of the general framework to handle Double Encryption.

THE SETTING. Here we still use the generic setting as stated in Section 3, but make an assumption on the metric σ. For a mu transcript τ and each user u_i of τ, let $\mathrm{Map}(i,\tau)$ be the induced su transcript for user u_i that consists of the $\textsc{Cons}(i,\cdot)$ queries/answers and $\textsc{Prim}(\cdot)$ queries/answers of τ. We require that for any mu transcript τ, if the $\textsc{Cons}$ queries in τ have data complexity σ, and those in each $\mathrm{Map}(i,\tau)$ have data complexity σ_i, then

$$\sum_i \sigma_i \le \sigma \ .$$

This requirement obviously holds if we let, for example, σ be the total length of the $\textsc{Cons}$ queries.

SUPER-ADDITIVITY. For a function $\delta : (\mathbb{N})^3 \to [0,1]$, we say that it is *super-additive* if

$$\delta(x,y_0,z_0) + \delta(x,y_1,z_1) \le \delta(x,y_0+y_1,z_0+z_1)$$

for every $x,y_0,y_1,z_0,z_1 \in \mathbb{N}$. In many schemes, the desired bounds (such as $\delta(p,q,\sigma) = \sigma^2/2^n$) are often super-additive.

THE TECHNIQUE. One begins by defining an undesirable property on su transcripts that involves only $\textsc{Cons}$ queries/answers. If a su transcript has this property then we say that it is *bad*, otherwise it is *good*.[3] A mu transcript τ is *nice* if there is no user u_i such that its induced su transcript $\mathrm{Map}(i,\tau)$ is bad. Let $\mathcal{N}$ be the set of nice mu transcripts. We require that there be a monotonic function δ

[3] In Section 3, we partitioned the set of bad su transcripts into $\mathcal{S}$ and $\mathcal{S}'$, and required that it is unlikely for the adversary to produce a bad transcript in $\mathcal{S}$. Here $\mathcal{S}$ is simply the empty set.

such that for any adversary A making p PRIM queries and q Cons queries of data complexity σ,

$$\Pr[\mathrm{Script}(A, \mathbf{S}_{\mathrm{ideal}}) \notin \mathcal{N}] \leq \delta(p, q, \sigma), \tag{19}$$

where for any system $\mathbf{S}$, $\mathrm{Script}(A, \mathbf{S})$ denotes the random variable for the transcript of the interaction of A and $\mathbf{S}$. Moreover, we require that there be a monotonic function ϵ' and a super-additive, monotonic function ϵ such that for any good su transcript τ of p PRIM queries and q CONS queries of data complexity σ,

$$\mathsf{ps}_{\mathrm{ideal}}(\tau) - \mathsf{ps}_{\mathrm{real}}(\tau) \leq (\epsilon(p, q, \sigma) + \epsilon'(p, q, \sigma)) \cdot \mathsf{ps}_{\mathrm{ideal}}(\tau) \ . \tag{20}$$

Lemma 3. *Assume that the systems* $\mathbf{S}_{\mathrm{real}}$ *and* $\mathbf{S}_{\mathrm{ideal}}$ *meet the conditions in Equations* (19) *and* (20)*. Then*

$$\mathsf{Adv}^{\mathrm{dist}}_{\mathbf{S}_{\mathrm{real}}, \mathbf{S}_{\mathrm{ideal}}}(A) \leq \delta(p, q, \sigma) + 2\epsilon(p + t\sigma, q, \sigma) + 2q \cdot \epsilon'(p + t\sigma, q, \sigma) \ .$$

Proof. For a su transcript τ of p PRIM queries and q CONS queries of data complexity σ, let

$$\mathrm{Rate}(\tau) = \epsilon(p, q, \sigma) + \epsilon'(p, q, \sigma) \ .$$

This function Rate is increasing, in the sense that if τ' contains all the query-answer pairs of τ, then $\mathrm{Rate}(\tau') \geq \mathrm{Rate}(\tau)$. To use Lemma 2, we need to establish the mu-boundedness of Rate. We claim that for any nice mu transcript τ of r users, using p PRIM queries and q CONS queries of data complexity σ,

$$\sum_{i=1}^{r} \mathrm{Rate}(\mathrm{Map}(i, \tau)) \leq \epsilon(p, q) + q\epsilon'(p, q) \ .$$

To justify this, suppose that τ_i contains q_i CONS queries of data-complexity σ_i. Then

$$\begin{aligned}
\sum_{i=1}^{r} \mathrm{Rate}(\mathrm{Map}(i, \tau)) &= \sum_{i=1}^{r} \epsilon(p, q_i, \sigma_i) + \epsilon'(p, q_i, \sigma_i) \\
&\leq \sum_{i=1}^{r} \epsilon(p, q_i, \sigma_i) + \epsilon'(p, q, \sigma) \\
&\leq \epsilon(p, q, \sigma) + r \cdot \epsilon'(p, q, \sigma) \leq \epsilon(p, q, \sigma) + q \cdot \epsilon'(p, q, \sigma) \ .
\end{aligned}$$

Finally, applying Lemma 2 for $\delta_0 = \delta$, $\delta_1 = \epsilon + q\epsilon'$, and $\delta_2 = 0$, leads to the claimed advantage. □

4.2 The specialized framework for Double Encryption and beyond

We now specialize the general framework into a more specific result that covers the case of Single Encryption, Double Encryption, and Key-Alternating Cipher (KAC) [11]. This result explains why these constructions, despite being somewhat similar in the structure, have different blowups when we move from su setting to mu one.

THE SETTING. Let $\Pi[E] : \mathcal{K} \times \{0,1\}^n \times \{0,1\}^n$ be a blockcipher construction built on top of an ideal blockcipher $E : \{0,1\}^k \times \{0,1\}^n \to \{0,1\}^n$ such that a single call to Π/Π^{-1} makes at most t calls to E/E^{-1}. Let $\mathbf{S}_{\text{real}}$ and $\mathbf{S}_{\text{ideal}}$ be two stateless systems implementing games $\text{Real}_{\Pi[E],\text{Sample}}$ and $\text{Rand}_{\Pi[E],\text{Sample}}$ in Fig. 1, respectively. We will measure adversaries' resources in terms of q (the number of ENC/DEC queries) and p (the number of PRIM/PRIMINV queries). A transcript recording the interaction between an adversary and a system $\mathbf{S} \in \{\mathbf{S}_{\text{ideal}}, \mathbf{S}_{\text{real}}\}$ contains the following:

- ENC/DEC queries: A query to $\text{ENC}(i,x)$ returning y is associated with an entry $(\mathsf{enc}, +, i, x, y)$. Likewise, a query to $\text{DEC}(i,y)$ returning x is associated with an entry $(\mathsf{enc}, -, i, x, y)$.
- PRIM/PRIMINV queries: A query to $\text{PRIM}(J,u)$ returning v is recorded in the transcript as $(\mathtt{prim}, +, J, u, v)$. Likewise, a query to $\text{PRIMINV}(J,v)$ returning u is associated with an entry $(\mathtt{prim}, -, J, u, v)$.

SUPER-ADDITIVITY AND BEYOND. For a function $\delta : (\mathbb{N})^2 \to [0,1]$, we say that it is *super-additive* if $\delta(x,y) + \delta(x,z) \le \delta(x, y+z)$, for every $x, y, z \in \mathbb{N}$. For real numbers $M > 0$ and $z \ge 0$, let $\text{Cost}(M,z) = \max\{M, z\}$ if $z > 1$, and $\text{Cost}(M,z) = M/\lg(M)$ if $z \le 1$.

