

Trellis: Robust and Scalable Metadata-private Anonymous Broadcast

Simon Langowski
MIT CSAIL

Sacha Servan-Schreiber
MIT CSAIL

Srinivas Devadas
MIT CSAIL

Abstract—Trellis is a mix-net based anonymous broadcast system with cryptographic security guarantees. Trellis can be used to anonymously publish documents or communicate with other users, all while assuming full network surveillance. In Trellis, users send messages through a set of servers in successive rounds. The servers mix and post the messages to a public bulletin board, hiding which users sent which messages.

Trellis hides all network-level metadata, remains robust to changing network conditions, guarantees availability to honest users, and scales with the number of mix servers. Trellis provides three to five orders of magnitude faster performance and better network robustness compared to Atom, the state-of-the-art anonymous broadcast system with a similar threat model.

In achieving these guarantees, Trellis contributes: (1) a simpler theoretical mixing analysis for a routing mix network constructed with a fraction of malicious servers, (2) anonymous routing tokens for verifiable random paths, and (3) lightweight blame protocols built on top of onion routing to identify and eliminate malicious parties.

We implement and evaluate Trellis in a networked deployment. With 64 servers located across four geographic regions, Trellis achieves a throughput of 220 bits per second with 100,000 users. With 128 servers, Trellis achieves a throughput of 320 bits per second. Trellis’s throughput is only 100 to 1000× slower compared to Tor (which has 6,000 servers and 2M daily users) and is therefore potentially deployable at a smaller “enterprise” scale. Our implementation is open-source.

I. INTRODUCTION

Communication on the internet is prone to large-scale surveillance and censorship [14, 38, 42, 51, 73, 93]. People trying to hide their communication from powerful (e.g., state-sponsored) adversaries rely on anonymity tools such as Tor [32] or I2P [98] to hide their identities [99, 100, 105]. With Tor, users proxy their traffic through a chain of servers which breaks the link between the identity of the user and the traffic destination. With I2P, users proxy traffic through a chain of peer nodes. While these tools serve an important role, and help obfuscate the identity of the user, they can be insufficient against determined adversaries capable of observing large swaths of traffic (e.g., malicious ISPs). Such adversaries (e.g., ISPs or nation states) can easily identify users and the traffic content through metadata such as packet timing, packet

size, and other identifying features—even when all traffic is encrypted [9, 13, 37, 45, 49, 58, 77, 82, 84, 96]. State-of-the-art attacks can deanonymize encrypted Tor traffic with upwards of 90% accuracy by analyzing the encrypted packet traffic [9, 45, 82].

This problem has motivated many systems [2, 19, 20, 35, 36, 60, 62, 78, 103, 109] for **anonymous broadcast**. Anonymous broadcast consists of anonymously posting a message to a public bulletin board and is the most general form of anonymous communication possible; it can be used for file sharing and to instantiate weaker primitives such as anonymous point-to-point communication [27, 61, 65, 66, 86, 92, 101, 104]. Anonymous broadcast provides *sender* anonymity: fully hiding which user sent which message. Most mix-net based systems, in contrast, only provide *relationship* anonymity [27, 61, 65, 66, 86, 92, 101, 104] (hiding who is communicating with whom). Other systems [19, 20, 78, 103, 109] exploit a *non-collusion* assumption (which inhibits scalability) to instantiate anonymous broadcast or just focus on recipient anonymity [6].

As with prior work, Trellis is metadata-private, which guarantees that sender anonymity is preserved in the face of an adversary monitoring the entire network. Trellis also provides horizontal scalability—the ability to increase throughput with the number of (possibly untrusted) servers participating in the network. Horizontal scalability is key to real-world efficiency: Tor and I2P are only capable of handling millions of users [71] on a daily basis thanks to similar scalability guarantees. Moreover, horizontal scalability is important for real-world security: the more users an anonymity system can support, the more “plausible deniability” each user in the system has.

The state-of-the-art system for anonymous broadcast with horizontal scalability is Atom [60]. Unfortunately, Atom suffers from practical barriers to deployment even though it is *theoretically* optimal. (See Section II-B for overview of related work.) For example, Atom achieves a throughput of roughly 0.01 bits per second and is six to eight orders of magnitude slower than practical systems such as Tor [71].

All other systems either focus on weaker anonymity notions (e.g., point-to-point communication [3, 5, 15, 61, 65, 66, 92, 101, 104]) or sacrifice scalability by assuming non-colluding servers [2, 19, 20, 35, 36, 59, 78, 109].

Trellis is designed for *fast* anonymous broadcast with *large messages*. Trellis is inspired by a long line of work [16, 60, 62, 89] which augments mix-nets to instantiate anonymous broadcast. Trellis innovates on these works in three important ways: (1) Trellis decouples the use of expensive cryptography required for system setup from the broadcasting phase, (2)




Fig. 1: Trellis arranges the set of servers into “layers,” (each server appears once per layer) to instantiate a mix-net. Users send messages through a random path which guarantees that all messages are mixed before reaching the public bulletin board, even when a subset of servers are malicious (colored in red).

Trellis is robust to malicious parties and server churn, and (3) Trellis provides message delivery guarantees even when a *majority*¹ of servers and users are malicious (note that performance degrades with a larger proportion of adversarial servers; see Section VIII). With these three key contributions, Trellis brings metadata-private anonymous broadcast a step closer to real-world deployment. However, we note that Trellis is still limited by challenges *outside* of the system design itself, such as requiring that all users remain online and contribute messages to ensure metadata privacy. These limitations will be present in any anonymous communication system design with strong anonymity guarantees [42]. Even so, in Section VIII, we find that Trellis is only 100 to 1000 \times slower compared to systems like Tor and I2P, both of which sacrifice strong anonymity in favor of concrete performance.

Ideas, challenges, and contributions. We start by using a mix-net topology known as a random routing network [60, 62] (see Section II for an overview of mix-nets). In such mix-nets, servers are organized into L layers (all servers appear in each layer). The user selects a random path through the servers and sends a message through the path for mixing. Adding more servers increases the throughput, which satisfies the horizontal scalability requirement. While this simple idea seems promising at first (and also proposed in prior work [60, 62]), there are several practical hurdles associated with it:

- Malicious users or servers might deviate from protocol, select non-random paths, or otherwise attack the system. The difficulty is detecting and blaming the responsible parties while preserving anonymity for honest users.
- Mix-nets that model permutation networks traditionally require that all servers are honest to achieve mixing guarantees. Atom [60] resolves this by emulating honest “servers” using random groups of servers. However, this technique is suboptimal and contributes to slow mixing times in practice.
- Real-world networks consisting of disparate servers are not perfect: server churn and elimination can occur frequently. Classic mix-nets are **not** designed to gracefully handle and adapt to changing network conditions on-the-fly.

In Trellis, we overcome the above challenges as follows.

Preventing deviation from protocol. To ensure availability and message delivery, we develop two new tools: anonymous

routing tokens (Section V-A) and boomerang encryption (Section IV-A). Routing tokens force all users to choose random paths through the mix-net. Boomerang encryption ensures that all servers along a path route messages correctly and enables efficient blaming of malicious servers when required. The combination of these two primitives allows us to port our mix-net to a setting with malicious users and servers actively deviating from protocol *without* compromising on anonymity.

Modeling a mix-net with malicious servers. Kwon [62] develops a novel theoretical mixing analysis for routing mix-nets constructed with some fraction of adversarial servers. In Trellis, we use a routing mix-net topology in conjunction with the analysis of Kwon [62] to avoid the inefficiencies of “emulating” honest servers (see Section V). Additionally, we derive a more straightforward analysis and bound on the number of mix layers required. Our mixing analysis may be of independent interest to other systems in this area.

Efficient blame, server elimination, and network healing. We design Trellis so that honest servers can efficiently detect failures and deviations from protocol, assign blame, and eliminate responsible parties from the network. Our blame protocols (described in Section VI) handle malicious servers and users without ever deanonymizing honest users. Trellis gracefully handles network changes (e.g., offline or eliminated servers) using proactive secret sharing (explained in Section VI), which permits on-the-fly state recovery following server churn.

Contributions. In summary, this paper contributes:

- 1) the design of Trellis: a novel system for anonymous broadcast providing cryptographic anonymity guarantees for users, network robustness, and concrete efficiency,
- 2) a simpler theoretical mixing analysis for routing mix-nets instantiated with a subset of mix servers that are malicious and colluding with a network adversary,
- 3) anonymous routing tokens: a new tool for anonymously enforcing random path selection in mix-nets, with conceivable applications to other networking systems,
- 4) boomerang encryption: a spin on *onion* encryption that allows for efficient proofs of delivery and blame assignment,
- 5) and an open-source implementation which we evaluate on a networked deployment, achieving a throughput of 320 bits per second with 100,000 users and 20 kB messages.

Limitations. Trellis shares the primary limitation of other metadata-private systems [20, 27, 60–62, 65, 78, 101, 103, 104]: it provides anonymity only among honest online users and therefore requires all users to contribute a message in each round—even if they have nothing to send—so as to hide all network metadata [42]. (Users who stop participating correlate themselves with patterns or content of messages that stop being output [24, 25, 110].) Trellis does not, however, preclude the use of any external solutions to this problem (e.g., [110]).

II. BACKGROUND AND RELATED WORK

We start by describing mix-nets and their guarantees in Section II-A. We then describe related work and existing use of mix-nets for anonymous communication in Section II-B.

¹In the case of a dishonest majority of servers, Trellis reconfigures itself into a new network containing a dishonest *minority*; see Section VI-C2.

A. Mix networks

In 1981, Chaum [16] developed the first mix network (mix-net) for the purpose of anonymous *email*. The idea behind Chaum’s mix-net is simple. A set of servers is arranged in a sequence. Each server has a public key. Each user recursively encrypts the user’s message under the sequence of server public keys. We call this recursive encryption a sequence of **envelopes** (a.k.a. onions [32]), with each envelope encrypted under all subsequent public keys in the sequence. All users send the first envelope to the first server in the sequence. Each server, in sequence, decrypts the envelopes it receives, randomly shuffles the decrypted envelopes, and forwards them to the next server in the sequence. This repeats for N layers. The use of onion encryption prevents intermediate servers from learning which messages they are routing (a server can only decrypt one layer of encryption). At the end of the sequence, the innermost envelope is decrypted to reveal the plaintext message. The mix-net guarantees that the ordering of the messages is uniformly random and independent from the input, which in turn guarantees unlinkability between each user and their submitted message. As such, classic mix-net architectures can be used for anonymous broadcast. Unfortunately, this architecture suffers from security and practicality challenges:

- 1) Mix-nets do not inherently guarantee metadata privacy: a network adversary observing all traffic can link messages to users through metadata, such as timing information.
- 2) Mix-nets are vulnerable to active attacks (e.g., dropped envelopes), which can be exploited by malicious servers to link users to their messages [70, 75, 79, 94, 108].
- 3) Sequential mix-nets do not provide scalability: adding more servers does **not** improve performance since *all* servers must process *all* envelopes.

Scalable mix-nets. Scalable systems built using mix-net architectures are designed around **parallel mix-nets** [44, 89], where users are assigned to a mix-net instantiated from a subset of servers in the network. Adding more servers thus proportionally increases the total capacity of the network. However, naively parallelizing mix-nets does not provide privacy, since a message may end up mixed with only a small subset of other users’ messages. To avoid this, parallel mix-nets must model a *random permutation network* [22, 106] (a theoretical framework for modeling mixing in a communication network) to guarantee that all messages are mixed [44, 62, 89]. This is the strategy taken in Atom [60].

B. Related work

In this section, we focus on comparing Trellis to other systems that achieve anonymous broadcast. We pay particular attention to mix-net-based systems as other architectures have incomparable threat models.

