



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, USA
{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 [6, 10, 16]

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.

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 [6, 10] can be combined to analyze the INDR security¹ of nonce-based encryption. For example, consider a toy scheme² 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 [15], 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 $b = 1$, ENC_1 returns real ciphertexts, and DEC_1 decrypts properly. For $b = 0$, instead, ENC_0 returns random ciphertexts, and DEC_0 decrypts only previous outputs from ENC_0 . It is

¹ Which measures the indistinguishability of ciphertexts from truly random ones.

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

important to use a combined definition, as it captures settings such as chosen-ciphertext attacks and padding-oracle attacks [17], which use a decryption oracle to break confidentiality.³

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 [15]) 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 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 Sect. 3.

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.

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

- (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 [11].

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* 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* DEC_b . If $b = 1$, this just applies the decryption algorithm of the AE scheme, using increasing nonces. If decryption fails, 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* ENC_0 . If the ciphertexts come in the wrong order, 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 ENC_1 with ENC_0 . The catch here is that when doing so, we need to simulate the 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 δ oldest ciphertexts which have not been delivered yet, for some δ – the key point here is that δ 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 [6]. 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, 8, 18] 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 [8, 18], 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 [6, 10]). 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 Sect. 2, we introduce our notation, basic definitions and cover some cryptographic background necessary for the paper. In Sect. 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 Sect. 4, we show that memory-tight reductions can be given for the combined confidentiality and integrity security of cryptographic channels. Using the result from Sect. 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 Sect. 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 CTEXT security. This establishes that our move to the channel setting for Sect. 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 ℓ blocks of length n (except x_ℓ 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 ε .

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 $y \leftarrow_s \mathcal{A}^{O_1, \dots}(x_1, \dots)$. 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 $\Pr[G]$ denote the probability that a game 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 G^{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 G^{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 $\text{Adv}_E^x(\mathcal{A}) = \Pr[G_{E,1}^x(\mathcal{A})] - \Pr[G_{E,0}^x(\mathcal{A})]$.

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 $G_{D,b}^{\text{sl}}$ shown in Fig. 1. In it, the adversary is given oracle access to either a random function or a random

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

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

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

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

permutation with domain/range $[D]$ and is trying to figure out which. We define $\text{Adv}_D^{\text{sl}}(\mathcal{A}) = \Pr[G_{D,b}^{\text{sl}}(\mathcal{A})] - \Pr[G_{D,b}^{\text{sl}}(\mathcal{A})]$.

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 ρ -method [12, 13]) 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.

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 [10] showed that $\text{Adv}_D^{\text{sl}}(q, S) \leq \sqrt{Sq/D}$ under a combinatorial conjecture. Later, Dinur [6] 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_H^{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 $\text{Adv}_H^{\text{axu}}(\mathcal{X}) = \Pr[G_H^{\text{axu}}(\mathcal{X})]$. Typically one makes use of a c -AXU hash which for all \mathcal{X} satisfy $\text{Adv}_H^{\text{axu}}(\mathcal{X}) \leq c \cdot (N_1 + N_2)/2^n$ where N_1 (resp. N_2) is the maximum

block length of any A (resp. C) output by \mathcal{A} . 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 [14]. 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 Sect. 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 Sect. 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 Sect. 5 shows it is not possible for general nonce-based encryption if the reduction behaves in a black-box manner.

Finally, in Sect. 3.3 we recall the CAU scheme by Bellare and Tackmann [4], an abstraction of GCM [11]. Following existing proofs [4, 9, 11] and using [6, 10], 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 NE is defined by algorithms NE.Kg , NE.D , and NE.E . Additionally it is associated with message space $\text{NE.M} \subseteq \{0, 1\}^*$ and nonce space NE.N .

The syntax of the algorithms is shown in Fig. 3. The key generation algorithm NE.Kg takes no input and returns key K . The encryption algorithm 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 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.

