



# Handling Adaptive Compromise for Practical Encryption Schemes

Joseph Jaeger<sup>1</sup>(✉) and Nirvan Tyagi<sup>2</sup>

<sup>1</sup> Paul G. Allen School of Computer Science & Engineering,  
University of Washington, Seattle, USA

[jsjaeger@cs.washington.edu](mailto:jsjaeger@cs.washington.edu)

<sup>2</sup> Cornell Tech, New York City, USA

[tyagi@cs.cornell.edu](mailto:tyagi@cs.cornell.edu)

**Abstract.** We provide a new definitional framework capturing the multi-user security of encryption schemes and pseudorandom functions in the face of adversaries that can adaptively compromise users' keys. We provide a sequence of results establishing the security of practical symmetric encryption schemes under adaptive compromise in the random oracle or ideal cipher model. The bulk of analysis complexity for adaptive compromise security is relegated to the analysis of lower-level primitives such as pseudorandom functions.

We apply our framework to give proofs of security for the BurnBox system for privacy in the face of border searches and the in-use searchable symmetric encryption scheme due to Cash et al. In both cases, prior analyses had bugs that our framework helps avoid.

**Keywords:** Symmetric cryptography · Ideal models · Adaptive security · Searchable symmetric encryption

## 1 Introduction

A classic question in cryptography has been dealing with adversaries that adaptively compromise particular parties, thereby learning their secrets. Consider a setting where parties use keys  $k_1, \dots, k_n$  to encrypt messages  $m_1, \dots, m_n$  to derive ciphertexts  $\text{Enc}(k_1, m_1), \dots, \text{Enc}(k_n, m_n)$ . An adversary obtains the ciphertexts and compromises a chosen subset of the parties to learn their keys. What can we say about the security of the messages encrypted by the keys that remain secret? Surprisingly, traditional approaches to formal security analysis, such as using encryption schemes that provide semantic security [19], fail to suffice for proving these messages' confidentiality. This problem was first discussed in the context of secure multiparty computation [10], and it arises in a variety of important cryptographic applications, as we explain below.

In this work, we introduce a new framework for formal analyses when security in the face of adaptive compromise is desired. Our approach provides a modular route towards analysis using idealized primitives (such as random oracles or ideal ciphers) for practical and in-use schemes. This modularity helps


us sidestep the pitfalls of prior ideal-model analyses that either invented new (less satisfying) ideal primitives, omitted proofs, or gave detailed but incorrect proofs. We exercise our framework across applications including searchable symmetric encryption (SSE), revocable cloud storage, and asymmetric password-authenticated key exchange (aPAKE). In particular, we provide full, correct proofs of security against adaptive adversaries for the Cash et al. [12] searchable symmetric encryption scheme that is used often in practice and the Burn-Box system [33] for dealing with compelled-access searches. We show that our new definitions imply the notion of equivocable encryption introduced to prove security of the OPAQUE [24] asymmetric password-authenticated key exchange protocol. More broadly, our framework can be applied to a wide variety of constructions [1, 2, 9, 13, 17, 20, 21, 25–29, 34].

**Current approaches to the “commitment problem”.** Our motivating applications have at their core an adaptive simulation-based model of security. Roughly speaking, they ask that no computationally bound adversary can distinguish between two worlds. In the first world, the adversary interacts with the scheme whose security is being measured. In the second world, the “ideal” world, the adversary’s queries are instead handled by a simulator that must make do with only limited information which represents allowable “leakage” about the queries the adversary has made so far. The common unifying factor between varying applications we consider is that the adversary can make queries resulting in being given a ciphertexts encrypting messages of its choosing, then with future queries adaptively choose to expose the secret keys underlying some of the ciphertexts. The leakage given to the simulator will not include the messages encrypted unless a query has been made to expose the corresponding key.

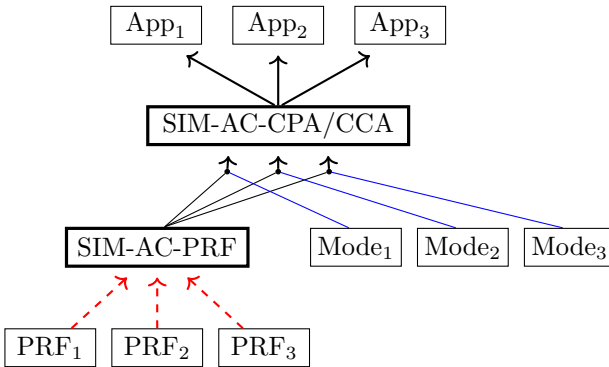
Proving security in this model, however, does not work based on standard assumptions of the underlying encryption scheme. The problem is that the simulator must commit to ciphertexts, revealing them to the adversary, before knowing the messages associated to them. Hence the commitment problem. Several prior approaches for proving positive security results exist.

One natural approach attempts to build special non-committing encryption schemes [10] that can be proven (in the standard model) to allow opening some a priori fixed ciphertext to a message. But these schemes are not practical, as they require key material at least as long as the underlying message. Another unsatisfying approach considers only non-adaptive security in which an attacker specifies all of its queries at the beginning of the game. This is one of the two approaches that were simultaneously taken by Cash et al. [12]. Here the simulator is given the leakage for all of these queries at once and generates a transcript of all of its response. This is unsatisfying because more is lost when switching from adaptive to non-adaptive security than just avoiding the commitment problem. It is an easy exercise to construct encryption schemes which are secure when all queries to it must be chosen ahead of time but are not secure even against key-recovery attacks when an adversary may adaptively choose its queries.

The primary approach used to avoid this is to use idealized models, which we can again split into two versions. The first is to use an idealized model of



**Fig. 1.** Old state of affairs. Red dashed lines correspond to implications proved through programming in an ideal model proof. A different programming proof is needed to prove an application secure for each pair of PRF and symmetric encryption mode. (Color figure online)



**Fig. 2.** New state of affairs. Red dashed lines correspond to implications proved through programming in an ideal model proof. New definitions are in bold boxes. Programming proofs are only needed to show each low level PRF construction meets SIM-AC-PRF. (Color figure online)

encryption. Examples of this include indistinguishable authenticated encryption [3] (IAE) or the ideal encryption model (IEM) of Tyagi et al. [33]. Security analyses in these models might not say much when one uses real encryption schemes, even when one is willing to use more established idealized models such as the ideal cipher model (ICM) or the random oracle model (ROM). One hope would be to use approaches such as indistinguishability [30] to modularly show that symmetric schemes sufficiently “behave like” ideal encryption, but this approach is unlikely to work for most encryption schemes used in practice [3].

The final approach, which is by far the most common in searchable symmetric encryption [1, 2, 9, 12, 13, 17, 20, 21, 25–29, 34], is to fix a *particular encryption scheme* and prove security with respect to it in the ICM or ROM. Typically encryption schemes are built as modes of operations of an underlying pseudo-random function (PRF) and this function (or its constituent parts) is what is modeled as an ideal function. The downside of this is represented in Fig. 1. On the top, we have the applications one would like to prove secure, and on the

bottom, we have the different modes of operation and PRFs that one might use. Using this approach means that for each application, we have to provide a separate ideal model proof for each different choice of a mode of operation and a PRF (represented by dotted red arrows in Fig. 1). If there are  $A$  applications,  $P$  PRFs, and  $M$  modes of operation one might consider using, then this requires  $A \cdot P \cdot M$  ideal model proofs in total, an unsatisfying state of affairs.

Moreover, the required ideal analysis can be tedious<sup>1</sup> and error-prone. This is presumably why only a few of the papers we found actually attempt to provide the full details of the ROM proof. We have identified bugs in all of the proofs that did provide the details. The lack of a full, valid proof among the fifteen papers we considered indicates the need for a more modular framework to capture this use of the random oracle. Our work provides such a framework, allowing the random oracle details to be abstracted away as a proof that only needs to be provided once. This framework provides definitions for use by other cryptographers that are simple to use, apply to practical encryption schemes, and allow showing adaptive security in well-studied models.

**Examples of the “commitment problem”.** We proceed by discussing the example applications where we will apply our framework.

*Revocable cloud storage and the compelled access setting.* We start with the recently introduced compelled access setting (CAS) [33]. Here one wants encryption tools that provide privacy in the face of an authority searching digital devices, e.g., government searches of mobile phones or laptops at border crossings. To protect against compelled access searches, the BurnBox tool [33] uses what they call revocable encryption. At its core, this has the system encrypt a user’s files  $m_1, \dots, m_n$  with independent keys  $k_1, \dots, k_n$ . Ciphertexts are stored on (adversarially visible) cloud storage. Before a search occurs, the user instructs the application to delete the keys corresponding to files that the user wishes to hide from the authority, thereby revoking their own access to them. The other keys and file contents are disclosed to the authority.

The formal security definition introduced by Tyagi et al. captures confidentiality for revoked files even in the face of adversarial choice of which files to revoke, meaning they want security in the face of adaptive compromises. This very naturally results in the commitment problem because the simulator can be forced to provide ciphertexts for files, but only later learn the contents of these files at the time of key revelation. At which point, it is supposed to give keys which properly decrypt these ciphertexts.

To address the commitment problem they introduced the IEM. This models symmetric encryption as perfect: every encryption query is associated to a freshly chosen random string as ciphertext, and decryption is only allowed on previously returned ciphertexts. Analyses in the IEM can commit to ciphertexts (when the adversary doesn’t know the key) and later open them to arbitrary messages. In their implementation, they used AES-GCM for encryption which cannot be

---

<sup>1</sup> Even more-so because SSE protocols often *also* run into the commitment problem with a PRF and need to model that using a random oracle as well.

thought of as indifferentiable from the IEM. Hence their proof can ultimately only provide heuristic evidence for the security of their implementation.

*Symmetric searchable encryption.* Our second motivating setting is symmetric searchable encryption (SSE), which has similar issues as that discussed above for BurnBox, but with added complexity. SSE handles the following problem: a client wants to offload storage of a database of documents to an untrusted server while maintaining the ability to perform keyword searches on the database. The keyword searches should not reveal the contents of the documents to the server. To enable efficient solutions, we allow queries to leak some partial information about documents. Security is formalized via a simulation-based definition [15], in which a simulator given only the allowed leakage must trick an adversary into believing it is interacting with the actual SSE construction. An adaptive adversary can submit keyword searches as a function of prior returned results. Proving security here establishes that the scheme only leaks what is allowed and nothing more. While the leakage itself has been shown to be damaging in various contexts [11, 22], our focus here is on the formal analyses showing that leakage-abuse attacks are the best possible ones.

A common approach for SSE can be summarized at a high level as follows. The client generates a sequence of key pairs  $(k_1, k'_1), \dots, (k_n, k'_n)$  for keywords  $w \in \{1, \dots, n\}$  represented as integers for simplicity. The first key  $k_w$  in each pair is used to encrypt the identifiers of documents containing  $w$ . The latter key  $k'_w$  is used as a pseudorandom function (PRF) key to derive pseudorandom locations to store the encryption of the document identifiers. When the client later wants to search for documents containing  $w$  it sends the associated  $(k_w, k'_w)$  keys to the server. The server then uses  $k'_w$  to re-derive the pseudorandom locations of the ciphertexts and uses  $k_w$  to decrypt them.

To prove adaptive security, the simulator for such a protocol runs into the commitment problem because it must commit to ciphertexts of the document identifiers before knowing what the identifiers are. Perhaps less obviously, a simulator also runs into a commitment issues with the PRF. To ensure security the simulated locations of ciphertexts must be random, but then when responding to a search query the simulator is on the hook to find a key for the PRF that “explains” the simulated locations. Papers on SSE typically address these issue by modeling the PRF as a random oracle and fixing a specific construction of an encryption scheme based on a random oracle. As noted earlier, this has resulted in a need for many tedious and error-prone proofs.

*Asymmetric password-authenticated key exchange and equivocal encryption.* In independent and concurrent work, Jarecki et al. updated the ePrint version of [24] to introduce the notion of equivocal encryption and use it to prove security of their asymmetric password-authenticated key exchange protocol OPAQUE. The definition of equivocal encryption is essentially a weakened version of our confidentiality definition, considering only single-user security and allowing only a single encryption query; whereas we consider multi-user security and arbitrarily many adaptively chosen queries. Since their definition is implied

by ours, our results will make rigorous their claim that “common encryption modes are equivocable under some idealized assumption”.

**A new approach.** We introduce a new framework for analyzing security in adaptive compromise scenarios. Our framework has a simple, but powerful recipe: augment traditional simulation-based, property-based security definitions to empower adversaries with the ability to perform adaptive compromise of secret keys. For symmetric encryption, for example, we convert the standard simulation-based, multi-user indistinguishability under chosen plaintext attack ( $\mu$ -IND-CPA) [4] to a game that includes the same adversarial powers, but adds an additional oracle for adaptively compromising individual user secret keys. Critical to our approach is (1) the use of simulators, which allows handling corruptions gracefully, and (2) incorporating handling of idealized models (e.g., the ROM or ICM). The latter is requisite for analyzing practical constructions.

We offer new definitions for multi-user CPA and CCA security of symmetric encryption, called SIM-AC-CPA (simulation-based security under adaptive corruption, chosen plaintext attack) and SIM-AC-CCA (chosen ciphertext attack). By restricting the classes of allowed simulators we can obtain stronger definitions (e.g., SIM-AC- $\$$  which requires that ciphertexts look like random strings).

