



Hiding in Plain Sight: Memory-Tight Proofs via Randomness Programming

Ashrujit Ghoshal¹(✉), Riddhi Ghosal², Joseph Jaeger³, and Stefano Tessaro¹

¹ Paul G. Allen School of Computer Science & Engineering,
University of Washington, Seattle, USA
{ashrujit,tessaro}@cs.washington.edu

² University of California, Los Angeles, USA
riddhi@cs.ucla.edu

³ Georgia Institute of Technology, Atlanta, GA, USA
josephjaeger@gatech.edu

Abstract. This paper continues the study of *memory-tight reductions* (Auerbach et al., CRYPTO '17). These are reductions that only incur minimal memory costs over those of the original adversary, allowing precise security statements for memory-bounded adversaries (under appropriate assumptions expressed in terms of adversary time and memory usage). Despite its importance, only a few techniques to achieve memory-tightness are known and impossibility results in prior works show that even basic, textbook reductions cannot be made memory-tight.

This paper introduces a new class of memory-tight reductions which leverage random strings in the interaction with the adversary to hide state information, thus shifting the memory costs to the adversary.

We exhibit this technique with several examples. We give memory-tight proofs for digital signatures allowing many forgery attempts when considering randomized message distributions or probabilistic RSA-FDH signatures specifically. We prove security of the authenticated encryption scheme Encrypt-then-PRF with a memory-tight reduction to the underlying encryption scheme. By considering specific schemes or restricted definitions we avoid generic impossibility results of Auerbach et al. (CRYPTO '17) and Ghoshal et al. (CRYPTO '20).

As a further case study, we consider the textbook equivalence of CCA-security for public-key encryption for one or multiple encryption queries. We show two qualitatively different memory-tight versions of this result, depending on the considered notion of CCA security.

Keywords: Provable security · Memory-tightness · Time-memory trade-offs

1 Introduction

The aim of concrete security proofs is to lower bound, as precisely as possible, the resources needed to break a cryptographic scheme of interest, under some plausible assumptions. The traditional resource used in provable security is *time*

complexity (as well as related metrics, like data complexity). Recent works [1, 7, 9–12, 15–17, 20–22] have focused on additionally taking the *memory* costs of the adversary into account. This is important, as the amount of available memory can seriously impact the feasibility of an attack.

This paper presents new techniques for *memory-tight reductions*, a notion introduced by Auerbach et al. [1] to relate the assumed time-memory hardness of an underlying computational problem to the security of a scheme. More precisely, the end goal is to prove, via a reduction, that any adversary running in time t and with s bits of memory can achieve at most advantage $\epsilon = \epsilon(t, s)$ in compromising a scheme, by assuming that some underlying computational problem can only be solved with advantage $\delta = \delta(t', s')$ by algorithms running in time t' and with memory s' . A memory-tight reduction guarantees that $s \approx s'$, and usually, we want this to be tight also according to other parameters, i.e., $t \approx t'$ and $\epsilon \approx \delta$.

Memory-tight reductions are of value whenever the underlying problem is (conjectured to) be *memory sensitive*, i.e., the time needed to solve it grows as the amount of memory available to the adversary is reduced. Examples of memory-sensitive problems include classical ones in the public-key setting, such as breaking RSA and factoring, lattice problems and LPN, solving discrete logarithms over finite fields,¹ as well as problems in the secret-key setting, such as finding k -way collisions (for $k > 2$), finding several collisions at once [11], and distinguishing random permutations from random functions [12, 17, 20].

Developing memory-tight reductions is not always easy, and can be (provably) impossible [1, 15, 16, 22]. This makes it fundamental to develop as many techniques as possible to obtain such reductions. In this paper, we identify a class of examples which admit a new kind of memory-tight reductions. Our approach relies on the availability of random strings exchanged between the adversary and the security game, and which the reduction can leverage to encode state which can be recovered from later queries of the adversary, without the need to store this information locally, and thus saving memory. (In particular, the burden of keeping this information remains on the adversary, which needs to reproduce this random string for this state information to be relevant.) We present these techniques abstractly in the next section, with the help of a motivating example, and then move on to an overview of our specific results.

1.1 Our Techniques - An Overview

As a motivating example, consider the standard UFCMA security notion for signatures. It is defined via a game where the attacker, given the verification key vk , obtains signatures for chosen messages $m_1, m_2, \dots$, after which it outputs a candidate message-signature pair (m^*, σ^*) , and wins if m^* was not signed before, and σ^* is valid for m^* . When ignoring memory, this notion is *tightly* equivalent to one (which we refer to as mUFCMA) that allows for an arbitrary number

¹ However, the discrete logarithm problem in elliptic-curve groups, or any other group in which the best-known attacks are generic, is not memory sensitive, since optimal memory-less attacks are known.

of “forgery attempts” for pairs (m^*, σ^*) , and the adversary wins if one of them succeeds in the above sense. This is convenient: we generally target mUFCMA, but only need to deal with proving the simpler UFCMA notion.

The classical reduction transforms any mUFCMA adversary into a roughly equally efficient UFCMA adversary, which wins with the same probability, by (1) simulating forgery queries using the verification key, and (2) outputting the first forgery query (m^*, σ^*) which validates and such that m^* is fresh. This reduction is however *not* memory-tight, as we need to ensure the freshness of m^* , which requires remembering the previously signed messages. ACFK [1] prove that this is in some sense necessary, by showing that a (restricted) class of reductions cannot be memory-tight via a reduction to streaming lower bounds.

OUR IDEA: EFFICIENT TAGGING. To illustrate our new technique, which we refer to as *efficient tagging*, imagine now that we only use the signature scheme to sign *random* messages $m_1, m_2, \dots, m_q \leftarrow^s \{0, 1\}^\ell$, and consider a corresponding variant of mUFCMA security, which we want to reduce to (plain) UFCMA security. This, intuitively, does not seem to help resolve the above issue, because random messages are hardest to compress.

However, what is important here is that the reduction is responsible for simulating the random messages, and can simulate them in special ways, and *program* them so that they encode state information. For instance, assume that the reduction has access to an *injective* random function $f : [q] \rightarrow \{0, 1\}^\ell$, with inverse f^{-1} , which can be simulated *succinctly* from a short key as a pseudorandom object. Then, the reduction to UFCMA can set $m_i \leftarrow f(i)$ for the i -th query, and upon simulating a forgery query for (m^*, σ^*) , the reduction checks whether $f^{-1}(m^*) \in [q]$ to learn whether m^* is a fresh signing query or not.

Of course, the simulation is not perfect: The original m_i ’s are not necessarily distinct (this can be handled via the classical “switching lemma”). Also, the reduction could miss a valid forgery if the adversary outputs m_i *before* it is given to the adversary, but this again only occurs with small probability.

INEFFICIENT TAGGING AND NON-TIME-TIGHT REDUCTIONS. In the above example, we can efficiently check that $f^{-1}(m^*) \in [q]$. However, in some cases we may not – again, consider an example where the messages to be signed are sampled as $m_i \leftarrow h(r_i)$, where h is a hard-to-invert function and r_i is random. Then, we could adapt our proof above by setting $m_i \leftarrow h(f(i))$, but now, to detect a prior signing query, we would have to check whether $m^* = h(f(i))$ for some $i \in [q]$, and this can only be done in linear time. The resulting UFCMA adversary runs in time $t' = t + \Theta(q_F \cdot q)$, where t is the running time of the original adversary and q_F is the number of forgery attempts. For example, if $q \approx q_F \approx t$, the reduction is not time tight, and the adversary runs in time $t' = O(t^2)$.

ARE NON-TIME-TIGHT REDUCTIONS USELESS? It turns out that such non-time-tight reductions can still be helpful to infer that breaking a scheme requires memory, although this ultimately depends on the concrete security of the problem targeted by the reduction. Say, for example, a reduction for a given scheme transforms a successful adversary running in time t and using memory s into an

adversary running in time t^2 and using memory s breaking discrete logarithms over $\mathbb{F}_p$, for a 4096-bit prime p . It turns out that if we have fewer than 2^{78} bits of memory, no known discrete logarithm algorithm is better than a generic one (i.e. runs in time better than 2^{2048}), which means that our non-time-tight reduction is still sufficient to infer security for any $s < 2^{78}$ as long as $t < 2^{1024}$.

MESSAGE ENCODING. At the highest level, what happens is that the reduction is in control of certain random values which we can exploit to hide state information which can later be uniquely recovered, since triggering a situation where the reduction needs to remember requires the adversary to actually give back to the reduction this value. In the above, this state information is simple, namely whether the query is old or not. But as we will show below, the paradigm can be used to store complex information – we refer to this technique as *message encoding*, and discuss an example below.

A NEW VIEWPOINT: $\mathcal{F}$ -ORACLE ADVERSARIES. In our technique described above we needed access to a large random injection, which we argue can be simulated pseudorandomly. Prior works have similarly used PRFs to pseudorandomly simulate random oracles [1, 6] with low memory. The fact that one needs to decide how to simulate such objects when stating a memory-tight reduction is rather inconvenient: different instantiations seemingly lead to quantitatively different reductions, although this fact does not appear to be a reflection of any particular reality. In this paper, we propose (and advocate for) what we believe to be the “right” viewpoint: Our reductions are stated in terms of $\mathcal{F}$ -oracle adversaries where $\mathcal{F}$ is a set of functions and such an adversary expects oracle access to a random $f \in \mathcal{F}$. Then, a memory-tightness theorem is obtained in one of two ways, by either (1) applying a generic lemma stating that f can be instantiated in low memory using an $\mathcal{F}$ -pseudorandom function, or (2) assuming that the use of f does not functionally increase the success chances of the adversary because f is independent of the problem instance being solved (this is provably the case for some information theoretic problems). In particular, (1) is more conservative than (2), but it is very likely that (2) is also a viable approach which leads to cleaner result – indeed, we do not expect any of the considered memory-sensitive problems to become easier given access to an oracle from any natural class $\mathcal{F}$ – e.g., Factoring does not become easier given access to a random injection.

1.2 Our Results

We now move to an overview of our results (summarized in Fig. 1) which exemplify different applications of the tagging and message-encoding techniques.

MULTI-CHALLENGE SECURITY OF DIGITAL SIGNATURES. Our first results consider the security of digital signatures in the face of multiple forgery attempts (i.e., challenge queries), generalizing the examples discussed above. We work with a notion we refer to as UFRMA (unforgeability under randomized message attack). This notion is parameterized by a *message distribution* D and when the attacker makes a signing query for m it receives a signature

of $m' = D(m; r)$ for a random r . If m and r can be extracted from m' , giving the notion xUFRMA (or mxUFRMA for many forgery attempts), we can generalize our efficient tagging approach above by having the reduction to UFCMA choose $r = f(m, i)$ where each $f(m, \cdot)$ is a random injection. This setting can capture, e.g., the signatures used in key exchange protocols like TLS 1.3 where the server signs a transcript which includes a random 256-bit nonce. A version of our inefficient tagging example works when only m can be extracted from m' (wUFRMA); we pick $r = f(m, i)$ and in verification of a forgery query perform the linear time check of whether $m^* = D(m; f(m, i))$ for some $i \in [q]$. This setting captures places where the message to be signed includes a fresh public key or ciphertext. This includes, for example, the use of signatures for signing certificates, in some key exchange protocols, and in signcryption.

We further prove mUFCMA security for particular schemes. First, we can randomize any digital signature scheme DS (obtaining a scheme we call RDS) by signing $m \| r$ for random r chosen by the signing algorithm and including r as part of the signature. An immediate implication of our mxUFCRA result is a tight reduction from the mUFCMA security of RDS to the UFCMA security of the underlying scheme. One particular instantiation of RDS is Probabilistic Full Domain Hash with RSA (RSA-PFDH) which was introduced by Coron [8] to provide a variant of Full Domain Hash [4] with an (advantage-) tighter security proof. Using our efficient tagging technique we obtain a fully tight proof of the *strong* mUFCMA security of RSA-PFDH from the RSA assumption.

In independent and concurrent work, Diemert, Gellert, Jager, and Lyu [10] studied the mUFCMA security of digital signature schemes. They also considered the RDS construction, proving that if DS can be proven strong UFCMA1 secure² with a restricted class of “canonical” memory-tight reductions then there is a memory-tight reduction for the strong mUFCMA security of RDS. This complements our result, showing memory-tight strong mUFCMA security of RDS based on a restricted class of schemes while our result proves memory-tight plain mUFCMA security based on any plain UFCMA scheme. They apply their RDS result to establish tight proofs for the strong mUFCMA security of RSA-PFDH (matching our direct proof in Theorem 3). as well as schemes based on lossy identification schemes and pairings.

