

# Pudding: Private User Discovery in Anonymity Networks

Ceren Kocaoğullar  
University of Cambridge

Daniel Hugenroth  
University of Cambridge

Martin Kleppmann  
TU Munich

Alastair R. Beresford  
University of Cambridge

**Abstract**—Anonymity networks allow messaging with metadata privacy, providing better privacy than popular encrypted messaging applications. However, contacting a user on an anonymity network currently requires knowing their public key or similar high-entropy information, as these systems lack a privacy-preserving mechanism for contacting a user via a short, human-readable username. Previous research suggests that this is a barrier to widespread adoption.

In this paper we propose Pudding, a novel private user discovery protocol that allows a user to be contacted on an anonymity network knowing only their email address. Our protocol hides contact relationships between users, prevents impersonation, and conceals which usernames are registered on the network. Pudding is Byzantine fault tolerant, remaining available and secure as long as less than one third of servers are crashed, unavailable, or malicious. It can be deployed on Loopix and Nym without changes to the underlying anonymity network protocol, and it supports mobile devices with intermittent network connectivity. We demonstrate the practicality of Pudding with a prototype using the Nym anonymity network. We also formally define the security and privacy goals of our protocol and conduct a thorough analysis to assess its compliance with these definitions.

## 1. Introduction

As more of our conversations take place online, protecting them becomes increasingly important. As a result, we have seen widespread deployment of end-to-end encryption in messaging apps in order to protect the confidentiality and integrity of our correspondence. Today’s secure messaging apps like Signal, iMessage, and WhatsApp allow users to contact their friends using only a low-entropy username, such as a phone number. In contrast, earlier systems such as PGP required users to explicitly manage cryptographic key material, which was a significant usability issue and a barrier to adoption [1]. It is likely that user discovery via phone number played a large part in why Signal, iMessage, and WhatsApp succeeded in gaining widespread adoption, while PGP has remained a niche tool used by experts.

While the current generation of messaging apps offer confidentiality of message contents against malicious servers and network adversaries as long as the users verify each other’s keys, these apps offer little *metadata privacy*. In particular, the service provider is able to track who is talking to whom, and at what times—information that can be sensitive [2]–[5]. Even Signal’s sealed sender mechanism [6], which

aims to hide the sender of the messages from the service provider, does not provide protection at the network level. In response, significant academic progress has been made in the design of practical anonymity networks to protect metadata privacy for messaging apps, such as Loopix [7], Vuvuzela [8], Talek [9], and Pung [10]. However, despite the fact that user discovery was important for widespread adoption of secure messaging, none of these networks consider user discovery as a part of their design. Instead, they defer the problem to future work. Yet, an automated means of supporting user discovery and key management is critical if we want to build usable messaging apps that reach a wide audience and support metadata privacy, as well as end-to-end encryption. Such a system needs to allow a messaging app to securely and automatically resolve user-readable names entered in a contact address book (such as phone numbers or email addresses) into appropriate network addresses and the necessary cryptographic key material required to send encrypted messages. Furthermore, such a system needs to do this while protecting metadata privacy.

The user discovery model found in the current generation of end-to-end encrypted messaging apps is unsuitable if we wish to provide metadata privacy, because this model undermines the metadata privacy that anonymity networks provide. Specifically, encrypted messaging apps such as Signal, iMessage, and WhatsApp store user public key material in trusted key directory servers, which users query with phone numbers to retrieve the associated public keys. This enables the servers to learn which contacts in a user’s address book are registered to the messaging app (e.g. Signal [11]), and in some cases even those that are not registered (e.g. WhatsApp [12]). Communicating with a key directory via an anonymity network may provide metadata privacy against a network observer, but it does not provide protection against malicious or compromised directory servers. A compromised directory server may try to deceive a user into talking to someone other than the intended user. Moreover, an adversary can learn the phone numbers of all registered users by querying the directory [13].

A relatively small body of literature is concerned with providing user discovery for anonymity networks (see §7). Two prominent private user discovery protocols are Alpenhorn [14] and UDM [15]. Neither of these protocols can hide whether a given username is a member of the network, and they cannot maintain availability in the presence of even a single faulty server. Moreover, Alpenhorn requires users to regularly download inboxes, whether or not they are performing a lookup. Consequently, this protocol can lead

to substantial monthly bandwidth usage, posing a particular challenge for mobile devices relying on cellular networks.

In Tor, users can access onion services, previously known as hidden services [16], [17], through special URLs. A Tor onion service URL is a 56-character pseudorandom string derived from a public key [16], making it difficult to remember or communicate verbally. One can create a shortcut for this URL, for example by using a DNS alias or a service that redirects from a short and memorable URL to the hidden service URL. However, this requires the service providing that redirection to be trusted to respond correctly.

In this paper, we propose a practical and metadata-private user discovery protocol named Pudding<sup>1</sup>, which has two main privacy properties. Firstly, Pudding hides contact relationships between users, so that neither a global network observer nor a compromised discovery server can determine who is talking to whom. We build our system as an application on top of Loopix [7], which is a practical anonymity network design that has been implemented and is currently operating at significant scale, with hundreds of mix nodes provided as part of the Nym<sup>2</sup> anonymity network.

Pudding also ensures that an adversarial user cannot learn whether a given username is registered on Pudding, except when authorised by the owner of that username. We achieve this by establishing a new mechanism to generate single-use reply blocks (SURBs), which are special network packets that allow a user to send a message to another user without knowing her network address or public key [18], [19]. This mechanism allows us to distribute the discovery operations among multiple servers (see §4.1).

Pudding also ensures that even when up to one third of the directory servers are unavailable or return incorrect responses, a user's request to initiate contact with another user resolves correctly and the system remains operational.

Users who want to be discovered through Pudding register and validate access to an email address. This allows messaging apps to automate the process of resolving user-readable names into the required network information and cryptographic key material to establish end-to-end encrypted communication. Moreover, it provides users with reassurance that the person they are talking to is the intended recipient, since the email addresses used come from their own contact address book. Furthermore, institutional email addresses (e.g. `famous.journalist@newspaper.com`) allow the sender of a message to assure themselves that they are talking to someone at a particular organisation.

Using a non-anonymous identifier such as a phone number or an email address may appear counter-intuitive in an anonymity network. However, our design does not require users to register in order to be able to discover other users. Therefore, for instance, a journalist may publish their Pudding username as part of their social media presence, and whistleblowers can use Pudding to discover and contact the journalist securely, without having to register to Pudding with their own email address.

1. A wordplay on the abbreviation of 'private user discovery' (PUD)  
2. <https://nymtech.net/>

In this paper, we make the following contributions:

- 1) We identify the need for a new *usable* and *private* user discovery mechanism in anonymous messaging systems and define usability and privacy requirements in this context (§2).
- 2) We propose a new private user discovery protocol named *Pudding* (§4) that conceals who is talking to whom, tolerates Byzantine or unavailable servers, does not require users to be constantly online, protects the information that a certain email address belongs to a member of the network or not, and supports authenticated usernames that can be short and memorable.
- 3) We propose a new method for generating single-use reply blocks (SURBs) [18], [19], which allows the servers to collectively respond to a discovery query without needing to coordinate. By comparing the responses from several nodes we provide resilience against malicious or compromised discovery nodes, as well as concealing registered usernames.
- 4) We show that our protocol is practical by evaluating a prototype over the Nym anonymity network. Registration has a mean latency of less than 9 seconds, and end-to-end discovery takes less than 13 seconds on average. These numbers are dominated by delays in the underlying anonymity network.

## 2. Problem definition

To establish contact with a user on an anonymity network, one needs to know some network-specific information about that user. We use the term *contact information* to refer to the information that is needed to make initial contact with a user through the anonymity network. Although the content of contact information differs among anonymity network schemes, it generally includes a user's public key. In some systems, such as Talek [9] and Vuvuzela [8], the public key is used as a user ID as well. In Loopix [7], the contact information is comprised of the public key, the IP address of the user's provider, and the account ID at that provider (see §3.1). Nym is similar to Loopix, only using a Nym-specific provider ID instead of an IP address.

We define *user discovery* as initiating contact with a user based on her username, either by obtaining her contact information, or by directly sending a message to the user without revealing her contact information. We believe that user discovery should be simple, so that it is accessible to all users. Inspired by trust establishment usability definitions from Unger et al.'s work [20], we identify two key usability requirements for user discovery. A *usable* discovery mechanism in an anonymity network should:

- 1) Require only a short, memorable, and hence low-entropy username to establish contact. The username can either be a 'pseudonym' that is unique only within the system, or an identifier that has meaning outside of the anonymity network and therefore may already be in the local address books of the user's friends (e.g. a phone number or an email address).

- 2) Allow the user discovery to happen without requiring the user to leave the anonymity network, because any communication outside the anonymity network risks revealing to a network adversary the very metadata that the anonymity network is trying to hide.

A private user discovery mechanism must not undermine the level of privacy provided by the anonymity network. Therefore, a user discovery system should ensure that:

- 1) The information “Alice is searching for Bob” is kept secret from everyone other than Alice and Bob. The providers of the discovery service and anonymity network, and any active or passive network adversaries, should not learn which users are communicating.
- 2) A user should not be able to learn whether a given username is registered, except by contacting the owner of that username and the recipient choosing to respond.

As we show in §7, there is no existing solution for user discovery that fits the usability and privacy needs of anonymity networks, is a lightweight design that is compatible with applications running on mobile devices, and provides fault tolerance.

### 3. System overview and goals

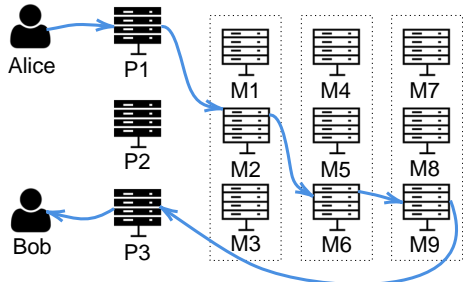


Figure 1: Diagram showing Alice sending Bob a message through a Loopix network with three layers of mix nodes ( $M1-M9$ ) and three provider nodes ( $P1, P2, P3$ ). The arrows show a possible route for Alice’s message. Each message passes through one mix node from each layer. Alice and Bob access the network via provider nodes, which are connected to all the mix nodes in the outermost layers.

#### 3.1. Setup

**Loopix network architecture.** Pudding builds upon the Loopix [7] design. Loopix is a prominent medium-latency mix network [21] that protects sender-receiver relationships from third parties, including global passive and active adversaries, enabling metadata-private email or instant messaging applications. Loopix hides communication patterns using cover traffic and Poisson mixing. The network consists of three types of nodes, as illustrated in Figure 1: (1) *user devices* are associated with a given user and provide

send/receive functionality for end-user applications; (2) multiple layers of *mix nodes* transmit real and cover messages; (3) *provider nodes* manage access to the network. Each user device is registered at a provider node and can only send and receive packets through its provider. A provider node can receive messages on the user’s behalf, even when the user device is offline. Loopix network packets use the Sphinx packet format, which we describe below.

Pudding is an application-layer protocol – it does not require any modification of the anonymity network for which it provides a user discovery service. Pudding can be built upon any anonymity network that allows messaging without explicitly knowing the recipient’s contact information (see *single-use reply block* below), such as Nym. In this paper, we build Pudding on Nym/Loopix, because this anonymity network design provides a high level of metadata privacy, is suitable for asynchronous messaging, and is designed to be feasible to use on mobile devices.

**Pudding nodes.** For our protocol, we introduce a new type of server, the *discovery node*, which provides the user discovery service. Discovery nodes can be independently managed and maintained, and no single entity should control more than one-third of the discovery nodes (see §3.2). There is a fixed set of  $n$  discovery nodes, where  $n = 4$  or  $n = 7$  would be reasonable choices (see *Security parameters* in §4). Each discovery node has a database that maps usernames to contact information. From a Loopix protocol point of view, discovery nodes act as user devices (although in practice, they may be co-located with mix nodes or provider nodes). This allows discovery nodes to be added to the network without changes to the Loopix protocol, without disturbing its security properties.