Anonymous broadcast using parallel mix-nets. Rackoff and Simon [89] put forth the idea of anonymous broadcast using parallel mix-nets based on permutation networks but do not build or evaluate their approach. Kwon et al. [60] are the first to use their ideas to build a prototype system called Atom. A problem with mix-nets instantiated as permutation networks is that they only guarantee mixing if all servers are honest. The insight in Atom is to use **anytrust groups**—sets of random

servers such that at least one server in each set is honest with high probability—to emulate a trusted network from groups of possibly malicious servers. That is, each “server” in the mix-net is emulated by an anytrust group selected at random. Instantiating the mix-net in this way comes at a high cost in terms of concrete performance. The problem is that Atom requires expensive zero-knowledge proofs to ensure servers in each group behave correctly. These quickly become computational bottlenecks limiting the throughput and preventing scaling to large message sizes (Kwon et al. [60] only evaluate Atom on up to 160 byte messages). Another problem is that emulating honest servers with anytrust groups wastes resources and introduces large communication overheads. These performance barriers make Atom impractical to deploy.

Quark is a followup proposal for anonymous broadcast that appears in Kwon’s PhD thesis [62] but is primarily of theoretical interest. Like Atom, Quark is designed on top of a random routing mix-net. However, Quark avoids emulating honest servers by developing a novel theoretical mixing analysis for a network containing malicious servers. Unfortunately, Quark does not prevent disruption attacks: any malicious user or server can cause the entire protocol to restart by dropping (or not sending) a message. Moreover, we find that Quark may not be secure: the blind Schnorr multi-signatures around which Quark is designed can be forged using parallel signing attacks [33, 80] (these attacks were discovered *after* Quark). In Trellis, we salvage the theoretical underpinning of Quark and use a routing mix-net as a starting foundation. We also contribute an improved mixing analysis for this routing mix-net when instantiated with malicious servers, inspired by the ideas of Kwon [62] (see Section V-E). Apart from the mix-net, all other components of Trellis are different.

Anonymous communication using mix-nets. Chaum [16], cMix [15], MCMix [5], Karaoke [65], Stadium [101], Vuvuzela [104], XRD [61], are all anonymous point-to-point communication systems built on top of mix-net architectures. Of these, Karaoke, Stadium, and XRD provide horizontal scalability. However, only XRD provides *cryptographic privacy*. Karaoke, Stadium, and Vuvuzela provide *differential privacy*, which introduces a host of practical challenges (see discussion of Kwon et al. [61, Section 1]). Other systems, such as Loopix [86], Streams [27], and Tor [32] are examples of *free-route* mix-nets [26, 40, 68, 76, 90], i.e., parallel mix-nets where users asynchronously route messages via a small subset of mix servers. Free-route mix-nets do not achieve metadata privacy in the face of network adversaries and are susceptible to traffic analysis and active attacks [9, 45, 82].

Other architectures. Other anonymous broadcast systems [2, 5, 7, 19, 20, 35, 72, 78, 109] are instantiated under different threat models. Dining cryptographer networks (DC-nets) and multi-party computation [2, 5, 35, 72] require two or more *non-colluding* servers to process messages and prevent malicious users from sending malformed content. Because of this, these techniques inherently do not provide horizontal scalability and impose high overheads on the servers. Other anonymous communication systems [4, 6, 17, 41, 59, 66, 67, 69, 92] focus on specific applications (e.g., group messaging [17] and voice-communication [4, 66, 67]) or focus on one-sided anonymity

(e.g., only recipient anonymity [6]). We refer the reader to the excellent survey of systems in Piotrowska’s PhD thesis [85].

III. SYSTEM OVERVIEW

In Section III-A, we describe the core ideas and intuition underpinning the overall design of Trellis. In Section III-B, we describe the threat model and guarantees achieved by Trellis.

A. Main ideas

Following other mix-net based systems [15, 27, 60, 62, 65, 66, 89, 92, 101, 104], Trellis is instantiated using a set of N servers and M users. The servers are organized into a network of L layers and users’ messages are routed by servers from one layer to the next (see Figure 1). Trellis is organized in two stages. A one-time path establishment stage followed by a repeated message broadcasting stage. In both stages, Trellis resolves failures and attacks by malicious users and servers through on-demand blame protocols (see Section VI).

Do once: Path establishment. We reference the steps in Figure 2. Path establishment proceeds layer-by-layer. ① Users obtain signed anonymous routing tokens (Section V-A) from an anytrust group of signers, which guarantees the resulting path assigned to the user is uniformly random. Each token issued for a layer determines the server at the next layer and contains all the necessary routing information needed to route messages between the two layers (details in Sections V and V-A). ② Users distribute the L tokens anonymously using boomerang encryption (Section IV-A). All servers obtain the necessary routing tokens without learning the full path.




Fig. 2: Users are assigned a server in each layer through a one-time setup protocol involving an anytrust group of signers.

Repeat: Anonymous broadcasting. We reference the steps in Figure 3. Anonymous broadcasting repeats indefinitely on the pre-established path. ① Each user sends their envelope (containing the user’s message) through the server assigned by the routing token, for each of the L layers. ② Servers collect all the envelopes, decrypt one layer of encryption, and group the resulting envelopes based on their destination server in the next layer. Servers shuffle and pad each group with dummy envelopes as necessary. Each group is then sent to the destination server in the next layer. ③ Once envelopes reach the final layer, they are guaranteed to be shuffled and unlinkable to the users that sent them, at which point the final layer of encryption is removed to reveal the plaintext message.

On-demand: Blame, elimination, and recovery. It could be the case that a server goes offline or acts maliciously, both during path establishment and the broadcasting stages. To deal




Fig. 3: Users send messages through the layers, based on their paths. Users can repeat this process indefinitely to broadcast many different messages.

with such failures (intentional and otherwise) we develop on-demand blame and recovery protocols that automatically reassign affected paths. Importantly, failure resolution is handled *locally* on the link between two layers and does not involve the users. (Failures caused by a malicious user cause the user to be eliminated by other means; see Section VI.) We follow the steps of Figure 4. Blame and recovery is invoked whenever an error is detected (e.g., wrong signature, dropped envelopes). ① A server detects an error in the received batch of envelopes in layer $(i+2)$ and proceeds to blame the server in layer $(i+1)$ that sent the batch. ② Servers evaluate evidence from both parties to decide which of the two is malicious and ③ vote for the honest party by providing a secret share of the malicious party’s secret state. The state is used by replacement server to take over the duties of the eliminated server.




Fig. 4: On-demand blame and recovery protocols are invoked by servers to automatically reassign affected paths and eliminate malicious servers from the network.

B. Threat model and assumptions

The threat model and assumptions underpinning Trellis mirror those of existing anonymous broadcast [60, 62] (and communication [5, 15, 61, 65, 66, 101, 104]) systems based on mix-nets that provide metadata-privacy guarantees.

Assumptions and adversarial model.

Observed network. We assume that all communication on the network is observed by a passive network adversary but that the communication *content* is encrypted using standard authenticated public key infrastructure (PKI)—e.g., TLS [91].

Malicious users and servers. We assume that a fixed fraction $0 \leq f < 1$ of servers and any number of users are malicious, may arbitrarily deviate from the protocol, and collude with the network adversary. Note: in the case of a dishonest majority ($f \geq 0.5$), our blame protocols in Section VI cause successive

network reconfigurations which converge to an *honest majority* (details in Section VI-C2).

Cryptography. We make use of standard cryptographic assumptions. We require digital signatures, symmetric-key encryption, and the computational Diffie-Hellman (CDH) assumption [31] (and variants of the CDH assumption for *gap* groups [56]). We also require pairings on elliptic curves [12, 53] and hash functions modeled as random oracles when constructing our routing tokens (Section V-A).

Liveness. We require two liveness assumptions: synchronous communication and user liveness. Both assumptions are standard in all prior work on metadata-private anonymous communication [2, 5, 6, 19, 20, 35, 60–62, 78, 101, 104, 109].

- *Synchronous communication:* Servers are expected to be able to receive all envelopes for each layer before processing the next layer. This is equivalent to assuming all servers have access to a global clock [95].
- *User liveness:* We assume that all users submit a message for every round—even if they have nothing to send—to prevent the risk deanonymization via intersection attacks [24, 25, 110]. We do **not** assume server liveness.

Non-goals. Trellis does not aim to prevent denial-of-service (DoS) attacks on the network or handle infrastructure failures. However, Trellis is fully compatible with existing solutions [18, 21] for network availability and reliability.

Trellis guarantees. Under the above threat model and assumptions, Trellis provides the following guarantees.

Robustness and availability. No subset of malicious servers or colluding users can disrupt availability for honest users.

Message delivery. All messages submitted by honest users are guaranteed to be delivered and result in a “receipt” that can be efficiently verified by each user.

Unlinkability. Fix any small $\epsilon > 0$ and any subset of honest users. After mixing, no message (or set of messages) can be linked with any honest user (or set of honest users) with probability ϵ over random guessing by the adversary.

Metadata privacy. The unlinkability guarantee holds even when all network communication is visible to the adversary who is controlling the malicious users and servers.

IV. BUILDING BLOCKS

In this section, we describe the cryptographic tools and concepts that we use to instantiate Trellis in Section V.

Cryptographic foundations. We briefly describe classic cryptographic primitives that we use as building blocks in Trellis. We point to excellent textbooks for more details on these standard primitives [11, 43, 55].

Diffie-Hellman key agreement. The Diffie-Hellman key agreement protocol [31] allows two parties to establish a shared secret key over an authenticated communication channel. Let

\mathbb{G} be a suitable group of prime order p with generator g .

$\text{DHKeyGen}(\mathbb{G}, g) \rightarrow (A, a):$	$\text{DHKeyAgree}(A, b) \rightarrow \text{sk}:$
1: Sample $a \in \mathbb{Z}_p$.	1: Output $\text{sk} := (A)^b$.
2: Output $(A := g^a, a)$.	

Observe that for all (A, a) and (B, b) sampled according to $\text{DHKeyGen}(\mathbb{G}, g)$, $\text{DHKeyAgree}(A, b) = \text{DHKeyAgree}(B, a)$. Thus, two parties derive the same secret key sk .

Digital signatures. A digital signature (DS) scheme consists of three algorithms ($\text{KeyGen}, \text{Sign}, \text{Verify}$) satisfying *correctness* and *unforgeability*. Informally, correctness requires Verify to output yes on all valid signatures and unforgeability states that no efficient adversary should be able to produce a forged signature under vk that is accepted by Verify .

Symmetric key encryption. An encryption scheme (e.g., AES [23, 81]) consists of three algorithms ($\text{KeyGen}, \text{Enc}, \text{Dec}$) satisfying correctness and IND-CCA security. Informally, correctness requires that $\text{Dec}_{\text{sk}}(\text{Enc}_{\text{sk}}(m)) = m$ for all messages m and secret keys sk sampled according to KeyGen . IND-CCA security requires the distribution of ciphertexts induced by any pair of messages to be computationally indistinguishable *and* authenticated (see Boneh and Shoup [11] for definitions).

Anytrust groups. Let κ be a statistical security parameter. Given a set of servers where some fraction $0 \leq f < 1$ are malicious, there exists an integer s such that the probability that any group of s random servers contains *at least* one honest server is $1 - 2^{-\kappa}$ (proof: it is enough to pick s such that $f^s < 2^{-\kappa}$) [60].

A. New tool: Boomerang encryption

A core challenge in a mix-net is ensuring that all envelopes get to the appropriate server on the path. A user complaining of protocol deviation (e.g., dropped messages) runs the risk of deanonymization. Boomerang encryption—a simple extension of onion encryption [16]—solves this problem by providing “receipts” that can be used by servers to *anonymously* blame on a user’s behalf. Boomerang encryption is designed to (1) ensure that all servers on a path receive messages addressed to them and (2) allow honest servers to trace the source of active attacks and assign blame accordingly.

Onion encryption and routing. Onion encryption [16, 32] involves recursively encrypting a message under a sequence of keys. Only the corresponding secret-key holders can recover the message by “peeling off” the layers of encryption one by one. Formally, given any symmetric-key encryption scheme with algorithm Enc , onion encryption is defined as:

$$\overrightarrow{\text{onion}}(\mathbf{sk}, m) := \text{Enc}_{\text{sk}_1}(\text{Enc}_{\text{sk}_2}(\dots \text{Enc}_{\text{sk}_\ell}(m) \dots)),$$

where $\mathbf{sk} := (\text{sk}_1, \dots, \text{sk}_\ell)$ and m is the message.

