Hecate: Abuse Reporting in Secure Messengers with Sealed Sender

Rawane Issa, Nicolas AlHaddad, and Mayank Varia
Boston University*, {ra1issa,nhaddad,varia}@bu.edu

Abstract

End-to-end encryption provides strong privacy protections to billions of people, but it also complicates efforts to moderate content that can seriously harm people. To address this concern, Tyagi et al. [CRYPTO 2019] introduced the concept of *asymmetric message franking* (AMF), which allows people to report abusive content to a moderator, while otherwise retaining end-to-end privacy by default and even compatibility with anonymous communication systems like Signal's sealed sender.

In this work, we provide a new construction for asymmetric message franking called Hecate that is faster, more secure, and introduces additional functionality compared to Tyagi et al. First, our construction uses fewer invocations of standardized crypto primitives and operates in the plain model. Second, on top of AMF's accountability and deniability requirements, we also add forward and backward secrecy. Third, we combine AMF with source tracing, another approach to content moderation that has previously been considered only in the setting of non-anonymous networks. Source tracing allows for messages to be forwarded, and a report only identifies the original source who created a message. To provide anonymity for senders and forwarders, we introduce a model of *AMF with preprocessing* whereby every client authenticates with the moderator out-of-band to receive a token that they later consume when sending a message anonymously.

1 Introduction

End-to-end encrypted messaging systems like Facebook Messenger, Signal, Telegram, Viber, and WhatsApp are used by billions of people [72] due to their powerful combination of cryptographic protections and ease of use. The security guarantees provided by encrypted messengers are both varied and valuable [69]: confidentiality and integrity from authenticated key exchange [16, 19, 49], deniability from the use of symmetric authenticated encryption [15, 28, 38], and forward and backward security via key evolution (aka ratcheting) [24, 39]. However, these very security guarantees complicate efforts by secure messaging platforms to investigate reports of abuse or disinformation campaigns, which can have serious consequences for individual people and collective society [10, 30, 64, 66, 71].

To address these concerns, the security research community has developed three methods to augment end-to-end messengers with privacy-respecting technologies to assist with content moderation: message franking, source tracing, and automated identification. First, *message franking* [28, 31, 38, 50, 67] allows recipients to manually report abusive messages with assurance that

*This material is based upon work supported by the National Science Foundation under Grants No. 1718135, 1739000, 1801564, 1915763, and 1931714, by the DARPA SIEVE program under Agreement No. HR00112020021, and by DARPA and the Naval Information Warfare Center (NIWC) under Contract No. N66001-15-C-4071.

unreported messages retain all guarantees of secure messengers, and reported messages are both accountable (the moderator correctly identifies the message's sender) and deniable (the moderator cannot prove this fact to anybody else). Second, *source tracing* [55, 68] allows the moderator to pinpoint the original source of a viral message rather than the person who forwarded the message to the eventual reporter. Finally, *automated identification* [11, 48, 51] proactively matches messages against a moderator-provided list of messages using a private (approximate) set membership test, with possible interventions like rate-limiting or warning labels in case of a match [65].

This work contributes a new construction called Hecate that simplifies, strengthens, and unifies the first two content moderation techniques: asymmetric message franking and source tracing. We do not consider automated identification, focusing instead on abuse reporting schemes that empower the people who receive messages to choose the action they wish to take [45, 56]. To provide context for our work, we describe the nascent space of message franking and source tracing in more detail before explaining our improvements.

1.1 Prior work

There exists a long line of research into the security of end-to-end encrypted messaging systems (EEMS) at both the protocol design and software implementation layers [4, 9, 22, 26, 41, 53]. Our work relies on these analyses in order to treat the underlying messaging protocol in a black-box manner and abstract away its details, so that we may focus on the additions provided by content moderation protocols themselves.

Message franking constructions involve four parties: a sender and receiver of a message, plus the platform providing the secure messaging service and a moderator who acts on abuse reports (see Figure 2). *Symmetric* message franking protocols are limited to the setting in which the platform and moderator are the same entity and have sufficient network-level visibility to pinpoint the sender of each message. At a high level, these constructions operate as follows: when a sender submits a ciphertext corresponding to the message m, the platform signs an attestation binding the sender's identity $\mathsf{id}_{\mathtt{src}}$ to a compact commitment $\mathsf{com}(\mathsf{m})$ provided by the sender in the clear. The receiver also sees this commitment (e.g., if it is part of a robust encryption scheme [1, 32, 33]) and can check whether it is correct, dropping the packet if it is malformed. Subsequently, the receiver can report the message as abusive by opening all [28, 31, 38] or part [50] of the commitment so that the platform can determine whether the message is abusive and take appropriate action.

The work of Tyagi et al. [67], which is the starting point for this paper and which we will henceforth refer to as TGLMR, removes the limitations from above. They define and construct an *asymmetric message franking* (AMF) scheme that can operate even if network metadata is partially or fully hidden from the platform. Specifically, AMF can operate when using Signal's sealed sender [62] or an anonymous communication system (e.g., [3, 25, 27, 70]) that hide the identity of the sender and/or receiver from even the platform.

Inspired by designated-verifier signatures [42, 60], the TGLMR construction requires the sender to make a Diffie-Hellman tuple $\langle g, g^{\mathsf{sk}_{\mathtt{src}}}, g^{\mathsf{sk}_{\mathtt{mod}}}, g^{\mathsf{sk}_{\mathtt{src}} \cdot \mathsf{sk}_{\mathtt{mod}}} \rangle$ involving the moderator's secret key and her own, as well as a non-interactive zero knowledge proof that the tuple is well-formed. This construction achieves accountability and deniability for the sender, but does not guarantee forward and backward security due to the use of long-lived secret keys. Moreover, it is complex and expensive to implement (see §4.2), and requires a non-falsifiable knowledge of exponent assumption in the random oracle model. Finally, TGLMR does not easily generalize to more complex conversation graphs that account for forwarding.

Construction	Features					Security Guarantees							
	Abuse reporting	Message forwarding	Source tracing	Trace info	Threshold report	Deniability	Forward security	Backward security	Unforgeability	Accountability	Confidentiality	Anonymity	Tree unlinkability
Signal	○	●	×	×	×	●	●	●	●	×	●	◐	●
Tyagi et al. [67]	●	○	×	src	○	●	○	○	●	●	●	●	×
Traceback [68]	●	●	●	path	○	◕	○	○	●	●	●	○	○
FACTS [51]	●	●	●	src	●	◕	◕	○	●	●	●	◐	○
Peale et al. [55]	●	●	●	src	○	●	◕	○	●	●	●	○	●
Hecate (*this work*)	●	●	●	src	○	●	●	●	●	●	●	●	○

●: fully provided, ◕ : provided but not proven, ◐: partially provided, ○: not provided, ×: not applicable

Table 1: A comparison of features and security properties provided by the Signal EEMS protocol as well as several abuse reporting constructions. Security properties are defined in §2.2 and §5. For the anonymity column, ◐ refers to providing anonymity at the level of Signal's sealed sender [62].

Another line of research investigates the ability for the moderator to trace the source of messages that might have been forwarded several times within an EEMS. Tyagi, Miers, and Ristenpart's [68] began this line of study with their Traceback scheme. This protocol reveals to the moderator the entire path from the original source to the reporter while formally guaranteeing notions of confidentiality and accountability to members of that path. However, their scheme imposes computational and storage burdens on the moderator; the required storage is proportional to the number of messages eligible to be traced. Moreover, it may not be desirable to reveal the entire forwarding path.

Two recent works provide *source tracing*, identifying only the original source of a reported message. First, Peale, Eskandarian, and Boneh [55] contribute a source tracing construction that inherits most security properties from the underlying EEMS (see Table 1). Using more more expensive crypto operations, a stronger variant of their construction is the only one to date to achieve tree unlinkability — namely, that a receiver who gets the same message twice cannot tell if they originate from the same or different source messages. Second, the FACTS scheme by Liu et al. [51] provides source tracing along with a threshold reporting scheme so that the moderator is only able to learn when sufficiently many complaints have been lodged against an abusive source client. However, none of these traceback or source tracing schemes [51, 55, 68] considers backward security as part of their security model. Also, none of them provides full anonymity of senders and receivers from the EEMS platform or moderator. For instance, FACTS is compatible with a network that provides one-sided anonymity, but it requires senders to identify themselves and request tokens from the moderator on the fly whenever they wish to send a message.

This leaves the following open question:

> *Can we design a protocol that simultaneously provides asymmetric message franking (AMF) and source tracing, achieves forward and backward security, maintains anonymity*

of senders and receivers to the extent provided by the underlying EEMS network, and only makes black-box use of standardized cryptography in the plain model?

In this work, we answer the question in the affirmative.

1.2 Our contributions

In this work, we provide a new definition and construction for asymmetric message franking (AMF) that is more general, more secure, and faster than previous work. To achieve this goal, we revisit the decision by TGLMR [67] to "restrict[] attention to non-interactive schemes for which franking, verification, and judging requires sending just a single message." On its face, this restriction seems natural because end-to-end encrypted messengers are designed to work asynchronously in situations with limited network connectivity, so one-round (online) protocols are desirable. However, this restriction also appears to direct the solution space toward expensive crypto tools like designated-verifier signatures and zero knowledge proofs.

Our core insight is to introduce an *AMF with preprocessing* model. Here, the online work of message franking and transmission still requires only one round of communication from the source to platform to receiver, as before. In addition, we allow the source and moderator to engage in data-independent communication beforehand (e.g., during off-peak hours when the source's device is connected to power and wifi) in order to produce *tokens* that can be consumed during the online phase. As with MPC [7] or PIR [8], we show that adding a preprocessing round to AMF allows for more efficient protocols, and in particular allows us to answer the open question from above.

Concretely, we contribute an AMF scheme called Hecate. Our construction contains a single data-independent preprocessing round in which the moderator sends a token to the sender, which is subsequently consumed in the traditional sender → platform → receiver communication flow of an end-to-end encrypted messaging system. There is an additional round of communication only if the receiver chooses to report an abusive message.

Preprocessing can be batched to produce many tokens at once, and it should be performed in advance (e.g., every night) rather than on the fly in order to avoid network-level traffic linking attacks [52]. With preprocessing, the communication path of reported messages begins and ends with the content moderator, so we can leverage techniques from (faster) symmetric message franking where the moderator can prepare a token (e.g., a symmetric encryption of the source's identity) that is only intelligible to its future self. Hecate also supports *source tracing*, in which receivers can opt to forward the message along to others, so the communication flow becomes a tree rooted at the original source.

A big challenge in our construction is to combine message forwarding with our AMF *backward security* requirement, which states that an attacker who previously (but no longer) controlled a source's device cannot blame the source for new messages. To our knowledge, this work is the first one to consider and formalize backward security within AMF. As we will discuss in more detail in §2, the challenge in combining AMF with backward security stems from the fact that immediate receivers of the message from the original source can rely on the backward security of the underlying encrypted messaging protocol to know that they're speaking with the source rather than the attacker, whereas indirect receivers cannot.

In summary, we make three contributions in this work.

- We formally define AMF with preprocessing (§3), generalizing the definition from TGLMR and adding forward and backward security. We include an optional extension that adds the

objectives of source tracing.

- We provide a construction called Hecate (§4) that provides AMF and source tracing. These goals are interrelated: the existence of message forwarding means that AMF must protect the anonymity of sources, forwarders, and receivers.
- We formalize and prove that Hecate achieves all of the security guarantees shown in Table 1. To the best of our knowledge, our work is the first to codify and achieve backward security for abuse reporting.

Before continuing, we wish to stress that any decision to use content moderation within end-to-end encrypted messengers requires weighing all of its potential benefits and risks. This includes an understanding of the limitations of Hecate (see §4.3 and §6) and prior works, as well as an assessment of the risk of abuse by or coercion of the moderator. This is a complex policy question whose discussion should involve computer scientists, but not only computer scientists. We take no stance on the policy question in this work; instead, we observe that these policy discussions are already ongoing [2, 17, 58] and that a sub-optimal understanding of the technological possibilities may push a service provider or nation-state policymakers toward a worse policy decision. We undertake this research in order to demonstrate the feasibility of alternatives to blunt privacy-inhibiting legislation.

2 Overview

In this section, we describe our objectives for an asymmetric message franking (AMF) system. We begin by describing the setting and threat model for our work, and then we provide a high-level description of the security requirements and a brief description of how our Hecate protocol will achieve them.

2.1 Setting and Threat Model

In this work, we consider an EEMS that might contain network-level anonymity protections such as Signal's sealed sender [62]. We focus exclusively on point-to-point two party communication within encrypted messengers; that said, our techniques translate directly to Signal's group messaging protocol as described in §4.3. Within the context of any single message transmission, we refer to the participating clients using the following terminology:

- The *source* who initially produced the message within the messaging platform.
- The *receivers* who receive the message, and can optionally decide to report it. If a receiver does choose to report a message, then we refer to this client as a *reporter*.
- The *forwarders* who are receivers that decide to send the message along to others.

Hence, each message communication flow has the structure of a tree rooted at the sender. Each client device can play different roles in the communication trees of different messages. We assume that clients only possess the computational power of a regular phone. In addition to the messenger clients, our model contains two (possibly separate, and more computationally powerful) entities that everyone can communicate with: (1) the platform that provides the messaging service, and

(2) the content moderator. We consider the platform and content moderator as possibly separate so that our model also captures settings where social media platforms outsource moderation tasks to other, more qualified organizations (e.g., Facebook's oversight board [54]).

For the most part, the parties in the system view all other parties as potentially malicious and colluding together. In particular, every party wishes to ensure confidentiality and integrity to the strongest extent possible, even if some or all of their counterparty, the platform, and the moderator are colluding against them. Put simply: we wish to retain all of the security goals that end-to-end encrypted messengers provide, as detailed in §2.2 and §5.

In this work, a malicious attacker has the power to compromise one or more parties, in which case it can observe these parties' local state (e.g., cryptographic keys) and run arbitrary software for the duration of their control of a victim's machine. A semi-honest party, by contrast, is assumed to perform all actions honestly, and the only objective against such a party is data minimization. We presume that the software implementing the encrypted messenger faithfully reproduces the intended specification so that the adversary cannot control the behavior of honest parties. Put another way, supply chain attacks and formal verification are out of scope of this work.

The parties' relationship toward the moderator is more subtle, and deserves further attention.

- The moderator and platform view each other as semi-honest. Looking ahead to our Hecate construction, the moderator trusts the accuracy of any timestamp applied by the platform; it need not trust the platform for any other purpose.

- Clients may choose to view the moderator as malicious, in which case it wants to be assured of limits on the moderator's power, or semi-honest, in which case it wants to be assured that the moderator can perform its role.

The objective of holding senders accountable for reported messages creates a tension with the security goals of end-to-end encrypted messengers. In particular, clients no longer receive confidentiality, deniability, sealed sender anonymity, or other privacy guarantees for reported messages. Moreover, an AMF scheme imposes a limit on forward security, because messages sent in the past now can be revealed to the moderator in the present. Our objective is to ensure security up to these fundamental limits. We emphasize that even if the moderator is malicious and colluding with some clients, *all of the security guarantees for end-to-end encrypted messaging continue to hold for all unreported messages communicated between non-colluding clients.* Moreover, even for reported messages, security holds against all other parties who are not colluding with the moderator.

Another tension exists between content moderation and network anonymity. For example, *sealed sender* is a feature introduced by the Signal protocol to hide the identity of the sender from the platform. It offers sender confidentiality and minimizes the amount of metadata stored by the platform. That said: if the sender can deny ever sending a message, then can we hold anyone responsible for sending an abusive message? TGLMR [67] resolved this problem through the use of zero-knowledge signatures; in this paper we contribute an alternative construction based solely on black-box use of standard crypto primitives.

2.2 Security goals

In an asymmetric message franking scheme, we aim to provide all of the security and privacy goals of encrypted messengers [22, 69]. Some goals are already consistent with content moderation, in which case AMF constructions can use these properties and must ensure that they don't weaken

them. To give some concrete examples for our Hecate protocol: we will not affect the underlying authentication of the messaging packet, the participant consistency, the system's causality preservation, etc. Conversely, some security goals are not fully compatible with content moderation, in which case we aim to make the smallest modification possible.

In this section, we summarize each security goal (also shown in Table 1) and describe the extent to which it is impacted by content moderation. These security goals apply to all clients who construct properly formatted messages that adhere to the encrypted messaging protocol, whether or not they are subsequently reported. That is: even though malicious parties in a crypto protocol receive no security guarantees, the mere act of sending a reported message does not render a client malicious.

- *Confidentiality.* Any users not involved in the creation, forwarding, or reporting of a message must not learn anything about the message contents, except for an upper bound on the message size.

- *Anonymity.* The AMF scheme should not allow the moderator, platform, and a receiving client to learn anything about the source and forwarding path of a message above and beyond what they would learn from the underlying EEMS. In particular, if a fully anonymous network communication system is used, then an AMF scheme should hide all message metadata. There are two exceptions: a receiver can learn the identity of the person who sent/forwarded the message to them, and the moderator can learn the source of a reported message.

- *Deniability.* If the moderator is honest and uncorrupted, then every user should be able to deny any non-moderator user's claim about the contents of a message. In particular, even if the message's recipient attempts to reveal the message contents, the source can deny it. Against a malicious moderator (and possibly a set of corrupted receivers), any user can deny the contents of a non-reported message. Reported messages on the other hand are deniable with respect to anyone other than the moderator. In particular, even if the moderator hands out their secret key material, the sender should still be able to deny having sent the message.

- *Forward security.* Adversaries that compromise user's state in the present should not be able to deduce anything about messages exchanged in the past. This goal does not apply to messages that happen to remain on the phone in the present, which can still be read *and reported*.

- *Backward security.* Once a client recovers from a compromise event, then (potentially after a short time delay) the adversary should not be able to break the confidentiality or integrity of messages exchanged after recovery. Moreover, the adversary should not be able to construct a new message that (if reported) would cause the moderator to blame the client. In particular, any preprocessing-related state compromised by the adversary becomes 'useless' after a recovery period.

- *Unforgeability.* Parties cannot send a message that appears as though it was sent by another party. If the platform (or anyone else) tampers with a message in transit, this deviation will be detected. Also, no honest party will accept receiving a malformed message.

- *Accountability.* If a message passes a receiver's verification check and is subsequently reported, then the moderator will trace it back to its original source. In more detail: nobody can

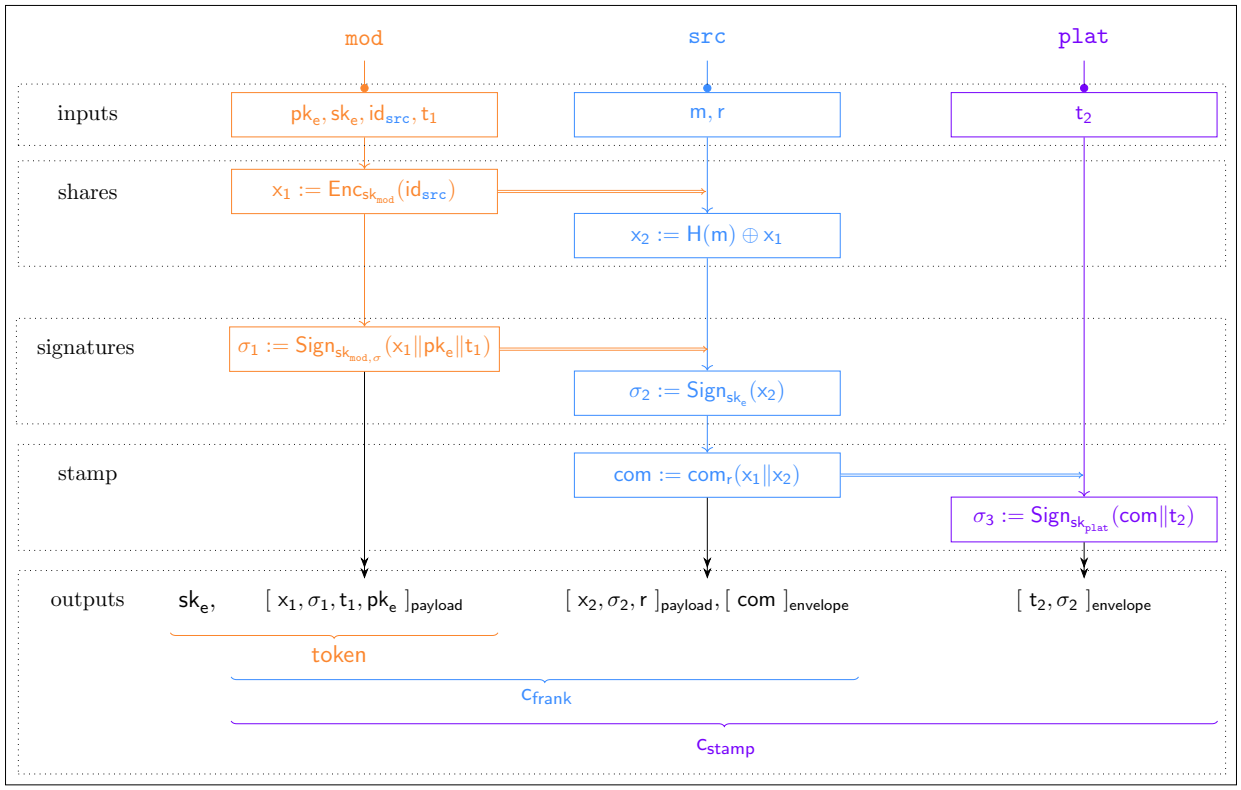

Figure 1: The construction of the different parts of a franked cipher. The outputs of the diagram correspond to each party's contribution to the eventual stamped cipher c_{stamp}. The source constructs the franked cipher c_{frank} using a token provided by the moderator during preprocessing. We denote by payload and envelope two different parts of the ciphertext as defined in the Signal sealed sender protocol [62]; the platform and receiver can read the envelope whereas only the receiver can read the payload.

falsely accuse someone who wasn't the source of a message, and the true source cannot evade detection and yet also have the message verified by the receiver.

2.3 Protocol Overview

In this section, we give a high level overview of our Hecate protocol in two stages (with and without message forwarding), along with an informal explanation of how it satisfies our security goals.

Hecate without forwarding. At a high level, our Hecate construction can be thought of as an interactive variant of designated-verifier signatures. Given a message m, the source constructs a 2-out-of-2 secret sharing, say $H(m) = x_1 \oplus x_2$. The objective of our protocol is for the moderator to bind x_1 with the source's identity (which on its own reveals nothing about m), and then for the source to bind x_1 to x_2 without using any long-lived keys. We show this procedure pictorially in Figure 1.

Since one of the two shares can be sampled even before m is known, during preprocessing the moderator selects x_1 as an encryption of the source's identity (which appears random to everyone

else), samples an ephemeral digital signature keypair $(\mathsf{sk}_e, \mathsf{pk}_e)$, and signs both x_1 and pk_e. The tuple of x_1, pk_e, and their signature constitute the preprocessing token token. During the online phase, the source uses the ephemeral key sk_e to sign x_2; we refer to the pair of x_2 and its signature as another token token.

The source provides both tokens to the receiver within the payload of an ordinary Signal packet, as shown at the top of Figure 2 (ignore the other elements of the franked message for now). Any receiver can check on its own whether the signatures are valid and the underlying values x_1 and x_2 combine to form the real message m that the receiver also gets from the underlying Signal communication; if verification fails, then the message is malformed, so it is dropped without displaying on the receiver's device. If a verified message is later reported, the two tokens together will convince the moderator that the source was the originator of message m.