*Symmetric encryption under adaptive compromise.* We then begin exercising our framework by first answering the question: Are practical, in-use symmetric encryption schemes secure in the face of adaptive compromises? We give positive results here, in idealized models. Taking an encrypt-then-MAC scheme such as AES in counter mode combined with HMAC [5] as an example, we could directly show SIM-AC-CCA security while modeling AES as an ideal cipher and HMAC as a random oracle (c.f., [16]). But this would lead to a rather complex proof, and we’d have to do similarly complex proofs for other encryption schemes.

Instead, we provide simple, modular proofs by lifting the underlying assumptions made about primitives (AES and HMAC) to hold in adaptive compromise scenarios. Specifically, we introduce a new security notion for pseudorandom functions under adaptive compromise attacks (SIM-AC-PRF). This adapts the standard multi-user PRF notion to also give adversaries the ability to adaptively compromise particular keys. Then we prove that AES and HMAC each achieve this notion in the ICM and ROM, respectively. The benefit is that these proofs encapsulate the complexity of ideal model programming proofs in the simpler context of SIM-AC-PRF (as opposed to SIM-AC-CCA).

The workflow when using our framework is represented by Fig. 2. Here PRFs are individually shown to achieve SIM-AC-PRF security in an ideal model. Then modes of operation are proven secure under the assumption that they use a SIM-AC-PRF secure PRF. Then each application is proven secure under the appropriate assumption of the encryption scheme used. This decreases the total number of proofs required to  $A + P + M$ , significantly fewer than the  $A \cdot P \cdot M$  required previously. Moreover, the complicated ideal model programming analysis (represented by red dashed arrows) is restricted to only appearing in the simplest of these proofs (analyzing of PRFs); it can then simply be “passed along” to the higher level proofs.

We can then show that for most CPA modes of operation (e.g., CBC mode or CTR mode), one can prove SIM-AC-CPA security assuming the underlying block cipher is SIM-AC-PRF. The core requirement is that the mode of operation enjoys a property that we call extraction security. This is a technical condition capturing the key security properties needed to prove that a mode of operation is SIM-AC-CPA assuming the underlying block cipher is SIM-AC-PRF. Moreover, we show that most existing (standard) proofs of IND-CPA security show, implicitly, the extraction security of the mode. Thus, we can easily establish adaptive compromise proofs given existing (standard) ones.

The above addresses only confidentiality. Luckily, integrity is inherited essentially for free from existing analysis. We generically show that SIM-AC-CPA security combined with the standard notion of ciphertext integrity implies SIM-AC-CCA security. Thus, one can prove encrypt-then-MAC is SIM-AC-CCA secure assuming the SIM-AC-CPA security of the encryption and the standard unforgeability security of the MAC. This is an easy adaptation of the standard proof [8] of encrypt-then-MAC.

*Applying the framework to high-level protocols.* Equipped with our new SIM-AC-CCA and SIM-AC-PRF security notions, we can return to our motivating task: providing positive security analyses of BurnBox and the Cash et al. SSE scheme.

We give a proof of BurnBox’s CAS security assuming the SIM-AC-CPA security of the underlying symmetric encryption scheme. Our proof is significantly simpler than the original analysis, avoiding specifically the nuanced programming logic that led to the bug in the original analysis. For the Cash et al. scheme we apply our SIM-AC-PRF definition and a key-private version of our SIM-AC-CPA definition. Their adaptive security claim was accompanied only by a brief proof sketch which fails to identify an important detail that need to be considered in the ROM analysis (see the full version of this paper [23]). Our proof handles this detail cleanly while being ultimately of comparable complexity to their non-adaptive security proof.

Unfortunately, these settings and constructions are inherently complicated. So even with the simplification provided by our analysis techniques there is not space to fit their analysis in the body of our paper; it has instead been relegated to the full version of this work. We choose this organization because our main contribution is the definition abstraction which we believe will be of use for future work, rather than the particular applications we chose to exhibit its use.

*Treatment of symmetric encryption.* In this work, we focus on randomized encryption, over more modern nonce-based variants because this was the form of encryption used by the applications we identified. In the full version of this paper [23], we extend our definitions to nonce-based encryption. The techniques we introduce for analyzing randomized symmetric encryption schemes should extend to nonce-based encryption schemes.

**Related works.** A related line of work is that of selective-opening attacks [7] which studies the security of asymmetric encryption schemes against compromises in a multi-sender setting (where coins underlying encryption may be com-



promised) or multi-receiver setting (where secret decryption keys may be compromised). Selective-opening definitions are typically formulated to aim for standard model (or non-programmable random oracle model) achievability and hence do not suffice for the applications we consider in this work.

The full version of this paper is available on ePrint [23].

## 2 Preliminaries

A list  $T$  of length  $n \in \mathbb{N}$  specifies an ordered sequence of elements  $T[1], T[2], \dots, T[n]$ . The operation  $T.\text{add}(x)$  appends  $x$  to this list by setting  $T[n+1] \leftarrow x$ . This making  $T$  a list of length  $n+1$ . We let  $|T|$  denote the length of  $T$ . The operation  $x \leftarrow T.\text{dq}()$  sets  $x$  equal to the last element of  $T$  and removes this element from  $T$ . In pseudocode lists are assumed to be initialized empty (i.e. have length 0). An empty list or table is denoted by  $[\cdot]$ . We sometimes use set notation with a list, e.g.  $x \in T$  is true if  $x = T[i]$  for any  $1 \leq i \leq |T|$ .

Let  $S$  and  $S'$  be two sets with  $|S| \leq |S'|$ . Then  $\text{Inj}(S, S')$  is the set of all injections from  $S$  to  $S'$ . We will sometimes abuse terminology and refer to functions with co-domain  $\{S : S \subseteq \{0, 1\}^*\}$  as sets. For  $n \in \mathbb{N}$  we define  $[n] = \{1, \dots, n\}$ .

The notation  $y \leftarrow^s A(x_1, x_2, \dots : \sigma)$  denotes the (randomized) execution of  $A$  with state  $\sigma$ . Changes that  $A$  makes to its input variable  $\sigma$  are maintained after  $A$ 's execution. For given  $x_1, x_2, \dots$  and  $\sigma$  we let  $[A(x_1, x_2, \dots : \sigma)]$  denote the set of possible outputs of  $A$  given these inputs.

We define security notions using pseudocode-based games. The pseudocode “Require bool” is shorthand for “If not bool then return  $\perp$ ”. We will sometimes use infinite loops defining variable  $x_u$  for all  $u \in \{0, 1\}^*$ . Such code is executed lazily; the code is initially skipped, then later if a variable  $x_u$  would be used, the necessary code to define it is first run. The pseudocode “ $\exists x \in X$  s.t.  $p(x)$ ” for some predicate  $p$  evaluates to the boolean value  $\bigvee_{x \in X} p(x)$ . If this is true, the variable  $x$  is set equal to the lexicographically first  $x \in X$  for which  $p(x)$  is true.

We use an asymptotic formalism. The security parameter is denoted  $\lambda$ . Our work is generally written in a way to allow concrete security bounds to be extracted easily. In security proofs we typically explicitly state how we will bound the advantage of an adversary by the advantages of reduction adversaries we build (and possibly other terms). Reduction adversaries and simulators are explicitly given in code (from which concrete statements about their efficiency can be obtained by observation).

Let  $f : \mathbb{N} \rightarrow \mathbb{N}$ . We say  $f$  is negligible if for all polynomials  $p$  there exists a  $\lambda_p \in \mathbb{N}$  such that  $f(\lambda) \leq 1/p(\lambda)$  for all  $\lambda \geq \lambda_p$ . We say  $f$  is super-polynomial if  $1/f$  is negligible. We say  $f$  is super-logarithmic if  $2^f$  is super-polynomial.

**Ideal primitives.** We will make liberal use of ideal primitives such as random oracles or ideal ciphers. An ideal primitive  $P$  specifies algorithms  $P.\text{Init}$  and  $P.\text{Prim}$ . The initialization algorithm has syntax  $\sigma_P \leftarrow^s P.\text{Init}(1^\lambda)$ . The stateful evaluation algorithm has syntax  $y \leftarrow^s P.\text{Prim}(1^\lambda, x : \sigma_P)$ . We sometimes use  $A^P$  as shorthand for giving algorithm  $A$  oracle access to  $P.\text{Prim}(1^\lambda, \cdot : \sigma_P)$ . Adversaries are often given access to  $P$  via an oracle  $\text{PRIM}$ .



Ideal primitives should be *stateless*. By this we mean that after  $\sigma_{\mathcal{P}}$  is output by  $\mathcal{P}.\text{Init}$ , it is never modified by  $\mathcal{P}.\text{Prim}$  (so we could have used the syntax  $y \leftarrow_{\$} \mathcal{P}.\text{Prim}(1^\lambda, x, \sigma_{\mathcal{P}})$ ). However, when written this way, ideal primitives are typically inefficient, e.g., for the random oracle model  $\sigma_{\mathcal{P}}$  would store a huge random table. Our security results will necessitate that  $\mathcal{P}$  be efficiently instantiated so we have adopted the stateful syntax to allow ideal primitives to be written in their efficient “lazily sampled” form. Despite this notational convenience, we will assume that any ideal primitive we reference is *essentially stateless*. By this, we mean that it could have been equivalently written to be stateless (if inefficient).<sup>2</sup>

The standard model is captured by the primitive  $\mathcal{P}_{\text{sm}}$  for which  $\mathcal{P}_{\text{sm}}.\text{Init}(1^\lambda)$  and  $\mathcal{P}_{\text{sm}}.\text{Prim}(1^\lambda, x : \sigma_{\mathcal{P}})$  always returns the empty string  $\varepsilon$ .

We define a random oracle that takes arbitrary input and produce variable length outputs. It is captured by the primitive  $\mathcal{P}_{\text{rom}}$  defined as follows.

$$\frac{\mathcal{P}_{\text{rom}}.\text{Init}(1^\lambda)}{\text{Return } [\cdot]} \quad \left| \quad \frac{\mathcal{P}_{\text{rom}}.\text{Prim}(1^\lambda, x : T)}{(x, l) \leftarrow x} \right.$$

$$\left. \begin{array}{l} \text{If } T[x, l] = \perp \text{ then } T[x, l] \leftarrow_{\$} \{0, 1\}^l \\ \text{Return } T[x, l] \end{array} \right.$$

The ideal-cipher model is parameterized by a block-length  $n : \mathbb{N} \rightarrow \mathbb{N}$  and captured by  $\mathcal{P}_{\text{icm}}^n$  defined as follows.<sup>3</sup>

$$\frac{\mathcal{P}_{\text{icm}}^n.\text{Init}(1^\lambda)}{\text{Return } ([\cdot], [\cdot])} \quad \left| \quad \frac{\mathcal{P}_{\text{icm}}^n.\text{Prim}(1^\lambda, x : (E, D))}{(\text{op}, K, y) \leftarrow x} \right.$$

$$\left. \begin{array}{l} \text{If } \text{op} = + \text{ then} \\ \quad \text{If } E[K, y] = \perp \text{ then} \\ \quad \quad z \leftarrow_{\$} \{0, 1\}^{n(\lambda)} \setminus \{ E[K, a] : a \in \{0, 1\}^{n(\lambda)} \} \\ \quad \quad E[K, y] \leftarrow z ; D[K, z] \leftarrow y \\ \quad \quad \text{Return } E[K, y] \\ \text{Else} \\ \quad \text{If } D[K, y] = \perp \text{ then} \\ \quad \quad z \leftarrow_{\$} \{0, 1\}^{n(\lambda)} \setminus \{ D[K, a] : a \in \{0, 1\}^{n(\lambda)} \} \\ \quad \quad D[K, y] \leftarrow z ; E[K, z] \leftarrow y \\ \quad \quad \text{Return } D[K, y] \end{array} \right.$$

It stores tables  $E$  and  $D$  which it uses to lazily sample a random permutation for each  $K$ , with  $E[K, \cdot]$  representing the forward evaluation and  $D[K, \cdot]$  its inverse. It parses its input as a tuple  $(\text{op}, K, y)$  where  $\text{op} \in \{+, -\}$  specifies the direction of evaluation and  $K \in \{0, 1\}^*$  and  $y \in \{0, 1\}^{n(\lambda)}$  specify the input.

Sometimes we construct a cryptographic primitive from multiple underlying cryptographic primitives which expect different ideal primitives. To capture this it will be useful to have a notion of combining ideal primitives. Let  $\mathcal{P}'$  and  $\mathcal{P}''$  be ideal primitives. We define their cross product  $\mathcal{P} = \mathcal{P}' \times \mathcal{P}''$  as follows.

<sup>2</sup> Without this restrictions an ideal primitive could behave in undesirable, contrived ways (e.g., on some special input outputting all prior inputs it has received).

<sup>3</sup> We will implicitly assume  $n(\lambda)$  can be computed in time polynomial in  $\lambda$ . We make similar implicit assumptions for other functions that parameterize constructions of cryptographic primitives.

$\frac{\text{P.Init}(1^\lambda)}{\sigma'_P \leftarrow_s \text{P}'.\text{Init}(1^\lambda)}$ $\sigma''_P \leftarrow_s \text{P}''.\text{Init}(1^\lambda)$ $\text{Return } (\sigma'_P, \sigma''_P)$	$\frac{\text{P.Prim}(1^\lambda, x : \sigma_P)}{(\sigma'_P, \sigma''_P) \leftarrow \sigma_P}$ $(d, x) \leftarrow x$ $\text{If } d = 1 \text{ then } y \leftarrow_s \text{P}'.\text{Prim}(1^\lambda, x : \sigma'_P)$ $\text{Else } y \leftarrow_s \text{P}''.\text{Prim}(1^\lambda, x : \sigma''_P)$ $\sigma_P \leftarrow (\sigma'_P, \sigma''_P)$ $\text{Return } y$
---	---

By our earlier convention  $A^{P' \times P''}$  is shorthand for giving algorithm  $A$  oracle access to  $\text{P.Prim}(1^\lambda, \cdot : \sigma_P)$ . In  $A$ 's code,  $B^{P'}$  denotes giving  $B$  oracle access to  $\text{P.Prim}(1^\lambda, (1, \cdot) : \sigma_P)$  and  $B^{P''}$  to denote giving  $B$  oracle access to  $\text{P.Prim}(1^\lambda, (2, \cdot) : \sigma_P)$ .