**Sphinx packet format.** In Pudding, user devices and discovery nodes communicate via Loopix (i.e. by routing each message through all mix layers) using the Sphinx [19] packet format. The one exception to this rule is the validation of email addresses, as discussed in §5. A Sphinx packet consists of two main components: a header and an encrypted payload. The header contains the metadata necessary to route and process the message correctly, i.e. to verify the integrity of the packet at each hop, to send the packet to the next node in the path, and to decrypt the outer-most payload layer. The layered encryption ensures that each node in the message path learns only the routing information to the next hop and a network observer cannot correlate a mix node’s incoming and outgoing packets. Sphinx packets are source-routed, meaning that the creator of the header decides and encodes the routing information into the packet [19].

**Single-use reply block (SURB).** Pudding uses a special Sphinx header called a *single-use reply block (SURB)* for some of its messages. A SURB is a Sphinx header that allows a user to send a message without learning the identity of the recipient [18], [19]. SURBs are typically used when user Alice wants to send a message to Bob and allow him to reply to her without revealing her own contact information. Alice constructs a Sphinx header with onion-encrypted routing information that goes through a sequence of mix nodes

and ends at herself. She can then include this SURB, which includes keys for the payload encryption, in her message to Bob. While SURBs are often used for anonymous replies to the original sender, they can be constructed by anyone with knowledge of the contact information of the destination.

**System assumptions.** We assume that all user devices know the public contact information for all discovery nodes, so that they can send anonymous messages to the discovery nodes. Provider nodes can supply this information to user devices.

### 3.2. Threat Model and Goals

**Security and privacy goals.** Our security and privacy goal definitions are inspired by AnoA [22], which is also the approach taken by Loopix [7]. In terms of notation, each user ( $U$ ) possesses a username ( $ID$ ) and contact information ( $\Delta$ ). We denote the  $ID$  and  $\Delta$  associated with a certain user  $U$  as  $(U : ID, \Delta)$ . We write  $\{U_i \rightarrow U_j\}$  to indicate that user  $U_i$  searches for user  $U_j$ . For now we state our goals informally, and we formalise them in §6.1.

**G1 Unlinkability:** The adversary cannot learn *who is searching for whom*. This goal ensures that the *sender-receiver third-party unlinkability* goal of Loopix is preserved. In other words, assuming a sufficiently large Loopix anonymity set, the adversary cannot tell the difference between  $\{U_1 \rightarrow U_*\}$  and  $\{U_2 \rightarrow U_*\}$ , for all honest users  $U_1$  and  $U_2$  chosen by the adversary. Essentially, when a user searches for  $U_*$ , the adversary learns  $U_*$  but not who sent the query.

**G2 Security against impersonation:** Assume a user  $U_1$  has registered contact information  $\Delta_1$  for a particular username  $ID_1$ . When another user subsequently looks up  $ID_1$ , the contact request should be sent to the registered contact information  $\Delta_1$  and no other user. In simpler terms, an adversary cannot cause any honest user to send a message intended for  $U_1$  to a  $\Delta_2 \neq \Delta_1$ .

**G3 External identity verification:** A user’s identity within the system should be verifiably associated with an external identity by the discovery nodes. For instance, this can be achieved by discovery nodes sending a code to an email address or phone number, assuming that the communication channel used to send the code is secure and the user’s device is trusted. This goal provides assurance to the user that she is talking to the correct person, defined as the owner of an email address. Let us assume that the user  $U$  is a user controlled by the adversary, and  $ID$  is an external identity, such as an email address at a trusted email service, that is *not* controlled by the adversary. We define external identity verification as the inability of the adversary to convince any honest discovery node into registering the user  $(U : ID, \Delta)$ .

**G4 Membership unobservability:** No user should be able to determine whether a username has been registered, unless the owner of that username decides to reveal this information. Membership unobservability is important

in situations where exposing the fact that someone is a member of a network may harm the user. For a username  $ID$  chosen by the adversary, we define *membership unobservability* as the inability of the adversary to distinguish with better than negligible probability whether there exists some  $U, \Delta$  such that  $(U : ID, \Delta)$  is a registered honest user.

**Threat Model.** For the *unlinkability (G1)*, *security against impersonation (G2)*, and *external identity verification (G3)* goals, we assume the following threat model. We assume that the adversary can join Pudding with any number of user devices, and may observe the network traffic globally, or interfere with it in arbitrary ways. We assume that at most  $f$  out of  $n = 3f + 1$  discovery nodes are faulty. A faulty node is either non-Byzantine (crashed, disconnected), or Byzantine (may maliciously deviate from the protocol). We assume that Byzantine nodes are controlled by the adversary, and that they may collude. Non-Byzantine nodes follow an honest-but-curious model [23]: they correctly follow the protocol, but any messages they receive may be leaked to the adversary. The set of non-Byzantine-faulty nodes may change over time, whereas a Byzantine node is assumed to remain Byzantine forever. The users and the non-faulty nodes do not need to know which nodes are faulty. This is a common approach for Byzantine fault-tolerant protocols [24].

We assume that the underlying Loopix construction protects against active and passive network adversaries who wish to learn which nodes are communicating with each other. We base these assumptions upon the Loopix threat model [7]. Lastly, we assume that (1) the email services employed for username verification ensure that an email originating from a specific email address contains a DKIM signature (see §5.1) corresponding to that email address, and (2) that email traffic is encrypted with TLS so that a network adversary cannot read or modify messages.

For the *membership unobservability (G4)* goal, we assume a weaker threat model, where the adversary controls any number of user devices but no mix, provider, or discovery nodes (an adversary with access to a single discovery node’s database can trivially determine which usernames are registered). Despite this weaker threat model, providing membership unobservability against user devices significantly limits the number of parties that have access to this information, and also prevents crawling the database of usernames, which has been an issue in existing encrypted messaging apps [13].

**Security parameters.** We assume  $n = 3f + 1$  discovery nodes for the following reason. Assume  $f$  discovery nodes are temporarily unavailable during registration, so  $n - f$  discovery nodes have a copy of the contact information for a given username. When searching for that username, a different set of  $f$  discovery nodes may be unavailable, but at least  $n - 2f \geq f + 1$  nodes will be both available and have a copy of the contact information of that username, which will be sufficient to allow discovery of this username.

## 4. User discovery

Padding user discovery allows a user device to make contact with the registered owner of a particular username. Registering to Padding with a username allows a user to become discoverable through Padding by that username, but a user who wants to search for another user does not need to register with Padding. We first discuss the discovery protocol, and defer the description of the registration protocol to §5.

### 4.1. Deterministic SURB generation

Padding’s *membership unobservability* goal (see §3) requires that a searcher cannot tell from the discovery nodes’ responses whether a given username has been registered. We achieve this by always returning a SURB for a discovery request. If the searched username is registered at the discovery nodes, the SURB will allow the searcher to send a message to the username’s owner. Otherwise, the result of the discovery query is a fake SURB that does not route to any real user, but simply causes the message to be dropped in the mix network. This provides membership unobservability because (i) a user cannot distinguish whether a SURB was generated with fake or real contact information, and (ii) a user cannot distinguish whether a message was dropped by the network, or the recipient user chose not to respond to the message. Although the final mix node in a packet’s route can tell whether a SURB is fake, as per §3.2, the threat model for membership unobservability only includes malicious users, not adversarial mix or discovery nodes.

SURB generation cannot simply be performed by a single discovery node, since otherwise that discovery node could impersonate any user by routing the searcher’s message to an adversary-controlled destination. Instead, we define a deterministic scheme for generating SURBs, so that all honest discovery nodes generate the same SURB for the same query. We then require that a user receives an identical SURB from at least  $f + 1$  distinct discovery nodes before sending a message using that SURB. Since we assume that a maximum of  $f$  discovery nodes may be Byzantine, there is no way that the adversary can cause the user to receive  $f + 1$  incorrect SURBs, and therefore any SURB received from at least  $f + 1$  discovery nodes must be correct.

**Traditional SURB generation.** The original SURB generation procedure for Sphinx [19] takes two inputs: (1) a destination contact information  $\Delta$  and (2) a sequence of mix nodes  $\{n_0, n_1, \dots, n_{v-1}\}$  through which the message will be routed, where  $v < r$  and  $r$  is the maximum number of mix nodes in a message route. The SURB generation procedure in Loopix needs randomness for two tasks: (1) picking a delay value  $d_i$  for the message to be delayed at each hop  $n_i$ , chosen from an exponential distribution so that each mix node behaves as a Poisson process; and (2) picking an element of a prime-order group, which is used by each mix node in the path to derive a shared secret key, known by the mix node and the SURB’s creator. Each mix node

uses this shared secret key to derive the keys for validating the integrity of the message, removing a layer of encryption, and learning the next hop in the message route [7], [19].

**Deterministic SURB generation.** Our SURB generation process is identical to the original procedure, with the exception that all sources of randomness are replaced with a pseudorandom generator seeded with  $\text{KDF}(\textit{nonce} \parallel \textit{username} \parallel k)$ , where *username* is the username being queried, *nonce* is a nonce in the query message, *k* is a long-lived secret that is shared among all the discovery nodes (set up when the discovery nodes are created), and KDF is a key derivation function. Instead of taking the sequence of mix nodes  $\{n_0, n_1, \dots, n_{v-1}\}$  (the message route) as an input, it also generates this sequence pseudorandomly from the KDF output. The inputs to our modified procedure are therefore the *username*, the *nonce*, and the destination contact information  $\Delta$ . Determinism ensures that all honest discovery nodes generate the same SURB for the same destination contact information and nonce. Since *k* is not known to users, a malicious user cannot run the same SURB generation algorithm to break membership unobservability.

While not explicitly outlined in our protocol, potential leaks of the *k* value can be mitigated through an epoch-based mechanism, where the discovery nodes periodically agree on fresh *k* values. Lookup failures during key rotation can be reduced by having overlapping epochs. The epoch number can also be incorporated into nonces, which avoids having to store nonces forever.

**Creating fake SURBs.** When a user queries a username that is not registered on the system, Padding nevertheless returns a valid SURB in order to provide membership unobservability. This SURB is generated in the same way as a SURB for a genuine user, except that the discovery nodes use a fixed false contact information  $\Delta_{\text{fake}}$  instead of  $\Delta$ . This  $\Delta_{\text{fake}}$  is a non-existent “black hole” destination, so that any messages sent to it are dropped by the mix network. The querying user cannot distinguish a SURB for the black hole from a SURB for real contact information, because the outermost layer of the SURB header is encrypted with the public key of the first mix node on the path [19], and thus cannot be decrypted by the user.

If a user makes queries for two different usernames that are both unregistered, the SURB responses will be constructed based on the same  $\Delta_{\text{fake}}$  destination. However, since the seed of the pseudorandom generator incorporates the queried username, the SURBs are encrypted with different keys, and therefore the querying user cannot tell whether the SURBs are based on the same or different contact information. Having a single, fixed  $\Delta_{\text{fake}}$  eliminates the need for coordination between discovery nodes.

The public key within  $\Delta_{\text{fake}}$  must be chosen in a way that the corresponding private key is not known by any party in the system. For elliptic curves, this can be achieved by running a hash-to-curve algorithm [25] with a simple constant input, such as the string “Padding fake identity”, to obtain a group element whose discrete logarithm is unknown.

## 4.2. Authenticating contact information

The SURB described in the previous section allows Alice to send a single message to Bob without Alice learning Bob’s contact information, and without Bob learning Alice’s identity. However, if Alice wishes to identify herself to Bob by including her username in her message, the SURB does not allow Bob to verify that the message did indeed come from Alice, nor does it allow Alice to verify that replies to her message did indeed come from Bob. Security against impersonation (G2) requires users to authenticate each other.

If membership unobservability were not a concern, the discovery nodes could return a user’s public key (which is part of their contact information  $\Delta$ ) when queried for a username. Alice and Bob could query the discovery servers for each other’s usernames to learn each other’s public keys, and then perform an authenticated key exchange protocol to establish a mutually authenticated channel.