Onion routing. Onion encryption is the backbone of onion routing, where a chain of servers route onion-encrypted messages, removing one layer of encryption at a time. Onion routing prevents any intermediate server from learning the plaintext message. We will use the term *envelope* to describe the intermediate ciphertexts of an onion-encrypted message.

Boomerang encryption and routing. Onion routing works in one “direction:” routing envelopes *forward* through the chain of servers such that only the destination server learns the plaintext message. In contrast, boomerang routing is bidirectional: it applies onion encryption *twice*, once in the forward direction and once in reverse (hence the name, *boomerang*). If we define reverse-onion encryption as:

$$\overleftarrow{\text{onion}}(\mathbf{sk}, m) := \text{Enc}_{\text{sk}_\ell}(\text{Enc}_{\text{sk}_{\ell-1}}(\dots \text{Enc}_{\text{sk}_1}(m) \dots)),$$

then boomerang encryption is defined as:

$$\overleftrightarrow{\text{boomerang}}(\mathbf{sk}, \vec{m}, \overleftarrow{m}) := \overrightarrow{\text{onion}}(\mathbf{sk}, \vec{m} \parallel \overleftarrow{\text{onion}}(\mathbf{sk}, \overleftarrow{m})),$$

where \vec{m} is the forward message and \overleftarrow{m} is the returned message. We note that the \overleftarrow{m} cannot be returned without first decrypting all envelopes in the forward onion chain.

Boomerang decryption as an anonymous proof-of-delivery. We make the following somewhat surprising observation: a successful boomerang *decryption* can be turned into a proof of message delivery when applied to onion routing.² Specifically, letting \overleftarrow{m} (the returned message) be a random nonce acts as a message delivery receipt: it guarantees that every server along the path decrypted the correct onion layer and forwarded it to the next server. To see why this holds, observe that in order to successfully recover \overleftarrow{m} , all envelopes in the forward onion must be decrypted to recover $\vec{m} \parallel \overleftarrow{\text{onion}}(\mathbf{sk}, \overleftarrow{m})$. Therefore, message \vec{m} must have been delivered to the server at the end of the path (i.e., the last envelope in the forward onion chain must have been decrypted).

V. THE TRELLIS SYSTEM

In this section, we provide further details on the path establishment and broadcasting protocols (overviewed in Section III). The formal protocol descriptions are provided in the full version [64]. We start by describing anonymous routing tokens (Section V-A), which are integral to the design of Trellis.

A. New tool: Anonymous routing tokens

Anonymous routing tokens (ARTs) are used in Trellis to determine and verify routing paths. Our construction of ARTs is inspired by anonymous tokens based on blind signatures [28, 34, 56], but has two key differences: namely, ARTs are signed by a *group* of signers and are publicly verifiable. Specifically, tokens are signed by an anytrust group. This guarantees that at least one honest signer is present when signing each token and forms the basis for integrity in Trellis. The blind signature serves two purposes: ensuring that (1) the payload encapsulated by the ART comes from an (anonymous) user, and (2) the path through the mix-net, as determined by the ART, is uniformly random (even for malicious users).

How our routing tokens work. Routing tokens are generated by users. Each token encapsulates the necessary information a server requires to route envelopes on a link in the mix-net. Each signer (a server in an anytrust group) holds a

partial signing key x_i associated with a global public signature verification key tvk_ℓ for the ℓ th layer. The user has each group member provide a partial signature for the blinded token. The user then unblinds and recovers the multi-signature with respect to tvk_ℓ by combining the partial signatures. If at least one signer is honest, then the signature is unforgeable and cannot be traced to the tokens seen during signing. The randomness of each token signature—which is uniform and independent provided at least one signer is honest—determines how envelopes are routed through the network.

Definition 1 (Anonymous Routing Tokens; ARTs). Fix a security parameter λ , integers $s, N > 1$, a common reference string crs generated using a trusted Setup process, and a set \mathbb{S} of N server identifiers. An ART scheme consists of algorithms KeyGen, NextServer, Verify, and an interactive protocol Sign instantiated between a user and s signers.

- KeyGen(crs) \rightarrow ($\text{tvk}, x_1, \dots, x_s$): takes as input the crs ; outputs a public key tvk and s partial signing keys.
- Sign(⟨User($\text{tvk}, \text{payload}$), Signers(x_1, \dots, x_s)⟩) \rightarrow (t, \perp). The user inputs a payload consisting of the server identity $S \in \mathbb{S}$, a **new** Diffie-Hellman message and a **new** digital signature verification key. The i th signer provides as input the partial signing key x_i . The user obtains a (signed) routing token t committing to the payload. The signers obtain no output (\perp). See Figure 5 for an overview.
- NextServer($\text{tvk}_\ell, \text{t}_\ell$) \rightarrow ($S_{\ell+1}$): takes as input the layer public key tvk_ℓ , the token t_ℓ ; outputs the identity of the next server $S_{\ell+1}$ deterministically chosen based on t_ℓ .
- Verify($\text{tvk}_\ell, S_{\ell+1}, \text{payload}_\ell, \text{t}_\ell$) \rightarrow yes or no: takes as input the public key for layer ℓ , the payload payload_ℓ , identity of the next server $S_{\ell+1}$, and signed token t_ℓ ; outputs yes if ($S_{\ell+1}$ and payload_ℓ are valid under t_ℓ and no otherwise.




Fig. 5: Overview of the ART signing protocol (Sign).

The above functionality must satisfy the standard properties of blind digital signatures [87] (and anonymous tokens [28, 56]): *completeness*, *unforgeability*, *unlinkability*, in addition to a new property we call *unpredictability*. Unforgeability ensures that if the ART verifies under the public key tvk , then the encapsulated payload came from some user and was not tampered with. Unlinkability for ARTs guarantees that the encapsulated payload is random and independent of any user (thus the Diffie-Hellman message and a signature verification key are sampled randomly). Finally, unpredictability requires that $S_{\ell+1}$ (as output by NextServer) be unpredictable given the payload. We formalize these properties in the full version [64].

1) ART construction: We realize an ART scheme satisfying Definition 1 following the blueprint of BLS signatures [12] and anonymous tokens [28, 56]. In the full version [64], we prove security of our construction under a variant of the computational Diffie-Hellman (CDH) assumption in gap groups [10, 12, 56]. Let \mathbb{G} and \mathbb{G}_T be groups of prime order p with generator g and equipped with an efficiently computable non-degenerate bilinear pairing $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ [11]. We let $\text{crs} := (\mathbb{Z}_p, \mathbb{G}, g) \leftarrow \text{Setup}(1^\lambda)$ for a security parameter λ .

²Note that proof-of-delivery requires the underlying encryption scheme to be IND-CCA secure (such as Encrypt-then-MAC [11] with AES).

KeyGen(crs):

- 1: $x_1, \dots, x_s \leftarrow_R \mathbb{Z}_p$
- 2: $x \leftarrow \sum_{i=1}^s x_i$
- 3: $\text{tvk} \leftarrow (g^{x_1}, \dots, g^{x_s}, g^x)$
- 4: **return** $(\text{tvk}, x_1, \dots, x_s)$

NextServer(tvk_ℓ, t_ℓ):

- 1: $S_{\ell+1} := \text{Hash}_{\mathbb{Z}_N}(\text{tvk}_\ell || t_\ell)$
- 2: **return** $S_{\ell+1}$

Token signing protocol. We instantiate the token signing protocol in one round between the user and each signer. The user begins by generating a fresh Diffie-Hellman key exchange message that will later be used to derive the encryption key for its ℓ th link. The user also generates a fresh digital signature key pair, which will be used to verify integrity of its encryptions on the ℓ th link.

Step 1 (user):

- 1: $(A_\ell, a_\ell) \leftarrow \text{DHKeyGen}(\mathbb{G}, g)$; $(\text{vk}_\ell, \text{sk}_\ell) \leftarrow \text{DS.KeyGen}(1^\lambda)$.
- 2: $r \leftarrow_R \mathbb{Z}_p$; $T_\ell \leftarrow \text{Hash}_{\mathbb{G}}(\underbrace{S_\ell || A_\ell || \text{vk}_\ell}_{\text{payload}_\ell})^{1/r}$.
- 3: **send** T_ℓ to all signers.

Step 2 (ith signer): Upon receiving T_ℓ , the i th signer partially signs the blind hash using their share x_i of the signing key:

- 1: $W_{\ell,i} \leftarrow T_\ell^{x_i}$.
- 2: **send** $W_{\ell,i}$ to the user.

Step 3 (user): The user combines the partial signatures and outputs an unblinded token:

- 1: $t \leftarrow (\prod_{i=1}^s W_i)^r$.

The user then checks that t_ℓ was signed correctly:

- 2: $S_{\ell+1} \leftarrow \text{NextServer}(\text{tvk}_\ell, t_\ell)$.
- 3: $\text{Verify}(\text{tvk}_\ell, S_{\ell+1}, \text{payload}_\ell, t_\ell) \stackrel{?}{=} \text{yes}$.

If Verify outputs no, then the user discards the token and outputs \perp . Otherwise, the user stores $(\text{payload}_\ell, t_\ell, \text{sk}_\ell, a_\ell)$.

Generalization. We remark that our construction of ARTs can be generalized to support more generic anonymous tokens such as the ones described in Kreuter et al. [56] and Davidson et al. [28] but instantiated with a *set* of signers rather than just one signer as in previous work [28, 56, 102]. Additionally, unlike prior constructions, ARTs are *publicly* verifiable. These features make ARTs potentially of independent interest. For simplicity, we restrict our definition of ARTs to Trellis.

B. Putting things together

We decouple the expensive cryptographic operations from the lightweight ones (similarly to cMix [15]). During key generation and path establishment, all necessary cryptographic key material is distributed via the ARTs to servers on the path. During the mixing phase, messages are routed through the established paths. This is a lightweight procedure: each server only needs to decrypt and shuffle envelopes. The mixing phase can be repeated indefinitely (with many different messages) and does not involve the setup overhead of the path establishment phase.

Verify($\text{tvk}_\ell, S_{\ell+1}, \text{payload}_\ell, t_\ell$):

- 1: parse $\text{tvk}_\ell = (g^{x_1}, \dots, g^{x_s}, g^x)$
- 2: $y \leftarrow \text{Hash}_{\mathbb{G}}(\text{payload}_\ell)$
- 3: $S'_{\ell+1} \leftarrow \text{NextServer}(\text{tvk}_\ell, t_\ell)$
- 4: **return** $e(t_\ell, g) = e(y, g^x)$
and $S'_{\ell+1} = S_{\ell+1}$

Key generation. Each server uses DHKeyGen and publishes the Diffie-Hellman public key to the PKI. Each server uses Feldman’s verifiable secret sharing (VSS) scheme [39] to share the Diffie-Hellman secret, for state recovery purposes (Section VI-C1). In addition, we use proactive secret sharing [48, 74] to generate shares of a common ART secret signing key for each anytrust group. With proactive secret sharing, multiple *independent* sets of shares are generated for the *same* secret. This ensures that each anytrust group holds shares of the secret key while preventing subsets of colluding servers across *different* groups from recovering the secret key.

ART generation. Each user is issued one ART per layer by a random anytrust group. Each user can only send (boomerang encrypted) messages through the network according to the path dictated by the ARTs. All invalid or duplicate envelopes are discarded by the honest servers. Servers that improperly route envelopes are blamed by the next honest server on the path (blame protocols are described in Section VI).

Routing tokens and path selection. During the path establishment protocol (described in further detail in Section V-C), each server (anonymously) receives ARTs for the links it is responsible for. Each token determines the next server on a link (and can be checked with ART.NextServer) which allows each server to (1) determine where to send envelopes in the subsequent layer and (2) anonymously verify path adherence (without knowledge of the full path, only the local link).