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

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

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 $\text{NE.D}(K, N, \text{NE.E}(K, N, M)) = M$.

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 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.

<p>Game $G_{NE,b}^{indr}(\mathcal{A})$</p> <p>$K \leftarrow \text{\\$ NE.Kg}$</p> <p>$b' \leftarrow \mathcal{A}^{ENC_b}$</p> <p>Return b'</p> <hr style="border: 0.5px solid black;"/> <p>Oracle $ENC_b(N, M)$</p> <p>$C_1 \leftarrow \text{NE.E}(K, N, M)$</p> <p>$C_0 \leftarrow \text{\\$ } \{0, 1\}^{NE.cl(M)}$</p> <p>$M[N, C_b] \leftarrow M$</p> <p>Return C_b</p>	<p>Game $G_{NE,b}^{ctxt-w}(\mathcal{A})$</p> <p>$K \leftarrow \text{\\$ NE.Kg}$</p> <p>$b' \leftarrow \mathcal{A}^{ENC_1, DEC_b^w}$</p> <p>Return b'</p> <hr style="border: 0.5px solid black;"/> <p>Game $G_{NE,b}^{ae-w}(\mathcal{A})$</p> <p>$K \leftarrow \text{\\$ NE.Kg}$</p> <p>$b' \leftarrow \mathcal{A}^{ENC_b, DEC_b^w}$</p> <p>Return b'</p>	<p>Oracle $DEC_b^w(N, C)$</p> <p>If $M[N, C] \neq \perp$ then</p> <p style="padding-left: 20px;">Return $M[N, C]$ if $w = 1$</p> <p style="padding-left: 20px;">Return \diamond if $w = 2$</p> <p style="padding-left: 20px;">Return \perp if $w = 3$</p> <p>$M_1 \leftarrow \text{NE.D}(K, N, C)$</p> <p>$M_0 \leftarrow \perp$</p> <p>Return M_b</p>
--	--	---

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 [6,10] 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 [10] 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 Sect. 4). Nonces may repeat in the latter case, but we can follow [6,10] 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 definition 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 $G_{NE,b}^{ae-w}$ 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 $G_{NE,b}^{ctxt-w}$ 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 $ENC(N, M)$ and then queries $DEC(N, C)$ 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 $x \in \{ae, ctxt\}$ and $w \in \{1, 2, 3\}$ we define the advantage of an adversary \mathcal{A} by $Adv_{NE}^{x-w}(\mathcal{A}) = \Pr[G_{NE,1}^{x-w}(\mathcal{A})] - \Pr[G_{NE,0}^{x-w}(\mathcal{A})]$. 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.⁴ 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 DEC_b^2 can simulate DEC_b^3 with low memory. If 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

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

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 [15]. 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 DEC_0^1 which seems to require storing the table $M[\cdot, \cdot]$.⁵ 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 Sect. 4.2, we make it memory tight when restricting attention to secure channels which only accept ciphertexts if they are received in order. In Sect. 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 Sect. 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) [11]. 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 [10] we obtain a proof of security for its nonce-respecting INDR security with an appealing time-memory tradeoff.

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 $N \parallel |0^{n-\text{CAU}[E, H].nl-1} \parallel 1$.

⁵ The standard reduction *would* be memory tight for $w = 3$.

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

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

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 ℓ 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 $T \parallel C$ is the full ciphertext output by encryption.

The decryption algorithm parses the input ciphertext as $T \parallel C$. 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 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}_{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}_{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}_{prf} never repeats queries because it allows us to apply the time-memory switching lemma from Sect. 2. This give us roughly,

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

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}_{prf} from repeating queries without storing the prior queries of \mathcal{A} .

Proof (Sketch). One constructs \mathcal{A}_{prf} to first set $L \leftarrow \text{EVAL}(0^n)$. Then it runs \mathcal{A} and simulates encryption queries by running 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}_{prf} is “real” encryptions when $b = 1$ and random strings when $b = 0$, so the claimed advantage bound follows. \square

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 $\text{H}_L(\mathcal{A}, C)$ 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}_{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 [6, 10] to get appealing time-memory tradeoffs when E is a PRP.