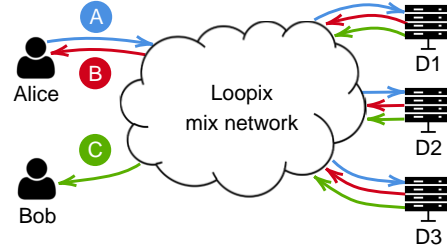
To enable mutual authentication while preserving membership unobservability, we use key-blinded signatures [26], [27], a technique also used by Tor to maintain anonymity of Tor onion services [16], [28]. (Not to be confused with blind signatures [29]: key-blinded signatures hide the unblinded public key from the verifier, whereas blind signatures hide the message from the signer.)

The scheme we use [26] works as follows. Let  $g$  be a generator of an elliptic curve group with prime order  $p$ . Let  $x \in \mathbb{Z}_p$  be a private key and let  $g^x$  be the corresponding public key. Let  $y \in \mathbb{Z}_p$  be a blinding factor chosen uniformly at random. One can then sign and verify digital signatures similarly to regular ECDSA or EdDSA signatures, using  $xy \pmod{p}$  as private key and  $g^{xy}$  as public key. Given a user’s public key  $g^x$ , a discovery node can therefore generate blinded public keys  $(g^x)^{y_1}, (g^x)^{y_2}, \dots$  that cannot be linked with each other, or with the unblinded public key  $g^x$ , without knowledge of the blinding factors  $y_1, y_2, \dots$  [27]. We write the blinded-key signing operation as  $\text{SignBK}(sk, bk, msg)$  where  $sk$  is the unblinded private key,  $bk$  is the blinding key, and  $msg$  is the data to be signed.

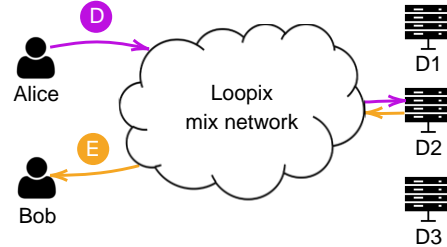
## 4.3. Discovery

User discovery in Pudding consists of three protocol phases: LOOKUP queries the database of users, CONTACTINIT allows the searcher to send an initial message to the queried user, and ADDFRIEND establishes an authenticated channel between the two users.

**LOOKUP protocol.** Let Alice be a user who is searching for Bob, and let  $ID_A$  and  $ID_B$  be their usernames (having a username is optional in Alice’s case). Note that we use “Alice/Bob” as shorthand for “Alice’s/Bob’s user device”. Alice and Bob are both Loopix users, with contact information  $\Delta_A$  and  $\Delta_B$  respectively. Each user’s contact information includes a long-term public key for which the user’s device holds the corresponding private key. Alice’s private key is  $x_A \in \mathbb{Z}_p$  and her public key is  $g^{x_A}$ , whereas Bob’s private and public keys are  $x_B$  and  $g^{x_B}$  respectively. We also assume



(a) During the LOOKUP protocol: (A) Alice sends a query message containing Bob’s username to each discovery node; (B) each discovery node replies with a SURB and a blinded public key for Bob; (C) each discovery node also sends Bob the blinding key.



(b) Then, Alice runs the CONTACTINIT protocol where: (D) Alice creates a Sphinx packet  $M_{init}$  using the SURB she received in LOOKUP. Alice then sends a message to an arbitrary discovery node, where  $M_{init}$  is the payload. (E) The discovery node places  $M_{init}$  in the network, which routes the message to Bob.

Figure 2: Alice discovering Bob (LOOKUP) and initiating contact with him (CONTACTINIT) through a Pudding setup with three discovery nodes ( $D1, D2, D3$ ).

that each discovery node has a private signing key, and that Alice and Bob know the corresponding public key for each discovery node (this information can be part of the network topology that a user downloads when first connecting to the Loopix network). The protocol is illustrated in Figure 2:

- 1) Alice generates a random *nonce*. For each discovery node  $i$ , Alice also generates a SURB  $S_i$  that allows that discovery node to reply to her. She then sends a discovery message  $(ID_B, nonce, S_i)$  to each discovery node via the anonymity network.
- 2) Upon receiving a discovery message, each discovery node checks if *nonce* has been used before; if so, it drops the message. This prevents a user from detecting when the contact data for another user changes by making repeated queries using the same nonce.
- 3) Each discovery node seeds a pseudorandom generator with  $\text{KDF}(nonce \parallel ID_B \parallel k)$ , where  $k$  is a secret shared with the other discovery nodes as described in §4.1. If the discovery node has a record for  $ID_B$  with contact information  $\Delta_B$  in its database, it creates a deterministic SURB  $surb$  using  $\Delta_B$  and the pseudorandom generator; if not, it generates the SURB using  $\Delta_{fake}$  instead. Each discovery node also obtains a blinding key for Bob  $y_B \in \mathbb{Z}_p$  from the pseudorandom generator, obtains Bob’s public key  $g^{x_B}$  from  $\Delta_B$ , and



computes the blinded public key  $bpk_B = (g^{x_B})^{y_B}$ . If  $ID_B$  is not registered, it computes a blinded key  $bpk_B = (pk_{fake})^{y_B}$  based on the public key  $pk_{fake}$  in  $\Delta_{fake}$ .

- 4) Each discovery node computes a signature  $sig_A = \text{Sign}_{sk}(nonce \parallel surb \parallel bpk_B)$  using its signing key  $sk$ , and sends  $(nonce, surb, bpk_B, sig_A)$  to Alice using the SURB  $S_i$  in Alice's discovery message.
- 5) If  $ID_B$  is registered with contact information  $\Delta_B$ , each discovery node also computes a signature  $sig_B = \text{Sign}_{sk}(nonce \parallel y_B)$  using its signing key  $sk$ , and sends  $(nonce, y_B, sig_B)$  to Bob using  $\Delta_B$ . This will allow Bob to generate signatures for the blinded public key.
- 6) For each message received, Alice checks that  $sig_A$  is a valid signature by one of the discovery nodes. She waits until she has received the same  $surb$  and  $bpk_B$ , along with the  $nonce$  she generated, from at least  $f+1$  distinct discovery nodes. If she does not receive them within some timeout, she discards all values and restarts the protocol.
- 7) For each message received, Bob checks that  $sig_B$  is a valid signature by one of the discovery nodes. If Bob receives  $f+1$  copies of the same blinding key  $y_B$  and  $nonce$  from distinct discovery nodes, he stores the two values locally in anticipation of a future contact request (see CONTACTINIT Step 3).

**CONTACTINIT protocol.** The SURB received from  $f+1$  discovery nodes during the LOOKUP protocol allows Alice to send a message to Bob anonymously. However, if Alice sends her initial message directly to Bob using this SURB, an adversary who controls a discovery node in addition to Alice's provider or the first mix node on the message path can recognise the SURB and link it to Alice. To prevent such linking attacks, the message is reflected via a different node. For simplicity, we choose a discovery node as the reflector. Alternatively, the system can incorporate separate, untrusted reflector nodes, or even allow other users to serve as message reflectors. Our protocol allows Alice to incorporate a pseudonym or a codeword in her initial message, enabling Bob to make an informed decision about whether to respond to her.

- 1) Alice generates an ephemeral Diffie-Hellman private key  $a \in \mathbb{Z}_p$ , public key  $g^a$ , and initial message key  $K_e = \text{KDF}((bpk_B)^a \parallel \text{"init key"})$  from the blinded public key received in LOOKUP Step 6. She then uses the SURB from LOOKUP Step 6 to create a Sphinx packet  $M_{init}$  [19]. The payload of  $M_{init}$  contains  $g^a$ , the nonce value she used when searching for Bob, as well as the following fields encrypted under  $K_e$ : (1) a SURB that allows Bob to respond to her, (2) an optional pseudonym or codeword for Bob, and (3.a) either her username  $ID_A$ , or (3.b) if she is not registered or does not want to share her username, Alice samples a blinding key  $y_A \in \mathbb{Z}_p$  and uses her long-term public key  $g^{x_A}$  to compute a blinded public key  $bpk_A = (g^{x_A})^{y_A}$ , which she includes in  $M_{init}$ . Alice then picks a discovery node and sends it a message via

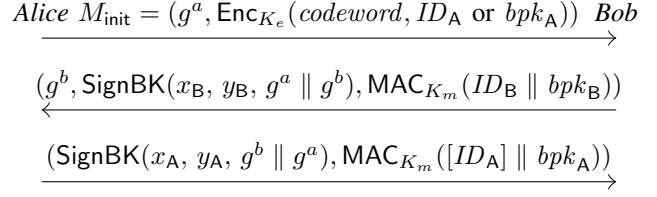


Figure 3: Summary of the ADDFRIEND protocol. Square brackets denote that including Alice's username is optional. Omitted from the diagram: each message also includes the nonce from Alice's original LOOKUP request (to identify messages belonging to the same protocol run) and a SURB that allows the other party to send a reply message; Bob's message to Alice also includes the nonce he used in his optional LOOKUP run for Alice's username.

the anonymity network that contains  $M_{init}$  as a payload.

- 2) The discovery node that receives Alice's message sends  $M_{init}$  through the anonymity network. It is routed to Bob using the SURB created by the discovery nodes.
- 3) When Bob receives  $M_{init}$ , he computes  $K_e = \text{KDF}((g^a)^{x_B y_B} \parallel \text{"init key"})$  using his private key  $x_B$  and the blinding key  $y_B$  received in LOOKUP Step 7, and then decrypts the rest of the message. Based on Alice's codeword he can either decide to reply to it, or ignore it. If Bob decides to reply, he proceeds to the ADDFRIEND protocol specified below. If Bob decides to ignore the message, Alice does not learn if  $ID_B$  is registered to Pudding or not.
- 4) If Alice does not receive a reply from Bob within some timeout, she retries by sending  $M_{init}$  through a different discovery node up to  $f+1$  times to route around faulty discovery nodes.

**ADDFRIEND protocol.** If Bob decides to accept Alice's contact request, the users need to establish an authenticated communication channel: Alice needs to ensure that she is indeed talking to the user registered under the username  $ID_B$ , and if Alice supplied her username, Bob needs to check that he is indeed talking to  $ID_A$ .

We achieve mutual authentication using the SIGMA authenticated key exchange protocol [30], adapted to support blinded-key signatures (see §4.2). We use SIGMA because it is fairly simple and proven correct [31], but another authenticated key exchange protocol could be substituted in its place. The ADDFRIEND protocol is shown in Figure 3.

- 1) Bob obtains a blinded public key  $bpk_A$  for Alice: (a) if  $M_{init}$  does not include Alice's username, Bob uses the  $bpk_A$  in this message; (b) if  $M_{init}$  includes  $ID_A$ , Bob obtains a blinded public key  $bpk_A = (g^{x_A})^{y_A}$  for Alice by running LOOKUP with input  $ID_A$  and using the value returned by at least  $f+1$  discovery nodes.
- 2) Bob uses the blinding key  $y_B$  he received from the discovery servers (LOOKUP Step 7) and his public key  $g^{x_B}$  to obtain his own blinded public key  $bpk_B = (g^{x_B})^{y_B}$ .
- 3) Bob generates an ephemeral Diffie-Hellman private key  $b \in \mathbb{Z}_p$  and public key  $g^b$ , and calculates the

MAC key  $K_m = \text{KDF}((g^a)^b \parallel \text{“MAC key”})$  using the value  $g^a$  from  $M_{\text{init}}$ . He uses  $K_m$  to calculate the message authentication code  $\text{MAC}_{K_m}(ID_B \parallel \text{bpk}_B)$ . Bob also generates  $\text{SignBK}(x_B, y_B, g^a \parallel g^b)$  using the key-blinded signature scheme [26]. Finally, he sends a message to Alice using the SURB in  $M_{\text{init}}$ , containing  $g^b$ , the MAC, the signature, and a SURB for Alice to reply. If he ran LOOKUP for  $ID_A$  in Step 1, Bob also includes the nonce he used in this LOOKUP.