Boomerang routing on a path. We use boomerang encryption to route envelopes through the mix-net. We assume that the boomerang encryption scheme is instantiated using IND-CCA secure encryption (where the shared encryption key is derived from the Diffie-Hellman message in the ART payload). Each intermediate layer ciphertext (i.e., envelope) is signed and can be verified by the server using the signature verification key given in the corresponding ART payload. See the formal protocol in the full version [64] for details. This gives us *authenticated* boomerang routing which ensures that: (1) each server receives every envelope for each link and (2) all envelopes on a link are valid authenticated encryptions under the associated link public keys. Any failure of these two properties (e.g., failed decryption, bad signatures, or missing envelopes) triggers a blame protocol (described in Section VI-A4).

Dummy envelopes. To prevent leaking how many messages are passed between servers in each layer (required to avoid leaking parts of the mixing permutation applied on the messages to the network adversary), each server batches the envelopes (intermediate boomerang ciphertexts) it sends to the destination server(s) in the next layer, padding the batch with “dummy” envelopes as needed. The dummy envelopes are then discarded by the receiving server(s).

Assigning blame. All communication between parties is additionally signed using a digital signature (DS) scheme for lightweight blaming of malicious parties. See Section VI for details on our blame protocols.

C. Do once: Path establishment

We now describe how we resolve the first challenge: anonymously establishing random paths. The difficulty is that

paths must be created in a way that (1) does not reveal the path to any party other than the user and (2) guarantees that all paths are uniformly random, even when generated by malicious users. Each server on the path must receive an ART associated with the path link but must not learn which user provided it. We solve this with an incremental approach: we inductively extend an existing path from one server to another while simultaneously forcing all users to choose random paths. Each user starts with a random path of length one, with the choice of this first server assigned by the first ART. This base case does not yet provide anonymity since the first server knows the user’s identity. We then proceed by induction as follows. Given a path of length ℓ ending at server S_ℓ , the user extends the path to length $\ell + 1$ using boomerang routing through $S_1 \rightarrow S_2 \rightarrow \dots \rightarrow S_\ell$ to communicate with the new server $S_{\ell+1}$. The user sends the routing information $S_{\ell+1}$ needs: the ART t_ℓ for the link $(S_\ell \rightarrow S_{\ell+1})$, the ART $t_{\ell+1}$ for the link $(S_{\ell+1} \rightarrow S_{\ell+2})$, and the associated Diffie-Hellman messages and signature verification keys. (The user also sends a signature on the ART $t_{\ell+1}$ under the verification key t_ℓ to prove that t_ℓ and $t_{\ell+1}$ are part of the same path.) Because the user communicates exclusively through the existing path and because ARTs are unlinkable, the path extension procedure reveals no information to any server. After receiving and verifying the tokens, $S_{\ell+1}$ routes the reverse boomerang envelopes backwards along the path. Each server verifies the authenticated encryption, which implicitly ensures that the user and all servers $S_\ell \rightarrow S_{\ell-1} \rightarrow \dots \rightarrow S_1$ know that $S_{\ell+1}$ received (and validated) the tokens. This process is repeated L times until the full path to S_L is established.

Anytrust group bookends. The Trellis mix-net consists of L mix-net layers “bookended” between two layers of anytrust groups (see Figure 6). The anytrust groups at the entry and exit layers ensure that the start and end of a mix-net path are always processed by an honest server, which prevents dropped messages and denial-of-service attacks. The anytrust group at the entry layer verifies the boomerang receipts on behalf of the user: when the boomerang receipt comes back, each member of the group can check the attached signature and ensure that the receipt was correctly decrypted. Additionally, the entry group stores the envelopes it is given so as to prevent denial-of-service attacks: if the first k servers in the mix-net are malicious and drop messages, the entry group can resend the envelopes after the adversarial servers are eliminated. The anytrust group at the exit layer outputs the messages to the world (the “broadcast” part of anonymous broadcast), ensuring that all messages make it out. During path establishment, these exit groups are at the “peak” of the boomerang and use a group decryption protocol to decrypt the L th boomerang envelopes. In this protocol, each group verifies the ART for layer L and retains the included signature verification key. (Note that the verification keys are unlinkable to the start of the path because of the mixing guarantees provided by the mix-net.) Then, during the repeated broadcast rounds, these exit groups check that one message is received for each of the verification keys they have, ensuring delivery of all messages. See the formal protocol in the full version [64] for details.

Why path establishment works. Both the incremental path extension and the boomerang message receipts are required to prevent dropped messages. The incremental approach to

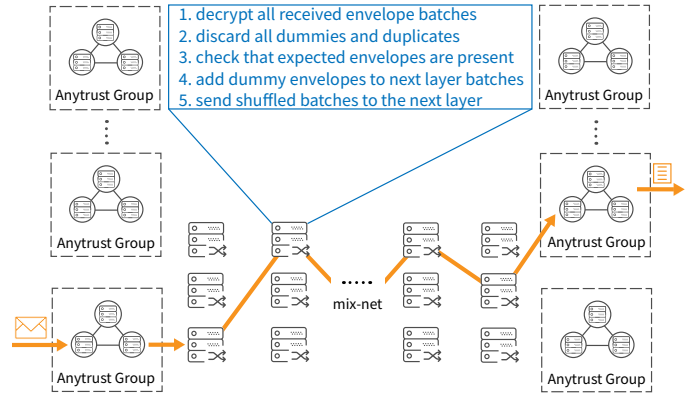


Fig. 6: Detailed overview of the Trellis mix-net “bookended” by anytrust groups to prevent denial-of-service attacks at the entry and exit points. Users send boomerang-encrypted messages through a pre-selected random path in the network. The anytrust group at the last layer ensures that all messages get published to the bulletin board.

path establishment enables immediate blame assignment and prevents errors from propagating. The message receipt is designed to allow servers on the existing path to assign blame on a user’s behalf if the new server deviates from protocol. Specifically, each server expects to receive an envelope for each established key, in both the forward and reverse direction. This gives rise to a simple blame protocol in the case of a dropped envelope: blame the previous server. The previous server, if innocent, should have reported the envelope missing during the last synchronous layer; because it did not do so, it must be malicious (see Section VI for details).

D. Repeat: Mixing rounds

The repeated broadcasting rounds use the established paths to mix (different) message anonymously and quickly. Messages in Trellis are of fixed size (e.g., 10kB), and all users must submit a message of that size, padding if necessary. Because mixing rounds can be repeated many times, users with larger messages can split their broadcasts across several rounds. The random routes guarantee that all messages are mixed before reaching the final layer (see mixing analysis in the next section).

E. Number of layers required

We now turn to analyzing the number of layers required to achieve a random permutation on the delivered messages. This analysis forms the crux of our security argument in Section VII. Previous theoretical work [46, 89] performs similar analyses in the context of *permutation networks*. Unfortunately, such analyses are only applicable to networks consisting exclusively of *honest* servers. The analysis of Kwon [62] is the first to consider networks containing a fraction of malicious servers. We improve the mixing analysis of Kwon [62]. We give a simpler analysis of mixing guarantees to bound the number of layers required in a random routing mix-net.

Modeling the network. We consider an information-theoretic view of an adversary observing the network and learning a fraction f of the permutation applied to the messages in each

	Number of users				
	100 K	500 K	1 M	5 M	10 M
$f = 0.1$	56	58	58	60	60
$f = 0.2$	86	88	89	92	93
$f = 0.3$	124	128	129	133	134
$f = 0.4$	179	184	186	191	194
$f = 0.5$	266	273	277	284	288
$f = 0.6$	419	431	436	448	453
$f = 0.7$	736	757	767	787	796
$f = 0.8$	1603	1649	1668	1715	1734
$f = 0.9$	6069	6244	6319	6494	6569

TABLE I: Number of layers required for mixing as a function of the malicious server fraction f and the number of users. See full version of the paper [64] for derivation.

layer (obtained from the servers that collude with the adversary in Trellis). We construct a Markov chain where each state consists of the server each message is in and a deck of M cards, where each card uniquely corresponds to a message. We make one step in the Markov chain with each layer-to-layer transition. The servers are chosen randomly, which corresponds to the random path assignment dictated by the ARTs. We shuffle the deck according to the hidden part of the permutation, summarized by the following observation: only messages sent between two *honest* servers (in adjacent layers) are mixed, meaning that messages passing between two *malicious* servers are not mixed.

We model the adversary’s view as a probability distribution over all possible permutations of the deck and message locations. At the start, the adversary knows who submitted each message, and can choose the starting state. We determine the probability distribution for subsequent layers through the Markov process. We note that the above process operates on the entire group of permutations, and so eventually converges to the uniform distribution where each deck permutation is equally likely (see [8, Lemma 15.1]). This corresponds to a uniformly random permutation of the messages.

Determining the number of layers. To determine the number of layers required, we need to find the “time” to convergence. Computing this value requires some additional analysis. This analysis, while straightforward, is somewhat tedious to derive (and is similar in nature to the analysis of Kwon [62]). For details, see the full version of the paper [64]. We report the number of layers as a function of the corruption fraction f in Table I. We note that the number of layers is logarithmic in M (the number of users) and is not dependent on N (the number of servers). This is because the probability of being routed through an honest server only depends on f [62].

VI. BLAME, ADVERSARY REMOVAL, AND RECOVERY

If all parties behave honestly, then all envelopes are present and well-formed. In this case, boomerang routing and mixing analysis (Section V-E) ensure that all messages are delivered and mixed (see security analysis; Section VII). However, an adversarial user or server may deviate from the protocol in arbitrary ways. To prevent system disruption, or worse, user deanonymization following active attacks, each (honest) server is expected to verify all envelopes it receives for layer ℓ *before* sending the shuffled batch of envelopes to layer $\ell+1$. If any of

the checks do not pass, the server must notify (all) the other servers before the protocol finishes layer ℓ (note that this server will not send messages for $\ell+1$ until the error is resolved, and so layer ℓ should not complete). This prevents errors from propagating through honest servers; honest servers blame immediately in the layer where a check fails and recovery protocols are initiated to eliminate the blamed server(s).

Blaming without deanonymization. We make an important observation which allows us to maintain the anonymity of honest users when blaming malicious servers: if we can show that at least one of the two servers on a link misbehaved, then we can reveal the envelopes along that link *without* compromising user anonymity. We formalize this in Claim 1.

Claim 1 (Malicious links provide no anonymity). Links between two servers, where at least one of the servers is an adversary, can be revealed without compromising anonymity.

Proof: Consider the view of an adversarial server on a link consisting of at least one malicious server. The adversary knows all encrypted envelopes on the link. Revealing any of this information (to all of the servers) does not further the adversary’s knowledge. Our analysis of Section V-E relies only on links between pairs of *honest* servers for this reason. ■

A. Blame protocols

In order to illustrate our blame protocols, we provide more details about what exactly is sent between servers (see the full version [64] for a formal description). First, all signatures sign the round, layer, source server (which may be different from the signer), destination server, protocol, length, and direction of travel. This will allow anyone to publicly verify which step of which protocol a message corresponds to. Second, servers sign the batch of envelopes and corresponding ART keys. Signing the concatenation of the batch ensures that no server can equivocate on the number of envelopes it sent for a given round and layer. These signatures can use any standard public-key signature scheme [11, 55].

We assume that all blame messages are sent to all servers (e.g., using a standard gossip protocol [63, 83]). If we assume an honest majority of servers ($f < 0.5$; see Section VI-C2 for why this assumption is without loss of generality), the majority vote assigns blame and selects a replacement server. In Section VII, we cover how blame protocols support our delivery guarantees. In short, we ensure that no honest server or honest user can ever be blamed through a blame protocol, and thus only adversaries are blamed. The blame protocols ensure that the tokens, signatures, and envelopes are present and well-formed, which will imply successful delivery.

1) Arbitration protocol: Our first blame protocol deals with a missing or incorrect server signature of the concatenated envelopes. The arbitration protocol begins when a server $S_{\ell+1}$ reports that the signature given by S_ℓ is missing (including if no envelopes were sent at all) or incorrect. In this case, we have either that S_ℓ is adversarial or $S_{\ell+1}$ is maliciously initiating the blame protocol. An anytrust group randomly chooses an arbitrator from the other servers (the honest member will contribute true randomness and ensure the selection is random). We then require S_ℓ to send messages to the arbitrator, denoted

$A_{\ell, \ell+1}$, who forwards them to $S_{\ell+1}$. The arbitrator also checks that the signature on the batch is correct (using S_ℓ 's public verification key). If messages are delivered successfully (from S_ℓ to $S_{\ell+1}$) through $A_{\ell, \ell+1}$, then $A_{\ell, \ell+1}$ continues to arbitrate for future rounds. The performance of Trellis is only minimally impacted by arbitration, as it simply adds an extra forwarding and signature check for one layer.