In Sect. 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 Sect. 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 σ^s and receiver state σ^r . The sending algorithm CH.S takes a sender state σ^s and message $M \in \text{CH.M}$. It returns updated sender state σ^s and a ciphertext C . The receiving algorithm CH.R takes a receiver state σ^r and a ciphertext C . It returns updated receiver state σ^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

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

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

to be run over a reliable transport layer and has been the standard notion for channels since the work of Bellare, Kohno, and Namprempre [2]. When a protocol (e.g. QUIC or DTLS) is run over an unreliable transport layer, then a *robust* channel is used instead [7]. 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 σ^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 .

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.⁶ 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 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 ENC_1 and

⁶ 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.

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

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

Game $G_{L,\delta}^{\text{it}}(\mathcal{A}_1, \mathcal{A}_2)$ $R \leftarrow_{\$} \{0, 1\}^L$ $(i, \sigma) \leftarrow_{\$} \mathcal{A}_1(R)$ $r \leftarrow_{\$} \mathcal{A}_2(i, \sigma, R[i : i-1])$ Return $r = R[i : i + \delta]$

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

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 δ -bits of the string using state σ . 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*

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

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 σ , 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.

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

The last inequality follows from there being at most $L \cdot 2^\delta$ choices for (i, σ) . \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}_δ 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 δ which determines how much local memory it uses. Using Lemma 1, our concrete advantage bound is expressed in terms of δ 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}_δ (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}_δ 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}_δ 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 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}_δ may use much more memory than \mathcal{A} .

Our proof of Theorem 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 see 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 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}_δ 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}_δ will store, then \mathcal{A}_δ 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}'_δ , \mathcal{A}''_δ , and S -bounded adversary $(\mathcal{A}_1, \mathcal{A}_2)$ and show that

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

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

CTXT TRANSITION. Let $\text{G}_0 = \text{G}_{\text{CH},1}^{\text{ch-ae}}(\mathcal{A})$ and $\text{G}_1 = \text{G}_{\text{CH},0}^{\text{ch-ctxt}}(\mathcal{A})$. Because $\text{G}_{\text{CH},1}^{\text{ch-ae}}(\mathcal{A})$ and $\text{G}_{\text{CH},1}^{\text{ch-ctxt}}(\mathcal{A})$ are identical games we have that $\Pr[\text{G}_0] - \Pr[\text{G}_1] = \text{Adv}_{\text{CH}}^{\text{ch-ctxt}}(\mathcal{A})$.

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 $\Pr[\text{G}_1] = \Pr[\text{G}_2]$.

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 δ 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,

$$\Pr[\text{G}_3] - \Pr[\text{G}_2] \leq \Pr[\text{G}_3 \text{ sets bad}].$$

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

Games $\boxed{\mathbb{G}_2}$, \mathbb{G}_3 , \mathbb{G}_4 , $\boxed{\mathbb{G}_5}$	Oracle ENC(M)	Oracle DEC(C)
$\text{flag} \leftarrow \text{true}$ $\text{sync} \leftarrow \text{true}$ $(\sigma^s, \sigma^r) \leftarrow \text{CH.Sg}$ $b' \leftarrow \text{A}^{\text{ENC,DEC}}$ Return $b' = 1$	$(\sigma^s, C) \leftarrow \text{CH.S}(\sigma^s, M)$ $C \leftarrow \{0, 1\}^{\text{CH.cl}(M)}$ $\text{M}_2.\text{add}(M); \text{C}_2.\text{add}(C)$ If flag then If $\ \text{C}\ + C < \delta$ then $\text{M}.\text{add}(M); \text{C}.\text{add}(C)$ Else $\text{flag} \leftarrow \text{false}$ Return C	$M' \leftarrow \text{M}.\text{dq}(); M_2 \leftarrow \text{M}_2.\text{dq}()$ $C' \leftarrow \text{C}.\text{dq}(); C_2 \leftarrow \text{C}_2.\text{dq}()$ If sync then If $C = C'$ then Return M' Elif $C = C_2$ then $\text{bad} \leftarrow \text{true}$ $\boxed{\text{Return } M_2}$ $\text{sync} \leftarrow \text{false}$ Return \perp

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.