THE TECHNIQUE. One begins by defining an undesirable property on su transcripts, which can involve both ENC/DEC and PRIM/PRIMINV queries/answers. If a su transcript has this property, we say that it is *bad*; otherwise it is *good*. Let $\mathcal{S}$ be the set of all bad su transcripts.[4] If a su transcript is not bad, we say that it is *good*. We demand that there be a monotonic function ϵ^* such that for any su adversary A that makes at most q ENC/DEC queries and p PRIM/PRIMINV queries,

$$\Pr[\text{Script}(A, \mathbf{S}_{\text{ideal}}) \in \mathcal{S}] \le \epsilon^*(p,q) \tag{21}$$

where for any system $\mathbf{S}$, $\text{Script}(A, \mathbf{S})$ denotes the random variable for the transcript of the interaction of A and $\mathbf{S}$.

For any transcript τ in which the adversary attacks just a single user, let $\text{Ent}(\tau)$ be the number of entries $(\mathtt{prim}, \cdot, \cdot, u, v)$ such that τ contains either an entry $(\mathsf{enc}, +, 1, \cdot, x)$ or an entry $(\mathsf{enc}, -, 1, x, \cdot)$, for some $x \in \{u, v\}$. Suppose that there are monotonic functions ϵ', ϵ'' and a monotonic, super-additive function ϵ such that, for any good su transcript of q queries to ENC/DEC, and p queries to PRIM/PRIMINV,

$$\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau) - \mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau) \le (\epsilon(p,q) + \epsilon'(p,q) \cdot \text{Ent}(\tau) + \epsilon''(p,q)) \cdot \mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau) \ . \tag{22}$$

If Equations (21) and (22) are met, then we say that $\Pi[E]$ has the $(\epsilon, \epsilon', \epsilon'', \epsilon^*)$-proximity property.

[4] In Section 3, we factored the set of bad su transcripts into two disjoint sets $\mathcal{S}$ and $\mathcal{S}'$, and required that it is unlikely for the adversary to produce a bad transcript in $\mathcal{S}$. Here $\mathcal{S}'$ is simply the empty set.

Note that $\mathrm{Ent}(\tau) \leq \min\{p, 2^{k+2}q\}$, where k is the key length of the primitive E. Thus $(\epsilon, \epsilon', \epsilon'', \epsilon^*)$-proximity immediately implies that for any adversary attacking a single user via q Enc/Dec queries and p Prim/PrimInv queries, its su advantage is at most $\epsilon(p,q) + \epsilon'(p,q) \cdot \min\{p, q \cdot 2^{k+2}\} + \epsilon''(p,q) + \epsilon^*(p,q)$. The following result bounds the mu security of $\Pi[E]$.

Lemma 4. *Assume that $\Pi[E]$ has the $(\epsilon, \epsilon', \epsilon'', \epsilon^*)$-proximity property. Then for any adversary A that makes at most q Enc/Dec queries, and p Prim/PrimInv queries,*

$$\mathsf{Adv}^{\pm\text{mu-prp}}_{\Pi[E],\mathrm{Sample}}(A) \leq 2^{-n} + 2\epsilon + 2q(\epsilon'' + \epsilon^*) + \mathrm{Cost}(4n, 8q/2^n) \cdot 10(p + qt)\epsilon',$$

where t is the number of calls to E/E^{-1} that a single call to Π/Π^{-1} makes, and functions $\epsilon, \epsilon', \epsilon'', \epsilon^*$ all take arguments $p + qt$ and q. □

Discussion. Recall that our technique dissects a su bound into three components: ϵ, $\epsilon' \cdot \min\{p, q \cdot 2^{k+2}\}$, and $(\epsilon'' + \epsilon^*)$. Lemma 4 above then lifts those to ϵ, $\mathrm{Cost}(4n, 8q/2^n) \cdot (p + qt) \cdot \epsilon'$, and $q \cdot (\epsilon'' + \epsilon^*)$, respectively, for the corresponding mu bound. This trisection captures different possibilities of security loss when one moves from su to mu security: (i) Key-Alternating Cipher (where ϵ is the dominant term in both the su and mu bounds) [19], (ii) Single Encryption (where $\epsilon'' + \epsilon^*$ and $q \cdot (\epsilon'' + \epsilon^*)$ are the dominant term in the su and mu bounds respectively), and (iii) Double Encryption (where $\epsilon' \cdot \min\{p, q \cdot 2^{k+2}\}$ and $\mathrm{Cost}(4n, 8q/2^n) \cdot (p + qt)\epsilon'$ are the dominant term in the su and mu bounds respectively).

Given a su analysis, there might be multiple choices for ϵ and ϵ''. However, recall that when we move from su to mu security, the former term remains the same, whereas the latter blows up with a factor q. Therefore, when we need to pinpoint ϵ and ϵ'', we will shift as much weight to ϵ as possible, and the optimal choice of ϵ will often be clear from the context and the best mu attacks. On the other hand, due to the q-blowup of ϵ'', one may need a very fine-grained su analysis to obtain a good mu bound.

The proof of Lemma 4. We want to show that Lemma 2 implies the claimed result. In order to do that, we need to define (i) function $\mathrm{Rate}(\tau)$ for su transcripts τ, and (ii) a *niceness* property for mu transcripts. The former is obvious: for a su transcript τ of p Prim/PrimInv queries and q Enc/Dec queries, let

$$\mathrm{Rate}(\tau) = \epsilon(p,q) + \epsilon'(p,q) \cdot \mathrm{Ent}(\tau) + \epsilon''(p,q) \ .$$

This function Rate is increasing, in the sense that if τ' contains all the query-answer pairs of τ then $\mathrm{Rate}(\tau') \geq \mathrm{Rate}(\tau)$. Next, let $d = \frac{5}{4}\mathrm{Cost}(4n, 8q/2^n)$. We say that a mu transcript τ in the support of $\mathrm{Script}(A, \mathbf{S}_{\mathrm{ideal}})$ is *nice* if it satisfies the following constraints:

- There are no d entries in τ of the form $(\mathsf{enc}, +, \cdot, \cdot, y), \ldots, (\mathsf{enc}, +, \cdot, \cdot, y)$.
- There are no d entries in τ of the form $(\mathsf{enc}, -, \cdot, x, \cdot), \ldots, (\mathsf{enc}, -, \cdot, x, \cdot)$.

Clearly, the definition of niceness involves only Enc/Dec query-answer pairs of τ. Let $\mathcal{N}$ be the set of nice mu transcripts. The following bounds the chance that A's transcript is not nice; the proof is in the full version of this paper.

Lemma 5. *For any adversary A that makes at most q Enc/Dec queries, and p Prim/PrimInv queries,*

$$\Pr[\mathrm{Script}(A, \mathbf{S}_{\mathrm{ideal}}) \notin \mathcal{N}] \leq \frac{1}{2^n} \ . \quad \square$$

To use Lemma 2, we need to establish the mu-boundedness of Rate. Specifically, we claim that, for any nice mu transcript τ of r users, using p Prim/PrimInv queries and q Enc/Dec queries,

$$\sum_{i=1}^{r} \mathrm{Rate}(\mathrm{Map}(i,\tau)) \leq \epsilon(p,q) + q\epsilon''(p,q) + 4dp\epsilon'(p,q) \ . \tag{23}$$

Then using Lemma 2 for $\delta_0 = 2^{-n}$ and $\delta_1 = \epsilon + q\epsilon'' + 4dp\epsilon'$ and $\delta_2 = \epsilon^*$ leads to our claimed result.

