

# Universally Composable End-to-End Secure Messaging: A Modular Analysis\*

Ran Canetti<sup>†</sup>    Palak Jain<sup>†</sup>    Marika Swanberg<sup>†</sup>    Mayank Varia<sup>†</sup>

May 19, 2023

## Abstract

We model and analyze the Signal end-to-end secure messaging protocol within the Universal Composability (UC) framework. Specifically:

- We formulate an ideal functionality that captures end-to-end secure messaging in a setting with Public Key Infrastructure (PKI) and an untrusted server, against an adversary that has full control over the network and can adaptively and momentarily compromise parties at any time, obtaining their entire internal states. Our analysis captures the forward secrecy and recovery-of-security properties of Signal and the conditions under which they break.
- We model the main components of the Signal architecture (PKI and long-term keys, the backbone continuous-key-exchange or “asymmetric ratchet,” epoch-level symmetric ratchets, authenticated encryption) as individual ideal functionalities. These components are realized and analyzed separately, and then composed using the UC and Global-State UC theorems.
- We show how the ideal functionalities representing these components can be realized using standard cryptographic primitives with minimal hardness assumptions.

Our modeling introduces additional innovations that enable arguing about the security of Signal, irrespective of the underlying communication medium, and facilitate the secure composition of dynamically generated modules that share state. These features, in conjunction with the basic modularity of the UC framework, will hopefully facilitate the use of both Signal-as-a-whole and its individual components within cryptographic applications.

---

\*This article is the full version of a published article at CRYPTO 2022, © IACR 2022. This material is based upon work supported by the National Science Foundation under Grants No. 1718135, 1763786, 1801564, 1915763, and 1931714, by the DARPA SIEVE program under Agreement No. HR00112020021, by DARPA and the Naval Information Warfare Center (NIWC) under Contract No. N66001-15-C-4071, and by a Sloan Foundation Research Award.

<sup>†</sup>Boston University, {canetti,palakj,marikas,varia}@bu.edu.

# Contents

<b>1</b>	<b>Introduction</b>	<b>4</b>
1.1	This Work . . . . .	5
1.2	Notes on the Ideal Secure Messaging Functionality, $\mathcal{F}_{\text{SM}}$ . . . . .	6
1.3	Modelling the Components of the Signal Architecture as UC Modules . . . . .	8
1.4	Streamlining our UC Analysis: Global Functionalities, Party Corruptions . . . . .	13
1.5	Related Work . . . . .	14
<b>2</b>	<b>Universally Composable Security</b>	<b>15</b>
2.1	Universal Composability: A Primer . . . . .	15
2.2	New Capabilities: Global Functionalities, Adversarially Provided Code . . . . .	17
<b>3</b>	<b>Modelling Secure Messaging</b>	<b>19</b>
3.1	Global Functionalities . . . . .	19
3.2	The Secure Messaging Functionality, $\mathcal{F}_{\text{SM}}$ . . . . .	19
<b>4</b>	<b>Overview of our Modular Decomposition</b>	<b>26</b>
<b>5</b>	<b>A Signal-style Secure Messaging Protocol: Realizing <math>\mathcal{F}_{\text{SM}}</math></b>	<b>32</b>
5.1	The Epoch Key Exchange Functionality $\mathcal{F}_{\text{eKE}}$ . . . . .	33
5.2	The Forward Secure Encryption Functionality $\mathcal{F}_{\text{fs\_aead}}$ . . . . .	37
5.3	The Signal Protocol, $\Pi_{\text{SGNL}}$ . . . . .	40
5.4	Security Analysis of $\Pi_{\text{SGNL}}$ . . . . .	42
<b>6</b>	<b>The Public Key Ratchet: Realizing <math>\mathcal{F}_{\text{eKE}}</math></b>	<b>48</b>
6.1	Protocol $\Pi_{\text{eKE}}$ . . . . .	48
6.2	Cascaded PRF-PRG (CPRFG) . . . . .	51
6.3	Security Analysis of $\Pi_{\text{eKE}}$ . . . . .	58
<b>7</b>	<b>Unidirectional Forward Secure Authenticated Channels: Realising <math>\mathcal{F}_{\text{fs\_aead}}</math></b>	<b>72</b>
7.1	The Message Key Exchange Functionality $\mathcal{F}_{\text{mKE}}$ . . . . .	72
7.2	The Single-Message Authenticated Encryption Functionality $\mathcal{F}_{\text{aead}}$ . . . . .	72
7.3	Protocol $\Pi_{\text{fs\_aead}}$ . . . . .	77
<b>8</b>	<b>The Symmetric Ratchet: Realising <math>\mathcal{F}_{\text{mKE}}</math></b>	<b>81</b>
<b>9</b>	<b>Authenticated Encryption for Single Messages: Realising <math>\mathcal{F}_{\text{aead}}</math></b>	<b>84</b>
	<b>References</b>	<b>89</b>

## List of Figures

1	Modeling and realizing secure messaging: The general subroutine structure . . . . .	10
2	The code library functionality, $\mathcal{F}_{lib}$ . . . . .	18
3	The Public-Key Directory Functionality, $\mathcal{F}_{DIR}$ . . . . .	20
4	The Programable Random Oracle Functionality, $\mathcal{F}_{pRO}$ . . . . .	20
5	The Long-Term Keys Module Functionality, $\mathcal{F}_{LTM}$ . . . . .	21
6	The Secure Messaging Functionality $\mathcal{F}_{SM}$ . . . . .	22
7	The Secure Messaging Functionality $\mathcal{F}_{SM}$ . . . . .	23
8	A pictorial rendition of Theorem 2 . . . . .	28
9	A pictorial rendition of Theorem 3 . . . . .	29
10	A pictorial rendition of Theorem 4 . . . . .	30
11	A pictorial rendition of Theorem 5 . . . . .	31
12	A pictorial rendition of Theorem 6 . . . . .	31
13	The Epoch Key Exchange Functionality, $\mathcal{F}_{eKE}$ . . . . .	38
14	The Epoch Key Exchange Functionality, $\mathcal{F}_{eKE}$ (continued) . . . . .	39
15	The Forward-Secure Encryption Functionality $\mathcal{F}_{fs\_aead}$ . . . . .	41
16	The Signal Protocol, $\Pi_{SIGNAL}$ . . . . .	43
17	Secure Messaging Simulator, $\mathcal{S}_{SM}$ . . . . .	44
18	Secure Messaging Simulator $\mathcal{S}_{SM}$ continued... . . . . .	45
19	Internal Adversary $\mathcal{I}_{SM}$ . . . . .	46
20	The Epoch Key Exchange Protocol $\Pi_{eKE}$ . . . . .	50
21	The Epoch Key Exchange Protocol $\Pi_{eKE}$ (continued) . . . . .	52
22	Cascaded PRF-PRG Security Game . . . . .	54
23	CPRFG Construction . . . . .	56
24	Cascaded PRF-PRG Simulator . . . . .	56
25	Internal adversarial code $\mathcal{I}_{eKE}$ . . . . .	59
26	Simulator $\mathcal{S}_{eKE}$ for realizing $\mathcal{F}_{eKE}$ . . . . .	60
27	Simulator $\mathcal{S}_{eKE}$ (continued) . . . . .	61
28	The Message Key Exchange Functionality $\mathcal{F}_{mKE}$ . . . . .	73
29	The Authenticated Encryption with Associated Data Functionality, $\mathcal{F}_{aead}$ . . . . .	74
30	The Forward-Secure Encryption $\Pi_{fs\_aead}$ . . . . .	75
31	Internal Adversary, $\mathcal{I}_{fs\_aead}$ . . . . .	76
32	Forward Secure Authenticated Encryption Simulator, $\mathcal{S}_{fs\_aead}$ . . . . .	78
33	$\mathcal{S}_{fs\_aead}$ Continued... . . . . .	79
34	The Message Key Exchange Protocol $\Pi_{mKE}$ . . . . .	82
35	Message Key Exchange Simulator, $\mathcal{S}_{mKE}$ . . . . .	83
36	The Authenticated Encryption with Associated Data Protocol, $\Pi_{aead}$ . . . . .	85
37	Internal adversarial code $\mathcal{I}_{aead}$ . . . . .	86
38	Authenticated Encryption with Associated Data Simulator, $\mathcal{S}_{aead}$ . . . . .	87

# 1 Introduction

Secure communication, namely allowing Alice and Bob to exchange messages securely, over an untrusted communication channel, without having to trust any intermediate component or party, is perhaps the quintessential cryptographic problem. Indeed, constructing and breaking secure communication protocols, as well as modeling security concerns and guarantees, providing a security analysis, and then breaking the modeling and analysis, has been a mainstay of cryptography since its early days.

Successful secure communication protocols have naturally been built to secure existing communication patterns. Indeed, IPsec has been designed to provide IP-layer end-to-end security for general peer-to-peer communication without the need to trust routers and other intermediaries, while SSL (which evolved into TLS) has been designed to secure client-server interactions, especially in the context of web browsing, and PGP has been designed to secure email communication.

Securing the communication over messaging applications poses a very different set of challenges, even for the case of pairwise communication (which is the focus of this work). First, the communicating parties do not typically have any direct communication connection and may not ever be online at the same time. Instead, they can communicate only via an untrusted server. Next, the communication may be intermittent and have large variability in volumes and level of interactivity. At the same time, a received message should be processed immediately and locally. Furthermore, connections may span very long periods of time, during which it is reasonable to assume that the endpoint devices would be periodically hacked or otherwise compromised – and hopefully later regain security.

The Signal protocol has been designed to give a response to these specific challenges of secure messaging, and in doing so it has revolutionized the concept of secure communication over the Internet in many ways. Built on top of predecessors like Off-The-Record [16], the Signal protocol is currently used to transmit hundreds of billions of messages per day [56].

Modeling the requirements of secure messaging in general, and analyzing the security properties of the Signal protocol in particular, has proved to be challenging and has inspired multiple analytical works [1–3, 7, 10, 11, 14, 17, 20, 29–37, 39–41, 53–55, 59–62, 64]. Some of these works directly address the Signal architecture and realization, whereas others propose new cryptographic primitives that are inspired by Signal’s various modules.

**The need for composable security analysis.** It is well documented that standalone security analyses of protocols (namely, analyses that only consider an execution of the protocol “in vitro”) are not always sufficient to capture the security of the protocol when used as a component within a larger system. This situation is particularly relevant to secure messaging and the Signal protocol. People typically participate concurrently in several conversations spanning several multi-platform chat services (e.g., smartphone and web), and the subtleties between a chat service and the underlying messaging protocol have led to network and systems security issues (e.g., [35, 36, 47]). For example, the Signal protocol is combined with other cryptographic protocols in WhatsApp [63] to perform abuse reporting or Status [57] and Slyo [58] to perform cryptocurrency transactions and Tor-style onion routing.

Moreover, Signal isn’t always employed as a single monolithic protocol. Rather, variations and subcomponents of the Signal protocol are used within the Noise protocol family [52], file sharing services like Keybase [42] (which performs less frequent ratcheting), and videoconferencing services like Zoom [45] (which isn’t concerned with asynchrony).

This state of affairs seems to call for a security analysis within a framework that allows for modular analysis and composable security guarantees. First steps in this direction were taken by

the work of Jost, Maurer, and Mularczyk [41] that defines an abstract ratcheting service within the Constructive Cryptography framework [48, 49], and concurrent work by Bienstock et al. [12] that formulates an ideal functionality of the Signal protocol within the UC framework (see Section 1.5 for details). However, neither of these works give a modular decomposition of Signal into its basic components (as described in [51].)

**The apparent non-modularity of Signal.** One of the main sticking points when modeling and analyzing Signal in a composable fashion is that the protocol purposefully breaks away from the traditional structure of a short-lived “key exchange” module followed by a longer-lived module that primarily encrypts and decrypts messages using symmetric authenticated encryption. Instead, it features an intricate “continuous key exchange” module where shared keys are continually being updated, in an effort to provide forward security (i.e., preventing an attacker from learning past messages), as well as enabling the parties to quickly regain security as soon as the attacker loses access. Furthermore, Signal’s process of updating the shared keys crucially depends on feedback from the “downstream” authenticated encryption module. This creates a seemingly inherent circularity between the key exchange and the authenticated encryption modules, and gets in the way of basing the security of Signal on traditional components such as authenticated symmetric encryption, authenticated key exchange, and key-derivation functions.

**Security of Signal in face of adaptive corruptions.** Another potentially thorny aspect of the security of secure messaging protocols (Signal included) is the need to protect against an adversary that decides whom and when to corrupt, adaptively, based on all the communication seen so far. Indeed, not only is standard semantic security not known to imply security in this setting: there exist encryption schemes that are semantically secure (under reasonable intractability assumptions) but completely break in such a setting [38].

## 1.1 This Work

This work proposes a modular analysis of the Signal protocol and its components using the language of universally composable (UC) security [21, 22]. We focus on modeling Signal at the level specified in their documentation [51], taking care to adhere to the abstractions within the specification and not limiting our analysis to any single choice of cipher suite.

We provide an ideal functionality,  $\mathcal{F}_{\text{SM}}$ , for secure messaging along with individual ideal functionalities that capture each module within Signal’s architecture. We then compose the modules to realize the top-level secure messaging functionality and demonstrate how to realize the modules in a manner consistent with the Signal specification [51]. Our instantiation achieves adaptive security against transient corruptions while making minimal use of the random oracle model.

This combination of composability and modularity makes Signal and its components conveniently plug-and-play: future analyses can easily re-purpose or swap out instantiations of the modules in this work without needing to redo most of the security analysis.

In the process of instantiating the key exchange module in the plain model, we propose a new abstraction for Signal’s continuous key derivation module, which we call a Cascaded PRF-PRG (CPRFG), and we show that this primitive suffices for Signal’s continuous key exchange module to achieve adaptive security. We also show how to construct CPRFGs from PRGs and puncturable PRFs. This new primitive may be useful as a building block in other protocols as well.

The rest of the Introduction is organized as follows. Section 1.2 presents and motivates our formulation of  $\mathcal{F}_{\text{SM}}$ . Section 1.3 presents and motivates the formulation of the individual modules,

and describes how these modules can be realized. Section 1.4 highlights some new uses of the UC framework that might be useful elsewhere. Section 1.5 discusses related work.

## 1.2 Notes on the Ideal Secure Messaging Functionality, $\mathcal{F}_{\text{SM}}$

We provide an ideal functionality  $\mathcal{F}_{\text{SM}}$  that captures end-to-end secure messaging, with some Signal-specific caveats. The goal here is to provide idealized security guarantees that will allow the analysis of existing protocols that use Signal, as well as to facilitate the use of Signal (or any protocol that realizes  $\mathcal{F}_{\text{SM}}$ ) as a component within other protocols in a security-preserving manner.

When a party asks to encrypt a message,  $\mathcal{F}_{\text{SM}}$  returns a string to the party that represents the encapsulated message. When a party asks to decrypt (and provides the representative string), the functionality checks whether the provided string matches a prior encapsulation, and returns the original message in case of a match. The encapsulation string is generated via adversarially provided code that doesn't get any information about the encapsulated message, thereby guaranteeing secrecy. If the provided representative string at the time of decapsulation does not match the original encapsulation exactly, then whether to decapsulate is also decided via adversarially provided code that doesn't get any information about the encapsulated message. This approach avoids restricting to only strong MACs (Message Authentication Codes).

**Simple user interface.** The above encapsulation and decapsulation requests are the only ways that a parent protocol interacts with  $\mathcal{F}_{\text{SM}}$ . This eliminates the need for the parent protocol to maintain session-related state such as epoch-ids or sequence numbers. In addition to simplicity, this modeling provides the guarantee that a badly designed parent protocol cannot harm the security of a protocol realising  $\mathcal{F}_{\text{SM}}$ .<sup>1</sup>

**Abstracting away network delivery.** The fact that  $\mathcal{F}_{\text{SM}}$  models a secure messaging scheme as a set of local algorithms (an encapsulation algorithm and a decapsulation one) substantially simplifies traditional UC modeling of secure communication. In traditional UC modeling, the communication medium is considered part of the service provided by the protocol, and the actual communication is abstracted away.

Moreover, by having  $\mathcal{F}_{\text{SM}}$  return to the parent protocol an actual string (that represents an idealized encapsulated message), our model allows the parent protocol to further process the string as needed, similar to what is done in existing systems. This abstraction grants more flexibility in how the secure communication is utilized and integrated within various applications and contexts.

**Immediate decryption.**  $\mathcal{F}_{\text{SM}}$  guarantees that message decapsulation requests are fulfilled locally on the receiver's machine, and are not susceptible to potential network delays. Furthermore, this holds even if only a subset of the messages arrive, and arrival is out of order (as formalized in [1]). To provide this guarantee within the UC framework, we introduce a mechanism that enables  $\mathcal{F}_{\text{SM}}$  to execute adversarially provided code, without enabling the adversary to prevent immediate fulfillment of a decapsulation request. See more details in Section 2.

**Modeling of PKI and long term keys.** We directly model Signal's specific design for the public keys and associated secret keys that are used to identify parties across multiple sessions. Specifically, we formulate a "PKI" functionality  $\mathcal{F}_{\text{DIR}}$  that models a public "bulletin board," which

---

<sup>1</sup>Furthermore, this modeling forces any protocol that realizes  $\mathcal{F}_{\text{SM}}$  to handle all aspects that are critical to security, such as, say, the timely deletion of sensitive information (keys, internal random choices, epoch identifiers).

stores the long-term, ephemeral, and one-time public keys associated with identities of parties. In addition, we model “long term private key” module  $\mathcal{F}_{\text{LTM}}$  for each identity. This module stores the private keys associated with the public keys of the corresponding party. Both functionalities are modeled as *global*, namely they are used as subroutines by multiple instances of  $\mathcal{F}_{\text{SM}}$ . This modeling is what allows to tie the two participants of a session to long-term identities. Similarly to [19, 28], we treat these modules as incorruptible. It is stressed, however, that, following the Signal architecture, our realization of  $\mathcal{F}_{\text{SM}}$  calls the  $\mathcal{F}_{\text{LTM}}$  module of each party exactly once, at the beginning of the session.

**Modelling corruption and recovery.** One of the main design goals of Signal is its resilience to recurring but transient break-ins. We facilitate the exposition of these properties by modelling corruption as an instantaneous event where the adversary learns the entire state of the corrupted party.<sup>2</sup>

The security guarantees for corruption and recovery are then specified as follows. When the adversary instructs  $\mathcal{F}_{\text{SM}}$  to corrupt a party, it is provided all the messages that have been sent to that party and were not yet received. In addition, the party is marked as compromised until a certain future point in the execution. While compromised, all the messages sent and received by the party are disclosed to the adversary, who can also instruct  $\mathcal{F}_{\text{SM}}$  to decapsulate ciphertexts to any plaintext of its choice. This captures the practical aspects of the Signal protocol, where as long as any one of the parties is compromised, neither party can securely authenticate incoming messages.

Forward secrecy guarantees that the adversary learns nothing about any messages that have been sent and received by the party until the point of corruption. Furthermore, the adversary obtains no information on the history of the session such as its duration or the long term identity of the peer. In  $\mathcal{F}_{\text{SM}}$ , this is guaranteed because corruption does not provide the adversary with any messages that were previously sent and successfully received.

On the other hand, the specific point by which a compromised party regains its security is Signal-specific and described in more detail within. After this point, the adversary no longer obtains the messages the messages sent and received by the parties; furthermore, the adversary can no longer instruct  $\mathcal{F}_{\text{SM}}$  to decapsulate forged ciphertexts.

**Resilience to adaptive corruptions.** All the security guarantees provided by  $\mathcal{F}_{\text{SM}}$  hold in the presence of an adversary that has access to the entire communication among the parties and adaptively decides when and whom to corrupt based on all the communication seen so far. In particular, we do not impose any restrictions on when a party can be corrupted.

**Signal-specific limitations.** The properties discussed so far relate to the general task of secure messaging. In addition,  $\mathcal{F}_{\text{SM}}$  incorporates the following two relaxations that represent known weaknesses that are specific to the Signal design.

First, Signal does not give parties a way to detect whether their peers have received forged messages in their name during corruption. (Such situations may occur when either party was corrupted in the past and then recovered.) This represents a known weakness of Signal [17, 35]. Consequently,  $\mathcal{F}_{\text{SM}}$  exhibits similar behavior.

---

<sup>2</sup>We don’t directly model “Byzantine” corruptions, where the adversary is allowed to destroy or modify the state or program of the corrupted party. Indeed, when the internal state is modified, the concept of “regaining security” of a party’s device following a break-in becomes hard to pin down and is left out of scope for this work. We stress, however, that in our setting, where the adversary has complete control over the communication, mere knowledge of the internal state of a party suffices for full impersonation of that party to its peer.



Second, as remarked in the Signal documentation [51], when one of the parties is compromised, an adversary can “fork” the messaging session. That is, the adversary can create a person-in-the-middle situation where both parties believe they are talking with each other in a joint session, but they are actually both talking with the adversary. Furthermore, this can remain the case indefinitely, even when no party is compromised anymore. (In fact, we know this situation is inherent in an unauthenticated network with transient attacks, at least without repeated use of a long-term uncompromised public key [24].) While such a situation is mentioned in the Signal design documents, pinpointing and analyzing the conditions under which forking occurs has not been formally done before our work and the concurrent work by Bienstock et al. [12]. In our modeling,  $\mathcal{F}_{\text{SM}}$  forks when one of the parties is compromised, and at the same time, the other party successfully decapsulates a forged incoming message with an “epoch id” that is different than the one used by the sender. In that case,  $\mathcal{F}_{\text{SM}}$  remains forked indefinitely, without any additional corruptions.

### 1.3 Modelling the Components of the Signal Architecture as UC Modules

Signal’s strong forward secrecy and recovery from compromise guarantees are obtained via an intricate mechanism where shared keys are continually being updated, and each key is used to encapsulate at most a single message.

To help keep the parties in sync regarding which key to use for a given message, the conversation is logically partitioned into sending epochs, where each sending epoch is associated with one of the two parties, and consists of all the messages sent by that party from the end of its previous sending epoch until the first time this party successfully decapsulates an incoming message that belongs to the peer’s latest sending epoch.

Within each sending epoch, the keys are pseudorandomly generated one after the other in a chain. The initial chaining key for each epoch is generated from a ‘root chain’ that ratchets forward every time a new sending epoch starts. Each ratcheting of the root chain involves a Diffie-Hellman key exchange; the resulting Diffie-Hellman secret is then used as input to the root ratchet (along with an existing chaining value). The public values of each such Diffie-Hellman exchange are piggybacked on the messages within the epoch and therefore authenticated using the same AEAD used for the data. Furthermore, these public values are used as unique identifiers of the sending epoch that each message is a part of. This mechanism allows the parties to keep in sync without storing any long-term information about the history of the session.

The Signal architecture document [51] de-composes the above mechanism into 3 main cryptographic modules, plus non-cryptographic code used to put these modules together. The modules are: (1) a symmetric authenticated encryption with associated data (AEAD) scheme that is applied to individual messages; (2) a symmetric key *ratcheting* mechanism to evolve the key between messages within an epoch; (3) an asymmetric key *ratcheting* (or “continuous key exchange”) mechanism to evolve the “root chain.” Since these modules are useful for applications beyond this particular protocol, we follow this partitioning and decompose Signal’s protocol into similar components. (Our partitioning into components is also inspired by that of Alwen et al. [1].)

We model the security of each component as an ideal functionality within the UC framework. (These are  $\mathcal{F}_{\text{aead}}$ ,  $\mathcal{F}_{\text{mKE}}$ ,  $\mathcal{F}_{\text{eKE}}$ , respectively.) This allows us to distill the properties provided by each module and demonstrate how they can be composed, along with the appropriate management code to obtain the desired functionality—namely to realize  $\mathcal{F}_{\text{SM}}$ . The management code (specifically, protocols  $\Pi_{\text{fs\_aead}}$  and  $\Pi_{\text{SGNL}}$ ), does not directly access any keying material. Indeed, these protocols realise their respective specifications, namely  $\mathcal{F}_{\text{fs\_aead}}$  and  $\mathcal{F}_{\text{SM}}$ , perfectly—see Theorems 2 and 4.

Before proceeding to describe the modules in more detail, we highlight the following apparent



circularity in the security dependence between these modules: the messages in each sending epoch need to be authenticated (by the AEAD in use) using a key  $k$  *that’s derived from part of the message itself*. This interdependence between message content and the authentication key might initially suggest that a modular security analysis along the above partitioning to modules would be impossible.

One way to get around this issue is to include the AEAD module (or at least the authentication of the first message in each epoch) within the key exchange module, thereby allowing  $\mathcal{F}_{\text{eKE}}$  to authenticate new epoch identifiers. However, this would alter the partitioning of modules as given in the Signal protocol. We thus keep the AEAD module separate, and explicitly model the fact that the  $\mathcal{F}_{\text{eKE}}$  does not determine the authenticity of new epoch identifiers. In our work,  $\mathcal{F}_{\text{eKE}}$  assigns a fresh pseudorandom secret key with each new epoch identifier, regardless of whether it is authentic or not. The determination of whether a new purported epoch identifier is authentic (or a forgery caused by an adversarially generated incoming message) is done elsewhere – specifically by the calling protocol, which receives feedback from an authenticated encryption functionality that uses the key  $k$  output by  $\mathcal{F}_{\text{eKE}}$  for this epoch identifier. This is a separation of tasks successfully breaks-up the circularity issue, allowing it to be addressed across multiple components in a modular fashion. We now proceed to provide a more detailed overview of our partitioning and the general protocol logic. See also Figure 1.

$\mathcal{F}_{\text{eKE}}$ . The core component of the protocol is the epoch key exchange functionality  $\mathcal{F}_{\text{eKE}}$ , which captures the generation of the initial shared secret key from the public information, as well as the continuous Diffie-Hellman protocol that generates the unique epoch identifiers and the “root chain” of secret keys. Whenever a party wishes to start a new epoch as a sender, it asks  $\mathcal{F}_{\text{eKE}}$  for a new epoch identifier, as well as an associated secret key. The receiving party of an epoch must present an epoch identifier, and is then given the associated secret key.

As mentioned, we allow the receiving party of a new epoch to present multiple potential epoch identifiers, and obtain a secret epoch key associated with each one of these identifiers. Furthermore, while only one of these keys is the one used by the sender for this epoch, all the other keys provided by  $\mathcal{F}_{\text{eKE}}$  are guaranteed to appear random and independent to the adversary. In other words,  $\mathcal{F}_{\text{eKE}}$  leaves it to the receiver to determine which of the candidate identifiers for the new epoch is the correct one. If  $\mathcal{F}_{\text{eKE}}$  recognizes, from observing the corruption activity and the generated epoch ids, that the session has forked, then it exposes the secret keys to the adversary. We postpone the discussion of realizing  $\mathcal{F}_{\text{eKE}}$  to the end of this section.

$\mathcal{F}_{\text{mKE}}$ . The per-epoch key chain is captured by an ideal functionality  $\mathcal{F}_{\text{mKE}}$  that is identified by an epoch-id, and generates, one at a time, a sequence of random symmetric keys associated with this epoch-id. The length of the chain is not a priori bounded; however, once  $\mathcal{F}_{\text{mKE}}$  receives an instruction to end the chain for a party, it complies.  $\mathcal{F}_{\text{mKE}}$  guarantees forward secrecy by making each key retrievable at most once by each party; that is, the key becomes inaccessible upon first retrieval, even for a corrupted party. However, it does not post-compromise security: once corrupted, all the future keys in the sequence are exposed to the adversary.

$\mathcal{F}_{\text{mKE}}$  is realized by a protocol,  $\Pi_{\text{mKE}}$ , that first calls  $\mathcal{F}_{\text{eKE}}$  with its current epoch-id, to obtain the initial chaining key associated with that epoch-id. The rest of the keys in this epoch are derived using a key derivation function (KDF) (of which Signal’s typical instantiation using HKDF is a special case.) The desired properties of the KDF are discussed at a high level in Signal’s documentation, we introduce the new primitive cascaded PRF-PRG to capture the necessary properties of this module.

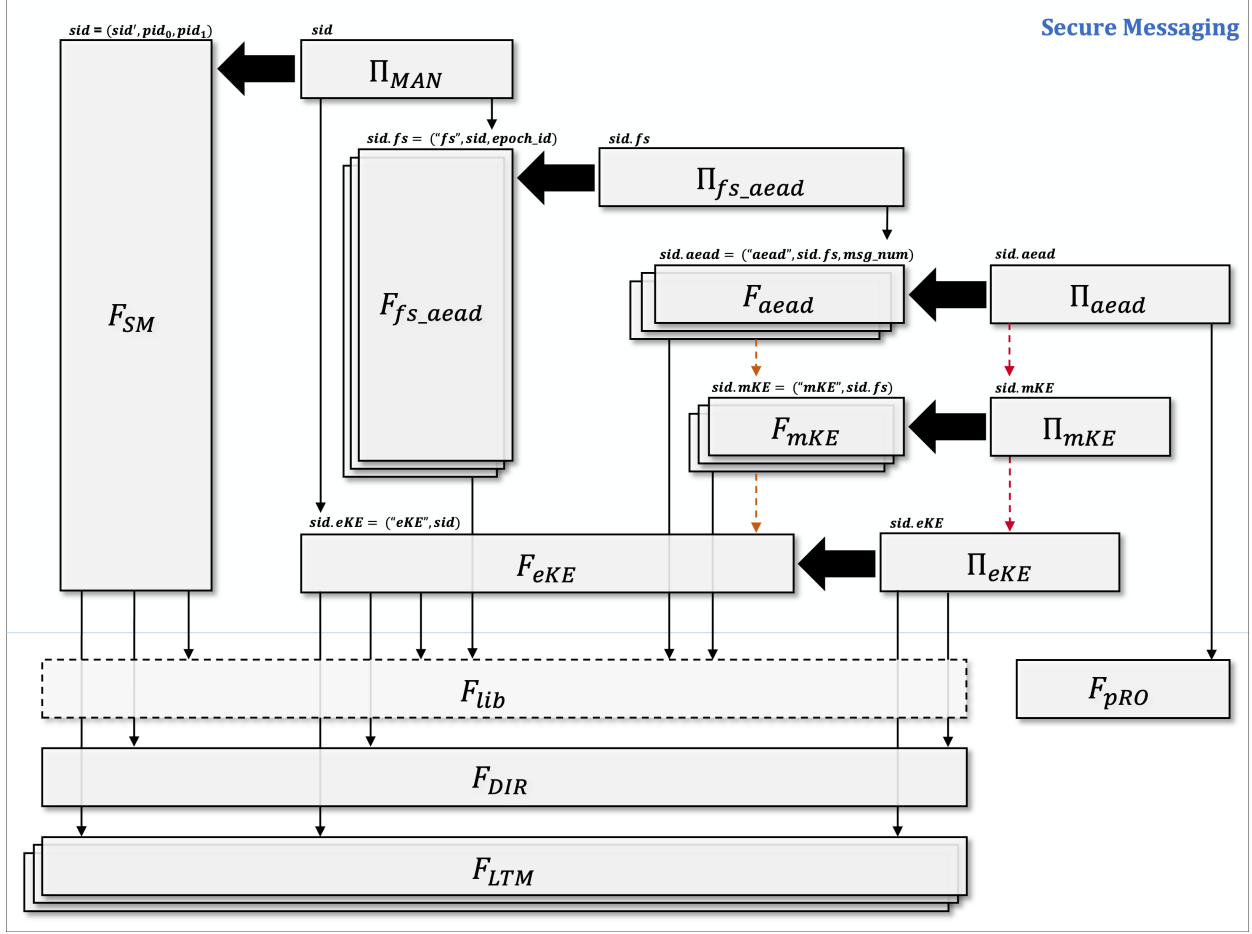


Figure 1: Modeling and realizing secure messaging: The general subroutine structure. Ideal functionalities are denoted by  $F$  and protocols by  $\Pi$ . Horizontal arrows denote realization, whereas vertical arrows denote subroutine calls; the red dashed arrows denote subroutine calls to  $F^\Pi$  (see Section 2.2 and 4.) Functionalities  $\mathcal{F}_{DIR}$ ,  $\mathcal{F}_{LTM}$ ,  $\mathcal{F}_{lib}$ ,  $\mathcal{F}_{PRO}$  are global with respect to  $\mathcal{F}_{SM}$  and all its sub-functionalities, whereas  $\mathcal{F}_{eKE}$  and  $\mathcal{F}_{mKE}$  are global only to a subset of functionalities.

Demonstrating that  $\Pi_{mKE}$  realizes  $\mathcal{F}_{mKE}$  is relatively straightforward, except for the need to address the fact that the same instance of  $\mathcal{F}_{eKE}$  is used by multiple instances of  $\Pi_{mKE}$ . Using the formalism of [5], we thus show that  $\Pi_{mKE}$  UC-realizes  $\mathcal{F}_{mKE}$  in the presence of a global  $\mathcal{F}_{eKE}$ .

$\mathcal{F}_{aead}$ . Authenticated encryption with associated data is captured by ideal functionality  $\mathcal{F}_{aead}$ , which provides a one-time ideal authenticated encryption service: the encrypting party calls  $\mathcal{F}_{aead}$  with a plaintext and a recipient identity, and obtains an opaque ciphertext. Once the recipient presents the ciphertext,  $\mathcal{F}_{aead}$  returns the plaintext. (The recipient is given the plaintext only once, even when corrupted.) The “associated data,” namely the public part of the authenticated message, is captured via the session identifier of  $\mathcal{F}_{aead}$ .

$\mathcal{F}_{aead}$  is realized via protocol  $\Pi_{aead}$ , which employs an authenticated encryption algorithm using a key obtained from  $\mathcal{F}_{mKE}$ . If we had opted to assert security against non-adaptive corruptions, any standard AEAD scheme would do. However, we strive to provide simulation-based security in the presence of fully adaptive corruptions, which is provably impossible in the plain model

whenever the key is shorter than the plaintext [50]. We get around this issue by realizing  $\mathcal{F}_{\text{aead}}$  in the programmable random oracle model. While we provide a very simple AEAD protocol in this model, many common block cipher-based AEADs can also realize  $\mathcal{F}_{\text{aead}}$  provided we model the block cipher as a programmable random oracle. It is stressed however that the random oracle is used *only* in the case of short keys and adaptive corruptions. In particular, when corruptions are non-adaptive or the plaintext is sufficiently short, our protocol continues to UC-realize  $\mathcal{F}_{\text{aead}}$  even when the random oracle is replaced by the identity function.

Since each instance of  $\mathcal{F}_{\text{mKE}}$  is used by multiple instances of  $\Pi_{\text{aead}}$ , we treat  $\mathcal{F}_{\text{mKE}}$  as a global functionality with respect to  $\Pi_{\text{aead}}$ . That is, we show that  $\Pi_{\text{aead}}$  UC-realizes  $\mathcal{F}_{\text{aead}}$  in the presence of (a global)  $\mathcal{F}_{\text{mKE}}$ .

$\mathcal{F}_{\text{fs\_aead}}$ . Functionality  $\mathcal{F}_{\text{fs\_aead}}$  is an abstraction of the management module that handles the encapsulation and decapsulation of all the messages within a single epoch. An instance of  $\mathcal{F}_{\text{fs\_aead}}$  is created by the main module of Signal whenever a new epoch is created, with session ID that contains the the identifier of this epoch.  $\mathcal{F}_{\text{fs\_aead}}$  then provides encapsulation and decapsulation services, akin to those of  $\mathcal{F}_{\text{aead}}$ , for all the messages in its epoch. In addition, once instructed by the main module that its epoch has ended,  $\mathcal{F}_{\text{aead}}$  no longer allows encapsulation of new messages — even when the party is corrupted.

$\mathcal{F}_{\text{fs\_aead}}$  is realized (perfectly, and in a straightforward way) by protocol  $\Pi_{\text{fs\_aead}}$  that calls multiple instances of  $\mathcal{F}_{\text{aead}}$ , plus an instance of  $\mathcal{F}_{\text{mKE}}$  for this epoch - where, again, the session ID of  $\mathcal{F}_{\text{mKE}}$  contains the current epoch id.

$\Pi_{\text{SGNL}}$ . At the highest level of abstraction, we have each of the two parties run protocol  $\Pi_{\text{SGNL}}$ . When initiating a session, or starting a new epoch within a session, (i.e., when encapsulating the first message in an epoch),  $\Pi_{\text{SGNL}}$  first calls  $\mathcal{F}_{\text{eKE}}$  to obtain the identifier of that epoch, then creates an instance of  $\mathcal{F}_{\text{fs\_aead}}$  for that epoch id and asks this instance to encapsulate the first message of the epoch. All subsequent messages of this epoch are encapsulated via the same instance of  $\mathcal{F}_{\text{fs\_aead}}$ .

On the receiver side, once  $\Pi_{\text{SGNL}}$  obtains an encapsulated message in a new epoch id, it creates an instance of  $\mathcal{F}_{\text{fs\_aead}}$  for that epoch id and asks this instance to decapsulate the message. It is stressed that the epoch id on the incoming message may well be a forgery; however in this case it is guaranteed that decapsulation will fail, since the peer has encapsulated this message with respect to a different epoch id, namely a different instance of  $\mathcal{F}_{\text{fs\_aead}}$ . (This is where the circular dependence breaks: even though the environment may invoke  $\Pi_{\text{SGNL}}$  on arbitrary incoming encapsulated message, along with related epoch ids,  $\mathcal{F}_{\text{fs\_aead}}$  is guaranteed to reject unless the encapsulated message uses the same epoch id as the as actual sender. Getting under the hood, this happens since the instances of  $\mathcal{F}_{\text{mKE}}$  that correspond to different epoch ids generate keys that are mutually pseudorandom.) It is emphasised that  $\Pi_{\text{SGNL}}$  is purely “management code” in the sense that it only handles idealized primitives and does not directly access cryptographic keying material. Commensurately, it UC-realizes  $\mathcal{F}_{\text{SM}}$  perfectly.

**Realizing  $\mathcal{F}_{\text{eKE}}$ .** Recall that  $\mathcal{F}_{\text{eKE}}$  is tasked to generate, at the beginning of each new epoch, multiple alternative keys for that epoch – a key for each potential epoch-id for that epoch. This should be done while preserving simulatability in the presence of adaptive corruptions.

Following the Signal architecture, the main component of the protocol that realizes  $\mathcal{F}_{\text{eKE}}$  is a key derivation function (KDF) that combines existing secret state, with new public information (namely the public Diffie-Hellman exponents, which also double-up as an epoch-id), and a new

shared key (the corresponding Diffie-Hellman secret), to obtain a new secret key associated with the given epoch-id, along with potential new local secret state for the KDF.

If the KDF is modeled as a random oracle then it is relatively straightforward to show that the resulting protocol UC-realizes  $\mathcal{F}_{\text{eKE}}$ . On the other extreme, it can be seen that no plain-model instantiation of the KDF module, with bounded-size local state, can possibly realize  $\mathcal{F}_{\text{eKE}}$  in our setting. Indeed, since the adversary can obtain unboundedly many alternative keys for a given epoch, where all keys are generated using the same bounded-size secret state, the Nielsen bound [50] applies.

We propose a middle-ground solution: we show how to instantiate the KDF via a plain-model primitive which we call a *cascaded PRF-PRG (CPRFG)* whose local state grows linearly with the number of keys requested from  $\mathcal{F}_{\text{eKE}}$  at the beginning of a given epoch. Once the epoch advances, the state shrinks back to its original size. Our instantiation uses standard primitives (pseudorandom generators and puncturable pseudorandom functions) and realises a marginally weaker functionality  $\mathcal{F}_{\text{eKE}}^*$  that differs in how it handles the problem of the adversary sending malformed packets to the receiver during a compromised epoch. Our name “cascaded PRF-PRG” is based on the conceptual similarity between our definition and the PRF-PRG introduced by Alwen et al. [1]; nevertheless, we stress that technically the primitives are quite different, as we elaborate below in Section 1.5.

**Modularity with weaker adaptivity.** Some earlier analyses of Signal (see Section 1.5) consider security against adversaries which are limited in their adaptivity. Our modular decomposition of Signal as presented here remains valid for such adversaries. That is, as long as the building blocks can withstand a certain level of adaptivity, the overall protocol remains secure.<sup>3</sup>

**The benefits of modularity** In conclusion, we emphasise the advantages of a fully composable and modular approach to analysing the Signal architecture for secure instant messaging. We mentioned earlier in this section that currently existing implementations of AEAD and the asymmetric ratchet could be shown to realize  $\mathcal{F}_{\text{aead}}$  and  $\mathcal{F}_{\text{eKE}}$  if one assumes a stronger reliance on the random oracle model. Concretely, one could (1) instantiate  $\mathcal{F}_{\text{aead}}$  using any CTR- or CBC-based encryption scheme under the assumption that the block cipher is a random oracle, and (2) instantiate  $\mathcal{F}_{\text{eKE}}$  using HKDF construction under the assumption that its HMAC subroutine is a random oracle. We emphasize that the modularity and generality of our interconnected ideal functionalities allow other instantiations of  $\Pi_{\text{aead}}$  and  $\Pi_{\text{eKE}}$  to be easily plugged in.

**Additional functionalities and methods.** Our modeling and instantiation of Signal also relies on the existence of several functionalities at the network level and on individual devices. For instance, we rely on the existence of a public key infrastructure  $\mathcal{F}_{\text{lib}}$ , which can be modeled as a global UC functionality as described in the next section. We also assume that each party’s corresponding secret key is stored in a long-term module  $\mathcal{F}_{\text{LTM}}$  that is protected at the system level from exfiltration even if the device is corrupted. Moreover, we presume that local devices have the ability to perform secure deletion, in order to erase old keys as part of the key ratcheting process.

One of the contributions of this work is to specify concretely what functionalities Signal relies upon, and in which components of the system. When a protocol is not specified clearly, an instantiation that adheres to the specified protocol could mess up implied aspects of the protocol (such as deletion of old keys) and, in doing so, open up avenues for attacks. The protocol specifications in

---

<sup>3</sup>For more details on modeling adaptivity levels of corruptions in the UC framework, see [22, Section 7.1 on Page 69].

this work are somewhat longer than expected because they address such concerns, whereas Signal’s specification [51] sometimes lacks such detailed guidance.

## 1.4 Streamlining our UC Analysis: Global Functionalities, Party Corruptions

We highlight two additional modeling and analytical techniques that we used to simplify the overall analysis. We hope that these would be useful elsewhere.

**Multiple levels of global state.** Our analysis makes extensive use of universal composition with global state (UCGS) within the plain UC model, as formulated and proven in [5]. Specifically:

- At the highest level we use UCGS to model three global modules available to all the other functionalities in our paper. (1) A single *global directory*  $\mathcal{F}_{\text{DIR}}$  that stores the public keys of all parties. (2) A *long term storage module*  $\mathcal{F}_{\text{LTM}}$  for each party that stores the private keys corresponding to that party’s globally available public keys. (2) A single *adversarial code library*  $\mathcal{F}_{\text{lib}}$  that stores the adversarially provided code for every functionality.
- Additionally, we use UCGS to model the two key exchange modules in our breakdown of the Signal architecture (This allows multiple short lived functionalities to ask for keys from a single long term module): (1) For each epoch of the conversation, an instance of the *message key exchange functionality*  $\mathcal{F}_{\text{mKE}}$  receives a request for each message key from the encryption module for that particular message. (2) For the entire conversation, a single *epoch key exchange functionality*  $\mathcal{F}_{\text{eKE}}$  receives requests for each epoch key from the message key exchange module  $\mathcal{F}_{\text{mKE}}$  for that particular epoch.
- Finally, we use UCGS to model the *programmable random oracle*  $\mathcal{F}_{\text{PRO}}$  that is used in the encryption module.

Furthermore, since our use of global state is not completely covered by the UCGS Theorem of [5], we extend that theorem so as to cover our use case. To see why our use case is not covered by the UCGS Theorem, recall that the UCGS Theorem states that if a protocol  $\Pi$  UC-realizes functionality  $\mathcal{F}$  in the presence of a globally accessible functionality  $\mathcal{G}$ , then for any protocol  $\rho$ , the protocol  $\rho^{\mathcal{F} \rightarrow \Pi}$  UC emulates  $\rho$  in the presence of  $\mathcal{G}$ .

While this theorem suffices for most of our uses, for our multi-level use of UCGS we would like to additionally show that  $\rho^{\mathcal{F} \rightarrow \Pi}$  UC emulates  $\rho$  even in the presence of  $\Pi_{\mathcal{G}}$ , where  $\Pi_{\mathcal{G}}$  is a protocol that UC-emulates  $\mathcal{G}$ . However, such implication is not true in general [6, 28].

We get around this problem by proving the following simple-but-useful lemma (Lemma 1 in Section 2): Assume that  $\Pi_{\mathcal{G}}$  UC-realizes  $\mathcal{G}$  via a simulator  $\mathcal{S}$ , then any protocol  $\rho$  that UC-realizes  $\mathcal{F}$  in the presence of  $\mathcal{G}^{\mathcal{S}}$  also UC-realizes  $\mathcal{F}$  in the presence of  $\Pi_{\mathcal{G}}$ , where  $\mathcal{G}^{\mathcal{S}}$  is the functionality that combines  $\mathcal{G}$  with  $\mathcal{S}$  in the natural way. We then show that, for the protocols  $\rho$  in this work, having access to  $\mathcal{G}^{\mathcal{S}}$  suffices for them to realise their respective functionalities  $\mathcal{F}$ .

**Multiple levels of corruptions.** The UC framework allows the adversary to adaptively and individually corrupt each party in each module within a composite protocol. While this is very general, it makes the handling of party-wise corruption events (where typically the internal states of multiple modules belonging to the party are exposed together) rather complex. We thus adopt a somewhat simpler modeling of party corruption: When the environment corrupts a protocol/-functionality belonging to a party, it obtains the state of the party corresponding to that module as well as all the sub-modules used within; (1) A corrupted module forwards the corruption notice

to all its subroutines. (2) Each subroutine responds with a local state for the corrupted party. (3) The module collects the local states of the subroutines together with its own and reports them to the environment. (If the corrupted module is a functionality, it asks its simulator to produce a local state corresponding to the corrupted party.) In addition to being simpler, this modeling provides a tighter correspondence between the real and ideal executions and is thus preferable whenever realizable (which is the case in this work).

## 1.5 Related Work

This section briefly surveys the state of the art for security analyses of the Signal architecture in particular and end-to-end secure messaging in general, highlighting the differences from and similarities to the present work.

There is a long line of research into the design and analysis of two-party Signal messaging, its subcomponents, and variants of the Signal architecture; this research builds upon decades of study into key exchange protocols (e.g., [8, 9, 26, 27]) and self-healing after corruption (e.g., [24, 32, 34]). Some of these secure messaging analyses purposely consider a limited notion of adaptive security in order to analyze instantiations of Signal based on standardized crypto primitives (e.g., [1, 10, 36, 40, 64]). Other works consider a strong threat model in which the adversary is malicious, fully adaptive, and can tamper with local state [4, 7, 39, 41, 53], which then intrinsically requires strong HIBE-like primitives that depart from the Signal specification. By contrast, we follow a middle ground in this work: our adversary is fully adaptive and has no restrictions on when it can corrupt a party, yet its corruptions are instantaneous and passive.

We stress that, while this work is inspired by the clear game-based modeling and analyses of Signal in works like Alwen et al. [1], our modeling differs in a number of significant ways. For one, our analysis provides a composable security guarantee. Furthermore, we directly model secrecy against a fully adaptive adversary that decides who and when to corrupt based on all the information seen so far. In contrast, Alwen et al. [1] guarantee secrecy only against a *selective* adversary that determines ahead of time who and when it will corrupt.

There are two prior works that perform composable analyses of Signal. In concurrent work to our own, Bienstock et al. [12] provide an alternative modeling of an ideal secure messaging within the UC framework and demonstrate how the Signal protocol can be modeled in a way that is shown to realize their formulation of ideal secure messaging. Like this work, they demonstrate several shortcomings of previous formulations, such as overlooking the effect of choosing keys too early or keeping them around for too long. Additionally, Jost, Maurer, and Mularczyk [41] conduct an analysis in the constructive cryptography framework. Their work provides a model for message transmission as well as one for ratcheting protocols.

That said, the ideal functionalities in [12] and [41] differ from our  $\mathcal{F}_{\text{SM}}$  in several ways.

- *Differences between our work and both of [12, 41]:* Their modeling does not account for the session initiation process, nor the PKI and long-term key modules that are an integral part of any secure messaging application. Additionally, they include the communication medium as part of the protocol, which (a) makes it harder to argue about immediate decryption and (b) means that an instantiation of Signal would have to include an entire TCP/IP stack, which weakens modularity and inhibits the use of Signal as a sub-routine within larger functionalities.
- *Additional differences with Bienstock et al. [12]:* While the modeling of the Signal protocol in [12] follows the traditional partitioning into continuous key exchange, epoch key derivation and authenticated encryption modules, it does not formalize this partitioning within the UC



framework as done in this work; commensurately, they model all key derivation modules as random oracles. Also, their modeling forces the “calling protocol” to keep track of the message IDs for the Secure Messaging functionality/protocol, and assumes uniqueness of the IDs which might create a security risk. On the other hand, [12] accounts for adversarial choice of randomness, which our modeling does not account for. They also propose and analyze an enhancement of the double ratchet structure, which they call the Triple Ratchet protocol, that helps parties regain security faster following a compromise event.

- *Additional differences with Jost, Maurer, and Mularczyk [41]:* To model ratcheting components in a modular fashion, [41] introduces a global event history defined for the entire real (or ideal) world, where a history is a list of events having happened at a module (e.g. a message being input by Alice or one having leaked to the adversary). The event history is visible to the environment, the resources, and the simulator. The security of a resource is then allowed to depend on the global event history. They make composable statements about continuous key agreement protocols in their model (a notion introduced by Alwen et al. [1]) by restricting the adversaries capabilities in the real world as a function of the global event history. Alternatively, for the case of unrestricted adversaries [41] provide a HIBE-based implementation which is quite different than that of Signal (and ours) and requires heavier cryptographic primitives. Additionally, they transform their HIBE-based protocol into one that is fully composable via a technique that requires a restriction on the number of messages a party can send before receiving a response from the other party.

**Comparison to the CRYPTO 2022 version of this work.** In addition to providing full specifications of the ideal functionalities and protocols, as well as full proofs of security, the current version of  $\mathcal{F}_{\text{SM}}$  provides a marginally weaker guarantee than the version in [25], in terms of recovery from session compromise. Specifically, the version of  $\mathcal{F}_{\text{SM}}$  in [25] guarantees that a session regains its security at the third epoch since the last party corruption event, unless the environment chooses to prevent recovery by delivering a specifically crafted bogus message to the party that expects to receive the first message in a new epoch. If such a message is delivered, the session is forked indefinitely.

