

A Universally Composable Framework for the Privacy of Email Ecosystems^{*}

Pyrros Chaidos¹, Olga Fourtounelli¹, Aggelos Kiayias^{2,3}, and Thomas Zacharias²

¹ National and Kapodistrian University of Athens, Greece

{pchaidos,folga}@di.uoa.gr

² The University of Edinburgh, UK

{akiayias,tzachari}@inf.ed.ac.uk

³ IOHK, UK

Abstract. Email communication is amongst the most prominent online activities, and as such, can put sensitive information at risk. It is thus of high importance that internet email applications are designed in a privacy-aware manner and analyzed under a rigorous threat model. The Snowden revelations (2013) suggest that such a model should feature a *global adversary*, in light of the observational tools available. Furthermore, the fact that protecting metadata can be of equal importance as protecting the communication context implies that end-to-end encryption may be necessary, but it is not sufficient.

With this in mind, we utilize the Universal Composability framework [Canetti, 2001] to introduce an expressive cryptographic model for email “ecosystems” that can formally and precisely capture various well-known privacy notions (unobservability, anonymity, unlinkability, etc.), by parameterizing the amount of leakage an ideal-world adversary (simulator) obtains from the email functionality.

Equipped with our framework, we present and analyze the security of two email constructions that follow different directions in terms of the efficiency vs. privacy tradeoff. The first one achieves optimal security (only the online/offline mode of the users is leaked), but it is mainly of theoretical interest; the second one is based on parallel mixing [Golle and Juels, 2004] and is more practical, while it achieves anonymity with respect to users that have similar amount of sending and receiving activity.

1 Introduction

During the last decade, internet users increasingly engage in interactions that put their sensitive information at risk. Social media, e-banking, e-mail, and e-government, are prominent cases where personal data are collected and processed in the web. To protect people’s personal data, it is important that applications intended for communication of such information over the internet are designed in a privacy-aware manner and analyzed under a rigorous threat model.

^{*} This work was supported by the European Union’s Horizon 2020 research and innovation programme under grant agreements No 653497 (project PANORAMIX).

The recent revelations by Snowden (2013) on massive surveillance of citizens’ internet interactions, confirmed researchers’ views that current technology is sufficient to provide adversaries with the tools to monitor the entire network. This was a turning point in that, henceforth, treating internet security and privacy in a threat model that considers a *global adversary* seems not only desirable, but imperative for the design of state-of-the-art cryptographic protocols.

As far as standard security is concerned, i.e., hiding the context between communicating internet users, there have been significant advancements on the aforementioned matter, mainly to due to wide deployment of end-to-end (E2E) encryption tools, even for some of the world’s most popular applications, such as WhatsApp, Viber, Facebook Messenger and Skype (over Signal). However, it is well understood that E2E encryption is not enough to protect the users’ metadata (e.g. users’ identities and location, or the communication time), that often can be of equal importance. The protection of metadata is studied in the context of *anonymous communications*, that were introduced by the seminal works of Chaum with the concept of mix-nets [10], followed by DC-nets a few years later [8]. A mix-net is a message transmission system that aims to decouple the relation of senders to receivers by using intermediate servers to re-encrypt and re-order messages. The operation of mix-nets relies on messages from A to B making intermediate stops in mix servers, with appropriate delay times so that multiple messages “meet” at each server. The server re-encrypts messages before forwarding them, thus breaking the link between incoming and outgoing messages. We will analyse a mix-based system in Sect. 6 and contrast its overhead to the more expensive broadcast solution in Sect. 5. Nowadays, the most scalable solutions of anonymous communications in the real-world rely on *onion-routing* [33], and mostly on the Tor anonymous browser [16]. Although very efficient and a major step forward for privacy-preserving technologies, it has been pointed out (e.g., [21, 32, 34]) that onion-routing can provide anonymity only against adversaries with local views with respect to the (three) relay routing nodes, whereas a global observer can easily derive the addresses of two entities that communicate over onion-routing applications. Towards the goal of communication anonymity against a global adversary [2, 9, 11–13, 24–26, 29, 35], various schemes have been proposed, and several recent ones achieving reasonable latency [1, 9, 25, 26, 29, 35].

Modeling privacy for email ecosystems. In this work, we focus on the study of privacy (as expressed via several anonymity-style notions cf. [27]) for email ecosystems. The reason why we choose to focus on the email case is threefold:

(i). Email is one of the most important aspects of internet communication, as email traffic is estimated to be in the order of $\sim 10^{11}$ messages per day, while there are approximately 2.5 billion accounts worldwide ⁴.

(ii). The actual network infrastructure of an email ecosystem has some special features that encourage a separate study from the general case of private messaging. Namely, the users dynamically register, go online/offline, and communicate,

⁴ <https://www.radicati.com/wp/wp-content/uploads/2014/10/Email-Market-2014-2018-Executive-Summary.pdf>

in a client-friendly environment and the management of the protocol execution is mainly handled by their service providers (SPs) that manage their inboxes. In turn, the client interface allows the user to log in/log out and while online, submit send and fetch requests to their SP. Moreover, adding a subsystem of mix-nodes which, in principle, are functionally different than the clients and the SPs, stratifies the observational power of the global adversary into three layers (i) the client \leftrightarrow SP channels, (ii) the SP \leftrightarrow mix-node channels, and (iii) the channels within the mix-node system. Under this real-world setting, exploring the feasibility and the trade off between efficiency and privacy for anonymous email routing poses restrictions on the expected secrecy, that would not be present in a generic peer-to-peer setting (e.g. users jointly engaging in an MPC execution).

(iii). To the best of our knowledge, there is no prior work on general modeling of email privacy in a computational model, that captures protocol flow under a composition of individual email messaging executions. The *Universal Composability (UC)* framework [6] is the ideal tool for such a modeling.

Contributions. Our contributions are as follows:

1. In Section 3, we introduce a framework for the formal study of email ecosystems in the real-ideal world paradigm of the UC model [6]. The real-world entities involved in our framework comprise the set of clients, the of SPs and the subsystem of mix-nodes; all entities are synchronized via a *global clock* functionality $\mathcal{G}_{\text{clock}}$ and communicate over an *authenticated channel functionality with bounded message delay* Δ_{net} , denoted by $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}$. In the ideal-world, an *email privacy functionality* $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}$ manages email traffic among dummy parties that forward their inputs. The functionality is parameterized by Δ_{net} and a *leakage function* Leak , defined over the history transcript, that formally expresses the underlying privacy notion the studied email ecosystem should satisfy. To illustrate the expressibility of our framework, in Section 4, we show how to formally capture intuitively well understood privacy notions by properly defining the leakage function. In particular, we express and study the relation of notions of anonymity, unlinkability, unobservability and pseudonymity defined in [27], as well as E2E encryption, and a notion we call *weak anonymity* that, although a relaxed version of standard anonymity (still stronger than E2E encryption), provides reasonable privacy guarantees given the setting.

2. In Section 5, we present and formally analyze a theoretical construction with quadratic communication overhead that we prove it achieves unobservability (i.e, only the online/offline mode of the clients is leaked), which we argue that it sets the optimal level of privacy that can be expected under the restrictions posed in our client-SP setting, even against a global adversary that only observes the network. As a result, the said construction shows that in principle, optimal privacy is feasible, while the challenge of every real-world email ecosystem is to balance the privacy vs. efficiency trade off.

3. In Sections 6 and 7 we analyze a construction similar to the classical parallel mix of Golle and Juels [18], to illustrate the expressiveness of our model in a more practice-oriented protocol. We focus on the UC simulation in Section

6, and in Section 7, we use Håstad’s matrix shuffle to model the permutation’s distribution. This in turn makes our analysis relevant to Atom [25], a state of the art anonymity system using similar permutation strategies. At the same time, as we only assume an adversary that is a global passive observer, Atom’s techniques to mitigate corruptions are complementary, even if orthogonal, to our work.

2 Background

2.1 Notation

We use λ as the security parameter and write $\text{negl}(\lambda)$ to denote that some function $f(\cdot)$ is negligible in λ . We write $[n]$ to denote the set $\{1, \dots, n\}$ and $[\![\cdot]\!]$ to denote a multiset. By $X \approx_\epsilon Y$, we denote that the random variable ensembles $\{X_\lambda\}_{\lambda \in \mathbb{N}}, \{Y_\lambda\}_{\lambda \in \mathbb{N}}$ are computationally indistinguishable with error $\epsilon(\cdot)$, i.e., for every probabilistic polynomial time (PPT) algorithm \mathcal{A} , it holds that

$$\left| \Pr[w \leftarrow X_\lambda : \mathcal{A}(w) = 1] - \Pr[w \leftarrow Y_\lambda : \mathcal{A}(w) = 1] \right| < \epsilon(\lambda).$$

We simply write $X \approx Y$ when the error ϵ is $\text{negl}(\lambda)$. The notation $x \stackrel{\$}{\leftarrow} S$ stands for x being sampled from the set S uniformly at random.

2.2 IND-CPA security of public-key encryption schemes

In our constructions, we utilize public-key encryption (PKE). We require that a PKE scheme $\text{PKE} = (\text{KeyGen}, \text{Enc}, \text{Dec})$ satisfies the property of *multiple challenge IND-CPA (m-IND-CPA) security*, which is equivalent to standard IND-CPA security (up to negligible error). We recall that m-IND-CPA with error $\epsilon(\lambda)$ dictates that any adversary \mathcal{B} that (a) obtains the public key, and (b) sends (polynomially many) challenge queries of the form (M_0, M_1) to the challenger receiving encryption of M_b , where b is the random bit of the challenger, can not guess b with more than $1/2 + \epsilon(\lambda)$ probability.

2.3 Related Work

Early works treating anonymity followed the intuitive definition of Pfitzmann and Khöntopp [28] as “the state of not being identifiable within a set of subjects”, and aimed to augment it by quantifying the degree of non-identifiability. One of the first efforts in that direction (predating [28]) was the concept of “ k -anonymity” by Samarati and Sweeney [30], that (in the context of databases) attempts to identify an individual produce at least k candidates.

In [15,31], anonymity is quantified by measuring the probability that a message M was sent by a user U . Thus, we are no longer interested only in the size of the set of candidates, but also their relative probabilities. This definition improved upon the “folklore” metric of only measuring the size of the subject set, even if the probability distribution on that set was highly non-uniform –e.g. [23].

The seminal work of Dwork [17] on Differential Privacy, while originating in the realm of databases, highlights and formalizes the strength of combining different pieces of seemingly privacy-respecting information to arrive at a privacy-impacting conclusion. Influenced in part by Differential Privacy, AnoA [3] is a game-based privacy analysis framework that is flexible enough to model various privacy concepts. Their approach is based on games between a challenger and an adversary who has control of the execution of the game, apart from a challenge message representing the scenarios the adversary is trying to distinguish.

In a different direction, the Universal Composability (UC) framework, [6] models security as a simulation not against an adversary, but a malicious environment, given strong control over the inputs of each party as well as a complete view of the system. This rigorous approach produces strong and composable security guarantees but is quite demanding in that the simulation must operate with the bare minimum of data (i.e. what we assume the protocol leaks). This precision in both simulation and leakage is a key motivation of this work.

On the other hand, state of the art anonymous communication solutions such as Loopix [29] which aims for high performance while maintaining strong anonymity properties, as well as unobservability, are analyzed under a weaker adversary. Moreover, Atom [25] is engineered to provide statistical indistinguishable shuffling with strong safeguards against malicious servers, but lacks formal proofs. In our work, we analyze a construction that shares a similar design (namely Håstad’s matrix shuffle), so that we are able to offer a suggested T value (i.e mix length) as a side contribution in Section 7. A key difference between Loopix and Atom is that Loopix uses a free routing approach (i.e a message’s path is determined by its sender) as opposed to allowing mix nodes to route messages. The first approach is more agreeable with high-efficiency solution aiming for a practical level of resilience against active adversaries while the second approach is easier to reason about but requires a passive adversary or measures such as NIZKs or trap messages to ensure correct behavior.

Camenish and Lysyanskaya [5] offer a treatment of onion routing in the Universal Composability model. The defining characteristic of onion routing, is that routing is entirely determined by the initial sender and is not influenced by the intervening nodes. As such, their analysis focuses on defining security with regards to the encryption, padding, structuring and layering of onions rather than the routing strategy itself. This is orthogonal to our approach: we focus on evaluating the anonymity of different mixing strategies under what we view as realistic requirements about the message encapsulation.

Wikström [36] covers the UC-security of a specific mix construction. His analysis is well-suited to voting but is hard to generalize over other use cases and performance parameters. In contrast, our work, while focusing on email, is more general and flexible in regards to leakage, timings and network topology.

In the work of Alexopoulos *et al.* [1], anonymity is studied in the concept of messaging via a stand-alone simulation-based model. Even though formally treated, anonymity in [1] is defined under a framework that is weaker than UC.

3 A UC framework for the privacy of email ecosystems

In this section, we present our UC framework for email privacy. As in standard UC approach, privacy will be defined via the inability of an environment \mathcal{Z} , that schedules the execution and provides the inputs, to distinguish between (i) a real-world execution of an email ecosystem \mathbb{E} in the presence of a (global passive) adversary \mathcal{A} and (ii) an ideal-world execution handled by an email privacy functionality interacting with a PPT simulator Sim . More specifically, we adjust our definitions to the *global UC* setting [7], by incorporating a global clock functionality (cf. [4, 22]) that facilitates synchronicity and is accessed by all parties, including the environment.

3.1 Entities and protocols of an email ecosystem

The entities that are involved in a private email “ecosystem” \mathbb{E} are the following:

- The *service providers (SPs)* $\text{SP}_1, \dots, \text{SP}_N$ that register users and are responsible for mailbox management and email transfer.
- The *clients* C_1, \dots, C_n that wish to exchange email messages and are registered to the SPs. For simplicity, we assume that each client is linked with only one SP that manages her mailbox. We write $C_\ell @ \text{SP}_i$ to denote that C_ℓ is registered to SP_i , where registration is done dynamically. We define the set $\mathbf{C}_i := \{C_\ell \mid C_\ell @ \text{SP}_i\}$ of all clients whose mailboxes SP_i is managing.
- The *mix node* subsystem \mathbf{MX} that consists of the mix nodes $\text{MX}_1, \dots, \text{MX}_m$ and is the core of the anonymous email routing mechanism.

