

# The Memory-Tightness of Authenticated Encryption\*

Ashrujit Ghoshal, Joseph Jaeger, and Stefano Tessaro

Paul G. Allen School of Computer Science & Engineering  
University of Washington, Seattle, US  
{ashrujit, jsjaeger, tessaro}@cs.washington.edu

**Abstract.** This paper initiates the study of the provable security of authenticated encryption (AE) in the memory-bounded setting. Recent works – Tessaro and Thiruvengadam (TCC '18), Jaeger and Tessaro (EUROCRYPT '19), and Dinur (EUROCRYPT '20) – focus on confidentiality, and look at schemes for which trade-offs between the attacker's memory and its data complexity are inherent. Here, we ask whether these results and techniques can be lifted to the full AE setting, which additionally asks for integrity.

We show both positive and negative results. On the positive side, we provide tight memory-sensitive bounds for the security of GCM and its generalization, CAU (Bellare and Tackmann, CRYPTO '16). Our bounds apply to a restricted case of AE security which abstracts the deployment within protocols like TLS, and rely on a new memory-tight reduction to corresponding restricted notions of confidentiality and integrity. In particular, our reduction uses an amount of memory which linearly depends on that of the given adversary, as opposed to only imposing a constant memory overhead as in earlier works (Auerbach et al., CRYPTO '17).

On the negative side, we show that a large class of black-box reductions cannot generically lift confidentiality and integrity security to a joint definition of AE security in a memory-tight way.

**Keywords:** Provable security, symmetric cryptography, time-memory trade-offs, memory-tightness

## 1 Introduction

Cryptographic attacks aim to use as little *memory* as possible. While some attacks are memoryless (e.g., for collision finding), others are subject to a *trade-off* – as the available memory decreases, the time and data complexities increase. A security proof (especially one in the spirit of *concrete* security) should tell us precisely *how* memory affects other complexity metrics. However, this is technically challenging, and consequently, security proofs ignored memory until recently.

This paper continues an ongoing line of works introducing memory limitations in provable security, and initiates the study of (nonce-based) *authenticated encryption* (AE) in the memory-bounded setting. Recent works [17,11,7] have shown memory-sensitive proofs of security for symmetric encryption, showing that trade-offs between memory and data complexities are inherent. These results, however, only deal with *confidentiality* of encryption – and one of the main contributions of this paper is to highlight the challenges of lifting them to the more complex setting of AE.

We discuss definitional aspects, and then shift our focus to *memory-tight reductions* [1] in the AE setting. We prove both *positive* and *negative* results. We introduce a new technique for memory-tight reductions to obtain tight memory-sensitive bounds for the AE-security of GCM in a setting that corresponds to its usage for establishing a secure channel. We also show that restricting AE security to specific settings is inherent for memory-tight reductions – indeed, we show that the common approach of lifting confidentiality and integrity guarantees into a combined notion of AE security (or of CCA security) fails in its most general form, at least with respect to a broad class of security reductions.

---

\* A preliminary version of this paper appears in the proceedings of CRYPTO 2020. This is the full version.

## 1.1 Context: Time-memory Trade-offs for AE

Let us start by setting the context and highlighting some of the challenges. First off, existing results [11,7] can be combined to analyze the INDR security<sup>1</sup> of nonce-based encryption. For example, consider a toy scheme<sup>2</sup> SE based on a block cipher E with block length  $n$  which encrypts  $M \in \{0,1\}^n$  with key  $K$  as

$$\text{SE.E}(K, N, M) = E_K(N) \oplus M .$$

Here,  $N$  is the nonce and INDR security should hold as long as no two messages are encrypted with the same nonce. One can show that for every adversary  $\mathcal{A}$  with time, data, and memory complexities  $t$ ,  $q$ , and  $S$ , respectively,

$$\text{Adv}_{\text{SE}}^{\text{indr}}(\mathcal{A}) \leq O\left(\frac{q \cdot S \cdot \log(q)}{2^n}\right) + \text{Adv}_{\text{E}}^{\text{prp}}(\mathcal{B}) , \quad (1)$$

where  $\mathcal{B}$  is an adversary against the security of E as a pseudorandom permutation (PRP), which has time and memory complexities (roughly)  $t$  and  $S$ , respectively, and makes  $q$  queries. In particular, if  $S < 2^{n/2}$ , then SE achieves beyond-birthday security  $q > 2^{n/2}$  with respect to data complexity.

OUR GOAL, IN MORE DETAIL. However, INDR security is rarely sufficient on its own – we want *fully secure* AE schemes which also satisfy (ciphertext) *integrity* (or CTXT security, for short). Following [16], we adopt a *single* AE security definition that incorporates *both* INDR and CTXT, by measuring indistinguishability of two oracle pairs  $(\text{ENC}_b, \text{DEC}_b)$  for  $b \in \{0,1\}$ . For  $b = 1$ ,  $\text{ENC}_1$  returns real ciphertexts, and  $\text{DEC}_1$  decrypts properly. For  $b = 0$ , instead,  $\text{ENC}_0$  returns random ciphertexts, and  $\text{DEC}_0$  decrypts only previous outputs from  $\text{ENC}_0$ . It is important to use a combined definition, as it captures settings such as chosen-ciphertext attacks and padding-oracle attacks [18], which use a decryption oracle to break confidentiality.<sup>3</sup>

LIFTING TRADE-OFFS. We want to prove a bound analogous to that of (1) for AE security, preserving in particular the existing space-time trade-off. The usual approach is to prove INDR and CTXT *individually*, and then combine them to show AE security. This makes sense because (1) we *know* how to prove tight trade-offs for INDR security, and (2) we may be able to prove stronger bounds on CTXT easily, even without memory restrictions. The classical statement (originally in [16]) is that for every adversary  $\mathcal{A}$ ,

$$\text{Adv}_{\text{SE}}^{\text{ae}}(\mathcal{A}) \leq \text{Adv}_{\text{SE}}^{\text{indr}}(\mathcal{B}) + \text{Adv}_{\text{SE}}^{\text{ctxt}}(\mathcal{C}) ,$$

for suitable adversaries  $\mathcal{B}$  and  $\mathcal{C}$ , with similar time and query complexities as those of  $\mathcal{A}$ . However, this is only helpful towards our goal if the reduction is *memory-tight*, in the sense Auerbach et al. (ACKF) [1], i.e.,  $\mathcal{B}$  and  $\mathcal{C}$ 's memory costs must not noticeably exceed those of  $\mathcal{A}$ . This is fundamental to preserve a time-memory trade-off like the one from (1).

Unfortunately, the standard proof is not memory-tight with respect to the INDR adversary  $\mathcal{B}$ , as it needs to simulate  $\text{DEC}_0$  which requires remembering prior ciphertexts. In a nutshell, we will show that the lack of memory-tightness is inherent, *but* the definition can be restricted enough for interesting deployment scenarios to actually allow for a memory-tight reduction.

DEFINITIONAL ISSUES. Several “without loss of generality” definitional equivalences are false in the memory-bounded setting. For example, INDR security holds as long as nonces do not repeat, but there are options to formalize this, e.g.: (A) The game enforces this by answering encryption queries repeating a nonce with  $\perp$ , unless the same message is re-encrypted, or (B) The adversary never repeats a nonce. If we do not care about memory, these two definitions are indeed equivalent, but if we do, then they are not. Indeed, the bound in (1) for our toy scheme can only be true for (B) – it is not hard to see that otherwise we can mount a memory-less distinguishing attack with  $q \approx 2^{n/2}$  queries. (The attack also works if  $\perp$  is returned even if we re-encrypt the same message.) We discuss definitions in detail in Section 3.

<sup>1</sup> Which measures the indistinguishability of ciphertexts from truly random ones.

<sup>2</sup> Our discussion can easily be extended to many schemes following the format of counter-mode encryption.

<sup>3</sup> While we target such a single definition of AE, we stress that our results would extend to considering CCA security as a target.

## 1.2 Positive Results

We provide a novel memory-tight reduction for the common case where AE is used to establish a secure communication *channel*, as in TLS. The key point is that in this setting, only certain restricted adversarial interactions can occur in the AE security game, i.e.:

- (1) Nonces are *implicit* – they are incremented as a counter.
- (2) The receiver *aborts upon the first decryption failure*. In particular, messages *need* to be delivered in the same order as they are encrypted.

Our memory-tight reduction is for an abstraction of this setting we refer to as a *channel*. (Although, for this introduction, we stick with the more conventional language of AE.) We apply our reduction to prove (tight) memory-sensitive bounds for a channel instantiated with the CAU scheme by Bellare and Tackmann [4], an abstraction of GCM [12].

**THE SECURITY GAME.** When restricting AE security to this setting, we can assume that the adversary  $\mathcal{A}$  can encrypt messages  $M_1, M_2, \dots$  and obtains ciphertexts  $C_1, C_2, \dots$  via an *encryption oracle*  $\text{ENC}_b$ , for  $b \in \{0, 1\}$ . When  $b = 1$ , the  $C_i$ 's are actual encryptions of the  $M_i$ 's (with increasing nonces), whereas when  $b = 0$ , they are truly random ciphertexts. The adversary is also given access to a *decryption oracle*  $\text{DEC}_b$ . If  $b = 1$ , this just applies the decryption algorithm of the AE scheme, using increasing nonces. If decryption fails,  $\text{DEC}_b$  responds to this and any future queries with  $\perp$ . For  $b = 0$ , the oracle responds with  $M_1, M_2, \dots$  as long as it is supplied the ciphertexts  $C_1, C_2, \dots$  *in the order they have been produced by*  $\text{ENC}_0$ . If the ciphertexts come in the wrong order,  $\text{DEC}_0$  responds to this and any future queries with  $\perp$ . The goal here is to distinguish  $(\text{ENC}_0, \text{DEC}_0)$  and  $(\text{ENC}_1, \text{DEC}_1)$ .

**PROOF IDEA.** In this channel setting, to obtain a memory-tight reduction from AE security to CTXT and INDR security, we first use CTXT security to replace the oracles  $(\text{ENC}_1, \text{DEC}_1)$  with  $(\text{ENC}_1, \text{DEC}_0)$ . (This step is easily seen to be memory-tight.) Next, we aim to use INDR security to replace  $\text{ENC}_1$  with  $\text{ENC}_0$ . The catch here is that when doing so, we need to simulate the  $\text{DEC}_0$  oracle in the INDR security game (which does not provide one). Again, this seems to require remembering all prior ciphertexts, thus preventing memory-tightness.

A key observation, however, is that ciphertexts are only accepted when arriving with the right order. For this reason, we will show (via an information-theoretic argument) that our reduction only needs to store the  $\delta$  oldest ciphertexts which have not been delivered yet, for some  $\delta$  – the key point here is that  $\delta$  can be chosen to depend (roughly linearly) on the memory of the adversary used by the reduction, so the overall memory of the constructed adversary is of the same magnitude of that of the AE adversary.

This is in contrast to existing memory-tight reductions in the literature which are (near) “memory-less”, i.e., the reduction adds a small memory overhead, *independent* of the memory of the adversary. Our reduction is the first example where the reduction uses memory in addition to that of the adversary, but the size of this memory is bounded in terms of the adversary’s memory complexity.

**APPLICATION TO CAU.** We apply our memory-tight reduction to show bounds for CAU (and hence GCM) in the communication channel setting. We refer to the resulting channel as **NCH**, and it is based on a block cipher  $E$ . We show that for every adversary  $\mathcal{A}$ , there exists  $\mathcal{B}$  such that

$$\text{Adv}_{\text{NCH}}^{\text{ch-ae}}(\mathcal{A}) \leq 4 \cdot \text{Adv}_{\text{E}}^{\text{prp}}(\mathcal{B}) + O\left(\frac{pqS}{2^n}\right), \quad (2)$$

where  $O(\cdot)$  hides a small constant,  $q$  and  $S$  are the data and memory complexities of  $\mathcal{A}$ , and  $p$  is an upper bound on the length of ciphertexts. Further,  $\mathcal{B}$  makes  $q \cdot p$  queries, and has time complexity similar to that of  $\mathcal{A}$ . Instrumental to our result here is Dinur’s Switching Lemma [7]. The main challenge is to prove a bound for CTXT security – our proof relies once again on similar techniques to our memory-tight reduction.

## 1.3 Negative Results

A meaningful question is whether we can give a memory-tight reduction beyond the setting of channels, and reduce AE security to INDR and CTXT security in the most general sense. Here, we show that this is unlikely by giving impossibility results for black-box reductions.

We consider reductions to INDR and CTXT which are restricted, but note that all prior impossibility results on memory-tight reductions [1,19,9] make similar or stronger restrictions. In particular, we require the reductions to simulate their encryption oracles “faithfully” to an AE adversary, i.e., if they answer an encryption query with a ciphertext  $C$ , the same query (1) has been asked to the encryption oracle available to the reduction and (2) it has returned  $C$ . This restriction is natural, and we are not aware of any reductions evading it.

**STRAIGHTLINE REDUCTIONS.** Our first result builds an (inefficient) adversary  $\mathcal{A}$  against AE security which no straightline reduction can use to (1) break CTXT security (regardless of the memory available to the reduction) or, more importantly, to (2) break INDR security (unless the reduction uses an amount of memory proportional to the query complexity of the adversary). Moreover,  $\mathcal{A}$  uses little memory, and thus our result implies impossibility even for “weakly memory-tight reductions” which adapt their memory usage (such as the one we give in this paper). This is unlike recent works [19,9], which only rule out reductions with memory independent of that of the adversary.

At a high level,  $\mathcal{A}$  forces the reduction to complete a memory-hard task before being useful. If the reduction succeeds,  $\mathcal{A}$  executes an (inefficient) procedure to break INDR security. (And importantly, this procedure does not help in breaking CTXT security!) More in detail, the first part of  $\mathcal{A}$ ’s execution consists of *challenge rounds*. In each of these rounds,  $\mathcal{A}$  encrypts random plaintexts  $M_1, \dots, M_u$ , which result in ciphertexts  $C_1, \dots, C_u$ , and also picks a random index  $i^* \in [u]$ . It then asks for the decryption of  $C_{i^*}$ , and checks whether the response equals  $M_{i^*}$ . If so, it moves to the next round, if not it aborts by doing something useless. Only if all rounds are successful  $\mathcal{A}$  proceeds to break INDR security. We use techniques borrowed from the setting of random oracles with auxiliary input (AI-ROM) [5] to prove that the probability that all rounds are successful decays exponentially as long as the reduction’s memory does not fit all of  $M_1, \dots, M_u$ .

**FULL REWINDING.** The restriction to straightline reductions seem too restrictive: After all, a reduction could (1) wait for a decryption query  $C_{i^*}$ , then (2) rewind the adversary to re-ask  $M_1, M_2, \dots$  until  $M_{i^*}$  is asked. The caveat is that our definition of INDR security does not allow for *re-asking* encryption queries (again, as pointed out above, such a notion would prevent us from using the results of [11,7]). Therefore, if we assume that all the reduction can do is remember (say)  $S$  plaintext-ciphertext pairs, the above adversary  $\mathcal{A}$  will fail to pass a challenge round with probability at least  $1 - S/u$ .

Still, this does not mean that rewinding cannot help when allowing more general adversarial strategies. While handling arbitrary rewinding appears to be out of reach, we make partial progress by extending our proof (and our construction of  $\mathcal{A}$ ) to show that “full” rewinding (i.e., re-running  $\mathcal{A}$  from the beginning) does not help. This is the same rewinding model considered in prior memory-tightness lower bounds [1]. However, in those results, one obtains a rewinding-memory trade-offs (in that reducing memory would require more rewinding). Here, our result is absolute, in the sense that if memory is too small, no amount of rewinding can help.

**Paper overview.** In Section 2, we introduce our notation, basic definitions and cover some cryptographic background necessary for the paper. In Section 3, we recall the standard definitions for the security notions of nonce-based encryption. We point out several nuances while defining security in the memory bounded setting. We conclude the section by giving a time-memory tradeoff for the INDR security of CAU. In Section 4, we show that memory-tight reductions can be given for the combined confidentiality and integrity security of cryptographic channels. Using the result from Section 3, we prove the security of a channel based on CAU. The resulting channel can be viewed as (a simplification of) the channel obtained when using GCM in TLS 1.3. In Section 5, we give impossibility results (for a natural restricted class of black-box reductions) for giving a memory-tight reduction from AE security to INDR and CTXT security. This establishes that our move to the channel setting for Section 4 was necessary for our positive result.

## 2 Definitions

Let  $\mathbb{N} = \{0, 1, 2, \dots\}$ . For  $D \in \mathbb{N}$ , let  $[D] = \{1, 2, \dots, D\}$ . If  $S$  and  $S'$  are finite sets, then  $\text{Fcs}(S, S')$  denotes the set of all functions  $F : S \rightarrow S'$  and  $\text{Perm}(S)$  denotes the set of all permutations on  $S$ . Picking an element

uniformly at random from  $S$  and assigning it to  $s$  is denoted by  $s \leftarrow S$ . The set of finite vectors with entries in  $S$  is  $S^*$  or  $(S)^*$ . Thus  $\{0, 1\}^*$  is the set of finite length strings.

If  $x \in \{0, 1\}^*$  is a string, then  $|x|$  denotes its bitlength. If  $n \in \mathbb{N}$  and  $x \in \{0, 1\}^*$ , then  $|x|_n = \max\{1, \lceil |x|/n \rceil\}$ . We let  $x_1 \dots x_\ell \leftarrow_n x$  denote setting  $\ell \leftarrow |x|_n$  and parsing  $x$  into  $\ell$  blocks of length  $n$  (except  $x_\ell$  which may have  $|x_\ell| < n$ ). We let  $x[:n]$  denote the first  $n$  bits of  $x$  and  $x[i:n]$  denote the  $i$ -th (exclusive) through  $n$ -th (inclusive) bits of  $x$ . We adopt the convention that if  $|x| < |x'|$  then  $x \oplus x' = x \oplus x'[:|x|]$ . The empty string is  $\varepsilon$ .

We will make use of queues which operate in first-in, first-out order. If  $Q$  is a queue then  $Q.\text{add}(M)$  adds  $M$  to the back of the queue and  $M \leftarrow Q.\text{dq}()$  removes the first element of the queue and assigns it to  $M$ . If the queue is empty, then  $M$  is assigned the value  $\perp \notin \{0, 1\}^*$  which is used to represent rejection or uninitialized values.

Algorithms are randomized when not specified otherwise. If  $\mathcal{A}$  is an algorithm, then  $y \leftarrow \mathcal{A}^{O_1, \dots}(x_1, \dots; r)$  denotes running  $\mathcal{A}$  on inputs  $x_1, \dots$  with coins  $r$  and access to the oracles  $O_1, \dots$  to produce output  $y$ . Performing this execution with a random  $r$  is denoted  $y \leftarrow \mathcal{A}^{O_1, \dots}(x_1, \dots)$ . The set of all possible outputs of  $\mathcal{A}$  when run with inputs  $x_1, \dots$  is  $[\mathcal{A}(x_1, \dots)]$ . The notation  $y \leftarrow O(x_1, \dots)$  is used for calling oracle  $O$  with inputs  $x_1, \dots$  and assigning its output to  $y$ . (Note, the code run by the oracle is not necessarily deterministic.)

We make regular use of pseudocode games inspired by the code-based framework of [3]. Examples of games can be found in Fig. 1. We let  $\Pr[\mathbf{G}]$  denote the probability that a game  $\mathbf{G}$  outputs **true**. Booleans are implicitly initialized to **false**, integers to 0, and all other types to  $\perp$ .

COMPLEXITY CONVENTIONS. When measuring the efficiency of an adversary we follow the standard convention used in studying memory-tightness [1] on measuring the local complexity of an adversary and not included the complexity of whatever game it interacts with. We primarily focus on the worst-case runtime (i.e. how much computation it performs in between making oracle queries) and memory complexity (i.e. how many bits of state it stores for local computation) of adversaries. Note that while these exclude the time and memory used within whatever oracles the adversary may call, we do include the time and memory used to write down an oracle query and receive the response.

## 2.1 Cryptographic Background

FUNCTION FAMILY. A function family is an efficiently computable function  $F : A \times B \rightarrow C$ , where  $A, B$ , and  $C$  are sets. A hash function is a family of functions. We often write  $F_K(\cdot)$  in place of  $F(K, \cdot)$ .

PSEUDORANDOM FUNCTION/PERMUTATION. Let  $E : \{0, 1\}^k \times \{0, 1\}^n \rightarrow \{0, 1\}^m$  be a function family. If  $n = m$  and  $E_K(\cdot)$  is a permutation for each  $K \in \{0, 1\}^k$ , then we say that  $E$  is a block-cipher. The primary security notions of interest for such functions are PRF and PRP security. The former is typically more useful in applications, but when  $E$  is a block-cipher we prefer to assume PRP security and use that to deduce PRF security.

These security notions are defined by games shown in Fig. 1. In  $\mathbf{G}^{\text{prp}}$ , the adversary is given access to either  $E_K(\cdot)$  for a random key or a random permutation  $P : \{0, 1\}^n \rightarrow \{0, 1\}^n$ . Game  $\mathbf{G}^{\text{prf}}$  is defined similarly except a random function  $F : \{0, 1\}^n \rightarrow \{0, 1\}^m$  is used in place of the permutation. For  $x \in \{\text{prp}, \text{prf}\}$ , we define the advantage of  $\mathcal{A}$  by  $\text{Adv}_{\mathbf{E}}^x(\mathcal{A}) = \Pr[\mathbf{G}_{\mathbf{E}, 1}^x(\mathcal{A})] - \Pr[\mathbf{G}_{\mathbf{E}, 0}^x(\mathcal{A})]$ .

SWITCHING LEMMA. A classic result in cryptography is the “switching lemma” which bounds how well an adversary can distinguish between a random function and a random permutation. Consider the game  $\mathbf{G}_{D, b}^{\text{sl}}$  shown in Fig. 1. In it, the adversary is given oracle access to either a random function or a random permutation with domain/range  $[D]$  and is trying to figure out which. We define  $\text{Adv}_D^{\text{sl}}(\mathcal{A}) = \Pr[\mathbf{G}_{D, b}^{\text{sl}}(\mathcal{A})] - \Pr[\mathbf{G}_{D, b}^{\text{sl}}(\mathcal{A})]$ .

The classic switching lemma shows  $\text{Adv}_D^{\text{sl}}(\mathcal{A}) \in O(q^2/D)$  where  $q$  is the number of queries made by  $\mathcal{A}$ . In general, bounding the memory-complexity of the attacker cannot be used to meaningfully improve this bound because a low-memory collision-finding attack (e.g., using Pollard’s  $\rho$ -method [13, 14]) achieves advantage  $\text{Adv}_D^{\text{sl}}(\mathcal{A}) \in \Omega(q^2/D)$ . However, as originally observed by Jaeger and Tessaro we *can* obtain better results when restricting attention to adversaries that never repeat any queries.

Game $G_{E,b}^{\text{prp}}(\mathcal{A})$	Game $G_{E,b}^{\text{prf}}(\mathcal{A})$	Game $G_{D,b}^{\text{sl}}(\mathcal{A})$
$K \leftarrow \{0, 1\}^k$	$K \leftarrow \{0, 1\}^k$	$F \leftarrow \text{Fcs}([D], [D])$
$P \leftarrow \text{Perm}(\{0, 1\}^n)$	$F \leftarrow \text{Fcs}(\{0, 1\}^n, \{0, 1\}^m)$	$P \leftarrow \text{Perm}([D])$
$b' \leftarrow \mathcal{A}^{\text{EVAL}_b}$	$b' \leftarrow \mathcal{A}^{\text{EVAL}_b}$	$b' \leftarrow \mathcal{A}^{\text{EVAL}_b}$
Return $b' = 1$	Return $b' = 1$	Return $b' = 1$
Oracle $\text{EVAL}_b(x)$	Oracle $\text{EVAL}_b(x)$	Oracle $\text{EVAL}_b(x)$
$y_1 \leftarrow E_K(x)$	$y_1 \leftarrow E_K(x)$	$y_1 \leftarrow F(x)$
$y_0 \leftarrow P(x)$	$y_0 \leftarrow F(x)$	$y_0 \leftarrow P(x)$
Return $y_b$	Return $y_b$	Return $y_b$

**Fig. 1.** Security games for PRF and PRP security of E and the switching lemma.

Game $G_{\text{H}}^{\text{axu}}(\mathcal{X})$
$((A_1, C_1), (A_2, C_2), Z) \leftarrow \mathcal{X}$
$L \leftarrow \{0, 1\}^k$
If $(A_1, C_1) = (A_2, C_2)$ then return <b>false</b>
Return $H_L(A_1, C_1) \oplus H_L(A_2, C_2) = Z$

**Fig. 2.** Security game for AXU security of H.

Let  $\text{Adv}_D^{\text{sl}}(q, S)$  denote the maximal value of  $\text{Adv}_D^{\text{sl}}(\mathcal{A})$  for all  $\mathcal{A}$  that are  $S$ -bounded and make  $q$  non-repeating queries to their oracle. Jaeger and Tessaro [11] showed that  $\text{Adv}_D^{\text{sl}}(q, S) \leq \sqrt{Sq/D}$  under a combinatorial conjecture. Later, Dinur [7] improved this to show that  $\text{Adv}_D^{\text{sl}}(q, S) \in O(Sq \log(q)/D)$ .