The current version of  $\mathcal{F}_{\text{SM}}$  allows the environment to deliver yet another form of bogus message to the party that expected the first message in a new epoch. However, in this case the recovery is only delayed so as to take three epochs from the delivery event. We view this weakening as having a secondary impact of the actual security of the protocol, yet it simplifies the analysis and prevents us from using stronger cryptographic hardness assumptions. See more details in Sections ?? The current version of  $\mathcal{F}_{\text{SM}}$  allows the environment to deliver yet another form of bogus message to the party that expected the first message in a new epoch. However, in this case the recovery is only delayed so as to take three epochs from the delivery event. We view this weakening as having a secondary impact of the actual security guarantees provided by our treatment, yet it simplifies the analysis and prevents us from using stronger cryptographic hardness assumptions. See more details in Sections 3.2 and 5.1.

## 2 Universally Composable Security

### 2.1 Universal Composability: A Primer

UC security [22] is an instantiation of the simulation-based security paradigm in which the real world execution of a protocol  $\Pi$  is compared with an idealized abstraction  $\mathcal{F}$ . The UC security



framework gives two special powers to the distinguisher (also known as the *environment*  $\text{Env}$ ) in order to provide maximum flexibility to distinguish  $\Pi$  from  $\mathcal{F}$ : direct interaction with either the protocol or the abstract specification (or, ideal functionality) by way of providing all inputs and obtaining all outputs, as well as interaction via pre-specified adversarial channels — either directly with the protocol or else with the specification, where the interaction is mediated by a special computational component called the simulator.

A protocol  $\Pi$  is deemed to be a *UC-realization* of the functionality  $\mathcal{F}$  if there exists an efficient simulator  $\mathcal{S}$  such that no environment can tell whether it is interacting with  $\Pi$ , or else with  $\mathcal{F}$  and  $\mathcal{S}$ , namely  $\text{exec}_{\mathcal{E},\Pi} \approx \text{exec}_{\mathcal{E},\mathcal{F},\mathcal{S}}$  for all polytime environments  $\mathcal{E}$ .

The UC model also formulates a stylized model of execution that is sufficiently general so as to capture most realistic computational systems. In this model, machines (which are the basic computational entity) can interact by sending messages to each other. Messages take the form of either input, or output, or “side information”, where the latter model either adversarial leakage of information from a machine, or adversarial influence on the behavior of the machine. Machines can create other machines dynamically during an execution of the system, and each new machine is given an identity that includes its own program and is accessible to the machine itself. When a machine sends input or output to another machine, the system lets the recipient machine know the full identity of the sender.

By convention, identities consist of two fields, called session ID (**sid**) and party ID (**pid**). All machines that have the same **sid** and same program  $\Pi$  are called a *session* of  $\Pi$ . These machines are also called the *main machines* of this instance. In this paper, we use the notation  $(\mathcal{F}, \text{sid})$  to denote the specific instance of the machine running the code of  $\mathcal{F}$  that has session id **sid**; this combination uniquely identifies a single machine. Within **sid**, many (but not all) of the functionalities in this work will include the **pid** of the parties that are permitted to invoke this session; this serves as a form of access control.

If machine  $A$  has sent input to machine  $B$  in an execution, or machine  $B$  sent output to machine  $A$ , then we say that  $B$  is a subroutine of  $A$ . Protocol session  $B$  is a subroutine of protocol session  $A$  if some machine in  $B$  is a subroutine of some machine in  $A$ . The extended session of some protocol session  $A$  in an execution of a system includes the transitive closure of all the protocol sessions under the subroutine relation starting from  $A$ .

Remarkably, the UC model of execution considers only a single (extended) instance of the protocol under consideration, leading to relative simplicity of the specification and analysis. Still, the UC framework provides the following generic composition theorem, called the *UC Theorem*: Suppose one proves that a protocol  $\Pi$  UC-realizes  $\mathcal{F}$ , and there exists another “hybrid” protocol  $\rho$  that makes (perhaps many) subroutine calls to functionality  $\mathcal{F}$ . Now, consider the protocol  $\rho^{\mathcal{F} \rightarrow \Pi}$  that replaces all instances of the ideal functionality  $\mathcal{F}$  with the real protocol  $\Pi$ . The composition guarantees that the instantiation  $\rho^{\mathcal{F} \rightarrow \Pi}$  is “just as secure” as the  $\rho$  itself, in the same sense defined above. In this case we say that  $\rho^{\mathcal{F} \rightarrow \Pi}$  UC-emulates  $\rho$ .

**UC with global subroutines.** Crucially, the UC theorem requires that both  $\mathcal{F}$  and  $\Pi$  are *subroutine respecting*. (A protocol is subroutine respecting if the only machines in any extended session of the protocol that take input from a machine that is not part of this extended session, or provides output to a machine that is not part of this extended session, are the main machines of this protocol session.)

While this requirement is both natural and essential, it does not allow for direct, “out of the box” application of the UC theorem in prevalent situations where one wants to decompose systems where multiple protocols (or multiple sessions of the same protocol) use some common construct

as subroutine. In the context of this work, examples include public-key infrastructure, a long-term memory module that is used by multiple sessions, a key generation protocol that is used in multiple epochs and multiple messages in an epoch, or a global construct modeling the random oracle.

First attempts to handle such situations involved extending the UC framework to explicitly allow for multiple sessions of protocols within the basic model of execution [23]. However, this resulted in additional complexity and incompatibility with the basic UC model. More recently, the following formalism has been shown to suffice for capturing universal composition with global subroutines within the basic UC framework [5]:

Say that protocol  $\Pi$  UC-realizes functionality  $\mathcal{F}$  *in the presence of global subroutine  $G$*  if there exists an efficient simulator  $\mathcal{S}$  such that no environment can tell whether it is interacting with  $\Pi$  and  $G$ , or else with  $\mathcal{F}$ ,  $G$ , and  $\mathcal{S}$ . Here  $G$  can be either a single machine or an entire protocol instance, where  $G$  can be a subroutine of  $\Pi$  or of  $\mathcal{F}$ , and at the same time take inputs directly from the environment and provide outputs directly to the environment<sup>4</sup>.

Now, consider protocol  $\rho$  that makes subroutine calls to functionality  $\mathcal{F}$ , and additionally also calls to  $G$ . Then the *UC With Global Subroutines Theorem* states that the protocol  $\rho^{\mathcal{F} \rightarrow \Pi}$ , that is identical to  $\rho$  except that all instances of the ideal functionality  $\mathcal{F}$  are replaced by instances of the real protocol  $\Pi$ , UC emulates  $\rho$  *in the presence of  $G$* . Note that in both  $\rho$  and in  $\rho^{\mathcal{F} \rightarrow \Pi}$ ,  $G$  may take input from and provide outputs to multiple instance of  $\Pi$  (or of  $\mathcal{F}$ ), of  $\rho$ , and also directly to the environment.

## 2.2 New Capabilities: Global Functionalities, Adversarially Provided Code

We describe two new modeling techniques that simplify our analysis, and may be of more general interest.

**Instantiating global functionalities.** The first technique relates to applying the UC theorem to global functionalities. Assume that we have a protocol  $\rho$  that UC-realises a functionality  $\mathcal{F}$  in the presence of a global functionality  $\mathcal{G}$ . Assume also that we have a protocol  $\Pi_{\mathcal{G}}$  that UC-realises the functionality  $\mathcal{G}$ . It may be tempting to deduce directly that  $\rho$  UC-realizes  $\mathcal{F}$  in the presence of  $\Pi_{\mathcal{G}}$ ; however, this implication is false in general [6, 28]. Still, the following implication does hold: If  $\Pi_{\mathcal{G}}$  UC-realises  $\mathcal{G}$ , and  $\rho$  UC-realizes  $\mathcal{F}$  in the presence of  $\Pi_{\mathcal{G}}$ , then  $\rho$  UC-realizes  $\mathcal{F}$  in the presence of  $\mathcal{G}$ .

We use this fact as follows: since  $\Pi_{\mathcal{G}}$  UC-realizes  $\mathcal{G}$ , there must exist a simulator  $\mathcal{S}$  such that no environment can distinguish between an interaction with  $\Pi_{\mathcal{G}}$  and an interaction with  $\mathcal{G}$  and  $\mathcal{S}$ . Now consider the machine  $\mathcal{G}^{\mathcal{S}}$  that represents the combination of  $\mathcal{G}$  and  $\mathcal{S}$  (the communication between  $\mathcal{G}$  and  $\mathcal{S}$  are now internal to the combined machine  $\mathcal{G}^{\mathcal{S}}$ ). We observe that  $\Pi_{\mathcal{G}}$  and  $\mathcal{G}^{\mathcal{S}}$  UC-emulate each other, more specifically  $\Pi_{\mathcal{G}}$  UC-emulates  $\mathcal{G}^{\mathcal{S}}$  *and in addition  $\mathcal{G}^{\mathcal{S}}$  UC-emulates  $\Pi_{\mathcal{G}}$* . Hence, instead of demonstrating that  $\rho$  UC-realizes  $\mathcal{F}$  in the presence of  $\Pi$ , it suffices to demonstrate that  $\rho$  UC-realizes  $\mathcal{F}$  in the presence of  $\mathcal{G}^{\mathcal{S}}$ . That is:

**Lemma 1** *Let  $\Pi_{\mathcal{G}}$  be a protocol that UC-realizes an ideal functionality  $\mathcal{G}$ , and let  $\mathcal{S}$  be a simulator that demonstrates this fact, i.e.  $\text{exec}_{\mathcal{E}, \Pi_{\mathcal{G}}} \approx \text{exec}_{\mathcal{E}, \mathcal{G}, \mathcal{S}}$ . Then protocols  $\Pi_{\mathcal{G}}$  and  $\mathcal{G}^{\mathcal{S}}$  UC-emulate each other. Consequently, for any protocol  $\rho$  and ideal functionality  $\mathcal{F}$  we have that  $\rho$  UC-realizes  $\mathcal{F}$  in the presence of  $\Pi_{\mathcal{G}}$  if and only if  $\rho$  UC-realizes  $\mathcal{F}$  in the presence of  $\mathcal{G}^{\mathcal{S}}$ .*

<sup>4</sup>Since the standard UC model of execution only considers an interaction of the environment with a single instance of some protocol, [5] first demonstrate that, without loss of generality, an instance of  $\mathcal{F}$  alongside  $G$  exhibits the same behavior as an instance of a “dummy protocol”  $\delta$  that simply runs  $\Pi$  alongside  $G$  as subroutines of  $\delta$ .

**Modeling don't-care code and immediate response.** The second modeling technique can actually be thought of as further formalization of a technique that has been used in several works for different purposes. Specifically, it is sometimes convenient to have an ideal functionality expect to obtain from the adversary a piece of code that will be later run by the functionality under some conditions. In this work we have the additional requirement that the adversarially provided code should be readily available for use by ideal functionalities immediately upon first invocation.

Specifically, the formalism proceeds as follows. We formulate an ideal functionality,  $\mathcal{F}_{\text{lib}}$  (see Figure 2), to be used as a global functionality in the system. The adversary can upload code to be used by ideal functionalities. (We call the uploaded code  $\mathcal{I}$  an internal adversary, as it can be thought of as an adversary run internally by the functionality without direct access to the environment.) More specifically, each uploaded code is associated with an identifier. Ideal functionalities can then ask  $\mathcal{F}_{\text{lib}}$  for the code associated with a given identifier.  $\mathcal{F}_{\text{lib}}$  then either responds with the uploaded code, or else returns an error message. For further convenience, we allow uploaded code to refer and use other uploaded pieces of code, referring to these pieces of code by their identifiers. (This is akin to “code linking” in standard software packages.)

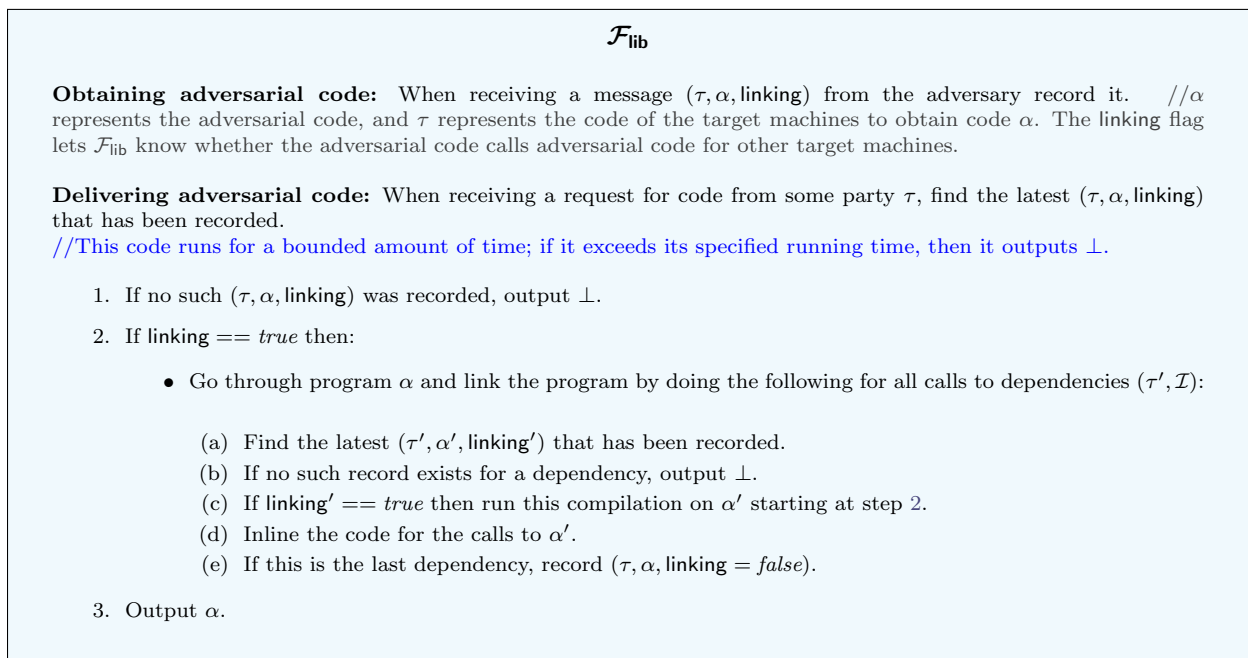


Figure 2: The code library functionality,  $\mathcal{F}_{\text{lib}}$

In this work this technique is used to model the immediate encryption and immediate decryption properties of secure messaging. The uploaded internal adversarial code is specific to the protocol that realizes the functionality, and essentially it acts as the ideal-world simulator during an honest execution. This ensures that the functionality does not need to wait for the adversary to encrypt or decrypt messages that are not corrupted. In cases where the message or ciphertext is corrupted, the fully adaptive adversary is called for input (for example, asking  $\mathcal{A}$  to encrypt a message or decrypt a ciphertext). The state of the static code  $\mathcal{I}$  is maintained across calls in a variable  $\text{state}_{\mathcal{I}}$ , and it is sent to the adversary upon corruption. Here the fact that  $\mathcal{F}_{\text{lib}}$  is global is crucial, in allowing the static code to be already defined at the time that the functionality is instantiated.

The linking feature of  $\mathcal{F}_{\text{lib}}$  becomes handy when writing simulators for protocols that (a) realize

an ideal functionality  $\mathcal{F}$  that expect adversarial code, and (b) make use of another ideal functionality  $\mathcal{F}'$  that also expects adversarial code. In such situations the code that the simulator will upload to  $\mathcal{F}_{\text{lib}}$  will link to the code that is to be uploaded by the adversary (or simulator) for  $\mathcal{F}'$ .

### 3 Modelling Secure Messaging

This section presents our overall modeling of secure messaging. Relying on the overview provided in Section 1.2, we dive right into the detailed descriptions of the main components. Section 3.1 describes the functionalities ( $\mathcal{F}_{\text{DIR}}$ ,  $\mathcal{F}_{\text{LTM}}$ , and  $\mathcal{F}_{\text{PRO}}$ ) that are global with respect to the top level secure messaging functionality  $\mathcal{F}_{\text{SM}}$  (other than  $\mathcal{F}_{\text{lib}}$  which was described on Page 18). Section 3.2 describes the the top-level secure messaging functionality  $\mathcal{F}_{\text{SM}}$ .

#### 3.1 Global Functionalities

In Section 2.2 we described the functionality  $\mathcal{F}_{\text{lib}}$  that we use to model the instant encryption and instant decryption properties of secure messaging. In this section we describe three other global subroutines used in our work. The global subroutines presented here are adaptations of well-studied UC functionalities.

$\mathcal{F}_{\text{DIR}}$ . First, we construct a public key infrastructure functionality called *the directory*  $\mathcal{F}_{\text{DIR}}$  (Fig. 3). This functionality allows each party (through their long-term module  $\mathcal{F}_{\text{LTM}}$ ) to upload their long-term public identity key  $\text{ik}$  to the directory. The directory also supports the execution of the triple Diffie-Hellman protocol that binds the session to the identity of the two participants. Specifically,  $\mathcal{F}_{\text{DIR}}$  allows the initial sender in a session of secure messaging to fetch the recipient’s long- and short-term keys as well as a unique one time key for the session.

$\mathcal{F}_{\text{PRO}}$ . Next, we design a programmable random oracle module  $\mathcal{F}_{\text{PRO}}$  (Fig. 4) that will be used to generate one-time keys during the symmetric ratcheting step. We need a random oracle for equivocation against an adaptive attacker, given that there is no bound on the number of message keys that a party might use within an epoch. Specifically, we use the following random oracle functionality from [18]. Anyone can query the random oracle, but only the (real or ideal world) adversary has the power to program it.

$\mathcal{F}_{\text{LTM}}$ . Finally, we design a module  $\mathcal{F}_{\text{LTM}}$  (Fig. 5) with two responsibilities: local storage of long-term public and private cryptographic key material, and performing computations that require access to the long term private keys of a party. Intuitively, one can think of  $\mathcal{F}_{\text{LTM}}$  as a trusted execution enclave or secure co-processor that performs the Diffie-Hellman operations associated with the long term keys. Looking ahead, our secure messaging protocol only invokes  $\mathcal{F}_{\text{LTM}}$  when establishing a new session of secure messaging; it is not invoked by ongoing communications. The functionality also has several methods to support the execution of Signal’s triple Diffie-Hellman protocol [46]. Concretely, each party can generate short-term rotating and one-time keys that limit the period of vulnerability if a long-term key is compromised.

#### 3.2 The Secure Messaging Functionality, $\mathcal{F}_{\text{SM}}$

This section presents our *secure instant messaging functionality*  $\mathcal{F}_{\text{SM}}$  (the complete details of which can be found in Figure 6.) This functionality takes two types of inputs from the peers that are using it, `SendMessage` and `ReceiveMessage`; `SendMessage` inputs are used to encapsulate messages

### $\mathcal{F}_{\text{DIR}}$

$\mathcal{F}_{\text{DIR}}$  has a fixed session ID, denoted  $\mathcal{F}_{\text{DIR}}$ .

**RecordKeys:** On input  $(\text{RecordKeys}, \text{pid}, \text{ik}_{\text{pid}}^{\text{pk}}, \text{rk}_{\text{pid}}^{\text{pk}})$  from  $(\mathcal{F}_{\text{LTM}}, \text{pid})$  do: //For simplicity  $\mathcal{F}_{\text{DIR}}$  records just one public key per pid. Multiple public keys per "user level party" can be handled by having multiple pid's per party.

1. Set  $\text{ik}_{\text{pid}}^{\text{pk}}$  and  $\text{rk}_{\text{pid}}^{\text{pk}}$  as the identity and rotating keys corresponding to  $\text{pid}$ , respectively, and
2. Output  $(\text{RecordKeys}, \text{pid}, \text{Success})$  to the caller.

**ReplaceRotatingKey:** On input  $(\text{ReplaceRotatingKey}, \text{pid}, \text{rk}_{\text{pid}}^{\text{pk}})$  from  $(\mathcal{F}_{\text{LTM}}, \text{pid})$ , replace the rotating key corresponding to  $\text{pid}$  with  $\text{rk}_{\text{pid}}^{\text{pk}}$  and Output  $(\text{ReplaceRotatingKey}, \text{pid}, \text{Success})$ .

**StoreOnetimeKeys:** On input  $(\text{StoreOnetimeKeys}, \text{pid}, \text{ls})$  from  $(\mathcal{F}_{\text{LTM}}, \text{pid})$ , do:

1. If the list  $\text{onetime\_keys}_{\text{pid}}$  corresponding to  $\text{pid}$  doesn't exist, then create it.
2. Append  $\text{ls}$  to  $\text{onetime\_keys}_{\text{pid}}$  and Output  $(\text{StoreOnetimeKeys}, \text{pid}, \text{Success})$  to the caller.

**GetInitKeys:** On input  $(\text{GetInitKeys}, \text{pid}_j, \text{pid}_i)$ :

// $\text{pid}_j$  is the responder and  $\text{pid}_i$  is the initiator.

1. If there is no entry for  $\text{pid}_j$  then output  $(\text{GetInitKeys}, \text{Fail})$  to the caller.
2. Choose the first key  $\text{ok}_{\text{pid}_j}^{\text{pk}}$  from the list  $\text{onetime\_keys}_{\text{pid}_j}$  (If the list is empty then let  $\text{ok}_{\text{pid}_j}^{\text{pk}} = \perp$ .)
3. Remove  $\text{ok}_{\text{pid}_j}^{\text{pk}}$  from the list  $\text{onetime\_keys}_{\text{pid}_j}$ .
4. Output  $(\text{GetInitKeys}, \text{pid}_i, \text{ik}_{\text{pid}_i}^{\text{pk}}, \text{rk}_{\text{pid}_i}^{\text{pk}}, \text{ok}_{\text{pid}_j}^{\text{pk}})$  to the caller.

**GetResponseKeys:** On input  $(\text{GetResponseKeys}, \text{pid}_i)$  from a machine with party id  $\text{pid}_j$ : // $\text{pid}_j$  is the responder.

1. Send  $(\text{GetResponseKeys}, \text{pid}_i, \text{ik}_{\text{pid}_i}^{\text{pk}})$  to the caller.

**GetRotatingKey:** On input  $(\text{GetRotatingKey}, \text{pid})$ , do: If there is no entry for  $\text{pid}$  then output  $(\text{GetRotatingKey}, \text{Fail})$  to the caller. Else  $(\text{GetRotatingKey}, \text{pid}, \text{rk}_{\text{pid}}^{\text{pk}})$  to the caller.

Figure 3: The Public-Key Directory Functionality,  $\mathcal{F}_{\text{DIR}}$

### $\mathcal{F}_{\text{PRO}}$

On input  $(\text{HashQuery}, m, \ell)$ :

1. If there is a record  $(m, h)$ 
  - If  $|h| \geq \ell$ : let  $h'$  be the first  $\ell$  bits of  $h$ . // $\mathcal{F}_{\text{PRO}}$  returns prefixes of already-computed entries.
  - If  $|h| < \ell$ : choose  $h_{\text{end}} \xleftarrow{\$} \{0, 1\}^{\ell - |h|}$ , let  $h' = h || h_{\text{end}}$ , and replace the record  $(m, h)$  with  $(m, h')$ .

Else choose  $h' \xleftarrow{\$} \{0, 1\}^{\ell}$  and record  $(m, h')$ .
2. Output  $(\text{HashQuery}, h')$  to the caller.

On message  $(\text{Program}, m, h)$  from the adversary:

1. If there is no record  $(m, h')$ , then record  $(m, h)$ . Send  $(\text{Program})$  to the adversary. //If  $m$  has already been queried then programming fails silently.

Figure 4: The Programable Random Oracle Functionality,  $\mathcal{F}_{\text{PRO}}$

## $\mathcal{F}_{\text{LTM}}$

$\mathcal{F}_{\text{LTM}}$  is parameterized by a specific activator program *Root*, a key derivation function *HKDF* and key generation function *keyGen()*, and an algebraic group  $\mathbb{G}$ . All algebraic operations are done in  $\mathbb{G}$ . The local session ID is of the form  $(\mathcal{F}_{\text{LTM}}, \text{pid})$ . Inputs from senders whose party ID is different than *pid* are ignored.

**Initialize:** On input (*Initialize*) from  $(\text{pid}, \text{Root})$  do: If this is not the first activation then end the activation. Else:

1. Create an empty list  $\text{onetime\_keys}_{\text{pid}} = []$ . Also, choose and record the key pairs  $(\text{ik}_{\text{pid}}^{\text{sk}}, \text{ik}_{\text{pid}}^{\text{pk}}), (\text{rk}_{\text{pid}}^{\text{sk}}, \text{rk}_{\text{pid}}^{\text{pk}}) \xleftarrow{\mathbb{S}} \text{keyGen}()$ , which will be called the party's identity key-pair and rotating key-pair, respectively.
2. Provide input (*RecordKeys*,  $\text{pid}, \text{ik}_{\text{pid}}^{\text{pk}}, \text{rk}_{\text{pid}}^{\text{pk}}$ ) to  $\mathcal{F}_{\text{DIR}}$ .

**UpdateRotatingKey:** On input (*UpdateRotatingKey*) from  $(\text{pid}, \text{Root})$ , do:

1. Replace the rotating key pair with a new key pair  $(\text{rk}_{\text{pid}}^{\text{sk}}, \text{rk}_{\text{pid}}^{\text{pk}}) \xleftarrow{\mathbb{S}} \text{keyGen}()$ .
2. Provide input (*ReplaceRotatingKey*,  $\text{pid}, \text{rk}_{\text{pid}}^{\text{pk}}$ ) to  $\mathcal{F}_{\text{DIR}}$ .

**GenOnetimeKeys:** On input (*GenOnetimeKeys*,  $\text{pid}, j$ ) from  $(\text{pid}, \text{Root})$ , do:

1. Choose  $j$  new key pairs  $(\text{ok}_1^{\text{sk}}, \text{ok}_1^{\text{pk}}), \dots, (\text{ok}_j^{\text{sk}}, \text{ok}_j^{\text{pk}}) \xleftarrow{\mathbb{S}} \text{keyGen}()$  and append them to  $\text{onetime\_keys}_{\text{pid}}$ .
2. Provide input (*StoreOnetimeKeys*,  $\text{pid}, \text{ok}_1^{\text{pk}}, \dots, \text{ok}_j^{\text{pk}}$ ) to  $\mathcal{F}_{\text{DIR}}$ .

**ConfirmRegistration:** On input (*ConfirmRegistration*) from  $(\mathcal{F}_{\text{SM}}, \text{pid})$  or  $(\Pi_{\text{SGNL}}, \text{pid})$ , do:

1. If *pid* has already called (*Initialize*), output (*ConfirmRegistration*, *Success*).
2. Otherwise output (*ConfirmRegistration*, *Fail*).

**ComputeSendingRootKey:** On input (*ComputeSendingRootKey*,  $\text{ik}_{\text{partner}}^{\text{pk}}, \text{rk}_{\text{partner}}^{\text{pk}}, \text{ok}_{\text{partner}}^{\text{pk}}$ ) from a machine with PID *pid* and code  $\Pi_{\text{eKE}}$ :

1. Choose an ephemeral key pair  $(\text{ek}^{\text{sk}}, \text{ek}^{\text{pk}}) \xleftarrow{\mathbb{S}} \text{keyGen}()$  and compute the following:
  - $DH_1 = (\text{rk}_{\text{partner}}^{\text{pk}})^{\text{ik}_{\text{pid}}^{\text{sk}}}$  //Here  $(a)^b$  denotes the exponentiation operation in the respective algebraic group.
  - $DH_2 = (\text{ik}_{\text{partner}}^{\text{pk}})^{\text{ek}^{\text{sk}}}$
  - $DH_3 = (\text{rk}_{\text{partner}}^{\text{pk}})^{\text{ek}^{\text{sk}}}$
  - $DH_4 = (\text{ok}_{\text{partner}}^{\text{pk}})^{\text{ek}^{\text{sk}}}$
2. Output (*ComputeSendingRootKey*,  $\text{HKDF}(DH_1 || DH_2 || DH_3 || DH_4), \text{ek}^{\text{pk}}$ ).

**ComputeReceivingRootKey:** On input (*ComputeReceivingRootKey*,  $\text{ik}_{\text{partner}}^{\text{pk}}, \text{ek}_{\text{partner}}^{\text{pk}}, \text{ok}^{\text{pk}}$ ) from  $(\Pi_{\text{eKE}}, \text{pid})$  do:

1. If list  $\text{onetime\_keys}_{\text{pid}}$  does not contain an entry  $(\text{ok}^{\text{sk}}, \text{ok}^{\text{pk}})$  for the given  $\text{ok}^{\text{pk}}$ , then output an error message to  $(\Pi_{\text{eKE}}, \text{pid})$ .
2. Else, delete the one-time key pair  $(\text{ok}^{\text{sk}}, \text{ok}^{\text{pk}})$  from the list and compute:
  - $DH_1 = (\text{ik}_{\text{partner}}^{\text{pk}})^{\text{rk}_{\text{pid}}^{\text{sk}}}$
  - $DH_2 = (\text{ek}_{\text{partner}}^{\text{pk}})^{\text{ik}_{\text{pid}}^{\text{sk}}}$
  - $DH_3 = (\text{ek}_{\text{partner}}^{\text{pk}})^{\text{rk}_{\text{pid}}^{\text{sk}}}$
  - $DH_4 = (\text{ek}_{\text{partner}}^{\text{pk}})^{\text{ok}^{\text{sk}}}$
3. Output (*ComputeReceivingRootKey*,  $\text{HKDF}(DH_1 || DH_2 || DH_3 || DH_4)$ ).

Figure 5: The Long-Term Keys Module Functionality,  $\mathcal{F}_{\text{LTM}}$

## $\mathcal{F}_{SM}$ (Part 1)

This functionality has a session id  $sid = (sid', pid_0, pid_1)$ . Inputs arriving from machines whose identity is neither  $pid_0$  nor  $pid_1$  are ignored. //For notational simplicity we assume some fixed interpretation of  $pid_0$  and  $pid_1$  as complete identities of the two calling machines.

**SendMessage:** On input (**SendMessage**,  $m$ ) from  $pid$ : //  $pid$  is the extended identity of a party's machine.

1. Let  $i$  be such that  $pid = pid_i$ .
2. **If initialized** not set do: //initialization
  - **If**  $pid \neq pid_0$ , end the activation. **Else**, send (**ConfirmRegistration**) to  $(\mathcal{F}_{LTM}, pid)$ .
  - Upon output (**ConfirmRegistration**,  $t$ ) from  $\mathcal{F}_{LTM}$ : **If**  $t = \text{Success}$ , send (**GetInitKeys**) to  $\mathcal{F}_{DIR}$ . **Else**, end the activation.
  - Upon receiving a response (**GetInitKeys**,  $pid_1, ik_1^{pk}, rk_1^{pk}, ok_1^{pk}$ ) from  $\mathcal{F}_{DIR}$ : **If**  $ok_1^{pk} \neq \perp$ , continue. **Else**, end the activation.
  - Initialize boolean variables **initialized** = *true*, **diverge\_parties** = *false*, integer variables (set to 0) **epoch\_num**<sub>0</sub>, **sent\_msgnum**<sub>0</sub>, **rcv\_msgnum**<sub>0</sub>, **N\_self**<sub>0</sub> = 0 and empty dictionaries **advControl**, **id\_dict**, **N\_dict** = {}.  
// **advControl** will record which parties are adversarially controlled in each epoch, **id\_dict** maps epoch id's to epoch numbers, and **N\_dict** will hold the number of messages sent in each epoch.
  - For all  $e \geq 0$ , set **N\_dict**[ $e$ ] =  $\infty$ . Set **advControl**[**epoch\_num**<sub>0</sub>] =  $\perp$ .
  - Initialize **state** <sub>$\mathcal{I}$</sub>  =  $\perp$  and call  $\mathcal{F}_{ib}$  to obtain the internal code  $\mathcal{I}$ .
3. Increment **sent\_msgnum** <sub>$i$</sub>  by 1.
4. **If** **leak**  $\in$  **advControl**[**epoch\_num** <sub>$i$</sub> ]  $\vee$  **diverge\_parties** = *true*, send backdoor message (**state** <sub>$\mathcal{I}$</sub> , **SendMessage**,  $pid, m$ ) to  $\mathcal{A}$ . **Else**, run  $\mathcal{I}(\text{state}_{\mathcal{I}}, \text{SendMessage}, pid, |m|)$ .
5. Upon obtaining (**state**' <sub>$\mathcal{I}$</sub> , **SendMessage**,  $pid, epoch\_id, c$ ) from  $\mathcal{A}$  or  $\mathcal{I}$ , continue.
6. Update **state** <sub>$\mathcal{I}$</sub>   $\leftarrow$  **state**' <sub>$\mathcal{I}$</sub> .
7. **If** **sent\_msgnum** <sub>$i$</sub>  == 1: //If this is the start of a new sending epoch store the returned epoch id.
  - **If** **epoch\_id**  $\notin$  **keys**(**id\_dict**), record **id\_dict**[**epoch\_id**] = **epoch\_num** <sub>$i$</sub> . **Else**, end the activation.
8. Set  $h = (\text{epoch\_id}, \text{sent\_msgnum}_i, N\_self_i)$ .  
//  $N\_self_i$  is the number of messages sent by  $pid_i$  in its previous sending epoch.
9. **If** **diverge\_parties** = *false*, record ( $pid, h, c, m$ ). **Else**, continue.  
//If the parties' states have diverged, then encrypted messages are no longer recorded.
10. Output (**SendMessage**,  $sid, pid, h, c$ ) to  $pid$ .

**Corrupt:** On input (**Corrupt**,  $pid$ ) from  $\text{Env}$ :

1. **If**  $pid \notin \{pid_0, pid_1\}$ , end the activation. **Else**, continue.
2. Initialize the list **corruptions** <sub>$i$</sub>  if it does not exist.
3. Append (**epoch\_num** <sub>$i$</sub> , **sent\_msg\_num** <sub>$i$</sub> , **received\_msg\_num** <sub>$i$</sub> ) to the list **corruptions** <sub>$i$</sub> .
4. For all epochs  $e \leq \text{epoch\_num}_i$ , set **advControl**[ $e$ ] = {**leak**, **Inject**} to allow the adversary to influence messages still in transit.
5. Set **advControl**[**epoch\_num** <sub>$i$</sub> +1], **advControl**[**epoch\_num** <sub>$i$</sub> +2] = {**leak**, **Inject**}. Set **advControl**[**epoch\_num** <sub>$i$</sub> +3] = {**leak**}.
6. Initialize a list **pending\_msgs** = []. For each record ( $pid_{1-i}, h, c, m$ ) do the following:  
//Make a list of all the messages still in transit to party  $pid_1$ .
  - (a) **If** there already was a successful **ReceiveMessage** for  $h$  (i.e there is a record (**Authenticate**,  $h, c', 1$ ) for some  $c'$ ), continue to the next record. **Else**, append ( $pid_{1-i}, h, c, m$ ) to **pending\_msgs**.
7. Send a request (**state** <sub>$\mathcal{I}$</sub> , **ReportState**,  $i, \text{pending\_msgs}$ ) to  $\mathcal{A}$ .
8. On receiving (**ReportState**,  $i, S$ ) from  $\mathcal{A}$ , output (**Corrupt**,  $S$ ) to  $\text{Env}$ .

(The rest of this functionality is in Fig. 7 on Page 23)

Figure 6: The Secure Messaging Functionality  $\mathcal{F}_{SM}$



## $\mathcal{F}_{SM}$ (Part 2)

(This functionality begins in Fig. 7 on Page 23)

**ReceiveMessage:** On receiving  $(\text{ReceiveMessage}, h = (\text{epoch\_id}, \text{msg\_num}, N), c)$  from  $\text{pid}$ , do:

1. Let  $i$  be such that  $\text{pid} = \text{pid}_i$ .
2. **If** this is the first **ReceiveMessage** request: //initialize the responder
  - **If**  $i = 1$ , continue. **Else**, end the activation.
  - Send **(ConfirmRegistration)** to  $(\mathcal{F}_{LTM}, \text{pid})$ .
  - Upon receiving the output **(ConfirmRegistration,  $t$ )** from  $\mathcal{F}_{LTM}$ : **If**  $t = \text{Success}$ , continue. **Else**, end activation.
  - Send **(GetResponseKeys,  $\text{pid}_0, \text{pid}_1$ )** to  $\mathcal{F}_{DIR}$ .
  - Upon receiving output **(GetResponseKeys,  $\text{pid}_0, \text{ik}_0^{\text{pk}}$ )** from  $\mathcal{F}_{DIR}$ , continue.
  - Initialize the variables  $\text{sent\_msgnum}_1, \text{rcv\_msgnum}_1 = 0$  and  $\text{epoch\_num}_1 = 1$ .
3. **If** there already was a successful **ReceiveMessage** for  $h$  (i.e there is a record **(Authenticate,  $h, c', 1$ )** for some  $c'$ ), or this ciphertext previously failed to authenticate (i.e. a record **(Authenticate,  $h, c, 0$ )** exists), output **(ReceiveMessage,  $h, c, \text{Fail}$ )** to  $\text{pid}$ .
4. **If**  $\text{epoch\_id} \in \text{keys}(\text{id\_dict})$ , set temporary variable  $\text{epoch\_num} = \text{id\_dict}[\text{epoch\_id}]$ . **Else**:
  - (a) **If**  $\text{sent\_msgnum}_i = 0$ , output **(ReceiveMessage,  $h, c, \text{Fail}$ )** to  $\text{pid}$ . **Else**, continue.  
//If  $\text{pid}$  is in a receiving state and hasn't sent any messages in its current sending epoch, it will not accept messages with a new epoch id.
  - (b) Set temporary variable  $\text{epoch\_num} = \text{epoch\_num}_i + 1$ .
  - (c) **If**  $\text{leak} \in \text{advControl}[\text{epoch\_num}]$ :<sup>a</sup>
    - For epochs  $e \in \{\text{epoch\_num}, \text{epoch\_num} + 1\}$ : Set  $\text{advControl}[e] = \{\text{leak}, \text{Inject}\}$ .
    - Add  $\text{leak}$  to  $\text{advControl}[\text{epoch\_num} + 2]$
5. **If**  $\text{msg\_num} > \text{N\_dict}[\text{epoch\_num}]$ , output **(ReceiveMessage,  $h, c, \text{Fail}$ )** to  $\text{pid}$   
//For  $\text{epoch\_num}$ 's that are not finished yet, the  $\text{N\_dict}$  returns a default value of  $\infty$ , so this check passes automatically.
6. **If**  $\text{diverge\_parties} = \text{false} \wedge \text{Inject} \notin \text{advControl}[\text{epoch\_num}]$ , run  $\mathcal{I}(\text{state}_{\mathcal{I}}, \text{Inject}, \text{pid}, h, c)$ . //honest case
7. **Else**, send backdoor message  $(\text{state}_{\mathcal{I}}, \text{Inject}, \text{pid}, h, c)$  to  $\mathcal{A}$ .
8. On receiving  $(\text{state}'_{\mathcal{I}}, \text{Inject}, h, c, v)$  from  $\mathcal{A}$  or  $\mathcal{I}$ :
  - Update  $\text{state}_{\mathcal{I}} \leftarrow \text{state}'_{\mathcal{I}}$ .
  - **If**  $v = \perp$ , record **(Authenticate,  $\text{pid}, h, c, 0$ )** and output **(ReceiveMessage,  $h, c, \text{Fail}$ )**. **Else**, continue.
  - **If**  $\text{diverge\_parties} = \text{false}$  and  $\text{Inject} \notin \text{advControl}[\text{epoch\_num}]$ : //honest case
    - **If** there is a record **(sender,  $h, c^*, m$ )** for header  $h$ , record **(Authenticate,  $h, c, 1$ )** and set  $m^* = m$ . **Else**, output **(ReceiveMessage,  $h, c, \text{Fail}$ )**. //allow authentication of a message with a different mac in the honest case.
  - **Else**: //compromised case
    - Record **(Authenticate,  $h, c, 1$ )**, and set  $m^* = v$ .
    - **If**  $\text{epoch\_id}$  does not appear as a key in  $\text{id\_dict}$  then set  $\text{diverge\_parties} = \text{true}$ . //diverge parties is being set here.
9. **If**  $\text{epoch\_num}_i < \text{epoch\_num}$ : //we only get to this step if decryption is successful
  - Set  $\text{N\_dict}[\text{epoch\_num} - 2] = N$ ,  $\text{epoch\_num}_i += 2$ ,  $N\_self_i = \text{sent\_msgnum}_i$ , and  $\text{sent\_msgnum}_i = 0$ .
10. Output **(ReceiveMessage,  $h, m^*$ )** to  $\text{pid}$ .

<sup>a</sup>This provision (which was missing in the [25] version) allows the environment to delay recovery from corruption by delivering messages with bogus epoch ids, this occurs even if these messages do not contain a payload that passes authentication.

Figure 7: The Secure Messaging Functionality  $\mathcal{F}_{SM}$

for sending to the peer, whereas `ReceiveMessage` inputs are used to decapsulate received packets. The functionality takes a few other ‘modeling inputs’ that are not part of the interface of protocols realising this functionality. The first of these, `Corrupt`, captures the meaning of party corruptions in our model. The others are a number of ‘side channel’ messages from the adversary, these are used to fine-tune the security guarantees within the `SendMessage` and `ReceiveMessage` interfaces.

An instance of  $\mathcal{F}_{\text{SM}}$  is created by some party (namely an ITM, or colloquially a machine) by way of sending the first `SendMessage` input to a machine whose code is  $\mathcal{F}_{\text{SM}}$  and whose session ID is `sid`. (Recall that ideal functionalities have null party identifier.) The creating party encodes its own identity, as well as the identity of the desired peer for the interaction, in the session ID. That is, it is expected that  $\text{sid} = (\text{sid}', \text{pid}_0, \text{pid}_1)$ , where  $\text{pid}_0$  is the identity of the creating party (ie, the *initiator*), and  $\text{pid}_1$  is the identity of the other party. It is stressed that there is no special ‘initiation’ input, namely the first `SendMessage` input already contains the message to be encapsulated.

**SendMessage.** Upon the first  $(\text{SendMessage}, m)$  activation (which is the first activation of the functionality overall)  $\mathcal{F}_{\text{SM}}$  runs its initialization by doing the following four things:

1. Verifying the existence of the instance of  $\mathcal{F}_{\text{LTM}}$  corresponding to the initiator  $\text{pid}_0$ .
2. Checking if the desired peer ( $\text{pid}_1$ ) has an available one-time key  $\text{ok}^{\text{pk}}$  registered with the directory  $\mathcal{F}_{\text{DIR}}$ .
3. Initiating variables that will record subsequent epoch identifiers, message numerals within each epoch, compromised epochs, etc.
4. Initializing internal adversarial state  $\text{state}_{\mathcal{I}}$  and calling  $\mathcal{F}_{\text{lib}}$  to obtain internal adversarial code  $\mathcal{I}$ .<sup>5</sup>

If the local state of  $\mathcal{F}_{\text{SM}}$  indicates that this `SendMessage` activation is the first one in a new epoch, then  $\mathcal{F}_{\text{SM}}$  will have already allowed the ideal-model adversary (namely, the simulator) to choose a new epoch identifier for this epoch within the `ReceiveMessage` interface. When the epoch is uncompromised,  $\mathcal{I}$  will have been run internally by  $\mathcal{F}_{\text{SM}}$  to allow the simulator to choose the new epoch identifier. Otherwise the epoch is compromised and  $\mathcal{F}_{\text{SM}}$  will have asked the simulator for a new epoch identifier and waited to receive a response. In both cases  $\mathcal{F}_{\text{SM}}$  will have verified that the newly chosen identifier is different than all previously used ones before continuing.

After initialization during the first activation and at the start of all subsequent activations,  $\mathcal{F}_{\text{SM}}$  lets the ideal-model adversary choose the ciphertext  $c$  that will correspond to  $m$ . There are two cases (**Case 1: Uncompromised sending epoch.**) The functionality allows the simulator to make the choice of ciphertext by running the code  $\mathcal{I}$  internally with only the length of  $m$  as input. (**Case 2: Compromised sending epoch.**) The message  $m$  is leaked in full to the simulator by  $\mathcal{F}_{\text{SM}}$  who then waits for the simulator to send back its chosen ciphertext  $c$ .

Finally,  $\mathcal{F}_{\text{SM}}$  records  $(m, c, h)$  where  $h$  is the ‘header information’ that includes the epoch identifier `epoch_id` of the message and the message number `msg_num` in the epoch. The output  $(c, h)$  is sent to  $\text{pid}_0$ . As long as the epoch ids are unique, no two records of encrypted messages have the same header information. Indeed, uniqueness is the only property that the epoch ids need to satisfy.

---

<sup>5</sup>When no party is compromised, the functionality never hands over control to the simulator. Instead, it allows adversarial choices by internally running the adversarially provided code  $\mathcal{I}$ . This enables the functionality to achieve immediate decryption.

**ReceiveMessage.** This input allows the receiving party to perform an “idealized authenticated decryption” operation, even though the ciphertext was generated without knowledge of the message and  $\mathcal{F}_{\text{SM}}$  itself has no keying material.

Upon the first (**ReceiveMessage**,  $c, h$ ) input,  $\mathcal{F}_{\text{SM}}$  perform initialization by (1) Verifying that this request is coming from  $\text{pid}_1$  (2) Verifying that the instance of  $\mathcal{F}_{\text{LTM}}$  that corresponds to  $\text{pid}_1$  exists, and (3) Verifying that  $\text{pid}_0$  is registered with  $\mathcal{F}_{\text{DIR}}$ .

After initialization during the first activation and on all subsequent activations, decryption proceeds with respect to the following cases:

1. if there is an encryption record  $(m, c, h)$  for the header  $h$  with the exact ciphertext  $c$ , then  $\mathcal{F}_{\text{SM}}$  returns the corresponding message  $m$  to  $\text{pid}$  regardless of whether the epoch is compromised.
2. Otherwise, if there is some encryption record  $(m, c', h)$  for the header  $h$  with a different ciphertext  $c'$  (presumably,  $c$  is a “mauled ciphertext”) then  $\mathcal{F}_{\text{SM}}$  gives the adversary the latitude to decide whether decryption should succeed. This behavior combines the standard EU-CMA guarantee for the underlying authentication scheme, combined with one-time decryption. There are two sub-cases **1) Uncompromised receiving epoch.** The functionality allows the simulator to decide if the ciphertext  $c$  successfully decrypts by running the code  $\mathcal{I}$  internally and providing it  $h$  and  $c$  as inputs. If the output of  $\mathcal{I}$  isn’t  $\perp$  then decryption succeeds and the message  $m$  is output. Hence immediate decryption is preserved. **2) Compromised receiving epoch.** The ideal-model adversary can cause the receiving party to accept any plaintext of its choosing. To this end, the functionality provides  $h$  and  $c$  as inputs to the adversary and waits for the adversary to return a plaintext  $m'$  to decrypt to. As long as  $m' \neq \perp$ , decryption succeeds.
3. Finally, if there is no record  $(m, c', h)$  for the header  $h$ , then  $\mathcal{F}_{\text{SM}}$  fails to decrypt unless the next receiving epoch of decrypting party is compromised. In the case that the next receiving epoch is compromised, the adversary may attempt a person-in-the-middle attack which succeeds if and only if the honest recipient accepts a new epoch created by the adversary rather than the honest sender. The functionality proceeds as in the previous case and gives the adversary the latitude to decide whether decryption should succeed and what the decrypted plaintext should be. If decryption succeeds then the party has now started receiving messages in a new epoch where its peer can never send a message and therefore the parties’ states and keys have diverged from each other. In that case,  $\mathcal{F}_{\text{SM}}$  notes that the session is *forked* — or in other words, that the parties have *diverged*. In this case, the functionality behaves as if both parties are compromised for the rest of the session. This divergence(i.e. fork) event represents a complete break of security of the session. If decryption doesn’t succeed then the compromise is extended to two epochs past the next receiving epoch of the decrypting party.

It is stressed that  $\mathcal{F}_{\text{SM}}$  only supports decrypting a message for each header once; whenever a successful decryption is output by  $\mathcal{F}_{\text{SM}}$  for some header  $h$ , the functionality will note this and will refuse to participate in all subsequent calls with the same header  $h$ . This ensures forward secrecy. In addition,  $\mathcal{F}_{\text{SM}}$  stores the exact number of messages sent in previous epochs so that honest parties can detect if there is an attempt to inject a message after the sender’s planned ending to an epoch.

Finally, if this happens to be the first successfully received message for party  $\text{pid}_i$  in the newest epoch, then  $\mathcal{F}_{\text{SM}}$  notes that the next **SendMessage** activation with sender  $\text{pid}$  will start a new epoch and allows the ideal-model adversary (namely, the simulator) to choose a new epoch identifier for this sending epoch as described earlier.

**Corrupt.** The response of  $\mathcal{F}_{\text{SM}}$  to party corruption inputs is the core of the security guarantees it provides. The goal is to bound the effect of exposure of the local states of the parties on the loss of security and, and provide guarantees as to how soon security of the communication is restored (if at all).

First, we note two ways in which our modeling of corruptions differs from the traditional. First, corruption is captured as an instantaneous exposure event, as opposed to having separate *corrupt* and *leave* events. Second, corruptions are modeled as inputs (coming from the environment, as opposed to the traditional UC modeling corruption as message coming from the adversary. Both of these changes simplify the mechanics and yield a cleaner model while not restricting the adversary’s power.

In  $\mathcal{F}_{\text{SM}}$  when a party is corrupted, the adversary learns nothing about all messages that have been sent and received by the party until the point of corruption. Furthermore, the adversary obtains no other information on the history of the session such as its duration. This guarantees forward secrecy of a protocol that realises  $\mathcal{F}_{\text{SM}}$ . However, any messages in transit at the time of corruption is sent to the ideal-model adversary instantly, who sends back a simulated state that can be output to the environment. We’ve already described the behavior of  $\mathcal{F}_{\text{SM}}$  within the `SendMessage` and `ReceiveMessage` interfaces while epochs are compromised.