- 4) If Alice included her username  $ID_A$  in  $M_{\text{init}}$ , she waits until she has received a blinding key  $y_A$  from at least  $f + 1$  distinct discovery nodes, as a result of Bob’s LOOKUP run at Step 1 of ADDFRIEND. She uses the nonce in Bob’s message to identify the blinding key. If Alice chose to remain anonymous, she uses the blinding key  $y_A$  that she generated at CONTACTINIT Step 1.
- 5) When Alice receives Bob’s message, she checks that the signature over  $g^a \parallel g^b$  is valid using the blinded public key  $\text{bpk}_B$  she received from the discovery nodes. She then calculates the MAC key  $K_m = \text{KDF}((g^b)^a \parallel \text{“MAC key”})$  and uses it to check the MAC over  $ID_B$  (the username she originally searched for) and the blinded public key  $\text{bpk}_B$  from the discovery nodes.
- 6) If all checks pass, Alice sends a reply to Bob containing a key-blinded signature  $\text{SignBK}(x_A, y_A, g^b \parallel g^a)$  computed from her private key  $x_A$  and the blinding key  $y_A$  from Step 4. If Alice sent her username in  $M_{\text{init}}$ , the message also includes  $\text{MAC}_{K_m}(ID_A \parallel \text{bpk}_A)$ , where  $\text{bpk}_A = g^{x_A y_A}$ . If Alice is anonymous, the message includes  $\text{MAC}_{K_m}(\text{bpk}_A)$ .
- 7) When Bob receives Alice’s message, he checks the signature using Alice’s blinded public key that he obtained in Step 1, and checks that the MAC over Alice’s blinded public key (and optionally username) is valid using key  $K_m$ . If these checks succeed, Alice and Bob have fully authenticated each other.
- 8) Alice and Bob can then derive a shared session key  $K_s = \text{KDF}(g^{ab} \parallel \text{“session key”})$ , and use an authenticated encryption scheme to secure their further communication (which may include sharing their contact information with each other).

## 5. Registration

A user must register to Pudding with a username to be discoverable through that username. However, lookups are possible without registration. We describe the Pudding registration protocol in this section.

### 5.1. Authentication with DKIM

The conventional way of checking whether a user owns a particular email address is by sending a random token to that address and asking the user to send that token back to the server, typically by clicking an HTTPS link in the email. However, since Pudding has multiple discovery servers, using traditional email address verification would cause complications. The first option to use traditional email

authentication in Pudding would be for discovery nodes to trust each other, such that one node performs the verification and the other nodes trust its verdict. This option does not meet Pudding’s security goals, since we assume that up to  $f$  discovery nodes may be Byzantine and hence not perform the email verification truthfully. The second option would be for each discovery node to perform its own email address verification, meaning the user would have to click  $n$  links in  $n$  separate emails. This is how Alphenhorn handles email verification for its users [14]. However, we believe that this method is inconvenient for users. We therefore propose a different method for checking whether the user owns an email address.

DomainKeys Identified Email (DKIM) [32] uses a digital signature in a header to confirm the origin of an email. The recipient can retrieve the public verification key that the domain owner has published as a DNS TXT record and use it to verify the DKIM signature. The signature confirms that the email was sent via an SMTP server authorised by the owner of the sender’s domain [33]. Assuming that this SMTP server has authenticated the user (the part of the email address before the ‘@’ sign), DKIM allows us to confirm that the sender of a message has not been spoofed.

In Pudding, one discovery node collects challenge values from all discovery nodes, and sends these challenges in a single email to the email address picked as the username. To confirm receipt, the registering user responds to this email using the reply function of their email client. Since email applications attach the original message to the response, and servers automatically sign the email with the domain’s DKIM key, the user typically does not have to type or copy-paste anything. The discovery node that receives the reply then shares it with all other nodes. This mechanism allows each discovery node to independently confirm that the user’s response email is (1) a genuine registration request (by checking if its challenge value is included in the email) and (2) sent by the owner of the email address (by checking the validity of the DKIM signature).

**Substitutability of DKIM.** Although this implementation uses DKIM, it is not the only viable verification method. For instance, if DKIM verification fails, the system can fall back to a method where each discovery node sends the user a verification link, with the usability downside of requiring multiple confirmation emails and link clicks.

### 5.2. Registering to Pudding

The user registration protocol REGISTER is described below and illustrated in Figure 4.

- 1) The user device joins the Loopix network by signing up to a provider node. This process involves creating a public/private key pair for secure communication over the network and obtaining the network topology from the provider (including the set of discovery servers and their public keys). The public key, along with the provider details and any other necessary information, constitutes the user’s contact information  $\Delta$ . The user






Figure 4: Schematic of user Alice registering to a Pudding setup with three discovery nodes ( $D1$ ,  $D2$ ,  $D3$ ). (A) Alice sends to the discovery nodes a message containing her email address, contact information, and the ID of a randomly-selected discovery node  $D_{\text{auth}}$ . (B)  $D_{\text{auth}}$ , ( $D3$  in this case), collects challenge values from all discovery nodes and sends Alice an email including these challenges, along with her contact information and username. (C) Alice replies to the email through an email service;  $D3$  receives Alice’s DKIM-signed reply email and sends it to the other discovery nodes. (D) Each discovery node verifies the DKIM signature and its own challenge’s presence in the email. Once confirmed, they register Alice and send her a confirmation message.

also chooses an email address that she controls as her username  $ID$ .

- 2) The user device randomly picks a discovery node  $D_{\text{auth}}$  as the node responsible for sending the verification email. The user then sends a *registration message* to all discovery nodes, which contains  $ID$ , the contact information  $\Delta$ , and the identifier of  $D_{\text{auth}}$ .
- 3) After receiving a registration message, each discovery node checks whether it has  $ID$  already registered, in which case it marks the registration request as invalid. Regardless, the discovery node generates a random challenge value.
- 4) Any discovery node that is not  $D_{\text{auth}}$  sends its challenge value and  $ID$  to  $D_{\text{auth}}$ . Meanwhile,  $D_{\text{auth}}$  waits until it has challenges from at least  $2f + 1$  distinct discovery nodes, including itself (see §3.2).  $D_{\text{auth}}$  then sends a *verification email* to the email address provided as  $ID$ , which contains all the challenges and Alice’s contact information  $\Delta$ . Including  $\Delta$  in the email prevents a malicious  $D_{\text{auth}}$  from providing a  $\Delta_{\text{bad}} \neq \Delta$  to the remaining discovery nodes.
- 5) If the user receives the email within a predetermined timeout period, she verifies that the included contact information  $\Delta$  is correct, and responds with an email that includes all challenge values<sup>3</sup>. Otherwise, the user

3. This process can be automated in practice, for example by including a link in the email, which redirects the user from the email app to the anonymous messaging app. After automatically checking that the  $\Delta$  is correct, the anonymous messaging app can use the OS features to automatically generate a reply email that contains all the email content and headers, and allow user to reply with the click of a button.

returns to Step 2 to try registering again with a new  $D_{\text{auth}}$ .

- 6) When  $D_{\text{auth}}$  receives the user’s response,  $D_{\text{auth}}$  forwards the email response verbatim, including all headers, to all other discovery nodes. Each discovery node independently checks if (1) the challenge it generated was included in the email, (2) the DKIM signature is valid, using the public key stored in DNS for the domain of  $ID$ , and (3) the registration request was not marked as invalid at Step 3. If any of these fail, the discovery node aborts the registration process.
- 7) Each discovery node that successfully completes the previous steps sends a signed *confirmation message* to all other discovery nodes, which includes Alice’s  $ID$  and  $\Delta$ . Once a discovery node has received messages from  $2f$  discovery nodes for Alice’s  $ID$ , it stores the link between  $ID$  and  $\Delta$  in its database. Waiting for confirmation from  $2f$  distinct discovery nodes ensures all correct discovery nodes register the same  $\Delta$ , even if  $D_{\text{auth}}$  is malicious. Each discovery node then sends a message via Loopix to the user’s contact information, confirming registration.
- 8) If the user device receives at least  $2f + 1$  confirmations within a predetermined timeout period, the user knows that she is successfully registered in Pudding.

If  $D_{\text{auth}}$  is malicious, it is possible with this protocol that some honest discovery nodes do not complete the registration. This can be avoided by disseminating  $(ID, \Delta)$  using a Byzantine reliable broadcast protocol [34], [35].

## 6. Evaluation

To evaluate our protocol, we compare it against our security goals (§6.1), discuss its limitations (§6.2), and present findings from implementing Pudding over Nym (§6.3).

### 6.1. Security analysis

We begin by formalising the goals **G1**...**G4** from §3.2 as security games. We denote the adversary as  $\mathcal{A}$  and challenger as  $\mathcal{C}$ .  $U_{\mathcal{C}}$  and  $U_{\mathcal{A}}$  refer to users controlled by the challenger and adversary, respectively.

**Unlinkability game  $Exp_{G1}$ :** The adversary  $\mathcal{A}$  controls up to  $f$  discovery nodes, the user  $U_{\mathcal{A}}$ , the mix nodes in all but one of the Loopix network layers, and the network between all of the nodes. The challenger  $\mathcal{C}$  controls users  $U_{\mathcal{C}}$ ,  $U_0$ , and  $U_1$  chosen by the adversary, and generates a random bit  $b \in_R \{0, 1\}$ .  $U_{\mathcal{C}}$  is registered to the discovery nodes with the username  $ID_{\mathcal{C}}$ .  $U_b$  runs the discovery protocol (LOOKUP, CONTACTINIT, and ADDFRIEND) with input  $ID_{\mathcal{C}}$ , and  $U_{\mathcal{C}}$  decides to respond in CONTACTINIT Step 3. Finally,  $\mathcal{A}$  outputs  $b'$ . The adversary wins if  $b = b'$ .

**$Exp_{G1}$  security arguments:** We claim that the adversary cannot win  $Exp_{G1}$  with a significantly higher probability than the chance of breaking sender-receiver unlinkability in the underlying Loopix network (which depends on the size of the anonymity set). The reason is threefold. (1)

An adversary performing passive or active attacks on the network links between nodes or adversary-controlled mix nodes is limited by the unlinkability property of Loopix. (2) Controlling discovery nodes does not confer the adversary an advantage, since all communication between  $U_b$  and those nodes is done via SURBs, which hide the identity of the recipient. (3) The potentially identifying information in  $M_{\text{init}}$  is encrypted with the blinded public key of  $U_C$ , which  $\mathcal{A}$  does not control. Appendix B.1 provides a more detailed discussion.

**Security against impersonation game  $Exp_{G2}$ :** There are three users:  $U_C$ ,  $U_A$ , and an honest user  $U_{\mathcal{H}}$  controlled by  $\mathcal{C}$ .  $U_{\mathcal{H}}$  registers using  $ID$ , which is an email address not controlled by  $\mathcal{A}$  and known to  $\mathcal{C}$ .  $\mathcal{A}$  controls up to  $f$  discovery nodes, any number of mix nodes and provider nodes, and the network between all of the nodes.  $\mathcal{A}$  chooses two messages  $m_0, m_1$  of the same length and reveals them to  $\mathcal{C}$ .  $\mathcal{C}$  generates a random bit  $b \in_R \{0, 1\}$ .  $U_C$  runs the discovery protocol with input  $ID$ , and once ADDFRIEND completes,  $U_C$  sends a message  $m_b$  encrypted with the session key  $K_s$ . Meanwhile, the adversary may inject arbitrary messages into the network.  $\mathcal{A}$  outputs  $b'$ , and wins the game if  $b' = b$ .

