

# An asynchronous provably-secure protocol for hidden services

Philippe Camacho and Fernando Krell  
{philippe.camacho,fernando.krell}@dreamlab.net

September 13, 2017

## Abstract

The client-server architecture is one of the most widely used in Internet for its simplicity and flexibility. In practice the server is assigned a public address so that its services can be consumed. This makes the server vulnerable to a number of attacks such as Distributed Denial of Service (DDoS), censorship from authoritarian governments or exploitation of software vulnerabilities.

In this work we propose an asynchronous protocol for allowing a client to issue requests to a server without leaking any information about the location of the server. In addition, our solution reveals limited information about the network topology, leaking the distance from the client to the corrupted participants.

We also provide a simulation-based security definition capturing the requirement described above. Our protocol is secure in the semi-honest model against any number of colluding participants. Moreover our solution is efficient as it requires  $O(N \cdot |M|)$  bits per client-server interaction where  $N$  is the number of participants and  $|M|$  is the number of bits of the message.

To the best of our knowledge our solution is the first asynchronous protocol that provides strong security guarantees.

# 1 Introduction

## 1.1 Motivation

The client-server architecture is one of the most widely used in Internet for its simplicity and flexibility. In practice the server is assigned a domain name and one or more IP addresses so that its services can be consumed. This makes the server vulnerable to a number of attacks such as DDoS, censorship from authoritarian governments or exploitation of software vulnerabilities. Thus it would be desirable to hide the location of the server in the network. By doing so an attacker will not be able to attack directly the host containing the server's code nor interrupt the execution of its services by non-technical means. While the literature is abundant on the topic of anonymous channels [6, 5, 17, 18], the problem of hiding the location of a server remains of great interest. TOR hidden services [7] is without a doubt the most popular alternative for this purpose. Unfortunately, the security provided by TOR is not guaranteed; in fact, several practical attacks have been discovered [15, 12, 20, 24].

We observe that simple solutions for the problem described above do not work. Standard end-to-end encryption is vulnerable to tracing the ciphertext across the network, and hence, an adversary that is powerful enough to corrupt several nodes is very likely to detect the origin or destination of the message. Other approaches like using multicast are not enough either since clients that are close to the server will notice that the response comes back within short time. The main challenge is to avoid nodes to distinguish whether the server is close or far away.

In this work we focus on solving the following problem. A set of clients wish to establish a communication with a server, yet we want to hide the location of this server in the network. We also expect the client's queries and server's responses to remain private.

At a high level our protocol implements two phases: (1) a client issues a *request* to the server, and then (2) the server returns a *response*. The first phase of the protocol is straightforward: the client encrypts the request using the public key of the server and then multicasts the message across the network. Note that the server must still forward the request as if it were any other node, otherwise its neighbors may infer its location. The second phase is much more complex because as mentioned above the client or other nodes could detect the presence of the server by a simple timing

attack. To circumvent this difficulty we introduce the following idea: we force all the nodes to behave as the server. We achieve this by relying on a secret sharing scheme where every participant holds a share of the response. To perform this split-and-reconstruct phase every node including the server generate a random share. Then all shares are propagated to the server. At this stage the server replaces its share by a value that enables to reconstruct the response. In order to improve performance, we use a spanning tree<sup>1</sup> over the network graph. This allows us to optimize multicast invocations and shares aggregation. We emphasize that our protocol is asynchronous, which means that participants do not rely a on shared clock to run the protocol, but rather acts upon the reception of neighbors messages. Unfortunately, asynchronism comes at price: Since nodes do not know when a participant initiates a request, it is impossible to hide the requesters activity. Hence our protocol leaks proximity information of the requester to other nodes.

## 1.2 Contributions

Our contributions are the following:

- To the best of our knowledge we provide the first simulation-based security definition capturing the requirement of hiding a server in a network. This definition considers the full interaction (request and response) between clients and server.
- To the best of our knowledge our protocol is the first to provide strong security guarantees in an asynchronous setting (see Figure 1).
- Our protocol is secure against any number of corrupted participants: Indeed if the adversary controls all nodes but two (including the server) it will not be able to guess the right location with probability better than  $\frac{1}{2}$ .
- Our protocol is efficient, requiring  $O(N \cdot |M|)$  bits per client-server interaction. Note that a linear communication complexity is required as nodes that do not communicate would implicitly leak some information about the location of the server.

---

<sup>1</sup>which we borrow from Dolev and Ostrovsky [8].

### 1.3 Related Work

While the problem of hiding the physical location of a server in a network is not exactly an anonymity problem (we do not want to hide the fact that a specific client connects to the server) the techniques and concepts we use are borrowed from the area of anonymity. Since Chaum’s two seminal papers on mixes [6, 5], a large body of work has been written in order to enable communications that do not reveal the identity of participants. An alternative to mixers for achieving anonymity has been introduced by Reiter *et al.* with a protocol named Crowds[19] and consists of using random paths among a set of “dummy” nodes a.k.a. *jondo* before reaching a specific destination (the server). In this protocol – contrary to our setting – the location of the server is public and the goal is to hide the clients. This solution is simple, efficient and provide some level of anonymity for the client. Beyond the protocol itself, the authors highlight some fundamental problems that arise with these types of constructions where traffic is routed through possible corrupted nodes: In particular, preserving the initiator’s anonymity turns out to be more complex than expected [23, 21]. Indeed in our case, we have to solve a similar problem where we must hide the location of the server during the phase of responding a request. Hordes [13] is an improvement to Crowds where the reply from the server is done using multicast. This change makes passive attacks consisting in tracing back messages harder while adding only a reasonable operational cost. While Crowds and Hordes do not aim to hide the server like we do, these protocols highlight the difficulty of hiding nodes in a network where the adversary controls a subset of the participants and can leverage traffic analysis. Another approach to establish anonymous channels between client and servers is onion routing [9]. An onion is obtained by encrypting the message in a layered fashion using the public keys of the nodes on a path from sender to receiver. By doing so, a node on the circuit will not be able to identify the original source, the final destination, nor the message itself. The most popular onion routing protocol is without a doubt TOR [7]. TOR not only enables to preserve the anonymity of clients but also provides a mechanism to hide the location of the server through a *rendez-vous* node where both client and server meet. Unfortunately, as in Crowds and Hordes, a number of practical attacks based on traffic analysis are possible [12, 20, 24, 16]: In particular if a node manages to be the first relay between the server and the *rendez-vous* node, it will likely detect the server presence [16]. In case managing a Public-Key Infrastructure is too complex, one can use Katti *et*