## 2.1 Standard Cryptographic Definitions

We recall standard cryptographic syntax and security notions.

**Symmetric encryption syntax.** A symmetric encryption scheme  $\text{SE}$  specifies algorithms  $\text{SE.Kg}$ ,  $\text{SE.Enc}$ , and  $\text{SE.Dec}$  as well as sets  $\text{SE.M}$ ,  $\text{SE.Out}$ , and  $\text{SE.K}$  representing the message, ciphertext, and key space respectively. The key generation algorithm has syntax  $K \leftarrow_s \text{SE.Kg}(1^\lambda)$ . The encryption algorithm has syntax  $c \leftarrow_s \text{SE.Enc}^P(1^\lambda, K, m)$ , where  $c \in \text{SE.Out}(\lambda, |m|)$  is required. The deterministic decryption algorithm and has syntax  $m \leftarrow \text{SE.Dec}^P(1^\lambda, K, c)$ . Rejection of  $c$  is represented by returning  $m = \perp$ . Informally, correctness requires that encryptions of messages in  $\text{SE.M}(\lambda)$  decrypt properly. We assume the boolean ( $m \in \text{SE.M}(\lambda)$ ) can be efficiently computed.

**Integrity of ciphertexts.** Integrity of ciphertext security is defined by the game  $\text{G}_{\text{SE}, \mathcal{A}_{\text{ctxt}}}^{\text{int-ctxt}}$  shown in Fig. 3. In the game, the attacker interacts with one of two “worlds” (determined by the bit  $b$ ) via its oracles  $\text{ENC}$ ,  $\text{PRIM}$ ,  $\text{EXP}$ , and  $\text{DEC}$ . The attacker’s goal is to determine which world it is interacting with.

$\text{Game } \frac{\text{G}_{\text{SE}, \mathcal{A}_{\text{ctxt}}}^{\text{int-ctxt}}(\lambda)}{\text{For } u \in \{0, 1\}^* \text{ do}}$ $K_u \leftarrow_s \text{SE.Kg}(1^\lambda)$ $\sigma_P \leftarrow_s \text{P.Init}(1^\lambda)$ $b \leftarrow_s \{0, 1\}$ $b' \leftarrow_s \mathcal{A}_{\text{ctxt}}^{\text{ENC}, \text{DEC}, \text{EXP}, \text{PRIM}}(1^\lambda)$ $\text{Return } (b = b')$ $\frac{\text{PRIM}(x)}{y \leftarrow_s \text{P.Prim}(1^\lambda, x : \sigma_P)}$ $\text{Return } y$	$\frac{\text{ENC}(u, m)}{\text{Require } m \in \text{SE.M}(\lambda)}$ $c \leftarrow_s \text{SE.Enc}^P(1^\lambda, K_u, m)$ $C_u.\text{add}(c)$ $\text{Return } c$ $\frac{\text{EXP}(u)}{X.\text{add}(u)}$ $\text{Return } K_u$	$\frac{\text{DEC}(u, c)}{\text{Require } u \notin X}$ $\text{Require } c \notin C_u$ $m_1 \leftarrow \text{SE.Dec}^P(1^\lambda, K_u, c)$ $m_0 \leftarrow \perp$ $\text{Return } m_b$
--	---	--

**Fig. 3.** Game defining multi-user CTXT security of  $\text{SE}$  in the face of exposures.

The PRIM oracle gives the attacker access to the ideal primitive  $P$ . The encryption oracle ENC takes as input a user  $u$  and message  $m$ , then returns the encryption of that message using the key of that user,  $K_u$ . Recall that by our convention each  $K_u$  is not sampled until needed. The exposure oracle EXP takes in  $u$  and then returns  $K_u$  to the attacker. The decryption oracle DEC is the only oracle whose behavior depends on the bit  $b$ . It takes as input a user  $u$  and ciphertext  $c$ . When  $b = 1$ , it will return the decryption of  $c$  using  $K_u$  while when  $b = 0$  it will always return  $\perp$ . Thus, the goal of the attacker is to forge a ciphertext which will decrypt to a non- $\perp$  value.

To prevent trivial attacker, we disallow querying a ciphertext to  $\text{DEC}(u, \cdot)$  if it came from  $\text{ENC}(u, \cdot)$  or if  $u$  was already exposed. This is captured by the “Require” statements in DEC using lists  $C_u$  and  $X$  (which store the ciphertexts returned by  $\text{ENC}(u, \cdot)$  and the users that have been exposed, respectively).

We define the advantage function  $\text{Adv}_{\text{SE}, P, \mathcal{A}_{\text{ctxt}}}^{\text{int-ctxt}}(\lambda) = 2 \Pr[\text{G}_{\text{SE}, P, \mathcal{A}_{\text{ctxt}}}^{\text{int-ctxt}}(\lambda)] - 1$ . We say SE is INT-CTXT secure with P if for all PPT  $\mathcal{A}_{\text{ctxt}}$ , the advantage  $\text{Adv}_{\text{SE}, P, \mathcal{A}_{\text{ctxt}}}^{\text{int-ctxt}}(\cdot)$  is negligible. INT-CTXT security is typically defined to only consider a single user and no exposures. Using a hybrid argument one can show that our definition of INT-CTXT security is implied by the more standard definition.

**Function family.** A family of functions  $F$  specifies algorithms  $F.\text{Kg}$  and  $F.\text{Ev}$  together with sets  $F.\text{Inp}$  and  $F.\text{Out}$ . The key generation algorithm has syntax  $K \leftarrow F.\text{Kg}^P(1^\lambda)$ . The evaluation algorithm is deterministic and has the syntax  $y \leftarrow F.\text{Ev}(1^\lambda, K, x)$ . It is required that for all  $\lambda \in \mathbb{N}$  and  $K \in [F.\text{Kg}(1^\lambda)]$  that  $F.\text{Ev}(1^\lambda, K, x) \in F.\text{Out}(\lambda)$  whenever  $x \in F.\text{Inp}(\lambda)$ . It is assumed that random elements of  $F.\text{Out}(\lambda)$  can be efficiently sampled.

Game $G_{F, P, \mathcal{A}}^{\text{ow}}(\lambda)$	$\text{PRIM}(x)$
$K \leftarrow F.\text{Kg}(1^\lambda); \sigma_P \leftarrow P.\text{Init}(1^\lambda)$	$y \leftarrow P.\text{Prim}(1^\lambda, x; \sigma_P)$
$x \leftarrow F.\text{Inp}(\lambda); y \leftarrow F.\text{Ev}^P(\lambda, K, x)$	Return $y$
$x' \leftarrow \mathcal{A}^{\text{PRIM}}(1^\lambda, K, y)$	
Return $(F.\text{Ev}^P(1^\lambda, K, x') = y)$	

**Fig. 4.** Game defining one-wayness of  $F$ .

**One-wayness.** The one-wayness of a family of functions  $F$  is given by the game  $G^{\text{ow}}$  shown in Fig. 4. The adversary is given a key  $K$  to  $F$  and the image  $y$  of a random point  $x$  in the domain. Its goal is to find a point with the same image. We define the advantage function  $\text{Adv}_{F, P, \mathcal{A}}^{\text{ow}}(\lambda) = \Pr[G_{F, P, \mathcal{A}}^{\text{ow}}(\lambda)]$  and say  $F$  is OW secure with  $P$  if  $\text{Adv}_{F, P, \mathcal{A}}^{\text{ow}}(\cdot)$  is negligible for all PPT  $\mathcal{A}$ .

**Security definitions.** In the body of this paper we sometimes informally reference other security notions for symmetric encryption schemes (IND-CPA, IND-CCA, IND-KP, IND- $\$$ ) and function families (PRF, UF-CMA). These definitions are recalled in the full version of this paper [23].

### 3 New Security Definitions for Symmetric Primitives

In this section we provide our definitions for the security of symmetric cryptographic primitives (namely randomized encryption and pseudorandom functions) against attackers able to adaptively compromise users' keys.

#### 3.1 Randomized Symmetric Encryption

We describe our security definitions for randomized symmetric encryption. We refer to them as SIM-AC-CPA and SIM-AC-CCA security. The definition of SIM-AC-CPA (resp. SIM-AC-CCA) security is a generalization of IND-CPA (IND-CCA) security to a multi-user setting in which some users' keys may be compromised by an attacker.

Consider game  $G^{\text{sim-ac-cpa}}$  shown in Fig. 5. It is parameterized by a symmetric encryption scheme SE, simulator S, ideal primitive P, and attacker  $\mathcal{A}_{\text{cpa}}$ . The attacker interacts with one of two "worlds" via its oracles ENC, EXP, and PRIM. The attacker's goal is to determine which world it is interacting with.

In the real world ( $b = 1$ ) the encryption oracle ENC takes  $(u, m)$  as input and returns an encryption of  $m$  using  $u$ 's key  $K_u$ . Oracle PRIM returns the output of the ideal primitive on input  $x$ . Oracle EXP returns  $u$ 's key  $K_u$  to the attacker.

In the ideal world ( $b = 0$ ), the return values of each of these oracles are instead chosen by a simulator S. In PRIM it is given the input provided to the oracle. In ENC it is given the name of the current user  $u$  and some leakage  $\ell$  about the message  $m$ . If  $u$  has not yet been exposed ( $u \notin X$ ) this leakage is just the length of the message. Otherwise the leakage is the message itself. The inputs and outputs of this oracle for a user  $u$  are stored in the lists  $M_u$  and  $C_u$  so they can be leaked to the simulator when  $\text{EXP}(u)$  is called.

We define  $\text{Adv}_{\text{SE}, \text{S}, \text{P}, \mathcal{A}_{\text{cpa}}}^{\text{sim-ac-cpa}}(\lambda) = 2 \Pr[G_{\text{SE}, \text{S}, \text{P}, \mathcal{A}_{\text{cpa}}}^{\text{sim-ac-cpa}}(\lambda)] - 1$ . We say SE is SIM-AC-CPA secure with P if for all PPT  $\mathcal{A}_{\text{cpa}}$  there exists a PPT S such that  $\text{Adv}_{\text{SE}, \text{S}, \text{P}, \mathcal{A}_{\text{cpa}}}^{\text{sim-ac-cpa}}(\cdot)$  is negligible. Intuitively, this definition captures that ciphertexts reveal nothing about the messages encrypted other than their length unless the encryption key is known to the attacker. In the full version of this paper [23], we show that SIM-AC-CPA security is impossible in the standard model. The proof is a simple application of the ideas of Nielsen [31].

SIM-AC-CCA security extends SIM-AC-CPA security by giving  $\mathcal{A}_{\text{cca}}$  access to a decryption oracle which takes as input  $(u, c)$ . In the real world, it returns the decryption of  $c$  using  $K_u$ . In the ideal world, the simulator simulates this. To prevent trivial attacks, the attacker is disallowed from querying  $(u, c)$  if  $c$  was returned from an earlier query  $\text{ENC}(u, m)$ . We define  $\text{Adv}_{\text{SE}, \text{S}, \text{P}, \mathcal{A}_{\text{cca}}}^{\text{sim-ac-cca}}(\lambda) = 2 \Pr[G_{\text{SE}, \text{S}, \text{P}, \mathcal{A}_{\text{cca}}}^{\text{sim-ac-cca}}(\lambda)] - 1$ . We say SE is SIM-AC-CCA secure with P if for all PPT  $\mathcal{A}_{\text{cca}}$  there exists a PPT S such that  $\text{Adv}_{\text{SE}, \text{S}, \text{P}, \mathcal{A}_{\text{cca}}}^{\text{sim-ac-cca}}(\cdot)$  is negligible.

**Simplifications.** It will be useful to keep in mind simplifications we can make to restrict the behavior of the adversary or simulator without loss of generality. They are applicable to all SIM-AC-style definitions we provide in this paper.