An email ecosystem \mathbb{E} has the two following phases:

■ **Initialization** is a setup phase where all SPs and mix nodes generate any possible private and public data, and commune their public data to a subset of the ecosystem’s entities.

■ **Execution** is a phase that comprises executions of the following protocols:

- The REGISTER protocol between client C_s and her service provider SP_i . For simplicity, we assume that registration can be done only once.
- The SEND protocol between client C_s and her service provider SP_i . In particular, C_s that wishes to send a message M to some client address $C_r @ \text{SP}_j$ authenticates to SP_i and provides her with an encoding $\text{Encode}(M, C_r @ \text{SP}_j)$ of $(M, C_r @ \text{SP}_j)$ (that may not necessarily include information about the sender). At the end of the protocol, $\text{Encode}(M, C_r @ \text{SP}_j)$ is at the outbox of $C_s @ \text{SP}_i$ managed by SP_i .
- The ROUTE protocol that is executed among $\text{SP}_1, \dots, \text{SP}_N$ and $\text{MX}_1, \dots, \text{MX}_m$. Namely, the encoded message $\text{Encode}(M, C_r @ \text{SP}_j)$ is forwarded to the \mathbf{MX} subsystem, which in turn delivers it to SP_j that manages the inbox of C_r .
- The RECEIVE protocol between client C_r and her service provider SP_j , where C_r can retrieve the messages from the inbox of $C_r @ \text{SP}_j$ via fetch requests.

Remark 1. In this work, we consider email solutions that follow the realistic client-side approach, where the client-side operations are relatively simple and do not include complex interaction with the other entities for the execution of heavy cryptographic primitives (e.g. pairwise secure MPC). As we will explain shortly, the client-friendly approach poses some limitations on the privacy level that the email ecosystem can achieve.

3.2 A global clock functionality

In our setting, the protocol flow within the email ecosystem \mathbb{E} advances in *time slots*, that could refer to any suitable time unit (e.g. ms). The entities of \mathbb{E} are synchronized via a *global clock functionality* $\mathcal{G}_{\text{clock}}$ that interacts with a set of parties \mathbf{P} , a set of functionalities \mathbf{F} , the UC environment \mathcal{Z} and the adversary \mathcal{A} . In the spirit of [4, 22], the functionality $\mathcal{G}_{\text{clock}}$, presented in Fig. 1, advances when all entities in \mathbf{P} and \mathbf{F} declare completion of their activity within the current time slot, whereas all entities have read access to it.

The global clock functionality $\mathcal{G}_{\text{clock}}(\mathbf{P}, \mathbf{F})$.

The functionality initializes the global clock variable as $\text{Cl} \leftarrow 0$ and the set of advanced parties as $L_{\text{adv}} \leftarrow \emptyset$.

- Upon receiving $(\text{sid}, \text{ADVANCE_CLOCK}, P)$ from $\mathcal{F} \in \mathbf{F}$ or $P \in \mathbf{P}$, if $P \notin L_{\text{adv}}$, then it adds P to L_{adv} and sends the message $(\text{sid}, \text{ADVANCE_ACK}, P)$ to \mathcal{F} or P , respectively, and notifies \mathcal{A} by forwarding $(\text{sid}, \text{ADVANCE_CLOCK}, P)$. If $L_{\text{adv}} = \mathbf{P}$, then it updates as $\text{Cl} \leftarrow \text{Cl} + 1$ and resets $L_{\text{adv}} \leftarrow \emptyset$.
- Upon receiving $(\text{sid}, \text{READ_CLOCK})$ from $X \in \mathbf{P} \cup \mathbf{F} \cup \{\mathcal{Z}, \mathcal{A}\}$, then it sends $(\text{sid}, \text{READ_CLOCK}, \text{Cl})$ to X .

Fig. 1. The global clock functionality $\mathcal{G}_{\text{clock}}(\mathbf{P}, \mathbf{F})$ interacting with the environment \mathcal{Z} and the adversary \mathcal{A} .

3.3 A UC definition of e-mail privacy

Let \mathbf{Ad} be the set of all valid email addresses linking the set of clients $\mathbf{C} = \{C_1, \dots, C_n\}$ and with their corresponding providers in $\mathbf{SP} = \{\text{SP}_1, \dots, \text{SP}_N\}$, i.e. $\mathbf{Ad} := \cup_{i \in [N]} \mathbf{C}_i$. We denote by \mathbf{P} the union $\mathbf{C} \cup \mathbf{SP} \cup \mathbf{MX}$.

The *history* of an email ecosystem execution that involves the entities in \mathbf{C}, \mathbf{SP} and \mathbf{MX} is a transcript of actions expressed as a list H , where each action entry of H is associated with a unique pointer ptr to this action. The leakage in each execution step, is expressed via a *leakage function* $\text{Leak}(\cdot, \cdot)$ that, when given as input (i) a pointer ptr and (ii) an execution history sequence H , outputs some leakage string z . Here, z could be \perp indicating no leakage to the adversary. This leakage may depend on the entry indexed by ptr as well as on entries recorded previously (i.e. prior than ptr).

We require that during a time slot, the environment sends a message for every party, even when the party is idle (inactive) for this slot, so that the clock can be advanced as described in Fig. 1.

The ideal world execution. In the ideal world, the protocol execution is managed by the *email privacy functionality* $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}(\mathbf{P})$, parameterized by the message delivery delay bound Δ_{net} and the leakage function $\text{Leak}(\cdot, \cdot)$, with access to $\mathcal{G}_{\text{clock}}$. The functionality $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}(\mathbf{P})$ consists of the **Initialization**, **Execution**, and **Clock advancement** phases, that informally are run as follows:

- At the **Initialization** phase, all the SPs in **SP** and mix nodes in **MX** provide $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}(\mathbf{P})$ with an initialization notification via public delayed output. The functionality proceeds to the **Execution** phase when all SPs and mix nodes are initialized. Note that in the ideal world, the SPs and the mix nodes remain idle after **Initialization** (besides messages intended for $\mathcal{G}_{\text{clock}}$), as privacy-preserving email routing is done by $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}(\mathbf{P})$. Their presence in the ideal setting is for consistency in terms of UC interface.

- At the **Execution** phase, $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}(\mathbf{P})$ manages the email traffic, as scheduled per time slot by the environment. During this phase, the clients may (dynamically) provide $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}(\mathbf{P})$ with registration, log in, log out, send or fetch requests. Upon receiving a request, $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}(\mathbf{P})$ updates the history by adding the request as a new entry associated with a unique pointer ptr , in a ‘pending’ mode. Then, it notifies the simulator **Sim** by attaching the corresponding leakage. The execution of a pending request which record is indexed by a pointer ptr is completed when $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}(\mathbf{P})$ receives an `ALLOW_EXEC` message paired with ptr from **Sim**.

Within a time slot T , each client may perform only one action that also implies a time advancement request to $\mathcal{G}_{\text{clock}}$. In order for the clock to advance all the other parties that performed no action (i.e., the SPs, the mix nodes and the clients that remained idle during T), send an explicit time advancement request to $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}(\mathbf{P})$. Besides, any party may submit clock reading requests arbitrarily. All the messages that are intended for $\mathcal{G}_{\text{clock}}$ are forwarded to it by $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}(\mathbf{P})$.

- At the **Clock advancement** phase, all parties have already submitted time advancement requests during time slot T , so $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}(\mathbf{P})$ takes the necessary final steps before proceeding to $T + 1$. In particular, $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}(\mathbf{P})$ completes the execution of all send and fetch requests that have been delayed for Δ_{net} steps (by **Sim**). This suggests that in the ideal-world, the delay in message delivery is upper bounded by Δ_{net} . Finally, $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}(\mathbf{P})$ informs **Sim** of the leakage derived from the aforementioned executions, advances its local time by 1 and reenters the **Execution** phase for time slot $T + 1$.

Formally, the email privacy functionality $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}(\mathbf{P})$ is described as follows:

Initialization on status ‘init’.

- $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}(\mathbf{P})$ sets its status to ‘init’. It initializes the set of valid addresses Ad , the set of active entities L_{act} , the set of clock-advanced entities L_{adv} , the *history* list H , and the set of leaked entries L_{leak} as empty.

- Upon receiving $(\text{sid}, \text{INIT})$ from a party $P \in \mathbf{SP} \cup \mathbf{MX}$, if $L_{\text{act}} \subsetneq \mathbf{SP} \cup \mathbf{MX}$, then it sends the message $(\text{sid}, \text{INIT}, P)$ to Sim.
- Upon receiving $(\text{sid}, \text{ALLOW_INIT}, P)$ from Sim, if $P \in (\mathbf{SP} \cup \mathbf{MX}) \setminus L_{\text{act}}$, then it adds P to L_{act} . If $L_{\text{act}} = \mathbf{SP} \cup \mathbf{MX}$, then it sends $(\text{sid}, \text{ready})$ to Sim.
- Upon receiving $(\text{sid}, \text{EXECUTE})$ from Sim, it sends $(\text{sid}, \text{READ_CLOCK})$ to $\mathcal{G}_{\text{clock}}$.
- Upon receiving $(\text{sid}, \text{READ_CLOCK}, \text{Cl})$ from $\mathcal{G}_{\text{clock}}$, it sets its clock as Cl and its status to ‘execute’, and sends the message $(\text{sid}, \text{start}, \text{Cl})$ to Sim.

Execution on status ‘execute’.

Registration:

- Upon receiving $(\text{sid}, \text{REGISTER}, @\text{SP}_i)$ from C_ℓ , if for every $j \in [N] : C_\ell @ \text{SP}_j \notin \mathbf{Ad}$ and $C_\ell \notin L_{\text{adv}}$, then
 1. It sends the message $(\text{sid}, \text{ADVANCE_CLOCK}, C_\ell)$ to $\mathcal{G}_{\text{clock}}$.
 2. Upon receiving $(\text{sid}, \text{ADVANCE_ACK}, C_\ell)$ from $\mathcal{G}_{\text{clock}}$, it adds C_ℓ to L_{adv} and the entry $(\text{ptr}, (\text{sid}, \text{Cl}, \text{REGISTER}, C_\ell @ \text{SP}_i), \text{‘pending’})$ to H .
 3. It sends the message $(\text{sid}, \text{ptr}, \text{Leak}(\text{ptr}, H))$ to Sim.
 4. Upon receiving $(\text{sid}, \text{ALLOW_EXEC}, \text{ptr})$ from Sim, if ptr refers to an entry of the form $(\text{ptr}, (\text{sid}, \text{Cl}, \text{REGISTER}, C_\ell @ \text{SP}_i), \text{‘pending’})$, then
 - (a) It adds $C_\ell @ \text{SP}_i$ to \mathbf{Ad} and L_{act} , and initializes a list $\text{Inbox}[C_\ell @ \text{SP}_i]$ as empty.
 - (b) It updates the entry as $(\text{ptr}, (\text{sid}, \text{Cl}', \text{REGISTER}, C_\ell @ \text{SP}_i), \text{‘(registered, Cl)’})$.
 - (c) It sends the message $(\text{sid}, \text{ptr}, \text{Leak}(\text{ptr}, H))$ to Sim.

Log in:

- Upon receiving $(\text{sid}, \text{ACTIVE}, @\text{SP}_i)$ from C_ℓ , if $C_\ell @ \text{SP}_i \in \mathbf{Ad}$ and $C_\ell \notin L_{\text{adv}}$,
 1. It sends the message $(\text{sid}, \text{ADVANCE_CLOCK}, C_\ell)$ to $\mathcal{G}_{\text{clock}}$.
 2. Upon receiving $(\text{sid}, \text{ADVANCE_ACK}, C_\ell)$ from $\mathcal{G}_{\text{clock}}$, it adds C_ℓ to L_{adv} and the entry $(\text{ptr}, (\text{sid}, \text{Cl}, \text{ACTIVE}, C_\ell @ \text{SP}_i), \text{‘pending’})$ to H .
 3. It sends the message $(\text{sid}, \text{ptr}, \text{Leak}(\text{ptr}, H))$ to Sim.
 4. Upon receiving $(\text{sid}, \text{ALLOW_EXEC}, \text{ptr})$ from Sim, if ptr refers to an entry of the form $(\text{ptr}, (\text{sid}, \text{Cl}', \text{ACTIVE}, C_\ell @ \text{SP}_i), \text{‘pending’})$, then
 - (a) If $C_\ell @ \text{SP}_i \notin L_{\text{act}}$, then it adds $C_\ell @ \text{SP}_i$ to L_{act} .
 - (b) It updates the entry as $(\text{ptr}, (\text{sid}, \text{Cl}', \text{ACTIVE}, C_\ell @ \text{SP}_i), \text{‘(logged in, Cl)’})$.
 - (c) It sends the message $(\text{sid}, \text{ptr}, \text{Leak}(\text{ptr}, H))$ to Sim.

Log out:

- Upon receiving $(\text{sid}, \text{INACTIVE}, @\text{SP}_i)$ from C_ℓ , if $C_\ell @ \text{SP}_i \in \mathbf{Ad}$ and $C_\ell \notin L_{\text{adv}}$, then
 1. It sends the message $(\text{sid}, \text{ADVANCE_CLOCK}, C_\ell)$ to $\mathcal{G}_{\text{clock}}$.
 2. Upon receiving $(\text{sid}, \text{ADVANCE_ACK}, C_\ell)$ from $\mathcal{G}_{\text{clock}}$, it adds C_ℓ to L_{adv} and the entry $(\text{ptr}, (\text{sid}, \text{Cl}, \text{INACTIVE}, C_\ell @ \text{SP}_i), \text{‘pending’})$ to H .
 3. It sends the message $(\text{sid}, \text{ptr}, \text{Leak}(\text{ptr}, H))$ to Sim.
 4. Upon receiving $(\text{sid}, \text{ALLOW_EXEC}, \text{ptr})$ from Sim, if ptr refers to an entry of the form $(\text{ptr}, (\text{sid}, \text{Cl}', \text{INACTIVE}, C_\ell @ \text{SP}_i), \text{‘pending’})$, then
 - (a) If $C_\ell @ \text{SP}_i \in L_{\text{act}}$, then it deletes $C_\ell @ \text{SP}_i$ from L_{act} .

- (b) It updates the entry as $(\text{ptr}, (\text{sid}, \text{Cl}', \text{INACTIVE}, C_\ell @ \text{SP}_i), '(\text{logged out}, \text{Cl})')$.
- (c) It sends the message $(\text{sid}, \text{ptr}, \text{Leak}(\text{ptr}, H))$ to Sim.

