



Universally Composable End-to-End Secure Messaging

Ran Canetti, Palak Jain^(✉), Marika Swanberg, and Mayank Varia

Boston University, Boston, USA
{canetti,palakj,marikas,varia}@bu.edu

Abstract. We model and analyze the Signal end-to-end messaging protocol within the UC framework. In particular:

- We formulate an ideal functionality that captures end-to-end secure messaging, in a setting with 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 and obtain their entire internal states. In particular 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 that 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 under minimal hardness assumptions.

Our modeling introduces additional innovations that enable arguing about the security of Signal irrespective of the underlying communication medium, as well as secure composition of dynamically generated modules that share state. These features, together 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.

Two other features of our modeling are the treatment of fully adaptive corruptions, and making minimal use of random oracle abstractions. In particular, we show how to realize continuous key exchange in the plain model, while preserving security against adaptive corruptions.

1 Introduction

Secure communication, namely allowing Alice and Bob to exchange messages securely, over an untrusted communication channel, without having to trust any

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.

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 [14], the Signal protocol is currently used to transmit hundreds of billions of messages per day [49].

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, 13, 15, 17, 25–33, 35–37, 46–48, 52–55, 57]. 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. Standalone security analyses of the Signal protocol are not always sufficient to capture the security of an entire messaging ecosystem that includes (components of) 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., [31, 32, 40]). For example, the Signal protocol is combined with other cryptographic protocols in WhatsApp [56] to perform abuse reporting or Status [50] and Slyo [51] 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 [45], file sharing services like Keybase [38] (which performs less frequent ratcheting), and videoconferencing services like Zoom [39] (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 [37] that defines an abstract ratcheting service within the Constructive Cryptography framework [41, 42], and concurrent work by Bienstock et al. [12] that formulates an ideal functionality of the Signal protocol within the UC framework (see Sect. 1.5 for details). However, neither of these works give a modular decomposition of Signal into its basic components (as described in [44].)

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 [34].

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 [18, 19]. We focus on modeling Signal at the level specified in their documentation [44] (i.e., not limited to any single choice, of cipher suite), taking care to adhere to the abstractions within the specification.

We provide an ideal functionality, \mathcal{F}_{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 [44]. 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, we propose a new abstraction for Signal’s continuous key derivation module, which we call a Cascaded PRF-PRG (CPRFG), and we show that it 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 may be of independent interest.

The rest of the Introduction is organized as follows. Section 1.2 presents and motivates our formulation of \mathcal{F}_{SM} . Section 1.3 presents and motivates the formulation of the individual modules, and describes how these modules can be realized. Section 1.5 discusses related work.

1.2 On the Ideal Secure Messaging Functionality, \mathcal{F}_{SM}

We provide an ideal functionality \mathcal{F}_{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 enable Signal (or any protocol that realizes \mathcal{F}_{SM}) to be readily usable as a component within other protocols in security-preserving manner.

When a party asks to encrypt a message, \mathcal{F}_{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.

Simple User Interface. The above encapsulation and decapsulation requests are the only ways that a parent protocol interacts with \mathcal{F}_{SM} . In particular, the parent protocol is not required to keep state related to the session, such as epochs or sequence numbers. In addition to simplicity, this imparts the additional guarantee that a badly designed parent protocol cannot harm the security of a protocol realising \mathcal{F}_{SM} .

Abstracting Away Network Delivery. The fact that \mathcal{F}_{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, where the communication medium is modeled as part of the service provided by the protocol and the actual communication is abstracted away.

Furthermore, the fact that \mathcal{F}_{SM} returns to the parent protocol an actual string (that represents an idealized encapsulated message) allows the parent protocol to further process the string as needed, similarly to what’s done in existing systems.

Immediate Decryption. \mathcal{F}_{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}_{SM} to execute adversarially provided code, without enabling the adversary to prevent immediate fulfillment of a decapsulation request. See more details in Sect. 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}_{DIR} that models a public “bulletin board,” which 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}_{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}_{SM} . This modeling is what allows to tie the two participants of a session to long-term identities. Similarly to [16, 24], we treat these modules as incorruptible. It is stressed, however, that, following the Signal architecture, our realization of \mathcal{F}_{SM} calls the \mathcal{F}_{LTM} module of each party exactly once, at the beginning of the session.

Modelling Corruptions. Resilience to recurring but transient break-ins is one of the main design goals of Signal. We facilitate the exposition of these properties as follows. First, we model corruption as an instantaneous event where the adversary learns the entire state of the corrupted party.

The security guarantees for corruption and recovery are then specified as follows. When the adversary instructs \mathcal{F}_{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}_{SM} to decapsulate ciphertexts to any plaintext of its choice. This captures the fact that 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}_{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}_{SM} to decapsulate forged ciphertexts.

Resilience to Adaptive Corruptions. All the security guarantees provided by \mathcal{F}_{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}_{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 [15,31]. Consequently, \mathcal{F}_{SM} exhibits a similar behavior.

Second, as remarked in the Signal documentation [44], 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, and yet 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 [20].) 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}_{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}_{SM} remains forked indefinitely, without any additional corruptions.

1.3 Realizing \mathcal{F}_{SM} , Modularly

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 [44] 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}_{mKE} , \mathcal{F}_{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}_{SM} . The management code (specifically, protocols $\Pi_{\text{fs_aead}}$ and Π_{SGNL}), does not directly access any keying material. Indeed, these protocols realise their respective specifications, namely $\mathcal{F}_{\text{fs_aead}}$ and \mathcal{F}_{SM} , perfectly—see Theorems 1 and 2.

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 the message itself*. Thus, modular security analysis along the above partitioning to modules might initially appear to be impossible.

The critical observation that allows us to proceed with modular decomposition is that the continuous key exchange module (which in our modeling corresponds to \mathcal{F}_{eKE}) need not determine the authenticity of new epoch identifiers. Rather, this module is only tasked to assign a fresh pseudorandom secret key with each new epoch identifier, be it 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 at the management level.

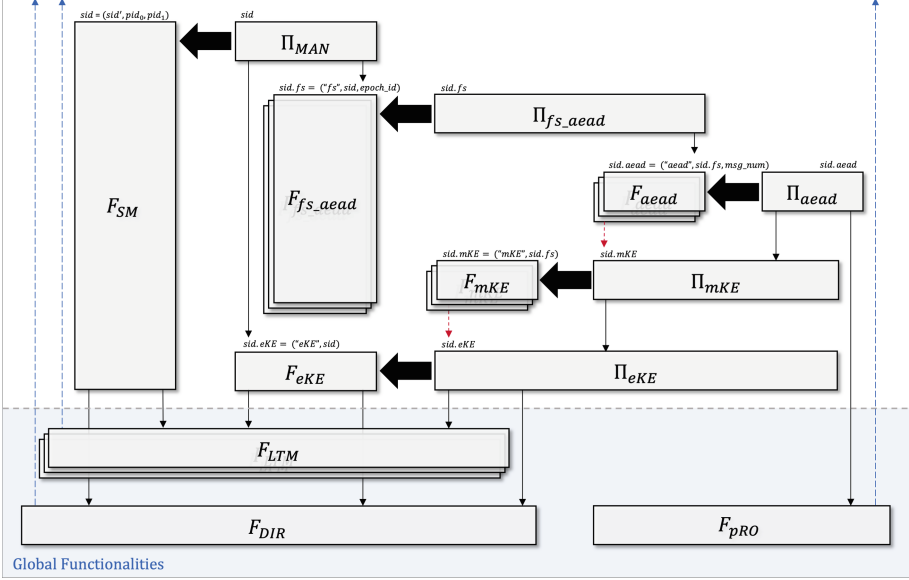


Fig. 1. Modeling and realizing secure messaging: The general subroutine structure. Ideal functionalities are denoted by F and protocols by Π . Thin vertical arrows denote subroutine calls, whereas thick horizontal arrows denote realization. Functionalities \mathcal{F}_{DIR} , \mathcal{F}_{LTM} , \mathcal{F}_{PRO} are global with respect to \mathcal{F}_{SM} , whereas \mathcal{F}_{eKE} (and Π_{eKE}) are global for \mathcal{F}_{mKE} , and each instance of \mathcal{F}_{mKE} (and the corresponding instance of Π_{mKE}) are global for \mathcal{F}_{aead} . (Color figure online)

We proceed to provide a more detailed overview of our partitioning and the general protocol logic. See also Fig. 1.

\mathcal{F}_{eKE} . The core component of the protocol is the epoch key exchange functionality \mathcal{F}_{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}_{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 keys provided by \mathcal{F}_{eKE} are guaranteed to appear random and independent to the adversary. In other words, \mathcal{F}_{eKE} leaves it to the receiver to determine which of the candidate identifiers for the new epoch is the correct one. (If \mathcal{F}_{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}_{eKE} to the end of this section.

\mathcal{F}_{mKE} . The per-epoch key chain is captured by an ideal functionality \mathcal{F}_{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}_{mKE} receives an instruction to end the chain for a party, it complies. \mathcal{F}_{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}_{mKE} is realized by a protocol, Π_{mKE} , that first calls \mathcal{F}_{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 generic length-doubling PRG (of which Signal’s typical instantiation using HKDF is a special case).