AUTHENTICATED ENCRYPTION SECURITY. Ghoshal, Jaeger, and Tessaro [15] have recently observed that in the context of authenticated encryption (AE), it is difficult to lift confidentiality of the scheme, in terms of INDR security, to full AE security, when additionally assuming ciphertext integrity, if we want to do so in a memory-tight way. This is well motivated, as several works establish tight time-memory trade-offs for INDR security [9, 12, 17, 20, 21], which we would like to lift to their AE security. The difficulty in the proof is that the INDR reduction must simulate a decryption oracle which rejects all ciphertexts *except those forwarded from an encryption query*. Recognizing these forwarded ciphertexts seems to require remembering state.

² The suffix ‘1’ indicates a variant of UFCMA security in which the adversary can only obtain a single signature per message. The security game always returning the same signature if the adversary repeats signature queries.

Assumption	Scheme	Result	Time	Memory	Advantage	New Technique
1UFCMA	Any	mxUFRMA	✓	✓	✓	Efficient tagging
	RDS	mUFCMA	✓	✓	✓	Efficient tagging
	Any	mwUFRMA	✗	✓	✓	Inefficient tagging
RSA	RSA-PFDH	mSUFCMA	✓	✓	✓	Efficient tagging
(PRF, INDR)	EtP	AE	(✓,✓)	(✗,✓)	(✓,✓)	Efficient tagging
1CCA	Any	mCCA	✗	✓	✗	Inefficient tagging
1\$CCA	Any	m\$CCA	✓	✓	✗	Message encoding

Fig. 1. Memory-tight reductions we provide. A 1 vs. an m prefix indicates whether one or many challenge queries are allowed. A ✓ vs. an ✗ indicates whether the reduction is tight with respect to that complexity metric. Reductions lacking tightness multiply running time/advantage by $O(q)$ or add $O(q)$ to the memory complexity, where q is the number of queries. An x vs. a w indicates whether the coins underlying the distribution of messages can be extracted from the message. RDS is randomization of any digital signature scheme by padding input messages with randomness. RSA-PFDH is probabilistic full-domain hash with RSA. EtP is the Encrypt-then-PRF AE construction.

Here, we give a different take and show that for specific schemes – in particular, those obtained by adding integrity via a PRF, following the lines of [3, 18, 19] – a memory-tight reduction *can* be given. Our INDR reduction is applied after arguing that the PRF looks like a random function f and thus forgeries are unlikely to occur. It uses f in a version of our efficient tagging technique to identify whether a ciphertext queried to decryption is fresh.³

CHOSEN CIPHERTEXT SECURITY: ONE TO MANY. A classical textbook result for public-key encryption shows that CCA-security against a *single* encryption query (1CCA) implies security against *multiple* queries (mCCA), with a quantitative advantage loss accounting to the number of such queries. ACFK [1] claim, incorrectly, that the associated reduction from 1CCA to mCCA is easy to make memory-tight, but this appears to be an oversight: No such reduction is known, and here we use our techniques to recover a memory-tight version of this result.

Let us consider concretely the “left-or-right” formulation of 1CCA/mCCA-security: The reduction from 1CCA to mCCA, given an adversary $\mathcal{A}$, picks a random $i \leftarrow^* [q]$ (where q is the number of encryption queries) and simulates the multi-query challenger to $\mathcal{A}$ by answering its first $i - 1$ encryption queries with an encryption of the left message, whereas the last $q - i$ queries are answered by encrypting the right message. Only the answer to the i -th query is answered by the single-query challenger. A problem arises when simulating the decryption queries: Indeed, we need to guarantee that a decryption query for *any* of the challenge ciphertexts $c_1^*, \dots, c_q^*$ returns an error $\perp$, yet this suggests that we seemingly need to remember the extra challenge ciphertexts c_j^* for $j \neq i$.

³ Ghoshal et al. [15] in fact described three variants of AE with different conventions for how decryption responds to non-fresh queries. By our results, memory-tight reductions to INDR are possible for two of the three variants.

We will resolve this in two ways. First, we give a new memory-tight reduction using the inefficient tagging method, with the same advantage loss as the original textbook reduction. Our reduction is non-time-tight, so may not be suitable for all situations. The main idea here is that we use the *randomness* used to generate the challenge ciphertext as our tag.

To obtain a reduction which is also tight with respect to time, we resort to the observation that changing to a stronger (but still commonly achieved) definition of CCA-security allows for different memory-tight reductions. We give in particular a memory-tight *and* time-tight reduction (with the usual factor q advantage loss) from the notion of $1\$CCA-m$ security to the notion of $m\$CCA-m$ security. These are variants of CCA security where (1) encryption queries are with respect to a *single* message, and return either the encryption of the message, or a random, independent ciphertext, and (2) decryption queries on a challenge ciphertext c_i^* returns the associated message.

Our reduction uses the full power of our message encoding approach, simulating random ciphertexts in a careful way which allows for recovering the associated challenge plaintext.

A FEW REMARKS. The above results on CCA security show us that the ability to give a memory-tight reduction is strongly coupled with definitional choices. In particular, different equivalent approaches to modeling the decryption oracle in the memory unbounded regime may not be equivalent in the memory-bounded setting. This means in particular that we need to exercise more care in choosing the right definition. We believe, for example, that the approach taken in $m\$CCA-m$ security is the more “natural” one (as it does not require artificially blocking the output of the decryption oracle, by always returning a message), but there may be contexts where other definitional choices are favored.

Another important lesson learnt from our AE result is that impossibility results, such as those in [1, 15, 16, 22], do not preclude positive results in form of memory-tight reductions, either by leveraging the structure of specific schemes, or by considering restricted security notions.

1.3 Paper Outline

Section 2 introduces notation, our computational model, and basic cryptographic background. Section 3 discusses our convention of using $\mathcal{F}$ -oracle adversaries. Section 4 gives our memory-tight reduction for digital signature schemes when many forgery attempts are allowed. In particular, the generic results are in Sect. 4.2, while the result specific to RSA-PFDH is in Sect. 4.5. Section 5 proves the security of Encrypt-then-PRF with a memory-tight reduction to the INDR security of the encryption scheme. Section 6 gives our results relating the one- and many-challenge query variants of CCA security. In particular, Sect. 6.1 gives our result for the traditional “left-vs.-right” notion and Sect. 6.2 gives our result for the “indistinguishable from random” variant. The full version of this paper [14] contains omitted proofs.

2 Preliminaries

Let $\mathbb{N} = \{0, 1, \dots\}$ and $[n] = \{1, \dots, n\}$ for $n \in \mathbb{N}$. If $x \in \{0, 1\}^*$ is a string, then $|x|$ denotes its length in bits. If S is a set, then $|S|$ denotes its size. We let $x \parallel y \parallel \dots$ denote an encoding of the strings $x, y, \dots$ from which the constituent strings can be unambiguously recovered. We identify bitstrings with integers in the standard way.

FUNCTIONS. Let T be a set (called the tweak set) and for each $t \in T$ let D_t and R_t be sets. Then $\text{Fcs}(T, D, R)$ denotes the set of all f such that for each $t \in T$, $f(t, \cdot)$ is a function from D_t to R_t . Similarly, $\text{Inj}(T, D, R)$ denotes the set of all f such that for each $t \in T$, $f(t, \cdot)$ is an injection from D_t to R_t . When D_t or R_t are independent of the choice of t we may omit the subscript.

If $f \in \text{Inj}(T, D, R)$, then its inverse f^{-1} is defined by $f^{-1}(t, f(t, x)) = x$ for all (t, x) and $f^{-1}(t, y) = \perp$ for $y \notin f(t, D_t)$. For such f we let $f^\pm$ denote the function defined by $f^\pm(+, x) = f(x)$ and $f^\pm(-, x) = f^{-1}(x)$. We let $\text{Inj}^\pm(T, D, R) = \{f^\pm : f \in \text{Inj}(T, D, R)\}$.

2.1 Computational Model

PSEUDOCODE. We regularly use pseudocode inspired by the code-based framework of [5]. We think of algorithms as randomized RAMs 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 . When the coins are implicit we write $\leftarrow_s$ in place of $\leftarrow$ and omit r .

We let $x \leftarrow_s \mathcal{D}$ denote sampling x according to the distribution $\mathcal{D}$. If $\mathcal{D}$ is a set, we overload notation and let $\mathcal{D}$ also denote the uniform distribution over elements of $\mathcal{D}$. The domain of $\mathcal{D}$ is denoted by $[\mathcal{D}]$.

Security notions are defined via games; for an example see Fig. 2. The probability that $\mathbf{G}$ outputs **true** is denoted $\Pr[\mathbf{G}]$. In proofs we sometimes define a sequence of “hybrid” games in one figure, using comments of the form “// $\mathbf{H}_{[i,j]}$.” A line of code commented thusly is only included in the hybrids $\mathbf{H}_k$ for $i \leq k < j$. (We are of course referring only to values of $k \in \mathbb{N}$.) By this convention to identify the differences between $\mathbf{H}_{k-1}$ and $\mathbf{H}_k$ one looks for comments $\mathbf{H}_{[i,k]}$ (code no longer included in the k -th hybrid) and $\mathbf{H}_{[k,j]}$ (code new to the k -th hybrid).

We let $\perp$ be a special symbol used to indicate rejection. If we do not explicitly include $\perp$ in a set, then $\perp$ is not contained in that set. If $\perp$ is an input to a function or algorithm, then we assume its output is $\perp$. We do not distinguish between $\perp$ and tuples $(\perp, \dots, \perp)$. Algorithms cannot query $\perp$ to their oracles.

COMPLEXITY MEASURES. To measure the complexity of algorithms we follow the conventions of measuring their local complexity, not including the complexity of whatever oracles they interact with. Local complexity was preferred by Auerbach et al. [1] for analyzing memory-limited adversaries so that analysis can be agnostic to minor details of security definitions’ implementations. We focus on worst-case runtime **Time**($\mathcal{A}$) and memory complexity **Mem**($\mathcal{A}$) (i.e.

how many bits of state it stores for local computation). These exclude the internal complexity of oracles queried by $\mathcal{A}$, but include the time and memory used to write the query and receive the response. If $\mathcal{A}$ expects access to n oracles then we let $\mathbf{Query}(\mathcal{A}) = (q_1, \dots, q_n)$ where q_i is an upper bound on the number of queries to its i -th oracle. (Here we index from left to right, so for $\mathcal{A}^{O_1, \dots, O_n}$ the i -th oracle is O_i .) If S is a scheme, then $\mathbf{Time}(S)$ and $\mathbf{Mem}(S)$ are the sums of the corresponding complexities over all of its algorithms. If G is a game, then we define $\mathbf{Time}(G)$ and $\mathbf{Mem}(G)$ to *exclude* the complexity of any adversaries embedded in the game.

2.2 Cryptographic Background

IDEAL MODELS. Some schemes we look at may be proven secure in ideal models (e.g. the random oracle or ideal cipher models). To capture this we can think of a scheme S as specifying a set of functions $S.I$. At the beginning of a security game a function h will be sampled from this set. The adversary and all algorithms of S are given oracle access to h .

FUNCTION FAMILIES. A family of functions F specifies, for each $K \in F.K$, an efficiently computable function $F_K \in F.F$. We refer to $F.F$ as the function space of F . Pseudorandom (PR) security of F is captured by the game defined in Fig. 2. It measures how F with a random key can be distinguished from a random function in $F.F$ via oracle access. We define $\mathbf{Adv}_F^{\text{pr}}(\mathcal{A}) = \Pr[\mathbf{G}_{F,1}^{\text{pr}}(\mathcal{A})] - \Pr[\mathbf{G}_{F,0}^{\text{pr}}(\mathcal{A})]$. The standard notions of (tweakable) pseudorandom functions/injections/permutations or strong injections/permutations are captured by appropriate choices of $F.F$.

Game $\mathbf{G}_{F,b}^{\text{pr}}(\mathcal{A})$	$\mathbf{Ev}(x)$
$h \leftarrow \$ F.I$	$y_1 \leftarrow F_K^h(x)$
$K \leftarrow \$ F.K$	$y_0 \leftarrow f(x)$
$f \leftarrow \$ F.F$	Return y_b
$b' \leftarrow \$ \mathcal{A}^{\text{Ev},h}$	
Return $b' = 1$	

Fig. 2. Security game capturing the pseudorandomness of function family F .

SWITCHING LEMMA. We make use of the following standard result which bounds how well a random function and a random injection can be distinguished.

Lemma 1 (Switching Lemma). *Fix T , D , and R . Let $N = \min_{t \in T} |R_t|$. Then for any adversary $\mathcal{A}$ with $q = \mathbf{Query}(\mathcal{A})$ we have that*

$$|\Pr[\mathcal{A}^f \Rightarrow 1] - \Pr[\mathcal{A}^g \Rightarrow 1]| \leq 0 \cdot q^2/N.$$

The probabilities are measured over the coins of $\mathcal{A}$, the uniform choice of f from $\mathbf{Fcs}(T, D, R)$, and the uniform choice of g from $\mathbf{Inj}(T, D, R)$.

