



HAL
open science

Spores: Stateless Predictive Onion Routing for E-Squads

Daniel Bosk, Yérom-David Bromberg, Sonja Buchegger, Adrien Luxey,
François Taïani

► **To cite this version:**

Daniel Bosk, Yérom-David Bromberg, Sonja Buchegger, Adrien Luxey, François Taïani. Spores: Stateless Predictive Onion Routing for E-Squads. 2019. hal-02885795

HAL Id: hal-02885795

<https://inria.hal.science/hal-02885795>

Preprint submitted on 1 Jul 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - ShareAlike 4.0 International License

SPORES: Stateless Predictive Onion Routing for E-Squads

Daniel Bosk
dbosk@kth.se
KTH Royal Institute of Technology
Stockholm, Sweden

Yérom-David Bromberg
david.bromberg@irisa.fr
Univ Rennes, CNRS, Inria, IRISA
Rennes, France

Sonja Buchegger
buc@kth.se
KTH Royal Institute of Technology
Stockholm, Sweden

Adrien Luxey
adrien.luxey@irisa.fr
Univ Rennes, CNRS, Inria, IRISA
Rennes, France

François Taïani
francois.taiani@irisa.fr
Univ Rennes, CNRS, Inria, IRISA
Rennes, France

Abstract

Mass surveillance of the population by state agencies and corporate parties is now a well-known fact. Journalists and whistle-blowers still lack means to circumvent global spying for the sake of their investigations. With SPORES, we propose a way for journalists and their sources to plan *a posteriori* file exchanges when they physically meet. We leverage on the multiplication of personal devices per capita to provide a lightweight, robust and fully anonymous decentralised file transfer protocol between users. SPORES hinges on our novel concept of *e-squads*: one’s personal devices, rendered intelligent by gossip communication protocols, can provide private and dependable services to their user. People’s *e-squads* are federated into a novel onion routing network, able to withstand the inherent unreliability of personal appliances while providing reliable routing. SPORES’ performances are competitive, and its privacy properties of the communication outperform state of the art onion routing strategies.

ACM Reference format:

Daniel Bosk, Yérom-David Bromberg, Sonja Buchegger, Adrien Luxey, and François Taïani. 2020. SPORES: Stateless Predictive Onion Routing for E-Squads. In *Proceedings of ACM Conference, Washington, DC, USA, July 2017 (Conference’17)*, 15 pages. DOI: 10.1145/nnnnnnn.nnnnnnn

1 Introduction

Recent years have been marked by multiple high-profile mass surveillance scandals, involving a diverse range of players, from state agencies [31, 34], to large technology firms [6], through start-ups with close links to academia [6, 27, 32, 54].

In this context, journalists and whistle-blowers must be extremely careful when sourcing or exchanging sensitive or damaging information, but they unfortunately still lack the technical means to fully circumvent corporate and governmental surveillance efforts. Although data encryption is often used as a first line of defense to protect confidential information, it is unfortunately insufficient on its own to fully protect the parties involved in a remote exchange of data [17, 20]. Without additional counter-measures, metadata such

as a user’ location and activity can usually still be tracked, thus revealing what each user shared with whom and when. The exposure of metadata seriously weakens user privacy; as famously emphasized by a former NSA and CIA director: “*We kill people based on metadata.*” [37].

Anonymity networks have been proposed to lower the risk of online spying. They notably hide the identity of the client, most often using mix networks [8] or onion routing [13, 16]. These approaches basically mangle user requests through a series of relay servers, effectively hiding the link between sender and receiver. To bootstrap such anonymity networks, relays need to be sufficiently numerous, available, distributed over many autonomous systems, and to provide enough bandwidth. For instance, Tor [13], the most popular onion routing implementation, relies on relays maintained by volunteer operators. On average, 6000 connected relays handle the traffic of 2 million online users¹. This small amount of relays relative to the userbase makes Tor particularly prone to attacks such as traffic analysis [47]. Tor’s infrastructure is also in part centralised, as 10 Directory Authorities (DAs) compute an hourly consensus listing online relays and providing means to contact them. The DAs constitute a single point of failure, and could be subverted, just like any other relay. We argue that embracing a peer-to-peer (P2P) strategy—where any participating user also acts as a relay—would notably improve onion routing security by drowning malevolent entities in an ocean of honest peers.