Protocol	Asynchronous	Collusion-resistant	Communication complexity
TOR [7]	YES	NO	$O(D^2 \cdot \kappa + D M )$
Dolev and Ostrovsky [8]	NO	Up to $\lfloor (N-1)/2 \rfloor$	$O(N \cdot  M )$
MPC-Hiding topology [1]	NO	YES	$O(\kappa(\kappa + \log N)N^9 \cdot  M )$
Our work	YES	YES	$O(N \cdot  M )$

Figure 1: **Comparison of protocols for hiding a node location.** In this table  $N$  is the number of participants,  $D$  is the diameter of the graph representing the network,  $|M|$  is the number of bits of the message and  $\kappa$  is the security parameter. TOR is not collusion resistant because some attacks can succeed with only two corrupted nodes [16]. We assume also that the length of the onion circuit should be proportional to the diameter  $D$  of the graph as otherwise traffic analysis attacks would become more effective. Regarding communication complexity, we do not take into account the setup phase occurring in Dolev and Ostrovsky’s construction and ours.

*al.*’s protocol [11] that relies on the idea of splitting the routing information in such a way that only the right nodes on the circuit are able to reconstruct it correctly. In our protocol we also leverage secret-sharing techniques, but for splitting and reconstructing the message only. Also our solution does not require a sender to control different nodes as in the onion slicing approach.

Early attempts to counter traffic analysis attacks were not practical as they assumed the existence of some broadcast channel or ad-hoc topology and required a synchronous execution [5, 18, 22]. The more general problem of hiding the topology of a network has been solved recently in the Secure multi-party computation setting [1, 14, 10]. However, these solutions involves a lot of communication and computational overhead. One of the most promising attempts for hiding the location of a server was due to Dolev and Ostrovsky [8]: Indeed our solution borrows some of the techniques of their work, in particular we also use spanning-trees to make the multicast communications more efficient. Nonetheless our solution has two major advantages: it is asynchronous and it is secure against any number of corrupted nodes.

In Figure 1 we compare our work with other proposals that allow arbitrary topologies. We observe that our construction is the most efficient with Dolev and Ostrovsky’s one [8] but is asynchronous and tolerates up to  $N - 2$  corrupted participants<sup>2</sup>.

---

<sup>2</sup> $N - 1$  corrupted participants is a trivial case where the adversary controls all nodes which are not the server.

## 1.4 Organization of the paper

This paper is organized as follows. Section 2 introduces definitions and notations. The abstract functionality capturing the secure interaction between client and server is introduced in Section 3. We describe our main protocol in Section 4, prove its security in Section 5 before concluding in Section 6.

## 2 Preliminaries

### 2.1 Definitions and notations

Let  $n \in \mathbb{N}$  be an integer, we denote by  $[n]$  the set  $\{1, 2, 3, \dots, n\}$ .

For a graph  $G = \langle V, E \rangle$  the distance  $d(u, v)$  between two vertices  $u$  and  $v$  is the length of the shortest path between  $u$  and  $v$ . Let  $(\mathbf{M}, \circ)$  be an abelian group and  $\kappa \in \mathbb{N}$  the security parameter. A (single-operation) homomorphic encryption scheme over message space  $\mathbf{M}$  is a tuple of algorithms  $\langle \text{Gen}, \text{Enc}, \text{Dec}, \text{Add} \rangle$  in which  $\langle \text{Gen}, \text{Enc}, \text{Dec} \rangle$  is a regular public-key encryption scheme and algorithm  $\text{Add}$  satisfy the following property: For every valid key-pair  $(\text{pk}, \text{sk}) \leftarrow \text{Gen}(1^\kappa)$ , and for every pair of messages  $m_1, m_2 \in \mathbf{M}$ :

$$\text{Dec}_{\text{sk}}(\text{Add}_{\text{pk}}(\text{Enc}_{\text{pk}}(m_1), \text{Enc}_{\text{pk}}(m_2))) = m_1 \circ m_2$$

For some arbitrary ciphertext set  $C = \{c_i = \text{Enc}_{\text{pk}}(m_i)\}_{i \in I}$ , we abuse notation by using  $\sum_{i \in I} c_i$  or  $\text{Enc}_{\text{pk}}(\sum_{i \in I} m_i)$  to denote the result of a sequential computation of  $\text{Add}_{\text{pk}}$  over  $C$ .

### 2.2 Modeling networks

We can think of a regular communication network as a graph  $G$ , composed by a set of nodes  $V$  and a set of edges  $E$  between them. Participants (nodes)  $v_i$  and  $v_j$  cannot communicate directly unless there is an edge  $(v_i, v_j)$  in  $E$ . To allow communication between distant participants, nodes can forward incoming messages to neighbor nodes following some protocol.

We use the approach of [10] in which the participants in the real protocol have access to a *network* functionality that allows to send message between neighbor participants. The network functionality is specified in Figure 2, and allows any participant to send messages to a neighbor at an arbitrary

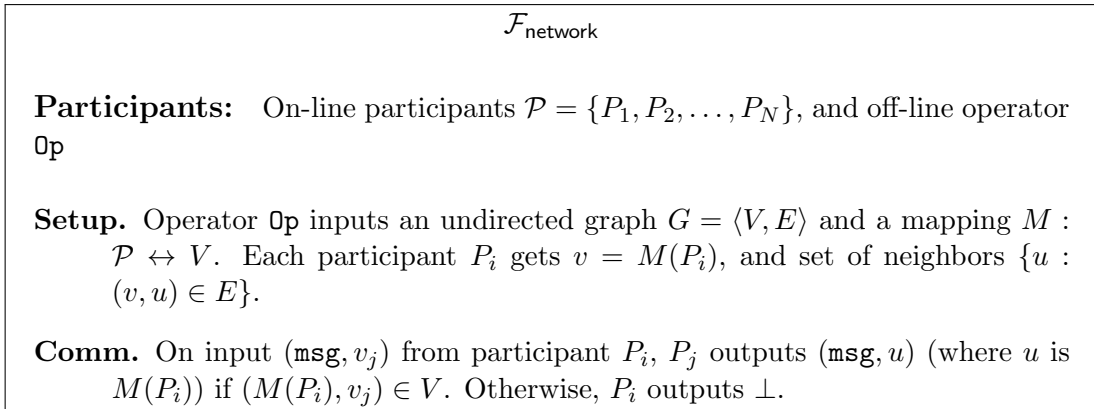


Figure 2: Physical Network Functionality

time<sup>3</sup>. It provides two services, **Setup** and **Comm**. On the setup phase, the communication graph is specified. This can be done by an off-line operator, or by the participant itself describing their neighbors (or their pseudonyms as inputs). The **Comm** service allows for neighbor participant to exchange messages.