Demonstrating that Π_{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 Π_{mKE} . Using the formalism of [5], we thus show that Π_{mKE} UC-realizes \mathcal{F}_{mKE} in the presence of a global \mathcal{F}_{eKE} .

$\mathcal{F}_{\text{aead}}$. Authenticated encryption with associated data is captured by ideal functionality $\mathcal{F}_{\text{aead}}$, which provides a one-time ideal authenticated encryption service: the encrypting party calls $\mathcal{F}_{\text{aead}}$ with a plaintext and a recipient identity, and obtains an opaque ciphertext. Once the recipient presents the ciphertext, $\mathcal{F}_{\text{aead}}$ 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 $\mathcal{F}_{\text{aead}}$.

$\mathcal{F}_{\text{aead}}$ is realized via protocol Π_{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 [43]. 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}_{mKE} is used by multiple instances of Π_{aead} , we treat \mathcal{F}_{mKE} as a global functionality with respect to Π_{aead} . That is, we show that Π_{aead} UC-realizes $\mathcal{F}_{\text{aead}}$ in the presence of (a global) \mathcal{F}_{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 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}_{mKE} for this epoch - where, again, the session ID of \mathcal{F}_{mKE} contains the current epoch ID.

Π_{SGNL} . At the highest level of abstraction, we have each of the two parties run protocol Π_{SGNL} . When initiating a session, or starting a new epoch within a session, (i.e., when encapsulating the first message in an epoch), Π_{SGNL} first calls \mathcal{F}_{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 Π_{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 Π_{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 actual sender. Getting under the hood, this happens since the instances of \mathcal{F}_{mKE} that correspond to different epoch IDs generate keys that are mutually pseudorandom.) IT is stressed that Π_{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}_{SM} perfectly.

Realizing \mathcal{F}_{eKE} . Recall that \mathcal{F}_{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}_{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 obtains 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}_{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}_{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 [43] applies.

We propose a middle-ground solution: we show how to instantiate the KDF via a plain-model primitive where the local state grows linearly with the number of keys requested from \mathcal{F}_{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. We also abstract the properties of our construction into a primitive which we call *cascaded pseudorandom function and generator (CPRFG)*, following a primitive of [1] that is used for a similar purpose. We stress however that technically the primitives are quite different; we elaborate in the related work section.

1.4 Streamlining UC Analysis

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, we use UCGS to model a global directory that holds the public keys of parties, as well as the long-term storage, within each party, of the secret keys associated with said public keys. Similarly, we use UCGS to model the fact that a single instance of \mathcal{F}_{eKE} is used by multiple instances of Π_{mKE} , and that a single instance of \mathcal{F}_{mKE} is used by multiple instances of Π_{aead} . The random oracle is also modeled as a global functionality.

To facilitate our multi-layer use of the UCGS theorem we also prove a simple-but-useful lemma that allows us to get around the following difficulty. Recall that the UCGS theorem allows demonstrating that protocol π UC-realizes functionality \mathcal{F} in the presence of some other ‘global’ functionality \mathcal{G} that takes inputs from π , \mathcal{F} , and also potentially directly from the environment. Furthermore, we would like to use *multiple levels* of UCGS: after showing that π UC-realizes \mathcal{F} in the presence of \mathcal{G} , we wish to argue that π UC-realizes \mathcal{F} in the presence of protocol γ , where γ is some protocol that UC-realizes \mathcal{G} . However, such implication is not true in general [6, 24].

Lemma 1 in Sect. 2 asserts that, if γ UC-realizes \mathcal{G} via some simulator \mathcal{S} , then for any π that UC-realizes \mathcal{F} in the presence of \mathcal{G} , it also holds that π UC-realizes \mathcal{F} in the presence of $\mathcal{G}_{\mathcal{S}}$, where $\mathcal{G}_{\mathcal{S}}$ be the functionality that combines \mathcal{G} and \mathcal{S} in the natural way. We then show that, for the protocols in this work, having access to $\mathcal{G}_{\mathcal{S}}$ suffices.

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 corruption events (where typically the internal states of multiple modules are exposed together) rather complex. We thus adopt a somewhat simpler modeling of party corruption: We let the environment directly corrupt parties and obtain their local states. Furthermore, a corrupted module forwards the corruption notice to all its subroutines and collects the local states of all to report to the environment. Ideal functionalities operate similarly, except that they ask their respective simulators for the appropriate simulated local states. In addition to being simpler, this modeling provide 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, 22, 23]) and self-healing after corruption (e.g., [20, 28, 30]). 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, 32, 36, 57]). Other works consider a strong threat model in which the adversary is malicious, fully adaptive, and can tamper with local state [4, 7, 35, 37, 46], 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. They also propose and analyze an enhancement of the double ratchet structure that helps parties regain security faster following a compromise event. Additionally, Jost, Maurer, and Mularczyk

[37] 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 [37] differ from our \mathcal{F}_{SM} in a number of ways. First, 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. Second, 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.

Additionally, Bienstock et al. [12] models all key derivation modules as random oracles rather than formalizing the partition of continuous key exchange components within the UC framework as done in this work, and their modeling forces the “calling protocol” to keep track of—and ensure uniqueness for—the message IDs for the Secure Messaging functionality/protocol, which might create a security risk. On the other hand, [12] accounts for adversarial choice of randomness, which our modeling does not account for. Also, Jost et al. [37] requires explicit modeling of a global event history (a list of events having happened at each module), restricts the real-world adversary’s events based on this global event history, and employs a HIBE-based implementation that is quite different than that of Signal (and ours) and requires heavier cryptographic primitives.

2 Universally Composable Security: New Capabilities

This work makes extensive use of UC with global subroutines, which allow analysis that a protocol Π UC-realizes functionality \mathcal{F} in the presence of a global subroutine G that is not subroutine-respecting. Due to space constraints, we defer a primer of UC security (with global subroutines) to the full version of this work [21]. In this space, we describe two new modeling techniques that simplify our analysis, and may be of more general interest.

The first technique relates to applying the UC theorem to global functionalities. As stated in Sect. 1.4, the analysis in this work requires the ability to apply composition with global state across multiple layers of Fig. 1. We prove the following lemma in the full version of this work [21].

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

The second technique is that, in order to model the immediate encryption and immediate decryption properties of secure messaging, we require the adversary to upload static code (which we call \mathcal{I}) to the global \mathcal{F}_{lib} —shown in Fig. 2—that the relevant functionality will run in honest cases of the execution. This static code is specific to the protocol that realizes the functionality, and essentially it

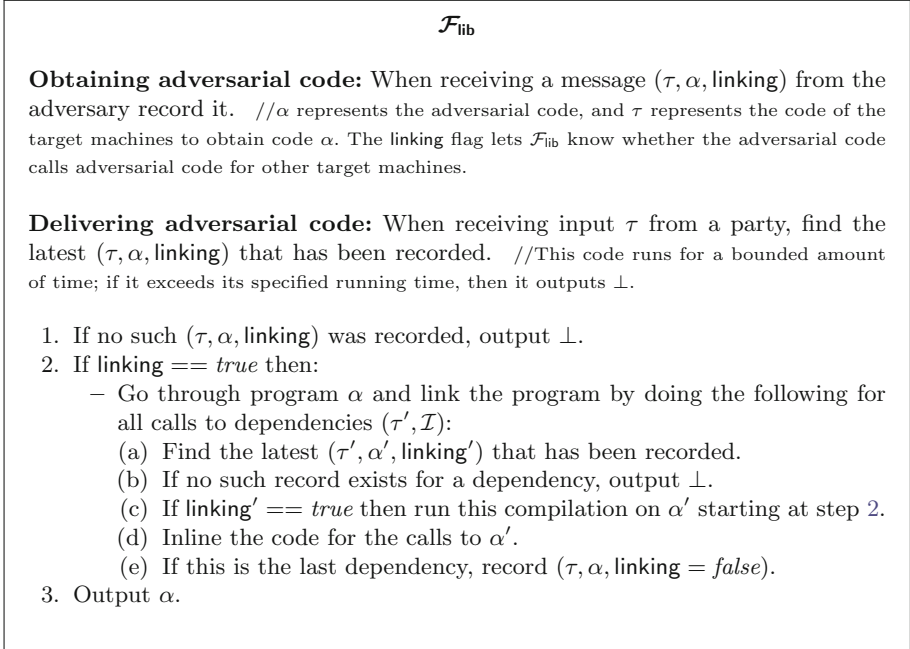


Fig. 2. The code library functionality, \mathcal{F}_{lib}

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 `state \mathcal{I}` , and it is sent to the adversary upon corruption. \mathcal{F}_{lib} is global because the static code must be defined at the time that the functionality is instantiated; however, the use of \mathcal{F}_{lib} specifically is mainly a matter of plumbing rather than a topic of conceptual importance.

3 Formal Modeling and Analysis

In this section, we showcase our modular, iterative process for decomposing the ideal secure messaging functionality \mathcal{F}_{SM} into a collection of functionalities and protocols that each address one specific purpose. After fully specifying \mathcal{F}_{SM} itself, we present its realization at the “second level” of Fig. 1 by Π_{SGNL} , $\mathcal{F}_{\text{fs_aead}}$, and \mathcal{F}_{eKE} , and at the “third level” by functionalities $\mathcal{F}_{\text{aead}}$ and \mathcal{F}_{mKE} .