An immediate application of the switching lemma is that if  $\mathcal{A}$  is an  $S$ -bounded adversary which makes  $q$  non-repeating queries to its oracle, then  $|\text{Adv}_E^{\text{prf}}(\mathcal{A}) - \text{Adv}_E^{\text{prp}}(\mathcal{A})| \leq \text{Adv}_D^{\text{sl}}(q, S)$  for any block-cipher E whose range has size  $D$ .

**AXU HASH FUNCTION.** Let  $H : \{0, 1\}^k \times (\{0, 1\}^* \times \{0, 1\}^*) \rightarrow \{0, 1\}^n$  be a hash function. Its almost XOR-universal (AXU) security is defined by the game  $G_{\text{H}}^{\text{axu}}$  shown in Fig. 2. In it, an adversary  $\mathcal{X}$  attempts to guess the xor of the output of H on two distinct inputs of its choosing for a random key  $L$ . We define  $\text{Adv}_{\text{H}}^{\text{axu}}(\mathcal{X}) = \Pr[G_{\text{H}}^{\text{axu}}(\mathcal{X})]$ . Typically one makes use of a  $c$ -AXU hash which for all  $\mathcal{X}$  satisfy  $\text{Adv}_{\text{H}}^{\text{axu}}(\mathcal{X}) \leq c \cdot (N_1 + N_2)/2^n$  where  $N_1$  (resp.  $N_2$ ) is the maximum block length of any  $A$  (resp.  $C$ ) output by  $\mathcal{X}$ . Note this is unconditional, so we will not have to worry about memory complexity when reducing to AXU security.

### 3 Nonce-based Encryption and Memory-boundedness

In this section we recall known definitions and results for nonce-based encryption [15]. We carefully consider how these change when we move to the memory-bounded setting. For example, as was previously noted by Auerbach, et al. [1], definitions which are tightly equivalent when the memory usage of adversaries is not bounded do not necessarily remain so with bounds on memory. So we will consider several variants of the definitions we are recalling and try to reason about which is the “correct” one to use. We additionally note some results which can be extended to give appealing time-memory tradeoffs in the memory-bounded setting and some for which this does not seem to be possible.

In Section 3.1, we discuss INDR security which measures the indistinguishability of ciphertexts from truly random ones. This security notion requires that the adversary be disallowed from repeating nonces. We discuss three conventions for capturing this which are tightly equivalent when ignoring memory restrictions, but observe they are no longer tightly equivalent with these restrictions. Based on these discussions, the rest of the paper focuses on the restricted class of adversaries that will never repeat nonces in their queries to encryption oracles. In Section 3.2, we discuss CTXT (integrity of ciphertexts) and AE security (combined INDR and CTXT) security. For these, the adversary must be disallowed from trivially winning by forwarding

ciphertexts from its encryption oracle to its decryption oracle. Again we discuss several conventions for this which are tightly equivalent when ignoring memory restrictions. Based on these discussions, the rest of the paper will use the convention that if an adversary queries  $(N, C)$  to its decryption oracle after receiving  $C$  from an encryption query for  $(N, M)$ , the oracle will respond with  $M$ . With our chosen conventions, it does not appear to be possible to prove that AE security is implied by INDR and CTXT security with a memory-tight reduction. The rest of the paper will focus on this (im)possibility. Section 4 shows it *is* possible in the restricted setting of secure channels while Section 5 shows it is not possible for general nonce-based encryption if the reduction behaves in a black-box manner.

Finally, in Section 3.3 we recall the CAU scheme by Bellare and Tackmann [4], an abstraction of GCM [12]. Following existing proofs [4,10,12] and using [11,7], we show that INDR security of CAU can be proven by a memory-tight reduction to PRP security with an appealing time-memory tradeoff and we informally discuss why such reductions seem impossible for CTXT or AE security.

**SYNTAX AND CORRECTNESS.** A (nonce-based) encryption scheme  $\text{NE}$  is defined by algorithms  $\text{NE.Kg}$ ,  $\text{NE.D}$ , and  $\text{NE.E}$ . Additionally it is associated with message space  $\text{NE.M} \subseteq \{0, 1\}^*$  and nonce space  $\text{NE.N}$ .

The syntax of the algorithms is shown in Fig. 3. The key generation algorithm  $\text{NE.Kg}$  takes no input and returns key  $K$ . The encryption algorithm  $\text{NE.E}$  takes key  $K$ , nonce  $N \in \text{NE.N}$ , and message  $M \in \text{NE.M}$ . It returns ciphertext  $C$ . The decryption algorithm  $\text{NE.D}$  takes key  $K$ , nonce  $N \in \text{NE.N}$ , and ciphertext  $C$ . It returns message  $M \in \text{NE.M} \cup \{\perp\}$ . When  $M = \perp$ , the ciphertext is rejected as invalid.

We additionally assume there is a ciphertext-length function  $\text{NE.cl} : \mathbb{N} \rightarrow \mathbb{N}$  such that for any  $K \in [\text{NE.Kg}]$ ,  $N \in \text{NE.N}$ , and  $M \in \text{NE.M}$  we have  $|C| = \text{NE.cl}(|M|)$  whenever  $C \leftarrow \text{NE.E}(K, N, M)$ . Typically, a nonce-based encryption scheme also takes associated data as input which is authenticated during encryption. Associated data does not meaningfully effect our results, so we have omitted it for simplicity of notation.

$\text{NE Syntax}$ $K \leftarrow \text{NE.Kg}$ $C \leftarrow \text{NE.E}(K, N, M)$ $M \leftarrow \text{NE.D}(K, N, C)$
---

Correctness of an encryption scheme requires for all  $K \in [\text{NE.Kg}]$ ,  $N \in \text{NE.N}$ , and  $M \in \text{NE.M}$  that  $\text{NE.D}(K, N, \text{NE.E}(K, N, M)) = M$ .

**Fig. 3.** Syntax of nonce-based encryption scheme.

### 3.1 Indistinguishability From Random (INDR) Security

The first security notion we will consider requires that ciphertexts output by the encryption scheme cannot be distinguished from ciphertexts chosen at random.

**DEFINITIONS.** Consider the game  $G_{\text{NE},b}^{\text{indr}}$  shown in Fig. 4. Here an adversary  $\mathcal{A}$  is given access to an encryption oracle  $\text{ENC}$  to which it can query a pair  $(N, M)$  and receive back either the encryption of message  $M$  with nonce  $N$  ( $b = 1$ ) or a random string of the appropriate length ( $b = 0$ ). The adversary outputs a bit trying to guess which of these two views it was given. We define  $\text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{A}) = \Pr[G_{\text{NE},1}^{\text{indr}}(\mathcal{A})] - \Pr[G_{\text{NE},0}^{\text{indr}}(\mathcal{A})]$ .

In defining security we must address how to handle the possibility of  $\mathcal{A}$  making multiple queries with the same nonce. Encryption schemes are typically designed under the assumption that the same nonce will not be used multiple times and may become completely insecure in the face of such nonce repetition. The primary convention we will adopt is to restrict attention to adversaries that will never repeat nonces in their encryption queries. We use the phrase “nonce-respecting INDR” to refer to security with respect to such adversaries.

An alternate approach would be to modify the code of the game to respond appropriately to queries where nonces repeat. One version of this, which we will refer to as INDR-R, would restrict attention to adversaries that will only repeat nonces when they also repeat the message queried to encryption. For this the game would be modified to keep track of all encryption queries that have been made so far. When it receives a repeated  $(N, M)$  pair, it simply returns the same  $C$  that it returned last time it saw that pair. A second version of this, which we will refer to as INDR-B, makes no restriction on the queries of the adversary. Instead, the game is modified to return  $\perp$  whenever the adversary makes a query with a nonce it has already used.

Game $G_{NE,b}^{\text{indr}}(\mathcal{A})$ $K \leftarrow_s \text{NE.Kg}$ $b' \leftarrow \mathcal{A}^{\text{ENC}_b}$ Return $b' = 1$	Game $G_{NE,b}^{\text{ctxt-}w}(\mathcal{A})$ $K \leftarrow_s \text{NE.Kg}$ $b' \leftarrow \mathcal{A}^{\text{ENC}_1, \text{DEC}_b^w}$ Return $b' = 1$	Oracle $\text{DEC}_b^w(N, C)$ If $M[N, C] \neq \perp$ then Return $M[N, C]$ if $w = 1$ Return $\diamond$ if $w = 2$ Return $\perp$ if $w = 3$ $M_1 \leftarrow \text{NE.D}(K, N, C)$ $M_0 \leftarrow \perp$ Return $M_b$
Oracle $\text{ENC}_b(N, M)$ $C_1 \leftarrow \text{NE.E}(K, N, M)$ $C_0 \leftarrow_s \{0, 1\}^{\text{NE.cl}( M )}$ $M[N, C_b] \leftarrow M$ Return $C_b$	Game $G_{NE,b}^{\text{ae-}w}(\mathcal{A})$ $K \leftarrow_s \text{NE.Kg}$ $b' \leftarrow \mathcal{A}^{\text{ENC}_b, \text{DEC}_b^w}$ Return $b' = 1$	

**Fig. 4.** Games defining INDR, CTXT- $w$ , and AE- $w$  security of NE for  $w \in \{1, 2, 3\}$ .

DISCUSSION. When memory is not an issue, all of these variants would be equivalent. Proving this follows by noting that an adversary can just remember all prior queries it has made and thus never need to repeat. This proof strategy is no longer available to us when we want to preserve the memory usage of adversaries. We focus on nonce-respecting INDR because it hits the sweet spot of being strong enough for common applications, yet weak enough that we know how to give provable time-memory trade-offs.

Because nonce-respecting INDR considers a strictly smaller class of adversaries than the other two and all of the games behave identically for this class of adversary it is tightly implied by the others. In fact, using ideas from [11,7] we can see that nonce-respecting INDR is strictly weaker. The toy encryption scheme SE considered in the introduction built from a block-cipher with block length  $n$  is vulnerable to low-memory collision-finding attacks with advantage  $\Omega(q^2/2^n)$  in the INDR-R and INDR-B settings, but no attacks can have advantage better than  $O(qs/2^n)$  in the nonce-respecting INDR setting. Here  $q$  and  $s$  refer to the number of queries and amount of memory used by the attackers, respectively. This underlies why the ideas of Jaeger and Tessaro [11] can be used to prove nonce-respecting INDR (but not INDR-R or INDR-B) time-memory trade-offs for natural counter-mode based encryption schemes. In most common uses of nonce-based encryption the nonces are incremented as a counter or picked uniformly at random. In the former case, nonces clearly never repeat so nonce-respecting INDR suffices (we will see this formally in Section 4). Nonces may repeat in the latter case, but we can follow [11,7] here and replace the uniform random values with random, non-repeating values so again nonce-respecting INDR suffices.

### 3.2 Security Beyond Confidentiality

INDR security only guarantees confidentiality of the messages against passive attackers. However, in practice, attackers may actively modify ciphertexts in transit. As such, it is important to consider security definitions that take this into account. We will consider integrity definitions and authenticated encryption definitions which simultaneously asked for integrity and confidentiality.

DEFINITIONS. Consider the other two games shown in Fig. 4. We will first focus on  $G_{NE,b}^{\text{ae-}w}$  which defines three variants of authenticated encryption security parameterized by  $w \in \{1, 2, 3\}$ . In this game, the adversary is given access to an encryption oracle and a decryption oracle. Its goal is to distinguish between a “real” and “ideal” world. In the real world ( $b = 1$ ) the oracles uses NE to encrypt messages and decrypt ciphertexts. In the ideal world ( $b = 0$ ) encryption returns random messages of the appropriate length and decryption returns  $\perp$ . For simplicity, we will restrict attention nonce-respecting adversaries which do not repeat nonces across encryption queries (as in nonce-respecting INDR security). Note there is no restriction placed on nonces used for decryption queries. Integrity of ciphertext security is defined by  $G_{NE,b}^{\text{ctxt-}w}$  which behaves similarly except the adversary is always given access to the real encryption algorithm.

The decryption oracle needs to prevent trivial attacks. If the adversary receives  $C$  from a query of  $\text{ENC}(N, M)$  and then queries  $\text{DEC}(N, C)$  it would receive  $M$  in the real world and  $\perp$  in the ideal world, making them easy to distinguish. We must adopt some convention for how the oracles behave when such a query is made to prevent this type of trivial attack. Towards this, the decryption oracle is parameterized by



the value  $w \in \{1, 2, 3\}$  corresponding to three different security notions. In all three, we use a table  $M[\cdot, \cdot]$  to detect when the adversary forwards encryption queries on to its decryption oracle. When  $w = 1$ , the decryption oracle returns  $M[N, C]$  in this case. When  $w = 2$ , it returns a special symbol  $\diamond$ . When  $w = 3$ , it returns the symbol  $\perp$  which is also used by the encryption scheme to represent rejection. For  $x \in \{\text{ae}, \text{ctxt}\}$  and  $w \in \{1, 2, 3\}$  we define the advantage of an adversary  $\mathcal{A}$  by  $\text{Adv}_{\text{NE}}^{x-w}(\mathcal{A}) = \Pr[\text{G}_{\text{NE},1}^{x-w}(\mathcal{A})] - \Pr[\text{G}_{\text{NE},0}^{x-w}(\mathcal{A})]$ . The corresponding security notions are referred to as AE- $w$  and CTXT- $w$ .

DISCUSSION. When memory usage is not an issue, the choice of  $w$  does not matter. We can without loss of generality assume that the adversary never makes one of these trivial attack queries because it could simply store the table  $M[\cdot, \cdot]$  for itself and simulate any such queries.<sup>4</sup> It’s not clear that this equivalence holds if we do not assume that storing  $M[\cdot, \cdot]$  is “free” for the adversary.

The only memory-tight implication we are aware of between these is that security for  $w = 2$  tightly implies security for  $w = 3$ . This follows because an adversary with access to  $\text{DEC}_b^2$  can simulate  $\text{DEC}_b^3$  with low memory. If  $\text{DEC}_b^2$  returns  $M = \diamond$  the adversary returns  $\perp$ , otherwise it does not modify  $M$ . All of the other implications we might want to show seem to require remembering all prior encryption queries to properly simulate DEC.

Ultimately, for heuristic reasons, we believe that  $w = 1$  is the “correct” choice and will focus on it in our later sections. The typical motivation behind chosen-ciphertext security notions is that in practice an attacker can often observe the behavior of the decrypting party to learn something about the message they received. There is no reason to think an attacker should only be able to do that for ciphertexts that have been modified, but not ciphertexts that have been unmodified. This is best captured by  $w = 1$ . The  $w = 2$  definition seems to posit that the adversary can distinguish between ciphertexts it forwarded on and ciphertexts that it modified (whether they were accepted or rejected) by observing the decrypting party’s behavior. The  $w = 3$  definition seems to posit that the adversary cannot learn anything about ciphertexts it forwards on unmodified, but can learn about other modified ciphertexts by observing the decrypting party’s behavior.

REVISITING A CLASSIC RESULT. A classic result, which has been shown for numerous styles of encryption, is that confidentiality and integrity together imply authenticated encryption [16]. However, this becomes more difficult for nonce-based encryption when we consider memory-tightness.

The classic proof that INDR and CTXT-1 security imply AE-1 security first replaces real decryption with  $\perp$  via a reduction to CTXT-1 security and then replace real encryption with random using INDR security. However, in this second step the reduction adversary would have to simulate the oracle  $\text{DEC}_0^1$  which seems to require storing the table  $M[\cdot, \cdot]$ .<sup>5</sup> This potentially requires using much more memory than the AE-1 adversary, losing the benefit of time-memory tradeoffs for INDR-R. The rest of the paper is dedicated to understanding this reduction. In Section 4.2, we make it memory tight when restricting attention to secure channels which only accept ciphertexts if they are received in order. In Section 5, we give negative results showing that for nonce-based encryption this reduction cannot be made memory tight (using a black-box reduction).

### 3.3 Security of the CAU Encryption Scheme

We conclude this section by considering the specific encryption scheme CAU for which we can prove INDR security with a time-memory tradeoff. We will use this scheme in Section 4 to show a time-memory tradeoff for the authenticated encryption security of a channel instantiated with it.

One of the most widely deployed encryption schemes is Galois Counter-Mode (GCM) [12]. Bellare and Tackmann [4] generalized it to the scheme CAU which constructs an encryption scheme from a block cipher E and hash function H. Using the techniques of Jaeger and Tessaro [11] we obtain a proof of security for its nonce-respecting INDR security with an appealing time-memory tradeoff.

<sup>4</sup> Restricting attention to adversaries which never make trivial attack queries is, indeed, a fourth way one could define security.

<sup>5</sup> The standard reduction *would* be memory tight for  $w = 3$ .

Algorithm CAU[E, H].E(K, N, M)	Algorithm CAU[E, H].D(K, N, T    C)
$Y \leftarrow \text{pad}(N)$	$L \leftarrow E_K(0^n); Y \leftarrow \text{pad}(N)$
$M_1 \dots M_\ell \leftarrow_n M$	$C_1 \dots C_\ell \leftarrow_n C$
For $i = 1, \dots, \ell$ do	$T' \leftarrow H_L(A, C) \oplus E_K(Y)$
$C_i \leftarrow M_i \oplus E_K(Y + i)$	If $T \neq T'$ then return $\perp$
$C \leftarrow C_1 \dots C_\ell$	For $i = 1, \dots, \ell$ do
$L \leftarrow E_K(0^n)$	$M_i \leftarrow C_i \oplus E_K(Y + i)$
$T \leftarrow H_L(A, C) \oplus E_K(Y)$	$M \leftarrow M_1 \dots M_\ell$
Return $T    C$	Return $M$

**Fig. 5.** Encryption scheme CAU parameterized by function family E (typically a block cipher) and hash function H. In the code,  $\text{pad}(N) = N || 0^m || 1$  for the appropriate choice of  $m$  and  $M_1 \dots M_\ell \leftarrow_n M$  splits  $M$  into  $n$ -bit blocks.

CONSTRUCTION. We recall the CAU construction of an encryption scheme. Fix a key length  $\text{CAU.kl} \in \mathbb{N}$ , a block length  $n = \text{CAU.bl} \in \mathbb{N}$ , and a nonce length  $\text{CAU.nl} < \text{CAU.bl}$ . Then let E be a function family with  $E : \{0, 1\}^{\text{CAU.kl}} \times \{0, 1\}^{\text{CAU.bl}} \rightarrow \{0, 1\}^{\text{CAU.bl}}$  and H be a function family with  $H : \{0, 1\}^{\text{CAU.bl}} \times (\{0, 1\}^* \times \{0, 1\}^*) \rightarrow \{0, 1\}^{\text{CAU.bl}}$ . The scheme constructed from E and H is denoted  $\text{CAU}[E, H]$ . Its message space  $\text{CAU}[E, H].M$  is the set of all strings of length at most  $n \cdot (2^{n-\text{CAU.nl}} - 1)$  and its nonce space  $\text{CAU}[E, H].N$  is the set  $\{0, 1\}^{\text{CAU.nl}}$ .

The algorithms of  $\text{CAU}[E, H]$  are shown in Fig. 5. The code uses  $\text{pad}(\cdot)$  to denote the padding function which on input  $N$  outputs  $N || 0^{n-\text{CAU}[E, H].nl-1} || 1$ . Since our simplified notation does not use associated data we instead assume there is a fixed associated data string  $A$  used with every message.

The encryption algorithm parses the input message into  $\ell$  blocks of length  $n$  (except for the last, which may be shorter) and pads the nonce to a string  $Y$  of length  $n$ . It encrypts the message using counter-mode encryption with  $Y + 1$  as the first counter. This gives it a partial ciphertext  $C$ . The authentication is inspired by a Carter-Wegman MAC. A key  $L$  for the hash function is obtained as  $L \leftarrow E_K(0^n)$ . This key is used to compute the tag  $T$  as  $T \leftarrow H_L(A, C) \oplus E_K(Y)$  and then  $T || C$  is the full ciphertext output by encryption.

The decryption algorithm parses the input ciphertext as  $T || C$ . It computes the correct tag  $T'$  for  $C$  by setting  $L \leftarrow E_K(0^n)$  and  $T' \leftarrow H_L(A, C) \oplus E_K(Y)$  (as was done in encryption). If  $T \neq T'$  the ciphertext is rejected by returning  $M = \perp$ . Otherwise the message  $M$  is obtained by counter-mode decrypting  $C$ .

INDR SECURITY OF CAU. The following theorem formalizes that CAU is nonce-respecting INDR secure assuming E is a secure PRF.

**Theorem 1.** *Let  $\mathcal{A}$  be an adversary against the nonce-respecting INDR security of  $\text{CAU}[E, H]$  that makes at most  $q$  oracle queries, each at most  $p \cdot \text{CAU.bl}$  bits long. Then we can construct a  $\mathcal{A}_{\text{prf}}$  such that*

$$\text{Adv}_{\text{CAU}[E, H]}^{\text{indr}}(\mathcal{A}) \leq \text{Adv}_{\text{E}}^{\text{prf}}(\mathcal{A}_{\text{prf}}) .$$

*Adversary  $\mathcal{A}_{\text{prf}}$  has runtime essentially that of  $\mathcal{A}$ , makes at most  $q(p + 1) + 1$  queries to its oracle, has memory/time complexity essentially that of  $\mathcal{A}$  and never repeats queries to its oracle.*

It is important that  $\mathcal{A}_{\text{prf}}$  never repeats queries because it allows us to apply the time-memory switching lemma from Section 2. This gives us roughly,

$$\text{Adv}_{\text{CAU}[E, H]}^{\text{indr}}(\mathcal{A}) \in \text{Adv}_{\text{E}}^{\text{pp}}(\mathcal{A}_{\text{prf}}) + O(S \cdot pq \cdot \log(pq)/2^n)$$

where  $S$  is a bound on the memory complexity of  $\mathcal{A}$ . For variants other than nonce-respecting INDR it would not be clear how to prevent  $\mathcal{A}_{\text{prf}}$  from repeating queries without storing the prior queries of  $\mathcal{A}$ .