Achieving our security goals. Many of our security guarantees derive strongly from the corresponding property of the underlying encrypted messenger protocol, so we focus on the most challenging goals here. Hecate provides accountability for the same reason as symmetric messaging franking schemes: the moderator created an authenticated encryption of the source's identity for its future self. Forward security holds because ephemeral signing key sk_e from the past were deleted before a compromise event in the present. Deniability can be shown in two parts: if the moderator's keys are breached then anyone can produce signatures for any choices of x_1 and x_2, and otherwise the source's identity is hidden within the encrypted token so anyone could have 'forged' signatures of an $(\mathsf{x}_1, \mathsf{x}_2)$ pair using her own tokens rather than those of the real source.

Backward (or post-compromise [24]) security is more challenging to address, and it is worth pausing for a moment to discuss what this guarantee means in the context of content moderation. If an adversary corrupts the source's phone, it *can* produce messages whose reports blame the source; this is inevitable. Our goal is to ensure that once the source recovers control of her phone, then (perhaps after a short delay δ) any new message produced by the adversary cannot implicate the honest source. To provide this guarantee within Hecate, the moderator includes a timestamp within its attestation to x_1, and receivers drop any message where this timestamp is too old. This ensures that an adversary cannot continue to use pre-processing tokens after the compromise event.

Hecate with message forwarding. Next, we allow forwarding of messages and consider source tracing, in which the moderator should identify only the original source of a reported message. For the most part, our construction is already amenable to source tracing: a forwarder can simply include the original source's tokens within a forwarded message rather than generating new tokens that would implicate herself. However, our timestamp-based solution to backward security now fails because the age of x_1 is insufficient to determine whether the original source had control of her cryptographic keys at the moment that the *original message* was sent (as opposed to the time of the forwarding).

As shown in Figure 2, we solve this problem by appending a timestamp time as the data traverses through the platform, so that receivers can check whether the timestamps on the preprocessing and sending stages are close in time to each other. This is sufficient because the original source's message x_2 inherits backward security from the underlying encrypted messenger, so to verify backward security it suffices to verify whether the pre-processed token (which contains the identity of the source to blame) was produced close in time to the message transmission. As shown on the bottom of Figure 2, timestamps for forwarded messages are disregarded; future recipients only care about

the timestamp from original source.

It only remains to bind the source timestamp to the message, so that it cannot be tampered later. Note that we cannot reveal x_2 to the platform, or else the platform and moderator together could recover the content of messages. Blind signatures are a possible solution to allow the platform to timestamp-and-sign obliviously, but constructions that only require one message received and sent by the platform require trusted setup [34], non-standard crypto assumptions [36, 37], or a concretely slow runtime with non-black-box reductions [35, 46]. Instead, we take advantage of the fact that the platform's actions need only be verified by recipients who already know x_1 and x_2, so it suffices for the platform to produce a signature σ of the current time time together with a commitment to the two shares. The corresponding decommitment randomness can be sent to the receiver within the encrypted messenger payload, so that the recipient can verify that it is well-formed.

3 Definitions

In this section, we present rigorous definitions for the cryptographic protocols that we use in this work. After briefly discussing the required cryptographic building blocks and our nearly black-box use of Signal, we detail a new definition for an asymmetric message franking scheme that generalizes TGLMR [67].

3.1 Crypto building blocks

This work uses four standard cryptographic building blocks that we use and adapt from Boneh-Shoup [14] and Katz-Lindell [47]. In what follows, we define the message space as $\mathcal{M} := \{0,1\}^*$, the key space as $\mathcal{K} := \{0,1\}^n$, the ciphertext space as $\mathcal{C} := \{0,1\}^*$, the randomness space $\mathcal{R} := \{0,1\}^n$ and the signature space as $\Sigma := \{0,1\}^n$, where n denotes the security parameter.

Definition 1. ***Commitment scheme.*** *A non-interactive commitment scheme is defined by two algorithms* Com *and* Vf*.*

- Com *is an algorithm that takes a random string* $\mathsf{r} \leftarrow\!\!\$\, \mathcal{R}$*, and a plaintext message* $m \in \mathcal{M}$ *and outputs a commitment* $\mathsf{com} := \mathsf{Com}(m, \mathsf{r})$*.*
- Vf *is an algorithm that takes a commitment* com*, a string* r *and a plaintext message* m *and checks if* $\mathsf{Vf}(m, \mathsf{com}, \mathsf{r}) := (\mathsf{Com}(m, \mathsf{r}) \stackrel{?}{=} \mathsf{com})$*.*

Definition 2. ***Binding Commitment.*** *A commitment scheme* $\pi = \{\mathsf{Com}, \mathsf{Vf}\}$ *is* computationally binding *if for all probabilistic polynomial time (PPT) adversaries* $\mathcal{A}$*, there is a negligible function* $\mathsf{negl}(n)$ *such that:*

$$\mathsf{Adv}_{\pi}^{\mathrm{binding}_{\mathrm{com}}}(\mathcal{A}) = \Pr[\mathsf{Com}(m, \mathsf{r}, \mathsf{param}) = \mathsf{Com}(m', \mathsf{r}', \mathsf{param}) \mid m \neq m'] \leq \mathsf{negl}(n).$$

Definition 3. ***Hiding Commitment.*** *Let* $\pi = \{\mathsf{Com}, \mathsf{Vf}\}$ *be a commitment scheme. Let* $\mathsf{Com}^{\mathcal{A}}_{\mathsf{hiding}}$ *be defined by the following experiment:*

- *The adversary* $\mathcal{A}$ *outputs a pair of messages* $\mathsf{m}_0, \mathsf{m}_1 \in \mathcal{M}$*.*
- *A uniform bit* $b \in \{0,1\}$ *and the randomness* $\mathsf{r} \leftarrow \{0,1\}^n$ *are chosen.*

- *The adversary $\mathcal{A}$ is given access to the commitment oracle $\mathsf{O}^{\mathsf{com-hiding}}$ which on messages m_0 and m_1 computes and returns the commitment $\mathsf{com} \leftarrow \mathsf{Com}(\mathsf{m}_b, \mathsf{r})$, where $\mathsf{Vf}(\mathsf{Com}(\mathsf{m}_b, \mathsf{r}), \mathsf{m}_b, \mathsf{r}) = 1$.*

- *The output of the experiment is 1 if $b' = b$ and 0 otherwise.*

A commitment scheme π is computationally hiding if for all PPT adversaries $\mathcal{A}$ there is a negligible function $\mathsf{negl}(n)$ such that:

$$\mathsf{Adv}_{\pi}^{\mathrm{hiding_{com}}}(\mathcal{A}) = \Pr[\mathsf{Com}_{\mathsf{hiding}}^{\mathcal{A}}(n) = 1] \leq \frac{1}{2} + \mathsf{negl}(n).$$

Definition 4. ***Encryption Scheme.*** *A private key encryption scheme is defined by three algorithms KGen, Enc and Dec over a finite message space $\mathcal{M}$.*

- KGen *is a probabilistic key generation algorithm that output a key pair $(\mathsf{pk}, \mathsf{sk})$ sampled uniformly at random from $\mathcal{K}$, where pk is defined as the public key and sk is defined as the secret key.*

- Enc *is the encryption algorithm that takes as an input sk and plaintext message $m \in \mathcal{M}$ and outputs $\mathsf{c} := \mathsf{Enc}_{\mathsf{sk}}(m)$ where $c \in \mathcal{C}$.*

- Dec *is the decryption algorithm that takes as an input pk and a ciphertext c in the ciphertext space $\mathcal{C}$ and outputs a plaintext message $m := \mathsf{Dec}_{\mathsf{pk}}(\mathsf{c})$ such that $\mathsf{c} := \mathsf{Enc}_{\mathsf{sk}}(m)$.*

Definition 5. ***CCA security.*** *Let $\pi = \{\mathsf{Enc}, \mathsf{Dec}, \mathsf{KGen}\}$ be an encryption scheme. Let $\mathsf{ENC}_{\mathsf{cca},\pi}^{\mathcal{A}}(n)$ denote the following experiment:*

- KGen *is run to obtain $(\mathsf{pk}, \mathsf{sk})$ and a uniform bit $b \in \{0, 1\}$ is chosen. The adversary $\mathcal{A}$ is given pk.*

- *The adversary $\mathcal{A}$ is given access to the encryption oracle $\mathsf{O}_{\mathsf{cca}}^{\mathsf{enc}}$ which, on messages $\mathsf{m}_0, \mathsf{m}_1$, outputs a ciphertext $c \leftarrow \mathsf{Enc}_{\mathsf{pk}}(\mathsf{m}_b)$.*

- *The adversary $\mathcal{A}$ is given access to the decryption oracle $\mathsf{O}_{\mathsf{cca}}^{\mathsf{decrypt}}$ which outputs the decrypted plaintext message m under sk when handed out a ciphertext c'.*

- *$\mathcal{A}$ continues to interact with the decryption and encryption oracles, but may not request a decryption of any ciphertext c returned by $\mathsf{O}_{\mathsf{cca}}^{\mathsf{enc}}$.*

- *Finally $\mathcal{A}$ output a bit b'. The output of the experiment is defined to be 1 if $b = b'$, and 0 otherwise.*

We say that π is secure under a chosen-ciphertext attack *(CCA) if for all PPT adversaries $\mathcal{A}$, there exists a negligble function $\mathsf{negl}(n)$ such that:*

$$\mathsf{Adv}_{\pi}^{\mathrm{enc_{cca}}} = \Pr[\mathsf{ENC}_{\mathsf{cca},\pi}^{\mathcal{A}}(n) = 1] \leq \frac{1}{2} + \mathsf{negl}(n).$$

Definition 6. ***CPA security.*** *Let* $\pi = \{\mathsf{Enc}, \mathsf{Dec}, \mathsf{KGen}\}$ *be an encryption scheme. Let* $\mathsf{ENC}^{\mathcal{A}}_{\mathsf{cpa},\pi}(n)$ *denote a similar experiment to* $\mathsf{ENC}^{\mathcal{A}}_{\mathsf{cca},\pi}(n)$ *where the adversary* $\mathcal{A}$ *only has access to the encryption oracle that rename as* $\mathsf{O}^{\mathsf{enc}}_{\mathsf{cpa}}$*. We say that* π *is* secure under a chosen-plaintext attack *(CPA) if for all PPT adversaries* $\mathcal{A}$*, there exists a negligble function* $\mathsf{negl}(n)$ *such that:*

$$\mathsf{Adv}^{\mathrm{enc_{cpa}}}_{\pi} = \Pr[\mathsf{ENC}^{\mathcal{A}}_{\mathsf{cpa},\pi}(n) = 1] \leq \frac{1}{2} + \mathsf{negl}(n).$$

Definition 7. ***Digital Signature Scheme.*** *A signature scheme is defined as the triple of algorithms* KGen*,* Sign*,* Vf *over the message space* $\mathcal{M}$ *and the signature space* Σ*.*

- KGen *is a probabilistic key generation algorithm that output a key pair* $(\mathsf{pk}, \mathsf{sk})$ *sampled from* $\mathcal{K}$*, where* pk *is the public verification key and* sk *is the secret signing key.*
- Sign *is the probabilistic signing algorithm that takes the signing key* sk *and a plaintext message* m *and outputs a signature* $\sigma \leftarrow \mathsf{Sign}_{\mathsf{sk}}(\mathsf{m})$*, where* $\sigma \in \Sigma$*.*
- Vf *is the deterministic verification algorithm which checks the signature* σ *against the plaintext message* m *and public key* pk *and outputs* $\perp$ *or* 1 *such that:*

$$\Pr[\mathsf{Vf}(\mathsf{pk}, \mathsf{m}, \mathsf{Sign}_{\mathsf{sk}}(\mathsf{m})) = 1] = 1.$$

Definition 8. ***EU-CMA security.*** *Let* $\pi = \{\mathsf{KGen}, \mathsf{Sign}, \mathsf{Vf}\}$ *denote a digital signature scheme. Let* $\mathsf{Sig}^{\mathcal{A}}_{\mathsf{eu-cma}}$ *be the experiment defined as:*

- KGen *is run to obtain* $(\mathsf{pk}, \mathsf{sk})$*.*
- *The adversary* $\mathcal{A}$ *is given* pk *and access to the signing oracle* $\mathsf{O}^{\mathsf{sign}}_{\mathsf{eu-cma}}$ *which on message* m *computes and outputs the signature* σ *of that message under the secret signing key* sk*. Let* $\mathcal{Q}$ *denote the set of all queries that* $\mathcal{A}$ *makes to* $\mathsf{O}^{\mathsf{sign}}_{\mathsf{eu-cma}}$*.*
- *The adversary* $\mathcal{A}$ *then outputs* (m', σ')*.*
- *The experiment outputs* 1 *if and only if* $\mathsf{Vf}_{\mathsf{pk}}(\mathsf{m}', \sigma') = 1$ *and* $m' \notin \mathcal{Q}$*, and* 0 *otherwise.*

We say that π *is existentially unforgeable under an adaptive chosen-message attack (EU-CMA) if for all PPT adveries* $\mathcal{A}$*, there is a negligible function* $\mathsf{negl}(n)$ *such that:*

$$\mathsf{Adv}^{\mathrm{sig_{eu\text{-}cma}}}_{\pi}(\mathcal{A}) = \Pr[\mathsf{Sig}^{\mathcal{A}}_{\mathsf{eu\text{-}cma}}(n) = 1] \leq \mathsf{negl}(n).$$

Definition 9. ***Hash Function.*** *A hash function with output length* l *is defined by two algorithms* Gen *and* H*.*

- Gen *is a probabilistic algorithm which outputs a key* $k \in \mathcal{K}$*.*
- H *is an algorithm which takes as input as key* k *and a string* $m \in \mathcal{M}$ *and outputs a string* $\mathsf{H}_k(m) \in \{0,1\}^{l(n)}$*.*

Definition 10. ***Collision Resistance.*** *Let* $\pi = \{\mathsf{Gen}, \mathsf{H}\}$ *denote a hash function. Let* $\mathsf{Hash}^{\mathcal{A}}_{\mathsf{coll}}$ *be the experiment defined as:*

- Gen *is run to obtain* k*.*

- *The adversary $\mathcal{A}$ is given access to the hashing oracle $\mathsf{O}^{\mathsf{hash}}$ which on input m returns $\mathsf{H}_k(\mathsf{m})$*
- *The adversary then outputs m_0 and m_1.*
- *The experiment outputs 1 if and only if $\mathsf{m}_0 \neq \mathsf{m}_1$ and $\mathsf{H}_k(\mathsf{m}_0) = \mathsf{H}_k(\mathsf{m}_1)$.*

We say that π is collision resistant if for all PPT adversaries $\mathcal{A}$ there is a negligible function $\mathsf{negl}(n)$ such that:

$$\mathsf{Adv}_{\pi}^{\mathrm{hash}_{\mathrm{coll}}}(\mathcal{A}) = \Pr[\mathsf{Hash}_{\mathsf{coll}}^{\mathcal{A}}(n) = 1] \leq \mathsf{negl}(n).$$

3.2 Modeling an End-to-end Encrypted Messaging System (EEMS)

End-to-end encrypted messaging systems (EEMS) using the Signal protocol [63] are a complex, delicate combination of the above cryptographic primitives. Starting with Cohn-Gordon et al. [22], there is a long line of research that analyzes the security of EEMS constructions such as the two-party Signal protocol itself (e.g., [4, 9, 44]), modified versions that provide provide stronger guarantees (e.g., [40, 43, 57]), and extensions to support group messaging (e.g., [18, 21, 59]).

For the purposes of this work, we wish to treat an EEMS as a black box and consider only an API-level description of its operation and underlying security guarantees. For this reason, we choose to follow the ideal functionality modeling of an EEMS within the recent work of Bienstock et al. [13] rather than the game-based definitions in the literature [4, 9]). Their ideal functionality $\mathcal{F}_{\mathrm{Signal}}$ models the creation, evolution, and destruction of communication 'sessions' between different pairs of parties, and keeps track of the long-term and ephemeral state that parties must hold for each operation.

We follow this abstract model, with three changes. First, we allow a sender to attach public information outside of a sealed sender envelope that is visible to the platform, as shown in Figure 2. Second, in order to support an anonymous network, we allow for inputs involving the parties' identities to be optional. Third, we add an explicit forgery method to highlight the fact that an EEMS achieves deniable authentication [29] (which is implicitly true in universally composable security models [20]): that is, either the sender or receiver can forge a transcript showing that a message originated with the other party.

Hence, our abstract model of Signal involves three methods. All of these methods implicitly use the state of the party (or parties) that participate in each method.

- $\mathsf{send}_{\mathsf{eems}}(\mathsf{m}^*; \mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}}) \to \mathsf{c}$: The underlying EEMS's end to end encrypted sending method, run by a source client $\mathsf{id}_{\mathtt{src}}$. It takes a message m^* as an input, and it sends a ciphertext c to the platform. This message m^* might contain payload and envelope components, similarly to how Signal's sealed sender operates. The sender also inputs their own identity $\mathsf{id}_{\mathtt{src}}$ and that of the intended recipient $\mathsf{id}_{\mathtt{rec}}$; we write them explicitly here, but for ease of notation we will omit one or both of these inputs when they are clear from context.

- $\mathsf{deliver}_{\mathsf{eems}}(\mathsf{c}; \mathsf{id}_{\mathtt{rec}}) \to \mathsf{m}^*$: The underlying EEMS's end to end delivery method. This is an interactive protocol in which the platform delivers a ciphertext c to the receiver $\mathsf{id}_{\mathtt{rec}}$. If this receiver was the intended target of a previous $\mathsf{send}_{\mathsf{eems}}$ that produced c, then they can decrypt using their local state to recover m^*. As above, we presume that $\mathsf{deliver}_{\mathsf{eems}}$ handles the payload and envelope of the message and splits c and m^* accordingly. Here, the decision about whether to include $\mathsf{id}_{\mathtt{rec}}$ as an input is less clear: for an anonymous communication channel it is important that the platform not know the receiver's identity but for non-anonymous

networks it may be required. We leave $\mathsf{id}_{\mathtt{rec}}$ as an optional parameter, and throughout this work we focus on the stronger setting in which the network is anonymous so this input is not provided.

- $\mathsf{forge}_{\mathsf{eems}}(\mathsf{m}^*; \mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}}) \to \mathsf{c}$: A forgery algorithm executed by a party $\mathsf{id}_{\mathtt{rec}}$ and requiring its state $\mathsf{state}_{\mathtt{rec}}$. It forges a transcript that looks as though the message m^* were sent by its counterparty $\mathsf{id}_{\mathtt{src}}$ in an EEMS communication, with a destination of $\mathsf{id}_{\mathtt{rec}}$. The parameters $\mathsf{id}_{\mathtt{src}}$ and $\mathsf{id}_{\mathtt{rec}}$ are optional for the same reasons as $\mathsf{send}_{\mathsf{eems}}$, and they will be omitted from this work.

3.3 Defining AMF with Preprocessing

Next, we present a rigorous definition for an asymmetric message franking system with preprocessing. This definition extends the one from TGLMR [67] in two ways. First, it includes an (optional) out-of-band communication between the moderator and sender, which results in a one-time *token* that is consumed when sending a message. Second, it is designed in a modular fashion so that it can be built on top of any EEMS that adheres to the model in §3.2.

Definition 11. *An* asymmetric message franking scheme with preprocessing $\mathsf{AMF} = (\mathsf{KGen}, \mathsf{TGen}, \mathsf{Frank}, \mathsf{Forward}, \mathsf{Stamp}, \mathsf{Inspect}, \mathsf{Verify}_{\mathtt{rec}}, \mathsf{Forge}_{\mathtt{mod}}, \mathsf{Forge}_{\mathtt{rec}})$ *is a tuple of algorithms called by different parties in the messaging ecosystem. We assume that each party has a unique identifier* id *provided by the underlying EEMS, and we define a state variable* state *for each party that contains all keys and tokens generated by the AMF scheme and the underlying EEMS, except for those that have previously been deleted. (Note that the user's* state *does not contain a transcript of prior message exchanged.) The algorithms operate as follows.*

- $(\mathsf{pk}, \mathsf{sk}) \leftarrow\!\!\$\ \mathsf{KGen}()$*: The key generation algorithm accessed by any party in the EEMS and used for creating (potentially multiple) cryptographic keys. The algorithm is at least run once at the beginning of time to setup the long-term key material for each party.*

- $\mathsf{TGen}(\mathsf{id}_{\mathtt{src}}, \mathsf{time}_{\mathtt{mod}}, \mathsf{sk}_{\mathtt{mod}}) \to \mathsf{token}$*: An algorithm run by the moderator periodically that provides a one-time token for use when sending a message. An honest moderator should only provide tokens to a participant that correspond to their actual identity* $\mathsf{id}_{\mathtt{src}}$*. It is assumed that the moderator can rely on the authentication of the existing EEMS to authenticate a user's identity* $\mathsf{id}_{\mathtt{src}}$ *before running* TGen*.*

- $\mathsf{Frank}(\mathsf{state}_{\mathtt{src}}, \mathsf{m}, \mathsf{id}_{\mathtt{rec}}, \mathsf{token}) \to \mathsf{m}_{\mathsf{frank}}$*: The message franking algorithm that allows a user with state* $\mathsf{state}_{\mathtt{src}}$ *to frank a plaintext message* m *that they wish to send to a receiver with id* $\mathsf{id}_{\mathtt{rec}}$*, using the* token *received in the preprocessing stage. The* $\mathsf{state}_{\mathtt{src}}$ *contains all key material produced by* KGen *and the underlying EEMS; we stress that the implementation of* Frank *might not use this state. The resulting franked message* $\mathsf{m}_{\mathsf{frank}}$ *can be sent to the platform using the* $\mathsf{send}_{\mathsf{eems}}$ *method.*

- $\mathsf{Stamp}(\mathsf{c}_{\mathsf{frank}}, \mathsf{sk}_{\mathtt{plat}}, \mathsf{time}) \to \mathsf{c}_{\mathsf{stamp}}$*: The stamping procedure run by the platform to authenticate and timestamp a franked cipher* $\mathsf{c}_{\mathsf{frank}}$*. The resulting stamped cipher* $\mathsf{c}_{\mathsf{stamp}}$ *can then be delivered to its intended recipient using the* $\mathsf{deliver}_{\mathsf{eems}}$ *method. We emphasize that* Stamp *does not have the sender or receiver's identity (even if* $\mathsf{deliver}_{\mathsf{eems}}$ *does).*

- $\mathsf{Forward}(\mathsf{m}_{\mathsf{frank}}, \mathsf{state}_{\mathtt{fwd}}, \mathsf{id}_{\mathtt{rec}}) \to \mathsf{m}_{\mathsf{frank}}'$: *The forwarding algorithm that allows a user to forward a franked message* $\mathsf{m}_{\mathsf{frank}}$ *to a new recipient* $\mathsf{id}_{\mathtt{rec}}$. *The resulting franked message can be sent using* $\mathsf{send}_{\mathsf{eems}}$*; in this way, the platform cannot distinguish between new and forwarded messages.*

- $\mathsf{Verify}_{\mathtt{rec}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{state}_{\mathtt{rec}}) \to (\mathsf{m}, \mathsf{report})$ *or* $\perp$*: The report construction algorithm that allows a user to validate a received franked message* $\mathsf{m}_{\mathsf{frank}}$ *with respect to its state* $\mathsf{state}_{\mathtt{rec}}$. *If valid,* $\mathsf{Verify}_{\mathtt{rec}}$ *returns the corresponding plaintext message and a string* report *that can be sent to the moderator if reporting an abusive message.*

- $\mathsf{Inspect}(\mathsf{report}, \mathsf{sk}_{\mathtt{mod}}) \to (\mathsf{id}_{\mathtt{src}}, \mathsf{m}, \mathsf{time})$ *or* $\perp$*: The inspection algorithm that allows a moderator to handle reported message* report *using their secret key* $\mathsf{sk}_{\mathtt{mod}}$ *by validating and possibly source tracing them. If the verification step succeeds, the moderator produces the id of the source* $\mathsf{id}_{\mathtt{src}}$*, the message contents* m*, and a timestamp of the message* time.