If the arbitrator itself has a dispute with S_ℓ or $S_{\ell+1}$, then it can determine which server is honest from its perspective. When the arbitrator is honest, it will never have a dispute with the honest server. Once an arbitrator decides who the malicious server is (because it has a dispute), a new arbitrator is randomly selected by an anytrust group. If there is an honest majority of servers, then there will be more disputes involving the dishonest server than there are involving the honest server. For a dishonest majority, we can track all server disputes with a trust graph [107]: a graph in which each edge represents if two servers trust each other. Although the initial disagreement breaks only one trust graph link, we can choose arbitrators who trust both parties. Then, each dispute between the arbitrator and a party breaks another trust graph link. Wan et al. [107] prove that such an algorithm eventually converges to the honest clique, who can then reform without the troublesome members.

Because of the synchronous communication assumption, messages are always delivered between honest parties in a timely manner, and so the arbitration protocol is only invoked with adversaries. Arbitration introduces a maximum overhead of $2\times$, as each message is forwarded one extra time.

2) Duplicate envelopes corresponding to the same ART:

There should be exactly one message per ART per round and layer. A server can blame a user by producing as evidence two different signed envelopes for the same round and layer. This evidence is submitted to all of the servers, who vote to remove the user if the evidence is valid (or to remove the server if the evidence is invalid). We set $N(1 - f)$ as the minimum threshold of votes required to remove a server, as described in Section VI-C1. All of the protocols below follow this format: a party will produce some publicly verifiable evidence, which the servers will check and vote on accordingly.

3) *Envelope decryption failure:* If an envelope fails to decrypt, then a server $S_{\ell+1}$ can request that the previous server S_ℓ be blamed. This means that we are in one of three cases:

- 1) a malicious user sent an ill-formed envelope,
- 2) S_ℓ tampered with the envelope, or
- 3) $S_{\ell+1}$ is maliciously invoking the blame protocol.

Since the potentially malicious user, S_ℓ , and $S_{\ell+1}$ all know both the envelope and the verifying ART token, $S_{\ell+1}$ can reveal these to prove the envelope is ill-formed (in a similar argument to Claim 1). In response, S_ℓ may claim it never sent the ill-formed envelope and token. It then requests that $S_{\ell+1}$ reveal the signed concatenation of envelopes on the link to prove that the ill-formed envelope and ART were sent by S_ℓ . In this case, either S_ℓ or $S_{\ell+1}$ is blamed.

Otherwise, S_ℓ agrees that the envelope was ill-formed with respect to the corresponding ART, and we won't blame $S_{\ell+1}$. We still need to decide if the user sent the ill-formed envelope, or S_ℓ tampered with it. If S_ℓ is honest, then it will be able to present a proof that it acted correctly. This proof involves

revealing each step taken by S_ℓ and proving that each step was executed correctly. To prove this, S_ℓ first reveals the ART t_ℓ and the encapsulated payload (containing a partial Diffie-Hellman message and verification key). Using a discrete log equivalence proof [28, 50, 56], S_ℓ proves correct decryption (without revealing its long-term secret key). Finally, it provides the user's signature on $t_{\ell+1}$ under the key from t_ℓ to show that $t_{\ell+1}$ is the correct next token. With these elements the other servers can verify the following:

- S_ℓ received an envelope that was well-formed (signed), and therefore S_ℓ was not required to blame earlier.
- The decryption of said envelope was computed correctly, but the output of the decryption was an ill-formed envelope.
- The keys are correct and came from the user who created both of the ARTs, as t_ℓ signs the payload.

With this information, it can only be that the user is to blame, and we apply the protocol in Section VI-C3 to deanonymize and remove the malicious user. This well-formed envelope is signed (anonymously) by the user under some key vk_ℓ , and so could only have been created by them. Similarly, the token $t_{\ell+1}$ is signed by vk_ℓ and so is the correct next token, and the correct $vk_{\ell+1}$. Therefore, the server has produced a proof that the decryption was performed correctly, but the decryption output is ill-formed.

4) *Missing envelope:* Each server expects to receive one envelope corresponding to each ART pair it received. In the event that an envelope is missing, $S_{\ell+1}$ blames server S_ℓ for not sending the envelope (and not reporting the envelope missing in the previous layer). We have two cases to consider:

- $S_{\ell+1}$ is maliciously invoking the blame protocol, or
- S_ℓ is malicious and dropped an envelope.

Again, by Claim 1, $S_{\ell+1}$ can reveal the signed concatenation of the envelopes on link $(S_\ell \rightarrow S_{\ell+1})$. Specifically, $S_{\ell+1}$ provides the signed batch (with the missing envelope) that it claims to have received from S_ℓ in addition to a signed batch from a previous round showing that S_ℓ sent an envelope with the ART in question. We can then apply the protocol described in Section VI-C1 to remove the offending server.

Remark 1 (An edge case). At the peak of the boomerang during path establishment, note the pattern of sending $(S_\ell \rightarrow S_{\ell+1} \rightarrow S_\ell)$. Because $S_{\ell+1}$ does not know which tokens it is supposed to receive, it cannot blame S_ℓ for not sending them. Therefore, we allow $S_{\ell+1}$ to respond to claims by S_ℓ (that it dropped a receipt) by complaining that S_ℓ never sent the ART in the first place. It does this by presenting the signed concatenation from $(S_\ell \rightarrow S_{\ell+1})$ which is secure by Claim 1.

B. Malicious anytrust group member

First, if one of the signers in an anytrust group incorrectly signs the ART during the signing protocol described in Section V-A, the user receives an invalid signature (which it locally determines is invalid using ART.Verify). In this case, the user can request that each signer (publicly) proves it computed the partial signature correctly. In our construction, each signer can do so efficiently by providing a zero-knowledge proof of discrete log equivalence [28, 50, 56], proving that $W_{\ell,i}$ was calculated correctly with respect to g^{x_i} . Second, during

path establishment, we require the anytrust group to compute a group decryption at the last layer. If decryption fails, each signer can prove it computed the partial decryption correctly with a zero-knowledge discrete log equivalence proof. Any signer that does not comply is eliminated (see Section VI-C1).

C. Removal and recovery protocols

Once we have identified the adversarial party, we can proceed to remove them from the system. To remove adversarial servers, we ensure that each server’s secret keys are t -out-of- n shared, using a verifiable secret sharing [88] (VSS) scheme, with $t > Nf$ at setup time. Blame decisions are made as a group, using the t -out-of- n secret key shares as votes. Using the shares allows the server receiving t (or more) shares (i.e., votes) to recover the offending server’s secrets and replace it.

1) *Server removal and state recovery*: When servers vote to blame a malicious server, a replacement server is chosen randomly from among the remaining servers by an anytrust group (as the arbitrator was in Section VI-A1). Each server sends its share of the blamed server’s secret key sk to the replacement server, who can use the VSS scheme to reconstruct sk if at least t servers agree that the server is at fault. Next, to reconstruct the routing information, the ARTs and corresponding information can be resent by the other servers. We note that the replacement server may additionally need to show that it did not receive an ART during this replay and pass blame for a missing message accordingly, which it can do using the signature of the replay batch. With the keys, the replacement can decrypt and forward the messages encrypted to the old server. Note that the replacement server does not recover the batch signature key. Instead, it signs batches with a new signing key which lets other servers know that envelopes came from the replacement rather than the eliminated server.

Remark 2 (Server churn). In the event of server churn, where an (honest or adversarial) server leaves the network and stops responding, we can also apply the state recovery protocol to select a replacement and continue.

2) *Server removal and state recovery with a dishonest majority*: If there is a dishonest majority, then the vote described in Section VI-C1 can deadlock, with neither side meeting the threshold required to find a replacement. However, note that all honest servers will vote together, and only adversaries will vote incorrectly. If we partition into a new instance along the voting lines, then the partition with the honest servers (who all voted together) will have a higher fraction of honest servers than before. If we let $h := 1 - f$ be the fraction of honest servers, then we will require no more than $\lceil 1/h \rceil - 1$ restarts after voting deadlocks in order to achieve an honest majority. For example, if $1/2 < f < 2/3$, we will need at most 1 restart. This is because in order to deadlock the vote, at least $1/3$ of the servers must vote incorrectly (all servers that vote incorrectly are adversaries), who will form one partition. The other partition will then consist of at most $2/3$ of the servers, including all of the (at least $1/3$) honest servers, giving it an honest majority. Since the evidence is publicly verifiable, it can also be verified by the users, who can decide which partition to join. The honest users will evaluate the evidence the same as the honest servers, and so they will choose to send messages to the honest partition.

3) *Blocking malicious users*: Once a user has been blamed, we can revoke their keys by having each server on the user’s path reveal the user’s ARTs. Specifically, a server can produce both of the user’s ARTs and the signature binding them together. Then, we can repeat this for each server on the path (by following the identities committed in the ARTs), to remove all of the user’s tokens from the servers’ lists. We assume that the set of users is fixed, as users who join late are inherently hard to anonymize due to statistical disclosure attacks (see Section III-B). Users that are late must wait until the next restart of the protocol to join.

VII. SECURITY ANALYSIS

In this section, we analyze the security properties of Trellis. We show that (1) messages are mixed and no metadata is leaked in the process, (2) messages are guaranteed to be delivered, and (3) blame protocols eliminate adversarial parties—users or servers—without compromising anonymity.

Claim 2 (Path establishment). All honest users establish a random path through the layers and each server on the path obtains a decryption key and a signature verification key.

Proof: Boomerang encryption verifies that the anonymous routing tokens are delivered to each server along the path. The blame and recovery protocols (Section VI) ensure this process completes. In particular, note that the user establishes a key with the final anytrust group. ■

Claim 3 (Path adherence). All messages submitted by honest users are routed through all honest servers along the chosen path determined by the ARTs.

Proof: Each honest user follows the ART signing protocol and obtains valid ARTs for each link. Suppose, towards contradiction, that the envelope sent by the user does **not** travel through an honest server, say S_ℓ , as dictated by the ART on the ℓ th layer. The presence of a correctly signed message in the final anytrust group at the last layer implies that all preceding onion layers were decrypted correctly. In turn, this means that S_ℓ decrypted a layer, which is a contradiction. ■

Claim 4 (Metadata-private shuffling). All envelopes sent between two honest servers on a link appear indistinguishable to an external observer.

Proof: The dummy messages ensure that each batch of envelopes exchanged between two honest servers is of fixed size B (and TLS prevents the network adversary from seeing the contents of the batch). Crucially, no malicious server or user can overload an honest-honest link such that there are more than B envelopes. This holds because (1) the envelopes expected by each server are uniquely associated with ARTs provided during path establishment and (2) no honest server would forward any envelope not uniquely associated with an ART. By (1) and (2), the size of the batch sent to the second honest server is always B (padded with dummies as needed). ■

Claim 5 (User anonymity). All messages contributed by honest users are mixed, resulting in a random permutation of all messages being output to the public bulletin board.

Proof: We combine Claim 3 and Claim 4. Since all honest messages travel through all honest servers, they travel through all honest-honest links. Since the messages travel along honest-honest links, these messages are shuffled. Then, by Section V-E all honest messages become unlinkable with their senders as the total variation distance of the applied permutation is ϵ -close to a truly random permutation of all messages. ■

Claim 6. An honest user is never blamed as a result of any blame protocol described in Section VI, except with negligible probability in the computational security parameter λ .