So all that’s left is describing the process by which the functionality recovers from a compromise. The point by which a compromised session becomes uncompromised is Signal-specific, and is determined by inspecting the sequence of encapsulation and decapsulation activations of that two parties. More specifically, Signal partitions the messages sent in a session into *sending epochs*, where each sending epoch is associated with one of the two parties, and consists of all the messages sent by that party from the end of its previous sending epoch until the first time this party successfully decapsulates an incoming message that belongs to the peer’s newest sending epoch. For a compromised session to become uncompromised, the corrupted party must start at least two new sending epochs since the last corruption event. Specifically, if party  $\text{pid}_i$  is corrupted when it is in sending epoch  $e$ , the resulting compromise lasts at least until epoch number  $e + 3$ . The compromise period may be extended further if a party receives a packet with a mauled header that corresponds to some compromised epoch  $e^*$ . In this case, one of two things can happen:

- **If decryption fails, the compromise extends until epoch  $e^* + 2$ .** This outcome corresponds to a weakness in the standalone security of  $\Pi_{\text{eKE}}$ , as it does not include any authentication of keys.<sup>6</sup>
- **If decryption succeeds, the functionality behaves as if both parties are compromised for the rest of the session.** The success of decryption corresponds to the party starting to receive messages in a new epoch where its peer can never send a message, causing the parties’ states and keys to diverge from each other. In this scenario,  $\mathcal{F}_{\text{SM}}$  notes that the session is *forked*—or in other words, that the parties’ states have *diverged*, and the functionality then behaves as if both parties are compromised for the rest of the session. This divergence (i.e., fork) event represents a complete break of security for the session.

## 4 Overview of our Modular Decomposition

This work provides a modular analysis of Signal’s protocol. In this section provide more details about our modular, iterative process for decomposing the Signal architecture into a collection of

---

<sup>6</sup>This provision was missing in the [25] version of this work.

ideal modules that each address a specific purpose. In essence, this section summarizes the theorems that collectively serve to instantiate  $\mathcal{F}_{\text{SM}}$ , and it serves as an organization for the rest of the paper.

While reading this section, we encourage readers to review Fig. 1 on page 10, which graphically depicts the various components and how they fit together.

**A Signal-style realization of  $\mathcal{F}_{\text{SM}}$ .** In our first instantiation depicted by the thick black arrow at the top left of Fig. 1, we decompose the Signal protocol into two ideal modules representing the interconnected components of the double ratchet architecture: (1) a long-lived  $\mathcal{F}_{\text{eKE}}$  component that models the asymmetric ratcheted key exchange in the Signal architecture and (2) multiple copies of short-lived  $\mathcal{F}_{\text{fs.aead}}$  components that model unidirectional forward secure authenticated channels, each representing the combination of authenticated encryption (with associated data) and the symmetric key ratchet from the Signal architecture, for handling the messages associated with a single epoch. Protocol  $\Pi_{\text{SGNL}}$  uses these modules to realize  $\mathcal{F}_{\text{SM}}$ , following the overall logic of Signal’s architecture

The long-lived key exchange module  $\mathcal{F}_{\text{eKE}}$  contributes two main features of the overall protocol: (1) healing from compromise through the incorporation of randomness, and (2) long-term forward secrecy. Meanwhile, the short-lived forward secure encryption modules  $\mathcal{F}_{\text{fs.aead}}$  each grant fine-grained forward secrecy within the epochs they are associated with. By relying on  $\mathcal{F}_{\text{eKE}}$ , the  $\mathcal{F}_{\text{fs.aead}}$  modules additionally enable the overall protocol to realise the long-term forward secrecy and healing properties endowed by  $\mathcal{F}_{\text{eKE}}$ . At this level of abstraction, the reliance of  $\mathcal{F}_{\text{fs.aead}}$  on  $\mathcal{F}_{\text{eKE}}$  is indirect (via the epoch identifiers.) However, at a lower level of abstraction, the security of this module will depend directly on  $\mathcal{F}_{\text{eKE}}$  since the chain keys provided by  $\mathcal{F}_{\text{eKE}}$  will be used eventually to derive the encapsulation and decapsulation keys for individual messages. Next, we will briefly describe these modules along with our main theorem. Full descriptions as well as a rigorous proof of the theorem can be found in Section 5.

- *Signal Protocol ( $\Pi_{\text{SGNL}}$ , Fig. 16):* The protocol  $\Pi_{\text{SGNL}}$  (see Section 5.3) is the top-level protocol that interfaces only with the ideal functionalities  $\mathcal{F}_{\text{eKE}}$  and  $\mathcal{F}_{\text{fs.aead}}$  and realises the functionality  $\mathcal{F}_{\text{SM}}$ . There are three primary takeaways from the design of  $\Pi_{\text{SGNL}}$ : 1) it has the same input-output API as our ideal functionality  $\mathcal{F}_{\text{SM}}$ , 2) it displays an idealized version of the double ratchet with clearly distinct roles for the two subroutines, and finally 3) it moves one level of abstraction closer to the specification of the real Signal architecture. Added features at this level of abstraction include key material stored within party states, explicit accounting for out-of-order messages by holding onto missed message keys, and epochs being identified directly by their `epoch_id` rather than an idealized `epoch_num` ordering.
- *Epoch Key Exchange ( $\mathcal{F}_{\text{eKE}}$ , Fig. 13):* The epoch key exchange functionality  $\mathcal{F}_{\text{eKE}}$  persists throughout the entire secure messaging session between two parties. It represents the root ratchet of the Signal protocol which comprises the public key “backbone” of continuous key agreement in the secure messaging architecture of Signal. At the level of abstraction where  $\Pi_{\text{SGNL}}$  realises  $\mathcal{F}_{\text{SM}}$ , The only role of the epoch key exchange functionality is to produce the epoch id’s that identify the communicating parties’ sending epochs. At this level the functionality interacts only with the protocol  $\Pi_{\text{SGNL}}$  and the global functionalities  $\mathcal{F}_{\text{lib}}$ ,  $\mathcal{F}_{\text{DIR}}$ ,  $\mathcal{F}_{\text{LTM}}$ . However,  $\mathcal{F}_{\text{eKE}}$  doesn’t only produce epoch identifiers. It also maps identifier pairs (`epoch_id0`, `epoch_id1`) to sending and receiving chain keys that can be leveraged by the symmetric ratchet to harness  $\mathcal{F}_{\text{eKE}}$ ’s properties of *healing from compromise* and *coarse-grained forward secrecy*. Looking ahead, this property of  $\mathcal{F}_{\text{eKE}}$  is used two levels of abstraction below this one, where the module  $\mathcal{F}_{\text{mKE}}$  is instantiated by a protocol  $\Pi_{\text{mKE}}$ . The message

key exchange protocol retrieves the chain keys produced by  $\mathcal{F}_{eKE}$  and uses them to produce encapsulation and decapsulation keys for individual messages.

- *Forward Secure Authenticated Encryption* ( $\mathcal{F}_{fs.aead}$ , Fig. 15): Each instance of the forward secure authenticated encryption functionality models the encryption and decryption of messages for a single epoch of the secure messaging system. In each epoch, one of the parties may only send messages while the other may only receive messages. As the name suggests, the forward secure encryption functionality provides forward security by allowing each message to be decrypted exactly once. Additionally, the sender’s (resp. receiver’s)  $\Pi_{SGNL}$  for the epoch may send a `StopEncrypting` (resp. `StopDecrypting`) request to the  $\mathcal{F}_{fs.aead}$  instance, after which the  $\mathcal{F}_{fs.aead}$  instance will no longer allow any more messages to be encrypted (resp. decrypted) by anyone. This way the functionality  $\mathcal{F}_{fs.aead}$  allows for the ‘expiry’ of the epoch it represents once the parties move on in the conversation.

In Section 5 we provide a rigorous specification for each of the modules  $\Pi_{SGNL}$ ,  $\mathcal{F}_{eKE}$ ,  $\mathcal{F}_{fs.aead}$  discussed above. We then prove Theorem 2 (Figure 8): we provide a concrete simulator  $\mathcal{S}_{SM}$  and show that (together with  $\mathcal{F}_{SM}$ ) it is perfectly indistinguishable from the  $\Pi_{SGNL}$  hybrid world.

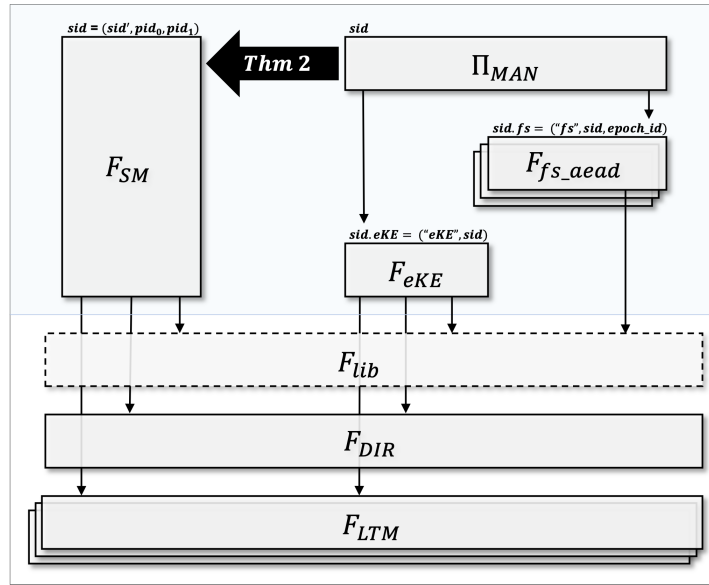


Figure 8: A pictorial rendition of Theorem 2

**Theorem 2** Protocol  $\Pi_{SGNL}$  (perfectly) UC-realizes the ideal functionality  $\mathcal{F}_{SM}$  in the presence of  $\mathcal{F}_{lib}$ ,  $\mathcal{F}_{DIR}$  and  $\mathcal{F}_{LTM}$ .

**Instantiating the public ratchet (realizing  $\mathcal{F}_{eKE}$ ).** The instantiation of  $\mathcal{F}_{eKE}$  in a modular fashion is one of the most delicate parts of this work (alongside the teasing apart of a modular functionality  $\mathcal{F}_{eKE}$  to represent the security provided by Signal’s public key ratchet). Signal’s public key ratchet is specified by  $\Pi_{eKE}$  in Section 6.1 of our work and the instantiation of  $\mathcal{F}_{eKE}$  is depicted in Figure 1 by the lowermost thick black arrow in the diagram. One main challenge, as observed by Alwen et al. [1] and others, is that the key derivation function (KDF) within the public ratchet ( $\Pi_{eKE}$ ), must maintain security if either of the previous root key or the newly generated

ephemeral keys are uncompromised. As described in the introduction, we point to an additional challenge to instantiating  $\mathcal{F}_{eKE}$ , namely that of asserting security at the presence of adaptively chosen key exposures.<sup>7</sup>

Alwen et al. formalized the guarantee they realised  $KDF$  must satisfy by way of constructing a new primitive: a PRF-PRG. In Section 6.2 we formulate a new primitive primitive called a *Cascaded PRF-PRG (CPRFG)* that extends the PRF-PRG concept to provide also the equivocation needed to handle adaptive state exposures (Def. 11). We then provide a plain model instantiation of this primitive based only on punctured PRFs (Thm. 12 in Section 6.2.3). Using this CPRFG construction, we are able to UC-realise the functionality  $\mathcal{F}_{eKE}$  (described in Section 6.3) The proof of Theorem 3 (Figure 9) is quite intricate; See Section 6.3 for the details.

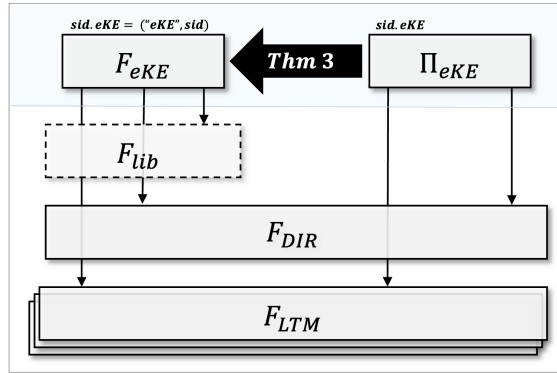


Figure 9: A pictorial rendition of Theorem 3

**Theorem 3** *Assume that  $KDF$  is a CPRFG with security parameter  $\lambda$ , that the DDH assumption holds in the group  $\mathbb{G}$  where  $|\mathbb{G}| \geq 2^\lambda$ . Then protocol  $\Pi_{eKE}$  UC-realizes the ideal functionality  $\mathcal{F}_{eKE}$  in the presence of global functionalities  $\mathcal{F}_{lib}$ ,  $\mathcal{F}_{DIR}$  and  $\mathcal{F}_{LTM}$ .*

### Instantiating unidirectional forward secure authenticated channels (realizing $\mathcal{F}_{fs\_aead}$ ).

In the top right part of Fig. 1, we decompose the symmetric key component of Signal into two smaller pieces: a message key exchange functionality  $\mathcal{F}_{mKE}$  that interfaces with the epoch key exchange to produce the symmetric chain keys, and a one-time-use authenticated encryption routine. This decomposition spans Sections 7 to 9 in the paper.

- *Message Key Exchange ( $\mathcal{F}_{mKE}$ , Fig. 28):* Each message key exchange functionality instance handles the key derivation for the symmetric ratchet for a particular epoch. Specifically, it provides `key_seed`'s to  $\Pi_{aead}$  instances that are then expanded to any length using the global random oracle  $\mathcal{F}_{pRO}$ . The functionality also closes epochs at a certain message number  $N$  when instructed to by  $\Pi_{fs\_aead}$  by generating all `key_seed`'s up to  $N$  and later disallowing the generation of any further key seeds for its epoch. The protocol ( $\Pi_{mKE}$ , Fig. 34) realizes  $\mathcal{F}_{mKE}$  by iteratively applying a length-doubling pseudorandom function to the `chain_key` provided by  $\Pi_{eKE}$  to generate `key_seed`'s. If it needs to skip message key seeds (for example, if the messages arrive out of order),  $\Pi_{eKE}$  applies the PRG several times until reaching the correct `key_seed`, meanwhile storing intermediate `key_seed`'s. To close an epoch at a particular message number  $N$ , it generates all message key seeds up to  $N$  and then deletes the `chain_key`. The functionality

<sup>7</sup>Note that it is relatively straightforward to achieve both these guarantees by modeling the  $KDF$  as a programmable random oracle.



(and protocol) enforces the forward security guarantee by: deleting key seeds for messages that have been retrieved, ensuring that message keys look independent and random from each other (using the PRG), and deleting the `chain_key` when the parties have moved on to future epochs (to prevent message injection in old epochs).

- *Authenticated Encryption with Associated Data* ( $\mathcal{F}_{\text{aead}}$ , Fig. 29): Each authenticated encryption functionality instance handles the encryption, decryption, and authentication of a particular message for a particular epoch and hands the ciphertext or message back to  $\Pi_{\text{fs\_aead}}$ . It gets a `key_seed` from  $\Pi_{\text{mKE}}$  and then asks the adversary to provide a ciphertext  $c$  (while leaking either  $|m|$  or  $m$  depending on whether the epoch is compromised). For decryption, if it gets the same ciphertext back,  $\mathcal{F}_{\text{aead}}$  returns message  $m$ . If it gets a different ciphertext  $c' \neq c$ , it asks the adversary whether it wants to inject a message. Depending on the corruption status,  $\mathcal{F}_{\text{aead}}$  will either return  $\mathcal{A}$ 's message or return  $m$ , or return a failure. The protocol ( $\Pi_{\text{aead}}$ , Fig. 36) realizes  $\mathcal{F}_{\text{aead}}$  by querying the random oracle  $\mathcal{F}_{\text{pRO}}$  on the `key_seed` to get the full `msg_key`. It then computes the ciphertext  $c$  using a One Time Pad and then a secure message authentication code to authenticate  $c$  as well as its `sid` (which contains information about the `pid`'s, `epoch_id`, `msg_num`, and such). To decrypt,  $\Pi_{\text{aead}}$  similarly gets the `key_seed` from  $\mathcal{F}_{\text{mKE}}$  and queries  $\mathcal{F}_{\text{pRO}}$  to expand it. Then,  $\Pi_{\text{aead}}$  decrypts the ciphertext only if the tag verifies.

Within Section 7, we rigorously define all of the functionalities described above and their associated instantiations as cryptographic protocols. We also formally prove that our hybrid world UC-realizes a real world containing only cryptographic protocols rather than functionalities (albeit still in the presence of the global subroutines).

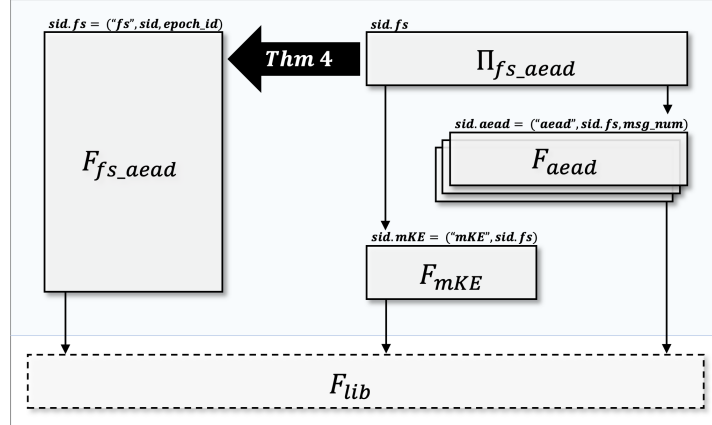


Figure 10: A pictorial rendition of Theorem 4

**Theorem 4** Protocol  $\Pi_{\text{fs\_aead}}$  (perfectly) UC-realizes the ideal functionality  $\mathcal{F}_{\text{fs\_aead}}$  in the presence of  $\mathcal{F}_{\text{lib}}$ .

**Instantiating the Symmetric Ratchet (realizing  $\mathcal{F}_{\text{mKE}}$ ).** By this point we have already described all of the functionalities in our model. As shown in Figure 1, it only remains to construct real-world protocols that realize each of them. Unlike with  $\mathcal{F}_{\text{eKE}}$ , our instantiations of symmetric functionalities are relatively straightforward. In Section 8, we provide a simple instantiation of  $\mathcal{F}_{\text{mKE}}$  based on Signal’s symmetric ratcheting algorithm.

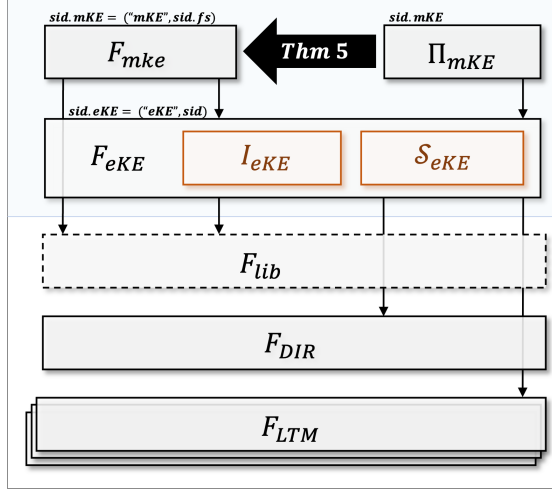


Figure 11: A pictorial rendition of Theorem 5

**Theorem 5** Assume that PRG is a secure length-doubling pseudorandom generator. Then protocol  $\Pi_{mKE}$  UC-realizes  $\mathcal{F}_{mKE}$  in the presence of global functionalities  $\mathcal{F}_{lib}$ ,  $\mathcal{F}_{DIR}$ ,  $\mathcal{F}_{LTM}$ , as well as  $\mathcal{F}_{eKE}^{\Pi_{eKE}}$ , where  $\mathcal{F}_{eKE}^{\Pi_{eKE}} = (\mathcal{I}_{eKE}, \mathcal{S}_{eKE}, \mathcal{F}_{eKE})$ .

**Instantiating authenticated encryption for single messages (realising  $\mathcal{F}_{aead}$ ).** Additionally,  $\mathcal{F}_{aead}$  is a slight variant of the *secure message transmission* functionality  $\mathcal{F}_{SMT}$  that has been analyzed in the original work of Canetti [21] that introduced the UC security framework. We show the following theorem in Section 9.

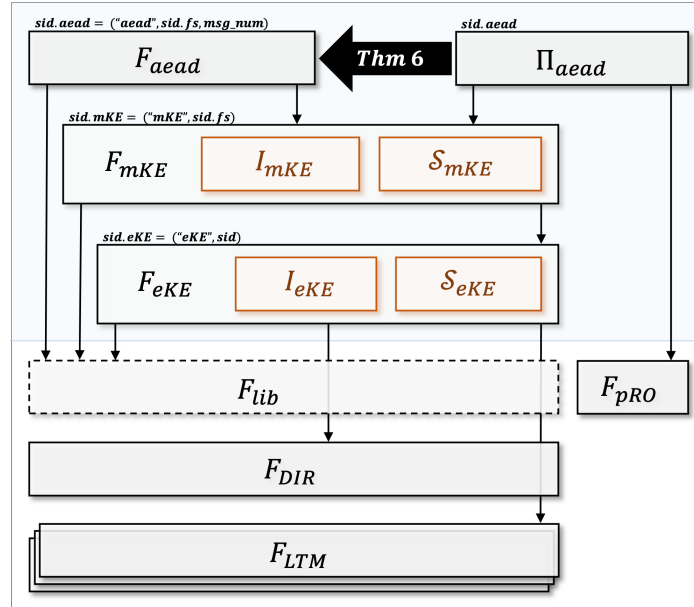


Figure 12: A pictorial rendition of Theorem 6

**Theorem 6** Assuming the unforgeability of (MAC, Verify), protocol  $\Pi_{aead}$  UC-realizes the ideal

functionality  $\mathcal{F}_{\text{aead}}$  in the presence of global functionalities  $\mathcal{F}_{\text{PRO}}, \mathcal{F}_{\text{lib}}, \mathcal{F}_{\text{DIR}}, \mathcal{F}_{\text{LTM}}$ , as well as  $\mathcal{F}_{\text{mKE}}^{\Pi_{\text{mKE}}}$ ,  $\mathcal{F}_{\text{eKE}}^{\Pi_{\text{eKE}}}$ , where  $\mathcal{F}_{\text{mKE}}^{\Pi_{\text{mKE}}} = (\mathcal{S}_{\text{mKE}}, \mathcal{F}_{\text{mKE}})$ , and  $\mathcal{F}_{\text{eKE}}^{\Pi_{\text{eKE}}} = (\mathcal{I}_{\text{eKE}}, \mathcal{S}_{\text{eKE}}, \mathcal{F}_{\text{eKE}})$ .

**Putting it all together (composition theorems).** It is left to put the pieces together and assert the security of the composed protocol as shown in Figure 1. We first use Theorems 3 and 5, together with Proposition 1 to deduce that  $\Pi_{\text{mKE}}$  UC-realizes  $\mathcal{F}_{\text{mKE}}$  in the presence of  $\Pi_{\text{eKE}}$ , as well as  $\mathcal{F}_{\text{LTM}}, \mathcal{F}_{\text{DIR}}$ :

**Corollary 7** *Assume that the KDF module used by  $\Pi_{\text{eKE}}$  is a CPRFG, that the DDH assumption holds for the group  $\mathbb{G}$  used in  $\mathcal{F}_{\text{LTM}}$  and  $\Pi_{\text{eKE}}$ , and that the PRG used in  $\Pi_{\text{mKE}}$  is a secure length-doubling pseudorandom generator. Then protocol  $\Pi_{\text{mKE}}$  UC-realizes  $\mathcal{F}_{\text{mKE}}$  in the presence of  $\Pi_{\text{eKE}}, \mathcal{F}_{\text{lib}}, \mathcal{F}_{\text{DIR}}, \mathcal{F}_{\text{LTM}}$ .*

Next we use Theorem 6 together with Corollary 7 and Lemma 1 to deduce that  $\Pi_{\text{aead}}$  UC-realizes  $\mathcal{F}_{\text{aead}}$  in the presence of  $\Pi_{\text{mKE}}, \Pi_{\text{eKE}}, \mathcal{F}_{\text{LTM}}, \mathcal{F}_{\text{DIR}}$ :

**Corollary 8** *Assume that (MAC, Verify) used in  $\Pi_{\text{aead}}$  is unforgeable, that the KDF module used by  $\Pi_{\text{eKE}}$  is a CPRFG, that the DDH assumption holds for the group  $\mathbb{G}$  used in  $\mathcal{F}_{\text{LTM}}$  and  $\Pi_{\text{eKE}}$ , and that the PRG used in  $\Pi_{\text{mKE}}$  is a secure length-doubling pseudorandom generator. Then protocol  $\Pi_{\text{aead}}$  UC-realizes  $\mathcal{F}_{\text{aead}}$  in the presence of  $\Pi_{\text{mKE}}, \Pi_{\text{eKE}}, \mathcal{F}_{\text{lib}}, \mathcal{F}_{\text{DIR}}, \mathcal{F}_{\text{LTM}}, \mathcal{F}_{\text{PRO}}$ .*

Now, recall that  $\Pi_{\text{fs\_aead}}^{\mathcal{F}_{\text{mKE}} \rightarrow \Pi_{\text{mKE}}, \mathcal{F}_{\text{aead}} \rightarrow \Pi_{\text{aead}}}$  is the protocol that's identical to  $\Pi_{\text{fs\_aead}}$ , except that calls to  $\mathcal{F}_{\text{mKE}}$  are replaced with calls to  $\Pi_{\text{mKE}}$  and calls to  $\mathcal{F}_{\text{aead}}$  are replaced with calls to  $\Pi_{\text{aead}}$ . Then Theorem 4, together with Corollary 8 and the UC with Global Subroutines (UCGS) theorem says that:

**Corollary 9** *Assume that (MAC, Verify) used in  $\Pi_{\text{aead}}$  is unforgeable, that the KDF module used by  $\Pi_{\text{eKE}}$  is a CPRFG, that the DDH assumption holds for the group  $\mathbb{G}$  used in  $\mathcal{F}_{\text{LTM}}$  and  $\Pi_{\text{eKE}}$ , and that the PRG used in  $\Pi_{\text{mKE}}$  is a secure length-doubling pseudorandom generator. Then protocol  $\Pi_{\text{fs\_aead}}^{\mathcal{F}_{\text{mKE}} \rightarrow \Pi_{\text{mKE}}, \mathcal{F}_{\text{aead}} \rightarrow \Pi_{\text{aead}}}$  UC-realizes  $\mathcal{F}_{\text{fs\_aead}}$  in the presence of  $\Pi_{\text{eKE}}, \mathcal{F}_{\text{lib}}, \mathcal{F}_{\text{DIR}}, \mathcal{F}_{\text{LTM}}, \mathcal{F}_{\text{PRO}}$ .*

Finally, recall that  $\Pi_{\text{SGNL}}^{\mathcal{F}_{\text{eKE}} \rightarrow \Pi_{\text{eKE}}, \mathcal{F}_{\text{fs\_aead}} \rightarrow \Pi_{\text{fs\_aead}}^{\mathcal{F}_{\text{mKE}} \rightarrow \Pi_{\text{mKE}}, \mathcal{F}_{\text{aead}} \rightarrow \Pi_{\text{aead}}}}$  is the protocol that's identical to  $\Pi_{\text{SGNL}}$ , except that calls to  $\mathcal{F}_{\text{eKE}}$  are replaced with calls to  $\Pi_{\text{eKE}}$  and calls to  $\mathcal{F}_{\text{fs\_aead}}$  are replaced with calls to  $\Pi_{\text{fs\_aead}}^{\mathcal{F}_{\text{mKE}} \rightarrow \Pi_{\text{mKE}}, \mathcal{F}_{\text{aead}} \rightarrow \Pi_{\text{aead}}}$ . Then Theorem 2, together with Corollary 9 and the UCGS theorem says that:

**Corollary 10** *Assume that (MAC, Verify) used in  $\Pi_{\text{aead}}$  is unforgeable, that the KDF module used by  $\Pi_{\text{eKE}}$  is a CPRFG, that the DDH assumption holds for the group  $\mathbb{G}$  used in  $\mathcal{F}_{\text{LTM}}$  and  $\Pi_{\text{eKE}}$ , and that the PRG used in  $\Pi_{\text{mKE}}$  is a secure length-doubling pseudorandom generator. Then protocol  $\Pi_{\text{SGNL}}^{\mathcal{F}_{\text{eKE}} \rightarrow \Pi_{\text{eKE}}, \mathcal{F}_{\text{fs\_aead}} \rightarrow \Pi_{\text{fs\_aead}}^{\mathcal{F}_{\text{mKE}} \rightarrow \Pi_{\text{mKE}}, \mathcal{F}_{\text{aead}} \rightarrow \Pi_{\text{aead}}}}$  UC-realizes  $\mathcal{F}_{\text{SM}}$  in the presence of  $\mathcal{F}_{\text{lib}}, \mathcal{F}_{\text{DIR}}, \mathcal{F}_{\text{LTM}}, \mathcal{F}_{\text{PRO}}$ .*

## 5 A Signal-style Secure Messaging Protocol: Realizing $\mathcal{F}_{\text{SM}}$

In this section we decompose Signal into two ideal modules ( $\mathcal{F}_{\text{eKE}}, \mathcal{F}_{\text{fs\_aead}}$ ) and a protocol that glues them together ( $\Pi_{\text{SGNL}}$ ). We then prove that these pieces together form an instantiation of the secure messaging functionality  $\mathcal{F}_{\text{SM}}$ . We begin by describing a long-lived component  $\mathcal{F}_{\text{eKE}}$  in Section 5.1 that models the public key ratchet in the Signal architecture. Next, we present a short-lived

component  $\mathcal{F}_{\text{fs\_aead}}$  in Section 5.2 that models a unidirectional forward secure authenticated channel. (This will later, in Section 7, be instantiated using authenticated encryption and a symmetric key ratchet component in the style of the Signal architecture.) In Section 5.3 we present the protocol  $\Pi_{\text{SGNL}}$  that glues together a long lived  $\mathcal{F}_{\text{eKE}}$  and several short lived  $\mathcal{F}_{\text{fs\_aead}}$ . Finally, in Section 5.4 we prove that the signal like protocol  $\Pi_{\text{SGNL}}$  realises the secure messaging functionality we presented in Section 3.2.

## 5.1 The Epoch Key Exchange Functionality $\mathcal{F}_{\text{eKE}}$

The fundamental security objective of the epoch key exchange functionality (see Figure 13 on page 38) is to provide recovery from an adversarial state corruption. This requires that the parties’ secret keys are updated periodically with new random values. Intrinsic to the security goal of post-compromise security and periodic updates is the concept of *epochs*: agreed-upon points in the conversation to re-randomize the secret keys. The functionality  $\mathcal{F}_{\text{eKE}}$  presented in this section is a long lived functionality that allows the parties to refresh their randomness every epoch. With this functionality we capture the properties provided to the Signal architecture by their public key ratchet.

An important usability feature of the public key ratchet from the Signal Architecture is that an online party can refresh their own randomness unilaterally, producing a new key that can be used for messaging without even requiring the other party to be online at the same time. Then, when the other party comes online, they are able to complete the re-randomization of the secret keys and therefore the healing from prior corruption events. To achieve this, their protocol enforces that the parties must take turns refreshing their randomness – ensuring that parties are able to assimilate the randomness introduced in the right order. When  $\text{pid}_i$  receives a ciphertext from  $\text{pid}_{1-i}$  (through  $\mathcal{F}_{\text{fs\_aead}}$ ) marked with an *epoch\_id* that it has not seen before,  $\text{pid}_i$  knows that  $\text{pid}_{1-i}$  has started a new epoch. The party  $\text{pid}_i$  then does a *tentative* ratcheting step using the *epoch\_id* to derive its new keys (without deleting its old keys right away). This new key must then be verified by  $\text{pid}_i$ . If the verification succeeds,  $\text{pid}_i$  *confirms* the new epoch and updates its state accordingly, otherwise the party deletes any temporary variables it computed and remains in the same sending and receiving epochs. Our functionality models the exact same behavior.

Each instance of the functionality  $\mathcal{F}_{\text{eKE}}$  has a parties  $\text{pid}_0, \text{pid}_1$  inherited from the overall protocol, where  $\text{pid}_0$  is the party that initiates the overall conversation. Additionally, each epoch within the instance has exactly one sender and one receiver, with the parties alternating between the two roles. (i.e.  $\text{pid}_0$  is the sender for all even numbered epochs and  $\text{pid}_1$  is the sender for all odd numbered epochs.) While the epoch numbers exist internally within the functionality, the epochs are only ever externally identified by the *epoch\_id* generated by the sending party.<sup>8</sup> An epoch id is a unique adversarially chosen value (in Signal, the *epoch\_id* is the sender’s public Diffie-Hellman key). The adversarial choice of epoch id models the fact that we have no requirements from this value beyond uniquely identifying epochs.

At initialization, the epoch key exchange functionality checks that parties are registered in the global directory functionality  $\mathcal{F}_{\text{DIR}}$  and that the responding party has fresh one time keys available. It then makes a call to its long term module functionality  $\mathcal{F}_{\text{LTM}}$  to bind both parties’ directory keys to the conversation. The directory and long term module functionalities  $\mathcal{F}_{\text{DIR}}$  and  $\mathcal{F}_{\text{LTM}}$  are used by to provide the identity binding automatically assumed in  $\mathcal{F}_{\text{eKE}}$ . Finally,  $\mathcal{F}_{\text{eKE}}$  makes a call to  $\mathcal{F}_{\text{lib}}$  to obtain the adversarial code  $\mathcal{I}$  that will be run internally by the functionality to allow

<sup>8</sup>This way of identifying the epochs provides the property of revealing as little as possible about the history of the communication. The Signal protocol achieves this property by requiring that the parties store only epoch ids and never track epoch numbers at all.

adversarial choice of all epoch ids without actually relinquishing control to the adversary unless a party is compromised. ( $\mathcal{I}$  will also be used to allow adversarial choice of receiving keys when the receiver inputs an incorrect epoch id but neither party is compromised.) This mechanism of internally running adversarial code within our functionalities allows us to model the immediate decryption property of the Signal protocol. Overall, the epoch key exchange functionality  $\mathcal{F}_{\text{eKE}}$  models the forward secrecy and post-compromise security guarantees of the public key ratchet in the Signal protocol, but *without providing any epoch keys* at the level of  $\Pi_{\text{SGNL}} + \mathcal{F}_{\text{fs\_aead}} + \mathcal{F}_{\text{mKE}}$ .

In addition to providing epoch ids, the epoch key exchange functionality  $\mathcal{F}_{\text{eKE}}$  also provides keys at the level where  $\mathcal{F}_{\text{fs\_aead}}$  and  $\mathcal{F}_{\text{mKE}}$  have been instantiated by  $\Pi_{\text{fs\_aead}}$  and  $\Pi_{\text{mKE}}$  respectively. At this level, when a  $\Pi_{\text{mKE}}$  protocol instance belonging to  $\text{pid}_i$  sends a `GetReceivingKey` request and a new `epoch_id` to the functionality, a new receiving key for the epoch is produced for this party. The receiving key will only match the sender’s key if the epoch id is the correct one. This way the functionality  $\mathcal{F}_{\text{eKE}}$  enables the calling party to identify that the provided receiving epoch id is not the correct one without directly communicating it. (In fact, as a sub-component of  $\Pi_{\text{SGNL}}$ ,  $\mathcal{F}_{\text{eKE}}$  receives such a request when the party  $\text{pid}_i$  is using an instantiation  $\Pi_{\text{fs\_aead}} + \Pi_{\text{mKE}} + \mathcal{F}_{\text{aead}}$  of  $\mathcal{F}_{\text{fs\_aead}}$  to check that the ciphertext and its associated information decrypts (and authenticates) successfully under the key produced for the received epoch id.)<sup>9</sup> If the verification succeeds (either by a protocol using the provided key or by a functionality directly checking whether the epoch id’s match),  $\mathcal{F}_{\text{eKE}}$  expects a `ConfirmRcvEpoch` request directly from the party; When it receives this request, it replaces the the party’s current receiving epoch id with the one that was confirmed. It then allows the adversarially provided code to choose a brand new sending epoch id for this party with which to send any future messages. The choice of id by the adversary simply represents the fact that this value is an identifier; the epoch id is able to provide the properties needed by the functionality  $\mathcal{F}_{\text{eKE}}$  without needing to be hidden or to be chosen by the functionality. If the verification fails, the functionality simply does not receive a `ConfirmRcvEpoch` request and does not update the relevant variables.

We now describe the input output behavior of the  $\mathcal{F}_{\text{eKE}}$  in detail. (The pseudocode for this functionality can be found in Fig. 13.) The epoch key exchange functionality takes inputs from and provide output to the corresponding two machines of  $\Pi_{\text{SGNL}}$ , namely  $(\Pi_{\text{SGNL}}, \text{sid}, \text{pid}_0)$  and  $(\Pi_{\text{SGNL}}, \text{sid}, \text{pid}_1)$ , (as well as  $(\Pi_{\text{mKE}}, \text{“mKE”}, \text{sid.fs} = (\text{“fs\_aead”}, \text{sid}' = (\text{sid}, \text{pid}_0), \text{epoch\_id}))$  and  $(\Pi_{\text{mKE}}, \text{“mKE”}, \text{sid.fs} = (\text{“fs\_aead”}, \text{sid}' = (\text{sid}, \text{pid}_1), \text{epoch\_id}))$  once  $\mathcal{F}_{\text{fs\_aead}}$  and its subcomponent  $\mathcal{F}_{\text{mKE}}$  have been instantiated by their respective protocols) and the adversary  $\mathcal{A}$ . Inputs coming from other sources are ignored. Recall that at first activation,  $\mathcal{F}_{\text{eKE}}$  checks that parties are registered in  $\mathcal{F}_{\text{DIR}}$ , it checks that the responding party has fresh one time keys available, and it makes a call to  $\mathcal{F}_{\text{lib}}$  to obtain the adversarial code  $\mathcal{I}$ . The adversarial code  $\mathcal{I}$  allows us to model the immediate decryption property of secure messaging protocols when parties have not been corrupted. In this case, the choices made by the adversary are fixed based on the protocol and are not allowed to be adaptive based on the honest keys chosen during the particular run of the architecture.

The functionality has four interfaces, which we describe next.

**Confirm Receiving Epoch** After initialization, the epoch key exchange functionality expects an input of the form  $(\text{ConfirmRcvEpoch}, \text{epoch\_id}^*)$  from  $(\Pi_{\text{SGNL}}, \text{sid}, \text{pid})$  only after a message has been successfully received using `epoch_id*` in a new receiving epoch. This input triggers the functionality to update its state according to the receiving epoch id being confirmed and provide a new adversarially chosen sending epoch id to the party. On the first activation  $\mathcal{F}_{\text{eKE}}$  gets the

<sup>9</sup>Note that even at the level where  $\mathcal{F}_{\text{fs\_aead}}$  is instantiated by  $\Pi_{\text{fs\_aead}} + \mathcal{F}_{\text{mKE}} + \mathcal{F}_{\text{aead}}$ , no keys are used and therefore this request is never sent to  $\mathcal{F}_{\text{eKE}}$ .

first receiving epoch id from the global directory  $\mathcal{F}_{\text{DIR}}$ . For all other epochs, the receiving epoch id is updated to the value `epoch_id*` from the request. The state update carried out within the functionality depends on whether `epoch_id*` matches the sending epoch id of the party that `pid` is talking to.

When `epoch_id*` matches the epoch identifier of the other party, then the functionality runs the adversarially provided code  $\mathcal{I}$  to choose a new epoch id for the party’s next sending epoch. Otherwise, if `epoch_id*` is *not* the current sending epoch identifier for the other party (for example, if the adversary provided an alternative epoch identifier under which the ciphertext successfully authenticates) the functionality no longer provides any guarantees. To model this outcome, the variable `diverge_parties` is set to true in this case and the adversary  $\mathcal{A}$  is allowed to choose all epoch ids and chain keys for both parties for the rest of time. In the protocol, if the parties’ epoch identifiers diverge in this way, their symmetric keys diverge as well and cannot be restored. Note that in the Signal protocol, the adversary can carry out a person-in-the-middle attack against both parties simultaneously forever after just a single state compromise.

**Get Sending Key** As mentioned earlier, the epoch keys generated by  $\mathcal{F}_{\text{eKE}}$  (and  $\Pi_{\text{eKE}}$ ) are used by the Signal protocol to derive many message keys, one for each message within the epoch (These keys don’t exist at the level of  $\Pi_{\text{SGNL}} + \mathcal{F}_{\text{eKE}} + \mathcal{F}_{\text{fs.aead}}$ . Hence, this will interface never be called at this level.) However, at the level where  $\mathcal{F}_{\text{fs.aead}}$  and its sub-component  $\mathcal{F}_{\text{mKE}}$  have both been instantiated, the functionality  $\mathcal{F}_{\text{eKE}}$  expects `GetSendingKey` requests from the parties’  $\Pi_{\text{mKE}}$  instances. For each epoch, the functionality  $\mathcal{F}_{\text{eKE}}$  expects one `GetSendingKey` request from a  $\Pi_{\text{mKE}}$  instance that belongs to the sender `pidi` for the epoch. When no parties are compromised in the corresponding epoch, the functionality responds to this request with a uniformly chosen value `sending_chain_keyi`, which it stores in its state. Otherwise, if any party is compromised during this epoch or if divergence has occurred,  $\mathcal{F}_{\text{eKE}}$  allows the adversary to choose the new sending chain key.

**Get Receiving Key** Similarly to the `GetSendingKey` interface, this interface will also never be called at the level of  $\Pi_{\text{SGNL}} + \mathcal{F}_{\text{eKE}} + \mathcal{F}_{\text{fs.aead}}$ . This is because the epoch keys and message keys do not exist at this level of abstraction. At the level where  $\mathcal{F}_{\text{fs.aead}}$  and its sub-component  $\mathcal{F}_{\text{mKE}}$  have both been instantiated, the functionality  $\mathcal{F}_{\text{eKE}}$  expects `GetReceivingKey` requests from the parties’  $\Pi_{\text{mKE}}$  instances. For each epoch, the functionality  $\mathcal{F}_{\text{eKE}}$  expects a non-zero number of `GetReceivingKey` requests from  $\Pi_{\text{mKE}}$  instances that belong to the receiver for the epoch. When either party is compromised for the epoch or when divergence has occurred,  $\mathcal{F}_{\text{eKE}}$  allows the adversary  $\mathcal{A}$  to choose the key that it will output. Otherwise, the output key depends on whether the epoch id of the  $\Pi_{\text{mKE}}$  instance matches the senders epoch id for the epoch. When the epoch ids match,  $\mathcal{F}_{\text{eKE}}$  responds to this request with the value `sending_chain_keyi`, which was provided to the sender of the epoch – thus preserving the symmetry of the keys. When the epoch ids don’t match,  $\mathcal{F}_{\text{eKE}}$  responds with a key chosen by the latest adversarial code  $\mathcal{I}$  uploaded by the adversary to  $\mathcal{F}_{\text{lib}}$  at the time of the first activation of the  $\mathcal{F}_{\text{eKE}}$  instance. Note importantly that  $\mathcal{I}$  will not have access to the key provided to the sender of the epoch. The functionality  $\mathcal{F}_{\text{eKE}}$  leaves it to the instantiation of  $\mathcal{F}_{\text{fs.aead}}$  to use the key chosen by  $\mathcal{I}$  to figure out that  $\Pi_{\text{mKE}}$ ’s epoch id should not be confirmed as the new receiving epoch identifier. This models the fact that in the Signal protocol, naive tampering with the sender’s public exponent will cause an unsuccessful temporary ratchet upon a failure of authentication (and later reversion to the previous epoch). In traditional instantiations of the Signal architecture, if  $\mathcal{A}$  corrupts the receiver after naive tampering attempts, the adversary can use the party’s stored root key to compute the `recv_chain_key`’s from these tampering events after

the fact, therefore running into the Nielsen bound and being unable to realise  $\mathcal{F}_{eKE}$  in the plain model. Looking ahead, we show that an instantiation of the Signal architecture with a KDF built from laconic primitives can circumvent this problem. It can realise  $\mathcal{F}_{eKE}$  in the plain model by removing the ability to compute receiving chain keys corresponding to epoch ids that resulted in failed ratcheting attempts.

**Corrupt** On receiving a **Corrupt** notification from a party’s  $\Pi_{\text{SGNL}}$  instance,  $\mathcal{F}_{eKE}$  sends a **ReportState** request to the adversary along with the  $(\text{epoch\_id}, \text{recv\_chain\_key}')$  pairs corresponding to **GetReceivingKey** calls made so far. Additionally, at the time of making a **ReportState** call to the adversary  $\mathcal{A}$ ,  $\mathcal{F}_{eKE}$  provides  $\mathcal{A}$  with the corrupted party’s newest receiving key if it hasn’t been successfully retrieved yet. In response, the adversary sends back some state  $S$  which is forwarded by  $\mathcal{F}_{eKE}$  to the  $\Pi_{\text{SGNL}}$  instance that made the corruption request. Before responding to the calling instance of  $\Pi_{\text{SGNL}}$ ,  $\mathcal{F}_{eKE}$  notes that the current epoch is corrupted, and adds the next 3 epochs to a list of compromised epochs until full post-compromise recovery may be achieved. This is because the recovery for a single compromise goes through the following phases:

1. **full state compromise (the party’s entire state is known during epochs  $e, e + 1$ .)** If a party is corrupted during sending epoch  $e$  then they will be a receiver in epoch  $e + 1$ . Since the party does not add any secret randomness to its state in epoch  $e + 1$ , the adversary still knows the party’s entire state at this point.
2. **sender’s randomness is updated (epoch  $e + 2$ )** In epoch  $e + 2$ , the corrupted party will update its randomness and send a value  $\text{epoch\_id}_{e+2}$  to their peer. Once their peer confirms the new epoch id, the adversary will be unable to tamper with the communication in a significant way; however, full deniability has not been restored.
3. **both parties’ randomness is updated (epoch  $e + 3$ )** In epoch  $e + 3$  the peer will update its randomness and send a value  $\text{epoch\_id}_{e+3}$  to the corrupted party. Once the corrupted party confirms the new epoch id, full deniability will be restored. However, since such a confirmation is evidenced by the start of epoch  $e + 4$ , this is the first uncompromised epoch after the corruption.

The compromise period may be extended further if a party receives a packet with a bogus epoch id that corresponds to some compromised epoch  $e^*$ . In this case, one of two things can happen:

- **If the party does not confirm the bogus epoch id, the compromise extends until epoch  $e^* + 2$ .** This outcome corresponds to a weakness in the standalone security of  $\Pi_{eKE}$ , as it does not include any authentication of keys.<sup>10</sup>
- **If the party *does* confirm the bogus epoch id, the functionality behaves as if both parties are compromised for the rest of the session.** The confirmation of a bogus epoch id corresponds to the divergence of the parties’ states and keys. In this scenario,  $\mathcal{F}_{eKE}$  notes that the parties have *diverged*, and the functionality then behaves as if both parties are compromised for the rest of the session.

---

<sup>10</sup>This provision was missing in the [25] version of this work.



### Remark

We remark here that one may easily replace the 3 epoch compromise here and in  $\mathcal{F}_{\text{SM}}$  with a variable for the time before healing. This would be especially useful since a Signal-like secure messaging protocol will inherit its healing time from the epoch key exchange that it uses. Using our modular analysis, one can easily swap out one epoch key exchange protocol that heals in  $r$  rounds with another that heals in a different number of rounds  $r'$  and show that the resulting  $\Pi_{\text{SGNL}} + \Pi_{\text{fs\_aead}} + \Pi'_{\text{eKE}}$  system will realise an  $\mathcal{F}_{\text{SM}}$  that enforces healing in a number of rounds that is greater than or equal to  $r'$ .

A simple modification to  $\Pi_{\text{eKE}}$  in fact allows healing in 2 rounds instead of 3. This modification is discussed in the remark on Page 51

## 5.2 The Forward Secure Encryption Functionality $\mathcal{F}_{\text{fs\_aead}}$

The forward secure authenticated encryption functionality (see Figure 15 on page 41) processes encryptions and decryptions for a single epoch (specified in  $\text{sid.fs}$ ) for the protocol  $\Pi_{\text{SGNL}}$ . As the name suggests,  $\mathcal{F}_{\text{fs\_aead}}$  enforces the forward security property for messages encrypted within an epoch. That is, on a state compromise of the receiver for the epoch, the adversary only gets  $(c, m)$  pairs for messages that are *in transit* from the sender to receiver, and it gets the power to replace ciphertexts in transit with authentic-looking ones. Note that once an epoch has been compromised, there is no recovery within the epoch; that is, the adversary retains the power to tamper with in-transit ciphertexts from the epoch until all have arrived at the receiver. Any ciphertexts that the receiver decrypted prior to state compromise are not available to  $\mathcal{A}$ ; this models the forward security property of Signal's symmetric chain.

**Encryption** On receiving an encryption request for a message  $m$ , it sends  $N$  (the number of messages sent in the previous sending epoch) and leaks either  $|m|$  or  $m$  to  $\mathcal{A}$  (depending on whether the epoch is compromised) and gets a ciphertext  $c$  in return, which it records along with  $m$ ,  $\text{msg\_num}$ ,  $N$ , and the leakage. Note that in the real protocol, the  $\text{msg\_num}$ ,  $\text{epoch\_id}$ , as well as  $N$  are authenticated but sent in the clear with each ciphertext.  $\mathcal{F}_{\text{fs\_aead}}$  then sends the ciphertext up to the protocol  $\Pi_{\text{SGNL}}$ .

**Decryption** When receiving a decrypt request for ciphertext  $c$  and message number  $\text{msg\_num}$ ,  $\mathcal{F}_{\text{fs\_aead}}$  checks whether the receiver has already successfully decrypted this  $\text{msg\_num}$ ; if so, the  $\text{msg\_num}$  was set to inaccessible and  $\mathcal{F}_{\text{fs\_aead}}$  will return a failure message to  $\Pi_{\text{SGNL}}$ . Next,  $\mathcal{F}_{\text{fs\_aead}}$  checks whether the ciphertext  $c$  previously failed authentication for  $\text{msg\_num}$ ; in this case, the functionality also outputs a failure message to  $\Pi_{\text{SGNL}}$ . If the decryption has not failed from the previous two cases, the functionality sends an **Inject** message to  $\mathcal{A}$ .

If the state of  $\mathcal{F}_{\text{fs\_aead}}$  is not compromised, then  $\mathcal{A}$  should only be able to **Inject** the true message  $m$  that was encrypted for  $\text{msg\_num}$ . In the honest setting (no state corruption), if the adversary returns  $\perp$  or there is no record of an encryption for  $\text{msg\_num}$ ,  $\mathcal{F}_{\text{fs\_aead}}$  returns a failure; otherwise, regardless of which message  $v$  the adversary returns,  $\mathcal{F}_{\text{fs\_aead}}$  sends  $m$  to  $\Pi_{\text{SGNL}}$ . This models the fact that without compromising a party, the real world adversary should not be able to produce ciphertexts that authenticate.