We are not the first to advocate more decentralisation, as many attempts at building P2P anonymous data-sharing networks [10, 15, 18, 35, 43] have been proposed in the last decades. Typically, such solutions leverage onion routing with added components to perform e.g. node discovery in a decentralised fashion. Alas, frequent dis/connections of participating peers (*churn*) hinder these systems’ performances [28], making them unusable in practice.

In this paper, we propose SPORES, a fully decentralised anonymous file exchange protocol, adapted from traditional onion routing. To sustain the unavoidable churn, we revisit

¹See Tor Metrics at <https://metrics.torproject.org/>.

P2P by leveraging machine learning in order to predict peers’ availability. Towards this goal, we make two assumptions on the peers: we consider that each participating device belongs to a particular user, and that each user owns several devices (encouraged by the multiplication of appliances per household [1]). We thus propose the concept of *e-squads*: an e-squad is constituted of a single user’s devices, that model their user’s behaviour by exchanging information through gossip messaging. Using their user behavioural model, each device can predict its future state of availability.

We use these estimates to propose a novel onion routing mechanism, called *Probabilistic Onion Routing* (POR). With POR, onion routes may include several candidate relays at each hop, such that a message can go through the route as long as one candidate is online per hop. Users employ the availability prediction to ensure that the routes they create will remain available with a good probability, without sacrificing their privacy. In addition, POR is stateless—all routing information is contained in the headers—enabling short-lived relay servers to pass on messages as soon as they join the network, eschewing any bootstrap phase.

Building on POR, SPORES enables two users to exchange a file in complete anonymity: firstly, the two users exchange file metadata and routing information *out-of-band* (using another communication channel than SPORES); then, using the routes they agreed upon, their respective e-squads collaborate to perform the file exchange through SPORES, without revealing their identities to the rest of the network. The proposed service is quite similar to OnionShare², without the security limitations of Tor, and without the need to spawn a web service prior to the exchange. We say that SPORES is an *a posteriori* file exchange service.

Our contributions are the following:

- We introduce the concept of *e-squads*, and build a intra-e-squad protocol, that allows devices owned by the same user to create user behavioural models and thus, to estimate their future availability.
- Based on the e-squad predictions, we introduce Probabilistic Onion Routing (POR), a onion routing protocol tailored for networks with high degrees of churn.
- We use POR to realise SPORES, an anonymous file transfer service. After an initial out-of-band exchange of metadata, two users can privately exchange a file. The transfer remains efficient despite the network’s unreliability, and ensures a better anonymity than existing onion routing approaches.

The remainder of this chapter is organised as follows: we first present our protocol and its sub-systems in section 2, before presenting our attack model and security properties in section 3. An evaluation of SPORES is proposed in section 4.

²An anonymous file exchange service backed by Tor’s hidden services, see <https://onionshare.org/>.

We make a review of the state of the art in section 5, and finally conclude the paper in section 6.

2 Our approach

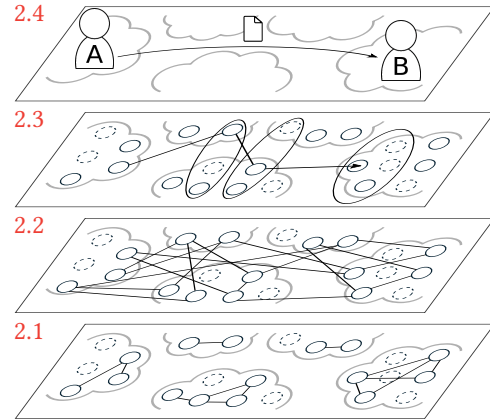


Figure 1. The subsystems constituting SPORES. From bottom to top: the e-squad overlays (Sec. 2.1) and the global overlay (2.2) enable Probabilistic Onion Routing (2.3), which lies at the core of the SPORES anonymous file exchange service (2.4).