<p>Game <math>G_{SE,S,P,\mathcal{A}_{cpa}}^{sim-ac-cpa}(\lambda)</math></p> <p>For <math>u \in \{0, 1\}^*</math> do  <math>K_u \leftarrow \text{SE.Kg}(1^\lambda)</math>  <math>\sigma_P \leftarrow \text{P.Init}(1^\lambda)</math>  <math>\sigma \leftarrow \text{S.Init}(1^\lambda)</math>  <math>b \leftarrow \{0, 1\}</math>  <math>b' \leftarrow \mathcal{A}_{cpa}^{ENC,EXP,PRIM}(1^\lambda)</math>  Return <math>(b = b')</math></p> <hr/> <p><u>PRIM(<math>x</math>)</u>  <math>y_1 \leftarrow \text{P.Prim}(1^\lambda, x : \sigma_P)</math>  <math>y_0 \leftarrow \text{S.Prim}(1^\lambda, x : \sigma)</math>  Return <math>y_b</math></p>	<p><u>ENC(<math>u, m</math>)</u>  Require <math>m \in \text{SE.M}(\lambda)</math>  If <math>u \notin X</math> then <math>\ell \leftarrow  m </math> else <math>\ell \leftarrow m</math>  <math>c_1 \leftarrow \text{SE.Enc}^P(1^\lambda, K_u, m)</math>  <math>c_0 \leftarrow \text{S.Enc}(1^\lambda, u, \ell : \sigma)</math>  <math>M_u.add(m) ; C_u.add(c_b)</math>  Return <math>c_b</math></p> <hr/> <p><u>EXP(<math>u</math>)</u>  <math>K_1 \leftarrow K_u</math>  <math>K_0 \leftarrow \text{S.Exp}(1^\lambda, u, M_u, C_u : \sigma)</math>  <math>X.add(u)</math>  Return <math>K_b</math></p>
<p>Game <math>G_{SE,S,P,\mathcal{A}_{cca}}^{sim-ac-cca}(\lambda)</math></p> <p>For <math>u \in \{0, 1\}^*</math> do  <math>K_u \leftarrow \text{SE.Kg}(1^\lambda)</math>  <math>\sigma_P \leftarrow \text{P.Init}(1^\lambda)</math>  <math>\sigma \leftarrow \text{S.Init}(1^\lambda)</math>  <math>b \leftarrow \{0, 1\}</math>  <math>b' \leftarrow \mathcal{A}_{cca}^{ENC,DEC,EXP,PRIM}(1^\lambda)</math>  Return <math>(b = b')</math></p>	<p><u>DEC(<math>u, c</math>)</u>  Require <math>c \notin C_u</math>  <math>m_1 \leftarrow \text{SE.Dec}^P(1^\lambda, K_u, c)</math>  <math>m_0 \leftarrow \text{S.Dec}(1^\lambda, u, c : \sigma)</math>  Return <math>m</math></p>

**Fig. 5.** Games defining SIM-AC-CPA and SIM-AC-CCA security of SE.

- If an oracle is deterministic in the real world, then we can assume that the adversary never repeats a query to this oracle or that the simulator always provides the same output to repeated queries.
- We can assume the adversary never makes a query to a user it has already exposed or that for such queries the simulator just runs the code of the real world (replacing calls to P with calls to S.Prim).
- We can assume the adversary always queries with  $u \in [u_\lambda]$  for some polynomial  $u_{(\cdot)}$  or that the simulator is agnostic to the particular strings used to reference users.
- We can assume that adversaries never make queries that fail “Require” statements. (All requirements of oracles we provide will be efficiently computable given the transcripts of queries the adversary has made.)

Proving these are slightly more subtle to establish than analogous simplifications would be in non-simulation-based games because of the order that algorithms are quantified in our security definitions. They all follow the same pattern though, so we sketch the second of these.

Suppose SE is SIM-AC-CPA secure for all adversaries that never make a call  $\text{ENC}(u, m)$  after having made a call  $\text{EXP}(u)$ , then we claim SE is SIM-AC-CPA secure. Let  $\mathcal{A}$  be an arbitrary adversary. Then we build a wrapper adversary  $\mathcal{A}'$  that simply forwards all of  $\mathcal{A}$ 's queries except for encryption queries made

for a user that has already been exposed. In these cases  $\mathcal{B}$  responds with the output of  $\text{SE.Enc}^{\text{PRIM}(\cdot)}(1^\lambda, K_u, m)$  (or  $\perp$  if  $m \notin \text{SE.M}(\lambda)$ ), where  $K_u$  is the key last returned from  $\text{EXP}(u)$ . Let  $S'$  be a simulator for  $\mathcal{A}'$ . Then we construct  $S$  for  $\mathcal{A}$  which responds exactly as  $S'$  would except in response to encryption queries made for a user that has already been exposed. In these cases  $S'$  responds with the output of  $\text{SE.Enc}^{S'.\text{Prim}(1^\lambda, \cdot; \sigma)}(1^\lambda, K_u, m)$ , where  $K_u$  is the key it last returned for  $\text{EXP}(u)$ . It is clear that  $\text{Adv}_{\text{SE}, S, P, \mathcal{A}}^{\text{sim-ac-cpa}}(\lambda) = \text{Adv}_{\text{SE}, S', P, \mathcal{A}'}^{\text{sim-ac-cpa}}(\lambda)$  because the view of  $\mathcal{A}$  is identical in the corresponding games.

**Stronger security notions.** It is common in the study of symmetric encryption primitives to study stronger security definitions than IND-CPA security. Most schemes instead aim directly for their output to be indistinguishable from random bits (IND- $\$$ ). This implies IND-CPA security and additional nice properties such as forms of key-privacy.

We can capture such notions by placing restrictions on the behavior of the simulator. Let  $S$  be a simulator (for which we think of  $S.\text{Enc}$  as being undefined) which additionally defines algorithms  $S.\text{Enc}_1$  and  $S.\text{Enc}_2$  as well as set  $S.\text{Out}$ . Then we define simulators  $S_k[S]$  and  $S_\$[S]$  to be identical to  $S$  except for the following encryption simulation algorithms.

$$\begin{array}{l|l} \hline S_k[S].\text{Enc}(1^\lambda, u, \ell : \sigma) & S_\$[S].\text{Enc}(1^\lambda, u, \ell : \sigma) \\ \hline \text{If } \ell \in \mathbb{N} \text{ then } c \leftarrow S.\text{Enc}_1(1^\lambda, \ell : \sigma) & \text{If } \ell \in \mathbb{N} \text{ then } c \leftarrow S.\text{Out}(\lambda, \ell) \\ \text{Else } c \leftarrow S.\text{Enc}_2(1^\lambda, u, \ell : \sigma) & \text{Else } c \leftarrow S.\text{Enc}_2(1^\lambda, u, \ell : \sigma) \\ \text{Return } c & \text{Return } c \\ \hline \end{array}$$

Checking  $\ell \in \mathbb{N}$  acts as a convenient way of verifying if the user being queried has been exposed yet. Because  $S.\text{Enc}_1(1^\lambda, \ell : \sigma)$  is not given  $u$  in  $S_k$ , the output of  $S_k$  is distributed identically for any unexposed users. The class of key-anonymous simulators  $\mathcal{S}_k$  is the set of all  $S_k[S]$  for some  $S$ . Similarly,  $S_\$$  always outputs a random bitstring as the ciphertext for any unexposed user. The class of random-ciphertext simulators  $\mathcal{S}_\$$  is the set of all  $S_\$[S]$  for some  $S$ . Note that  $\mathcal{S}_\$ \subset \mathcal{S}_k$ .

We say  $\text{SE}$  is SIM-AC-KP secure with  $P$  if for all PPT  $\mathcal{A}_{\text{cpa}}$  there exists a PPT  $S \in \mathcal{S}_k$  such that  $\text{Adv}_{\text{SE}, S, P, \mathcal{A}_{\text{cpa}}}^{\text{sim-ac-cpa}}(\cdot)$  is negligible. We say that  $\text{SE}$  is SIM-AC- $\$$  secure with  $P$  if for all PPT  $\mathcal{A}_{\text{cpa}}$  there exists a PPT  $S \in \mathcal{S}_\$$  such that  $\text{Adv}_{\text{SE}, S, P, \mathcal{A}_{\text{cpa}}}^{\text{sim-ac-cpa}}(\cdot)$  is negligible. It is straightforward to see that SIM-AC- $\$$  security implies SIM-AC-KP security which in turn implies SIM-AC-CPA security. Standard counterexamples will show that these implications do not hold in the other direction.

It is sometimes useful to define security in an all-in-one style, introduced by Rogaway and Shrimpton [32], which simultaneously requires IND- $\$$  security and INT-CTXT security. In our framework we can define  $\mathcal{S}_\perp$  as the class of IND-CCA simulators which always return  $\perp$  for decryption queries to unexposed users. Then we say  $\text{SE}$  is SIM-AC-AE secure with  $P$  if for all PPT  $\mathcal{A}_{\text{cca}}$  there exists a PPT  $S \in \mathcal{S}_\$ \cap \mathcal{S}_\perp$  such that  $\text{Adv}_{\text{SE}, S, P, \mathcal{A}_{\text{cca}}}^{\text{sim-ac-cca}}(\cdot)$  is negligible.

### 3.2 Pseudorandom Functions

Typically a symmetric encryption scheme will use a PRF as one of their basic building blocks. For modularity, it will be useful to provide a simulation-based security definition for PRFs in the face of active compromises. In Sect. 6, we show our PRF definition can be applied to construct a SIM-AC secure symmetric encryption scheme. Additionally, in the full version of this paper [23], we show that our definition is of independent use by using it to prove the adaptive security of a searchable symmetric encryption scheme introduced by Cash et al. [12].

Game $G_{F,S,P,\mathcal{A}_{\text{prf}}}^{\text{sim-ac-prf}}(\lambda)$	$\text{Ev}(u, x)$
For $u \in \{0, 1\}^*$ do	$y_1 \leftarrow F.\text{Ev}^P(1^\lambda, K_u, x)$
$K_u \leftarrow F.\text{Kg}(1^\lambda)$	If $u \notin X$ then
$\sigma_P \leftarrow S.\text{Init}(1^\lambda)$	If $T_u[x] = \perp$ then $y_0 \leftarrow F.\text{Out}(\lambda)$
$\sigma \leftarrow S.\text{Init}(1^\lambda)$	Else $y_0 \leftarrow T_u[x]$
$b \leftarrow \{0, 1\}$	Else
$b' \leftarrow \mathcal{A}_{\text{prf}}^{\text{EV,EXP,PRIM}}(1^\lambda)$	$y_0 \leftarrow S.\text{Ev}(1^\lambda, u, x : \sigma)$
Return $b = b'$	$T_u[x] \leftarrow y_0$
$\text{PRIM}(x)$	Return $y_b$
$y_1 \leftarrow P.\text{Prim}(1^\lambda, x : \sigma_P)$	$\text{EXP}(u)$
$y_0 \leftarrow S.\text{Prim}(1^\lambda, x : \sigma)$	$K_1 \leftarrow K_u$
Return $y_b$	$K_0 \leftarrow S.\text{Exp}(1^\lambda, u, T_u : \sigma)$
	$X.\text{add}(u)$
	Return $K_b$

**Fig. 6.** Game defining multi-user PRF security of  $F$  in the face of exposures.

The game  $G_{F,S,P,\mathcal{A}_{\text{prf}}}^{\text{sim-ac-prf}}$  is shown in Fig. 6. In the real world,  $\text{Ev}$  gives adversary  $\mathcal{A}_{\text{prf}}$  the real output of  $F$ . In the ideal world,  $\text{Ev}$ 's output is chosen at random (and stored in the table  $T_u$ ), unless  $u$  has already been exposed in which case simulator  $S$  chooses the output. The table  $T_u$  is given to  $S$  when an exposure of  $u$  happens so it can output a key consistent with prior  $\text{Ev}$  queries; we assume it is easy to iterate over all  $(x, T_u[x])$  pairs for which  $T_u[x]$  is not  $\perp$ . We define  $\text{Adv}_{F,S,P,\mathcal{A}_{\text{prf}}}^{\text{sim-ac-prf}}(\lambda) = 2 \Pr[G_{F,S,P,\mathcal{A}_{\text{prf}}}^{\text{sim-ac-prf}}(\lambda)] - 1$ . We say  $F$  is SIM-AC-PRF secure with  $P$  if for all PPT  $\mathcal{A}_{\text{prf}}$  there exists a PPT  $S$  such that  $\text{Adv}_{F,S,P,\mathcal{A}_{\text{prf}}}^{\text{sim-ac-prf}}(\cdot)$  is negligible.

## 4 Applications

The value of our definitions stems from their usability in proving the security of protocols constructed from symmetric encryption and pseudorandom functions. In this section, we discuss the application our definitions to simplify and modularize existing security results of Cash et al. [12] and Tyagi et al. [33], and how they imply the notion of equivocal encryption introduced by Jarecki et al. [24].



#### 4.1 Asymmetric Password-Authenticated Key Exchange: OPAQUE

Password-authenticated key exchange (PAKE) protocols allow a client and a server with a shared password to establish a shared key resistant to offline guessing attacks. *Asymmetric* PAKE (aPAKE) further considers security in the case of server compromise, meaning that the server must store some secure representation of the password, rather than the password itself.

OPAQUE [24] is an aPAKE protocol currently being considered for standardization by the IETF. At a high level, OPAQUE is constructed from an oblivious pseudorandom function (OPRF) and an authenticated key exchange protocol (AKE). User key material for the AKE protocol is stored encrypted under a password-derived key from an OPRF. Key exchange proceeds in two steps: (1) the user rederives the encryption key by running the OPRF protocol with the server on their password, then (2) retrieves and decrypts the AKE keys from the server-held ciphertext and proceeds with the AKE protocol. The “commitment problem” arises when an adversary compromises the server state and then later compromises a user password.

Game $G_{SE,S,P,A}^{\text{eqv}}(\lambda)$	$\text{PRIM}(x)$
$\sigma_P \leftarrow \text{P.Init}(1^\lambda)$	$y_1 \leftarrow \text{P.Prim}(1^\lambda, x : \sigma_P)$
$\sigma \leftarrow \text{S.Init}(1^\lambda)$	$y_0 \leftarrow \text{S.Prim}(1^\lambda, x : \sigma)$
$b \leftarrow \{0, 1\}$	Return $y_b$
$(m, \sigma_A) \leftarrow \mathcal{A}_1^{\text{PRIM}}(1^\lambda)$	
$K_1 \leftarrow \text{SE.Kg}(1^\lambda)$	
$c_1 \leftarrow \text{SE.Enc}^P(1^\lambda, K_1, m)$	
$c_0 \leftarrow \text{S.Enc}(1^\lambda,  m  : \sigma)$	
$K_0 \leftarrow \text{S.Exp}(1^\lambda, m : \sigma)$	
$b' \leftarrow \mathcal{A}_2^{\text{PRIM}}(1^\lambda, c_b, K_b, \sigma_A)$	
Return $(b = b')$	

**Fig. 7.** Game defining EQV security of SE.

**Comparison to equivocal encryption.** To prove security of their scheme, Jarecki et al. independently propose a weaker version of SIM-AC-CPA that they call equivocal encryption (EQV). Consider game  $G^{\text{eqv}}$  defined in Fig. 7. An encryption scheme SE is *equivocal* if for any PPT adversary  $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ , there exists a simulator S, such that the advantage function  $\text{Adv}_{SE,S,P,A}^{\text{eqv}}(\lambda) = 2\Pr[G_{SE,S,P,A}^{\text{eqv}}(\lambda)] - 1$  is negligible. The [24] definition does not specify how to incorporate the ideal model, so we make a reasonable assumption.