We will use this functionality as the basic mechanism to send message throughout the network. For simplicity, we require that **Setup** is called before any **Comm** service can be processed.

### 2.3 Multicast protocol

In this section we describe a simple multicast protocol that uses functionality  $\mathcal{F}_{\text{network}}$  as its basic communication mechanism. We assume that a trusted party has already instantiated the network functionality, and hence each participant knows the vertex label associated with its neighbor for functionality  $\mathcal{F}_{\text{network}}$ . When a participant issues a multicast, it sends the message to its neighbor using functionality  $\mathcal{F}_{\text{network}}$ . Each participant, upon reception of a multicast message, first check if the message has not been seen before. In this case, it forwards the message to its neighbors and outputs the message. Jumping ahead, our main protocol will use this functionality on a subgraph

---

<sup>3</sup>The network functionality of [10] is rather different in the sense that all participant call it at same time, and all have message to all its neighbors.

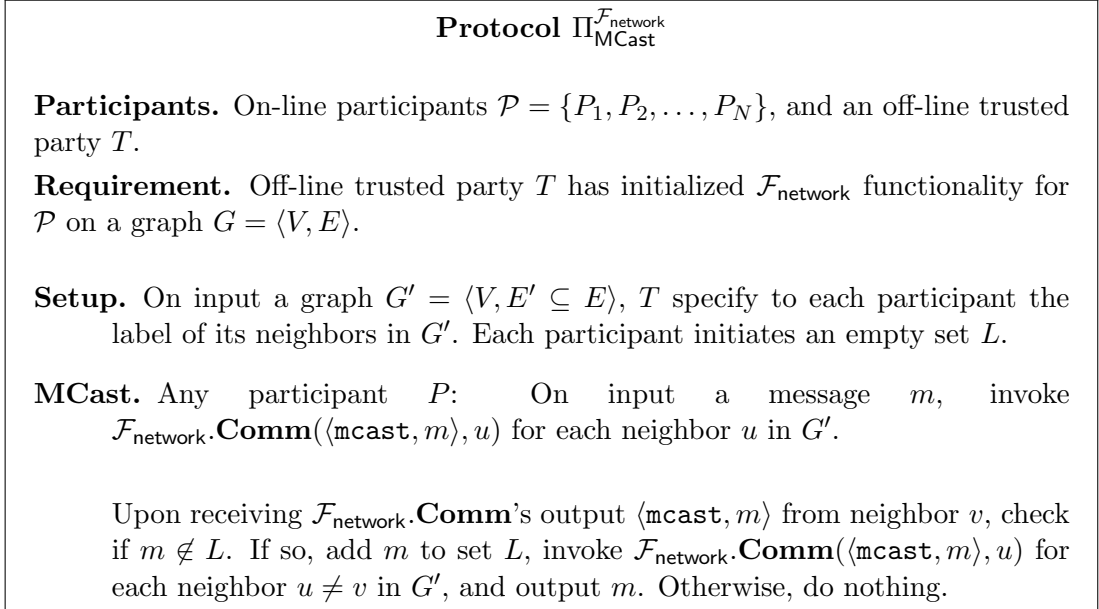


Figure 3:  $\Pi_{\text{MCast}}^{\mathcal{F}_{\text{network}}}$

of the network graph to efficiently broadcast the client's encrypted requests.

## 2.4 Security definition

As standard in cryptographic protocols, we define security in terms of a real-versus-ideal world procedures. That is, we first specify a desired functionality for our protocol, and then claim that a protocol computing the functionality is *secure* if its real-world execution realizes an ideal procedure in which the participants get their outputs by sending their inputs to a trusted party computing the functionality on behalf of them. More specifically, we say that our protocol *privately computes* the functionality if whatever can be computed by adversary interacting in the real execution of the protocol, can also be obtained with only inputs and outputs of the corrupted participants in the ideal execution.

We now formally provide a security definition for *semi-honest static* adversaries. In what follows we let algorithms  $\text{Sim}$ ,  $\text{Adv}$ , and  $\mathcal{Z}$  be state-full.

$\text{Ideal}_{\mathcal{Z}, \text{Sim}}^{\mathcal{F}}(\kappa)$ : 1) Run  $\mathcal{Z}(1^\kappa)$  to produce participant inputs  $\{\text{in}_j\}_{j \in [N]}$  and ad-



versary input  $\text{in}_{\text{sim}}$ . 2) Run  $\text{Sim}(1^\kappa, \text{in}_{\text{sim}})$  to get the index set of corrupted parties  $I_C \subseteq [N]$ . 3) Run  $\text{Sim}(\{\text{in}_k\}_{k \in I_C})$  to obtain modified input  $\{\text{in}'_k\}_{k \in I_C}$  for the corrupted parties. 4) Call functionality  $\mathcal{F}$  on previous inputs to obtain output  $\{\text{out}_j\}_{j \in [N]}$ . 5) Run  $\text{Sim}(\{\text{out}_k\}_{k \in I_C})$  to get adversary's output  $\text{out}_{\text{sim}}$ . 6) Run  $\mathcal{Z}(\{\text{out}_j\}_{j \in [N] \setminus I_C}, \text{out}_{\text{sim}})$  to obtain output bit  $b$ . 7) Return  $b$  as the output of the ideal-world execution.

$\text{Real}_{\mathcal{Z}, \text{Adv}}^\Pi(\kappa)$ : 1) Run  $\mathcal{Z}(1^\kappa)$  to produce participant inputs  $\{\text{in}_j\}_{j \in [N]}$  and adversary input  $\text{in}_{\text{adv}}$ . 2) Run  $\text{Adv}(1^\kappa, \text{in}_{\text{adv}})$  to get set of corrupted parties  $I_C \subseteq [N]$ . 3) Run  $\text{Adv}(\{\text{in}_k\}_{k \in I_C})$  to obtain modified input  $\{\text{in}'_k\}_{k \in I_C}$  for the corrupted parties. 4) Execute protocol  $\Pi$  with previously computed inputs, saving the view of every corrupted participant,  $\{\text{view}_k\}_{k \in I_C}$ . When every participant finishes the protocol execution, recollect output of every uncorrupted participants,  $\{\text{out}_j\}_{j \in [N] \setminus I_C}$ . 5) Run  $\text{Adv}(\{\text{view}_k\}_{k \in I_C})$  to get adversary's output  $\text{out}_{\text{adv}}$ . 6) Run  $\mathcal{Z}(\{\text{out}_j\}_{j \in [N] \setminus I_C}, \text{out}_{\text{adv}})$  to obtain output bit  $b$ . 7) Return  $b$  as the output of the real-world execution.