Send:

- Upon receiving $(\text{sid}, \text{SEND}, \langle C_s @ \text{SP}_i, M, C_r @ \text{SP}_j \rangle)$ from C_s , if $C_s @ \text{SP}_i, C_r @ \text{SP}_j \in \mathbf{Ad}$ and $C_s \in L_{\text{act}} \setminus L_{\text{adv}}$, then
 1. It sends the message $(\text{sid}, \text{ADVANCE_CLOCK}, C_s)$ to $\mathcal{G}_{\text{clock}}$.
 2. Upon receiving $(\text{sid}, \text{ADVANCE_ACK}, C_s)$ from $\mathcal{G}_{\text{clock}}$, it adds C_s in L_{adv} and the entry $(\text{ptr}, (\text{sid}, \text{Cl}, \text{SEND}, \langle C_s @ \text{SP}_i, M, C_r @ \text{SP}_j \rangle), '(\text{pending})')$ to H .
 3. It sends the message $(\text{sid}, \text{ptr}, \text{Leak}(\text{ptr}, H))$ to Sim.
 4. Upon receiving $(\text{sid}, \text{ALLOW_EXEC}, \text{ptr})$ from Sim, if ptr refers to an entry $(\text{sid}, \text{Cl}', \text{SEND}, \langle C_s @ \text{SP}_i, M, C_r @ \text{SP}_j \rangle)$ with status 'pending', then
 - (a) It adds $(\text{sid}, \text{Cl}', \text{SEND}, \langle C_s @ \text{SP}_i, M, C_r @ \text{SP}_j \rangle)$ to $\text{Inbox}[C_r @ \text{SP}_j]$.
 - (b) It updates as $(\text{ptr}, (\text{sid}, \text{Cl}', \text{SEND}, \langle C_s @ \text{SP}_i, M, C_r @ \text{SP}_j \rangle), '(\text{sent}, \text{Cl})')$.
 - (c) It sends the message $(\text{sid}, \text{ptr}, \text{Leak}(\text{ptr}, H))$ to Sim.

Fetch:

- Upon receiving $(\text{FETCH}, \text{sid}, C_r @ \text{SP}_j)$ from C_r , if $C_r @ \text{SP}_j \in \mathbf{Ad}$ and $C_r \in L_{\text{act}} \setminus L_{\text{adv}}$, then
 1. It sends the message $(\text{sid}, \text{ADVANCE_CLOCK}, C_r)$ to $\mathcal{G}_{\text{clock}}$.
 2. Upon receiving $(\text{sid}, \text{ADVANCE_ACK}, C_r)$ from $\mathcal{G}_{\text{clock}}$, it adds C_s in L_{adv} and the entry $(\text{ptr}, (\text{sid}, \text{Cl}, \text{FETCH}, C_r @ \text{SP}_j), '(\text{pending})')$ to H .
 3. It sends the message $(\text{sid}, \text{ptr}, \text{Leak}(\text{ptr}, H))$ to Sim.
 4. Upon receiving $(\text{sid}, \text{ALLOW_EXEC}, \text{ptr})$ from Sim, if ptr refers to an entry of the form $(\text{sid}, \text{Cl}', \text{FETCH}, C_r @ \text{SP}_j)$ with status 'pending', then
 - (a) It sends the message $(\text{sid}, \text{Inbox}[C_r @ \text{SP}_j])$ to C_r .
 - (b) It updates the entry as $(\text{ptr}, (\text{sid}, \text{Cl}', \text{FETCH}, C_r @ \text{SP}_j), '(\text{fetched}, \text{Cl})')$.
 - (c) It resets $\text{Inbox}[C_r @ \text{SP}_j]$ as empty.
 - (d) It sends the message $(\text{sid}, \text{ptr}, \text{Leak}(\text{ptr}, H))$ to Sim.

Clock reading:

- Upon receiving $(\text{sid}, \text{READ_CLOCK})$ from a party $P \in \mathbf{P}$, then
 1. It sends the message $(\text{sid}, \text{READ_CLOCK})$ to $\mathcal{G}_{\text{clock}}$.
 2. On receiving $(\text{sid}, \text{READ_CLOCK}, \text{Cl})$ from $\mathcal{G}_{\text{clock}}$ it adds $(\text{ptr}, (\text{sid}, \text{Cl}, \text{READ_CLOCK}, P))$ to H , sending $(\text{sid}, \text{READ_CLOCK}, \text{Cl})$ to P .
 3. It sends the message $(\text{sid}, \text{ptr}, \text{Leak}(\text{ptr}, H))$ to Sim.

Clock advance:

- Upon receiving $(\text{sid}, \text{ADVANCE_CLOCK})$ from a party $P \in \mathbf{P} \setminus L_{\text{adv}}$, then
 1. It sends the message $(\text{sid}, \text{ADVANCE_CLOCK}, P)$ to $\mathcal{G}_{\text{clock}}$.
 2. Upon receiving $(\text{sid}, \text{ADVANCE_ACK}, P)$ from $\mathcal{G}_{\text{clock}}$, it adds P in L_{adv} and $(\text{ptr}, (\text{sid}, \text{Cl}, \text{ADVANCE_CLOCK}, P))$ to H .
 3. It sends the message $(\text{sid}, \text{ptr}, \text{Leak}(\text{ptr}, H))$ to Sim.
 4. If $L_{\text{adv}} = \mathbf{P}$, then it sets its status to 'advance' and proceeds to the **Clock advancement** phase below.

Clock advancement on status 'advance':

- Upon setting its status to 'advance':

1. For every history entry of the form $(\text{sid}, \text{Cl}', \text{SEND}, \langle C_s @ \text{SP}_i, M, C_r @ \text{SP}_j \rangle)$ with status ‘pending’ such that $\text{Cl} - \text{Cl}' = \Delta_{\text{net}}$, it adds this entry to $\text{Inbox}[C_r @ \text{SP}_j]$ and updates the entry’s status to ‘(sent, Cl)’.
2. For every history entry of the form $(\text{sid}, \text{Cl}', \text{FETCH}, C_r @ \text{SP}_j)$ with status ‘pending’ such that $\text{Cl} - \text{Cl}' = \Delta_{\text{net}}$, it sends the message $(\text{sid}, \text{Inbox}[C_r @ \text{SP}_j])$ to C_r , resets the list $\text{Inbox}[C_r @ \text{SP}_j]$ as input and updates the entry’s status to ‘(fetched, Cl)’.
3. It sends the message $(\text{sid}, \text{ptr}, \text{Leak}(\text{ptr}, H))$ to Sim .
4. It finalizes execution for the current slot as follows:
 - (a) It advances its time by $\text{Cl} \leftarrow \text{Cl} + 1$.
 - (b) It adds $(\text{ptr}, (\text{sid}, \text{CLOCK_ADVANCED}))$ to H .
 - (c) It reverts its status to ‘execute’ and resets L_{adv} to empty.
 - (d) It sends the message $(\text{sid}, \text{CLOCK_ADVANCED})$ to Sim .

We denote by $\text{EXEC}_{\text{Sim}, \mathcal{Z}, \mathcal{G}_{\text{clock}}}^{\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}}[\mathbf{P}](\lambda)$, the output of the environment \mathcal{Z} in an ideal-world execution of $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}(\mathbf{P})$ under the presence of Sim .

The authenticated channel functionality $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$.

The functionality initializes a list of pending messages L_{pend} as empty.

- Upon receiving $(\text{sid}, \text{CHANNEL}, M, P')$ from $P \in \mathbf{P}$, then
 1. It sends the message $(\text{sid}, \text{READ_CLOCK})$ to $\mathcal{G}_{\text{clock}}(\mathbf{P})$.
 2. Upon receiving $(\text{sid}, \text{READ_CLOCK}, \text{Cl})$ to $\mathcal{G}_{\text{clock}}(\mathbf{P})$, it picks a unique pointer ptr and stores the entry $(\text{ptr}, (\text{sid}, \text{Cl}, \text{CHANNEL}, P, M, P'))$ to L_{pend} .
 3. It sends the message $(\text{ptr}, (\text{sid}, \text{CHANNEL}, P, M, P'))$ to \mathcal{A} .
- Upon receiving $(\text{sid}, \text{ALLOW_CHANNEL}, \text{ptr}')$ from \mathcal{A} , if there is an entry $(\text{ptr}', (\text{sid}, \text{Cl}', \text{CHANNEL}, P, M, P'))$ in L_{pend} , then it sends the message (sid, M, P) to P' and deletes $(\text{ptr}', (\text{sid}, \text{Cl}', \text{CHANNEL}, P, M, P'))$ from L_{pend} .
- Upon any activation from a party $P \in \mathbf{P}$ or \mathcal{A} as above,
 1. It sends the message $(\text{sid}, \text{READ_CLOCK})$ to $\mathcal{G}_{\text{clock}}(\mathbf{P})$.
 2. Upon receiving $(\text{sid}, \text{READ_CLOCK}, \text{Cl})$ to $\mathcal{G}_{\text{clock}}(\mathbf{P})$, it parses L_{pend} . For every entry $(\text{ptr}', (\text{sid}, \text{Cl}', \text{CHANNEL}, P, M, P'))$ s.t. $\text{Cl} - \text{Cl}' = \Delta_{\text{net}}$, it sends the message (sid, M, P) to P' and deletes $(\text{ptr}', (\text{sid}, \text{Cl}', \text{CHANNEL}, P, M, P'))$ from L_{pend} .

Fig. 2. The authenticated channel functionality $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$ interacting with the adversary \mathcal{A} .

The $(\mathcal{G}_{\text{clock}}, \mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}})$ -hybrid world execution. In the real world email ecosystem \mathbb{E} , the clients, the SPs and the mix nodes interact according to the protocols’ guidelines and the environment’s instructions. The message delivery is executed via the functionality $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$ described in Fig. 2 that captures the notion of an authenticated channel, upon which a maximum delivery delay Δ_{net} can be

imposed. Clock advancement is done via calls to $\mathcal{G}_{\text{clock}}$, which interacts with all entities and $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}$.

We denote by $\text{EXEC}_{\mathcal{A}, \mathcal{Z}, \mathcal{G}_{\text{clock}}}^{\mathbb{E}^{\mathcal{G}_{\text{clock}}, \mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}}}[\mathbf{P}](\lambda)$ the output of the environment \mathcal{Z} in an execution of $\mathbb{E}^{\mathcal{G}_{\text{clock}}, \mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}}$ under the presence of \mathcal{A} .

The UC definition of a private email ecosystem is provided below.

Definition 1 (UC Email Privacy). *Let $\Delta_{\text{net}}, \epsilon$ be non-negative values. Let \mathbb{E} be an email ecosystem with client set $\mathbf{C} = C_1, \dots, C_n$, service provider set $\mathbf{SP} = \text{SP}_1, \dots, \text{SP}_N$ and mix node set $\mathbf{MX} = \text{MX}_1, \dots, \text{MX}_m$. Let $\mathbf{P} := \mathbf{C} \cup \mathbf{SP} \cup \mathbf{MX}$. We say that $\mathbb{E}^{\mathcal{G}_{\text{clock}}, \mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}}$ achieves statistical (resp. computational) ϵ -privacy with respect to leakage (Leak) and message delay Δ_{net} , if for every unbounded (resp. PPT) global passive adversary \mathcal{A} , there is a PPT simulator Sim such that for every PPT environment \mathcal{Z} , it holds that*

$$\text{EXEC}_{\text{Sim}, \mathcal{Z}, \mathcal{G}_{\text{clock}}}^{\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}}[\mathbf{P}](\lambda) \approx_{\epsilon} \text{EXEC}_{\mathcal{A}, \mathcal{Z}, \mathcal{G}_{\text{clock}}}^{\mathbb{E}^{\mathcal{G}_{\text{clock}}, \mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}}}[\mathbf{P}](\lambda).$$

4 Formalizing privacy notions via types of leakage functions

In [27], Pfitzmann and Hansen provide definitions for anonymity, unlinkability, unobservability and pseudonymity. Even though outside the context of a formal framework, the definitions in this seminal work have served as a reference point by researchers for the understanding of privacy notions. In this section, we formally express the said (yet not only these) notions by carefully specifying a corresponding leakage function.

Basic leakage sets. Below, we define some useful sets that will enable the succinct description of the various leakage functions that we will introduce. In our formalization, leakage will derive from the history entries that are in a ‘pending’ mode. This is due to technical reasons, as the ideal-world simulator Sim (cf. Section 3.3) must be aware of the actions to be taken by the email privacy functionality $\mathcal{F}_{\text{priv}}^{\text{Leak}, \Delta_{\text{net}}}(\mathbf{P})$ before allowing their execution, so that it can simulate the real-world run in an indistinguishable manner. In the following, the symbol $*$ denotes a wildcard, and $\text{ptr}' \leq \text{ptr}$ denotes that entry indexed with pointer ptr' was added earlier than the entry with pointer ptr .

– The *active address set for H by pointer ptr* :

$$\begin{aligned} \text{Act}_{\text{ptr}}[H] =: & \left\{ C_{\ell} @ \text{SP}_i \mid \exists \text{ptr}' \leq \text{ptr} : \left[[(\text{ptr}', (\text{sid}, *, \text{ACTIVE}, C_{\ell} @ \text{SP}_i), \text{'pending'}) \in H] \vee \right. \right. \\ & \vee [(\text{ptr}', (\text{sid}, *, \text{REGISTER}, C_{\ell} @ \text{SP}_i), \text{'pending'}) \in H] \left. \right] \wedge \\ & \wedge \left[\forall \text{ptr}'' : \text{ptr}' \leq \text{ptr}'' \leq \text{ptr} \Rightarrow (\text{ptr}'', (\text{sid}, *, \text{INACTIVE}, C_{\ell} @ \text{SP}_i), \text{'pending'}) \notin H \right] \left. \right\}. \end{aligned}$$

Note. To simplify the notation and terminology that follows, we consider as active all the addresses that are in a pending registration status.

– The *sender set for H by pointer ptr* :

$$\mathbf{S}_{\text{ptr}}[H] := \left\{ C_s @ \text{SP}_i \mid \exists \text{ptr}' \leq \text{ptr} : (\text{ptr}', (\text{sid}, *, \text{SEND}, \langle C_s @ \text{SP}_i, *, * \rangle), \text{'pending'}) \in H \right\}.$$

– The *sender multiset for H by pointer ptr* , denoted by $\llbracket \mathbf{S}_{\text{ptr}} \rrbracket [H]$, is defined analogously. The difference with $\mathbf{S}_{\text{ptr}}[H]$ is that the cardinality of the pending SEND messages provided by $C_s @ \text{SP}_i$ is attached.