We now justify Equation (23). Suppose that in τ, the adversary uses q_i Enc/Dec queries for the i-th user. Then

$$\begin{aligned}
\sum_{i=1}^{r} \mathrm{Rate}(\mathrm{Map}(i,\tau)) &= \sum_{i=1}^{r} \Big(\epsilon(p,q_i) + \epsilon''(p,q_i) + \mathrm{Ent}(\mathrm{Map}(i,\tau)) \cdot \epsilon'(p,q_i) \Big) \\
&\leq \epsilon(p,q) + r\epsilon''(p,q) + \sum_{i=1}^{r} \mathrm{Ent}(\mathrm{Map}(i,\tau)) \cdot \epsilon'(p,q) \\
&\leq \epsilon(p,q) + q\epsilon''(p,q) + \sum_{i=1}^{r} \mathrm{Ent}(\mathrm{Map}(i,\tau)) \cdot \epsilon'(p,q),
\end{aligned}$$

where the first inequality is due to the superadditivity of ϵ and the monotone of ϵ' and ϵ''. Thus to justify (23), what's left is to prove that

$$\sum_{i=1}^{r} \mathrm{Ent}(\mathrm{Map}(i,\tau)) \leq 4dp \ .$$

Since τ is nice, for each entry $(\mathtt{prim}, \cdot, \cdot, u, v)$, there are at most $4d$ entries $(\mathsf{enc}, \cdot, \cdot, \cdot, x)$ or $(\mathsf{dec}, \cdot, \cdot, x, \cdot)$, for $x \in \{u, v\}$. Since each Enc/Dec entry belongs to exactly one user, for each Prim/PrimInv entry of τ, there are at most $4d$ indices i such that $\mathrm{Ent}(\mathrm{Map}(i,\tau))$ counts this entry, and thus summing over p Prim/PrimInv entries of τ gives us

$$\sum_{i=1}^{r} \mathrm{Ent}(\mathrm{Map}(i,\tau)) \leq 4dp$$

as claimed.

5 Exact Multi-user Security of Double Encryption

5.1 Results and Discussion

RESULTS. In this section, we give an exact mu security bound of Double Encryption via the specialized framework in Section 4.2; the key-sampling algorithm is uniform. While it is relatively easy to give an exact su security bound of Double Encryption [2,14], giving a good $(\epsilon, \epsilon', \epsilon'', \epsilon^*)$-proximity bound, as in Lemma 6 below, requires a much more fine-grained analysis. The proof, given in Section 5.2, is based on the expectation method of Hoang and Tessaro [19].

Lemma 6. *Let $n \geq 16$ and $k \geq 1$ be integers, and let $E : \{0,1\}^k \times \{0,1\}^n \to \{0,1\}^n$ be a blockcipher. Then $\mathrm{DE}[E]$ satisfies the $(\epsilon, \epsilon', \epsilon'', \epsilon^*)$-proximity property, with $\epsilon(p,q) = \frac{2q}{2^{k+n/2}} + \frac{3qB^2+2Bpq}{2^{2k}}$, $\epsilon'(p,q) = \frac{2p}{2^{2k}}$, $\epsilon''(p,q) = \frac{5Bp}{2^{2k}}$, and $\epsilon^*(p,q) = \frac{1}{2^{k+n}}$, where $B = \frac{5}{4} \cdot \mathrm{Cost}(4n + 2k, 8q/2^n)$.* □

From Lemma 4 and Lemma 6, we immediately obtain the following result.

Theorem 1 (Mu security of Double Encryption). *Let $n, k \in \mathbb{N}$ be integers, and let $E : \{0,1\}^k \times \{0,1\}^n \to \{0,1\}^n$ be a blockcipher. Then for any adversary making only q ENC/DEC queries and p PRIM/PRIMINV queries,*

$$\mathsf{Adv}^{\pm\text{mu-prp}}_{\mathrm{DE}[E]}(A) \leq \frac{1}{2^n} + \frac{5q}{2^{k+n/2}} + \frac{6qB^2 + 222BQ^2}{2^{2k}}$$

where $B = \frac{5}{4} \cdot \mathrm{Cost}(4n + 2k, 8q/2^n)$ and $Q = \max\{p, q\}$. □

DISCUSSION. Admittedly, the bound in Theorem 1 looks complicated. However, for the "usual" setting $n \geq k \geq 16$ and $q \leq \frac{2^k}{8}$, the bound can be simplified to $\mathsf{Adv}^{\pm\text{mu-prp}}_{\mathrm{DE}[E]}(A) \leq \frac{1}{2^n} + \frac{(n+5)q}{2^{1.5k}} + \frac{1554nQ^2}{\lg(4n)\cdot 2^{2k}}$. On the other hand, recall that the classical su bound of $\mathrm{DE}[E]$ by Aiello et al. [2] is $\mathsf{Adv}^{\pm\text{prp}}_{\mathrm{DE}[E]}(A) \leq \frac{p^2}{2^{2k}}$. If we apply the hybrid argument to this, we will get the following inferior bound $\mathsf{Adv}^{\pm\text{mu-prp}}_{\mathrm{DE}[E]}(A) \leq \frac{q(p+2q)^2}{2^{2k}}$. While this bound is enough to show that Double Encryption squarely beats Single Encryption in mu security,[5] it is much worse than the bound in Theorem 1, as illustrated in Fig. 2.

5.2 Proof of Lemma 6

It is convenient to assume without loss of generality that the adversary doesn't make redundant queries. Our proof borrows the overall approach used by Hoang and Tessaro [19] for key-alternating ciphers. We begin with some high-level setup.

[5] Recall that $\mathsf{Adv}^{\pm\text{prp}}_{E}(A) \leq \frac{p}{2^k}$ and $\mathsf{Adv}^{\pm\text{mu-prp}}_{E}(A) \leq \frac{p(p+q)}{2^k}$ for an adversary A making only q ENC/DEC queries and p PRIM/PRIMINV queries.

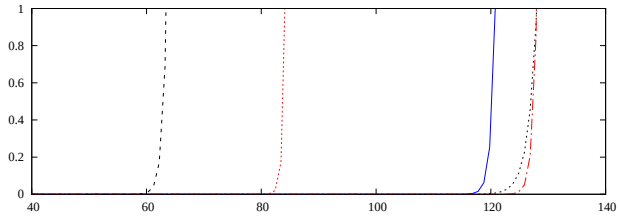

Fig. 2: **Mu and su security of Single and Double Encryption on AES.** From left to right: the mu bound of Single Encryption, the naive mu bound of Double Encryption via the hybrid argument, the mu bound of Double Encryption via Theorem 1, the su bound of Single Encryption, and the classical su bound of Double Encryption by Aiello et al. [2]. We set $p = q$ and $n = k = 128$. The x-axis gives the log (base 2) of p, and the y-axis gives the security bounds.