**Definition 1.** *A protocol  $\Pi$  privately computes functionality  $\mathcal{F}$  if for every PPT algorithm  $\text{Adv}$ , there exists a PPT algorithm  $\text{Sim}$  such that for every PPT algorithm  $\mathcal{Z}$  the random variables  $\text{Ideal}_{\mathcal{Z}, \text{Sim}}^\mathcal{F}(1^\kappa)$  and  $\text{Real}_{\mathcal{Z}, \text{Adv}}^\Pi(1^\kappa)$  are computationally indistinguishable, for all sufficiently long  $\kappa$ .*

In our work it is sufficient to show a PPT simulator  $\text{Sim}$  that can produce a view that is computationally indistinguishable from the corrupted participants view. Then, the simulator can run  $\mathcal{A}$  to produce a simulated output to  $\mathcal{Z}$ .

We slightly modify the ideal world to include a leakage function,  $\mathcal{L}$ , whose output is leaked to the simulator  $\text{Sim}$ . This leakage model the fact the protocol may reveal some partial private information to the adversary (for example, the length of the messages to encrypt). It also allows for the specification of trade-offs between protocol features or efficiency and security. This leakage information is added to the simulator's input on step 5.

### 3 Request Response Functionality

The functionality is executed between a set of participant  $\mathcal{P} = \{P_1, P_2, P_3, \dots\}$ . A server node, which we denote as  $S$ , provides an arbitrary polynomial-time requests-response service for all participants. A protocol realizing this

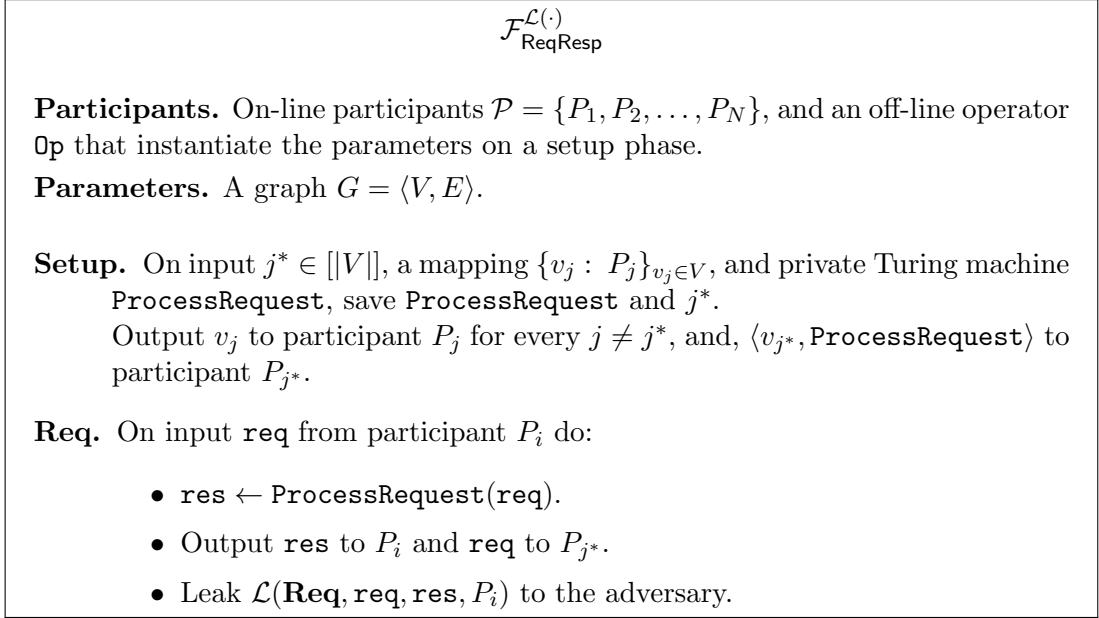


Figure 4: Hidden-Server Request-Response functionality  $\mathcal{F}_{\text{ReqResp}}$  over an incomplete network with leakage profile  $\mathcal{L}(\cdot)$ .

functionality needs to hide which of the participant is the server node. A secondary desired goal is to also protect the requests and the responses (including the request issuer and destination nodes). Hence, our functionality allows for the specification of a leakage profile over the request and response; however, we do not allow any leakage on who the server is or any secret information on what it is computing.

The functionality is parametrized by a public graph. On a setup phase, the operator participant  $\text{Op}$  specifies the server node, its service Turing machine  $\text{ProcessRequest}$ , and a mapping between graph nodes and participants. As a result of this setup phase, every node gets its graph label, and the server node gets the Turing machine  $\text{ProcessRequest}$ .

## 4 Protocol Design

### 4.1 Overview

For a set of participants  $\mathcal{P} = \{P_1, P_2, \dots, P_N\}$  communicating over an arbitrary network graph  $G$ , the goal of our protocol is to hide the location of a server  $\mathcal{S} = P_{j^*}$  in  $G$  while enabling other participants to consume its services. The main difficulty is to make it impossible for an adversary to leverage timing information to obtain (or estimate) the distance between  $\mathcal{S}$  and some other corrupted nodes in  $G$ .

The protocol proceeds in two high level steps. The first step corresponds to enabling a client  $P_i$  to send a request  $\mathbf{req}$  to the server  $\mathcal{S}$ . This step can be easily implemented using a multicast protocol (see Section 2.3): The client encrypts  $\mathbf{req}$  using  $\mathcal{S}$ 's public key and multicasts the ciphertext  $c = \mathbf{Enc}_{\mathbf{pk}^{\mathcal{S}}}(\mathbf{req})$ . Indeed,  $\mathcal{S}$ 's location is not leaked<sup>4</sup>.

The second step consists of letting the server  $\mathcal{S}$  to send the response  $\mathbf{res}$  back to  $P_i$ . This turns out to be more challenging. Indeed, proceeding as in the first step is not secure since nodes that are close to  $\mathcal{S}$  would detect  $\mathcal{S}$ 's activity and be able to deduce its location or some information about it (as for example the subnet that contains  $\mathcal{S}$ ). In order to circumvent this difficulty we introduce the following high level idea: each node  $P_j$  sends a random share  $s_j$  to the server  $\mathcal{S}$  (including the  $\mathcal{S}$  itself). The server will obtain all the shares  $\{\mathbf{share}_j\}_{j \neq i}$  and recompute its share  $\mathbf{share}_{j^*}$  so that combination of all shares reconstruct to  $\mathbf{res}$ . Then, all the participants send their shares to the requester  $P_i$ , and finally,  $P_i$  reconstruct and outputs the response.