- $\mathsf{Forge}_{\mathtt{mod}}(\mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{m}, \mathsf{sk}_{\mathtt{mod}}) \to \mathsf{m}_{\mathsf{frank}}$*: For deniability, this forgery protocol allows a moderator with secret key* $\mathsf{sk}_{\mathtt{mod}}$ *to forge a franked message with plaintext* m *on behalf of a user with id* $\mathsf{id}_{\mathtt{src}}$ *and with an intended recipient with id* $\mathsf{id}_{\mathtt{rec}}$.

- $\mathsf{Forge}_{\mathtt{rec}}([\mathsf{id}_{\mathtt{src}}]_{\mathsf{opt}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{m}, \mathsf{state}_{\mathtt{rec}}) \to \mathsf{c}_{\mathsf{frank}}$*: For deniability, this forgery algorithm allows a receiver with id* $\mathsf{id}_{\mathtt{rec}}$ *and state* $\mathsf{state}_{\mathtt{rec}}$ *to forge a franked ciphertext as though the message* m *was transmitted through the EEMS by the sender* $\mathsf{id}_{\mathtt{src}}$ *to the receiver* $\mathsf{id}_{\mathtt{rec}}$. *Note that* $\mathsf{id}_{\mathtt{src}}$ *is an optional parameter and may not be needed by systems that support anonymous messaging. In this work, we omit it from the presentation of this work since we are aiming for the highest level of anonymity.*

We say that an AMF scheme with preprocessing is secure if all computationally bounded attackers have negligible advantage in winning the confidentiality, deniability, accountability, unforgeability, and backward secrecy games. These games are nuanced to describe, so rather than doing so here, we defer our discussion to the security analysis in §5.

4 Constructing Hecate

In this section, we describe the Hecate construction in detail, analyze its complexity, and remark on some limits to our construction.

4.1 Protocol Details

As per Definition 11, Hecate has eight algorithms. We describe them within this section, with the full specification provided in Figure 3. We defer discussing the forgery algorithms to the section on deniability, as these are proof artifacts and are not actual elements of the construction.

The key generation algorithm KGen initializes a few long-term keys: the moderator samples an authenticated encryption key and both the moderator and platform sample a digital signature key pair. One strength of our construction is that individual parties do not need any key information (above and beyond their existing Signal keys), which simplifies our security analysis of forward and backward security. Due to its simplicity, we omit KGen from Figure 3.

In the token generation algorithm TGen, the moderator creates a batch of tokens for users at specific time intervals. Each token provides users with:

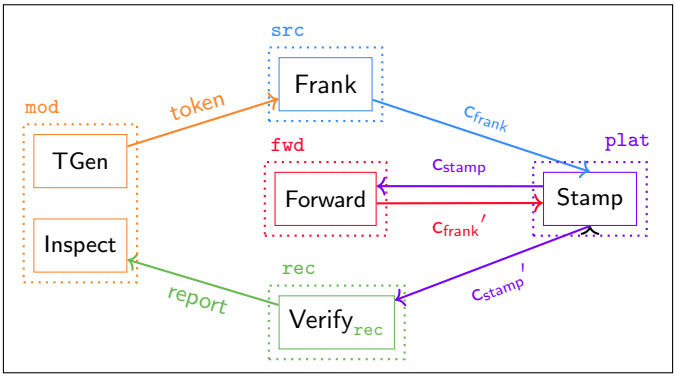

Figure 2: Diagram of Hecate's data flow for a message m through, from the source (top) to the forwarder (middle) and then to a reporting receiver (bottom). The source constructs $\mathsf{c}_{\mathsf{frank}}$ with an encrypted payload and an envelope shown in Fig. 1; it is securely transmitted and stamped by the platform to produce $\mathsf{c}_{\mathsf{stamp}}$. The forwarder constructs $\mathsf{c}_{\mathsf{frank}}'$ whose payload contains the source's token and stamp; in this case the receiver disregards the stamp applied by the platform. Note that token denotes the tuple $\langle \mathsf{x}_1, \mathsf{pk}_\mathsf{e}, \mathsf{sk}_\mathsf{e}, \mathsf{t}_1, \sigma_1 \rangle$, $\mathsf{c}_{\mathsf{frank}}$ denotes $\langle \mathsf{x}_2, \sigma_2, \mathsf{r}, \mathsf{com} \rangle$ together with everything in the token except sk_e, and $\mathsf{c}_{\mathsf{stamp}}$ denotes $\langle \mathsf{c}_{\mathsf{frank}}, \sigma_3, \mathsf{t}_2 \rangle$

1. TGen [mod → src]

$(\mathsf{pk}_{\mathtt{mod}}, \mathsf{sk}_{\mathtt{mod}}) \leftarrow\$ \mathsf{KGen}(1^n)$
$(\mathsf{pk}_{\mathtt{mod},\sigma}, \mathsf{sk}_{\mathtt{mod},\sigma}) \leftarrow\$ \mathsf{KGen}(1^n)$
$\mathsf{token} := \mathsf{construct}_{\mathsf{token}}(\mathsf{sk}_{\mathtt{mod}}, \mathsf{id}_{\mathtt{src}})$
return token

2a. Frank [src → plat]

$\mathsf{m}_{\mathsf{frank}} := \mathsf{construct}_{\mathsf{frank}}(\mathsf{token}, \mathsf{m})$
$\mathsf{c}_{\mathsf{frank}} := \mathsf{send}_{\mathsf{eems}}(\mathsf{m}_{\mathsf{frank}})$
return $\mathsf{c}_{\mathsf{frank}}$

2b. Forward [fwd → rec]

$\mathsf{m}_{\mathsf{frank}}' := \mathsf{construct}_{\mathsf{fwd}}(\mathsf{m}_{\mathsf{frank}})$
$\mathsf{c}_{\mathsf{frank}} := \mathsf{send}_{\mathsf{eems}}(\mathsf{m}_{\mathsf{frank}}')$
return $\mathsf{c}_{\mathsf{frank}}$

3. Stamp [plat → rec]

$(\mathsf{pk}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}}) \leftarrow\$ \mathsf{KGen}(1^n)$
$\mathsf{c}_{\mathsf{stamp}} := \mathsf{stamp}_{\mathsf{time}}(\mathsf{c}_{\mathsf{frank}})$
$\mathsf{c}_{\mathsf{stamp}}' := \mathsf{send}_{\mathsf{eems}}(\mathsf{c}_{\mathsf{stamp}})$
return $\mathsf{c}_{\mathsf{stamp}}'$

4. Verify$_{\mathtt{rec}}$ [rec → mod/rec]

$\mathsf{m}_{\mathsf{frank}} := \mathsf{deliver}_{\mathsf{eems}}(\mathsf{c}_{\mathsf{stamp}})$
$\mathsf{m}_{\mathsf{frank}} := \mathsf{move}_{\mathsf{stamp}}(\mathsf{m}_{\mathsf{frank}})$
$(\mathsf{m}, \mathsf{report}) := \mathsf{verifyRec}(\mathsf{m}_{\mathsf{frank}})$
if $(\mathsf{m}, \mathsf{report}) \stackrel{?}{=} \perp$:
  return $\perp$
return $(\mathsf{m}, \mathsf{report})$

5. Inspect [mod]

if verifyMsg(report) :
  return $(\mathsf{Dec}(\mathsf{report}.\mathsf{x}_1), \mathsf{report}.\mathsf{t}_2)$
return $\perp$

Figure 3: Hecate's construction. In the notation $[a \to b]$, a executes the method and sends the returned value to b. Note that the plat relays messages between the src and the rec. When the receiver receives a message, they may elect to just display a well formed message, forward it or report it. In other words, Forward is assumed to be preceded by Verify$_{\mathtt{rec}}$ and is omitted from this figure for simplicity.

$\mathsf{construct}_{\mathsf{token}}(\mathsf{sk}_{\mathtt{mod}}, \mathsf{id}_{\mathtt{src}})$

1 : $(\mathsf{pk}_e, \mathsf{sk}_e) \leftarrow\$ \mathsf{KGen}(1^n)$
2 : $\mathsf{t}_1 := \mathsf{time}()$
3 : $\mathsf{x}_1 := \mathsf{Enc}_{\mathsf{sk}_{\mathtt{mod}}}(\mathsf{id}_{\mathtt{src}})$
4 : $\sigma_1 := \mathsf{Sign}_{\mathsf{sk}_{\mathtt{mod}},\sigma}(\mathsf{x}_1 \| \mathsf{pk}_e \| \mathsf{t}_1)$
5 : $\mathsf{token} := (\mathsf{x}_1, \mathsf{t}_1, \sigma_1, (\mathsf{pk}_e, \mathsf{sk}_e))$
6 : **return** token

$\mathsf{construct}_{\mathsf{frank}}(\mathsf{m}, \mathsf{token})$

1 : $\mathsf{r} \leftarrow\$ \{0,1\}^n$
2 : $(\mathsf{x}_1, \mathsf{t}_1, \sigma_1, (\mathsf{pk}_e, \mathsf{sk}_e)) = \mathsf{token}$
3 : $\mathsf{x}_2 := split(\mathsf{x}_1, H(\mathsf{m}))$
4 : $\sigma_2 := \mathsf{Sign}_{\mathsf{sk}_e}(\mathsf{x}_2)$
5 : $\mathsf{com} := \mathsf{com}_{\mathsf{r}}(\mathsf{x}_1 \| \mathsf{x}_2)$
6 : $\mathsf{envelope} := \mathsf{com}$
7 : $\mathsf{payload} := (\mathsf{x}_1, \mathsf{x}_2, \mathsf{r}, \mathsf{t}_1, \sigma_1, \sigma_2, \mathsf{pk}_e)$
8 : $\mathsf{m}_{\mathsf{frank}} := (\mathsf{payload}, \mathsf{envelope})$
9 : **return** $\mathsf{m}_{\mathsf{frank}}$

$\mathsf{construct}_{\mathsf{fwd}}(\mathsf{m}_{\mathsf{frank}})$

1 : $\mathsf{m}_{\mathsf{frank}} := \mathsf{move}_{\mathsf{stamp}}(\mathsf{m}_{\mathsf{frank}})$
2 : $\mathsf{envelope} \leftarrow\$ \{0,1\}^n$
3 : **return** $\mathsf{m}_{\mathsf{frank}}$

$\mathsf{stamp}_{\mathsf{time}}(\mathsf{c}_{\mathsf{frank}}, \mathsf{sk}_{\mathtt{plat}})$

1 : $\mathsf{t}_2 = \mathsf{time}()$
2 : $\sigma_3 := \mathsf{Sign}_{\mathsf{sk}_{\mathtt{plat}}}(\mathsf{com} \| \mathsf{t}_2)$
3 : $\mathsf{c}_{\mathsf{stamp}}.\mathsf{envelope} := (\mathsf{com} \| \mathsf{t}_2 \| \sigma_3)$
4 : $\mathsf{c}_{\mathsf{stamp}}.\mathsf{payload} := \mathsf{c}_{\mathsf{frank}}.\mathsf{payload}$
5 : **return** $\mathsf{c}_{\mathsf{stamp}}$

$\mathsf{verifyRec}(\mathsf{m}_{\mathsf{frank}})$

1 : **if** $\mathsf{verifyMsg}(\mathsf{m}_{\mathsf{frank}})$:
2 : **return** $\perp$
3 : $\mathsf{report} := \mathsf{m}_{\mathsf{frank}}$
4 : **return** $(\mathsf{m}_{\mathsf{frank}}.\mathsf{m}, \mathsf{report})$

$\mathsf{verifyMsg}(\mathsf{report})$

1 : $b_1 := \mathsf{verifyToken}(\mathsf{report})$
2 : $b_2 := \mathsf{verifyCommit}(\mathsf{report})$
3 : $b_3 := \mathsf{verifyExpiry}(\mathsf{report})$
4 : **return** $b_1 \wedge b_2 \wedge b_3$

$\mathsf{verifyToken}(\mathsf{report})$

1 : $reveal := \mathsf{open}(\mathsf{x}_1, \mathsf{x}_2)$
2 : $b_1 := (reveal \stackrel{?}{=} \mathsf{H}(\mathsf{m}))$
3 : $b_2 := \mathsf{Vf}_{\mathsf{pk}_{\mathtt{mod}}}(\mathsf{x}_1 \| \mathsf{pk}_e \| \mathsf{t}_1, \sigma_1)$
4 : $b_3 := \mathsf{Vf}_{\mathsf{pk}_e}(\mathsf{x}_2, \sigma_2)$
5 : **return** $b_1 \wedge b_2 \wedge b_3$

$\mathsf{verifyExpiry}(\mathsf{m}_{\mathsf{frank}})$

1 : $b := |\mathsf{t}_1 - \mathsf{t}_2| \stackrel{?}{<} \mathsf{expiry}$
2 : **return** b

$\mathsf{verifyCommit}(\mathsf{m}_{\mathsf{frank}})$

1 : $b_1 := \mathsf{Vf}(\mathsf{x}_1 \| \mathsf{x}_2, \mathsf{com}, \mathsf{r})$
2 : $b_2 := \mathsf{Vf}_{\mathsf{pk}_{\mathtt{plat}}}(\mathsf{com} \| \mathsf{t}_2, \sigma_3)$
3 : **return** $b_1 \wedge b_2$

$\mathsf{move}_{\mathsf{stamp}}(\mathsf{m}_{\mathsf{frank}})$

1 : // check if $\mathsf{m}_{\mathsf{frank}}$ already contains
2 : // a platform stamp
3 : **if** $\mathsf{stamp} \stackrel{?}{\notin} \mathsf{payload}$:
4 : $\mathsf{payload} := \mathsf{payload} \| \mathsf{envelope}$
5 : **return** $\mathsf{m}_{\mathsf{frank}}$

Figure 4: Hecate's subroutines. We omit writing out attribute access notation when it is obvious from the context (i.e. com for instance is a short hand for $\mathsf{c}_{\mathsf{frank}}.\mathsf{com}$).

- Ephemeral session keys $(\mathsf{pk}_e, \mathsf{sk}_e)$ that they can use to sign their message. None of the keys tie to users' long-term key material, thus giving the sender plausible deniability and confidentiality with respect to other users.
- A dual purpose randomized encryption $\mathsf{x}_1 := \mathsf{Enc}_{\mathsf{sk}_{\mathtt{mod}}}(\mathsf{id}_{\mathtt{src}})$ of the user's identity $\mathsf{id}_{\mathtt{src}}$ under the moderator's secret key $\mathsf{sk}_{\mathtt{mod}}$ that enforces accountability with respect to the moderator, confidentiality with respect to other user, and provides token integrity. The later property is ensured by having the sender create a share x_2 that along with x_1 reconstructs to a hash of the sent message.
- A timestamp t_1 that provides backward security.
- A signature σ_1 of the entire token that guarantees integrity and unforgeability of the token. σ_1 is signed with the moderator secret signing keys $(\mathsf{pk}_{\mathtt{mod},\sigma}, \mathsf{sk}_{\mathtt{mod},\sigma})$.

The Frank method is executed every time the source wishes to send a message, and produces a franked cipher $\mathsf{c}_{\mathsf{frank}}$ that is relayed to the platform via the underlying sending method $\mathsf{send}_{\mathsf{eems}}$. This procedure requires an input plaintext message m from the source and consumes a single token at a time. The sender begins by unpacking x_1 from the token and computes x_2 such that these variables constitute a 2-out-of-2 sharing of $\mathsf{H}(\mathsf{m})$. Next, x_2 is signed via the ephemeral keys in the original token to produce σ_2. Collectively, x_2, σ_2, and elements of the pre-processing token (excluding the secret ephemeral key) will constitute the payload of the franked message. Then, the sender creates a commitment com of $\mathsf{x}_1 \| \mathsf{x}_2$ using the randomness r. The user then pushes com onto the envelope of the franked message and appends r to its payload. Finally, the sender sends out the resulting franked message to the platform via the underlying sending method $\mathsf{send}_{\mathsf{eems}}$ to stamp and relay the produced franked cipher back to its destined receiver. The constructed franked message $\mathsf{m}_{\mathsf{frank}}$ has the property of: (1) binding both the online and pre-processing stages via x_2 and com, (2) being signed in an ephemeral way that allows the receiver to check the well-formness of the message while still providing the sender with deniability guarantees with respect to anyone other than the moderator.

In Stamp, the platform timestamps and signs the envelope of franked ciphers before relaying the resulting franked cipher to its intended recipient. By doing so, the clients are guaranteed that no moderator token is indefinitely used after a compromise to blame them for unsent messages.

On reception, the receiver executes Verify to validate the signatures, timestamps expiration date, packet integrity, and check the envelope commitments against the inner tokens. If a message fails the integrity check, the receiver drops the packet and the application never displays the plaintext message. Otherwise, Verify generates the message report report which can be sent out to the moderator, and the plaintext message m that can be displayed on the user's phone instead. In Hecate, report is entirely made out of the franked message $\mathsf{m}_{\mathsf{frank}}$. When a moderator receives report, they locally run the Inspect method which performs the same verification procedure as the recipient, and if successful, decrypts the source's identity from the ciphertext x_1 within the token.

Verified messages can alternatively be forwarded using the optional Forward method. There are two differences between Forward and Frank: the forwarder creates a nonsensical commitment outside the Signal envelope and moves the true commitment and signed timestamp into the payload of the franked message. As a consequence, Frank and Forward payloads are indistinguishable to the platform but distinguishable by the receiver. The receiver applies an identical verification procedure for forwarded messages after it has handled the envelope of the franked message.

	KeyGen	Sign	Verify
TGen (moderator)	1	1	0
Frank (source)	0	1	0
Stamp (platform)	0	1	0
Verify (receiver)	0	0	3
Forward (forwarder)	0	0	0
$\mathsf{Verify}_{\mathtt{rec}}$ (moderator)	0	0	3

Table 2: For each of the interactive algorithms within Hecate (except for the one-time KGen at setup), we count the number of public-key digital signature operations computed by each party. In this table, we only include the additional cryptographic operations for AMF, on top of the computation required for the underlying Signal protocol. By contrast, each method requires at least 11 modular exponentiations in TGLMR [67].

In this section, we describe the operation of our AMF protocol Hecate in detail. Then, we compare the performance of our scheme to TGLMR [67].

4.2 Performance Analysis

The Hecate protocol adds a small overhead per party on top of the underlying Signal protocol. Our construction uses only the crypto primitives listed in §3, standardized crypto primitives in a black box manner, all of which have high-performance instantiations that have been heavily benchmarked. In Table 2, we list the number of times each party invokes a public-key digital signature method, since that is the most expensive primitive within our construction.

Looking closely at our protocol, one can see that the moderator spends for token creation 1 keygen operation to produce the ephemeral key pair and 1 signature to sign the public ephemeral key with the identity of the sender. For processing reported messages, the moderator spends 3 verifications to check that the two shares are not tampered with and have the right timestamps followed by one decryption to hold the right source accountable. As for sender, she spends 1 signature to sign the second share. And lastly for the receiver, she spends three verifications to check that the two shares are not tampered with and have the right timestamps.

To put this in perspective, the prior work of TGLMR [67] required at least 11 modular exponentiations per operation. Concretely, using TGLMR's choice of Schnorr signatures as our underlying digital signature, a sender would spend approximately 11 times less (2.4 ms vs 29.1 ms) to sign a message of size 4KB instantiated over NIST elliptic curve groups P-521 according to the benchmarks provided within [67]. The improvement is due to the fact that the sender is using a vanilla Schnorr signature to sign his share.

4.3 Limitations

In this section, we discuss a few limitations of our scheme, thus clearly defining the space of solutions that we are after.

Reporting Benign Messages. Our construction allows receivers to report messages that may later be deemed to be non-abusive. While it might be possible to require the receiver to prove to an honest moderator that the message they are reporting is actually abusive, this question is

incredibly delicate and is therefore out of scope for this and all prior works on end-to-end abuse reporting. In the special case that the receiver is colluding with the moderator, we remark that (a) there is little that can be done to prevent false reports, and (b) the overall leakage is no worse than what end-to-end encrypted messengers would already reveal to this colluding set.

Distinguishing Forwarded vs. Original Messages. In our construction, receivers are able to distinguish between sent and forwarded messages. While this may be a desirable feature in certain cases, it is still a leakage in our system.

Forwarding Cycle Linkability. If the forwarding path of a message contains a cycle, i.e. a receiver receives the same forwarded message multiple time, then she can directly correlate the forwards and specifically tell that they originated from the same source. In this case, the receiver can tell that the forwarders have a common ancestry along the forwarding path. Additionally, she can infer that there exists some number of users that connect them in their network of friends. This is an inherent weakness in our protocol as a result of forwarding the same tokens per message that we do not attempt to protect against. We argue that in practice this leakage may be admissible as most user's "friend networks" are connected.

Forwarding Tree Up-Rooting. Receiver's of a message may "up-root" its forwarding sub-tree by acting as the original senders of that message instead of forwarding it themselves. In general, we do not believe that this poses a concern as users do not have the incentive to incriminate themselves with bad messages. Moreover, prior work [68] suggests that this problem warrants an application side solution, if any, and not a cryptographic one that would restrict users from copying received messages and acting as their creators.

5 Security Analysis

In this section, we formally define the security properties of asymmetric message franking (AMF) schemes with preprocessing, and we prove that Hecate guarantees them.

5.1 Deniability

Deniability states that *a sender should always be able to deny that they sent a particular message to anyone, except to the moderator when a message is reported.* Deniability could hold with respect to a colluding moderator and receivers, as shown in the $\mathsf{DENM}_{\mathsf{b}}^{\mathcal{A}}$ game in Figure 6, or against malicious receivers who are not colluding with an honest moderator, which corresponds to the $\mathsf{DENR}_{\mathsf{b}}^{\mathcal{A}}$ game in Figure 6.