**$Exp_{G2}$  security arguments:** We claim that the adversary has a negligible advantage over a random guess. (1)  $U_C$  only accepts one response per discovery node per LOOKUP, since responses are signed and linked with  $U_C$ 's nonce. Moreover,  $\mathcal{A}$  cannot modify the LOOKUP responses from the honest discovery nodes, since the Sphinx packet format ensures integrity [19]. Therefore,  $U_C$  will not reach the threshold of  $f+1$  responses required to accept the adversary's responses, and therefore the SURB and blinded public key used by  $U_C$  will be the one generated by the honest discovery nodes using the contact information of  $U_{\mathcal{H}}$ . (2) The adversary knows the nonce and SURBs for  $U_C$ , and can therefore send fake ADDFRIEND protocol messages to  $U_C$  and  $U_{\mathcal{H}}$ . However, due to the existential unforgeability of the key-blinded signature scheme, the adversary cannot forge a signature that  $U_C$  will accept as valid for  $U_{\mathcal{H}}$ 's blinded public key (as produced by the honest discovery nodes). Therefore, the computation of  $K_m$  at ADDFRIEND Step 5 will only be performed using the  $g^b$  sent by  $U_{\mathcal{H}}$ , and therefore  $\mathcal{A}$  does not learn  $K_s$ . Even though the adversary can trick  $U_{\mathcal{H}}$  into sending the ciphertext of  $m_b$  to  $U_A$  by replacing the SURB in ADDFRIEND protocol messages, by the IND-CPA property of the symmetric cipher used to encrypt  $m_b$ , the adversary has only negligible advantage over guessing  $b$ . See Appendix B.2 for further details.

**External identity verification game  $Exp_{G3}$ :**  $\mathcal{A}$  controls  $f$  discovery nodes out of  $3f+1$ , as well as the user  $U_A$ .  $\mathcal{C}$  controls the remaining  $2f+1$  discovery nodes.  $ID$  is an email address not controlled by  $\mathcal{A}$ . Per our threat model for **G3** (§3.2), the email service used for identity verification is modelled as an oracle, guaranteeing that an email originating from an email address includes a DKIM signature that corresponds to that particular email address.  $U_A$  runs the REGISTER protocol for  $ID$ .  $\mathcal{A}$  wins the game if any of the  $2f+1$  challenger-controlled discovery nodes registers  $ID$ .

**$Exp_{G3}$  security arguments:** We claim that the adversary has a negligible probability of winning game  $Exp_{G3}$ . In order to successfully complete the REGISTER protocol for  $ID$ ,  $\mathcal{A}$  would have to forge a digital signature and pass the checks in REGISTER Step 6 for any honest discovery node.

**Membership unobservability game  $Exp_{G4}$ :**  $\mathcal{A}$  controls an arbitrary number of users  $U_{A0}, U_{A1}, \dots$ , but  $\mathcal{A}$  does not have access to the internal states of any discovery or provider nodes, as per our threat model for **G4** (§3.2).  $ID_0$  and  $ID_1$  are initially unregistered usernames known to both  $\mathcal{A}$  and  $\mathcal{C}$ .  $\mathcal{C}$  controls user  $U_C$ , generates a random bit  $b \in_R \{0, 1\}$ , and runs the REGISTER protocol to register ( $U_C : ID_b, \Delta$ ) for some arbitrary  $\Delta$  not controlled by  $\mathcal{A}$ .  $\mathcal{A}$  then runs the discovery protocol for  $ID_0$  and  $ID_1$  from any adversary-controlled users arbitrarily many times, but  $U_C$  does not respond to any of those contact requests. Finally,  $\mathcal{A}$  outputs  $b'$  and wins the game if  $b = b'$ .

**$Exp_{G4}$  security arguments:** We claim that  $\mathcal{A}$  has a negligible advantage over a random guess. Since  $U_C$  does not respond to  $M_{\text{init}}$ , the only messages received by the adversary are the discovery node responses. These messages contain the searcher's nonce, a SURB, a blinded public key, and a signature. The nonce and signature do not depend on whether the username is registered; the SURB is computationally indistinguishable from a uniform random bit string [19]; and the key-blinding scheme we use ensures that the blinded public key for real contact information is indistinguishable from that for  $\Delta_{\text{fake}}$  [26].

## 6.2. Discussion and limitations

**Email service provider.** There are two attacks that an adversary who controls a corrupted email service provider can perform. (1) An adversary who controls the user's SMTP server, or who can intercept unencrypted SMTP traffic, can register with Pudding on behalf of any email address on that server. (2) DKIM works at the *domain level*: it does not strictly verify that the part of the email address before the '@' sign is correct, but that an email has been signed by an email server that holds the private key for a particular domain [33]. Therefore, if the adversary has access to the domain's signing key, the adversary can forge emails with valid DKIM signatures, which appear to be sent from any email address hosted by that domain.

**Email traffic.** Even if the attacker does not control an email service provider, she can try to learn if a user has registered to Pudding with a certain email address by observing the network traffic between the user device and an email service provider, or between  $D_{\text{auth}}$  and the email service provider (e.g. for personal domains). This attack is outside our threat model.

**Malicious Loopix providers.** A Loopix provider node receives messages on behalf of its users while the users are offline, and allows users to download those messages when they next connect to the provider. This design is advantageous for mobile devices that are frequently offline,

but it has the disadvantage that a malicious provider can learn the number of messages received over time by each of its users (but not their origin). Some messaging patterns, such as receiving  $3f + 1$  messages in short succession, could indicate that the user has just performed a LOOKUP request. By linking the timing of such message receipt patterns with the timing of LOOKUP requests at a colluding discovery node, an adversary can make statistical estimates of which users might be searching for each other. This attack can be mitigated by ensuring that users receive sufficient (real or cover) traffic, and by increasing the size of the anonymity set by having more users submitting LOOKUP requests.

**Spam.** Spam is usually detected by looking at message contents. However, end-to-end encrypted messaging apps can reduce spam by detecting high volumes of automated messages coming from the same source and directed at a high number of users [36]. Since Pudding does not have access to such metadata, we should use other measures to prevent spam. One strict yet effective measure can be dropping the message if the person is not in your contact address book. Another measure could be making spamming expensive; for example, Nym requires its users to purchase bandwidth credentials.

**Pseudonymous usernames.** It may seem that Pudding should support pseudonymous identities for users who do not want to tie a non-anonymous identifier, such as an email address, to their anonymity network presence. However, removing external verification from registration also eliminates any chance of achieving membership unobservability. If anybody can register any username, an adversary can check whether username *ID* is registered by attempting to register *ID*, and then using another device to discover *ID* and send a message to it. If the adversary receives her own message, then *ID* was not previously registered. If the adversary does not receive her own message, even after several retries, then *ID* must have already been registered.

Consequently, in a seemingly paradoxical manner, tying anonymity network accounts to external identities such as email addresses provides *better* privacy than using short pseudonymous identities, in terms of hiding membership of the network. We have therefore chosen not to include pseudonymous usernames as a registration option in Pudding. However, we do not require users to register to discover others or send messages. Only one party in a conversation needs to register with Pudding in order to initiate a two-way conversation.

**Online / offline patterns.** Loopix hides communication patterns, however the provider nodes can still determine whether a user device is online or offline. If device pairs transition between online and offline states in a correlated fashion, it can potentially be used to infer that two users are communicating. However, this type of intersection attack is not unique to Loopix, but rather a concern in all anonymity networks [37]. This vulnerability is also acknowledged by other private user discovery protocols, as highlighted in Alpenhorn [14] and UDM [15].

**Usability.** In this work, we base Pudding’s usability traits upon previous work. We argue that enabling user discovery through usernames increases usability, because previous research shows that users are not good at generating and exchanging information needed to establish communication over an encrypted channel [1], [38], [39]. Using email addresses as usernames increases usability, based on the fact that all popular encrypted messaging apps allow users to find their friends through identifiers that already exist in their address books. The requirement to allow user discovery without leaving the network draws from research showing that authentication ceremonies of encrypted messaging apps have poor usability [40], [41]. In such ceremonies, users share public keys or verification codes manually with each other via QR code, near-field communication (NFC), or through speech [20]. Although during registration to Pudding the user has to leave the anonymity network to respond to a verification email, this only happens once for each username registration, and the user typically does not need to take any action other than hitting the reply button. Further work is required to measure the specific usability improvement that our protocol presents.

### 6.3. Performance evaluation

We have implemented a prototype of Pudding on Nym to demonstrate that our protocol is practical. The prototype is implemented in 3000 lines of Rust code using the Nym SDK<sup>4</sup>. We also modified the SDK and the `sphinx-packet` crate to expose and modify internal methods that allow us to create deterministic SURBs and access the user’s secret keys. The changes are contained in a `.patch` file with less than 2000 lines. Importantly, no modifications of Nym infrastructure (mix nodes, provider nodes, or protocols) are required. As such, all our experiments run on the general main net and share the same anonymity set as existing Nym users. We have chosen Nym because it is an operational anonymity network running on hundreds of nodes, and it does not currently provide a user discovery mechanism. All source code and SDK modifications are available in our open-source repository under an MIT license: <https://github.com/ckocaogullar/pudding-protocol>.

The Nym SDK manages SURBs when sending messages to other recipients using their Nym address. When sending a message from a user  $U_{\text{sender}}$  to another user  $U_{\text{recipient}}$ , the Nym client transparently includes SURBs so that  $U_{\text{recipient}}$  can reply without knowing  $U_{\text{sender}}$ ’s address. As messages are exchanged,  $U_{\text{recipient}}$  requests more SURBs from  $U_{\text{sender}}$  when her stock runs low. This mechanism is used in our protocol when the user communicates with the discovery servers. However, when using our deterministic SURBs (§4.1) both parties are anonymous towards each other and we have to make sure ourselves that the responses include sufficiently many SURBs.

In our evaluation, all clients run on a machine with an 8-core CPU (Intel Xeon Skylake) in a commercial data cen-

4. <https://nymtech.net/docs/sdk/rust.html>

ter. During an initialisation phase, the test program creates ephemeral clients with randomly chosen gateways for the users and discovery servers. All clients execute in a multi-threaded asynchronous runtime inside the same process. However, they do not share any information out-of-band once started. The DKIM verification step is simulated by another Nym client that replies with a signature immediately, effectively excluding the time taken by the user to receive and reply to an email.

In this evaluation, we are interested in the latency distributions for the Pudding sub-protocols. Specifically, we are interested in four metrics: latency of the REGISTER protocol (§5), latency of the LOOKUP protocol (§4.3), and the latency for the combination of the CONTACTINIT and ADDFRIEND protocols in both the anonymous and non-anonymous variant. Since LOOKUP is required for CONTACTINIT, these protocols are executed as part of the same scenario. Therefore, we execute three different scenarios: a registration scenario, an anonymous discovery (LOOKUP +CONTACTINIT +ADDFRIEND) scenario, and a non-anonymous discovery scenario. We run each scenario for 10 minutes and repeat it 6 times to account for outliers, such as unlucky choice of slow gateway nodes. We run configurations with 20 client nodes and a varying number of  $n = \{4, 7, 10\}$  discovery nodes.

For the REGISTER scenario, each client starts a new registration attempt every 30 seconds after an initial random pause. Similarly, for the other scenarios every client sends a LOOKUP request for a randomly chosen other user every 30 seconds and proceeds then with the CONTACTINIT and ADDFRIEND steps. For the discovery scenarios, all clients are pre-registered with the discovery nodes such that the results are not influenced by pending registration operations. We stop the clock for the REGISTER operation once the user receives  $2f + 1$  confirmations, and for the LOOKUP operation once the searcher receives  $f + 1$  confirmations. The combination of CONTACTINIT + ADDFRIEND is measured from the time when the searcher sends her first message to receiving the confirmation from the searchee, at which point the searcher and searchee can send messages to each other.

We have verified that the number of users (up to the maximum that our test bed can support) has negligible impact on the latency distributions. We also noted occasional instances where some messages were lost or nodes experienced temporary disconnections from the Nym network. Since these cases are automatically covered by the fault-tolerance of our protocol, we observed little difference between executions with failed nodes and those without.