– The *message-sender set for H by pointer ptr* :

$$\mathbf{MS}_{\text{ptr}}[H] := \left\{ (M, C_s @ \text{SP}_i) \mid \exists \text{ptr}' \leq \text{ptr} : (\text{ptr}', (\text{sid}, *, \text{SEND}, \langle C_s @ \text{SP}_i, M, * \rangle), \text{'pending'}) \in H \right\}.$$

– The *recipient set for H by pointer ptr* :

$$\mathbf{R}_{\text{ptr}}[H] := \left\{ C_r @ \text{SP}_j \mid \exists \text{ptr}' \leq \text{ptr} : (\text{ptr}', (\text{sid}, *, \text{SEND}, \langle *, *, C_r @ \text{SP}_j \rangle), \text{'pending'}) \in H \right\}.$$

– The *recipient multiset for H at time slot T* , denoted by $\llbracket \mathbf{R}_{\text{ptr}} \rrbracket [H]$, is defined analogously. The difference with $\mathbf{R}_{\text{ptr}}[H]$ is that the cardinality of the pending SEND messages intended for $C_r @ \text{SP}_j$ is attached.

– The *message-recipient set for H by pointer ptr* :

$$\mathbf{MR}_{\text{ptr}}[H] := \left\{ (M, C_r @ \text{SP}_j) \mid \exists \text{ptr}' \leq \text{ptr} : (\text{ptr}', (\text{sid}, *, \text{SEND}, \langle *, M, C_r @ \text{SP}_j \rangle), \text{'pending'}) \in H \right\}.$$

– The *set of fetching clients for H by pointer ptr*

$$\mathbf{F}_{\text{ptr}}[H] := \left\{ C_r @ \text{SP}_j \mid \exists \text{ptr}' \leq \text{ptr} : (\text{ptr}', (\text{sid}, *, \text{FETCH}, C_r @ \text{SP}_j), \text{'pending'}) \in H \right\}.$$

Unobservability. Unobservability is the state where “the messages are not discernible from random noise”. Here, we focus on the case of *relationship unobservability*, that we will refer to unobservability for brevity, where within the set of all possible sender-recipient-pairs, a message is exchanged in any relationship. Hence, in our setting the *unobservability set* is the set of the users that are online, i.e. only the “activity bit”. As a result, we can define the *unobservability leakage function* $\text{Leak}_{\text{unob}}$ as the active address set:

$$\text{Leak}_{\text{unob}}(\text{ptr}, H) := \text{Act}_{\text{ptr}}[H]. \quad (1)$$

Remark 2 (Unobservability a golden standard from email privacy). In our UC formalization of e-mail ecosystems, we consider a dynamic scenario where the clients register, go online/offline and make custom fetch requests, which is consistent with the real-world dynamics of email communication. It is easy

to see that in such a setting the clients’ online/offline status may be leaked to a global observer. E.g., the environment may provide send requests to offline clients and notify the global adversary that provided the said requests, so that the latter can check the activity of those clients. Hence, in our framework, unobservability as defined in Eq. (1), sets a “golden standard” for optimal privacy. In Section 5, we show that this golden standard is feasible in principle. Namely, we describe a theoretical construction with quadratic communication complexity and we prove it achieves unobservability. As a result, that construction sets one extreme point in the privacy vs. efficiency trade off for the client-server email infrastructure, the other being a simple and fast network with no security enhancements. Clearly, the challenge of every email construction is to balance the said trade off between these two extreme points.

We conclude our remark noting that a higher level privacy (e.g., no leakage at all) could be possible if we considered an alternative setting where the email addresses are a priori given, the clients are always online and mail delivery is via continuous push by the SPs. However, we believe that such a setting is restrictive for formally capturing what is an email ecosystem in general.

Anonymity. According to [27], *anonymity* “is the state of being not identifiable within a set of subjects, the anonymity set”. In the email scenario, a sender (resp. recipient) should be anonymous within the set of potential senders (resp. recipients), i.e. the *sender* (resp. *recipient*) *anonymity set*. In addition, anonymity sets may change over time, which in our framework is done via global clock advancement and per slot. We recall from the discussion in Remark 2 that in our setting, the anonymity sets are restricted within the set of online users.

We define the predicate $\text{End}(\cdot, \cdot)$ over the pointers and history transcripts to denote that a pointer ptr refers to the last history entry before the functionality enters the **Clock advancement** phase in order to finalize execution for the running time slot. By the above, we define the *anonymity leakage function*, $\text{Leak}_{\text{anon}}$, as follows:

$$\text{Leak}_{\text{anon}}(\text{ptr}, H) := \begin{cases} (\mathbf{S}_{\text{ptr}}[H], \mathbf{R}_{\text{ptr}}[H], \text{Act}_{\text{ptr}}[H]), & \text{if } \text{End}(\text{ptr}, H) \\ \text{Act}_{\text{ptr}}[H], & \text{otherwise} \end{cases} \quad (2)$$

Unlinkability. Unlinkability of items of interest (e.g. subjects, messages, etc.) means that “the ability of the attacker to relate these items does not increase by observing the system”. Here, we provide an example of unlinkability from the sender side, where the message and its intended recipient can not be related to the original sender. We define the *sender-side unlinkability leakage function* $\text{Leak}_{\text{s.unlink}}$ as follows:

$$\text{Leak}_{\text{s.unlink}}(\text{ptr}, H) := \begin{cases} (\mathbf{S}_{\text{ptr}}[H], \mathbf{MR}_{\text{ptr}}[H], \text{Act}_{\text{ptr}}[H]), & \text{if } \text{End}(\text{ptr}, H) \\ \text{Act}_{\text{ptr}}[H], & \text{otherwise} \end{cases} \quad (3)$$

Alternatively, we may define unlinkability from the recipient side via the function

$$\text{Leak}_{\text{r.unlink}}(\text{ptr}, H) := \begin{cases} (\mathbf{MS}_{\text{ptr}}[H], \mathbf{R}_{\text{ptr}}[H], \text{Act}_{\text{ptr}}[H]), & \text{if } \text{End}(\text{ptr}, H) \\ \text{Act}_{\text{ptr}}[H], & \text{otherwise} \end{cases}$$

Pseudonymity. According to [27] “being pseudonymous is the state of using a pseudonym as ID”. To capture pseudonymity, we may slightly abuse definition and consider leakage as a randomized function (or program). Namely, the functionality initially chooses a random permutation π over the set of clients \mathbf{C} , and the pseudonym of each client C_ℓ is $\pi(C_\ell) \in [n]$. We denote by $\pi[H]$ the “pseudonymized history” w.r.t. to π , i.e. in every entry of H we replace C_ℓ by $\pi(C_\ell)$. Clearly, in our infrastructure, the clients remain pseudonymous among the set of clients that are registered to the same SP. We define the *pseudonymity leakage function* as follows:

$$\text{Leak}_{\text{pseudon}}(\text{ptr}, H) := \pi[H], \quad \text{where } \pi \stackrel{\$}{\leftarrow} \{f \mid f : \mathbf{C} \rightarrow [n]\}. \quad (4)$$

Besides anonymity, unlinkability, unobservability and pseudonymity defined in [27], other meaningful notions of privacy can be formally expressed in our framework. We present two such notions below.

Weak anonymity. We define *weak anonymity*, as the privacy notion where the number of messages that a client sends or receives and her fetching activity is leaked. In this weaker notion, the anonymity set for a sender (resp. recipient) consists of the subset of senders (resp. recipients) that are associated with the same number of pending messages. In addition, now the leakage for sender anonymity set is gradually released according to the protocol scheduling, whereas the recipient anonymity set still is leaked “per slot”. The *weak anonymity leakage function*, $\text{Leak}_{\text{w.anon}}$, is defined via the sender and recipient multisets as follows:

$$\text{Leak}_{\text{w.anon}}(\text{ptr}, H) := \begin{cases} (\llbracket \mathbf{S}_{\text{ptr}} \rrbracket[H], \llbracket \mathbf{R}_{\text{ptr}} \rrbracket[H], \mathbf{F}_{\text{ptr}}[H], \text{Act}_{\text{ptr}}[H]), & \text{if } \text{End}(\text{ptr}, H) \\ (\llbracket \mathbf{S}_{\text{ptr}} \rrbracket[H], \mathbf{F}_{\text{ptr}}[H], \text{Act}_{\text{ptr}}[H]), & \text{otherwise} \end{cases} \quad (5)$$

Remark 3. Even though not a very strong privacy notion, weak anonymity supports a reasonable level of privacy for email realizations that aim at a manageable overhead and practical use. Indeed, observe that if we can not tolerate to blow up the ecosystem’s complexity by requiring some form of cover traffic (which is a plausible requirement in practical scenarios), then a global adversary monitoring the client-SP channel can easily infer the number of sent/received messages over this channel. Moreover, one may informally argue that in case the email users do not vary significantly in terms of their sending and fetching activity (or at least they can be grouped into large enough sets of similar activity), weak anonymity and standard anonymity are not far. In Section 6, we present an efficient weakly anonymous email construction based on parallel mixing [18, 19].

End-to-end encryption. The standard notion of *end-to-end encryption*, now applied in many internet applications (e.g., Signal, WhatsApp, Viber, Facebook Messenger, Skype), suggests context hiding of M in the communication of the

end users (up to the message length $|M|$), in our case the sender and the recipient. Hence, we define the *end-to-end leakage function* Leak_{e2e} as shown below.

$$\text{Leak}_{\text{e2e}} := \left(\text{Act}_{\text{ptr}}[H], \left\{ (C_s @ \text{SP}_i, |M|, C_r @ \text{SP}_j) \mid \exists \text{ptr}' \leq \text{ptr} : \right. \right. \\ \left. \left. (\text{ptr}', (\text{sid}, *, \text{SEND}, \langle C_s @ \text{SP}_i, M, C_r @ \text{SP}_j \rangle), \text{'pending'}) \in H \right\} \right). \quad (6)$$

Relation between privacy notions. Observe that the relation between two privacy notions can be deduced via their corresponding leakage functions. Namely, if for every (ptr, H) a PPT adversary given the output of leakage function $\text{Leak}_1(\text{ptr}, H)$ can derive the output of some other leakage functions $\text{Leak}_2(\text{ptr}, H)$, then $\text{Leak}_2(\cdot, \cdot)$ refers to a stronger notion of privacy than $\text{Leak}_1(\cdot, \cdot)$. In Fig. 3, given the definitions of $\text{Leak}_{\text{unob}}$, $\text{Leak}_{\text{anon}}$, $\text{Leak}_{\text{s.unlink}}/\text{Leak}_{\text{r.unlink}}$, $\text{Leak}_{\text{w.anon}}$, Leak_{e2e} above we relate the respective notions in an intuitively consistent way.

$$\begin{array}{ccccc} \text{Unobservability} & \implies & \text{Anonymity} & \implies & \text{Sender/Recipient-side Unlinkability} \\ & & \downarrow & & \\ & & \text{Weak anonymity} & \implies & \text{E2E encryption} \end{array}$$

Fig. 3. Relations between privacy notions. By $A \implies B$, we denote that notion A is stronger than notion B .

Remark 4. We observe that pseudonymity can not be compared to any of the notions in Fig. 3. Indeed, even for the stronger notion of unobservability, having the set of active addresses is not enough information to derive the pseudonyms. Conversely, having the entire email activity pseudonymized, is not enough information to derive the active clients' real identities. In addition, we can combine pseudonymity with some other privacy notion and result in a new 'pseudonymized' version of the latter (e.g. pseudonymous unobservability/anonymity/etc.). It is easy to see that the new notions can also be expressed via suitable (randomized) leakage functions, by applying a random permutation on the clients' identities and then define leakage as in the original corresponding leakage function, up to this permutation. E.g., for $\pi \xleftarrow{\$} \{f \mid f : \mathbf{C} \rightarrow \mathbf{C}\}$, "pseudonymized unobservability" could be expressed via the leakage function

$$\text{Leak}_{\text{ps.unob}}(\text{ptr}, H) := \left\{ \pi(C_\ell) @ \text{SP}_i \mid C_\ell @ \text{SP}_i \in \text{Act}_{\text{ptr}}[H] \right\}.$$

Remark 5. As our E2E leakage does not cover fetch information, strictly speaking the implication from Weak anonymity to E2E encryption only holds if the fetch behavior is either known in advance (e.g. because of the system specification) or irrelevant. One could also opt to add the additional leakage to the E2E definition, but we believe there is little practical value in doing so.

5 An email ecosystem with optimal privacy

We present an email ecosystem, denoted by \mathbb{E}_{comp} , that achieves privacy at an optimal level at the cost of high (quadratic) communication complexity. Specifically, in each time slot all SPs in \mathbb{E}_{comp} communicate with *complete connectivity* and always pad the right amount of dummy traffic, so that the activity of their registered clients is unobservable by a third party, leaking nothing more than that they are online (logged in). In addition, end-to-end communication between the clients is done via encryption layers by utilizing a public key encryption scheme $\text{PKE} = (\text{KeyGen}, \text{Enc}, \text{Dec})$. The encryption layers are structured according to the network route

$$\text{Sender} \longrightarrow \text{Sender's SP} \longrightarrow \text{Receiver's SP} \longrightarrow \text{Receiver}$$

To support unobservability, the online clients who do not send an actual message during some round provide their SPs with a dummy ciphertext.