In more detail, each game provides the adversary $\mathcal{A}$ with polynomially-many queries to an oracle $\mathsf{O}_{\mathsf{b}}^{\mathsf{DENM}}$ or $\mathsf{O}_{\mathsf{b}}^{\mathsf{DENR}}$, respectively. For each query, the adversary chooses a plaintext message m and the sender $\mathsf{id}_{\mathtt{src}}$ and corrupted receiver $\mathsf{id}_{\mathtt{rec}_{\mathcal{A}}}$ of that message. Both oracles behave similarly: depending on the parameterized choice bit b, the oracle will either forge a message as a corrupted moderator/receiver as if originating from $\mathsf{id}_{\mathtt{src}}$, or ask the honest source $\mathsf{id}_{\mathtt{src}}$ to produce it themselves. Deniability requires that no adversary can distinguish between forged and real messages, even with access to the secret keys of malicious parties. In other words, we want to show that a user can always repudiate having sent a message even when the moderator/receiver provides access

$\mathsf{send}_{\mathsf{amf}}(\mathsf{m}, \mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}})$

1 : $\mathsf{state}_{\mathtt{src}} := \mathsf{retrieve}_{\mathsf{state}}(\mathsf{id}_{\mathtt{src}})$
2 : *fetch* token *from* $\mathsf{state}_{\mathtt{src}}$
3 : $\mathsf{m}_{\mathsf{frank}} := \mathsf{Frank}(\mathsf{state}_{\mathtt{src}}, \mathsf{m}, \mathsf{id}_{\mathtt{rec}}, \mathsf{token})$
4 : $\mathsf{c}_{\mathsf{frank}} := \mathsf{send}_{\mathsf{eems}}(\mathsf{m}_{\mathsf{frank}})$
5 : **return** $\mathsf{c}_{\mathsf{frank}}$

$\mathsf{fwd}_{\mathsf{amf}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{fwd}}, \mathsf{id}_{\mathtt{rec}})$

1 : $\mathsf{state}_{\mathtt{fwd}} := \mathsf{retrieve}_{\mathsf{state}}(\mathsf{id}_{\mathtt{fwd}})$
2 : $\mathsf{m}_{\mathsf{frank}}{}' := \mathsf{Forward}(\mathsf{state}_{\mathtt{fwd}}, \mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}})$
3 : $\mathsf{c}_{\mathsf{frank}} := \mathsf{send}_{\mathsf{eems}}(\mathsf{m}_{\mathsf{frank}}{}')$
4 : **return** $\mathsf{c}_{\mathsf{frank}}$

$\mathsf{receive}_{\mathsf{amf}}(\mathsf{c}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}})$

1 : $\mathsf{c}_{\mathsf{stamp}} := \mathsf{Stamp}(\mathsf{c}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}})$
2 : $\mathsf{m}_{\mathsf{frank}} := \mathsf{deliver}_{\mathsf{eems}}(\mathsf{c}_{\mathsf{stamp}})$
3 : $\mathsf{state}_{\mathtt{rec}} := \mathsf{retrieve}_{\mathsf{state}}(\mathsf{id}_{\mathtt{rec}})$
4 : **if** $\mathsf{Verify}_{\mathtt{rec}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{state}_{\mathtt{rec}}) \stackrel{?}{=} 0$:
5 : **return** $\perp$
6 : **return** $\mathsf{m}_{\mathsf{frank}}$

$\mathsf{O}^{\mathsf{corrupt}}(\mathsf{id})$

1 : // global_t is only relevant in BAC
2 : $\mathsf{corrupted} = \mathsf{corrupted} \cup (\mathsf{id}, \mathsf{global}_t)$
3 : $\mathsf{state}_{\mathsf{id}} := \mathsf{retrieve}_{\mathsf{state}}(\mathsf{id})$
4 : **return** $\mathsf{state}_{\mathsf{id}}$

$\mathsf{O}^{\mathsf{request}}(\mathsf{id})$

1 : // global_t is only relevant in BAC
2 : **if** $(\mathsf{id}, \mathsf{global}_t) \in \mathsf{corrupted}$:
3 : *create a batch of* d tokens *using* $\mathsf{TGen}()$ *for* id
4 : $\mathsf{T} := \mathsf{T} \cup \mathsf{token}$
5 : $\mathsf{global}_t := \mathsf{global}_t + 1$
6 : **return** d tokens
7 : **return** $\perp$

Figure 5: Game subroutines and oracles used in several games. Here, d is a fixed parameter known to the moderator.

to their secret key material. The knowledge of the original source of a message is non-transferable in that sense. Additionally, $\mathsf{O}_{\mathsf{b}}^{\mathsf{DENR}}$ provides an interesting guarantee: since the receiver forgery $\mathsf{Forge}_{\mathtt{rec}}$ does not depend on the original sender of a message in any way, then it can be called by *any* user even if they were not participating in the forwarding path of that message. This is a strong claim since the number of possible senders of any particular message in Hecate is now as large as the number of users in the EEMS.

In Hecate, $\mathsf{Forge}_{\mathtt{rec}}$ allows users to forge messages by using their own pre-processing tokens and constructing their own franked messages that they send back to themselves. The receiver does not have any more capabilities than any other user, and in particular the sender, without the secret key of the moderator. In other words, the only thing that a receiver can do is construct the franked message themselves. In Hecate, tokens and franked messages are not bound to the sender's long term key material and the origin of franked messages as a result is indistinguishable without the secret key of the moderator.

In $\mathsf{Forge}_{\mathtt{mod}}$ on the other hand, the moderator can forge messages by using their secret key to produce tokens for any user identity of their choosing, constructing franked messages on their behalf and sending the resulting message to the receiver. This is again a result of how Hecate does not bind the user's long term key material to a franked message.

$\mathsf{DENM}_{\mathsf{b}}^{\mathcal{A}}$

1 : $s_1, s_2 \leftarrow\$\, \mathcal{A}$
2 : $(\mathsf{pk}_{\mathtt{mod}}, \mathsf{sk}_{\mathtt{mod}}) \leftarrow\$\, \mathsf{KGen}(s_1)$
3 : $(\mathsf{pk}_{\mathtt{mod},\sigma}, \mathsf{sk}_{\mathtt{mod},\sigma}) \leftarrow\$\, \mathsf{KGen}(s_2)$
4 : $(\mathsf{pk}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}}) \leftarrow\$\, \mathsf{KGen}(1^n)$
5 : $b' \leftarrow\$\, \mathcal{A}^{\mathsf{O}_{\mathsf{b}}^{\mathsf{DENM}}}(\mathsf{sk}_{\mathtt{mod}}, \mathsf{sk}_{\mathtt{mod},\sigma})$
6 : **return** b'

$\mathsf{DENR}_{\mathsf{b}}^{\mathcal{A}}$

1 : $(\mathsf{pk}_{\mathtt{mod}}, \mathsf{sk}_{\mathtt{mod}}) \leftarrow\$\, \mathsf{KGen}(1^n)$
2 : $(\mathsf{pk}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}}) \leftarrow\$\, \mathsf{KGen}(1^n)$
3 : $b' \leftarrow\$\, \mathcal{A}^{\mathsf{O}_{\mathsf{b}}^{\mathsf{DENR}}}$
4 : **return** b'

$\mathsf{Forge}_{\mathtt{mod}}(\mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{m}, \mathsf{sk}_{\mathtt{mod}})$

1 : $\mathsf{token} := \mathsf{construct}_{\mathsf{token}}(\mathsf{sk}_{\mathtt{mod}}, \mathsf{id}_{\mathtt{src}})$
2 : $\mathsf{m}_{\mathsf{frank}} := \mathsf{construct}_{\mathsf{frank}}(\mathsf{m}, \mathsf{token})$
3 : **return** $\mathsf{m}_{\mathsf{frank}}$

$\mathsf{Forge}_{\mathtt{rec}}(\mathsf{m}, \mathsf{state}_{\mathtt{rec}})$

1 : *fetch* token *from* $\mathsf{state}_{\mathtt{rec}}$
2 : $\mathsf{m}_{\mathsf{frank}} := \mathsf{construct}_{\mathsf{frank}}(\mathsf{m}, \mathsf{token})$
3 : $\mathsf{c}_{\mathsf{frank}} := \mathsf{forge}_{\mathsf{eems}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{state}_{\mathtt{rec}})$
4 : **return** $\mathsf{c}_{\mathsf{frank}}$

$\mathsf{O}_{\mathsf{b}}^{\mathsf{DENM}}(\mathsf{m}, \mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}})$

1 : **if** $b = 0$:
2 : $\quad \mathsf{m}_{\mathsf{frank}} := \mathsf{Forge}_{\mathtt{mod}}(\mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{m}, \mathsf{sk}_{\mathtt{mod}})$
3 : $\quad \mathsf{c}_{\mathsf{frank}} := \mathsf{send}_{\mathsf{eems}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}})$
4 : $\quad \mathsf{m}_{\mathsf{frank}}{}' := \mathsf{receive}_{\mathsf{amf}}(\mathsf{c}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}, \mathsf{sk}_{\mathtt{plat}})$
5 : **else** :
6 : $\quad \mathsf{c}_{\mathsf{frank}} := \mathsf{send}_{\mathsf{amf}}(\mathsf{m}, \mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}})$
7 : $\quad \mathsf{m}_{\mathsf{frank}}{}' := \mathsf{receive}_{\mathsf{amf}}(\mathsf{c}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}, \mathsf{sk}_{\mathtt{plat}})$
8 : **return** $\mathsf{m}_{\mathsf{frank}}{}'$

$\mathsf{O}_{\mathsf{b}}^{\mathsf{DENR}}(\mathsf{m}, \mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}})$

1 : **if** $b = 0$:
2 : $\quad$ *fetch* $\mathsf{state}_{\mathtt{rec}}$
3 : $\quad \mathsf{c}_{\mathsf{frank}} := \mathsf{Forge}_{\mathtt{rec}}(\mathsf{m}, \mathsf{state}_{\mathtt{rec}})$
4 : $\quad \mathsf{m}_{\mathsf{frank}}{}' := \mathsf{receive}_{\mathsf{amf}}(\mathsf{c}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}, \mathsf{sk}_{\mathtt{plat}})$
5 : **else** :
6 : $\quad \mathsf{c}_{\mathsf{frank}} := \mathsf{send}_{\mathsf{amf}}(\mathsf{m}, \mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}})$
7 : $\quad \mathsf{m}_{\mathsf{frank}}{}' := \mathsf{receive}_{\mathsf{amf}}(\mathsf{c}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}, \mathsf{sk}_{\mathtt{plat}})$
8 : **return** $\mathsf{m}_{\mathsf{frank}}{}'$

Figure 6: The security games for Deniability.

5.1.1 Formalizing moderator deniability

Definition 12. *We define the advantage of the adversary in the* $\mathsf{DENM}_{\mathsf{b}}^{\mathcal{A}}$ *game for* Hecate *as:*

$$\mathsf{Adv}^{\mathrm{deniability}_{\mathtt{mod}}}_{\mathsf{Hecate}}(\mathcal{A}) = \left|\Pr\left[\mathsf{DENM}_1^{\mathcal{A}} = 1\right] - \Pr\left[\mathsf{DENM}_0^{\mathcal{A}} = 1\right]\right|.$$

We say that a scheme S is deniable with respect to the moderator if for all PPT adversaries $\mathcal{A}$:

$$\mathsf{Adv}^{\mathrm{deniability}_{\mathtt{mod}}}_{S}(\mathcal{A}) = 0.$$

Theorem 5.1. *Our construction* Hecate *is deniable against a moderator. Moreover the advantage of $\mathcal{A}$ is:*

$$\mathsf{Adv}^{\mathrm{deniability}_{\mathtt{mod}}}_{\mathsf{Hecate}}(\mathcal{A}) = 0.$$

The essence of this theorem is the claim that that Hecate's real send routine is indistinguishable from the moderator's forgery. Intuitively, Hecate achieves moderator deniability because Hecate implements algorithms TGen, Frank and Forward without ever using the user's long term key materials and instead relying on ephemeral keys generated by the moderator themselves. Additionally, the preprocessing token relies on an encryption and signature by the moderator in a way that is not directly bound to the message. This claim holds even against a distinguisher who also has the

$\mathsf{send}_0(n)$	$\mathsf{send}_1(n)$	$\mathsf{send}_2(n)$
1 : $\mathsf{state}_{\mathtt{src}} := \mathsf{retrieve}_{\mathsf{state}}(\mathsf{id}_{\mathtt{src}})$	$\mathsf{construct}_{\mathsf{token}}(\mathsf{sk}_{\mathtt{mod}}, \mathsf{id}_{\mathtt{src}})$	$\mathsf{construct}_{\mathsf{token}}(\mathsf{sk}_{\mathtt{mod}}, \mathsf{id}_{\mathtt{src}})$
2 : *fetch* token *from* $\mathsf{state}_{\mathtt{src}}$	$\mathsf{m}_{\mathsf{frank}} := \mathsf{Frank}(\mathsf{state}_{\mathtt{src}}, \mathsf{m}, \mathsf{id}_{\mathtt{rec}}, \mathsf{token})$	$\mathsf{m}_{\mathsf{frank}} := \mathsf{construct}_{\mathsf{frank}}(\mathsf{m}, \mathsf{token})$
3 : $\mathsf{m}_{\mathsf{frank}} := \mathsf{Frank}(\mathsf{state}_{\mathtt{src}}, \mathsf{m}, \mathsf{id}_{\mathtt{rec}}, \mathsf{token})$	$\mathsf{c}_{\mathsf{frank}} := \mathsf{send}_{\mathsf{eems}}(\mathsf{m}_{\mathsf{frank}})$	$\mathsf{c}_{\mathsf{frank}} := \mathsf{send}_{\mathsf{eems}}(\mathsf{m}_{\mathsf{frank}})$
4 : $\mathsf{c}_{\mathsf{frank}} := \mathsf{send}_{\mathsf{eems}}(\mathsf{m}_{\mathsf{frank}})$	**return** $\mathsf{c}_{\mathsf{frank}}$	**return** $\mathsf{c}_{\mathsf{frank}}$
5 : **return** $\mathsf{c}_{\mathsf{frank}}$		

Figure 7: The hybrid steps modifying $\mathsf{send}_{\mathsf{amf}}$ in the Moderator Deniability game

moderator's secret key – that is, if the moderator chooses to leak their own keys in an attempt to convince the rest of the world about the actions of a sender.

Proof. We show via a series of hybrids that $\mathsf{DENM}_0^{\mathcal{A}} \overset{c}{\approx} \mathsf{DENM}_1^{\mathcal{A}}$ (Figure 6). This effectively boils down to showing that $\mathsf{O}_0^{\mathsf{DENM}} \overset{c}{\approx} \mathsf{O}_1^{\mathsf{DENM}}$, and hence that $\mathsf{send}_{\mathsf{amf}}$ is computationally indistinguishable from $\mathsf{Forge}_{\mathtt{mod}}$ and $\mathsf{send}_{\mathsf{eems}}$ (Figure 6). We start with the $\mathsf{send}_{\mathsf{amf}}$ subroutine.

Game$_1$: In $\mathsf{send}_{\mathsf{amf}}$, we replace "*fetch* token *from* $\mathsf{state}_{\mathtt{src}}$" with the moderator token construction method $\mathsf{construct}_{\mathsf{token}}$ on the source's id $\mathsf{id}_{\mathtt{src}}$. We can do so because the moderator in Hecate does not require any information from a user $\mathsf{id}_{\mathtt{src}}$ in order to construct a token on their behalf. The adversary cannot observe the authentication that occurs between the sender and the moderator since the oracle is acting on behalf of the moderator in this game.

Game$_2$: In $\mathsf{send}_{\mathsf{amf}}$, we can disregard the state passed to Frank and replace it with its instantiation $\mathsf{construct}_{\mathsf{frank}}$. Similarly to $\mathsf{construct}_{\mathsf{token}}$ (and hence TGen), $\mathsf{construct}_{\mathsf{frank}}$ does not require the state of the sender to construct the franked message.

Notice that the resulting game from the prior series of hybrid has transformed $\mathsf{send}_{\mathsf{amf}}$ to look exactly like $\mathsf{Forge}_{\mathtt{mod}}$. The resulting game is identical to $\mathsf{DENM}_0^{\mathcal{A}}$, where only the branch corresponding to $b = 0$ is executed. □

5.1.2 Formalizing receiver deniability

Definition 13. *We define the advantage of the adversary in the* $\mathsf{DENR}_{\mathsf{b}}^{\mathcal{A}}$ *game for* Hecate *as:*

$$\mathsf{Adv}^{\mathsf{deniability}_{\mathtt{rec}}}_{\mathsf{Hecate}}(\mathcal{A}) = \left|\Pr\left[\mathsf{DENR}_1^{\mathcal{A}} = 1\right] - \Pr\left[\mathsf{DENR}_0^{\mathcal{A}} = 1\right]\right|.$$

We say that a scheme S is deniable with respect to a receiver if for all PPT adversaries $\mathcal{A}$:

$$\mathsf{Adv}^{\mathsf{deniability}_{\mathtt{rec}}}_{S}(\mathcal{A}) \leq \mathsf{negl}(n).$$

Theorem 5.2. *Our construction* Hecate *is deniable against a malicious receiver. Moreover the advantage of $\mathcal{A}$ is:*

$$\mathsf{Adv}^{\mathsf{deniability}_{\mathtt{rec}}}_{\mathsf{Hecate}}(\mathcal{A}) \leq \mathsf{Adv}^{\mathsf{eems}_{\mathsf{deniability}}}_{\mathcal{E}}(\mathcal{A}) + \mathsf{Adv}^{\mathsf{enc}_{\mathsf{cpa}}}_{\mathcal{E}}(\mathcal{A}).$$

Proof. The main difference with moderator deniability is that the adversary has to perform a forgery without the secret keys of the moderator. Intuitively, a forger can use her own tokens to create a franked message of her choosing and claim that it came from another source. Users with no access to the moderator's secret key should not be able to verify her claim without breaking the underlying encryption schemes.

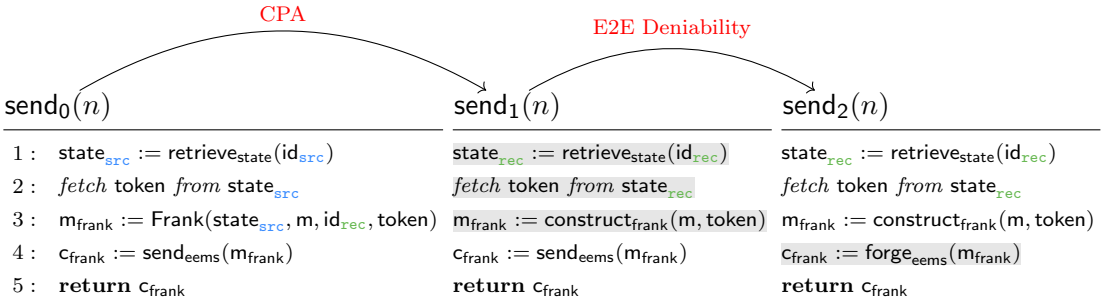

Figure 8: The hybrid steps modifying $\mathsf{send}_{\mathsf{amf}}$ in the Receiver Deniability game

We show via a series of hybrids that $\mathsf{DENR}_0^{\mathcal{A}} \overset{c}{\approx} \mathsf{DENR}_1^{\mathcal{A}}$, and more precisely show that $\mathsf{O}_0^{\mathsf{DENR}} \overset{c}{\approx} \mathsf{O}_1^{\mathsf{DENR}}$, and hence that $\mathsf{send}_{\mathsf{amf}}$ is computationally indistinguishable from $\mathsf{Forge}_{\mathtt{rec}}$ (Figure 6).

Game_0: We start with $\mathsf{O}_1^{\mathsf{DENR}}$, we focus on the if-else branch corresponding to $b = 1$ since its the only one executed. We can disregard the other branch.

Game_1: In $\mathsf{O}_1^{\mathsf{DENR}}$, we replace $\mathsf{state}_{\mathtt{src}}$ in the $\mathsf{send}_{\mathsf{amf}}$ subroutine with $\mathsf{state}_{\mathtt{rec}}$ (lines 1, 2, 3). This is equivalent to replacing $\mathsf{send}_{\mathsf{amf}}$ with its instantiation $\mathsf{construct}_{\mathsf{frank}}$. Since the adversary does not have access to the moderator's secret key, then $\mathcal{A}$ has negligible advantage in distinguishing between different x_1 values in the constructed pre-processing token and the franked message. Specifically, they cannot distinguish between an encryption of $\mathsf{id}_{\mathtt{src}}$ and $\mathsf{id}_{\mathtt{rec}}$ without breaking the CPA security of the symmetric key encryption scheme used by the moderator during preprocessing. We formally show this by constructing an adversary $\mathcal{B}$ that can break the CPA security game $\mathsf{ENC}_{\mathsf{cpa}}^{\mathcal{A}}$. When adversary $\mathcal{A}$ queries oracle $\mathsf{O}_1^{\mathsf{DENR}}$, $\mathcal{B}$ queries $\mathsf{O}_{\mathsf{cpa}}^{\mathsf{enc}}$ with $\mathsf{id}_{\mathtt{src}}$ and $\mathsf{id}_{\mathtt{rec}}$, constructs the franked message with the resulting cipher-text. When $\mathcal{A}$ submits a choice bit b, $\mathcal{B}$ returns the same bit to $\mathsf{ENC}_{\mathsf{cpa}}^{\mathcal{A}}$ and succeeds if and only if $\mathcal{A}$ can distinguish between Games 0 and 1.

Game_2: We replace $\mathsf{send}_{\mathsf{eems}}$ with $\mathsf{forge}_{\mathsf{eems}}$ since the underlying EEMS provides receiver deniability and the receiver hence can forge a channel with themselves.

The resulting game is identical to $\mathsf{DENR}_0^{\mathcal{A}}$, where only the branch corresponding to $b = 0$ is executed. □

5.2 Anonymity

The anonymity property of the scheme models each user's restricted access to a particular message and its metadata. No user should be able to learn anything beyond what they would learn in the underlying encrypted messenger's environment.