Proof: A user is only blamed if either (1) there exist two envelopes for the same round and layer (Section VI) or (2) if there is an ill-formed envelope at some layer ℓ (Section VI-A3). Because an honest user’s envelopes are always well-formed, (1) and (2) do not apply to honest users. Furthermore, a user cannot be blamed by a malicious invocation of Section VI-C3 (user removal) because each ART used to trace back the path to the user is signed by the verification key in the next ART, and an honest user would never sign a different user’s ART. Hence, if an honest user is blamed then a malicious party was capable of forging the honest user’s signature, which can only happen with negligible probability. ■

Claim 7. An honest server is never blamed as a result of a blame protocol described in Section VI, except with negligible probability in the computational security parameter λ .

Proof: A server is only blamed if (1) there is a (signed) batch of envelopes showing deviation from the protocol or (2) the server refuses to produce a valid proof of correct decryption. (1) and (2) do not apply to honest servers who follow the protocol. Hence, if an honest server is blamed then a malicious party was capable of forging the honest server’s signature, which can only happen with negligible probability. ■

VIII. IMPLEMENTATION AND EVALUATION

Implementation. We implement Trellis in Go 1.17. Our implementation is open source [1] and consists of approximately 11,000 lines of code. We implement distributed key generation [54] and proactive secret-sharing [30] using the Kyber library [29]. We use the BLS12-381 [12, 111] elliptic curve to implement ARTs. Symmetric-key encryption is instantiated using AES [23, 81] and digital signatures using EdDSA [52]. Diffie-Hellman key agreement is performed over the ed25519 elliptic curve to match EdDSA. For our empirical comparison to Atom, we use their open source implementation [57].

Environment. We distribute the servers evenly among four geographic regions: Oregon, Virginia, Frankfurt, and Stockholm, to roughly match the distribution of Tor servers [71]. We use the general purpose `m5.xlarge` instances (quad-core; 16 GiB of RAM). We measured median round trip ping times of around 150 ms, with rare latency spikes of up to one second.

Due to high international networking costs, we run most of our evaluations on a simulated network environment informed from a real network deployment. In Figure 7, we report the latency of Trellis (broadcast round) under a 100 Mbps and 200 Mbps real network and compare latency to our simulated

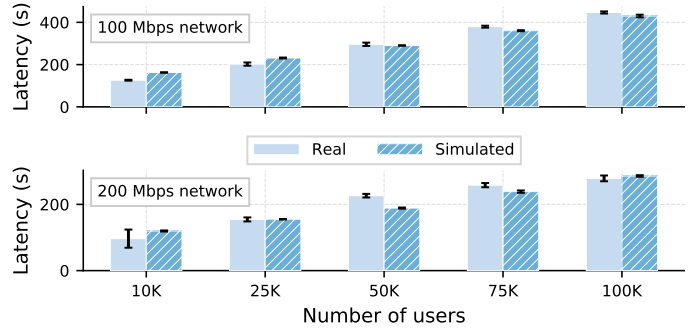


Fig. 7: Comparison between broadcasting round latency on an international (Oregon, Virginia, Frankfurt, and Stockholm) network and our simulated network. **Top:** 100 Mbps bandwidth connections. **Bottom:** 200 Mbps bandwidth connections.

network deployment. We ensure that the simulated network accurately approximates the real network by evaluating the broadcasting rounds (on pre-established paths). We observe modest differences in latency (less than 20%) between the real and simulated networks.

We use the Linux NetEm [47] tool to model latency and packet loss rates for each geographic region. We apply a Pareto distribution to simulate the infrequent, but large, latency spikes observed in the real network. Properly simulating these latency spikes is important because the all-to-all nature of the mixnet in Trellis is sensitive to the *highest* latency observed across all links. As such, tail latencies can significantly impact performance and lead to high variance in broadcast latency.

Systems optimizations. In the full version, we describe a few basic system-level and network optimizations that we incorporate into our implementation [64].

Parameters. We set $\epsilon := 2^{-64}$ for the total variational distance of the mixing. (We also set the size of each anytrust group so that the probability that *any* anytrust group has all adversaries is at most 2^{-64} .) Note that both of these parameters set the *statistical* security of Trellis. To simplify our evaluation, we fix the message size to 10 kB (over $60\times$ larger compared to the 160 B messages evaluated in Atom [60]). We vary the number of users M , number of servers N , and the fraction of malicious servers f (which influences the number of layers L and anytrust group size s).

A. Evaluation

The goals of our evaluation are to:

- determine the processing overheads of Trellis on the user and on the servers through microbenchmarks,
- measure the network overhead as a function of the number of users and fraction of malicious servers in the mix-net,
- evaluate the overhead of dummy envelopes between layers,
- evaluate how Trellis scales with additional servers, and
- compare Trellis to Atom—the state-of-the-art mix-net based system for metadata-private anonymous broadcast.

Computational and network overheads. The computational overhead of the broadcasting round is minimal given the relatively lightweight cryptography required: AES-decryption of




Fig. 8: **Left:** Path establishment with different fractions of malicious servers (f). **Right:** Broadcast round scaling with number of users M (messages), message sizes, and different adversary fractions. Both experiments are run on a simulated 200 Mbps network configuration with $N = 128$ servers. Shaded regions represent a 95% confidence interval (mostly invisible).




Fig. 9: Latency decreases with the number of servers added to the network (until reaching a constant overhead) with 200 Mbps bandwidth cap. Number of messages is $M = 2,000,000$ and message size is set to 1 kB. Shaded regions represent a 95% confidence interval.




Fig. 10: Dummy envelopes between layers incur a large overhead due to increased bandwidth requirements. Broadcast round (with and without dummy envelopes) evaluated on the real network with 200Mbps bandwidth cap, $f = 0.2$, and message size set to 10 kB. Shaded regions represent a 95% confidence interval.

envelopes and digital signature verification. All the computationally expensive cryptographic operations (e.g., pairings and token generation) are precomputed during path establishment. These computational workloads scale with the number of servers, as we show in Figure 9. However, the network latency incurred by routing through L layers does not scale with the number of servers and thus the scalability (for a given number of users) eventually plateaus. Furthermore, as N^2 approaches M , the overhead of dummy envelopes increases, as we must always send at least one dummy envelope between each pair of servers. This results in a limit on the horizontal scalability of Trellis. To illustrate this point further, in Figure 10, we compare a broadcast round in Trellis on a real network deployment with and without dummy envelopes. We find that dummy envelopes account for roughly $2.5\times$ in latency overhead.

Path establishment. Path establishment can be seen as running consecutive broadcasting rounds with $\ell = 1, 2, \dots, L$ layers in each consecutive round. Therefore, the asymptotic time complexity for path establishment is $O(L^2)$, and the concrete time it takes follows accordingly (see Figure 8). Path establishment can take one to five hours, depending on parameters. We note that path establishment depends only on the number of users, and not on the message size or number of subsequent broadcast rounds. We report micro-benchmarks for ARTs below:

- ART.PartSign: 0.08 ms (per token). // Overhead on each signer.
- ART.Verify: 1.09 ms (per token). // Token signature verification.

Broadcasting rounds. Once the path establishment completes, broadcasting rounds can be run indefinitely to broadcast messages (even following server churn and elimination). In Figure 8, we report the latency of the broadcasting rounds on four different message sizes (ranging between 1 kB and 20 kB) and fraction of adversarial servers. The broadcasting rounds are $10\text{-}120\times$ faster compared to a path establishment round. With $f = 0.2$ and 100,000 users, broadcasting a 10 kB message incurs approximately four minutes of latency. We report micro-benchmarks for the decryption and signing (performed by servers on each link):

- AES.Dec: $16\ \mu\text{s}$ for 10 kB. // Envelope decryption.
- DS.Verify: $71\ \mu\text{s}$ for 10 kB. // Signature verification.

Blame protocols. The primary computational overhead of blame protocols is generating a proof-of-decryption. In Trellis, this is done using a discrete-log equivalence proof (DLEQ), as explained in Section VI. We run a microbenchmark to determine the proving (and verification) time of each DLEQ:

- DLEQ.Prove: $152\ \mu\text{s}$. // Proving signing and correct decryption.
- DLEQ.Verify: $310\ \mu\text{s}$. // Proof verification done by server (or user).

The low overheads of the DLEQ proofs make our blame protocols lightweight. Similarly, arbitration (in case of a dispute) incurs one extra hop, which translates to an overall latency increase of 150 ms (in our network configuration).

Throughput. In Table II, we report the throughput (in bits per second) with 32, 64, and 128 servers and 100,000 users. The throughput decreases with more users and increases with more servers but eventually plateaus (see Figure 9 for scaling). In comparison, Tor achieves around 400,000 bits per second [71].

Trellis Throughput (bits/s)			
	32	64	128
$f = 0.1$	234 ± 9	435 ± 6	534 ± 15
$f = 0.2$	169 ± 3	220 ± 4	324 ± 8
$f = 0.3$	84 ± 1	168 ± 2	222 ± 5

TABLE II: Throughput of Trellis’s broadcast as a function of the adversary fraction f and number of servers for 20 kB messages. Includes 95% confidence interval.

B. Comparison to Atom

To the best of our knowledge, Atom [60] is the state-of-the-art system for metadata-private anonymous broadcast. Compared to Atom, Trellis achieves four orders of magnitude lower latency (and proportionally higher throughput). Trellis scales easily to support large message sizes (e.g., in the kilobytes); in contrast, Atom does not, due to the zero-knowledge proofs and message passing within each anytrust group. We evaluate Atom on 32 B messages and multiply out to obtain the estimate time required to broadcast a 1 MB file (Kwon et al. [60] explicitly state that latency grows linearly with larger messages). We report the result for varying number of users and match Atom’s $f = 0.2$ in Figure 11.




Fig. 11: Latency of Trellis vs. Atom to broadcast a 1 MB file. We vary the number of users and set $f = 0.2$ for both systems (we run Atom on 32 B messages and extrapolate) with 200 Mbps bandwidth cap.

C. Discussion and practical considerations

By decoupling path establishment and expensive operations from the message size, Trellis handles large messages better than all existing metadata-private anonymous broadcast systems, especially as the number of users increases. While Trellis is targeted to asynchronous applications (e.g., email, document uploads, surveys, etc.) due to concrete performance limitations (notably, high latency), further reducing latency and bandwidth overheads would require sacrificing security guarantees. By assuming that an adversary has a full view of the entire network, guaranteeing metadata privacy is imperative to preserving user anonymity. All metadata-private systems, including Trellis, will inherently require high bandwidth overheads so as to maintain a uniform view of all network traffic.

In Trellis, to achieve metadata privacy, we must introduce dummy envelopes between servers and fix the set of participating users. Our evaluation in Section VIII demonstrates that Trellis pushes metadata-private anonymous broadcast to its limit by being bandwidth—rather than computationally—bottlenecked. As a consequence, any further improvements in performance (beyond minor optimizations) will require sacrificing metadata privacy. Specifically, there are three possible

security compromises that can be made: (1) removing dummy envelopes between layers, (2) allowing user churn, and (3) reducing the number of layers in the mixnet. All three of these compromises are made by systems like Tor so as to scale to millions of daily users and achieve *real-time* (synchronous) communication (e.g., web browsing). We briefly elaborate on the consequences of applying (1) and (2) on the practical security guarantees of Trellis. Applying (3) would result in messages not being mixed and therefore voids all guarantees.

- Not using dummy envelopes would maintain all the guarantees of Section III-B, except for metadata privacy. In return, however, removing dummy envelopes would reduce the bandwidth overhead on the servers and improves the overall scalability of Trellis (see Figure 10). Of course, in practice it may be reasonable to assume a weaker adversary that is only capable of corrupting a subset of the servers and without a view of the entire network. Under this (weaker) threat model, the metadata privacy guarantee of Trellis can be sacrificed in favor of performance (without impacting the cryptographic anonymity guarantees of Trellis).
- Choosing to let users come and go makes Trellis (and all anonymous communication systems) vulnerable to intersection attacks [24, 25, 110]. Online (or offline) users correlate themselves with the message output patterns, making it possible to link messages back to the set of users that sent them. In practice, however, with sufficiently many users and a high churn rate (e.g., 50% of users churn in each round) the effectiveness of such attacks can be mitigated. Other methods, such as buddy sets [110], can also be applied to reduce the effectiveness of intersection attacks.