Recent papers [11, 17, 20] have given improved versions of the switching lemma for adversaries with bounded memory complexity, as long as it does not repeat oracle queries. In our application of the switching lemma the adversary's memory complexity is too large for these bounds to provide any improvement.

OTHER PRIMITIVES. We recall relevant syntax and security definitions for digital signatures, nonce-based encryption, and public key encryption schemes in the sections where we consider them (Sects. 4, 5, and 6 respectively).

3 Adversaries with Access to Random Functions

This paper proposes and adopts what we consider to be a better formalism to deal with memory-tight reductions. Namely, all of our reductions will require access to some variety of large random functions which it will query on a small number of inputs (specifically uniformly random functions and invertible random injections). That is, our reduction adversaries can be written in the form shown of the left below, for some set of functions $\mathcal{F}$ and algorithm $\mathcal{A}_2$. (On the right is a pseudorandom version of $\mathcal{A}$ which we will discuss momentarily.)

Adversary $\mathcal{A}^O(in)$	Adversary $\mathcal{A}_F^O(in)$
$f \leftarrow_s \mathcal{F}$	$K \leftarrow_s \mathbf{F.K}$
$out \leftarrow_s \mathcal{A}_2^{O,f}(in)$	$out \leftarrow_s \mathcal{A}_2^{O,\mathbf{F}_K}(in)$
Return out	Return out

We refer to such an $\mathcal{A}$ as an $\mathcal{F}$ -oracle adversary. In this section we will generally discuss such adversaries, rather than separately providing the discussion for such adversaries each time we apply them.

The time and memory complexity of any $\mathcal{F}$ -oracle adversary must include the complexity of sampling, storing, and evaluating f . This will be significant if $\mathcal{F}$ is large. However, as we will argue, this additional state and time should be assumed to not significantly increase the advantage of $\mathcal{A}$. As such, we will define the *reduced complexity* of $\mathcal{A}$ by

$$\mathbf{Time}^*(\mathcal{A}) = \mathbf{Time}(\mathcal{A}_2) \text{ and } \mathbf{Mem}^*(\mathcal{A}) = \mathbf{Mem}(\mathcal{A}_2).$$

Later we state theorems in terms of reduced complexity.

PSEUDORANDOM REPLACEMENT. The most conservative justification of $\mathcal{F}$ -oracle adversaries is to bound how much the oracle can help by replacing it with a pseudorandom version. This was the approach taken by Auerbach et al. [1] when they used pseudorandom functions for purposes such as emulating random oracles and storing the coins required by an adversary with low memory, and has been adopted by follow-up work [7, 10, 22]. If $\mathbf{F}$ is a function family with $\mathbf{F.F} = \mathcal{F}$, then the adversary $\mathcal{A}_F$ we gave above does exactly this. It replaces $\mathcal{A}_2$'s oracle access to f with access to $\mathbf{F}_K$ for a random K . The following lemma is straightforward.

Lemma 2. *Let $\mathcal{A}$ be an $\mathcal{F}$ -oracle adversary for a game $\mathbf{G}$. Then for any function family $\mathbf{F}$ with $\mathbf{F.F} = \mathcal{F}$ we can define a pseudorandomness adversary $\mathcal{A}_i$ such that*

$$\begin{aligned} \Pr[\mathbf{G}(\mathcal{A})] &\leq \Pr[\mathbf{G}(\mathcal{A}_F)] + \mathbf{Adv}_F^{\text{Pr}}(\mathcal{A}_i), & \mathbf{Time}(\mathcal{A}_i) &= \mathbf{Time}^*(\mathcal{A}) + \mathbf{Time}(\mathbf{G}(\mathcal{A})), \\ \mathbf{Query}(\mathcal{A}_i) &= q, \text{ and} & \mathbf{Mem}(\mathcal{A}_i) &= \mathbf{Mem}^*(\mathcal{A}) + \mathbf{Mem}(\mathbf{G}(\mathcal{A})). \end{aligned}$$

Here q is an upper bound on the number of queries $\mathcal{A}_2$ makes to its second oracle.

Note that the complexity of $\mathcal{A}_F$ is given by $\mathbf{Time}(\mathcal{A}_F) = \mathbf{Time}^*(\mathcal{A}) + q \cdot \mathbf{Time}(F)$ and $\mathbf{Mem}(\mathcal{A}_F) = \mathbf{Mem}^*(\mathcal{A}) + \mathbf{Mem}(F)$. Thus the existence of an appropriate pseudorandom F ensures that the memory and time complexity excluded by $\mathbf{Time}^*$ and $\mathbf{Mem}^*$ cannot significantly aid an adversary. In the use of this technique by Auerbach et al. [1] the reduction $\mathcal{A}_i$ was memory-tight. Note this is not strictly necessary as long as we are willing to assume the existence of F with sufficient security as a function of attackers' time and query complexities without regard to memory complexity.

We could have combined Lemma 2 with any of our coming theorems to obtain bounds in terms of $\mathbf{Time}$ and $\mathbf{Mem}$, rather than their reduced version. However we find the use of reduced complexity cleaner as it simplifies our theorems, allowing us to focus on the conceptual core of the proofs without having to repeat the rote step of replacing random objects with pseudorandom ones.

When combining the lemma with a theorem, game G would correspond to the security game played by the reduction adversary. For our theorems, that game will have low time and memory overhead over that of $\mathcal{A}$, so the application of the lemma would be time- and memory-tight. That said, the tightness of this is less important than the tightness of the other components of the theorem we would apply it to. Note that the definition of $\mathcal{A}_i$ is *independent* of the choice of F . Consequently, we can always choose F with a very high security threshold to counteract any looseness in the lemma. In the full version [14], we summarize the $\mathcal{F}$ used in our theorems and how they could be pseudorandomly instantiated.

ASSUMED INDEPENDENCE. As a second observation why the storage of f may not help $\mathcal{A}$, note that f is completely “independent” of the problem $\mathcal{A}$ is trying to solve (as specified by *in* and the behavior of O). In various settings it seems likely that such independent state does not help. For example, it would be very surprising (or even a breakthrough) to show a better factoring or lattice algorithm given access to a random function f from a natural set. Indeed, cryptanalytic work often makes use of random oracles without significant comment (from which other types of random functions can be constructed).

INFORMATION THEORETIC SETTINGS. In some information theoretic settings, the “independence” of f from the problem can be made rigorous. Information theoretic results are typically depending only on the query complexity of the attacker or its memory usage, *ignoring code size*. In such settings, we expect bounds of the form $\mathbf{Adv}(\mathcal{A}) \leq \epsilon(\mathbf{Mem}(\mathcal{A}), \mathbf{Query}(\mathcal{A}))$ for some function ϵ . Because this bound *does not* depend on the code size of $\mathcal{A}$, if $\mathcal{A}$ is an $\mathcal{F}$ -oracle adversary we should be able to prove $\mathbf{Adv}(\mathcal{A}) \leq \epsilon(\mathbf{Mem}^*(\mathcal{A}), \mathbf{Query}(\mathcal{A}))$ by a coin-fixing argument in which we fix the random choice of function ahead of time and embed it in the description of the adversary. This is, for example, the case for the recent time-memory tradeoffs shown for distinguishing between a random function and a random injection without repeating queries [11, 17, 20]. A coin-fixing readily shows that these tradeoffs hold when using $\mathbf{Mem}^*(\mathcal{A})$ in place of $\mathbf{Mem}(\mathcal{A})$, $\mathbf{Query}(\mathcal{A})$.

4 Multi-challenge Security of Digital Signature Schemes

In the context of memory-tightness, the security of digital signature schemes has been considered in several works [1, 10, 22]. The standard security notion for signatures asks the attacker, given examples, to come up with a forged signature on a fresh message. A straightforward proof shows (in the standard setting where memory efficiency is not a concern) that the security notion is equivalent whether the attacker is allowed one or many forgery attempts. However, Auerbach et al. [1] proved an impossibility result showing that a (certain form of black-box) reduction cannot be time, memory, and advantage tight. The difficulty faced by the reduction is in distinguishing between when the adversary has produced a novel forgery and when it is simply repeating a signature that it was given.

In this section we show a few ways that security against many forgery attempts (i.e., multiple challenges) can be proven to follow from security against a single forgery (i.e., a single challenge) in a memory-tight manner. Our first results consider a variant definition of digital signature security we introduce (called UFRMA) in which the adversary has only partial control over the messages being signed. Using our new techniques, we show that single challenge UFCMA security implies multi-challenge UFRMA security in a memory-tight manner (for some practically relevant distributions over messages). We also consider the security of the RSA full domain hash digital signature scheme. Auerbach et al. [1] gave a memory-, but not advantage-tight proof of the security of the standard version of this scheme in the single challenge setting. By considering a probabilistic variant of the scheme introduced by Coron [8] we are able to provide a memory-, time-, and advantage-tight proof of the many-forgery SUFCMA security of the variant.

4.1 Syntax and Security

DIGITAL SIGNATURE SYNTAX. A digital signature scheme DS specifies a key generation algorithm DS.K , a signing algorithm DS.Sign , and a verification algorithm DS.Ver . The syntaxes of these algorithms are shown in Fig. 3.

We capture ideal models by providing DS.Sign and DS.Ver with oracle access to a function h drawn at random from the set DS.I . When relevant we let DS.M denote the set of messages it accepts.

The verification and signing keys are respectively denoted by vk and sk . The message to be signed is m , the signature produced is σ , and the decision is $d \in \{\text{true}, \text{false}\}$. Correctness requires $\text{DS.Ver}^h(vk, m, \sigma) = \text{true}$ for all $h \in \text{DS.I}$, all $(vk, sk) \in [\text{DS.K}]$, all $m \in \text{DS.M}$, and all $\sigma \in [\text{DS.Sign}^h(sk, m)]$.

DS Syntax

$$h \leftarrow \text{DS.I}$$

$$(vk, sk) \leftarrow \text{DS.K}$$

$$\sigma \leftarrow \text{DS.Sign}^h(sk, m)$$

$$d \leftarrow \text{DS.Ver}^h(vk, m, \sigma)$$

Fig. 3. Syntax of digital signature scheme.

Game	$\mathbf{G}_{\text{DS}}^{\text{ufcma}}(\mathcal{A})$, $\mathbf{G}_{\text{DS},\text{D}}^{\text{ufрма}}(\mathcal{A})$	$\text{SIGN}(m)$	$\text{RSIGN}(m)$
$h \leftarrow \$ \text{DS.I}$		$\mathcal{S} \leftarrow \mathcal{S} \cup \{m\}$	$r \leftarrow \$ \text{D.R}$
$(vk, sk) \leftarrow \$ \text{DS.K}$		$\sigma \leftarrow \$ \text{DS.Sign}^h(sk, m)$	$m' \leftarrow \text{D.S}(m; r)$
$\mathcal{S} \leftarrow \emptyset$		Return σ	$\mathcal{S} \leftarrow \mathcal{S} \cup \{m'\}$
win $\leftarrow$ false		$\text{FORGE}(m^*, \sigma^*)$	$\sigma \leftarrow \$ \text{DS.Sign}^h(sk, m')$
Run $\mathcal{A}^{\text{SIGN, FORGE}, h}(vk)$		If $m^* \notin \mathcal{S}$:	Return (σ, r)
Run $\mathcal{A}^{\text{RSIGN, FORGE}, h}(vk)$		If $\text{DS.Ver}^h(vk, m, \sigma)$:	
Return win		win $\leftarrow$ true	

Fig. 4. Security games capturing the unforgeability of a digital signature scheme.

MESSAGE DISTRIBUTION SYNTAX. One of the security notions we consider for digital signature schemes will be parameterized by a message distribution via which the adversary is given incomplete control over the messages which are signed. A message distribution D specifies sampling algorithm $D.S$ which samples an output message m' based on parameters m given as input (written $m' \leftarrow \$ D.S(m)$). The parameters m must be drawn from a set $D.M$, which we typically leave implicit. When making the randomness of the sampling algorithm explicit we let $D.R$ be the set from which its randomness is drawn and write $m' \leftarrow D.S(m; r)$. If there exists an extraction algorithm $D.X$ such that $D.X(D.S(m; r)) = (m, r)$ for all m, r then we say D is *extractable*. If $D.X(D.S(m; r)) = m$ for all m, r then we say D is *weakly extractable*. We assume that $D.X(m') = \perp$ if $m' \neq D.S(m; r)$ for all m, r . We define the min-entropy of D as

$$D.H_{\infty} = -\lg \max_m \Pr[r \leftarrow \$ D.R : D.S(m; r) = m'] .$$

UNFORGEABILITY SECURITY. The unforgeability security notions we consider are defined in Fig. 4. The standard notion of UFCMA (unforgeability under chosen message attack) security is captured by $\mathbf{G}_{\text{DS}}^{\text{ufcma}}$ which includes the boxed but not the highlighted code, giving the adversary access to a regular signing oracle SIGN . The goal of the adversary is to query FORGE with a valid signature σ^* of a message m^* which was not previously included in a signing query (as stored by the set $\mathcal{S}$). We define $\text{Adv}_{\text{DS}}^{\text{ufcma}}(\mathcal{A}) = \Pr[\mathbf{G}_{\text{DS}}^{\text{ufcma}}(\mathcal{A})]$.