Even though certainly impractical, \mathbb{E}_{comp} sets a “golden standard” of privacy according to the discussion in Remark 2 that efficient constructions refer to in order to balance the privacy vs. efficiency trade off. **Description of \mathbb{E}_{comp} .** The email ecosystem \mathbb{E}_{comp} operates under a known delay bound Δ_{net} . Throughout the description of \mathbb{E}_{comp} , we assume that the following simplifications: (a) all ciphertexts are of the same length. By $\text{Enc}_{[P]}(M)$, we denote the encryption of M under P 's public key, and (b) all computations require one time slot.⁵

The phases of \mathbb{E}_{comp} are as follows:

■ Initialization:

- On input $(\text{sid}, \text{INIT})$, a service provider SP_i that is not yet initialized, runs $\text{KeyGen}(1^\lambda)$ to generate a private and a public key pair $(\text{sk}_{\text{SP}_i}, \text{pk}_{\text{SP}_i})$. Then, it initializes its list of setup entities, denoted by $L_{\text{setup}}^{\text{SP}_i}$, as the pair $(\text{pk}_{\text{SP}_i}, \text{SP}_i)$, implying that at first SP_i is only aware of itself. In addition, SP_i initializes its list of valid addresses, denoted by Ad_{SP_i} , as empty. Finally, it broadcasts the message $(\text{sid}, \text{CHANNEL}, (\text{setup}, \text{pk}_{\text{SP}_i}), \text{SP}_j)$ to $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$ for every $j \in [N] \setminus \{i\}$, so that all other SPs receive its public key.
- Upon receiving $(\text{sid}, (\text{setup}, \text{pk}_{\text{SP}_j}, \text{SP}_j))$ from $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$, SP_i adds $((\text{pk}_{\text{SP}_j}, \text{SP}_j))$ to $L_{\text{setup}}^{\text{SP}_i}$. When $L_{\text{setup}}^{\text{SP}_i}$ contains all SPs, the SP_i sets its status to ‘execute’, and only then it processes messages of the **Execution** phase described below.

■ Execution:

Registration:

- On input $(\text{sid}, \text{REGISTER}, @\text{SP}_i)$, if C_ℓ is not registered to any SP and has not yet sent a message $(\text{sid}, \text{ADVANCE_CLOCK}, C_\ell)$, then:
 1. C_ℓ sends the message $(\text{sid}, \text{ADVANCE_CLOCK}, C_\ell)$ to $\mathcal{G}_{\text{clock}}$.

⁵ As it will become clear by the ecosystem’s description, the above simplifications do not harm generality essentially. Namely, (a) can be reached via padding, while (b) leads to similar analysis as requiring a computational time upper bound.

2. Upon receiving $(\text{sid}, \text{ADVANCE_ACK}, C_\ell)$ from $\mathcal{G}_{\text{clock}}$, C_ℓ runs $\text{KeyGen}(1^\lambda)$ to generate a private and a public key pair $(\text{sk}_\ell, \text{pk}_\ell)$. It also initializes her list of setup entities, L_{setup}^ℓ as the pair (pk_ℓ, C_ℓ) , and her list of valid addresses, \mathbf{Ad}_ℓ as empty. Then, she sends the message $(\text{sid}, \text{CHANNEL}, (\text{register}, \text{pk}_\ell, \text{SP}_i))$ to $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$.
3. Upon receiving $(\text{sid}, (\text{register}, \text{pk}_\ell), C_\ell)$ from $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$, SP_i checks that $(C_\ell, \cdot) \notin L_{\text{setup}}^{\text{SP}_i}$ and that pk_ℓ is a valid public key, and if so, then it adds (pk_ℓ, C_ℓ) to $L_{\text{setup}}^{\text{SP}_i}$ and $C_\ell @ \text{SP}_i$ to $\mathbf{Ad}_{\text{SP}_i}$. Next, it updates other SPs and its registered clients by broadcasting the message $(\text{sid}, \text{CHANNEL}, (\text{setup}, \text{pk}_{C_\ell}, C_\ell), P)$ to $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$ for every $P \in (\mathbf{SP} \setminus \{\text{SP}_i\}) \cup \mathbf{C}_i$. It also sends the message $(\text{sid}, \text{CHANNEL}, (\text{setup}, \{\text{pk}_P, P\}_{P \in L_{\text{setup}}^{\text{SP}_i}}, \mathbf{Ad}_{\text{SP}_i}), C_\ell)$, updating C_ℓ with all the valid public keys and addresses it knows so far. Finally, it initializes the inbox $\text{Inbox}[C_\ell @ \text{SP}_i]$ of C_ℓ .
4. Upon receiving $(\text{sid}, (\text{setup}, \text{pk}_\ell, C_\ell), \text{SP}_i)$ from $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$, SP_j checks that $(C_\ell, \cdot) \notin L_{\text{setup}}^{\text{SP}_j}$ and that pk_ℓ is a valid public key, and if so, then it adds (pk_ℓ, C_ℓ) to $L_{\text{setup}}^{\text{SP}_j}$ and $C_\ell @ \text{SP}_i$ to $\mathbf{Ad}_{\text{SP}_j}$. It also adds it adds C_ℓ to its set of active users, denoted by $L_{\text{act}}^{\text{SP}_j}$ and initialized as empty. Next, it updates its registered clients by broadcasting the message $(\text{sid}, \text{CHANNEL}, (\text{setup}, \text{pk}_{C_\ell}, C_\ell), C)$ to $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$ for every $C \in \mathbf{C}_j$.
5. Upon receiving $(\text{sid}, (\text{setup}, \{\text{pk}_P, P\}_{P \in L_{\text{setup}}^{\text{SP}_i}}, \mathbf{Ad}_{\text{SP}_i}), \text{SP}_i)$ from $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$, the client C_ℓ , newly registered to SP_i , checks that all public keys are valid. If the check is successful, then C_ℓ adds $\{\text{pk}_P, P\}_{P \in L_{\text{setup}}^{\text{SP}_i}}$ to L_{setup}^i and sets $\mathbf{Ad}_\ell \leftarrow \mathbf{Ad}_{\text{SP}_i}$. Thus, from this point, C_ℓ is aware of the public information of all SPs and all registered clients up to now. In addition, it sets its status as logged in to SP_i .
6. Upon receiving $(\text{sid}, (\text{setup}, \text{pk}_{C_t}, \text{SP}_j), \text{SP}_i)$ from $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$, client C_ℓ (now already registered to SP_i) checks the validity of pk_{C_t} , and if so, then she adds (pk_t, C_t) to L_{setup}^ℓ and $C_t @ \text{SP}_j$ to \mathbf{Ad}_ℓ .

Log in:

- On input $(\text{sid}, \text{ACTIVE}, @\text{SP}_i)$, if C_ℓ is not logged in, $C_\ell @ \text{SP}_i$ is her valid address, and has not yet sent a message $(\text{sid}, \text{ADVANCE_CLOCK}, C_\ell)$, then:
 1. C_ℓ sends the message $(\text{sid}, \text{ADVANCE_CLOCK}, C_\ell)$ to $\mathcal{G}_{\text{clock}}$.
 2. Upon receiving $(\text{sid}, \text{ADVANCE_ACK}, C_\ell)$ from $\mathcal{G}_{\text{clock}}$, C_ℓ “logs in” by sending $(\text{sid}, \text{CHANNEL}, \text{Enc}_{[\text{SP}_i]}(\text{ACTIVE}), \text{SP}_i)$ to $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$.
 3. Upon receiving $(\text{sid}, \text{Enc}_{[\text{SP}_i]}(\text{ACTIVE}), C_\ell)$ from $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$, SP_i decrypts as $(\text{sid}, \text{ACTIVE}, C_\ell)$ and checks that $C_\ell @ \text{SP}_i \in \mathbf{Ad}_{\text{SP}_i}$. If so, then it adds C_ℓ to $L_{\text{act}}^{\text{SP}_i}$.

Log out:

- On input $(\text{sid}, \text{INACTIVE}, @\text{SP}_i)$, if C_ℓ is logged in, $C_\ell @ \text{SP}_i$ is her valid address, and has not yet sent a message $(\text{sid}, \text{ADVANCE_CLOCK}, C_\ell)$, then:
 1. C_ℓ sends the message $(\text{sid}, \text{ADVANCE_CLOCK}, C_\ell)$ to $\mathcal{G}_{\text{clock}}$.
 2. Upon receiving $(\text{sid}, \text{ADVANCE_ACK}, C_\ell)$ from $\mathcal{G}_{\text{clock}}$, C_ℓ “logs out” by sending $(\text{sid}, \text{CHANNEL}, \text{Enc}_{[\text{SP}_i]}(\text{INACTIVE}), \text{SP}_i)$ to $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$.

3. Upon receiving $(\text{sid}, \text{Enc}_{[\text{SP}_i]}(\text{INACTIVE}), C_\ell)$ from $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$, SP_i decrypts as $(\text{sid}, \text{INACTIVE}, C_\ell)$ and checks that $C_\ell @ \text{SP}_i \in \mathbf{Ad}_{\text{SP}_i}$. If so, then it removes C_ℓ from $L_{\text{act}}^{\text{SP}_i}$.

Send:

- On input $(\text{sid}, \text{SEND}, \langle C_s @ \text{SP}_i, M, C_r @ \text{SP}_j \rangle)$, if C_s is logged in to SP_i and has not yet sent a message $(\text{sid}, \text{ADVANCE_CLOCK}, C_s)$, then:
 1. C_s sends the message $(\text{sid}, \text{ADVANCE_CLOCK}, C_s)$ to $\mathcal{G}_{\text{clock}}$.
 2. Upon receiving $(\text{sid}, \text{ADVANCE_ACK}, C_s)$ from $\mathcal{G}_{\text{clock}}$, C_s encrypts the message M into layers and provides $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$ with the layered encryption

$$(\text{sid}, \text{CHANNEL}, \text{Enc}_{[\text{SP}_i]}(\text{Enc}_{[\text{SP}_j]}(C_r @ \text{SP}_j, \text{Enc}_{[C_r]}(M))), \text{SP}_i)$$

3. Upon receiving $(\text{sid}, \text{Enc}_{[\text{SP}_i]}(\text{Enc}_{[\text{SP}_j]}(C_r @ \text{SP}_j, \text{Enc}_{[C_r]}(M))), C_s)$ from $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$, SP_i checks that $C_s @ \text{SP}_i \in \mathbf{Ad}_{\text{SP}_i}$. If so, then it decrypts the first layer with sk_{SP_i} and adds $(\text{sid}, C_s @ \text{SP}_i, \text{Enc}_{[\text{SP}_j]}(C_r @ \text{SP}_j, \text{Enc}_{[C_r]}(M)))$ to its set of messages pending to be sent, denoted by $L_{\text{send}}^{\text{SP}_i}$ and initialized as empty.

Fetch:

- On input $(\text{sid}, \text{FETCH}, C_r @ \text{SP}_j)$, if C_r is logged in to SP_j and has not yet sent a message $(\text{sid}, \text{ADVANCE_CLOCK}, C_r)$:
 1. C_r sends the message $(\text{sid}, \text{ADVANCE_CLOCK}, C_r)$ to $\mathcal{G}_{\text{clock}}$.
 2. Upon receiving $(\text{sid}, \text{ADVANCE_ACK}, C_r)$ from $\mathcal{G}_{\text{clock}}$, C_r sends the message $(\text{sid}, \text{CHANNEL}, \text{Enc}_{[\text{SP}_j]}(\text{FETCH}), \text{SP}_j)$ to $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$.
 3. Upon receiving $(\text{sid}, \text{Enc}_{[\text{SP}_j]}(\text{FETCH}), C_r)$ from $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$, C_r checks that $C_r @ \text{SP}_j \in \mathbf{Ad}_{\text{SP}_j}$. If so, then she decrypts and adds $\text{Inbox}[C_r @ \text{SP}_j]$ to her set of inboxes which messages are pending to be pushed, denoted by $L_{\text{push}}^{\text{SP}_j}$.
 4. Upon receiving $(\text{sid}, E_{r,1}, \dots, E_{r,n}, \text{SP}_j)$ from $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$ (see below), if C_r is registered to SP_j and has sent a $(\text{sid}, \text{FETCH}, C_r @ \text{SP}_j)$ request, then she decrypts all ciphertexts and stores the ones that are not dummy, i.e. they correspond to actual mail messages with her as recipient. Otherwise, she discards $(\text{sid}, E_{r,1}, \dots, E_{r,n}, \text{SP}_j)$.

Clock reading:

- On input $(\text{sid}, \text{READ_CLOCK})$, the entity $P \in \mathbf{C} \cup \mathbf{SP}$ sends the message $(\text{sid}, \text{READ_CLOCK})$ to $\mathcal{G}_{\text{clock}}$. Upon receiving $(\text{sid}, \text{READ_CLOCK}, \text{Cl})$ from $\mathcal{G}_{\text{clock}}$, P stores Cl as its local time.

Clock advance (for clients):

- On input $(\text{sid}, \text{ADVANCE_CLOCK})$, if the client C_ℓ is logged in to SP_i and has not yet sent a message $(\text{sid}, \text{ADVANCE_CLOCK}, \text{SP}_i)$, then she executes the following steps:
 1. C_ℓ sends the message $(\text{sid}, \text{ADVANCE_CLOCK}, C_\ell)$ to $\mathcal{G}_{\text{clock}}$.
 2. Upon receiving $(\text{sid}, \text{ADVANCE_ACK}, C_\ell)$ from $\mathcal{G}_{\text{clock}}$, then she sends a dummy message $(\text{sid}, \text{Enc}_{[\text{SP}_i]}(\text{null}))$ to SP_i via $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$ (in turn, SP_i will discard the received null upon decryption).

Clock advance (for SPs):

- On input $(\text{sid}, \text{ADVANCE_CLOCK})$, if SP_i has not yet sent a message $(\text{sid}, \text{ADVANCE_CLOCK}, \text{SP}_i)$, then it executes the following steps:
 1. SP_i sends the message $(\text{sid}, \text{ADVANCE_CLOCK}, \text{SP}_i)$ to $\mathcal{G}_{\text{clock}}$.
 2. Upon receiving $(\text{sid}, \text{ADVANCE_ACK}, \text{SP}_i)$ from $\mathcal{G}_{\text{clock}}$, for every address $C_s @ \text{SP}_i \in \mathbf{Ad}_{\text{SP}_i}$:
 - If there is a message $(\text{sid}, C_s @ \text{SP}_i, \text{Enc}_{[\text{SP}_j]}(C_r @ \text{SP}_j, \text{Enc}_{[C_r]}(M)))$ in $L_{\text{send}}^{\text{SP}_i}$, then SP_i broadcasts $(\text{sid}, \text{Enc}_{[\text{SP}_j]}(C_r @ \text{SP}_j, \text{Enc}_{[C_r]}(M)))$ to all SPs via $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$, and removes the message from $L_{\text{send}}^{\text{SP}_i}$.
 - If there is no such message for $C_s @ \text{SP}_i$ but $C_s \in L_{\text{act}}^i$, then SP_i broadcasts a dummy message $(\text{sid}, \text{Enc}_{[\text{SP}_j]}(\text{null}))$ under its own key.
 3. Upon receiving a message $(\text{sid}, \tilde{E}, \text{SP}_i)$ from $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$, SP_j checks whether \tilde{E} is a ciphertext under its public key that decrypts as a pair of a valid address $C_r @ \text{SP}_j$ along with an (encrypted) message E . If so, then it adds E to $\text{Inbox}[C_r @ \text{SP}_j]$.
 4. When $L_{\text{fin},k}^{\text{SP}_j}$ contains all SPs, then for every address $C_r @ \text{SP}_j$:
 - If $\text{Inbox}[C_r @ \text{SP}_j] \in L_{\text{push}}^{\text{SP}_j}$, then SP_j forwards all messages $E_{r,1}, \dots, E_{r,n_r}$ in $\text{Inbox}[C_r @ \text{SP}_j]$ to C_r along with $n - n_r$ dummy ciphertexts under C_r 's public key, empties $\text{Inbox}[C_r @ \text{SP}_j]$ and removes it from $L_{\text{push}}^{\text{SP}_j}$.
 - If $\text{Inbox}[C_r @ \text{SP}_j] \notin L_{\text{push}}^{\text{SP}_j}$ but $C_r \in L_{\text{act}}^{\text{SP}_j}$, then SP_j forwards n dummy encryptions of ‘null’ to C_r , under her public key.
- Thus, in any case, if C_r is active, then SP_j sends a message of the form $(\text{sid}, E_{r,1}, \dots, E_{r,n})$ to C_r via $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$.