Note that EQV is a weaker version of SIM-AC-CPA in that it allows for only one user and only one encryption query. Showing SIM-AC-CPA implies EQV can be done with a simple wrapper reduction in which the output of  $\mathcal{A}_1$  from EQV is forwarded to the encryption oracle of SIM-AC-CPA. Since

EQV allows for only one encryption query, we can further show that EQV does not imply SIM-AC-CPA. Consider a scheme that uses a key  $K = (K_1, K_2)$  and constructs ciphertexts as  $(K_1, \text{Enc}_{K_2}(m))$  unless  $m = K_1$ , in which case it is formed as  $(K_2, \text{Enc}_{K_2}(m))$ . Such a scheme could be secure with respect to EQV but will not be secure in a game that allows multiple encryption queries. Interestingly, showing that our multi-user SIM-AC-CPA notion is implied by its single-user version through a hybrid argument is not straightforward due to managing inconsistencies in simulator state between hybrid steps. We have not been able to prove this result and leave it open for future work. Thus, even if EQV was extended to allow multiple encryption queries, it still may not be widely applicable to situations that require multiple users.

Ultimately, our work fills in the claim of Jarecki et al. that “common encryption modes are equivocable under some idealized assumption”.

## 4.2 Searchable Symmetric Encryption

In the full version of this paper [23], we show that our symmetric encryption and PRF security definitions are useful for proving the security of searchable searchable symmetric encryption (SSE) schemes. An SSE scheme allows a client with a database of documents to store them in encrypted form on a server while still being able to perform keyword searches on these documents.

As a concrete example, we consider Cash et al. [12] which proved non-adaptive security of an SSE scheme when using a PRF and an IND- $\$$  secure encryption scheme and claimed adaptive security when the PRF is replaced with a random oracle and the encryption scheme is replaced with a specific random-oracle-based scheme. We will prove their adaptive result, this time assuming the family of functions is SIM-AC-PRF secure and the encryption scheme is SIM-AC-KP secure. This makes the result more modular because one is no longer restricted to use their specific choices of a PRF and encryption scheme constructed from a random oracle. As a concrete benefit of this, their choice of encryption scheme does not provide INT-CTXT security. To replace the scheme with one that does would require a separate proof while our proof allows the user to choose their favorite INT-CTXT secure scheme without requiring any additional proofs (assuming that scheme is SIM-AC-CPA secure).

Our proof is roughly as complex as their non-adaptive proof; it consists of three similar reductions to the security of the underlying primitives. Without our definitions, a full adaptive proof would have been a technically detailed (though “standard” and not conceptually difficult) proof because it would have to deal with programming the random oracle. Perhaps because of this, the authors of [12] only provided a sketch of the result, arguing that it follows from the same ideas as their non-adaptive proof plus programming the random oracle to maintain consistency. They claim, “[t]he only defects in the simulation occur when an adversary manages to query the random oracle with a key before it is revealed”. This is technically insufficient; a defect also occurs if the same key is sampled multiple times by the simulator (analogously to parts of our proofs for Theorem 3 and Theorem 4). In our SSE proof, we need not address these details because

programming the ideal primitive is handled by the assumed simulation security of the underlying primitives.

A large number of other works on SSE have used analogous techniques of constructing a PRF and/or encryption scheme from a random oracle to achieve adaptive security [1, 2, 9, 12, 13, 17, 20, 21, 25–29, 34]. As we discuss in the full version of this paper [23], these papers all similarly elided the details of the random oracle programming proof and/or made mistakes in writing these details. The mistakes are individually small and not difficult to fix, but their prevalence indicates the value our definitions can provide to modularize and simplify the proofs in these works. We chose to analyze the Cash et al. scheme to highlight the application of our definitions because it was the simplest construction requiring both SIM-AC-PRF and SIM-AC-KP secure and because their thorough non-adaptive proof served as a useful starting point from which to build our proof.

### 4.3 Self-revocable Encrypted Cloud Storage: BurnBox

In the full version of this paper [23], we consider the BurnBox construction of a self-revocable cloud storage scheme proposed by Tyagi et al. [33]. Its goal is to help provide privacy in the face of an authority searching digital devices, e.g., searches of mobile phones or laptops at border crossings. In their proposed scheme a user stores encrypted version of their files on cloud storage. At any point in time they are able to temporarily revoke their own access to these files. Thereby an authority searching their device is unable to learn the content of these files despite their possession of all the secrets stored on the user’s device.

Proving security of their scheme in their security model necessitates solving the “commitment problem.” A simulator is forced to simulate the attacker’s view by providing ciphertexts for files that it does not know the contents of, then later produce a plausible looking key which decrypts the files properly when told the contents. To resolve this issue in their security they modeled the symmetric encryption scheme in the ideal encryption model (which they introduced for this purpose). We are able to recover their result assuming the SIM-AC-CPA security of the encryption scheme. This provides rigorous justification for the use of practically-used encryption schemes which cannot necessarily be thought of as well modeled by the ideal encryption model (e.g. AES-GCM which they used in their prototype implementation). Moreover, the proof we obtain is simpler than the original proof of Tyagi et al. because we do not have to reason about the programming of the ideal encryption model. The original proof has a bug in this programming which we discuss in the full version of this paper [23].

## 5 Symmetric Encryption Security Results

In this section, we show that important existing results about the security of symmetric encryption schemes “carry over” to our new definitions. These results (together with our results in the next section) form the foundation of our claim

that encryption schemes used in practice can be considered to achieve SIM-AC-AE security when their underlying components are properly idealized. First, we show that SIM-AC-CPA and INT-CTXT security imply SIM-AC-CCA security. Then we show that the classic Encrypt-then-MAC scheme achieves SIM-AC-CCA security. Each of these results are, conceptually, a straightforward extension of their standard proof. Finally, we show that random oracles and ideal ciphers are SIM-AC-PRF secure and ideal encryption [33] is SIM-AC-AE secure.

**CPA and CTXT imply CCA.** The following theorem captures that SIM-AC-CPA and INT-CTXT security imply SIM-AC-CCA security. Bellare and Namprempre [8] showed the analogous result for IND-CPA and IND-CCA security.

**Theorem 1.** *If SE is SIM-AC-CPA and INT-CTXT secure with P, then SE is SIM-AC-CCA secure with P.*

*Proof (Sketch).* Here we sketch the main ideas of the proof. The full details are provided in the full version of this paper [23].

The SIM-AC-CCA simulator we provide is parameterized by a SIM-AC-CPA simulator  $S_{\text{cpa}}$ . As state it stores  $\sigma$  of  $S_{\text{cpa}}$  and keeps each  $K_u$  that is has returned to exposure queries. For PRIM, ENC, and EXP queries it simply runs  $S_{\text{cpa}}$ . For DEC queries it does one of two things. If  $u$  has already been exposed it uses the key it previously returned to run the actual decryption algorithm (with oracle access to  $S_{\text{cpa}}$ 's emulation of P) and returns the result. Otherwise it assumes the adversary has failed at producing a forgery and simply returns  $\perp$ . (Note this means we have SIM-AC-AE security if SE is SIM-AC- $\$$  secure.)

The SIM-AC-CPA security of SE ensures that the adversary cannot differentiate between the real and ideal world queries to PRIM, ENC, and EXP. The INT-CTXT security of SE does the same for the DEC queries. In the full proof we show that  $\text{Adv}_{\text{SE}, S_{\text{cca}}, P, \mathcal{A}_{\text{cca}}}^{\text{sim-ac-cca}}(\lambda) \leq \text{Adv}_{\text{SE}, P, \mathcal{A}_{\text{ctxt}}}^{\text{int-ctxt}}(\lambda) + \text{Adv}_{\text{SE}, S_{\text{cpa}}, P, \mathcal{A}_{\text{cpa}}}^{\text{sim-ac-cpa}}(\lambda)$ .  $\square$

**Encrypt-then-MAC.** Let SE be an encryption scheme. Let F be a family of functions for which  $F.\text{Inp}(\lambda) = \{0, 1\}^*$ . Then the Encrypt-then-MAC encryption scheme using SE and F is denoted  $\text{EtM}[\text{SE}, F]$ . Its message space is defined as  $\text{EtM}[\text{SE}, F].M(\lambda) = \text{SE}.M(\lambda)$ . If SE expects access to ideal primitive  $P_1$  and F expects access to ideal primitive  $P_2$  then  $\text{EtM}[\text{SE}, F]$  expects access to  $P_1 \times P_2$ . The key-generation algorithm  $\text{EtM}[\text{SE}, F].\text{Kg}$  returns  $K = (K_{\text{SE}}, K_{\text{F}})$  where  $K_{\text{SE}}$  was sampled with  $\text{SE}.\text{Kg}(1^\lambda)$  and  $K_{\text{F}}$  was sampled with  $F.\text{Kg}(1^\lambda)$ . Algorithms  $\text{EtM}[\text{SE}, F].\text{Enc}$ , and  $\text{EtM}[\text{SE}, F].\text{Dec}$  are defined as follows.

$$\left. \begin{array}{l} \text{EtM}[\text{SE}, F].\text{Enc}^{P_1 \times P_2}(1^\lambda, K, m) \\ \overline{(K_{\text{SE}}, K_{\text{F}}) \leftarrow K} \\ c_{\text{SE}} \leftarrow \text{SE}.\text{Enc}^{P_1}(1^\lambda, K_{\text{SE}}, m) \\ \tau \leftarrow F.\text{Ev}^{P_2}(1^\lambda, K_{\text{F}}, c_{\text{SE}}) \\ \text{Return } (c_{\text{SE}}, \tau) \end{array} \right| \left. \begin{array}{l} \text{EtM}[\text{SE}, F].\text{Dec}^{P_1 \times P_2}(1^\lambda, K, (c_{\text{SE}}, \tau)) \\ \overline{(K_{\text{SE}}, K_{\text{F}}) \leftarrow K} \\ \text{If } \tau \neq F.\text{Ev}^{P_2}(1^\lambda, K_{\text{F}}, c_{\text{SE}}) \text{ then return } \perp \\ m \leftarrow \text{SE}.\text{Dec}^{P_1}(1^\lambda, K_{\text{SE}}, c_{\text{SE}}) \\ \text{Return } m \end{array} \right.$$

The following theorem establishes that the generic composition result of Bellare and Namprempre [8] holds with our simulation-based definitions of security. We sketch its straightforward proof in the full version of this paper [23].

**Theorem 2.** *Let SE be an encryption scheme. Let F be a family of functions for which  $F.\text{Inp}(\lambda) = \{0, 1\}^*$ . If SE is SIM-AC-CPA secure with  $P_1$  and F is UF-CMA secure with  $P_2$ , then  $\text{EtM}[\text{SE}, F]$  is SIM-AC-CCA secure with  $P_1 \times P_2$ .*

**Random oracles are good PRFs.** We show that a SIM-AC-PRF secure family of functions can be constructed simply in the random oracle model. Consider R defined as follows. It is parameterized by a key-length function  $R.\text{kl} : \mathbb{N} \rightarrow \mathbb{N}$  and output length function  $R.\text{ol} : \mathbb{N} \rightarrow \mathbb{N}$ . It has input set  $R.\text{Inp}(\lambda) = \{0, 1\}^*$  and output set  $R.\text{Out}(\lambda) = \{0, 1\}^{R.\text{ol}(\lambda)}$ .

$$\begin{array}{l|l} R.\text{Kg}(1^\lambda) & R.\text{Ev}^P(1^\lambda, K, x) \\ K \leftarrow_s \{0, 1\}^{R.\text{kl}(\lambda)} & y \leftarrow P((K \| x, R.\text{ol}(\lambda))) \\ \text{Return } K & \text{Return } y \end{array}$$

**Theorem 3.** *R is SIM-AC-PRF secure with  $P_{\text{rom}}$  if  $R.\text{kl}$  is super-logarithmic.*

Concretely, in our proof we provide a simulator  $S_{\text{prf}}$  for which we show that,

$$\text{Adv}_{R, S_{\text{prf}}, P_{\text{rom}}, \mathcal{A}_{\text{prf}}}^{\text{sim-ac-prf}}(\lambda) \leq \frac{u_\lambda^2 + p_\lambda u_\lambda}{2^{R.\text{kl}(\lambda)}}$$

where  $u_\lambda$  is an upper bound on the number of users that  $\mathcal{A}_{\text{prf}}$  queries to and  $p_\lambda$  is an upper bound on the number of PRIM queries that  $\mathcal{A}_{\text{prf}}$  makes.

This theorem captures the random oracle programming implicit in the adaptive security claims of the numerous SSE papers we have identified that used a random oracle like a PRF to achieve adaptive security [1, 2, 9, 12, 13, 17, 20, 21, 25–29, 34]. Of these works, most chose to elide the details of establishing that the adversary cannot detect the random oracle programming, likely considering them simple and/or standard. Despite this, we have identified bugs in all of the proofs that did provide more details. We discuss these bugs in more detail in the full version of this paper [23].