*Proof (Sketch).* One constructs  $\mathcal{A}_{\text{prf}}$  to first set  $L \leftarrow \text{EVAL}(0^n)$ . Then it runs  $\mathcal{A}$  and simulates encryption queries by running  $\text{CAU.E}$  while using its EVAL oracle in place of  $E_K$ . It does not recompute  $L$  each time because it has already computed it. Its final output is whatever  $\mathcal{A}$  outputs. One can verify that the view of  $\mathcal{A}$  when simulated by  $\mathcal{A}_{\text{prf}}$  is “real” encryptions when  $b = 1$  and random strings when  $b = 0$ , so the claimed advantage bound follows.  $\square$

<b>CH Syntax</b> $(\sigma^s, \sigma^r) \leftarrow \text{CH.Sg}$ $(\sigma^s, C) \leftarrow \text{CH.S}(\sigma^s, M)$ $(\sigma^r, M) \leftarrow \text{CH.R}(\sigma^r, C)$	<b>Game <math>G_{\text{CH},b}^{\text{ch-corr}}(\mathcal{A})</math></b> $(\sigma^s, \sigma^r) \leftarrow \text{CH.Sg}$ $b' \leftarrow \mathcal{A}^{\text{ENCDEC}}$ Return $b' = 1$	<b>Oracle <math>\text{ENCDEC}_b(M_0)</math></b> $(\sigma^s, C) \leftarrow \text{CH.S}(\sigma^s, M_0)$ $(\sigma^r, M_1) \leftarrow \text{CH.R}(\sigma^r, C)$ Return $M_b$
--	--	---

**Fig. 6.** Left: Syntax of channel algorithms. Right: Channel correctness game.

CTXT/AE SECURITY OF CAU. It does not appear to be possible to give a similar time-memory trade-off for the CTXT or AE security of CAU. The standard analysis of either of these first uses PRF security to replace the output of E with random. It then argues that the adversary’s view is independent of the  $H_L(A, C)$  values produced in encryption so that it can apply the security of H. For  $x = \text{ae}$  or  $x = \text{ctxt}$  this would give a bound of the form,

$$\text{Adv}_{\text{CAU}[\text{E},\text{H}]}^{x-1}(\mathcal{A}) = \text{Adv}_{\text{E}}^{\text{prf}}(\mathcal{A}_{\text{prf}}) + \text{Adv}_{\text{H}}^{\text{axu}}(\mathcal{X}).$$

However, this PRF adversary  $\mathcal{A}_{\text{prf}}$  needs to simulate a decryption oracle to  $\mathcal{A}$ . The natural ways of doing this (remembering all prior encryption queries or using EVAL to run decryption) either require significant use of memory or repeating queries to EVAL. This prevents us from applying the switching lemmas of [11,7] to get appealing time-memory tradeoffs when E is a PRP.

In Section 4.3, we will use a new technique for memory-tight reductions to prove that using CAU in a channel can provide (the channel equivalent of) CTXT security (and thus AE security from Section 4.2).

## 4 Memory-tight Reductions for Cryptographic Channels

In this section we show that memory-tight reductions can be given for the combined confidentiality and integrity security of cryptographic channels. These are a form of stateful encryption which provide the guarantee that messages cannot be duplicated or reordered, in addition to the typical confidentiality and integrity goals of encryption.

### 4.1 Syntax and Security Notions

SYNTAX AND CORRECTNESS. A (cryptographic) channel CH specifies algorithms CH.Sg, CH.S, and CH.R along with message space  $\text{CH.M} \subseteq \{0, 1\}^*$ . The syntax of these algorithms is shown in Fig. 6. The state generation algorithm CH.Sg takes no input. It returns sender state  $\sigma^s$  and receiver state  $\sigma^r$ . The sending algorithm CH.S takes a sender state  $\sigma^s$  and message  $M \in \text{CH.M}$ . It returns updated sender state  $\sigma^s$  and a ciphertext  $C$ . The receiving algorithm CH.R takes a receiver state  $\sigma^r$  and a ciphertext  $C$ . It returns updated receiver state  $\sigma^r$  and a message  $M \in \text{CH.M} \cup \{\perp\}$ . When  $M = \perp$ , this represents the receiver rejecting the message as invalid.

A channel is expected to never again return  $M \neq \perp$  after if it has rejected a message. This models the behavior of protocols such as TLS which are assumed to be run over a reliable transport layer and has been the standard notion for channels since the work of Bellare, Kohno, and Namprempe [2]. When a protocol (e.g. QUIC or DTLS) is run over an unreliable transport layer, then a *robust* channel is used instead [8]. We leave memory-tight proofs of security for robust channels as an interesting direction for future work.

We typically assume there is a ciphertext-length function  $\text{CH.cl} : \mathbb{N} \rightarrow \mathbb{N}$  such that for any  $M \in \text{CH.M}$  and state  $\sigma^s$ , we have  $\Pr[|C| = \text{CH.cl}(|M|) : (\sigma^s, C) \leftarrow \text{CH.S}(\sigma^s, M)] = 1$ .

Correctness requires that if the receiver is given the ciphertexts sent by the sender in order and without modification then the receiver will output the same sequence of messages that were sent. One way to formalize this is via the game  $G_{\text{CH},b}^{\text{ch-corr}}$  shown in Fig. 6. We define  $\text{Adv}_{\text{CH}}^{\text{ch-corr}}(\mathcal{A}) = \Pr[G_{\text{CH},1}^{\text{ch-corr}}(\mathcal{A})] - \Pr[G_{\text{CH},0}^{\text{ch-corr}}(\mathcal{A})]$ . Perfect correctness requires that  $\text{Adv}_{\text{CH}}^{\text{ch-corr}}(\mathcal{A}) = 0$  for all (even unbounded)  $\mathcal{A}$ . This implies that the  $M_1$  output by CH.R always equals  $M_0$ .

<p>Game <math>G_{\text{CH},b}^{\text{ch-ae}}(\mathcal{A})</math></p> <p>sync <math>\leftarrow</math> true</p> <p><math>(\sigma^s, \sigma^r) \leftarrow</math> CH.Sg</p> <p><math>b' \leftarrow</math> <math>\mathcal{A}^{\text{ENC}_b, \text{DEC}_b}</math></p> <p>Return <math>b' = 1</math></p>	<p>Game <math>G_{\text{CH},b}^{\text{ch-ctxt}}(\mathcal{A})</math></p> <p>sync <math>\leftarrow</math> true</p> <p><math>(\sigma^s, \sigma^r) \leftarrow</math> CH.Sg</p> <p><math>b' \leftarrow</math> <math>\mathcal{A}^{\text{ENC}_1, \text{DEC}_b}</math></p> <p>Return <math>b' = 1</math></p>	<p>Oracle <math>\text{DEC}_b(C)</math></p> <p><math>(\sigma^r, M_1) \leftarrow</math> CH.R(<math>\sigma^r, C</math>)</p> <p><math>M_0 \leftarrow \perp</math></p> <p><math>M' \leftarrow</math> M.dq()</p> <p><math>C' \leftarrow</math> C.dq()</p> <p>If sync then</p> <p>  If <math>C = C'</math> then</p> <p>    Return <math>M'</math></p> <p>  sync <math>\leftarrow</math> false</p> <p>Return <math>M_b</math></p>
<p>Game <math>G_{\text{CH},b}^{\text{ch-indr}}(\mathcal{A})</math></p> <p>sync <math>\leftarrow</math> true</p> <p><math>(\sigma^s, \sigma^r) \leftarrow</math> CH.Sg</p> <p><math>b' \leftarrow</math> <math>\mathcal{A}^{\text{ENC}_b}</math></p> <p>Return <math>b' = 1</math></p>	<p>Oracle <math>\text{ENC}_b(M)</math></p> <p><math>(\sigma^s, C_1) \leftarrow</math> CH.S(<math>\sigma^s, M</math>)</p> <p><math>C_0 \leftarrow</math> <math>\{0, 1\}^{\text{CH.cl}( M )}</math></p> <p>M.add(<math>M</math>); C.add(<math>C_b</math>)</p> <p>Return <math>C_b</math></p>	

**Fig. 7.** Games defining the INDR, CTXT, and AE security of a channel.

<p>Game <math>G_{L,\delta}^{\text{it}}(\mathcal{A}_1, \mathcal{A}_2)</math></p> <p><math>R \leftarrow</math> <math>\{0, 1\}^L</math></p> <p><math>(i, \sigma) \leftarrow</math> <math>\mathcal{A}_1(R)</math></p> <p><math>r \leftarrow</math> <math>\mathcal{A}_2(i, \sigma, R[i : i-1])</math></p> <p>Return <math>r = R[i : i + \delta]</math></p>
---

**Fig. 8.** Information theoretic game in which  $\mathcal{A}$  tries to remember a  $\delta$  bit sequence in an  $L$ -bit random string.

SECURITY DEFINITIONS. We consider indistinguishability from random, integrity of ciphertext, and authenticated encryption security for channels just like we did for nonce based encryption.

Authenticated encryption security of a channel CH is defined by game  $G_{\text{CH},b}^{\text{ch-ae}}$  defined in Fig. 7. In it the adversary is given access to an encryption oracle and a decryption oracle. The adversary’s goal is to distinguish between a “real” and “ideal” world. In the real world ( $b = 1$ ) the oracles use CH to encrypt messages and decrypt ciphertexts. In the ideal world ( $b = 0$ ) encryption returns random messages of the appropriate length and decryption returns  $\perp$ . In both worlds, as long as the adversary’s queries to decryption have consisted of the outputs of encryption in the correct order, the oracles are considered in sync and decryption just returns the appropriate message that was queried to encryption.<sup>6</sup> After the first time the adversary queries something else, the oracles are out of sync and will never be in sync again (so DEC will always return  $M_b$ ).

Authenticated encryption security is a combined confidentiality and integrity notion. We can also define separate notions. INDR security is defined by the game  $G_{\text{CH},b}^{\text{ch-indr}}$  which is the same as  $G_{\text{CH},b}^{\text{ch-ae}}$  except the adversary is only given oracle access to  $\text{ENC}_b$ . CTXT security is defined by the game  $G_{\text{CH},b}^{\text{ch-ctxt}}$  which is the same as  $G_{\text{CH},b}^{\text{ch-ae}}$  except the adversary is given oracle access to  $\text{ENC}_1$  and  $\text{DEC}_b$ . These games are given explicitly in Fig. 7. We define the advantage of  $\mathcal{A}$  by  $\text{Adv}_{\text{CH}}^x(\mathcal{A}) = \Pr[G_{\text{CH},1}^x(\mathcal{A})] - \Pr[G_{\text{CH},0}^x(\mathcal{A})]$  for  $x \in \{\text{ch-ae}, \text{ch-indr}, \text{ch-ctxt}\}$ .

## 4.2 Confidentiality and Integrity Imply Authenticated Encryption

We will show that INDR security plus CTXT security imply AE security using a memory-tight reduction. While the normal proof that INDR and CTXT security suffice to imply AE security is not particularly difficult, it uses a non-memory tight reduction to INDR security. Making the proof memory tight will require more involved analysis.

INFORMATION THEORETIC LEMMA. Before proceeding to the proof, we first will provide a simple information theoretic lemma that will be a useful subcomponent of that proof. Consider the game  $G_{L,\delta}^{\text{it}}$  shown in Fig. 8. In it, an adversary is given a length  $L$  string  $R$  and tries to choose an index  $i$  for which it is able to remember the next  $\delta$ -bits of the string using state  $\sigma$ . We say that an adversary  $(\mathcal{A}_1, \mathcal{A}_2)$  is  $S$ -bounded if  $|\sigma| = S$  always. We define  $\text{Adv}_{L,\delta}^{\text{it}}(\mathcal{A}_1, \mathcal{A}_2) = \Pr[G_{L,\delta}^{\text{it}}(\mathcal{A}_1, \mathcal{A}_2)]$ .

**Lemma 1.** *Let  $L, \delta, S \in \mathbb{N}$ . Let  $(\mathcal{A}_1, \mathcal{A}_2)$  be an  $S$ -bounded adversary. Then*

$$\text{Adv}_{L,\delta}^{\text{it}}(\mathcal{A}_1, \mathcal{A}_2) \leq L \cdot 2^S / 2^\delta .$$

<sup>6</sup> This matches the convention of CTXT-1 and AE-1 for encryption schemes. We believe it to be “correct” for the same reasons discussed for those definitions.

*Proof.* Let  $L, \delta, S, \mathcal{A}_1, \mathcal{A}_2$  be defined as in the theorem statement. Without loss of generality we can assume that  $\mathcal{A}_1$  and  $\mathcal{A}_2$  are deterministic. Then for any fixed choice of  $i$  and  $\sigma$ , the probability that  $\mathcal{A}_2(i, \sigma, R[: i - 1]) = R[i : i + \delta]$  will be exactly  $1/2^\delta$ . Then we can calculate as follows.

$$\begin{aligned} \Pr[\mathsf{G}_{L,n}^{\text{it}}(\mathcal{A}_1, \mathcal{A}_2)] &\leq \Pr_R[\exists i, \sigma \text{ s.t. } R[i : i + \delta] = \mathcal{A}_2(i, \sigma, R[: i - 1])] \\ &\leq \sum_{i, \sigma} \Pr[R[i : i + \delta] = \mathcal{A}_2(i, \sigma, R[: i - 1])] \\ &= \sum_{i, \sigma} 1/2^\delta \leq L \cdot 2^S / 2^\delta . \end{aligned}$$

The last inequality follows from there being at most  $L \cdot 2^S$  choices for  $(i, \sigma)$ .  $\square$

**SECURITY RESULT.** Now we can proceed to our security result showing that AE security can be implied by INDR and CTXT security in a memory-tight manner. The technical crux of the result is the reduction adversary  $\mathcal{A}_\delta$  which simulates the view of an AE adversary  $\mathcal{A}$  to attack the INDR security of the channel. In our theorem statement this reduction adversary is parameterized by a variable  $\delta$  which determines how much local memory it uses. Using Lemma 1, our concrete advantage bound is expressed in terms of  $\delta$  and establishes that the reduction can be successful with this value not much larger than the local memory of  $\mathcal{A}$ .

**Theorem 2.** *Let CH be a cryptographic channel. Let  $\mathcal{A}$  be an adversary with memory complexity  $S$  and making at most  $q$  queries to its ENC oracle, each of which returns a ciphertext of length at most  $x$ . Then for any  $\delta \in \mathbb{N}$  we can build an adversary  $\mathcal{A}_\delta$  (described in the proof) such that*

$$\text{Adv}_{\text{CH}}^{\text{ch-ae}}(\mathcal{A}) \leq \text{Adv}_{\text{CH}}^{\text{ch-ctxt}}(\mathcal{A}) + 2 \cdot \text{Adv}_{\text{CH}}^{\text{ch-indr}}(\mathcal{A}_\delta) + 2q \cdot x \cdot 2^S / 2^\delta .$$

*Adversary  $\mathcal{A}_\delta$  has running time approximately that of  $\mathcal{A}$  and uses about  $S + 2\delta$  bits of state.*

Setting  $\delta = S + \log(qx) + \kappa$  makes the last term about  $1/2^\kappa$  while limiting the memory usage of  $\mathcal{A}_\delta$  to only  $2S + 2\log(qx) + 2\kappa$ .

The standard way of proving that INDR security and CTXT security imply AE security would first use CTXT security to transition from a world in which  $\mathcal{A}$  is given oracle access to  $(\text{ENC}_1, \text{DEC}_1)$  to a world in which  $\mathcal{A}$  is given oracle access to  $(\text{ENC}_1, \text{DEC}_0)$ . Then INDR security would be used to transition to  $\mathcal{A}$  being given oracle access to  $(\text{ENC}_0, \text{DEC}_0)$ . The issue in our setting with this proof arises in the second step. The INDR reduction adversary needs to simulate  $\text{DEC}_0$  for  $\mathcal{A}$ . The natural way of doing so requires storing the entirety of the tables  $\mathbf{M}$  and  $\mathbf{C}$  which means that  $\mathcal{A}_\delta$  may use much more memory than  $\mathcal{A}$ .

Our proof of Thm. 2 follows this same general proof flow, but uses a more involved analysis for the reduction to INDR security. In particular, we make use of the following insight: If  $\mathcal{A}$  has memory complexity  $S$  but cannot distinguish the ciphertexts it sees from random (because of INDR security), then from Lemma 1 it cannot remember many more than  $S$  of the ciphertext bits that it has received from ENC but not yet forwarded to DEC.

If  $\mathcal{A}$  ever queries a ciphertext which is not the next ciphertext in  $\mathbf{C}$ , then  $\text{DEC}_0$  oracle will never again return anything other than  $\perp$ . Because we can assume that  $\mathcal{A}$  will be unable to remember too many bits of ciphertext, we can just have our reduction adversary  $\mathcal{A}_\delta$  remember a few more bits of ciphertext than  $\mathcal{A}$  can. If the total length of ciphertext that  $\mathcal{A}$  has received from its encryption oracle, but not forwarded on to its decryption oracle ever exceeds the amount that  $\mathcal{A}_\delta$  will store, then  $\mathcal{A}_\delta$  assumes  $\mathcal{A}$  must have forgotten some intermediate ciphertext before that point, allowing the reduction to cease storing future ciphertexts because `sync` will be `false` before that point.

*Proof.* We will construct INDR adversaries  $\mathcal{A}'_\delta, \mathcal{A}''_\delta$ , and  $S$ -bounded adversary  $(\mathcal{A}_1, \mathcal{A}_2)$  and show that

$$\text{Adv}_{\text{CH}}^{\text{ch-ae}}(\mathcal{A}) \leq \text{Adv}_{\text{CH}}^{\text{ch-ctxt}}(\mathcal{A}) + \text{Adv}_{\text{CH}}^{\text{ch-indr}}(\mathcal{A}'_\delta) + 2 \cdot \text{Adv}_{q,x,\delta}^{\text{it}}(\mathcal{A}_1, \mathcal{A}_2) + \text{Adv}_{\text{CH}}^{\text{ch-indr}}(\mathcal{A}''_\delta) .$$

The stated theorem then follows by applying Lemma 1 and constructing the adversary  $\mathcal{A}_\delta$  which runs either  $\mathcal{A}'_\delta$  or  $\mathcal{A}''_\delta$  (chosen at random) and outputs whatever that adversary does. The resulting  $\mathcal{A}_\delta$  will satisfy the efficiency constraints stated in the theorem statement. We will prove this bound via a sequence of transformations that slowly change  $\mathsf{G}_{\text{CH},1}^{\text{ch-ae}}$  to  $\mathsf{G}_{\text{CH},0}^{\text{ch-ae}}$ .

<p>Games <math>\boxed{G_2}, G_3, G_4, \boxed{G_5}</math></p> <p><b>flag</b> <math>\leftarrow</math> <b>true</b></p> <p><b>sync</b> <math>\leftarrow</math> <b>true</b></p> <p><math>(\sigma^s, \sigma^r) \leftarrow_s \text{CH.Sg}</math></p> <p><math>b' \leftarrow_s \mathcal{A}^{\text{ENC,DEC}}</math></p> <p>Return <math>b' = 1</math></p>	<p>Oracle ENC(<math>M</math>)</p> <p><math>(\sigma^s, C) \leftarrow_s \text{CH.S}(\sigma^s, M)</math></p> <p><math>C \leftarrow_s \{0, 1\}^{\text{CH.ct}( M )}</math></p> <p><math>\mathbf{M}_2.\text{add}(M); \mathbf{C}_2.\text{add}(C)</math></p> <p>If <b>flag</b> then</p> <p>  If <math>\ \mathbf{C}\  +  C  &lt; \delta</math> then</p> <p>    <math>\mathbf{M}.\text{add}(M); \mathbf{C}.\text{add}(C)</math></p> <p>  Else</p> <p>    <b>flag</b> <math>\leftarrow</math> <b>false</b></p> <p>Return <math>C</math></p>	<p>Oracle DEC(<math>C</math>)</p> <p><math>M' \leftarrow \mathbf{M}.\text{dq}(); M_2 \leftarrow \mathbf{M}_2.\text{dq}()</math></p> <p><math>C' \leftarrow \mathbf{C}.\text{dq}(); C_2 \leftarrow \mathbf{C}_2.\text{dq}()</math></p> <p>If <b>sync</b> then</p> <p>  If <math>C = C'</math> then</p> <p>    Return <math>M'</math></p> <p>  Elif <math>C = C_2</math> then</p> <p>    <b>bad</b> <math>\leftarrow</math> <b>true</b></p> <p>    <math>\boxed{\text{Return } M_2}</math></p> <p>  <b>sync</b> <math>\leftarrow</math> <b>false</b></p> <p>Return <math>\perp</math></p>
--	--	---

**Fig. 9.** Hybrid games for proof of Theorem 2. Highlighted code is only included in highlighted games. Boxed code is only included in boxed games.

CTXT TRANSITION. Let  $G_0 = G_{\text{CH},1}^{\text{ch-ae}}(\mathcal{A})$  and  $G_1 = G_{\text{CH},0}^{\text{ch-ctxt}}(\mathcal{A})$ . Because  $G_{\text{CH},1}^{\text{ch-ae}}(\mathcal{A})$  and  $G_{\text{CH},1}^{\text{ch-ctxt}}(\mathcal{A})$  are identical games we have that  $\Pr[G_0] - \Pr[G_1] = \text{Adv}_{\text{CH}}^{\text{ch-ctxt}}(\mathcal{A})$ .

TRANSITION TO LIMITED MEMORY GAME. Next we want to transition to a version of  $G_1$  that stores a bounded amount of local state. Consider the games  $G_2$  and  $G_3$  shown in Fig. 9. The tables  $\mathbf{M}_2$  and  $\mathbf{C}_2$  track the messages and ciphertexts as in the real game. Because of this  $\Pr[G_1] = \Pr[G_2]$ .

In the transition to  $G_3$  we are going to stop using these tables and instead solely rely on the tables  $\mathbf{M}$  and  $\mathbf{C}$ . With these tables, if the total number of bits of ciphertexts that would be stored in  $\mathbf{C}$  exceeds  $\delta$  then we permanently stop adding elements to these tables – we assume that the adversary will cause **sync** to be set to **false** at some point earlier in the game. Note that up until this point the tables  $(\mathbf{M}_2, \mathbf{C}_2)$  and  $(\mathbf{M}, \mathbf{C})$  are used identically. The two games only differ in the boxed code in DEC which returns  $M_2$  if the adversary has queried a ciphertext stored in  $\mathbf{C}_2$  that was not stored in  $\mathbf{C}$ . Hence, these games are identical-until-bad so the Fundamental Lemma of Game Playing [3] gives,

$$\Pr[G_3] - \Pr[G_2] \leq \Pr[G_3 \text{ sets bad}].$$

We want to apply Lemma 1 to bound the probability that **bad** is set. To do so we need to be able to treat the ciphertexts as random strings. Thus we defer the analysis of the probability that it occurs until after applying INDR security.

INDR TRANSITION. Now consider the game  $G_4$ . It is identical to  $G_3$  except that the ciphertexts returned by ENC are chosen at random instead of using CH. We can transition to this game using a reduction to INDR security. It is important here that our reduction adversary will not need to use too much memory because of the way that we have limited the memory needed for  $G_3$ .

Consider the adversaries  $\mathcal{A}'_\delta$  and  $\mathcal{A}''_\delta$  shown in Fig. 10. Highlighted code is only included in the latter adversary.

Adversary  $\mathcal{A}'_\delta$  uses its ENC oracle to present  $\mathcal{A}$  with a view identical to  $G_3$  if  $b = 1$  and identical to  $G_4$  if  $b = 0$ . Note here that the tables  $(\mathbf{M}_2, \mathbf{C}_2)$  do not effect the view of  $\mathcal{A}$  in either of these game, allowing  $\mathcal{A}'_\delta$  not to have to store them. We have that  $\Pr[G_{\text{CH},1}^{\text{ch-indr}}(\mathcal{A}'_\delta)] = \Pr[G_3]$  and  $\Pr[G_{\text{CH},0}^{\text{ch-indr}}(\mathcal{A}'_\delta)] = \Pr[G_4]$ . In other words,  $\text{Adv}_{\text{CH}}^{\text{ch-indr}}(\mathcal{A}'_\delta) = \Pr[G_3] - \Pr[G_4]$ .

Adversary  $\mathcal{A}''_\delta$  instead uses its INDR oracle to simulate the view of  $\mathcal{A}$ , but returns 1 if the flag **bad** would have been set. Because this can only be set by the first ciphertext not stored in  $\mathbf{C}$  we only need to be able to simulate the games up until that point. So we store this extra ciphertext and put an  $*$  in  $\mathbf{C}$  so that in DEC we know when we have reached the relevant point. We have that  $\Pr[G_{\text{CH},1}^{\text{ch-indr}}(\mathcal{A}''_\delta)] = \Pr[G_3 \text{ sets bad}]$  and  $\Pr[G_{\text{CH},0}^{\text{ch-indr}}(\mathcal{A}''_\delta)] = \Pr[G_4 \text{ sets bad}]$ . In other words,  $\Pr[G_3 \text{ sets bad}] \leq \text{Adv}_{\text{CH}}^{\text{ch-indr}}(\mathcal{A}''_\delta) + \Pr[G_4 \text{ sets bad}]$ .