Our new security notion UFRMA (unforgeability under randomized message attack) is captured by the game $\mathbf{G}_{\text{DS},\text{D}}^{\text{ufрма}}$ which is parameterized by a message distribution D . In this game the adversary is instead given access to the randomized signing oracle RSIGN where the message to be signed is chosen by D . Note that the coins used by D are returned to the adversary along with the signature. Otherwise this game matches that of UFCMA security. We define $\text{Adv}_{\text{DS},\text{D}}^{\text{ufрма}}(\mathcal{A}) = \Pr[\mathbf{G}_{\text{DS},\text{D}}^{\text{ufрма}}(\mathcal{A})]$.

We will relate the advantage of attacks making only a single forgery attempt and those making many such attempts. When wanting to make the distinction explicit we prefix the abbreviation of a security notion with an ‘m’ or ‘1’. Strong UFCMA security, denoted SUFCMA, is captured by modifying $\mathbf{G}_{\text{DS}}^{\text{ufcma}}$ to

store the tuple (σ, m) in $\mathcal{S}$ in SIGN and checking $(m^*, \sigma^*) \notin \mathcal{S}$ in FORGE. We denote this by G^{sufcma} and the corresponding advantage by $\text{Adv}^{\text{sufcma}}$. We define SUFRMA, $G^{\text{sufрма}}$, and $\text{Adv}^{\text{sufрма}}$ analogously. We write xUFRMA when assuming that D is extractable and wUFRMA when assuming it is weakly extractable.

4.2 Multi-challenge Security for Extractable Message Distributions

The first applications we show for our techniques are generic methods of tightly implying security of a digital signature scheme against multiple forgery attempts (i.e., multi-challenge security). Recall that Auerbach et al. [1] gave a lower bound showing that a black-box reduction proving that single UFCMA security implies many UFCMA cannot be made memory-tight and time-tight. We avoid this in two ways; first by considering mUFRMA, rather than mUFCMA, security and then by considering a particular choice of digital signature scheme.

HIGH-LEVEL IDEA. The primary difficulty of a tight proof that 1UFCMA security implies mUFCMA security is that a successful mUFCMA attacker may have made many FORGE queries which verify correctly, one of which is a valid forgery and the rest of which were just forwarded from its SIGN oracle. A 1UFCMA reduction must then somehow be able to identify which of the queries is the true forgery so it can forward this to its own FORGE oracle.

The technical core of the coming proof for mUFRMA is that our reduction adversary will use the random coins of the message distribution D to signal things to its future self. In particular, when $\mathcal{A}_r$ makes a query $\text{RSIGN}(m)$, the reduction will choose coins for $D.S$ via $r \leftarrow f(m, i)$ where i is a counter which is incremented with each query and f is a random tweakable function/injection. The coins then act as a sort of authentication tag for m . On a later $\text{FORGE}(m^*, \sigma^*)$ query, if $m^* = D.S(m; r)$ where $r = f(m, i)$ for some $i \in [q_{\text{SIGN}}]$ the reduction can safely assume this message was signed by an earlier RSIGN query.

When D is fully extractable, we can perform the requisite check for FORGE by having f be an injection. We extract m and r from m^* and then compute $i \leftarrow f^{-1}(m, r)$. This is the strategy used in Theorem 1. If we assume only that D is weakly extractable, we can extract m if D has a sufficient amount of entropy, and then individually check if $D.S(m; f(m, i))$ holds for each choice of i . This reduction strategy obtains the same advantage at the cost of an extra runtime being needed to iterate over the possible choices of i in FORGE.

EXTRACTABLE MESSAGE DISTRIBUTION. If the message distribution D is extractable, the following theorem captures that 1UFCMA security tightly implies mUFRMA security. The proof makes use of our efficient tagging technique.

Theorem 1 (1UFCMA $\Rightarrow$ mxUFRMA). *Let DS be a digital signature scheme and D be an extractable message distribution. Let $\mathcal{A}_r$ be an adversary with $(q_{\text{SIGN}}, q_{\text{FORGE}}, q_h) = \text{Query}(\mathcal{A}_r)$ and assume $q_{\text{SIGN}} \leq 0.5|D.R|$. Let $\mathcal{A}_u$ be the*

Adversary $\mathcal{A}_u^{\text{SIGN, FORGE, h}}(vk)$	SIMRSIGN(m)	SIMFORGE(m^*, σ^*)
$i \leftarrow 0$	$i \leftarrow i + 1$	$(m, r) \leftarrow \text{D.X}(m^*)$
$f \leftarrow \text{Inj}^\pm(\text{DS.M}, [q_{\text{SIGN}}], \text{D.R})$	$r \leftarrow f(m, i)$	If $f^{-1}(m, r) \notin [q_{\text{SIGN}}]$:
Run $\mathcal{A}_r^{\text{SIMRSIGN, SIMFORGE, h}}(vk)$	$m' \leftarrow \text{D.S}(m; r)$	If $\text{DS.Ver}^h(vk, m^*, \sigma^*)$:
	$\sigma \leftarrow \text{SIGN}(m')$	Query FORGE(m^*, σ^*)
	Return (σ, r)	Halt execution

Fig. 5. Adversary $\mathcal{A}_u$ used in proof of Theorem 1.

$\text{Inj}^\pm(\text{DS.M}, [q_{\text{SIGN}}], \text{D.R})$ -oracle adversary shown in Fig. 5. Then,

$$\text{Adv}_{\text{DS}, \text{D}}^{\text{ufirma}}(\mathcal{A}_r) \leq \text{Adv}_{\text{DS}}^{\text{ufcma}}(\mathcal{A}_u) + (0.5 \cdot q_{\text{SIGN}}^2 + 2 \cdot q_{\text{SIGN}} \cdot q_{\text{FORGE}})/|\text{D.R}|$$

$$\mathbf{Query}(\mathcal{A}_u) = (q_{\text{SIGN}}, 1, q_h + q_{\text{FORGE}} \cdot \mathbf{Query}(\text{DS}))$$

$$\mathbf{Time}^*(\mathcal{A}_u) = \mathbf{Time}(\mathcal{A}_r) + q_{\text{SIGN}} \cdot \mathbf{Time}(\text{D}) + q_{\text{FORGE}}(\mathbf{Time}(\text{D}) + \mathbf{Time}(\text{DS}))$$

$$\mathbf{Mem}^*(\mathcal{A}_u) = \mathbf{Mem}(\mathcal{A}_r) + \mathbf{Mem}(\text{D}) + \mathbf{Mem}(\text{DS}) + \lg(q_{\text{SIGN}}).$$

This is time-tight because $\mathbf{Time}(\mathcal{A}_r) \in \Omega(q_{\text{SIGN}} + q_{\text{FORGE}})$ must hold and $\mathbf{Time}(\text{D})$ and $\mathbf{Time}(\text{DS})$ will be small. This is memory-tight because $\mathbf{Mem}(\text{D})$, $\mathbf{Mem}(\text{DS})$, and $\lg(q_{\text{SIGN}})$ will be small.

The main idea of $\mathcal{A}_u$ is using the output of an invertible random injection f on the message and a counter as coins instead of sampling them uniformly at random when answering RSIGN queries. Since D is fully extractable, during a FORGE query on m^* , we can extract $(m, r) \leftarrow \text{D.X}(m^*)$ and use the fact that f is invertible to compute $f^{-1}(m, r)$ and check if the index is in $[q_{\text{SIGN}}]$. This is used to avoid remembering $\mathcal{S}$. If $m^* \in \mathcal{S}$, and $(m, r) \leftarrow \text{D.X}(m^*)$, then there exists $j \in [q_{\text{SIGN}}]$ such that $r = f(m, j)$ – so the check passes. We can argue that if $m^* \notin \mathcal{S}$, our check is unlikely to pass. We give the formal proof of this theorem in Sect. 4.3. It applies the switching lemma to argue the use of f cannot be distinguished from honestly sampling r with advantage better than $0.5 \cdot q_{\text{SIGN}}^2/|\text{D.R}|$ and shows that the probability of falsely making the check pass is bounded by $2q_{\text{SIGN}}q_{\text{FORGE}}/|\text{D.R}|$.

We would not be able to use the technique in this proof to prove mxSUFRMA from 1SUFCMA in a memory-tight way. In particular, since the coins r of the message distribution are chosen before σ is known, our trick of using r to signal freshness of a forgery query does not work for a message-signature pair.

4.3 Proof of Theorem 1 (1UFCMA $\Rightarrow$ mUFRMA)

Proof. We consider a sequence of hybrids H_0 through H_4 defined in Fig. 6. When examining these hybrids recall our conventions regarding “ $//\text{H}_{[i,j]}$ ” comments described in Sect. 2.1. Of these hybrids we will make the following claims, which establish the upper bound on the advantage of $\mathcal{A}_r$ claimed in the proof.

Games H_h for $0 \leq h \leq 4$	
$h \leftarrow \$ \text{DS.I}$ $(vk, sk) \leftarrow \$ \text{DS.K}; \mathcal{S} \leftarrow \emptyset$ $\text{win} \leftarrow \text{false}$ $i \leftarrow 0; I[\cdot] \leftarrow \emptyset$ $f \leftarrow \$ \text{Fcs}(\text{DS.M}, [q_{\text{SIGN}}], \text{D.R}) // \text{H}_{[1,2]}$ $f \leftarrow \$ \text{Inj}^\pm(\text{DS.M}, [q_{\text{SIGN}}], \text{D.R}) // \text{H}_{[2,\infty]}$ $\text{Run } \mathcal{A}_r^{\text{RSIGN, FORGE}, h}(vk)$ Return win	$\text{RSIGN}(m)$ $r \leftarrow \$ \text{D.R} // \text{H}_{[0,1]}$ $i \leftarrow i + 1 // \text{H}_{[1,\infty]}$ $r \leftarrow f(m, i) // \text{H}_{[1,\infty]}$ $I[m] \leftarrow I[m] \cup \{i\} // \text{H}_{[3,4]}$ $m' \leftarrow \text{D.S}(m; r)$ $\mathcal{S} \leftarrow \mathcal{S} \cup \{m'\} // \text{H}_{[0,3]}$ $\sigma \leftarrow \$ \text{DS.Sign}^h(sk, m')$ $\text{Return } (\sigma, r)$ $\text{FORGE}(m^*, \sigma^*)$ $(m, r) \leftarrow \text{D.X}(m^*)$ $\text{If } m^* \notin \mathcal{S}: // \text{H}_{[0,3]}$ $\text{If } f^{-1}(m, r) \notin I[m]: // \text{H}_{[3,4]}$ $\text{If } f^{-1}(m, r) \notin [q_{\text{SIGN}}]: // \text{H}_{[4,\infty]}$ $\quad \text{If } \text{DS.Ver}^h(vk, m^*, \sigma^*):$ $\quad \quad \text{win} \leftarrow \text{true}$

Fig. 6. Hybrid games used in proof of Theorem 1.

1. $\Pr[\text{G}_{\text{DS}, \text{D}}^{\text{ufma}}(\mathcal{A}_r)] = \Pr[\text{H}_0] = \Pr[\text{H}_1]$
2. $\Pr[\text{H}_1] \leq \Pr[\text{H}_2] + 0.5 \cdot q_{\text{SIGN}}^2 / |\text{D.R}|$
3. $\Pr[\text{H}_2] = \Pr[\text{H}_3]$
4. $\Pr[\text{H}_3] \leq \Pr[\text{H}_4] + 2q_{\text{SIGN}}q_{\text{FORGE}}/|\text{D.R}|$
5. $\Pr[\text{H}_4] = \text{Adv}_{\text{DS}}^{\text{ufcma}}(\mathcal{A}_u)$

TRANSITION H_0 TO H_1 . The hybrid H_0 is simply a copy of the game G^{ufma} . (We also added code to initialize variables i and $I[\cdot]$ that will be used in later hybrids.) Hence $\Pr[\text{G}^{\text{ufma}}(\mathcal{A}_r)] = \Pr[\text{H}_0]$. In hybrid H_1 , we replace the random sampling of r for D in RSIGN with the output of a random function f applied to m , using a counter i to provide domain separation between different queries. This method of choosing r is equivalent, so $\Pr[\text{H}_0] = \Pr[\text{H}_1]$.