Due to space constraints, we only give brief descriptions of these functionalities here and defer more exposition to the full version of this work [21]. The full version also contains rigorous specifications of the remaining protocols (on

the far right of Fig. 1) and all underlying global functionalities (in blue at the bottom of Fig. 1), along with proofs of all theorems in this section.

Secure Messaging Functionality. Our top-level functionality \mathcal{F}_{SM} can be found in Figs. 3 to 4. It takes two types of inputs: `SendMessage` is used to encapsulate a message for sending to the peer, whereas `ReceiveMessage` is used to decapsulate a received message. We also have a `Corrupt` input; this is a ‘modeling input’ that is used to capture party corruption. In addition, \mathcal{F}_{SM} takes a number of ‘side channel’ messages from the adversary which are used to fine-tune the security guarantees. It relies on three global functionalities whose specifications are provided in the full version of this work [21]: \mathcal{F}_{DIR} representing the directory of public keys, \mathcal{F}_{LTM} representing the long-term key storage within a party, and a programmable random oracle \mathcal{F}_{pRO} .

The Double Ratchet. In our first layer, we decompose \mathcal{F}_{SM} into two components that model the interconnected pieces of the double ratchet: a public key exchange component \mathcal{F}_{eKE} and a symmetric key authenticated encryption component $\mathcal{F}_{\text{fs_aead}}$. These components are ‘glued’ together with a manager protocol Π_{SGNL} .

There are three primary takeaways from the design of Π_{SGNL} (Fig. 5): it has the same input-output API as our ideal functionality \mathcal{F}_{SM} , it displays a idealized version of the double ratchet with clearly distinct roles for the two ratcheting sub-routines, and finally it moves closer toward realism. 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.

The epoch key exchange functionality \mathcal{F}_{eKE} (Figs. 6 to 7) comprises the public key “backbone” of the secure messaging continuous key agreement. The functionality is persistent during the entire session, mapping $(\text{epoch_id}_0, \text{epoch_id}_1)$ pairs to sending and receiving chain keys for the symmetric ratchet. It also provides recovery from a state compromise (aka, post-compromise security).

The forward secure authenticated encryption functionality $\mathcal{F}_{\text{fs_aead}}$ (Figs. 8 to 9) models the symmetric key ratchet for secure messaging. Each $\mathcal{F}_{\text{fs_aead}}$ instance handles the encryption and decryption of messages for a single epoch. The protocol $\Pi_{\text{fs_aead}}$ realizes $\mathcal{F}_{\text{fs_aead}}$ by outsourcing authenticated encryption and decryption of each message to separate $\mathcal{F}_{\text{aead}}$ instances, described below.

Theorem 1. *Protocol Π_{SGNL} (perfectly) UC-realizes the ideal functionality \mathcal{F}_{SM} in the presence of \mathcal{F}_{DIR} and \mathcal{F}_{LTM} .*

The Symmetric Ratchet: Realizing $\mathcal{F}_{\text{fs_aead}}$. Next, we decompose the symmetric key component of Signal into two smaller pieces: a one-time-use authenticated encryption routine $\mathcal{F}_{\text{aead}}$, and a message key exchange functionality \mathcal{F}_{mKE} that interfaces with the epoch key exchange to produce the symmetric chain keys.

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

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 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.

It also has *internal adversary code* $\mathcal{I} = \mathcal{I}_{\text{sm}}$. We initialize the state for \mathcal{I} to be $\text{state}_{\mathcal{I}} = \perp$.

Sending messages: On receiving ($\text{SendMessage}, m$) from pid do: //Here pid is an extended identity of a machine.

1. If initialized not set do: //initialization
 - If $\text{pid} \neq \text{pid}_0$, end the activation. Otherwise, send ($\text{ConfirmRegistration}$) to $(\mathcal{F}_{\text{LTM}}, \text{pid})$.
 - Upon output ($\text{ConfirmRegistration}, t$) from \mathcal{F}_{LTM} , if $t = \text{Fail}$ end the activation. Else input (GetInitKeys) to \mathcal{F}_{DIR} .
 - Upon output ($\text{GetInitKeys}, \text{pid}_1, \text{ik}_1^{\text{pk}}, \text{rk}_1^{\text{pk}}, \text{ok}_1^{\text{pk}}$) from \mathcal{F}_{DIR} : if $\text{ok}_1^{\text{pk}} = \perp$, end the activation. Else:
 - Set $\text{initialized}, \text{epoch_num}_0 = 0, \text{sent_msgnum}_0 = 0, \text{rcv_msgnum}_0 = 0, N_{\text{self}_0} = 0, \text{diverge_parties} = \text{false}$.
 - Create the dictionaries $\text{advControl} = \{\}$, $\text{id_dict} = \{\}$, and $N_{\text{dict}} = \{\}$. Initialize $\text{advControl}[\text{epoch_num}_0] = \perp$ and $\text{advControl}[e] = \infty$ for all $e \geq 0$. //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.
 - Call \mathcal{F}_{ib} with input \mathcal{F}_{SM} to obtain the internal code \mathcal{I} .
2. Let i be such that $\text{pid} = \text{pid}_i$. Increment sent_msgnum_i by 1.
3. If $\text{leak} \in \text{advControl}[\text{epoch_num}_i]$ or $\text{diverge_parties} = \text{true}$: Send a backdoor message ($\text{state}_{\mathcal{I}}, \text{SendMessage}, \text{pid}, m$) to \mathcal{A} .
4. Else ($\text{leak} \notin \text{advControl}[\text{epoch_num}_i]$ and $\text{diverge_parties} = \text{false}$): Run $\mathcal{I}(\text{state}_{\mathcal{I}}, \text{SendMessage}, \text{pid}, |m|)$
5. Upon obtaining ($\text{state}'_{\mathcal{I}}, \text{SendMessage}, \text{pid}, \text{epoch_id}, c$) from \mathcal{A} or \mathcal{I} do:
 - Update $\text{state}_{\mathcal{I}} \leftarrow \text{state}'_{\mathcal{I}}$.
 - If $\text{sent_msgnum}_i = 1$: If epoch_id equals any of the keys in the dictionary id_dict then end the activation. Else record $\text{id_dict}[\text{epoch_id}] = \text{epoch_num}_i$.
 - Set $h = (\text{epoch_id}, \text{sent_msgnum}_i, N_{\text{self}_i})$. // N_{self_i} holds the # of messages sent by pid_i in its previous sending epoch.
 - If $\text{diverge_parties} = \text{false}$ then record (pid, h, c, m) . //If the parties' states have diverged, then encrypted messages are no longer recorded.
 - Output ($\text{SendMessage}, \text{sid}, \text{pid}, h, c$) to pid .

Corrupt: On receiving a ($\text{Corrupt}, \text{pid}_i$) request from Env for $\text{pid}_i \in \{\text{pid}_0, \text{pid}_1\}$, do:

1. Append $(\text{epoch_num}_i, \text{sent_msg_num}_i, \text{received_msg_num}_i)$ to the list corruptions_i .
2. For all epochs $e \leq \text{epoch_num}_i$, set $\text{advControl}[e] = \{\text{leak}, \text{inject}\}$ to allow the adversary to influence messages still in transit.
3. Create a list pending_msgs with all records of the form $(\text{pid}_{1-i}, h, c, m)$ corresponding to headers for which there is no record ($\text{Authenticate}, \text{pid}_{1-i}, h, -, 1$) (these are the messages that were not decrypted yet).
4. Send a request ($\text{state}_{\mathcal{I}}, \text{ReportState}, \text{pid}_i, \text{pending_msgs}$) to \mathcal{A} .
5. On receiving a state ($\text{ReportState}, \text{pid}_i, S$) from \mathcal{A} , send S to Env .

Fig. 3. The Secure Messaging Functionality \mathcal{F}_{SM}

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

The local session ID is parsed as $sid = (sid', pid_0, pid_1)$. 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.

It also has *internal adversary code* $\mathcal{I} = \mathcal{I}_{sm}$. We initialize the state for \mathcal{I} to be $state_{\mathcal{I}} = \perp$.