**Registration and discovery over Nym are fast.** Our results are presented as cumulative distributions in Figure 5, while detailed numerical values can be found in Appendix A. All operations have mean latencies below 13 seconds. We therefore believe that both registration and discovery latency are suitable for real-world usage.

For this paragraph, we consider the scenario with  $n = 10$  discovery servers. The mean latency for REGISTER is




Figure 5: Latency of the Pudding sub-protocols represented as CDFs for different numbers of discovery nodes. Note that the x-axis is log-scaled. See also Table 1 in Appendix A.

8.04 seconds ( $p_{90} = 10.02$  s). These numbers are lower for the  $n = 4$  configuration. This is because, as  $n$  increases, the client has to wait for a larger number of replies which increases the chance of having to wait for a message with large mix node delays. The mean LOOKUP latency is 3.44 seconds ( $p_{90} = 4.41$  s). It is similar for other  $n$  as there are no “synchronization barriers”, where a single discovery node has to collect message from the others. The latency for the combined CONTACTINIT and ADDFRIEND step differs between the anonymous and non-anonymous variants. For the former, the mean latency is 8.21 seconds ( $p_{90} = 5.70$  s) while the non-anonymous variant has a mean latency of 12.34 seconds ( $p_{90} = 11.86$  s). The non-anonymous variant is expected to be slower, because the searchee has to perform an extra communication round with the discovery servers.

**Pudding’s performance is influenced by the anonymity network’s throughput limits.** Under high load, the latencies can exhibit significant outliers. This is because

Nym clients send data at a fixed rate (adding cover traffic if necessary) to prevent traffic analysis. Consequently, sending more messages increases the expected queuing delay. This effect is especially pronounced for the discovery requests, where the responses from the discovery nodes exceed the size limit of a single payload in Nym. As a result, each response to the searcher must be split up into multiple Nym packets (in our implementation we need up to ten).

In a real deployment, we anticipate that the queuing delay on user devices will be small, because registration and discovery requests are sent rarely – occurring only once during app installation for registration and once during initial communication with a friend for discovery. However, the discovery nodes may need to handle a high rate of requests and responses. This can be addressed by running multiple instances of the Nym client at each discovery node, and load-balancing messages across them. For this, the list of discovery nodes may contain multiple Nym identifiers for each, and clients randomly pick one when sending a message to that node.

## 7. Related work

We know of six systems that provide user discovery while hiding who is communicating with whom: Alpenhorn [14], UDM [15], PROUD [42], OUTOPIA [43], Arke [44], and Apres [45].

*Alpenhorn* [14] is based on the Vuvuzela [8] anonymity network, and therefore provides a high level of metadata privacy. Alpenhorn has a round-based system that requires users to periodically download large inboxes rather than accessing individual messages on demand. As a result, this scheme has high bandwidth costs, which is especially challenging for devices relying on cellular networks. For instance, with one million users, each user device needs to download a 7.5 MB mailbox, which exceeds the allowance of many cellular data plans [14]. The alternative of delaying mailbox downloads until WiFi is available can increase the lookup latency. Furthermore, if a user device fails to participate in a round where a lookup request was initiated for that user, the request will be lost. Moreover, Alpenhorn cannot remain operational if a single server crashes, goes offline or otherwise does not follow the protocol. Pudding meanwhile, is fault-tolerant and fast; it works by exchanging small messages rather than large mailboxes.

*UDM* [15] aims to hide contact relationships and protect users' external public identities from being linked to their private network identities. Compared to Pudding, UDM does not provide *membership unobservability* (G4) or *external identity verification* (G3). In addition, UDM does not prevent an adversary from deducing contact relationships by passively observing network patterns. Lastly, UDM does not provide fault tolerance by default, and discusses relevant measures as extensions to the current model.

*PROUD* [42] is a DNS-based private user discovery system that does not provide traffic analysis protection and requires users to exchange public keys out-of-band. Similarly, *OUTOPIA* [43] requires users to exchange out-of-band

information such as public keys and “long hard-to-guess names”, and is therefore not a usable discovery mechanism. Neither of these schemes meet our usability goals.

*Arke* [44] is a contact discovery protocol based on an unlinkable handshake mechanism, built atop an identity-based non-interactive key exchange. While it provides Byzantine fault tolerance like Pudding, Arke does not offer inherent unlinkability against passive adversaries observing network traffic. Also notably, Arke requires mutual interest for contact initiation, making it unsuitable for scenarios where one party seeks to contact another unilaterally, such as a whistleblower reaching out to a journalist.

*Apres* [45] allows users to find and message each other privately, providing unlikability and online unobservability. However, Apres requires each pair of users to communicate out-of-band and determine a pairwise introduction ID, which must remain secret and therefore also be hard-to-guess. In contrast, Pudding users have a single non-secret ID.

Private discovery of network entities is relevant for many types of anonymous and other communication networks. Below, we detail some prominent techniques and systems used for various private discovery purposes:

*Private Set Intersection (PSI)* protocols allow two or more parties to find the intersection of sets without learning anything else about each other's sets, except maybe their sizes [46]. PSI is a suitable solution for privately determining the intersection of the contacts in a user's address book and the users registered in the communication system [47]. However, PSI is not directly applicable to retrieving contact information.

*Private Information Retrieval (PIR)* protocols allow querying items from a server without revealing the query to the server [10]. Demmler et al. used a technique combining PIR with PSI to privately find which users in someone's contact address book are registered in a messaging system [46]. However, this technique does not achieve *membership unobservability* (G4). DP5 [48] also aims to solve a related but different problem: privately learning the online status of one's contacts. DP5 also does not provide *membership unobservability* (G4), and it requires users to exchange private keys out-of-band.

*Trusted Execution Environments (TEEs)* are used for protecting the confidentiality of data *in-use*. In the user discovery context, Signal uses Intel Software Guard Extensions (SGX) to privately find the intersection of the user's address book and the users registered to the messenger service [11]. TEEs can also be integrated into Pudding discovery nodes, to reduce the level of trust associated with discovery node providers.

*Private DNS solutions* including DNSCrypt [49], DNSCurve [50], Confidential DNS [51] and Private DNS [52], focus on protecting DNS queries and responses from being read or modified by unauthorised parties. These mechanisms are not suitable for private user discovery in anonymity networks, since they do not prevent the DNS servers from linking users with their queries. More recent approaches like Oblivious DNS [53], DNS for Tor [54], and

distributed DNS [55] are not resistant against global passive adversaries.

*Distributed Hash Tables (DHTs)* allow peers in the system to search for data objects using the keys associated with them. As with DNS, most DHT implementations do not provide protection against linking users with queries. The ones that aim to provide this property, including Salsa [56], AP3 [57], NISAN [58], Torsk [59] are susceptible to information leakage, as illustrated by numerous studies [58], [60]–[62].

## 8. Conclusions

Anonymity networks that allow messaging, such as Loopix [7] and Nym, currently lack a usable and privacy-preserving user discovery mechanism. Previous research suggests [1] this might impede widespread adoption. In this work, we have described Pudding, a novel private user discovery protocol that hides contact relationships, prevents impersonation, and provides fault-tolerance. Our protocol also provides membership unobservability against malicious users using a new deterministic single-use reply block (SURB) generation mechanism. Users register to Pudding with their email address. We verify email ownership of through a mechanism based on DomainKey Identified Mails (DKIM). A messaging app can then automatically search for the friends of the user by using the email addresses of friends stored in her address book, confident that the recipient is in control of the associated email inbox. Pudding is an application-layer protocol; it can be implemented as an overlay network on top of any anonymity network that supports SURBs. We have implemented our protocol, and evaluated its performance on Nym anonymity network. Our evaluation shows that Pudding is scalable, practical, and lightweight, achieving mean latencies of under 4 seconds for lookup, 13 seconds for initiating contact, and 8.5 seconds for registration—numbers significantly influenced by the delays in the underlying anonymity network.

## Acknowledgements

We would like to thank Alberto Sonnino and the anonymous reviewers for their valuable feedback. Special thanks to Furkan Saygın Şener for his creativity in naming the protocol. Ceren Kocaoğullar is supported by King’s College, Cambridge and the Cambridge Trust. She was supported by Nokia Bell Labs during the initial stages of this research project. Daniel Hugenroth is supported by Nokia Bell Labs. Martin Kleppmann is funded by the Volkswagen Foundation and crowdfunding supporters including Rohit Kulshreshtha, Mintter, and SoftwareMill.

## References

[1] A. Whitten and J. D. Tygar, “Why Johnny can’t encrypt: A usability evaluation of PGP 5.0.” in *USENIX Security*, 1999.

[2] M. Srivastava, M. Murgia, and H. Murphy, “The secret mission to bolster Ukraine’s cyber defences ahead of Russia’s invasion,” *Financial Times*, Mar. 2022, (accessed 2021-08-27). [Online]. Available: <https://perma.cc/3YZC-863H>

[3] M. Strohmeier, M. Smith, V. Lenders, and I. Martinovic, “The real first class? Inferring confidential corporate mergers and government relations from air traffic communication,” in *EuroS&P*, 2018.

[4] “Fitness app Strava lights up staff at military bases,” *BBC News*, Jan. 2018, (accessed 2021-08-27). [Online]. Available: <https://perma.cc/ZWS2-5VTC>

[5] A. Toler, “Hunting the hunters: How we identified Navalny’s FSB stalkers,” *Bellingcat.com*, Dec. 2020, (accessed 2021-08-27). [Online]. Available: <https://perma.cc/25JU-V4A4>

[6] J. Lund, “Technology preview: Sealed sender for Signal,” *Signal blog*, Oct. 2018, (accessed 2021-05-09). [Online]. Available: <https://perma.cc/6HC9-Y44E>

[7] A. M. Piotrowska, J. Hayes, T. Elahi, S. Meiser, and G. Danezis, “The Loopix anonymity system,” in *USENIX Security*, 2017.

[8] J. Van Den Hooff, D. Lazar, M. Zaharia, and N. Zeldovich, “Vuvuzela: Scalable private messaging resistant to traffic analysis,” in *ACM SOSP*, 2015.

[9] R. Cheng, W. Scott, E. Masserova, I. Zhang, V. Goyal, T. Anderson, A. Krishnamurthy, and B. Parno, “Talek: Private group messaging with hidden access patterns,” in *Annual Computer Security Applications Conference (ACSAC)*, 2020.

[10] S. Angel and S. Setty, “Unobservable communication over fully untrusted infrastructure,” in *USENIX OSDI*, 2016.

[11] M. Marlinspike, “Technology preview: Private contact discovery for Signal,” *Signal blog*, Sep. 2017, (accessed 2021-09-02). [Online]. Available: <https://perma.cc/N3GF-8WD6>

[12] WhatsApp Ltd, “WhatsApp privacy policy,” (accessed 2021-09-02). [Online]. Available: <https://perma.cc/CC39-2CFT>

[13] C. Hagen, C. Weinert, C. Sendner, A. Dmitrienko, and T. Schneider, “All the numbers are US: Large-scale abuse of contact discovery in mobile messengers,” in *NDS*, 2021.

[14] D. Lazar and N. Zeldovich, “Alpenhorn: Bootstrapping secure communication without leaking metadata,” in *USENIX OSDI*, 2016.

[15] D. Chaum, M. Yaksetig, A. T. Sherman, and J. de Ruiter, “UDM: Private user discovery with minimal information disclosure,” *Cryptologia*, vol. 46, no. 4, pp. 347–379, 2022.

[16] “Tor rendezvous specification - version 3,” *Tor Project*, (accessed 2021-09-02). [Online]. Available: <https://gitweb.torproject.org/torspec.git/tree/rend-spec-v3.txt>

[17] R. Dingledine, N. Mathewson, and P. Syverson, “Tor: The second-generation onion router,” *Naval Research Lab Washington DC, Tech. Rep.*, 2004.

[18] G. Danezis, R. Dingledine, and N. Mathewson, “Mixminion: Design of a type III anonymous remailer protocol,” in *IEEE S&P*, 2003.

[19] G. Danezis and I. Goldberg, “Sphinx: A compact and provably secure mix format,” in *IEEE S&P*, 2009.