TRANSITION H_1 TO H_2 . In hybrid H_2 we replace the random function with a random injection. This modifies the behavior of the game only in that values of r are guaranteed not to repeat across different signing queries that used the same message. There are at most q_{SIGN} invocations of f , so the switching lemma (Lemma 1) tells us that $\Pr[\text{H}_1] \leq \Pr[\text{H}_2] + 0.5 \cdot q_{\text{SIGN}}^2 / |\text{D.R}|$.

TRANSITION H_2 TO H_3 . In hybrid H_3 , we replace the check whether $m^* \notin \mathcal{S}$ in oracle FORGE with a check if $f^{-1}(m, r) \notin I[m]$ where $(m, r) = \text{D.X}(m^*)$. Here $I[\cdot]$ is a new table introduced into the game. In RSIGN , code was added which uses $I[m]$ to store each of the counter values for which $\mathcal{A}_r$ made a signing query for m . Hence $f^{-1}(m, r)$ will be in $I[m]$ iff m^* is in $\mathcal{S}$ and so $\Pr[\text{H}_2] = \Pr[\text{H}_3]$.

TRANSITION H_3 TO H_4 . In the final transition to hybrid H_4 we replace the FORGE check $f^{-1}(m, r) \notin I[m]$ with $f^{-1}(m, r) \notin [q_{\text{SIGN}}]$. This *does* change behavior if $\mathcal{A}_r$ ever makes a successful forgery query for $m^* = \text{D.S}(m; f(m, i))$ without its i -th signing query having used the message m . This would require guessing $f(m, i)$ for

some $i \in [q_{\text{SIGN}}] \setminus I[m]$. We can bound the probability of this ever occurring by a union bound over the FORGE queries made by $\mathcal{A}_r$. Consider the set $f(m, [q_{\text{SIGN}}] \setminus I[m]) = \{f(m, i) : i \in [q_{\text{SIGN}}] \setminus I[m]\}$. It has size at most q_{SIGN} . Because f is a random injection it is uniform subset of the set $\text{D.R} \setminus f(m, I[m])$ (which has size at least $|\text{D.R}| - q_{\text{SIGN}}$). Hence the probability of any particular query triggering this different behavior is at most $q_{\text{SIGN}} / (|\text{D.R}| - q_{\text{SIGN}}) \leq 2q_{\text{SIGN}} / |\text{D.R}|$. Applying the union bound gives us $\Pr[H_3] \leq \Pr[H_4] + 2q_{\text{SIGN}} \cdot q_{\text{FORGE}} / |\text{D.R}|$.

REDUCTION TO UFCMA. We complete the proof using adversary $\mathcal{A}_u$ from Fig. 5 which simulates hybrid H_4 and succeeds whenever $\mathcal{A}_r$ would. The adversary $\mathcal{A}_u$ samples f at random from $\text{Inj}(\text{DS.M}, [q_{\text{SIGN}}], \text{D.R})$. When run on input vk , it runs $\mathcal{A}_r$ on the same input. It gives $\mathcal{A}_r$ direct access to h . To simulate a query $\text{RSIGN}(m)$, it computes $m' \leftarrow \text{D.S}(m; f(m, i))$, increments i , and queries $\text{SIGN}(m')$, returning the result to $\mathcal{A}_r$. On a query $\text{FORGE}(m^*, \sigma^*)$, it computes $(m, r) \leftarrow \text{D.X}(m^*)$. If $f^{-1}(r) \notin [q_{\text{SIGN}}]$ and $\text{DS.Ver}(vk, m^*, \sigma^*) = \text{true}$ then it queries its own oracle with (m^*, σ^*) and halts. Otherwise it ignores the query.

If adversary $\mathcal{A}_u$ ever makes a FORGE query, it will succeed. It ensured that (m^*, σ^*) is verified correctly and $f^{-1}(r) \notin [q_{\text{SIGN}}]$ ensures that it has not previously made a SIGN query for m^* . If $\mathcal{A}_r$ would have succeeded in hybrid H_4 , its winning query will cause $\mathcal{A}_u$ to make a FORGE query. Hence, we have $\Pr[H_4] = \text{Adv}_{\text{DS}}^{\text{ufcma}}(\mathcal{A}_u)$.

4.4 Applications and Weakly Extractable Variant

We discuss some applications of Theorem 1. This includes scenarios where extractable message distributions are used and proving security of digital signature schemes when their messages are padded with randomness. Additionally, we give a variant of the theorem when the underlying message distribution is only weakly extractable. The resulting reduction is memory- but not time-tight.

EXAMPLE EXTRACTABLE DISTRIBUTIONS. The simplest extractable distribution does not accept parameters as input and simply outputs its randomness as the message. Security with respect to this is the standard notion of security against random message attacks which was originally introduced by Even, Goldreich, and Micali [13].

Extractable distributions arise naturally when the messages being signed include random values. For example, protocols often include random nonces in messages that are signed. In TLS 1.3, for example, when the server is responding to the Client Hello Message it signs a transcript of the conversation up until that point which includes a 256-bit nonce just chosen by the server. We could think of the security for this setting being captured by an extractable distribution D_{tls} that takes as input message parameter m that specifies all of the transcript other than the nonce and sets the nonce to its randomness $r \in \{0, 1\}^{256}$.

PADDING SCHEMES WITH RANDOMNESS. Using Theorem 1, we can see that augmenting any digital signature scheme by appending auxiliary randomness will give us a memory-tight reduction from the mUFCMA security of the augmented scheme to the 1UFCMA security of the original scheme.

Let DS be a digital signature scheme and $\mathbf{R}$ be a set. We define $\text{RDS}[\text{DS}, \mathbf{R}]$ by having $\text{RDS}[\text{DS}, \mathbf{R}].\text{Sign}(sk, m)$ do “ $r \leftarrow \mathbf{R}$; Return $\text{DS}.\text{Sign}(sk, m \parallel r) \parallel r$ ” and having $\text{RDS}[\text{DS}, \mathbf{R}].\text{Ver}(vk, m, \sigma')$ do “ $\sigma \parallel r \leftarrow \sigma'$; Return $\text{DS}.\text{Ver}(vk, m \parallel r, \sigma)$.” We also define a related message distribution $\text{RD}[\mathbf{R}]$ by $\text{RD}[\mathbf{R}].\mathbf{R} = \mathbf{R}$ and $\text{RD}[\mathbf{R}].\mathbf{S}(m; r) = m \parallel r$. Clearly it is extractable.

The following reduces the mUFCMA security of RDS to the mUFRMA security of DS . Theorem 1 can in turn be used to reduce this to the 1UFCMA security of DS . It also relates the mSUFCMA security of RDS to the mSUFRMA security of DS . We note this because if DS has unique signatures, then its mSUFRMA and mUFRMA security are identical and hence UFCMA security of DS implies mSUFCMA security of RDS in a memory-tight way.

Theorem 2. *Let DS be a digital signature scheme and $\mathbf{R}$ be a set. Then for any $\mathcal{A}_u$ we can construct $\mathcal{A}_r$ such that $\text{Adv}_{\text{RDS}[\text{DS}, \mathbf{R}]}^{\text{ufcma}}(\mathcal{A}_u) = \text{Adv}_{\text{DS}, \text{RD}[\mathbf{R}]}^{\text{ufcma}}(\mathcal{A}_r)$. It additionally holds that $\text{Adv}_{\text{RDS}[\text{DS}, \mathbf{R}]}^{\text{sufcma}}(\mathcal{A}_u) = \text{Adv}_{\text{DS}, \text{RD}[\mathbf{R}]}^{\text{sufcma}}(\mathcal{A}_r)$. Adversary $\mathcal{A}_r$ has essentially the same complexity as $\mathcal{A}_u$.*

Proof (Sketch). The proof of this is straightforward. If $\mathcal{A}_u$ queries $\text{SIGN}(m)$, then $\mathcal{A}_r$ queries $\text{SIGN}(m)$ and receives (σ, r) and returns $\sigma \parallel r$ to $\mathcal{A}_u$. If $\mathcal{A}_u$ queries $\text{FORGE}(m^*, \sigma^* \parallel r^*)$, then $\mathcal{A}_r$ queries $\text{FORGE}(m^* \parallel r^*, \sigma^*)$. Note that $\mathcal{A}_r$ wins whenever $\mathcal{A}_u$ would. $\square$

In independent and concurrent work, Diemert, Gellert, Jager, and Lyu [10] also considered RDS , proving that if DS can be proven SUFCMA1 secure (in this notion the game records its responses to signature queries and repeats them if the adversary repeats a query) with a restricted class of “canonical” memory-tight reductions, then there is a memory-tight reduction for the mSUFCMA security of RDS . This complements our results as they use a more restrictive assumption to prove mSUFCMA while we use a generic assumption to prove mUFCMA.

In the full version [14], we further show that if $\mathbf{D}$ is only *weakly* extractable (but still has high entropy), then we can prove a variant of Theorem 1 with a less efficient reduction. In particular, the running time of the reduction has an additional term of $q_{\text{FORGE}} \cdot q_{\text{SIGN}} \cdot \mathbf{Time}(\mathbf{D.S})$. This difference arises because rather than extracting r and computing $j \leftarrow f^{-1}(m, r)$ in FORGE we instead need to iterate over the possible values of $f(m, j)$ to check for consistency. Thus the proof for this is an instance of our inefficient tagging technique.

4.5 mSUFCMA Security of RSA-PFDH

The RSA-based Probabilistic Full-Domain Hash (RSA-PFDH) scheme, originally introduced by Coron [8], can be viewed as the result of applying the $\text{RDS}[\cdot]$ transform to RSA-based Full-Domain Hash [4] (RSA-FDH). Auerbach et al. [1] gave a memory-, but not time-tight reduction from the 1UFCMA security of RSA-FDH to the one-wayness of RSA. Applying Theorems 1 and 2 would give a memory-, but not time-tight reduction for the security of RSA-PFDH.

By a careful combination of our tagging technique with the proof ideas of Coron and of Auerbach et al. we can analyze RSA-PFDH directly. We prove the following result – a time, memory, and advantage tight reduction for the security of RSA-PFDH. The theorem is properly formalized and proven in the full version [14].

Theorem 3 (Informal, mSUFCMA security of RSA-PFDH).

Define $\text{RSA} := \text{RDS}[\text{RSA-FDH}, \{0,1\}^{\ell}]$. Let $\mathcal{A}_m$ be an adversary with $(q_{\text{SIGN}}, q_{\text{FORGE}}, q_h) = \text{Query}(\mathcal{A}_m)$. Then we can construct an adversary $\mathcal{A}_{\text{RSA}}$ against one-wayness of ℓ -bit RSA such that

$$\text{Adv}_{\text{RSA}}^{\text{sufcma}}(\mathcal{A}_m) \leq \text{Adv}^{\text{ow-rsa}}(\mathcal{A}_{\text{RSA}}) + (0.5 \cdot q_{\text{SIGN}}^2 + 2 \cdot q_{\text{SIGN}} \cdot q_{\text{FORGE}})/2^{\ell}$$

The running time and memory of $\mathcal{A}_{\text{RSA}}$ is roughly the same as $\mathcal{A}_m$.

In concurrent work, Diemert, Gellert, Jager, and Lyu [10] also give a time, memory, and advantage tight reduction for RSA-PFDH via a different proof.

5 AE Security of Encrypt-then-PRF

For nonce-based secret-key encryption schemes, we often want Authenticated Encryption (AE) security which simultaneously asks for confidentiality and ciphertext integrity. The common approach to prove AE security of a nonce-based encryption scheme is to give separate reductions to the indistinguishability of its ciphertexts from truly random ones (INDR security) and its ciphertext integrity. Ghoshal et al. [15] proved an impossibility result showing that a (certain form of black-box) reduction from AE security to INDR security and ciphertext integrity cannot be memory-tight. Making the INDR part memory-tight is of particular interest because of results which establish tight time-memory trade-offs for INDR security [9, 12, 17, 20, 21].

In this section we look at a particular scheme which we refer to as Encrypt-then-PRF. Given a nonce-based encryption scheme NE that only has INDR security, one generic way to construct a new encryption scheme NE' which also achieves ciphertext integrity is to use a PRF and let the ciphertext of NE' be the concatenation of the ciphertext of NE and a tag which is the evaluation of the PRF on the ciphertext and the nonce.

We show that in the context of Encrypt-then-PRF, for two of the notions of AE security introduced in [15], we can give a memory-tight reduction to the INDR security of the underlying encryption scheme and a non-memory-tight reduction to the security of the PRF. This shows that we can bypass the generic impossibility result of [15] if we consider specific constructions of nonce-based authenticated encryption schemes. In more detail, the impossibility result of [15] rules out lifting the INDR security of a scheme to full AE security in a memory tight way, when additionally assuming ciphertext integrity for a generic scheme. Here, we show that for the specific case of Encrypt-then-PRF schemes, lifting the INDR security of the encryption scheme to full AE security of Encrypt-then-PRF is possible in a memory-tight way, assuming security of the PRF.