Receiving messages: On receiving (ReceiveMessage, $h = (\text{epoch_id}, \text{msg_num}, N)$, c) from pid , do:

1. Let i be such that $pid = pid_i$.
2. If this is the first ReceiveMessage request: If $i = 0$ then end the activation. Else ($pid = pid_1$), initialize the responder:
 - Send (ConfirmRegistration) to (\mathcal{F}_{LTM}, pid) .
 - Upon receiving the output (ConfirmRegistration, t) from \mathcal{F}_{LTM} : If $t = \text{Fail}$ then end activation. Else provide input (GetResponseKeys, pid_0, pid_1) to \mathcal{F}_{DIR} .
 - Upon receiving output (GetResponseKeys, pid_0, ik_0^{pk}) from \mathcal{F}_{DIR} , set $\text{epoch_num}_1 = 1$, $\text{sent_msgnum}_1 = 0$, and $\text{rcv_msgnum}_1 = 0$.
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, Fail) to pid .
4. If epoch_id appears as a key in id_dict , set $\text{epoch_num} = \text{id_dict}[\text{epoch_id}]$.
Else: //this is a new epoch id that hasn't been generated within SendMessage
 - If $\text{sent_msgnum}_i = 0$, output (ReceiveMessage, h, c, Fail) to pid . //pid is in a receiving state and hasn't sent any messages in its current sending epoch, so it should not be accepting messages with a new epoch id.
 - Otherwise set $\text{epoch_num} = \text{epoch_num}_i + 1$.
//this temporary variable will never be made permanent if decryption is unsuccessful.
5. If $\text{msg_num} > N_{\text{dict}}[\text{epoch_num}]$, output (ReceiveMessage, h, c, Fail) to pid
//For epoch_num 's that are not finished yet, the N_{dict} returns a default value of ∞ , so this check passes automatically.
6. If ($\text{diverge_parties} = \text{false}$ and $\text{inject} \notin \text{advControl}[\text{epoch_num}]$): Run $\mathcal{I}(\text{state}_{\mathcal{I}}, \text{inject}, pid, h, c)$ //Honest Case
7. Else: send backdoor message ($\text{state}_{\mathcal{I}}, \text{inject}, pid, h, c$) to \mathcal{A} // \mathcal{F}_{SM} is asking the adversary for advice on how to decrypt c .
8. On receiving ($\text{state}'_{\mathcal{I}}, \text{inject}, h, c, v$) from \mathcal{A} or \mathcal{I} update $\text{state}_{\mathcal{I}} \leftarrow \text{state}'_{\mathcal{I}}$ and:
If (sender, h, c, m) is recorded then record (Authenticate, $pid, h, c, 1$) and set $m^* = m$. Else:
 - If $v = \perp$: record (Authenticate, $pid, h, c, 0$) and output (ReceiveMessage, h, c, Fail).
 - If $v \neq \perp$ and $\text{diverge_parties} = \text{false}$ and $\text{inject} \notin \text{advControl}[\text{epoch_num}]$, then:
 - If there is no record (sender, h, c^*, m) for header h , output (ReceiveMessage, h, c, Fail). //since h contains N , this value will match the view of the sender if this check succeeds.
 - Else (there is such a record), record (Authenticate, $h, c, 1$) and set $m^* = m$.
//allowing for authenticating a message with a different mac
 - If $v \neq \perp$ and ($\text{diverge_parties} = \text{true}$ or $\text{inject} \in \text{advControl}[\text{epoch_num}]$), then:
 - Record (Authenticate, $h, c, 1$), and set $m^* = v$.
 - If epoch_id does not appear as a key in id_dict then set $\text{diverge_parties} = \text{true}$.
//diverge parties is being set here.
9. If $\text{epoch_num}_i < \text{epoch_num}$, do: //we only get to this step if decryption is successful
 - Set $N_{\text{dict}}[\text{epoch_num} - 2] = N$, $\text{epoch_num}_i += 2$, $N_{\text{self}_i} = \text{sent_msgnum}_i$, and $\text{sent_msgnum}_i = 0$.
 - if $\text{diverge_parties} = \text{false}$ then:
 - If $\text{advControl}[\text{epoch_num} - 1] = \{\text{leak}, \text{inject}\}$ and $\text{epoch_num}_i \notin \text{corruptions}_i$, then set $\text{advControl}[\text{epoch_num}] = \{\text{leak}\}$. //Corruption status is changed if this is the other party's first new sending epoch that involves a fresh epoch id generated after corruption.
 - If $\text{advControl}[\text{epoch_num} - 1] = \{\text{leak}\}$, then set $\text{advControl}[\text{epoch_num}] = \perp$.
10. Output (ReceiveMessage, h, m^*) to pid .

Fig. 4. The Secure Messaging Functionality \mathcal{F}_{SM}

Π_{SGNL}

SendMessage: Upon receiving input (**SendMessage**, m) from pid , do:

1. If this is the first activation do: //initialization for the initiator of the session
 - Parse the local session id sid to retrieve the party identifiers ($\text{pid}_0, \text{pid}_1$) for the initiator and responder. If pid_0 is different from either the local party identifier pid , or the party identifier of pid , end the activation.
 - Initialize $\text{epoch_id}_{\text{self}} = \perp$, $\text{epoch_id}_{\text{partner}} = \perp$, $\text{sent_msg_num} = 0$, $N_{\text{last}} = 0$.
 - Provide input (**ConfirmReceivingEpoch**, \perp) to $(\mathcal{F}_{\text{eKE}}, \text{sid.eKE})$.
 - On receiving (**ConfirmReceivingEpoch**, epoch_id) from $(\mathcal{F}_{\text{eKE}}, \text{sid.eKE})$, set $\text{epoch_id}_{\text{self}} = \text{epoch_id}$.
 - Initialize a list $\text{receiving_epochs} = []$.
2. Provide input (**Encrypt**, m, N_{last}) to $(\mathcal{F}_{\text{fs_aead}}, \text{sid.fs})$, where $\text{sid.fs} = (\text{sid}, \text{epoch_id}_{\text{self}})$.
// $\mathcal{F}_{\text{fs_aead}}$ already knows epoch_id and msg_num
3. On receiving (**Encrypt**, c, N_{last}) from $(\mathcal{F}_{\text{fs_aead}}, \text{sid.fs})$, delete m , increment $\text{sent_msg_num} + = 1$, output (**SendMessage**, sid, h, c) to pid , where $h = (\text{epoch_id}_{\text{self}}, \text{sent_msg_num}, N_{\text{last}})$.

ReceiveMessage: Upon receiving (**ReceiveMessage**, $h = (\text{epoch_id}, \text{msg_num}, N), c$) from pid :

1. If this is the first activation then do: //initialization for the responder of the session
 - Parse the local session identifier sid to retrieve the party identifiers ($\text{pid}_0, \text{pid}_1$) for the initiator and responder. If pid_1 is different from either the local party identifier, or the party identifier for pid , then end the activation.
 - Initialize $\text{epoch_id}_{\text{self}} = \perp$, $\text{epoch_id}_{\text{partner}} = \perp$, $\text{sent_msg_num} = 0$ 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 pid . //Otherwise, v is the decrypted message
4. While $\text{msg_num} > \text{received_msg_num}$:
//note down any expected messages
 - Append received_msg_num to the entry $\text{missed_msgs}[\text{epoch_id}]$.
 - Increment $\text{received_msg_num} + = 1$.
5. If 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 epoch_id from missed_msgs.keys .
6. Else ($\text{msg_num} \notin \text{missed_msgs}[\text{epoch_id}]$):
 - If $\text{epoch_id} = \text{epoch_id}_{\text{partner}}$ 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 $(\mathcal{F}_{\text{fs_aead}}, (\text{sid}, \text{epoch_id}_{\text{partner}}))$. //‘Closing’ the $\mathcal{F}_{\text{fs_aead}}$ for the last epoch.
 - On receiving (**StopDecrypting**, **Success**), update $\text{epoch_id}_{\text{partner}} = \text{epoch_id}$, and send (**StopEncrypting**) to $(\mathcal{F}_{\text{fs_aead}}, (\text{sid}, \text{epoch_id}_{\text{self}}))$.
 - On receiving (**StopEncrypting**, **Success**), send (**ConfirmReceivingEpoch**, epoch_id) to $(\mathcal{F}_{\text{eKE}}, \text{sid.eKE})$.
 - On receiving (**ConfirmReceivingEpoch**, epoch_id^*), update $\text{epoch_id}_{\text{self}} = \text{epoch_id}^*$, $N_{\text{last}} = \text{sent_msg_num}$, and $\text{sent_msg_num} = 0$.
7. Output (**ReceiveMessage**, h, c, v) to pid while deleting the decrypted message v .

Corruption: Upon receiving (**Corrupt**, 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_{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**, pid_i, S) to Env .