5.2.1 Anonymity with respect to the receiver

This security property guarantees that *receivers should not be able to learn any other member of the forwarding path of a message beyond their direct neighbors.* In this property, we assume that senders and forwarders of a message are honest and wish to hide themselves from non-neighboring recipients in the presence of an honest moderator. We model this property in the $\mathsf{ANON}_{\mathtt{rec},\mathsf{b}}^{\mathcal{A}}$ game. The adversary can send and forward messages between parties using $\mathsf{O}_{\mathsf{anon},\mathsf{b}}^{\mathsf{send}}$, $\mathsf{O}_{\mathsf{anon},\mathsf{b}}^{\mathsf{fwd}}$ and $\mathsf{O}_{\mathsf{anon}}^{\mathsf{deliver}}$ and is provided with the resulting franked message $\mathsf{m}_{\mathsf{frank}}$. In this game, we do not attempt to hide chat participants from on another. To that end, both $\mathsf{O}_{\mathsf{anon},\mathsf{b}}^{\mathsf{send}}$ and $\mathsf{O}_{\mathsf{anon},\mathsf{b}}^{\mathsf{fwd}}$ call $\mathsf{check}_{\mathsf{topology}}$ to ensure that the adversary provided the same pairs of senders and recipients when either of the provided

$\mathsf{ANON}^{\mathcal{A}}_{\mathtt{rec},\mathsf{b}}$

$s \leftarrow\$ \mathcal{A}$
$(\mathsf{pk}_{\mathtt{mod}}, \mathsf{sk}_{\mathtt{mod}}) \leftarrow\$ \mathsf{KGen}(1^n)$
$(\mathsf{pk}_{\mathtt{mod},\sigma}, \mathsf{sk}_{\mathtt{mod},\sigma}) \leftarrow\$ \mathsf{KGen}(1^n)$
$(\mathsf{pk}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}}) \leftarrow\$ \mathsf{KGen}(s)$
$b' := \mathcal{A}^{\mathsf{O}^{\mathsf{send}}_{\mathsf{b}}, \mathsf{O}^{\mathsf{fwd}}_{\mathsf{b}}, \mathsf{O}^{\mathsf{deliver}}}(\mathsf{sk}_{\mathtt{plat}})$
return b'

$\mathsf{ANON}^{\mathcal{A}}_{\mathtt{mod},\mathsf{b}}$

$s_1, s_2, s_3 \leftarrow\$ \mathcal{A}$
$(\mathsf{pk}_{\mathtt{mod}}, \mathsf{sk}_{\mathtt{mod}}) \leftarrow\$ \mathsf{KGen}(s_1)$
$(\mathsf{pk}_{\mathtt{mod},\sigma}, \mathsf{sk}_{\mathtt{mod},\sigma}) \leftarrow\$ \mathsf{KGen}(s_2)$
$(\mathsf{pk}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}}) \leftarrow\$ \mathsf{KGen}(s_3)$
$b' := \mathcal{A}^{\mathsf{O}^{\mathsf{fwd}}_{\mathsf{b}}, \mathsf{O}^{\mathsf{deliver}}}(\mathsf{sk}_{\mathtt{mod}}, \mathsf{sk}_{\mathtt{mod},\sigma}, \mathsf{sk}_{\mathtt{plat}})$
return b'

$\mathsf{O}^{\mathsf{deliver}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}})$

$\mathsf{m}_{\mathsf{frank}} := \mathsf{receive}_{\mathsf{amf}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}})$
$\mathsf{R} := \mathsf{R} \cup \langle \mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{id}_{\mathtt{rec}} \rangle$
return $\mathsf{m}_{\mathsf{frank}}$

$\mathsf{check}_{\mathsf{received}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}0}, \mathsf{id}_{\mathtt{rec}1})$

if $\langle \mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}0}, \mathsf{id}_{\mathtt{rec}1} \rangle \notin \mathsf{R}$:
  return $\perp$
return 1

$\mathsf{O}^{\mathsf{send}}_{\mathsf{anon},\mathsf{b}}(\mathsf{m}, \langle \mathsf{id}_{\mathtt{src}0}, \mathsf{id}_{\mathtt{rec}0} \rangle, \langle \mathsf{id}_{\mathtt{src}1}, \mathsf{id}_{\mathtt{rec}1} \rangle, \mathsf{time}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}})$

if $\mathsf{check}_{\mathsf{topology}}(\langle \mathsf{id}_{\mathtt{src}0}, \mathsf{id}_{\mathtt{rec}0} \rangle, \langle \mathsf{id}_{\mathtt{src}1}, \mathsf{id}_{\mathtt{src}1} \rangle) = \perp$:
  return $\perp$
$\mathsf{c}_{\mathsf{frank}} = \mathsf{send}_{\mathsf{amf}}(\mathsf{m}, \mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}})$
$\mathsf{m}_{\mathsf{frank}} := \mathsf{receive}_{\mathsf{amf}}(\mathsf{c}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}, \mathsf{sk}_{\mathtt{plat}})$
$\mathsf{R} := \mathsf{R} \cup \langle \mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}0}, \mathsf{id}_{\mathtt{rec}1} \rangle$
return $\mathsf{m}_{\mathsf{frank}}$

$\mathsf{O}^{\mathsf{fwd}}_{\mathsf{anon},\mathsf{b}}(\mathsf{m}_{\mathsf{frank}}, \langle \mathsf{id}_{\mathtt{fwd}0}, \mathsf{id}_{\mathtt{rec}0} \rangle, \langle \mathsf{id}_{\mathtt{fwd}1}, \mathsf{id}_{\mathtt{rec}1} \rangle, \mathsf{time}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}})$

if $\mathsf{check}_{\mathsf{topology}}(\langle \mathsf{id}_{\mathtt{fwd}0}, \mathsf{id}_{\mathtt{rec}0} \rangle, \langle \mathsf{id}_{\mathtt{fwd}1}, \mathsf{id}_{\mathtt{rec}1} \rangle) = \perp$:
  return $\perp$
if $\mathsf{check}_{\mathsf{received}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}0}, \mathsf{id}_{\mathtt{rec}1}) = \perp$:
  return $\perp$
$\mathsf{c}_{\mathsf{frank}} = \mathsf{fwd}_{\mathsf{amf}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}})$
$\mathsf{m}_{\mathsf{frank}}{}' := \mathsf{receive}_{\mathsf{amf}}(\mathsf{c}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}, \mathsf{sk}_{\mathtt{plat}})$
$\mathsf{R} := \mathsf{R} \cup \langle \mathsf{m}_{\mathsf{frank}}{}', \mathsf{id}_{\mathtt{rec}0}, \mathsf{id}_{\mathtt{rec}1} \rangle$
return $\mathsf{m}_{\mathsf{frank}}{}'$

$\mathsf{check}_{\mathsf{topology}}(\langle \mathsf{id}_{\mathtt{src}0}, \mathsf{id}_{\mathtt{rec}0} \rangle, \langle \mathsf{id}_{\mathtt{src}1}, \mathsf{id}_{\mathtt{rec}1} \rangle)$

if $\mathsf{id}_{\mathtt{src}0} \vee \mathsf{id}_{\mathtt{src}1} \in \mathsf{corrupted}$:
  return $\perp$
if $\mathsf{id}_{\mathtt{rec}0} \vee \mathsf{id}_{\mathtt{rec}1} \in \mathsf{corrupted}$:
  if $\mathsf{id}_{\mathtt{rec}0} \neq \mathsf{id}_{\mathtt{rec}1} \vee \mathsf{id}_{\mathtt{src}0} \neq \mathsf{id}_{\mathtt{src}1}$:
    return $\perp$
return *true*

Figure 9: The security games for Anonymity with respect to the Receiver and Moderator.

receivers is corrupted. Without this check, $\mathcal{A}$ would trivially win the game by inspecting which of their provided correspondents sent the message or who received it. Additionally, $\mathsf{check}_{\mathsf{topology}}$ ensures that $\mathsf{O}^{\mathsf{send}}_{\mathsf{anon,b}}$ and $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{anon,b}}$ handle messages relayed from honest senders and forwarders. $\mathsf{O}^{\mathsf{deliver}}_{\mathsf{anon}}$ on the other hand allows $\mathcal{A}$ to send franked messages from corrupted nodes to an honest receiver. In $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{anon,b}}$, we also check that the queried $\mathsf{m}_{\mathsf{frank}}$ was initially received by either of the provided forwarders using the $\mathsf{check}_{\mathsf{received}}$ method in order to model the actual behavior of messages forwarders. Both $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{anon,b}}$ and $\mathsf{O}^{\mathsf{deliver}}_{\mathsf{anon}}$ check that the provided franked message are consistent with the recipient's state by calling Vf in the $\mathsf{receive}_{\mathsf{amf}}$ subroutine, thus eliminating any trivial wins that arise from malformed messages.

All three oracles, along with the corruption oracles $\mathsf{O}^{\mathsf{corrupt}}$ and $\mathsf{O}^{\mathsf{request}}$, allow the adversary to adaptively build any two message paths of their choice (modulo the topological restrictions) and receive the transcript of the chosen path, effectively encompassing the full power of a malicious recipient that may intercept messages along the path. The adversary is tasked with guessing which of the two message paths, with honest sources/roots, was chosen by the game. If they fail to distinguish between them, then they would have failed to determine the original sender of that message and the anonymity of that user and honest forwarder along the path is preserved.

Definition 14 (Anonymity w.r.t. receivers)**.** *We define the advantage of the adversary in the* $\mathsf{ANON}^{\mathcal{A}}_{\mathtt{rec}}$ *game for a source tracing and asymmetric message franking scheme* S *as:*

$$\mathsf{Adv}^{\mathrm{anonymity}_{\mathtt{rec}}}_{S}(\mathcal{A}) = \left|\Pr\left[\mathsf{ANON}^{\mathcal{A}}_{\mathtt{rec},1} = 1\right] - \Pr\left[\mathsf{ANON}^{\mathcal{A}}_{\mathtt{rec},0} = 1\right]\right|.$$

We say that a scheme S *is anonymous w.r.t. receivers if for all PPT adversaries* $\mathcal{A}$*:*

$$\mathsf{Adv}^{\mathrm{anonymity}_{\mathtt{rec}}}_{S}(\mathcal{A}) \leq \mathsf{negl}(n).$$

Theorem 5.3. *Our scheme* Hecate *is anonymous w.r.t. receiver. Moreover the advantage of* $\mathcal{A}$ *is:*

$$\mathsf{Adv}^{\mathrm{anonymity}_{\mathtt{rec}}}_{\mathsf{Hecate}} \leq \mathsf{Adv}^{\mathrm{enc}_{\mathrm{cpa}}}_{\mathsf{Hecate}}.$$

Informally, this theorem holds because the preprocessing tokens in Hecate only contain any information about the original sender's identity in encrypted form; without access to the moderator's secret key, a receiver can't distinguish between tokens that originate from different senders. Additionally, Hecate stores no information about forwarders of a message at all, thereby guaranteeing their anonymity as well. We provide a rigorous proof of this theorem below.

Proof. At a high level, Hecate guarantees sender anonymity for the same reason it achieves receiver deniability: with no access to the moderator's secret key, message recipients cannot tell who the originator of a message is without breaking the underlying encryption scheme. Additionally, since the commitment scheme used for envelope commitments is hiding, then access to the platforms secret key does not reveal anything about the sender of a message. On the other hand, forwarding franked messages in Hecate does not utilize the forwarder's state or identity. In other words, no attribute in any franked message can be traced back to a forwarder guaranteeing forwarder anonymity.

We show via a series of hybrids that $\mathsf{ANON}^{\mathcal{A}}_{\mathtt{rec},0} \overset{c}{\approx} \mathsf{ANON}^{\mathcal{A}}_{\mathtt{rec},1}$.

Game$_0$: We start with $\mathsf{ANON}^{\mathcal{A}}_{\mathtt{rec},\mathsf{b}}$ with $b = 0$.

Game$_1$: We replace $\mathsf{state}_{\mathtt{src},0}$ with $\mathsf{state}_{\mathtt{src},1}$ in the $\mathsf{send}_{\mathsf{amf}}$ subroutine that is called by $\mathsf{O}^{\mathsf{send}}_{\mathsf{anon,b}}$. The moderator in Hecate binds a sender to a token by encrypting their identity on line 3 in

Figure 10: The hybrid steps of $\mathsf{send}_{\mathsf{amf}}$ in the $\mathsf{ANON}^{\mathcal{A}}_{\mathtt{rec},\mathsf{b}}$ game. We refer the reader to $\mathsf{ANON}^{\mathcal{A}}_{\mathtt{mod},\mathsf{b}}$ hybrids for $\mathsf{fwd}_{\mathsf{amf}}$'s hybrids.

$\mathsf{construct}_{\mathsf{token}}$ (Figure 4) and generating the sub-token x_1. Without the moderator's key, the adversary cannot decrypt the identity of the sender within x_1 in the token constructed in $\mathsf{send}_{\mathsf{amf}}$ (Figure 5) on line 2 , without breaking the CPA security of the underlying encryption scheme. The formalism here is similar to Game 1 of $\mathsf{DEN}^{\mathcal{A}}_{\mathtt{rec}}$.

Game_2: We replace $\mathsf{state}_{\mathtt{src},0}$ with $\mathsf{state}_{\mathtt{src},1}$ in Frank in the $\mathsf{send}_{\mathsf{amf}}$ subroutine that is called by $\mathsf{O}^{\mathsf{send}}_{\mathsf{anon},\mathsf{b}}$. The user's state and long term keys are never used during the construction of the franked message via $\mathsf{construct}_{\mathsf{frank}}$ in Hecate. $\mathsf{m}_{\mathsf{frank}}$ is only bound to a particular user by x_1 which we have discussed and handled in the previous hybrid.

Game_3: We replace $\mathsf{id}_{\mathtt{rec}0}$ with $\mathsf{id}_{\mathtt{rec}1}$ in both $\mathsf{O}^{\mathsf{send}}_{\mathsf{anon},0}$ and $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{anon},0}$. In $\mathsf{O}^{\mathsf{send}}_{\mathsf{anon}}$, $\mathsf{check}_{\mathsf{topology}}$ enforces that $\mathsf{id}_{\mathtt{src}0} = \mathsf{id}_{\mathtt{src}1}$ and $\mathsf{id}_{\mathtt{rec}0} = \mathsf{id}_{\mathtt{rec}1}$ when either of the receivers is corrupted (and similarly for $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{anon}}$). If both receivers are honest, then we can make this replacement because Hecate does not use any information related to the receiver of a message when constructing or forwarding a franked message (see $\mathsf{construct}_{\mathsf{frank}}$ and $\mathsf{construct}_{\mathsf{fwd}}$ respectively) and the adversary cannot hence distinguish between both Games 2 and 3.

We have shown that $\mathsf{O}^{\mathsf{send}}_{\mathsf{anon},0} \overset{c}{\approx} \mathsf{O}^{\mathsf{send}}_{\mathsf{anon},1}$ in Hecate since both oracles are now identical. The next series of hybrids are similar to the ones we have already seen.

Game_4: We replace $\mathsf{state}_{\mathtt{fwd},0}$ with $\mathsf{state}_{\mathtt{fwd},0}$ in $\mathsf{Forward}$ in the $\mathsf{fwd}_{\mathsf{amf}}$ subroutine that is called by $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{anon},\mathsf{b}}$. The reason we can do so is two folds: (1) Hecate does not use the forwarder states in constructing the franked message, (2) when the receiver of a forwarded message is corrupted,

$\mathsf{check}_{\mathsf{topology}}$ enforces that $\mathsf{id}_{\mathtt{fwd},0} = \mathsf{id}_{\mathtt{fwd},1}$ and $\mathsf{id}_{\mathtt{rec}0} = \mathsf{id}_{\mathtt{rec}1}$. In either case, the source or forwarder of a message have to be honest.

The resulting game is identical to $\mathsf{ANON}^{\mathcal{A}}_{\mathtt{rec},1}$. We conclude that Hecate is sender and forwarder anonymous and the advantage of the adversary is equal to that of the $\mathsf{ENC}^{\mathcal{A}}_{\mathsf{cca}}$. □

5.2.2 Anonymity with respect to the moderator

This property ensures that *the moderator should not be able to learn members of the forwarding path of a reported message beyond the neighbors of colluding receivers and the reported source.* Here, honest forwarders want to be assured that, when their direct contacts are honest, only their neighboring recipients know that they forwarded a specific message. Since the moderator can directly trace the source of a franked message after receiving it, we can additionally assume that for all intents and purpose the sender of a message in this game is also colluding with them. We show this property via the $\mathsf{ANON}^{\mathcal{A}}_{\mathtt{mod},1}$ game, where the adversary now only has access to $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{anon},\mathsf{b}}$, $\mathsf{O}^{\mathsf{deliver}}_{\mathsf{anon},\mathsf{b}}$, $\mathsf{O}^{\mathsf{corrupt}}$, $\mathsf{O}^{\mathsf{request}}$ oracles. Contrary to the prior property, the adversary now chooses the moderator's secret keys and controls the root of the message path. They must now, as a result, use $\mathsf{O}^{\mathsf{deliver}}_{\mathsf{anon},\mathsf{b}}$ to deliver messages between the compromised sender and honest recipients. If, by the end of the game, the adversary cannot guess the correct message path chosen by the game then they could not have distinguished between the different honest forwarders provided in each message path. The adversary in this game is strictly stronger than the one in $\mathsf{ANON}^{\mathcal{A}}_{\mathtt{rec},\mathsf{b}}$ because of the gained source tracing capability from the moderators secret keys, and hence implies the anonymity of forwarders in the presence of malicious receivers and an honest moderator. However, $\mathsf{ANON}^{\mathcal{A}}_{\mathtt{rec},\mathsf{b}}$ independently makes that guarantee because of the way $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{anon},\mathsf{b}}$ handles interactions between honest users. By the end of $\mathsf{ANON}^{\mathcal{A}}_{\mathtt{rec},\mathsf{b}}$, if the adversary could not guess the message path chosen by the game, then they could not distinguish between honest forwarders with honest neighbors as well.

Definition 15 (Anonymity w.r.t. moderator)**.** *We define the advantage of the adversary in the* $\mathsf{ANON}^{\mathcal{A}}_{\mathtt{mod}}$ *game for a source tracing and asymmetric message franking scheme S as:*

$$\mathsf{Adv}^{\mathrm{anonymity}_{\mathtt{mod}}}_{S}(\mathcal{A}) = \left|\Pr\left[\mathsf{ANON}^{\mathcal{A}}_{\mathtt{mod},1} = 1\right] - \Pr\left[\mathsf{ANON}^{\mathcal{A}}_{\mathtt{mod},0} = 1\right]\right|.$$

We say that a scheme S is anonymous w.r.t. receivers if for all PPT adversaries $\mathcal{A}$:

$$\mathsf{Adv}^{\mathrm{anonymity}_{\mathtt{mod}}}_{S}(\mathcal{A}) \leq \mathsf{negl}(n).$$

Theorem 5.4. *Our scheme* Hecate *is anonymous w.r.t. moderator. Moreover the advantage of $\mathcal{A}$ is:*

$$\mathsf{Adv}^{\mathrm{anonymity}_{\mathtt{mod}}}_{\mathsf{Hecate}}(\mathcal{A}) = 0.$$

Proof. In this proof we need to show that a corrupted moderator cannot learn the forwarding path of a message beyond senders/forwarders/receivers that she has already corrupted and their neighbors. The reader is referred to the similar receiver anonymity proof for more details.

We can show via a series of hybrids that $\mathsf{ANON}^{\mathcal{A}}_{\mathtt{mod},0} \overset{\mathrm{c}}{\approx} \mathsf{ANON}^{\mathcal{A}}_{\mathtt{mod},1}$.

Game_0: We start with $\mathsf{ANON}^{\mathcal{A}}_{\mathtt{mod},0}$ with $b = 0$.

Game_1: We replace $\mathsf{state}_{\mathtt{fwd},0}$ with $\mathsf{state}_{\mathtt{fwd},1}$ in Forward in the $\mathsf{fwd}_{\mathsf{amf}}$ subroutine that is called by $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{fwd},\mathsf{b}}$. The reason we can do so is two folds: (1) Hecate does not use the forwarder states

send_0

1 : $\mathsf{state}_{\mathtt{fwd},0} := \mathsf{retrieve}_{\mathsf{state}}(\mathsf{id}_{\mathtt{fwd},0})$
2 : $\mathsf{m}_{\mathsf{frank}}' := \mathsf{Forward}(\mathsf{state}_{\mathtt{fwd},0}, \mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec},0})$
3 : $\mathsf{c}_{\mathsf{frank}} := \mathsf{send}_{\mathsf{eems}}(\mathsf{m}_{\mathsf{frank}}')$
4 : **return** $\mathsf{c}_{\mathsf{frank}}$

send_1

1 : $\mathsf{state}_{\mathtt{fwd},1} := \mathsf{retrieve}_{\mathsf{state}}(\mathsf{id}_{\mathtt{fwd},1})$
2 : $\mathsf{m}_{\mathsf{frank}}' := \mathsf{Forward}(\mathsf{state}_{\mathtt{fwd},1}, \mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec},0})$
3 : $\mathsf{c}_{\mathsf{frank}} := \mathsf{send}_{\mathsf{eems}}(\mathsf{m}_{\mathsf{frank}}')$
4 : **return** $\mathsf{c}_{\mathsf{frank}}$

send_2

1 : $\mathsf{state}_{\mathtt{fwd},1} := \mathsf{retrieve}_{\mathsf{state}}(\mathsf{id}_{\mathtt{fwd},1})$
2 : $\mathsf{m}_{\mathsf{frank}}' := \mathsf{Forward}(\mathsf{state}_{\mathtt{fwd},1}, \mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec},1})$
3 : $\mathsf{c}_{\mathsf{frank}} := \mathsf{send}_{\mathsf{eems}}(\mathsf{m}_{\mathsf{frank}}')$
4 : **return** $\mathsf{c}_{\mathsf{frank}}$

Figure 11: The hybrid steps of $\mathsf{send}_{\mathsf{amf}}$ in the $\mathsf{ANON}^{\mathcal{A}}_{\mathtt{mod},\mathsf{b}}$ game. We refer the reader to the moderator anonymity game for $\mathsf{fwd}_{\mathsf{amf}}$'s hybrids.

in constructing the franked message, (2) when the receiver of a forwarded message is corrupted, $\mathsf{check}_{\mathsf{topology}}$ enforces that $\mathsf{id}_{\mathtt{fwd},0} = \mathsf{id}_{\mathtt{fwd},1}$ and (3) that both forwarders passed to $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{anon},0}$ are honest.

Game$_3$: We can replace $\mathsf{id}_{\mathtt{rec}0}$ with $\mathsf{id}_{\mathtt{rec}1}$ in the $\mathsf{fwd}_{\mathsf{amf}}$ and $\mathsf{receive}_{\mathsf{amf}}$ subroutines that are called by $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{anon},\mathsf{b}}$. $\mathsf{check}_{\mathsf{topology}}$ enforces that $\mathsf{id}_{\mathtt{rec}0} = \mathsf{id}_{\mathtt{rec}1}$ when either of the receivers is corrupted (and similarly for $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{anon},0}$). If both receivers are honest, then we can make this replacement because Hecate does not use any information related to the receiver of a message when constructing or forwarding a franked message (see $\mathsf{construct}_{\mathsf{fwd}}$) and the adversary cannot hence distinguish between both Games 2 and 3.

We can conclude that $\mathsf{ANON}^{\mathcal{A}}_{\mathtt{mod},0}$ becomes indistinguishable from $\mathsf{ANON}^{\mathcal{A}}_{\mathtt{mod},1}$. □

Note that our construction and security games do not consider forwarding graphs (i.e. trees with cycles). In those cases, users can identify that the same message was forwarded to them multiple times, a property called *tree linkability* in prior work [55]. Additionally, we allow users (but not the platform or the moderator) to distinguish between a sent and a forwarded message as is the case in several messaging system. We think an exciting opportunity for future work is to combine the ideas in this paper with the tree unlinkability scheme by Peale et al. [55].

5.3 Forward Secrecy and Message Confidentiality

Message confidentiality dictates that *any party not involved in the creation, reception or reporting of a message should not be able to learn anything about the message.* Moreover, forward secrecy guarantees that *corrupted users should be guaranteed confidentiality of all their messages and interactions prior to the time of compromise.* Recall that for the purpose of this work, we consider the state of users to consist entirely of their key material and their tokens. Put another way, our claims in this section rely on the message confidentiality and forward secrecy of the underlying EEMS, and they can only guarantee confidentiality for messages that have been securely deleted from the local device prior to the compromise event.