Game $G_{NE,b}^{\text{indr}}(\mathcal{A})$ $K \leftarrow \$ \text{NE.K}$ $b' \leftarrow \mathcal{A}^{\text{ENC}_b}$ Return $b' = 1$	$\text{ENC}_b(n, m)$ $c_1 \leftarrow \text{NE.E}(K, n, m)$ $c_0 \leftarrow \$ \{0, 1\}^{\text{NE.cl}(m)}$ $M[n, c_b] \leftarrow m$ Return c_b	$\text{DEC}_b^w(n, c)$ If $M[n, c] \neq \perp$: Return $M[n, c]$ if $w = \mathfrak{m}$ Return $\diamond$ if $w = \diamond$ Return $\perp$ if $w = \perp$ $m_1 \leftarrow \text{NE.D}(k, n, c)$ $m_0 \leftarrow \perp$ Return m_b
Game $G_{NE,b}^{\text{ae-w}}(\mathcal{A})$ $K \leftarrow \$ \text{NE.K}$ $b' \leftarrow \mathcal{A}^{\text{ENC}_b, \text{DEC}_b^w}$ Return $b' = 1$		

Fig. 7. Games defining INDR and AE- w security of NE for $w \in \{\mathfrak{m}, \diamond, \perp\}$.

5.1 Syntax and Security Definitions

NONCE-BASED ENCRYPTION. A nonce-based (secret-key) encryption scheme NE specifies algorithms NE.K, NE.E, and NE.D. It specifies message space NE.M and nonce space NE.N. The syntax of the algorithms is shown in Fig. 8. The secret key is denoted by K , the message is m , the nonce is n , and the ciphertext is c . The decryption algorithm may return $m = \perp$ to indicate rejection of the ciphertext. Correctness requires for all $K \in [\text{NE.K}]$, $n \in \text{NE.N}$, and $m \in \text{NE.M}$ that $\text{NE.D}(K, n, \text{NE.E}(K, n, m)) = m$. We assume there is a ciphertext-length function $\text{NE.cl} : \mathbb{N} \rightarrow \mathbb{N}$ such that for all $K \in [\text{NE.K}]$, $n \in \text{NE.N}$, and $m \in \text{NE.M}$ we have $|c| = \text{NE.cl}(|m|)$ where $c \leftarrow \text{NE.E}(K, n, m)$. We define $\text{NE.C} = \bigcup_{m \in \text{NE.M}} \{0, 1\}^{\text{NE.cl}(|m|)}$. Typically, a nonce-based encryption scheme also takes associated data as input which is authenticated during encryption. This does not meaningfully affect our proof, so we omit it for simplicity.

ENCRYPT-THEN-PRF. In this section we consider the Encrypt-then-PRF construction of a nonce-based encryption scheme, due to Rogaway [19]. Namprempe et al. [18] gave a more extensive exploration of the many ways to construct an AEAD encryption scheme via generic composition. Given nonce-based encryption scheme NE and function family F, we define $\text{EtP}[\text{NE}, F]$

by the following algorithms. We refer to the t component of the ciphertext returned by $\text{EtP}[\text{NE}, F].E$ as the “tag” below. When including associated data, it would be input to F.

NE Syntax
$K \leftarrow \$ \text{NE.K}$
$c \leftarrow \text{NE.E}(K, n, m)$
$m \leftarrow \text{NE.D}(K, n, c)$

Fig. 8. Syntax of (nonce-based) secret-key encryption scheme.

$\text{EtP}[\text{NE}, \text{F}].\text{K}$	$\text{EtP}[\text{NE}, \text{F}].\text{E}(K, n, m)$	$\text{EtP}[\text{NE}, \text{F}].\text{D}(K, n, c)$
$K \leftarrow \$ \text{NE}.K$	$(K, K') \leftarrow K$	$(K, K') \leftarrow K; (c', t) \leftarrow c$
$K' \leftarrow \$ \text{F}.K$	$c' \leftarrow \text{NE}.E(K, n, m)$	If $t = \text{F}_{K'}(n, c')$:
Return (K, K')	$t \leftarrow \text{F}_{K'}(n, c')$	Return $\text{NE}.D(K, n, m)$
	Return (c', t)	Return $\perp$

Our security result will analyze the authenticated security of EtP assuming NE has ciphertexts indistinguishable from random ciphertexts and F is pseudorandom. Let us recall these security notions.

INDISTINGUISHABILITY FROM RANDOM (INDR) SECURITY. This security notion requires that ciphertexts output by the encryption scheme cannot be distinguished from random strings. Consider the game $\text{G}_{\text{NE},b}^{\text{indr}}$ defined in Fig. 7. Here an adversary $\mathcal{A}$ is given access to an encryption oracle ENC_b to which it can query a pair (n, m) and receive an honest encryption of message m with nonce n if $b = 1$ or a random string of the appropriate length if $b = 0$. We restrict attention to “valid” adversaries that never repeat the nonce n across different encryption queries. We define $\text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{A}) = \Pr[\text{G}_{\text{NE},1}^{\text{indr}}(\mathcal{A})] - \Pr[\text{G}_{\text{NE},0}^{\text{indr}}(\mathcal{A})]$.

AUTHENTICATED ENCRYPTION (AE) SECURITY. AE security simultaneously asks for integrity and confidentiality. Consider the games $\text{G}_{\text{NE},b}^{\text{ae}-w}$ which defines three variants of authenticated encryption security parameterized by $w \in \{\mathfrak{m}, \diamond, \perp\}$ shown in Fig. 7. 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 use 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 again restrict attention nonce-respecting adversaries which do not repeat nonces across encryption queries. (Note that there is no restriction placed on nonces used for decryption queries.)

The decryption oracle is parameterized by the value $w \in \{\mathfrak{m}, \diamond, \perp\}$ 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 $w = \mathfrak{m}$, the decryption oracle returns $M[n, c]$. When $w = \diamond$, it returns a special symbol $\diamond$. When $w = \perp$, it returns the symbol $\perp$ which is also used by the encryption scheme to represent rejection. For $w \in \{\mathfrak{m}, \diamond, \perp\}$ we define the advantage of an adversary $\mathcal{A}$ by $\text{Adv}_{\text{NE}}^{\text{ae}-w}(\mathcal{A}) = \Pr[\text{G}_{\text{NE},1}^{\text{ae}-w}(\mathcal{A})] - \Pr[\text{G}_{\text{NE},0}^{\text{ae}-w}(\mathcal{A})]$.

DISCUSSION OF VARIANTS. This choice of considering three variants of the definition follows the same choice made by Ghoshal et al. [15]. First off, we note that *if there are no restrictions on the memory of the adversary, all the three definitions are tightly equivalent*. An adversary can simply remember its past encryption queries and answers, and without loss of generality never make a decryption query on the answer of an encryption query. In the memory restricted setting these definitions *no longer appear to be equivalent*. The only known implication is that $w = \diamond$ security tightly implies $w = \perp$ security. Other implications seem to require remembering all encryption queries to properly simulate the decryption oracle. In Sect. 6 we parameterize public-key encryption CCA definitions similarly. This discussion applies to those definitions as well.

Ghoshal et al. argued that $w = \mathfrak{m}$ is the “correct” definition. They argue that chosen ciphertext security is intended to capture the power of an adversary that can observe the behavior of a decrypting party. Both the $w = \perp$ and $w = \diamond$ definitions restrict what the adversary learns about this behavior when honestly generated ciphertexts are forwarded, which does not seem to model anything about real use of encryption. The $w = \mathfrak{m}$ definition avoids this unnatural restriction.

We provide some technical context for this philosophical argument. In the full version [14] we give memory-tight proofs for the security of encryption schemes constructed with the KEM/DEM paradigm with $w = \mathfrak{m}$ and noting this does not seem possible for the other choices of w . In this section and Sect. 6 we prove the AE/CCA- w security of encryption schemes for differing choices of w . We view this as a general exploration of what results are possible with memory-tight proofs. A proof which works for some w , but not others helps build some understanding of how these notions related.

5.2 Security Result

Now we give a proof of the AE- $\diamond$ security of EtP[NE, F]. In particular we provide a memory-tight reduction to the INDR security of NE and a non-memory-tight reduction to the security of F. Such a result is useful if a time-memory tradeoff is known for NE and F is sufficiently secure even against high-memory attackers.

Theorem 4 (Security of EtP). *Let NE be a nonce-based encryption scheme and F be a family of function with $F.F = \text{Fcs}(\text{NE.N}, \text{NE.C}, \{0, 1\}^\tau)$ for $\tau \in \mathbb{N}$. Let $\mathcal{A}_a$ be an AE- $\diamond$ adversary with $(q_{\text{ENC}}, q_{\text{DEC}}) = \text{Query}(\mathcal{A}_a)$. Define adversaries $\mathcal{A}_p$ and $\mathcal{A}_r$ as shown in Fig. 9. Then,*

$$\begin{aligned} \text{Adv}_{\text{EtP}[\text{NE}, \text{F}]}^{\text{ae-}\diamond}(\mathcal{A}_a) &\leq \text{Adv}_{\text{F}}^{\text{pr}}(\mathcal{A}_p) + \text{Adv}_{\text{NE}}^{\text{indr}}(\mathcal{A}_r) + 2q_{\text{DEC}}/2^\tau \\ \text{Query}(\mathcal{A}_p) &= q_{\text{ENC}} + q_{\text{DEC}} & \text{Query}(\mathcal{A}_r) &= q_{\text{ENC}} \\ \text{Time}(\mathcal{A}_p) &= \text{Time}(\text{G}_{\text{EtP}[\text{NE}, \text{F}]}^{\text{ae-}\diamond}(\mathcal{A}_a)) & \text{Time}^*(\mathcal{A}_r) &= \text{Time}(\mathcal{A}_a) \\ \text{Mem}(\mathcal{A}_p) &= \text{Mem}(\text{G}_{\text{EtP}[\text{NE}, \text{F}]}^{\text{ae-}\diamond}(\mathcal{A}_a)) & \text{Mem}^*(\mathcal{A}_r) &= \text{Mem}(\mathcal{A}_a). \end{aligned}$$

Adversary $\mathcal{A}_r$ is an F.F-oracle adversary.

The standard (not memory-tight) proof of the security of EtP begins identically to our proof; we start in $\text{G}_{\text{EtP}[\text{NE}, \text{F}], 1}^{\text{ae-}\diamond}$ replace the use of F with a truly random function (using $\mathcal{A}_p$) and then information theoretically argue that the attacker shall be incapable of creating any forgeries. In the standard proof we would transition to a game where the decryption oracle is exactly that of $\text{DEC}_0^\diamond$, i.e. it always returns $\perp$ when $M[n, c] = \perp$. Then we reduce to the security of NE to replace the generated ciphertexts with random. However this standard reduction will not be memory-tight because the attacker must store the table $M[\cdot, \cdot]$ to know whether it should return $\diamond$ or $\perp$ when simulating decryption queries.⁴

⁴ Note this *would* be memory-tight for AE- $\perp$ security.

Adversary $\mathcal{A}_p^{\text{EV}}$	$\text{SIMENC}(n, m)$	$\text{SIMDEC}(n, c)$
$K \leftarrow \$\text{NE.K}$	$c' \leftarrow \text{NE.E}(K, n, m)$	If $M[n, c] \neq \perp$: Return $\diamond$
$b' \leftarrow \mathcal{A}_a^{\text{SIMENC}, \text{SIMDEC}}$	$t \leftarrow \text{EV}(n, c')$	$(t, c') \leftarrow c$
Return b'	$c \leftarrow (t, c')$	If $t = \text{EV}(n, c')$:
	$M[n, c] \leftarrow m$	Return $\text{NE.D}(k, n, c')$
	Return c	Return $\perp$

Adversary $\mathcal{A}_r^{\text{ENC}}$	$\text{SIMENC}(n, m)$	$\text{SIMDEC}(n, c)$
$f \leftarrow \$\text{Fcs}(\text{NE.N}, \text{NE.M}, \{0, 1\}^\tau)$	$c' \leftarrow \text{ENC}(n, m)$	$(t, c') \leftarrow c$
$b' \leftarrow \mathcal{A}_a^{\text{SIMENC}, \text{SIMDEC}}$	$t \leftarrow f(n, c')$	If $t = f(n, c')$:
Return b'	$c \leftarrow (t, c')$	Return $\diamond$
	Return c	Return $\perp$

Fig. 9. Adversaries used for proof of Theorem 4.