In the case that a state compromise has occurred, if  $\mathcal{A}$  returns some  $v \neq \perp$ ,  $\mathcal{F}_{\text{fs\_aead}}$  marks  $\text{msg\_num}$  as unavailable and sends  $v$  up to  $\Pi_{\text{SGNL}}$ . This models the power that the adversary has after a state compromise (of either party) to tamper with the sender's ciphertexts to produce authentic-looking ciphertexts. Note that  $\mathcal{F}_{\text{fs\_aead}}$  never recovers from a state compromise; thus, the adversary maintains the power to tamper with the ciphertexts for the epoch as long as there are messages from the epoch in transit.

## $\mathcal{F}_{eKE}$

This functionality has a session id  $\text{sid.eKE}$  that takes the following format:  $\text{sid.eKE} = (\text{"eKE"}, \text{sid})$ . Inputs arriving from machines whose identity is neither  $\text{pid}_0$  nor  $\text{pid}_1$  are ignored. //For notational simplicity we assume some fixed interpretation of  $\text{pid}_0$  and  $\text{pid}_1$  as complete identities of the two calling machines.

//The following method is also used by the sender of epoch 0 to start the conversation.

**ConfirmReceivingEpoch:** On input  $(\text{ConfirmReceivingEpoch}, \text{epoch\_id}^*)$  from  $(\Pi_{\text{SGNL}}, \text{sid}, \text{pid}_i)$ :

1. **If** this is the first activation, initialize conversation using  $\mathcal{F}_{\text{DIR}}, \mathcal{F}_{\text{LTM}}$  and get internal code from  $\mathcal{F}_{\text{lib}}$ :
  - (a) Parse  $\text{sid}$  to retrieve two party ids  $(\text{pid}_0, \text{pid}_1)$  for the initiator and responder parties and store them. If  $\text{pid}_0 \neq \text{pid}_i$  end the activation.
  - (b) Provide input  $(\text{GetInitKeys}, \text{pid}_1, \text{pid}_0)$  to  $(\mathcal{F}_{\text{DIR}})$ .
  - (c) Upon receiving output  $(\text{GetInitKeys}, \text{ik}_1^{\text{pk}}, \text{rk}_1^{\text{pk}}, \text{ok}_1^{\text{pk}})$  from  $(\mathcal{F}_{\text{DIR}})$ :
    - i. **If**  $\text{ok}_{\text{pid}_1}^{\text{pk}} = \perp$  then output  $(\text{ConfirmReceivingEpoch}, \text{Fail})$ .  
//Don't start the conversation if the one time keys belonging to the other party have run out.
    - ii. Initialize empty lists  $\text{corruptions}_0, \text{corruptions}_1, \text{compromised\_epochs}$ .
    - iii. Set  $\text{epoch\_id\_partner}_0 = \text{epoch\_id\_self}_1 = \text{ok}_{\text{pid}_1}^{\text{pk}}$ ,  $\text{epoch\_num}_0 = -2$ , and  $\text{epoch\_num}_1 = -1$ .
    - iv. Send  $(\text{ComputeSendingRootKey}, \text{ik}_1^{\text{pk}}, \text{rk}_1^{\text{pk}}, \text{ok}_1^{\text{pk}})$  to  $\mathcal{F}_{\text{LTM}}$ .
  - (d) On receiving  $(\text{ComputeSendingRootKey}, k, \text{ek}^{\text{pk}})$  from  $\mathcal{F}_{\text{LTM}}$ , continue.
  - (e) Initialize  $\text{state}_{\mathcal{I}} = \perp$ , call  $\mathcal{F}_{\text{lib}}$  to obtain internal adversarial code  $\mathcal{I}_{eKE}$ .
2. **Else** (this is not first activation):
  - (a) Set  $\text{epoch\_id\_partner}_i = \text{epoch\_id}^*$ .
  - (b) **If**  $\text{epoch\_id}^* \neq \text{epoch\_id\_self}_{1-i}$ : Set  $\text{diverge\_parties} = \text{true}$ , run step 3 of **Corrupt** to set  $\text{recv\_chain\_key}^*$  and leakage, and send  $(\text{ReportState}, \text{state}_{\mathcal{I}}, i, \text{recv\_chain\_key}^*, \text{leakage})$  to the adversary. On receiving a response, continue. //diverge\_parties is set here  
//If this may be the first divergence, make sure the simulator has control.
3. **If**  $\text{epoch\_num}_i + 2 \in \text{compromised\_epochs}$  or  $\text{diverge\_parties}$ : Send backdoor message  $(\text{GenEpochId}, i, \text{epoch\_id}^*)$  to the adversary. **Else** run  $\mathcal{I}(\text{state}_{\mathcal{I}}, \text{GenEpochId}, i, \text{epoch\_id}^*)$ .  
//If the parties are diverged or compromised send a backdoor message to the adversary, otherwise run the internal adversarial code.
4. Upon receiving  $(\text{GenEpochId}, \text{state}'_{\mathcal{I}}, i, \text{epoch\_id})$  from  $\mathcal{A}$  or from  $\mathcal{I}$ , update  $\text{state}_{\mathcal{I}} \leftarrow \text{state}'_{\mathcal{I}}$  and do the following:
  - (a) **If**  $\text{epoch\_id}$  is the same as any previous invocation of **GenEpochId**, end the activation.
  - (b) Set  $\text{epoch\_id\_self}_i = \text{epoch\_id}$ ,  $\text{epoch\_num\_dict}[\text{epoch\_id}] = \text{epoch\_num}_i$ ,  $\text{got\_sending\_key}_i = \text{false}$ , and  $\text{epoch\_num}_i += 2$ . //save the next sending epoch id.
  - (c) Output  $(\text{ConfirmReceivingEpoch}, \text{epoch\_id\_self}_i)$  to  $(\Pi_{\text{SGNL}}, \text{sid}, \text{pid}_i)$ .

**GetSendingKey:** On receiving input  $(\text{GetSendingKey})$  from  $(\Pi_{mKE}, \text{sid.mKE}, \text{pid})$ :

1. Parse  $\text{sid.mKE} = (\text{"mKE"}, \text{"fs\_aead"}, \text{sid}, \text{epoch\_id})$ . **If**  $\text{epoch\_id} \neq \text{epoch\_id\_self}_i$  end the activation; **else** set  $i$  such that  $\text{pid} = \text{pid}_i$ .
2. **If**  $\text{got\_sending\_key}_i = \text{true}$  or **ConfirmReceivingEpoch** has never been run successfully then end the activation.  
//the functionality isn't initialized or the sending key for the current epoch has already been retrieved
3. Sample  $\text{sending\_chain\_key}_i \xleftarrow{\$} \mathcal{K}_{ep}$  from the key distribution. //In the honest case, the key is not known to the adversary. Otherwise the key will get overwritten in the following step.
4. **If**  $\text{diverge\_parties} = \text{true}$ , or  $\text{epoch\_num}_i \in \text{compromised\_epochs}$  send backdoor message  $(\text{GetSendingKey}, i)$  to the adversary; on receiving backdoor message  $(\text{GetSendingKey}, i, K_{\text{send}})$  from  $\mathcal{A}$  set  $\text{sending\_chain\_key}_i = K_{\text{send}}$ .  
//If the parties are diverged or compromised let the adversary choose  $\text{sending\_chain\_key}_i$ .
5. Set  $\text{got\_sending\_key}_i = \text{true}$  and output  $(\text{GetSendingKey}, \text{sending\_chain\_key}_i)$ .

(The rest of this functionality is in Fig. 14 on Page 39)

Figure 13: The Epoch Key Exchange Functionality,  $\mathcal{F}_{eKE}$

### $\mathcal{F}_{eKE}$ continued...

(This functionality begins in Fig. 13 on Page 38)

**GetReceivingKey:** On receiving input  $(\text{GetReceivingKey}, \text{epoch\_id})$  from  $(\Pi_{mKE}, \text{sid}, mKE, \text{pid})$ :

1. **If**  $\text{pid} \notin \{\text{pid}_0, \text{pid}_1\}$  then end this activation. **Else**, set  $i$  such that  $\text{pid} = \text{pid}_i$ .
2. **If**  $\text{sending\_chain\_key}_{1-i}$  has been deleted or **ConfirmReceivingEpoch** has never been run successfully end the activation.
3. Parse  $\text{epoch\_id} = (\text{epoch\_id}', \text{ek}_j^{\text{pk}}, \text{ok}_{i \leftarrow j}^{\text{pk}})$  and  $\text{sid}, mKE = ("mKE", "fs\_aad", \text{sid}, \text{epoch\_id}')$ . **If**  $\text{epoch\_id}' \neq \text{epoch\_id}$  end the activation.
4. **If** this is the first call to **GetReceivingKey**:
  - (a) **If**  $\text{pid}_1 \neq \text{pid}_i$ , then end the activation.
  - (b) Send  $(\text{GetResponseKeys}, \text{pid}_{1-i})$  to  $\mathcal{F}_{\text{DIR}}$ .
  - (c) Upon receiving  $(\text{GetResponseKeys}, \text{ik}_j^{\text{pk}})$ , send input  $(\text{ComputeReceivingRootKey}, \text{ik}_j^{\text{pk}}, \text{ek}_j^{\text{pk}}, \text{ok}_{i \leftarrow j}^{\text{pk}})$  to  $\mathcal{F}_{\text{LTM}}$ .
5. **If**  $\text{diverge\_parties} = \text{true}$  or  $\text{epoch\_num}_i + 1 \in \text{compromised\_epochs}$ : //Let  $\mathcal{A}$  choose key
  - (a) **If**  $\text{epoch\_id} \neq \text{epoch\_id\_self}_{1-i}$  add  $\text{epoch\_num}_i + 2, \text{epoch\_num}_i + 3$  to  $\text{compromised\_epochs}$ .<sup>a</sup>
  - (b) Send  $(\text{GetReceivingKey}, i, \text{epoch\_id})$  to the adversary.
  - (c) Upon receiving  $(\text{GetReceivingKey}, i, \text{epoch\_id}, \text{recv\_chain\_key}^*)$  from  $\mathcal{A}$ , output  $(\text{GetReceivingKey}, \text{recv\_chain\_key}^*)$ .
6. **Else**, if  $\text{epoch\_id} \neq \text{epoch\_id\_self}_{1-i}$ : //No corruptions or divergence but epoch\_id doesn't match epoch\_id\_{1-i}.
  - (a) Sample  $\text{recv\_chain\_key}_i \xleftarrow{\$} \mathcal{K}_{ep}$ .
  - (b) Add  $(\text{epoch\_id}, \text{recv\_chain\_key})$  to  $\text{receive\_attempts}[\text{epoch\_num}]$ .
  - (c) Output  $(\text{GetReceivingKey}, \text{recv\_chain\_key})$ .
7. **Else**, output  $(\text{GetReceivingKey}, \text{sending\_chain\_key}_{1-i})$ . //Expected case.

**Corrupt:** On receiving a **(Corrupt)** request from  $(\Pi_{\text{SGNL}}, \text{sid}, \text{pid}_i)$  for  $i \in \{0, 1\}$  do:

1. Add  $\text{epoch\_id\_self}_i$  to the list  $\text{corruptions}_i$ .
2. Add  $\text{epoch\_num}_i, \text{epoch\_num}_i + 1, \dots, \text{epoch\_num}_i + 3$  to the list  $\text{compromised\_epochs}$ .  
//Compromise goes through the following stages: fully compromised for 2-3 epochs, sender randomness updated, both parties' randomness updated.  
// $\text{epoch\_num}_i + 5$  is protected by forward secrecy even in the case of re-corruption.
3. Let  $j$  be such that  $\text{epoch\_num}_j > \text{epoch\_num}_{1-j}$  and:
  - (a) **If**  $\text{got\_sending\_key}_j = \text{false}$  set  $\text{recv\_chain\_key}^* = \perp$ . **Else** set  $\text{recv\_chain\_key}^* = \text{sending\_chain\_key}_j$ .
  - (b) **If**  $\text{epoch\_num}_j \in \text{receive\_attempts.keys}$  then set  $\text{leakage} = \text{receive\_attempts}[\text{epoch\_num}_j]$ . **Else** set  $\text{leakage} = []$ .
4. Send  $(\text{ReportState}, \text{state}_{\mathcal{I}}, i, \text{recv\_chain\_key}^*, \text{leakage})$  to the adversary.
5. Upon receiving  $(\text{ReportState}, i, S)$  from  $\mathcal{A}$ , output  $(\text{Corrupt}, S)$  to  $(\Pi_{\text{SGNL}}, \text{sid}, \text{pid}_i)$ .

---

<sup>a</sup>This provision (which was missing in the [25] version) allows the environment to delay recovery from corruption by delivering messages with bogus epoch ids, this occurs even if these messages do not contain a payload that passes authentication.

Figure 14: The Epoch Key Exchange Functionality,  $\mathcal{F}_{eKE}$  (continued)

In all of the above cases, if a decryption was successful,  $\mathcal{F}_{\text{fs\_aead}}$  marks that `msg_num` as unavailable for future decryption attempts. This models the forward security property of the symmetric ratchet in the Signal protocol. Note that  $\mathcal{F}_{\text{fs\_aead}}$  doesn't mark messages as inaccessible upon unsuccessful decryption. We chose this to prevent denial-of-service attacks that would prevent honest parties from decrypting messages sent to them.

**Stop Encrypting** When receiving a `StopEncrypting` request from the sender for the epoch,  $\mathcal{F}_{\text{fs\_aead}}$  notes this and blocks all future encryptions for the epoch. This prevents an adversary from compromising the sender and injecting additional messages after the sender has moved to a new epoch.

**Stop Decrypting** When receiving a `(StopDecrypting, msg_num*)` request from the receiver,  $\mathcal{F}_{\text{fs\_aead}}$  marks all message numbers larger than `msg_num*` as inaccessible, thereby preventing their decryption. This prevents an adversary from compromising the receiver of the epoch and injecting additional messages in the epoch after the receiver has advanced to a new epoch.

**Corruption** On receiving a state compromise notification from above (specifically,  $\Pi_{\text{SGNL}}$ ), if the receiver is the corrupted party,  $\mathcal{F}_{\text{fs\_aead}}$  leaks to the adversary all message/ciphertext pairs that are *in transit* from the sender to receiver. This models the fact that in the Signal protocol, the receiver has keys for out-of-order messages stored in its state until they have all arrived. On the other hand, if the sender is corrupted, no sent messages are leaked to the adversary, since in the Signal protocol the sender does not store any keys for message it has already sent. (Either way, the adversary will be able to read all messages sent after corruption.) The adversary (or simulator) then returns a constructed state for the party: the message keys for messages in transit, along with the chain key if it has not been deleted. This state is passed back up to  $\Pi_{\text{SGNL}}$ . Since this functionality corresponds to a single epoch, there is no provision for recovery of security properties after a corruption.

### 5.3 The Signal Protocol, $\Pi_{\text{SGNL}}$

Protocol  $\Pi_{\text{SGNL}}$  (see Figure 16 on page 43) is the top-level protocol that takes input commands from a party, communicates with  $\mathcal{F}_{\text{eKE}}$  and  $\mathcal{F}_{\text{fs\_aead}}$ , and returns outputs to the party to coordinate the encryption and decryption of messages for the duration of the conversation session between two parties. The ciphertexts are transferred between parties via the environment, which has full control over the network. Next we describe the three interfaces to  $\Pi_{\text{SGNL}}$  (which have identical API's as  $\mathcal{F}_{\text{SM}}$ ).

**Send Message** When receiving a `SendMessage` request from party `pid`,  $\Pi_{\text{SGNL}}$  first initializes  $\mathcal{F}_{\text{eKE}}$  if necessary to get the first `epoch_id`, and sends an `Encrypt` request with message  $m$  and  $N_{\text{last}}$  (the number of messages that were sent in the party's previous sending epoch) to the  $\mathcal{F}_{\text{fs\_aead}}$  instance for the current epoch. On receiving a ciphertext  $c$  from  $\mathcal{F}_{\text{fs\_aead}}$ , the protocol  $\Pi_{\text{SGNL}}$  deletes  $m$ , increments the number of messages sent, and outputs  $c$  along with a header  $h = (\text{epoch\_id}_{\text{self}}, \text{sent\_msg\_num}, N_{\text{last}})$  to `pid`.

**Receive Message** On receiving a `ReceiveMessage` command with ciphertext  $c$  and header  $h = (\text{epoch\_id}, \text{msg\_num}, N)$  from a party with id `pid`, the protocol  $\Pi_{\text{SGNL}}$  first initializes the

### $\mathcal{F}_{fs\_aead}$

This functionality processes encryptions and decryptions for a *single* epoch and has session id  $sid.fs$ , where  $sid.fs = ("fs\_aead", sid = (sid', (pid_0, pid_1)), epoch\_id)$ . Inputs arriving from machines other than  $pid_0, pid_1$  are ignored. //For notational simplicity we assume a fixed interpretation of  $pid_0$  and  $pid_1$  as complete identities of the two calling machines.

**Encrypt:** On receiving input  $(\text{Encrypt}, m, N)$  from  $(\Pi_{SGNL}, sid, pid)$  do:

1. If this is the first activation then:
  - Let  $i$  be such that  $pid = pid_i$ . Initialize  $msg\_num = 0$  and sender  $sender = i$ .
  - Initialize  $state_{\mathcal{I}} = \perp$ , call  $\mathcal{F}_{lib}$  to obtain the internal adversarial code  $\mathcal{I}$ .
2. Verify that  $sid$  matches the one in the local state and  $pid = pid_b$ , otherwise end the activation.
3. If the sender has deleted the ability to encrypt messages, then end the activation.
4. Increment  $msg\_num = msg\_num + 1$ .
5. If  $IsCorrupt? = false$ : Run  $\mathcal{I}(state_{\mathcal{I}}, \text{Encrypt}, pid, N, |m|)$ .
6. If  $IsCorrupt? = true$ : Send a backdoor message  $(state_{\mathcal{I}}, \text{Encrypt}, pid, msg\_num, N, m)$  to  $\mathcal{A}$ .
7. On obtaining the output  $(state_{\mathcal{I}}, \text{Encrypt}, pid, c, msg\_num, N)$  from  $\mathcal{I}$  or  $\mathcal{A}$ , update  $state_{\mathcal{I}}$  and record  $(m, c, msg\_num, N)$ .
8. Output  $(\text{Encrypt}, c)$  to  $(\Pi_{SGNL}, sid, pid)$ .

**Decrypt:** On receiving  $(\text{Decrypt}, c, msg\_num, N)$  from  $(\Pi_{SGNL}, sid, pid)$  do:

1. Verify that  $sid$  matches the one in the local state and  $pid = pid_{1-sender}$ , otherwise end the activation.  
//end the activation if the decrypt request is not from the receiving party
2. If  $msg\_num$  is set as inaccessible, or there is a record  $(\text{Authenticate}, c, msg\_num, N, 0)$ , then output  $(\text{Decrypt}, c, msg\_num, N, \text{Fail})$  to  $(\Pi_{SGNL}, sid, pid)$ .
3. If  $IsCorrupt? = false$ :
  - Run  $\mathcal{I}(state_{\mathcal{I}}, \text{Authenticate}, pid, c, msg\_num, N)$  and obtain updated state  $state_{\mathcal{I}}$  and output  $(\text{Authenticate}, pid, c, msg\_num, N, v)$ .
  - If  $v = \perp$ , then record  $(\text{Authenticate}, c, msg\_num, N, 0)$  and output  $(\text{Decrypt}, c, msg\_num, N, \text{Fail})$  to  $(\Pi_{SGNL}, sid, pid)$ .
  - Otherwise, mark  $msg\_num$  as inaccessible and output  $(\text{Decrypt}, c, msg\_num, N, m)$  to  $(\Pi_{SGNL}, sid, pid)$ .
4. Else ( $IsCorrupt? = true$ ):
  - Send  $(state_{\mathcal{I}}, \text{Inject}, pid, c, msg\_num, N)$  to  $\mathcal{A}$ .
  - On receiving  $(state_{\mathcal{I}}, \text{Inject}, pid, c, msg\_num, N, v)$  from  $\mathcal{A}$ , update  $state_{\mathcal{I}}$  and do:
    - If  $v = \perp$ , record  $(\text{Authenticate}, c, msg\_num, N, 0)$  and output  $(\text{Decrypt}, c, msg\_num, N, \text{Fail})$ .
    - Else, then mark  $msg\_num$  as inaccessible and output  $(\text{Decrypt}, c, msg\_num, N, v)$  to  $(\Pi_{SGNL}, sid, pid)$ .

**StopEncrypting:** On receiving  $(\text{StopEncrypting})$  from  $(\Pi_{SGNL}, sid, pid)$  do:

1. If  $sid$  doesn't match the one in the local state, if  $pid \neq pid_{sender}$ , or if this is the first activation: end the activation.
2. Otherwise, note that  $pid_i$  has deleted the ability to encrypt future messages. Output  $(\text{StopEncrypting}, \text{Success})$ .

**StopDecrypting:** On receiving  $(\text{StopDecrypting}, msg\_num^*)$  from  $(\Pi_{SGNL}, sid, pid)$  do:

1. If  $sid$  doesn't match the one in the local state,  $pid \neq pid_{1-sender}$ , or no messages have been successfully decrypted by  $pid_i$ : end the activation.
2. Mark all  $msg\_num > msg\_num^*$  as inaccessible, and output  $(\text{StopDecrypting}, \text{Success})$  to  $(\Pi_{SGNL}, sid, pid)$ .

**Corrupt:** On receiving  $(\text{Corrupt}, pid)$  from  $(\Pi_{SGNL}, sid, pid)$ :

1. Record  $(\text{Corrupt}, pid)$  and set  $IsCorrupt? = true$ .
2. If  $pid = pid_{1-sender}$  ( $pid$  is the receiver), let  $leakage = \{(pid_{sender}, h = (epoch\_id, msg\_num, N), c, m)\}$  be the set of all messages sent by  $pid_{sender}$  which are not marked as inaccessible.
3. Otherwise ( $pid$  is the sender), set  $leakage = \emptyset$ .
4. Send  $(\text{ReportState}, state_{\mathcal{I}}, pid, leakage)$  to  $\mathcal{A}$ .
5. Upon receiving a response  $(\text{ReportState}, state_{\mathcal{I}}, pid, S)$  from  $\mathcal{A}$ , send  $S$  to  $(\Pi_{SGNL}, sid, pid)$ .

receiver’s state if necessary. It then sends a `Decrypt` request for ciphertext  $c$  to the  $\mathcal{F}_{\text{fs\_aead}}$  instance corresponding to the value `epoch_id` in header  $h$ . On receiving a response from  $\mathcal{F}_{\text{fs\_aead}}$ , if decryption failed,  $\Pi_{\text{SGNL}}$  outputs a failure message. Otherwise, it updates the list of `msg_num`’s that were skipped in the epoch corresponding to the value `epoch_id`. If the value `epoch_id` is new (i.e. no messages have been received by this party under this value before) then  $\Pi_{\text{SGNL}}$  closes its last sending and receiving epochs (by sending a `StopDecrypting` request to the  $\mathcal{F}_{\text{fs\_aead}}$  instance corresponding to `epoch_id`<sub>partner</sub>, and a `StopEncrypting` request to the  $\mathcal{F}_{\text{fs\_aead}}$  instance corresponding to `epoch_id`<sub>self</sub>.) The protocol  $\Pi_{\text{SGNL}}$  then updates `epoch_id`<sub>partner</sub> = `epoch_id` and sends (`ConfirmReceivingEpoch, epoch_id`) to  $\mathcal{F}_{\text{eKE}}$  to ratchet forward and receive a new `epoch_id`\* for its next sending epoch. Finally,  $\Pi_{\text{SGNL}}$  deletes the decrypted message  $v$  returned by  $\mathcal{F}_{\text{fs\_aead}}$  and outputs the ciphertext  $c$ , message  $v$ , and header  $h$  to `pid`.

**Corruption** The protocol  $\Pi_{\text{SGNL}}$  has one additional interface, a `Corrupt` interface that is accessible only to `Env`. This interface is not part of the real protocol, but is included only for UC-modelling purposes. On a corruption from the environment,  $\Pi_{\text{SGNL}}$  sends `Corrupt` notifications to  $\mathcal{F}_{\text{eKE}}$  and to every  $\mathcal{F}_{\text{fs\_aead}}$  instance that has messages in transit. These sub-functionalities report their internal states to  $\Pi_{\text{SGNL}}$  who forwards the union of their states up to `Env`.

## 5.4 Security Analysis of $\Pi_{\text{SGNL}}$

In this section we prove Theorem 2 which says that  $\Pi_{\text{SGNL}}$ ,  $\mathcal{F}_{\text{eKE}}$ , and  $\mathcal{F}_{\text{fs\_aead}}$  together UC-realize  $\mathcal{F}_{\text{SM}}$  in the presence of global functionalities  $\mathcal{F}_{\text{lib}}$ ,  $\mathcal{F}_{\text{LTM}}$ , and  $\mathcal{F}_{\text{DIR}}$ .

As a reminder, the claim that “ $A$  UC-realizes  $B$  in the presence of  $C$ ” means that the environment’s views are indistinguishable when interacting with  $A$  or  $B$ , together with their respective adversaries and a global subroutine  $C$ . We refer readers to Section 2 for a more detailed primer on the universally composable security framework.

**Theorem 2** *Protocol  $\Pi_{\text{SGNL}}$  (perfectly) UC-realizes the ideal functionality  $\mathcal{F}_{\text{SM}}$  in the presence of  $\mathcal{F}_{\text{lib}}$ ,  $\mathcal{F}_{\text{DIR}}$  and  $\mathcal{F}_{\text{LTM}}$ .*

**Proof:** To prove Theorem 2, we first construct the simulator  $\mathcal{S}_{\text{SM}}$  and the internal code  $\mathcal{I}_{\text{SM}}$ . Then, we argue that the environment `Env` has an identically distributed view in its interaction with  $\Pi_{\text{SGNL}} + \mathcal{F}_{\text{fs\_aead}} + \mathcal{F}_{\text{eKE}}$  as it does in its interaction with  $\mathcal{F}_{\text{SM}} + \mathcal{S}_{\text{SM}} + \mathcal{I}_{\text{SM}}$ .

In an interaction between `Env` and  $\mathcal{F}_{\text{SM}}$ , the simulator  $\mathcal{S}_{\text{SM}}$  and internal adversarial code  $\mathcal{I}_{\text{SM}}$  are provided with only the information that  $\mathcal{F}_{\text{SM}}$  gives to its ideal process adversary and internal adversarial code. When the session is not compromised, the functionality never calls the simulator  $\mathcal{S}_{\text{SM}}$ . Instead, the only adversarial choices are made when the functionality runs the internal code  $\mathcal{I}_{\text{SM}}$ . The objective of  $\mathcal{S}_{\text{SM}}$  and  $\mathcal{I}_{\text{SM}}$  is to respond in such a way that simulates the artifacts that would be generated if `Env` were interacting with  $\Pi_{\text{SGNL}}$ . The detailed versions of  $\mathcal{I}_{\text{SM}}$  and  $\mathcal{S}_{\text{SM}}$  can be found in Figure 19 (Page 46) and Figure 17 (Page 44) respectively.

Before arguing that the adversary’s view is identical in the two scenarios, we describe where  $\mathcal{S}_{\text{SM}}$  and  $\mathcal{I}_{\text{SM}}$  are called by  $\mathcal{F}_{\text{SM}}$  within each of its methods:

- Within `SendMessage`, as long as initialization has been properly performed then  $\mathcal{F}_{\text{SM}}$  will generate tuple (`SendMessage, pid,  $\ell$` ) and either run the code  $\mathcal{I}_{\text{SM}}$  on it or send it to  $\mathcal{S}_{\text{SM}}$  based on whether the parties are compromised, diverged, or neither. When the parties are neither compromised nor diverged,  $\ell = |m|$  and the code  $\mathcal{I}_{\text{SM}}$  is run on (`SendMessage, pid,  $\ell$` ). Otherwise,  $\ell = m$  and the tuple is sent to  $\mathcal{S}_{\text{SM}}$ .



## $\Pi$ SGNL

**SendMessage:** Upon receiving input (`SendMessage`,  $m$ ) from  $\text{pid}$ , do:

1. If this is the first activation do: //initialization for the initiator of the session
  - Parse the local session id  $\text{sid}$  to retrieve the party identifiers ( $\text{pid}_0, \text{pid}_1$ ) for the initiator and responder. If  $\text{pid}_0$  is different from either the local party identifier  $\text{pid}$ , or the party identifier of  $\text{pid}$ , end the activation.
  - Initialize  $\text{epoch\_id\_self} = \perp$ ,  $\text{epoch\_id\_partner} = \perp$ ,  $\text{sent\_msg\_num} = 0$ ,  $N_{\text{last}} = 0$ .
  - Provide input (`ConfirmReceivingEpoch`,  $\perp$ ) to  $(\mathcal{F}_{\text{eKE}}, \text{sid.eKE})$ .
  - On receiving (`ConfirmReceivingEpoch`,  $\text{epoch\_id}$ ) from  $(\mathcal{F}_{\text{eKE}}, \text{sid.eKE})$ , set  $\text{epoch\_id\_self} = \text{epoch\_id}$ .
  - Initialize a list  $\text{receiving\_epochs} = []$ .
2. Provide input (`Encrypt`,  $m, N_{\text{last}}$ ) to  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid.fs})$ , where  $\text{sid.fs} = (\text{sid}, \text{epoch\_id\_self})$ . // $\mathcal{F}_{\text{fs\_aead}}$  already knows  $\text{epoch\_id}$  and  $\text{msg\_num}$
3. On receiving (`Encrypt`,  $c, N_{\text{last}}$ ) from  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid.fs})$ , delete  $m$ , increment  $\text{sent\_msg\_num} += 1$ , output (`SendMessage`,  $\text{sid}, h, c$ ) to  $\text{pid}$ , where  $h = (\text{epoch\_id\_self}, \text{sent\_msg\_num}, N_{\text{last}})$ .

**ReceiveMessage:** Upon receiving (`ReceiveMessage`,  $h = (\text{epoch\_id}, \text{msg\_num}, N), c$ ) from  $\text{pid}$ :

1. If this is the first activation then do: //initialization for the responder of the session
  - Parse the local session identifier  $\text{sid}$  to retrieve the party identifiers ( $\text{pid}_0, \text{pid}_1$ ) for the initiator and responder. If  $\text{pid}_1$  is different from either the local party identifier, or the party identifier for  $\text{pid}$ , then end the activation.
  - Initialize  $\text{epoch\_id\_self} = \perp$ ,  $\text{epoch\_id\_partner} = \perp$ ,  $\text{sent\_msg\_num} = 0$  and  $N_{\text{last}} = 0$ ,  $\text{received\_msg\_num} = 0$ .
  - Initialize a dictionary  $\text{missed\_msgs} = \{\}$  and a list  $\text{receiving\_epochs} = []$ .
2. Provide input (`Decrypt`,  $c, \text{msg\_num}, N$ ) to  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid.fs} = (\text{sid}, \text{epoch\_id}))$ .
3. Upon receiving (`Decrypt`,  $c, \text{msg\_num}, N, v$ ) from  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid.fs})$ : if  $v = \text{Fail}$  then send (`ReceiveMessage`,  $h, \text{ad}, \text{Fail}$ ) to  $\text{pid}$ . //Otherwise,  $v$  is the decrypted message
4. While  $\text{msg\_num} > \text{received\_msg\_num}$ :  
//note down any expected messages
  - Append  $\text{received\_msg\_num}$  to the entry  $\text{missed\_msgs}[\text{epoch\_id}]$ .
  - Increment  $\text{received\_msg\_num} += 1$ .
5. If  $\text{msg\_num}$  is in the entry  $\text{missed\_msgs}[\text{epoch\_id}]$ :
  - remove it from the list.
  - If the entry  $\text{missed\_msgs}[\text{epoch\_id}]$  is now an empty list then remove  $\text{epoch\_id}$  from  $\text{missed\_msgs.keys}$ .
6. Else ( $\text{msg\_num} \notin \text{missed\_msgs}[\text{epoch\_id}]$ ):
  - If  $\text{epoch\_id} = \text{epoch\_id\_partner}$  or  $\text{sent\_msg\_num} = 0$ , output (`ReceiveMessage`,  $h, c, \perp$ ). Otherwise continue. //Starting new epoch-ratchet forward
  - Append the numbers  $\text{received\_msg\_num}, \dots, N$  to the entry  $\text{missed\_msgs}[\text{epoch\_id}]$ .
  - Send (`StopDecrypting`,  $N$ ) to  $(\mathcal{F}_{\text{fs\_aead}}, (\text{sid}, \text{epoch\_id\_partner}))$ . //‘Closing’ the  $\mathcal{F}_{\text{fs\_aead}}$  for the last epoch.
  - On receiving (`StopDecrypting`, `Success`), update  $\text{epoch\_id\_partner} = \text{epoch\_id}$ , and send (`StopEncrypting`) to  $(\mathcal{F}_{\text{fs\_aead}}, (\text{sid}, \text{epoch\_id\_self}))$ .
  - On receiving (`StopEncrypting`, `Success`), send (`ConfirmReceivingEpoch`,  $\text{epoch\_id}$ ) to  $(\mathcal{F}_{\text{eKE}}, \text{sid.eKE})$ .
  - On receiving (`ConfirmReceivingEpoch`,  $\text{epoch\_id}^*$ ), update  $\text{epoch\_id\_self} = \text{epoch\_id}^*$ ,  $N_{\text{last}} = \text{sent\_msg\_num}$ , and  $\text{sent\_msg\_num} = 0$ .
7. Output (`ReceiveMessage`,  $h, c, v$ ) to  $\text{pid}$  while deleting the decrypted message  $v$ .

**Corruption:** Upon receiving (`Corrupt`,  $\text{pid}$ ) from Env:

//Note that the **Corrupt** interface is not part of the “real” protocol; it is only included for modelling purposes.

1. Initialize a list  $S$  and send (`Corrupt`) as input to  $(\mathcal{F}_{\text{eKE}}, \text{sid.eKE} = \text{“eKE”}, \text{sid})$ .
2. On receiving (`Corrupt`,  $S_{\text{eKE}}$ ) from  $(\mathcal{F}_{\text{eKE}}, \text{sid.eKE} = \text{“eKE”}, \text{sid})$ , add it to  $S$  and continue. //now corrupt individual  $\mathcal{F}_{\text{fs\_aead}}$  instances.
3. For  $\text{epoch\_id} \in \text{missed\_msgs.keys}$  do:
  - Send (`Corrupt`) as input to  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid.fs} = (\text{“fs\_aead”}, \text{sid}, \text{epoch\_id}))$ .
  - On receiving  $S_{\text{epoch\_id}}$ , add it to  $S$ .
4. Output (`Corrupt`,  $\text{pid}_i, S$ ) to Env.



## $\mathcal{S}_{SM}$

**At first activation:** Send  $(\mathcal{F}_{SM}, \mathcal{I})$  to  $\mathcal{F}_{lib}$ .

**SendMessage:** On receiving  $(state_{\mathcal{I}}, \text{SendMessage}, pid, m)$  from  $(\mathcal{F}_{SM}, sid)$  do:

1. Set  $i$  such that  $pid = pid_i$ . //this is the identity of the sender
2. If this is the first invocation of **SendMessage** do:
  - Initialize  $diverge\_parties = false$ ,  $injectable = false$ ,  $corrupted\_party = \perp$ ,  $sent\_msg\_num_0 = 0$ , and  $sent\_msg\_num_1 = 0$ .
  - Create empty stacks  $sent\_ids_0 = []$  and  $sent\_ids_1 = []$ .
  - Create empty sets  $injectable\_ids_0 = \emptyset$  and  $injectable\_ids_1 = \emptyset$ .
  - Parse  $state.eKE$  from  $state_{\mathcal{I}}$
  - Send  $(state.eKE, \text{GenEpochId}, i, \perp)$  to Env (in the name of  $(\mathcal{F}_{eKE}, sid.eKE)$ ).
  - Upon receiving  $(state, \text{GenEpochId}, i, epoch\_id)$  from Env, update  $state.eKE \leftarrow state$  and push  $epoch\_id$  onto the stack  $sent\_ids_0$ .
3. Set  $epoch\_id$  equal to the top of the stack  $sent\_ids_i$ , and increment  $sent\_msg\_num_i += 1$ .
4. Parse  $state.fs.epoch\_id$  from  $state_{\mathcal{I}}$ .
5. Send  $(state.fs.epoch\_id, \text{Encrypt}, pid, sent\_msg\_num_i, m)$  to Env (in the name of  $(\mathcal{F}_{fs\_aead}, sid.fs)$ , where  $sid.fs = ("fs\_aead", sid, epoch\_id)$ )
6. On receiving  $(state, \text{Encrypt}, pid, sent\_msg\_num, c)$  from Env:
  - Update  $state.fs.epoch\_id \leftarrow state$
  - Add  $(h, c) \in \text{sentheaders}$ .
  - Update  $\text{sentheaders}$  and  $sent\_ids_0, sent\_ids_1$  in  $state_{\mathcal{I}}$
  - Send  $(state_{\mathcal{I}}, \text{SendMessage}, pid, epoch\_id, c)$  to  $(\mathcal{F}_{SM}, sid)$

**Inject:** On receiving  $(state_{\mathcal{I}}, \text{Inject}, pid, h, c)$  from  $(\mathcal{F}_{SM}, sid)$  do:

1. Set  $i$  such that  $pid = pid_i$ . //this is the identity of the attempted sender, and  $pid_{1-i}$  is the receiver
2. Parse  $h = (epoch\_id, msg\_num, N)$ .
3. Parse  $state.fs.epoch\_id$  from  $state_{\mathcal{I}}$ .
4. Send  $(state.fs.epoch\_id, \text{Inject}, pid, msg\_num, c)$  to Env (in the name of  $(\mathcal{F}_{fs\_aead}, sid.fs)$  where  $sid.fs = ("fs\_aead", sid, epoch\_id)$ ).
5. On receiving  $(state, \text{Inject}, v)$  from Env, update  $state.fs.epoch\_id \leftarrow state$ .
6. If  $(h, c) \notin \text{sentheaders} \wedge v = \perp$ : output  $(\text{Inject}, h, c, \perp)$  to  $(\mathcal{F}_{SM}, sid)$ . //in this case  $\mathcal{F}_{SM}$  should output **Fail**, otherwise decryption succeeds
7. If  $epoch\_id \neq sent\_ids_i$  then set  $diverge\_parties = true$ . //this epoch id was generated by the adversary rather than the sender, and it caused a divergence
8. If this is the first successfully received message with this  $epoch\_id$  do: //received first message from the sender's newest epoch
  - Parse  $state.eKE$  from  $state_{\mathcal{I}}$
  - Send a message  $(state.eKE, \text{GenEpochId}, i, epoch\_id)$  to Env (in the name of  $(\mathcal{F}_{eKE}, sid.eKE)$ )
  - On receiving  $(state, \text{GenEpochId}, i, epoch\_id^*)$ , update  $state.eKE \leftarrow state$  and add  $epoch\_id^*$  to the stack  $sent\_ids_{1-i}$  and add  $epoch\_id^*$  to  $state_{\mathcal{I}}$ . //this will be party  $i$ 's next  $epoch\_id$  when it next sends a message
  - If  $epoch\_id = sent\_ids_i.top$  and  $injectable = true$  and  $diverge\_parties = false$  and  $corrupted\_party = i$  then: set  $injectable = false$  and  $corrupted\_party = \perp$ . //if party  $i$  succeeds in establishing a new sending epoch, the adversary can no longer inject
9. Output  $(state_{\mathcal{I}}, \text{Inject}, h, c, v)$  to  $(\mathcal{F}_{SM}, sid)$ .

(The rest of this simulator is in Fig. 18 on Page 45)

Figure 17: Secure Messaging Simulator,  $\mathcal{S}_{SM}$

### $\mathcal{S}_{SM}$ continued...

(This simulator begins in Fig. 17 on Page 44)

**ReportState:** On receiving  $(\text{ReportState}, \text{pid}, \text{pending\_msgs})$  from  $(\mathcal{F}_{SM}, \text{sid})$  do:

1. Set  $i$  such that  $\text{pid} = \text{pid}_i$ . //this is the identity of the corrupted party
2. Set  $\text{corrupted\_party} = i$ ,  $\text{injectable\_ids}_0 = \text{sent\_ids}_0$ ,  $\text{injectable\_ids}_1 = \text{sent\_ids}_1$ , and initialize an empty list  $S_i$ .
3. Set  $\text{recv\_chain\_key} \xleftarrow{\$} \mathcal{K}_{ep}$  from the key distribution.
4. Send  $(\text{state.eKE}, \text{ReportState}, i, \text{recv\_chain\_key})$  to Env (in the name of  $(\mathcal{F}_{eKE}, \text{sid.eKE})$ ).
5. Upon receiving  $(\text{ReportState}, i, S^*)$  from Env, add  $S^*$  to  $S_i$ .
6. For all  $\text{epoch\_id}^*$  such that there exists a header  $h \in \text{pending\_msgs}$  containing  $\text{epoch\_id}^*$ :
  - Send  $(\text{state.fs.epoch\_id}^*, \text{Corrupt}, \text{pid}, \text{leakage})$  to Env (in the name of  $(\mathcal{F}_{fs\_aead}, \text{sid.fs})$  where  $\text{sid.fs} = (\text{"fs\_aead"}, \text{sid}, \text{epoch\_id}^*)$ ).
  - Upon receiving a response  $(\text{Corrupt}, \text{pid}, S^*)$  from Env, add  $S^*$  to the set  $S_i$ .
7. Output  $(\text{state}_{\mathcal{I}}, \text{ReportState}, \text{pid}_i, S_i)$  to  $(\mathcal{F}_{SM}, \text{sid})$ .

Figure 18: Secure Messaging Simulator  $\mathcal{S}_{SM}$  continued...

- Within **ReceiveMessage**, after performing several input validation checks (e.g., that the epoch/message header hasn't been used before)  $\mathcal{F}_{SM}$  will generate a tuple  $(\text{state.fs.epoch\_id}, \text{Inject}, \text{pid}, \text{msg})$ . When the parties are not compromised or diverged,  $\mathcal{F}_{SM}$  runs the code  $\mathcal{I}_{SM}$  on the tuple to decide whether the message is authentic. Otherwise, it sends the tuple to  $\mathcal{S}_{SM}$  to allow it to decide whether the message is authentic or even run a rushing attack to change the message contents. If this is the first successfully received message of a new epoch then  $\mathcal{I}_{SM}/\mathcal{S}_{SM}$  additionally generates a new **epoch\_id** for the recipient party to use when it sends its next message.
- If the corruption of a party is requested by environment, then  $\mathcal{F}_{SM}$  calls the simulator's **ReportState** method to generate a simulated state for the party which it sends to the environment in response.

In the remainder of this proof, we describe why the actions of  $\mathcal{S}_{SM}$  and  $\mathcal{I}_{SM}$  ensure that Env's view in its interaction with  $\mathcal{F}_{SM}$  is identically distributed (when treated as a random variable) to its view when interacting with the real protocol  $\Pi_{SGNL}$ .

This argument is divided into two cases based on whether the session is compromised or not. When the session is not compromised, the proof proceeds via induction over the steps of the internal code  $\mathcal{I}_{SM}$  where we argue that each individual action taken within the internal code maintains the indistinguishability property between the environment's view in the real and ideal worlds. Likewise, when the session is compromised, the proof proceeds by induction over the steps of the simulator. Observing that corruption, uncorruption, and divergence occur at the same times in  $\mathcal{F}_{SM}$  and  $\Pi_{SGNL}$ , one can see that checking the above cases suffices to complete the argument that the environment's view is identical in the real and ideal cases.

Note that the simulator  $\mathcal{S}_{SM}$  may assume that the first call made by the environment is to **SendMessage** and that all subsequent calls to **ReceiveMessage** use  $(\text{epoch\_id}, \text{msg\_num})$  headers that haven't been used before and that have valid **epoch\_id**. If these constraints do not hold, then we observe by inspection that both  $\mathcal{F}_{SM}$  and  $\Pi_{SGNL}$  terminate before ever invoking the ideal-world

## $\mathcal{I}_{SM}$

This internal adversary is only called when no parties are compromised.

**SendMessage:** On receiving  $(state_{\mathcal{I}}, \text{SendMessage}, \text{pid}, |m|)$  from  $(\mathcal{F}_{SM}, \text{sid})$  do:

1. Set  $i$  such that  $\text{pid} = \text{pid}_i$ . //this is the identity of the sender
2. If this is the first invocation of **SendMessage** do:
  - Initialize  $\text{sent\_msg\_num}_0 = 0$ , and  $\text{sent\_msg\_num}_1 = 0$ .
  - Create empty stacks  $\text{sent\_ids}_0 = []$  and  $\text{sent\_ids}_1 = []$ .
  - Create empty sets  $\text{injectable\_ids}_0 = \emptyset$  and  $\text{injectable\_ids}_1 = \emptyset$ .
  - Parse  $state.eKE$  from  $state_{\mathcal{I}}$
  - Run  $\mathcal{I}_{eKE}(state.eKE, \text{GenEpochId}, i, \perp)$
  - Upon receiving  $(state', \text{GenEpochId}, i, \text{epoch\_id})$  from  $Env$ , update  $state.eKE \leftarrow state'$  and push  $\text{epoch\_id}$  onto the stack  $\text{sent\_ids}_0$ .
3. Set  $\text{epoch\_id}$  equal to the top of the stack  $\text{sent\_ids}_i$ , and increment  $\text{sent\_msg\_num}_i += 1$ .
4. Parse  $state.fs.epoch\_id$  from  $state_{\mathcal{I}}$ .
5. Run  $\mathcal{I}_{fsaead}(state.fs.epoch\_id, \text{Encrypt}, \text{pid}, \text{sent\_msg\_num}_i, |m|)$  for specifically  $(\mathcal{I}_{fsaead}, \text{sid}.fs)$ , where  $\text{sid}.fs = ("fs.aead", \text{sid}, \text{epoch\_id})$ .
6. On receiving  $(state, \text{Encrypt}, \text{pid}, \text{sent\_msg\_num}, c)$  from  $\mathcal{I}_{fsaead}$ :
  - Update  $state.fs.epoch\_id \leftarrow state$
  - Add  $(h, c) \in \text{sentheaders}$
  - Update  $\text{sentheaders}$  and  $\text{sent\_ids}_0, \text{sent\_ids}_1$  in  $state_{\mathcal{I}}$
  - Send  $(state_{\mathcal{I}}, \text{SendMessage}, \text{pid}, \text{epoch\_id}, c)$  to  $(\mathcal{F}_{SM}, \text{sid})$

**Inject:** On receiving  $(state_{\mathcal{I}}, \text{Inject}, \text{pid}, h, c)$  from  $(\mathcal{F}_{SM}, \text{sid})$  do:

1. Set  $i$  such that  $\text{pid} = \text{pid}_i$ . //this is the identity of the attempted sender, and  $\text{pid}_{1-i}$  is the receiver
2. Parse  $h = (\text{epoch\_id}, \text{msg\_num}, N)$ , read in  $\text{sentheaders}$  and  $\text{sent\_ids}_0, \text{sent\_ids}_1$  from  $state_{\mathcal{I}}$ , and read in  $state.fs.epoch\_id$  from  $state_{\mathcal{I}}$
3. Run  $\mathcal{I}_{fsaead}(state.fs.epoch\_id, \text{Inject}, \text{pid}, \text{msg\_num}, c)$  specifically for  $(\mathcal{F}_{fs.aead}, \text{sid}.fs)$  where  $\text{sid}.fs = ("fs.aead", \text{sid}, \text{epoch\_id})$ .
4. On receiving  $(state, \text{Inject}, v)$  from  $\mathcal{I}$ , update  $state.fs.epoch\_id \leftarrow state$ .
5. If  $(h, c) \notin \text{sentheaders} \wedge (v = \perp \vee (\nexists c^* \text{ s.t. } (h, c^*) \in \text{sentheaders}))$ : output  $(state_{\mathcal{I}}, \text{Inject}, h, c, \perp)$  to  $(\mathcal{F}_{SM}, \text{sid})$ . //in this case  $\mathcal{F}_{SM}$  should output **Fail**, otherwise decryption succeeds
6. If  $\text{epoch\_id} \neq \text{sent\_ids}_i$  then set  $\text{diverge\_parties} = \text{true}$ . //this epoch id was generated by the adversary rather than the sender, and it caused a divergence
7. If this is the first successfully received message with this  $\text{epoch\_id}$  do: //received first message from the sender's newest epoch
  - Parse  $state.eKE$  from  $state_{\mathcal{I}}$
  - Run  $\mathcal{I}_{eKE}(state.eKE, \text{GenEpochId}, i, \text{epoch\_id})$
  - On receiving  $(state', \text{GenEpochId}, i, \text{epoch\_id}^*)$ , update  $state.eKE \leftarrow state'$ , add  $\text{epoch\_id}^*$  to the stack  $\text{sent\_ids}_{1-i}$  and add  $\text{epoch\_id}^*$  to  $state_{\mathcal{I}}$ . //this will be party  $i$ 's next  $\text{epoch\_id}$  when it next sends a message
8. Output  $(state_{\mathcal{I}}, \text{Inject}, h, c, v)$  to  $(\mathcal{F}_{SM}, \text{sid})$ .

Figure 19: Internal Adversary  $\mathcal{I}_{SM}$

adversary  $\mathcal{S}_{\text{SM}}$  or the real-world adversary  $\mathcal{A}$ , respectively.

Without loss of generality, we restrict our attention to a deterministic environment  $\text{Env}$  and a dummy adversary  $\mathcal{A}$ . Consequently, there are only two sources of randomness in the entire execution: first, the choice of  $\text{ik}$  and  $\text{rk}$  within  $\mathcal{F}_{\text{DIR}}$ , which influences the epoch key generated in  $\mathcal{F}_{\text{eKE}}$ , and second, the choice of epoch key  $\text{recv\_chain\_key}$  within the view of a corrupted receiver.