Privacy of \mathbb{E}_{comp} . To prove the privacy of \mathbb{E}_{comp} , we require that the underlying public key encryption scheme $\text{PKE} = (\text{KeyGen}, \text{Enc}, \text{Dec})$ satisfies m-IND-CPA, as specified in Section 2.2. In the following theorem, we prove that \mathbb{E}_{comp} only leaks the ‘‘activity bit’’ of the clients formally expressed by the leakage function $\text{Leak}_{\text{unob}}(\cdot, \cdot)$ defined in Eq. (1).

Theorem 1. *Let \mathbb{E}_{comp} with clients $\mathbf{C} = \{C_1, \dots, C_n\}$ and service providers $\mathbf{SP} = \text{SP}_1, \dots, \text{SP}_N$ be implemented over the PKE scheme $\text{PKE} = (\text{KeyGen}, \text{Enc}, \text{Dec})$ that achieves m-IND-CPA security with error $\epsilon(\lambda)$. Then, $\mathbb{E}_{\text{comp}}^{\mathcal{G}_{\text{clock}}, \mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}}$ achieves computational $2(n + N)\epsilon(\lambda)$ -privacy for message delay Δ_{net} with respect to the unobservability leakage function defined below*

$$\text{Leak}_{\text{unob}}(\text{ptr}, H) := \text{Act}_{\text{ptr}}[H] .$$

Proof. Let \mathcal{A} be a global passive PPT adversary against $\mathbb{E}_{\text{comp}}^{\mathcal{G}_{\text{clock}}, \mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}}$. We begin by constructing a simulator Sim for \mathcal{A} as shown below.

Constructing a simulator for \mathcal{A} . The ideal adversary Sim for \mathcal{A} that for any environment \mathcal{Z} , simulates an execution of \mathbb{E}_{comp} as follows:

Simulating interaction between \mathcal{Z} and \mathcal{A} .

- Upon receiving a message (sid, M) from \mathcal{Z} , it forwards (sid, M) to \mathcal{A} playing the role of a simulated environment.
- Upon receiving a message (sid, M) from \mathcal{A} intended for the environment, it forwards (sid, M) to \mathcal{Z} .

Achieving synchronicity.

- Upon receiving any message from $\mathcal{F}_{\text{priv}}^{\text{Leak}_{\text{unob}}, \Delta_{\text{net}}}(\mathbf{P})$, Sim sends the message $(\text{sid}, \text{READ_CLOCK})$ to $\mathcal{G}_{\text{clock}}$. Upon receiving $(\text{sid}, \text{READ_CLOCK}, \text{Cl})$ from $\mathcal{G}_{\text{clock}}$, it stores Cl as the global time of the real-world simulation. This way, Sim simulates an execution where the simulated entities are synchronized with respective actual ones in the ideal-world.

Simulating real-world message delivery.

- Upon receiving a leakage message of the form $(\text{sid}, (\text{ptr}, M))$ (possibly $M = \perp$) from $\mathcal{F}_{\text{priv}}^{\text{Leak}_{\text{unob}}, \Delta_{\text{net}}}(\mathbf{P})$, Sim knows that this message refers to some command (register/active/inactive/send/fetch) that in the real-world protocol is realized via communication between a client and her SP. Since in the simulation Sim also plays the role of $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$ in the eyes of \mathcal{A} , it must be consistent with the bounded delays (up to Δ_{net}) that \mathcal{A} imposes on message communication. To achieve this consistency, Sim keeps record of the simulated message \tilde{M} that sends to the simulated $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$ and is associated with ptr . Whenever the message delivery of \tilde{M} is allowed, either by \mathcal{A} or automatically when Δ_{net} delay has passed, Sim sends the message $(\text{sid}, \text{ALLOW_EXEC}, \text{ptr})$ to $\mathcal{F}_{\text{priv}}^{\text{Leak}_{\text{unob}}, \Delta_{\text{net}}}(\mathbf{P})$.

Simulating Initialization.

- Upon receiving $(\text{sid}, \text{INIT}, \text{SP}_i)$ from $\mathcal{F}_{\text{priv}}^{\text{Leak}_{\text{unob}}, \Delta_{\text{net}}}(\mathbf{P})$, it runs $\text{Gen}(1^\lambda)$ on behalf of SP_i to generate a pair of a private and a public key pair $(\text{sk}_{\text{SP}_i}, \text{pk}_{\text{SP}_i})$. Then, it broadcasts the message $(\text{sid}, \text{CHANNEL}, (\text{setup}, \text{pk}_{\text{SP}_i}), \text{SP}_j)$ to every $j \in [N] \setminus \{i\}$, also simulating the role of $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$. Observe that since \mathcal{A} is global and passive, the execution will always initiate upon \mathcal{Z} 's request. Then, Sim sends the message $(\text{sid}, \text{ALLOW_INIT}, \text{SP}_i)$ to $\mathcal{F}_{\text{priv}}^{\text{Leak}_{\text{unob}}, \Delta_{\text{net}}}(\mathbf{P})$.
- Upon receiving $(\text{sid}, \text{ready})$ from $\mathcal{F}_{\text{priv}}^{\text{Leak}_{\text{unob}}, \Delta_{\text{net}}}(\mathbf{P})$, if all simulated SPs have initialized by generating and broadcasting their keys, then it sends $(\text{sid}, \text{EXECUTE})$ to $\mathcal{F}_{\text{priv}}^{\text{Leak}_{\text{unob}}, \Delta_{\text{net}}}(\mathbf{P})$. Otherwise, it aborts simulation.

Simulating Execution.

Whenever the environment sends a register/active/inactive/send/fetch/clock advance command to a dummy party P that forwards it to $\mathcal{F}_{\text{priv}}^{\text{Leak}_{\text{unob}}, \Delta_{\text{net}}}(\mathbf{P})$, Sim obtains (i) an $(\text{sid}, \text{ADVANCE_CLOCK}, P)$ notification from $\mathcal{G}_{\text{clock}}$, and (ii) the leakage of the form $(\text{sid}, \text{ptr}, \text{Act}_{\text{ptr}}[H])$ from $\mathcal{F}_{\text{priv}}^{\text{Leak}_{\text{unob}}, \Delta_{\text{net}}}(\mathbf{P})$. Namely, Sim obtains the sequence of clock advances and the transcript of activations/deactivations. We describe how using this information, Sim simulates execution:

- Upon receiving $(\text{sid}, \text{ADVANCE_CLOCK}, C_\ell)$ and $(\text{sid}, \text{ptr}, \text{Act}_{\text{ptr}}[H])$, then:
 - Playing the role of the global clock, Sim sends a simulated notification $(\text{sid}, \text{ADVANCE_CLOCK}, C_\ell)$ to \mathcal{A} .
 - If $C_\ell @ \text{SP}_i$ is in $\text{Act}_{\text{ptr}}[H]$ and $(\text{sid}, \text{ptr}, \text{Act}_{\text{ptr}}[H])$ is the first entry that $C_\ell @ \text{SP}_i$ is activated, then Sim deduces that this refers to a registration command (Recall that for simplicity we included the pending registration commands in the set of active addresses). In this case, Sim runs the registration protocol between C_ℓ and SP_i exactly as in the description of \mathbb{E}_{comp} , except that it replaces the ciphertext contents with ‘null’ messages. When $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}$ delivers the message, Sim sends the message $(\text{sid}, \text{ALLOW_EXEC}, \text{ptr})$ to $\mathcal{F}_{\text{priv}}^{\text{Leak}_{\text{unob}}, \Delta_{\text{net}}}(\mathbf{P})$.
 - If $C_\ell @ \text{SP}_i$ is in $\text{Act}_{\text{ptr}}[H]$ and is registered but not yet logged in, then Sim deduces that this refers to an active or a clock advance command. In either of these cases, Sim simulates execution by sending a dummy ciphertext $(\text{sid}, \text{CHANNEL}, \text{Enc}_{[\text{SP}_i]}(\text{null}), \text{SP}_i)$ to the simulated $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}$. When $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}$ delivers the message, Sim sends the message $(\text{sid}, \text{ALLOW_EXEC}, \text{ptr})$ to $\mathcal{F}_{\text{priv}}^{\text{Leak}_{\text{unob}}, \Delta_{\text{net}}}(\mathbf{P})$.
 - If $C_\ell @ \text{SP}_i$ is in $\text{Act}_{\text{ptr}}[H]$ and is registered and already logged in, then Sim deduces that this refers to either a inactive, send, fetch or a clock advance command. In either of these cases, Sim simulates execution by sending a dummy ciphertext $(\text{sid}, \text{CHANNEL}, \text{Enc}_{[\text{SP}_i]}(\text{null}), \text{SP}_i)$ as above.
 - If $C_\ell @ \text{SP}_i$ is not in $\text{Act}_{\text{ptr}}[H]$, then Sim deduces that $C_\ell @ \text{SP}_i$ is inactive and takes no further action.
- Upon receiving $(\text{sid}, \text{ADVANCE_CLOCK}, \text{SP}_i)$ and $(\text{sid}, \text{ptr}, \text{Act}_{\text{ptr}}[H])$:
 - Playing the role of the global clock, Sim sends a simulated notification $(\text{sid}, \text{ADVANCE_CLOCK}, \text{SP}_i)$ to \mathcal{A} .
 - For every address $C_s @ \text{SP}_i \in \mathbf{Ad}_{\text{SP}_i}$, it broadcasts a dummy message $(\text{sid}, \text{CHANNEL}, \text{Enc}_{[\text{SP}_i]}(\text{null}), \text{SP}_i)$ to all other SPs. Then, it sends the message $(\text{sid}, \text{ALLOW_EXEC}, \text{ptr})$ to $\mathcal{F}_{\text{priv}}^{\text{Leak}_{\text{unob}}, \Delta_{\text{net}}}(\mathbf{P})$.

Reducing privacy to m-IND-CPA security. We prove the privacy of $\mathbb{E}_{\text{comp}}^{\mathcal{G}_{\text{clock}}, \mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}}$ via a reduction to the m-IND-CPA security with error ϵ of the underlying public key encryption scheme $\text{PKE} = (\text{KeyGen}, \text{Enc}, \text{Dec})$, which is assumed in the theorem’s statement. Our reduction works as follows: Let \mathcal{A} be a real-world adversary and \mathcal{Z} be an environment. First, we order the clients and servers as parties P_1, \dots, P_{n+N} . Then, we construct a sequence of “hybrid” m-IND-CPA adversaries $\mathcal{B}_1, \dots, \mathcal{B}_{n+N}$, where \mathcal{B}_{j^*} executes the following steps:

1. It receives a public key pk from the m-IND-CPA challenger.
2. It generates the parties P_1, \dots, P_{n+N} and simulates an execution of $\mathbb{E}_{\text{comp}}^{\mathcal{G}_{\text{clock}}, \mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}}$ conducted by \mathcal{Z} and under the presence of \mathcal{A} , also playing the role of $\mathcal{G}_{\text{clock}}, \mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}$. The simulation differs from an actual execution as shown below:
 - (a) Upon initialization of a party P_j : if $P_j \neq P_{j^*}$, then \mathcal{B}_{j^*} honestly generates a fresh key pair $(\text{sk}_j, \text{pk}_j)$. If $P_j = P_{j^*}$, then it sets $\text{pk}_{j^*} := \text{pk}$.

- (b) When a party P_i must send an encrypted message M under the public key of P_j (note it may be the case that $P_i = P_j$) via $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}$:
- If $j < j^*$, then \mathcal{B}_{j^*} sends an encryption of M under pk_j .
 - If $j = j^*$, then it sends a challenge pair $(M_0, M_1) := (\text{null}, M)$ to the m-IND-CPA challenger. Upon receiving a ciphertext $\text{Enc}_{[P_{j^*}]}(M_b)$, where b is the m-IND-CPA challenge bit, it sends $\text{Enc}_{[P_{j^*}]}(M_b)$ to P_{j^*} .
 - If $j > j^*$, then it sends an encryption of null under pk_j .
- (c) Since \mathcal{A} is passive, all parties are honest, thus \mathcal{B}_{j^*} is completely aware of the plaintext-ciphertext correspondence. Therefore, when P_i encrypts M under P_j 's public key to a ciphertext $\text{Enc}_{[P_j]}(M)$, \mathcal{B}_{j^*} proceeds as if P_j had indeed decrypted this ciphertext to M .
3. It returns the output of \mathcal{Z} .