Instead we first transition to a world where F has been replaced by the random function f and DEC always returns $\diamond$ when given a ciphertext with a correct tag. (Which we can do because either $M[n, c] \neq \perp$ held or the attacker managed to guess a random tag, which is unlikely.) Now we can make our INDR reduction memory-tight. It forwards encryption queries to its encryption oracle and then uses its own function f to create the tag. For decryption queries it checks $f(n, c') = t$, returning $\diamond$ if so and $\perp$ otherwise. Then we can finally conclude by switching to the decryption oracle $\text{DEC}_0^\diamond$ by arguing that noticing this change requires guessing a random tag.

The full proof is given in the full version [14].

It does not seem possible to extend this proof technique to AE-m security because the tag would be too short to embed values of m we need to remember.

6 Chosen Ciphertext Security of Public Key Encryption

Now we apply our techniques to give memory-tight reductions between single- and multi-challenge notions of chosen-ciphertext security. The standard reduction bounds the advantage of an adversary making q_{ENC} encryption queries by q_{ENC} times the advantage of an adversary making 1 query. The reduction requires memory linear in q_{ENC} and so is not memory-tight.⁵ In Sect. 6.1, we consider the most common “left-or-right” definition of CCA security and introduce three different variants (mirroring the three notions for AE security in Sect. 5). We give a memory-tight reduction between single- and multi-challenge security for two of

⁵ Auerbach et al. [1] stated that this reduction is memory-tight for both CPA and CCA security. While the former is correct, the latter depends on the definition of CCA. In personally communication with Auerbach et al. [2], they concurred that their claim was incorrect for their intended definition of CCA security ($w = \diamond$) but pointed out that it does work for an “exclusion” variant, $w = E$, which we discuss in the full version [14].

Game $G_{\text{PKE},b}^{\text{cca-}w}(\mathcal{A})$	$\text{ENC}_b(m_0, m_1)$	$\text{DEC}^w(c)$
$(ek, dk) \leftarrow \$ \text{PKE.K}$ $b' \leftarrow \mathcal{A}^{\text{ENC}_b, \text{DEC}}(ek)$ Return $b' = 1$	$// m_0 = m_1 $ $c_0 \leftarrow \$ \text{PKE.E}(ek, m_0)$ $c_1 \leftarrow \$ \text{PKE.E}(ek, m_1)$ $M[c_b] \leftarrow m_1$ Return c_b	If $M[c] \neq \perp$: Return $M[c]$ if $w = \mathfrak{m}$ Return $\diamond$ if $w = \diamond$ Return $\perp$ if $w = \perp$ $m \leftarrow \text{PKE.D}(dk, c)$ Return m

Fig. 11. Game defining CCA- w security of PKE for $w \in \{\mathfrak{m}, \diamond, \perp\}$.

the three variants ($\diamond$ and $\perp$), but the reduction is not time-tight. In Sect. 6.2, we look at the CCA security variant that requires ciphertexts be indistinguishable from random. We give a memory-tight and time-tight reduction between single- and multi-challenge security for all three variants of this notion.

PUBLIC KEY ENCRYPTION. A public key encryption scheme PKE specifies algorithms PKE.K, PKE.E, and PKE.D. The syntax of these algorithms is shown in Fig. 10. The key generation algorithm PKE.K returns encryption key ek and decryption key dk . The encryption algorithm PKE.E encrypts message m with ek to produce a ciphertext c . We write $\text{PKE.E}(ek, m; r)$ when making random coins $r \in \text{PKE.R}$ explicit. The decryption algorithm decrypts c with dk to produce m . The decryption algorithm may output $m = \perp$ to indicate rejection.

PKE Syntax
$(ek, dk) \leftarrow \$ \text{PKE.K}$
$c \leftarrow \$ \text{PKE.E}(ek, m)$
$m \leftarrow \text{PKE.D}(dk, c)$

Fig. 10. Syntax of a public key encryption scheme PKE.

Correctness requires that $\text{PKE.D}(dk, c) = m$ for all $(ek, dk) \in [\text{PKE.K}]$, all m , and all $c \in [\text{PKE.E}(ek, m)]$. We define the min-entropy of PKE as

$$\text{PKE.H}_\infty = -\lg \max_{m, ek, c} \Pr[r \leftarrow \$ \text{PKE.R} : \text{PKE.E}(ek, m; r) = c].$$

6.1 Left-or-Right CCA Security of PKE

LEFT-OR-RIGHT CCA SECURITY. In this section, we consider the left-or-right definition of CCA-security most commonly used in the literature. For $w \in \{\mathfrak{m}, \diamond, \perp\}$ we denote this as CCA- w ⁶ and the corresponding security game $G_{\text{PKE},b}^{\text{cca-}w}$ is defined in Fig. 11. The adversary gets the encryption key ek and has access to an encryption and a decryption oracle. The encryption oracle takes in messages m_0 and m_1 and encrypts m_b where b is the secret bit. The decryption oracle returns the decryption of a ciphertext, unless the ciphertext was previously returned by an encryption query. This is tracked by table M .

⁶ The discussion in Sect. 5 about the choice to have three variants of the definitions is applicable here as well.

Adversary $\mathcal{A}_1^{\text{ENC,DEC}}(ek)$	$\text{SIMENC}(m_0, m_1)$	$\text{SIMDEC}(c)$
$k \leftarrow \$[q_{\text{ENC}}]$	$i \leftarrow i + 1$	If $c = c^*$: Return $\diamond$
$i \leftarrow 0$	For $d \in \{0, 1\}$:	$m \leftarrow \text{DEC}(c)$
$D_{(\cdot)} \leftarrow \{0, 1\}^{(\cdot)} \times [q_{\text{ENC}}]$	$r_d \leftarrow f(m_d , (m_d, i))$	If $m = \perp$: Return $\perp$
$f \leftarrow \$ \text{Fcs}(\mathbb{N}, D, \text{PKE.R})$	$c_d \leftarrow \text{PKE.E}(ek, m_d; r_d)$	For $j \in [i]$ do:
$b' \leftarrow \mathcal{A}_m^{\text{SIMENC, SIMDEC}}(ek)$	If $i < k$: $c \leftarrow c_1$	If $m \in \{m_0^*, m_1^*\}$ and $j = k$:
Return b'	If $i = k$:	Skip to next j
	$c \leftarrow \text{ENC}(m_0, m_1)$	$r \leftarrow f(m , (m, j))$
	$c^* \leftarrow c$	If $\text{PKE.E}(ek, m; r) = c$:
	$(m_0^*, m_1^*) \leftarrow (m_0, m_1)$	Return $\diamond$
	If $i > k$: $c \leftarrow c_0$	Return m
	Return c	

Fig. 12. Adversary $\mathcal{A}_1$ for Theorem 5.

When $w = \mathfrak{m}$, the decryption oracle returns $M[c]$ which is m_1 from the earlier encryption query. When $w = \diamond$, it returns $\diamond$. When $w = \perp$, it returns $\perp$ which is also used by the encryption scheme to represent rejection. The advantage of an adversary $\mathcal{A}$ against the CCA- w security of PKE is defined as $\text{Adv}_{\text{PKE}}^{\text{cca-}w}(\mathcal{A}) = \Pr[\text{G}_{\text{PKE},1}^{\text{cca-}w}(\mathcal{A})] - \Pr[\text{G}_{\text{PKE},0}^{\text{cca-}w}(\mathcal{A})]$.

The goal of this section is to relate the advantage of attacks making only a single encryption query and those making many such queries. When wanting to make the distinction explicit we may use the adjectives ‘many’ and ‘single’ or prefix the abbreviation of a security notion with an ‘m’ or ‘1’.

1CCA- $\diamond$ IMPLIES mCCA- $\diamond$. The following theorem gives a memory-tight reduction establishing that CCA- $\diamond$ security against adversaries making one encryption query implies security for an arbitrary number of queries. The proof makes use of our inefficient tagging technique. The reduction performs a hybrid over the encryption queries of the original adversary and is thus not advantage-tight.

Theorem 5. (1CCA- $\diamond \Rightarrow$ mCCA- $\diamond$). *Let PKE be a public key encryption scheme. Let $\mathcal{A}_m$ be an adversary with $(q_{\text{ENC}}, q_{\text{DEC}}) = \text{Query}(\mathcal{A}_m)$. Define $D_{(\cdot)}$ by $D_n = \{0, 1\}^n \times [q_{\text{ENC}}]$. Let $\mathcal{A}_1$ be the $\text{Fcs}(\mathbb{N}, D, \text{PKE.R})$ -oracle adversary shown in Fig. 12. Then,*

$$\text{Adv}_{\text{PKE}}^{\text{cca-}\diamond}(\mathcal{A}_m) \leq q_{\text{ENC}} \cdot \text{Adv}_{\text{PKE}}^{\text{cca-}\diamond}(\mathcal{A}_1) + 4 \cdot q_{\text{ENC}} \cdot q_{\text{DEC}} / 2^{\text{PKE.H}\infty}$$

$$\text{Query}(\mathcal{A}_1) = (1, q_{\text{DEC}})$$

$$\text{Time}^*(\mathcal{A}_1) = O(\text{Time}(\mathcal{A}_m)) + q_{\text{ENC}}(q_{\text{DEC}} + 1)\text{Time}(\text{PKE})$$

$$\text{Mem}^*(\mathcal{A}_1) = O(\text{Mem}(\mathcal{A}_m)) + \text{Mem}(\text{PKE}) + \lg q_{\text{ENC}}.$$

The standard (non-memory-tight) reduction against 1CCA security picks an index $k \in [q_{\text{ENC}}]$ where q_{ENC} is the number of encryption queries made by $\mathcal{A}_m$. It runs $\mathcal{A}_m$, simulating encryption queries as follows. For the first $k - 1$ encryption queries, it answers with an encryption of m_1 , for the k -th encryption query it forwards the query to its own encryption oracle, and the rest of the queries it

Game $G_{\text{PKE},b}^{\text{cca-}w}(\mathcal{A})$	$\text{ENC}_b(m)$	$\text{DEC}^w(c)$
$(ek, dk) \leftarrow \text{PKE.K}$ $b' \leftarrow \mathcal{A}^{\text{ENC}_b, \text{DEC}^w}(ek)$ Return $b' = 1$	$c_1 \leftarrow \text{PKE.E}(ek, m)$ $c_0 \leftarrow \text{PKE.C}(ek, m)$ $M[c_b] \leftarrow m$ Return c_b	If $M[c] \neq \perp$: Return $M[c]$ if $w = \mathfrak{m}$ Return $\diamond$ if $w = \diamond$ Return $\perp$ if $w = \perp$ $m \leftarrow \text{PKE.D}(dk, c)$ Return m

Fig. 13. Game defining $\text{\$CCA-}w$ security of PKE for $w \in \{\mathfrak{m}, \diamond, \perp\}$.

answers with an encryption of m_0 . To answer the decryption queries, the reduction returns $\diamond$ if it was ever queried the ciphertext for a previous encryption query. Otherwise, it forwards the query to its own decryption oracle. Finally, the reduction adversary outputs whatever $\mathcal{A}_{\mathfrak{m}}$ outputs. Standard hybrid analysis shows that if the advantage of $\mathcal{A}_{\mathfrak{m}}$ is ϵ , then the advantage of the reduction adversary is ϵ/q_{ENC} . Simulating decryption queries required remembering all prior encryption queries and hence the reduction is not memory-tight.

We give an adversary $\mathcal{A}_1$ in Fig. 12 that is very similar to the reduction just described, but avoids remembering prior encryption queries. The main idea is that it picks the coins when encrypting m_0 or m_1 locally as the output of a random function f applied to the message and a counter. This allows $\mathcal{A}_1$ to detect whether a ciphertext c queried to the decryption oracle is one it answered to an earlier encryption query as follows: it first asks for the decryption of c from its own decryption oracle and receives m . Then it iterates over all counter values for which encryption queries have been made so far and checks if c was the encryption of m using the output of f on m and the counter as coins. If any of these checks succeed it returns $\diamond$, otherwise it returns m . If c was the answer of an encryption query $\mathcal{A}_1$ detects it successfully. The probability that $\mathcal{A}_1$ returns $\diamond$ for a decryption query when it should not is small. We give the formal proof of Theorem 5 in the full version [14] where we use a sequence of hybrid games to transition from $G_{\text{PKE},b}^{\text{cca-}\diamond}$ to a hybrid game that is simulated by $\mathcal{A}_1$.

Notice that the additional memory overhead for $\mathcal{A}_1$ is just that required to store a counter, run PKE.E , and store (c^*, m_0^*, m_1^*) . However, there is an *increase* in runtime by $q_{\text{ENC}} \cdot q_{\text{DEC}} \cdot \text{Time}(\text{PKE})$ because of the iteration over the counters to answer decryption queries. As discussed in the introduction, such reductions may be useful when the best attack for the underlying problem with low memory requires significantly more running time than the best attack with high memory.