**Case 1: neither party is compromised.** The functionality  $\mathcal{F}_{\text{SM}}$  notifies the internal code when any message is sent or received, as long as it passes the input validation checks described above. During `ReceiveMessage`, the internal code is given a message header ( $\text{epoch\_id}$ ,  $\text{msg\_num}$ ,  $N$ ) together with a ciphertext  $c$ , and it is permitted to attempt to inject a message. In the analogous `ReceiveMessage` routine in the real world, the protocol  $\Pi_{\text{SGNL}}$  invokes the specific instance of  $\mathcal{F}_{\text{fs\_aead}}$  corresponding to  $\text{epoch\_id}$ , which then sends a notification to the internal code  $\mathcal{I}_{\text{fsaead}}$  asking for the desired message to inject. In the ideal world, the internal code  $\mathcal{I}_{\text{SM}}$  sends this exact backdoor message to the internal code  $\mathcal{I}_{\text{fsaead}}$  and retrieves a value  $v$ . It is straightforward to confirm by inspection that the logic in  $\mathcal{I}_{\text{SM}}$  matches the input-validation logic used within the real world:  $v$  is ignored if  $c$  is a valid ciphertext created by a previous invocation to `SendMessage`, and otherwise the message is authenticated if and only if  $v \neq \perp$ . Finally, if this is the first successfully received message of a new epoch, then the real world protocol  $\Pi_{\text{SGNL}}$  invokes  $\mathcal{F}_{\text{eKE}}$  to perform a public ratchet; the internal code  $\mathcal{I}_{\text{SM}}$  emulates the corresponding backdoor message that  $\mathcal{F}_{\text{eKE}}$  sends to its internal adversarial code  $\mathcal{I}_{\text{eKE}}$  to receive the  $\text{epoch\_id}$  for the new epoch. Hence, the real and ideal worlds move to a new epoch in lockstep. The internal code also records all  $\text{epoch\_ids}$  for later use during `SendMessage`, as described below. It is straightforward to check that the code of  $\mathcal{F}_{\text{SM}}$  and  $\Pi_{\text{SGNL}}$  identically perform other checks that do not involve the (ideal or real world) adversary at all, such as refusing to decrypt a message whose header  $h$  has already been used or that is invalid because its  $\text{msg\_num}$  is larger than expected. As a result, the views of the environment in both scenarios contains the same backdoor messages, as well as the same outputs since they are deterministically derived from the environment’s own responses to the backdoor messages.

During `SendMessage`, the internal adversarial code is given the identity of the sending party  $\text{pid}$  and the length of the desired message  $\ell = |m|$ . In the ideal world, by the time that  $\mathcal{F}_{\text{SM}}$  has invoked  $\mathcal{I}_{\text{SM}}$ , it has properly initialized `SendMessage` and is using  $\mathcal{I}_{\text{SM}}$  to generate a ciphertext so that it can complete the message transmission. Note that  $\mathcal{I}_{\text{SM}}$  generates the ciphertext by invoking the code  $\mathcal{I}_{\text{fsaead}}$  with state  $\text{state.fs.epoch\_id}$  that corresponds to the newest epochid stored in the state  $\text{state}_{\mathcal{I}}$  given to  $\mathcal{I}_{\text{SM}}$ . For the corresponding call to `SendMessage` in the real world,  $\Pi_{\text{SGNL}}$  performs the same initialization and then invokes the specific instance of  $\mathcal{F}_{\text{fs\_aead}}$  corresponding to the latest  $\text{epoch\_id}$ , which in turn uses the internal code  $\mathcal{I}_{\text{fsaead}}$  directly to generate the ciphertext. Because  $\mathcal{I}_{\text{SM}}$  has  $\ell$  and the newest  $\text{epoch\_id}$  (from previous calls to `ReceiveMessage`), the views of the environment in both scenarios are the same.

Finally, it is simple to observe by inspection that internal state variables like  $\text{msg\_num}$  and  $N$  remain in sync as well.

**Case 2: one or both parties are compromised.** During the interval where the adversary can tamper with the communication (i.e.,  $\text{Inject} \in \text{advControl}$ ) or if the parties’ root chains have diverged (i.e.,  $\text{diverge\_parties} = \text{true}$ ), the simulator works analogously to the way  $\mathcal{I}_{\text{SM}}$  does in the case described above with some small differences. Specifically, the simulator  $\mathcal{S}_{\text{SM}}$  sends backdoor messages to the environment (addressed to  $\mathcal{S}_{\text{fsaead}}$ ) in the situations where the internal code  $\mathcal{I}_{\text{SM}}$  would have run the code  $\mathcal{I}_{\text{fsaead}}$ , there is a small difference in how the returned value is processed in the case of `Inject`.

The difference between  $\mathcal{I}_{\text{SM}}$  and  $\mathcal{S}_{\text{SM}}$  in `SendMessage` is the following: While  $\mathcal{I}_{\text{SM}}$  invokes  $\mathcal{I}_{\text{fsaead}}$  with the length  $|m|$ , the simulator  $\mathcal{S}_{\text{SM}}$  sends a nearly identical backdoor message with the length  $|m|$  replaced by message  $m$ , this backdoor message is marked for  $\mathcal{S}_{\text{fsaead}}$  corresponding to the latest epoch id stored in `state $\mathcal{I}$` .

During `Inject`, the difference between  $\mathcal{S}_{\text{SM}}$  and  $\mathcal{S}_{\text{SM}}$  is in how the returned value from the environment (resp.  $\mathcal{I}_{\text{fsaead}}$ ) is processed. The simulator  $\mathcal{S}_{\text{SM}}$  sends a backdoor message to the environment that's identical to the input that  $\mathcal{I}_{\text{SM}}$  runs  $\mathcal{I}_{\text{fsaead}}$  with. On receiving a response,  $\mathcal{S}_{\text{SM}}$  must compute an input validation predicate that's different than the one in  $\mathcal{I}_{\text{SM}}$ , because in this case whenever the ciphertext  $c$  isn't exactly the one generated in the corresponding `Encrypt` call, both  $\mathcal{F}_{\text{SM}}$  and  $\Pi_{\text{SGNL}}$  allow the adversary to inject any message of its choice.

To complete the proof, it only remains to check that corruption and uncorruption happen at the same times in the real and ideal worlds and that the environment receives the same state in response to a `ReportState` command.

Corruption is initiated by the environment itself; both  $\Pi_{\text{SGNL}}$  and  $\mathcal{F}_{\text{SM}}$  (with  $\mathcal{S}_{\text{SM}}$ ) respond immediately to this request by corrupting parties and reporting state back to the environment. In response to a `ReportState` command sent to  $\mathcal{F}_{\text{SM}}$ , the simulator  $\mathcal{S}_{\text{SM}}$  constructs the corrupted party's view by using  $\mathcal{F}_{\text{eKE}}$  and all pertinent  $\mathcal{F}_{\text{fsaead}}$  in exactly the same way as  $\Pi_{\text{SGNL}}$  does, with only one exception. Because  $\Pi_{\text{SGNL}}$  does not expose to the environment the commands `GetSendingKey` and `GetReceivingKey` within its  $\mathcal{F}_{\text{eKE}}$  subroutine, it suffices for the simulator to sample `recv_chain_key` independently (from the same distribution as  $\mathcal{F}_{\text{eKE}}$  does) because `Env` cannot make the queries needed to test consistency with the underlying state held within  $\mathcal{F}_{\text{eKE}}$ .

Additionally, it is simple to observe by inspection that  $\mathcal{F}_{\text{SM}}$  and  $\Pi_{\text{SGNL}}$  uncorrupt a party and mark the parties as diverged at the same times.

□

## 6 The Public Key Ratchet: Realizing $\mathcal{F}_{\text{eKE}}$

This section describes a protocol  $\Pi_{\text{eKE}}$  (Fig. 20) that UC-realizes  $\mathcal{F}_{\text{eKE}}$  (Fig. 13). This protocol mirrors the public key ratchet from the Signal protocol – in particular it uses values from a continuous Diffie-Hellman ratchet as inputs to a key derivation function (KDF) chain. The properties of a KDF chain are captured by a new primitive, Cascaded PRF-PRG (CPRFG) (Section 6.2) that is realisable in the plain model and may be of independent interest. In Section 6.3 it is shown that  $\Pi_{\text{eKE}}$  (instantiated with a CPRFG) UC-realizes  $\mathcal{F}_{\text{eKE}}$  in the plain model (without a random oracle) in the presence of the global functionalities  $\mathcal{F}_{\text{DIR}}$ ,  $\mathcal{F}_{\text{lib}}$ , and  $\mathcal{F}_{\text{LTM}}$ .

### 6.1 Protocol $\Pi_{\text{eKE}}$

This section describes a protocol  $\Pi_{\text{eKE}}$  (Fig. 20) that UC-realizes the epoch key exchange functionality  $\mathcal{F}_{\text{eKE}}$  (Fig. 13). The protocol  $\Pi_{\text{eKE}}$ , which persists for the entire duration of the secure messaging session between two parties, mirrors the public key ratchet from the Signal protocol. It uses values from a continuous Diffie-Hellman ratchet as inputs to a key derivation function (KDF) chain. The epoch key exchange protocol  $\Pi_{\text{eKE}}$  (Section 6.1) has a security tradeoff: (1) When the KDF used in  $\Pi_{\text{eKE}}$  is modeled as a random oracle, the protocol realises  $\mathcal{F}_{\text{eKE}}$  as is, this allows the overall protocol  $\Pi_{\text{SGNL}}$  to realise  $\mathcal{F}_{\text{eKE}}$  as is. Allowing an adversary to send malformed packets during corruption means that the adversary can get extra information about the parties secret Diffie-Hellman exponents:

Say that Alice and Bob have secret exponents  $a$  and  $b$  respectively. Alice tried to share her new epoch identifier  $g^a$  with Bob, but Bob has not yet received this information. The adversary can

now send as many malformed epoch identifiers  $g^{a'}$  to the Bob as it chooses and get back the value  $KDF(g^{a'b})$  for each. Additionally  $g^{a'}$  can depend on the honest epoch identifier  $g^a$  of Alice.<sup>11</sup>

This is the reason that the adversary to  $\mathcal{F}_{\text{eKE}}$  can extend the healing time after corruption by sending malformed packets to parties during the first epoch after recovery. Without step 5a  $\mathcal{F}_{\text{eKE}}$  could not be realised by  $\Pi_{\text{eKE}}$  in the plain model without specific strong properties being enforced on  $KDF$ .<sup>12</sup>

Recall that each of the two parties will run their own copy of epoch key-exchange protocol  $\Pi_{\text{eKE}}$ . This is different from the ideal world where the parties share a copy of the epoch key exchange functionality  $\mathcal{F}_{\text{eKE}}$  that they both interact with. One cosmetic impact of this is that each party's protocol must be initialized on first activation. The functionality  $\mathcal{F}_{\text{eKE}}$  is initialized just once, on the first call to `ConfirmReceivingEpoch` by the initiator of the conversation. Since there are two copies of the protocol, each must be initialized separately. The initiator's instance of the protocol  $\Pi_{\text{eKE}}$  is also initialized on the first call to `ConfirmReceivingEpoch` by the initiator. However, the responder's instance of  $\Pi_{\text{eKE}}$  is initialized on the first call to `GetReceivingKey` by the responder since this is the first action the responder must take.

The protocol has two basic components: (1) A continuous Diffie-Hellman ratchet between the two parties which provides both forward secrecy and healing from compromise. This is the place where parties can add fresh randomness into the system. (2) A KDF-chain that provides immediate decryption for the secure messaging system. This chain leverages the healing property of the Diffie-Hellman ratchet and extends the forward secrecy property. We briefly describe each of the methods of  $\Pi_{\text{eKE}}$  below. The full details can be found in Figure 20 on Page 50.

Each of the parties involved in the protocol execution run an instance of  $\Pi_{\text{eKE}}$ . Just like in  $\mathcal{F}_{\text{eKE}}$  the parties take turns starting epochs in which they will send messages. They achieve this interleaving by updating their sending randomness (via a new Diffie-Hellman pair) as soon as they get confirmation of a new Diffie-Hellman pair being used by the other party. This means that most of the work happens when the parties successfully receive a message in the other party's newest sending epoch. In that vein, let's start by briefly describing the `GetReceivingKey` method.

**GetReceivingKey** When a party receives a new public Diffie-Hellman value `epoch_id` from its partner, it runs the method `(GetReceivingKey, epoch_id)` to produce a tentative `rcv_chain_key`. This key is then confirmed to be the correct value (or confirmed to be wrong) by the party out of band of the epoch key exchange protocol.

**ConfirmReceivingEpoch** If a new public Diffie-Hellman value `epoch_id` produces a key that the party confirms to be correct (with this confirmation being out of band of the epoch key exchange protocol), then the party will run the `ConfirmReceivingEpoch` method to update the both the Diffie-Hellman ratchet and the KDF-chain accordingly. (This update entails the overwriting of the partner's old epoch id with the new one and the overwriting of the old `root_key` with the new one output by  $KDF$ .) The party also knows at this point that it must update its own sending randomness – so this method additionally updates both the Diffie-Hellman ratchet and the KDF chain again according to a randomly chosen Diffie-Hellman pair `(epoch_id', epoch_key')`, this update also produces a new `sending_chain_key` value (output by  $KDF$ ) which is stored temporarily until the

<sup>11</sup>We note that in the context of the full secure messaging protocol the adversary does not actually see these chain keys. It instead only learns if decryption succeeds or not.

<sup>12</sup>Another option for weakening the functionality is to assume that the adversary is completely inactive during the recovery epoch, or even during the entire healing period as is done in the work by Alwen et al. [1], we prefer instead to characterize the exact additional power gained by the adversary.

## $\Pi_{eKE}$

This protocol has a party id  $pid$  and session id  $sid.eKE$  of the form:  $sid.eKE = ("eKE", sid)$  where  $sid = (sid', pid_0, pid_1)$ .

This protocol uses protocols  $keyGen$  and  $KDF$  as subroutines: (1)  $keyGen$  chooses a random Diffie-Hellman exponent  $epoch\_key \xleftarrow{\$} |\mathbb{G}|$  for a known group  $\mathbb{G}$  and sets  $epoch\_id = g^{epoch\_key}$ . It then outputs  $(epoch\_key, epoch\_id)$ . (2) The protocol  $KDF$  is a cascaded PRF-PRG.

//The method **ConfirmReceivingEpoch** is also used by the initiator of the conversation to start the first sending epoch (epoch 0) before it has any receiving epochs to confirm.

**ConfirmReceivingEpoch:** On input  $(ConfirmReceivingEpoch, epoch\_id^*)$  from  $(\Pi_{SGNL}, sid, pid')$ :

1. If  $pid' \neq pid$ , then end the activation. Let  $i$  be such that  $pid = pid_i$ .
2. Set  $temp\_epoch\_id\_partner_i = epoch\_id^*$ .
3. If this is the first activation:
  - Initialize state variables  $(root\_key_a, root\_key_c), epoch\_id, epoch\_key, sending\_chain\_key = \perp$ .
  - Send  $(GetInitKeys, pid_{1-i}, pid_i)$  to  $\mathcal{F}_{DIR}$ .
  - Upon receiving  $(GetInitKeys, ik_j^{pk}, rk_j^{pk}, ok_{j \leftarrow i}^{pk})$ , if  $ok_{j \leftarrow i}^{pk} = \perp$  then output  $(ConfirmReceivingEpoch, Fail)$ . Otherwise, send input  $(ComputeSendingRootKey, ik_j^{pk}, rk_j^{pk}, ok_{j \leftarrow i}^{pk})$  to  $\mathcal{F}_{LTM}$ .
  - Upon receiving  $(ComputeSendingRootKey, s = (s_a, s_c), ek_i^{pk})$ , set  $(root\_key_a, root\_key_c) = (s_a, s_c)$ .
  - Run the subroutine **Compute Sending Chain Key**.
  - Erase  $ek_i^{pk}$  and output  $(ConfirmReceivingEpoch, epoch\_id_{self} || ek_i^{pk} || ok_{j \leftarrow i}^{pk})$  to  $(\Pi_{SGNL}, sid, pid_i)$ .
4. Else (this is not the first activation):
  - Run the steps in **Compute Sending Chain Key**.
  - Output  $(ConfirmReceivingEpoch, epoch\_id_{self})$  to  $(\Pi_{SGNL}, sid, pid_i)$ .

**GetSendingKey:** On receiving input  $(GetSendingKey)$  from  $(\Pi_{mKE}, sid.mKE, pid')$ :

1. If  $pid' \neq pid$ , or if  $sending\_chain\_key$  has already been erased, end the activation.
2. Output  $(GetSendingKey, sending\_chain\_key)$  and erase  $sending\_chain\_key$ .

**GetReceivingKey:** On receiving input  $(GetReceivingKey, epoch\_id)$  from  $(\Pi_{mKE}, sid, pid')$ :

1. If  $pid' \neq pid$ , then end the activation. Otherwise, let  $i$  be such that  $pid = pid_i$ .
2. Set  $temp\_epoch\_id\_partner = epoch\_id$ .
3. If this is the first activation:
  - Initialize state variables  $(root\_key_a, root\_key_c), epoch\_id, epoch\_key, sending\_chain\_key = \perp$ .
  - Parse  $epoch\_id = (epoch\_id', ek_j^{pk}, ok_{i \leftarrow j}^{pk})$  and set  $temp\_epoch\_id\_partner = epoch\_id'$
  - Send  $(GetResponseKeys, pid_{1-i})$  to  $\mathcal{F}_{DIR}$ .
  - Upon receiving  $(GetResponseKeys, ik_j^{pk})$ , send input  $(ComputeReceivingRootKey, ik_j^{pk}, ek_j^{pk}, ok_{i \leftarrow j}^{pk})$  to  $\mathcal{F}_{LTM}$ .
  - Upon receiving  $(ComputeReceivingRootKey, s = (s_a, s_c))$ , set  $(root\_key_a, root\_key_c) = (s_a, s_c)$ .
4. Run the subroutine **Compute Receiving Chain Key**.
5. Output  $(GetReceivingKey, temp\_recv\_chain\_key)$  and erase  $temp\_recv\_chain\_key$ .

(The rest of this protocol is in Fig. 21 on Page 52)

Figure 20: The Epoch Key Exchange Protocol  $\Pi_{eKE}$



party retrieves it using the `GetSendingKey` method. (This second update also entails overwriting of old value. This time, the party’s old Diffie-Hellman pair is overwritten along with the old `root_key`.)

**GetSendingKey** This method simply outputs the stored `sending_chain_key` value and then deletes it. If the value has already been deleted then it does nothing.

Now we briefly discuss the initialization of the protocol for each party and the corrupt method that exists only for record keeping purposes.

**Initiator Initialization** The first time the initiator runs the `ConfirmReceivingEpoch` method of  $\Pi_{\text{eKE}}$ , the method must initialize the KDF-chain and the Diffie-Hellman ratchet. It then updates both using a randomly chosen Diffie-Hellman pair (`epoch_id'`, `epoch_key'`), and temporarily store the produced `sending_chain_key` value till the party retrieves it using the `GetSendingKey` method. To initialize the Diffie-Hellman Ratchet, the party retrieves the responder’s keys from directory functionality  $\mathcal{F}_{\text{DIR}}$  using the `GetInitKeys` method. It can then use the `ComputeSendingRootKey` method of its long-term memory functionality  $\mathcal{F}_{\text{LTM}}$  to run a triple Diffie-Hellman on both parties’ keys. This initializes the Diffie-Hellman ratchet and KDF-chain, it also binds the conversation to the longterm identity keys of the two parties.

**Responder Initialization** The first time the responder runs the `GetReceivingKey` method of  $\Pi_{\text{eKE}}$ , the method must initialize the KDF-chain and Diffie-Hellman ratchet of the responder in much the same way as the initialization of the Initiator that happens on the its first call to `ConfirmReceivingEpoch`. After this, the steps of the `GetReceivingKey` method are run like they will be for the rest of the conversation.

**Corrupt** The corrupt method in Figure 20 defines the model of corruption we are considering. On corruption,  $\Pi_{\text{eKE}}$  returns its internal state containing: (`epoch_key`, `epoch_idself`, `epoch_idpartner`, `root_key`). Note that, as with the other protocols, the `Corrupt` method is not a “real” interface, but is only record keeping for the purposes of the model.

**Remark**

We remark that a simple modification to  $\Pi_{\text{eKE}}$ , also mentioned in [1], allows for the protocol to heal even faster with a small increase in communication. This modified protocol can realize a functionality  $\mathcal{F}_{\text{eKE}}$  that heals in 2 rounds instead of 3. Such a protocol can easily be substituted in place of the current  $\Pi_{\text{eKE}}$  to show that the resulting  $\Pi_{\text{SGNL}} + \Pi_{\text{aead}} + \Pi'_{\text{eKE}}$  system will realise functionality  $\mathcal{F}_{\text{SM}}$  that heals from corruption in  $\geq 2$  rounds. This modification has each party use two Diffie-Hellman pairs when an epoch turns over, one for the party’s new sending epoch, and one for the other party’s next sending epoch when it responds. This allows both parties to delete the Diffie-Hellman exponents corresponding to their sending epochs as soon as they begin.

Bienstock et al. [13] provide a different modification, which they call the ‘triple ratchet’, that achieves similarly fast healing from corruption without the additional communication overhead.

## 6.2 Cascaded PRF-PRG (CPRFG)

A crucial component of the epoch key exchange protocol  $\Pi_{\text{eKE}}$  (fully specified later in Fig. 20 on Page 50) is a stateful key derivation function (KDF). The KDF consists of a secret state  $s = (s_a, s_c)$ , and two deterministic functions, `KDF.Advance` and `KDF.Compute`, which take an values  $s_a$ ,  $s_c$  respectively along with externally provided input  $r$ , and do the following:

## $\Pi_{eKE}$ continued...

(This protocol begins in Fig. 20 on Page 50)

**Corrupt:** On receiving (**Corrupt**) from  $(\Pi_{SGNL}, \text{sid}, \text{pid}_i)$ : return  $(\text{epoch\_key}, \text{epoch\_id}_{\text{self}}, \text{epoch\_id}_{\text{partner}}, (\text{root\_key}_a, \text{root\_key}_c))$  to  $(\Pi_{SGNL}, \text{sid}, \text{pid}_i)$

//Note that the **Corrupt** interface is not part of the “real” protocol; it simply serves to formalize what the adversary sees on corruption.

//Below are subroutines used in the interfaces above. This is where the cascaded PRF-PRG protocol *KDF* is used.

### Compute Sending Chain Key:

1. If this is the first activation start at step 5.
2. Compute a root input  $\text{root\_input} = \text{Exp}(\text{temp\_epoch\_id}_{\text{partner}}, \text{epoch\_key}_{\text{self}})$ .
3. Compute the new root key  $KDF.\text{Advance}(\text{root\_key}_a, \text{root\_input}) = (\text{root\_key}'_a, \text{root\_key}'_c)$  and update the value  $(\text{root\_key}_a, \text{root\_key}_c) = (\text{root\_key}'_a, \text{root\_key}'_c)$ .
4. Generate a key pair  $(\text{epoch\_key}_{\text{self}}, \text{epoch\_id}_{\text{self}}) \leftarrow \text{keyGen}()$ .
5. Compute a new root input  $\text{root\_input} = \text{Exp}(\text{temp\_epoch\_id}_{\text{partner}}, \text{epoch\_key}_{\text{self}})$ .
6. Compute the sending chain key  $KDF.\text{Compute}(\text{root\_key}_c, \text{root\_input}) = (\text{root\_key}'_c, \text{sending\_chain\_key})$  and update  $\text{root\_key}_c = \text{root\_key}'_c$ .
7. Erase  $\text{root\_input}$ . //The old root key is overwritten and therefore erased. The old sending chain key was already erased.

### Compute Receiving Chain Key:

1. Compute  $\text{root\_input} = \text{Exp}(\text{temp\_epoch\_id}_{\text{partner}}, \text{epoch\_key}_{\text{self}})$ .
2. Compute  $KDF.\text{Compute}(\text{root\_key}_c, \text{root\_input}) = (\text{root\_key}'_c, \text{temp\_rcv\_chain\_key})$  and update the value  $\text{root\_key}_c = \text{root\_key}'_c$ .
3. Erase  $\text{root\_input}$ . //The old root key is overwritten and therefore already erased.

Figure 21: The Epoch Key Exchange Protocol  $\Pi_{eKE}$  (continued)

- **KDF.Advance**( $s_a, r$ ): Update the state  $s = (s_a, s'_c)$  according to the input  $s_a, r$  to generate an updated state  $s^* = (s_a^*, s'_c)$ . (Within the Signal protocol, an application of **KDF.Advance** corresponds to advancing an epoch.)
- **KDF.Compute**( $s_c, r$ ): Generate a key  $k$  for a potential next epoch corresponding to a given the input  $s_c, r$ , along with an updated state  $s'_c$ . (Within the Signal protocol, an application of **KDF.Compute** corresponds to computing a receiving chain key  $k$  *without advancing the epoch*. Note that we allow an update to the state even in this case; this is a crucial ingredient in providing security against adaptive corruptions.)

In this subsection we define a new primitive, **Cascaded PRF-PRG (CPRFG)**, that captures the security requirements from this KDF. Additionally, we present a candidate construction of a **CPRFG** in the standard model. (The HKDF function [44] used in the Signal application also meets this definition of a **CPRFG**, albeit in the random oracle model.)

### 6.2.1 Defining the CPRFG

The definition of a **Cascaded PRF-PRG (CPRFG)** extends the PRF-PRNG model from Alwen et al. [1] to capture the requirements from a KDF-chain under adaptive corruptions. The benefit of the CPRFG definition is that it directly requires the exact adaptive properties needed from a KDF-chain as specified in the Signal architecture [51] while the PRF-PRNG definition from [1] has to be stringent enough to provide the right adaptive guarantees when used in a chain. This definition follows the real/ideal paradigm. Specifically, it specifies an ideal game and a real game, and requires existence of a simulator  $\mathcal{S}_{\text{KDF}}$  such that no adversary can distinguish whether it is interacting in a real game that represents an execution of an actual *KDF*, or in an ideal game with simulator  $\mathcal{S}_{\text{KDF}}$ , where the ideal game represents the expected behavior of the system that exhibits the properties described above.

In both games, the adversary can repeatedly make one of four queries: (**Compute**,  $r$ ), (**Advance**,  $r$ ), (**Compromise**,  $r$ ), or (**Recover**) which are answered as defined in Figure 22 below.

We note that, while the definition does follow the real/ideal paradigm, we found it more convenient not to frame it within the UC framework. In particular, this allows us to directly require that **KDF.Compute** and **KDF.Advance** be deterministic functions, and to make the queries in the game different than the normal API of **KDF.Compute** and **KDF.Advance**.

A formal definition follows in Theorem 11 with the full game details shown in Fig. 22.

**Definition 11 (Cascaded PRF-PRG (CPRFG))** *Let  $m$  be a polynomial,  $\mathcal{O}_{\text{CPRFG}}$  be an oracle that runs either the real or ideal version of the CPRFG security game (Figure 22, pg Page 54),  $\lambda \in \mathbb{N}$  be the security parameter, and  $|R_\lambda| \geq 2^\lambda$  be an input space. Let the module *KDF* have a secret state  $(s_a, s_c) \in \{0, 1\}^\lambda$  and deterministic functions (**KDF.Compute**, **KDF.Advance**) such that: **KDF.Advance** takes inputs  $s_a, r \in \{0, 1\}^\lambda \times R_\lambda$  and produces an output  $(s_a^*, s'_c) \in \{0, 1\}^\lambda \times \{0, 1\}^\lambda$ . **KDF.Compute** takes inputs  $s_c, r \in \{0, 1\}^\lambda \times R_\lambda$  and produces an output  $s'_c, k \in \{0, 1\}^\lambda \times \{0, 1\}^{m(\lambda)}$ .*

*KDF is a **cascaded PRF-PRG (CPRFG)** if  $\exists$  probabilistic polynomial time (PPT) algorithm  $\mathcal{S}_{\text{KDF}}$  such that any PPT machine  $\mathcal{A}$  interacting with oracle  $\mathcal{O}_{\text{CPRFG}}$  can distinguish whether the oracle is running the real game or the ideal game with advantage at most negligible in  $\lambda$ .*

### 6.2.2 Properties of the CPRFG

To provide intuition for this new primitive, in this section we provide an informal discussion of some of the properties that it guarantees.

### Cascaded PRF-PRG Security Game

Let  $m$  be a polynomial,  $\lambda \in \mathbb{N}$  a security parameter. This figure describes the Cascaded PRF-PRG security game for a KDF with secret state  $(s_a, s_c) \in \{0, 1\}^\lambda \times \{0, 1\}^\lambda$  and deterministic functions ( $KDF.\mathbf{Compute}$ ,  $KDF.\mathbf{Advance}$ ) such that:  $KDF.\mathbf{Advance}$  takes inputs  $s_a, r \in \{0, 1\}^\lambda \times R_\lambda$  and produces an output  $(s_a^*, s_c^*) \in \{0, 1\}^\lambda \times \{0, 1\}^\lambda$ .  $KDF.\mathbf{Compute}$  takes inputs  $s_c, r \in \{0, 1\}^\lambda \times R_\lambda$  and produces an output  $s'_c, k \in \{0, 1\}^\lambda \times \{0, 1\}^{m(\lambda)}$ .

#### Real game:

- On Initialization: Oracle  $\mathcal{O}_{\text{CPRFG}}$  is initialized with uniformly sampled state  $s = (s_a, s_c) \xleftarrow{\$} \{0, 1\}^\lambda$ .
- On input (Compute,  $r$ ):  $\mathcal{O}_{\text{CPRFG}}$  runs  $KDF.\mathbf{Compute}(s_c, r) = (s'_c, k)$ , outputs  $k$ , changes state  $s = s'$ .
- On input (Advance,  $r$ ):  $\mathcal{O}_{\text{CPRFG}}$  runs  $KDF.\mathbf{Advance}(s_a, r) = (s_a^*, s_c^*)$ , and changes state  $(s_a, s_c) = (s_a^*, s_c^*)$ .
- On input (Compromise,  $r^*$ ):  $\mathcal{O}_{\text{CPRFG}}$  outputs the old state  $s$ . Mark the state as compromised.
- On input (Recover):  $\mathcal{O}_{\text{CPRFG}}$  samples  $r' \xleftarrow{\$} R_\lambda$  uniformly at random, computes  $KDF.\mathbf{Compute}(s_c, r') = s'_c, k$ , and  $KDF.\mathbf{Advance}(s_a, r') = (s_a^*, s_c^*)$ , and changes state  $s = (s_a^*, s_c^*)$ . Mark the state as recovered and output  $k$ .

#### Ideal game:

- On Initialization: Oracle  $\mathcal{O}_{\text{CPRFG}}$  is initialized with a uniformly sampled state  $F \xleftarrow{\$} \{f : R_\lambda \rightarrow \{0, 1\}^{m(\lambda)}\}$ .
- On input (Compute,  $r$ ):
  1. **If** the state is compromised it outputs  $\mathcal{S}(\text{Compute}, r) = k$  and adds the tuple  $(r, k)$  to the list `computed_values`.
  2. **Else** (the state is not compromised), it outputs  $F(r) = k$  and adds the tuple  $(r, k)$  to the list `computed_values`.
- On input (Advance,  $r$ ):
  1. The oracle  $\mathcal{O}_{\text{CPRFG}}$  resets `computed_values` = {}.
  2. **If** the state isn't compromised,  $\mathcal{O}_{\text{CPRFG}}$  updates its state to a new function  $F \xleftarrow{\$} \{f : R_\lambda \rightarrow \{0, 1\}^{m(\lambda)}\}$  sampled uniformly at random.
  3. **Else** (the state is compromised),  $\mathcal{O}_{\text{CPRFG}}$  runs  $\mathcal{S}(\text{Advance}, r)$ .
- On input (Compromise,  $r^*$ ): **If** the state is not compromised then  $\mathcal{O}_{\text{CPRFG}}$  marks the state as compromised and outputs  $\mathcal{S}(\text{computed\_values}, (r^*, F(r^*)))$  (or  $\mathcal{S}(\text{computed\_values}, \perp)$  if  $r^* = \perp$ ), where `computed_values` correspond to all the Compute queries made by  $\mathcal{A}$  since the last Advance or Recover query and  $F$  is the current state. //If the state is compromised, do nothing.
- On input (Recover):  $\mathcal{O}_{\text{CPRFG}}$  marks the state as recovered. It samples  $k \xleftarrow{\$} \{0, 1\}^\lambda$  uniformly at random and updates its state to a new function  $F \xleftarrow{\$} \{f : R_\lambda \rightarrow \{0, 1\}^{m(\lambda)}\}$  sampled uniformly at random. It runs  $\mathcal{S}(\text{Recover})$ . Reset `computed_values`  $\leftarrow \square$ .

Figure 22: Cascaded PRF-PRG Security Game

**Property 1:** As long as the initial state  $s = (s_a, s_c)$  is random and secret, any sequence of calls to  $KDF.\text{Compute}$  that does not contain a call to  $KDF.\text{Advance}$  should behave like a sequence of calls to a random function. That is, for each new, adversarially chosen value of  $r$ , the key  $k$  generated by running  $KDF.\text{Compute}(s_c, r)$  should look like a fresh random key. Furthermore, two calls with the same input  $r$  should result with the same  $k$ . That is, regardless of the sequence of  $KDF.\text{Compute}$  queries run before the call with input  $r$ , the output key  $k$  should be the same.

(This property is intended to ensure that even though an adversary for  $\Pi_{\text{eKE}}$  may cause a receiver to compute epoch keys using  $r$ 's that don't match the one used by the sender, the adversary cannot actually guess what any of the resulting epoch keys will be. Furthermore, the consistency of the outputs to  $KDF.\text{Compute}$  queries ensures that two parties that have the same initial state  $(s_a, s_c)$  and each make a different sequence of  $KDF.\text{Compute}$  queries, will get the same output key  $k$  for any query  $r$  contained in both sequences.)

This property holds for the following reasons: (1) **Compute** queries in the ideal game (from Figure 22), are answered via a function  $F$  chosen uniformly at random as long as the state is not compromised. (2) The outputs of the randomly chosen function  $F$  on inputs  $r$  are fixed regardless of the order of **Compute** queries (till an **Advance** occurs and a new random function is chosen). (3) Finally, when a call to  $(\text{Compute}, r)$  is answered with a key  $k$ , the tuple  $(r, k)$  is stored till the next **Advance** query to make sure all calls to  $(\text{Compute}, r)$  till then are answered with the same key  $k$ .

**Property 2:** As long as either  $s_a$  or  $r$  is both random and secret, running  $KDF.\text{Advance}(s_a, r)$  updates the state to a new pseudorandom state  $(s_a^*, s_c^*)$ . This in turn causes  $KDF.\text{Compute}$  to behave like a fresh random function. Furthermore, a party that starts with some initial state  $s = (s_a, s_c)$  and runs some sequence of  $KDF.\text{Compute}$  queries that lead to a state  $s' = (s_a, s'_c)$ , will always get the same resulting state  $s^*$  when it queries  $KDF.\text{Advance}$  for some input  $r$ , regardless of the sequence of  $KDF.\text{Compute}$  queries that it made before advancing to lead to the state  $s'$ .

(This property provides forward secrecy of any keys computed before the last time  $KDF.\text{Advance}$  was run. It also allows parties who can agree on a fresh random and secret  $r$  to recover from the corruption. Additionally, the consistency property ensures that even though an adversary for  $\Pi_{\text{eKE}}$  can cause parties to compute keys for several incorrect randomizer values, the adversary still can't make the parties go out of sync without corrupting them.)

In the case where the state is not compromised, the ideal game samples a fresh random function on receiving  $(\text{Advance}, \text{root\_input})$ . On the other hand, the case where the state is compromised running  $KDF.\text{Advance}(s_a, r')$  with a secret random value  $r'$  is indistinguishable from updating the state to a newly chosen random function because the actions for  $(\text{Recover})$  in the real game look indistinguishable from those for  $(\text{Recover})$  in the ideal game. Finally, consistency of the resulting state  $s^*$  with respect to the input  $r$  holds because  $KDF.\text{Advance}$  is deterministic with respect to its inputs  $(s_a, r)$  and queries to  $KDF.\text{Compute}$  don't change  $s_a$ .

**Property 3:** To demonstrate the security properties of a CPRFG against break-ins we would like to show that: (1) Exposing the current state of the KDF does not expose any keys generated before the last **Advance**. (2) At any point during the execution, the current local state of the system is simulatable given only the inputs and outputs of  $KDF.\text{Compute}$  since the last  $KDF.\text{Advance}$ .

(This guarantees that an adversary does not obtain any computational advantage over what is allowed in an ideal execution where the generated keys are truly random.)

This property is modeled in the ideal  $(\text{Compromise}, \text{root\_input}^*)$  operation in which the simulator is given only the inputs and outputs since the last advance and produces a state  $s = (s_a, s_c)$  that is indistinguishable from that seen during compromise in the real world.

### 6.2.3 Constructing the CPRFG

In this section, we provide a construction for the CPRFG in the standard model. The construction requires a puncturable pseudorandom function [15] along with standard crypto primitives.

**CPRFG Construction**

The construction is parameterized by a puncturable PRF  $\{f\}$ , PRF  $g$ , and length-tripling PRG  $h$ .

**KDF.Compute** $(s_c, r)$ :

1. Parse  $s_c = (s_1, s_2)$ , and further parse  $s_1 = (\widehat{s}_1, (a_1, b_1), \dots, (a_k, b_k))$
2. If  $r = a_i$  for some  $i$ : set  $\widehat{f}_{s_1}(\text{root\_input}) = b_i$  and  $s'_c = s_c$ .
3. Else:
  - (a)  $\widehat{s}'_1 \leftarrow \text{Puncture}(\widehat{s}_1, r)$
  - (b) Update  $s'_1 = (\widehat{s}'_1, (a_1, b_1), \dots, (a_k, b_k), (r, f_{\widehat{s}'_1}(r)))$
  - (c) Compute  $\widehat{f}_{s_1}(r) \leftarrow f_{\widehat{s}'_1}(r)$
4. Output  $(s'_c, \widehat{f}_{s_1}(r) \oplus s_2)$

**KDF.Advance** $(s_a, r)$ :

1. Output  $(s_a, s_c) \leftarrow h(g_{s_a}(r))$

Figure 23: CPRFG Construction

**Theorem 12** *Assume that  $\{f\}$  is a puncturable PRF,  $\{g\}$  is a PRP, and  $h$  is a PRG. Then the KDF construction in Figure 23 is a cascaded PRF-PRG.*

**Cascaded PRF-PRG Simulator**

On input (Compute,  $r$ ) run the honest **KDF.Compute** $(s_c, r)$  from Figure 23 and update  $s_c$  accordingly.  
 On input (Advance,  $r$ ) run the honest **KDF.Advance** $(s_a, r)$  from Figure 23 and update  $s = (s_a, s_c)$  accordingly.  
 On input (computed\_values,  $(r, k)$ ) or (computed\_values,  $\perp$ ):

- Sample  $s_1, s_a \xleftarrow{\$} \{0, 1\}^\lambda$  independently and uniformly at random.
- Parse **computed\_values** =  $((a_1, b_1), \dots, (a_k, b_k))$ .
- Let  $\widehat{s}_1$  be the result of puncturing  $s_1$  at points  $a_1, \dots, a_k$ .
- Set  $s'_1 = (\widehat{s}_1, (a_1, b_1), \dots, (a_k, b_k))$ .
- If  $(r, k) \neq \perp$ , set  $s_2 = k \oplus \widehat{f}_{s'_1}(r)$ .
- Finally, set  $s_c = (s'_1, s_2)$  and output  $s = (s_a, s_c)$ .

On input (Recover): delete state  $s$ .

Figure 24: Cascaded PRF-PRG Simulator

**Proof:** Consider the simulator  $\mathcal{S}_{\text{KDF}}$  in Figure 24 (on Page 56). We argue that an adversary  $\mathcal{A}$  that distinguishes between the real and ideal games can be used to break either the fact that  $\{f\}$  is a puncturable PRF (as in [15]), or the fact that  $\{g\}$  is a PRF, or the PRG property of  $h$ .

The proof uses the following hybrid argument. Let the hybrid experiment  $H_i^1$  consist of the ideal CPRFG game up to and including the  $i$ -th instance of (Advance,  $\cdot$ ) or (Recover) in the CPRFG

game; all following calls to the game are real. Next, let  $H_i^2$  consist of the ideal CPRFG game up to the  $i$ -th instance of (Advance,  $\cdot$ ) or (Recover) and the subsequent calls to the CPRFG game before the next advance; only then do all subsequent calls act as in the real game. Specifically, if epoch  $i$  was not corrupted, the  $(i + 1)$ -th (Advance,  $\cdot$ ) will use a random `root_key3` to key  $g$ .

We will begin with the first case. Suppose there is some  $\mathcal{A}$  that can distinguish  $H_i^1$  from  $H_i^2$  (which differ only in whether the operations computed during the  $i$ -th epoch are real or ideal). Then, clearly epoch  $i$  is not compromised because the simulator ensures that the two hybrids are identical (the simulation is perfect). Then, construct the following adversary  $\mathcal{A}_{PPRF}$  against the puncturable PRF property of  $f$ :

1. Emulate the ideal CPRFG game to  $\mathcal{A}$  up to and including the  $i$ -th Advance.
2. Until the next Advance or Compromise:
  - (a) Initialize a list of computed values `computed_values`
  - (b) Let  $\mathcal{K}$  be the set of possible values for  $r$ .
  - (c) For inputs of the form (Compute,  $r$ ), for some  $r$  from  $\mathcal{A}$ :  
Query  

$$k \leftarrow PPRF.Challenge(r)$$
and return  $k$ ; Also, set  $\mathcal{K} = \mathcal{K} \setminus \{r\}$  to puncture this input.
3. If we exited Step 2 with a (Compromise,  $r^*$ ) for some  $r^*$ , query  
 $s_1 \leftarrow PPRF.Constrain(\mathcal{K} \cup \{r^*\})$ , sample  $s'_2$  and  $s_a$  uniformly and set  $s_2 = PPRF.Challenge(r^*) \oplus s'_2$ . Output  $s = ((s_1, s_2), s_a)$ .
4. If we exited Step 2 with a (Advance,  $r$ ) for some  $r$ , sample  $(s_a, s_c) \xleftarrow{\$} \{0, 1\}^\lambda$  uniformly.
5. Emulate the real CPRFG game with the state initialized to  $s$  for the remaining calls from  $\mathcal{A}$ .
6. If  $\mathcal{A}$  outputs " $H_i^1$ ", then output "pseudorandom function". Otherwise ( $\mathcal{A}$  outputs  $H_i^2$ ), output "random function".

Notice that if the function in the PPRF game is pseudorandom, then hybrid  $H_i^1$  is exactly emulated to  $\mathcal{A}$ , because the PPRF key is truly random; otherwise,  $H_i^2$  is perfectly emulated to  $\mathcal{A}$ . Thus,  $\mathcal{A}$ 's distinguishing probability is identical to  $\mathcal{A}_{PRF\_PRG}$ 's.

Now, consider the case that  $\mathcal{A}$  can distinguish between hybrids  $H_{i-1}^2$  and  $H_i^1$  (which differ only on whether the  $i$ -th Advance is real or ideal). Clearly then epoch  $i$  is not corrupt because the simulator ensures that the two hybrids are identical (the simulation is perfect) Then, we will construct the following PRF-PRG adversary  $\mathcal{A}_{PRF\_PRG}$ :

1. Emulate the ideal CPRFG game to  $\mathcal{A}$  up to the  $i$ -th Advance.
2. On the  $i$ -th (Advance,  $\cdot$ ) or Recover:
  - (a) If (Advance,  $r$ ) for some  $r$ :  $(s_a, s_c) \leftarrow h(PRF.Challenge(r))$ .
  - (b) Else (Recover): the PRF key  $k$  is known to  $\mathcal{A}$ , but the  $r$  is uniformly random. Set  
 $k \leftarrow PRG.Challenge()$ .
3. For the remaining calls from  $\mathcal{A}$ , emulate the real CPRFG game with state initialized to  
 $s = (s_a, s_c)$ .



4. If  $\mathcal{A}$  outputs “ $H_{i-1}^2$ ”, then output “pseudorandom”. Otherwise ( $\mathcal{A}$  outputs “ $H_i^1$ ”), output “truly random”.

Notice that if the functions in the PRF or PRG games are pseudorandom, then we have perfectly emulated hybrid  $H_{i-1}^2$  to  $\mathcal{A}$ ; otherwise, if the functions in the PRF or PRG games are random, we have perfectly emulated hybrid  $H_i^1$ . Thus, the distinguishing advantage of  $\mathcal{A}$  is the same as that of  $\mathcal{A}_{PRF\_PRG}$ .

Lastly, since the CPRFG adversary  $\mathcal{A}$  runs in polynomial time in the security parameter, there can be at most a polynomial number of **Advance** and **Recover** events—and thus, hybrids. Overall,  $\mathcal{A}$ ’s distinguishing advantage over all hybrids remains negligible.  $\square$

Observe that the size of the state of above CPRFG grows roughly linearly with the number of applications of **KDF.Compute** between two consecutive applications of **KDF.Advance**, which may in principle lend to a denial of service attack on the protocol. However, we argue that a linear growth in space is unavoidable in plain model constructions – not only for the notion of CPRFG, but also for realizing  $\mathcal{F}_{eKE}$  in the presence of security against adaptive corruptions [50]. Furthermore, it may be reasonable to mitigate such denial of service attacks by imposing an a priori bound on the number of failed attempts to advance an epoch before the protocol raises an alarm to the user.

### 6.3 Security Analysis of $\Pi_{eKE}$

**Proof:**The proof of Theorem 3 proceeds in two parts. The first part is the construction of the simulator  $\mathcal{S}_{eKE}$  and the internal code  $\mathcal{I}_{eKE}$ . The second part of the proof is a hybrid argument used to show the following indistinguishability (for the previously constructed  $\mathcal{S}_{eKE}, \mathcal{I}_{eKE}$ ):

$$(\mathcal{F}_{lib}, \mathcal{F}_{DIR}, \mathcal{F}_{LTM}, \mathcal{F}_{eKE}^*, \mathcal{S}_{eKE}, \mathcal{I}_{eKE}) \approx (\mathcal{F}_{lib}, \mathcal{F}_{DIR}, \mathcal{F}_{LTM}, \Pi_{eKE}).$$

When neither party has ever been corrupted, the functionality never calls the simulator. Instead, the only adversarial choices are made when the functionality runs the internal code  $\mathcal{I}_{eKE}$  to receive new epoch ids. The internal adversarial code for this proof simply samples  $(\text{epoch\_key}, \text{epoch\_id})$  pairs uniformly at random for each new epoch. The detailed version of  $\mathcal{I}_{eKE}$  can be found in Fig. 25 on Page 59.

Once a party has been corrupted, the functionality will begin to call the simulator  $\mathcal{S}_{eKE}$ . (The detailed simulator  $\mathcal{S}_{eKE}$  can be found in Figs. 26 to 27 on Page 60.) The simulator is first called by the functionality at the time of corruption. At this time, the simulator must return a simulated state **state** for the corrupted party which will be passed to the environment by the functionality. The receipt of party state **state** simulates a passive cloning attack. Then, during the healing period when the parties are still compromised from the state corruption, the simulator must perform two actions: (1) provide **sending\_chain\_keys** and **recv\_chain\_keys** for the parties to  $\mathcal{F}_{eKE}$  until compromise ends (maybe never if the adversary chooses to person-in-the-middle the parties forever!), and (2) continue to return snapshots of the current party’s state on new corruptions in a way that seems consistent with the previously produced values, (even if the party is still compromised from a previous corruption.)

To understand how the simulator goes about this, imagine that at corruption, the simulator creates dummy parties  $\mathcal{P}_0, \mathcal{P}_1$ . These dummy parties run the code of  $\Pi_{eKE}$  based on the state that the simulator provides them. If the simulator can produce a ‘convincing state’ for the corrupted party at the time of corruption, then simply running the parties as in the honest protocol and updating them based on the inputs provided by the environment will look indistinguishable from

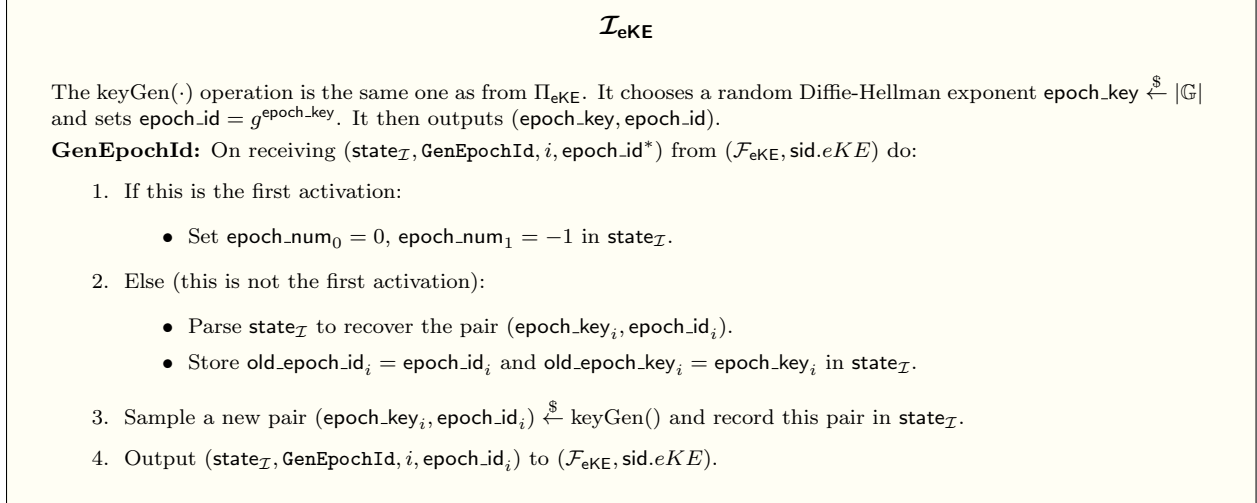


Figure 25: Internal adversarial code  $\mathcal{I}_{eKE}$

how an honest corruption would go. To flesh out this intuition, let's start by discussing how the simulator will handle a **ReportState** request from  $\Pi_{eKE}$ .

**ReportState:** When the functionality  $\mathcal{F}_{eKE}$  receives a corruption notification from its calling protocol, it will send a request of the form  $(\text{ReportState}, i, \text{recv\_chain\_key})$  to the simulator. The simulator must return 'the state of  $\text{pid}_i$ '; this state will be returned to the calling protocol and in turn to the adversary.

The only variables stored in a party's state in  $\Pi_{eKE}$  are  $\text{epoch\_key}_i, \text{epoch\_id}_i, \text{epoch\_id}_{1-i}, s = s_a || s_c$ .