We provide a combined definition of message confidentiality and forward security in Figure 12. Message confidentiality is ensured because the $\mathsf{CONF\text{-}FS}^{\mathcal{A}}_{\mathsf{b}}$ game requires that content moderation does not break CCA security. Additionally, forward security is ensured because the adversary in

$\mathsf{CONF\text{-}FS}_{\mathsf{b}}^{\mathcal{A}}$

$s_1, s_2 \leftarrow\$ \mathcal{A}$
$(\mathsf{pk}_{\mathtt{mod}}, \mathsf{sk}_{\mathtt{mod}}) \leftarrow\$ \mathsf{KGen}(s_1)$
$(\mathsf{pk}_{\mathtt{mod},\sigma}, \mathsf{sk}_{\mathtt{mod},\sigma}) \leftarrow\$ \mathsf{KGen}(s_2)$
$\mathsf{F} := \emptyset$
$b' := \mathcal{A}^{\mathsf{O}^*}(\mathsf{sk}_{\mathtt{mod}}, \mathsf{sk}_{\mathtt{mod},\sigma})$
return b'

$\mathsf{O}^{\mathsf{decrypt}}_{\mathsf{conf\text{-}fs}}(\mathsf{c}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}})$

if $\mathsf{c}_{\mathsf{frank}} \in \mathsf{C}$:
  return $\perp$
$\mathsf{m}_{\mathsf{frank}} := \mathsf{receive}_{\mathsf{amf}}(\mathsf{c}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}, \mathsf{sk}_{\mathtt{plat}})$
return $\mathsf{m}_{\mathsf{frank}}$

$\mathsf{O}^{\mathsf{send}}_{\mathsf{conf\text{-}fs,b}}(\mathsf{m}_0, \mathsf{m}_1, \mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}})$

if $\forall \mathsf{id} \in \{\mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}}\}, \mathsf{id} \in \mathsf{corrupted}$:
  return $\perp$
$\mathsf{c}_{\mathsf{frank}} := \mathsf{send}_{\mathsf{amf}}(\mathsf{m}_b, \mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}})$
$\mathsf{C} := \mathsf{C} \cup \mathsf{c}_{\mathsf{frank}}$
return $\mathsf{c}_{\mathsf{frank}}$

$\mathsf{O}^{\mathsf{fwd}}_{\mathsf{conf\text{-}fs,b}}(\mathsf{m}_{\mathsf{frank}0}, \mathsf{m}_{\mathsf{frank}1}, \mathsf{id}_{\mathtt{fwd}}, \mathsf{id}_{\mathtt{rec}})$

if $\forall \mathsf{id} \in \{\mathsf{id}_{\mathtt{fwd}}, \mathsf{id}_{\mathtt{rec}}\}, \mathsf{id} \in \mathsf{corrupted}$:
  return $\perp$
$\mathsf{c}_{\mathsf{frank}} := \mathsf{fwd}_{\mathsf{amf}}(\mathsf{m}_{\mathsf{frank}\,b}, \mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}})$
return $\mathsf{c}_{\mathsf{frank}}$

Figure 12: The security games for Message Confidentiality. Where O^* refers to $\mathsf{O}^{\mathsf{send}}_{\mathsf{conf\text{-}fs,b}}$, $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{conf\text{-}fs,b}}$, $\mathsf{O}^{\mathsf{decrypt}}$, $\mathsf{O}^{\mathsf{corrupt}}$ and $\mathsf{O}^{\mathsf{request}}$.

the $\mathsf{CONF\text{-}FS}_{\mathsf{b}}^{\mathcal{A}}$ game is allowed to corrupt any user of their choice, in particular they can corrupt users who had previously honestly interacted using $\mathsf{O}^{\mathsf{send}}_{\mathsf{conf\text{-}fs,b}}$ and $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{conf\text{-}fs,b}}$. The game requires that the adversary cannot learn anything about their previously exchanged honest messages, their prior keys, or states.

In this game, we note an important type difference between the franked messages $\mathsf{m}_{\mathsf{frank}}$ returned by $\mathsf{O}^{\mathsf{decrypt}}_{\mathsf{conf\text{-}fs}}$ on one hand, and the franked cipher $\mathsf{c}_{\mathsf{frank}}$ returned by $\mathsf{O}^{\mathsf{send}}_{\mathsf{conf\text{-}fs,b}}$ and $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{conf\text{-}fs,b}}$ on the other. $\mathsf{O}^{\mathsf{send}}_{\mathsf{conf\text{-}fs,b}}$ and $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{conf\text{-}fs,b}}$ produce a franked cipher that is handed out to the platform for stamping before it gets relayed back to the receiver. In other words, the platform cannot read any part of $\mathsf{c}_{\mathsf{frank}}$ that is not intended for it. $\mathsf{O}^{\mathsf{decrypt}}_{\mathsf{conf\text{-}fs}}$ returns the franked message after it has been delivered and hence decrypted by the receiver. If a content moderation scheme does not handle the distinction between the franked message and franked cipher properly, by say appending the id of the sender to the envelope, then the adversary should be able to easily win the game.

Definition 16 (Message Confidentiality)**.** *We define the advantage of the adversary in the* $\mathsf{CONF\text{-}FS}_{\mathsf{b}}^{\mathcal{A}}$ *game for a source tracking and asymmetric message franking scheme* S *as:*

$$\mathsf{Adv}_{S}^{\mathrm{confidentiality_m}}(\mathcal{A}) = \left|\Pr\left[\mathsf{CONF\text{-}FS}_{1}^{\mathcal{A}} = 1\right] - \Pr\left[\mathsf{CONF\text{-}FS}_{0}^{\mathcal{A}} = 1\right]\right|.$$

We say that a scheme S *is message confidential if for all PPT adversaries* $\mathcal{A}$*::*

$$\mathsf{Adv}_{S}^{\text{conf-fs}} \leq \mathsf{negl}(n).$$

Theorem 5.5. *Our scheme* Hecate *is message confidential. Moreover the advantage of* $\mathcal{A}$ *is:*

$$\mathsf{Adv}^{\text{conf-fs}}_{\mathsf{Hecate}} \leq \mathsf{Adv}^{\mathrm{hiding_{com}}}_{\mathsf{Hecate}} + \mathsf{Adv}^{\mathrm{enc_{cca}}}_{\mathsf{Hecate}}.$$

Hecate constructs franked messages by appending tokens to the payload of the message (see $\mathsf{construct}_{\mathsf{frank}}$ in figure 4), and by adding a commitment and timestamp to its envelope (see $\mathsf{stamp}_{\mathsf{time}}$ Figure 4). The tokens are encrypted alongside the plaintext message. The identifying content of the commitment is encrypted, and we rely on the hiding properties of the commitment scheme and

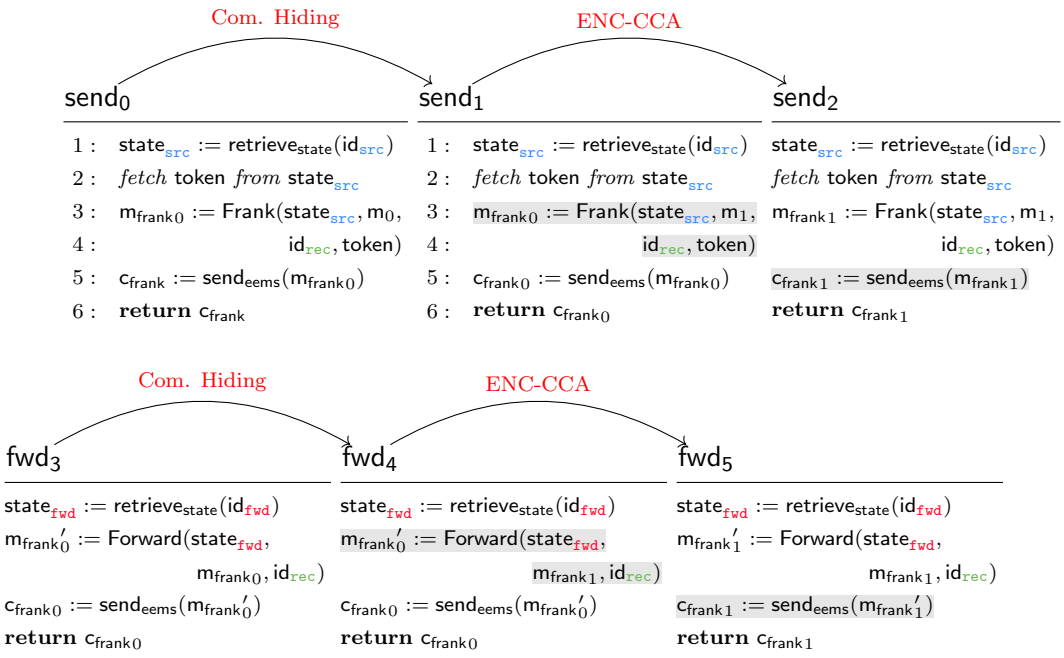

Figure 13: The hybrid steps of the $\mathsf{CONF\text{-}FS}_{\mathsf{b}}^{\mathcal{A}}$ game

the security of the connection that exists between parties and the platform. No part of a Hecate franked message can therefore break this security property. We prove this theorem in detail in what follows.

Proof. In this proof, we show that Hecate does not break the CCA security of the underlying cryptographic scheme.

Game$_0$: We start with $\mathsf{CONF\text{-}FS}_0^{\mathcal{A}}$.

Game$_1$: In $\mathsf{send}_{\mathsf{amf}}$ (called by $\mathsf{O}^{\mathsf{send}}_{\mathsf{conf\text{-}fs,b}}$), we replace m_0 with m_1 in Frank on line 3. Since the adversary does not have access to r, then they cannot only decommit com with negligible property equal to breaking the secrecy of the commitment scheme. We construct $\mathcal{B}$ that can uses an $\mathcal{A}$ that can distinguish between both games, to break the secrecy of the commitment scheme. When $\mathcal{A}$ calls $\mathsf{O}^{\mathsf{send}}_{\mathsf{b}}$ on messages m_0 and m_1, $\mathcal{B}$ queries $\mathsf{O}^{\mathsf{com\text{-}hiding}}$ with both messages and uses the returned com to construct the franked cipher on behalf of the moderator and the users $\mathsf{c}_{\mathsf{frank}}$. If $\mathcal{A}$ succeeds in picking a choice bit, then $\mathcal{B}$ will pick the same choice bit in $\mathsf{Com}^{\mathcal{A}}_{\mathsf{hiding}}$ and will win with at least the same advantage.

Game$_2$: In $\mathsf{send}_{\mathsf{amf}}$ (called by $\mathsf{O}^{\mathsf{send}}_{\mathsf{conf\text{-}fs,b}}$), we replace $\mathsf{m}_{\mathsf{frank}0}$ with $\mathsf{m}_{\mathsf{frank}1}$ in Frank. Since the adversary does not have access to the source $\mathsf{id}_{\mathtt{src}}$'s encryption secret keys, then they can only notice this change with advantage equal to breaking the CCA security $\mathsf{ENC}^{\mathcal{A}}_{\mathsf{cca}}$ of $\mathsf{send}_{\mathsf{eems}}$ and decrypting the contents of $\mathsf{c}_{\mathsf{frank}}$. When $\mathcal{A}$ calls $\mathsf{O}^{\mathsf{send}}_{\mathsf{conf\text{-}fs,b}}$ on messages $\mathsf{m}_{\mathsf{frank}0}$ and $\mathsf{m}_{\mathsf{frank}1}$, $\mathcal{B}$ queries $\mathsf{O}^{\mathsf{enc}}_{\mathsf{cca}}$ with both messages and uses the returned cipher text $\mathsf{c}_{\mathsf{frank}}$ to construct the franked message on behalf of the moderator and the users. When $\mathcal{A}$ calls $\mathsf{O}^{\mathsf{decrypt}}$ on ciphertext c, $\mathcal{B}$ behaves similarly and requests the decryption of the cipher-text from $\mathsf{O}^{\mathsf{decrypt}}_{\mathsf{cca}}$. If $\mathcal{A}$ succeeds in picking a choice bit, then $\mathcal{B}$ will pick the same choice bit in $\mathsf{ENC}^{\mathcal{A}}_{\mathsf{cca}}$ and will win with at least the same advantage.

Game$_3$: In $\mathsf{fwd}_{\mathsf{amf}}$ (called by $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{conf\text{-}fs,b}}$), we replace $\mathsf{m}_{\mathsf{frank}0}$ with $\mathsf{m}_{\mathsf{frank}1}$ in Forward on line 3. When a message is forwarded in Hecate, the commitment com in the original franked message is

replaced with a random commitment and the adversary cannot distinguish this change.

Game_4: In $\mathsf{fwd}_{\mathsf{amf}}$ (called by $\mathsf{O}^{\mathsf{fwd}}_{\mathsf{conf\text{-}fs},\mathsf{b}}$), we replace $\mathsf{m}_{\mathsf{frank}_0}'$ with $\mathsf{m}_{\mathsf{frank}_1}'$ in Forward. Since the adversary does not have access to the source $\mathsf{id}_{\mathtt{src}}$'s encryption secret keys, then they can only notice this change with advantage equal to breaking the CCA security of $\mathsf{send}_{\mathsf{eems}}$ and decrypting the contents of $\mathsf{c}_{\mathsf{frank}}$. The proof is similar to the on in Game_1.

The resulting $\mathsf{CONF\text{-}FS}_0^{\mathcal{A}}$ is identical to $\mathsf{CONF\text{-}FS}_1^{\mathcal{A}}$ and the adversary can only win the $\mathsf{CONF\text{-}FS}_{\mathsf{b}}^{\mathcal{A}}$ game with advantage equal to twice the advantage of breaking the CCA security of the underlying encryption scheme and the hiding property of the commitment scheme. □

5.4 Unforgeability and Accountability

$\mathsf{ACC}^{\mathcal{A}}$

1 : $s \leftarrow\$\, \mathcal{A}$
2 : $(\mathsf{pk}_{\mathtt{mod}}, \mathsf{sk}_{\mathtt{mod}}) \leftarrow\$\, \mathsf{KGen}(1^n)$
3 : $(\mathsf{pk}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}}) \leftarrow\$\, \mathsf{KGen}(s)$
4 : $\mathsf{report} \leftarrow\$\, \mathcal{A}^{\mathsf{O}^*}(\mathsf{sk}_{\mathtt{plat}})$
5 : $(\mathsf{id}, \mathsf{time}, \mathsf{m}) := \mathsf{Inspect}(\mathsf{m}_{\mathsf{frank}}, \mathsf{sk}_{\mathtt{mod}})$
6 : **if** $\mathsf{id} \notin \{\perp, \mathsf{corrupted}\} \wedge \mathsf{m} \notin \mathsf{M}$:
7 : **return** 1
8 : **return** 0

$\mathsf{O}^{\mathsf{send}}_{\mathsf{acc}}(\mathsf{m}, \mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}})$

1 : $\mathsf{c}_{\mathsf{frank}} := \mathsf{send}_{\mathsf{amf}}(\mathsf{m}, \mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}})$
2 : $\mathsf{m}_{\mathsf{frank}} := \mathsf{receive}_{\mathsf{amf}}(\mathsf{c}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}, \mathsf{sk}_{\mathtt{plat}})$
3 : $\mathsf{F} := \mathsf{F} \cup \mathsf{m}_{\mathsf{frank}}$
4 : **return** $\mathsf{m}_{\mathsf{frank}}$

$\mathsf{O}^{\mathsf{fwd}}_{\mathsf{acc}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}})$

1 : $\mathsf{c}_{\mathsf{frank}} := \mathsf{fwd}_{\mathsf{amf}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{fwd}}, \mathsf{id}_{\mathtt{rec}})$
2 : $\mathsf{m}_{\mathsf{frank}}{}' := \mathsf{receive}_{\mathsf{amf}}(\mathsf{c}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}, \mathsf{sk}_{\mathtt{plat}})$
3 : **return** $\mathsf{m}_{\mathsf{frank}}{}'$

$\mathsf{O}^{\mathsf{deliver}}_{\mathsf{acc}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}})$

1 : $\mathsf{m}_{\mathsf{frank}} := \mathsf{receive}_{\mathsf{amf}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}})$
2 : **return** $\mathsf{m}_{\mathsf{frank}}$

Figure 14: The security games for Accountability. Where O^* refers to $\mathsf{O}^{\mathsf{send}}$,$\mathsf{O}^{\mathsf{fwd}}$, $\mathsf{O}^{\mathsf{deliver}}$,$\mathsf{O}^{\mathsf{request}}$ and $\mathsf{O}^{\mathsf{corrupt}}$

The unforgeability and accountability properties describe a scheme's ability to *bind senders to well-formed messages while guaranteeing that no user can be accused of sending a message that they did not send.* These properties go hand in hand as well formed messages are necessarily bound to their original sender and cannot hence be attributed to any other source. Note that these properties should hold with respect to an honest moderator who handles source tracing reported messages. In other words, the adversary in the unforgeability and accountability game attempts to create a message and fool the moderator into believing that it came from a different user.

We model this property via the $\mathsf{ACC}^{\mathcal{A}}$ security game. The adversary starts by making polynomially many queries to $\mathsf{O}^{\mathsf{send}}$, $\mathsf{O}^{\mathsf{fwd}}$, $\mathsf{O}^{\mathsf{deliver}}$, $\mathsf{O}^{\mathsf{corrupt}}$, and $\mathsf{O}^{\mathsf{request}}$ that collectively allow $\mathcal{A}$ to build any message path of their choice and receive the resulting franked messages at each node within that path. Afterward, the adversary is tasked with producing a franked message that: (1) can be reported back to an existing user (line 5 in $\mathsf{ACC}^{\mathcal{A}}$) (2) traces back to an uncorrupted party (line 6), and (3) was not previously created during the challenge phase (line 6). In producing a message that can pass these predicates, the adversary can effectively produce new messages that can be traced back to other users. Note that who the message traces back to, beyond being an existing honest

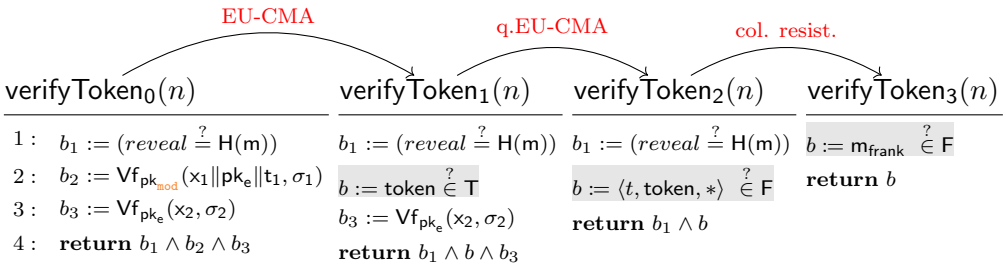

Figure 15: The hybrid steps modifying verifyToken in the Accountability game

user, is inconsequential: if the adversary cannot find anyone else to blame them for a message, regardless of who they are, then they cannot avoid accountability.

Definition 17 (Accountability)**.** *We define the advantage of the adversary in the ACC game for a scheme S as:*

$$\mathsf{Adv}_{S}^{\text{accountability}}(\mathcal{A}) = \Pr\left[\mathsf{ACC}^{\mathcal{A}} = 1\right].$$

We say that a scheme S holds users accountable if for all PPT adversaries $\mathcal{A}$::

$$\mathsf{Adv}_{S}^{\text{accountability}}(\mathcal{A}) \leq \mathsf{negl}(n).$$

Theorem 5.6. *Our scheme* Hecate *holds users accountable. Concretely, for any adversary $\mathcal{A}$ that makes at most q queries to its $\mathsf{O}^{\mathsf{send}}$ oracle, the advantage of $\mathcal{A}$ is at most:*

$$\mathsf{Adv}_{\mathsf{Hecate}}^{\text{accountability}}(\mathcal{A}) \leq (q+1) \cdot \mathsf{Adv}_{\mathcal{S}}^{\text{sig}_{\text{eu-cma}}}(\mathcal{A}) + \mathsf{Adv}_{\mathcal{H}}^{\text{hash}_{\text{coll}}}(\mathcal{A}).$$

Proof. We show how the $\mathsf{ACC}^{\mathcal{A}}$ game can only be won in Hecate with negligible probability. We note that the game's win condition requires a well-formed $\mathsf{m}_{\mathsf{frank}}$ that traces back to an existing id (line 6) and that contains a plain-text message that was not previously submitted and stored in $\mathsf{state}_{\mathsf{chal}}$(line 6). $\mathsf{ACC}^{\mathcal{A}}$ verifies the message integrity via $\mathsf{Verify}_{\mathtt{rec}}$ (on line 5), which returns an id in the case of a well formed message and $\perp$ otherwise. $\mathsf{Verify}_{\mathtt{rec}}$ is composed of verifyMsg (as seen in Figure 3 of Hecate's construction) which itself makes three separate calls to verifyExpiry, verifyCommit and verifyToken (line 1-2) based on which it determines the validity of the message. For a franked message $\mathsf{m}_{\mathsf{frank}}$ to be viable, all three functions must evaluate to true. verifyExpiry and verifyCommit are envelope commitment verification steps and are hence inconsequential for the accountability property. It's sufficient for the sake of this proof to show that verifyToken can never return true in Hecate. We handle the other two verification checks in backward security.

In order for verifyToken to return true, all three clauses lines (2-4) must also evaluate to true.

Game$_0$: This game is equivalent to verifyToken, since this is the only method in verifyMsg (Inspect's instantiation in Hecate) that is relevant for Hecate.

Game$_1$: We replace line 3 of verifyToken with a new boolean predicate $b \coloneqq \mathsf{token} \overset{?}{\in} \mathsf{T}$ that checks whether a token was generated by $\mathsf{O}^{\mathsf{request}}$, and hence by a legitimate call to TGen. Note that the only available way for the adversary to retrieve pre-processing tokens is through a call to $\mathsf{O}^{\mathsf{request}}$ with the id of a corrupted party and that every such token is stored in T on line 4. Since the adversary does not have access to the moderator's secret key, then $\mathcal{A}$ has a negligible advantage

in distinguishing this change by breaking the EU-CMA of the underlying signature scheme, forging the signature of the moderator and creating their own token without using $\mathsf{O}^{\mathsf{request}}$. We formally show this by assuming that there exists an adversary $\mathcal{A}$ that can distinguish between Games 0 and 1, and showing how to construct an adversary $\mathcal{B}$ that breaks EU-CMA. We primarily show how $\mathcal{B}$ constructs $\mathsf{O}^{\mathsf{request}}$, all other oracles outputs can be trivially generated by $\mathcal{B}$ in Hecate since they are ephemeral and only depend on the pre-processing token and additionally do not depend on the moderator's secret key. When $\mathcal{A}$ requests a pre-processing token, $\mathcal{B}$ locally construct x_1, t_1 and $(\mathsf{pk}_\mathsf{e}, \mathsf{sk}_\mathsf{e})$, then signs their concatenation by making a query to $\mathsf{Sig}^{\mathcal{A}}_{\mathsf{eu-cma}}$. $\mathcal{B}$ then sends the resulting token to $\mathcal{A}$ and waits until $\mathcal{A}$ submits an $\mathsf{m}_{\mathsf{frank}}$. $\mathcal{B}$ additionally checks franked messages delivered using $\mathsf{O}^{\mathsf{deliver}}_{\mathsf{acc}}$ that pass the verification step and the predicates on line 6 in the accountability game. If $\mathcal{A}$ does not fail, $\mathcal{B}$ can strip $\mathsf{m}_{\mathsf{frank}}$ of everything except the moderator's pre-processring signature σ_1 and submit that to the EU-CMA game. Since this $\mathcal{A}$ is required to submitted a new franked message originating from uncorrupted parties, then the moderator signature in $\mathsf{m}_{\mathsf{frank}}$ will not correspond to any of outputs of $\mathsf{Sig}^{\mathcal{A}}_{\mathsf{eu-cma}}$, and $\mathcal{B}$ will win the EU-CMA games when $\mathcal{A}$ does.