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 \mathbb{G}_3 .

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

Adversary \mathcal{A}'_δ uses its ENC oracle to present \mathcal{A} with a view identical to \mathbb{G}_3 if $b = 1$ and identical to \mathbb{G}_4 if $b = 0$. Note here that the tables (M_2, C_2) do not effect the view of \mathcal{A} in either of these game, allowing \mathcal{A}_δ not to have to store them. We have that $\Pr[\text{G}_{\text{CH},1}^{\text{ch-indr}}(\mathcal{A}'_\delta)] = \Pr[\mathbb{G}_3]$ and $\Pr[\text{G}_{\text{CH},0}^{\text{ch-indr}}(\mathcal{A}'_\delta)] = \Pr[\mathbb{G}_4]$. In other words, $\text{Adv}_{\text{CH}}^{\text{ch-indr}}(\mathcal{A}'_\delta) = \Pr[\mathbb{G}_3] - \Pr[\mathbb{G}_4]$.

Adversary \mathcal{A}''_δ 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 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 C so that in DEC we know when we have reached the relevant point. We have that $\Pr[\text{G}_{\text{CH},1}^{\text{ch-indr}}(\mathcal{A}''_\delta)] = \Pr[\mathbb{G}_3 \text{ sets bad}]$ and $\Pr[\text{G}_{\text{CH},0}^{\text{ch-indr}}(\mathcal{A}''_\delta)] = \Pr[\mathbb{G}_4 \text{ sets bad}]$. In other words, $\Pr[\mathbb{G}_3 \text{ sets bad}] \leq \text{Adv}_{\text{CH}}^{\text{ch-indr}}(\mathcal{A}''_\delta) + \Pr[\mathbb{G}_4 \text{ sets bad}]$.

FINAL TRANSITION. The final transition is from \mathbb{G}_4 to \mathbb{G}_5 . These two games are identical-until-bad as can be seen in DEC. Because of this we have that

$$\Pr[\mathbb{G}_4] - \Pr[\mathbb{G}_5] \leq \Pr[\mathbb{G}_4 \text{ sets bad}].$$

Using all of M_2 and C_2 instead of just M and C makes \mathbb{G}_5 identical to $\text{G}_{\text{CH},0}^{\text{ch-ae}}$.

BOUNDING PROBABILITY OF bad . We conclude by bounding the probability \mathbb{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 \mathbb{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 σ with i chosen so the next δ bits of C and c are the values of R for \mathcal{A}_2 to guess.

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

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

Then \mathcal{A}_2 will resume executing \mathcal{A} using σ . 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 \mathbf{C} (otherwise \mathcal{A}_2 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 δ bits long, \mathcal{A}_2 outputs that as its guess.

We can see that when `bad` would be set in \mathbf{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 $\Pr[\mathbf{G}_4 \text{ sets bad}] \leq \text{Adv}_{q,x,\delta}^{\text{it}}(\mathcal{A}_1, \mathcal{A}_2)$ 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 [10] it is easy to give such examples for INDR security.