SPORES is constituted of several sub-systems, as depicted in Fig. 1. At the root of our proposal is predictive routing, which is made possible by having every peer publish predictions about their future connectivity. Each peer in the network is a device owned by an individual, who possesses several appliances. Their appliances exchange information about their user through a *private e-squad overlay* (Sec. 2.1); there is one private overlay per participating user. Devices use this information to build a model of their user, and make predictions about their future availability. This estimate is regularly published by each device, along with their network address and public key, on the *global overlay* (Sec. 2.2). This overlay enables peer discovery at the scale of the whole network. Using the above information, any device can intelligently build Probabilistic Onion Routes (PORs, see Sec. 2.3), a new kind of onion route featuring several relays per hop, in order to maximise the route’s availability despite the relays’ churn. Finally, Sec. 2.4 presents the anonymous file exchange protocol in itself, SPORES, that anonymises a file transfer through PORs, while making use of each user’s e-squad for increased dependability.

2.1 The private e-squad overlay

An e-squad overlay is constituted only of devices owned by the same user. Its role is to make any user-related information available to the whole set of devices.

2.1.1 Sharing the user’s behaviour

The e-squad overlay is based on the SPRINKLER Gossiper algorithm [30], extended with acknowledgements [5] to

better resist churn. As in SPRINKLER, we assume each user owns a set of devices \mathcal{D} . The user’s activity is an ever-growing sequence $S = \{r_1, \dots, r_i, \dots\}$ of interactions r_i . Each device d initially only knows about interactions that took place on it, S_d . Through the e-squad overlay, all nodes of the e-squad share their local interactions to obtain $S = \bigcup_{d \in \mathcal{D}} S_d$. Interactions are timestamped, and totally ordered.

Contrarily to SPRINKLER, an interaction can be either a device usage event or a file exchange event. A file exchange is tied to a single device (sender or receiver), and all the e-squad needs to know what device is involved in which file exchange. To this end, an interaction r is constituted of the following fields:

$$r = (ts, d, \text{typ}, f) \in \mathbb{R} \times \mathcal{D} \times \mathcal{T} \times \mathcal{F}$$

such that: $ts \in \mathbb{R}$ is the interaction timestamp, $d \in \mathcal{D}$ is the descriptor (see Sec. 2.2) for the device where the interaction r took place, $\mathcal{T} = \{\text{USE}, \text{DL}, \text{UL}\}$ is the set of interaction *types* (resp. device usage, new file download, or new file download). When $\text{typ} = \text{DL}$ (resp. UL), $f \in \mathcal{F}$ contains the unique ID of the file that just started downloading (resp. uploading) on d . When $\text{typ} = \text{USE}$, it means that device d was connected at time ts . Devices issue a USE message when they are grabbed, and every T seconds while they remain connected.

2.1.2 Modelling the user’s behaviour

Given the sequence of devices’ usage $S^U = \{r \in S, r.\text{typ} = \text{USE}\}$, each device needs to compute its own probability $P_i(d)$ of staying online in the near future, before advertising it.

First of all, using only S^U , each device builds an availability sequence $X = X_1, \dots, X_i, \dots$, where X_i contains the set of online devices during the interval $[t_i, t_{i+1}[$ (see Eq. 1). The observation sequence has a period of T : $\forall i, t_{i+1} = t_i + T$. The sequence X can be represented as a 2D sparse matrix of booleans.

$$X_i(d) = 1 \iff \exists r \in S^U, r.d = d \wedge t_i \leq r.ts < t_{i+1}. \quad (1)$$

Now, to predict $P_i(d)$, we consider that the stochastic process X follows the Markov property: ‘the future only depends on the present, not on the past’. We use the hypothesis in Eq. 2.