To be clear, we do not claim that any of the SSE schemes studied in these works are insecure. The prevalence of this issue speaks to the difficulty of properly accounting for the details in an ideal model programming proof. Our SIM-AC-PRF notion provides a convenient intermediate definition via which these higher-level protocols could have been proved secure without having to deal with the tedious details of a random oracle programming proof.

*Proof (Sketch).* Here we sketch the main ideas of the proof. The full details are provided in the full version of this paper [23]. The SIM-AC-PRF simulator works as follows. For PRIM queries it just emulates  $P_{\text{rom}}$  using a table  $T$ . For EV queries, it just runs  $R.\text{Ev}$  honestly with the key it previously returned for the given user. For EXP queries (on an unexposed user) it picks a random key for this user and sets  $T$  to be consistent with values in the table  $T_u$  it is given. This simulation is only detectable by an attacker that makes a query to the random oracle with some key that is later chosen by the simulator in response to an

exposure or if the simulator happened to chose the same key for two different users.<sup>4</sup> These events happen with negligible probability.  $\square$

**Ideal ciphers are good PRFs.** One of the most commonly used PRFs is AES so it would be useful to think of it as being SIM-AC-PRF secure; however, due to its invertible nature we cannot realistically model it as a random oracle and refer to the above theorem. Instead, AES is often modeled as an ideal cipher. Let  $B.kl : \mathbb{N} \rightarrow \mathbb{N}$  be given and consider  $B$  defined as follows. It has input set  $B.Inp(\lambda) = \{0, 1\}^{n(\lambda)}$  and output set  $B.Out(\lambda) = \{0, 1\}^{n(\lambda)}$ .

$$\begin{array}{l|l} B.Kg(1^\lambda) & B.Ev^P(1^\lambda, K, x) \\ \hline K \leftarrow_{\$} \{0, 1\}^{B.kl(\lambda)} & y \leftarrow P((+, K, x)) \\ \text{Return } K & \text{Return } y \end{array}$$

The following establishes that an ideal cipher is SIM-AC-PRF secure.

**Theorem 4.**  $B$  is SIM-AC-PRF secure with  $P_{icm}^n$  if  $B.kl, n$  are super-logarithmic.

Concretely, in our proof we provide a simulator  $S_{prf}$  for which we show that,

$$\text{Adv}_{B, S_{prf}, P_{icm}^n, \mathcal{A}_{prf}}^{\text{sim-ac-prf}}(\lambda) \leq \frac{u_\lambda^2 + p_\lambda u_\lambda}{2^{B.kl(\lambda)}} + \frac{q_\lambda^2}{2^{n(\lambda)+1}}$$

where  $u_\lambda$  is an upper bound on the number of users that  $\mathcal{A}_{prf}$  queries to,  $p_\lambda$  is an upper bound on the number of PRIM queries that  $\mathcal{A}_{prf}$  makes, and  $q_\lambda$  is an upper bound on the number of Ev queries that  $\mathcal{A}_{prf}$  makes.

The proof of this theorem follows the same general pattern as the proof that a random oracle is SIM-AC-PRF secure (Theorem 3). It only needs to extend the ideas of this prior result slightly to apply a birthday bound so that we can treat the values of  $P_{icm}^n$  as being sampled with replacement. It works best to process this step last so we do not have to consider the order in which queries are made. The proof is given in the full version of this paper [23].

**Ideal encryption model.** In the full version of this paper [23], we recall the ideal encryption model used in the analysis of Tyagi et al. [33] and show that it gives a SIM-AC-AE secure encryption scheme. While doing so, we identify and show how to fix a bug in their proof which used this model.

## 6 Security of Modes of Operation

In the previous section, we showed that existing analysis of the integrity of a symmetric encryption scheme carries over to our simulation setting to lift SIM-AC-CPA security to SIM-AC-CCA security. It would be convenient to be able to

<sup>4</sup> The latter of these points is the subtle issue that does not have appear to have been identified in *any* of the SSE papers that were (implicitly) using a random oracle as a SIM-AC-PRF.

similarly prove that existing IND-CPA security of an encryption scheme suffices to imply SIM-AC-CPA security. Unfortunately, we cannot possibly hope for this to be the case. We know that IND-CPA security can be achieved in the standard model (assuming one-way functions exist), but SIM-AC-CPA security necessarily requires the use of ideal models.

For any typical encryption scheme we could figure out the appropriate way to idealize its underlying components and then write a programming proof to establish security. This would likely be detail intensive and prone to mistakes. We can improve on this by noting that typical symmetric encryption schemes are built as modes of operation using an underlying PRF. We can aim to prove security more modularly by assuming the SIM-AC-PRF security of the underlying family of functions. This alleviates the detail-intensiveness of the proof because the ideal model programming has already been handled in the assumption of SIM-AC-PRF security; it can simply be “passed” along to the new analysis.

In this section, we will show that we can do *even better* than that. We will restrict attention to modes of operation which are IND- $\$$  secure when built from a PRF and satisfy a special extractability property we define in Sect. 6.1 (which standard examples of modes of operation do). Then, in Sect. 6.2, we establish a generic proof framework to elevate an existing IND- $\$$  security proof to a SIM-AC- $\$$  security proof, by showing that existing proofs of IND- $\$$  security tend to (implicitly) prove that the scheme satisfies our extractability property. Finally, in Sect. 6.3 we discuss how the techniques of this section can be extended to other constructions not captured by our formalism, but also note the existence of a (contrived) mode of operation which is IND- $\$$  secure with any secure PRF, but is never SIM-AC- $\$$  secure.

## 6.1 Modes of Operation and Extractability

We first need to have a formalism capturing what a mode of operation is. Our formalism does not capture all possible modes of operation, but does seem to capture most constructions that are of practical interest and would not be hard to modify to capture other constructions.

A mode of operation SE specifies efficient algorithms SE.Kg, SE.Enc, and SE.Dec as well as sets SE.M, SE.Out, SE.FInp, and SE.FOut. For any family of functions F with F.Inp = SE.FInp and F.Out = SE.FOut, it defines a symmetric encryption scheme SE[F] as follows.

$$\begin{array}{l|l|l} \text{SE[F].Kg}(1^\lambda) & \text{SE[F].Enc}^P(1^\lambda, K, m) & \text{SE[F].Dec}^P(1^\lambda, K, c) \\ \hline K_F \leftarrow_s \text{F.Kg}(1^\lambda) & (K_{SE}, K_F) \leftarrow K & (K_{SE}, K_F) \leftarrow K \\ K_{SE} \leftarrow_s \text{SE.Kg}(1^\lambda) & c \leftarrow_s \text{SE.Enc}^{\text{F}_{K_F}^P}(1^\lambda, K_{SE}, m) & m \leftarrow \text{SE.Dec}^{\text{F}_{K_F}^P}(1^\lambda, K_{SE}, c) \\ \text{Return } (K_{SE}, K_F) & \text{Return } c & \text{Return } m \end{array}$$

The superscript  $\text{F}_{K_F}^P$  is shorthand for oracle access to  $\text{F.Ev}^P(1^\lambda, K_F, \cdot)$ . It is required that  $\text{SE[F].M} = \text{SE.M}$ . Moreover, for a given  $\lambda \in \mathbb{N}$  the encryption of a message  $m \in \text{SE.M}(\lambda)$  must always be in  $\text{SE.Out}(\lambda, |m|)$ .



Suppose we want to prove that SE is SIM-AC- $\$$  whenever F is SIM-AC-PRF. The natural way to do so is to build our simulator S from the encryption scheme from the given simulator  $S_F$  for F. In PRIM we can simply have S.Prim run  $S_F$ .Prim. In ENC the ciphertext is chosen at random if the user has not been exposed, otherwise we can simply run SE.Enc but use  $S_F$ .Ev in place of  $F_{K_F}$ . This just leaves EXP, here we are given a list of ciphertexts for the user and need to output a key to “explain” them. A natural approach is to randomly pick our own  $K_{SE}$  and use  $S_F$ .Exp to chose  $K_F$ . Doing so requires giving  $S_F$  a list of input and outputs to the family of function. Intuitively, it seems we want to be able to “extract” a list of input-outputs pairs for F that explain our ciphertexts.

**Extractability.** A mode of operation is extractable if it additionally specifies an efficient extraction algorithm SE.Ext satisfying a correctness and uniformity property we now define. The extraction algorithm SE.Ext has syntax  $(\mathbf{y}, r) \leftarrow_s \text{SE.Ext}(1^\lambda, K_{SE}, c, m)$ . The goal of this algorithm is to “extract” a sequence of responses  $\mathbf{y}$  by F and a string of randomness  $r$  that explains how message  $m$  could be encrypted to ciphertext  $c$  when using key  $K_{SE}$ . We formally define correctness by the following game. It is assumed that SE.Ext provides outputs of the appropriate lengths to make this code well-defined. Extraction correctness of SE requires that  $\Pr[G_{SE,m}^{\text{corr}}(1^\lambda)] = 1$  for all  $\lambda \in \mathbb{N}$  and  $m \in \text{SE.M}(\lambda)$ .

<p>Game <math>G_{SE,m}^{\text{corr}}(1^\lambda)</math></p> <p><math>K_{SE} \leftarrow_s \text{SE.Kg}(1^\lambda)</math></p> <p><math>c \leftarrow_s \text{SE.Out}(\lambda,  m )</math></p> <p><math>(\mathbf{y}, r) \leftarrow_s \text{SE.Ext}(1^\lambda, K_{SE}, c, m)</math></p> <p><math>i \leftarrow 0</math></p> <p><math>c' \leftarrow \text{SE.Enc}^{\text{RF}}(1^\lambda, K_{SE}, m; r)</math></p> <p>Return <math>c = c'</math></p> <p><math>\text{RF}(x)</math></p> <p><math>i \leftarrow i + 1</math></p> <p>Return <math>\mathbf{y}[i]</math></p>	<p><u>Distribution 1</u></p> <p><math>c \leftarrow_s \text{SE.Out}(\lambda,  m )</math></p> <p><math>(\mathbf{y}, r) \leftarrow_s \text{SE.Ext}(1^\lambda, K_{SE}, c, m)</math></p> <p>Return <math>(\mathbf{y}, r)</math></p> <p><u>Distribution 2</u></p> <p>For <math>i = 1, \dots, q(\lambda,  m )</math> do</p> <p style="padding-left: 20px;"><math>\mathbf{y}[i] \leftarrow_s \text{SE.Out}(\lambda)</math></p> <p><math>r \leftarrow_s \{0, 1\}^{l(\lambda,  m )}</math></p> <p>Return <math>(\mathbf{y}, r)</math></p>
--	---

We will also require a uniformity property of SE.Ext. Specifically we require that its output be uniformly random whenever  $c$  is. Formally, there must exist  $q, l : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$  such that the two distributions on the right above are equivalent for all  $\lambda \in \mathbb{N}$ ,  $m \in \text{SE.M}(\lambda)$ , and  $K_{SE} \in [\text{SE.Kg}(1^\lambda)]$ .<sup>5</sup>

**Extraction security.** A core step in our proof will require an additional property of SE which we will now define. Roughly, the desired property is that if SE.Ext is repeatedly used to explain randomly chosen ciphertexts an adversary cannot notice if it causes inconsistent values to be returned to SE.Enc.

Formally, consider the game  $G^{\text{ind-ac-ext}}$  shown in Fig. 8. In it, a key is chosen for each user and then the adversary is given access to an encryption oracle. In this oracle a random ciphertext is sampled. Then SE.Ext is run to provide vector

<sup>5</sup> Computational relaxations of our uniformity and correctness property would suffice for our results, but seem to be unnecessary for any “natural” modes of operation.

Game $G_{SE, \mathcal{A}}^{\text{ind-ac-ext}}(\lambda)$	$\text{ENC}(u, m)$	$\text{RF}(u, x) // \text{private}$
For $u \in \{0, 1\}^*$ do	Require $m \in \text{SE.M}(\lambda)$	$i \leftarrow i + 1$
$K_{SE, u} \leftarrow_s \text{SE.Kg}(1^\lambda)$	Require $u \notin X$	If $T_u[x] \neq \perp$ then
$b \leftarrow_s \{0, 1\}$	$c \leftarrow_s \text{SE.Out}(\lambda,  m )$	If $b = 1$ then
$b' \leftarrow_s \mathcal{A}^{\text{ENC, EXP}}(1^\lambda)$	$(\mathbf{y}, r) \leftarrow_s \text{SE.Ext}(1^\lambda, K_{SE, u}, c, m)$	$\mathbf{y}[i] \leftarrow T_u[x]$
Return $(b = b')$	$i \leftarrow 0$	$T_u[x] \leftarrow \mathbf{y}[i]$
$\text{EXP}(u)$	$c \leftarrow \text{SE.Enc}^{\text{RF}(u, \cdot)}(1^\lambda, K_{SE, u}, m; r)$	Return $T_u[x]$
$\bar{X}.\text{add}(u)$	Return $c$	
Return $(K_{SE, u}, T_u)$		

**Fig. 8.** Game defining IND-AC-EXT security of SE. Note that the adversary is not given oracle access to the “private” oracle RF.

$\mathbf{y}$  and coins  $r$  which explain this ciphertext with respect to the queried message. Finally,  $\text{SE.Enc}$  is run with coins  $r$  and access to an oracle RF whose behavior depends on the chosen  $\mathbf{y}$ . The ciphertext it outputs is returned to the adversary.