ASSUMPTIONS ON THE TRANSCRIPT. We consider an arbitrary fixed transcript τ which contains q ENC / DEC queries and p PRIM/PRIMINV queries. Moreover, for a transcript τ, we also denote (following [14])

$$\mathrm{Fwd}(\tau) = \max_{y \in \{0,1\}^n} \left|\{(J, x) \mid (\mathtt{prim}, +, J, x, y) \in \tau\}\right| ,$$
$$\mathrm{Bwd}(\tau) = \max_{x \in \{0,1\}^n} \left|\{(J, y) \mid (\mathtt{prim}, -, J, x, y) \in \tau\}\right| .$$

Recall that to establish $(\epsilon, \epsilon', \epsilon'', \epsilon^*)$-proximity, we have to define bad transcripts. A transcript is bad if either $\mathrm{Fwd}(\tau) > B$ or $\mathrm{Bwd}(\tau) > B$, where

$$B := \frac{5}{4} \cdot \mathrm{Cost}(4n + 2k, 8p/2^n) .$$

Let $\mathcal{S}$ be the set of all bad transcripts. The following bounds the chance that the adversary produces a bad transcript; the proof is in the full version of this paper.

Lemma 7. *For any adversary A that makes p PRIM/PRIMINV queries and q ENC/DEC queries,*

$$\Pr[\mathrm{Script}(A, \mathbf{S}_{\mathrm{ideal}}) \in \mathcal{S}] \leq \frac{1}{2^{n+k}} \quad . \quad \square$$

From now on, we assume that additionally $\tau \notin \mathcal{S}$. We shall use the expectation method to prove the claimed bound of the gap $\mathsf{p}\mathbf{s}_{\mathbf{S}_{\mathrm{ideal}}}(\tau) - \mathsf{p}\mathbf{s}_{\mathbf{S}_{\mathrm{real}}}(\tau)$. We begin with some combinatorial results on the transcript.

TYPE-1 CHAINS. Consider a pair of entries $(\mathtt{prim}, \cdot, \cdot, x_1, y_1), (\mathtt{prim}, \cdot, \cdot, x_2, y_2)$ in τ such that $y_1 = x_2$. We say that it is a *positive type-1 chain* if there's an entry $(\mathsf{enc}, +, x_1, \cdot)$ in τ. We say that it is a *negative type-1 chain* if there's an entry $(\mathsf{enc}, -, \cdot, y_2)$. The following lemma bounds the number of type-1 chains; the proof is in Appendix A.

Lemma 8. *The number of type-1 chains is at most* $4Bp + 2B^2q + 2Bpq$. □

TYPE-2 CHAINS. Consider a pair of entries $(\mathtt{prim}, \cdot, \cdot, x_1, y_1), (\mathtt{prim}, \cdot, \cdot, x_2, y_2)$. We say that it is a *positive type-2 chain* if there's an entry $(\mathsf{enc}, +, x_1, y_2)$ in τ. We say that it is a *negative type-2 chain* if there's an entry $(\mathsf{enc}, -, x_1, y_2)$ in τ. The following lemma bounds the number of type-2 chains; the proof is in Appendix B.

Lemma 9. *The number of type-2 chains is at most* $2p \cdot \mathrm{Ent}(\tau)$. □

GOOD AND BAD KEYS. We shall use the expectation method. Let S be the random variable for the key. The key-space $\mathcal{K}$ is $(\{0,1\}^k)^2$ and S is uniformly distributed over $\mathcal{K}$. For each key vector $s = (K_1, K_2) \in \mathcal{K}$ and each $i \in \{1, 2\}$, let $p_i[s]$ be the number of queries $(\mathtt{prim}, \cdot, K_i, \cdot, \cdot)$ in τ.

Definition 1 (Good and bad keys). *We say that a key* $s = (K_1, K_2)$ *is* bad *if one the following happens:*

(i) $K_1 = K_2$ *and* $p_1[s] \geq 1$*, or*
(ii) $K_1 \neq K_2$*,* $p_1[s] \geq 1$ *and* $p_2[s] \geq 2^n/4$*, or*
(iii) $K_1 \neq K_2$*,* $p_1[s] \geq 2^n/4$ *and* $p_2[s] \geq 1$*, or*
(iv) $K_1 \neq K_2$ *and there's a (type-1 or 2) chain* $(\mathtt{prim}, \cdot, K_1, \cdot, \cdot), (\mathtt{prim}, \cdot, K_2, \cdot, \cdot)$.

If a key is not bad then we say that it is good*. Let* Γ_{bad} *be the set of bad keys, and let* $\Gamma_{\mathrm{good}} = \mathcal{K} \backslash \Gamma_{\mathrm{bad}}$.

We first bound the probability that S is bad. First, the chance that S satisfies condition (i) above is at most $\frac{p}{2^{2k}}$. Next, we say that a subkey $J \in \{0,1\}^k$ is *heavy* if there are at least $2^n/4$ entries $(\mathtt{prim}, \cdot, J, \cdot, \cdot)$ in τ. Since there are at most $4p/2^n$ heavy subkeys, the chance that S satisfies condition (ii) above is at most $\frac{4p/2^n}{2^k} \cdot \frac{p}{2^k} = \frac{4p^2}{2^{2k+n}}$. Likewise, the chance that S satisfies condition (ii) above is at most $\frac{4p^2}{2^{2k+n}}$. From Lemma 8 and Lemma 9, there are at most $2p \cdot \mathrm{Ent}(\tau) + 4Bp + 2qB^2 + 2Bpq$ chains, and thus the chance that S satisfies condition (iii) above is at most $\frac{2p \cdot \mathrm{Ent}(\tau) + 4Bp + 2qB^2 + 2Bpq}{2^{2k}}$. Summing up

$$\begin{aligned}\Pr[S \text{ is bad}] &\leq \frac{p + 8p^2/2^n}{2^{2k}} + \frac{2p \cdot \mathrm{Ent}(\tau) + 4Bp + 2qB^2 + 2Bpq}{2^{2k}} \\ &\leq \frac{2p \cdot \mathrm{Ent}(\tau) + 5Bp + 2qB^2 + 2Bpq}{2^{2k}} \ .\end{aligned}$$

Next, recall that in the expectation method, one needs to find a non-negative function $g : \mathcal{K} \to [0, \infty)$ such that $g(s)$ bounds $1 - \mathsf{p}_{\mathbf{S}_0}(\tau, s)/\mathsf{p}_{\mathbf{S}_1}(\tau, s)$ for all $s \in \Gamma_{\mathrm{good}}$. Let U be the subset of Γ_{good} such that for any $(K_1, K_2) \in U$, we have $K_1 = K_2$. We will define $g(s)$ such that $g(s) = 2q/2^{n/2}$ for every $s \in U$, and $g(s) = \frac{4q \cdot p_1[s] \cdot p_2[s]}{2^{2n}}$ for every $s \in \mathcal{K} \backslash U$. We will show that $g(s)$ bounds

$1 - \mathsf{p}_{\mathbf{S}_0}(\tau, s)/\mathsf{p}_{\mathbf{S}_1}(\tau, s)$ later. Then

$$\begin{aligned}\mathbf{E}[g(S)] &= \frac{1}{2^{2k}}\Big(\sum_{s\in U} g(s) + \sum_{s\in\mathcal{K}\setminus U} g(s)\Big) \\ &= \frac{1}{2^{2k}}\Big(\sum_{s\in U} \frac{q}{2^{n/2}} + \sum_{s\in\mathcal{K}\setminus U} \frac{4qp_1[s]p_2[s]}{2^{2n}}\Big) \\ &\le \frac{1}{2^{2k}}\Big(\frac{q2^k}{2^{n/2}} + \frac{4qp^2}{2^{2n}}\Big) \le \frac{q}{2^{k+n/2}} + \frac{qB^2}{2^{2k}} \ .\end{aligned}$$

Then from Lemma 1,

$$\begin{aligned}\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau) - \mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau) &\le \Big(\Pr[S \text{ is bad}] + \mathbf{E}[g(S)]\Big)\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau) \\ &\le \Big(\frac{2q}{2^{k+n/2}} + \frac{2p\cdot \mathrm{Ent}(\tau) + 5Bp + 3qB^2 + 2Bpq}{2^{2k}}\Big)\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau).\end{aligned}$$

We now show that $g(s)$ indeed bounds $1-\mathsf{p}_{\mathbf{S}_0}(\tau, s)/\mathsf{p}_{\mathbf{S}_1}(\tau, s)$ for every $s \in \Gamma_{\text{good}}$. We consider the following cases, depending on whether $s \in \Gamma_{\text{good}}\setminus U$ or $s \in U$.