As a result, the probability for d to be online in the near future only depends on its probability to stay online after the current round $X_i = x$. To estimate this probability, we simply count³ the number of times the current situation x

led to a situation where d was also online (Eq. 3):

$$\begin{aligned} P_i(d) &= P[X_{i+1}(d) = 1 \mid X_i = x, \dots, X_0 = x_0] \\ &= P[X_{i+1}(d) = 1 \mid X_i = x] \\ &= \frac{\left| \{X_j \in X, X_j = x \wedge X_{j+1}(d) = 1\}_{0 \leq j < i} \right|}{\left| \{X_j \in X, X_j = x\}_{0 \leq j < i} \right|} \end{aligned} \quad (2)$$

$$= \frac{\left| \{X_j \in X, X_j = x \wedge X_{j+1}(d) = 1\}_{0 \leq j < i} \right|}{\left| \{X_j \in X, X_j = x\}_{0 \leq j < i} \right|} \quad (3)$$

Given the high dimensionality of the state space, it might happen that x was never seen before, leading to an undefined $P_i(d)$. In such a case, we estimate the probability that d stays online two turns in a row as fallback.

2.2 The global overlay

To create PORS, each device needs to know some other online devices’ descriptors. For a device d , a descriptor contains its address $@_d$, its public key pk_d , and its estimated probability of remaining online $P_i(d)$. d also knows its own private key sk_d , that it uses to decipher messages encrypted with pk_d .

Given the decentralised nature of SPORES, we cannot rely on a central registry of online peers as e.g. Tor does. We use instead a global Random Peer Sampling (RPS) service [23, 51]. Essentially, each node maintains a view \mathcal{V}_{RPS} containing $l_{\mathcal{V}}$ other devices’ descriptors. Every T_{RPS} seconds, the view is updated as follows: a device d pops the oldest descriptor d' from its view, then swaps a predefined number of l_{gossip} elements from \mathcal{V}_{RPS} with d' . Both devices add a fresh descriptor of themselves to the view exchange. If d' was offline, its descriptor is simply removed from d ’s view, with no further modification to \mathcal{V}_{RPS} .

This allows for two things: firstly, each device’s view contains a constantly changing random sample of participating devices; secondly, stale descriptors get removed from one’s view after a bounded time, such that \mathcal{V}_{RPS} mostly contains online devices’ descriptors.

Given their epidemic nature, RPS services are very sensitive to Byzantine attacks, where malicious nodes gossip bad views in order to disrupt the randomness of the neighbourhood graph. Several proposals overcome this limitation, sometimes by relying on a trusted third-party [2], sometimes by computing a reputation of the peers [4, 24]. We leverage on the latter, so as to remain entirely decentralised.

2.3 PORS: Probabilistic Onion Routes

Legacy Tor primer Onion routing makes connections between a client (say Alice) and their correspondent (Bob) go through two or more servers (or relays) before reaching their destination. With Tor, to create a route, Alice randomly picks three relays to constitute the path, and incrementally establishes TLS connections to each of them through the

³Because we work with low-probability events observed with small amounts of data, there is a possibility that an event never occurs in X . To counter that, we apply *add-one smoothing* [46] while computing probabilities. We left this engineering optimization out of the demonstration for clarity.