[20] N. Unger, S. Dechand, J. Bonneau, S. Fahl, H. Perl, I. Goldberg, and M. Smith, “SoK: secure messaging,” in *IEEE S&P*, 2015.

[21] D. L. Chaum, “Untraceable electronic mail, return addresses, and digital pseudonyms,” *Communications of the ACM*, vol. 24, no. 2, pp. 84–90, 1981.

[22] M. Backes, A. Kate, P. Manoharan, S. Meiser, and E. Mohammadi, “AnoA: A framework for analyzing anonymous communication protocols,” in *IEEE Computer Security Foundations Symposium*, 2013.

[23] A. Paverd, A. Martin, and I. Brown, “Modelling and automatically analysing privacy properties for honest-but-curious adversaries,” *University of Oxford, Tech. Rep.*, 2014. [Online]. Available: <https://ajpaverd.org/publications/casper-privacy-report.pdf>



- [24] L. Lamport, R. Shostak, and M. Pease, "The Byzantine generals problem," *ACM Transactions on Programming Languages and Systems*, vol. 4, no. 3, pp. 382–401, Jul. 1982.
- [25] A. Faz-Hernandez, S. Scott, N. Sullivan, R. S. Wahby, and C. A. Wood, "Hashing to elliptic curves," IETF Internet-Draft, 2023. [Online]. Available: <https://datatracker.ietf.org/doc/draft-irtf-cfrg-has-h-to-curve/>
- [26] F. Denis, E. Eaton, T. Lepoint, and C. A. Wood, "Key blinding for signature schemes," IETF Internet-Draft, 2023. [Online]. Available: <https://datatracker.ietf.org/doc/draft-irtf-cfrg-signature-key-blinding/>
- [27] E. Eaton, D. Stebila, and R. Stracovsky, "Post-quantum key-blinding for authentication in anonymity networks," in *LATINCRYPT*, 2021.
- [28] N. Hopper, "Proving security of Tor's hidden service identity blinding protocol," Tor project, Tech. Rep., 2013.
- [29] D. Chaum, "Blind signatures for untraceable payments," in *CRYPTO*, 1983.
- [30] H. Krawczyk, "SIGMA: The 'SIGn-and-MAC' approach to authenticated Diffie-Hellman and its use in the IKE-protocols," in *CRYPTO*, 2003.
- [31] R. Canetti and H. Krawczyk, "Security analysis of IKE's signature-based key-exchange protocol," in *CRYPTO*, 2002.
- [32] B. Leiba and J. Fenton, "DomainKeys Identified Mail (DKIM): Using digital signatures for domain verification," in *Conference on Email and Anti-Spam*, 2007.
- [33] D. Crocker, T. Hansen, and M. Kucherawy, "DomainKeys Identified Mail (DKIM) signatures," RFC6376, 2011.
- [34] G. Bracha, "Asynchronous Byzantine agreement protocols," *Information and Computation*, vol. 75, no. 2, pp. 130–143, 1987.
- [35] C. Cachin, R. Guerraoui, and L. Rodrigues, *Introduction to Reliable and Secure Distributed Programming*, 2nd ed. Springer, 2011.
- [36] M. Jones, "How WhatsApp reduced spam while launching end-to-end encryption," in *USENIX Enigma*, 2017.
- [37] G. Danezis and A. Serjantov, "Statistical disclosure or intersection attacks on anonymity systems," in *International Workshop on Information Hiding*, 2004.
- [38] S. Ruoti, J. Andersen, D. Zappala, and K. Seamons, "Why Johnny still, still can't encrypt: Evaluating the usability of a modern PGP client," *arXiv preprint arXiv:1510.08555*, 2015.
- [39] S. Sheng, L. Broderick, C. A. Koranda, and J. J. Hyland, "Why Johnny still can't encrypt: evaluating the usability of email encryption software," in *ACM SOUPS*, 2006.
- [40] E. Vaziripour, J. Wu, M. O'Neill, J. Whitehead, S. Heidbrink, K. Seamons, and D. Zappala, "Is that you, Alice? A usability study of the authentication ceremony of secure messaging applications," in *SOUPS*, 2017.
- [41] A. Herzberg and H. Leibowitz, "Can Johnny finally encrypt? Evaluating E2E-encryption in popular IM applications," in *STAST*, 2016, pp. 17–28.
- [42] P. Papadopoulos, A. A. Chariton, E. Athanasopoulos, and E. P. Markatos, "Where's Wally? How to privately discover your friends on the Internet," in *ACM ASIACCS*, 2018.
- [43] P. Papadopoulos, A. A. Chariton, M. Pachilakis, and E. P. Markatos, "OUTOPIA: private user discovery on the internet," in *EuroSec*, 2022.
- [44] N. Mohnblatt, A. Sonnino, K. Gurkan, and P. Jovanovic, "Arke: Scalable and Byzantine fault tolerant privacy-preserving contact discovery," Cryptology ePrint Archive, 2023. [Online]. Available: <https://eprint.iacr.org/2023/1218>
- [45] B. Laurie, "Apres-a system for anonymous presence," Tech. Rep., 2004.
- [46] D. Demmler, P. Rindal, M. Rosulek, and N. Trieu, "PIR-PSI: scaling private contact discovery," in *PETS*, 2018.
- [47] D. Kales, C. Rechberger, T. Schneider, M. Senker, and C. Weinert, "Mobile private contact discovery at scale," in *USENIX Security*, 2019.
- [48] N. Borisov, G. Danezis, and I. Goldberg, "DP5: A private presence service," in *PoPETS*, 2015.
- [49] OpenDNS, "DNSCrypt," accessed 2021-01-07. [Online]. Available: <https://perma.cc/UDW6-7HYA>
- [50] M. Dempsy, "DNSCurve: Link-level security for the Domain Name System," IETF Internet-Draft, 2009. [Online]. Available: <https://tools.ietf.org/id/draft-dempsey-dnscurve-00.html>
- [51] W. Wijngaards and G. Wiley, "Confidential DNS," IETF Internet-Draft, 2015. [Online]. Available: <https://tools.ietf.org/html/draft-wijn-gaards-dnsop-confidentialdns-03>
- [52] P. Hallam-Baker, "Private DNS," IETF Internet-Draft, 2014. [Online]. Available: <https://tools.ietf.org/html/draft-hallambaker-privatedns-00>
- [53] P. Schmitt, A. Edmundson, A. Mankin, and N. Feamster, "Oblivious DNS: Practical privacy for DNS queries," in *PETS*, 2019.
- [54] M. Sayrafi, "Introducing DNS resolver for Tor," The Cloudflare Blog, May 2018, (accessed 2021-05-14). [Online]. Available: <https://perma.cc/NJY9-GDMD>
- [55] N. P. Hoang, I. Lin, S. Ghavamnia, and M. Polychronakis, "K-resolver: towards decentralizing encrypted DNS resolution," *arXiv preprint arXiv:2001.08901*, 2020.
- [56] A. Nambiar and M. Wright, "Salsa: a structured approach to large-scale anonymity," in *ACM CCS*, 2006.
- [57] A. Mislove, G. Oberoi, A. Post, C. Reis, P. Druschel, and D. S. Wallach, "AP3: Cooperative, decentralized anonymous communication," in *ACM SIGOPS European Workshop*, 2004.
- [58] A. Panchenko, S. Richter, and A. Rache, "Nisan: network information service for anonymization networks," in *ACM CCS*, 2009.
- [59] J. McLachlan, A. Tran, N. Hopper, and Y. Kim, "Scalable onion routing with Torsk," in *ACM CCS*, 2009.
- [60] P. Mittal and N. Borisov, "Information leaks in structured peer-to-peer anonymous communication systems," *ACM Transactions on Information and System Security*, no. 1, pp. 1–28, 2012.
- [61] Q. Wang, P. Mittal, and N. Borisov, "In search of an anonymous and secure lookup: attacks on structured peer-to-peer anonymous communication systems," in *ACM CCS*, 2010.
- [62] G. Danezis and P. Syverson, "Bridging and fingerprinting: Epistemic attacks on route selection," in *PETS*, 2008.



## Appendix A. Performance measurement data

We provide more detailed data for our evaluation measurements from §6.3 in Table 1. For each configuration we provide the mean latency and the 50-th, 90-th and 95-th percentile of the latency distribution.

# Discovery Nodes	Latency [s]			
	mean	$p_{50}$	$p_{90}$	$p_{95}$
<i>Registration</i>				
4	6.48	5.54	8.42	11.01
7	6.74	6.18	8.35	9.74
10	8.04	7.44	10.02	11.37
<i>Lookup</i>				
4	3.26	2.48	3.97	6.45
7	3.00	2.51	3.81	5.73
10	3.44	2.77	4.41	7.08
<i>ContactInit + AddFriend (anonymous)</i>				
4	8.77	3.66	9.47	30.44
7	10.58	3.43	4.87	6.01
10	8.21	3.51	5.70	14.21
<i>ContactInit + AddFriend (with verify)</i>				
4	11.51	6.52	11.16	32.06
7	10.00	6.23	8.55	10.35
10	12.34	6.75	11.86	38.07

TABLE 1: Registration and discovery requests latency over Nym for different numbers of discovery nodes.

## Appendix B. Security proof sketches

### B.1. Unlinkability (G1)

**Lemma B.1.** The adversary’s advantage in winning  $Exp_{G1}$  from observing or performing active attacks on the network links between nodes, or from controlling mix nodes, can be reduced to the probability of breaking sender-receiver unlinkability in the underlying Loopix network, provided that at least one mix node on the path of every message is honest.

All messages exchanged between the challenger’s users and the adversary-controlled nodes are sent via Loopix [7], which provides *sender-receiver third-party unlinkability* under a threat model that includes a global passive adversary (GPA), and protection against  $(n - 1)$  active attacks. Analytical and empirical analysis of this unlinkability property of Loopix shows that given appropriate parameters are chosen for the network (i.e. at least 3 mix node layers, message sending rate to delay parameter ratio  $\lambda/\mu \geq 2$ ), and assuming that there is a sufficient number of active Loopix users who are not controlled by the adversary to provide a large enough anonymity set, it is very unlikely that the adversary can identify which pairs of Loopix users are communicating by observing all traffic within the network

and also controlling a portion of corrupt mix nodes [7]. If the adversary were to control all mix nodes on the path of a message between  $U_b$  and  $U_c$ , she would be able to trace the message and thereby link sender and receiver; however, assuming at least one honest mix node on the path prevents this attack. Moreover, since the Sphinx packet format [19] used by Loopix ensures integrity, any packets that are modified by the adversary are simply dropped, which prevents an active network adversary from modifying the responses from discovery nodes to a user.

**Lemma B.2.** Assume the adversary controls up to  $f$  discovery nodes, any number of mix nodes, and the network between the nodes. When an honest user  $U_c$  performs a LOOKUP, the SURB and blinded public key she accepts at Step 6 of LOOKUP are those generated by the honest discovery nodes.

$U_c$  waits to receive identical responses from at least  $f + 1$  distinct discovery nodes before accepting those responses. Adversary-controlled discovery nodes may send multiple responses to the same LOOKUP request in order to try to reach that threshold (for example, the adversary may save a SURB from one LOOKUP request, and use it to send two responses to a later LOOKUP request; alternatively, the adversary may guess the searcher’s username, look up their contact information in the database, and send them Loopix messages directly). However, discovery nodes’ responses are signed,  $U_c$  knows the public keys of the discovery nodes (which it obtains as part of creating their user account), and we assume the adversary cannot forge signatures from the honest discovery nodes with greater than negligible probability.  $U_c$  therefore only accepts one response per LOOKUP nonce per discovery node signing key. Moreover, any malicious mix nodes or active network attackers cannot modify the responses from honest discovery nodes without invalidating the signature. Thus, the adversary-generated responses are unable to reach the  $f + 1$  threshold, and therefore if some SURB and blinded public key are accepted by  $U_c$ , they must have been generated by honest discovery nodes.