Fig. 5. The Signal Protocol, Π_{SGNL}

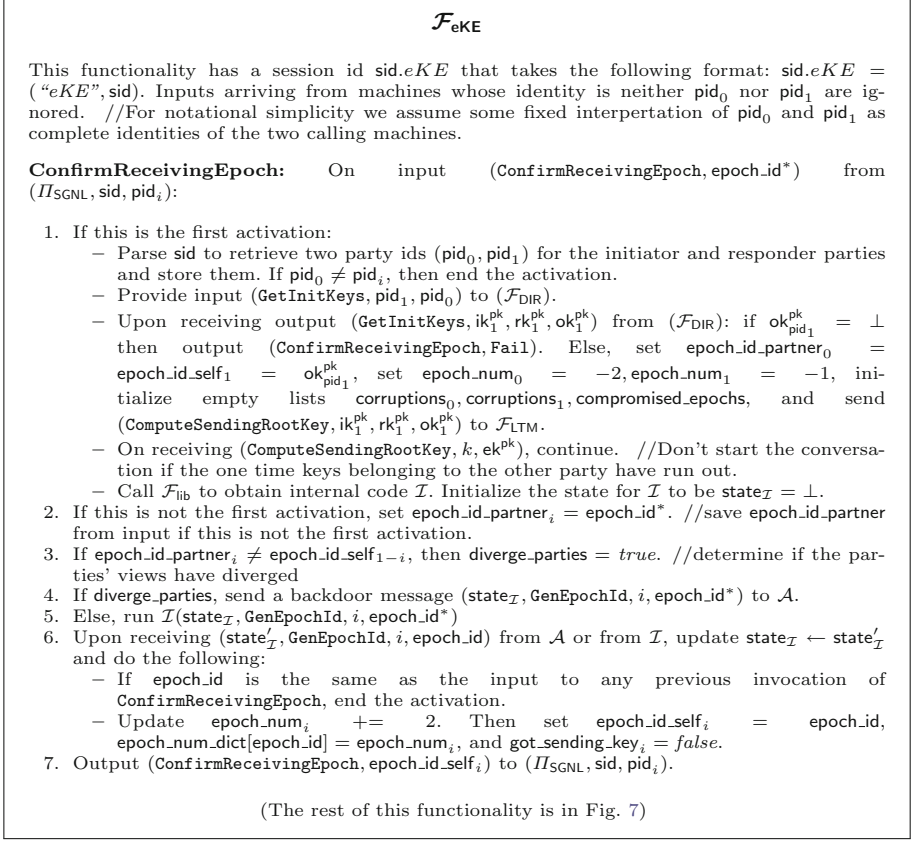


Fig. 6. The Epoch Key Exchange Functionality, \mathcal{F}_{eKE}

Each authenticated encryption functionality instance \mathcal{F}_{aead} (Fig. 10) handles the encryption, decryption, and authentication of a particular message for a particular epoch. It hands the ciphertext or message back to Π_{fs_aead} .

Theorem 2. *Protocol Π_{fs_aead} (perfectly) UC-realizes \mathcal{F}_{fs_aead} , in the presence of $\mathcal{F}_{DIR}, \mathcal{F}_{LTM}, \mathcal{F}_{PRO}$, and $\mathcal{F}_{eKE}^{\Pi_{eKE}} = (\mathcal{S}_{eKE}, \mathcal{F}_{eKE})$.*

Note that the simulator \mathcal{S}_{eKE} , along with a proof of this theorem, are deferred to the full version of this work [21].

Each instance of the message key exchange functionality \mathcal{F}_{mKE} (Fig. 11) handles the key derivation for the symmetric ratchet for a particular epoch. Specifically, it provides key_seed 's to Π_{aead} instances that are then expanded to any length using the global random oracle \mathcal{F}_{PRO} . When instructed, it also closes epochs at a certain message number N by generating all key_seed 's up to N and later disallowing the generation of any further key seeds for its epoch.

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

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

1. Set i such that $\text{pid} = \text{pid}_i$.
2. If **ConfirmReceivingEpoch** has never been run successfully (i.e. epoch_id_self_0 hasn't been initialized) or $\text{got_sending_key}_i = \text{true}$, 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 $\text{epoch_num}_i \in \text{compromised_epochs}$, then:
 - Send backdoor message $(\text{state}_{\mathcal{I}}, \text{GetSendingKey}, i)$ to \mathcal{A}
 - On receiving backdoor message $(\text{state}'_{\mathcal{I}}, \text{GetSendingKey}, i, K_{\text{send}})$ from \mathcal{A} , update $\text{state}_{\mathcal{I}} \leftarrow \text{state}'_{\mathcal{I}}$ and set $\text{sending_chain_key}_i = K_{\text{send}}$.
5. Set $\text{got_sending_key}_i = \text{true}$ and output $(\text{GetSendingKey}, \text{sending_chain_key}_i)$.

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

1. If $\text{pid} \notin \{\text{pid}_0, \text{pid}_1\}$ then end this activation. Otherwise, set i such that $\text{pid} = \text{pid}_i$.
2. If **ConfirmReceivingEpoch** has never been run successfully (i.e. epoch_id_self_0 hasn't been initialized) or $\text{sending_chain_key}_{1-i}$ has been deleted then end the activation.
3. If this is the first activation:
 - Initialize state variables $\text{root_key}, \text{epoch_id}, \text{epoch_key}, \text{sending_chain_key} = \perp$.
 - Parse $\text{epoch_id} = (\text{epoch_id}', \text{ek}_j^{\text{pk}}, \text{ok}_{i \leftarrow j}^{\text{pk}})$ and set $\text{temp_epoch_id_partner} = \text{epoch_id}'$
 - Send $(\text{GetResponseKeys}, \text{pid}_{1-i})$ to \mathcal{F}_{DIR} .
 - 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}_{LTM} .
 - Upon receiving $(\text{ComputeReceivingRootKey}, k)$, call \mathcal{F}_{lib} to obtain internal code \mathcal{I} . Initialize the state for \mathcal{I} to be $\text{state}_{\mathcal{I}} = \perp$.
4. If $\text{diverge_parties} = \text{true}$ or $\text{epoch_id} \neq \text{epoch_id_self}_{1-i}$: //Let \mathcal{A} choose key
 - Send $(\text{state}_{\mathcal{I}}, \text{GetReceivingKey}, i, \text{epoch_id})$ to \mathcal{A}
 - Upon receiving $(\text{state}'_{\mathcal{I}}, \text{GetReceivingKey}, i, \text{epoch_id}, \text{recv_chain_key}^*)$ from \mathcal{A} , update $\text{state}_{\mathcal{I}} \leftarrow \text{state}'_{\mathcal{I}}$.
 - If $\text{diverge_parties} = \text{false}$ and $\text{epoch_num}_i + 1 \notin \text{compromised_epochs}$, add epoch_id to $\text{receive_attempts}[\text{epoch_num}]$.
5. Else ($\text{diverge_parties} = \text{false}$ and $\text{epoch_id} = \text{epoch_id_self}_{1-i}$), set $\text{recv_chain_key}_i = \text{sending_chain_key}_{1-i}$ //Expected case
6. Output $(\text{GetReceivingKey}, \text{recv_chain_key}_i)$ and erase recv_chain_key_i .

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

- Add epoch_id_self_i to the list corruptions_i .
- Add $\text{epoch_num}_i, \text{epoch_num}_i + 1, \text{epoch_num}_i + 2, \text{epoch_num}_i + 3$ to the list $\text{compromised_epochs}$. //We need the compromise to go through the following stages: fully compromised, sender randomness updated, both parties' randomness updated.
- Initialize an empty list $\text{leak} = []$ and a variable $\text{recv_chain_key} = \perp$.
- If $\text{epoch_num}_{1-i} > \text{epoch_num}_i$:
 - Set $\text{recv_chain_key} = \text{sending_chain_key}_{1-i}$.
 - If $\text{epoch_num}_{1-i} \in \text{receive_attempts.keys}$ then set $\text{leak} = \text{receive_attempts}[\text{epoch_num}_{1-i}]$
- Send $(\text{ReportState}, \text{state}_{\mathcal{I}}, i, \text{recv_chain_key}_i, \text{leak})$ to \mathcal{A} .
- Upon receiving $(\text{ReportState}, i, S)$ from \mathcal{A} , output $(\text{Corrupt}, S)$ to $(\Pi_{\text{SGNL}}, \text{sid}, \text{pid}_i)$.

Fig. 7. The Epoch Key Exchange Functionality, \mathcal{F}_{eKE} (continued)

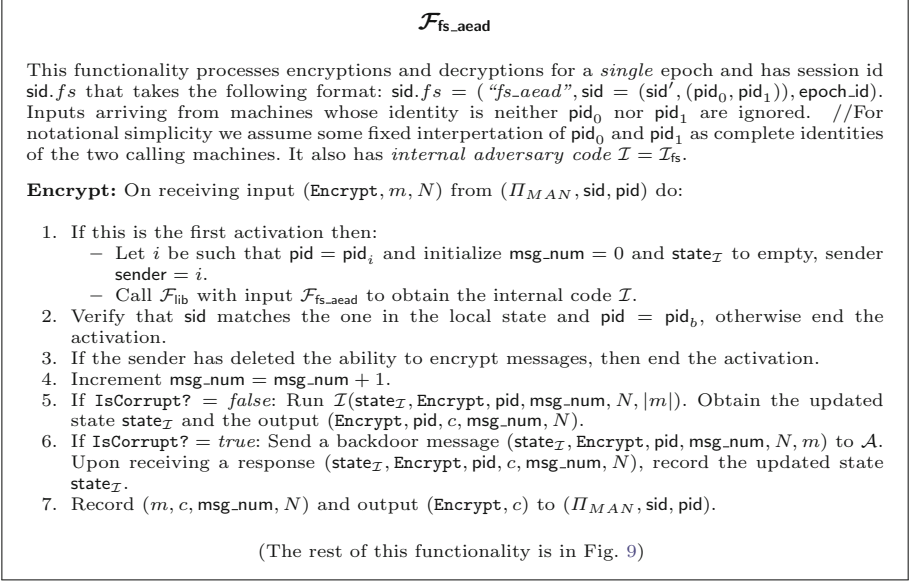


Fig. 8. The Forward-Secure Encryption Functionality $\mathcal{F}_{\text{fs_aead}}$

Theorem 3. *Assume that PRG is a secure length-doubling pseudorandom generator. Then protocol Π_{mKE} UC-realizes \mathcal{F}_{mKE} in the presence of \mathcal{F}_{DIR} , \mathcal{F}_{LTM} , and $\mathcal{F}_{\text{eKE}}^{\Pi_{\text{eKE}}} = (\mathcal{S}_{\text{eKE}}, \mathcal{F}_{\text{eKE}})$.*

The Public Ratchet: Realizing \mathcal{F}_{eKE} . By this point we have already described all of the functionalities in our model. As shown in Fig. 1, it only remains to construct real-world protocols that realize each of them. We defer a description of Π_{aead} to the full version [21] and show only the result here.