Game$_2$: Let $t := (\mathsf{x}_2, \sigma_2)$, i.e. the subset of the franked message constructed during the online stage. We replace and combine lines 3 and 4 with the boolean predicate $b \coloneqq (\langle t, \mathsf{token}, * \rangle \overset{?}{\in} \mathsf{F})$ that checks that the immutable tuple $\langle t, \mathsf{token} \rangle$ is a subset of some franked message in F. Recall that $\mathsf{O}^{\mathsf{send}}_{\mathsf{acc}}$ is the adversary's only way to instruct honest parties to construct and send messages. All such messages are saved in F on line 3. Since the adversary does not have access to honest parties' ephemeral keys, then $\mathcal{A}$ has negligible advantage in distinguishing between the original signature verification and the predicate we replaced it with by breaking the EU-CMA of the underlying signature scheme, forging the signature of the an honest party and creating their own franked message originating from that user without using $\mathsf{O}^{\mathsf{send}}_{\mathsf{acc}}$. We formally show this by assuming that there exists an adversary $\mathcal{A}$ that can distinguish between Games 1 and 2, and showing how to construct an adversary $\mathcal{B}$ that breaks EU-CMA. Let q be an upper bound on the number of queries $\mathcal{A}$ can make to $\mathsf{O}^{\mathsf{send}}_{\mathsf{acc}}$. We construct an $\mathcal{B}$ that has access to q different EU-CMA games (with their own $\mathsf{O}^{\mathsf{sign}}_{\mathsf{eu-cma}}$) for q ephemeral key pairs (where $q \in \mathsf{poly}(\lambda)$) and that will try to win at least one of these games. $\mathcal{B}$ starts by sampling and fixing the moderator's secret key and uses it to sign pre-processing tokens. When $\mathcal{A}$ queries $\mathsf{O}^{\mathsf{send}}_{\mathsf{acc}}$, $\mathcal{B}$ picks one of the unused q ephemeral public keys, constructs the pre-processing token for that public key by randomly sampling its secret ephemeral key, and generates the rest of the franked message by asking the $\mathsf{O}^{\mathsf{sign}}_{\mathsf{eu-cma}}$ associated with that public key to sign x_2. Note that $\mathcal{B}$ can trivially generate all other fields in the franked message since they have access to the moderator's secret key. Additionally, $\mathcal{B}$ can entirely act on behalf of the moderator when $\mathcal{A}$ calls $\mathsf{O}^{\mathsf{request}}$. $\mathcal{B}$ then waits until $\mathcal{A}$ submits an franked messages or delivers a well formed $\mathsf{m}_{\mathsf{frank}}$ using $\mathsf{O}^{\mathsf{deliver}}_{\mathsf{acc}}$ that pass predicates on line 6. $\mathcal{B}$ can win the EU-CMA game $\mathsf{Sig}^{\mathcal{A}}_{\mathsf{eu-cma}}$ by stripping that $\mathsf{m}_{\mathsf{frank}}$ of everything except its online signature σ_2 and its ephemeral public key pk_e that it submits to the corresponding $\mathsf{Sig}^{\mathcal{A}}_{\mathsf{eu-cma}}$ game. If $\mathcal{A}$ succeeds at distinguishing Games 1 and 2 using $\mathsf{m}_{\mathsf{frank}}$, then $\mathcal{B}$ will win $\mathsf{Sig}^{\mathcal{A}}_{\mathsf{eu-cma}}$ since $\mathsf{m}_{\mathsf{frank}}$ will need to be a new franked message originating from an uncorrupted party and could not have have been generated by $\mathsf{O}^{\mathsf{sign}}_{\mathsf{eu-cma}}$. Since each $\mathsf{Sig}^{\mathcal{A}}_{\mathsf{eu-cma}}$ game is defined by a separate pair of ephemeral keys generated i.i.d., then these games are independent of one another and by the union bound the advantage of the adversary is at most equal to $q \times \mathsf{Adv}^{\mathsf{sig}_{\mathsf{eu-cma}}}_{\mathcal{A}}(n)$.

Game$_3$: Finally, we can replace and combine all three lines in $\mathsf{verifyToken}$ with the boolean predicate $b \coloneqq \langle \mathsf{token}, t, m \rangle \overset{?}{\in} \mathsf{M}$, which is in essence equivalent to $b \coloneqq \mathsf{m}_{\mathsf{frank}} \overset{?}{\in} \mathsf{M}$. The adversary

can distinguish the change made with negligible advantage equal to the likelihood of finding a collision $m' \notin \mathsf{M}$ of the collision resistant hash function H, such that $\exists \mathsf{m}_{\mathsf{frank}}.m \in \mathsf{M}, \mathsf{H}(m) = \mathsf{H}(m')$. Let's assume that there exists an adversary $\mathcal{A}$ that can distinguish between Games 2 and 3. We construct an adversary $\mathcal{B}$ that will try to win the collision resistance hash game. When $\mathcal{A}$ queries $\mathsf{O}^{\mathsf{send}}_{\mathsf{acc}}$, $\mathcal{B}$ requests the hash of m from $\mathsf{O}^{\mathsf{hash}}$ of the collision resistance hash game, then constructs the rest of the franked message honestly. Note that all other oracles can be run by $\mathcal{B}$ since they act on behalf of the moderator and honest users. $\mathcal{B}$ then waits until $\mathcal{A}$ submits an franked messages or delivers a well formed $\mathsf{m}_{\mathsf{frank}}$ using $\mathsf{O}^{\mathsf{deliver}}_{\mathsf{acc}}$ that pass predicates on line 6. If $\mathcal{A}$ succeeds at distinguishing Games 2 and 3 using $\mathsf{m}_{\mathsf{frank}}$, then $\mathcal{B}$ will win the collision resistance hash game since the only way $\mathcal{A}$ can distinguish between both games is if finds a collision m' of $\mathsf{H}(\mathsf{m})$ that was not stored in M.

Now notice that there can be no franked message $\mathsf{m}_{\mathsf{frank}}$ that can satisfy the hybrid predicate $\mathsf{m}_{\mathsf{frank}} \stackrel{?}{\in} \mathsf{M}$ and the winning condition of the game $\mathsf{id} \notin \{\bot, \mathsf{corrupted}\} \wedge \mathsf{m}_{\mathsf{frank}}.\mathsf{m} \notin \mathsf{M}$, since we have shown that the adversary has to necessarily submit a stored franked message in M. The adversary cannot therefore construct a franked message $\mathsf{m}_{\mathsf{frank}}$ that can win this game. We can reduce the attacker's advantage in winning the accountability game to the sum of its advantage in breaking the collision-resistance or $(q+1)$ the EU-CMA games. □

5.5 Backward Security

$\mathsf{BAC}^{\mathcal{A}}_{\delta}$

1 : $(\mathsf{pk}_{\mathtt{mod}}, \mathsf{sk}_{\mathtt{mod}}) \leftarrow\$ \mathsf{KGen}(1^n)$
2 : $(\mathsf{pk}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}}) \leftarrow\$ \mathsf{KGen}(1^n)$
3 : $\mathcal{A}^{\mathsf{O}^*}$
4 : $\mathsf{corrupted} := \emptyset, \mathsf{F} := \emptyset$
5 : $\mathsf{m}_{\mathsf{chal}} \leftarrow\$ \mathcal{M}$
6 : $\mathsf{recovery}_{\mathsf{start},t} := \mathsf{global}_t$
7 : $\mathsf{global}_t := \mathsf{global}_t + \delta$
8 : $\mathsf{recovery}_{\mathsf{delay},t} := \mathsf{global}_t$
9 : $\mathsf{m}_{\mathsf{frank}} \leftarrow \mathcal{A}^{\mathsf{O}^*}(\mathsf{m}_{\mathsf{chal}})$
10 : **if** $\mathsf{check}_{\mathsf{report}}(\mathsf{m}_{\mathsf{frank}})$:
11 : **return** 1
12 : **return** 0

$\mathsf{check}_{\mathsf{report}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{m}_{\mathsf{chal}})$

1 : $(\mathsf{m}, \mathsf{report}) := \mathsf{Verify}_{\mathtt{rec}}(\mathsf{m}_{\mathsf{frank}})$
2 : $(\mathsf{id}_{\mathtt{src}}, \mathsf{time}, \mathsf{m}) := \mathsf{Inspect}(\mathsf{report})$
3 : **if** $\mathsf{id}_{\mathtt{src}} \stackrel{?}{\neq} \bot \wedge (\mathsf{id}_{\mathtt{src}}, *) \stackrel{?}{\notin} \mathsf{corrupted} \wedge \mathsf{m}_{\mathsf{frank}} \notin \mathsf{F}$:
4 : **if** $\mathsf{m} \stackrel{?}{=} \mathsf{m}_{\mathsf{chal}} \vee \mathsf{time} > \mathsf{recovery}_{\mathsf{delay},t}$:
5 : **return** 1
6 : **return** 0

$\mathsf{O}^{\mathsf{send}}_{\mathsf{bs}}(\mathsf{m}, \mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}})$

1 : $\mathsf{c}_{\mathsf{frank}} := \mathsf{send}_{\mathsf{amf}}(\mathsf{m}, \mathsf{id}_{\mathtt{src}}, \mathsf{id}_{\mathtt{rec}})$
2 : $\mathsf{m}_{\mathsf{frank}} := \mathsf{receive}_{\mathsf{amf}}(\mathsf{c}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}, \mathsf{sk}_{\mathtt{plat}})$
3 : $\mathsf{F} := \mathsf{F} \cup \mathsf{m}_{\mathsf{frank}}$
4 : **return** $\mathsf{m}_{\mathsf{frank}}$

$\mathsf{O}^{\mathsf{fwd}}_{\mathsf{bs}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{fwd}}, \mathsf{id}_{\mathtt{rec}})$

1 : $\mathsf{c}_{\mathsf{frank}} := \mathsf{fwd}_{\mathsf{amf}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{fwd}}, \mathsf{id}_{\mathtt{rec}})$
2 : $\mathsf{m}_{\mathsf{frank}}{}' := \mathsf{receive}_{\mathsf{amf}}(\mathsf{c}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}, \mathsf{sk}_{\mathtt{plat}})$
3 : **return** $\mathsf{m}_{\mathsf{frank}}{}'$

$\mathsf{O}^{\mathsf{deliver}}_{\mathsf{bs}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}})$

1 : $\mathsf{m}_{\mathsf{frank}} := \mathsf{receive}_{\mathsf{amf}}(\mathsf{m}_{\mathsf{frank}}, \mathsf{id}_{\mathtt{rec}}, \mathsf{time}_{\mathtt{plat}}, \mathsf{sk}_{\mathtt{plat}})$
2 : **return** $\mathsf{m}_{\mathsf{frank}}$

Figure 16: The security games for Backward Secrecy, where O^* denotes to $\mathsf{O}^{\mathsf{send}}$, $\mathsf{O}^{\mathsf{fwd}}$, $\mathsf{O}^{\mathsf{deliver}}$, $\mathsf{O}^{\mathsf{corrupt}}$ and $\mathsf{O}^{\mathsf{request}}$.

Backward Security requires that an adversary that controlled the state and keys of a device pre-comprise should not be able to benefit from them after the device recovers. In this work we define post-compromise recovery as the adversary's inability to: (1) craft new messages from a recovered user, (2) claim that during-compromise messages were sent out (not forwarded) after the compromise period. To this goal, we design the game $\mathsf{BAC}_{\delta}^{\mathcal{A}}$. We provide the adversary with oracle access to $\mathsf{O}^{\mathsf{corrupt}}$, $\mathsf{O}^{\mathsf{request}}$, $\mathsf{O}^{\mathsf{send}}$, $\mathsf{O}^{\mathsf{fwd}}$, $\mathsf{O}^{\mathsf{deliver}}$ that allow them to perform any possible interaction between users.

Similarly to prior games, the adversary can corrupt users using $\mathsf{O}^{\mathsf{corrupt}}$ and request their pre-processing tokens via $\mathsf{O}^{\mathsf{request}}$. However contrary to prior games, handling time in $\mathsf{BAC}_{\delta}^{\mathcal{A}}$ is necessary since backward security is a property of the corruption recovery period. In $\mathsf{BAC}_{\delta}^{\mathcal{A}}$, corruption is not indefinite and we define it with respect to the global time variable $\mathsf{global}_{\mathsf{t}}$. When $\mathcal{A}$ calls $\mathsf{O}^{\mathsf{corrupt}}$, the oracle stores the id of the corrupted user along with the time of corruption at $\mathsf{global}_{\mathsf{t}}$ in the set $\mathsf{corrupted}$. When $\mathcal{A}$ calls $\mathsf{O}^{\mathsf{request}}$ to request pre-processing tokens for corrupted users, the oracle first checks that the requested user is in $\mathsf{corrupted}$ at the current global time and increments the global time if that check succeeds. In other words, the global time is incremented each time a pre-processing token is returned to $\mathcal{A}$, effectively requiring them to re-corrupt users at the new global time. By defining time in this manner, we capture how an adversary can no longer impersonate recovered users and request tokens on their behalf.

The adversary can additionally send and forward messages between any two users regardless of their corruption status via $\mathsf{O}^{\mathsf{send}}$, $\mathsf{O}^{\mathsf{fwd}}$ and $\mathsf{O}^{\mathsf{deliver}}$. $\mathsf{O}^{\mathsf{send}}$ in particular grants the adversary the power to request that an honest user creates a new franked message that the oracle stores in the set F. Otherwise, with access to $\mathsf{O}^{\mathsf{request}}$, the adversary can create franked messages on their own before sending/forwarding them along. $\mathsf{O}^{\mathsf{fwd}}$ allows them to forward a franked message between honest users, and $\mathsf{O}^{\mathsf{deliver}}$ allows them to send/forward messages from a corrupted user to an honest one. With access to all these oracles, we model an all powerful adversary that can fully control and build a forwarding tree of their choice.

The $\mathsf{BAC}_{\delta}^{\mathcal{A}}$ games proceeds in three phases.

The first phase marks the pre-recovery phase where the adversary is given access to all aforementioned oracles and can interact with users as they see fit (line 3).

When the adversary is done, the second phase begins and the game master sets up the post-recovery period by: (1) sampling a challenge $\mathsf{m}_{\mathsf{chal}}$ uniformly at random from the message space excluding plain-text messages sent in the first phase (line 5), (2) saving the time at which the first phase ended in $\mathsf{recovery}_{\mathsf{start,t}}$ (line 6), (3) and advancing the current time by a grace period delta to mark the delayed beginning of the recovery period $\mathsf{recovery}_{\mathsf{delay,t}}$ (line 8).

The final phase marks the post-recovery period when the adversary is given $\mathsf{m}_{\mathsf{chal}}$ and asked to provide a franked message $\mathsf{m}_{\mathsf{frank}}$ that is well formed, traces back to an honest user and either: (1) contains the plain-text challenge or (2) is timestamped after the the delayed recovery period, after they are done interacting with the provided oracles. The former two possibilities provided to the adversary encompass this works notion of backward security: the adversary should neither (1) be able to produce a message that they did not think of during compromise, (2) nor should they circulate a message that was not timestamped during the time of compromise. The later point essentially captures how abuse reporting systems may deem messages that were sent during known data breach periods as unaccountable and possibly invalid. This is especially important in an era where the number of cyber-attack campaigns is on a rise, and where the adversary may attempt to create many messages during the compromise period and delay reporting them or further circulating

them until an opportune time after the recovery period begins. This is especially critical in the case where a public official's device is hacked and when the time at which their leaked messages were sent can influence public opinion.

Note that the game master resets both the $\mathsf{corrupted}$ and F sets on line 4 since they are not useful in the first phase when the recovery period has not yet begun. Both sets will allow the game master to evaluate the response of $\mathcal{A}$ with respect to the winning condition. Since backward security is a property of corruption recovery, the winning condition is itself a function of that period. In this essence, we define recovery after some fixed δ has passed on line 8, to model how recovery in practice is not instantaneous and may require a grace period.

Definition 18. *We define the advantage of the adversary in the* $\mathsf{BAC}_{\delta}^{\mathcal{A}}$ *game for a scheme* S *as:*

$$\mathsf{Adv}_{S}^{\mathrm{bac}}(\mathcal{A}) = \Pr\left[\mathsf{BAC}_{\delta}^{\mathcal{A}} = 1\right].$$

We say that a scheme S *is post-compromise secure if for all PPT adversaries* $\mathcal{A}$*:*

$$\mathsf{Adv}_{S}^{\mathrm{bac}}(\mathcal{A}) \leq \mathsf{negl}(n).$$

Theorem 5.7. *Our scheme* Hecate *is backward (aka post-compromise) secure. The advantage of the adversary* $\mathcal{A}$ *is:*

$$\mathsf{Adv}_{\mathsf{Hecate}}^{\mathrm{bac}}(\mathcal{A}) \leq \mathsf{Adv}^{\mathrm{accountability}}(\mathcal{A}) + \mathsf{Adv}^{\mathrm{sig}_{\text{eu-cma}}}(\mathcal{A}) + \mathsf{Adv}^{\mathrm{binding}_{\mathrm{com}}}(\mathcal{A}).$$

Proof. The winning condition of the game requires the chosen $\mathsf{m}_{\mathsf{frank}}$ to trace back to an honest user. $\mathsf{m}_{\mathsf{frank}}$ may then either contain the challenge $\mathsf{m}_{\mathsf{chal}}$ chosen by the game master, or be timestamped after the recovery period delay $\mathsf{recovery}_{\mathsf{delay,t}}$ on line 8. The game requires $\mathsf{m}_{\mathsf{frank}}$ to be traceable by the moderator via $\mathsf{Inspect}$ in $\mathsf{check}_{\mathsf{report}}$. In Hecate, $\mathsf{Inspect}$ is a conjunction of three separate verification steps: $\mathsf{verifyExpiry}$ that checks the expiry of the different time components of the franked message, $\mathsf{verifyCommit}$ that checks the envelope commitment, and $\mathsf{verifyToken}$ which checks all other parts of the franked message and which we discuss in depth in the accountability game.

In this proof we distinguish between the payload of the franked message that we denote by $\mathsf{m}_{\mathsf{frank}}.\mathsf{payload}$ and its envelope $\mathsf{m}_{\mathsf{frank}}.\mathsf{envelope}$. In what follows, $\mathsf{m}_{\mathsf{frank}}$ is equivalent to ($\mathsf{m}_{\mathsf{frank}}.\mathsf{payload}$, $\mathsf{m}_{\mathsf{frank}}.\mathsf{envelope}$), $\mathsf{m}_{\mathsf{frank}}.\mathsf{envelope}$ is equivalent to (com, σ_2, t_2). We use the regular expression operator $*$ to denote that a field can take any value and is not specified by a specific game reduction. Each game hybrid will allow us to specify different pieces of the eventual franked message.

Game_1: We replace $\mathsf{verifyToken}$ with checking that $\mathsf{m}_{\mathsf{frank}}$'s preprocessing token is either $(\mathsf{token}, *) \overset{?}{\in} \mathsf{F}$ or $\mathsf{token} \overset{?}{\in} \mathsf{T}$ (where token is a shorthand for $\mathsf{m}_{\mathsf{frank}}.\mathsf{payload}.\mathsf{token}$). The adversary can distinguish this change with negligible probability equal to the probability of winning the EU-CMA game since the adversary does not have access to the moderator's secret key and has therefore a negligible chance in constructing well formed tokens for a user of their choice locally. Recall that T is the set of corrupted users tokens constructed by $\mathsf{O}^{\mathsf{request}}$ and that $(\mathsf{token}, *) \in \mathsf{F}$ refers to the set of honest user tokens constructed by $\mathsf{O}_{\mathsf{bs}}^{\mathsf{send}}$. We refer the reader to Game 1 in the accountability game for more details.

Game_2: We parametrize $\mathsf{expiry} := \delta$ and replace $\mathsf{verifyExpiry}$ with $(\mathsf{token}, *) \overset{?}{\in} \mathsf{F}$, i.e. we drop the $\mathsf{token} \overset{?}{\in} \mathsf{T}$ from the disjunction in Game_1. $\mathsf{verifyExpiry}$ requires that $|\mathsf{t}_2 - \mathsf{t}_1| < \mathsf{expiry}$, i.e. that $\mathsf{m}_{\mathsf{frank}}$'s moderator and platform time stamps are within a set expiry time from one another (line

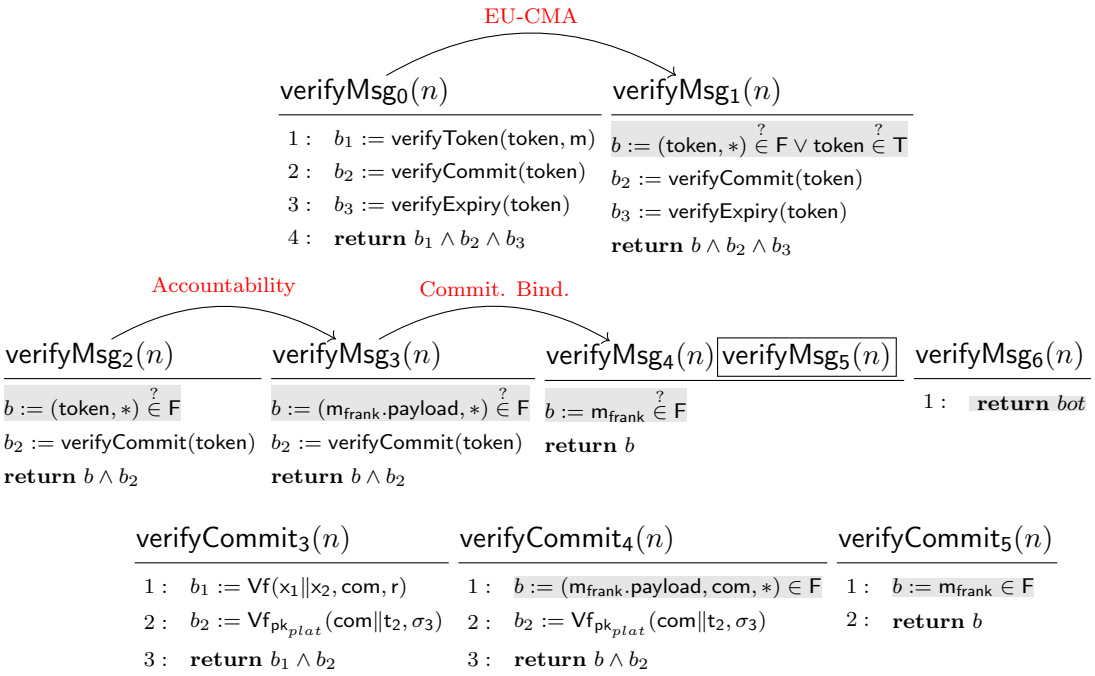

Figure 17: The hybrid steps modifying verifyMsg and verifyCommit in the Backward Security game. We use the shorthand token to refer to $\mathsf{m_{frank}.payload.token}$ and use the regular expression wildcard operator "$*$" throughout.

1). Note that the winning condition of the game requires $\mathsf{id_{src}} \notin \mathsf{corrupted}$. In other words the only possible way for token to have been created by a call $\mathsf{O^{corrupt}}$ and $\mathsf{O^{request}}$ must happen prior to the beginning of the recovery time, i.e. $\mathsf{token.t_1} < \mathsf{recovery_{start,t}}$. Recall that users can only be corrupt for one epoch at the current $\mathsf{global_t}$ before being considered as honest and requiring the adversary to call $\mathsf{O^{corrupt}}$ and $\mathsf{O^{request}}$ if they wish to corrupt them again. This definition of corruption is enforced by $\mathsf{O^{request}}$: if the id passed is not corrupted at the current $\mathsf{global_t}$, the oracle returns $\perp$, otherwise it returns the appropriate token and increments $\mathsf{global_t}$. Additionally, since the game master increments the global time by δ after $\mathsf{recovery_{start,t}}$ on line 8 in $\mathsf{BAC}^{\mathcal{A}}_{\delta}$ then $\mathsf{envelope.t_2} > \mathsf{recovery_{start,t}} + \delta$ (where envelope is a shorthand for $\mathsf{m_{frank}.envelope}$). In other words, $\mathsf{token} \notin \mathsf{T}$ since otherwise $|\mathsf{envelope.t_2} - \mathsf{t_1}| > \mathsf{expiry}$ and verifyExpiry would return $\perp$.