**Lemma B.3.** An adversary controlling up to  $f$  discovery nodes, and the mix nodes in all but one of the Loopix network layers, gains only a negligible advantage in winning  $Exp_{G1}$ .

Let us consider all of the messages that the adversary-controlled nodes receive from the challenger. In LOOKUP, each adversary-controlled discovery node receives a fresh random nonce, the searched username  $ID_C$ , and a SURB for replying to  $U_b$ . In CONTACTINIT, an adversary-controlled discovery node may receive  $M_{init}$ , whose mix header is the SURB generated by discovery nodes, and whose payload is  $g^a$ , the same nonce, and a ciphertext encrypted with the blinded public key of  $U_c$ . The adversary cannot learn any information from these messages that would increase its advantage in winning the game. The Sphinx packet format used by Loopix does not reveal anything about the sender of a message. The SURB for  $U_b$  is cryptographically

indistinguishable from a uniform random bit string [19], and hence does not reveal its destination. Knowing that  $ID_C$  is the target of the search does not help distinguish whether  $U_0$  or  $U_1$  is the searcher, and the nonce and  $g^a$  are fresh random values.

By Lemma B.2, the blinded public key  $bpk_B$  that  $U_b$  uses to encrypt  $M_{init}$  is generated by the honest discovery nodes by blinding the public key in the contact information for  $U_C$ , and the corresponding private key is not known to  $\mathcal{A}$ . In the case of a lookup of a nonexistent username,  $M_{init}$  is encrypted using the public key in  $\Delta_{fake}$ , which is chosen so that no party knows the corresponding private key, as noted in §4.1. In either case, the ciphertext in  $M_{init}$  cannot be decrypted by the adversary except with negligible probability, and therefore the adversary does not learn the username of  $U_b$  or the codeword that  $U_b$  included in the message.

In the case where the adversary controls  $U_b$ 's provider node or the first mix node along the path of  $M_{init}$ , the adversary would be able to recognise the mix header of  $M_{init}$  and link it to  $U_b$  if  $U_b$  were to place it in the network directly. However, by reflecting  $M_{init}$  off a discovery node, with the communication between  $U_b$  and the mix node going via Loopix, the adversary becomes unable to link  $M_{init}$  to any particular sender (even if the reflecting discovery node is controlled by the adversary).

**Lemma B.4.** An adversary who controls up to  $f$  discovery nodes, and who performs an active attack on the CONTACTINIT and ADDFRIEND protocols, gains only a negligible advantage in winning  $Exp_{G1}$ .

Lemma B.3 already shows that the adversary gains no advantage from the contents of  $M_{init}$ . The other two messages sent during the ADDFRIEND protocol (from  $U_C$  to  $U_b$  in Step 3, and from  $U_b$  to  $U_C$  in Step 6) are not normally seen by the adversary (except in encrypted form travelling through the mix network). However, through an active attack, the adversary could inject its own SURBs into a protocol run, and thereby  $\mathcal{A}$  could trick  $U_C$  and  $U_b$  into sending these messages to  $U_A$  instead of each other.

Intercepting the messages in this way gives the adversary only negligible advantage in winning  $Exp_{G1}$ . Both messages contain only a key-blinded signature and a MAC, and the message from  $U_C$  to  $U_b$  additionally contains a Diffie-Hellman key  $g^b$ . The signature is unlinkable, i.e. the adversary cannot distinguish between two signatures produced from two separate signing keys, and two signatures produced from the same signing key but with different blinding keys [26]; since our protocol generates a different blinding key on every protocol run with overwhelming probability, this means that the signature cannot be linked to any particular user. The MAC is based on a key that is not known to  $\mathcal{A}$  (see Lemma B.6), and therefore the adversary cannot distinguish it from a random string. And finally,  $g^b$  is a fresh random value. Therefore, even an active attack that redirects messages to the adversary does not impact unlinkability. (We examine active attacks on impersonation in Lemma B.6.)

**Theorem B.5.** The adversary's advantage over a random guess in winning the unlinkability game  $Exp_{G1}$  can be reduced to the probability of breaking sender-receiver unlinkability in the underlying Loopix network.

The adversary gains only negligible advantage from inspecting message contents or returning bad responses at the discovery nodes (Lemma B.3), or from performing active attacks on the authenticated key exchange protocol (Lemma B.4). This leaves observing or manipulating the traffic between nodes as the only remaining attack, through which the adversary can win  $Exp_{G1}$  only by breaking sender-receiver unlinkability in the underlying Loopix network (Lemma B.1).

## B.2. Security against impersonation (G2)

**Lemma B.6.** An adversary who controls up to  $f$  discovery nodes, and who performs an active attack on the CONTACTINIT and ADDFRIEND protocols, has only a negligible chance of obtaining the session key  $K_s$  in  $Exp_{G2}$ .

By controlling at least one discovery node, the adversary learns the username  $ID$  that  $U_C$  is searching for, the nonce, the contact information  $\Delta_{\mathcal{H}}$  for user  $U_{\mathcal{H}}$ , one or more SURBs that allow  $\mathcal{A}$  to send messages to  $U_C$  without knowing  $U_C$ 's identity, the SURB for constructing  $M_{init}$ , the blinded public key  $bpk_B$  for  $U_{\mathcal{H}}$ , and the blinding key  $y_B$  for  $U_{\mathcal{H}}$ . If  $U_C$  picks an adversary-controlled discovery node as reflector,  $\mathcal{A}$  additionally learns the Diffie-Hellman key  $g^a$  generated by  $U_C$  and the ciphertext in  $M_{init}$ . The adversary can use this information to inject modified versions of  $M_{init}$  and the ADDFRIEND protocol messages (see Figure 3) into the network.

First consider the  $M_{init}$  message.  $\mathcal{A}$  can construct her own version  $M'_{init}$  of this message, containing the nonce from  $U_C$ . If she uses  $U_C$ 's DH key  $g^a$ , her chance of constructing a modified ciphertext that is correctly authenticated under  $K_e = \text{KDF}((bpk_B)^a \parallel \text{"init key"})$  is negligible without knowledge of  $a$ ; therefore she must either forward the unmodified ciphertext or replace  $g^a$  with a new DH key  $g^{a'}$  (where  $a' \in \mathbb{Z}_p$ ) and generate a new ciphertext. If she chooses a new DH key,  $\mathcal{A}$  can include any username and/or codeword of her choice in the new ciphertext (but not the codeword in  $M_{init}$ , which by Lemma B.3 she cannot decrypt), and she can include a SURB that will cause  $U_{\mathcal{H}}$ 's reply to be routed back to  $\mathcal{A}$ .

$U_{\mathcal{H}}$  treats the adversary's  $M'_{init}$  like any other contact request, and decides whether to respond based on the username and/or codeword in the message. If  $U_{\mathcal{H}}$  chooses to respond and the SURB in  $M'_{init}$  is routed to  $\mathcal{A}$ , the adversary learns  $g^b$ , a blinded-key signature over the Diffie-Hellman terms, and a MAC over  $U_{\mathcal{H}}$ 's identity.  $\mathcal{A}$  can also let the response from  $U_{\mathcal{H}}$  go directly to  $U_C$  without modification. However, if  $\mathcal{A}$  replaced  $g^a$  with  $g^{a'}$  in  $M'_{init}$ , and the response from  $U_{\mathcal{H}}$  is unmodified, then  $U_C$  will reject the signature, and the active attack on  $M_{init}$  leads to an abort of the protocol.

The adversary could also inject an altered version of the message from  $U_{\mathcal{H}}$  to  $U_C$  (ADDFRIEND Step 3) into the network, using one of the SURBs for  $U_C$ . However,  $U_{\mathcal{H}}$  will only accept a message containing a signature that validates with the blinded public key  $bpk_B$ , which by Lemma B.2 is generated by honest discovery nodes and is not under the adversary’s control. Due to the existential unforgeability of the key-blinded signature scheme,  $\mathcal{A}$  has a negligible chance of generating a valid key-blinded signature on this message without knowing  $U_{\mathcal{H}}$ ’s private key  $x_B$ .  $\mathcal{A}$  could replay a signature and  $g^b$  Diffie-Hellman term from another protocol run, but in that case  $\mathcal{A}$  would not know  $b$ , hence  $\mathcal{A}$  cannot compute  $K_m$ , and hence  $\mathcal{A}$  has a negligible chance of forging the MAC on this message. In any case, the adversary cannot inject an altered version of this message that is accepted by  $U_C$ .

Finally, the adversary could inject an altered version of the final message from  $U_C$  to  $U_{\mathcal{H}}$  (ADDFRIEND Step 6). If  $M_{\text{init}}$  was previously unaltered,  $U_{\mathcal{H}}$  will accept this message only with a key-blinded signature that is validated with  $U_C$ ’s blinded public key  $bpk_A$  (which was either included unaltered in  $M_{\text{init}}$ , or which was obtained from  $U_C$ ’s username in  $M_{\text{init}}$  via LOOKUP, which by Lemma B.2 is generated by honest discovery nodes and is not under the adversary’s control); by a similar argument to the previous paragraph,  $\mathcal{A}$  cannot construct a modified message where both the signature and the MAC are valid.

If  $M_{\text{init}}$  was previously replaced by  $M'_{\text{init}}$ , then the last message in the ADDFRIEND protocol must have also been generated by  $\mathcal{A}$ , since  $U_C$  would abort the protocol and would not generate such a message. In this case, since the username or blinded public key and the Diffie-Hellman term  $g^{a'}$  in  $M'_{\text{init}}$  are under the adversary’s control,  $\mathcal{A}$  is able to produce a valid key-blinded signature and MAC for this message. However, the session key derived from this protocol run is  $K'_s = \text{KDF}(g^{a'b} \parallel \text{“session key”})$ , whereas  $U_C$  computed the session key as  $K_s = \text{KDF}(g^{ab} \parallel \text{“session key”})$ .  $K_s$  and  $K'_s$  are different with overwhelming probability, and  $\mathcal{A}$  cannot compute  $K_s$  since she does not know  $a$ .

**Theorem B.7.** The adversary cannot win the impersonation game  $\text{Exp}_{G2}$  with a significantly better chance than random guessing.

By Lemma B.6, the adversary does not obtain the session key  $K_s$ . Therefore, when  $U_C$  encrypts  $m_b$  with  $K_s$  using a symmetric cipher that provides indistinguishability under chosen plaintext attack,  $\mathcal{A}$  has only negligible advantage over randomly guessing  $b$ .

## Appendix C. Meta-Review

The following meta-review was prepared by the program committee for the 2024 IEEE Symposium on Security and Privacy (S&P) as part of the review process as detailed in the call for papers.

### C.1. Summary

This paper proposes Pudding, a user-discovery mechanism for mix networks. User names are email addresses, and users register their public key and (mixnet) location information with discovery nodes, verified by email. When wishing to contact another user, the users reach out to the discovery nodes to fetch a single-use reply block for the destination user, along with that user’s blinded public key. Pudding provides user discovery with unlinkability, security against impersonation, external identity verification, and membership unobservability. The authors implement Pudding in Rust on top of the Nym network and show that it is useful in practice.

### C.2. Scientific Contributions

- Creates a New Tool to Enable Future Science
- Provides a Valuable Step Forward in an Established Field

### C.3. Reasons for Acceptance

- This paper provides a valuable step forward in an established field. User discovery is an important aspect of, but is often missing from, anonymous communication systems. Pudding provides this functionality for Loopix, a prominent mixnet design.
- This paper creates a new tool to enable future science. The authors wrote an implementation of Pudding in Rust that works with the Nym anonymity network that the authors will release open-source under an MIT license.

### C.4. Noteworthy Concerns

Much trust is placed in the Pudding discovery nodes. These nodes:

- Know who is registered in the system
- Know the location (i.e., provider identity) of each registered user
- Know when every user query is made

As a result, membership unobservability is provided only against adversary-controlled clients but not servers. This strong trust assumption may not be tolerable in practice; future work should consider ways to relax the trust assumption placed on these nodes.