<p>Adversary <math>\mathcal{A}'_3^{\text{ENC}}</math></p> <p>flag <math>\leftarrow</math> true</p> <p>sync <math>\leftarrow</math> true</p> <p><math>b' \leftarrow_{\\$} \mathcal{A}^{\text{SIMENC, SIMDEC}}</math></p> <p>Return <math>b' = 1</math></p>	<p>Oracle SIMENC(<math>M</math>)</p> <p>flag <math>\leftarrow</math> true</p> <p>sync <math>\leftarrow</math> true</p> <p><math>C \leftarrow_{\\$} \text{ENC}(M)</math></p> <p>If flag then</p> <p>  If <math>\ C\  +  C  &lt; \delta</math> then</p> <p>    M.add(<math>M</math>); C.add(<math>C</math>)</p> <p>  Else</p> <p>    C.add(*); <math>C^* \leftarrow C</math></p> <p>    flag <math>\leftarrow</math> false</p> <p>Return <math>C</math></p>	<p>Oracle SIMDEC(<math>C</math>)</p> <p><math>M' \leftarrow \text{M.dq}()</math>; <math>C' \leftarrow \text{C.dq}()</math></p> <p>If sync then</p> <p>  If <math>C = C'</math> then</p> <p>    Return <math>M'</math></p> <p>  Elif <math>C' = *</math> and <math>C = C^*</math> then</p> <p>    <b>abort</b>(1)</p> <p>  sync <math>\leftarrow</math> false</p> <p>Return <math>\perp</math></p>
<p>Adversary <math>\mathcal{A}''_3^{\text{ENC}}</math></p> <p>flag <math>\leftarrow</math> true</p> <p>sync <math>\leftarrow</math> true</p> <p><math>\mathcal{A}^{\text{SIMENC, SIMDEC}}</math></p> <p>Return 0</p>		

Fig. 10. INDR adversaries for proof of Theorem 2. Highlighting indicates code that is only used by adversary  $\mathcal{A}'_3$ .

FINAL TRANSITION. The final transition is from  $G_4$  to  $G_5$ . These two games are identical-until-bad as can be seen in DEC. Because of this we have that

$$\Pr[G_4] - \Pr[G_5] \leq \Pr[G_4 \text{ sets bad}].$$

Using all of  $M_2$  and  $C_2$  instead of just  $M$  and  $C$  makes  $G_5$  identical to  $G_{\text{CH},0}^{\text{ch-ae}}$ .

BOUNDING PROBABILITY OF bad. We conclude by bounding the probability  $G_4$  sets bad via a reduction to our information theoretic analysis. Consider the  $S$ -bounded  $(\mathcal{A}_1, \mathcal{A}_2)$  that behaves as follows. First,  $\mathcal{A}_1$  internally simulates the view of  $\mathcal{A}$  in  $G_4$  using the coins for  $\mathcal{A}$  which maximize the probability of bad and using the bits of  $R$  as the ciphertext bits returned by encryption. If  $\mathcal{A}$  causes flag to be set to false,  $\mathcal{A}_1$  will halt and output the current state of  $\mathcal{A}$  as  $\sigma$  with  $i$  chosen so the next  $\delta$  bits of  $C$  and  $c$  are the values of  $R$  for  $\mathcal{A}_2$  to guess.

Then  $\mathcal{A}_2$  will resume executing  $\mathcal{A}$  using  $\sigma$ . When  $\mathcal{A}$  makes encryption queries it will just make up its own responses. When  $\mathcal{A}$  makes a decryption query for a ciphertext  $C$  then  $\mathcal{A}_2$  will concatenate it into its guess  $r$ . It just assumes this was the correct next ciphertext that should have been stored in  $C$  (otherwise  $\mathcal{A}$  would fail in setting bad). To determine which  $M$  to return for this query,  $\mathcal{A}_2$  re-runs  $\mathcal{A}$  from the beginning using the same coins  $\mathcal{A}_1$  used. It uses its given prefix of  $R$  and the current value of  $r$  to respond to encryption queries until it reaches the encryption query corresponding to the current decryption query. Whatever message  $\mathcal{A}$  queried for this encryption query is then returned for the decryption query. Once  $r$  is  $\delta$  bits long,  $\mathcal{A}_2$  outputs that as its guess.

We can see that when bad would be set in  $G_4$ , the view of  $\mathcal{A}$  is perfectly simulated up until that point and  $\mathcal{A}_2$  will guess  $r$  correctly. This gives us  $\Pr[G_4 \text{ sets bad}] \leq \text{Adv}_{q,x,\delta}^{\text{it}}(\mathcal{A}_1, \mathcal{A}_2)$  as desired.

Combining all the bounds we have shown completes the proof.  $\square$

### 4.3 AE Security of a TLS 1.3-like Channel

We have shown that the AE security of a channel can be reduced to its constituent INDR and CTXT security in a way that preserves memory complexity. This is, of course, only meaningful if we have channels for which we can give provable time-memory tradeoffs for their INDR and CTXT security. Using the ideas of Jaeger and Tessaro [11] it is easy to give such examples for INDR security.

Using the ideas from proof of Thm. 2 we will prove the security of a channel based on GCM (or more generally CAU). The resulting channel can be viewed as a (simplified) version of the channel obtained by using GCM in TLS 1.3.

THE CONSTRUCTION. The construction we consider is a straightforward construction of a channel from a nonce-based encryption scheme NE by using a counter for the nonce. The INDR security of this channel follows easily from the nonce-respecting INDR security of NE. Proving integrity of the channel from the

$\overline{\text{NCH[NE].Sg}}$ $K \leftarrow_s \text{NE.Kg}$ $N \leftarrow_s \text{NE.N}$ Return $((K, N), (K, N))$	$\overline{\text{NCH[NE].S}((K, N), M)}$ $N \leftarrow N + 1$ $C \leftarrow \text{NE.E}(K, N, M)$ Return $((K, N), C)$	$\overline{\text{NCH[NE].R}((K, N), C)}$ If $N = \perp$ then Return $((\perp, \perp), \perp)$ $N \leftarrow N + 1$ $M \leftarrow \text{NE.D}(K, N, C)$ If $M = \perp$ then Return $((\perp, \perp), \perp)$ Return $((K, N), M)$
--	---	---

**Fig. 11.** Algorithms of channel  $\text{NCH[NE]}$  constructed from encryption scheme  $\text{NE}$ .

integrity of  $\text{NE}$  is possible, but of limited applicability since we do not have examples of encryption schemes with proven time-memory tradeoffs for integrity. We will instead only show integrity for the specific case that  $\text{NE} = \text{CAU}$ .

The channel  $\text{NCH[NE]}$  is parameterized by an encryption scheme  $\text{NE}$ . It has  $\text{NCH[NE].M} = \text{NE.M}$ . We assume that  $\text{NE.N}$  can be interpreted as a cyclic group written using additive notation. Its algorithms are shown in Fig. 11. State generation sets the state of both parties equal to a shared random key and nonce. Encryption increments the nonce and uses  $\text{NE}$  to encrypt the message with the current nonce. Decryption increments the nonce and uses  $\text{NE}$  to decrypt the ciphertext with the current nonce. If the ciphertext is rejected ( $M = \perp$ ), the receiver will replace its state with  $\perp$ 's. Henceforth it will reject all ciphertexts it receives (via the first line which checks if  $N = \perp$  already holds).

**INDR SECURITY.** The INDR security of  $\text{NCH[NE]}$  follows easily from nonce-respecting INDR security of  $\text{NE}$ . This is captured by the following theorem.

**Theorem 3.** *Let  $\mathcal{A}$  be an adversary against the INDR security of  $\text{NCH[NE]}$  that makes less than  $|\text{NE.N}|$  oracle queries. Then we can construct  $\mathcal{B}$  such that*

$$\text{Adv}_{\text{NCH[NE]}}^{\text{ch-indr}}(\mathcal{A}) \leq \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{B}).$$

*Adversary  $\mathcal{B}$  has complexity comparable to that of  $\mathcal{A}$  and is nonce-respecting.*

*Proof (Sketch).* Adversary  $\mathcal{B}$  picks  $N$  at random and then starts executing  $\mathcal{A}$ . Whenever  $\mathcal{A}$  makes a  $\text{ENC}(M)$  query,  $\mathcal{B}$  increments  $N$ , queries  $C \leftarrow \text{ENC}(N, M)$ , and returns  $C$  to  $\mathcal{A}$ . Adversary  $\mathcal{B}$  outputs whatever  $\mathcal{A}$  does. Verifying the claims made about this adversary is straightforward.  $\square$

**CTXT SECURITY.** For CTXT security we need to focus our attention on the particular construction of  $\text{NCH[NE]}$  obtained when using the encryption scheme  $\text{NE} = \text{CAU[E, H]}$  for some function families  $\text{E}$  and  $\text{H}$ .

In our proof, we will take advantage of the fact that the adversary can essentially only make a single forgery attempt. If it fails at this attempt, then the state of the decryption algorithm can be erased and it will henceforth always return  $\perp$ . Because  $\text{CAU}$  uses a Carter-Wegman style MAC we have to first use the PRF security of  $\text{E}$  to hide the values of  $H_L(A, C)$  used in encryption queries. To get our desired state-aware results we need to make sure that our PRF reduction does not use much more memory than the original adversary. This creates an issue similar to what we saw in Section 4 where it can be difficult to simulate the values returned by  $\text{DEC}$ . This issue is resolved by adjusting the proof technique used to establish Thm. 2 where we exploit the fact that ciphertexts look random to assume that  $\mathcal{A}$  cannot remember too many ciphertexts.

**Theorem 4.** *Let  $\text{NE} = \text{CAU[E, H]}$  for some  $\text{E}$  and  $\text{H}$ . Let  $\mathcal{A}$  be a nonce-respecting adversary against the CTXT security of  $\text{NCH[NE]}$  with memory complexity  $S$  that makes at most  $q \leq 2^{\text{CAU.nl}} - 1$  encryption queries, each of which returns a ciphertext of length at most  $x$ . Then for any  $\delta \in \mathbb{N}$  we can construct an adversary  $\mathcal{A}_{\text{prf}}$  such that*

$$\text{Adv}_{\text{NCH[NE]}}^{\text{ch-ctxt}}(\mathcal{A}) \leq 2 \cdot \text{Adv}_{\text{E}}^{\text{prf}}(\mathcal{A}_{\text{prf}}) + \text{Adv}_{\text{H}}^{\text{axu}}(\mathcal{X}) + q \cdot x \cdot 2^S / 2^\delta.$$

*Adversary  $\mathcal{A}_{\text{prf}}$  has running time approximately that of  $\mathcal{A}$  and uses about  $S + 2\delta$  bits of state. It makes at most  $q(x/n + 2) + 1$  non-repeating queries to its oracle.*



Reduction $\mathcal{R}[\mathcal{A}]^{\circ}$	Oracle $\text{REnc}(N, M)$
$\sigma \leftarrow \mathcal{R}^{\circ}.\text{Init}$	$(\sigma, \text{rf}, C) \leftarrow \mathcal{R}^{\circ}.\text{SimEnc}(\sigma, N, M)$
$i \leftarrow 0$	If $\text{rf}$ then goto NEXT
While $i \leq \mathcal{R}.\text{rew}$ do	Return $C$
$b \leftarrow \mathcal{A}^{\text{REnc, RDec}}$	Oracle $\text{RDec}(N, C)$
$\sigma \leftarrow \mathcal{R}^{\circ}.\text{Upd}(\sigma, b)$	$(\sigma, \text{rf}, M) \leftarrow \mathcal{R}^{\circ}.\text{SimDec}(\sigma, N, C)$
NEXT: $i \leftarrow i + 1$	If $\text{rf}$ then goto NEXT
Return $\mathcal{R}^{\circ}.\text{Fin}(\sigma)$	Return $M$

**Fig. 12.** Syntax of a black-box reductions  $\mathcal{R}$  running AE-1 adversary  $\mathcal{A}$ . We represent the oracles  $\mathcal{R}$  has access to collectively as  $\mathcal{O}$ .

The proof is given in Appendix F. As with Thm. 1, the PRF adversary we give never repeats queries so we can apply the switching lemma to obtain a bound using PRP security of E. Here it is important that the memory of  $\mathcal{A}_{\text{prf}}$  is not much more than that of  $\mathcal{A}$ . Assuming  $|\mathcal{A}| < p$ , setting  $\delta \approx S + n$ , and assuming  $S > n$  we can combine all of our theorems so far to obtain a bound of

$$\text{Adv}_{\text{NCH}[\text{NE}]}^{\text{ch-ae}}(\mathcal{A}) \leq 4 \cdot \text{Adv}_{\text{E}}^{\text{prf}}(\mathcal{B}) + O\left(\frac{Spq \log(pq) + c(p+q)}{2^n}\right)$$

for a  $\mathcal{B}$  with comparable efficiency to  $\mathcal{A}$  and assuming H is  $c$ -AXU.

## 5 Negative Results for Memory-tight AE Reductions

In this section we give impossibility results for giving a memory-tight reduction (for a natural restricted class of black-box reductions) from AE-1 security to nonce-respecting INDR and CTXT-1 security. This establishes that our restriction to the channel setting for Section 4 was necessary for our positive results.

**BLACK-BOX REDUCTIONS.** A reduction  $\mathcal{R}$  maps an adversary  $\mathcal{A}$  to an adversary  $\mathcal{R}[\mathcal{A}]$ . We consider reductions that run an AE-1 adversary  $\mathcal{A}$  in a black-box manner as shown in Fig. 12. It starts with initial state  $\sigma$  output by  $\mathcal{R}.\text{Init}$ . The parameter  $\mathcal{R}.\text{rew}$  determines how many times  $\mathcal{R}$  will perform a full rewind of  $\mathcal{A}$ . Then it runs  $\mathcal{A}$  while simulating its encryption and decryption oracles. For every encryption query,  $\mathcal{R}$  runs  $\mathcal{R}.\text{SimEnc}$  with the query and its state as input to produce the updated state, a flag  $\text{rf}$ , and a ciphertext. If the flag  $\text{rf}$  is **true**, then  $\mathcal{R}$  starts running  $\mathcal{A}$  from the beginning again. Otherwise, it answers with the query answer  $\mathcal{R}.\text{SimEnc}$  returned. Decryption queries are handled analogously. If  $\mathcal{R}$  did not rewind  $\mathcal{A}$  before  $\mathcal{A}$  finished its execution, then it runs  $\mathcal{R}.\text{Upd}$  on  $\mathcal{A}$ 's output to updates its state and starts running  $\mathcal{A}$  from the beginning if has not already rewinded  $\mathcal{R}.\text{rew}$  times. Finally,  $\mathcal{R}$  outputs whatever  $\mathcal{R}.\text{Fin}(\sigma)$  returns. The following definition captures some restrictions we will place on reductions.

**Definition 1.** *Let  $\mathcal{R}$  be a reduction using the syntax from Fig. 12. It is full-rewinding if  $\mathcal{R}.\text{rew} > 0$  or straightline if  $\mathcal{R}.\text{rew} = 0$ . It is nonce-respecting if  $\mathcal{R}[\mathcal{A}]$  is nonce-respecting when  $\mathcal{A}$  is nonce-respecting. It is faithful if  $\mathcal{R}[\mathcal{A}]$  answers encryption queries of  $\mathcal{A}$  consistent with its own encryption oracle, i.e.,  $\mathcal{R}$  responds with  $C$  on an encryption query made on  $(N, M)$ , only if it previously queried its own encryption oracle with  $(N, M)$  and received  $C$  as the answer.*

**ADDITIONAL NOTATION.** We fix an understood nonce-based encryption scheme NE for which we assume that  $\{0, 1\}^{\text{ml}} \subseteq \text{NE.M}$ . We also assume  $\mathbb{N} \times \mathbb{N} \subseteq \text{NE.N}$  and we use  $\mathbb{N} = \text{NE.N}$  as shorthand. We assume that  $[\text{NE.Kg}] = \{0, 1\}^{\text{kl}}$ . We let  $\mathbb{C} = \{0, 1\}^{\text{NE.cl(ml)}}$ . We also introduce some new notation for the complexity of an algorithm  $\mathcal{A}$ . First,  $\text{Mem}(\mathcal{A})$  is defined as the number of bits of memory that  $\mathcal{A}$  uses. The total number of queries to its oracles is  $\text{Query}(\mathcal{A})$ , and the number of computation steps  $\text{Time}(\mathcal{A})$ . For a reduction  $\mathcal{R}$  we use  $\text{Mem}(\mathcal{R})$  to denote the number of bits of memory that  $\mathcal{R}$  uses in addition any memory of the adversary it runs.

<p>Game <math>G_{u,m,P}^{\text{bf-it-chl}}(\mathcal{B}_1, \mathcal{B}_2)</math></p> <p><math>(I, Q) \leftarrow \mathcal{B}_1 //  I  =  Q  = P</math></p> <p>For <math>j \in [u]</math> do</p> <p style="padding-left: 2em;"><math>M_j \leftarrow \{0, 1\}^m</math></p> <p>For <math>j \in [P]</math> do</p> <p style="padding-left: 2em;"><math>M_{I[j]} \leftarrow Q[j]</math></p> <p><math>j^* \leftarrow [u]</math></p> <p><math>M \leftarrow \mathcal{B}_2(I, Q, j^*)</math></p> <p>If <math>M_{j^*} \neq M</math> then return <b>false</b></p> <p>Return <b>true</b></p>	<p>Adversary <math>\mathcal{D}'_1(\perp, M_1, \dots, M_u)</math></p> <p>While(<b>true</b>)</p> <p style="padding-left: 2em;"><math>\text{win} \leftarrow 1; \phi \leftarrow \perp</math></p> <p style="padding-left: 2em;">For <math>k \in [i-1]</math> do</p> <p style="padding-left: 4em;">For <math>j \in [u]</math> do</p> <p style="padding-left: 6em;"><math>M'_j \leftarrow \{0, 1\}^m</math></p> <p style="padding-left: 4em;"><math>\phi \leftarrow \mathcal{D}_1^*(\phi, M'_1, \dots, M'_u)</math></p> <p style="padding-left: 4em;"><math>j^* \leftarrow [u]</math></p> <p style="padding-left: 4em;"><math>(\phi, M') \leftarrow \mathcal{D}_2^*(\phi, j^*)</math></p> <p style="padding-left: 4em;">If <math>M' \neq M_{j^*}</math> then <math>\text{win} \leftarrow 0</math></p> <p style="padding-left: 2em;">If <math>\text{win} = 1</math> then break</p> <p><math>\phi \leftarrow \mathcal{D}_1^*(\phi, M_1, \dots, M_u)</math></p> <p>Return <math>\phi</math></p>
---	---

**Fig. 14.** Game  $G_{u,m,P}^{\text{bf-it-chl}}$  and adversary  $\mathcal{D}'_1$  used in the proof of Lemma 2.

INFORMATION THEORETIC LEMMA. We give a lemma that will be a useful sub-component of our proofs. It pertains to game  $G_{u,m}^{\text{it-chl-}r}$  in Fig. 13. It is an  $r$ -round game, played by a two-stage adversary  $(\mathcal{D}_1, \mathcal{D}_2)$ . In each round,  $\mathcal{D}_1$  gets state  $\sigma$  from the prior round, along with  $u$  random strings  $M_1, \dots, M_u$  each of length  $m$ . Adversary  $\mathcal{D}_1$  outputs state  $\sigma$  which is input to  $\mathcal{D}_2$  along with a randomly sampled index  $j^*$  from  $[u]$ . Then  $\mathcal{D}_2$  outputs a string  $M$  and state  $\sigma$  that is passed to  $\mathcal{D}_1$  in the next round. If  $M = M_{j^*}$ , we say that  $(\mathcal{D}_1, \mathcal{D}_2)$  has answered the challenge of this round correctly. If  $(\mathcal{D}_1, \mathcal{D}_2)$  answers all the  $r$  challenges correctly, the game returns **true**. Otherwise it returns **false**. We define  $\text{Adv}_{u,m}^{\text{it-chl-}r}(\mathcal{D}_1, \mathcal{D}_2) = \Pr[G_{u,m}^{\text{it-chl-}r}(\mathcal{D}_1, \mathcal{D}_2)]$ . Adversary  $(\mathcal{D}_1, \mathcal{D}_2)$  is  $S$ -bounded if the state output by  $\mathcal{D}_1$  is at most  $S$  bits long. We can prove the following.

<p>Game <math>G_{u,m}^{\text{it-chl-}r}(\mathcal{D}_1, \mathcal{D}_2)</math></p> <p><math>\sigma \leftarrow \perp</math></p> <p>For <math>i \in [r]</math> do</p> <p style="padding-left: 2em;">For <math>j \in [u]</math> do</p> <p style="padding-left: 4em;"><math>M_j \leftarrow \{0, 1\}^m</math></p> <p style="padding-left: 2em;"><math>\sigma \leftarrow \mathcal{D}_1(\sigma, M_1, \dots, M_u)</math></p> <p style="padding-left: 2em;"><math>j^* \leftarrow [u]</math></p> <p style="padding-left: 2em;"><math>(\sigma, M) \leftarrow \mathcal{D}_2(\sigma, j^*)</math></p> <p style="padding-left: 2em;">If <math>M_{j^*} \neq M</math> then return <b>false</b></p> <p>Return <b>true</b></p>
---

**Fig. 13.** Information theoretic game played by adversary  $(\mathcal{D}_1, \mathcal{D}_2)$ .

**Lemma 2.** *If  $(\mathcal{D}_1, \mathcal{D}_2)$  is  $S$ -bounded, then*

$$\text{Adv}_{u,m}^{\text{it-chl-}r}(\mathcal{D}_1, \mathcal{D}_2) \leq \left( \frac{2(S+m)}{u} + \frac{3}{2^m} \right)^r.$$

The proof reduces to the  $r = 1$  case and analyzes this case using techniques from the AI-ROM setting [5].

*Proof.* We split this lemma into the following two claims. The first claim bounds advantage of any  $S$ -bounded adversary playing the one round version of the game ( $r = 1$ ). The second step bounds the advantage of an adversary playing the full  $r$  rounds in using an adversary playing only one round.

*Claim.* If  $(\mathcal{D}_1, \mathcal{D}_2)$  is  $S$ -bounded, then  $\text{Adv}_{u,m}^{\text{it-chl-}1}(\mathcal{D}_1, \mathcal{D}_2) \leq \frac{2S+2m}{u} + \frac{3}{2^m}$

*Claim.* If  $\text{Adv}_{u,m}^{\text{it-chl-}1}(\mathcal{D}_1, \mathcal{D}_2) \leq \varepsilon$  for all  $S$ -bounded  $(\mathcal{D}_1, \mathcal{D}_2)$ , then  $\text{Adv}_{u,m}^{\text{it-chl-}r}(\mathcal{D}_1^*, \mathcal{D}_2^*) \leq \varepsilon^r$  for all  $r$  and all  $S$ -bounded  $(\mathcal{D}_1^*, \mathcal{D}_2^*)$ .

FIRST CLAIM. In order to show the upper bound on  $\text{Adv}_{u,m}^{\text{it-chl-}1}(\mathcal{D}_1, \mathcal{D}_2)$ , we first define the game  $G_{u,m,P}^{\text{bf-it-chl}}$  in Fig. 14. It is the ‘‘Bit-Fixed Random Oracle’’ [5] version of  $G_{u,m}^{\text{it-chl-}1}$ . (The sequence of  $u$  random strings that are input to the adversary can be viewed as a random oracle  $[u] \rightarrow \{0, 1\}^m$ .) This game is played by a two-stage adversary  $(\mathcal{B}_1, \mathcal{B}_2)$ . The first stage  $\mathcal{B}_1$  outputs two lists  $I$  and  $Q$ , each of length  $P$ . All the entries

of  $I$  are in  $[u]$  and all the entries of  $Q$  are in  $\{0, 1\}^m$ . Then  $u$  bit-strings  $M_1, \dots, M_u$  are sampled uniformly at random from  $\{0, 1\}^m$ . For each  $j \in [P]$ , the bit-string  $M_i$  is overwritten by  $Q[j]$ , where  $i = I[j]$ . The second stage adversary  $\mathcal{B}_2$  is given the lists and a randomly sampled  $j^*$  as input. It returns  $M \in \{0, 1\}^m$ . The game returns **true** if  $M = M_{j^*}$  and **false** otherwise. We define  $\text{Adv}_{u,m,P}^{\text{bf-it-chl}}(\mathcal{B}_1, \mathcal{B}_2) = \Pr[\text{G}_{u,m,P}^{\text{bf-it-chl}}(\mathcal{B}_1, \mathcal{B}_2)]$ .