Note that  $\text{epoch\_key}_i, \text{epoch\_id}_i$ , and  $\text{epoch\_id}_{1-i}$  were chosen by the adversarial code  $\mathcal{I}_{eKE}$  just like they would be in the real protocol and can be provided as-is. So, the only remaining question is what value  $s = s_a || s_c$  should the simulator provide. Remember that any sending or receiving chain key not yet provided by  $\mathcal{F}_{eKE}$  will be chosen by the simulator in the future using these dummy parties who simply run the honest protocol on this provided state. So, the  $s_a || s_c$  produced here only needs to take account past chain keys. Fortunately, because  $KDF$  is a cascaded PRF-PRG, the current  $s_a || s_c$  in the state of a party will be unrelated to all previous keys provided by  $\mathcal{F}_{eKE}$  in most cases and can simply be chosen at random. The only case where  $s$  is related to a previous output of  $\mathcal{F}_{eKE}$  is when the instance of protocol  $\Pi_{SGNL}$  for  $\text{pid}_{1-i}$  has already started a sending epoch whose chain key  $\text{recv\_chain\_key}$  has not yet been retrieved from  $\mathcal{F}_{eKE}$  by the receiving party's instance  $(\Pi_{SGNL}, \text{sid}, \text{pid}_i)$ . In this case, the provided  $(s_a || s_c)$  must satisfy  $KDF.\text{Compute}(s_c, \text{epoch\_id}_{1-i}^{\text{epoch\_key}_i}) = -, \text{recv\_chain\_key}$ . In these cases only, the functionality will provide  $\text{recv\_chain\_key}$  to the simulator at the time of making the request. On receiving a request containing  $\text{recv\_chain\_key}$ , the simulator will invoke the KDF simulator  $\mathcal{S}_{KDF}$  that proves the security of  $KDF$ ; since the  $KDF$  protocol is a secure CPRFG, such a simulator  $\mathcal{S}_{KDF}$  must exist.

After  $\text{pid}_i$  is corrupted, the simulator will be able to provide all the keys for the parties by running the instructions for  $\Pi_{eKE}$  within the dummy parties and using the  $\text{epoch\_id}^*$  provided in the **GenEpochId** requests to know how to ratchet forward to the receiving epochs for each party. Once the functionality 'heals' from the corruption, it will stop asking the simulator for keys unless the parties have diverged. If divergence doesn't occur, the simulator will end the execution of the

## Simulator $\mathcal{S}_{eKE}$ for realizing $\mathcal{F}_{eKE}$

$\mathcal{S}_{eKE}$  runs only if the parties are diverged, or if the current epoch is compromised (i.e. we're within the quarantine period for a corruption.)

The  $\text{keyGen}(\cdot)$  and  $KDF$  components are the same ones from  $\Pi_{eKE}$ : (1)  $\text{keyGen}$  chooses a random Diffie-Hellman pair ( $\text{epoch\_key}$ ,  $\text{epoch\_id}$ ). (2) The protocol  $KDF$  is a CPRFG (Cascaded PRF-PRG).

**At first activation:** Send  $(\mathcal{F}_{eKE}, \mathcal{I}_{eKE})$  to  $\mathcal{F}_{lib}$ .

*//GetSendingKey and GetReceivingKey are used as subroutines to answer ReportState and GenEpochId requests.*

**GetSendingKey:** On receiving  $(\text{GetSendingKey}, i)$  from  $(\mathcal{F}_{eKE}, \text{sid.eKE})$  do:

1. If  $\text{View}_i.\text{sending\_chain\_key}$  exists:
  - Output  $(\text{GetSendingKey}, i, \text{View}_i.\text{sending\_chain\_key})$  and delete  $\text{View}_i.\text{sending\_chain\_key}$ .
2. Otherwise end the activation.

**GetReceivingKey:** On receiving  $(\text{GetReceivingKey}, i, \text{epoch\_id})$  from  $(\mathcal{F}_{eKE}, \text{sid.eKE})$  do:

1. Set  $\text{View}_i.\text{temp\_epoch\_id}_{\text{partner}} = \text{epoch\_id}$
2. Compute  $\text{View}_i.\text{root\_input} = \text{Exp}(\text{View}_i.\text{temp\_epoch\_id}_{\text{partner}}, \text{View}_i.\text{epoch\_key}_{\text{self}})$ .
3. If  $\text{diverge\_parties} = \text{false}$ :
  - If  $\text{epoch\_id} \neq \text{View}_{1-i}.\text{epoch\_id}_{\text{self}}$ : Add  $\text{epoch\_num}_{1-i} + 1$ ,  $\text{epoch\_num}_{1-i} + 2$  to  $\text{compromised\_epochs}$ .
  - Run the KDF simulator  $\mathcal{S}_{KDF}(\text{Compute}, \text{View}_i.\text{root\_input})$  and set the output as  $\text{View}_i.\text{temp\_recv\_chain\_key}$ .
4. Else ( $\text{diverge\_parties} = \text{true}$ ):
  - Set  $\text{View}_i.\text{temp\_recv\_chain\_key} = KDF.\text{Compute}(\text{View}_i.\text{root\_key}, \text{View}_i.\text{root\_input})$ .
5. Erase  $\text{View}_i.\text{root\_input}$  and  $\text{View}_i.\text{temp\_epoch\_id}_{\text{partner}}$ .
6. Output  $(\text{GetReceivingKey}, i, \text{epoch\_id}, \text{View}_i.\text{temp\_recv\_chain\_key})$  to  $(\mathcal{F}_{eKE}, \text{sid.eKE})$  and delete  $\text{View}_i.\text{temp\_recv\_chain\_key}$ .

**ReportState:** On receiving  $(\text{ReportState}, \text{state}_{\mathcal{I}}, \text{pid}_i, \text{recv\_chain\_key}^*, \text{leakage})$  from  $(\mathcal{F}_{eKE}, \text{sid.eKE})$  do:

*//This method is run every time  $\mathcal{F}_{eKE}$  is informed that a party has been corrupted.*

*//recv\_chain\_key\* =  $\perp$  if and only if  $\text{epoch\_num}_{1-i} < \text{epoch\_num}_i$ .*

1. Add  $\text{epoch\_num}_i, \text{epoch\_num}_i + 1, \text{epoch\_num}_i + 2, \text{epoch\_num}_i + 3$  to the list  $\text{compromised\_epochs}$  in  $\text{state}_{\mathcal{I}}$ .
2. If  $\text{View}_0, \text{View}_1$  already exist, output  $(\text{ReportState}, \text{pid}_i, \text{View}_i)$  to  $(\mathcal{F}_{eKE}, \text{sid.eKE})$ . Otherwise, continue to create them.
3. Create two dictionaries  $\text{View}_0, \text{View}_1$ .  
*//This and following steps will run only if the parties weren't compromised when ReportState was called.*
4. From  $\text{state}_{\mathcal{I}}$  get the variables  $\text{epoch\_id}_0, \text{epoch\_id}_1, \text{epoch\_key}_0, \text{epoch\_key}_1, \text{old\_epoch\_id}_0, \text{old\_epoch\_id}_1$ .
5. Set the following values in the view objects  $\text{View}_0, \text{View}_1$ :
  - (a) Set  $\text{View}_0.\text{epoch\_id}_{\text{self}} = \text{epoch\_id}_0, \text{View}_0.\text{epoch\_key}_{\text{self}} = \text{epoch\_key}_0$ .
  - (b) Set  $\text{View}_1.\text{epoch\_id}_{\text{self}} = \text{epoch\_id}_1, \text{View}_1.\text{epoch\_key}_{\text{self}} = \text{epoch\_key}_1$ .
  - (c) Let  $j \in \{0, 1\}$  be such that  $\text{pid}_j$  started the latest sending epoch. (i.e  $\text{epoch\_num}_j > \text{epoch\_num}_{1-j}$ .)
    - i. Set  $\text{View}_j.\text{epoch\_id}_{\text{partner}} = \text{epoch\_id}_{1-j}$ .
    - ii. Set  $\text{View}_{1-j}.\text{epoch\_id}_{\text{partner}} = \text{old\_epoch\_id}_j$ .
    - iii. Let  $\text{root\_input}^* = \text{Exp}(\text{epoch\_id}_j, \text{epoch\_key}_{1-j})$ .
    - iv. Run  $\mathcal{S}_{KDF}(\text{Compromise}, \text{leakage}, (\text{root\_input}^*, \text{recv\_chain\_key}^*))$  to get an output  $\text{root\_key}$ .
    - v. Set  $\text{View}_{1-j}.\text{root\_key} = \text{root\_key}$ .
    - vi. If  $\text{recv\_chain\_key}^* \neq \perp$ , set  $\text{View}_j.\text{sending\_chain\_key} = \text{recv\_chain\_key}^*$ .
    - vii. Set  $\text{View}_j.\text{root\_key} = KDF.\text{Advance}(\text{root\_key}, \text{root\_input}^*)$ . *//Only the receiver tells  $\mathcal{S}_{KDF}$  to advance. Compute sender's root key locally.*
6. Output  $(\text{ReportState}, \text{pid}_i, \text{View}_i = \{\text{epoch\_key}_{\text{self}}, \text{epoch\_id}_{\text{self}}, \text{epoch\_id}_{\text{partner}}, \text{root\_key}\})$  to  $(\mathcal{F}_{eKE}, \text{sid.eKE})$ .

(The rest of this simulator is in Fig. 27 on Page 61)

Figure 26: Simulator  $\mathcal{S}_{eKE}$  for realizing  $\mathcal{F}_{eKE}$

### Simulator $\mathcal{S}_{eKE}$ continued...

(This simulator begins in Fig. 26 on Page 60)

**GenEpochId:** On receiving  $(\text{GenEpochId}, i, \text{epoch\_id}^*)$  from  $(\mathcal{F}_{eKE}, \text{sid.eKE})$  do:

1. Set  $\text{epoch\_num}_i += 2$ .
2. Set  $\text{View}_i.\text{epoch\_id}_{\text{partner}} = \text{epoch\_id}^*$ .
3. Compute  $\text{root\_input} = \text{Exp}(\text{epoch\_id}^*, \text{View}_i.\text{epoch\_key}_{\text{self}})$ .
4. If  $\text{diverge\_parties} = \text{false}$ :
  - (a) Run  $\mathcal{S}_{KDF}(\text{Advance}, \text{root\_input})$ . //The receiver will tell  $\mathcal{S}_{KDF}$  to advance its epoch.  
//Next, we check if the states have diverged.
  - (b) Compute  $s = \text{KDF}.\text{Advance}(\text{View}_i.\text{root\_key}, \text{root\_input})$  and set  $\text{View}_i.\text{root\_key} = s$ .
  - (c) If  $\text{View}_i.\text{root\_key} \neq \text{View}_{1-i}.\text{root\_key}$  then set  $\text{diverge\_parties} = \text{true}$ . //divergence occurs here
5. Else ( $\text{diverge\_parties} = \text{true}$ ):
  - Set  $\text{View}_i.\text{root\_key} = \text{KDF}.\text{Advance}(\text{View}_i.\text{root\_key}, \text{root\_input})$ .
6. If  $\text{epoch\_num}_i \notin \text{compromised\_epochs}$  and  $\text{diverge\_parties} \neq \text{true}$ :
 

//If the views are in sync and  $\text{epoch\_num}_i + 1$  is the first uncompromised epoch after corruption.  
//Delete the view objects and generate a new epoch id.

  - (a) Sample  $(\text{View}_i.\text{epoch\_key}_{\text{self}}, \text{View}_i.\text{epoch\_id}_{\text{self}}) \xleftarrow{\$} \text{keyGen}()$ .
  - (b) Let  $j \in \{0, 1\}$  be such that  $\text{pid}_j$  started the latest sending epoch. (i.e  $\text{epoch\_num}_j > \text{epoch\_num}_{1-j}$ .)
  - (c) Set the following values in  $\text{state}_{\mathcal{T}}$ :
    - i.  $\text{epoch\_id}_j = \text{View}_j.\text{epoch\_id}_{\text{self}}$  and  $\text{epoch\_key}_j = \text{View}_j.\text{epoch\_key}_{\text{self}}$ .
    - ii.  $\text{epoch\_id}_{1-j} = \text{View}_{1-j}.\text{epoch\_id}_{\text{self}}$  and  $\text{epoch\_key}_{1-j} = \text{View}_{1-j}.\text{epoch\_key}_{\text{self}}$ .
    - iii.  $\text{old\_epoch\_id}_j = \text{View}_{1-j}.\text{epoch\_id}_{\text{partner}}$  and  $\text{old\_epoch\_id}_{1-j} = \perp$ .
  - (d) Delete  $\text{View}_0, \text{View}_1$  and output  $(\text{GenEpochId}, \text{state}_{\mathcal{T}}, i, \text{epoch\_id}_i)$  to  $(\mathcal{F}_{eKE}, \text{sid.eKE})$ .
7. Otherwise ( $\text{epoch\_num}_i \in \text{compromised\_epochs}$  or  $\text{diverge\_parties} = \text{true}$ ):
 

//If the parties are still compromised or the views have diverged.  
//Continue to update  $\text{View}_i$  according to  $\Pi_{eKE}$ .

  - (a) Sample  $(\text{View}_i.\text{epoch\_key}_{\text{self}}, \text{View}_i.\text{epoch\_id}_{\text{self}}) \xleftarrow{\$} \text{keyGen}()$ .
  - (b) Compute  $\text{root\_input} = \text{Exp}(\text{epoch\_id}^*, \text{View}_i.\text{epoch\_key}_{\text{self}})$ .
  - (c) If  $\text{diverge\_parties} = \text{false}$ , run  $\mathcal{S}_{KDF}(\text{Compute}, \text{root\_input})$  and set  $\text{View}_i.\text{sending\_chain\_key}$  to be its output.
  - (d) Else ( $\text{diverge\_parties} = \text{true}$ ), set  $\text{View}_i.\text{sending\_chain\_key} = \text{KDF}.\text{Compute}(\text{View}_i.\text{root\_key}, \text{root\_input})$ .
  - (e) Set  $\text{View}_i.\text{root\_key} = \text{KDF}.\text{Advance}(\text{View}_i.\text{root\_key}, \text{root\_input})$ .
  - (f) Erase  $\text{root\_input}$ .
  - (g) Output  $(\text{GenEpochId}, i, \text{View}_i.\text{epoch\_id}_{\text{self}})$  to  $(\mathcal{F}_{eKE}, \text{sid.eKE})$ .

Figure 27: Simulator  $\mathcal{S}_{eKE}$  (continued)

dummy parties. However, in the case of divergence, it must continue to run the dummy parties to provide chain keys.

### Analysis of $\mathcal{S}_{\text{eKE}}$ .

To demonstrate the validity of  $\mathcal{S}_{\text{eKE}}, \mathcal{I}_{\text{eKE}}$ , One must prove that that the ideal functionality  $\mathcal{F}_{\text{eKE}}$  with access to the simulator  $\mathcal{S}_{\text{eKE}}$  and internal code  $\mathcal{I}_{\text{eKE}}$  is indistinguishable from the real world. This indistinguishability is proven via the following hybrid executions  $H_0, \dots, H_5$  that bridge the gap between a real-world execution  $H_0$  (namely an execution with  $\Pi_{\text{eKE}}, \mathcal{F}_{\text{DIR}}, \mathcal{F}_{\text{LTM}}$ ) and an ideal execution  $H_5$  (namely an execution with  $\mathcal{F}_{\text{eKE}}, \mathcal{S}_{\text{eKE}}, \mathcal{F}_{\text{DIR}}, \mathcal{F}_{\text{LTM}}$ ). The full hybrid argument follows this brief description of each hybrid:

- **Hybrid  $H_0$  (Ideal World):** This is the ideal world. Here you can think of  $\mathcal{S}_{\text{eKE}} + \mathcal{I}_{\text{eKE}} + \mathcal{F}_{\text{eKE}}$  all together as one combined unit with combined state.
- **Hybrid  $H_1$ :** In this hybrid the main change is to the method of choosing chain keys *when both parties are uncompromised*. In hybrid  $H_0$  these keys are chosen uniformly at random from  $\mathcal{K}_{ep}$  while in this hybrid they're chosen via a function  $F : \mathbb{G} \rightarrow \mathcal{K}_{ep}$  chosen uniformly at random for the corresponding epoch. the inputs to the random function are the values `root.input` chosen uniformly at random from the group  $\mathbb{G}$ (as in the CPRFG game). The adversary's view in this hybrid is identically distributed as in the previous hybrid.
- **Hybrid  $H_2$ :** This hybrid again changes the method of choosing chain keys *when both parties are uncompromised*. In this case,  $H_2$  transitions from using a version of the ideal CPRFG game to using the protocol *KDF*. In particular, the *KDF.Compute* and *KDF.Advance* methods of the *KDF* are now used to compute the chain keys and update the value  $s = s_a || s_c$  respectively instead of choosing a new random function for each epoch. The root inputs and initial root key are chosen identically to the previous hybrid. This hybrid is indistinguishable to  $H_1$  since *KDF* is a secure CascadedPRF-PRNG.
- **Hybrid  $H_3$ :** This hybrid again changes the method of choosing chain keys *when both parties are uncompromised*. In particular, we replace the random choices of `root.inputs` with the Diffie-Hellman keys computed as specified in the protocol  $\Pi_{\text{eKE}}$  over a series of hybrids. Let  $M = \text{max\_epochnum}$ , then the hybrids are denoted by  $H_3^1, \dots, H_3^M$  where  $H_3^k$  is the hybrid where the root inputs for the first  $k$  epochs are all using the specification of  $\Pi_{\text{eKE}}$ . The hybrid  $H_3^M$  is computationally indistinguishable from  $H_2$  based on the *DDH* assumption on the group  $\mathbb{G}$ .
- **Hybrid  $H_4$ :** In this hybrid the random choice of initial root key is replaced by the value computed using  $\mathcal{F}_{\text{LTM}}$ . The indistinguishability of this hybrid is also by reduction to the DDH assumption on group  $\mathbb{G}$  since  $\mathcal{F}_{\text{LTM}}$  computes the value  $s = s_a || s_c$  using the the X3DH protocol.
- **Hybrid  $H_5$ : (Real World)** This hybrid is the real world. The only real difference between this hybrid and the previous one is the fact that the parties run separate instances of *KDF* instead of a joint instance. This hybrid is indistinguishable from the previous one because: (1) *KDF* consists of deterministic functions, (2) the two parties run *KDF* on identical inputs to the ones they use in  $H_4$ .

## Note

In hybrids  $H_0$ ,  $H_1$ , and  $H_2$  we display the lines of the simulator and functionality that are relevant to the current hybrid. This note explains the colors used in the display boxes; **black text existed in the previous hybrid and does not change in the current hybrid**; **brown text was added in the current hybrid**; struck out grey text was erased in the current hybrid.

### 6.3.1 Hybrid $H_0$ (Ideal World)

This is the combination of the ideal functionality  $\mathcal{F}_{\text{eKE}}$ , the internal code  $\mathcal{I}_{\text{eKE}}$ , and the simulator  $\mathcal{S}_{\text{eKE}}$  all in the presence of  $\mathcal{F}_{\text{DIR}}$ ,  $\mathcal{F}_{\text{LTM}}$ . Here you can consider  $\mathcal{S}_{\text{eKE}} + \mathcal{I}_{\text{eKE}} + \mathcal{F}_{\text{eKE}}$  all together as one combined unit (with combined state) since this combined unit must be proven indistinguishable from the real protocol  $\Pi_{\text{eKE}}$ . In other words, this hybrid is exactly the ideal world, we are simply considering the code of  $\mathcal{S}_{\text{eKE}}$  and  $\mathcal{I}_{\text{eKE}}$  to be directly inlined into the functionality  $\mathcal{F}_{\text{eKE}}$  as shown below:

#### $\mathcal{S}_{\text{eKE}}$ (Hybrid 0)

~~At first activation: Send  $(\mathcal{F}_{\text{eKE}}, \mathcal{I}_{\text{eKE}})$  to  $\mathcal{F}_{\text{lib}}$ .~~

#### $\mathcal{F}_{\text{eKE}}$ (Hybrid 0)

##### ConfirmReceivingEpoch:

1. (e) Initialize  $\text{state}_{\mathcal{I}} = \perp$ , call  $\mathcal{F}_{\text{lib}}$  with input “eKE” to obtain internal adversarial code  $\mathcal{I}_{\text{eKE}}$ .
2. (b) **If  $\text{epoch\_id}^* \neq \text{epoch\_id\_self}_{1-i}$ : Set  $\text{diverge\_parties} = \text{true}$ , run step 3 of **Corrupt** to set  $\text{recv\_chain\_key}^*$  and leakage. Run  $\mathcal{S}_{\text{eKE}}(\text{ReportState}, \text{state}_{\mathcal{I}}, i, \text{recv\_chain\_key}^*, \text{leakage})$  and discard the output.** ~~and send  $(\text{ReportState}, \text{state}_{\mathcal{I}}, i, \text{recv\_chain\_key}^*, \text{leakage})$  to the adversary. On receiving a response, continue.~~
3. **If  $\text{epoch\_num}_i + 2 \in \text{compromised\_epochs}$  or  $\text{diverge\_parties}$ : Run  $\mathcal{S}_{\text{eKE}}(\text{GenEpochId}, i, \text{epoch\_id}^*)$ . Send backdoor message  $(\text{GenEpochId}, i, \text{epoch\_id}^*)$  to the adversary. Else run  $\mathcal{I}(\text{state}_{\mathcal{I}}, \text{GenEpochId}, i, \text{epoch\_id}^*)$ .**

##### GetSendingKey:

4. **If  $\text{diverge\_parties} = \text{true}$ , or  $\text{epoch\_num}_i \in \text{compromised\_epochs}$  run  $\mathcal{S}_{\text{eKE}}(\text{GetSendingKey}, i)$ , on producing output  $(\text{GetSendingKey}, i, K_{\text{send}})$ , set  $\text{sending\_chain\_key}_i = K_{\text{send}}$ .** ~~send backdoor message  $(\text{GetSendingKey}, i)$  to the adversary; on receiving backdoor message  $(\text{GetSendingKey}, i, K_{\text{send}})$  from  $\mathcal{A}$  set  $\text{sending\_chain\_key}_i = K_{\text{send}}$ .~~

##### GetReceivingKey:

5. (a) **If  $\text{epoch\_id} \neq \text{epoch\_id\_self}_{1-i}$ , add  $\text{epoch\_num}_i + 2$  to  $\text{compromised\_epochs}$ .**
- (b) **Run  $\mathcal{S}_{\text{eKE}}(\text{GetReceivingKey}, i, \text{epoch\_id})$ . On producing output  $(\text{GetReceivingKey}, i, \text{epoch\_id}, \text{recv\_chain\_key}^*)$ , output  $(\text{GetReceivingKey}, \text{recv\_chain\_key}^*)$ .** ~~Send  $(\text{GetReceivingKey}, i, \text{epoch\_id})$  to the adversary.~~
- (c) **Upon receiving  $(\text{GetReceivingKey}, i, \text{epoch\_id}, \text{recv\_chain\_key}^*)$  from  $\mathcal{A}$ , output  $(\text{GetReceivingKey}, \text{recv\_chain\_key}^*)$ .**

##### Corrupt:

4. **Run  $\mathcal{S}_{\text{eKE}}(\text{ReportState}, \text{state}_{\mathcal{I}}, i, \text{recv\_chain\_key}^*, \text{leakage})$ . On producing the output  $(\text{ReportState}, i, \text{View}_i)$  output  $(\text{Corrupt}, \text{View}_i)$ .** ~~Send  $(\text{ReportState}, \text{state}_{\mathcal{I}}, i, \text{recv\_chain\_key}^*, \text{leakage})$  to the adversary.~~
5. **Upon receiving  $(\text{ReportState}, i, \text{View}_i)$  from  $\mathcal{A}$ , output  $(\text{Corrupt}, \text{View}_i)$  to  $(\Pi_{\text{SGL}}, \text{sid}, \text{pid}_i)$ .**

### 6.3.2 Hybrid $H_1$

The changes made in this step are very important for setting up the next few hybrids. The two main changes between this hybrid and the previous one are:

1. This hybrid chooses the initial root key completely at random in Step 1d of  $\mathcal{F}_{eKE}$ , ignoring the output of  $\mathcal{F}_{LTM}$ .
2. When parties are uncompromised, this hybrid chooses root inputs randomly from  $G$  and runs those through a function  $F : \mathbb{G} \rightarrow \mathcal{K}_{ep}$  (chosen uniformly at random for the corresponding epoch.)

Note the exception in the choice of `root_input` that happens for the very first sending chain key chosen after the end of a compromise. The random choice of `root_input` in this case corresponds to the `Recover` interface in the ‘real execution’ of the CPRFG security game. The properties provided to the `root_input` via the DDH assumption on group  $\mathbb{G}$  are only required for secure recovery. In a later hybrid, this random choice of `root_input` in the case of recovery is replaced with the computation of a Diffie-Hellman key as in  $\Pi_{eKE}$  via a reduction to the DDH assumption on group  $\mathbb{G}$ .

Note that the view of the environment in this hybrid is identically distributed to its view in the previous hybrid: The first change doesn’t impact the view of the adversary at all in this hybrid since the initial root key isn’t used anywhere in the rest of this hybrid. Moreover, from the security of the X3DH protocol, the initial root key computed by  $\mathcal{F}_{LTM}$  is indistinguishable from a randomly chosen initial root key. This will be important in hybrid  $H_3$  where a CPRFG protocol  $KDF$  is used to choose the chainkeys.

The edits to the code of  $\mathcal{S}_{eKE} + \mathcal{I}_{eKE} + \mathcal{F}_{eKE}$  are presented in detail below. In particular, the only edits in this hybrid are to the code of  $\mathcal{F}_{eKE}$ . Recall that this functionality can now access the internal states of  $\mathcal{I}_{eKE}$  and  $\mathcal{S}_{eKE}$ ; the values `epoch_idi`, `epoch_keyi` referred to below belong to  $\mathcal{I}_{eKE}$ ’s state. These values are used to compute `root_input` for every epoch except when recovering from compromise.



## $\mathcal{F}_{\text{eKE}}$ (Hybrid 1)

### ConfirmReceivingEpoch:

1. (d) On receiving (ComputeSendingRootKey,  $k, \text{ek}^{\text{pk}}$ ) from  $\mathcal{F}_{\text{LTM}}$ , sample a value  $s_a || s_c \xleftarrow{\$} \{0, 1\}^n$ . ~~continue.~~
4. (c) Choose a new random function  $F : \mathbb{G} \rightarrow \mathcal{K}_{ep}$ .  
(d) Output (ConfirmReceivingEpoch, epoch\_id\_self <sub>$i$</sub> ) to ( $\Pi_{\text{SGNL}}$ , sid, pid <sub>$i$</sub> ).

### GetSendingKey:

3. **If** there is a tuple (epoch\_id <sub>$i$</sub> , epoch\_id <sub>$i-1$</sub> , chain\_key) in the list computed\_values, set sending\_chain\_key <sub>$i$</sub>  = chain\_key. **Else** :  
Sample sending\_chain\_key <sub>$i$</sub>   $\xleftarrow{\$} \mathcal{K}_{ep}$  from the key distribution.
  - (a) **If** epoch\_num <sub>$i$</sub>  - 1  $\notin$  compromised\_epochs compute root\_input = Exp(epoch\_id <sub>$i-1$</sub> , epoch\_key <sub>$i$</sub> ). **Else** choose a random value root\_input  $\xleftarrow{\$} \mathbb{G}$  and store the tuple (epoch\_num <sub>$i$</sub> , root\_input).
  - (b) Compute sending\_chain\_key <sub>$i$</sub>  =  $F(\text{root\_input})$  and store the tuple (epoch\_id <sub>$i$</sub> , epoch\_id <sub>$i-1$</sub> , sending\_chain\_key <sub>$i$</sub> ) in the list computed\_values.

### GetReceivingKey:

6. (a) **If** there is a tuple  $(\text{epoch\_id}, \text{epoch\_id}_i, \text{chain\_key})$  in the list `computed_values`, set  $\text{recv\_chain\_key}_i = \text{chain\_key}$ .  
 $\text{Sample } \text{recv\_chain\_key}_i \xleftarrow{\$} \mathcal{K}_{ep}$ .
- (b) **Else:**
  - i. Compute  $\text{root\_input} = \text{Exp}(\text{epoch\_id}, \text{epoch\_key}_i)$ .
  - ii. Compute  $\text{recv\_chain\_key}_i = F(\text{root\_input})$  and store the tuple  $(\text{epoch\_id}, \text{epoch\_id}_i, \text{recv\_chain\_key}_i)$  in the list `computed_values`.

Note that the changes here closely map to the ideal interface in the CPRFG security game when  $b = 1$ . (In the next hybrid, a secure CPRFG  $KDF$  will be used instead.)

**Lemma 13** *The view of the environment in  $H_0$  is identically distributed to that in hybrid  $H_1$ .*

**Proof:** As mentioned earlier, the first change doesn't impact the view of the adversary at all in this hybrid since the initial root key isn't used anywhere in the rest of this hybrid. Additionally, the outputs of a function  $F : \mathbb{G} \rightarrow \mathcal{K}_{ep}$  chosen uniformly at random are identically distribution to  $\text{chain\_key} \xleftarrow{\$} \mathcal{K}_{ep}$  chosen uniformly at random, this remains true both when `root_input` is sampled randomly and when `root_input` is known.  $\square$

### 6.3.3 Hybrid $H_2$ (This hybrid uses a CPRFG protocol $KDF$ to choose the output keys.)

This hybrid transitions from using a version of the ideal CPRFG game +  $\mathcal{S}_{KDF}$  to using the actual protocol. In particular, the  $KDF.\text{Compute}$  and  $KDF.\text{Advance}$  methods of the  $KDF$  are now used to compute the chain keys and update the value  $s = s_a || s_c$  for each party instead functions chosen uniformly at random for each epoch.

The detailed edits to the version of  $\mathcal{S}_{eKE} + \mathcal{I}_{eKE} + \mathcal{F}_{eKE}$  from the previous hybrid are presented below. The edits in the first box pertain to the same lines of code in  $\mathcal{F}_{eKE}$  as were edited in the previous hybrid. Following that are the edits to  $\mathcal{S}_{eKE}$  which replace the use of  $\mathcal{S}_{KDF}$  with the use of the protocol  $KDF$ . As before, the values  $\text{epoch\_id}_i, \text{epoch\_key}_i$  for  $i \in \{0, 1\}$  belong to  $\mathcal{I}_{eKE}$ 's state. These values are used to compute `root_input` for every epoch except for during recovery from compromise.

#### $\mathcal{F}_{eKE}$ (Hybrid 2)

##### ConfirmReceivingEpoch:

1. (d) On receiving  $(\text{ComputeSendingRootKey}, k, \text{ek}^{\text{pk}})$  from  $\mathcal{F}_{LTM}$ , sample a value  $s_a || s_c \xleftarrow{\$} \{0, 1\}^n$ .
4. (c) **If**  $\text{epoch\_num}_i - 2 \notin \text{compromised\_epochs}$ , compute  $\text{root\_input} = \text{Exp}(\text{epoch\_id}_{1-i}, \text{epoch\_key}_i)$ . **Else**, get `root_input` from the tuple  $(\text{epoch\_num}_i, \text{root\_input})$ .  
 $\text{Choose a new random function } F : \mathbb{G} \rightarrow \mathcal{K}_{ep}$ .
- (d) Run  $KDF.\text{Advance}(s_a, \text{root\_input}) = s'$ , and change state  $s = s'$ .

##### GetSendingKey:

3. **If** there is a tuple  $(\text{epoch\_id}_i, \text{epoch\_id}_{i-1}, \text{chain\_key})$  in the list `computed_values`, set  $\text{sending\_chain\_key}_i = \text{chain\_key}$ . **Else:**
  - (a) **If**  $\text{epoch\_num}_i - 1 \notin \text{compromised\_epochs}$  compute  $\text{root\_input} = \text{Exp}(\text{epoch\_id}_{1-i}, \text{epoch\_key}_i)$ . **Else** choose a random value  $\text{root\_input} \xleftarrow{\$} \mathbb{G}$  and store the tuple  $(\text{epoch\_num}_i, \text{root\_input})$ .
  - (b) Compute  $\text{sending\_chain\_key}_i = F(\text{root\_input})$  and  $s$   
Run  $KDF.\text{Compute}(s_c, \text{root\_input}) = (s'_c, \text{sending\_chain\_key}_i)$  and change state  $s_c = s'_c$ . Store the tuple  $(\text{epoch\_id}_i, \text{epoch\_id}_{i-1}, \text{sending\_chain\_key}_i)$  in the list `computed_values`.

**GetReceivingKey:**

6. (a) **If** there is a tuple  $(\text{epoch\_id}, \text{epoch\_id}_i, \text{chain\_key})$  in the list `computed_values`, set  $\text{recv\_chain\_key}_i = \text{chain\_key}$ .
- (b) **Else:**
  - i. Compute  $\text{root\_input} = \text{Exp}(\text{epoch\_id}, \text{epoch\_key}_i)$ .
  - ii. Compute  $\text{recv\_chain\_key}_i = F(\text{root\_input})$  and **Run**  $KDF.\text{Compute}(s_c, \text{root\_input}) = (s'_c, \text{recv\_chain\_key}_i)$  and update the value  $s_c = s'_c$ . Store the tuple  $(\text{epoch\_id}, \text{epoch\_id}_i, \text{recv\_chain\_key}_i)$  in the list `computed_values`.

 **$\mathcal{S}_{\text{eKE}}$  (Hybrid 2)****GetReceivingKey:**

3. If `diverge_parties = false`:
  - Run the KDF simulator  $\mathcal{S}_{KDF}(\text{Compute}, \text{View}_i.\text{root\_input})$ . Set the output as  $\text{View}_i.\text{temp\_recv\_chain\_key}$ .
4. Else (`diverge_parties = true`):
  - Set  $\text{View}_i.\text{temp\_recv\_chain\_key} = KDF.\text{Compute}(\text{View}_i.\text{root\_key}, \text{View}_i.\text{root\_input})$ .

**ReportState:**

5. (c) iv. **Run**  $\mathcal{S}_{KDF}(\text{Compromise}, \text{leakage}, (\text{root\_input}^*, \text{recv\_chain\_key}^*))$  to get an output `root_key`.  
v. Set  $\text{View}_{1-j}.\text{root\_key} = s_a || s_c \text{root\_key}$ .

**GenEpochId:**

4. (a) **Run**  $\mathcal{S}_{KDF}(\text{Advance}, \text{root\_input})$ .
6. (c) Set  $s_a || s_c = \text{View}_{1-j}.\text{root\_key}$ .
- (d) Delete the objects  $\text{View}_0, \text{View}_1$  and output  $(\text{GenEpochId}, \text{state}_{\mathcal{I}}, i, \text{epoch\_id}_i)$  to  $(\mathcal{F}_{\text{eKE}}, \text{sid.eKE})$ .
7. (c) Let  $s_a || s_c = \text{View}_i.\text{root\_key}$  and compute  $KDF.\text{Compute}(s_c, \text{View}_i.\text{root\_input}) = s'_c, k$ .  
If `diverge_parties = false`, run  $\mathcal{S}_{KDF}(\text{Compute}, \text{root\_input})$  and set  $\text{View}_i.\text{sending\_chain\_key}$  to be its output.
- (d) Update  $\text{View}_i.\text{sending\_chain\_key} = k$  and  $\text{View}_i.\text{root\_key} = s_a || s'_c$ .  
Else (`diverge_parties = true`), set  $\text{View}_i.\text{sending\_chain\_key} = KDF.\text{Compute}(\text{View}_i.\text{root\_key}, \text{root\_input})$ .

This hybrid is indistinguishable from the previous one because our  $KDF$  is a secure CPRFG. This is proved via a reduction to the security of the  $KDF$  protocol. An environment that can distinguish this hybrid from the previous one can be used to build an adversary that wins the CPRFG security game from Figure 22.

**Lemma 14** *Assume that  $KDF$  is a CPRFG. Then the view of the environment in hybrid  $H_1$  is computationally indistinguishable from that in hybrid  $H_2$ .*

**A:  $\mathcal{F}_{\text{eKE}}$  (proof of Lemma 14)****ConfirmReceivingEpoch:**

1. (d) On receiving  $(\text{ComputeSendingRootKey}, k, \text{ek}^{\text{pk}})$  from  $\mathcal{F}_{\text{LTM}}$ , initialise  $\mathcal{O}_{\text{CPRFG}}$ .  
sample a value  $s_a || s_c \xleftarrow{\$} \{0, 1\}^n$ .
4. (c) **If**  $\text{epoch\_num}_i - 2 \notin \text{compromised\_epochs}$ , compute  $\text{root\_input} = \text{Exp}(\text{epoch\_id}_{1-i}, \text{epoch\_key}_i)$ . **Else**, get `root_input` from the tuple  $(\text{epoch\_num}_i, \text{root\_input})$ .
- (d) Send  $(\text{Advance}, \text{root\_input})$  to  $\mathcal{O}_{\text{CPRFG}}$ .

Run  $KDF.Advance(s_a, root\_input) = s'$ , and change state  $s = s'$ .

//If  $epoch\_num_i \in compromised\_epochs$ , recovery just occurred so the epoch was advanced by the other party calling the `GetSendingKey` method.

**GetSendingKey:**

3. If there is a tuple  $(epoch\_id_i, epoch\_id_{i-1}, chain\_key)$  in the list `computed_values`, set  $sending\_chain\_key_i = chain\_key$ . **Else:**
  - (a) If  $epoch\_num_i - 1 \notin compromised\_epochs$  compute  $root\_input = Exp(epoch\_id_{i-1}, epoch\_key_i)$ . and send `(Compute, root_input)` to the oracle  $\mathcal{O}_{CPRFG}$ . **Else** send `(Recover)` to the oracle  $\mathcal{O}_{CPRFG}$ . choose a random value  $root\_input \xleftarrow{\$} \mathbb{G}$  and store the tuple  $(epoch\_num_i, root\_input)$ .
  - (b) Run  $KDF.Compute(s_c, root\_input) = (s'_c, sending\_chain\_key_i)$  and change state  $s_c = s'_c$ . **On getting a response**  $sending\_chain\_key_i$ , store the tuple  $(epoch\_id_i, epoch\_id_{i-1}, sending\_chain\_key_i)$  in the list `computed_values`.

**GetReceivingKey:**

6. (a) If there is a tuple  $(epoch\_id, epoch\_id_i, chain\_key)$  in the list `computed_values`, set  $recv\_chain\_key_i = chain\_key$ .
- (b) **Else:**
  - i. Compute  $root\_input = Exp(epoch\_id, epoch\_key_i)$ .
  - ii. If  $epoch\_id_i \notin compromised\_epochs$ , then send `(Compute, root_input)` to the oracle  $\mathcal{O}_{CPRFG}$  for the CPRFG security game and get an output  $chain\_key$ . **Else**, run  $KDF.Compute(s_c, root\_input) = (s'_c, recv\_chain\_key_i)$  and update the value  $s_c = s'_c$ . Store the tuple  $(epoch\_id, epoch\_id_i, recv\_chain\_key_i)$  in the list `computed_values`.

**A:  $S_{eKE}$  (proof of Lemma 14)**

**GetReceivingKey:**

3. If `diverge_parties = false`:
  - Send `(Compute, Viewi.root_input)` to  $\mathcal{O}_{CPRFG}$  Run the KDF simulator  $S_{KDF}(Compute, View_i.root\_input)$  and set the output as  $View_i.temp\_recv\_chain\_key$ .
4. Else (`diverge_parties = true`):
  - set  $s_a || s_c = View_i.root\_key$  and compute  $KDF.Compute(s_c, View_i.root\_input) = s'_c, k$ . Set  $View_i.temp\_recv\_chain\_key = KDF.Compute(View_i.root\_key, View_i.root\_input)$ .
  - Update  $View_i.temp\_recv\_chain\_key = k$  and  $View_i.root\_key = s_a || s'_c$ .

**ReportState:**

5. (c) iv. Send `(Compromise, root_input*)` to  $\mathcal{O}_{CPRFG}$  and get back a state  $s = s_a || s_c$ . Run  $S_{KDF}(Compromise, leakage, (root\_input^*, recv\_chain\_key^*))$  to get an output  $root\_key$ .
- v. Set  $View_{1-j}.root\_key = s_a || s_c || root\_key$ .

**GenEpochId:**

4. (a) Send `(Advance, root_input)` to  $\mathcal{O}_{CPRFG}$ . Run  $S_{KDF}(Advance, root\_input)$ .
6. (c) Set  $s_a || s_c = View_{1-j}.root\_key$ .
- (d) Delete the objects  $View_0, View_1$  and output `(GenEpochId, stateT, i, epoch_idi)` to  $(\mathcal{F}_{eKE}, sid.eKE)$ .
7. (c) If `diverge_parties = false`, send `(Compute, root_input)` to  $\mathcal{O}_{CPRFG}$  Run  $S_{KDF}(Compute, root\_input)$  and set  $View_i.sending\_chain\_key$  to be its output.
- (d) Else, set  $s_a || s_c = View_i.root\_key$  and compute  $KDF.Compute(s_c, View_i.root\_input) = s'_c, k$ . set  $View_i.sending\_chain\_key = KDF.Compute(View_i.root\_key, root\_input)$ .
- (e) Update  $View_i.sending\_chain\_key = k$  and  $View_i.root\_key = s_a || s'_c$ .

**Proof:** Let  $\mathcal{D}$  be an environment that distinguishes between hybrids  $H_1$  and  $H_2$  with advantage greater than negligible in  $\lambda$ , let the oracle  $\mathcal{O}_{\text{CPRFG}}$  be as in Definition 11 (an oracle that runs either the real or the ideal version of the CPRFG game from Figure 22 on Page 54). The oracle  $\mathcal{D}$  will be used to construct an adversary  $\mathcal{A}$  (described below) that can distinguish whether the oracle  $\mathcal{O}_{\text{CPRFG}}$  is running the real game or the ideal game with advantage greater than negligible in  $\lambda$ . The adversary  $\mathcal{A}$  maps messages from the distinguishing environment  $\mathcal{D}$  into messages for the CPRFG game and vice versa. The environment  $\mathcal{D}$ 's view in its interaction with the adversary  $\mathcal{A}$  exactly matches its view in either hybrid  $H_1$  or hybrid  $H_2$  depending on whether  $\mathcal{O}_{\text{CPRFG}}$  represents the ‘ideal game’ or the ‘real game’ respectively. Therefore, to acquire the exact same distinguishing probability as  $\mathcal{D}$ ,  $\mathcal{A}$  mimics the distinguishing decisions of  $\mathcal{D}$  exactly. That is, if at any point  $\mathcal{D}$  stops and produces an output,  $\mathcal{A}$  will produce the same output.

The adversary  $\mathcal{A}$  responds to requests from the distinguishing environment  $\mathcal{D}$  by running the instructions for  $\mathcal{F}_{\text{eKE}} + \mathcal{S}_{\text{eKE}} + \mathcal{I}_{\text{eKE}}$  from  $H_2$  with some minor changes where it queries  $\mathcal{O}_{\text{CPRFG}}$  to choose the sending and receiving chain keys that it outputs. The tricky part of replacing the choice of chain keys from  $H_2$  with calls to  $\mathcal{O}_{\text{CPRFG}}$  is the decision of when to send an `Advance` query to  $\mathcal{O}_{\text{CPRFG}}$ .

Most of the time, the `Advance` queries (in  $\mathcal{A}$ 's code) ensure that the epoch of  $\mathcal{O}_{\text{CPRFG}}$  is consistent with the epoch of the party that is behind, i.e the party that is not the latest sender. This allows  $\mathcal{A}$  to respond to `GetReceivingKey` requests for this party by making `Compute` queries to  $\mathcal{O}_{\text{CPRFG}}$ . However, when recovery from compromise occurs,  $\mathcal{A}$  must send a `Recover` query to  $\mathcal{O}_{\text{CPRFG}}$  to get the parties' new `sending_chain_key`, this necessarily advances the epoch of  $\mathcal{O}_{\text{CPRFG}}$  to be consistent with the latest sender. However due to the preceding corruption,  $\mathcal{A}$  has enough information in these situations to respond directly to `GetReceivingKey` requests for the party that is behind.

When the oracle  $\mathcal{O}_{\text{CPRFG}}$  corresponds to the ‘real game,’  $\mathcal{A}$  is simply executing hybrid  $H_2$ . Correspondingly, when the oracle  $\mathcal{O}_{\text{CPRFG}}$  corresponds to the ‘ideal game,’  $\mathcal{A}$  is simply executing hybrid  $H_1$ . Since  $KDF$  is a CPRFG, there is a simulator  $\mathcal{S}_{KDF}$  for which the ‘real execution’ and ‘ideal execution’ of the CPRFG game are computationally indistinguishable. For this same simulator then, hybrids  $H_1$  and  $H_2$  are therefore computationally indistinguishable.  $\square$

### 6.3.4 Hybrid $H_3$

This hybrid again changes the method of choosing chain keys *when both parties are uncompromised*. In particular, we replace the random choices of `root_inputs` with the Diffie-Hellman keys computed as specified in the protocol  $\Pi_{\text{eKE}}$  over a series of hybrids. Let  $\mathbb{G}, q, g$  be fixed, let  $M = \text{max\_epochnum}$ , then the hybrids are denoted by  $H_3^1, \dots, H_3^M$  where  $H_3^k$  is the hybrid where the root inputs for the first  $k$  epochs are DDH values computed using the specification of  $\Pi_{\text{eKE}}$ . All remaining root inputs are randomly sampled as in hybrid  $H_2$ . The hybrid  $H_3^M$  is computationally indistinguishable from  $H_2$  based on the *DDH* assumption on the group  $\mathbb{G}$ .

**Lemma 15** *Assume that the DDH assumption holds for group  $\mathbb{G}$ . Then the view of the environment in hybrid  $H_2$  is computationally indistinguishable from that of the last hybrid in series  $H_3$ .*

Lemma 15 will be proven via reduction to the following distinguishing game: The adversary  $\mathcal{A}$  is given a challenge triple  $\mathbb{G}, q, g, (g^x, g^y, g^z)$ . If  $b = 0$  then  $z \stackrel{\$}{\leftarrow} |\mathbb{G}|$ , if  $b = 1$  then  $z = x \cdot y$ .  $\mathcal{A}$  must distinguish whether  $b = 0$  or  $b = 1$ . Because the DDH assumption holds for group  $\mathbb{G}$ , no PPT adversary  $\mathcal{A}$  can distinguish whether  $b = 0$  or  $b = 1$  with advantage greater than negligible in the security parameter  $\lambda$ .

Say that there exists a distinguishing environment  $\mathcal{D}$  that can distinguish between  $H_2$  and  $H_3$  with non-negligible probability. Then  $\mathcal{D}$  can be used by an adversary  $\mathcal{A}$  to win the DDH distinguishing game with non-negligible advantage. The adversary  $\mathcal{A}$  is constructed differently on a case by case basis but in both cases the adversary  $\mathcal{A}$  will respond to requests from  $\mathcal{D}$  by running the instructions for  $\mathcal{F}_{\text{eKE}} + \mathcal{S}_{\text{eKE}} + \mathcal{I}_{\text{eKE}}$  from  $H_3$  up to epoch  $k - 1$  and from  $H_2$  starting epoch  $k + 1$  for some value  $k$ . Also in both cases  $\mathcal{A}$  will choose the root input for epoch  $k$  to be  $g^z$  from the challenge triple  $(g^x, g^y, g^z)$ , and the epoch ids for epochs  $k - 1$  and  $k$  to be  $g^x$  and  $g^y$  respectively. (Since  $\mathcal{A}$  not doesn't have the secret exponent for the epoch  $g^x$ , it chooses the root input for epoch  $k - 1$  to be  $(g^x)a$  where  $a$  is the known secret exponent for epoch  $k - 2$ .)

**Instructions for  $\mathcal{A}$  in both cases:**

1. If there is some environment  $\mathcal{D}$  that successfully distinguishes between hybrids  $H_2$  and  $H_3$  with non-negligible probability in the parameter  $\lambda$ , then it can also distinguish neighboring hybrids  $H_3^{k-1}$  and  $H_3^k$  for at least one  $k$ . First, find such a  $k$ .
2. When the parties are compromised, simply follow the code of hybrid  $H_2$  (which is the same as hybrid  $H_3^M$ .)
3. When the parties are not compromised follow the instructions below to choose root inputs for all epochs and to choose epoch ids for epochs  $k - 1$  and  $k$ . All the rest of the code of hybrids  $H_2$  and  $H_3^M$  is the same, so follow it exactly.
  - Up to epoch  $k$ , choose root inputs to be real Diffie-Hellman keys.
  - During epoch  $k$ , choose root input  $g^z$  from the challenge triple  $(g^x, g^y, g^z)$ . Correspondingly, choose epoch ids  $g^x, g^y$  for epochs  $k - 1$  and  $k$  respectively. (choose the root input for epoch  $k - 1$  to be  $(g^x)a$  where  $a$  is the known secret exponent for epoch  $k - 2$ )
  - For each epoch starting  $k + 1$ , randomly sample a fresh root input  $h \xleftarrow{\$} \mathbb{G}$ .

**Case 1:** The probability that epoch  $k$  is compromised by  $\mathcal{D}$  is higher in one of the two hybrids  $H_3^{k-1}$  and  $H_3^k$  by an amount that is non-negligible in  $\lambda$ .

If epoch  $k$  is compromised by  $\mathcal{D}$  significantly more often in one of the two hybrids, then the adversary can win the DDH distinguishing game by choosing its outputs as below:

- 3 If  $\mathcal{D}$  takes any actions that cause epoch  $k$  to be compromised, stop and output 1, otherwise output a randomly chosen bit at the end of the execution. (The actions that could cause epoch  $k$  to be compromised include: 1) corrupting the party in epochs `epoch_num* + 1, epoch_num*, epoch_num* - 1, epoch_num* - 2, epoch_num* - 3`, 2) diverging the parties 3) sending bogus epoch\_ids to the receiver of epoch  $k - 2$  or  $k - 1$  while the epoch is compromised.)

$$\begin{aligned} & \Pr[\mathcal{A}(g^x, g^y, g^{xy}) = 1] - \Pr[\mathcal{A}(g^x, g^y, g^z) = 1] \\ &= \Pr[\text{epoch } k \text{ compromised in } H_3^{k-1}] - \Pr[\text{epoch } k \text{ compromised in } H_3^k] \end{aligned}$$

**Case 2:** Now, if the probability that epoch  $k$  is compromised  $\mathcal{D}$  in hybrid  $H_3^{k-1}$  is within a negligible distance of the probability that epoch  $k$  is compromised in hybrid  $H_3^k$ , then this can be used by  $\mathcal{A}$  to win the DDH distinguishing game by choosing its outputs as follows:

- 3 If  $\mathcal{D}$  takes any actions that cause epoch  $k$  to be compromised, then stop and output a randomly chosen bit, otherwise wait till  $\mathcal{D}$  stops, and then produce the same output that  $\mathcal{D}$  produces.