Theorem 4. *Assuming the unforgeability of (MAC, Verify), protocol Π_{aead} UC-realizes the ideal functionality $\mathcal{F}_{\text{aead}}$ in the presence of \mathcal{F}_{pRO} , as well as $F_{\text{mKE}}^{\Pi} = (\mathcal{S}_{\text{mKE}}, \mathcal{F}_{\text{mKE}})$, $\mathcal{F}_{\text{eKE}}^{\Pi_{\text{eKE}}} = (\mathcal{S}_{\text{eKE}}, \mathcal{F}_{\text{eKE}})$, \mathcal{F}_{DIR} , \mathcal{F}_{LTM} , \mathcal{F}_{lib} .*

Instantiating \mathcal{F}_{eKE} via Π_{eKE} (Figs. 12 to 13) is subtle. The main challenge, as observed by Alwen et al. [1] and others, is that the key derivation module within the public ratchet must maintain security if either of the previous root key or the newly generated ephemeral keys are uncompromised. Alwen et al. formalized this guarantee by way of constructing a new primitive: a PRF-PRG. In this work, we make two important improvements upon this construction. First, we contribute

$\mathcal{F}_{\text{fs_aead}}$ continued...

Decrypt: On receiving $(\text{Decrypt}, c, \text{msg_num}, N)$ from $(\mathcal{I}_{\text{SGNL}}, \text{sid}, \text{pid})$ do:

1. Verify that sid matches the one in the local state and $\text{pid} = \text{pid}_{1\text{-sender}}$, 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 $(\text{Authenticate}, c, \text{msg_num}, N, 0)$, then output $(\text{Decrypt}, c, \text{msg_num}, N, \text{Fail})$ to $(\mathcal{I}_{\text{SGNL}}, \text{sid}, \text{pid})$.
3. If $\text{IsCorrupt?} = \text{false}$:
 - Run $\mathcal{I}(\text{state}_{\mathcal{I}}, \text{Authenticate}, \text{pid}, c, \text{msg_num}, N)$ and obtain updated state $\text{state}_{\mathcal{I}}$ and output $(\text{Authenticate}, \text{pid}, c, \text{msg_num}, N, v)$.
 - If $v = \perp$, then record $(\text{Authenticate}, c, \text{msg_num}, N, 0)$ and output $(\text{Decrypt}, c, \text{msg_num}, N, \text{Fail})$ to $(\mathcal{I}_{\text{SGNL}}, \text{sid}, \text{pid})$.
 - Otherwise, mark msg_num as inaccessible and output $(\text{Decrypt}, c, \text{msg_num}, N, m)$ to $(\mathcal{I}_{\text{SGNL}}, \text{sid}, \text{pid})$.
4. Else $(\text{IsCorrupt?} = \text{true})$:
 - Send $(\text{state}_{\mathcal{I}}, \text{inject}, \text{pid}, c, \text{msg_num}, N)$ to \mathcal{A} .
 - On receiving the updated $\text{state}_{\mathcal{I}}$ and (inject, v) from \mathcal{A} , do:
 - If $v = \perp$, record $(\text{Authenticate}, c, \text{msg_num}, N, 0)$ and output $(\text{Decrypt}, c, \text{msg_num}, N, \text{Fail})$.
 - Else, then mark msg_num as inaccessible and output $(\text{Decrypt}, c, \text{msg_num}, N, v)$ to $(\mathcal{I}_{\text{SGNL}}, \text{sid}, \text{pid})$.

StopEncrypting: On receiving (StopEncrypting) from $(\mathcal{I}_{\text{SGNL}}, \text{sid}, \text{pid})$ do:

1. If sid doesn't match the one in the local state, if $\text{pid} \neq \text{pid}_{\text{sender}}$, 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 $(\text{StopDecrypting}, \text{msg_num}^*)$ from $(\mathcal{I}_{\text{SGNL}}, \text{sid}, \text{pid})$ do:

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

Corrupt: On receiving $(\text{Corrupt}, \text{pid})$ from $(\mathcal{I}_{\text{SGNL}}, \text{sid}, \text{pid})$:

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

Fig. 9. The Forward-Secure Encryption Functionality $\mathcal{F}_{\text{fs_aead}}$ (continued)

a new construction called a *Cascaded PRF-PRG* that allows for equivocation, in order to maintain security against adaptive adversaries. Second, we provide an instantiation in the plain model based on punctured PRFs.

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

This functionality has a session id $\text{sid.aead} = (\text{"aead"}, \text{sid.fs}, \text{msg_num})$ where $\text{sid.fs} = (\text{"fs_aead"}, \text{sid} = (\text{sid}', \text{pid}_0, \text{pid}_1), \text{epoch.id})$. It also has *internal adversary code* $\mathcal{I} = \mathcal{I}_{\text{aead}}$.

We initialize the state for \mathcal{I} to be $\text{state}_{\mathcal{I}} = \perp$.

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

1. If this is not the first encryption request or $\text{pid} \notin (\text{pid}_0, \text{pid}_1)$, end the activation. Let i be such that $\text{pid} = \text{pid}_i$.
2. Provide input $(\text{RetrieveKey}, \text{pid})$ to $(\mathcal{F}_{\text{mKE}}, \text{sid.mKE}, \text{pid})$
3. 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.
 - Else if $\text{IsCorrupt?} = \text{false}$:
 - Call \mathcal{F}_{lib} with input $\mathcal{F}_{\text{aead}}$ to obtain the internal code \mathcal{I} .
 - Run $\mathcal{I}(\text{state}_{\mathcal{I}}, \text{Encrypt}, \text{pid}, |m|, N)$, obtain the updated state $\text{state}_{\mathcal{I}}$, and the output $(\text{Encrypt}, \text{pid}, c)$.
 - Record the tuple $(c, k, m, N, 1)$, record i as the sender, and set $\text{ready2decrypt} = \text{true}$.
 - Else ($k \neq \perp$ and $\text{IsCorrupt?} = \text{true}$):
 - Send a backdoor message $(\text{state}_{\mathcal{I}}, \text{Encrypt}, \text{pid}, m, N)$ to \mathcal{A} .
 - Upon receiving a response $(\text{state}_{\mathcal{I}}, \text{Encrypt}, \text{pid}, c)$, record the updated state $\text{state}_{\mathcal{I}}$, record the tuple $(c, k, m, N, 1)$, record i as the sender, and set $\text{ready2decrypt} = \text{true}$.
- 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 there hasn't been a successful encryption request, if $\text{pid} \neq \text{pid}_{1-i}$, or if $\text{ready2decrypt} = \text{false}$ then output $(\text{Decrypt}, \text{Fail})$ to $(\Pi_{\text{fs_aead}}, \text{sid.fs}, \text{pid})$.
2. Provide input $(\text{RetrieveKey}, \text{pid})$ to $(\mathcal{F}_{\text{mKE}}, \text{sid.mKE}, \text{pid})$.
3. Upon obtaining a response $(\text{RetrieveKey}, \text{pid}, k)$ from \mathcal{F}_{mKE} : If $k = \perp$ then output $(\text{Decrypt}, \text{Fail})$. //Failure of decryption can occur for an honest receiver so we need an explicit failure notification.
4. If there is a record $(c, k, m, N, 1)$, note $\text{ready2decrypt} = \text{false}$ and output $(\text{Decrypt}, m)$ to $(\Pi_{\text{fs_aead}}, \text{sid.fs}, \text{pid})$.
5. If there is a record $(c, N, 0)$, output $(\text{Decrypt}, \text{Fail})$ to $(\Pi_{\text{fs_aead}}, \text{sid.fs}, \text{pid})$.
6. If $\text{IsCorrupt?} = \text{false}$
 - Run $\mathcal{I}(\text{state}_{\mathcal{I}}, \text{Authenticate}, \text{pid}, c, N)$ and obtain updated state $\text{state}_{\mathcal{I}}$ and a value v from \mathcal{I} .
 - If $v = \perp$ then record $(c, N, 0)$, and output $(\text{Decrypt}, \text{Fail})$.
 - Otherwise, note $\text{ready2decrypt} = \text{false}$, and output $(\text{Decrypt}, m)$.
7. Else ($\text{IsCorrupt?} = \text{true}$)
 - Send backdoor message $(\text{state}_{\mathcal{I}}, \text{inject}, \text{pid}, c, N)$ to \mathcal{A} .
 - Upon receiving response $(\text{inject}, \text{pid}, c, N, v)$ from \mathcal{A} continue.
 - If $v = \perp$ then record $(c, N, 0)$, and output $(\text{Decrypt}, \text{Fail})$.
 - Else, note $\text{ready2decrypt} = \text{false}$, and output $(\text{Decrypt}, v)$.