We claim that  $\text{Adv}_{u,m,P}^{\text{bf-it-chl}}(\mathcal{B}_1, \mathcal{B}_2) \leq \frac{P}{u} + \frac{1}{2^m}$ . The first term captures the probability that  $j^*$  is in  $I$ . The second term captures the probability that  $\mathcal{B}_2$  guesses the random string  $M_{j^*} \in \{0, 1\}^m$  when  $j^*$  is not in  $I$ . Since  $(\mathcal{D}_1, \mathcal{D}_2)$  is  $S$ -bounded, Theorem 6 of [5] gives that for any  $P \in \mathbb{N}$  and every  $\gamma > 0$ , if  $P \geq S + \log \gamma^{-1}$  then there exists an adversary  $(\mathcal{B}_1, \mathcal{B}_2)$  such that

$$\text{Adv}_{u,m}^{\text{it-chl-1}}(\mathcal{D}_1, \mathcal{D}_2) \leq 2 \cdot \text{Adv}_{u,m,P}^{\text{bf-it-chl}}(\mathcal{B}_1, \mathcal{B}_2) + \gamma.$$

Setting  $\gamma = \frac{1}{2^m}$  and  $P = S + m$ , gives the claimed bound.

**SECOND CLAIM.** Let us assume, towards contradiction, that there exists  $S$ -bounded  $(\mathcal{D}_1^*, \mathcal{D}_2^*)$  such that  $\text{Adv}_{u,m}^{\text{it-chl-r}}(\mathcal{D}_1^*, \mathcal{D}_2^*) > \varepsilon^r \neq 0$ . Then we can construct an adversary  $(\mathcal{D}'_1, \mathcal{D}'_2)$  such that  $\text{Adv}_{u,m}^{\text{it-chl-1}}(\mathcal{D}'_1, \mathcal{D}'_2) > \varepsilon$ . We give  $\mathcal{D}'_1$  in Fig. 14 and let  $\mathcal{D}'_2 = \mathcal{D}_2^*$ . We say  $(\mathcal{D}'_1, \mathcal{D}'_2)$  wins round  $j$  if  $\text{G}_{u,m}^{\text{it-chl-r}}$  did not return **false** before the  $(j+1)^{\text{st}}$  iteration. It follows that,

$$\text{Adv}_{u,m}^{\text{it-chl-r}}(\mathcal{D}'_1, \mathcal{D}'_2) = \prod_{j=1}^r \Pr[(\mathcal{D}'_1, \mathcal{D}'_2) \text{ wins round } j \mid (\mathcal{D}'_1, \mathcal{D}'_2) \text{ wins round } j-1].$$

Since  $\text{Adv}_{u,m}^{\text{it-chl-r}}(\mathcal{D}'_1, \mathcal{D}'_2) > \varepsilon^r$ , there must exist some  $i \in [r]$ , such that,

$$\Pr[(\mathcal{D}'_1, \mathcal{D}'_2) \text{ wins round } i \mid (\mathcal{D}'_1, \mathcal{D}'_2) \text{ wins round } i-1] > \varepsilon.$$

Adversary  $\mathcal{D}'_1$  keeps running  $(\mathcal{D}_1^*, \mathcal{D}_2^*)$  until it wins round  $i-1$  on randomly generated  $M_{i,j}$ 's. Once this occurs  $\mathcal{D}'_1$  runs  $\mathcal{D}_1^*$  on its own inputs and returns the state returned by  $\mathcal{D}_1^*$ . Adversary  $\mathcal{D}'_2$  runs  $\mathcal{D}_2^*$  on its own inputs and returns whatever  $\mathcal{D}_2^*$  returns. Note that  $(\mathcal{D}'_1, \mathcal{D}'_2)$  is  $S$ -bounded because  $(\mathcal{D}_1^*, \mathcal{D}_2^*)$  is. Moreover,

$$\Pr[\text{G}_{u,m}^{\text{it-chl-1}}(\mathcal{D}'_1, \mathcal{D}'_2)] = \Pr[(\mathcal{D}_1^*, \mathcal{D}_2^*) \text{ win round } i \mid (\mathcal{D}_1^*, \mathcal{D}_2^*) \text{ win round } i-1].$$

Thus  $\text{Adv}_{u,m}^{\text{it-chl-r}}(\mathcal{D}'_1, \mathcal{D}'_2) > \varepsilon$ , which contradicts our assumption.  $\square$

## 5.1 Memory Lower Bound for Straightline Reductions

Our first theorem shows that it is not possible to give memory-tight, straightline reductions proving the AE-1 security of an encryption scheme from its INDR and CTXT-1 security. (As the theorem statement is somewhat complicated, we will describe how to interpret it below.)

**Theorem 5 (Impossibility for straightline reductions).** *Let NE be a nonce-based encryption scheme. Fix  $u, r \in \mathbb{N}$  and define the nonce-respecting adversary  $\mathcal{A}$  as shown in Fig. 16. Let  $\mathcal{R}$  be a straightline, nonce-respecting, faithful black-box reduction from AE-1 to nonce-respecting INDR with  $\text{Mem}(\mathcal{R}) = S$ . Let  $\mathcal{R}'$  be a straightline, nonce-respecting, faithful reduction from AE-1 to CTXT-1. Then, we can construct adversaries  $\mathcal{C}$  and  $\mathcal{W}$  such that,*

$$(i) \text{ Adv}_{\text{NE}}^{\text{ae-1}}(\mathcal{A}) \geq 1 - \frac{2^{\text{kl}}}{2^{2\text{NE.cl(ml)}}},$$

$$(ii) \text{ Adv}_{\text{NE}}^{\text{indr}}(\mathcal{R}[\mathcal{A}]) \leq 2 \cdot \left( \frac{2(S + \log r + \text{kl}) + 2\text{ml}}{u} + \frac{3}{2^{\text{ml}}} \right)^r + 4 \cdot \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}),$$

$$(iii) \text{ Adv}_{\text{NE}}^{\text{ctxt-1}}(\mathcal{R}'[\mathcal{A}]) \leq \text{Adv}_{\text{NE}}^{\text{ctxt-1}}(\mathcal{R}'[\mathcal{W}]).$$

Moreover,  $\mathcal{A}$  satisfies  $\text{Query}(\mathcal{A}) = (u+1) \cdot r + 2$  and  $\text{Mem}(\mathcal{A}) \leq 2\text{kl} + 2\text{ml} + 2\text{NE.cl(ml)} + 2 \log |\mathbb{N}| + \log u \cdot r$ . Also  $\mathcal{C}$  and  $\mathcal{W}$  satisfy  $\text{Query}(\mathcal{C}) < \text{Query}(\mathcal{R}) + \text{Query}(\mathcal{A})$ ,  $\text{Time}(\mathcal{C}) \in O(\text{Query}(\mathcal{A}) + \text{Time}(\mathcal{R}))$ , and  $\text{Query}(\mathcal{W}) = \text{Query}(\mathcal{A})$  and  $\text{Time}(\mathcal{W}) \in O(\text{Query}(\mathcal{A}))$ .

<p>Games <math>G_b^0, G^1</math></p> <p><math>b \leftarrow 1</math></p> <p><math>K^* \leftarrow \text{NE.Kg}(); \sigma \leftarrow \mathcal{R}^{\text{ENC}_b}.\text{Init}</math></p> <p>For <math>i \in [r]</math> do</p> <p>  <math>j^* \leftarrow [u]</math></p> <p>  For <math>j \in [u]</math> do</p> <p>    <math>M_j \leftarrow \{0, 1\}^{\text{ml}}; N_j \leftarrow (i, j)</math></p> <p>    <math>(\sigma, \cdot, C_j) \leftarrow \mathcal{R}^{\text{ENC}_b}.\text{SIMENC}(\sigma, N_j, M_j)</math></p> <p>    <math>(\sigma, \cdot, M) \leftarrow \mathcal{R}^{\text{ENC}_b}.\text{SIMDEC}(\sigma, N_{j^*}, C_{j^*})</math></p> <p>    If <math>M \neq M_{j^*}</math> then return <b>false</b></p> <p>  Return <b>true</b></p>	<p>Oracle <math>\text{ENC}_b(N, M)</math></p> <p><math>C_0 \leftarrow \mathcal{C}</math></p> <p><math>C_1 \leftarrow \text{NE.E}(K^*, N, M)</math></p> <p><math>C_1 \leftarrow \text{NE.E}(K^*, N, 0^{\text{ml}})</math></p> <p>Return <math>C_b</math></p>
--	---

**Fig. 15.** Games  $G^1$  and  $G_b^0$  for  $b \in \{0, 1\}$ . Highlighted code is only included in  $G^1$ .

<p>Adversaries <math>\mathcal{A}^{\text{ENC,DEC}}, \boxed{\mathcal{S}^{\text{ENC,DEC}}}, \mathcal{W}^{\text{ENC,DEC}}</math></p> <p>For <math>i \in [r]</math> do</p> <p>  <math>j^* \leftarrow [u]</math></p> <p>  For <math>j \in [u]</math> do</p> <p>    <math>M_j \leftarrow \{0, 1\}^{\text{ml}}</math></p> <p>    <math>N_j \leftarrow (i, j)</math></p> <p>    <math>C_j \leftarrow \text{ENC}(N_j, M_j)</math></p> <p>  <math>M \leftarrow \text{DEC}(N_{j^*}, C_{j^*})</math></p> <p>  If <math>M \neq M_{j^*}</math> then return 1</p> <p>  <b>bad</b> <math>\leftarrow \text{true}</math></p> <p>  Return <math>\mathcal{B}^{\text{ENC}}; \boxed{\text{Return } 0};</math> Return <math>\mathcal{E}^{\text{ENC}}</math></p>	<p>Adversaries <math>\mathcal{B}^{\text{ENC}}, \mathcal{E}^{\text{ENC}}</math></p> <p><math>M_1 \leftarrow \{0, 1\}^{\text{ml}}</math></p> <p><math>M_2 \leftarrow \{0, 1\}^{\text{ml}}</math></p> <p><math>N_1 \leftarrow (0, 1); N_2 \leftarrow (0, 2)</math></p> <p><math>C_1 \leftarrow \text{ENC}(N_1, M_1)</math></p> <p><math>C_2 \leftarrow \text{ENC}(N_2, M_2)</math></p> <p>For <math>K \in \{0, 1\}^{\text{kl}}</math> do</p> <p>  <math>\text{eq}_1 \leftarrow (\text{NE.E}(K, N_1, M_1) = C_1)</math></p> <p>  <math>\text{eq}_2 \leftarrow (\text{NE.E}(K, N_2, M_2) = C_2)</math></p> <p>  If <math>\text{eq}_1</math> and <math>\text{eq}_2</math> then return 1</p> <p>Return 0; Return 1</p>
---	---

**Fig. 16.** Adversaries against the AE-1 security of NE. Boxed code is only included in  $\mathcal{S}$ . Highlighted code is only included in  $\mathcal{A}$  and  $\mathcal{B}$ .

To interpret this theorem, assume that the parameters of NE are such that the advantage of  $\mathcal{A}$  is essentially one. Hence, a successful pair of reductions  $\mathcal{R}$  and  $\mathcal{R}'$  would need at least one of  $\mathcal{R}[\mathcal{A}]$  or  $\mathcal{R}'[\mathcal{A}]$  to have high advantage. For memory-tight  $\mathcal{R}$  and  $\mathcal{R}'$  we expect there to be linear functions  $f_1$  and  $f_2$  such that their local computation time and memory usage when interacting with an adversary  $\mathcal{A}$  would be bounded by  $f_1(q_{\mathcal{A}})$  and  $f_2(s_{\mathcal{A}})$  where  $q_{\mathcal{A}} = \text{Query}(\mathcal{A})$  and  $s_{\mathcal{A}} = \text{Mem}(\mathcal{A})$ .

Suppose this was the case. Then we can fix upper bounds for  $\log(u)$  and  $\log(r)$ , determining the memory usage of  $\mathcal{A}$  and hence  $f_1(s_{\mathcal{A}}) = S$ . Now we can pick reasonable  $u$  and  $r$  such that,  $2 \cdot \left( \frac{2(S + \log r + \text{kl}) + 2\text{ml}}{u} + \frac{3}{2^{\text{ml}}} \right)^r$  is very small (by say making the inside of the parenthesis less than  $1/2$  and setting  $r = 128$ ). Then, for one of the reductions to have high advantage, one of  $\mathcal{C}$  or  $\mathcal{R}'[\mathcal{W}]$  would have to have high advantage. But the efficiencies of these are bounded as small functions of the query complexity of  $\mathcal{A}$  (rather than its local runtime) so cannot be too large. But then assuming security of NE prevents any of them from having high advantage.

*Proof.* Consider the adversary  $\mathcal{A}$  in Fig. 16 against AE-1 security of NE. Note that it is nonce-respecting. It has a challenge phase followed by an invocation of  $\mathcal{B}$ . Each iteration of the challenge phase consists of  $\mathcal{A}$  making  $u$  encryption queries with unique nonces and making one decryption query on one of the  $u$  ciphertexts it received as answers chosen uniformly at random with its corresponding nonce. If the answer of the decryption query is not consistent with the prior encryption query,  $\mathcal{A}$  returns 1. There are  $r$  iterations of the challenge phase. If these are all passed,  $\mathcal{A}$  runs adversary  $\mathcal{B}$  (shown on the right) with its ENC oracle and outputs whatever  $\mathcal{B}$  outputs. From the code of  $\mathcal{A}$  we can see that it makes  $r \cdot u + 2$  encryption queries,

$r$  decryption queries, and satisfies

$$\text{Mem}(\mathcal{A}) \leq 2kl + 2ml + 2\text{NE.cl}(ml) + 2 \log |\mathbf{N}| + \log u \cdot r .$$

To prove the theorem we need to separately establish the three advantage claims (and corresponding statements about the efficiency of various algorithms). For the first claim, note that  $\text{Adv}_{\text{NE}}^{\text{ae-1}}(\mathcal{A}) = \text{Adv}_{\text{NE}}^{\text{ae-1}}(\mathcal{B})$  because  $M$  will always equal  $M_{j^*}$  when  $\mathcal{A}$  is playing  $\text{G}_{\text{NE},b}^{\text{ae-1}}$ . The simple analysis giving the needed bound on  $\text{Adv}_{\text{NE}}^{\text{ae-1}}(\mathcal{B})$  is deferred to Appendix B.

For the third claim, consider adversary  $\mathcal{W}$  defined as shown in Fig. 16. It is identical to  $\mathcal{A}$ , except that it calls  $\mathcal{E}$ , which is similar to  $\mathcal{B}$  but always returns 1. Because  $\mathcal{R}'$  is faithful,  $\mathcal{B}$  would never return 0 when run by  $\mathcal{R}'[\mathcal{A}]$  playing  $\text{G}_{\text{NE},b}^{\text{ctxt-1}}$  so  $\text{Adv}_{\text{NE}}^{\text{ctxt-1}}(\mathcal{R}'[\mathcal{A}]) = \text{Adv}_{\text{NE}}^{\text{ctxt-1}}(\mathcal{R}'[\mathcal{W}])$  holds trivially.

We spend the rest of the proof establishing the second claim. Consider the adversary  $\mathcal{S}$  in Fig. 16. It behaves identically to  $\mathcal{A}$  until the flag `bad` is set. Using the Fundamental Lemma of Game Playing [3], we can obtain for  $b \in \{0, 1\}$  that

$$\left| \Pr \left[ \text{G}_{\text{NE},b}^{\text{indr}}(\mathcal{R}[\mathcal{A}]) \right] - \Pr \left[ \text{G}_{\text{NE},b}^{\text{indr}}(\mathcal{R}[\mathcal{S}]) \right] \right| \leq \Pr \left[ \mathcal{R}[\mathcal{A}] \text{ sets bad in } \text{G}_{\text{NE},b}^{\text{indr}} \right] .$$

Consider the games  $\text{G}_b^0$  for  $b \in \{0, 1\}$  in Fig. 15. In it, we assume that  $\mathcal{R}$  always outputs `rf = false` since it is straightline. Note that  $\text{G}_b^0$  simulates the challenge phase of  $\mathcal{A}$  and the game  $\text{G}_b^{\text{indr}}$  to  $\mathcal{R}$  perfectly, so it returns `true` whenever  $\mathcal{R}[\mathcal{A}]$  would set `bad` is set in  $\text{G}_b^{\text{indr}}$ . From this we can show

$$\text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{R}[\mathcal{A}]) \leq \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{R}[\mathcal{S}]) + \Pr \left[ \text{G}_0^0 \right] + \Pr \left[ \text{G}_1^0 \right] . \quad (3)$$

Now consider the game  $\text{G}^1$  defined in the same figure. It is identical to either  $\text{G}_b^0$  except that it answers all encryption queries with the encryption of the message  $0^{\text{ml}}$ . We now state two lemmas which give bounds on both  $\Pr \left[ \text{G}_b^0 \right]$ 's via  $\text{G}^1$ . First, in Lemma 3, we use that the INDR security of NE implies  $\text{G}^1$ 's encryption oracle is indistinguishable from those in either  $\text{G}_b^0$  to transition to  $\text{G}^1$ . Next, in Lemma 4 we give a bound on  $\Pr \left[ \text{G}^1 \right]$  which was obtained by using  $\mathcal{R}$  to construct an adversary for  $\text{G}_{u,\text{ml}}^{\text{it-chl-r}}$  and bounding its advantage with Lemma 2. The proofs of these lemmas are deferred to Appendix C.1 and Appendix D, respectively.

**Lemma 3.** *There exist adversaries  $\mathcal{C}_1$  and  $\mathcal{C}_2$  such that*

$$\begin{aligned} \Pr \left[ \text{G}_1^0 \right] &\leq \Pr \left[ \text{G}_0^0 \right] + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_1), \\ \Pr \left[ \text{G}_0^0 \right] &\leq \Pr \left[ \text{G}^1 \right] + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_2) \end{aligned}$$

where  $\text{G}_b^0$  and  $\text{G}^1$  are defined as in Fig. 15. Moreover  $\text{Query}(\mathcal{C}_1) < \text{Query}(\mathcal{R}) + \text{Query}(\mathcal{A})$  and  $\text{Time}(\mathcal{C}_1) \in O(\text{Query}(\mathcal{A}) + \text{Time}(\mathcal{R}))$ . Adversary  $\mathcal{C}_2$ 's complexity is the same.

**Lemma 4.** *If  $\mathcal{R}$  is a straightline, nonce-respecting, faithful black-box reduction from AE-1 to nonce-respecting INDR with  $\text{Mem}(\mathcal{R}) = S$ . Then,*

$$\Pr \left[ \text{G}^1 \right] \leq \left( \frac{2(S + \log r + kl) + 2ml}{u} + \frac{3}{2^{\text{ml}}} \right)^r$$

where  $\text{G}^1$  is defined as in Fig. 15.

Applying these lemmas to equation (3) gives

$$\text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{R}[\mathcal{A}]) \leq 2 \cdot \left( \frac{2(S + \log r + kl) + 2ml}{u} + \frac{3}{2^{\text{ml}}} \right)^r + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{R}[\mathcal{S}]) + 2 \cdot \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_1) + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_2) .$$

To complete the proof, we combine the three INDR adversaries  $\mathcal{R}[\mathcal{S}]$ ,  $\mathcal{C}_1$ , and  $\mathcal{C}_2$ . Let  $\mathcal{C}$  be the INDR adversary that randomly chooses one of  $\mathcal{R}[\mathcal{S}]$ ,  $\mathcal{C}_1$ , or  $\mathcal{C}_2$  (with probabilities 1/4, 1/2, and 1/4, respectively) then runs the adversary it chose, outputting whatever that adversary does. Simple calculations give

$$4 \cdot \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}) = \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{R}[\mathcal{S}]) + 2 \cdot \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_1) + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_2) .$$

The claimed complexity of  $\mathcal{C}$  follows from that of  $\mathcal{R}[\mathcal{S}]$ ,  $\mathcal{C}_1$ , and  $\mathcal{C}_2$ . □

## 5.2 Memory Lower Bound for Full-Rewinding Reductions

We can extend our result to cover full-rewinding reductions as captured by the following theorem. Its interpretation works similarly to that of Thm. 5.

**Theorem 6 (Impossibility for full-rewinding reductions).** *Let NE be a nonce-based encryption scheme. Fix  $u, r, c \in \mathbb{N}$ . We can construct a nonce-respecting adversary  $\mathcal{A}$  such that for all full-rewinding, nonce-respecting, restricted reductions  $\mathcal{R}$  from AE-1 to nonce-respecting INDR with  $\text{Mem}(\mathcal{R}) = S$  and all full-rewinding, nonce-respecting, restricted reductions  $\mathcal{R}'$  from AE-1 to CTXT-1 there exist adversaries  $\mathcal{C}$  and  $\mathcal{W}$  such that,*

- (i)  $\text{Adv}_{\text{NE}}^{\text{ae-1}}(\mathcal{A}) \geq 1 - \frac{2^{\text{kl}}}{2^{2\text{NE.cl(ml)}}}$ ,
- (ii)  $\text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{R}[\mathcal{A}]) \leq 2 \cdot \left( \frac{2(S + \log r + \text{kl}) + 2\text{ml}}{u} + \frac{3}{2^{\text{ml}}} \right)^r + \frac{2^{S+1}}{2^{c \cdot \text{NE.cl(ml)}}} + 6 \cdot \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C})$ ,
- (iii)  $\text{Adv}_{\text{NE}}^{\text{ctxt-1}}(\mathcal{R}'[\mathcal{A}]) \leq \text{Adv}_{\text{NE}}^{\text{ctxt-1}}(\mathcal{R}'[\mathcal{W}])$ .

Moreover,  $\mathcal{A}$  satisfies  $\text{Query}(\mathcal{A}) = c + (u+1) \cdot r + 2$  and  $\text{Mem}(\mathcal{A}) \leq 2\text{kl} + 2\text{ml} + 2\text{NE.cl(ml)} + 2 \log |\mathbb{N}| + \log u \cdot r$ . Also  $\mathcal{C}$  and  $\mathcal{W}$  satisfy  $\text{Query}(\mathcal{C}) < \text{Query}(\mathcal{R}) + \text{Query}(\mathcal{A})$ ,  $\text{Time}(\mathcal{C}) \in O(\text{Query}(\mathcal{A}) + \text{Time}(\mathcal{R}))$ ,  $\text{Query}(\mathcal{W}) = \text{Query}(\mathcal{A})$ , and  $\text{Time}(\mathcal{W}) \in O(\text{Query}(\mathcal{A}))$ .

The proof of this result has been deferred to Appendix A. We give a very brief intuition about how this impossibility proof proceeds. We define a new adversary that is similar to  $\mathcal{A}$  used for the proof of Thm. 5, but has an additional “buffer” phase before the challenge phase. In the buffer phase, it makes  $c$  encryption queries on a fixed message  $0^{\text{ml}}$  using different nonces. The key idea is that if the reduction rewinds the adversary after going past the buffer phase and still manages to pass the challenge phase, it must have remembered the  $c$  ciphertexts. Because these  $c$  ciphertexts look random (from the INDR security of NE), the memory of the reduction has to grow with  $c$ . This rules out low memory reductions that pass the challenge phase after rewinding the adversary after going past the buffer phase. As in the previous section, we can show that if a reduction cannot pass the challenge phase, it cannot have a high advantage of breaking INDR security. If the reduction does not rewind after going past the buffer phase, we can bound its advantage analogously to the straightline reduction case.

## 6 Conclusions

Our work gives memory-sensitive bounds for the security of a particular construction of a channel and shows the difficulty of providing such bounds for encryption schemes. It leaves open a number of interesting questions including: (i) whether memory-sensitive bounds can be given for other practical examples of channels, (ii) whether analogous results can be shown for any robust channels [8], and (iii) whether memory-sensitive bounds can be extended to the multi-user setting.

**Acknowledgements.** This work was partially supported by NSF grants CNS-1719146, CNS-1553758 (CA-REER), and by a Sloan Research Fellowship.

## References

1. Benedikt Auerbach, David Cash, Manuel Fersch, and Eike Kiltz. Memory-tight reductions. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 101–132. Springer, Heidelberg, August 2017.
2. Mihir Bellare, Tadayoshi Kohno, and Chanathip Namprempre. Breaking and provably repairing the ssh authenticated encryption scheme: A case study of the encode-then-encrypt-and-mac paradigm. *ACM Transactions on Information and System Security (TISSEC)*, 7(2):206–241, 2004.

3. Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Heidelberg, May / June 2006.
4. Mihir Bellare and Björn Tackmann. The multi-user security of authenticated encryption: AES-GCM in TLS 1.3. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 247–276. Springer, Heidelberg, August 2016.
5. Sandro Coretti, Yevgeniy Dodis, Siyao Guo, and John P. Steinberger. Random oracles and non-uniformity. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 227–258. Springer, Heidelberg, April / May 2018.
6. Anindya De, Luca Trevisan, and Madhur Tulsiani. Time space tradeoffs for attacks against one-way functions and PRGs. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 649–665. Springer, Heidelberg, August 2010.
7. Itai Dinur. On the streaming indistinguishability of a random permutation and a random function. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 433–460. Springer, Heidelberg, May 2020.
8. Marc Fischlin, Felix Günther, and Christian Janson. Robust channels: Handling unreliable networks in the record layers of quic and dtls 1.3. Cryptology ePrint Archive, Report 2020/718, 2020. <https://eprint.iacr.org/2020/718>.
9. Ashrujit Ghoshal and Stefano Tessaro. On the memory-tightness of hashed ElGamal. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 33–62. Springer, Heidelberg, May 2020.
10. Tetsu Iwata, Keisuke Ohashi, and Kazuhiko Minematsu. Breaking and repairing GCM security proofs. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 31–49. Springer, Heidelberg, August 2012.
11. Joseph Jaeger and Stefano Tessaro. Tight time-memory trade-offs for symmetric encryption. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 467–497. Springer, Heidelberg, May 2019.
12. David A. McGrew and John Viega. The security and performance of the Galois/counter mode (GCM) of operation. In Anne Canteaut and Kapalee Viswanathan, editors, *INDOCRYPT 2004*, volume 3348 of *LNCS*, pages 343–355. Springer, Heidelberg, December 2004.
13. J. M. Pollard. A monte carlo method for factorization. *BIT Numerical Mathematics*, 15(3):331–334, Sep 1975.
14. Jean-Jacques Quisquater and Jean-Paul Deslaucille. How easy is collision search. New results and applications to DES. In Gilles Brassard, editor, *CRYPTO’89*, volume 435 of *LNCS*, pages 408–413. Springer, Heidelberg, August 1990.
15. Phillip Rogaway. Nonce-based symmetric encryption. In Bimal K. Roy and Willi Meier, editors, *FSE 2004*, volume 3017 of *LNCS*, pages 348–359. Springer, Heidelberg, February 2004.
16. Phillip Rogaway and Thomas Shrimpton. A provable-security treatment of the key-wrap problem. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 373–390. Springer, Heidelberg, May / June 2006.
17. Stefano Tessaro and Aishwarya Thiruvengadam. Provable time-memory trade-offs: Symmetric cryptography against memory-bounded adversaries. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part I*, volume 11239 of *LNCS*, pages 3–32. Springer, Heidelberg, November 2018.
18. Serge Vaudenay. Security flaws induced by CBC padding - applications to SSL, IPSEC, WTLS... In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 534–546. Springer, Heidelberg, April / May 2002.
19. Yuyu Wang, Takahiro Matsuda, Goichiro Hanaoka, and Keisuke Tanaka. Memory lower bounds of reductions revisited. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 61–90. Springer, Heidelberg, April / May 2018.

## A Proof of Theorem 6 (Impossibility for Full-Rewinding Reductions)

In this section we prove Thm. 6 which was given in Section 5.2.

Consider the adversary  $\mathcal{A}$  shown in Fig. 17 playing against AE-1 security of NE. It has a buffer phase, followed by a challenge phase, followed by an invocation of adversary  $\mathcal{B}$ . In the buffer phase,  $\mathcal{A}$  makes  $c$

<p>Adversaries <math>\mathcal{A}^{\text{ENC,DEC}}, \mathcal{S}^{\text{ENC,DEC}}, \mathcal{W}^{\text{ENC,DEC}}</math></p> <pre> For <math>i \in [c]</math> do // <b>Buffer Phase</b>   <math>N_i \leftarrow (r + 1, i)</math>   <math>C_i \leftarrow \text{ENC}(N_i, 0^{\text{ml}})</math> For <math>i \in [r]</math> do // <b>Challenge Phase</b>   <math>j^* \leftarrow_s [u]</math>   For <math>j \in [u]</math> do     <math>M_j \leftarrow_s \{0, 1\}^{\text{ml}}</math>     <math>N_j \leftarrow (i, j)</math>     <math>C_j \leftarrow \text{ENC}(N_j, M_j)</math>   <math>M \leftarrow \text{DEC}(N_{j^*}, C_{j^*})</math>   If <math>M \neq M_{j^*}</math> then return 1 <b>bad</b> <math>\leftarrow</math> <b>true</b> Return <math>\mathcal{B}^{\text{ENC}}</math>; Return 0; Return <math>\mathcal{E}^{\text{ENC}}</math> </pre>	<p>Adversaries <math>\mathcal{B}^{\text{ENC}}, \mathcal{E}^{\text{ENC}}</math></p> <pre> <math>M_1 \leftarrow_s \{0, 1\}^{\text{ml}}</math> <math>M_2 \leftarrow_s \{0, 1\}^{\text{ml}}</math> <math>N_1 \leftarrow (0, 1); N_2 \leftarrow (0, 2)</math> <math>C_1 \leftarrow \text{ENC}(N_1, M_1)</math> <math>C_2 \leftarrow \text{ENC}(N_2, M_2)</math> For <math>K \in \{0, 1\}^{\text{kl}}</math> do   <math>\text{eq}_1 \leftarrow (\text{NE.E}(K, N_1, M_1) = C_1)</math>   <math>\text{eq}_2 \leftarrow (\text{NE.E}(K, N_2, M_2) = C_2)</math>   If <math>\text{eq}_1</math> and <math>\text{eq}_2</math> then return 1 Return 0; Return 1 </pre>
---	--

**Fig. 17.** Adversaries against the AE-1 security of NE used for proof of Thm. 6. Highlighted code is only included in  $\mathcal{A}$ . Boxed code is only included in  $\mathcal{S}$ . Adversaries  $\mathcal{B}$  (used by  $\mathcal{A}$ ) and  $\mathcal{E}$  (used by  $\mathcal{W}$ ) are repeated from Fig. 16.

encryption queries on a fixed message  $0^{\text{ml}}$  using different nonces. The next phase is the challenge phase. Each iteration of the challenge phase consists of  $\mathcal{A}$  making  $u$  encryption queries to and asking for the decryption of one randomly chosen ciphertext from the ones it received as the answer of the encryption queries. If the answer of the decryption query is not consistent with the corresponding encryption query, then  $\mathcal{A}$  halts and returns 1. There are  $r$  iterations of the challenge phase. If all of these phases are passed, then  $\mathcal{A}$  runs adversary  $\mathcal{B}$  which is defined in the same figure (identically to the  $\mathcal{B}$  that was defined in Fig. 16). From the code of  $\mathcal{A}$  we can see that it makes  $r \cdot u + c + 2$  encryption queries,  $r$  decryption queries, and satisfies  $\text{Mem}(\mathcal{A}_2) \leq 2\text{kl} + 2\text{ml} + 2\text{NE.cl}(\text{ml}) + 2 \log |\mathcal{N}| + \log u \cdot r$ .

To prove the theorem we need to separately establish the three advantage claims (and corresponding statements about the efficiency of various algorithms).

**FIRST AND THIRD CLAIMS.** For the first claim, note that  $\text{Adv}_{\text{NE}}^{\text{ae-1}}(\mathcal{A}) = \text{Adv}_{\text{NE}}^{\text{ae-1}}(\mathcal{B})$  because  $M$  will always equal  $M_{j^*}$  when  $\mathcal{A}$  is playing  $\text{G}_{\text{NE},b}^{\text{ae-1}}$ . The simple analysis showing that  $\text{Adv}_{\text{NE}}^{\text{ae-1}}(\mathcal{B}) \geq 1 - 2^{\text{kl}}/2^{2\text{NE.cl}(\text{ml})}$  is given in Appendix B.

For the third claim, consider adversary  $\mathcal{W}$  defined as shown in Fig. 17. It is identical to  $\mathcal{A}$ , except that it calls  $\mathcal{E}$  which is defined in the same figure (identically to the  $\mathcal{E}$  that was defined Fig. 16), which is similar to  $\mathcal{B}$  but always returns 1. Because  $\mathcal{R}'$  is faithful,  $\mathcal{B}$  would never return 0 when run by  $\mathcal{R}'[\mathcal{A}]$  playing  $\text{G}_{\text{NE},b}^{\text{ctxt-1}}$  so  $\text{Adv}_{\text{NE}}^{\text{ctxt-1}}(\mathcal{R}'[\mathcal{A}]) = \text{Adv}_{\text{NE}}^{\text{ctxt-1}}(\mathcal{R}'[\mathcal{W}])$  holds trivially. The claims on its complexity follow from its similarity to  $\mathcal{A}$ .

**SECOND CLAIM.** We spend the rest of the proof establishing the second claim. Consider the adversary  $\mathcal{S}$  in Fig. 17. It behaves identically to  $\mathcal{A}$  until the flag **bad** is set. Using the Fundamental Lemma of Game Playing [3], we obtain for  $b \in \{0, 1\}$  that

$$\left| \Pr \left[ \text{G}_{\text{NE},b}^{\text{indr}}(\mathcal{R}[\mathcal{A}]) \right] - \Pr \left[ \text{G}_{\text{NE},b}^{\text{indr}}(\mathcal{R}[\mathcal{S}]) \right] \right| \leq \Pr \left[ \mathcal{R}[\mathcal{A}] \text{ sets bad in } \text{G}_{\text{NE},b}^{\text{indr}} \right].$$

We will use the sequence of games shown in Fig. 18 to bound these probabilities that  $\mathcal{R}[\mathcal{A}]$  sets **bad**. Start by considering the games  $\text{G}_b^0$  for  $b \in \{0, 1\}$ . They were obtained by plugging the code of  $\mathcal{R}[\mathcal{A}]$  into  $\text{G}_{\text{NE},b}^{\text{indr}}$ , then changing the code to output **true** if  $\mathcal{R}$  passes all of  $\mathcal{A}$ 's challenges and output **false** otherwise (and adding a flag **crw** that we will use later). In other words,  $\text{G}_b^0$  outputs **true** whenever  $\mathcal{R}[\mathcal{A}]$  would set **bad** in  $\text{G}_{\text{NE},b}^{\text{indr}}$ . This gives us that

$$\text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{R}[\mathcal{A}]) \leq \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{R}[\mathcal{S}]) + \Pr \left[ \text{G}_0^0 \right] + \Pr \left[ \text{G}_1^0 \right]. \quad (4)$$



Games $G_b^0$ , $G^1$ , $H^2$	Games $G^2$ , $G^3$
<pre> <b><math>\overline{b} \leftarrow 1</math></b> <math>K^* \leftarrow \mathcal{R}. \text{NE.Kg}(); k \leftarrow 0; \sigma \leftarrow \mathcal{R}^{\text{ENC}_b}. \text{Init}</math> While <math>k \leq \mathcal{R}. \text{rew}</math> do   For <math>i \in [c]</math> do // <b>Buffer Phase</b>     <math>N_i \leftarrow (r + 1, i)</math>     <math>(\sigma, \text{rf}, C_i) \leftarrow \mathcal{R}^{\text{ENC}_b}. \text{SimEnc}(\sigma, N_i, 0^{\text{ml}})</math>     If <b>rf</b> then goto NEXT2   For <math>i \in [r]</math> do // <b>Challenge Phase</b>     <math>j^* \leftarrow \mathcal{R}[u]</math>     For <math>j \in [u]</math> do       <math>M_j \leftarrow \mathcal{R}\{0, 1\}^{\text{ml}}</math>       <math>N_j \leftarrow (i, j)</math>       <math>(\sigma, \text{rf}, C_j) \leftarrow \mathcal{R}^{\text{ENC}_b}. \text{SimEnc}(\sigma, N_j, M_j)</math>       If <b>rf</b> then goto NEXT     <math>(\sigma, \text{rf}, M) \leftarrow \mathcal{R}^{\text{ENC}_b}. \text{SimDec}(\sigma, N_{j^*}, C_{j^*})</math>     If <b>rf</b> then goto NEXT     If <math>M \neq M_{j^*}</math> then       <math>\sigma \leftarrow \mathcal{R}^{\text{ENC}_b}. \text{Upd}(\sigma, 1)</math>       Goto NEXT   Return <b>true</b>   NEXT: <math>\text{crw} \leftarrow \text{true}; k \leftarrow \mathcal{R}. \text{rew} + 1</math>   NEXT2: <math>k \leftarrow k + 1</math> Return <b>false</b>  Oracle <math>\text{ENC}_b(N, M)</math> <math>C_0 \leftarrow \mathcal{R}C</math> <math>C_1 \leftarrow \text{NE.E}(K^*, N, M)</math> <math>C_1 \leftarrow \text{NE.E}(K^*, N, 0^{\text{ml}})</math> Return <math>C_b</math> </pre>	<pre> <math>\overline{b} \leftarrow 1</math> <math>K^* \leftarrow \mathcal{R}. \text{NE.Kg}(); k \leftarrow 0; \sigma \leftarrow \mathcal{R}^{\text{ENC}_b}. \text{Init}</math> While <math>k \leq \mathcal{R}. \text{rew}</math> do   For <math>i \in [c]</math> do // <b>Buffer Phase</b>     <math>N_i \leftarrow (r + 1, i)</math>     <math>(\sigma, \text{rf}, C_i) \leftarrow \mathcal{R}^{\text{ENC}_b}. \text{SimEnc}(\sigma, N_i, 0^{\text{ml}})</math>     If <b>rf</b> then goto NEXT2   For <math>i \in [r]</math> do // <b>Challenge Phase</b>     <math>j^* \leftarrow \mathcal{R}[u]</math>     For <math>j \in [u]</math> do       <math>M_j \leftarrow \mathcal{R}\{0, 1\}^{\text{ml}}</math>       <math>N_j \leftarrow (i, j)</math>       <math>(\sigma, \text{rf}, C_j) \leftarrow \mathcal{R}^{\text{ENC}_b}. \text{SimEnc}(\sigma, N_j, M_j)</math>       If <b>rf</b> then goto NEXT     <math>(\sigma, \text{rf}, M) \leftarrow \mathcal{R}^{\text{ENC}_b}. \text{SimDec}(\sigma, N_{j^*}, C_{j^*})</math>     If <b>rf</b> then goto NEXT     If <math>M \neq M_{j^*}</math> then       <math>\sigma \leftarrow \mathcal{R}^{\text{ENC}_b}. \text{Upd}(\sigma, 1)</math>       Goto NEXT   Return <b>true</b> <math>\wedge</math> <b>crw</b>   NEXT: <math>\text{crw} \leftarrow \text{true}</math>   NEXT2: <math>k \leftarrow k + 1</math> Return <b>false</b>  Oracle <math>\text{ENC}_b(N, M)</math> <math>C_1 \leftarrow \text{NE.E}(K^*, N, 0^{\text{ml}})</math> <math>C_1 \leftarrow \mathcal{R}C</math> Return <math>C_b</math> </pre>

**Fig. 18.** Sequence of games used in proof Theorem 6. On the left, boxed code is *not* included in  $G^1$  and highlighted code is only included in  $H^2$ . On the right, highlighted code is included in both games (indicating where they differ from  $G^1$ ) and boxed code is only included in  $G^3$ .

Now consider the game  $G^1$  defined in the same figure. It is identical to either  $G_b^0$  except that it answers all encryption queries with the encryption of the message  $0^{\text{ml}}$ . The following lemma allows us to bound the  $\Pr[G_b^0]$ 's using  $\Pr[G^1]$  and the advantage of two adversaries attacking the INDR security of NE. Its proof is given in Section C.2.

**Lemma 5.** *There exist adversaries  $C_1$  and  $C_2$  such that*

$$\Pr[G_1^0] \leq \Pr[G_0^0] + \text{Adv}_{\text{NE}}^{\text{indr}}(C_1) \text{ and } \Pr[G_0^0] \leq \Pr[G^1] + \text{Adv}_{\text{NE}}^{\text{indr}}(C_2)$$

where  $G_b^0$  and  $G^1$  are defined as in Fig. 18. Moreover  $\text{Query}(C_1) < \text{Query}(\mathcal{R}) + \text{Query}(\mathcal{A})$  and  $\text{Time}(C_1) \in O(\text{Query}(\mathcal{A}) + \text{Time}(\mathcal{R}))$ . Adversary  $C_2$ 's complexity is the same.

Next we are going to bound  $\Pr[G^1] \leq \Pr[H^2] + \Pr[G^2]$ . Game  $H^2$  does not let  $\mathcal{R}$  rewind again if it has already passed the buffer, so it captures the possibility that  $\mathcal{R}$  passes the challenges on its first attempt past the buffer. Game  $G^2$  only outputs **true** if  $\mathcal{R}$  has passed the buffer, then rewinded, and then passed the buffer again before successfully completing the challenge. To make this formal, let  $\text{crw}$  denote the event that  $\text{crw} = \text{true}$  at the end of the execution of a game and  $\overline{\text{crw}}$  to denote the opposite event in the following probability statements. We claim that,

$$\Pr[G^1 \wedge \overline{\text{crw}}] \leq \Pr[H^2] \text{ and } \Pr[G^1 \wedge \text{crw}] \leq \Pr[G^2]. \quad (5)$$

Adversary $\mathcal{D}_1(\phi, M_1, \dots, M_u)$	Adversary $\mathcal{D}_2(\phi, j^*)$
If $\phi = \perp$ then $K^* \leftarrow \text{NE.Kg}; k \leftarrow 0; \sigma \leftarrow \mathcal{R}^{\text{ENC}}.\text{Init}$ While $k \leq \mathcal{R}.\text{rew}$ do For $i \in [c]$ do // <b>Buffer Phase</b> $N_i \leftarrow (r + 1, i)$ $(\sigma, \text{rf}, C_i) \leftarrow \mathcal{R}^{\text{ENC}}.\text{SimEnc}(\sigma, N_i, M_0)$ If <b>rf</b> then goto NEXT2 Goto BUF-END NEXT2: $k \leftarrow k + 1$ BUF-END: $\phi \leftarrow (1, K^*, \sigma)$ $(i, K^*, \sigma) \leftarrow \phi_1$ For $j \in [u]$ do $N_j \leftarrow (i, j)$ $(\sigma, \cdot, C_j) \leftarrow \mathcal{R}^{\text{ENC}}.\text{SimEnc}(\sigma, N_j, M_j)$ $\phi \leftarrow (i, K^*, \sigma)$ Return $\phi$	$(i, K^*, \sigma) \leftarrow \phi_2$ $N_{j^*} \leftarrow (i, j^*)$ $C_{j^*} \leftarrow \text{NE.E}(K^*, N_{j^*}, 0^{\text{ml}})$ $(\sigma, \cdot, M) \leftarrow \mathcal{R}^{\text{ENC}}.\text{SimDec}(N_{j^*}, C_{j^*})$ $\phi \leftarrow (i + 1, K^*, \sigma)$ Return $(\phi, M)$ <hr/> Oracle $\text{ENC}(N, M)$ Return $\text{NE.E}(K^*, N, 0^{\text{ml}})$

**Fig. 19.** Adversary  $(\mathcal{D}_1, \mathcal{D}_2)$  used to prove Lemma 7.

The former of these holds because  $G^1$  and  $H^2$  are identical until  $\text{crw}$  has been set (and hence  $\Pr[G^1 \wedge \overline{\text{crw}}] = \Pr[H^1 \wedge \overline{\text{crw}}]$ ). The latter holds because  $G^2$  is identical to  $G^1$  except it only outputs **true** if  $G^1$  would *and*  $\text{crw} = \text{true}$ .

Because  $\mathcal{R}$  only has one attempt to pass the challenge phase in  $H^2$ , we can directly bound the probability that it succeeds in this by using  $\mathcal{R}$  to construct an adversary for  $G_{u, \text{ml}}^{\text{it-chl-r}}$  and bounding its advantage with Lemma 2. This gives the following lemma.

**Lemma 6.** *Let  $\mathcal{R}$  be a full-rewinding, nonce-respecting, faithful black-box reduction from AE-1 to nonce-respecting INDR with  $\text{Mem}(\mathcal{R}) = S$ . Then,*

$$\Pr[H^2] \leq \left( \frac{2(S + \log r + \text{kl}) + 2\text{ml}}{u} + \frac{3}{2^{\text{ml}}} \right)^r$$

where  $H^2$  is defined as in Fig. 18.

*Proof.* Consider the adversary  $(\mathcal{D}_1, \mathcal{D}_2)$  for  $G_{u, \text{ml}}^{\text{it-chl-r}}$  defined in Fig. 19. It assumes that  $\mathcal{R}$  will succeed in  $H^2$  and tries to use this to succeed in its own game. The first time  $\mathcal{D}_1$  is executed (i.e. when  $\phi = \perp$ ) it simulates the buffer phase, allowing  $\mathcal{R}$  to rewind up to  $\mathcal{R}.\text{rew}$  times until  $\mathcal{R}$  completes the buffer phase (if  $\mathcal{R}$  never completes the buffer phase, reaching the code “Goto BUF-END”, then it couldn’t have won in  $H^2$  so we do not need to worry about the behavior of  $\mathcal{D}_1$  in this case).

After this, and in all subsequent executions,  $\mathcal{D}_1$  uses its input as the messages given to  $\mathcal{R}$  in one iteration of the challenge. We let  $\mathcal{D}_1$  ignore  $\mathcal{R}$ ’s output **rf** here because if  $\mathcal{R}$  ever output **rf** = **true** then it would be unable to win that execution of  $H^2$ . Then  $\mathcal{D}_1$  halts, passing the current state of  $\mathcal{R}$  to  $\mathcal{D}_2$  which uses  $\mathcal{R}$  to attempt to answer its challenge.

The view of  $\mathcal{R}$  is perfectly simulated for any execution of  $H^2$  in which it would have succeeded, so  $\Pr[H^2] \leq \text{Adv}_{u, \text{ml}}^{\text{it-chl-r}}(\mathcal{D}_1, \mathcal{D}_2)$ . The state used by  $(\mathcal{D}_1, \mathcal{D}_2)$  is a tuple  $\phi = (i, K^*, \sigma)$  where  $|i| = \log r$ ,  $|K^*| = \text{kl}$ , and  $|\sigma| = S$  so  $(\mathcal{D}_1, \mathcal{D}_2)$  is  $S + \log r + \text{kl}$ -bounded. Applying Lemma 2 gives the claimed bound.  $\square$

We now state two lemmas which give us a bound on  $\Pr[G^2]$  via  $G^3$ . First, in Lemma 7 we use the INDR security of NE to transition from  $G^2$  to  $G^3$  which is identical, except the encryption oracle returns random ciphertexts to  $\mathcal{R}$ . Its proof is given in Appendix C.3. Now consider  $G^3$ . The reduction  $\mathcal{R}$  must pass the

buffer phase, then rewind, then pass the buffer phase again to win this game. Because it is faithful, passing the buffer phase the second time requires  $\mathcal{R}$  to remember  $c$  random ciphertexts. In Lemma 8, we use a compression argument (via a lemma from [6]) to give an information theoretic bound on  $\Pr[G^3]$ . Its proof is given in Appendix E.

**Lemma 7.** *There exists an adversary  $\mathcal{C}_3$  such that*

$$\Pr[G^2] \leq \Pr[G^3] + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_3)$$

where  $G^2$  and  $G^3$  are defined as in Fig. 18. Moreover,  $\text{Query}(\mathcal{C}_3) < \text{Query}(\mathcal{R}) + \text{Query}(\mathcal{A})$  and  $\text{Time}(\mathcal{C}_3) \in O(\text{Query}(\mathcal{A}) + \text{Time}(\mathcal{R}))$ .

**Lemma 8.** *Let  $\mathcal{R}$  be a full-rewinding, nonce-respecting, faithful black-box reduction from AE-1 to nonce-respecting INDR with  $\text{Mem}(\mathcal{R}) = S$ . Then,*

$$\Pr[G^3] \leq \frac{2^S}{|C|^c} = \frac{2^S}{2^{c \cdot \text{NE.cl}(\text{ml})}}$$

where  $G^3$  is defined as in Fig. 18.