Future work. The theoretical mixing analysis of Kwon [62] and our analysis in Section V-E assumes that the same number of envelopes are exchanged on each link, hence servers need to pad each envelope batch with dummy envelopes. One direction for improving the performance of Trellis would be to incorporate *different sized* batches into the mixing analysis and avoid adding dummy envelopes in Trellis (at the cost of increasing the number of layers). Doing so may reduce latency and improve the horizontal scalability of Trellis, as illustrated in Figure 10. Another avenue could be proving a tighter bound on the number of layers required for mixing. Additionally, eliminating the need for digital signatures (currently required for our blame protocols) would concretely improve efficiency of mixing by a small factor, which may have a cumulative impact when the number of layers becomes very large. For example, it is conceivable that message authentication codes (MACs) [11] could be suitable replacements to our digital signatures and result in smaller boomerang envelopes.

IX. CONCLUSIONS

Scalability and support for large messages plays a key role in many applications. For example, web browsing requires megabytes of data transfer (the median web page in 2021 was 2MB in size [97]). Additionally, in the real world, scalability and concrete efficiency play an important role in *anonymity* as well: the more users there are, the larger the anonymity set, and the more plausible deniability each user obtains.

Trellis is the first system to support large messages and horizontal scalability, advancing the state-of-the-art for anony-

mous broadcast. Our prototype of Trellis is potentially efficient enough to deploy at a small “enterprise” scale with a few hundred thousand users. In such a deployment, Trellis would incur a few hours of overhead to broadcast megabyte-sized files (e.g., PDF documents). There is still an efficiency gap to bridge in terms of performance when compared to systems like Tor and I2P, which provide weaker anonymity guarantees but better performance. However, Trellis is a step toward closing this gap.

ACKNOWLEDGEMENTS

We thank Jun Wan for helping us come up with the blame replacement strategy and answering many questions we had related to consensus. We thank Kyle Hogan, Manon Revel, and Mayuri Sridhar for reading several drafts of this paper and giving us insightful feedback on how to improve it. We thank Albert Kwon for many helpful discussions pertaining to Atom and Quark as well as giving us useful pointers for improving the mixing analysis. Finally, we would like to thank the NDSS 2023 reviewers and our shepherd Qiang Tang for many insightful comments and valuable suggestions.

REFERENCES

- [1] Source code for Trellis. <https://github.com/SimonLangowski/trellis>, 2022.
- [2] Ittai Abraham, Benny Pinkas, and Avishay Yanai. Blinder: Scalable, robust anonymous committed broadcast. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security, CCS '20*, pages 1233–1252, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450370899. doi: 10.1145/3372297.3417261. URL <https://doi.org/10.1145/3372297.3417261>.
- [3] Ben Adida and Douglas Wikström. Offline/online mixing. In *International Colloquium on Automata, Languages, and Programming*, pages 484–495. Springer, 2007.
- [4] Ishtiyaque Ahmad, Yuntian Yang, Divyakant Agrawal, Amr El Abbadi, and Trinabh Gupta. Adra: Metadata-private voice communication over fully untrusted infrastructure. In *15th USENIX Symposium on Operating Systems Design and Implementation (OSDI 21)*, 2021.
- [5] Nikolaos Alexopoulos, Aggelos Kiayias, Riivo Talviste, and Thomas Zacharias. MCMix: Anonymous messaging via secure multiparty computation. In *26th USENIX Security Symposium (USENIX Security 17)*, pages 1217–1234, 2017.
- [6] Sebastian Angel and Srinath Setty. Unobservable communication over fully untrusted infrastructure. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 551–569, 2016.
- [7] Ludovic Barman, Italo Dacosta, Mahdi Zamani, Ennan Zhai, Apostolos Pyrgelis, Bryan Ford, Joan Feigenbaum, and Jean-Pierre Hubaux. Prifi: Low-latency anonymity for organizational networks. *Proc. Priv. Enhancing Technol.*, 2020(4):24–47, 2020. doi: 10.2478/popets-2020-0061. URL <https://doi.org/10.2478/popets-2020-0061>.
- [8] Ehrhard Behrends. *Introduction to Markov chains : with special emphasis on rapid mixing*. Advanced lectures in mathematics. Vieweg, Braunschweig/Wiesbaden, 2000. ISBN 3528069864.
- [9] Sanjit Bhat, David Lu, Albert Kwon, and Srinivas Devadas. Var-CNN: A data-efficient website fingerprinting attack based on deep learning. *Proceedings on Privacy Enhancing Technologies*, 2019(4):292–310, 2019.
- [10] Alexandra Boldyreva. Threshold signatures, multisignatures and blind signatures based on the gap-Diffie-Hellman-group signature scheme. In *International Workshop on Public Key Cryptography*, pages 31–46. Springer, 2003.
- [11] Dan Boneh and Victor Shoup. A graduate course in applied cryptography. *Draft 0.5*, 2020.
- [12] Dan Boneh, Ben Lynn, and Hovav Shacham. Short signatures from the Weil pairing. In *International conference on the theory and application of cryptology and information security*, pages 514–532. Springer, 2001.
- [13] Nikita Borisov, George Danezis, Prateek Mittal, and Parisa Tabriz. Denial of service or denial of security? In *Proceedings of the 14th ACM Conference on Computer and Communications Security*, pages 92–102, 2007.
- [14] Bryan Burrough, Sarah Ellison, and Suzanna Andrews. The Snowden saga: A shadowland of secrets and light. *Vanity Fair*, 2014. URL <https://www.vanityfair.com/news/politics/2014/05/edward-snowden-politics-interview>. Accessed September 2022.
- [15] David Chaum, Debajyoti Das, Farid Javani, Aniket Kate, Anna Krasnova, Joeri De Ruiter, and Alan T Sherman. cMix: Mixing with minimal real-time asymmetric cryptographic operations. In *International conference on applied cryptography and network security*, pages 557–578. Springer, 2017.
- [16] David L Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–90, 1981.
- [17] Raymond Cheng, William Scott, Elisaweta Masserova, Irene Zhang, Vipul Goyal, Thomas Anderson, Arvind Krishnamurthy, and Bryan Parno. Talek: Private group messaging with hidden access patterns. In *Annual Computer Security Applications Conference*, pages 84–99, 2020.
- [18] Cloudflare. Comprehensive DDoS protection. <https://www.cloudflare.com/ddos/>, 2021. Accessed September 2022.
- [19] Henry Corrigan-Gibbs and Bryan Ford. Dissent: Accountable anonymous group messaging. In *Proceedings of the 17th ACM Conference on Computer and Communications Security*, pages 340–350. ACM, 2010.
- [20] Henry Corrigan-Gibbs, Dan Boneh, and David Mazières. Riposte: An anonymous messaging system handling millions of users. In *2015 IEEE Symposium on Security and Privacy*, pages 321–338. IEEE, 2015.
- [21] Cybersecurity and Infrastructure Security Agency. Security tip (ST04-015): Understanding denial-of-service attacks. <https://www.cisa.gov/uscert/ncas/tips/ST04-015/>, 2021. Accessed September 2022.
- [22] Artur Czumaj and Berthold Vöcking. Thorp shuffling, butterflies, and non-Markovian couplings. In *International Colloquium on Automata, Languages, and Programming*, pages 344–355. Springer, 2014.

- [23] Joan Daemen and Vincent Rijmen. *The design of Rijndael: AES — the Advanced Encryption Standard*. Springer-Verlag, 2002. ISBN 3-540-42580-2.
- [24] George Danezis. Statistical disclosure attacks. In *IFIP International Information Security Conference*, pages 421–426. Springer, 2003.
- [25] George Danezis and Andrei Serjantov. Statistical disclosure or intersection attacks on anonymity systems. In *International Workshop on Information Hiding*, pages 293–308. Springer, 2004.
- [26] George Danezis, Roger Dingledine, and Nick Mathewson. Mixminion: Design of a type III anonymous remailer protocol. In *2003 Symposium on Security and Privacy, 2003.*, pages 2–15. IEEE, 2003.
- [27] Debajyoti Das, Sebastian Meiser, Esfandiar Mohammadi, and Aniket Kate. Divide and funnel: a scaling technique for mix-networks. *Cryptology ePrint Archive*, 2021.
- [28] Alex Davidson, Ian Goldberg, Nick Sullivan, George Tankersley, and Filippo Valsorda. Privacy Pass: Bypassing internet challenges anonymously. *Proc. Priv. Enhancing Technol.*, 2018(3):164–180, 2018.
- [29] DEDIS. Dedis advanced crypto library for Go. <https://github.com/dedis/kyber>, 2022. Accessed August 2022.
- [30] Yvo Desmedt and Sushil Jajodia. Redistributing secret shares to new access structures and its applications, 1997.
- [31] Whitfield Diffie and Martin Hellman. New directions in cryptography. *IEEE transactions on Information Theory*, 22(6):644–654, 1976.
- [32] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. Technical report, Naval Research Lab Washington DC, 2004.
- [33] Manu Drijvers, Kasra Edalatnejad, Bryan Ford, Eike Kiltz, Julian Loss, Gregory Neven, and Igors Stepanovs. On the security of two-round multi-signatures. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 1084–1101, 2019. doi: 10.1109/SP.2019.00050.
- [34] Saba Eskandarian. Fast privacy-preserving punch cards. *Proceedings on Privacy Enhancing Technologies*, 2021(3):289–307, 2021. doi: doi:10.2478/popets-2021-0048. URL <https://doi.org/10.2478/popets-2021-0048>.
- [35] Saba Eskandarian and Dan Boneh. Clarion: Anonymous communication from multiparty shuffling protocols. In *Proceedings of the Network and Distributed Systems Security Symposium*, 2022.
- [36] Saba Eskandarian, Henry Corrigan-Gibbs, Matei Zaharia, and Dan Boneh. Express: Lowering the cost of metadata-hiding communication with cryptographic privacy. In *30th USENIX Security Symposium (USENIX Security 21)*, Vancouver, B.C., August 2021. USENIX Association.
- [37] Nathan S Evans, Roger Dingledine, and Christian Grothoff. A practical congestion attack on Tor using long paths. In *USENIX Security Symposium*, pages 33–50, 2009.
- [38] Ernesto Falcon. The FCC must update ISP privacy rules. <https://www.eff.org/deeplinks/2016/05/fcc>, 2016. Accessed September 2022.
- [39] Paul Feldman. A practical scheme for non-interactive verifiable secret sharing. In *28th Annual Symposium on Foundations of Computer Science (sfcs 1987)*, pages 427–438. IEEE, 1987.
- [40] Michael J Freedman and Robert Morris. Tarzan: A peer-to-peer anonymizing network layer. In *Proceedings of the 9th ACM Conference on Computer and Communications Security*, pages 193–206, 2002.
- [41] Nethanel Gelernter, Amir Herzberg, and Hemi Lebowitz. Two cents for strong anonymity: The anonymous post-office protocol. In *International Conference on Cryptology and Network Security*, pages 390–412. Springer, 2017.
- [42] Yossi Gilad. Metadata-private communication for the 99%. *Communications of the ACM*, 62(9):86–93, 2019.
- [43] Oded Goldreich. *Foundations of cryptography: a primer*, volume 1. Now Publishers Inc, 2005.
- [44] Philippe Golle and Ari Juels. Parallel mixing. In *Proceedings of the 11th ACM conference on Computer and communications security*, pages 220–226, 2004.
- [45] Maohua Guo and Jinlong Fei. Website fingerprinting attacks based on homology analysis. *Security and Communication Networks*, 2021, 2021.
- [46] Johan Håstad. The square lattice shuffle. *Random Structures and Algorithms*, 29(4):466–474, 2006.
- [47] Stephen Hemminger et al. Network emulation with NetEm. In *Linux conf au*, volume 5, page 2005. Citeseer, 2005.
- [48] Amir Herzberg, Stanisław Jarecki, Hugo Krawczyk, and Moti Yung. Proactive secret sharing or: How to cope with perpetual leakage. In *annual international cryptology conference*, pages 339–352. Springer, 1995.
- [49] Nicholas Hopper, Eugene Y Vasserman, and Eric Chan-Tin. How much anonymity does network latency leak? *ACM Transactions on Information and System Security (TISSEC)*, 13(2):1–28, 2010.
- [50] Stanisław Jarecki, Aggelos Kiayias, and Hugo Krawczyk. Round-optimal password-protected secret sharing and T-PAKE in the password-only model. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 233–253. Springer, 2014.
- [51] Rebecca Jeschke. Internet advocates call on ISPs to commit to basic user privacy protections. <https://www.eff.org/deeplinks/2021/03/internet-advocates-call-isps-commit-basic-user-privacy-protections>, 2021. Accessed September 2022.
- [52] Simon Josefsson and Ilari Liusvaara. Edwards-curve digital signature algorithm (EdDSA). *RFC*, 8032:1–60, 2017.
- [53] Antoine Joux. A one round protocol for tripartite Diffie–Hellman. In *International algorithmic number theory symposium*, pages 385–393. Springer, 2000.
- [54] Aniket Kate and Ian Goldberg. Distributed private-key generators for identity-based cryptography. In *International Conference on Security and Cryptography for Networks*, pages 436–453. Springer, 2010.
- [55] Jonathan Katz and Yehuda Lindell. *Introduction to modern cryptography*. CRC press, 2020.
- [56] Ben Kreuter, Tancrede Lepoint, Michele Orrù, and Mariana Raykova. Anonymous tokens with private metadata bit. In *Annual International Cryptology Conference*, pages 308–336. Springer, 2020.
- [57] Albert Kwon. Source code for Atom. <https://github.com/kwonalbert/atom>, 2017. Accessed August 2022.