Since shares on the last step reconstruct the response, it is clear that they need to be encrypted under  $P_i$ 's public-key. As the initial shares sent to the server reconstruct to a random value, it is tempting to send these in plaintext. However, an adversary that controls the requester can see the shares both times, and therefore notice when a share was updated, inferring information on  $\mathcal{S}$ 's location.

We take the approach of [8] and restrict the communication to an (arbitrary) spanning tree on the network graph. This allows us to efficiently communicate the messages on all phases. In particular, we use the following mechanism to send the shares to  $\mathcal{S}$  and  $P_i$ : First, the shares are sent up to the root node of the spanning tree, and then the root node multicasts the shares down

---

<sup>4</sup>Note that messages needs to be forwarded once – and only once – to neighbors, even when the message has arrived to its destination

the tree. By using  $n$ -out-of- $n$  information-theoretic secret sharing, we note that nor the server or the requester need to know every individual share. In fact, they only need to learn the final secret. Our idea, hence, is to use homomorphic encryption on the shares, and have each internal node to “add-up” its share to the shares computed by its children, and then send a single result up the tree (rather than the individual shares of every node in its subtree). The root node then obtains an encrypted secret, which is sent down the tree to reach the server or the requester. This efficient procedure allow our final protocol to have linear communication complexity, and is formally described in Section 4.2.

Our full protocol implementing functionality  $\mathcal{F}_{\text{ReqResp}}$  is specified in Section 4.3.

## 4.2 Encrypted Share Reconstruction Protocol

In this section we describe an important sub-protocol of our solution. This protocol, denoted  $\Pi_{\text{ESR}}$ , allows to efficiently and privately reconstruct a secret out of each participant share. In a nutshell, each party encrypts its share under the public-key of the recipient, and sends the ciphertext up into a spanning tree of the network graph. The participant at the root node of this tree can homomorphically compute the encrypted secret, and then send the result down the tree to reach the recipient. We do this efficiently in the following way: Each internal node privately reconstructs part of the secret by homomorphically combining its encrypted share with the ciphertext obtained from its children. Hence, each internal node needs to send a single ciphertext up the tree. Furthermore, we use  $n$ -out-of- $n$  information-theoretic secret sharing so that we only need a single homomorphic operation for the encryption scheme. Protocol  $\Pi_{\text{ESR}}$  is specified in Figure 5.

## 4.3 Request-Response Server Protocol

In this section we introduce an  $\mathcal{F}_{\text{network}}$ -hybrid protocol achieving functionality  $\mathcal{F}_{\text{ReqResp}}$ .

Our protocol is divided in an off-line setup phase and three on-line phases. In the setup phase, a trusted party  $T$  chooses a server participant  $\mathcal{S}$  and generates for it a key-pair  $(\text{pk}^{\mathcal{S}}, \text{sk}^{\mathcal{S}})$ .  $T$  also chooses an arbitrary rooted spanning tree instantiate protocols  $\Pi_{\text{MCast}}$  and  $\Pi_{\text{ESR}}$ .

On the first on-line phase, the requester  $P_i$  encrypts its query  $\text{req}$  under

### Protocol $\Pi_{\text{ESR}}^{\mathcal{F}_{\text{network}}}$

**Participants.** On-line participants  $\mathcal{P} = \{P_1, P_2, \dots, P_N\}$ , and an off-line trusted party  $T$ .

**Parameters.** An homomorphic encryption scheme  $\mathcal{H} = \langle \text{Gen}, \text{Enc}, \text{Dec}, \text{Add} \rangle$ .

**Requirement.** Off-line trusted party  $T$  has initialized  $\mathcal{F}_{\text{network}}$  functionality for  $\mathcal{P}$  on a graph  $G = \langle V, E \rangle$ .

**Setup.** On input a spanning tree  $\text{ST} = \langle \text{root} \in V, E_{\text{ST}} \subset E \rangle$  over  $G$ ,  $T$  specifies to every participant its parent  $p$  and children set **children** on  $\text{ST}$ .

**SendUp.** Any participant  $P \in \mathcal{P}$ : On input a message  $m$ , public key  $\text{pk}$ , and session id  $\text{sid}$ , compute  $c = \text{Enc}(m, \text{pk})$  and store  $\langle \text{sid}, c \rangle$ .

If  $P_i$  has no children jump to  $\star$ .

Upon receiving  $\mathcal{F}_{\text{network}} \cdot \mathbf{Comm}$ 's output  $\langle \text{sid}, \text{up}, c' \rangle$  from children  $u$ , use  $\text{sid}$  to get  $c$  and update it to  $\text{Add}_{\text{pk}}(c, c')$ . If all children have submitted their **up** message, then jump to  $\star$ .

$\star$ : If  $P$  is the root of the tree, output  $\langle c, \text{sid} \rangle$ . Otherwise, invoke  $\mathcal{F}_{\text{network}} \cdot \mathbf{Comm}(\langle \text{sid}, \text{up}, c \rangle, p)$ .

**SendDown.** Participant  $P \in \mathcal{P}$  (root of the tree): On input a message  $c$  and session id  $\text{sid}$ , invoke  $\Pi_{\text{MCast}}^{\mathcal{F}_{\text{network}}} \cdot \mathbf{MCast}(\langle \text{sid}, \text{down}, c \rangle)$ , and output  $\langle c, \text{sid} \rangle$ .

(Any participant). Upon receiving  $\Pi_{\text{MCast}}^{\mathcal{F}_{\text{network}}} \cdot \mathbf{MCast}$ 's output  $\langle \text{sid}, \text{down}, c \rangle$ , output  $\langle c, \text{sid} \rangle$

Figure 5:  $\Pi_{\text{ESR}}^{\mathcal{F}_{\text{network}}}$

the server's public key, and uses protocol  $\Pi_{\text{MCast}}$  to *multicast* propagate the ciphertext across the network.

Then, on the second on-line phase every participant (including the server) generates a random string (used as a share for the response) and sends it to the server using protocol  $\Pi_{\text{ESR}}$ . Upon receiving the combined shares  $cs = \sum_{j \neq i} \text{share}_j$ ,  $\mathcal{S}$  recomputes its share  $\text{share}_{j^*}$  as  $\text{res} - (cs - \text{share}_{j^*})$  so that the reconstruction procedure outputs the response  $\text{res}$ .