We combine (in order) Equation 4, Lemma 5, Equation 5, and Lemma 7 to obtain

$$\begin{aligned} \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{R}[\mathcal{A}]) &\leq \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{R}[\mathcal{S}]) + \Pr[G_0^0] + \Pr[G_1^0] \\ &\leq \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{R}[\mathcal{S}]) + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_1) + 2 \cdot \Pr[G_0^0] \\ &\leq \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{R}[\mathcal{S}]) + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_1) + 2 \cdot (\text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_2) + \Pr[G^1]) \\ &\leq \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{R}[\mathcal{S}]) + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_1) + 2 \cdot (\text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_2) + \Pr[H^2] + \Pr[G^2]) \\ &\leq \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{R}[\mathcal{S}]) + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_1) + 2 \cdot (\text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_2) + \Pr[H^2] + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_3) + \Pr[G^3]). \end{aligned}$$

To complete the proof, we apply our information theoretic bounds from Lemma 7 and Lemma 8, then combine the four INDR adversaries  $\mathcal{R}[\mathcal{S}]$ ,  $\mathcal{C}_1$ ,  $\mathcal{C}_2$ , and  $\mathcal{C}_3$ . Let  $\mathcal{C}$  be the INDR adversary that randomly chooses one of  $\mathcal{R}[\mathcal{S}]$ ,  $\mathcal{C}_1$ ,  $\mathcal{C}_2$ , and  $\mathcal{C}_3$  (with probabilities  $1/6$ ,  $1/6$ ,  $2/6$ , and  $2/6$ , respectively), then runs the adversary it chose, outputting whatever that adversary does. Simple calculations give

$$6 \cdot \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}) = \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{R}[\mathcal{S}]) + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_1) + 2 \cdot \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_2) + 2 \cdot \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_3).$$

The claimed complexity of  $\mathcal{C}$  follows from that of  $\mathcal{R}[\mathcal{S}]$ ,  $\mathcal{C}_1$ ,  $\mathcal{C}_2$ , and  $\mathcal{C}_3$ . □

## B Analysis of Brute-Force Adversary's Advantage

We prove a bound on the advantage of the adversary  $\mathcal{B}$  we use in the proofs of Thm. 5 and Thm. 6.

**Lemma 9.** *Fix  $\text{kl}, \text{ml} \in \mathbb{N}$ . Let NE be a nonce-based encryption scheme for which  $\{0, 1\}^{\text{ml}} \subseteq \text{NE.M}$ ,  $[\text{NE.Kg}] = \{0, 1\}^{\text{kl}}$ , and  $\{(0, 1), (0, 2)\} \subseteq \text{NE.N}$ . Then,*

$$\text{Adv}_{\text{NE}}^{\text{ae-1}}(\mathcal{B}) \geq 1 - 2^{\text{kl}}/2^{2\text{NE.cl}(\text{ml})}.$$

for  $\mathcal{B}$  as defined in Fig. 16.

*Proof.* Adversary  $\mathcal{B}$  makes two encryption queries then performs a brute-force to check if there is an encryption key consistent with those queries, returning 1 if so, and 0 otherwise. In  $G_1^{\text{ae-1}}$  there will always be a consistent key so  $\mathcal{B}$  will always return 1. For fixed  $M_1, M_2$ , and  $N_1 \neq N_2$  there are at most  $2^{\text{kl}}$  values the ordered pair  $(\text{NE.E}(K, N_1, M_1), \text{NE.E}(K, N_2, M_2))$  can take (one for each  $K \in \{0, 1\}^{\text{kl}}$ ). Moreover, there are  $2^{2\text{NE.cl}(\text{ml})}$  different values of  $(C_1, C_2)$  that might be chosen at random by the encryption oracle, so  $\mathcal{B}$  will return 1 with probability at most  $2^{\text{kl}}/2^{2\text{NE.cl}(\text{ml})}$  when interacting with  $G_{\text{NE},0}^{\text{ae-1}}$ . Combining these observations gives the desired bound. □

Adversary $\mathcal{C}_1^{\text{ENC}}, \mathcal{C}_2^{\text{ENC}}$ $\sigma \leftarrow \mathcal{R}^{\text{SIMENC}}.\text{Init}$ For $i \in [r]$ do $j^* \leftarrow [u]$ For $j \in [u]$ do $M_j \leftarrow \{0, 1\}^{\text{ml}}$ $N_j \leftarrow (i, j)$ $(\sigma, \cdot, C_j) \leftarrow \mathcal{R}^{\text{SIMENC}}.\text{SimEnc}(\sigma, N_j, M_j)$ $(\sigma, \cdot, M) \leftarrow \mathcal{R}^{\text{SIMENC}}.\text{SimDec}(\sigma, N_{j^*}, C_{j^*})$ If $M \neq M_{j^*}$ then return 0; <span style="border: 1px solid black; padding: 2px;">return 1</span> Return 1; <span style="border: 1px solid black; padding: 2px;">Return 0</span>	Oracle $\text{SIMENC}(N, M)$ <span style="border: 1px solid black; padding: 2px;">Return <math>\text{ENC}(N, M)</math></span> <span style="border: 1px solid black; padding: 2px;">Return <math>\text{ENC}(N, 0^{\text{ml}})</math></span>
--	--

**Fig. 20.** Adversaries  $\mathcal{C}_1$  and  $\mathcal{C}_2$  used for Lemma 3. The highlighted code is present only in  $\mathcal{C}_1$  and the boxed code is present only in  $\mathcal{C}_2$ .

## C Proofs of Lemma 3, Lemma 5, and Lemma 7 (INDR Reductions)

### C.1 Proof of Lemma 3

Consider the adversary  $\mathcal{C}_1$  shown in Fig. 20. It was obtained by copying the code of  $\mathcal{G}_0^0$  and  $\mathcal{G}_1^0$  (from Fig. 15), then modifying the code to use its own ENC oracle to respond to encryption queries from  $\mathcal{R}$  and, to return 1 whenever the games would output `true` and 0 otherwise. Consequently,  $\mathcal{C}_1$  playing  $\mathcal{G}_{\text{NE},0}^{\text{indr}}$  (resp.  $\mathcal{G}_{\text{NE},1}^{\text{indr}}$ ) perfectly simulates  $\mathcal{G}_0^0$  (resp.  $\mathcal{G}_1^0$ ) to  $\mathcal{R}$  and we have

$$\Pr[\mathcal{G}_0^0] = \Pr[\mathcal{G}_{\text{NE},0}^{\text{indr}}(\mathcal{C}_1)] \quad \text{and} \quad \Pr[\mathcal{G}_1^0] = \Pr[\mathcal{G}_{\text{NE},1}^{\text{indr}}(\mathcal{C}_1)].$$

It follows that,

$$\Pr[\mathcal{G}_1^0] \leq \Pr[\mathcal{G}_0^0] + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_1).$$

The claims on  $\mathcal{C}_1$ 's complexity can be verified from its code.

Now consider the adversary  $\mathcal{C}_2$  shown in Fig. 20. It was obtained by copying the code of  $\mathcal{G}_0^0$  and  $\mathcal{G}_1^1$  (from Fig. 15), then modifying the code to use its own ENC oracle to respond to encryption queries from  $\mathcal{R}$ , and to return 0 whenever the games would outputs `true` and 1 otherwise. Consequently,  $\mathcal{C}_2$  playing  $\mathcal{G}_{\text{NE},0}^{\text{indr}}$  (resp.  $\mathcal{G}_{\text{NE},1}^{\text{indr}}$ ) perfectly simulates  $\mathcal{G}_0^0$  (resp.  $\mathcal{G}_1^1$ ) to  $\mathcal{R}$  and we have

$$1 - \Pr[\mathcal{G}_0^0] = \Pr[\mathcal{G}_{\text{NE},0}^{\text{indr}}(\mathcal{C}_2)] \quad \text{and} \quad 1 - \Pr[\mathcal{G}_1^1] = \Pr[\mathcal{G}_{\text{NE},1}^{\text{indr}}(\mathcal{C}_2)].$$

It follows that,

$$\Pr[\mathcal{G}_0^0] \leq \Pr[\mathcal{G}_1^1] + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_2).$$

The claims on  $\mathcal{C}_2$ 's complexity can be verified from its code. □

### C.2 Proof of Lemma 5

Consider the adversary  $\mathcal{C}_1$  shown in Fig. 21. It was obtained by copying the code of  $\mathcal{G}_0^0$  and  $\mathcal{G}_1^0$  (from Fig. 18), then modifying the code to use its own ENC oracle to respond to encryption queries from  $\mathcal{R}$ , and to return 1 whenever the games would output `true` and 0 otherwise. Consequently,  $\mathcal{C}_1$  playing  $\mathcal{G}_{\text{NE},0}^{\text{indr}}$  (resp.  $\mathcal{G}_{\text{NE},1}^{\text{indr}}$ ) perfectly simulates  $\mathcal{G}_0^0$  (resp.  $\mathcal{G}_1^0$ ) to  $\mathcal{R}$  and we have

$$\Pr[\mathcal{G}_0^0] = \Pr[\mathcal{G}_{\text{NE},0}^{\text{indr}}(\mathcal{C}_1)] \quad \text{and} \quad \Pr[\mathcal{G}_1^0] = \Pr[\mathcal{G}_{\text{NE},1}^{\text{indr}}(\mathcal{C}_1)].$$

<p>Adversaries <math>\mathcal{C}_1^{\text{ENC}}, \mathcal{C}_2^{\text{ENC}}</math></p> <p><math>k \leftarrow 0; \sigma \leftarrow \mathcal{R}^{\text{SIMENC}}.\text{Init}</math></p> <p>While <math>k \leq \mathcal{R}.\text{rew}</math> do</p> <p>  For <math>i \in [c]</math> do // <b>Buffer Phase</b></p> <p>    <math>N_i \leftarrow (r + 1, i)</math></p> <p>    <math>(\sigma, \text{rf}, C_i) \leftarrow \mathcal{R}^{\text{SIMENC}}.\text{SimEnc}(\sigma, N_i, 0^{\text{ml}})</math></p> <p>    If rf then goto NEXT2</p> <p>  For <math>i \in [r]</math> do // <b>Challenge Phase</b></p> <p>    <math>j^* \leftarrow [u]</math></p> <p>    For <math>j \in [u]</math> do</p> <p>      <math>M_j \leftarrow \{0, 1\}^{\text{ml}}</math></p> <p>      <math>N_j \leftarrow (i, j)</math></p> <p>      <math>(\sigma, \text{rf}, C_j) \leftarrow \mathcal{R}^{\text{SIMENC}}.\text{SimEnc}(\sigma, N_j, M_j)</math></p> <p>      If rf then goto NEXT</p> <p>    <math>(\sigma, \text{rf}, M) \leftarrow \mathcal{R}^{\text{SIMENC}}.\text{SimDec}(\sigma, N_{j^*}, C_{j^*})</math></p> <p>    If rf then goto NEXT</p> <p>    If <math>M \neq M_{j^*}</math> then</p> <p>      <math>\sigma \leftarrow \mathcal{R}^{\text{ENC}}.\text{Upd}(\sigma, 1)</math></p> <p>      Goto NEXT</p> <p>  Return 1; <span style="border: 1px solid black; padding: 2px;">Return 0</span></p> <p>  NEXT: <math>\text{crw} \leftarrow \text{true}</math></p> <p>  NEXT2: <math>k \leftarrow k + 1</math></p> <p>Return 0; <span style="border: 1px solid black; padding: 2px;">Return 1</span></p> <hr/> <p>Oracle <math>\text{SIMENC}(N, M)</math></p> <p><math>C \leftarrow \text{ENC}(N, M)</math></p> <p><span style="border: 1px solid black; padding: 2px;"><math>C \leftarrow \text{ENC}(N, 0^{\text{ml}})</math></span></p> <p>Return <math>C</math></p>	<p>Adversary <math>\mathcal{C}_3^{\text{ENC}}</math></p> <p><math>k \leftarrow 0; \sigma \leftarrow \mathcal{R}^{\text{SIMENC}}.\text{Init}</math></p> <p>While <math>k \leq \mathcal{R}.\text{rew}</math> do</p> <p>  For <math>i \in [c]</math> do // <b>Buffer Phase</b></p> <p>    <math>N_i \leftarrow (r + 1, i)</math></p> <p>    <math>(\sigma, \text{rf}, C_i) \leftarrow \mathcal{R}^{\text{SIMENC}}.\text{SimEnc}(\sigma, N_i, 0^{\text{ml}})</math></p> <p>    If rf then goto NEXT2</p> <p>  For <math>i \in [r]</math> do // <b>Challenge Phase</b></p> <p>    <math>j^* \leftarrow [u]</math></p> <p>    For <math>j \in [u]</math> do</p> <p>      <math>M_j \leftarrow \{0, 1\}^{\text{ml}}</math></p> <p>      <math>N_j \leftarrow (i, j)</math></p> <p>      <math>(\sigma, \text{rf}, C_j) \leftarrow \mathcal{R}^{\text{SIMENC}}.\text{SimEnc}(\sigma, N_j, M_j)</math></p> <p>      If rf then goto NEXT</p> <p>    <math>(\sigma, \text{rf}, M) \leftarrow \mathcal{R}^{\text{SIMENC}}.\text{SimDec}(\sigma, N_{j^*}, C_{j^*})</math></p> <p>    If rf then goto NEXT</p> <p>    If <math>M \neq M_{j^*}</math> then</p> <p>      <math>\sigma \leftarrow \mathcal{R}^{\text{SIMENC}}.\text{Upd}(\sigma, 1)</math></p> <p>      Goto NEXT</p> <p>  If crw then return 1 else return 0</p> <p>  NEXT: <math>\text{crw} \leftarrow \text{true}</math></p> <p>  NEXT2: <math>k \leftarrow k + 1</math></p> <p>Return 0</p> <hr/> <p>Oracle <math>\text{SIMENC}(N, M)</math></p> <p><math>C \leftarrow \text{ENC}(N, 0^{\text{ml}})</math></p> <p>Return <math>C</math></p>
--	---

**Fig. 21.** Adversaries  $\mathcal{C}_1$  and  $\mathcal{C}_2$  for Lemma 5 and adversary  $\mathcal{C}_3$  for Lemma 7. Highlighted code is only included in  $\mathcal{C}_1$  and boxed code is only included in  $\mathcal{C}_2$ .

It follows that,

$$\Pr[G_1^0] \leq \Pr[G_0^0] + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_1).$$

The claims on  $\mathcal{C}_1$ 's complexity can be verified from its code.

Now consider the adversary  $\mathcal{C}_2$  shown in Fig. 21. It was obtained by copying the code of  $G_0^0$  and  $G^1$  (from Fig. 18), then modifying the code to use its own ENC oracle to respond to encryption queries from  $\mathcal{R}$ , and to return 0 whenever the games would output **true** and 1 otherwise. Consequently,  $\mathcal{C}_2$  playing  $G_{\text{NE},0}^{\text{indr}}$  (resp.  $G_{\text{NE},1}^{\text{indr}}$ ) perfectly simulates  $G_0^0$  (resp.  $G^1$ ) to  $\mathcal{R}$  and we have

$$1 - \Pr[G_0^0] = \Pr[G_{\text{NE},0}^{\text{indr}}(\mathcal{C}_2)] \quad \text{and} \quad 1 - \Pr[G^1] = \Pr[G_{\text{NE},1}^{\text{indr}}(\mathcal{C}_2)].$$

It follows that,

$$\Pr[G_0^0] \leq \Pr[G^1] + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_2).$$

The claims on  $\mathcal{C}_2$ 's complexity can be verified from its code. □

### C.3 Proof of Lemma 7

*Proof.* Consider the adversary  $\mathcal{C}_3$  shown in Fig. 21. It was obtained by copying the code of  $G^2$  and  $G^3$  (from Fig. 18), then modifying the code to use its own ENC oracle to respond to encryption queries from  $\mathcal{R}$ , and

Adversary $\mathcal{D}_1(\phi, M_1, \dots, M_u)$	Adversary $\mathcal{D}_2(\phi, j^*)$
If $\phi = \perp$ then	$(i, K^*, \sigma) \leftarrow \phi$
$i \leftarrow 1; K^* \leftarrow \text{NE.Kg}; \sigma \leftarrow \mathcal{R}^{\text{ENC}}.\text{Init}$	$N_{j^*} \leftarrow (i, j^*)$
Else	$C_{j^*} \leftarrow \text{NE.E}(K^*, N_{j^*}, 0^{\text{ml}})$
$(i, K^*, \sigma) \leftarrow \phi$	$(\sigma, \cdot, M) \leftarrow \mathcal{R}^{\text{ENC}}.\text{SimDec}(N_{j^*}, M_{j^*})$
Foreach $j \in [u]$ do	$\phi \leftarrow (i + 1, K^*, \sigma)$
$N_j \leftarrow (i, j)$	Return $(\phi, M)$
$(\sigma, \cdot, C_j) \leftarrow \mathcal{R}^{\text{ENC}}.\text{SimEnc}(\sigma, N_{i,j}, M_j)$	Oracle $\text{ENC}(N, M)$
$\phi \leftarrow (i, K^*, \sigma)$	Return $\text{NE.E}(K^*, N, 0^{\text{ml}})$
Return $\phi$	

**Fig. 22.** Adversary  $(\mathcal{D}_1, \mathcal{D}_2)$  used in the proof of Lemma 4.

to return 1 whenever the games would output **true** and 0 otherwise. Consequently,  $\mathcal{C}_3$  playing  $\mathsf{G}_{\text{NE},0}^{\text{indr}}$  (resp.  $\mathsf{G}_{\text{NE},1}^{\text{indr}}$ ) perfectly simulates  $\mathsf{G}^3$  (resp.  $\mathsf{G}^2$ ) to  $\mathcal{R}$  and we have

$$\Pr[\mathsf{G}^3] = \Pr[\mathsf{G}_{\text{NE},0}^{\text{indr}}(\mathcal{C}_3)] \quad \text{and} \quad \Pr[\mathsf{G}^2] = \Pr[\mathsf{G}_{\text{NE},1}^{\text{indr}}(\mathcal{C}_3)] .$$

It follows that,

$$\Pr[\mathsf{G}^2] \leq \Pr[\mathsf{G}^3] + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C}_3) .$$

The claims on  $\mathcal{C}_3$ 's complexity can be verified from its code.  $\square$

## D Proof of Lemma 4

Consider the adversary  $(\mathcal{D}_1, \mathcal{D}_2)$  shown in Fig. 22. Adversary  $\mathcal{D}_1$  samples a key for NE and starts running  $\mathcal{R}$  from the beginning in the first round of  $\mathsf{G}_{u,\text{ml}}^{\text{it-chl-}r}$  (when  $\phi = \perp$ ) and from the state at which  $\mathcal{R}$  was last stopped in later rounds. It responds to the encryption queries of  $\mathcal{R}$  as in  $\mathsf{G}^1$  (Fig. 15). Adversary  $\mathcal{D}_1$  makes  $u$  encryption queries which are distributed identically to the queries made to  $\mathcal{R}$  in one iteration of  $\mathsf{G}^1$  since the inputs to  $\mathcal{D}_1$  are sampled uniformly at random in  $\mathsf{G}_{u,\text{ml}}^{\text{it-chl-}r}$ . It then stops running  $\mathcal{R}$  and outputs its state  $\phi$  which consists of the state  $\sigma$  of  $\mathcal{R}$  along with the iteration number  $i$  and the key  $K^*$ . Adversary  $\mathcal{D}_2$  resumes running  $\mathcal{R}$  from the state it was last stopped and makes a decryption query which is identically distributed to the decryption query made in  $\mathsf{G}^1$  because  $j^*$  is chosen at random and independently of the  $M_i$ 's in  $\mathsf{G}_{u,\text{ml}}^{\text{it-chl-}r}$ . It outputs the message  $M$  it received as answer to the decryption query along with a state  $\phi$  which consists of the current state of  $\mathcal{R}$ , the next iteration number  $i + 1$ , and the key  $K^*$ .

It follows from their that during the  $r$  iterations of  $\mathsf{G}_{u,\text{ml}}^{\text{it-chl-}r}$ , the view of  $\mathcal{R}$  provided by  $(\mathcal{D}_1, \mathcal{D}_2)$  perfectly matches its view in  $\mathsf{G}^1$ . Note that  $\mathsf{G}_{u,\text{ml}}^{\text{it-chl-}r}(\mathcal{D}_1, \mathcal{D}_2)$  will output **true** whenever  $\mathsf{G}^1$  would output **true**, so  $\Pr[\mathsf{G}^1] \leq \text{Adv}_{u,\text{ml}}^{\text{it-chl-}r}(\mathcal{D}_1, \mathcal{D}_2)$ .

The state of  $(\mathcal{D}_1, \mathcal{D}_2)$  is always a tuple  $(i, K^*, \sigma)$  for which  $|i| = \log r$ ,  $|K^*| = \text{kl}$ , and  $|\sigma| = S$  all hold. Hence  $(\mathcal{D}_1, \mathcal{D}_2)$  is  $(S + \log r + \text{kl})$ -bounded and the desired bound follows by applying Lemma 2.  $\square$

## E Proof of Lemma 8 (Compression Argument)

For this proof we use the compression lemma given in [6], which we state here as a proposition. It quantitatively formalizes that it is impossible to compress a random element in set  $\mathcal{X}$  to a string that than  $\log |\mathcal{X}|$  bits long, even relative to a random string.

**Proposition 1.** *Let Encode be a randomized map from  $\mathcal{X}$  to  $\mathcal{Y}$  and let Decode be a randomized map from  $\mathcal{Y}$  to  $\mathcal{X}$  such that*

$$\Pr[x \leftarrow \mathcal{X} : \text{Decode}(\text{Encode}(x)) = x] \geq \varepsilon ,$$

*then  $\log |\mathcal{Y}| \geq \log |\mathcal{X}| - \log(1/\varepsilon)$ .*

<pre> Algorithm Encode(<math>(T_1, \dots, T_c)</math>) <math>k \leftarrow 0</math>; <math>\sigma \leftarrow \mathcal{R}^{\text{ENC}}.\text{Init}</math> While <math>k \leq \mathcal{R}.\text{rew}</math> do   For <math>i \in [c]</math> do // <b>Buffer Phase</b>     <math>N_i \leftarrow (r + 1, i)</math>     <math>(\sigma, \text{rf}, C_i) \leftarrow \mathcal{R}^{\text{ENC}}.\text{SimEnc}(\sigma, N_i, 0^{\text{ml}})</math>     If rf then goto NEXT2   For <math>i \in [r]</math> do // <b>Challenge Phase</b>     <math>j^* \leftarrow [u]</math>     For <math>j \in [u]</math> do       <math>M_j \leftarrow \{0, 1\}^{\text{ml}}</math>       <math>N_j \leftarrow (i, j)</math>       <math>(\sigma, \text{rf}, C_j) \leftarrow \mathcal{R}^{\text{ENC}}.\text{SimEnc}(\sigma, N_j, M_j)</math>       If rf then return <math>\sigma</math>     <math>(\sigma, \text{rf}, M) \leftarrow \mathcal{R}^{\text{ENC}}.\text{SimDec}(\sigma, N_{j^*}, C_{j^*})</math>     If rf then return <math>\sigma</math>     If <math>M \neq M_{j^*}</math> then       <math>\sigma \leftarrow \mathcal{R}^{\text{ENC}}.\text{Upd}(\sigma, 1)</math>     Return <math>\sigma</math>   Return <math>\sigma</math>   NEXT: <math>\text{crw} \leftarrow \text{true}</math>   NEXT2: <math>k \leftarrow k + 1</math> Return <math>\sigma</math>  Oracle ENC(<math>N, M</math>) For <math>i \in [c]</math> do   If <math>(N, M) = ((r + 1, i), 0^{\text{ml}})</math> then return <math>T_i</math> <math>C \leftarrow \mathcal{C}</math> Return <math>C</math> </pre>	<pre> Algorithm Decode(<math>\sigma</math>) <math>k \leftarrow 0</math> While <math>k \leq \mathcal{R}.\text{rew}</math> do   For <math>i \in [c]</math> do // <b>Buffer Phase</b>     <math>N_i \leftarrow (r + 1, i)</math>     <math>(\sigma, \text{rf}, C_i^*) \leftarrow \mathcal{R}^{\text{ENC}}.\text{SimEnc}(\sigma, N_i, 0^{\text{ml}})</math>     If rf then goto NEXT2   NEXT2: <math>k \leftarrow k + 1</math> Return <math>(C_1^*, \dots, C_c^*)</math>  Oracle ENC(<math>N, M</math>) <math>C \leftarrow \mathcal{C}</math> Return <math>C</math> </pre>
--	--

**Fig. 23.** Procedures Encode and Decode. Highlighting indicates interesting code changes from  $G^2$ .

To apply this proposition we use the reduction  $\mathcal{R}$  to construct procedures Encode and Decode which attempt to encode an element of  $\mathcal{X} = \mathcal{C}^c$  with a state  $\sigma \in \{0, 1\}^S = \mathcal{Y}$  and succeed whenever  $G^3$  (Fig. 18) would output **true**. They are shown in Fig. 23.