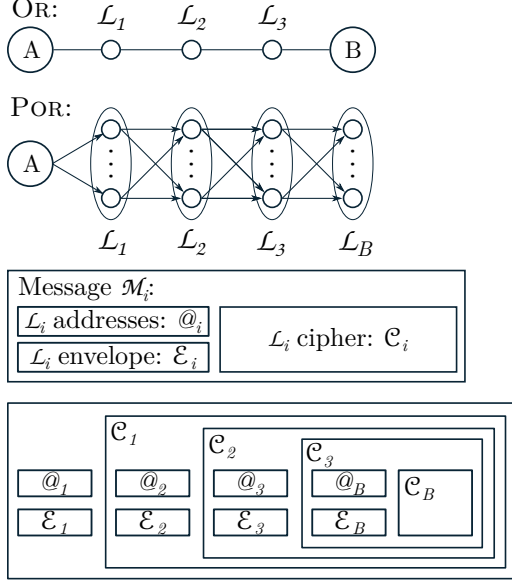


Figure 2. In Onion Routing (OR), each layer \mathcal{L} is constituted of only one node. In Probabilistic Onion Routing (POR), there are several *candidate* nodes that each message can go through at each layer. The bottom part of the plot shows the format of a POR message, and the message effectively sent by Alice to relays in \mathcal{L}_1 .

route. Once the route is established, it constitutes a persistent two-way TCP stream, although the traffic is internally chunked into fixed-size messages (or cells). Cells contain a header and a payload, that are encrypted altogether by the client several times: once per relay. Upon reception of a cell from the sender to its destination, each relay deciphers it using the encryption keys negotiated during the TLS connection bootstrap. Bob finally receives the message originally written by Alice, and can answer back on the same pipe. Messages on this direction are incrementally encrypted by the relays, such that Alice receives Bob’s message hidden under three layers of encryption. She decrypts it using the keys that were negotiated with the relays during the connection establishment.

The anonymizing property stems from the fact that each hop \mathcal{L}_i only knows the address of the previous relay \mathcal{L}_{i-1} (that sent the message) and the address of the next \mathcal{L}_{i+1} (determined at the connection’s establishment). Given that routes contain two or more hops, no intermediary knows both the sender and the receiver of a message, thus making the communication anonymous.

Several relays per layer The basic idea of Probabilistic Onion Routing (POR) is depicted in Fig. 2: each message sent from Alice to Bob may pass through multiple candidate nodes at each hop, instead of only one in traditional onion routing.

In onion routing, when any of the relays becomes unavailable, the route is broken and a new one needs to be

Algorithm 1 The Broadcast Encrypt/Decrypt algorithms

```

1: function BE( $\mathcal{P}, pk_{\mathcal{L}}$ )
2:    $k \leftarrow$  random symmetric key
3:    $C \leftarrow SE(\mathcal{P}, k)$ 
4:    $\mathcal{E} \leftarrow \{AE(k, pk)\}_{pk \in pk_{\mathcal{L}}}$ 
5:   return  $\mathcal{E}, C$ 

1: function BD( $\mathcal{E}, C, sk$ )
2:   for  $e \in \mathcal{E}$  do
3:      $k \leftarrow AD(e, sk)$ 
4:     if  $k \neq \perp$  then
5:       return  $SD(C, k)$ 
6:   return  $\perp$ 

```

created. The prime interest of PORs is that they are resilient to intermediaries churn: we only need one online relay per layer for the route to function. In practice, when a node from layer \mathcal{L}_i has a message to transmit, it tries sending it to each device in layer \mathcal{L}_{i+1} in random order, until it succeeds or all attempts fail. In the latter case, the message is dropped.

In contrast to Tor, PORs do not create TLS connections, which would be inapplicable with several nodes per hop. Instead, routes are stateless: all routing information is contained inside an encrypted message.

POR messages Fig. 2 shows the format of POR headers (along with an example of a full message as sent by Alice in Fig. 2). A POR message \mathcal{M}_i , as received by any member of \mathcal{L}_i , is constituted of three parts, $\mathcal{M}_i = (@_i, \mathcal{E}_i, C_i)$:

- $@_i$: The addresses of all members of the current layer \mathcal{L}_i , used by nodes of \mathcal{L}_{i-1} to forward \mathcal{M}_i .
- \mathcal{E}_i : An envelope, destined to \mathcal{L}_i , that will allow them to decipher the cipher C_i .
- C_i : A cipher, that can be deciphered by any member of \mathcal{L}_i using \mathcal{E}_i . It can unravel into another POR message \mathcal{M}_{i+1} for the next layer \mathcal{L}_{i+1} , or into an application payload once the message reached its destination.