When  $b = 0$ , this oracle simply returns the entries of  $\mathbf{y}$ , one at a time. The value returned for an input  $x$  is stored as  $T_u[x]$ . The behavior when  $b = 1$  is similar except that if an input  $x$  to RF is ever repeated for a user  $u$ , then the value stored in  $T_u[x]$  is used instead of the corresponding entry of  $\mathbf{y}$ . The attacker’s goal is to distinguish between these two cases.

The adversary may choose to expose any user  $u$ , learning  $K_{SE, u}$  and  $T_u$ . After doing so it is no longer able to make  $\text{ENC}$  queries to that user (as captured by the second “Require” statement in  $\text{ENC}$ ). Note that by the uniformity of  $\text{SE.Ext}$  we could instead think of  $\mathbf{y}$  and  $r$  as simply being picked at random without  $\text{SE.Ext}$  being run, but we believe the current framing is conceptually more clear.

We define  $\text{Adv}_{SE, \mathcal{A}}^{\text{ind-ac-ext}}(\lambda) = 2 \Pr[G_{SE, \mathcal{A}}^{\text{ind-ac-ext}}] - 1$  and say that SE is IND-AC-EXT secure if  $\text{Adv}_{SE, \mathcal{A}}^{\text{ind-ac-ext}}(\cdot)$  is negligible for all PPT  $\mathcal{A}$ . This notion will be used for an important step of the coming security proof. Of the properties required from an extraction algorithm it is typically the most difficult to verify.

**Example Modes.** As a simple example, we can consider counter-mode encryption. Let  $\text{CTR.ol}, \text{CTR.il} : \mathbb{N} \rightarrow \mathbb{N}$  be fixed and the latter be super-logarithmic. Then CTR is defined as follows. Its key generation algorithm,  $\text{CTR.Kg}$ , always returns  $\varepsilon$ . Its sets are defined by

$$\begin{aligned} \text{CTR.M}(\lambda) &= (\{0, 1\}^{\text{CTR.ol}(\lambda)})^*, & \text{CTR.Out}(\lambda, l) &= \{0, 1\}^{l + \text{CTR.il}(\lambda)} \\ \text{CTR.FInp}(\lambda) &= \{0, 1\}^{\text{CTR.il}(\lambda)}, & \text{CTR.FOut}(\lambda) &= \{0, 1\}^{\text{CTR.ol}(\lambda)}. \end{aligned}$$

Algorithms  $\text{CTR.Enc}$ ,  $\text{CTR.Dec}$ , and  $\text{CTR.Ext}$  are defined below where  $+$  is addition modulo  $2^{\text{CTR.il}(\lambda)}$  with elements of  $\{0, 1\}^{\text{CTR.il}(\lambda)}$  interpreted as integers.

$\text{CTR.Enc}^O(1^\lambda, K_{\text{SE}}, m)$ $c_0 \leftarrow_{\$} \{0, 1\}^{\text{CTR.il}(\lambda)}$ For $i = 1, \dots,  m _{\text{CTR.ol}(\lambda)}$ $c_i \leftarrow m_i \oplus O(c_0 + i)$ Return $c$	$\text{CTR.Dec}^O(1^\lambda, K_{\text{SE}}, c)$ $c_0 \parallel c' \leftarrow c$ For $i = 1, \dots,  c' _{\text{CTR.ol}(\lambda)}$ $m_i \leftarrow c_i \oplus O(c_0 + i)$ Return $m$	$\text{CTR.Ext}(1^\lambda, K, c, m)$ $r \leftarrow c_0$ For $i = 1, \dots,  m _{\text{F.ol}(\lambda)}$ $\mathbf{y}[i] \leftarrow m_i \oplus c_i$ Return $(\mathbf{y}, r)$
---	--	--

It is clear that  $\text{CTR.Ext}$  is correct and that its outputs are distributed uniformly when  $c$  is picked at random. The IND-AC-EXT security of CTR follows from the probabilistic analysis done in existing proofs of security for CTR, such as the proof of Bellare, Desai, Jokipii, and Rogaway [6]. The standard analysis simply bounds the probability that any of the values  $r_1 + 1, \dots, r_1 + l_1, r_2 + 1, \dots, r_2 + l_2, \dots, r_q + 1, \dots, r_q + l_q$  collide when the  $r_i$ 's are picked uniformly and the  $l_i$ 's are adaptively chosen (before the corresponding  $r_i$  is chosen).

Other IND-AC-EXT secure modes of operation include cipher-block chaining (CBC), cipher feedback (CFB), and output feedback (OFB).

## 6.2 Extractability Implies SIM-AC- $\$$ Security

Finally, we can state the main result of this section, that IND-AC-EXT security of an extractable mode of operation implies SIM-AC- $\$$  security.

**Theorem 5.** *Let SE be an extractable mode of operation which is IND-AC-EXT secure. Then  $\text{SE}[F]$  is SIM-AC- $\$$  secure with  $\text{P}$  whenever  $F$  is SIM-AC-PRF secure with  $\text{P}$  and satisfies  $F.\text{Inp} = \text{SE.FInp}$  and  $F.\text{Out} = \text{SE.FOut}$ .*

The full proof is given in the full version of this paper [23]. It considers a sequence of games which transition from the real world of  $\text{G}^{\text{sim-ac-cpa}}$  to the ideal world (using a simulator we specify). In the first transition we use the security of  $F$  to replace  $\text{SE}$ 's oracle access to it with oracle access to a lazily-sampled random function (or simulation by a given simulator  $\text{S}_{\text{prf}}$  if the corresponding user has been exposed). Next we modify the game so that (for unexposed users) ciphertexts are chosen at random and then explained by  $\text{SE.Ext}$ . Then  $\text{SE.Enc}$  is run with the chosen random coins and oracle access to this explanation (except for whenever a repeat query is made) to produce a modified ciphertext which is returned. The uniformity of  $\text{SE.Ext}$  ensures this game is identical to the prior game. Then we apply the IND-AC-EXT security of  $\text{SE}$  so that the oracle given to  $\text{SE.Enc}$  is not kept consistent on repeated queries. The correctness of  $\text{SE.Ext}$  gives that the output of  $\text{SE.Enc}$  is equal to the  $c$  that was sampled at random. We provide simulator  $\text{S}_{\text{§}}$  that simulates this game perfectly. It runs  $\text{S}_{\text{prf}}$  whenever the game would. On an exposure it generate the table  $T_u$  for  $\text{S}_{\text{prf}}$  by running  $\text{SE.Ext}$  on ciphertexts to obtain explanatory outputs of the PRF.

Concretely, in the proof we construct adversaries  $\mathcal{A}_{\text{prf}}$  and  $\mathcal{A}_{\text{ext}}$  along with simulator  $\text{S}_{\text{cpa}}$  for which we show

$$\text{Adv}_{\text{SE}[F], \text{S}_{\text{§}}[\text{S}_{\text{cpa}}], \text{P}, \mathcal{A}_{\text{cpa}}}^{\text{sim-ac-cpa}}(\lambda) \leq \text{Adv}_{F, \text{S}_{\text{prf}}, \text{P}, \mathcal{A}_{\text{prf}}}^{\text{sim-ac-prf}}(\lambda) + \text{Adv}_{\text{SE}, \mathcal{A}_{\text{ext}}}^{\text{ind-ac-ext}}(\lambda).$$

In the full version of this paper [23], we show that a variant of IND-AC-EXT security without exposures (which we call IND-EXT) necessarily holds if  $\text{SE}[F]$  is

single-user IND- $\$$  secure for all single-user PRF secure  $F$ 's. Moreover, we identify that the typical way that IND-EXT security is shown in security proofs for SE is by proving a slightly stronger property which *will* suffice to imply IND-AC-EXT security. Thereby, one can obtain a SIM-AC- $\$$  security proof from a IND- $\$$  security proof by using the information theoretic core of the existing proof.

### 6.3 Extensions and a Counter-Example Construction

**Simple extensions.** For encryption schemes not covered by our formalism, it will often be easy to extend the underlying ideas to cover the scheme. Suppose SE uses two distinct function families as PRFs, one could extend our mode of operation syntax to cover this by giving two separate PRF oracles to the encryption and decryption oracles. Then security would follow if there is an extraction algorithm satisfies analogous properties which explains outputs for both of the oracles. The proof would just require an additional step in which the second SIM-AC-PRF is replaced with simulation, as in our transition between games  $G_0$  and  $G_1$ .

One can analogously prove the SIM-AC- $\$$  security of the Encrypt-then-PRF construction, where instead of a second SIM-AC-PRF function family we have a SIM-AC- $\$$  encryption scheme. From random ciphertexts it is straightforward to extract the required output of the function family and encryption scheme.

We can also extend the analysis to cover GCM when its nonces chosen uniformly at random. It is not captured by our current syntax because the encryption algorithm always applies the PRF to the all-zero string to derive a sub-key for a hash function. It is straightforward to extend our extraction ideas to allow consistency on this PRF query while maintaining our general proof technique.

**Non-extractable counterexample.** We showed our general security result for extractable modes of operations and described how to extend it for some simple variants. One might optimistically hope that SIM-AC- $\$$  security would hold for any IND- $\$$  secure mode of operation (when a SIM-AC-PRF secure function family is used). Unfortunately, we can show that this is not the case. We can provide an example mode of operation which is IND- $\$$  secure when using a PRF, but not SIM-AC-CPA secure for any choice of function family. It will be clear that this mode of operation is not extractable, as required by our earlier theorem.

Fix  $n : \mathbb{N} \rightarrow \mathbb{N}$ . Let  $G$  be a function family that is OW secure with  $P_{sm}$  and for which  $G.Kg(1^\lambda)$  always returns  $\varepsilon$  and  $G.Ev(1^\lambda, \varepsilon, \cdot)$  is always a permutation on  $\{0, 1\}^{n(\lambda)}$ . Such a  $G$  as a one-way permutation on  $n$ -bits. From  $G$  we construct our counterexample  $CX$ . It has sets  $CX.Out(\lambda, l) = \{0, 1\}^{l+n(\lambda)}$  and  $CX.M(\lambda) = CX.FInp(\lambda) = CX.FOut(\lambda) = \{0, 1\}^{n(\lambda)}$ . Key generation is given by  $CX.Kg = G.Kg$ . Encryption and decryption are given as follows.

$$\begin{array}{l|l}
 \underline{CX.Enc^O(1^\lambda, K_{SE}, m)} & \underline{CX.Dec^O(1^\lambda, K_{SE}, c)} \\
 c_0 \leftarrow_{\$} \{0, 1\}^{n(\lambda)} & c_0 \parallel c_1 \leftarrow c \\
 y \leftarrow G.Ev^\varepsilon(1^\lambda, \varepsilon, O(c_0)) & y \leftarrow G.Ev^\varepsilon(1^\lambda, \varepsilon, O(c_0)) \\
 c_1 \leftarrow y \oplus m & m \leftarrow y \oplus c_1 \\
 \text{Return } c & \text{Return } m
 \end{array}$$

Above, the superscript  $\varepsilon$  is used as shorthand for the oracle that always returns  $\varepsilon$ . Note that this is exactly the behavior of  $\mathbf{G}$ 's expected ideal primitive  $\mathbf{P}_{\text{sm}}$ . This counterexample uses the ideas originally introduced by Fischlin et al. [18] to construct non-programmable random oracles by exploiting a one-way permutation. The construction is not extractable because doing so would require being able to invert the one-way permutation. The following theorem formally establishes that this is a counterexample.

**Theorem 6.** *Fix  $n : \mathbb{N} \rightarrow \mathbb{N}$ . Let  $\mathbf{G}$  be a one-way permutation on  $n$ -bits. Let  $\mathbf{F}$  be a family of functions with  $\mathbf{F}.\text{Out}(\lambda) = \mathbf{F}.\text{Inp}(\lambda) = \{0, 1\}^{n(\lambda)}$  and  $\mathbf{P}$  be an ideal primitive. Then  $\text{CX}[\mathbf{F}]$  is IND- $\$$  secure with  $\mathbf{P}$  if  $\mathbf{F}$  is PRF secure with  $\mathbf{P}$ . However,  $\text{CX}[\mathbf{F}]$  is not SIM-AC-CPA secure with  $\mathbf{P}$ .*

*Proof (Sketch).* That  $\text{CX}[\mathbf{F}]$  is IND- $\$$  secure when  $\mathbf{F}$  is PRF secure follows from, e.g., the standard security proof for CTR plus the observation that a permutation applied to a PRF is still a PRF. For the negative result, let  $\mathbf{S}$  be any simulator and consider the following SIM-AC-CPA adversary  $\mathcal{A}_{\text{cpa}}$  and OW adversary  $\mathcal{A}$ .