Case 1: $s \in \Gamma_{\text{good}}\setminus U$. For this case, we have to consider two sub-cases, depending on whether $q \le N/4$ or not.

Case 1.1: $q \le N/4$. Let $s = (K_1, K_2)$. Since $s \in \Gamma_{\text{good}}\setminus U$, we must have $K_1 \ne K_2$. We now use the following result of Chen and Steinberger [12]. (Their proof considered key-alternating ciphers, but we note that we are restricting ourselves to the setting $K_1 \ne K_2$, and and their proof also applies to the special case that all sub-keys of the key-alternating cipher are 0^n, which is equivalent to our setting here.)

Lemma 10. *[12] Assume that $p_1[s], p_2[s], q < 2^n/2$. Then*

$$1 - \frac{\mathsf{p}_{\mathbf{S}_0}(\tau, s)}{\mathsf{p}_{\mathbf{S}_1}(\tau, s)} \le \frac{q\cdot p_1[s]\cdot p_2[s]}{(2^n - q - p_1[s])(2^n - q - p_2[s])} \ . \quad \square$$

From Lemma 10, since $p_1[s], p_2[s], q \le 2^n/4$,

$$1 - \frac{\mathsf{p}_{\mathbf{S}_0}(\tau, s)}{\mathsf{p}_{\mathbf{S}_1}(\tau, s)} \le \frac{4q\cdot p_1[s]\cdot p_2[s]}{2^{2n}} = g(s) \ .$$

Case 1.2: $N/4 < q \le N$. Let Z be the random variable for the additional $(N-q)$ Enc queries that τ lacks. For we write $\mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau, s, z)$ to be the probability that $\mathbf{S}_{\text{real}}$ answers queries according to τ, and that $S = s$ and $Z = z$. In this case $\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau, s, z)$ is the probability that $\mathbf{S}_{\text{ideal}}$ behaves according to the entries in (τ, z), and $S \leftarrow\!\!\$\ \{0,1\}^{2k}$ agrees with s. We now show that $\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau, s, z) \le \mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau, s, z)$ for all choices of z such that $\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau, s, z) > 0$, and thus

$$\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau, s) - \mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau, s) \le \sum_z \mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau, s, z) - \mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau, s, z) \le 0 \le g(s) \ .$$

Let $s = (K_1, K_2)$ and $a = p_1[s]$ and $b = p_1[s] + p_2[s] < 2^n$. As $s \in \Gamma_{\text{good}} \backslash U$, the entries in (τ, z) consist of the following categories:

(1) $(\mathsf{enc}, \cdot, 1, x_1, y_1), \ldots, (\mathsf{enc}, \cdot, 1, x_{2^n}, y_{2^n})$,
(2) $(\mathtt{prim}, \cdot, K_1, x_1, u_1), \ldots, (\mathtt{prim}, \cdot, K_1, x_a, u_a)$ and $(\mathtt{prim}, \cdot, K_2, u_{a+1}, y_{a+1}), \ldots,$ $(\mathtt{prim}, \cdot, K_2, u_b, y_b)$, and
(3) $(\mathtt{prim}, \cdot, J, \cdot, \cdot)$, with $J \notin \{K_1, K_2\}$.

Hence $\mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau, s, z)$ is the probability of the following events:

(i) If we make queries in category (3) above, we get the answers provided by τ.
(ii) $S \leftarrow\!\!\$\ \{0,1\}^{2k}$ agrees with s.
(iii) For any $i \in \{1, \ldots, a+b\}$, querying $\textsc{Prim}(K_1, x_i)$ in $\mathbf{S}_{\text{real}}$ yields u_i, and querying $\textsc{PrimInv}(K_2, y_i)$ in $\mathbf{S}_{\text{real}}$ yields u_i. Moreover, for any $j \in \{b+1, \ldots, 2^n\}$, in $\mathbf{S}_{\text{real}}$, the output of $\textsc{Prim}(K_1, x_j)$ is the same as the output of $\textsc{PrimInv}(K_2, y_j)$.

Note that the three events above are independent, and the first two are independent of the system. On the other hand, $\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau, s, z)$ is likewise the probability of events (i), (ii), and the following

(iv) For any $i \in \{1, \ldots, a\}$, querying $\textsc{Prim}(K_1, x_i)$ in $\mathbf{S}_{\text{ideal}}$ yields u_i. For any $i \in \{a+1, \ldots, b\}$, querying $\textsc{PrimInv}(K_2, y_i)$ in $\mathbf{S}_{\text{ideal}}$ yields u_i. Moreover, for any $j \in \{1, \ldots, 2^n\}$, querying $\textsc{Enc}(1, x_j)$ yields y_j.

Again, note that events (i), (ii), and (iv) are independent. Hence we need only show that the probability that event (iii) happens is at least the probability that event (iv) happens. The chance that event (iii) is

$$\frac{1}{(2^n)! \cdot 2^n(2^n-1)(2^n-a-b)}$$

whereas the chance that event (iv) happens is

$$\frac{1}{(2^n)! \cdot 2^n(2^n-1)\cdots(2^n-a) \cdot 2^n(2^n-1)\cdots(2^n-b)} \ .$$

Hence the probability that event (iii) happens is indeed at least the probability that event (iv) happens.

Case 2: $s \in U$. Then $p_1[s] = 0$. Clearly if $q \geq 2^{n/2-1}$ then the claim vacuously holds. Assume that $q < 2^{n/2-1}$. Let $s = (K_1, K_1)$. Let the Enc/Dec entries in τ be $(\mathsf{enc}, \cdot, 1, x_1, y_1), \ldots, (\mathsf{enc}, \cdot, 1, x_q, y_q)$. Note that τ doesn't contain any entry $(\mathtt{prim}, \cdot, K_1, \cdot, \cdot)$. Then $\mathsf{p}_{\mathbf{S}_{\text{ideal}}}(\tau, s)$ is the probability of the following events:

(a) $S \leftarrow\!\!\$\ \{0,1\}^{2k}$ agrees with s.
(b) If we make Prim/PrimInv queries in τ, we get the answers provided by τ.
(c) $\mathbf{S}_{\text{ideal}}$ behaves according to the Enc/Dec queries in τ.