Given the description of \mathcal{B}_{j^*} , $j^* = 1, \dots, n + N$, we make the following observations:

- *The limit case $j^* = 1$* : if $b = 0$, then \mathcal{B}_1 replaces all real-world communication with encryptions of ‘null’, exactly as Sim does in its simulation. Thus, we have that

$$\Pr [\mathcal{B}_1 = 1 \mid b = 0] = \text{EXEC}_{\text{Sim}, \mathcal{Z}, \mathcal{G}_{\text{clock}}}^{\mathcal{F}_{\text{priv}}^{\text{Leak}_{\text{unob}}, \Delta_{\text{net}}}} [\mathbf{P}](\lambda) . \quad (7)$$

- *The hybrid step*: for every $1 \leq j^* < n + N$, the adversaries \mathcal{B}_{j^*} and \mathcal{B}_{j^*+1} have the same behavior regarding the parties P_j , where $j \neq j^*, j^* + 1$. In addition, if the m-IND-CPA challenge bit b is 1, then \mathcal{B}_{j^*} (i) respects the encryptions of P_j^* (hence, of every P_j , for $j \leq j^*$) and (ii) replaces with null any plaintext intended for P_j , for $j \geq j^* + 1$. Observe that this is exactly the behavior of \mathcal{B}_{j^*+1} , if $b = 0$. Therefore, it holds that

$$\Pr [\mathcal{B}_{j^*} = 1 \mid b = 1] = \Pr [\mathcal{B}_{j^*+1} = 1 \mid b = 0] . \quad (8)$$

- *The limit case $j^* = n + N$* : if $b = 1$, then \mathcal{B}_{n+N} executes real-world communication respecting the environments’ instructions and inputs. Thus, we have that

$$\Pr [\mathcal{B}_{n+N} = 1 \mid b = 1] = \text{EXEC}_{\mathcal{A}, \mathcal{Z}, \mathcal{G}_{\text{clock}}}^{\mathbb{E}_{\text{clock}}, \mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}} [\mathbf{P}](\lambda) . \quad (9)$$

Consequently, by Eq. (7) and the m-IND-CPA security of PKE, we have that for every $j^* \in [n + N]$, it holds that

$$\begin{aligned} & \left| \Pr [\mathcal{B}_{j^*} = 1 \mid b = 1] - \Pr [\mathcal{B}_{j^*} = 1 \mid b = 0] \right| = \\ & = \left| \Pr [\mathcal{B}_{j^*} = 1 \mid b = 1] - (1 - \Pr [\mathcal{B}_{j^*} = 0 \mid b = 0]) \right| \leq \\ & \leq \left| 2 \cdot \Pr [(\mathcal{B}_{j^*} = 1) \wedge (b = 1)] + 2 \cdot \Pr [(\mathcal{B}_{j^*} = 0) \wedge (b = 0)] - 1 \right| = \\ & = \left| 2 \cdot \Pr [\mathcal{B}_{j^*}(1^\lambda) \text{ breaks PKE}] - 1 \right| \leq \left| 2 \cdot (1/2 + \epsilon(\lambda)) - 1 \right| = 2\epsilon(\lambda) . \end{aligned} \quad (10)$$

Finally, by Eq. (7), (8),(9), and (10), we get that

$$\left| \text{EXEC}_{\text{Sim}, \mathcal{Z}, \mathcal{G}_{\text{clock}}}^{\mathcal{F}_{\text{priv}}^{\text{Leak}_{\text{unob}}, \Delta_{\text{net}}}}[\mathbf{P}](\lambda) - \text{EXEC}_{\mathcal{A}, \mathcal{Z}, \mathcal{G}_{\text{clock}}}^{\mathbb{E}^{\mathcal{G}_{\text{clock}}, \mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}}}[\mathbf{P}](\lambda) \right| \leq 2(n + N)\epsilon(\lambda)$$

which completes the proof. \square

6 A parallel mix email ecosystem with t strata

We will now describe a design to be used for routing messages between various users, based on parallel mixing [18, 19]. A parallel mix is a design that borrows characteristics from stratified mixes i.e mixes where servers are grouped in sets called *strata*, and routing is restricted so that each stratum except the first only receives messages from the previous one and each stratum except the last only forwards messages to the next (the first and last strata operate as the entry and exit points respectively). In parallel mixing routing is determined by the servers themselves in the interest of symmetry and predictability in performance and security. All t strata consist of σ nodes each. We use $\text{MX}_{i,j}$ to indicate the j -th server in stratum i , and let $\text{MX} = \{\text{MX}_{i,j} | i \leq t, j \leq \sigma\}$. We use $\mathbf{P} = (\mathbf{C} \cup \mathbf{SP} \cup \text{MX})$ to denote the set of all involved parties. We use a set of assumptions similar to those of section 5, specifically: (a) all communication is executed via $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{P})$ as described in Fig. 2; (b) all messages have the same size (i.e messages are padded ahead of time); (c) all computations complete within one unit slot; (d) each client is assigned to exactly one address.

As we assume a passive adversary and no corruptions, we are able to use a simple layering of encryptions instead of a more complex onion scheme. In practice one may wish to use a scheme such as Sphinx [14] or a variant thereof.

■ **Initialization:** Nodes of the same stratum share stratum-specific keying material. In practice, because of the long structure of the mixnet, and the large number of nodes involved, we might have that the same entities will be running multiple servers across different strata. We can thus regain some robustness by excluding some entities from each stratum so that each entity is absent from at least one stratum. Alternatively, we may use per-node keys and allow free routing, at the cost of slower (in terms of rounds) convergence to a random permutation .

- On input $(\text{sid}, \text{INIT})$, a party $P \in \mathbf{P}$ that is not yet initialised, runs $\text{Gen}(1^\lambda)$ to generate a pair of a private and a public key pair $(\text{sk}_P, \text{pk}_P)$. Then, it broadcasts the message $(\text{sid}, (\text{init}, \text{pk}_P), P)$ to all clients and SPs by sending $(\text{sid}, (\text{init}, \text{pk}_P), P')$ to $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{G}[\mathbb{P}])$, for every $P' \in \mathbf{P} \setminus \{P\}$.
- When SP_i has received $(\text{sid}, (\text{init}, \text{pk}_{\text{SP}_j}, \text{SP}_j))$ for every $i \in [N] \setminus \{j\}$, then begins the engagement in the email message exchange with its assigned clients and the other SPs.
- When $\text{MX}_{i,1}$ has received $(\text{sid}, (\text{init}, \text{pk}_S, S))$ for every $\text{MX}_{i,j}, j > 1$, it runs $\text{Gen}(1^\lambda)$ to generate stratum key pair $(\text{sk}_i, \text{pk}_i)$. Then, it broadcasts the message $(\text{sid}, (\text{init}, \text{pk}_i), \text{MX}_{i,1})$ to all parties P' outside stratum i by sending $(\text{sid}, (\text{init}, \text{pk}_i), P')$ to $\mathcal{F}_{\text{auth}}^{\Delta_{\text{net}}}(\mathbf{G}[\mathbf{P}])$. For parties P'' in stratum i it sends $(\text{sid}, (\text{init}, (\text{pk}_i, \text{sk}_i)), P'')$ instead.

■ **Execution:** Our mixnet operates in rounds. A round consists of $t+2$ subrounds, each consisting of $t_{sub} \geq \Delta_{net} + 1$ timeslots. We assume timing information is publicly available. During each subround, messages are only sent during the first timeslot. The remaining timeslots exist to ensure that even delayed messages are delivered before the next subround. To simplify notation we will introduce three functions on the clock value Cl :

Namely, we define (i) $round(Cl) := \lfloor \frac{Cl}{t_{sub}(t+2)} \rfloor$, (ii) $sub(Cl) := \lfloor \frac{Cl}{t_{sub}} \rfloor$, and (iii) $slot(Cl) := Cl \bmod t_{sub}$. Essentially, at clock Cl we are in slot $slot(Cl)$ of subround $sub(Cl)$. We also assume that using the above functions use `READ_CLOCK` to determine the current value of Cl .

Registration is handled as in Section 5. Messages are routed through the mixnet as follows:

- Messages from clients are queued by their SPs until the round begins.
- Once a round begins, in sub-round 0, clients send their messages to the SPs. In sub-round 1, each SP uniformly randomly selects a server in the first stratum to receive each message.
- In the sub-round 2 (3), first-stratum (second) servers tally up their incoming messages and pad them to a multiple of σ . They shuffle them and send $\frac{1}{\sigma}$ of them to each 2nd-stratum (3rd) server. No padding is required afterwards.
- In sub-round i , where $4 \leq i \leq t+1$, the servers of stratum $i-1$ shuffle their received messages and send $\frac{1}{\sigma}$ of them to each server in stratum $i+1$.
- At the end of sub-round $t+2$, the SPs move messages from their input buffers to user inboxes. .

We will now formally describe our system. Note that some inputs will only have effect when given during particular sub-rounds or when given to certain parties (e.g. only Clients). As in the previous section, $Enc_{[X]}(Y)$ denotes the encryption of Y under X 's public key. For brevity, we use $Enc_{[x,y]}(m)$ to denote $Enc_{[x]}(Enc_{[y]}(m))$.

- $C_s \in \mathbf{C}$ On input $(sid, SEND, \langle C_s@SP_i, M, C_r@SP_j \rangle)$, if C_s is not registered with an SP_i and $subround(Cl) = 0$ and $slot(Cl) = 0$, the client sets $reg = round(Cl)$ and runs the registration operation from Section 5.
- $C_s \in \mathbf{C}$ On input $(sid, SEND, \langle C_s@SP_i, M, C_r@SP_j \rangle)$, if C_s is logged in to SP_i , she prepares the message $(sid, Enc_{[SP_i]}(C_s@SP_i, Enc_{[SP_j]}(C_r@SP_j, Enc_{[C_r]}(M))))$ to be sent to SP_i . If, in addition the $sub(Cl)$ and $slot(Cl)$ are both 0 and $round(Cl) > reg$, all prepared messages are sent to SP_i .
- $C_r \in \mathbf{C}$ On input $(sid, FETCH, C_r@SP_j)$, if C_r is logged in to SP_j , it sends the message $(sid, C_r@SP_j, Enc_{[SP_j]}(FETCH))$ to SP_j which, if $C_r@SP_j$ is a valid address, it decrypts and forwards all messages $E_{r,1}, \dots, E_{r,n_r}$ in $Inbox[C_r@SP_j]$ to C_r , and empties $Inbox[C_r@SP_j]$.
- $C_r \in \mathbf{C}$ Upon receiving $(sid, E_{r,1}, \dots, E_{r,n})$ from SP_j and if C_r has sent a $(sid, FETCH, C_r@SP_j)$ request, C_r decrypts all ciphertexts and stores the ones that are not 0, i.e. they correspond to non-dummy mail messages.

- $P \in \mathbf{P}$ On input $(\text{sid}, \text{READ_CLOCK})$, the entity $P \in \mathbf{P}$ sends the message $(\text{sid}, \text{READ_CLOCK})$ to $\mathcal{G}_{\text{clock}}$. Upon receiving $(\text{sid}, \text{READ_CLOCK}, \text{Cl})$ from $\mathcal{G}_{\text{clock}}$, P stores Cl as its local time and forwards the message $(\text{sid}, \text{READ_CLOCK}, \text{Cl})$ to the environment.
- $\text{SP}_i \in \mathbf{SP}$ On input $(\text{sid}, \text{Enc}_{[\text{SP}_i]}(C_s @ \text{SP}_i, \text{Enc}_{[\text{SP}_j]}(C_r @ \text{SP}_j, \text{Enc}_{[C_r]}(M))))$, it checks that $C_s @ \text{SP}_i \in \mathbf{Ad}$ and if so, then it decrypts and adds $(\text{sid}, C_s @ \text{SP}_i, \text{Enc}_{[\text{SP}_j]}(C_r @ \text{SP}_j, \text{Enc}_{[C_r]}(M)))$ to its set of messages pending to be sent, denoted by L_{send}^i .
- $\text{SP}_j \in \mathbf{SP}$ Upon receiving a message $(\text{sid}, \text{Enc}_{[\text{SP}_j]}(\cdot, \cdot))$ from some $\text{MX}_{x,y}$, SP_j checks whether $x = t$, and if the content is a ciphertext under its public key that decrypts as a valid address $C_r @ \text{SP}_j$ along with a ciphertext E . If so, then it adds E to $\text{B}[C_r @ \text{SP}_j]$.
- $\text{MX}_{1,j} \in \mathbf{S}$ On receiving $(\text{sid}, \text{Enc}_{[1,\dots,t]}(C_r @ \text{SP}_j, \text{Enc}_{[C_r]}(M)), X)$, it checks that $X \in \mathbf{SP}$ and if so, it decrypts it and adds $(\text{sid}, \text{Enc}_{[2,\dots,t]}(C_r @ \text{SP}_j, \text{Enc}_{[C_r]}(M)))$ to its set of messages pending to be sent, denoted by L_{send}^i .
- $\text{MX}_{k+1,j} \in \mathbf{S}$ On receiving $(\text{sid}, \text{Enc}_{[k,\dots,t]}(C_r @ \text{SP}_j, \text{Enc}_{[C_r]}(M)), X)$, it checks that $X = \text{MX}_{k,x}$ for some x and if so, it decrypts it and adds $(\text{sid}, \text{Enc}_{[k+2,\dots,t]}(C_r @ \text{SP}_j, \text{Enc}_{[C_r]}(M)))$ to its set of messages pending to be sent, denoted by L_{send}^i . If $k = t - 1$, it instead adds $(\text{sid}, \text{Enc}_{[SP_j]}(C_r @ \text{SP}_j, \text{Enc}_{[C_r]}(M)))$ to the list.
- $P \in \mathbf{P}$ On input $(\text{sid}, \text{ADVANCE_CLOCK})$, the entity $P \in \mathbf{P}$ sends the message $(\text{sid}, \text{ADVANCE_CLOCK})$ to $\mathcal{G}_{\text{clock}}$.
- $\text{SP}_i \in \mathbf{SP}$ On input $(\text{sid}, \text{ADVANCE_CLOCK})$, If $\text{sub}(\text{Cl}) = 1$ and $\text{slot}(\text{Cl}) = 0$, for each message $(\text{sid}, C_s @ \text{SP}_i, \text{Enc}_{[\text{SP}_j]}(C_r @ \text{SP}_j, \text{Enc}_{[C_r]}(M)))$ in L_{send}^i , then SP_i sends $(\text{sid}, \text{SP}_i, \text{Enc}_{[1,\dots,t]}(C_r @ \text{SP}_j, \text{Enc}_{[C_r]}(M)))$ to a randomly selected $\text{MX}_{1,j}$ and removes the message from L_{send}^i . Finally, it sends the message $(\text{sid}, \text{ADVANCE_CLOCK})$ to $\mathcal{G}_{\text{clock}}$.
- $\text{MX}_{k,j} \in \mathbf{S}$ On input $(\text{sid}, \text{ADVANCE_CLOCK})$, If $\text{sub}(\text{Cl}) \neq k + 1$ or $\text{slot}(\text{Cl}) \neq 0$, send the message $(\text{sid}, \text{ADVANCE_CLOCK})$ to $\mathcal{G}_{\text{clock}}$ and return. Otherwise, if $k = 1$ or $k = 2$, $\text{MX}_{k,j}$ pads the list L_{send}^i with $(\text{sid}, \text{Enc}_{[k+1,\dots,t]}(0))$ so that its length is a multiple of σ . The list is then shuffled randomly. For each message $(\text{sid}, \text{Enc}_{[k+1,\dots,t]}(C_r @ \text{SP}_j, \text{Enc}_{[C_r]}(M)))$ in L_{send}^i , then $\text{MX}_{k,j}$ sends $(\text{sid}, \text{MX}_{k,j}, \text{Enc}_{[k+1,\dots,t]}(C_r @ \text{SP}_j, \text{Enc}_{[C_r]}(M)))$ to server $\text{MX}_{k+1,j \bmod \sigma}$, where j is the message's position on the list, and removes the message from L_{send}^i . Finally, it sends the message $(\text{sid}, \text{ADVANCE_CLOCK})$ to $\mathcal{G}_{\text{clock}}$.
- $\text{MX}_{t,j} \in \mathbf{S}$ On input $(\text{sid}, \text{ADVANCE_CLOCK})$, If $\text{sub}(\text{Cl}) = t + 1$ and $\text{slot}(\text{Cl}) = 0$, for each message $(\text{sid}, \text{Enc}_{[\text{SP}_j]}(C_r @ \text{SP}_j, \text{Enc}_{[C_r]}(M)))$ in L_{send}^i , $\text{MX}_{t,j}$ forwards it to SP_j . Finally it sends the message $(\text{sid}, \text{ADVANCE_CLOCK})$ to $\mathcal{G}_{\text{clock}}$.
- $\text{SP}_j \in \mathbf{SP}$ On input $(\text{sid}, \text{ADVANCE_CLOCK})$, If $\text{sub}(\text{Cl}) = t + 2$ and $\text{slot}(\text{Cl}) = 0$, it moves the contents of every buffer $\text{B}[C_r @ \text{SP}_j]$ to the corresponding inbox $\text{Inbox}[C_r @ \text{SP}_j]$. Finally it sends the message $(\text{sid}, \text{ADVANCE_CLOCK})$ to $\mathcal{G}_{\text{clock}}$.