Let the event that epoch  $k$  is compromised be  $\text{cprm}$  and let the event that epoch  $k$  is not compromised be  $\neg\text{cprm}$ . Note that when epoch  $k$  is compromised then the hybrids  $H_3^{k-1}$  and  $H_3^k$  are identical. So, an environment  $\mathcal{D}$  that corrupts epoch  $k$  roughly the same amount in hybrids  $H_3^k$  and  $H_3^{k+1}$  must have that  $\Pr[\mathcal{D}(H_3^{k+1}) = 1 \mid \text{cprm}] - \Pr[\mathcal{D}(H_3^k) = 1 \mid \text{cprm}] = \text{negl}(\lambda)$  for some negligible function  $\text{negl}(\cdot)$ . Therefore,

$$\begin{aligned} & \Pr[\mathcal{A}(g^x, g^y, g^{xy}) = 1] - \Pr[\mathcal{A}(g^x, g^y, g^z) = 1] \\ &= \Pr[\mathcal{D}(H_3^{k+1}) = 1 \mid \neg\text{cprm}] - \Pr[\mathcal{D}(H_3^k) = 1 \mid \neg\text{cprm}] \\ &= \Pr[\mathcal{D}(H_3^{k+1}) = 1] - \Pr[\mathcal{D}(H_3^k) = 1] - \text{negl}(\lambda) \end{aligned}$$

Given that epoch  $k$  is not compromised, the environment's view is exactly as in  $H_3^{k-1}$  when  $b = 0$  and  $H_3^k$  when  $b = 1$ . So if  $\mathcal{D}$  distinguishes between  $h_3^{k-1}$  and  $H_3^k$  with non-negligible probability then so does  $\mathcal{A}$ .

### 6.3.5 Hybrid $H_4$

In this hybrid the randomly chosen initial root keys are replaced by root keys computed using  $\mathcal{F}_{\text{DIR}}$ ,  $\mathcal{F}_{\text{LTM}}$  as in  $\Pi_{\text{eKE}}$ . While the functionality  $\mathcal{F}_{\text{eKE}}$  automatically provides a binding between party id's and epoch keys up to the first compromise, the protocol  $\Pi_{\text{eKE}}$  realizes this binding via access to  $\mathcal{F}_{\text{LTM}}$ ,  $\mathcal{F}_{\text{DIR}}$ . Also, this initial root key is indistinguishable from a random choice because  $(\mathcal{F}_{\text{LTM}}, \text{pid})$  is not corruptible and it hides all Diffie-Hellman exponents belonging to  $\text{pid}$ . In summary, this hybrid is indistinguishable from the previous one because:

- $\mathcal{F}_{\text{LTM}}$  hides the exponents for all the group elements it chooses.
- The DDH property applies to the group that  $\mathcal{F}_{\text{LTM}}$  chooses elements from.
- $\mathcal{F}_{\text{DIR}}$  reliably provide parties with each others' public diffie hellman halves  $\text{ik}, \text{rk}$ .
- $\mathcal{F}_{\text{DIR}}$  reliably provides pairs of parties with the same public diffie hellman half  $\text{ok}$ .
- The provided  $\text{ok}$  is unique to a pair of parties.

**Lemma 16** *Assume that the DDH assumption holds for group  $\mathbb{G}$ . Then the view of the environment in hybrid  $H_4$  is computationally indistinguishable from that in  $H_3$ .*

### 6.3.6 Hybrid $H_5$ (Real World)

Finally, we can now separate the epoch key exchange instances that the two parties interact with by splitting the functionality into two instances. In the hybrids before this, if either party is corrupted or confirms an incorrect epoch id then control is handed over to the simulator who runs separate instances of the protocol for each party anyway. So the interesting case is when the parties' views of the  $\text{epoch\_id}$ 's so far agree. In this case, agreement for all previous hybrids is by fiat. Here, agreement holds because of the consistency of  $\text{KDF.Compute}$  and because  $\text{KDF.Advance}$  is a deterministic function of  $s_a, r$ .

**Lemma 17** *The view of the environment in hybrid  $H_5$  is identically distributed to that in  $H_4$ .*

□



## 7 Unidirectional Forward Secure Authenticated Channels: Realising $\mathcal{F}_{\text{fs\_aead}}$

In this section, we prove that protocol  $\Pi_{\text{fs\_aead}}$  (which uses ideal functionalities  $\mathcal{F}_{\text{mKE}}$  and  $\mathcal{F}_{\text{aead}}$  as subroutines) UC-realizes the functionality  $\mathcal{F}_{\text{fs\_aead}}$  (which is described in section Section 5.2).

The protocol  $\Pi_{\text{fs\_aead}}$  and the proof that it realises  $\mathcal{F}_{\text{fs\_aead}}$  are presented in Section 7.3. Prior to that, the component functionalities  $\mathcal{F}_{\text{mKE}}$  and  $\mathcal{F}_{\text{aead}}$  are defined in Sections 7.1 and 7.2 respectively.

### 7.1 The Message Key Exchange Functionality $\mathcal{F}_{\text{mKE}}$

In this section, we describe the ideal message key exchange functionality  $\mathcal{F}_{\text{mKE}}$  (as described previously in Section 4). This functionality, which is shown in full detail in Figure 28 on page 73, provides message key management for a single epoch, mainly by keeping track of which message keys have been generated and retrieved and which message keys have been exposed during a state compromise. In the context of  $\Pi_{\text{fs\_aead}}$ ,  $\mathcal{F}_{\text{mKE}}$  provides `key_seed`'s to  $\mathcal{F}_{\text{aead}}$  instances for each message within the epoch that  $\mathcal{F}_{\text{mKE}}$  represents. For each `msg_num`, the functionality allows at most one (`RetrieveKey`, `pid`) request for each party `pid`. If a second request is made then it simply ends the activation. That is the functionality will provide the `key_seed` at most once for the encryption and for the decryption of each message in the epoch. The functionality also closes epochs by allowing the  $\Pi_{\text{fs\_aead}}$  instance for each party to make a request (`StopKeys`,  $N$ ). Once such a request is made by a party, the functionality will disallow any `RetrieveKey` requests made on behalf of that party for `msg_num`  $> N$ .

### 7.2 The Single-Message Authenticated Encryption Functionality $\mathcal{F}_{\text{aead}}$

In this section, we define an single-message authenticated encryption functionality  $\mathcal{F}_{\text{aead}}$ , with the full details shown in Figure 29 on page 74). Each  $\mathcal{F}_{\text{aead}}$  instance handles the encryption, decryption, and authentication of a particular message for a particular epoch and hands the ciphertext or message back to  $\Pi_{\text{fs\_aead}}$ .

**Encryption:** On receiving an encryption request, while not corrupted,  $\mathcal{F}_{\text{aead}}$  first consults the message key exchange functionality  $\mathcal{F}_{\text{mKE}}$  to decide whether encryption of this message is possible for the party `pid`. It then runs adversarial code  $\mathcal{I}_{\text{aead}}$  with the message length  $|m|$  to choose the corresponding ciphertext  $c$ . It then simply stores the message and adversarially-generated ciphertext in a table. While corrupted, the functionality instead provides the message  $m$  to the adversary and allows it to choose a ciphertext directly with this knowledge.

$$\mathcal{F}_{\text{mKE}}$$

This functionality has a session id  $\text{sid.mKE}$  that takes the following format:  $\text{sid.mKE} = (\text{"mKE"}, \text{sid.fs})$ . Where  $\text{sid.fs} = (\text{"fs_aead"}, \text{sid}, \text{epoch\_id})$ . The local session ID is parsed as  $\text{sid} = (\text{sid}', \text{pid}_0, \text{pid}_1)$ . Inputs arriving from machines whose identity is neither  $\text{pid}_0$  nor  $\text{pid}_1$  are ignored. //For notational simplicity we assume some fixed interpretation of  $\text{pid}_0$  and  $\text{pid}_1$  as complete identities of the two calling machines.

This functionality is parametrized by a seed length  $\lambda$

**RetrieveKey:** On receiving  $(\text{RetrieveKey}, \text{pid})$  from  $(\Pi_{\text{aead}}, \text{sid.aead})$ , where  $\text{sid.aead} = (\text{"aead"}, \text{sid.fs}, \text{msg\_num})$ , or from  $\mathcal{F}_{\text{aead}}$ :

1. If this is the first activation,
  - Parse  $\text{sid}$  to recover the two party ids  $(\text{pid}_0, \text{pid}_1)$ .
  - Initialize dictionary  $\text{key\_dict}$  and variables  $\text{IsCorrupt?} = \text{false}$ ,  $\text{msg\_num}_0, \text{msg\_num}_1 = 0$ .  
//If this epoch is corrupted then this isn't the first activation because there was already a  $(\text{Corrupt})$  request.  
// $\text{msg\_num}_i$  is the largest message number whose key has been successfully retrieved by  $\text{pid}_i$ .
2. If there is record  $(\text{Retrieved}, i, \text{msg\_num})$  or a record  $(\text{StopKeys}, i, N)$  for  $N < \text{msg\_num}$ , end the activation.
3. If  $\text{IsCorrupt?} = \text{false}$ :
  - If  $\text{msg\_num} \in \text{key\_dict.keys}$ , set  $k = \text{key\_dict}[\text{msg\_num}]$ .
  - Else  $(\text{msg\_num} \notin \text{key\_dict.keys})$ , set  $k \xleftarrow{\$} \{0, 1\}^\lambda$ .
4. Else  $(\text{IsCorrupt?} = \text{true})$ :
  - Send  $(\text{RetrieveKey}, \text{pid}, \text{msg\_num})$  to the the adversary.
  - Upon receiving  $(\text{RetrieveKey}, \text{pid}, k)$  from the adversary, continue.
5. Store  $\text{key\_dict}[\text{msg\_num}] = k$ .
6. If  $\text{msg\_num} > \text{msg\_num}_i$ , set  $\text{msg\_num}_i = \text{msg\_num}$ . // $\text{msg\_num}_i$  is the largest message number whose key has been successfully retrieved by  $\text{pid}_i$ .
7. Record  $(\text{Retrieved}, i, \text{msg\_num})$  and output  $(\text{RetrieveKey}, \text{pid}, k)$  to  $(\Pi_{\text{aead}}, \text{sid.aead})$ .

**StopKeys:** On receiving  $(\text{StopKeys}, N)$  from  $(\Pi_{\text{fs\_aead}}, \text{sid.fs} = (\text{"fs\_aead"}, \text{sid}, \text{epoch\_id}, b), \text{pid})$ ,

- Run steps 2-6 of **RetrieveKey** for all  $\text{msg\_num}$  such that  $\text{msg\_num}_i < \text{msg\_num} \leq N$ .
- Record  $(\text{StopKeys}, i, N)$  and output  $(\text{StopKeys}, \text{Success})$ .

**Corruption:** On receiving  $(\text{Corrupt})$  from  $(\Pi_{\text{fs\_aead}}, \text{sid.fs} = (\text{"fs\_aead"}, \text{sid}, \text{epoch\_id}, b), \text{pid})$ :

1. Let  $i$  be such that  $\text{pid} = \text{pid}_i$ .
2. Set  $\text{IsCorrupt?} = \text{true}$ , create empty lists  $\text{keys\_in\_transit}, \text{pending\_msgs}$ .
3. For all  $\text{msg\_num} \in \text{key\_dict.keys}$ , if there is no record  $(\text{Retrieved}, i, \text{msg\_num})$  then append  $(\text{msg\_num}, \text{key\_dict}[\text{msg\_num}])$  to  $\text{keys\_in\_transit}$  and append  $\text{msg\_num}$  to  $\text{pending\_msgs}$ .
4. If there is a record  $(\text{StopKeys}, i, N)$  set  $\text{stop\_request} = (\text{StopKeys}, i, N)$ , Else set  $\text{stop\_request} = \perp$ .
5. Send  $(\text{ReportState}, i, \text{keys\_in\_transit}, \text{msg\_num}_0, \text{msg\_num}_1, \text{stop\_request})$  to  $\mathcal{A}$ .
6. On receiving a response  $(\text{ReportState}, i, S)$  from  $\mathcal{A}$ , output  $(\text{Corrupt}, S)$  to  $(\Pi_{\text{fs\_aead}}, \text{sid.fs}, \text{pid}_i)$ .

Figure 28: The Message Key Exchange Functionality  $\mathcal{F}_{\text{mKE}}$

### $\mathcal{F}_{\text{aead}}$

This functionality has a session id  $\text{sid.aead} = (\text{"aead"}, \text{sid.fs}, \text{msg\_num})$  where  $\text{sid.fs} = (\text{"fs\_aead"}, \text{sid} = (\text{sid}', \text{pid}_0, \text{pid}_1), \text{epoch\_id})$ . Inputs arriving from machines whose identity is neither  $\text{pid}_0$  nor  $\text{pid}_1$  are ignored. //For notational simplicity we assume some fixed interpretation of  $\text{pid}_0$  and  $\text{pid}_1$  as complete identities of the two calling machines.

**Encryption:** On receiving  $(\text{Encrypt}, m, N)$  from  $(\Pi_{\text{fs\_aead}}, \text{sid.fs}, \text{pid})$ :

1. Let  $i$  be such that  $\text{pid} = \text{pid}_i$ .
2. If this is not the first encryption request end the activation.
3. **If**  $\text{IsCorrupt?} = \text{true}$ , Send a backdoor message  $(\text{state}_{\mathcal{I}}, \text{Encrypt}, \text{pid}, m, N)$  to  $\mathcal{A}$ .
4. **Else** ( $\text{IsCorrupt?} = \text{false}$ ):
  - Provide input  $(\text{RetrieveKey}, \text{pid})$  to  $(\mathcal{F}_{\text{mKE}}, \text{sid.mKE}, \text{pid})$ .
  - Upon receiving output  $(\text{RetrieveKey}, \text{pid}, k)$  from  $(\mathcal{F}_{\text{mKE}}, \text{sid.mKE}, \text{pid})$ , if  $k = \perp$  then end the activation. //The key is not available.
  - Initialize  $\text{state}_{\mathcal{I}} = \perp$ , call  $\mathcal{F}_{\text{lib}}$  to obtain the internal code  $\mathcal{I}$ , and run  $\mathcal{I}(\text{state}_{\mathcal{I}}, \text{Encrypt}, \text{pid}, |m|, N)$ .
5. Upon obtaining  $(\text{state}_{\mathcal{I}}, \text{Encrypt}, \text{pid}, c)$  from  $\mathcal{A}$  or  $\mathcal{I}$ , record the tuple  $(c, m, N, 1)$ , record the sender  $\text{S} = i$ , and set  $\text{ready2decrypt} = \text{true}$ .
6. Output  $(\text{Encrypt}, c)$  to  $(\Pi_{\text{fs\_aead}}, \text{sid.fs}, \text{pid})$ .

**Decryption:** On receiving  $(\text{Decrypt}, c, N)$  from  $(\Pi_{\text{fs\_aead}}, \text{sid.fs}, \text{pid})$ :

1. **If**  $\text{IsCorrupt?} = \text{false}$ :
  - If there isn't a stored key  $k$ , provide input  $(\text{RetrieveKey}, \text{pid})$  to  $(\mathcal{F}_{\text{mKE}}, \text{sid.mKE}, \text{pid})$ . Upon obtaining a response  $(\text{RetrieveKey}, \text{pid}, k)$  from  $\mathcal{F}_{\text{mKE}}$ : Store  $k$ .
  - If there hasn't been a successful encryption request or  $\text{pid} \neq \text{pid}_{\text{S}-i}$ , output  $(\text{Decrypt}, \text{Fail})$  to  $(\Pi_{\text{fs\_aead}}, \text{sid.fs}, \text{pid})$ .
2. If  $k = \perp$ , if  $\text{ready2decrypt} = \text{false}$ , or if there is a record  $(c, N, 0)$ , output  $(\text{Decrypt}, \text{Fail})$  to  $(\Pi_{\text{fs\_aead}}, \text{sid.fs}, \text{pid})$ . //Failure of decryption can occur for an honest receiver so we need an explicit failure notification.
3. If there is a record  $(c, m, N, 1)$ , note  $\text{ready2decrypt} = \text{false}$  and output  $(\text{Decrypt}, m)$ .
4. **If**  $\text{IsCorrupt?} = \text{false}$ :
  - Run  $\mathcal{I}(\text{state}_{\mathcal{I}}, \text{Authenticate}, \text{pid}, c, N)$  to obtain the output  $(\text{state}_{\mathcal{I}}, v)$ .
  - **If**  $v = \perp$  then record  $(c, N, 0)$ , and output  $(\text{Decrypt}, \text{Fail})$ . **Else**, note  $\text{ready2decrypt} = \text{false}$ , and output  $(\text{Decrypt}, m)$ .
5. **Else** ( $\text{IsCorrupt?} = \text{true}$ ):
  - Send backdoor message  $(\text{state}_{\mathcal{I}}, \text{Inject}, \text{pid}, c, N)$  to  $\mathcal{A}$ .
  - Upon receiving response  $(\text{state}_{\mathcal{I}}, \text{Inject}, \text{pid}, c, N, v)$  from  $\mathcal{A}$ : **If**  $v = \perp$  then record  $(c, N, 0)$ , and output  $(\text{Decrypt}, \text{Fail})$ . **Else** output  $(\text{Decrypt}, v)$  to  $(\Pi_{\text{fs\_aead}}, \text{sid.fs}, \text{pid})$ .

**Corruption:** On receiving  $(\text{Corrupt})$  from  $(\Pi_{\text{fs\_aead}}, \text{sid.fs}, \text{pid})$ :

1. Provide input  $(\text{RetrieveKey}, \text{pid})$  to  $(\mathcal{F}_{\text{mKE}}, \text{sid.mKE}, \text{pid})$ .
2. Upon obtaining a response  $(\text{RetrieveKey}, \text{pid}, k)$  from  $\mathcal{F}_{\text{mKE}}$ , store  $k$ .
3. If a message  $m$  was successfully decrypted then set  $m^* = m$ , otherwise, set  $m^* = \perp$ .
4. Set  $\text{IsCorrupt?}$  to  $\text{true}$ , and send  $(\text{ReportState}, \text{state}_{\mathcal{I}}, \text{pid}, k, m^*)$  to  $\mathcal{A}$ .
5. Upon receiving a response  $(\text{ReportState}, \text{pid}, S)$  from  $\mathcal{A}$ , send  $(\text{Corrupt}, S)$  to  $(\Pi_{\text{fs\_aead}}, \text{sid.fs})$ .

Figure 29: The Authenticated Encryption with Associated Data Functionality,  $\mathcal{F}_{\text{aead}}$

## $\Pi_{fs\_aead}$

This protocol is active during a *single* epoch and has session id  $sid.fs$  that takes the following format:  $sid.fs = ("fs\_aead", sid, epoch\_id)$ .

**Encrypt:** On receiving input  $(\text{Encrypt}, m, N)$  from  $(\Pi_{SGNL}, sid, pid)$  do:

1. Check that  $sid$  matches the one in the local session ID, and that  $pid$  matches the local party id, else abort.
2. If this is the first activation then initialize  $curr\_msg\_num = 0$ .
3. Increment  $curr\_msg\_num + = 1$ .
4. Send  $(\text{Encrypt}, m, N)$  to  $(\mathcal{F}_{aead}, sid.aead = ("aead", sid.fs, msg\_num))$  and delete  $m$ .
5. Upon receiving  $(\text{Encrypt}, c)$ , output  $(\text{Encrypt}, c)$  to  $(\Pi_{SGNL}, sid, pid)$ .

**Decrypt:** On receiving  $(\text{Decrypt}, c, msg\_num, N)$  from  $(\Pi_{SGNL}, sid, pid)$  do:

1. Check that  $sid$  matches the one in the local session ID, and that  $pid$  matches the local party id, else abort.
2. If this is the first activation then initialize  $curr\_msg\_num = 0$ ,  $pending\_msgs = []$ .
3. Send  $(\text{Decrypt}, c, N)$  to  $(\mathcal{F}_{aead}, sid.aead = ("aead", sid.fs, msg\_num))$ .
4. Upon receiving  $(\text{Decrypt}, v)$ , if  $v = \text{Fail}$  then output  $(\text{Decrypt}, \text{Fail})$  to  $(\Pi_{SGNL}, sid, pid)$ .
5. Otherwise ( $v \neq \text{Fail}$ ):
  - While  $curr\_msg\_num < msg\_num$ :
    - Increment  $curr\_msg\_num + = 1$ .
    - Add  $curr\_msg\_num$  to  $pending\_msgs$ .
  - Remove  $msg\_num$  from  $pending\_msgs$  and output  $(\text{Decrypt}, v)$  to  $(\Pi_{SGNL}, sid, pid)$ .

**StopEncrypting:** On receiving  $(\text{StopEncrypting})$  from  $(\Pi_{SGNL}, sid, pid)$  do:

1. Check that  $sid$  matches the one in the local session ID, and that  $pid$  matches the local party id, else abort.
2. Send  $(\text{StopKeys}, msg\_num)$  to  $(\mathcal{F}_{mKE}, sid.mKE)$ .
3. On receiving  $(\text{StopKeys}, \text{Success})$ , output  $(\text{StopEncrypting}, \text{Success})$  to  $(\Pi_{SGNL}, sid, pid)$ .

**StopDecrypting:** On receiving  $(\text{StopDecrypting}, msg\_num^*)$  from  $(\Pi_{SGNL}, sid, pid)$  do:

1. Check that  $sid$  matches the one in the local session ID, and that  $pid$  matches the local party id, else abort.
2. Send  $(\text{StopKeys}, msg\_num^*)$  to  $(\mathcal{F}_{mKE}, sid.mKE)$ .
3. On receiving  $(\text{StopKeys}, \text{Success})$ :
  - While  $curr\_msg\_num < msg\_num^*$ :
    - Increment  $curr\_msg\_num + = 1$ .
    - Add  $curr\_msg\_num$  to  $pending\_msgs$ .
  - Output  $(\text{StopDecrypting}, \text{Success})$  to  $(\Pi_{SGNL}, sid, pid)$ .

**Corruption:** On receiving  $(\text{Corrupt})$  from  $(\Pi_{SGNL}, sid, pid)$ :

*//Note that the **Corrupt** interface is not part of the "real" protocol; it is only included for UC-modelling purposes.*

1. Check that  $sid$  matches the one in the local session ID, and that  $pid$  matches the local party id, else abort.
2. Initialize a state object  $S$  and add  $curr\_msg\_num$ ,  $pending\_msgs$  to it.
3. Send  $(\text{Corrupt})$  to  $(\mathcal{F}_{mKE}, sid.mKE = ("mKE'', sid.fs))$ .
4. On receiving  $(\text{Corrupt}, S_{mKE})$ , add it to  $S$  and do the following.
5. For each record  $msg\_num \in pending\_msgs$ :
  - Send  $(\text{Corrupt}, pid_i)$  to  $(\mathcal{F}_{aead}, sid.aead = ("aead'', sid.fs, msg\_num))$
  - On receiving a response  $(\text{Corrupt}, pid_i, S_{msg\_num})$ , add  $S_{msg\_num}$  to  $S$ .
6. Output  $(\text{Corrupt}, S)$  to  $(\Pi_{SGNL}, sid, pid)$ .

Figure 30: The Forward-Secure Encryption  $\Pi_{fs\_aead}$

$$\mathcal{I}_{\text{fsaead}}$$

**Encrypt:** On receiving  $(\text{state}_{\mathcal{I}}, \text{Encrypt}, \text{pid}, N, |m|)$  from  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid}, fs)$  do:

1. Let  $i$  be such that  $\text{pid} = \text{pid}_i$ .
2. **If**  $\text{state}_{\mathcal{I}} = \perp$ :
  - (a) Initialize objects  $\text{View}_0$ ,  $\text{View}_1$ , and dictionary records.
  - (b) Set  $\text{View}_0.\text{curr\_msg\_num} = \text{View}_1.\text{curr\_msg\_num} = 0$ .
  - (c) Set  $\text{View}_0.\text{pending\_msgs} = \text{View}_1.\text{pending\_msgs} = []$
3. **Else**, parse the objects  $\text{View}_0$ ,  $\text{View}_1$ , and the dictionary records from  $\text{state}_{\mathcal{I}}$ .
4. Set  $\text{msg\_num} = \text{View}_i.\text{curr\_msg\_num} + 1$  and update  $\text{View}_i.\text{curr\_msg\_num} + = 1$ .
5. If  $\text{msg\_num\_state\_aead}$  exists in  $\text{state}_{\mathcal{I}}$  then parse it otherwise initialize  $\text{msg\_num\_state\_aead} = \perp$ .
6. Run  $\mathcal{I}_{\text{aead}}(\text{msg\_num\_state\_aead}, \text{Encrypt}, \text{pid}, |m|, N)$ .
7. Upon receiving output  $(\text{msg\_num\_state\_aead}, \text{Encrypt}, \text{pid}, c)$ , store  $(\ell = |m|, \text{pid}, \text{ready2decrypt} = \text{true}), (c, N, 1)$  in  $\text{records}[\text{msg\_num}]$ .
8. Update or store for the first time the values  $\text{View}_0$ ,  $\text{View}_1$ ,  $\text{records}$ ,  $\text{msg\_num\_state\_aead}$ , in  $\text{state}_{\mathcal{I}}$
9. Output  $(\text{state}_{\mathcal{I}}, \text{Encrypt}, \text{pid}, \text{msg\_num}, N, c)$  to  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid}, fs)$ .

**Authenticate:** On receiving  $(\text{state}_{\mathcal{I}}, \text{Authenticate}, \text{pid}, c, \text{msg\_num}, N)$  from  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid}, fs)$  do:

1. Let  $i$  be such that  $\text{pid} = \text{pid}_i$ .
2. Parse the dictionary records from  $\text{state}_{\mathcal{I}}$ :
  - (a) **If**  $\text{records}[\text{msg\_num}] = \perp$  then output  $(\text{state}_{\mathcal{I}}, \text{Authenticate}, \text{pid}, c, \text{msg\_num}, N, \perp)$  to  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid}, fs)$ .
  - (b) **Else** get  $(\ell, \text{pid}, \text{ready2decrypt})$  and any record of the form  $(c, N, v)$  from  $\text{records}[\text{msg\_num}]$ .
3. **If**  $\text{ready2decrypt} = \text{false}$  in the retrieved record, output  $(\text{state}_{\mathcal{I}}, \text{Authenticate}, \text{pid}, c, \text{msg\_num}, N, \perp)$  to  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid}, fs)$ .
4. **If** there is no record  $(c, n, v)$ :
  - Run  $\mathcal{I}_{\text{aead}}(\text{msg\_num\_state\_aead}, \text{Authenticate}, \text{pid}, c, N)$ . On obtaining the output  $(\text{msg\_num\_state\_aead}, \text{Authenticate}, \text{pid}, c, N, v)$ , append  $(c, N, v)$  to  $\text{records}[\text{msg\_num}]$  and update records in  $\text{state}_{\mathcal{I}}$ .
5. **If**  $v = 0$  then output  $(\text{state}_{\mathcal{I}}, \text{Authenticate}, \text{pid}, c, \text{msg\_num}, N, \perp)$  to  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid}, fs)$ . **Else** continue.
6. **While**  $\text{View}_i.\text{curr\_msg\_num} \leq \text{msg\_num}$  do: //we only get to this step if decryption succeeds.
  - Increment  $\text{View}_i.\text{curr\_msg\_num} + = 1$ .
  - Append  $\text{View}_i.\text{curr\_msg\_num}$  to  $\text{View}_i.\text{pending\_msgs}$ .
7. Output  $(\text{state}_{\mathcal{I}}, \text{Authenticate}, \text{pid}, c, \text{msg\_num}, N, 1)$  to  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid}, fs)$ .

Figure 31: Internal Adversary,  $\mathcal{I}_{\text{fsaead}}$

**Decryption:** Once a successful decryption occurs, an instance of the functionality  $\mathcal{F}_{\text{aead}}$  will never try to decrypt a ciphertext again. While not corrupted, the encryption functionality  $\mathcal{F}_{\text{aead}}$  will only ever output the originally encrypted message  $m$  or a failure notification. On receiving a decryption request, if the instance is not corrupted, then the functionality first consults with the message key exchange functionality  $\mathcal{F}_{\text{mKE}}$  to decide whether decryption of this message is possible for the party  $\text{pid}$  at all, if not then it outputs a failure notification. (If the instance has been corrupted already then this key has already been retrieved and sent to  $\mathcal{A}$  by the functionality.) Otherwise, if the ciphertext provided matches the ciphertext  $c$  chosen at the time of corruption then

$\mathcal{F}_{\text{aead}}$  always outputs  $m$  in both the honest and corrupted case. In general  $\mathcal{F}_{\text{aead}}$  keeps consistency with its prior outputs, always outputting a failure notification for ciphertexts it previously refused to decrypt. If the ciphertext does not match the one produced at the time of encryption, and if decryption of this ciphertext has never been attempted before, then based on whether corruption has occurred, the functionality will either run the adversarial code to decide if the ciphertext authenticates or it will allow the adversary to decrypt the ciphertext to any value of its choosing (up to consistency of outputs.).

### 7.3 Protocol $\Pi_{\text{fs\_aead}}$

As outlined in Section 4, protocol  $\Pi_{\text{fs\_aead}}$  makes straightforward use of multiple instances of  $\mathcal{F}_{\text{aead}}$ , along with  $\mathcal{F}_{\text{mKE}}$ . It is presented in Figure 30 on page 75. In this subsection we prove Theorem 4.

**Theorem 4** *Protocol  $\Pi_{\text{fs\_aead}}$  (perfectly) UC-realizes the ideal functionality  $\mathcal{F}_{\text{fs\_aead}}$  in the presence of  $\mathcal{F}_{\text{lib}}$ .*

The proof of Theorem 4 takes up the rest of this section. For legibility, we first provide a high-level overview of the proof, and then we provide the details in the two cases where the parties are and aren't corrupted.

**Proof:** To prove Theorem 4, we first construct the simulator  $\mathcal{S}_{\text{fsaead}}$  and the internal code  $\mathcal{I}_{\text{fsaead}}$ . Then, we argue that the environment  $\text{Env}$  has an identically distributed view in its interaction with  $\Pi_{\text{fs\_aead}} + \mathcal{F}_{\text{mKE}} + \mathcal{F}_{\text{aead}}$  as it does in its interaction with  $\mathcal{F}_{\text{fs\_aead}} + \mathcal{S}_{\text{fsaead}} + \mathcal{I}_{\text{fsaead}}$ .

The objective of  $\mathcal{I}_{\text{fsaead}}$  is to use  $\mathcal{I}_{\text{aead}}$  to provide ciphertexts that match the distribution of ciphertexts in the real world. This is necessary because we do not restrict the ciphertexts provided by  $\mathcal{F}_{\text{fs\_aead}}$  to any particular distribution or format. (The detailed specification of this internal code can be found in Figure 31 on page 76.)

The objective of the ideal process adversary  $\mathcal{S}_{\text{fsaead}}$  is to simulate the interactions that would take place between the environment and the protocol  $\Pi_{\text{fs\_aead}}$  once a party has been corrupted. Since  $\Pi_{\text{fs\_aead}}$ 's security properties derive entirely from its component functionalities  $\mathcal{F}_{\text{mKE}}$  and  $\mathcal{F}_{\text{aead}}$ , the view of an environment  $\text{Env}$  will be identically distributed in the real and ideal scenarios. When the simulator  $\mathcal{S}_{\text{fsaead}}$  receives messages from  $\mathcal{F}_{\text{fs\_aead}}$  in the ideal world, it takes the actions specified in the pseudocode in Figure 32 on page 78.

Observe that the APIs of  $\mathcal{F}_{\text{fs\_aead}}$  and  $\Pi_{\text{fs\_aead}}$  are identical: they each have 5 methods. In the remainder of the proof, we compare the actions taken by  $\Pi_{\text{fs\_aead}}$  in the real world with the actions taken by  $\mathcal{F}_{\text{fs\_aead}}$  together with  $\mathcal{I}_{\text{fsaead}}$  and  $\mathcal{S}_{\text{fsaead}}$  in the ideal world. As in the proof of Theorem 2, this proof holds by induction over the actions taken by the internal adversarial code  $\mathcal{I}_{\text{fsaead}}$  and the simulator  $\mathcal{S}_{\text{fsaead}}$ . We show that each action leaves the distribution in the real and ideal worlds the same: that is, each action of  $\mathcal{S}_{\text{fsaead}}$  and  $\mathcal{I}_{\text{fsaead}}$  maintains the property that the view of the environment  $\text{Env}$  is identical when interacting with either the ideal functionality  $\mathcal{F}_{\text{fs\_aead}}$  or the real protocol  $\Pi_{\text{fs\_aead}}$ .

Without loss of generality, we restrict our attention to a dummy adversary  $\mathcal{A}$  and a deterministic environment  $\text{Env}$ . As a consequence, the entire execution is deterministic, since the message keys themselves are never used for encryption in  $\mathcal{F}_{\text{aead}}$ .

## $\mathcal{S}_{fsaead}$

**At first activation:** Send  $(\mathcal{F}_{fs\_aead}, \mathcal{I})$  to  $\mathcal{F}_{lib}$ .

**ReportState:** On receiving  $(state_{\mathcal{I}}, ReportState, pid, leakage)$  for  $pid_i \in \{pid_0, pid_1\}$  from  $(\mathcal{F}_{fs\_aead}, sid.fs)$ :

1. Initialize  $keys\_in\_transit = []$  and a simulated state object  $state_{pid}$ .
2. **If**  $IsCorrupt? = true$ , skip to step 4. **Else**, set  $IsCorrupt? = true$  and continue.
3. **If**  $state_{\mathcal{I}} = \perp$ :
  - (a) Initialize objects  $View_0, View_1$ , and dictionary records.
  - (b) Set  $View_0.curr\_msg\_num = View_1.curr\_msg\_num = 0$ .
  - (c) Set  $View_0.pending\_msgs = View_1.pending\_msgs = []$
4. **Else**, parse the objects  $View_0, View_1$ , and the dictionary records from  $state_{\mathcal{I}}$  and store them. Also parse and store any records  $msg\_num\_state_{aead}$ .
5. Add  $View_i.curr\_msg\_num$  and  $View_i.pending\_msgs$  to  $state_{pid}$ .
6. For each  $h = (epoch\_id, msg\_num, N), c, m \in leakage$ :
  - (a) Pick a new key  $k_{msg\_num} \xleftarrow{\$} \{0, 1\}^\lambda$  (where  $\lambda$  is the key length that parametrizes  $\mathcal{F}_{mKE}$ ).
  - (b) Append the tuple  $(msg\_num, k_{msg\_num})$  to  $keys\_in\_transit$ .
  - (c) If  $msg\_num\_state_{aead}$  doesn't exist, initialize  $msg\_num\_state_{aead} = \perp$ .
  - (d) Send  $(msg\_num\_state_{aead}, ReportState, pid, k_{msg\_num}, m)$  to  $\mathcal{A}$  on behalf of  $(\mathcal{F}_{aead}, sid.aead = ("aead", sid.fs, msg\_num))$ .
  - (e) On receiving a response  $(ReportState, pid, S)$ , add  $S$  to  $state_{pid}$ .
7. **If** there is a record  $(StopKeys, i, N)$  set  $stop\_request = (StopKeys, i, N)$ , **Else** set  $stop\_request = \perp$ .
8. Send  $(ReportState, i, keys\_in\_transit, View_0.curr\_msg\_num, View_1.curr\_msg\_num, stop\_request)$ . to  $\mathcal{A}$  on behalf of  $(\mathcal{F}_{mKE}, sid.mKE = ("mKE", sid.fs))$
9. On receiving a response  $(ReportState, i, state_{mKE})$  from  $\mathcal{A}$ , add it to the state  $state_{pid}$ .
10. Output  $(state_{\mathcal{I}}, ReportState, state_{pid})$  to  $(\mathcal{F}_{fs\_aead}, sid.fs)$ .

**Encrypt:** On receiving  $(state_{\mathcal{I}}, Encrypt, pid, N, m)$  from  $(\mathcal{F}_{fs\_aead}, sid.fs)$  do:

1. Let  $i$  be such that  $pid = pid_i$ .
2. Set  $msg\_num = View_i.curr\_msg\_num + 1$  and update  $View_i.curr\_msg\_num + = 1$ .
3. If  $msg\_num\_state_{aead}$  exists in  $state_{\mathcal{I}}$  then parse it otherwise initialize  $msg\_num\_state_{aead} = \perp$ .
4. Send a backdoor message  $(RetrieveKey, pid, msg\_num)$  to  $\mathcal{A}$  on behalf of  $(\mathcal{F}_{mKE}, sid.mKE = ("mKE", sid.fs))$  and await a response  $(RetrieveKey, pid, k)$ . On receiving the response, continue.
5. Send a backdoor message  $(msg\_num\_state_{aead}, ReportState, pid, k)$  to  $\mathcal{A}$  on behalf of  $\mathcal{F}_{aead}$  with  $sid.aead = ("aead", sid.fs, msg\_num)$  and await a response.
6. Send a backdoor message  $(msg\_num\_state_{aead}, Encrypt, pid, m, N)$  to  $\mathcal{A}$  on behalf of  $\mathcal{F}_{aead}$  with  $sid.aead = ("aead", sid.fs, msg\_num)$ .
7. On receiving a response  $(msg\_num\_state_{aead}, Encrypt, pid, c)$  from  $\mathcal{A}$ , store  $(\ell = |m|, pid, ready2decrypt = true), (c, N, 1)$  in  $records[msg\_num]$ .
8. Update or store for the first time the values  $View_0, View_1, records, msg\_num\_state_{aead}$ , in  $state_{\mathcal{I}}$
9. Output  $(state_{\mathcal{I}}, Encrypt, pid, msg\_num, N, c)$  to  $(\mathcal{F}_{fs\_aead}, sid.fs)$ .

(The rest of this simulator is in Fig. 33 on Page 79.)

Figure 32: Forward Secure Authenticated Encryption Simulator,  $\mathcal{S}_{fsaead}$



### $\mathcal{S}_{\text{fsaead}}$ Continued...

(This simulator begins in Fig. 33 on Page 79.)

**Authenticate:** On receiving  $(\text{state}_{\mathcal{I}}, \text{Authenticate}, \text{pid}, c, \text{msg\_num}, N)$  from  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid}.fs)$  do:

1. Let  $i$  be such that  $\text{pid} = \text{pid}_i$ .
2. Provide input  $()$
3. Parse the dictionary records from  $\text{state}_{\mathcal{I}}$ :
  - (a) **If**  $\text{records}[\text{msg\_num}] = \perp$  then output  $(\text{state}_{\mathcal{I}}, \text{Authenticate}, \text{pid}, c, \text{msg\_num}, N, \perp)$  to  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid}.fs)$ .
  - (b) **Else** get  $(\ell, \text{pid}, \text{ready2decrypt})$  and any record of the form  $(c, N, v)$  from  $\text{records}[\text{msg\_num}]$ .
4. **If**  $\text{ready2decrypt} = \text{false}$  in the retrieved record, output  $(\text{state}_{\mathcal{I}}, \text{Authenticate}, \text{pid}, c, \text{msg\_num}, N, \perp)$  to  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid}.fs)$ .
5. **If** there is no record  $(c, n, v)$ :
  - Send a backdoor message  $(\text{RetrieveKey}, \text{pid}, \text{msg\_num})$  to  $\mathcal{A}$  on behalf of  $(\mathcal{F}_{\text{mKE}}, \text{sid}.mKE = (\text{"mKE"}, \text{sid}.fs))$  and await a response  $(\text{RetrieveKey}, \text{pid}, k)$ .
  - **If**  $k = \perp$  then output  $(\text{state}_{\mathcal{I}}, \text{Authenticate}, \text{pid}, c, \text{msg\_num}, N, \perp)$  to  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid}.fs)$ .
  - **Else** a backdoor message  $(\text{msg\_num\_state\_aead}, \text{ReportState}, \text{pid}, k)$  to  $\mathcal{A}$  on behalf of  $\mathcal{F}_{\text{aead}}$  with  $\text{sid}.aead = (\text{"aead"}, \text{sid}.fs, \text{msg\_num})$  and await a response.
  - Send a backdoor message  $(\text{msg\_num\_state\_aead}, \text{Inject}, \text{pid}, c, N)$  to  $\mathcal{A}$  on behalf of  $\mathcal{F}_{\text{aead}}$  with  $\text{sid}.aead = (\text{"aead"}, \text{sid}.fs, \text{msg\_num})$ .
  - On obtaining the output  $(\text{msg\_num\_state\_aead}, \text{Inject}, \text{pid}, c, N, v)$ , append  $(c, N, v)$  to  $\text{records}[\text{msg\_num}]$  and update records in  $\text{state}_{\mathcal{I}}$ .
6. **If**  $v = 0$  then output  $(\text{state}_{\mathcal{I}}, \text{Inject}, \text{pid}, c, \text{msg\_num}, N, \perp)$  to  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid}.fs)$ . **Else** output  $(\text{state}_{\mathcal{I}}, \text{Authenticate}, \text{pid}, c, \text{msg\_num}, N, v)$  to  $(\mathcal{F}_{\text{fs\_aead}}, \text{sid}.fs)$ .

Figure 33:  $\mathcal{S}_{\text{fsaead}}$  Continued...

### Case 1: No Corruptions Have Occurred

We begin by examining the action of the simulator and internal code when all parties are honest and uncorrupted. We show via induction that the real and ideal worlds remain indistinguishable after each action.

**Encrypt.** In the real world,  $\Pi_{\text{fs\_aead}}$  outsources the encryption of each message to a separate  $\mathcal{F}_{\text{aead}}$  instance  $(\mathcal{F}_{\text{aead}}, \text{sid}.aead = (\text{"aead"}, \text{sid}.fs, \text{msg\_num}))$  that is parametrized by the its  $\text{msg\_num}$ . The functionality  $\mathcal{F}_{\text{aead}}$  then sends a  $(\text{RetrieveKey}, \text{pid})$  request to  $\Pi_{\text{mKE}}$  to check whether the message key seed for the message is available to party  $\text{pid}$ . If the message key is available, then the functionality  $\mathcal{F}_{\text{aead}}$  runs the internal code  $\mathcal{I}_{\text{aead}}$  with message length  $|m|$  to produce a ciphertext. Meanwhile, the functionality  $\mathcal{F}_{\text{mKE}}$  samples a key at random for each  $\text{msg\_num}$  and provides each key exactly once for each party (up to any **StopKeys** requests made by the parties).

In the ideal world,  $\mathcal{F}_{\text{fs\_aead}}$  makes the same checks as  $\mathcal{F}_{\text{mKE}}$  to make sure that the key for the message is available to party  $\text{pid}$ . It then runs  $\mathcal{I}_{\text{aead}}$  with message length  $|m|$  to choose the ciphertext, just like  $\mathcal{F}_{\text{aead}}$  does in the real world.

**Decrypt.** In the real world,  $\Pi_{\text{fs\_aead}}$  again uses the appropriate  $\mathcal{F}_{\text{aead}}$  instance for decrypting ciphertexts. In this honest case  $\mathcal{F}_{\text{aead}}$  will only ever output the originally encrypted message  $m$

or a failure request. Specifically, it will never output a new message  $m'$  that is different in any way to the originally encrypted message  $m$  for that `msg_num`. The authenticate encryption  $\mathcal{F}_{\text{aead}}$  sends a `(RetrieveKey, pid)` request to  $\Pi_{\text{mKE}}$  to check whether the message key seed is available to the requesting party. If  $\Pi_{\text{mKE}}$  says that the key is available and if the ciphertext  $c$  is exactly the one that  $\mathcal{F}_{\text{aead}}$  sent on encryption, then  $\mathcal{F}_{\text{aead}}$  outputs the message  $m$  to  $\Pi_{\text{fs\_aead}}$ . Otherwise, if the ciphertext is different, then  $\mathcal{F}_{\text{aead}}$  runs internal code  $\mathcal{I}_{\text{aead}}$  to decide whether the message authenticates. Depending on the output  $v$  of the internal code,  $\Pi_{\text{fs\_aead}}$  also outputs either the message or a failure notification.

In the ideal world,  $\mathcal{F}_{\text{fs\_aead}}$  again makes the same checks as  $\mathcal{F}_{\text{mKE}}$  to ensure that the decryption key would have been available to the requester. If the key is available, then to decide whether to successfully decrypt,  $\mathcal{F}_{\text{fs\_aead}}$  sends an `Authenticate` request to  $\mathcal{I}_{\text{fs\_aead}}$  just like  $\mathcal{F}_{\text{aead}}$  does in the real world.

## Case 2: One or Both Parties Have Been Corrupted

Next, we show that the real and ideal worlds remain indistinguishable at the moment of corruption, and also after subsequent encryption and decryption operations are executed.

**Corruption.** The protocol  $\Pi_{\text{fs\_aead}}$ , on receiving a `Corrupt` request from  $\Pi_{\text{SGNL}}$ , sends `Corrupt` to  $\mathcal{F}_{\text{mKE}}$ , which leaks to  $\mathcal{A}$  a list of message seed keys in transit as well as the current message number for the corrupted party and the chain key. On receiving a state  $S_{\text{mKE}}$  from  $\mathcal{A}$ ,  $\mathcal{F}_{\text{mKE}}$  reports a list of all messages still in transit and  $S_{\text{mKE}}$ . Then,  $\Pi_{\text{fs\_aead}}$  corrupts each  $\mathcal{F}_{\text{aead}}$  instance that has a pending message, each of which leaks its message and key seed to  $\mathcal{A}$ . When  $\mathcal{F}_{\text{aead}}$  gets a state  $S_{\text{msg\_num}}$  back from  $\mathcal{A}$ , it reports this to  $\Pi_{\text{fs\_aead}}$ .  $\Pi_{\text{fs\_aead}}$  then reports all of the states  $S_{\text{mKE}}$  and  $S_{\text{msg\_num}}$ 's along with the list of pending messages and current message number to  $\Pi_{\text{SGNL}}$ .

The ideal functionality, on receiving `Corrupt`, sends a `ReportState` request to  $\mathcal{S}_{\text{fs\_aead}}$  and includes a list of the ciphertexts and messages that are in transit. Since no message key seeds are used in the ideal world, the ideal functionality simply chooses random keys for any messages still in transit to the corrupted party (as  $\mathcal{F}_{\text{mKE}}$  would have done) and uses these keys to send the necessary `ReportState` requests to  $\mathcal{S}_{\text{mKE}}$  and  $\mathcal{S}_{\text{aead}}$ . The messages that need to be leaked are simply stored in  $\mathcal{F}_{\text{fs\_aead}}$ .

**Encrypt.** In the real world,  $\Pi_{\text{fs\_aead}}$  outsources the encryption of each message to a separate  $\mathcal{F}_{\text{aead}}$  instance ( $\mathcal{F}_{\text{aead}}, \text{sid.aead} = (\text{"aead"}, \text{sid.fs}, \text{msg\_num})$ ) that is parametrized by the its `msg_num`. The functionality  $\mathcal{F}_{\text{aead}}$  then sends a `(RetrieveKey, pid)` request to  $\Pi_{\text{mKE}}$  to check whether the message key seed for the message is available to party `pid`. If the message key is available, then the functionality  $\mathcal{F}_{\text{aead}}$  send the message  $m$  to the simulator  $\mathcal{S}_{\text{aead}}$  and allows it to choose the ciphertext  $c$ . Meanwhile, the functionality  $\mathcal{F}_{\text{mKE}}$  allows the simulator  $\mathcal{S}_{\text{mKE}}$  to choose the required key as long as a key for this `msg_num` hasn't already been retrieved and there isn't a `stopkeys` request for a number smaller than `msg_num`.

In the ideal world,  $\mathcal{F}_{\text{fs\_aead}}$  also allows  $\mathcal{S}_{\text{mKE}}$  to choose the required key under the same conditions. It then sends the message  $m$  to the simulator  $\mathcal{S}_{\text{aead}}$  to allow it to choose the ciphertext. It is straightforward to validate that these actions look just like in the real world.

**Decrypt.** In the real world,  $\Pi_{\text{fs\_aead}}$  again uses the appropriate  $\mathcal{F}_{\text{aead}}$  instance for decrypting ciphertexts. The authenticate encryption  $\mathcal{F}_{\text{aead}}$  sends a `(RetrieveKey, pid)` request to  $\Pi_{\text{mKE}}$  to check whether the message key seed is available to the requesting party. If  $\Pi_{\text{mKE}}$  says that the key isn't available which happens if the key has already been retrieved or if `(StopKeys, N)` was run run

for  $N < \text{msg\_num}$  then a failure notification is output. Otherwise, if the ciphertext  $c$  is exactly the one that  $\mathcal{F}_{\text{aead}}$  sent on encryption, then  $\mathcal{F}_{\text{aead}}$  outputs the message  $m$  to  $\Pi_{\text{fs\_aead}}$ . Otherwise, if the ciphertext is different, then  $\mathcal{F}_{\text{aead}}$  allows the simulator  $\mathcal{S}_{\text{aead}}$  decide what the ciphertext decrypts to. Successful decryption is allowed only once even in the case of corruption.

In the ideal world,  $\mathcal{F}_{\text{fs\_aead}}$  again makes the same checks as  $\mathcal{F}_{\text{mKE}}$  to ensure that the decryption key would have been available to the requester. If the key is unavailable then  $\mathcal{F}_{\text{fs\_aead}}$  outputs a failure notification. If the key is available and the ciphertext matches the one at encryption, then again  $\mathcal{F}_{\text{fs\_aead}}$  takes an executive decision and outputs the message  $m$  that was originally encrypted. Otherwise, to decide how to decrypt,  $\mathcal{F}_{\text{fs\_aead}}$  sends an `Inject` request to  $\mathcal{S}_{\text{fsaead}}$  just like  $\mathcal{F}_{\text{aead}}$  does in the real world.

## Both Cases 1 and 2

Finally, we must consider the remaining two methods within the API for forward-secure AEAD: the stop encrypting and stop decrypting methods. Our analysis of these two methods is the same whether any parties have been previously corrupted or not.

**Stop Encrypting** In the real world protocol, a `StopEncrypting` request causes  $\Pi_{\text{fs\_aead}}$  to send a `StopKeys` request to  $\mathcal{F}_{\text{mKE}}$  for the current message number. Then  $\mathcal{F}_{\text{mKE}}$  will note that the sending epoch is closed, thereby disallowing the sender (or adversary) from being able to encrypt more messages in the epoch.

In the ideal world functionality,  $\mathcal{F}_{\text{fs\_aead}}$  simply notes that the sender has deleted the ability to encrypt any more messages for the epoch, also preventing the sender from encrypting more messages for the epoch.