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

1. End the activation if $\text{pid} \notin (\text{pid}_0, \text{pid}_1)$. Otherwise, let i be such that $\text{pid} = \text{pid}_i$. Else, set IsCorrupt? to true , and send $(\text{ReportState}, \text{state}_{\mathcal{I}})$ to \mathcal{A} .
2. Upon receiving a response $(\text{ReportState}, \text{pid}, S)$ from \mathcal{A} , send $(\text{Corrupt}, S)$ to $(\Pi_{\text{fs_aead}}, \text{sid.fs})$.

Fig. 10. The Authenticated Encryption with Associated Data Functionality, $\mathcal{F}_{\text{aead}}$

\mathcal{F}_{mKE}

This functionality has a session id 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 pid_0 nor pid_1 are ignored.

This functionality is parametrized by a seed length λ

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 $\mathcal{F}_{\text{aead}}$ if $\text{IsCorrupt?} = \text{true}$:

1. If this is the first activation,
 - Initialize dictionary key_dict and variables $\text{IsCorrupt?} = \text{false}$, msg_num_0 , $\text{msg_num}_1 = 0$.
 - Parse sid to recover the two party ids $(\text{pid}_0, \text{pid}_1)$.
2. If $\text{pid} \notin \{\text{pid}_0, \text{pid}_1\}$ then end this activation.
3. End the activation if there is record $(\text{Retrieved}, i, \text{msg_num})$ or a record $(\text{StopKeys}, i, N)$ for $N < \text{msg_num}$.
4. 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 $k \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$.
5. 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.
6. Store $\text{key_dict}[\text{msg_num}] = k$.
7. If $\text{msg_num} > \text{msg_num}_i$, set $\text{msg_num}_i = \text{msg_num}$. // msg_num_i is the largest successfully retrieved message by party i .
8. 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 $(\Pi_{\text{fs_aead}}, \text{sid.fs}) = (\text{"fs_aead"}, \text{sid}, \text{epoch.id}, b)$, pid ,

- Run steps 3-7 of **RetrieveKey** for all 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 (Corrupt) from $(\Pi_{\text{fs_aead}}, \text{sid.fs} = (\text{"fs_aead"}, \text{sid}, \text{epoch.id}, b))$, pid :

1. Let i be such that $\text{pid} = \text{pid}_i$.
2. Set $\text{IsCorrupt?} = \text{true}$, create empty lists keys_in_transit , pending_msgs , and initialize $\text{chain_key} \stackrel{\$}{\leftarrow} \{0, 1\}^\lambda$.
3. If $\text{msg_num}_i = 0$ send $(\text{GetReceivingKey}, \text{epoch.id})$ to $(\Pi_{\text{eKE}}, \text{sid.eKE}, \text{pid})$.
//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. If the receiver has not retrieved any keys, we get the chain key from Π_{eKE} to provide to the simulator so that it matches the real world.
//Also note that we only get corrupted if the sender has already initialized this box \implies the sender's msg_num will never be 0.
4. On receiving $(\text{GetReceivingKey}, \text{recv_chain_key})$ set $\text{chain_key} = \text{recv_chain_key}$.
5. 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 keys_in_transit and append msg_num to pending_msgs .
6. If there is a record $(\text{StopKeys}, i, N)$ then let $\text{chain_key} = \perp$.
7. Send $(\text{ReportState}, i, \text{keys_in_transit}, \text{msg_num}_i, \text{chain_key})$ to \mathcal{A} .
8. On receiving a response $(\text{ReportState}, i, S)$ from \mathcal{A} :
 - Output $(\text{Corrupt}, \text{pending_msgs}, S)$ to $(\Pi_{\text{fs_aead}}, \text{sid.fs}, \text{pid}_i)$.

Fig. 11. The Message Key Exchange Functionality \mathcal{F}_{mKE}

Π_{eKE}

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

keyGen chooses a random Diffie-Hellman exponent $epoch_key \xleftarrow{\$} |G|$ for a known group G and sets $epoch_id = g^{epoch_key}$.

ConfirmReceivingEpoch: On input $(ConfirmReceivingEpoch, epoch_id^*)$ from (Π_{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 $root_key, epoch_id, epoch_key, sending_chain_key = \perp$.
 - Send $(GetInitKeys, pid_{1-i}, pid_i)$ to \mathcal{F}_{DIR} .
 - Upon receiving $(GetInitKeys, ik_j^{pk}, rk_j^{pk}, ok_{j \leftarrow i}^{pk})$, send input $(ComputeSendingRootKey, ik_j^{pk}, rk_j^{pk}, ok_{j \leftarrow i}^{pk})$ to \mathcal{F}_{LTM} .
 - Upon receiving $(ComputeSendingRootKey, k, ek_i^{pk})$, set $root_key = k$.
 - Do steps 4-6 of **Compute Sending Chain Key**.
 - Erase ek_i^{pk} and output $(ConfirmReceivingEpoch, epoch_id_{self} || ek_i^{pk} || ok_{j \leftarrow i}^{pk})$ to (Π_{SGNL}, sid, pid_i) .
4. Else (this is not the first activation):
 - **Compute Sending Chain Key**.
 - Output $(ConfirmReceivingEpoch, epoch_id_{self})$ to (Π_{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 (Π_{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 $root_key, epoch_id, epoch_key, sending_chain_key = \perp$.
 - 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 $(ComputeReceivingRootKey, k)$, set $root_key = k$.
4. **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. 13)

Fig. 12. The Epoch Key Exchange Protocol Π_{eKE}

Π_{eKE} continued...

Corrupt: On receiving (Corrupt) from $(\Pi_{SGNL}, \text{sid}, \text{pid}_i)$: return $(\text{epoch_key}, \text{epoch_id_self}, \text{epoch_id_partner}, \text{root_key})$ to $(\Pi_{SGNL}, \text{sid}, \text{pid}_i)$
//Note that the **Corrupt** interface is not part of the “real” protocol; it is only included for UC-modelling purposes.
//Below are subroutines used in the interfaces above. The calls to **Advance** and **Compute** use Π_{KDF} that is a cascaded PRF – PRNG

Compute Sending Chain Key:

1. Compute $\text{root_input} = \text{Exp}(\text{temp_epoch_id_partner}, \text{epoch_key_self})$.
2. Compute $(\text{root_key}) = \Pi_{KDF}.\text{Advance}(\text{root_key}, \text{root_input})$.
3. Generate a key pair $(\text{epoch_key_self}, \text{epoch_id_self}) \leftarrow \text{keyGen}()$.
4. Compute the next input $\text{root_input} = \text{Exp}(\text{temp_epoch_id_partner}, \text{epoch_key_self})$.
5. Compute $(\text{root_key}, \text{sending_chain_key}) = \Pi_{KDF}.\text{Compute}(\text{root_key}, \text{root_input})$.
6. Finally, $\text{advance}(\text{root_key}) = \Pi_{KDF}.\text{Advance}(\text{root_key}, \text{root_input})$
7. Erase 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 $\text{root_input} = \text{Exp}(\text{temp_epoch_id_partner}, \text{epoch_key_self})$.
2. Compute $(\text{root_key}, \text{temp_recv_chain_key}) = \Pi_{KDF}.\text{Compute}(\text{root_key}, \text{root_input})$.
3. Erase root_input . //The old root key is overwritten and therefore already erased.

Fig. 13. The Epoch Key Exchange Protocol Π_{eKE} (continued)

Theorem 5. Assume that $KDF : \{0, 1\}^n \rightarrow \{0, 1\}^{2n}$ is a CPRFG, that the DDH assumption holds in the group G . Then protocol Π_{eKE} UC-realizes the ideal functionality \mathcal{F}_{eKE} in the presence of global functionalities \mathcal{F}_{DIR} and \mathcal{F}_{LTM} .

Cascaded PRF-PRG. The goal of this primitive is to formalize the requirements required of a key derivation function (KDF) to adhere to the Signal specification [44] in the adaptive setting. We consider a stateful key derivation function (KDF) with two algorithms. Algorithm $\text{Compute}(\text{root_key}, \text{root_input}) = (\text{chain_key}, \text{root_key}')$ is given a state root_key and randomizer root_input , and computes a chaining key chain_key and an updated state $\text{root_key}'$ (discarding the old state). Algorithm $\text{Advance}(\text{root_key}, \text{root_input}) = \text{root_key}'$ is given a state root_key and randomizer root_input , and updates the state for a new epoch.

Our *cascaded PRF-PRG* definition requires that the KDF be secure against an adversary who can repeatedly execute methods (**Compute**) and (**Advance**) at will, and who can also obtain the module’s local state at any time. More specifically, we require existence of a simulator such that no adversary can distinguish an interaction with the scheme from an ideal interaction (Fig. 14) where the keys are truly random and the exposed state is generated by the simulator.