Efficiency & Delivery times The overhead of the padding is an $O\left(\frac{\sigma^2}{m}\right)$ multiplicative increase in the messages sent, where m is the number of messages sent,

which we expect to be low for typical use cases. Disregarding padding messages, the cost to deliver a single email, is $3 + t$ messages compared to 3 in the insecure case (sender to SP_s to SP_r to receiver) or $1 + s \cdot n$ for the “golden standard” solution of Section 5. While in principle this is identical to a cascade (i.e. single server per stratum) solution, in practice a parallel mix requires a larger t value. The load per mix server is $\frac{m}{\sigma}$ messages, compared to m in a cascade.

The encryption overhead depends on the specifics of the cryptosystem. While naive encryption might cause an exponential blow-up, solutions based on hybrid encryption, or onioning solutions such as Sphinx can reduce the overhead to a small linear factor. Delivery latency is also directly proportional to the length of the mixnet. We note that latency can be significantly reduced by pipelining (i.e. allowing messages to be sent at the end of every subround rather than at the end of the first round only), but we opt to describe the base version for clarity.

Security Here, we will show that the system described above is secure under the weak anonymity definition and leakage function $\text{Leak}_{\text{w.anon}}(\text{ptr}, H)$, defined in Eq. (5). For convenience, we will assume that one timeslot maps to one round.

Theorem 2. *The parallel mix of Section 6, using t strata of σ servers to deliver m messages is $m^{1 - \lfloor \frac{t-1}{2} \rfloor \frac{1}{4}} 4^{\lfloor \frac{t-1}{2} \rfloor \frac{1}{2}} \log m^{\lfloor \frac{t-1}{2} \rfloor \frac{1}{4}} + 2|\mathbf{P}|_{\epsilon_E}$ weakly anonymous assuming Enc is ϵ_E m -IND-CPA secure.*

Proof. Due to space considerations, we postpone the proof to the full version of this work and sketch the main strategy to cover the difference between a real and simulated execution. We first utilize the m -IND-CPA security of the encryption and a series of hybrid games to replace message contents with dummies. To complete the proof, we use Theorem 3 (based on [20]) to show that replacing the final routing of messages with a random allocation that respects the leakage (i.e. messages per server) is statistically indistinguishable.

7 The Combinatorics of Parallel Mixing

Many of the works analysing parallel mixing investigate the probability distribution of a single message traversing the network. This is satisfactory for some definitions of anonymity but not for our modelling of a global adversary under universal composability. In our model, the environment determines the sender and receiver of each message, so it is not sufficient to argue that any one message is successfully shuffled (i.e. has a uniformly random exit point from the network).

To illustrate, assume messages are represented by a deck of n playing cards, and further assume that our mixnet operates by simply “cutting” the deck once, in secret (i.e. choosing $k \in \{0..n-1\}$, and placing the first k cards at the bottom of the deck in their original order). It is trivial to simulate drawing a single card from a deck shuffled this way, by sampling a random card. However, once a card has been drawn, subsequent draws are determined by the initial order. The environment knows the initial order because it set it, but the simulator does not, and the simulation fails.

Our approach will be to show that parallel mixing after a number of rounds produces a random permutation on the list of input messages, thus allowing the simulator to produce the list of output messages by sampling a random permutation of the recipients, independent of the senders (which is crucial as it does not know the relation between the two).

We will model parallel mixing as a generalisation of the square lattice shuffle of Håstad [20]. In a square lattice shuffle, $n = m^2$ cards are arranged in an $m \times m$ matrix, and shuffled as follows: in odd rounds each row is shuffled by an independently uniformly random sampled permutation. In even rounds, the same happens to columns. It is simple to check that t iterations of this process map directly to a t -stratum parallel mix with m servers per stratum, each with capacity m : we label odd strata as “rows” and even strata as “columns”, where the i -th server corresponds to the i -th row (column). The mapping is then completed by noting the result of an odd round is that each row randomly contributes one of its elements to each column, and vice-versa for even rounds.

Thus Håstad’s results are applicable to parallel mixing. A second observation is that because parties are assumed honest, we can assign multiple rows or columns to one party without invalidating the bounds. We thus reproduce Theorem 3.6 from [20] and explain how it applies in our construction.

Theorem 3 (Håstad [20], Theorem 3.6). *Let Π_t be the distribution defined by t iterations of lattice shuffling on m objects. Then*

$$\Delta(\Pi_t, U_m) \leq O(m^{1-\lfloor \frac{t-1}{2} \rfloor \frac{1}{4}} \log m^{\lfloor \frac{t-1}{2} \rfloor \frac{1}{2}})$$

Closer examination of the proof, and assuming $m > 81$ enables us to dismiss the big-O and obtain:

$$\Delta(\Pi_t, U_m) \leq m^{1-\lfloor \frac{t-1}{2} \rfloor \frac{1}{4}} 1.5^{\lfloor \frac{t-1}{2} \rfloor} \log m^{\lfloor \frac{t-1}{2} \rfloor \frac{1}{2}}$$

This in turn implies

Corollary 1. *For $m > 10^6$, 31 rounds of lattice shuffling are statistically $\frac{1}{m}$ close to uniform.*

The theorem’s proof also gives us insight in the effect of compromised servers in a stratum: as coupling takes place over 3 iterations (or 2 with the assumption that another honest iteration will follow), we must allow that a single compromised stratum essentially shortens our network by 3 strata at the worst case.

7.1 A brief discussion on convergence speed

The bounds stated above describe a parallel mix with many small servers. One would expect the situation to improve when examining fewer, larger servers. In that direction, we expect a generalization of Håstad’s result to yield a tighter bound. That would be of value as there are few competing designs for random permutation networks suited to anonymous communication [25].

The core of Håstad’s analysis is about the probability of “coupling” two permutations that start out differing by a single transposition, after 2 rounds of shuffling. A first observation is that with “large” servers, the probability that the transposition lies in one server (and thus the coupling is immediate) becomes significant, improving convergence. A second, is that the probability of a missed coupling is inversely proportional to the number of elements per server which again implies improved convergence. We believe that a bound of $m^{1-\lfloor \frac{t-1}{2} \rfloor \frac{1}{2}} 1.5^{\lfloor \frac{t-1}{2} \rfloor} \log m^{\lfloor \frac{t-1}{2} \rfloor \frac{\sigma-1}{\sqrt{\sigma}} \lfloor \frac{t-1}{2} \rfloor}$ is possible, which would approximately halve the rounds required for the bound to reach $\frac{1}{m}$, when σ is small, e.g. 17 rounds for $\sigma = 4$, $m > 300.000$. However, we consider the specifics outside the scope of this work, and leave the question of statistical bounds for parallel mixing open for further research.

References

1. Alexopoulos, N., Kiayias, A., Talviste, R., Zacharias, T.: MCMix: Anonymous messaging via secure multiparty computation. In: USENIX (2017)
2. Angel, S., Setty, S.: Unobservable communication over fully untrusted infrastructure. In: OSDI (2016)
3. Backes, M., Kate, A., Manoharan, P., Meiser, S., Mohammadi, E.: Anoa: A framework for analyzing anonymous communication protocols. In: CSF (2013)
4. Badertscher, C., Maurer, U., Tschudi, D., Zikas, V.: Bitcoin as a transaction ledger: A composable treatment. In: CRYPTO (2017)
5. Camenisch, J., Lysyanskaya, A.: A formal treatment of onion routing. In: CRYPTO. pp. 169–187. Springer (2005)
6. Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: Foundations of Computer Science. IEEE (2001)
7. Canetti, R., Dodis, Y., Pass, R., Walfish, S.: Universally composable security with global setup. In: TCC (2007)
8. Chaum, D.: The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of cryptology* **1**(1), 65–75 (1988)
9. Chaum, D., Das, D., Javani, F., Kate, A., Krasnova, A., de Ruiter, J., Sherman, A.T.: cmix: Mixing with minimal real-time asymmetric cryptographic operations. In: ACNS. pp. 557–578 (2017)
10. Chaum, D.L.: Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM* **24**(2), 84–90 (1981)
11. Corrigan-Gibbs, H., Boneh, D., Mazières, D.: Riposte: An anonymous messaging system handling millions of users. In: Security and Privacy (2015)
12. Corrigan-Gibbs, H., Ford, B.: Dissent: accountable anonymous group messaging. In: CCS. pp. 340–350 (2010)
13. Danezis, G., Dingleline, R., Mathewson, N.: Mixminion: Design of a type III anonymous remailer protocol. In: Security and Privacy. pp. 2–15 (2003)
14. Danezis, G., Goldberg, I.: Sphinx: A compact and provably secure mix format. In: Security and Privacy (2009)
15. Díaz, C., Seys, S., Claessens, J., Preneel, B.: Towards measuring anonymity. In: PETS (2002)
16. Dingleline, R., Mathewson, N., Syverson, P.: Tor: The second-generation onion router. Tech. rep., DTIC Document (2004)

17. Dwork, C.: Differential privacy. In: Automata, Languages and Programming. pp. 1–12 (2006)
18. Golle, P., Juels, A.: Parallel mixing. In: CCS. pp. 220–226. ACM (2004)
19. Goodrich, M.T., Mitzenmacher, M.: Anonymous card shuffling and its applications to parallel mixnets. In: Automata, Languages, and Programming. pp. 549–560. Springer (2012)
20. Håstad, J.: The square lattice shuffle. *Random Structures & Algorithms* **29**(4), 466–474 (2006)
21. Johnson, A., Wacek, C., Jansen, R., Sherr, M., Syverson, P.: Users get routed: Traffic correlation on tor by realistic adversaries. In: CCS. pp. 337–348 (2013)
22. Katz, J., Maurer, U., Tackmann, B., Zikas, V.: Universally composable synchronous computation. In: TCC (2013)
23. Kesdogan, D., Egner, J., Büschkes, R.: Stop-and-go-mixes providing probabilistic anonymity in an open system. In: International Workshop on Information Hiding. pp. 83–98. Springer (1998)
24. Kotzanikolaou, P., Chatzisofofroniou, G., Burmester, M.: Broadcast anonymous routing (BAR): scalable real-time anonymous communication. *Int. J. Inf. Sec.* **16**(3), 313–326 (2017)
25. Kwon, A., Corrigan-Gibbs, H., Devadas, S., Ford, B.: Atom: Horizontally scaling strong anonymity. In: SOSP. ACM (2017)
26. Kwon, A., Lazar, D., Devadas, S., Ford, B.: Riffle: An efficient communication system with strong anonymity. *PoPETS* **2016**(2), 115–134 (2015)
27. Pfizmann, A., Hansen, M.: Anonymity, Unlinkability, Unobservability, Pseudonymity, and Identity management – a consolidated proposal for terminology. Version v0.25 (December 2005)
28. Pfizmann, A., Köhntopp, M.: Anonymity, unobservability, and pseudonymity—a proposal for terminology. In: Designing privacy enhancing technologies. Springer (2001)
29. Piotrowska, A., Hayes, J., Elahi, T., Danezis, G., Meiser, S.: The loopix anonymity system. In: USENIX (2017)
30. Samarati, P., Sweeney, L.: Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression. In: Security and Privacy (1998)
31. Serjantov, A., Danezis, G.: Towards an information theoretic metric for anonymity. In: Privacy Enhancing Technologies. pp. 41–53. Springer (2002)
32. Shmatikov, V., Wang, M.: Timing analysis in low-latency mix networks: Attacks and defenses. In: ESORICS. pp. 18–33 (2006)
33. Syverson, P.F., Goldschlag, D.M., Reed, M.G.: Anonymous connections and onion routing. In: Security and Privacy. pp. 44–54 (1997)
34. Syverson, P.F., Tsudik, G., Reed, M.G., Landwehr, C.E.: Towards an analysis of onion routing security. In: Workshop on Design Issues in Anonymity and Unobservability. pp. 96–114 (2000)
35. Van Den Hooff, J., Lazar, D., Zaharia, M., Zeldovich, N.: Vuvuzela: Scalable private messaging resistant to traffic analysis. In: SOSP. pp. 137–152 (2015)
36. Wikström, D.: A universally composable mix-net. In: TCC. pp. 317–335. Springer (2004)