On the third on-line phase, every participant  $P_j$  use  $\Pi_{\text{ESR}}$  to send its  $\text{share}_j$  (encrypted under  $P_i$ 's public key), so that the response can be homomorphically reconstructed and sent to  $P_i$ .  $P_i$  decrypts and output the response.

Notice that these three phases can be executed in a pipeline. In fact, each encrypted share sent on the second on-line phase can be sent as soon as the participant sees the request multicast message issued by  $P_i$  on the first phase. Similarly, each participant can send its share in the third phase as soon as the participant sees the multicast-down message issued by the root node in the second phase. Therefore, our protocol is *asynchronous*.

We also note that the initial multicast of the encrypted request leaks the direction towards the requester node to each participant. Therefore, the

encrypted response on the third phase, can be sent efficiently from the root to the requester. In fact, when a participant receives the request message from neighbor  $u$ , this is saved so that at the final phase, each participant knows where to send the encrypted response.

Since all participants act according to the same communication pattern, and all messages are encrypted, our protocol does not reveal the location of the server, nor the request or response.

Our protocol is formally described in Figure 6.

#### 4.4 Variants of the protocol

**Avoiding an off-line trusted party.** Protocol 6 relies on a trusted party to set up the initial parameters of each participant. By using state-of-the-art topology-hiding secure computation protocols [14, 10, 2, 1] we can achieve a secure distributed setup without any trusted party.

**Precomputing shares using PRG.** It is possible to simplify the protocol described in Figure 6 by having the server computing the other participant shares locally. In practice, all the participants would receive a secret seed  $R_j$  to generate its seed, and the server receives the secret seeds of every participant. This means that the second on-line phase of the protocol can be removed, and hence save  $2N$  in communication complexity and  $N$  homomorphic operation. The other steps remain unchanged.

**Response recipient.** Our protocol can be modified so that the recipient of the response can be any arbitrary participant (or set of participants). This is achieved as follows: (a) the client chooses the public key of another participant as the session public key, and (b) because the location of the recipient is not necessarily known, the root node multicasts the encrypted response down the tree instead of sending it directly to the originator of the request.

**Avoiding the use of the spanning tree.** In a practical environment, the spanning tree could affect the resilience of the protocol and can be hard to maintain or configure. In such a scenario, the steps (**SendUp,SendDown**) can be replaced by multicast operations of the shares for each participant. Note that in addition to increase the communication complexity, this change introduces a new challenge: The nodes need to wait for the server to decrypt all its shares before the shares for the client can be multicasted. In order to keep the protocol asynchronous, we solve this issue by *precomputing the shares using PRG* (see above) so that nodes are not forced to delay the release

## Protocol $\Pi_{\text{ReqResp}}^{\mathcal{F}_{\text{network}}}$

**Participants.** On-line  $\mathcal{P} = \{P_1, P_2, \dots, P_N\}$ , and an off-line trusted party  $T$ .

**Parameters.** A security parameter  $\kappa$  and an homomorphic encryption scheme  $\mathcal{H} = \langle \text{Gen}, \text{Enc}, \text{Dec}, \text{Add} \rangle$ .

**Requirement.** Off-line trusted party  $T$  has initialized  $\mathcal{F}_{\text{network}}$  functionality for  $\mathcal{P}$  on a graph  $G$ .

**Setup.** a)  $T$  chooses a server participant  $\mathcal{S} \in \mathcal{P}$  and an arbitrary spanning tree  $\text{ST}$  on graph  $G$ . b)  $T$  instantiate protocols  $\Pi_{\text{MCast}}^{\mathcal{F}_{\text{network}}}$  and  $\Pi_{\text{ESR}}^{\mathcal{F}_{\text{network}}}$  on input  $\text{ST}$ , c)  $T$  generates server's key pair  $(\text{pk}^{\mathcal{S}}, \text{sk}^{\mathcal{S}}) \leftarrow \text{Gen}(1^\kappa)$  and securely distributes  $\mathcal{S}$ 's public key  $\text{pk}^{\mathcal{S}}$  to every participant. d) Finally,  $T$  securely sends  $(\text{sk}^{\mathcal{S}}, \text{pk}^{\mathcal{S}})$  and a Turing Machine **ProcessRequest** to  $\mathcal{S}$ .

(In what follows, let  $j^*$  denote the index of  $\mathcal{S}$  in  $\mathcal{P}$ .)

**Req.** On input **req**, participant  $P_i$  chooses a session key-pair  $(\text{pk}_{\text{sid}}, \text{sk}_{\text{sid}})$  and a fresh session id **sid**, and invokes multicast protocol  $\Pi_{\text{MCast}}^{\mathcal{F}_{\text{network}}}.\text{MCast}(\langle \text{request\_to\_server}, \text{sid}, \text{Enc}_{\text{pk}^{\mathcal{S}}}(\text{req}), \text{pk}_{\text{sid}} \rangle)$  over the spanning tree.