Cascaded PRF-PRG Security Game

Security game for a KDF (Compute, Advance) with domain $\{0, 1\}^n$ for the initial secret state, domain $\{0, 1\}^{m(n)}$ for the chaining keys and domain $\{R_n\}_{n \in N}$ for the randomizer, and a simulator S :

Real game:

- Oracle O is initialized with random state $s \leftarrow \{0, 1\}^n$.
- On input (Compute, root_input): O runs $\text{Compute}(s, \text{root_input}) = (\text{chain_key}, \text{root_key}')$, outputs chain_key and changes state $s = \text{root_key}'$.
- On input (Advance): O chooses $\text{root_input}' \leftarrow R_n$ at random, runs $\text{Advance}(s, \text{root_input}') = \text{root_key}'$, and changes state $s = \text{root_key}'$.
- On input (Expose-Advance): O outputs the old state s , chooses $\text{root_input}' \leftarrow R_n$ at random, computes $\text{Advance}(s, \text{root_input}') = \text{root_key}'$, and changes state $s = \text{root_key}'$.

Ideal game:

- Oracle O is initialized with a state consisting of a random function $F : R_n \rightarrow \{0, 1\}^m$.
- On input (Compute, root_input): O outputs $F(\text{root_input}) = \text{chain_key}$.
- On input (Advance): O updates its state to a new random function $F : R_n \rightarrow \{0, 1\}^m$.
- On input (Expose-Advance): O outputs $S((\text{root_input}_1, F(\text{root_input}_1)), \dots, (\text{root_input}_k, F(\text{root_input}_k)))$, where $\text{root_input}_1, \dots, \text{root_input}_k$ are all the queries made by \mathcal{A} since the last Advance query and F is the currently used random function. Finally O updates its state to a new random function $F : R_n \rightarrow \{0, 1\}^m$.

Fig. 14. Cascaded PRF-PRG Security Game

Definition 1 (Cascaded PRF-PRG (CPRFG)). A KDF (Compute, Advance) is a cascaded PRF-PRG (CPRFG) if there exist polytime algorithm S such that any polytime oracle machine \mathcal{A} can distinguish between the real and ideal interactions described in Fig. 14 only with advantage that is negligible in n .

In the full version of this work [21], we show two constructions of a Cascaded PRF-PRG: a straightforward one based on a programmable random oracle, and a non-trivial construction in the plain model based on a puncturable PRF and a PRF-PRG (in the style of Alwen et al. [1]).

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, Cham (2019). https://doi.org/10.1007/978-3-030-17653-2_5
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, Cham (2020). https://doi.org/10.1007/978-3-030-56784-2_9
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, November 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, Cham (2020). https://doi.org/10.1007/978-3-030-64378-2_10
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, Cham (2020). https://doi.org/10.1007/978-3-030-64381-2_1
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, Cham (2021). https://doi.org/10.1007/978-3-030-90453-1_22
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, Cham (2020). https://doi.org/10.1007/978-3-030-64840-4_21
8. Bellare, M., Rogaway, P.: Entity authentication and key distribution. In: Stinson, D.R. (ed.) CRYPTO 1993. LNCS, vol. 773, pp. 232–249. Springer, Heidelberg (1994). https://doi.org/10.1007/3-540-48329-2_21
9. Bellare, M., Rogaway, P.: Provably secure session key distribution: the three party case. In: 27th ACM STOC, pp. 57–66. ACM Press, May/June 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, Cham (2017). https://doi.org/10.1007/978-3-319-63697-9_21
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, Cham (2020). https://doi.org/10.1007/978-3-030-64378-2_8
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. 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)
14. Borisov, N., Goldberg, I., Brewer, E.A.: Off-the-record communication, or, why not to use PGP. In: WPES, pp. 77–84. ACM (2004)
15. Caforio, A., Durak, F.B., Vaudenay, S.: Beyond security and efficiency: on-demand ratcheting with security awareness. In: Garay, J.A. (ed.) PKC 2021, Part II. LNCS, vol. 12711, pp. 649–677. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-75248-4_23
16. 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 (2016). https://doi.org/10.1007/978-3-662-49387-8_10
17. Champion, S., Devigne, J., Duguey, C., Fouque, P.-A.: Multi-device for signal. In: Conti, M., Zhou, J., Casalicchio, E., Spognardi, A. (eds.) ACNS 2020, Part II. LNCS, vol. 12147, pp. 167–187. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-57878-7_9
18. Canetti, R.: Universally composable security: a new paradigm for cryptographic protocols. In: 42nd FOCS, pp. 136–145. IEEE Computer Society Press, October 2001

19. Canetti, R.: Universally composable security. *J. ACM* **67**(5), 28:1–28:94 (2020)
20. Canetti, R., Halevi, S., Herzberg, A.: Maintaining authenticated communication in the presence of break-ins. *J. Cryptol.* **13**(1), 61–105 (2000)
21. Canetti, R., Jain, P., Swanberg, M., Varia, M.: Universally composable end-to-end secure messaging. *Cryptology ePrint Archive*, Report 2022/376 (2022). <https://eprint.iacr.org/2022/376>
22. 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 (2001). https://doi.org/10.1007/3-540-44987-6_28
23. 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 (2002). https://doi.org/10.1007/3-540-46035-7_22
24. 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 (2016). https://doi.org/10.1007/978-3-662-49387-8_11
25. 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, November 2020
26. 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 2020*. LNCS, vol. 12282, pp. 588–605. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-61078-4_33
27. 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)
28. Cohn-Gordon, K., Cremers, C., Dowling, B., Garratt, L., Stebila, D.: A formal security analysis of the signal messaging protocol. *J. Cryptol.* **33**(4), 1914–1983 (2020)
29. 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, October 2018
30. 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)
31. 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, November 2020
32. Durak, F.B., Vaudenay, S.: Bidirectional asynchronous ratcheted key agreement with linear complexity. In: Attrapadung, N., Yagi, T. (eds.) *IWSEC 2019*. LNCS, vol. 11689, pp. 343–362. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26834-3_20
33. 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.A. (ed.) *PKC 2021, Part II*. LNCS, vol. 12711, pp. 410–440. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-75248-4_15

34. Hofheinz, D., Rao, V., Wichs, D.: Standard security does not imply indistinguishability under selective opening. In: Hirt, M., Smith, A. (eds.) TCC 2016, Part II. LNCS, vol. 9986, pp. 121–145. Springer, Heidelberg (2016). https://doi.org/10.1007/978-3-662-53644-5_5
35. 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, Cham (2018). https://doi.org/10.1007/978-3-319-96884-1_2
36. 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, Cham (2019). https://doi.org/10.1007/978-3-030-17653-2_6
37. 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, Cham (2019). https://doi.org/10.1007/978-3-030-36033-7_7
38. Keybase blog: New cryptographic tools on keybase (2020). <https://keybase.io/blog/crypto>
39. Krohn, M.: Zoom rolling out end-to-end encryption offering (2020). <https://blog.zoom.us/zoom-rolling-out-end-to-end-encryption-offering/>
40. Martiny, I., Kaptchuk, G., Aviv, A.J., Roche, D.S., Wustrow, E.: Improving signal’s sealed sender. In: NDSS. The Internet Society (2021)
41. Maurer, U.: Constructive cryptography – a primer. In: Sion, R. (ed.) FC 2010. LNCS, vol. 6052, p. 1. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14577-3_1
42. Maurer, U.: Constructive cryptography – a new paradigm for security definitions and proofs. In: Mödersheim, S., Palamidessi, C. (eds.) TOSCA 2011. LNCS, vol. 6993, pp. 33–56. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-27375-9_3
43. Nielsen, J.B.: Separating random oracle proofs from complexity theoretic proofs: the non-committing encryption case. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, pp. 111–126. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45708-9_8
44. Open Whisper Systems: Technical information: Specifications and software libraries for developers (2016). <https://signal.org/docs/>
45. Perrin, T.: The noise protocol framework (2018). <https://noiseprotocol.org/noise.html>
46. 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, Cham (2018). https://doi.org/10.1007/978-3-319-96884-1_1
47. 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)
48. Rotem, L., Segev, G.: Out-of-band authentication in group messaging: computational, statistical, optimal. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018. LNCS, vol. 10991, pp. 63–89. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96884-1_3
49. Singh, M.: Whatsapp is now delivering roughly 100 billion messages a day (2020). <https://techcrunch.com/2020/10/29/whatsapp-is-now-delivering-roughly-100-billion-messages-a-day/>
50. Status: Private, secure communication (2022). <https://status.im>

51. Sylo: Comms for the metaverse (2022). <https://sylo.io>
52. Unger, N., et al.: SoK: secure messaging. In: 2015 IEEE Symposium on Security and Privacy, pp. 232–249. IEEE Computer Society Press, May 2015
53. 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, October 2015
54. Unger, N., Goldberg, I.: Improved strongly deniable authenticated key exchanges for secure messaging. *PoPETs* **2018**(1), 21–66 (2018)
55. Vatandas, N., Gennaro, R., Ithurburn, B., Krawczyk, H.: On the cryptographic deniability of the signal protocol. In: Conti, M., Zhou, J., Casalicchio, E., Spognardi, A. (eds.) ACNS 2020, Part II. LNCS, vol. 12147, pp. 188–209. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-57878-7_10
56. WhatsApp LLC: About end-to-end encryption (2021). <https://faq.whatsapp.com/general/security-and-privacy/end-to-end-encryption/>
57. Yan, H., Vaudenay, S.: Symmetric asynchronous ratcheted communication with associated data. In: Aoki, K., Kanaoka, A. (eds.) IWSEC 2020. LNCS, vol. 12231, pp. 184–204. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-58208-1_11