The procedure Encode was obtained by copying the code of  $G^3$ , then making two modifications. First, whenever  $\mathcal{R}$  makes a ENC oracle query corresponding to the  $i^{\text{th}}$  buffer query, Encode responds using its input string  $T_i \in \mathcal{C}$  (the other queries are answered at random as in  $G^3$ ). Second, Encode returns the state  $\sigma$  of  $\mathcal{R}$  whenever  $\text{crw}$  would be set in  $G^3$  (i.e. whenever there was a goto NEXT) or whenever  $G^3$  would halt and return an output.

The procedure Decode uses  $\sigma$  to run  $\mathcal{R}$  and see if it would ever again pass the buffer phase. If  $\mathcal{R}$  does so, Decode does not need to continue simulating its view correctly because it already has the required values. Because Decode does not know how many times  $\mathcal{R}$  has already rewinded, it simply allows  $\mathcal{R}$  to rewind  $\mathcal{R}.\text{rew}$  times. The final output of Decode consists of the strings  $C_1^*, \dots, C_c^*$  that  $\mathcal{R}$  responded to buffer queries with.

We argue that Decode answers correctly if  $G^3$  would return **true**. Since  $G^3$  returns **true** only if  $\text{crw}$  is set, we only look at executions of Encode and Decode where  $\text{crw}$  would have been set in  $G^3$ . Observe that Encode simulates  $G^3$  to  $\mathcal{R}$  perfectly right up until  $\text{crw}$  is set in  $G^3$  (the simulation of ENC is perfect because  $T_1, \dots, T_c$  are distributed uniformly at random in  $\mathcal{C}$ ). Then Decode perfectly simulates  $G^3$  to  $\mathcal{R}$  from the point  $\text{crw}$  is set until  $\mathcal{R}$  passes the buffer phase (or until  $G^3$  would have halted without  $\mathcal{R}$  passing the buffer again). While Encode and Decode may end up running  $\mathcal{R}$  for more than  $\mathcal{R}.\text{rew}$  iterations,  $G^3$  would have returned **false** if it has completed  $\mathcal{R}.\text{rew}$  iterations and  $\mathcal{R}$  had never passed the buffer phase a second time, so we need not worry about incorrectness simulation after the first  $\mathcal{R}.\text{rew}$  iterations.

Games $G_0, G_1$	Oracle $\text{DEC}(C)$	Algorithm $\text{E}(M)$	Algorithm $\text{D}(T \  C)$
$\text{sync} \leftarrow \text{true}$	If $N_d = \perp$ then return $\perp$	$N_e \leftarrow N_e + 1$	$Y \leftarrow \text{pad}(N_d)$
$N_e \leftarrow_s \{0, 1\}^{\text{CAU.nl}}$	$N_d \leftarrow N_d + 1$	$Y \leftarrow \text{pad}(N_e)$	$C_1 \dots C_\ell \leftarrow_n C$
$N_d \leftarrow N_e$	$M' \leftarrow \text{M.dq}()$	$M_1 \dots M_\ell \leftarrow_n M$	$T' \leftarrow H_L(A, C) \oplus E_K(Y)$
$K \leftarrow_s \{0, 1\}^{\text{CAU.kl}}$	$C' \leftarrow \text{C.dq}()$	For $i = 1, \dots, \ell$ do	If $T = T'$ then
$L \leftarrow E_K(0^{\text{CAU.bl}})$	If $\text{sync}$ then	$C_i \leftarrow M_i \oplus E_K(Y + i)$	$\text{bad}_{\text{forge}} \leftarrow \text{true}$
$b' \leftarrow_s \mathcal{A}^{\text{ENC,DEC}}$	If $C = C'$ then	$C \leftarrow C_1 \dots C_\ell$	$N_d \leftarrow \perp$
Return $b' = 1$	Return $M'$	$T \leftarrow H_L(A, C) \oplus E_K(Y)$	Return $\perp$
Oracle $\text{ENC}(M)$	$\text{sync} \leftarrow \text{false}$	Return $((K, N), T \  C)$	For $i = 1, \dots, \ell$ do
$C \leftarrow \text{E}(M)$	$M_1 \leftarrow \text{D}(C)$		$M_i \leftarrow C_i \oplus E_K(Y + i)$
$\text{M.add}(M); \text{C.add}(C)$	$M_0 \leftarrow \perp$		$M \leftarrow M_1 \dots M_\ell$
Return $C$	Return $M_b$		Return $M$
			$N_d \leftarrow \perp$
			Return $\perp$

**Fig. 24.** Games  $G_0$  and  $G_1$  introducing  $\text{bad}_{\text{forge}}$  for proof of Thm. 4. Highlighted code is only included in  $G_1$ .

Now, because  $\mathcal{R}$  is faithful, it must have queried the buffer queries to ENC before passing the buffer phase in Encode and on every rewind, it answers with the same answers it had given previously. Hence, if  $\mathcal{R}$  passes the buffer phase in both Encode and Decode, then it must have responded with  $T_1, \dots, T_c$  in Decode. Consequently, if  $G^3$  would return **true**, then  $T_1, \dots, T_c$  will be output by Decode. So, we have that

$$\Pr[(T_1, \dots, T_c) \leftarrow_s C^c : \text{Decode}(\text{Encode}((T_1, \dots, T_c))) = (T_1, \dots, T_c)] \geq \Pr[G^3].$$

Letting  $\Pr[G^3] = \varepsilon$  and applying Proposition 1 gives  $\log 2^S \geq \log |C|^c - \log 1/\varepsilon$ . Then rearranging the terms gives  $\varepsilon \leq 2^S / |C|^c$ , which concludes the proof.

## F Proof of Thm. 4 (NCH[CAU[E, H]] is CTXT secure)

In this section we provide a proof of Thm. 4 which gave a concrete bound on the CTXT security of NCH[NE] for  $\text{NE} = \text{CAU}[\text{E}, \text{H}]$ .

*Proof.* Consider the game  $G_0$  shown in Fig. 24. It was obtained by plugging the code of NCH[NE] into  $G_{\text{NCH}[\text{NE}],1}^{\text{ae}}$  and making some minor organizational changes (e.g.  $L$  is computed once for all time and encryptions/decryptions are performed in the separate algorithms E and D). The game  $G_1$  is the same except that when a forgery would occur (i.e. when  $\text{bad}_{\text{forge}}$  is set) the oracle D sets the nonce to  $\perp$  and returns  $\perp$  anyway. Equivalently, this game always returns  $\perp$  once  $\text{sync}$  is false. Because of this we have that  $G_1$  is identical to  $G_{\text{NCH}[\text{NE}],0}^{\text{ae}}$ . Thus we have that

$$\text{Adv}_{\text{NCH}[\text{NE}]}^{\text{ch-ctxt}}(\mathcal{A}) = \Pr[G_0] - \Pr[G_1] \leq \Pr[G_1 \text{ sets } \text{bad}_{\text{forge}}].$$

The rest of the proof is dedicated to bounding this probability.

TRANSITIONING TO RANDOM. Recall that CAU's authentication is inspired by a Carter-Wegman MAC. Thus we ultimately expect to reduce security to the AXU security of H. However, to do so we first will need to replace the output of E with random so that the key  $L$  looks random and the encryption queries before the forgery attempt have not revealed anything about  $L$ . Naturally, this consists of a reduction to the PRF security of E, but this is somewhat difficult in our setting. We'd like the reduction to be memory tight and not repeat PRF queries. This makes it difficult to simulate DEC while  $\text{sync} = \text{true}$ .

Ultimately, the approach we will take here closely mirrors that used in our proof of Thm. 2. We will assume that the adversary  $\mathcal{A}$  cannot remember too many bits of ciphertext at a time because the ciphertexts should look random to it. If there are ever more than that many bits of ciphertext that  $\mathcal{A}$  has received from



Games $\boxed{H_0}$ , $H_1$ , $H_2$ , $H_3$	Oracle $\text{DEC}(C)$	Algorithm $\text{E}(M)$
$\text{flag} \leftarrow \text{true}$	$N_d \leftarrow N_d + 1$	$N_e \leftarrow N_e + 1$
$\text{sync} \leftarrow \text{true}$	$M' \leftarrow \mathbf{M}.\text{dq}()$	$Y \leftarrow \text{pad}(N_e)$
$N_e \leftarrow \$_\{0, 1\}^{\text{CAU.nl}}$	$C' \leftarrow \mathbf{C}.\text{dq}()$	$M_1 \dots M_\ell \leftarrow_n M$
$N_d \leftarrow N_e$	$M_2 \leftarrow \mathbf{M}_2.\text{dq}()$	$C_1 \dots C_\ell \leftarrow \$_\{0, 1\}^{ M }$
$K \leftarrow \$_\{0, 1\}^{\text{CAU.kl}}$	$C_2 \leftarrow \mathbf{C}_2.\text{dq}()$	For $i = 1, \dots, \ell - 1$ do
$L \leftarrow \text{RF}(0^{\text{CAU.bl}})$	If $C = C'$ then	$F[Y + i] \leftarrow C_i \oplus M_i$
Run $\mathcal{A}^{\text{ENC,DEC}}$	Return $M'$	$C_i \leftarrow M_i \oplus \text{RF}(Y + i)$
Return <b>false</b>	Elif $C = C_2$ then	$C \leftarrow C_1 \dots C_\ell$
Oracle $\text{ENC}(M)$	<b>bad</b> $\leftarrow$ <b>false</b>	$T \leftarrow \$_\{0, 1\}^{\text{CAU.bl}}$
$C \leftarrow \text{E}(M)$	$\boxed{\text{Return } M_2}$	$F[Y] \leftarrow T \oplus H_L(A, C)$
$\mathbf{M}_2.\text{add}(M); \mathbf{C}_2.\text{add}(C)$	<b>abort</b> ( <b>false</b> )	$T \leftarrow H_L(A, C) \oplus \text{RF}(Y)$
If <b>flag</b> then	$T \parallel C \leftarrow C$	Return $T \parallel C$
If $\ C\  +  C  < \delta$ then	$Y \leftarrow \text{pad}(N_d)$	Oracle $\text{RF}(x)$
$\mathbf{M}.\text{add}(M); \mathbf{C}.\text{add}(C)$	$X \leftarrow H_L(A, C)$	If $F[x] = \perp$ then
Else	If $T = X \oplus \text{RF}(Y)$ then	$F[x] \leftarrow \$_\{0, 1\}^{\text{CAU.bl}}$
<b>flag</b> $\leftarrow$ <b>false</b>	<b>bad</b> <sub>forge</sub> $\leftarrow$ <b>true</b>	$y \leftarrow F[x]$
Return $C$	<b>abort</b> ( <b>true</b> )	$y \leftarrow \mathbf{E}_K(x)$
	<b>abort</b> ( <b>false</b> )	Return $y$

**Fig. 25.** Games used to analyze  $\text{bad}_{\text{forge}}$  for proof of Thm. 4. Boxed or highlighted code is only included in correspondingly indicated game(s).

ENC, but not forwarded to DEC then we can stop remembering additional ciphertext bits in our table and assume the sender and receiver will be out of sync before  $\mathcal{A}$  has forwarded on all the ciphertext bits.

To start this analysis consider the game  $H_0$  shown in Fig. 25 for which we claim  $\Pr[G_1 \text{ sets } \text{bad}_{\text{forge}}] = \Pr[H_0]$ . Note that  $H_0$  returns  $\text{bad}_{\text{forge}}$  rather than something depending on the bit guessed by  $\mathcal{A}$ . To obtain  $H_0$  from  $G_1$  we have made a few syntactic changes that do not effect the probability  $\text{bad}_{\text{forge}}$  is set. First, the use of tables surrounding **sync** has been re-organized to use two copies of the table as was done in proving Thm. 2. The flag **sync** has been removed because we will simply abort whenever it would be changed. Next, all executions of  $\mathbf{E}_K$  has been forwarded to an oracle  $\text{RF}$  which simply runs  $\text{E}$ . Finally, we have added code to  $\text{DEC}$  to make the game abort and output  $\text{bad}_{\text{forge}}$  rather than return  $\perp$  and continue execution (after a failed attempt to set  $\text{bad}_{\text{forge}}$  in  $G_1$  the nonce  $N_d$  would be set to  $\perp$ , preventing  $\text{bad}_{\text{forge}}$  from ever being set again in the future).

The game  $H_1$  is the analog of  $G_3$  in the proof of Thm. 2. It differs from  $H_0$  by never using  $\mathbf{M}_2$  to determine the output of  $\text{DEC}$ , thereby limiting how much memory is needed to simulate it. Games  $H_0$  and  $H_1$  are identical-until-bad so we have that

$$\Pr[H_0] - \Pr[H_1] \leq \Pr[H_1 \text{ sets bad}].$$

PRF TRANSITION. Having moved to a game that can be simulated with low memory we can now apply PRF security to replace the output of  $\text{E}$  with random. Consider the game  $H_2$ . It differs from  $H_1$  in that the highlighted code has been removed from  $\text{RF}$  which now implements a (lazily sampled) random function. In Fig. 26, we give PRF adversaries  $\mathcal{A}'_{\text{prf}}$  and  $\mathcal{A}''_{\text{prf}}$  satisfying

$$\Pr[H_1] - \Pr[H_2] \leq \text{Adv}_{\text{E}}^{\text{prf}}(\mathcal{A}'_{\text{prf}}) \text{ and } \Pr[H_1 \text{ sets bad}] - \Pr[H_2 \text{ sets bad}] \leq \text{Adv}_{\text{E}}^{\text{prf}}(\mathcal{A}''_{\text{prf}}).$$

These adversaries simulate the view of  $\mathcal{A}$  in  $H_1$  or  $H_2$  by using its  $\text{EVAL}$  oracle whenever  $\text{E}$  would be executed, the first outputting 1 whenever  $\text{bad}_{\text{forge}}$  would be set and the second outputting 1 whenever **bad** would be set. They store one extra ciphertext  $C^*$  beyond what  $\mathbf{C}$  stores to know when **bad** occurs. We want to be careful that neither adversary repeats queries to  $\text{EVAL}$ . This is immediate for  $\mathcal{A}''$ . When  $\mathcal{A}'$  checks if the tag of a

<p>Adversaries <math>\mathcal{A}_{\text{prf}}^{\text{EVAL}}, \mathcal{A}_{\text{prf}}^{\text{EVAL}}</math></p> <p><b>flag</b> <math>\leftarrow</math> <b>true</b></p> <p><math>N_e \leftarrow_{\\$} \{0, 1\}^{\text{CAU.nl}}</math></p> <p><math>N_d \leftarrow N_e</math></p> <p><math>L \leftarrow \text{EVAL}(0^{\text{CAU.bl}})</math></p> <p>Run <math>\mathcal{A}^{\text{SIMENC, SIMDEC}}</math></p> <p>Return 0</p> <hr/> <p>Algorithm <math>\text{E}(M)</math></p> <p><math>N_e \leftarrow N_e + 1</math></p> <p><math>Y \leftarrow \text{pad}(N_e)</math></p> <p><math>M_1 \dots M_\ell \leftarrow_n M</math></p> <p>For <math>i = 1, \dots, \ell - 1</math> do</p> <p style="padding-left: 2em;"><math>Z_i \leftarrow \text{EVAL}(Y + i)</math></p> <p style="padding-left: 2em;"><math>C_i \leftarrow M_i \oplus Z_i</math></p> <p><math>C \leftarrow C_1 \dots C_\ell</math></p> <p><math>X \leftarrow H_L(A, C)</math></p> <p><math>T \leftarrow X \oplus \text{EVAL}(Y)</math></p> <p>Return <math>T \parallel C</math></p>	<p>Oracle <math>\text{SIMENC}(M)</math></p> <p><math>C \leftarrow \text{E}(M)</math></p> <p>If <b>flag</b> then</p> <p style="padding-left: 2em;">If <math>\ C\  +  C  &lt; \delta</math> then</p> <p style="padding-left: 4em;"><b>M.add</b>(<math>M</math>)</p> <p style="padding-left: 4em;"><b>C.add</b>(<math>C</math>)</p> <p style="padding-left: 2em;">Else</p> <p style="padding-left: 4em;"><b>C.add</b>(<math>*</math>)</p> <p style="padding-left: 4em;"><math>C^* \leftarrow C</math></p> <p style="padding-left: 4em;"><b>flag</b> <math>\leftarrow</math> <b>false</b></p> <p>Return <math>C</math></p>	<p>Oracle <math>\text{SIMDEC}(C)</math></p> <p><math>N_d \leftarrow N_d + 1</math></p> <p><math>M' \leftarrow \text{M.dq}(); C' \leftarrow \text{C.dq}()</math></p> <p>If <math>C = C'</math> then</p> <p style="padding-left: 2em;">Return <math>M'</math></p> <p>Elif <math>C' = *</math> and <math>C = C^*</math> then</p> <p style="padding-left: 2em;"><b>abort</b>(1)</p> <p style="padding-left: 2em;"><b>abort</b>(0)</p> <p style="padding-left: 2em;"><b>abort</b>(0)</p> <p><math>Y \leftarrow \text{pad}(N_d)</math></p> <p><math>T \parallel C \leftarrow C</math></p> <p><math>X \leftarrow H_L(A, C)</math></p> <p>If <math>C' = *</math> then <math>C' \leftarrow C^*</math></p> <p>If <math>C' \neq \perp</math> then</p> <p style="padding-left: 2em;"><math>T' \parallel C' \leftarrow C'</math></p> <p style="padding-left: 2em;"><math>X' \leftarrow H_L(A, C')</math></p> <p style="padding-left: 2em;"><math>Z \leftarrow T' \oplus X'</math></p> <p>Else</p> <p style="padding-left: 2em;"><math>Z \leftarrow \text{EVAL}(Y)</math></p> <p>If <math>T = X \oplus Z</math> then</p> <p style="padding-left: 2em;"><b>abort</b>(1)</p> <p style="padding-left: 2em;"><b>abort</b>(0)</p>
---	--	--

**Fig. 26.** PRF adversaries for proof of Thm. 4. Highlighted code is only used by the highlighted adversary.

ciphertext is correct in SIMDEC we have to be careful not to create a repeat query. This would be possible if  $N_e \geq N_d$ , but in that case we have the corresponding ciphertext stored from which we can re-derive the appropriate output of EVAL (referred to as  $Z$ ).

REWRITING TRANSITION. Our next transition (to  $H_3$ ) is merely a conceptual change which does not effect the behavior at the game. This gives

$$\Pr[H_2] - \Pr[H_3] = 0 \text{ and } \Pr[H_2 \text{ sets bad}] - \Pr[H_3 \text{ sets bad}] = 0.$$

The change occurs solely in E where, rather than creating the output ciphertext by xor-ing things with entries of the table  $F$  (via RF), we instead pick these outputs at random and use them to define  $F$  to match. Note that this results in a uniform distribution for  $F$  either way as long as  $F$  had not already been defined at that point. This is necessarily the case; E never repeats queries to RF because  $N_e$  is always always increasing and if DEC makes a query to RF, then the game will abort immediately afterwards.

BOUNDING PROBABILITY OF **bad**. At this point we are in a game where the adversary is given ciphertexts which are picked uniformly at random. To set **bad** is must at some point remember more than  $\delta$  bits of ciphertext that it delayed sending to DEC. In the exact same way we saw in the proof of Thm. 2 we can create an  $(\mathcal{A}_1, \mathcal{A}_2)$  that succeeds in  $G_{q,x,\delta}^{\text{it}}$  with the same probability that **bad** is set, giving

$$\Pr[H_3 \text{ sets bad}] \leq \text{Adv}_{q,x,\delta}^{\text{it}}(\mathcal{A}_1, \mathcal{A}_2).$$

This advantage is then bounded by  $q \cdot x \cdot 2^S / 2^\delta$  using Lemma 1.

FINAL BOUND. Finally we can conclude the proof by analyzing  $\Pr[H_3]$ . This reduces naturally to the AXU security of the hash function H. Note that ENC returns random bits, so when  $\mathcal{A}$  makes its one attempt to set **bad**<sub>forge</sub> its view is independent of  $L$ . Assuming  $N_e \geq N_d$  (the  $N_e < N_d$  case is strictly harder) at the time of this attempt, the adversary has seen a random tag  $T$  and random  $C$  from which  $F[\text{pad}(N_d)]$  was defined to equal  $T \oplus H_L(A, C)$ . Letting  $T^* \parallel C$  represent its final query, it will win if  $T^* = H_L(A, C^*) \oplus F[\text{pad}(N_d)]$ .

Adversary $\mathcal{A}$	Oracle $\text{SIMDEC}(C)$	Algorithm $\text{E}(M)$
$\text{flag} \leftarrow \text{true}$	$N_d \leftarrow N_d + 1$	$N_e \leftarrow N_e + 1$
$N_e \leftarrow \$_\{0, 1\}^{\text{CAU.nl}}$	$M' \leftarrow \mathbf{M.dq}(); C' \leftarrow \mathbf{C.dq}()$	$C \leftarrow \$_\{0, 1\}^{ M }$
$N_d \leftarrow N_e$	$M_2 \leftarrow \mathbf{M}_2.\text{dq}(); C_2 \leftarrow \mathbf{C}_2.\text{dq}()$	$T \leftarrow \$_\{0, 1\}^{\text{CAU.bl}}$
Run $\mathcal{A}^{\text{SIMENC}, \text{SIMDEC}}$	If $C = C'$ then return $M'$	Return $T \parallel C$
Return $\perp$	Elif $C = C_2$ then <b>abort</b> ( $\perp$ )	
Oracle $\text{SIMENC}(M)$	$T \parallel C \leftarrow C$	
$C \leftarrow \text{E}(M)$	If $C_2 \neq \perp$ then	
$\mathbf{M}_2.\text{add}(M); \mathbf{C}_2.\text{add}(C)$	$T_2 \parallel C_2 \leftarrow C_2$	
If <b>flag</b> then	Else	
If $\ \mathbf{C}\  +  C  < \delta$ then	$T_2 \leftarrow \$_\{0, 1\}^{\text{CAU.bl}}$	
$\mathbf{M}.\text{add}(M); \mathbf{C}.\text{add}(C)$	$C_2 \leftarrow C \oplus 1^{ C }$	
Else	$x \leftarrow (A, C)$	
<b>flag</b> $\leftarrow \text{false}$	$x_2 \leftarrow (A, C_2)$	
Return $C$	<b>abort</b> ( $x_1, x_2, T \oplus T_2$ )	

**Fig. 27.** Adversary against the AXU security of H.

Plugging in for  $F$  and re-arranging gives  $T^* \oplus T = H_L(A, C^*) \oplus H_L(A, C)$ . The adversary  $\mathcal{A}$  shown in Fig. 27 captures this idea. It simulates the view of  $\mathcal{A}$  with random strings until a query is made to DEC that has the potential of setting  $\text{bad}_{\text{forge}}$ , in which case it provides (In the case that  $N_d > N_e$  it just picks  $T_2$  at random and sets  $C_2$  to anything other than  $C$ . This gives the correct probability distribution because  $\mathcal{A}$ 's view would have been independent of  $F[\text{pad}(N_d)]$  at that point and it would have needed to guess  $H_L(A, C^*) \oplus F[\text{pad}(N_d)]$  to set the flag.)

The above reasoning tells us that

$$\Pr[\text{H}_3] = \text{Adv}_{\text{H}}^{\text{axu}}(\mathcal{A}).$$

Completing the proof is then a simply matter of combining all of the equations we have established. The PRF adversary  $\mathcal{A}$  is defined to randomly choose either  $\mathcal{A}'_{\text{prf}}$  and  $\mathcal{A}''_{\text{prf}}$  (each with probability 1/2) and then run the adversary it chose, outputting whatever that adversary does.  $\square$