**Response phase.** Every participant  $P_j$  (including  $\mathcal{S}$ ):

1. Upon receiving  $\Pi_{\text{MCast}}^{\mathcal{F}_{\text{network}}}.\text{MCast}$ 's output  $\langle \text{request\_to\_server}, \text{sid}, C, \text{pk}_{\text{sid}} \rangle$  from neighbor  $u$ , pick a random share  $\text{share}_j$  and invoke  $\Pi_{\text{ESR}}.\text{SendUp}(\text{share}_j, \text{pk}^{\mathcal{S}}, \text{sid})$ , and store  $\langle \text{sid}, \text{pk}_{\text{sid}}, \text{share}_j, u \rangle$ . In addition, if  $P_j$  is  $\mathcal{S}$ , compute  $\text{req} = \text{Dec}_{\text{sk}^{\mathcal{S}}}(C)$  and  $\text{res} \leftarrow \text{ProcessRequest}(\text{req})$ , and store  $\langle \text{sid}, \text{req}, \text{res} \rangle$ .
2. (Root node) Upon receiving  $\Pi_{\text{ESR}}.\text{SendUp}$ 's output  $(C = \text{Enc}_{\text{pk}^{\mathcal{S}}}(\sum_{j \neq i} \text{share}_j), \text{sid})$ , invoke  $\Pi_{\text{ESR}}.\text{SendDown}(C, \text{sid})$ .
3. Upon receiving  $\Pi_{\text{ESR}}.\text{SendDown}$ 's output  $(C, \text{sid})$ , use **sid** to get  $\text{pk}_{\text{sid}}$  and  $\text{share}_j$  from local storage, and:
  - If  $P_j$  is  $\mathcal{S} = P_{j^*}$ , decrypt  $C$  to get  $\text{share}_{\text{sum}} = \sum_{j \neq i} \text{share}_j$ , and update  $\text{share}_{j^*}$  to  $\text{res} - (\text{share}_{\text{sum}} - \text{share}_{j^*})$ .
  - Invoke  $\Pi_{\text{ESR}}.\text{SendUp}(\text{share}_j, \text{pk}_{\text{sid}}, \text{sid})$
4. (Root node) Upon receiving  $\Pi_{\text{ESR}}.\text{SendUp}$ 's output  $(C = \text{Enc}_{\text{pk}_{\text{sid}}}(\sum_{j \neq i} \text{share}_j), \text{sid})$ , use **sid** to get neighbor label  $u$ , and invoke  $\mathcal{F}_{\text{network}}.\text{Comm}(\langle C, \text{sid} \rangle, u)$ .
5. Upon receiving  $\mathcal{F}_{\text{network}}.\text{Comm}$ 's output  $(C, \text{sid})$  do:
  - If  $P_j$  is  $P_i$ , use **sid** to get  $\text{sk}_{\text{sid}}$  from local storage, and output  $\text{res} \leftarrow \text{Dec}_{\text{sk}_{\text{sid}}}(C)$ .
  - Otherwise, use **sid** to get  $u$ , and invoke  $\mathcal{F}_{\text{network}}.\text{Comm}(C, u)$ .

Figure 6:  $\Pi_{\text{ReqResp}}^{\mathcal{F}_{\text{network}}}$

of their shares to the client.

## 5 Proof of Security

We begin by specifying the private information revealed by protocol  $\Pi_{\text{ReqResp}}$ .

**Leakage 1.**  $\mathcal{L}(G, \text{ST}, M, P_i, \mathcal{C})$  On input a graph  $G = \langle V, E \rangle$ , a spanning tree  $\text{ST} = \langle \text{root} \in V, T \subset E \rangle$  over  $G$ , a mapping  $M := \mathcal{P} \leftrightarrow V$ , a requester participant  $P_i \in \mathcal{P}$ , and a set of corrupted participants  $\mathcal{C} \subset \mathcal{P}$ , output, for each  $P$  in  $\mathcal{C}$ , the distance and direction (edge to children or parent) from  $M(P)$  to  $M(P_i)$  in  $\text{ST}$ , its depth (distance to  $\text{ST}$ 's root node), and the height of each of its children nodes (distance to further leaf on subtree).

**Theorem 1.** Let  $\mathcal{H} = \langle \text{Gen}, \text{Enc}, \text{Dec}, \text{Add} \rangle$  be semantically secure homomorphic public-key encryption scheme. Then, protocol  $\Pi_{\text{ReqResp}}$  privately realizes functionality  $\mathcal{F}_{\text{ReqResp}}$  in the  $\mathcal{F}_{\text{network}}$  hybrid model under Leakage 1.

In the following proof we analyze the case in which the server is not corrupted and there is at least one other honest node (otherwise, the location of the server node is leaked anyway). In the case where the server is corrupted, there is no secret beyond the network graph structure, and hence simulation becomes simpler (and all the interesting simulation steps are included in the proof).

*Proof.* Let  $\mathcal{C}$  be the set of corrupted participants, and  $\mathcal{H}$  be the public key encryption scheme as in protocol  $\Pi_{\text{ReqResp}}$ . We next specify the ideal adversary behavior on each of the protocol phases.

**Simulating Setup** In the setup phase, the corrupted participants only receive their key-pairs, the server's public key  $\text{pk}^{\mathcal{S}}$ .

1. Instantiate network functionality  $\mathcal{F}_{\text{network}}$  using graph  $G$  for the participant set.
2. Generate additional server public key  $\text{pk}^{\mathcal{S}}$ .
3. For each corrupted party, assign its spanning tree edges (to children and parent) and  $\text{pk}^{\mathcal{S}}$ .

**Simulating Req** Let  $P_i \in \mathcal{C}$  be the requester, and  $P_{j^*} = \mathcal{S} \notin \mathcal{C}$  be the server participant. Simulation proceeds as follows:

1. Sample session id  $\text{sid}$ , key-pair  $(\text{sk}_i, \text{pk}_i)$ .
2. If  $P_i$  is corrupted, then upon receiving input  $\text{req}$  from  $P_i$ , simulate the real adversary execution to get the (possibly) updated request  $\text{req}'$ . Send  $\text{req}'$  to  $P_i$  as its input and get its output  $\text{res}$ . Otherwise, set  $\text{res}$  to  $0^\ell$ .
3. Using distance and direction from corrupted participants to  $P_i$ , simulate a  $P_i$  started multicast protocol on spanning tree with message



- $\langle \text{request\_to\_server}, \text{sid}, \text{Enc}_{\text{pk}^s}(0^\ell), \text{pk}_{\text{sid}} \rangle$  where  $\text{sid}$  and  $\text{pk}_{\text{sid}}$  are fresh values. (That is, the corrupted participants get  $\langle \text{request\_to\_server}, \text{sid}, \text{Enc}_{\text{pk}^s}(0^\ell), \text{pk}_i \rangle$  at the right moment and through the expected graph edge.)
4. Simulate `to_server_UP` messages by assigning a random share  $\text{share}_j$  to each corrupted participant, and assigning an arbitrary share to the honest children of each corrupted participant. Then, the simulation is done by giving message  $\langle \text{sid}, \text{to\_server\_UP}, S \rangle$  at the right moment from the correct children, where  $S = \text{Enc}_{\text{pk}^s}(\text{share})$ .
  5. Use corrupted participant depth to simulate the `to_server_DOWN` message by giving message  $\langle \text{sid}, \text{to\_server\_DOWN}, S \rangle$  to each corrupted participant at the right moment. If the root of the tree is corrupted, then  $S$  must match the homomorphic computed value of the sum of the nodes shares. Otherwise,  $S$  can contain a dummy value.
  6. Simulate each participant sending the `to_requester_UP` message were shares are identical as in step 4, except the honest participants, whose share are updated so that reconstruction produces  $\text{res}$ .
  7. Simulate `to_requester_DOWN` by sending  $\langle \text{sid}, \text{to\_requester\_DOWN}, C \rangle$  to corrupted participants in the path root to requester, where  $C = \text{Enc}_{\text{pk}_{\text{sid}}}(\text{res})$ .