Note that the three events above are independent, and the first two are independent of the system. On the other hand, $\mathsf{p}_{\mathbf{S}_{\text{real}}}(\tau, s)$ is at least the probability of events (a), (b), and the following:

(d) For every $i \in \{1, \ldots, q\}$, if we query $\textsc{Prim}(K, x_i)$, we will get an answer $z_i \notin \{x_1, y_1, \ldots, x_q, y_q\}$, and the strings $z_1, \ldots, z_q$ are distinct. Moreover, if we query $\textsc{Prim}(K, z_i)$, we will get y_i.

Again, events (a), (b), and (d) are independent. Hence we only need to show that, $\Pr[\text{Event (d)}] \geq (1 - 2q/2^{n/2}) \Pr[\text{Event (c)}]$. Note that event (c) happens with probability

$$\frac{1}{2^n(2^n-1)\cdots(2^n-q+1)},$$

whereas event (d) happens with probability

$$\Big(\prod_{i=0}^{q-1} \frac{2^n-2q-i}{2^n-i}\Big) \frac{1}{(2^n-q)\cdots(2^n-2q+1)} \ .$$

Hence

$$\begin{aligned}\frac{\Pr[\text{Event (d)}]}{\Pr[\text{Event (c)}]} &= \prod_{i=0}^{q-1} \frac{2^n-2q-i}{2^n-q-i} = \prod_{i=0}^{q-1}\Big(1 - \frac{q}{2^n-q-i}\Big) \\ &\geq 1 - \sum_{i=0}^{q-1} \frac{q}{2^n-q-i} \geq 1 - \frac{q^2}{2^n-2q} \geq 1 - \frac{2q}{2^{n/2}},\end{aligned}$$

where the first inequality is due to the fact that $(1-x)(1-y) \geq 1-x-y$ for any $x, y \geq 0$, and the last inequality is due to the assumption that $q < 2^{n/2-1}$. This concludes the proof.

6 Matching Attacks

In this section, we give matching attacks for both Single Encryption and Double Encryption, in which the adversary uses $\Theta(q)$ $\textsc{Enc}/\textsc{Dec}$ queries and $\Theta(p)$ $\textsc{Prim}/\textsc{PrimInv}$ queries. Our attack on Single Encryption generalizes Biham's work [10] for all choices of the parameters p and q. For Double Encryption, recall that one can launch a su attack with advantage $\frac{p^2}{2^{2k}}$, and Biham's key-collision attack [10] gives a mu attack with advantage $\frac{q^2}{2^{2k}}$. Thus those attacks already give matching bounds in the usual case $n \geq k$ (such as DES or AES). Hence for Double Encryption, the only interesting setting to find matching attacks is where $n \ll k$ (such as Format-Preserving Encryption or MISTY-1). We however only know how to give matching attacks for this setting if the adversary is given all the keys after it finishes querying, which is the model in our security proof and many prior works [14, 16]. Our attack yields the bound around $p^2 s/2^{2k}$, where $s = \max\{\lfloor n/8 \lg(n) \rfloor, q/2^n\}$, which is much better than the two known attacks

above. We leave as an open problem to find matching attacks for $n \ll k$ without key revelation.

A USEFUL INEQUALITY. In the attacks, we often need to make use of the following technical result.

Lemma 11. *Let $r \geq 1$ be an integer and $0 < a \leq 1/r$. Then $(1-a)^r \leq 1 - ar/2$.*

Proof. Clearly the claimed inequality holds for $r = 1$, and thus we need only consider $r \geq 2$. Let $f(x) = xr/2 + (1-x)^r - 1$. Our goal is to show that $f(a) \leq 0$. The derivative and second derivative of the function f are $f'(x) = \frac{r}{2} - r(1-x)^{r-1}$ and $f''(x) = \frac{1}{2} + r(r-1)(1-x)^{r-2}$ respectively. Since $f''(x) > 0$ for all $x \in [0, 1/r]$, the function $f'(x)$ is strictly increasing. We claim that $f(a) \leq \max\{f(0), f(1/r)\}$. Since $f(0) = 0$ and

$$f(1/r) = \frac{1}{2} + (1 - 1/r)^r - 1 \leq \frac{1}{e} - \frac{1}{2} < 0,$$

we have $f(a) \leq 0$. To justify the claim above, note that if $f'(1/r) < 0$ then function f is decreasing, and thus $f(a) \leq f(0) = \max\{f(0), f(1/r)\}$. If $f'(1/r) \geq 0$, since function f' is strictly increasing and $f'(0) = -r/2 < 0$, there exists $b \in [0, 1/r]$ such that $f'(x) < 0$ for every $x \in [0, b)$ and $f'(x) \geq 0$ for every $x \in [b, 1/r]$. Hence function f is decreasing in $[0, b)$ and increasing in $[b, 1/r]$, and thus $f(a) \leq \max\{f(0), f(1/r)\}$. □

6.1 Attacking Single Encryption

Let $d = \lceil \frac{k+2}{n-1} \rceil$ and assume that $d \leq 2^{n-1}$, which holds for all practical values of n and k. Then

$$2^n(2^n - 1) \cdots (2^n - d + 1) \geq (2^{n-1})^d \geq 2^{k+2} \ .$$

For all practical values of n and k, the value d will be very small. For example, if $n = 64$ and $k = 56$ (meaning DES parameters), we have $d = 1$. Or if $n = k = 128$ (meaning AES parameters), we have $d = 2$. Let $p, q \in \mathbb{N}$ such that $pq \leq 2^k$. Consider the following adversary A. It picks arbitrary distinct $x_1, \ldots, x_d \in \{0,1\}^n$ and queries $\textsc{Enc}(i, x_\ell)$ to get answer $y_{i,\ell}$, for every $i \in \{1, \ldots, q\}$ and $\ell \in \{1, \ldots, d\}$. It then picks p arbitrary distinct keys $K_1, \ldots, K_p \in \{0,1\}^k$ and queries $E(K_j, x_\ell)$ to get answer $z_{j,\ell}$, for every $j \in \{1, \ldots, p\}$ and $\ell \in \{1, \ldots, d\}$. If there are i and j such that $y_{i,\ell} = z_{j,\ell}$ for every $\ell \in \{1, \ldots, d\}$ then the adversary outputs 1, otherwise it outputs 0. In the real game, from Lemma 11, the chance that the adversary outputs 1 is

$$1 - \left(1 - \frac{p}{2^k}\right)^q \geq \frac{pq}{2^{k+1}} \ .$$

In the ideal game, the chance that it outputs 1 is

$$\frac{pq}{2^n(2^n-1)\cdots(2^n-d+1)} \le \frac{pq}{2^{k+2}} \quad .$$

Hence $\mathsf{Adv}_E^{\pm\text{mu-prp}}(A) \ge \frac{pq}{2^{k+2}}$, and the adversary uses dq ENC queries and dp PRIM queries.

6.2 Attacking Double Encryption

Here we assume that $16 \le n < k$, and aim to achieve advantage $p^2 s/2^{2k}$, where $s = \max\{\lfloor n/8\lg(n)\rfloor, q/2^n\}$. We have the following restrictions on the parameters p and q:

- Since there are attacks with advantage $Q^2/2^{2k}$, where $Q = \max\{p,q\}$, we need only consider $2^n/n \le q \le 2^k$.
- Since using $p \approx 2^k/\sqrt{s}$ is already enough to get a constant advantage, without loss of generality, we can assume that $p \le 2^k/\sqrt{s}$.