- [58] Albert Kwon, Mashael AlSabah, David Lazar, Marc Dacier, and Srinivas Devadas. Circuit fingerprinting attacks: Passive deanonymization of Tor hidden services. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 287–302, 2015.
- [59] Albert Kwon, David Lazar, Srinivas Devadas, and Bryan Ford. Riffle: An efficient communication system with strong anonymity. *Proc. Priv. Enhancing Technol.*, 2016 (2):115–134, 2016.
- [60] Albert Kwon, Henry Corrigan-Gibbs, Srinivas Devadas, and Bryan Ford. Atom: Horizontally scaling strong anonymity. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pages 406–422, 2017.
- [61] Albert Kwon, David Lu, and Srinivas Devadas. XRD: Scalable messaging system with cryptographic privacy. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*, pages 759–776, 2020.
- [62] Young Hyun Kwon. *Towards anonymous and metadata private communication at Internet scale*. PhD thesis, Massachusetts Institute of Technology, 2019.
- [63] Leslie Lamport. Paxos made simple. *ACM SIGACT News (Distributed Computing Column) 32, 4 (Whole Number 121, December 2001)*, pages 51–58, 2001.
- [64] Simon Langowski, Sacha Servan-Schreiber, and Srinivas Devadas. Trellis: Robust and scalable metadata-private anonymous broadcast. *Cryptology ePrint Archive*, 2022.
- [65] David Lazar, Yossi Gilad, and Nickolai Zeldovich. Karaoke: Distributed private messaging immune to passive traffic analysis. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*, pages 711–725, 2018.
- [66] David Lazar, Yossi Gilad, and Nickolai Zeldovich. Yodel: Strong metadata security for voice calls. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles*, pages 211–224, 2019.
- [67] Stevens Le Blond, David Choffnes, Wenxuan Zhou, Peter Druschel, Hitesh Ballani, and Paul Francis. Towards efficient traffic-analysis resistant anonymity networks. *ACM SIGCOMM Computer Communication Review*, 43 (4):303–314, 2013.
- [68] Stevens Le Blond, David Choffnes, Wenxuan Zhou, Peter Druschel, Hitesh Ballani, and Paul Francis. Towards efficient traffic-analysis resistant anonymity networks. *ACM SIGCOMM Computer Communication Review*, 43 (4):303–314, 2013.
- [69] Stevens Le Blond, David Choffnes, William Caldwell, Peter Druschel, and Nicholas Merritt. Herd: A scalable, traffic analysis resistant anonymity network for VoIP systems. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*, pages 639–652, 2015.
- [70] Brian N Levine, Michael K Reiter, Chenxi Wang, and Matthew Wright. Timing attacks in low-latency mix systems. In *International Conference on Financial Cryptography*, pages 251–265. Springer, 2004.
- [71] Karsten Loesing, Steven J. Murdoch, and Roger Dingledine. A case study on measuring statistical data in the Tor anonymity network. In *Proceedings of the Workshop on Ethics in Computer Security Research (WECSR 2010)*, LNCS. Springer, January 2010.
- [72] Donghang Lu and Aniket Kate. Rpm: Robust anonymity at scale. *Cryptology ePrint Archive*, 2022.
- [73] Ewen MacAskill and Gabriel Dance. NSA files: decoded. *The Guardian*, 1, 2013.
- [74] Sai Krishna Deepak Maram, Fan Zhang, Lun Wang, Andrew Low, Yupeng Zhang, Ari Juels, and Dawn Song. Churp: dynamic-committee proactive secret sharing. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 2369–2386, 2019.
- [75] Nick Mathewson and Roger Dingledine. Practical traffic analysis: Extending and resisting statistical disclosure. In *International Workshop on Privacy Enhancing Technologies*, pages 17–34. Springer, 2004.
- [76] Prateek Mittal and Nikita Borisov. Shadowwalker: peer-to-peer anonymous communication using redundant structured topologies. In *Proceedings of the 16th ACM conference on Computer and communications security*, pages 161–172, 2009.
- [77] Prateek Mittal, Ahmed Khurshid, Joshua Juen, Matthew Caesar, and Nikita Borisov. Stealthy traffic analysis of low-latency anonymous communication using throughput fingerprinting. In *Proceedings of the 18th ACM conference on Computer and communications security*, pages 215–226, 2011.
- [78] Zachary Newman, Sacha Servan-Schreiber, and Srinivas Devadas. Spectrum: High-bandwidth anonymous broadcast. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, pages 229–248, 2022.
- [79] Lan Nguyen and Rei Safavi-Naini. Breaking and mending resilient mix-nets. In *International Workshop on Privacy Enhancing Technologies*, pages 66–80. Springer, 2003.
- [80] Jonas Nick, Tim Ruffing, and Yannick Seurin. MuSig2: simple two-round Schnorr multi-signatures. In *Annual International Cryptology Conference*, pages 189–221. Springer, 2021.
- [81] NIST. Specification for the Advanced Encryption Standard (AES). Federal Information Processing Standards Publication 197, 2001. URL <http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf>.
- [82] Se Eun Oh, Nate Mathews, Mohammad Saidur Rahman, Matthew Wright, and Nicholas Hopper. GANDaLF: GAN for data-limited fingerprinting. *Proc. Priv. Enhancing Technol.*, 2021(2):305–322, 2021.
- [83] Diego Ongaro and John Ousterhout. In search of an understandable consensus algorithm. In *2014 USENIX Annual Technical Conference (USENIX ATC 14)*, pages 305–319, 2014.
- [84] Lasse Overlier and Paul Syverson. Locating hidden servers. In *2006 IEEE Symposium on Security and Privacy (S&P’06)*, pages 15–114. IEEE, 2006.
- [85] Ania M Piotrowska. *Low-latency mix networks for anonymous communication*. PhD thesis, UCL (University College London), 2020.
- [86] Ania M Piotrowska, Jamie Hayes, Tariq Elahi, Sebastian Meiser, and George Danezis. The loopix anonymity system. In *26th USENIX Security Symposium (USENIX Security 17)*, pages 1199–1216, 2017.
- [87] David Pointcheval and Jacques Stern. Provably secure blind signature schemes. In *International Conference on the Theory and Application of Cryptology and Infor-*

- mation Security, pages 252–265. Springer, 1996.
- [88] Tal Rabin and Michael Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority. In *Proceedings of the twenty-first annual ACM symposium on Theory of computing*, pages 73–85, 1989.
- [89] Charles Rackoff and Daniel R Simon. Cryptographic defense against traffic analysis. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 672–681, 1993.
- [90] Michael K Reiter and Aviel D Rubin. Anonymous web transactions with crowds. *Communications of the ACM*, 42(2):32–48, 1999.
- [91] Eric Rescorla and Tim Dierks. The transport layer security (TLS) protocol version 1.3. *RFC*, 2018.
- [92] David Schatz, Michael Rossberg, and Guenter Schaefer. Hydra: Practical metadata security for contact discovery, messaging, and dialing. In *ICISSP*, pages 191–203, 2021.
- [93] Adam Schwartz, Andrew Crocker, and Kit Walsh. EFF to court: Broadband privacy law passes first amendment muster. <https://www.eff.org/deeplinks/2020/05/eff-court-broadband-privacy-law-passes-first-amendment-muster>, 2020. Accessed September 2022.
- [94] Vitaly Shmatikov and Ming-Hsiu Wang. Timing analysis in low-latency mix networks: Attacks and defenses. In *European Symposium on Research in Computer Security*, pages 18–33. Springer, 2006.
- [95] Robert Shostak, Marshall Pease, and Leslie Lamport. The byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982.
- [96] Payap Sirinam, Nate Mathews, Mohammad Saidur Rahman, and Matthew Wright. Triplet fingerprinting: More practical and portable website fingerprinting with n-shot learning. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 1131–1148, 2019.
- [97] John Teague. 2021 Page Weight: The web almanac by HTTP Archive. <https://almanac.httparchive.org/en/2021/page-weight>, 2021. HTTP Archive.
- [98] The Invisible Internet Project. I2P anonymous network, 2022. URL <https://geti2p.net/en/>. Accessed September 2022.
- [99] The New York Times. Got a confidential news tip? <https://www.nytimes.com/tips>, 2022. Accessed September 2022.
- [100] The Wall Street Journal. Got a tip? <https://www.wsj.com/tips>, 2022. Accessed September 2022.
- [101] Nirvan Tyagi, Yossi Gilad, Derek Leung, Matei Zaharia, and Nikolai Zeldovich. Stadium: A distributed metadata-private messaging system. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pages 423–440, 2017.
- [102] Nirvan Tyagi, Sofía Celi, Thomas Ristenpart, Nick Sullivan, Stefano Tessaro, and Christopher A Wood. A fast and simple partially oblivious PRF, with applications. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 674–705. Springer, 2022.
- [103] Adithya Vadapalli, Kyle Storrier, and Ryan Henry. Sabre: Sender-anonymous messaging with fast audits. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 1953–1970. IEEE, 2022.
- [104] Jelle Van Den Hooff, David Lazar, Matei Zaharia, and Nikolai Zeldovich. Vuvuzela: Scalable private messaging resistant to traffic analysis. In *Proceedings of the 25th Symposium on Operating Systems Principles*, pages 137–152, 2015.
- [105] Von Spiegel Staff. Inside the NSA’s war on internet security. *Der Spiegel*, 2014. URL <https://www.spiegel.de/international/germany/inside-the-nsa-s-war-on-internet-security-a-1010361.html>. Accessed September 2022.
- [106] Abraham Waksman. A permutation network. *Journal of the ACM (JACM)*, 15(1):159–163, 1968.
- [107] Jun Wan, Hanshen Xiao, Elaine Shi, and Srinivas Devadas. Expected constant round byzantine broadcast under dishonest majority. In *Theory of Cryptography Conference*, pages 381–411. Springer, 2020.
- [108] Douglas Wikström. Five practical attacks for “optimistic mixing for exit-polls”. In *International Workshop on Selected Areas in Cryptography*, pages 160–174. Springer, 2003.
- [109] David Isaac Wolinsky, Henry Corrigan-Gibbs, Bryan Ford, and Aaron Johnson. Dissent in numbers: Making strong anonymity scale. In *10th USENIX Symposium on Operating Systems Design and Implementation (OSDI 12)*, pages 179–182, 2012.
- [110] David Isaac Wolinsky, Ewa Syta, and Bryan Ford. Hang with your buddies to resist intersection attacks. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, pages 1153–1166, 2013.
- [111] Shoko Yonezawa, Tetsutaro Kobayashi, and Tsunekazu Saito. Pairing-friendly curves. *Network Working Group. Internet-Draft. January*, 2019.