The same proof strategy would work essentially unchanged for $\text{CCA-}\perp$. For $\text{CCA-}\mathfrak{m}$, the strategy does not suffice. If the adversary queries the decryption oracle on a ciphertext c which was an answer to a previous query for (m_0, m_1) the oracle needs to return m_1 even if c is an encryption of m_0 .

6.2 Indistinguishable from Random CCA Security of PKE

We saw in the previous section that we could have a memory-tight reduction from $\text{mCCA-}\diamond$ to $\text{1CCA-}\diamond$; however, the reduction is not tight with respect to running time. In this section, we show that for a different formalization of CCA security, we can indeed have a memory-tight and time-tight reduction between many- and single-challenge variants.

CIPHERTEXT AND ENCRYPTION KEY SPACE. Before describing the indistinguishable from random formalization of CCA security, we need to make some assumptions on PKE. We define the encryption keyspace by $\text{PKE.Ek} = \{ek : (ek, dk) \in \text{PKE.K}\}$. We assume for each $ek \in \text{PKE.Ek}$ and allowed message length $n \in \mathbb{N}$ there is a set $\text{PKE.C}(ek, n)$ such that $\text{PKE.E}(ek, m; r) \in \text{PKE.C}(ek, |m|)$ always holds. Let $\text{PKE.C}^{-1}(ek, c)$ returns n such that $c \in \text{PKE.C}(ek, n)$. Correctness implies that $\text{PKE.C}(ek, n)$ and $\text{PKE.C}(ek, n')$ are disjoint for $n \neq n'$.

INDISTINGUISHABLE FROM RANDOM CIPHERTEXT CCA SECURITY. The security notion we will consider in this section is captured by the game $\text{G}^{\text{cca-}w}_{\text{PKE}}$ shown in Fig. 13. It requires that ciphertexts output by the encryption scheme cannot be distinguished from ciphertexts chosen at random even given access to a decryption oracle. The adversary gets the encryption key ek and has access to an encryption oracle ENC and a decryption oracle DEC . The adversary needs to distinguish the following real and ideal worlds: in the real world, a query to ENC with a message m returns an encryption of m under ek , while in the ideal world, the same query returns a uniformly random element of $\text{PKE.C}(ek, |m|)$. The decryption oracle DEC^w acts exactly as the corresponding oracle in $\text{G}^{\text{cca-}w}_{\text{PKE}}$.⁷ The advantage of an adversary $\mathcal{A}$ against the $\text{\$CCA-}w$ security of PKE is defined as $\text{Adv}^{\text{\$cca-}w}_{\text{PKE}}(\mathcal{A}) = \Pr[\text{G}^{\text{\$cca-}w}_{\text{PKE},1}(\mathcal{A})] - \Pr[\text{G}^{\text{\$cca-}w}_{\text{PKE},0}(\mathcal{A})]$.

1\\$CCA-m IMPLIES m\\$CCA-m. The following theorem captures a memory-tight reduction establishing that $\text{1\$CCA-m}$ security implies $\text{m\$CCA-m}$ security. The proof makes use of our message encoding technique.

Theorem 6 ($\text{1\$CCA-m} \Rightarrow \text{m\$CCA-m}$). *Let PKE be a public key encryption scheme. Let τ satisfy $|\text{PKE.C}(ek, n)| \geq 2^n \cdot 2^\tau$ for all n, ek . Let $\mathcal{A}_m$ be an adversary with $(q_{\text{ENC}}, q_{\text{DEC}}, q_h) = \text{Query}(\mathcal{A}_m)$ and assume $q_{\text{ENC}} + q_{\text{DEC}} \leq 0.5 \cdot 2^\tau$. Let $\mathcal{F} = \text{Inj}^\pm(T, D, R)$ where T, D , and R are defined by $T = \mathbb{N} \times \text{PKE.Ek}$, $D_{n,ek} = \{0, 1\}^n \times [q_{\text{ENC}}]$ and $R_{n,ek} = \text{PKE.C}(ek, n)$. Let $\mathcal{A}_1$ be the $\mathcal{F}$ -oracle adversary defined in Fig. 14. Then,*

$$\text{Adv}^{\text{\$cca-}m}_{\text{PKE}}(\mathcal{A}_m) \leq q_{\text{ENC}} \cdot \text{Adv}^{\text{\$cca-}m}_{\text{PKE}}(\mathcal{A}_1) + 8q_{\text{ENC}}q_{\text{DEC}}/2^\tau + 5q_{\text{ENC}}^2/2^\tau$$

$$\text{Query}(\mathcal{A}_1) = (1, q_{\text{DEC}}, q_h)$$

$$\text{Time}^*(\mathcal{A}_1) = O(\text{Time}(\mathcal{A}_m)) + q_{\text{ENC}} \text{Time}(\text{PKE})$$

$$\text{Mem}^*(\mathcal{A}_1) = O(\text{Mem}(\mathcal{A}_m)) + \text{Mem}(\text{PKE}) \lg q_{\text{ENC}}.$$

⁷ As mentioned, the discussion in Sect. 5 about the three variants definitions is applicable here as well. In the full version [14] we give an example where we can prove CCA security of a KEM/DEM scheme in the memory restricted setting, but only if we use the $w = \text{m}$ definition.

Adversary $\mathcal{A}_1^{\text{ENC,DEC}}(ek)$	$\text{SIMENC}(m)$	$\text{SIMDEC}(c)$
$// 0 \leq h \leq 2$	$i \leftarrow i + 1$	If $c = c^*$: Return m^*
$k \leftarrow \$ [q_{\text{ENC}}]$	$c_1 \leftarrow \$ \text{PKE.E}(ek, m)$	$n \leftarrow \text{PKE.C}^{-1}(ek, c)$
$i \leftarrow 0$	$c_0 \leftarrow f(m , ek), (m, i)$	$(m, j) \leftarrow f^{-1}(n, ek), c$
$f \leftarrow \$ \text{Inj}^\pm(T, D, R)$	If $i < k$: $c \leftarrow c_1$	If $m \neq \perp$ and $k \leq j \leq i$:
$b' \leftarrow \mathcal{A}_m^{\text{SIMENC, SIMDEC}}(ek)$	If $i = k$:	If $(m, j) = (m^*, k)$:
Return b'	$c \leftarrow \text{ENC}(m)$	Skip next line
	$(c^*, m^*) \leftarrow (c, m)$	Return m
	If $i > k$: $c \leftarrow c_0$	$m \leftarrow \text{DEC}(c)$
	Return c	Return m

Fig. 14. Adversary $\mathcal{A}_1$ for Theorem 6.

The standard (non-memory-tight) reduction against 1\$CCA security that runs an m\$CCA adversary $\mathcal{A}_m$ works in a similar manner as the standard reduction from an 1CCA adversary and an mCCA adversary that we described in Sect. 6.1. Again here, simulating decryption queries requires remembering all the answers of the encryption queries, and hence the reduction is not memory-tight.

We give an adversary $\mathcal{A}_1$ in Fig. 14 that is very similar to the standard reduction, but avoids remembering all the answers of the encryption queries. The main idea here is picking the ciphertext c_0 as the output of a random injective function f evaluated on the message and a counter, instead of sampling it uniformly at random. This way of picking the c_0 allows $\mathcal{A}_1$ detect whether a ciphertext c queried to the decryption oracle was the answer to an earlier encryption query as follows: it first checks if the inverse of f on the ciphertext is defined (i.e., not $\perp$), it returns the message part of the inverse. Otherwise it asks for the decryption of the ciphertext to its own decryption oracle and returns the answer. Using our assumption on the size of $\text{PKE.C}(ek, n)$, we can argue that except with small probability, $\mathcal{A}_1$ simulates the decryption oracle correctly. We give the formal proof in the full version [14] where we use a sequence of hybrid games to transition from $G_{\text{PKE}, b}^{\text{scca-m}}$ to a game that is perfectly simulated by $\mathcal{A}_1$.

The additional memory overhead for $\mathcal{A}_1$ is only a counter. Moreover, there is no increase in the running time of $\mathcal{A}_1$ unlike the adversary in Theorem 5.

EXTENSION TO \$CCA- $\diamond$, \$CCA- $\perp$. We can prove the same result for \$CCA- $\diamond$, \$CCA- $\perp$ but the adversary would not be tight with respect to running time. The adversary in these cases would pick the coins for encrypting m (to compute c_1) like the adversary in Theorem 5. This would require iterating over counters to answer decryption queries and hence lead to looseness with respect to running time. We omit the theorems for these notions because they would not involve any new ideas beyond those presented in Theorems 5 and 6.

Acknowledgements. Ashrujit Ghoshal, Joseph Jaeger, and Stefano Tessaro were partially supported by NSF grants CNS-1930117 (CAREER), CNS-1926324, CNS-2026774, a Sloan Research Fellowship, and a JP Morgan Faculty Award. Joseph Jaeger's work was done while at the University of Washington.

References

1. Auerbach, B., Cash, D., Fersch, M., Kiltz, E.: Memory-tight reductions. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017. LNCS, vol. 10401, pp. 101–132. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63688-7_4
2. Auerbach, B., Cash, Fersch, M., Kiltz, E.: Personal communication (2021)
3. Bellare, M., Namprempre, C.: Authenticated encryption: relations among notions and analysis of the generic composition paradigm. *J. Cryptol.* **21**(4), 469–491 (2008)
4. Bellare, M., Rogaway, P.: Random oracles are practical: a paradigm for designing efficient protocols. In: Denning, D.E., Pyle, R., Ganesan, R., Sandhu, R.S., Ashby, V. (eds.) ACM CCS, vol. 93, pp. 62–73. ACM Press (1993)
5. 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
6. Bernstein, D.J.: Extending the salsa20 nonce. In: Workshop record of Symmetric Key Encryption Workshop, vol. 2011. Citeseer (2011)
7. Bhattacharyya, R.: Memory-tight reductions for practical key encapsulation mechanisms. In: Kiayias, A., Kohlweiss, M., Wallden, P., Zikas, V. (eds.) PKC 2020. LNCS, vol. 12110, pp. 249–278. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45374-9_9
8. Coron, J.-S.: Optimal security proofs for PSS and other signature schemes. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 272–287. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-46035-7_18
9. Dai, W., Tessaro, S., Zhang, X.: Super-linear time-memory trade-offs for symmetric encryption. In: Pass, R., Pietrzak, K. (eds.) TCC 2020. LNCS, vol. 12552, pp. 335–365. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-64381-2_12
10. Diemert, D., Gellert, K., Jager, T., Lyu, L.: Digital signatures with memory-tight security in the multi-challenge setting. In: Tibouchi, M., Wang, H. (eds.) ASIACRYPT 2021. LNCS, vol. 13093, pp. 403–433. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-92068-5_14
11. Dinur, I.: On the streaming indistinguishability of a random permutation and a random function. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020. LNCS, vol. 12106, pp. 433–460. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45724-2_15
12. Dinur, I.: Tight time-space lower bounds for finding multiple collision pairs and their applications. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020. LNCS, vol. 12105, pp. 405–434. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45721-1_15
13. Even, S., Goldreich, O., Micali, S.: On-line/off-line digital signatures. *J. Cryptol.* **9**(1), 35–67 (1996). <https://doi.org/10.1007/BF02254791>
14. Ghoshal, A., Ghosal, R., Jaeger, J., Tessaro, S.: Hiding in plain sight: memory-tight proofs via randomness programming. *Cryptology ePrint Archive*, Report 2021/1409 (2021). <https://eprint.iacr.org/2021/1409>
15. Ghoshal, A., Jaeger, J., Tessaro, S.: The memory-tightness of authenticated encryption. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020. LNCS, vol. 12170, pp. 127–156. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-56784-2_5
16. Ghoshal, A., Tessaro, S.: On the memory-tightness of hashed elgamal. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020. LNCS, vol. 12106, pp. 33–62. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45724-2_2

17. Jaeger, J., Tessaro, S.: Tight time-memory trade-offs for symmetric encryption. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019. LNCS, vol. 11476, pp. 467–497. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17653-2_16
18. Namprepmpre, C., Rogaway, P., Shrimpton, T.: Reconsidering generic composition. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 257–274. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-55220-5_15
19. 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
20. Shahaf, I., Ordentlich, O., Segev, G.: An information-theoretic proof of the streaming switching lemma for symmetric encryption. In: 2020 IEEE International Symposium on Information Theory (ISIT), pp. 858–863 (2020)
21. Tessaro, S., Thiruvengadam, A.: Provable time-memory trade-offs: symmetric cryptography against memory-bounded adversaries. In: Beimel, A., Dziembowski, S. (eds.) TCC 2018. LNCS, vol. 11239, pp. 3–32. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-03807-6_1
22. Wang, Y., Matsuda, T., Hanaoka, G., Tanaka, K.: Memory lower bounds of reductions revisited. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018. LNCS, vol. 10820, pp. 61–90. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78381-9_3