The simulation above is perfect in terms of communication patterns (timing, length and type of messages). This is because the simulator uses the leakage profile to deliver the message to the corrupted participants at the right time and through the exact graph edge. Hence, the security of the protocol relies on the ability to simulate the content of the messages seen by the corrupted nodes. We next analyze the content by message type:

- Request multicast. If the request is known to the simulator, it can produce a ciphertext identically distributed to the real message. Otherwise, the simulator produces a dummy ciphertext (computationally indistinguishable to the real message by the security of the encryption scheme).
- `to_server_UP` messages. There is no secret information to simulate. Hence, the simulator produces ciphertexts identically distributed to the real protocol messages.
- `to_server_DOWN` message. Same as above.
- `to_requester_UP` and `to_requester_DOWN` messages. Here the shares corresponding to honest participant are updated so that the reconstruction produces  $\text{res}$ . In the worst case that the adversary controls  $P_i$ , then

it can decrypt the shares. However, these cannot be correlated with the shares sent to the server, since these are encrypted under the key of the server. In addition, shares are uniformly distributed,  $(n - 1)$ -wise independent, and they reconstruct to the same valid output `res`. Hence, the simulated shares in plaintext cannot be distinguished from the ones used in the real execution.

A simple hybrid-argument<sup>5</sup> over the security of the encryption scheme proves that the real and simulated views are computationally indistinguishable. □

## 6 Conclusion

We have introduced a new protocol that enables to hide a server in a network. This protocol has several advantages over previous proposals: it is efficient, asynchronous and collusion-resistant. To the best of our knowledge this is the first solution with these characteristics.

We believe that this work is an important step towards designing practical and provably secure systems that enable to hide relevant meta-data (such as the identity or location of participants) in a controllable way. Future work directions include improving the robustness of the protocol in order to handle adaptive and active adversaries. We also believe that the current construction can be proven secure in the UC framework[4].

## References

- [1] Adi Akavia, Rio LaVigne, and Tal Moran. Topology-hiding computation on all graphs. Cryptology ePrint Archive, Report 2017/296, 2017. <http://eprint.iacr.org/2017/296>.
- [2] Adi Akavia and Tal Moran. Topology-hiding computation beyond logarithmic diameter. In *Advances in Cryptology - EUROCRYPT 2017* -

---

<sup>5</sup>changing at each hybrid step the honest participant updated shares in the `to_requester.UP` messages from the ideal distribution to the corresponding ciphertext on the real distribution. Note that the fact we are in the multi-user setting (a message is encrypted under two different public keys) can be reduced to the single-user setting (standard IND-CPA security definition)[3].

*36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 - May 4, 2017, Proceedings, Part III*, pages 609–637, 2017.

- [3] Mihir Bellare, Alexandra Boldyreva, and Silvio Micali. Public-key encryption in a multi-user setting: Security proofs and improvements. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 259–274. Springer, 2000.
- [4] R. Canetti. Universally Composable Security: A New Paradigm for Cryptographic Protocols. page 136, oct 2001.
- [5] David Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *Journal of cryptology*, 1(1):65–75, 1988.
- [6] David L Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–90, 1981.
- [7] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. Technical report, Naval Research Lab Washington DC, 2004.
- [8] Shlomi Dolev and Rafail Ostrovsky. Xor-trees for efficient anonymous multicast and reception. *ACM Trans. Inf. Syst. Secur.*, 3(2):63–84, 2000.
- [9] David Goldschlag, Michael Reed, and Paul Syverson. Onion routing. *Communications of the ACM*, 42(2):39–41, 1999.
- [10] Martin Hirt, Ueli Maurer, Daniel Tschudi, and Vassilis Zikas. Network-hiding communication and applications to multi-party protocols. Cryptology ePrint Archive, Report 2016/556, 2016. <http://eprint.iacr.org/2016/556>.
- [11] Sachin Katti, Dina Katabi, and Katarzyna Puchala. Slicing the onion: Anonymous routing without pki. 2005.
- [12] Brian N Levine, Michael K Reiter, Chenxi Wang, and Matthew Wright. Timing attacks in low-latency mix systems. In *International Conference on Financial Cryptography*, pages 251–265. Springer, 2004.
- [13] Brian Neil Levine and Clay Shields. Hordes: a multicast based protocol for anonymity1. *Journal of Computer Security*, 10(3):213–240, 2002.
- [14] Tal Moran, Ilan Orlov, and Silas Richelson. *Topology-Hiding Computation*, pages 159–181. Springer Berlin Heidelberg, Berlin, Heidelberg, 2015.

- [15] Steven J Murdoch and George Danezis. Low-cost traffic analysis of tor. In *Security and Privacy, 2005 IEEE Symposium on*, pages 183–195. IEEE, 2005.
- [16] Lasse Overlier and Paul Syverson. Locating hidden servers. In *Security and Privacy, 2006 IEEE Symposium on*, pages 15–pp. IEEE, 2006.
- [17] Andreas Pfitzmann, Birgit Pfitzmann, and Michael Waidner. Isdn-mixes: Untraceable communication with very small bandwidth overhead. In *Kommunikation in verteilten Systemen*, pages 451–463. Springer, 1991.
- [18] Charles Rackoff and Daniel R Simon. Cryptographic defense against traffic analysis. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 672–681. ACM, 1993.
- [19] Michael K Reiter and Aviel D Rubin. Crowds: Anonymity for web transactions. *ACM transactions on information and system security (TISSEC)*, 1(1):66–92, 1998.
- [20] Andrei Serjantov and Peter Sewell. Passive attack analysis for connection-based anonymity systems. *Lecture notes in computer science*, 2808:116–131, 2003.
- [21] Vitaly Shmatikov. Probabilistic analysis of anonymity. In *Computer Security Foundations Workshop, 2002. Proceedings. 15th IEEE*, pages 119–128. IEEE, 2002.
- [22] Michael Waidner. Unconditional sender and recipient untraceability in spite of active attacks. In *Workshop on the Theory and Application of Cryptographic Techniques*, pages 302–319. Springer, 1989.
- [23] Matthew K Wright, Micah Adler, Brian Neil Levine, and Clay Shields. An analysis of the degradation of anonymous protocols. In *NDSS*, volume 2, pages 39–50, 2002.
- [24] Ye Zhu, Xinwen Fu, Bryan Graham, Riccardo Bettati, and Wei Zhao. On flow correlation attacks and countermeasures in mix networks. In *International Workshop on Privacy Enhancing Technologies*, pages 207–225. Springer, 2004.