Using the ideas from proof of Theorem 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

$\text{NCH}[\text{NE}].\text{Sg}$ $K \leftarrow_s \text{NE.Kg}$ $N \leftarrow_s \text{NE.N}$ Return $((K, N), (K, N))$	$\text{NCH}[\text{NE}].\text{S}((K, N), M)$ $N \leftarrow N + 1$ $C \leftarrow \text{NE.E}(K, N, M)$ Return $((K, N), C)$	$\text{NCH}[\text{NE}].\text{R}((K, N), C)$ 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 $\text{NCH}[\text{NE}]$ constructed from encryption scheme NE .

the integrity of 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 $\text{NCH}[\text{NE}]$ is parameterized by an encryption scheme NE . It has $\text{NCH}[\text{NE}].M = \text{NE.M}$. We assume that 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 NE to encrypt the message with the current nonce. Decryption increments the nonce and uses 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 $\text{NCH}[\text{NE}]$ follows easily from nonce-respecting INDR security of NE . This is captured by the following theorem.

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

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

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 $\text{NCH}[\text{NE}]$ obtained when using the encryption scheme $\text{NE} = \text{CAU}[\text{E}, \text{H}]$ for some function families E and 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 CAU uses a Carter-Wegman style MAC we have to first use the PRF

security of 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 Sect. 4 where it can be difficult to simulate the values returned by DEC . This issue is resolved by adjusting the proof technique used to establish Theorem 2 where we exploit the fact that ciphertexts look random to assume that \mathcal{A} cannot remember too many ciphertexts.

Theorem 4. *Let $NE = CAU[E, H]$ for some E and H . Let \mathcal{A} be a nonce-respecting adversary against the CTXT security of $NCH[NE]$ with memory complexity S that makes at most $q \leq 2^{CAU.nl} - 1$ 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}_{prf} such that*

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

Adversary \mathcal{A}_{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.

The proof is given in the full version. As with Theorem 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}_{prf} is not much more than that of \mathcal{A} . Assuming $|A| < p$, setting $\delta \approx S + n$, and assuming $S > n$ we can combine all of our theorems so far to obtain a bound of

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

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 Sect. 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 σ 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 rf , and a ciphertext. If the flag 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

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

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 $\text{N} = \text{NE.N}$ as shorthand. We assume that $[\text{NE.Kg}] = \{0, 1\}^{\text{kl}}$. We let $\text{C} = \{0, 1\}^{\text{NE.cl(ml)}}$. 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.

Game $\text{G}_{u,m}^{\text{it-chl-}r}(\mathcal{D}_1, \mathcal{D}_2)$
$\sigma \leftarrow \perp$
For $i \in [r]$ do
For $j \in [u]$ do
$M_j \leftarrow \{0, 1\}^m$
$\sigma \leftarrow \mathcal{D}_1(\sigma, M_1, \dots, M_u)$
$j^* \leftarrow [u]$
$(\sigma, M) \leftarrow \mathcal{D}_2(\sigma, j^*)$
If $M_{j^*} \neq M$ then return false
Return true

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

INFORMATION THEORETIC LEMMA. We give a lemma that will be a useful sub-component of our proofs. It pertains to game $\text{G}_{u,m}^{\text{it-chl-}r}$ 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 σ from the prior round, along with u random strings M_1, \dots, M_u each of length m . Adversary \mathcal{D}_1 outputs state σ 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 σ 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 $\text{Adv}_{u,m}^{\text{it-chl-}r}(\mathcal{D}_1, \mathcal{D}_2) = \Pr[\text{G}_{u,m}^{\text{it-chl-}r}(\mathcal{D}_1, \mathcal{D}_2)]$. 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.

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, deferred to the full version, goes via a reduction to the $r = 1$ case which is analyzed using techniques from the AI-ROM setting [5].

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. 15. 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) $\text{Adv}_{\text{NE}}^{\text{ae-1}}(\mathcal{A}) \geq 1 - \frac{2^{\text{kl}}}{2^{2\text{NE.cl}(\text{ml})}}$,
- (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 $\text{Mem}(\mathcal{A}) \leq 2\text{kl} + 2\text{ml} + 2\text{NE.cl}(\text{ml}) + 2 \log |\mathbb{N}| + \log u \cdot r$. 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 $\text{Time}(\mathcal{W}) \in O(\text{Query}(\mathcal{A}))$.