Moreover, recall that we are in the model where the keys are given to the adversary after it finishes querying.

THE ATTACK. For every $i \in \{1,\ldots,q\}$, query $(i,0^n)$ to ENC to receive answer y_i. View each string in $\{0,1\}^n$ as a bin, and querying $\text{ENC}(i,0^n)$ means throwing a ball to those 2^n bins at random. Let y be the bin of the most balls, and let S be the set of indices i such that the corresponding ball falls into bin y. The following lemma gives a strong concentration bound on $|S|$ in both the real and ideal games; see Appendix C for the proof.

Lemma 12. *Let $n \ge 16$ and $q \ge 2^n/n$ be integers. Consider throwing q balls to 2^n bins at random. Let X denote the random variable for the number of balls in the bin of most balls. Then*

$$\Pr\left[X \ge \max\{\lfloor n/8\lg(n)\rfloor, q/2^n\}\right] \ge 1 - 2^{-n/3} \quad . \quad \square$$

Next, if $|S| < s$ then the output a random guess to get advantage 0. If $|S| \ge s$, which happens with probability at least $1-2^{-n/3}$ according to Lemma 12, then adapt the meet-in-the-middle attack as follows. Pick distinct keys $J_1,\ldots,J_{2p} \in \{0,1\}^k$, and query $\text{PRIM}(J_i,x)$ and $\text{PRIMINV}(J_{i+p},y)$ to get answer u_i and v_i respectively. When the keys are given, check if there are some $i,j \in \{1,\ldots,p\}$ and $\ell \in S$ such that (J_i, J_{j+p}) is the key of user ℓ. If such a triple (i,j,ℓ) exists then output 1 if and only if $u_i = v_j$.

ANALYSES. Suppose that $|S| \ge s$. Then the chance that there are $i,j \in \{1,\ldots,p\}$ and $\ell \in S$ such that (J_i, J_{j+p}) is the key of user ℓ is

$$1-(1-p^2/2^{2k})^{|S|} \ge 1-(1-p^2/2^{2k})^s \ge \frac{p^2 s}{2^{2k+1}},$$

where the last inequality is due to Lemma 11. If this pair exists then in the ideal game, the conditional probability that $v_i = u_i$ is $1/2^n$, whereas in the real game, $v_i = u_i$ with conditional probability 1. Putting this all together, the adversary wins with advantage at least

$$(1 - 2^{-n/3})(1 - 2^{-n}) \cdot \frac{p^2 s}{2^{2k+1}} \geq \max\{\lfloor n/8 \lg(n) \rfloor, q/2^n\} \cdot \frac{p^2}{3 \cdot 2^{2k}} \quad .$$

Discussion. What's the problem if we are not given keys at the end of the querying process? Now we have many pairs (i, j) such that $u_i = v_i$. One such pair will yield the key (J_i, J_{j+p}) for some user, but we don't know which user. Moreover, there are too many pairs (i, j)—one can show that in the ideal world, there are on average $O(p^2/2^n)$ such pairs—and most of them are just false positives.

Acknowledgments. This research was partially supported by NSF grants CNS-1423566, CNS-1528178, CNS-1553758 (CAREER), IIS-152804, by a Hellman Fellowship, and by the Glen and Susanne Culler Chair.

References

1. ANSI X9.52: Triple data encryption algorithm modes of operation, 1998.
2. W. Aiello, M. Bellare, G. Di Crescenzo, and R. Venkatesan. Security amplification by composition: The case of doubly-iterated, ideal ciphers. In H. Krawczyk, editor, *CRYPTO'98*, volume 1462 of *LNCS*, pages 390–407. Springer, Heidelberg, Aug. 1998.
3. E. Andreeva, J. Daemen, B. Mennink, and G. V. Assche. Security of keyed sponge constructions using a modular proof approach. In G. Leander, editor, *FSE 2015*, volume 9054 of *LNCS*, pages 364–384. Springer, Heidelberg, Mar. 2015.
4. M. Bellare, D. J. Bernstein, and S. Tessaro. Hash-function based PRFs: AMAC and its multi-user security. In M. Fischlin and J.-S. Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 566–595. Springer, Heidelberg, May 2016.
5. M. Bellare, A. Boldyreva, and S. Micali. Public-key encryption in a multi-user setting: Security proofs and improvements. In B. Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 259–274. Springer, Heidelberg, May 2000.
6. M. Bellare, T. Ristenpart, P. Rogaway, and T. Stegers. Format-preserving encryption. In M. J. Jacobson Jr., V. Rijmen, and R. Safavi-Naini, editors, *SAC 2009*, volume 5867 of *LNCS*, pages 295–312. Springer, Heidelberg, Aug. 2009.
7. M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In S. Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Heidelberg, May / June 2006.
8. M. Bellare and B. Tackmann. The multi-user security of authenticated encryption: AES-GCM in TLS 1.3. In M. Robshaw and J. Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 247–276. Springer, Heidelberg, Aug. 2016.
9. D. J. Bernstein. How to stretch random functions: The security of protected counter sums. *Journal of Cryptology*, 12(3):185–192, 1999.
10. E. Biham. How to decrypt or even substitute DES-encrypted messages in 2^{28} steps. *Information Processing Letters*, 84(3):117–124, 2002.

11. A. Bogdanov, L. R. Knudsen, G. Leander, F.-X. Standaert, J. P. Steinberger, and E. Tischhauser. Key-alternating ciphers in a provable setting: Encryption using a small number of public permutations - (extended abstract). In D. Pointcheval and T. Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 45–62. Springer, Heidelberg, Apr. 2012.
12. S. Chen and J. P. Steinberger. Tight security bounds for key-alternating ciphers. In P. Q. Nguyen and E. Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 327–350. Springer, Heidelberg, May 2014.
13. Y. Dai, J. Lee, B. Mennink, and J. P. Steinberger. The security of multiple encryption in the ideal cipher model. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 20–38. Springer, Heidelberg, Aug. 2014.
14. Y. Dai and J. Steinberger. Tight security bounds for multiple encryption. Cryptology ePrint Archive, Report 2014/096, 2014. http://eprint.iacr.org/2014/096.
15. P. Gaži. Plain versus randomized cascading-based key-length extension for block ciphers. In R. Canetti and J. A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 551–570. Springer, Heidelberg, Aug. 2013.
16. P. Gazi, J. Lee, Y. Seurin, J. P. Steinberger, and S. Tessaro. Relaxing full-codebook security: A refined analysis of key-length extension schemes. In G. Leander, editor, *FSE 2015*, volume 9054 of *LNCS*, pages 319–341. Springer, Heidelberg, Mar. 2015.
17. P. Gaži and U. M. Maurer. Cascade encryption revisited. In M. Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 37–51. Springer, Heidelberg, Dec. 2009.
18. P. Gaži and S. Tessaro. Efficient and optimally secure key-length extension for block ciphers via randomized cascading. In D. Pointcheval and T. Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 63–80. Springer, Heidelberg, Apr. 2012.
19. V. T. Hoang and S. Tessaro. Key-alternating ciphers and key-length extension: Exact bounds and multi-user security. In *Advances in Cryptology — CRYPTO 2016*, pages 3–32, 2016.
20. J. Lee. Towards key-length extension with optimal security: Cascade encryption and xor-cascade encryption. In T. Johansson and P. Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 405–425. Springer, Heidelberg, May 2013.
21. M. Luby and C. Rackoff. Pseudo-random permutation generators and cryptographic composition. In J. Hartmanis, editor, *Proceedings of the 18th Annual ACM Symposium on Theory of Computing, May 28-30, 1986, Berkeley, California, USA*, pages 356–363. ACM, 1986.
22. U. M. Maurer. Indistinguishability of random systems. In L. R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 110–132. Springer, Heidelberg, Apr. / May 2002.
23. D. A. McGrew and J. Viega. The security and performance of the Galois/counter mode (GCM) of operation. In A. Canteaut and K. Viswanathan, editors, *INDOCRYPT 2004*, volume 3348 of *LNCS*, pages 343–355. Springer, Heidelberg, Dec. 2004.
24. N. Mouha and A. Luykx. Multi-key security: The Even-Mansour construction revisited. In R. Gennaro and M. J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 209–223. Springer, Heidelberg, Aug. 2015.
25. S. A. Myers. On the development of block-ciphers and pseudo-random function generators using the composition and xor operators, 1999.