$\mathcal{A}_{\text{cpa}}^{\text{ENC, EXP, PRIM}}(1^\lambda)$ $m \leftarrow_{\mathbf{S}} \{0, 1\}^{n(\lambda)}$ $c_0 \parallel c_1 \leftarrow \text{ENC}(1, m)$ $y \leftarrow c_1 \oplus m$ $(K_{\text{SE}}, K_{\text{F}}) \leftarrow \text{EXP}(1)$ $x \leftarrow \mathbf{F}.\text{Ev}^{\text{PRIM}}(1^\lambda, K_{\text{F}}, c_0)$ $\text{If } \mathbf{G}.\text{Ev}^\varepsilon(1^\lambda, \varepsilon, x) = y \text{ then return } 1$ $\text{Return } 0$	$\mathcal{A}^{\text{PRIM}}(1^\lambda, K, y)$ $\sigma \leftarrow \mathbf{S}.\text{Init}(1^\lambda)$ $c_0 \parallel c_1 \leftarrow_{\mathbf{S}} \mathbf{S}.\text{Enc}(1^\lambda, 1, n(\lambda) : \sigma)$ $m \leftarrow c_1 \oplus y$ $M.\text{add}(m) ; C.\text{add}(c_0 \parallel c_1)$ $(K_{\text{SE}}, K_{\text{F}}) \leftarrow_{\mathbf{S}} \mathbf{S}.\text{Exp}(1^\lambda, 1, M, C : \sigma)$ $x \leftarrow \mathbf{F}.\text{Ev}^{\mathbf{S}.\text{Prim}(1^\lambda, : : \sigma)}(1^\lambda, K_{\text{F}}, c_0)$ $\text{Return } x$
---	---

Adversary  $\mathcal{A}_{\text{cpa}}$  queries for the encryption of a random message. Then it exposes the corresponding users and uses the given key to calculate the input-output pair this claims for  $\mathbf{G}$ . If indeed, this is a valid pair it returns 1, otherwise it returns 0. When  $b = 1$ , note that  $\mathcal{A}_{\text{cpa}}$  will always return 1. Intuitively, when  $b = 0$ , adversary  $\mathcal{A}_{\text{cpa}}$  should almost never return 1 because from the perspective of the simulator  $\mathbf{S}$  it looks like  $y$  was chosen at random, so finding a pre-image for it requires breaking the security of  $\mathbf{G}$ .

This intuition is captured by the adversary  $\mathcal{A}$ . It simulates the view  $\mathbf{S}$  would see when run for  $\mathcal{A}$ , except instead of picking  $m$  at random it waits until after running  $\mathbf{S}.\text{Enc}$  and sets  $m \leftarrow c_1 \oplus y$  where  $y$  is the  $\mathbf{G}$  image it was given as input. Note that  $y$  is a uniformly random string because  $\mathbf{G}$  is a permutation and  $\mathbf{S}$  is only given the length of the message at this point. Thus, this re-ordering of the calculation of  $m$  does not change the view of  $\mathbf{S}$ . By asking  $\mathbf{S}$  for the appropriate key and running  $\mathbf{F}.\text{Ev}$ , the adversary obtains a potential pre-image for  $y$ .

Simple calculations give  $\text{Adv}_{\text{SE, S, P}, \mathcal{A}_{\text{cpa}}}^{\text{sim-ac-cpa}}(\lambda) = 1 - \text{Adv}_{\mathbf{G}, \mathcal{P}, \mathcal{A}}^{\text{ow}}(\lambda)$ . The latter advantage is negligible from the security of  $\mathbf{G}$ , so the former is non-negligible.  $\square$

**Extensions to PRFs.** It is often useful to construct a PRF  $\mathbf{H}$  with large input domains from a PRF  $\mathbf{F}$  with smaller input domains. The smaller PRF  $\mathbf{F}$  is often thought of as being reasonably modeled by a random oracle or ideal cipher. If the

larger construction  $H$  is an indifferentiable construction of a random oracle [14, 30], then we can apply Theorem 3 to obtain the SIM-AC-PRF security of  $H$ .

In the case that  $H$  is not indifferentiable, one can often use techniques similar to the above to lift a PRF security proof for  $H$  to a SIM-AC-PRF security proof for  $H$  whenever  $F$  is SIM-AC-PRF secure. Implicit in the existing security proof there will often be a way of “explaining” a random output of  $H$  with random outputs by  $F$ . On exposure queries, the simulator for  $H$  would extract these explanations and feed them to the existing simulator for  $F$  to obtain the key to output. For primitive queries, it would just run the  $F$  simulator and for evaluation queries after exposure it would just run  $H$  using the  $F$  simulator in place of  $F$ .

**Acknowledgments.** We thank Thomas Ristenpart for insightful discussions and helpful contributions in the earlier stage of this project. Jaeger was supported in part by NSF grant CNS-1717640, NSF grant CNS-1719146, and a Sloan Research Fellowship. Tyagi was supported in part by NSF grant CNS-1704296.

## References

1. Asharov, G., Naor, M., Segev, G., Shahaf, I.: Searchable symmetric encryption: optimal locality in linear space via two-dimensional balanced allocations. In: Wichs, D., Mansour, Y. (eds.) 48th ACM STOC, Cambridge, MA, USA, 18–21 June 2016, pp. 1101–1114. ACM Press (2016)
2. Asharov, G., Segev, G., Shahaf, I.: Tight tradeoffs in searchable symmetric encryption. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part I. LNCS, vol. 10991, pp. 407–436. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-96884-1\\_14](https://doi.org/10.1007/978-3-319-96884-1_14)
3. Barbosa, M., Farshim, P.: Indifferentiable authenticated encryption. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part I. LNCS, vol. 10991, pp. 187–220. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-96884-1\\_7](https://doi.org/10.1007/978-3-319-96884-1_7)
4. Bellare, M., Boldyreva, A., Micali, S.: Public-key encryption in a multi-user setting: security proofs and improvements. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, pp. 259–274. Springer, Heidelberg (2000). [https://doi.org/10.1007/3-540-45539-6\\_18](https://doi.org/10.1007/3-540-45539-6_18)
5. Bellare, M., Canetti, R., Krawczyk, H.: Keying hash functions for message authentication. In: Kobitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 1–15. Springer, Heidelberg (1996). [https://doi.org/10.1007/3-540-68697-5\\_1](https://doi.org/10.1007/3-540-68697-5_1)
6. Bellare, M., Desai, A., Jokipii, E., Rogaway, P.: A concrete security treatment of symmetric encryption. In: 38th FOCS, Miami Beach, Florida, 19–22 October 1997, pp. 394–403. IEEE Computer Society Press (1997)
7. Bellare, M., Hofheinz, D., Yilek, S.: Possibility and impossibility results for encryption and commitment secure under selective opening. In: Joux, A. (ed.) EUROCRYPT 2009. LNCS, vol. 5479, pp. 1–35. Springer, Heidelberg (2009). [https://doi.org/10.1007/978-3-642-01001-9\\_1](https://doi.org/10.1007/978-3-642-01001-9_1)
8. Bellare, M., Namprempre, C.: Authenticated encryption: relations among notions and analysis of the generic composition paradigm. In: Okamoto, T. (ed.) ASIACRYPT 2000. LNCS, vol. 1976, pp. 531–545. Springer, Heidelberg (2000). [https://doi.org/10.1007/3-540-44448-3\\_41](https://doi.org/10.1007/3-540-44448-3_41)

9. Bost, R.:  $\Sigma\phi\phi\sigma$ : forward secure searchable encryption. In: Weippl, E.R., Katzenbeisser, S., Kruegel, C., Myers, A.C., Halevi, S. (eds.) ACM CCS 2016, Vienna, Austria, 24–28 October 2016. ACM Press (2016)
10. Canetti, R., Feige, U., Goldreich, O., Naor, M.: Adaptively secure multi-party computation. In: 28th ACM STOC, Philadelphia, PA, USA, 22–24 May 1996, pp. 639–648. ACM Press (1996)
11. Cash, D., Grubbs, P., Perry, J., Ristenpart, T.: Leakage-abuse attacks against searchable encryption. In: Ray, I., Li, N., Kruegel, C. (eds.) ACM CCS 2015, Denver, CO, USA, 12–16 October 2015, pp. 668–679. ACM Press (2015)
12. Cash, D., et al.: Dynamic searchable encryption in very-large databases: data structures and implementation. In: NDSS 2014, San Diego, CA, USA, 23–26 February 2014. The Internet Society (2014)
13. Cash, D., Tessaro, S.: The locality of searchable symmetric encryption. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 351–368. Springer, Heidelberg (2014). [https://doi.org/10.1007/978-3-642-55220-5\\_20](https://doi.org/10.1007/978-3-642-55220-5_20)
14. Coron, J.-S., Dodis, Y., Malinaud, C., Puniya, P.: Merkle-Damgård revisited: how to construct a hash function. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 430–448. Springer, Heidelberg (2005). [https://doi.org/10.1007/11535218\\_26](https://doi.org/10.1007/11535218_26)
15. Curtmola, R., Garay, J.A., Kamara, S., Ostrovsky, R.: Searchable symmetric encryption: improved definitions and efficient constructions. In: Juels, A., Wright, R.N., De Capitani di Vimercati, S. (eds.) ACM CCS 2006, Alexandria, Virginia, USA, 30 October–3 November 2006, pp. 79–88. ACM Press (2006)
16. Dodis, Y., Ristenpart, T., Steinberger, J., Tessaro, S.: To hash or not to hash again? (In)differentiability results for  $H^2$  and HMAC. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 348–366. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-32009-5\\_21](https://doi.org/10.1007/978-3-642-32009-5_21)
17. Etemad, M., Küpçü, A., Papamanthou, C., Evans, D.: Efficient dynamic searchable encryption with forward privacy. PoPETs **2018**(1), 5–20 (2018)
18. Fischlin, M., Lehmann, A., Ristenpart, T., Shrimpton, T., Stam, M., Tessaro, S.: Random Oracles with(out) programmability. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 303–320. Springer, Heidelberg (2010). [https://doi.org/10.1007/978-3-642-17373-8\\_18](https://doi.org/10.1007/978-3-642-17373-8_18)
19. Goldwasser, S., Micali, S.: Probabilistic encryption. J. Comput. Syst. Sci. **28**(2), 270–299 (1984)
20. Hahn, F., Kerschbaum, F.: Searchable encryption with secure and efficient updates. In: Ahn, G.-J., Yung, M., Li, N. (eds.) ACM CCS 2014, Scottsdale, AZ, USA, 3–7 November 2014, pp. 310–320. ACM Press (2014)
21. Hu, S., Cai, C., Wang, Q., Wang, C., Luo, X., Ren, K.: Searching an encrypted cloud meets blockchain: a decentralized, reliable and fair realization. In: IEEE INFOCOM 2018 - IEEE Conference on Computer Communications, pp. 792–800, April 2018
22. Islam, M.S., Kuzu, M., Kantarcioglu, M.: Access pattern disclosure on searchable encryption: ramification, attack and mitigation. In: NDSS 2012, San Diego, CA, USA, 5–8 February 2012. The Internet Society (2012)
23. Jaeger, J., Tyagi, N.: Handling adaptive compromise for practical encryption schemes. Cryptology ePrint Archive, Report 2020/765 (2020). <http://eprint.iacr.org/2020/765>
24. Jarecki, S., Krawczyk, H., Xu, J.: OPAQUE: an asymmetric PAKE protocol secure against pre-computation attacks. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part III. LNCS, vol. 10822, pp. 456–486. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-78372-7\\_15](https://doi.org/10.1007/978-3-319-78372-7_15)



25. Kamara, S., Papamanthou, C.: Parallel and dynamic searchable symmetric encryption. In: Sadeghi, A.-R. (ed.) FC 2013. LNCS, vol. 7859, pp. 258–274. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-39884-1\\_22](https://doi.org/10.1007/978-3-642-39884-1_22)
26. Kamara, S., Papamanthou, C., Roeder, T.: Dynamic searchable symmetric encryption. In: Yu, T., Danezis, G., Gligor, V.D. (eds.) ACM CCS 2012, Raleigh, NC, USA, 16–18 October 2012, pp. 965–976. ACM Press (2012)
27. Kim, K.S., Kim, M., Lee, D., Park, J.H., Kim, W.-H.: Forward secure dynamic searchable symmetric encryption with efficient updates. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017, Dallas, TX, USA, 31 October–2 November 2017, pp. 1449–1463. ACM Press (2017)
28. Li, J., et al.: Searchable symmetric encryption with forward search privacy. IEEE Trans. Dependable Secure Comput., 1 (2019, early access). <https://ieeexplore.ieee.org/document/8621026>
29. Liu, Q., Tian, Y., Wu, J., Peng, T., Wang, G.: Enabling verifiable and dynamic ranked search over outsourced data. IEEE Trans. Services Comput., 1 (2019, early access). <https://ieeexplore.ieee.org/document/8734776>
30. Maurer, U., Renner, R., Holenstein, C.: Indifferentiability, impossibility results on reductions, and applications to the random Oracle methodology. In: Naor, M. (ed.) TCC 2004. LNCS, vol. 2951, pp. 21–39. Springer, Heidelberg (2004). [https://doi.org/10.1007/978-3-540-24638-1\\_2](https://doi.org/10.1007/978-3-540-24638-1_2)
31. Nielsen, J.B.: Separating random Oracle proofs from complexity theoretic proofs: the non-committing encryption case. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 111–126. Springer, Heidelberg (2002). [https://doi.org/10.1007/3-540-45708-9\\_8](https://doi.org/10.1007/3-540-45708-9_8)
32. Rogaway, P., Shrimpton, T.: A provable-security treatment of the key-wrap problem. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 373–390. Springer, Heidelberg (2006). [https://doi.org/10.1007/11761679\\_23](https://doi.org/10.1007/11761679_23)
33. Tyagi, N., Mughees, M.H., Ristenpart, T., Miers, I.: BurnBox: self-revocable encryption in a world of compelled access. In: Enck, W., Felt, A.P. (eds.) USENIX Security 2018, Baltimore, MD, USA, 15–17 August 2018, pp. 445–461. USENIX Association (2018)
34. Zuo, C., Sun, S.-F., Liu, J.K., Shao, J., Pieprzyk, J.: Dynamic searchable symmetric encryption schemes supporting range queries with forward (and backward) security. In: Lopez, J., Zhou, J., Soriano, M. (eds.) ESORICS 2018, Part II. LNCS, vol. 11099, pp. 228–246. Springer, Cham (2018). [https://doi.org/10.1007/978-3-319-98989-1\\_12](https://doi.org/10.1007/978-3-319-98989-1_12)