Game$_3$: We replace verifyToken with checking that $(\mathsf{m_{frank}.payload}, *) \overset{?}{\in} \mathsf{F}$, where payload refers to all elements of the franked message that are not on the envelope of the message. The adversary can distinguish this change with negligible probability equal to the advantage of the accountability game since they do not have access to the secret ephemeral keys of users for token constructed with $\mathsf{O^{send}}$. We refer the reader to that proof for more details and not that both Games 1 and 3 are jointly upper bounded by the advantage of the accountability game.

Game$_4$: In verifyCommit, we replace line 1 with $b_1 := (\mathsf{m_{frank}.payload}, \mathsf{com}, *) \in \mathsf{F}$. We show that the adversary can distinguish between Games 1 and 2 with negligible advantage equal to the probability of breaking the binding property of the commitment scheme. Let's assume that there exists an adversary $\mathcal{A}$ that can distinguish both games, we construct an adversary $\mathcal{B}$ that can break the commitment game. Whenever $\mathcal{A}$ calls either $\mathsf{O^{send}_{bs}}$, $\mathsf{O^{fwd}_{bs}}$ or $\mathsf{O^{deliver}_{bs}}$, $\mathcal{B}$ constructs/modifies the

franked message honestly on behalf of the moderator, the platform and the users and returns the result back to $\mathcal{A}$. Note that $\mathcal{B}$ stores every single such message. When $\mathcal{A}$ returns a challenge $\mathsf{m}_{\mathsf{frank}}$ back to $\mathcal{B}$ or successfully delivers an $\mathsf{m}_{\mathsf{frank}}$ to an honest user that was not previously logged in F using $\mathsf{O}_{\mathsf{bs}}^{\mathsf{deliver}}$, $\mathcal{B}$ strips the resulting $\mathsf{m}_{\mathsf{frank}}$ of everything except the commitment and its corresponding decommitment, and submits these values along with the original decommitment stored in F to the commitment game. If $\mathcal{A}$ succeeds then they will have necessarily submitted a decommitment that is different than the original one stored in F and $\mathcal{B}$ will win its corresponding game.

Game_5: In $\mathsf{verifyCommit}$, we replace line 1 with $b_2 := (\mathsf{m}_{\mathsf{frank}}.\mathsf{payload}, \mathsf{com}, \mathsf{t}_2, \sigma_2) \in \mathsf{F}$, which effectively implies replacing $\mathsf{verifyCommit}$ with $(\mathsf{m}_{\mathsf{frank}}.\mathsf{payload}, \mathsf{m}_{\mathsf{frank}}.\mathsf{envelope}) \in \mathsf{F}$ (i.e. $\mathsf{m}_{\mathsf{frank}} \in \mathsf{F}$). We can show similarly to Game 1 in the accountability game that, without the platform's secret key, the adversary can distinguish between Games 2 and 3 with negligible advantage equal to the probability of breaking EU-CMA security property of the underlying signature scheme.

Game_6: We replace $\mathsf{verifyMsg}$ with $\perp$ since we have shown that $\mathsf{m}_{\mathsf{frank}} \in \mathsf{F}$ and the game's winning condition requires otherwise.

We can therefore conclude that the adversary has negligible advantage in winning the $\mathsf{BAC}_{\delta}^{\mathcal{A}}$ game in Hecate equal to:

$$\mathsf{Adv}^{\mathrm{bac}}_{\mathsf{Hecate}}(\mathcal{A}) \leq \mathsf{Adv}^{\mathrm{accountability}}_{\mathsf{Hecate}}(\mathcal{A}) + \mathsf{Adv}^{\mathrm{sig}_{\text{eu-cma}}}(\mathcal{A}) + \mathsf{Adv}^{\mathrm{binding}_{\mathrm{com}}}(\mathcal{A}). \qquad \square$$

6 Conclusion and Future Work

In this work, we constructed the first protocol that combines aspects of asynchronous message franking and source tracing. Our construction also requires less concrete runtime than the prior work on AMF. Along the way, we generalized the formal definitions of AMF to the interactive setting with preprocessing.

Looking ahead, we believe there exist at least five possible avenues of future research in the space of privacy-respecting content moderation.

Content Censorship. Content moderation systems can in general be adversely used for censorship purposes. Questions surrounding what constitutes misinformation or a "bad" message fall outside the scope of this work and fall in the realm of policy making and social media regulation. We believe however that it may be interesting to federate the role of the moderator in: (1) constructing what constitutes bad messages, (2) verifying reports, (3) taking actions with respect to flagged contents and users.

Group Messaging. In this work we focus exclusively on point-to-point two party communication within encrypted messengers. Our definitions and the Hecate construction can be ported in a straightforward manner to to Signal's group messaging protocol, in which broadcasts to a group of size N are implemented via N individual point-to-point messages, after a server-assisted consensus protocol to determine the group [21]. There also exist several recent works and an IETF standardization effort that explores sub-linear ends-to-ends encrypted group chats [5, 6, 12, 23, 61]; these proposals do not fit within our model from §2. We leave questions surrounding different group messaging constructions to future work. We do not yet know if there are any efficiency gains and added benefits to federating content moderation among different members of a group.

Super Spreaders. A recent line of work [64,73] on misinformation spread in social media suggests a distinction between honest users who forward misinformation and malicious actors that act as super spreaders of misinformation. Since honest user may not be able to distinguish between real

and fake information that they receive, they can mistakenly forward or send misinformation content without ever realizing it. Super spreaders on the other are adversarially creating or spreading bad content. We propose in future work to look into aggregate malicious behavior in order to potentially discern between each type of user.

Partial opening. In this work, we restrict our attention in this work to a scenario where the receiver must report all (or none) of a message to the moderator. It may be possible to achieve partial opening to the moderator using generic zero knowledge proofs or the specific techniques of Leontiadis and Vaudenay [50].

Stronger Notions of Backward Security. Backward security makes no guarantees with respect to anything created during the time of compromise. In the context of content moderation, this implies that the adversary can blame users for old compromised messages. We encourage future research into ways to limit the damage of adversarial moderation reports or allow honest parties to correct the record post-recovery.

Supporting Multiple Devices. Throughout our presentation, Hecate considers laptop and phone sessions for the same user independently; each device receives separate tokens for the same user identity. We leave for future work the question of supporting multiple devices per user.

Ensuring System Security. Finally, we emphasize that our investigation into abuse reporting has been primarily through a cryptographic lens, and as a result does not capture all aspects of security. For instance, many of our crypto definitions assume that clients already have sufficient preprocessing tokens in hand; careful attention is required when implementing Hecate to ensure that an adversary cannot obtain a side channel by influencing when preprocessing is run. We encourage cryptographers, systems security researchers, usability experts, and domain specialists to investigate whether and how to integrate Hecate (or any abuse reporting mechanism) into an end-to-end encrypted messaging system in a matter that promotes online trust, safety, and security.

References

[1] Michel Abdalla, Mihir Bellare, and Gregory Neven. Robust encryption. In Daniele Micciancio, editor, *TCC 2010*, volume 5978 of *LNCS*, pages 480–497. Springer, Heidelberg, February 2010.

[2] Surabhi Agarwal. India proposes alpha-numeric hash to track WhatsApp chat. https://economictimes.indiatimes.com/tech/technology/govt-proposes-alpha-numeric-hash-to-track-whatsapp-chat/articleshow/81638939.cms, 2021.

[3] Nikolaos Alexopoulos, Aggelos Kiayias, Riivo Talviste, and Thomas Zacharias. MCMix: Anonymous messaging via secure multiparty computation. In Engin Kirda and Thomas Ristenpart, editors, *USENIX Security 2017*, pages 1217–1234. USENIX Association, August 2017.

[4] Joël Alwen, Sandro Coretti, and Yevgeniy Dodis. The double ratchet: Security notions, proofs, and modularization for the Signal protocol. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 129–158. Springer, Heidelberg, May 2019.

[5] Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis. Security analysis and improvements for the IETF MLS standard for group messaging. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 248–277. Springer, Heidelberg, August 2020.

[6] Joël Alwen, Sandro Coretti, Daniel Jost, and Marta Mularczyk. Continuous group key agreement with active security. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part II*, volume 12551 of *LNCS*, pages 261–290. Springer, Heidelberg, November 2020.

[7] Donald Beaver. Efficient multiparty protocols using circuit randomization. In Joan Feigenbaum, editor, *CRYPTO'91*, volume 576 of *LNCS*, pages 420–432. Springer, Heidelberg, August 1992.

[8] Amos Beimel, Yuval Ishai, and Tal Malkin. Reducing the servers' computation in private information retrieval: PIR with preprocessing. *Journal of Cryptology*, 17(2):125–151, March 2004.

[9] Mihir Bellare, Asha Camper Singh, Joseph Jaeger, Maya Nyayapati, and Igors Stepanovs. Ratcheted encryption and key exchange: The security of messaging. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part III*, volume 10403 of *LNCS*, pages 619–650. Springer, Heidelberg, August 2017.

[10] Luca Belli. Whatsapp skewed brazilian election, proving social media's danger to democracy. https://theconversation.com/whatsapp-skewed-brazilian-election-proving-social-medias-danger-to-democracy-106476, 2018.

[11] Abhishek Bhowmick, Dan Boneh, Steve Myers, Kunal Talwar, and Karl Tarbe. The apple psi system. https://www.apple.com/child-safety/pdf/Apple_PSI_System_Security_Protocol_and_Analysis.pdf, 2021.

[12] Alexander Bienstock, Yevgeniy Dodis, and Paul Rösler. On the price of concurrency in group ratcheting protocols. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part II*, volume 12551 of *LNCS*, pages 198–228. Springer, Heidelberg, November 2020.

[13] Alexander Bienstock, Jaiden Fairoze, Sanjam Garg, Pratyay Mukherjee, and Srinivasan Raghuraman. What is the exact security of the Signal protocol? https://cs.nyu.edu/~afb383/publication/uc_signal/uc_signal.pdf, 2021.

[14] Dan Boneh and Victor Shoup. A graduate course in applied cryptography. https://toc.cryptobook.us/book.pdf, 2020.

[15] Nikita Borisov, Ian Goldberg, and Eric A. Brewer. Off-the-record communication, or, why not to use PGP. In *WPES*, pages 77–84. ACM, 2004.

[16] Colin Boyd, Anish Mathuria, and Douglas Stebila. *Protocols for Authentication and Key Establishment, Second Edition*. Information Security and Cryptography. Springer, 2020.

[17] Brazilian fake news draft bill no. 2.630, of 2020. https://docs.google.com/document/d/1MHMDHsVJBi45PI1R5lAyoLmZvZk8eULHisYFqGy9X2s/edit, 2020.

[18] Sébastien Campion, Julien Devigne, Céline Duguey, and Pierre-Alain Fouque. Multi-device for signal. In Mauro Conti, Jianying Zhou, Emiliano Casalicchio, and Angelo Spognardi, editors, *ACNS 20, Part II*, volume 12147 of *LNCS*, pages 167–187. Springer, Heidelberg, October 2020.

[19] Ran Canetti and Hugo Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 453–474. Springer, Heidelberg, May 2001.

[20] Ran Canetti, Daniel Shahaf, and Margarita Vald. Universally composable authentication and key-exchange with global PKI. In Chen-Mou Cheng, Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang, editors, *PKC 2016, Part II*, volume 9615 of *LNCS*, pages 265–296. Springer, Heidelberg, March 2016.

[21] Melissa Chase, Trevor Perrin, and Greg Zaverucha. The signal private group system and anonymous credentials supporting efficient verifiable encryption. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 20*, pages 1445–1459. ACM Press, November 2020.

[22] Katriel Cohn-Gordon, Cas Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. A formal security analysis of the signal messaging protocol. *Journal of Cryptology*, 33(4):1914–1983, 2020.

[23] Katriel Cohn-Gordon, Cas Cremers, Luke Garratt, Jon Millican, and Kevin Milner. On ends-to-ends encryption: Asynchronous group messaging with strong security guarantees. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 1802–1819. ACM Press, October 2018.

[24] Katriel Cohn-Gordon, Cas J. F. Cremers, and Luke Garratt. On post-compromise security. In *CSF*, pages 164–178. IEEE Computer Society, 2016.

[25] Henry Corrigan-Gibbs, Dan Boneh, and David Mazières. Riposte: An anonymous messaging system handling millions of users. In *2015 IEEE Symposium on Security and Privacy*, pages 321–338. IEEE Computer Society Press, May 2015.

[26] Cas Cremers, Jaiden Fairoze, Benjamin Kiesl, and Aurora Naska. Clone detection in secure messaging: Improving post-compromise security in practice. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 20*, pages 1481–1495. ACM Press, November 2020.

[27] Roger Dingledine, Nick Mathewson, and Paul F. Syverson. Tor: The second-generation onion router. In Matt Blaze, editor, *USENIX Security 2004*, pages 303–320. USENIX Association, August 2004.

[28] Yevgeniy Dodis, Paul Grubbs, Thomas Ristenpart, and Joanne Woodage. Fast message franking: From invisible salamanders to encryptment. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 155–186. Springer, Heidelberg, August 2018.

[29] Yevgeniy Dodis, Jonathan Katz, Adam Smith, and Shabsi Walfish. Composability and on-line deniability of authentication. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 146–162. Springer, Heidelberg, March 2009.

[30] Elizabeth Dwoskin and Annie Gowen. On WhatsApp, fake news is fast – and can be fatal. https://www.washingtonpost.com/business/economy/on-whatsapp-fake-news-is-fast--and-can-be-fatal/2018/07/23/a2dd7112-8ebf-11e8-bcd5-9d911c784c38_story.html, July 2018. Accessed: 09-28-2020.

[31] Facebook. Messenger secret conversations: Technical whitepaper (version 2.0). https://about.fb.com/wp-content/uploads/2016/07/messenger-secret-conversations-technical-whitepaper.pdf, 2017.

[32] Pooya Farshim, Benoît Libert, Kenneth G. Paterson, and Elizabeth A. Quaglia. Robust encryption, revisited. In Kaoru Kurosawa and Goichiro Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 352–368. Springer, Heidelberg, February / March 2013.

[33] Pooya Farshim, Claudio Orlandi, and Răzvan Roşie. Security of symmetric primitives under incorrect usage of keys. *IACR Trans. Symm. Cryptol.*, 2017(1):449–473, 2017.

[34] Marc Fischlin. Round-optimal composable blind signatures in the common reference string model. In Cynthia Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 60–77. Springer, Heidelberg, August 2006.

[35] Marc Fischlin and Dominique Schröder. On the impossibility of three-move blind signature schemes. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 197–215. Springer, Heidelberg, May / June 2010.

[36] Georg Fuchsbauer, Christian Hanser, Chethan Kamath, and Daniel Slamanig. Practical round-optimal blind signatures in the standard model from weaker assumptions. In Vassilis Zikas and Roberto De Prisco, editors, *SCN 16*, volume 9841 of *LNCS*, pages 391–408. Springer, Heidelberg, August / September 2016.

[37] Georg Fuchsbauer, Christian Hanser, and Daniel Slamanig. Practical round-optimal blind signatures in the standard model. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 233–253. Springer, Heidelberg, August 2015.

[38] Paul Grubbs, Jiahui Lu, and Thomas Ristenpart. Message franking via committing authenticated encryption. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part III*, volume 10403 of *LNCS*, pages 66–97. Springer, Heidelberg, August 2017.

[39] Christoph G. Günther. An identity-based key-exchange protocol. In Jean-Jacques Quisquater and Joos Vandewalle, editors, *EUROCRYPT'89*, volume 434 of *LNCS*, pages 29–37. Springer, Heidelberg, April 1990.

[40] Joseph Jaeger and Igors Stepanovs. Optimal channel security against fine-grained state compromise: The safety of messaging. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 33–62. Springer, Heidelberg, August 2018.

[41] Jakob Jakobsen and Claudio Orlandi. On the CCA (in)security of mtproto. In *SPSM@CCS*, pages 113–116. ACM, 2016.

[42] Markus Jakobsson, Kazue Sako, and Russell Impagliazzo. Designated verifier proofs and their applications. In Ueli M. Maurer, editor, *EUROCRYPT'96*, volume 1070 of *LNCS*, pages 143–154. Springer, Heidelberg, May 1996.

[43] Daniel Jost, Ueli Maurer, and Marta Mularczyk. Efficient ratcheting: Almost-optimal guarantees for secure messaging. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 159–188. Springer, Heidelberg, May 2019.

[44] Daniel Jost, Ueli Maurer, and Marta Mularczyk. A unified and composable take on ratcheting. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part II*, volume 11892 of *LNCS*, pages 180–210. Springer, Heidelberg, December 2019.

[45] Seny Kamara, Mallory Knodel, Emma Llansó, Greg Nojeim, Lucy Qin, Dhanaraj Thakur, and Caitlin Vogus. Outside looking in: Approaches to content moderation in end-to-end encrypted systems. https://cdt.org/wp-content/uploads/2021/08/CDT-Outside-Looking-In-Approaches-to-Content-Moderation-in-End-to-End-Encrypted-Systems.pdf, 2021.

[46] Shuichi Katsumata, Ryo Nishimaki, Shota Yamada, and Takashi Yamakawa. Round-optimal blind signatures in the plain model from classical and quantum standard assumptions. In *EUROCRYPT*, Lecture Notes in Computer Science, 2021.

[47] Jonathan Katz and Yehuda Lindell. *Introduction to modern cryptography*. CRC press, 2020.

[48] Anunay Kulshrestha and Jonathan Mayer. Identifying harmful media in end-to-end encrypted communication: Efficient private membership computation. In *USENIX Security Symposium*. USENIX Association, 2021.

[49] Brian A. LaMacchia, Kristin Lauter, and Anton Mityagin. Stronger security of authenticated key exchange. In Willy Susilo, Joseph K. Liu, and Yi Mu, editors, *ProvSec 2007*, volume 4784 of *LNCS*, pages 1–16. Springer, Heidelberg, November 2007.

[50] Iraklis Leontiadis and Serge Vaudenay. Private message franking with after opening privacy. Cryptology ePrint Archive, Report 2018/938, 2018. https://eprint.iacr.org/2018/938.

[51] Linsheng Liu, Daniel S. Roche, Austin Theriault, and Arkady Yerukhimovich. Fighting fake news in encrypted messaging with the fuzzy anonymous complaint tally system (FACTS). *CoRR*, abs/2109.04559, 2021.

[52] Ian Martiny, Gabriel Kaptchuk, Adam Aviv, Dan Roche, and Eric Wustrow. Improving Signal's sealed sender. In *NDSS*. The Internet Society, 2021.

[53] Marino Miculan and Nicola Vitacolonna. Automated symbolic verification of telegram's mtproto 2.0. *CoRR*, abs/2012.03141, 2020.

[54] Oversight Board. Ensuring respect for free expression, through independent judgment. https://oversightboard.com/, 2021.

[55] Charlotte Peale, Saba Eskandarian, and Dan Boneh. Secure source-tracking for encrypted messaging. In *CCS*. ACM, 2021.

[56] Riana Pfefferkorn. Content-oblivious trust and safety techniques: Results from a survey of online service providers. https://ssrn.com/abstract=3920031, 2021.

[57] Bertram Poettering and Paul Rösler. Towards bidirectional ratcheted key exchange. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part I*, volume 10991 of *LNCS*, pages 3–32. Springer, Heidelberg, August 2018.

[58] Newley Purnell and Jeff Horowitz. WhatsApp says it filed suit in India to prevent tracing of encrypted messages. https://www.wsj.com/articles/whatsapp-says-it-filed-suit-in-india-to-prevent-tracing-of-encrypted-messages-11622000307, 2021.

[59] Paul Rösler, Christian Mainka, and Jörg Schwenk. More is less: On the end-to-end security of group chats in signal, whatsapp, and threema. In *EuroS&P*, pages 415–429. IEEE, 2018.

[60] Shahrokh Saeednia, Steve Kremer, and Olivier Markowitch. An efficient strong designated verifier signature scheme. In Jong In Lim and Dong Hoon Lee, editors, *ICISC 03*, volume 2971 of *LNCS*, pages 40–54. Springer, Heidelberg, November 2004.

[61] Michael Schliep and Nicholas Hopper. End-to-end secure mobile group messaging with conversation integrity and deniability. In *WPES@CCS*, pages 55–73. ACM, 2019.

[62] Signal. Technology preview: Sealed sender for signal. https://signal.org/blog/sealed-sender/, 2018.

[63] Signal. Technical information. https://signal.org/docs/, 2021.

[64] Kate Starbird. Online rumors, misinformation and disinformation: The perfect storm of covid-19 and election2020. In *Enigma 2021*. USENIX Association, February 2021.

[65] Li Q. Tay, Mark J. Hurlstone, Tim Kurz, and Ullrich K. H. Ecker. A comparison of prebunking and debunking interventions for implied versus explicit misinformation. PsyArXiv, 2021.

[66] Kurt Thomas, Devdatta Akhawe, Michael Bailey, Dan Boneh, Elie Bursztein, Sunny Consolvo, Nicola Dell, Zakir Durumeric, Patrick Gage Kelley, Deepak Kumar, Damon McCoy, Sarah Meiklejohn, Thomas Ristenpart, and Gianluca Stringhini. Sok: Hate, harassment, and the changing landscape of online abuse. https://research.google/pubs/pub49786.pdf, 2021.

[67] Nirvan Tyagi, Paul Grubbs, Julia Len, Ian Miers, and Thomas Ristenpart. Asymmetric message franking: Content moderation for metadata-private end-to-end encryption. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 222–250. Springer, Heidelberg, August 2019.

[68] Nirvan Tyagi, Ian Miers, and Thomas Ristenpart. Traceback for end-to-end encrypted messaging. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 413–430. ACM Press, November 2019.

[69] Nik Unger, Sergej Dechand, Joseph Bonneau, Sascha Fahl, Henning Perl, Ian Goldberg, and Matthew Smith. SoK: Secure messaging. In *2015 IEEE Symposium on Security and Privacy*, pages 232–249. IEEE Computer Society Press, May 2015.

[70] Jelle van den Hooff, David Lazar, Matei Zaharia, and Nickolai Zeldovich. Vuvuzela: scalable private messaging resistant to traffic analysis. In *SOSP*, pages 137–152. ACM, 2015.

[71] Soroush Vosoughi, Deb Roy, and Sinan Aral. The spread of true and false news online. *Science*, 359(6380):1146–1151, 2018.

[72] WhatsApp. Two billion users – connecting the world privately. https://blog.whatsapp.com/two-billion-users-connecting-the-world-privately/, 2020.

[73] Liang Wu, Fred Morstatter, Kathleen M Carley, and Huan Liu. Misinformation in social media: definition, manipulation, and detection. *ACM SIGKDD Explorations Newsletter*, 21(2):80–90, 2019.