26. J. Patarin. The "coefficients H" technique (invited talk). In R. M. Avanzi, L. Keliher, and F. Sica, editors, *SAC 2008*, volume 5381 of *LNCS*, pages 328–345. Springer, Heidelberg, Aug. 2009.
27. M. Raab and A. Steger. "Balls into bins" – a simple and tight analysis. In *RANDOM 1998*, pages 159–170. Springer, 1998.
28. S. Tessaro. Security amplification for the cascade of arbitrarily weak PRPs: Tight bounds via the interactive hardcore lemma. In Y. Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 37–54. Springer, Heidelberg, Mar. 2011.
29. S. Tessaro. Optimally secure block ciphers from ideal primitives. In T. Iwata and J. H. Cheon, editors, *ASIACRYPT 2015, Part II*, volume 9453 of *LNCS*, pages 437–462. Springer, Heidelberg, Nov. / Dec. 2015.

A Proof of Lemma 8

We claim that the number of positive type-1 chains is at most $2Bp + B^2q + Bpq$. By symmetry, the number of negative type-1 chains is also at most $2Bp + B^2q + Bpq$. Hence the total number of type-1 chains is at most $4Bp + 2B^2q + 2Bpq$.

To justify the claim above, consider a positive type-1 chain $(\texttt{prim}, \cdot, \cdot, x_1, y_1)$, $(\texttt{prim}, \cdot, \cdot, x_2, y_2)$. There are four ways to assign the signs $+/-$ to the entries. Fix a specific way to assign the signs. We consider the following cases.
Case 1: Both entries have sign $-$. Then there are at most Bq choices for the first entry, since $\tau \notin \mathcal{S}$. Moreover, once the first entry is fixed, there are only B choices for the second entry. Thus in this case, the total number of positive type-1 chains is at most B^2q.
Case 2: Both entries have sign $+$. There are at most p choices for the last entry. Moreover, once the last entry is fixed, there are at most B choices for the first entry. Thus in this case, the total number of positive type-1 chains is at most Bp.
Case 2: The first entry has sign $-$ and the second sign $+$. There are at most Bq choices for the first entry and p choices for the last one. Thus in this case, the total number of positive type-1 chains is at most Bpq.
Case 4: The first entries has sign $+$ and the second sign $-$. Then there are at most p choices for the first entry. Moreover, once the first entry is fixed, there are at most B choices for the last entry. Thus in this case, the total number of positive type-1 chains is at most Bp.

Summing up, the total number of positive type-1 chains is at most $2Bp + B^2q + Bpq$.

B Proof of Lemma 9

We claim that the number of negative type-2 chains is at most $p \cdot \mathrm{Ent}(\tau)$. By symmetry, the number of positive type-2 chains is also at most $p \cdot \mathrm{Ent}(\tau)$. Hence the total number of type-2 chains is at most $2p \cdot \mathrm{Ent}(\tau)$.

To justify the claim above, consider a negative type-2 chain $(\texttt{prim}, \cdot, \cdot, x_1, y_1)$, $(\texttt{prim}, \cdot, \cdot, x_2, y_2)$. Then there are at most $\mathrm{Ent}(\tau)$ choices for the first entry, and p choices for the last entry. Thus the total number of negative type-2 chains is at most $p \cdot \mathrm{Ent}(\tau)$.

C Proof of Lemma 12

Let $s = \lfloor n/8\lg(n) \rfloor$. Clearly $X \geq q/2^n$, hence we only need to consider the case that $q/2^n \leq s$. Our proof will closely follow the second-moment method in classic balls-into-bins papers [27]. For any $i \in \{1, \ldots, 2^n\}$, the chance that the i-th bin has at least s balls is

$$\binom{q}{s}\frac{1}{(2^n)^s}\left(1-\frac{1}{2^n}\right)^{q-s} \geq \left(\frac{q}{s}\right)^s \frac{1}{(2^n)^s}\left(1-\frac{1}{2^n}\right)^{q} \geq \left(\frac{q/2^n}{s}\right)^s \cdot e^{-q/2^n} \geq n^{-2s} \cdot e^{-q/2^n} \ .$$

Moreover,

$$n^{-2s} \cdot e^{-q/2^n} \geq 2^{-2s\lg(n)} \cdot 2^{-1.5q/2^n} \geq 2^{-n/4-\frac{1.5n}{8\lg(n)}} \geq 2^{-n/3} \ .$$

Let Y_i be the Bernoulli random variable such that $Y_i = 1$ if and only if the i-th bin has at least s balls. Then $\mathbf{E}[Y_i] = \Pr[Y_i = 1] \geq 2^{-n/3}$. Let $Y = Y_1 + \cdots + Y_{2^n}$, and thus

$$\mathbf{E}[Y] = \mathbf{E}[Y_1] + \cdots + \mathbf{E}[Y_{2^n}] \geq 2^{2n/3} \ .$$

Since

$$\Pr[X \geq s] = \Pr[Y \geq 1] = 1 - \Pr[Y = 0] \geq 1 - \Pr\Big[|Y - \mathbf{E}[Y]| \geq \mathbf{E}[Y]\Big],$$

what's left is to show that $\Pr\Big[|Y - \mathbf{E}[Y]| \geq \mathbf{E}[Y]\Big] \leq 2^{-n/3}$. By Chebyshev's inequality,

$$\Pr\Big[|Y - \mathbf{E}[Y]| \geq \mathbf{E}[Y]\Big] \leq \frac{\mathbf{Var}[Y]}{(\mathbf{E}[Y])^2} \leq \frac{\mathbf{Var}[Y]}{2^{4n/3}} \ .$$

It then suffices to show that $\mathbf{Var}[Y] \leq 2^n$. On the one hand, for any $i \neq j$, each Y_i and Y_j are negatively correlated, as some bin having more balls means that it is less likely for another bin to be so. Therefore, each covariance $\mathbf{Cov}(Y_i, Y_i)$ is at most 0. On the other hand, since each Y_i is a Bernoulli random variable, $(Y_i)^2 = Y_i$, and thus

$$\mathbf{Var}[Y_i] \leq \mathbf{E}[(Y_i)^2] = \mathbf{E}[Y_i] \leq 1 \ .$$

Hence

$$\mathbf{Var}[Y] = \sum_{i=1}^{2^n} \mathbf{Var}[Y_i] + \sum_{i \neq j} \mathbf{Cov}(Y_i, Y_j) \leq 2^n \ .$$