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

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

Fig. 14. Games G^1 and G_b^0 for $b \in \{0, 1\}$. Highlighted code is only included in G^1 .

Adversaries $\mathcal{A}^{\text{ENC,DEC}}, \mathcal{S}^{\text{ENC,DEC}}, \mathcal{W}^{\text{ENC,DEC}}$	Adversaries $\mathcal{B}^{\text{ENC}}, \mathcal{E}^{\text{ENC}}$
For $i \in [r]$ do	$M_1 \leftarrow \{0, 1\}^{\text{ml}}$
$j^* \leftarrow [u]$	$M_2 \leftarrow \{0, 1\}^{\text{ml}}$
For $j \in [u]$ do	$N_1 \leftarrow (0, 1); N_2 \leftarrow (0, 2)$
$M_j \leftarrow \{0, 1\}^{\text{ml}}$	$C_1 \leftarrow \text{ENC}(N_1, M_1)$
$N_j \leftarrow (i, j)$	$C_2 \leftarrow \text{ENC}(N_2, M_2)$
$C_j \leftarrow \text{ENC}(N_j, M_j)$	For $K \in \{0, 1\}^{\text{kl}}$ do
$M \leftarrow \text{DEC}(N_{j^*}, C_{j^*})$	$\text{eq}_1 \leftarrow (\mathcal{NE}.\text{E}(K, N_1, M_1) = C_1)$
If $M \neq M_{j^*}$ then return 1	$\text{eq}_2 \leftarrow (\mathcal{NE}.\text{E}(K, N_2, M_2) = C_2)$
bad \leftarrow true	If eq_1 and eq_2 then return 1
Return \mathcal{B}^{ENC} ; Return 0 ; Return \mathcal{E}^{ENC}	Return 0; Return 1

Fig. 15. 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} .

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. 15 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

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

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 the full version.

For the third claim, consider adversary \mathcal{W} defined as shown in Fig. 15. 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. 15. It behaves identically to \mathcal{A} until the flag `bad` is set. Using the Fundamental Lemma of Game Playing [3], we can obtain for $b \in \{0, 1\}$

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

Consider the games G_b^0 for $b \in \{0, 1\}$ in Fig. 14. In it, we assume that \mathcal{R} always outputs `rf = false` since it is straightline. Note that G_b^0 simulates the challenge phase of \mathcal{A} and the game G_b^{indr} to \mathcal{R} perfectly, so it returns `true` whenever $\mathcal{R}[\mathcal{A}]$ would set `bad` is set in G_b^{indr} . From this we can show

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

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^{ml} . We now state two lemmas which give bounds on both $\Pr [\text{G}_b^0]$'s via G^1 . First, in Lemma 3, we use that the INDR security of NE implies G^1 's encryption oracle is indistinguishable from those in either G_b^0 to transition to G^1 . Next, in Lemma 4 we give a bound on $\Pr [\text{G}^1]$ 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 the full version.

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

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

where G_b^0 and G^1 are defined as in Fig. 14. 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,*

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

where G^1 is defined as in Fig. 14.

Applying these lemmas to Eq. 3 gives

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

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 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 . \square

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 Theorem 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) $\text{Adv}_{\text{NE}}^{\text{ae-1}}(\mathcal{A}) \geq 1 - \frac{2^{\text{kl}}}{2^{2\text{NE.cl}(\text{ml})}}$,
- (ii) $\text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{R}[\mathcal{A}]) \leq 2 \cdot \left(\frac{2(S + \log r + \text{kl}) + 2\text{ml}}{u} + \frac{3}{2^{\text{ml}}} \right)^r + \frac{2^{S+1}}{2^{c \cdot \text{NE.cl}(\text{ml})}} + 6 \cdot \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{C})$,
- (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 $\text{Mem}(\mathcal{A}) \leq 2kl + 2ml + 2\text{NE.cl}(ml) + 2 \log |\mathbf{N}| + \log u \cdot r$. 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 $\text{Time}(\mathcal{W}) \in O(\text{Query}(\mathcal{A}))$.