By getting rid of TLS connections in favour of header-based routes, POR enables *stateless* routing: no prior communication is needed with relays to establish onion routes, they simply decipher any received message, and read their header to forward them to the next layer. This is particularly interesting for short-lived nodes such as seldom connected personal devices as we target: they can participate in the system as soon as they join, without any bootstrap phase. Their disconnection does not mandate a new route construction.

On the other hand, PORs are connectionless one-way channels (UDP-like), and the message is not fixed in size due to the lack of re-encryption between each hop. In particular, POR does not guarantee messages integrity nor order (as each cell potentially travels through a different path). It is the role of the upper abstraction layer (e.g. our file exchange protocol SPORES) to guarantee reliable & ordered transmission.

Algorithm 2 The Message Encrypt/Decrypt algorithms

```
1: function  $ME(\mathcal{P}, \mathbf{L})$ 
2:    $M \leftarrow \mathcal{P}$ 
3:   for  $\mathcal{L} \in \text{reverse}(\mathbf{L})$  do
4:      $M.\mathcal{E}, M.C \leftarrow BE(M, \mathcal{L}.pk)$ 
5:      $M.@ \leftarrow \mathcal{L}.@$ 
6:   return  $M$ 
1: function  $MD(M, sk)$ 
2:   return  $BD(M.\mathcal{E}, M.C, sk)$ 
```

Cryptographic primitives The encrypted message C_i containing the addresses of the next layer needs to be decipherable by any of the current layer \mathcal{L}_i 's members, and only by them. This cryptographic scheme is coined Broadcast Encryption (BE) [14]. We derive our encryption process from Hybrid Encryption [48] (as used in PGP), where a message M is encrypted into a cipher C using a unique symmetric key k (e.g. using AES). Each member of the group \mathcal{L}_i must be given this key, which is the purpose of the envelope \mathcal{E} . It contains the concatenation of k encrypted with each member's public key (using e.g. RSA). Upon reception of a ciphered message (\mathcal{E}, C) , a peer attempts to decrypt each portion of the envelope with its private key, until it succeeds (and gets k to decrypt C) or fails.

We write down our broadcast encryption/decryption algorithms in algorithm 1, and its application to our message cryptography in algorithm 2.

Let $C \leftarrow SE(\mathcal{P}, k)$ and $\mathcal{P} \leftarrow SD(C, k)$ be symmetric primitives for encrypting/decrypting an arbitrary payload \mathcal{P} with key k , such that SD returns \perp on decryption failure.

Let $C \leftarrow AE(\mathcal{P}, pk)$ and $\mathcal{P} \leftarrow AD(C, sk)$ be asymmetric ones for encrypting/decrypting a payload \mathcal{P} with the public key pk (resp. secret key sk), such that AD returns \perp on failure.

$(\mathcal{E}, C) \leftarrow BE(\mathcal{P}, pk_{\mathcal{L}})$ – Given a layer \mathcal{L} 's public keys $pk_{\mathcal{L}}$, and an arbitrary payload \mathcal{P} to encrypt, BE broadcast encrypts \mathcal{P} by outputting an envelope \mathcal{E} containing a symmetric key k encrypted with each $pk \in pk_{\mathcal{L}}$, and a ciphertext C containing the payload encrypted with k .

$\mathcal{P} \leftarrow BD(\mathcal{E}, C, sk)$ – Given an envelope \mathcal{E} , a ciphertext C and a secret key sk , BD broadcast decrypts the payload \mathcal{P} into the expected plaintext, or \perp if the decryption fails.

In algorithm 2, we write $\mathcal{L}.pk$ and $\mathcal{L}.@$ to refer to a layer's nodes' public keys and addresses. $\mathbf{L} = \{\mathcal{L}_i\}_i$ is an array of layers. $\text{reverse}(\mathbf{L})$ means we iterate on \mathbf{L} in reverse order (starting from the last element).