**Stop Decrypting** When  $\Pi_{\text{fs\_aead}}$  receives a `(StopDecrypting, N)` request from above, it sends `(StopKeys, N)` to the  $\mathcal{F}_{\text{mKE}}$  instance for the epoch. Then,  $\mathcal{F}_{\text{mKE}}$  makes a note that the epoch is closed for the receiver at message number  $N$ . This prevents the receiver (or adversary) from decrypting any messages past  $N$  for the epoch.

The ideal world functionality  $\mathcal{F}_{\text{fs\_aead}}$  simply mimics this behavior by directly marking all messages beyond  $N$  as inaccessible and returns control to  $\Pi_{\text{SGNL}}$ .

In conclusion,  $\mathcal{S}_{\text{fsaead}}$  gives an *exact* emulation of the real world for every action taken by `Env` and `A` in the ideal world, so the views of `Env` and `A` is identical to the real world. This completes the proof of Theorem 4.  $\square$

## 8 The Symmetric Ratchet: Realising $\mathcal{F}_{\text{mKE}}$

In this section, we prove that the message key exchange protocol  $\Pi_{\text{mKE}}$  associated with a particular epoch, realises the functionality  $\mathcal{F}_{\text{mKE}}$  for that epoch in the presence of the long-lived global epoch key exchange functionality  $\mathcal{F}_{\text{eKE}}$ . As outlined in Section 4, protocol  $\Pi_{\text{mKE}}$  (presented in Figure 28 on page 73) implements the symmetric keychain for the epoch specified in its session ID. This is done by obtaining the base key for the chain from  $\mathcal{F}_{\text{eKE}}$  and then extending the chain as needed. Importantly, the keys on the main symmetric chain are never directly given as output; rather the outputs are keys derived from the chain keys. This structure allows the key derivation function to only be a (length doubling) pseudorandom number generator.

Before reading the following proof that  $\Pi_{\text{mKE}}$  (Figure 34) realises  $\mathcal{F}_{\text{mKE}}$  in the presence of  $\mathcal{F}_{\text{eKE}}$ , one must understand that such a proof may not continue to hold in general when  $\mathcal{F}_{\text{eKE}}$  is realised by

some protocol  $\Pi_{eKE}$ . However, as explained in Section 2.2, one can instead prove that  $\Pi_{mKE}$  realises  $\mathcal{F}_{mKE}$  in the presence of  $\mathcal{F}_{eKE}^{\Pi_{eKE}}$  for some protocol  $\Pi_{eKE}$  that realises  $\mathcal{F}_{eKE}$ . This is the approach that we follow in this section.



Figure 34: The Message Key Exchange Protocol  $\Pi_{mKE}$

**Theorem 5** *Assume that  $PRG$  is a secure length-doubling pseudorandom generator. Then protocol  $\Pi_{mKE}$  UC-realizes  $\mathcal{F}_{mKE}$  in the presence of global functionalities  $\mathcal{F}_{lib}, \mathcal{F}_{DIR}, \mathcal{F}_{LTM}$ , as well as  $\mathcal{F}_{eKE}^{\Pi_{eKE}}$ , where  $\mathcal{F}_{eKE}^{\Pi_{eKE}} = (\mathcal{I}_{eKE}, \mathcal{S}_{eKE}, \mathcal{F}_{eKE})$ .*

**Proof:** We construct an ideal-process adversary  $\mathcal{S}_{mKE}$  in the figure on Fig. 35 that interacts with functionality  $\mathcal{F}_{mKE}$ . The objective of  $\mathcal{S}_{mKE}$  is to simulate the interactions that would take place between the environment and the protocol  $\Pi_{mKE}$  (in the presence of  $\mathcal{F}_{eKE} + \mathcal{S}_{eKE} + \mathcal{I}_{eKE}$ ), so that the views of the environment  $Env$  are computationally indistinguishable in the real and ideal scenarios.

### $\mathcal{S}_{\text{mKE}}$

The global session id  $\text{sid}$  encodes a ciphersuite, including the PRG (used by  $\Pi_{\text{mKE}}$ ), and the length  $\lambda$  of key seeds.

**ReportState:** On receiving  $(\text{ReportState}, i, \text{keys\_in\_transit}, \text{msg\_num}_i, \text{stop\_request})$  from  $(\mathcal{F}_{\text{mKE}}, \text{sid.mKE})$  do:

1. If  $\text{msg\_num}_i \neq 0$  or  $\text{stop\_request} \neq \perp$  output  $(\text{ReportState}, S = \perp)$  to  $(\mathcal{F}_{\text{mKE}}, \text{sid.mKE})$ .  
//Note that  $\mathcal{F}_{\text{mKE}}$  only gets corrupted if the sender has already initialized it. This means that the sender's  $\text{msg\_num}$  will never be 0.
2. **Else If**  $\text{msg\_num}_i \neq 0$  (i.e.  $\text{stop\_request} = \perp$ ), initialize  $\text{chain\_key} \xleftarrow{\$} \{0, 1\}^\lambda$ .  
//The chain key is selected at random unless the receiver is corrupted before retrieving any keys for the epoch, this is because later chain keys should be unrelated to the initial one due to the *PRG* property.
3. **Else** ( $\text{msg\_num}_i = 0$  and  $\text{stop\_request} \neq \perp$ ), send  $(\text{GetReceivingKey}, \text{epoch\_id})$  to  $\mathcal{F}_{\text{eKE}}^{\Pi_{\text{eKE}}}$ . On receiving a response  $(\text{GetReceivingKey}, \text{recv\_chain\_key})$  set  $\text{chain\_key} = \text{recv\_chain\_key}$ .  
// $(\mathcal{F}_{\text{eKE}}^{\Pi_{\text{eKE}}} = (\mathcal{S}_{\text{eKE}}, \mathcal{I}_{\text{eKE}}, \mathcal{F}_{\text{eKE}}), \mathcal{F}_{\text{lib}}, \mathcal{F}_{\text{DIR}}, \mathcal{F}_{\text{LTM}})$   
//If the receiver has not retrieved any keys, we get the chain key from  $\mathcal{F}_{\text{eKE}}^{\Pi_{\text{eKE}}}$  so that it matches the key the sender used in the real world.
4. Initialize a dictionary  $\text{seeds} = \{\}$ .
5. For  $(\text{msg\_num}, k) \in \text{keys\_in\_transit}$  such that  $\text{msg\_num} < \text{msg\_num}_i$ :
  - Store  $\text{seeds}[\text{msg\_num}] = k$
6. Let  $\text{curr\_msg\_num} = \text{msg\_num}_i$ , and let the  $\text{curr\_chain\_key} = \text{chain\_key}$ .
7. While there is some  $(\text{msg\_num}, k) \in \text{keys\_in\_transit}$  such that  $\text{msg\_num} > \text{msg\_num}_i$ :
  - $\text{curr\_msg\_num} += 1$
  - Let  $(\text{curr\_chain\_key}, \text{key\_seed}) = \text{PRG}(\text{curr\_chain\_key})$ .
  - Store  $\text{seeds}[\text{curr\_msg\_num}] = \text{key\_seed}$ .
  - Set  $\text{latest\_seed\_num} = \text{curr\_msg\_num}$ .
  - Delete  $(\text{msg\_num}, k)$  from  $\text{keys\_in\_transit}$ .
8. Delete local variable  $\text{curr\_msg\_num}$ .
9. Output  $(\text{ReportState}, S = \{\text{chain\_key}, \text{seeds}\})$ .

**RetrieveKey:** On receiving  $(\text{RetrieveKey}, \text{pid}, \text{msg\_num})$  from  $(\mathcal{F}_{\text{mKE}}, \text{sid.mKE})$  do:

//This happens only if  $\text{IsCorrupt?} = \text{true}$ . In particular,  $\mathcal{S}_{\text{mKE}}$  has already gotten a **ReportState** directive.

1. If  $\text{msg\_num} \leq \text{latest\_seed\_num}$  then output  $(\text{RetrieveKey}, \text{pid}, k = \text{seeds}[\text{msg\_num}])$  to  $(\mathcal{F}_{\text{mKE}}, \text{sid.mKE})$ .  
//The only keys not in this dictionary are keys that were already retrieved by both parties at the time of corruption.
2. If  $\text{msg\_num} > \text{latest\_seed\_num}$ , initialize  $j = \text{latest\_seed\_num} + 1$ .
3. While  $j < \text{msg\_num}$ :
  - $(\text{curr\_chain\_key}, \text{key\_seed}) \leftarrow \text{PRG}(\text{curr\_chain\_key})$ .
  - Store  $\text{seeds}[j] = \text{key\_seed}, j += 1$ .
4. Output  $(\text{RetrieveKey}, \text{pid}, k = \text{seeds}[\text{msg\_num}])$  to  $(\mathcal{F}_{\text{mKE}}, \text{sid.mKE})$ .

Figure 35: Message Key Exchange Simulator,  $\mathcal{S}_{\text{mKE}}$

When the simulator  $\mathcal{S}_{\text{mKE}}$  receives messages from  $\mathcal{F}_{\text{mKE}}$  in the ideal world, it takes the actions specified in the pseudocode on Fig. 35.

Note that the real world adversary sees no keys before a compromise. If the epoch has been compromised, the real world adversary gets all key seeds for messages that are in transit. The real world adversary also gains the ability to compute all future key seeds available to the party (if `StopKeys` has not been called).

At such a compromise, the ideal world functionality  $\mathcal{F}_{\text{mKE}}$  provides the simulator  $\mathcal{S}_{\text{mKE}}$  with all the message numbers for which pending keys should be stored in the party’s state, the party’s current message number, as well as a chain key. The provided chain key is chosen at random in most cases, except for the case in which the party has not retrieved a single key in the epoch. In this case, the functionality provides  $\mathcal{S}_{\text{mKE}}$  with the initial chain key that a party would retrieve from the combined  $\mathcal{F}_{\text{eKE}} + \mathcal{S}_{\text{eKE}} + \mathcal{I}_{\text{eKE}}$  for this epoch to provide indistinguishability from the real world. The simulator  $\mathcal{S}_{\text{mKE}}$  then produces the keyseeds that are “in transit” by sampling them uniformly at random; it computes the future key seeds honestly using the `chain_key` provided to it during compromise. Since  $\mathcal{S}_{\text{mKE}}$  generates future keys just like  $\Pi_{\text{mKE}}$  would and otherwise uses the keys produced at compromise, if the state produced by the simulator at compromise is indistinguishable from the real world, then this proof is complete.

Note that if a uniformly random input is run through a PRG and then the output of the PRG run through the PRG again – and chained like this polynomially many times, the tuple containing all the outputs is still pseudorandom and therefore indistinguishable from outputs chosen independently at random. Therefore, the state produced at compromise is indistinguishable from the real world.  $\square$

## 9 Authenticated Encryption for Single Messages: Realising $\mathcal{F}_{\text{aead}}$

As outlined in Section 4, the ideal authenticated encryption functionality  $\mathcal{F}_{\text{aead}}$  is realized by way of a specific symmetric authenticated encryption scheme, which obtains its secret key from  $\mathcal{F}_{\text{mKE}}$ , and provides security against adaptive corruptions in the programmable random oracle model (which is captured by way of  $\mathcal{F}_{\text{pRO}}$ ).

If corruptions were not adaptive, any authenticated encryption scheme would suffice here; in particular, there would have been no need to resort to the random oracle model. In fact, this would have remained true even if the overall corruption structure was adaptive, as long as the adversary does not learn the keys that correspond to messages that were sent to the corrupted party but not yet received.

However, asserting full-fledged security in our model requires coming up with a simulation process that first generates a ciphertext  $c$ , and is then given an arbitrary message  $m$  and asked to generate a key  $k$  such that  $\text{Dec}(k, c) = m$ . While such schemes exist, they require having a key that is longer than the total length of the messages encrypted with that key. Furthermore, impossibility holds even when authentication is not required: There do not exist adaptively secure encryption schemes in the plain model where the key is shorter than the message [50].

We circumvent this impossibility by resorting to the random oracle model (again, using ideas from [50]). Specifically, we employ a simple Encrypt-then-MAC scheme [43] where the encryption is simply a one-time-pad, and the random oracle is used to expand the key to the length needed for the MAC algorithm, plus the length of the message.

We note that when the message is shorter than the overall keylength minus the length of the MAC key, the above scheme is adaptively secure even in the plain model. Consequently, in situations where there is a known bound on the total length of messages sent in each epoch, our solution is

fully secure in the plain model.

Protocol  $\Pi_{\text{aead}}$  is presented in Figure 36 on page 85.

**$\Pi_{\text{aead}}$**

This protocol has a session id  $\text{sid.aead} = (\text{"aead"}, \text{sid.fs}, \text{msg\_num})$  where  $\text{sid.fs} = (\text{"fs\_aead"}, \text{sid} = (\text{sid}', \text{pid}_0, \text{pid}_1), \text{epoch\_id})$  and party id  $\text{pid}$ .

It uses a message authentication code (MAC, Verify) with key length  $\lambda$ .

**Encrypt:** On receiving  $(\text{Encrypt}, m, N)$  from  $(\Pi_{\text{fs\_aead}}, \text{sid.fs}, \text{pid})$ :

1. If this is not the first activation or  $\text{pid} \notin (\text{pid}_0, \text{pid}_1)$ , end the activation.
2. Provide input  $(\text{RetrieveKey}, \text{pid})$  to  $(\Pi_{\text{mKE}}, \text{sid.mKE}, \text{pid})$ .
3. Upon receiving output  $(\text{RetrieveKey}, k)$  from  $(\Pi_{\text{mKE}}, \text{sid.mKE}, \text{pid})$ :
  - Let  $\ell = |m| + \lambda$ .
  - Send  $(\text{HashQuery}, k, \ell)$  to  $\mathcal{F}_{\text{pRO}}$ .
  - Upon receiving the output  $(\text{HashQuery}, \text{msg\_key})$ , parse  $\text{msg\_key} = k_{\text{otp}} || k_{\text{mac}}$ , where the  $k_{\text{otp}}$  has length  $|m|$ , and  $k_{\text{mac}}$  has length  $\lambda$ .
  - Compute ciphertext  $c' = k_{\text{otp}} \oplus m$
  - Compute tag  $t = \text{MAC}(k_{\text{mac}}, (c', \text{sid.aead}, N))$ .
  - Finally, set  $c = (c', t)$ .
  - Delete  $\text{msg\_key}, k, m$ , and  $c$  and output  $(\text{Encrypt}, c)$  to  $(\Pi_{\text{fs\_aead}}, \text{sid.fs}, \text{pid})$ .
4. If the response from  $(\Pi_{\text{mKE}}, \text{sid.mKE}, \text{pid})$  is  $(\text{RetrieveKey}, \text{Fail})$ , then output  $(\text{Encrypt}, \text{Fail})$  to  $(\Pi_{\text{fs\_aead}}, \text{sid.fs}, \text{pid})$ .

**Decryption:** On receiving  $(\text{Decrypt}, c = (c', t), N)$  from  $(\Pi_{\text{fs\_aead}}, \text{sid.fs}, \text{pid})$ :

1. If this is not the first activation or  $\text{pid} \notin (\text{pid}_0, \text{pid}_1)$ , end the activation.
2. Provide input  $(\text{RetrieveKey}, \text{pid})$  to  $(\Pi_{\text{mKE}}, \text{sid.mKE}, \text{pid})$ .
3. Upon receiving output  $(\text{RetrieveKey}, k)$  from  $(\Pi_{\text{mKE}}, \text{sid.mKE}, \text{pid})$ :
  - Let  $\ell = |m| + \lambda$ .
  - Send  $(\text{HashQuery}, k, \ell)$  to  $\mathcal{F}_{\text{pRO}}$ .
  - Upon receiving the output  $(\text{HashQuery}, \text{msg\_key})$ , parse  $\text{msg\_key} = k_{\text{otp}} || k_{\text{mac}}$ , where the  $k_{\text{otp}}$  has length  $|m|$ , and  $k_{\text{mac}}$  has length  $\lambda$ .
  - If  $\text{Verify}(k_{\text{mac}}, t, (c', \text{sid.aead}, N)) \neq 1$ , then output  $(\text{Decrypt}, \text{Fail})$  to  $(\Pi_{\text{fs\_aead}}, \text{sid.fs}, \text{pid})$ .
  - Else (the tag is valid), compute message  $m = k_{\text{otp}} \oplus c'$ .
  - Delete  $\text{msg\_key}, k, m$ , and  $c$  and output  $(\text{Decrypt}, m)$  to  $(\Pi_{\text{fs\_aead}}, \text{sid.fs}, \text{pid})$ .
4. If the response from  $(\Pi_{\text{mKE}}, \text{sid.mKE}, \text{pid})$  is  $(\text{RetrieveKey}, \text{Fail})$ , then output  $(\text{Decrypt}, \text{Fail})$  to  $(\Pi_{\text{fs\_aead}}, \text{sid.fs}, \text{pid})$ .

**Corruption:** On receiving  $(\text{Corrupt})$  from  $(\Pi_{\text{fs\_aead}}, \text{sid.fs}, \text{pid})$ :  
//Note that the **Corrupt** interface is not part of the “real” protocol; it is only included for UC-modelling purposes.

1. Output  $(\text{Corrupt}, S = \perp)$  to  $(\Pi_{\text{fs\_aead}}, \text{sid.fs}, \text{pid})$ . //  $\Pi_{\text{aead}}$  has no persistent state.

Figure 36: The Authenticated Encryption with Associated Data Protocol,  $\Pi_{\text{aead}}$



**Theorem 6** Assuming the unforgeability of  $(\text{MAC}, \text{Verify})$ , protocol  $\Pi_{\text{aead}}$  UC-realizes the ideal functionality  $\mathcal{F}_{\text{aead}}$  in the presence of global functionalities  $\mathcal{F}_{\text{pRO}}, \mathcal{F}_{\text{lib}}, \mathcal{F}_{\text{DIR}}, \mathcal{F}_{\text{LTM}}$ , as well as  $\mathcal{F}_{\text{mKE}}^{\Pi_{\text{mKE}}}$ ,  $\mathcal{F}_{\text{eKE}}^{\Pi_{\text{eKE}}}$ , where  $\mathcal{F}_{\text{mKE}}^{\Pi_{\text{mKE}}} = (\mathcal{S}_{\text{mKE}}, \mathcal{F}_{\text{mKE}})$ , and  $\mathcal{F}_{\text{eKE}}^{\Pi_{\text{eKE}}} = (\mathcal{I}_{\text{eKE}}, \mathcal{S}_{\text{eKE}}, \mathcal{F}_{\text{eKE}})$ .

$\mathcal{I}_{\text{aead}}$

The internal code has a ciphersuite, including the MAC protocol  $(\text{MAC}, \text{Verify})$ , the encryption protocol  $\text{OTP}$ , and the length of the MAC keys  $\lambda$ .

**To compute**  $\mathcal{I}(\text{state}_{\mathcal{I}}, \text{Encrypt}, \text{pid}, |m|, N)$ :

1. Set  $\ell = |m|$  and  $m = 0^\ell$ .
2. Choose a random message key  $\text{msg\_key} \xleftarrow{\$} \{0, 1\}^{\ell+\lambda}$ .
3. Parse  $\text{msg\_key} = \text{k}_{\text{otp}} \parallel \text{k}_{\text{mac}}$  where  $|\text{k}_{\text{otp}}| = \ell$  and  $|\text{k}_{\text{mac}}| = \lambda$ .
4. Compute ciphertext  $c' = \text{k}_{\text{otp}} \oplus m$  and tag  $t = \text{MAC}(\text{k}_{\text{mac}}, (c', \text{sid.aead}, N))$ .
5. Finally, set  $c = (c', t)$ , and record  $(\text{msg\_key}, m, c, N)$  in  $\text{state}_{\mathcal{I}}$ .
6. Return  $(\text{state}_{\mathcal{I}}, \text{Encrypt}, \text{pid}, c)$ .

**To compute**  $\mathcal{I}(\text{state}_{\mathcal{I}}, \text{Authenticate}, \text{pid}, c, N)$ :

1. Parse  $c^* = (c', t')$
2. Let  $\ell = |c'|$ .
3. If there is no record  $(\text{msg\_key}, m, c = (c', t), N)$  in  $\text{state}_{\mathcal{I}}$  then end the activation.
4. Else, parse  $\text{msg\_key} = \text{k}_{\text{otp}} \parallel \text{k}_{\text{mac}}$  where  $|\text{k}_{\text{otp}}| = |m|$  and  $|\text{k}_{\text{mac}}| = \lambda$ .
5. If  $\text{Verify}(\text{k}_{\text{mac}}, t', (c', \text{sid.aead}, N)) = 0$ : set  $v = \perp$ .
6. Else  $(\text{Verify}(\text{k}_{\text{mac}}, t', (c', \text{sid.aead}, N)) = 1)$ : set  $v = m$ .
7. Return  $(\text{state}_{\mathcal{I}}, v)$ .

Figure 37: Internal adversarial code  $\mathcal{I}_{\text{aead}}$

**Proof:**

The ideal-process adversary  $\mathcal{S}_{\text{aead}}$  is presented in Figure 38 and the internal adversarial code  $\mathcal{I}_{\text{aead}}$  is presented in Figure 37. To demonstrate the validity of  $\mathcal{S}_{\text{aead}}$ , we show that no environment that has global access to  $\mathcal{F}_{\text{pRO}}, \mathcal{F}_{\text{mKE}}^{\Pi_{\text{mKE}}} = (\mathcal{S}_{\text{mKE}}, \mathcal{F}_{\text{mKE}}), \mathcal{F}_{\text{eKE}}^{\Pi_{\text{eKE}}} = (\mathcal{I}_{\text{eKE}}, \mathcal{S}_{\text{eKE}}, \mathcal{F}_{\text{eKE}}), \mathcal{F}_{\text{DIR}}, \mathcal{F}_{\text{LTM}}, \mathcal{F}_{\text{lib}}$ , can tell whether it is interacting with  $\Pi_{\text{aead}}$ , or else with  $\mathcal{F}_{\text{aead}}, \mathcal{I}_{\text{aead}},$  and  $\mathcal{S}_{\text{aead}}$ .

This is done as follows: We first argue that, conditioned on two bad events not happening, the simulation is perfect. Then, we show that the bad events happen with negligible probability. The first, **Forge**, is the event that the environment produces a verifying message tag pair  $(m', t')$  for a fresh message  $m'$ , when no party has been corrupted. The second, **Collision**, is the bad event that a message seed key provided by  $\mathcal{F}_{\text{mKE}}^{\Pi_{\text{mKE}}}$  collides with an input to the random oracle that was previously programmed by the environment or simulator.

### $\mathcal{S}_{\text{aead}}$

**At first activation:** Send  $(\mathcal{F}_{\text{aead}}, \mathcal{I})$  to  $\mathcal{F}_{\text{lib}}$ .

**On receiving**  $(\text{state}_{\mathcal{I}}, \text{ReportState}, \text{pid}, k, m^*)$  **from**  $(\mathcal{F}_{\text{aead}}, \text{sid.aead})$ :

1. Initialize simulated state  $\text{state}_{\text{pid}} = \{k\}$  and store the key  $k$ .
2. If  $k \neq \perp$  and there is a record  $(c = (c', t), \text{msg\_key}, m, N, b)$  in  $\text{state}_{\mathcal{I}}$ :
  - Parse  $\text{msg\_key} = k_{\text{otp}} || k_{\text{mac}}$  where  $|k_{\text{otp}}| = |m^*|$ .
  - Let  $k_{\text{otp}}^* = c' \oplus m^*$  and  $\text{msg\_key}^* = k_{\text{otp}}^* || k_{\text{mac}}$ .
  - Send a backdoor message  $(\text{Program}, k, \text{msg\_key}^*)$  to  $\mathcal{F}_{\text{PRO}}$ .
  - On receiving  $(\text{Program})$ , continue.
3. Output  $(\text{state}_{\mathcal{I}}, \text{ReportState}, \text{pid}, \text{state}_{\text{pid}})$  to  $(\mathcal{F}_{\text{aead}}, \text{sid.aead})$ .

**On receiving**  $(\text{state}_{\mathcal{I}}, \text{Encrypt}, \text{pid}, m, N)$ : //This is called in the case that either party has been corrupted and no ciphertext was released yet. In this case  $\mathcal{S}_{\text{aead}}$  simply runs the instructions for  $\Pi_{\text{aead}}$ .

1. If the stored key  $k = \perp$ , then end the activation.
2. Set  $\ell = |m|$ .
3. Send  $(\text{HashQuery}, k, \ell + \lambda)$  to  $\mathcal{F}_{\text{PRO}}$ . //Note that  $k$  is sent to  $\mathcal{S}_{\text{aead}}$  during  $\text{ReportState}$
4. Upon receiving the output  $(\text{HashQuery}, \text{msg\_key})$ , set  $\text{msg\_key} = k_{\text{otp}} || k_{\text{mac}}$  where  $|k_{\text{otp}}| = \ell$  and  $|k_{\text{mac}}| = \lambda$ .
5. Compute ciphertext  $c' = k_{\text{otp}} \oplus m$  and tag  $t = \text{MAC}(k_{\text{mac}}, (c', \text{sid.aead}, N))$ .
6. Finally, set  $c = (c', t)$  and record  $(\text{msg\_key}, m, c, N)$  in  $\text{state}_{\mathcal{I}}$ .
7. Output  $(\text{state}_{\mathcal{I}}, \text{Encrypt}, \text{pid}, c)$  to  $(\mathcal{F}_{\text{aead}}, \text{sid.aead})$ .

**On receiving**  $(\text{state}_{\mathcal{I}}, \text{Inject}, \text{pid}, c^*, N)$  **from**  $(\mathcal{F}_{\text{aead}}, \text{sid.aead})$ :

//This is called in the case that either party has been corrupted and no output was generated yet.  $\mathcal{S}_{\text{aead}}$  simply runs the instructions for  $\Pi_{\text{aead}}$ .

1. Parse  $c^* = (c', t')$
2. Let  $\ell = |c'|$ .
3. Send  $(\text{HashQuery}, k, \ell + \lambda)$  to  $\mathcal{F}_{\text{PRO}}$ . //Note that  $k$  is sent to  $\mathcal{S}_{\text{aead}}$  during  $\text{ReportState}$
4. Upon receiving the output  $(\text{HashQuery}, \text{msg\_key}')$ , continue.
5. Parse  $\text{msg\_key}' = k_{\text{otp}} || k_{\text{mac}}$  where  $|k_{\text{otp}}| = \ell$ ,  $|k_{\text{mac}}| = \lambda$ .
6. If  $\text{Verify}(k_{\text{mac}}, t', (c', \text{sid.aead}, N)) = 0$ : set  $v = \perp$ .
7. Else  $(\text{Verify}(k_{\text{mac}}, t', (c', \text{sid.aead}, N)) = 1)$ : compute  $v = k_{\text{otp}} \oplus c'$ .
8. Output  $(\text{state}_{\mathcal{I}}, \text{Inject}, \text{pid}, c, N, v)$  to  $(\mathcal{F}_{\text{aead}}, \text{sid.aead})$ .

Figure 38: Authenticated Encryption with Associated Data Simulator,  $\mathcal{S}_{\text{aead}}$

### MAC Forgery Game

We define MAC forgery in the standard chosen-plaintext setting:

1. A key  $k_{\text{mac}} \xleftarrow{\$} \mathcal{K}_{\text{mac}}$  is sampled uniformly.
2. The adversary  $\mathcal{A}$  is given access to the MAC oracle  $\text{MAC}(k_{\text{mac}}, \cdot)$  which on message  $m$  computes and outputs the tag  $t$  of that message under the MAC. Let  $Q$  denote the set of all queries that  $\mathcal{A}$  makes to  $\text{MAC}(k_{\text{mac}}, \cdot)$ .
3. The adversary  $\mathcal{A}$  then outputs  $(m', t')$
4. The event  $\text{Forge}_1$  occurs if and only if  $\text{Verify}(k_{\text{mac}}, m', t') = 1$  and  $m' \notin Q$ .

We will call a successful forgery event of this kind  $\text{Forge}$ .

### Second Forgery Game

Now we define a different type of forgery:

1. Two keys  $k_{\text{mac}}, k_{\text{mac}}' \xleftarrow{\$} \mathcal{K}_{\text{mac}}$  are sampled uniformly and independently.
2. The adversary  $\mathcal{A}$  is given access to the MAC oracle  $\text{MAC}(k_{\text{mac}}, \cdot)$  which on message  $m$  computes and outputs the tag  $t$  of that message under the MAC.
3. The adversary  $\mathcal{A}$  is given access to the Verify oracle  $\text{Verify}(k_{\text{mac}}', \cdot, \cdot)$  which on input  $(m', t')$  outputs whether the message tag pair verifies.
4. The adversary  $\mathcal{A}$  then outputs  $(m', t')$
5. The event  $\text{Forge}_2$  occurs if and only if  $\text{Verify}(k_{\text{mac}}', m', t') = 1$ .

We will call a successful forgery event of this kind  $\text{Forge}_2$ .

This experiment models the setting where the two parties have diverged and thus have independent MAC keys. Lastly, we define  $\text{Collision}$  to be the event that  $\text{Env}$  already programmed  $\mathcal{F}_{\text{pRO}}$  on input  $k$ .

**Lemma 18** *If neither  $\text{Forge}$  nor  $\text{Collision}$  events happen during the executions, then the simulation by  $\mathcal{S}_{\text{aead}}$  and  $\mathcal{I}_{\text{aead}}$  is perfect.*

**Proof:** Assume that the bad events  $\text{Forge}$  and  $\text{Collision}$  do not occur during the executions. We will prove that, for all environments  $\text{Env}$  and adversaries  $\mathcal{A}$ , the functionality  $\mathcal{F}_{\text{aead}}$  together with simulator  $\mathcal{S}_{\text{aead}}$  and internal adversarial code  $\mathcal{I}_{\text{aead}}$  perfectly simulates the real-world views of  $\text{Env}$  and  $\mathcal{A}$  when they interact with  $\Pi_{\text{aead}}$ .

**Encryption** Observe that encryption in the ideal world occurs exactly as in the real world, except that the  $\mathcal{I}_{\text{aead}}$  module in  $\mathcal{F}_{\text{aead}}$  chooses a random key under which to encrypt the all 0's message of the correct length using the one time pad to produce a ciphertext  $c$ . It then authenticates the produced ciphertext using the randomly chosen  $k_{\text{mac}}$ . Later, when a corruption occurs, the simulator can easily use the leaked message  $m$  to compute a key  $k_{\text{otp}} = c \oplus m$  that will decrypt the ciphertext to the correct message. It then uses the leaked key-seed and programs  $\mathcal{F}_{\text{pRO}}$  to output the key  $k_{\text{otp}} \parallel k_{\text{mac}}$  on this seed.

**Decryption** In the real world,  $\Pi_{\text{aead}}$  retrieves the message key seed from  $\Pi_{\text{mKE}}$ , and if the key is available,  $\Pi_{\text{aead}}$  queries  $\mathcal{F}_{\text{pRO}}$  to get the expanded  $\text{msg\_key}$ . If the tag  $t$  verifies, then it decrypts the ciphertext using  $\text{msg\_key}$ ; otherwise, it returns a failure message. If the adversary has compromised

the state of  $\Pi_{\text{aead}}$ , then  $\Pi_{\text{mKE}}$  for the same epoch is compromised as well, and  $\mathcal{A}$  will get the `key_seed` (which it can expand to inject ciphertexts that will authenticate).

In the ideal world,  $\mathcal{F}_{\text{aead}}$  retrieves the message key from  $\Pi_{\text{mKE}}$  and if the key is available,  $\mathcal{F}_{\text{aead}}$  returns the message  $m$  that it encrypted to  $c = (c', t)$  in the case that the other party asks to decrypt  $c$ . This case is identical to the real world, by definition. If, on the other hand,  $\mathcal{F}_{\text{aead}}$  gets a different ciphertext  $c^* \neq c$ , then  $\mathcal{F}_{\text{aead}}$  either runs the internal adversarial module  $\mathcal{I}_{\text{aead}}$  on  $(\text{Authenticate}, c^*)$  or calls the adversary on  $(\text{Inject}, c^*)$ .

If it is not corrupted then  $\mathcal{F}_{\text{aead}}$  runs the internal adversarial code to see if it wants to authenticate this ciphertext  $c^*$ , and waits for a response value. ( $\mathcal{I}_{\text{aead}}$  uses the  $k_{\text{mac}}$  it used during encryption to verify the tag in  $c^*$ ; if the tag fails to verify for  $c^*$ , then it returns  $v = \perp$ .) In the case that the epoch is not compromised,  $\mathcal{F}_{\text{aead}}$  will check that  $\mathcal{I}_{\text{aead}}$ 's returned value  $v \neq \perp$ . If  $v \neq \perp$ , then  $\mathcal{F}_{\text{aead}}$  will output the original message  $m$ . Assuming that  $\text{Forge}_1$  and  $\text{Forge}_2$  do not occur, the original  $m$  will be returned if and only if  $c$  is the input ciphertext.

In the case that the epoch is compromised,  $\mathcal{F}_{\text{aead}}$  calls  $\mathcal{S}_{\text{aead}}$  which uses the message key seed  $k$  that it received from  $\mathcal{F}_{\text{aead}}$  to query the random oracle  $\mathcal{F}_{\text{pRO}}$  to expand the key. It then verifies the tag and outputs the decryption of the ciphertext (which may be different from  $m$ ). In this case,  $\mathcal{F}_{\text{aead}}$  outputs the injected message from  $\mathcal{S}_{\text{aead}}$ . This is identical to the powers of the real-world adversary after a state compromise.

**Corruption** Notice that the protocol  $\Pi_{\text{aead}}$  has no persistent state besides its `sid` and whether it has been activated already. The message, ciphertext, tag, and keys are all deleted after its activation. Accordingly,  $\mathcal{S}_{\text{aead}}$  (and thus  $\mathcal{F}_{\text{aead}}$ ) returns no state upon corruption.

The main job of  $\mathcal{S}_{\text{aead}}$  on corruption is to equivocate on the ciphertext it provided during encryption by programming the random oracle. Importantly, encryption occurs at most *once* in  $\mathcal{F}_{\text{aead}}$  (and  $\Pi_{\text{aead}}$ ). Thus, the random oracle is programmed at most once, and since we are assuming that the bad event `Collide` (that  $\mathcal{F}_{\text{pRO}}$  was already programmed or queried on input  $k$ ) does not occur, the message seed key will be programmable.  $\mathcal{S}_{\text{aead}}$  receives the correct message  $m^*$  along with the key seed  $k$  from  $\mathcal{F}_{\text{aead}}$ , after which it computes the message key from the ciphertext  $c$  it generated and  $m^*$  (and re-uses the MAC key  $k_{\text{otp}}$ ) and programs these in  $\mathcal{F}_{\text{pRO}}$ .  $\square$

So, all that's left to argue is why `Forge` and `Collision` do not occur. Notice that before corruptions the message key seeds provided by  $F_{\text{mKE}}^{\Pi}$  are uniformly random, so event `Forge` corresponds to the standard security game for message authentication codes. In particular, assuming that we have an environment that successfully induces a `Forge` event with non-negligible probability, we can construct a forger against  $(\text{MAC}, \text{Verify})$ , breaking our assumption. Thus, `Forge` happens with negligible probability.

Next, we discuss the `Collision` event. Since the key seeds provided by  $F_{\text{mKE}}^{\Pi}$  are uniformly random and independent, and the space of inputs to  $\mathcal{F}_{\text{pRO}}$  is exponential in the security parameter  $\delta$ , the probability that any two of them collide is negligible in the security parameter. Furthermore, since the environment runs in polynomial time with respect to the security parameter  $\delta$ , the environment can only program a polynomial number of inputs. So, for a seed length that is  $\theta(\delta)$ , the probability that the adversary programs an input that causes a collision is roughly  $2^{\theta(\delta)} - \text{poly}(\delta)$ .  $\square$

## References

- [1] Alwen, J., Coretti, S., Dodis, Y.: The double ratchet: Security notions, proofs, and modularization for the Signal protocol. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part I.

- LNCS, vol. 11476, pp. 129–158. Springer, Heidelberg (May 2019)
- [2] Alwen, J., Coretti, S., Dodis, Y., Tselekounis, Y.: Security analysis and improvements for the IETF MLS standard for group messaging. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020, Part I. LNCS, vol. 12170, pp. 248–277. Springer, Heidelberg (Aug 2020)
  - [3] Alwen, J., Coretti, S., Dodis, Y., Tselekounis, Y.: Modular design of secure group messaging protocols and the security of MLS. In: Vigna, G., Shi, E. (eds.) ACM CCS 2021. pp. 1463–1483. ACM Press (Nov 2021)
  - [4] Alwen, J., Coretti, S., Jost, D., Mularczyk, M.: Continuous group key agreement with active security. In: Pass, R., Pietrzak, K. (eds.) TCC 2020, Part II. LNCS, vol. 12551, pp. 261–290. Springer, Heidelberg (Nov 2020)
  - [5] Badertscher, C., Canetti, R., Hesse, J., Tackmann, B., Zikas, V.: Universal composition with global subroutines: Capturing global setup within plain UC. In: Pass, R., Pietrzak, K. (eds.) TCC 2020, Part III. LNCS, vol. 12552, pp. 1–30. Springer, Heidelberg (Nov 2020)
  - [6] Badertscher, C., Hesse, J., Zikas, V.: On the (ir)replaceability of global setups, or how (not) to use a global ledger. In: Nissim, K., Waters, B. (eds.) TCC 2021, Part II. LNCS, vol. 13043, pp. 626–657. Springer, Heidelberg (Nov 2021)
  - [7] Balli, F., Rösler, P., Vaudenay, S.: Determining the core primitive for optimally secure ratcheting. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020, Part III. LNCS, vol. 12493, pp. 621–650. Springer, Heidelberg (Dec 2020)
  - [8] Bellare, M., Rogaway, P.: Entity authentication and key distribution. In: Stinson, D.R. (ed.) CRYPTO’93. LNCS, vol. 773, pp. 232–249. Springer, Heidelberg (Aug 1994)
  - [9] Bellare, M., Rogaway, P.: Provably secure session key distribution: The three party case. In: 27th ACM STOC. pp. 57–66. ACM Press (May / Jun 1995)
  - [10] Bellare, M., Singh, A.C., Jaeger, J., Nyayapati, M., Stepanovs, I.: Ratcheted encryption and key exchange: The security of messaging. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part III. LNCS, vol. 10403, pp. 619–650. Springer, Heidelberg (Aug 2017)
  - [11] Bienstock, A., Dodis, Y., Rösler, P.: On the price of concurrency in group ratcheting protocols. In: Pass, R., Pietrzak, K. (eds.) TCC 2020, Part II. LNCS, vol. 12551, pp. 198–228. Springer, Heidelberg (Nov 2020)
  - [12] Bienstock, A., Fairuze, J., Garg, S., Mukherjee, P., Raghuraman, S.: A more complete analysis of the signal double ratchet algorithm. Cryptology ePrint Archive, Report 2022/355 (2022), <https://eprint.iacr.org/2022/355>
  - [13] Bienstock, A., Fairuze, J., Garg, S., Mukherjee, P., Raghuraman, S.: A more complete analysis of the Signal double ratchet algorithm. In: Dodis, Y., Shrimpton, T. (eds.) CRYPTO 2022, Part I. LNCS, vol. 13507, pp. 784–813. Springer, Heidelberg (Aug 2022)
  - [14] Blazy, O., Bossuat, A., Bultel, X., Fouque, P., Onete, C., Pagnin, E.: SAID: reshaping signal into an identity-based asynchronous messaging protocol with authenticated ratcheting. In: EuroS&P. pp. 294–309. IEEE (2019)

- [15] Boneh, D., Waters, B.: Constrained pseudorandom functions and their applications. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013, Part II. LNCS, vol. 8270, pp. 280–300. Springer, Heidelberg (Dec 2013)
- [16] Borisov, N., Goldberg, I., Brewer, E.A.: Off-the-record communication, or, why not to use PGP. In: WPES. pp. 77–84. ACM (2004)
- [17] Caforio, A., Durak, F.B., Vaudenay, S.: Beyond security and efficiency: On-demand ratcheting with security awareness. In: Garay, J. (ed.) PKC 2021, Part II. LNCS, vol. 12711, pp. 649–677. Springer, Heidelberg (May 2021)
- [18] Camenisch, J., Drijvers, M., Gagliardoni, T., Lehmann, A., Neven, G.: The wonderful world of global random oracles. Cryptology ePrint Archive, Report 2018/165 (2018), <https://eprint.iacr.org/2018/165>
- [19] Camenisch, J., Drijvers, M., Lehmann, A.: Universally composable direct anonymous attestation. In: Cheng, C.M., Chung, K.M., Persiano, G., Yang, B.Y. (eds.) PKC 2016, Part II. LNCS, vol. 9615, pp. 234–264. Springer, Heidelberg (Mar 2016)
- [20] Champion, S., Devigne, J., Duguey, C., Fouque, P.A.: Multi-device for signal. In: Conti, M., Zhou, J., Casalicchio, E., Spognardi, A. (eds.) ACNS 20, Part II. LNCS, vol. 12147, pp. 167–187. Springer, Heidelberg (Oct 2020)
- [21] Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: 42nd FOCS. pp. 136–145. IEEE Computer Society Press (Oct 2001)
- [22] Canetti, R.: Universally composable security. *J. ACM* 67(5), 28:1–28:94 (2020)
- [23] Canetti, R., Dodis, Y., Pass, R., Walfish, S.: Universally composable security with global setup. In: Vadhan, S.P. (ed.) TCC 2007. LNCS, vol. 4392, pp. 61–85. Springer, Heidelberg (Feb 2007)
- [24] Canetti, R., Halevi, S., Herzberg, A.: Maintaining authenticated communication in the presence of break-ins. *J. Cryptol.* 13(1), 61–105 (2000)
- [25] Canetti, R., Jain, P., Swanberg, M., Varia, M.: Universally composable end-to-end secure messaging. In: Dodis, Y., Shrimpton, T. (eds.) CRYPTO 2022, Part II. LNCS, vol. 13508, pp. 3–33. Springer, Heidelberg (Aug 2022)
- [26] Canetti, R., Krawczyk, H.: Analysis of key-exchange protocols and their use for building secure channels. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 453–474. Springer, Heidelberg (May 2001)
- [27] Canetti, R., Krawczyk, H.: Universally composable notions of key exchange and secure channels. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 337–351. Springer, Heidelberg (Apr / May 2002)
- [28] Canetti, R., Shahaf, D., Vald, M.: Universally composable authentication and key-exchange with global PKI. In: Cheng, C.M., Chung, K.M., Persiano, G., Yang, B.Y. (eds.) PKC 2016, Part II. LNCS, vol. 9615, pp. 265–296. Springer, Heidelberg (Mar 2016)
- [29] Chase, M., Perrin, T., Zaverucha, G.: The signal private group system and anonymous credentials supporting efficient verifiable encryption. In: Ligatti, J., Ou, X., Katz, J., Vigna, G. (eds.) ACM CCS 2020. pp. 1445–1459. ACM Press (Nov 2020)

- [30] Chen, K., Chen, J.: Anonymous end to end encryption group messaging protocol based on asynchronous ratchet tree. In: Meng, W., Gollmann, D., Jensen, C.D., Zhou, J. (eds.) ICICS 20. LNCS, vol. 11999, pp. 588–605. Springer, Heidelberg (Aug 2020)
- [31] Cohn-Gordon, K., Cremers, C., Dowling, B., Garratt, L., Stebila, D.: A formal security analysis of the signal messaging protocol. In: EuroS&P. pp. 451–466. IEEE (2017)
- [32] Cohn-Gordon, K., Cremers, C., Dowling, B., Garratt, L., Stebila, D.: A formal security analysis of the signal messaging protocol. *Journal of Cryptology* 33(4), 1914–1983 (Oct 2020)
- [33] Cohn-Gordon, K., Cremers, C., Garratt, L., Millican, J., Milner, K.: On ends-to-ends encryption: Asynchronous group messaging with strong security guarantees. In: Lie, D., Mannan, M., Backes, M., Wang, X. (eds.) ACM CCS 2018. pp. 1802–1819. ACM Press (Oct 2018)
- [34] Cohn-Gordon, K., Cremers, C.J.F., Garratt, L.: On post-compromise security. In: Hicks, M., Köpf, B. (eds.) CSF 2016 Computer Security Foundations Symposium. pp. 164–178. IEEE Computer Society Press (2016)
- [35] Cremers, C., Fairoze, J., Kiesl, B., Naska, A.: Clone detection in secure messaging: Improving post-compromise security in practice. In: Ligatti, J., Ou, X., Katz, J., Vigna, G. (eds.) ACM CCS 2020. pp. 1481–1495. ACM Press (Nov 2020)
- [36] Durak, F.B., Vaudenay, S.: Bidirectional asynchronous ratcheted key agreement with linear complexity. In: Attrapadung, N., Yagi, T. (eds.) IWSEC 19. LNCS, vol. 11689, pp. 343–362. Springer, Heidelberg (Aug 2019)
- [37] Hashimoto, K., Katsumata, S., Kwiatkowski, K., Prest, T.: An efficient and generic construction for signal’s handshake (X3DH): Post-quantum, state leakage secure, and deniable. In: Garay, J. (ed.) PKC 2021, Part II. LNCS, vol. 12711, pp. 410–440. Springer, Heidelberg (May 2021)
- [38] Hofheinz, D., Rao, V., Wichs, D.: Standard security does not imply indistinguishability under selective opening. In: Hirt, M., Smith, A.D. (eds.) TCC 2016-B, Part II. LNCS, vol. 9986, pp. 121–145. Springer, Heidelberg (Oct / Nov 2016)
- [39] Jaeger, J., Stepanovs, I.: Optimal channel security against fine-grained state compromise: The safety of messaging. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part I. LNCS, vol. 10991, pp. 33–62. Springer, Heidelberg (Aug 2018)
- [40] Jost, D., Maurer, U., Mularczyk, M.: Efficient ratcheting: Almost-optimal guarantees for secure messaging. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part I. LNCS, vol. 11476, pp. 159–188. Springer, Heidelberg (May 2019)
- [41] Jost, D., Maurer, U., Mularczyk, M.: A unified and composable take on ratcheting. In: Hofheinz, D., Rosen, A. (eds.) TCC 2019, Part II. LNCS, vol. 11892, pp. 180–210. Springer, Heidelberg (Dec 2019)
- [42] Keybase blog: New cryptographic tools on keybase. <https://keybase.io/blog/crypto> (2020)
- [43] Krawczyk, H.: The order of encryption and authentication for protecting communications (or: How secure is ssl?). In: Kilian, J. (ed.) Advances in Cryptology - CRYPTO 2001, 21st Annual



- International Cryptology Conference, Santa Barbara, California, USA, August 19-23, 2001, Proceedings. Lecture Notes in Computer Science, vol. 2139, pp. 310–331. Springer (2001), [https://doi.org/10.1007/3-540-44647-8\\_19](https://doi.org/10.1007/3-540-44647-8_19)
- [44] Krawczyk, H.: Cryptographic extraction and key derivation: The HKDF scheme. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 631–648. Springer, Heidelberg (Aug 2010)
- [45] Krohn, M.: Zoom rolling out end-to-end encryption offering. <https://blog.zoom.us/zoom-rolling-out-end-to-end-encryption-offering/> (2020)
- [46] Marlinspike, M., Perrin, T.: The X3DH key agreement protocol. <https://signal.org/docs/specifications/x3dh/> (2016)
- [47] Martiny, I., Kaptchuk, G., Aviv, A.J., Roche, D.S., Wustrow, E.: Improving signal’s sealed sender. In: NDSS 2021. The Internet Society (Feb 2021)
- [48] Maurer, U.: Constructive cryptography - a primer (invited paper). In: Sion, R. (ed.) FC 2010. LNCS, vol. 6052, p. 1. Springer, Heidelberg (Jan 2010)
- [49] Maurer, U.: Constructive cryptography - A new paradigm for security definitions and proofs. In: TOSCA. Lecture Notes in Computer Science, vol. 6993, pp. 33–56. Springer (2011)
- [50] 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 (Aug 2002)
- [51] Open Whisper Systems: Technical information: Specifications and software libraries for developers. <https://signal.org/docs/> (2016)
- [52] Perrin, T.: The noise protocol framework. <https://noiseprotocol.org/noise.html> (2018)
- [53] Poettering, B., Rösler, P.: Towards bidirectional ratcheted key exchange. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part I. LNCS, vol. 10991, pp. 3–32. Springer, Heidelberg (Aug 2018)
- [54] Rösler, P., Mainka, C., Schwenk, J.: More is less: On the end-to-end security of group chats in signal, whatsapp, and threema. In: EuroS&P. pp. 415–429. IEEE (2018)
- [55] Rotem, L., Segev, G.: Out-of-band authentication in group messaging: Computational, statistical, optimal. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part I. LNCS, vol. 10991, pp. 63–89. Springer, Heidelberg (Aug 2018)
- [56] Singh, M.: Whatsapp is now delivering roughly 100 billion messages a day. <https://techcrunch.com/2020/10/29/whatsapp-is-now-delivering-roughly-100-billion-messages-a-day/> (2020)
- [57] Status: Private, secure communication. <https://status.im> (2022)
- [58] Sylo: Comms for the metaverse. <https://sylo.io> (2022)
- [59] Unger, N., Dechand, S., Bonneau, J., Fahl, S., Perl, H., Goldberg, I., Smith, M.: SoK: Secure messaging. In: 2015 IEEE Symposium on Security and Privacy. pp. 232–249. IEEE Computer Society Press (May 2015)

- [60] Unger, N., Goldberg, I.: Deniable key exchanges for secure messaging. In: Ray, I., Li, N., Kruegel, C. (eds.) ACM CCS 2015. pp. 1211–1223. ACM Press (Oct 2015)
- [61] Unger, N., Goldberg, I.: Improved strongly deniable authenticated key exchanges for secure messaging. PoPETs 2018(1), 21–66 (Jan 2018)
- [62] Vatandas, N., Gennaro, R., Ithurnburn, B., Krawczyk, H.: On the cryptographic deniability of the Signal protocol. In: Conti, M., Zhou, J., Casalicchio, E., Spognardi, A. (eds.) ACNS 20, Part II. LNCS, vol. 12147, pp. 188–209. Springer, Heidelberg (Oct 2020)
- [63] WhatsApp LLC: About end-to-end encryption. <https://faq.whatsapp.com/general/security-and-privacy/end-to-end-encryption/> (2021)
- [64] Yan, H., Vaudenay, S.: Symmetric asynchronous ratcheted communication with associated data. In: Aoki, K., Kanaoka, A. (eds.) IWSEC 20. LNCS, vol. 12231, pp. 184–204. Springer, Heidelberg (Sep 2020)