For interests of space, the proof of this result has been deferred to the full version. 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 Theorem 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^{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 [7], 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 (CAREER), and by a Sloan Research Fellowship.

References

1. Auerbach, B., Cash, D., Fersch, M., Kiltz, E.: Memory-tight reductions. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part 1. LNCS, vol. 10401, pp. 101–132. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63688-7_4
2. Bellare, M., Kohno, T., Namprempre, C.: Breaking and provably repairing the ssh authenticated encryption scheme: a case study of the encode-then-encrypt-and-mac paradigm. ACM Trans. Inf. Syst. Secur. (TISSEC) 7(2), 206–241 (2004)
3. Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 409–426. Springer, Heidelberg (2006). https://doi.org/10.1007/11761679_25

4. Bellare, M., Tackmann, B.: The multi-user security of authenticated encryption: AES-GCM in TLS 1.3. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part I. LNCS, vol. 9814, pp. 247–276. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53018-4_10
5. Coretti, S., Dodis, Y., Guo, S., Steinberger, J.: Random oracles and non-uniformity. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part I. LNCS, vol. 10820, pp. 227–258. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78381-9_9
6. Dinur, I.: On the streaming indistinguishability of a random permutation and a random function. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part II. LNCS, vol. 12106, pp. 433–460. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45724-2_15
7. Fischlin, M., Günther, F., Janson, C.: 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>
8. Ghoshal, A., Tessaro, S.: On the memory-tightness of hashed ElGamal. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part II. LNCS, vol. 12106, pp. 33–62. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45724-2_2
9. Iwata, T., Ohashi, K., Minematsu, K.: Breaking and repairing GCM security proofs. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 31–49. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32009-5_3
10. Jaeger, J., Tessaro, S.: Tight time-memory trade-offs for symmetric encryption. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part I. LNCS, vol. 11476, pp. 467–497. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17653-2_16
11. McGrew, D.A., Viega, J.: The security and performance of the Galois/Counter Mode (GCM) of operation. In: Canteaut, A., Viswanathan, K. (eds.) INDOCRYPT 2004. LNCS, vol. 3348, pp. 343–355. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-30556-9_27
12. Pollard, J.M.: A Monte Carlo method for factorization. BIT Numer. Math. **15**(3), 331–334 (1975)
13. Quisquater, J.-J., Delescaille, J.-P.: How easy is collision search. New results and applications to DES. In: Brassard, G. (ed.) CRYPTO 1989. LNCS, vol. 435, pp. 408–413. Springer, New York (1990). https://doi.org/10.1007/0-387-34805-0_38
14. Rogaway, P.: Nonce-based symmetric encryption. In: Roy, B., Meier, W. (eds.) FSE 2004. LNCS, vol. 3017, pp. 348–358. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-25937-4_22
15. Rogaway, P., Shrimpton, T.: A provable-security treatment of the key-wrap problem. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 373–390. Springer, Heidelberg (2006). https://doi.org/10.1007/11761679_23
16. Tessaro, S., Thiruvengadam, A.: Provable time-memory trade-offs: symmetric cryptography against memory-bounded adversaries. In: Beimel, A., Dziembowski, S. (eds.) TCC 2018, Part I. LNCS, vol. 11239, pp. 3–32. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-03807-6_1
17. Vaudenay, S.: Security flaws induced by CBC padding - applications to SSL, IPSEC, WTLS. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 534–545. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-46035-7_35
18. Wang, Y., Matsuda, T., Hanaoka, G., Tanaka, K.: Memory lower bounds of reductions revisited. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part I. LNCS, vol. 10820, pp. 61–90. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78381-9_3