$M \leftarrow ME(\mathcal{P}, \mathbf{L})$ – Given a payload \mathcal{P} (that can be a POR message), and an array of layers \mathbf{L} , ME recursively encrypts the output message M for each $\mathcal{L} \in \mathbf{L}$ starting from the last.

$\mathcal{P} \leftarrow MD(M, sk)$ – MD attempts to decrypt the message M using the secret key sk . It either returns a payload \mathcal{P} or \perp on failure.

Forwarding messages We finally display the message reception and forwarding procedure in algorithm 3, that

Algorithm 3 Receiving a POR message on d

```
1: on receive  $M$ 
2:    $\mathcal{P} \leftarrow MD(M, sk_d)$ 
3:   if  $\mathcal{P} = \perp$  then
4:     return  $\perp$  ▷ Decryption failed
5:   else if  $\mathcal{P} \neq \text{POR message}$  then
6:     process  $\mathcal{P}$  ▷ I am the recipient
7:   else
8:     Forward( $\mathcal{P}$ )
1: function Forward( $M$ )
2:   for  $@ \in \text{random}(M.@)$  do
3:     if send( $@, M, T_{\text{out}}$ ) then
4:       return  $\top$  ▷ Success
5:   return  $\perp$  ▷ All layer offline: drop message
```

runs on any node participating in SPORES (we consider a device d). Upon reception of a message M , d first attempts to decipher it using its secret key sk_d . Three cases are possible: either the decryption fails, which constitutes an error—the message is dropped; either the output is not a POR message, in which case the message is destined to d ; either d decryption unraveled another message, in which case d forwards it to the next layer using Forward function. This procedure iterates over each address $@$ in $M.@$ in random order, and attempts to send the message to $@$. The send function called at line 3 takes three parameters: the recipient's address, the message to send, and a timeout duration. T_{out} is a configuration parameter, usually below a second. If the send call succeeds, the message is duly forwarded. If the 'for' loop returns without any successful attempt, all the next layer is considered offline, and the message is dropped. Note that it takes $\#\mathcal{L} \times T_{\text{out}}$ seconds to drop a message when the next layer is offline.

2.4 SPORES: File exchanges through POR

We now have all the building blocks to perform anonymous file transfers using e-squads. In this section, we first present how two users agree upon probabilistic onion routes for their exchange, including the intelligent selection to maximise the routes' availability; finally, we discuss the file exchange protocol built atop PORs.

Routes creation Fig. 3 depicts the creation process of a route between our beloved Alice (uploader) and Bob (the receiver). As already mentioned, this process takes place out-of-band (on another communication channel such as Near Field Communication (NFC), LAN, Bluetooth, carrier pigeon, or else). The initialisation serves two purposes: to provide Bob with the exchanged file metadata (we come back to it in the next paragraphs), and to decide upon the PORs that will be used throughout the transfer.

Since PORs are one-way only, Alice and Bob need to agree upon two routes: a *forward* one, from Alice to Bob, that will carry file chunks, and a *backward* route, from Bob to Alice,

Algorithm 4 Route initialisation between device d_A uploading file f (on the left) and d_B downloading it (on the right).

<pre> 1: function INITUPLOAD(f, θ) Step ❶: 2: $fd \leftarrow \text{BuildFileDescriptor}(f)$ 3: $sq \leftarrow \{r.d, r.d \neq d_A\}_{r \in S}$ 4: $\mathcal{BL}_4 \leftarrow \text{PickLayer}(sq, \theta) \cup \{d_A\}$ 5: $\mathcal{BL}_3 \leftarrow \text{PickLayer}(\mathcal{V}_{\text{RPS}}, \theta)$ 6: $\mathcal{BR}_1 \leftarrow [\mathcal{BL}_3, \mathcal{BL}_4]$ 7: send $\langle fd, \mathcal{BR}_1 \rangle$ to d_B </pre>	$\rightarrow \rightarrow \rightarrow \rightarrow$	<pre> 7: on receive $\langle fd, \mathcal{BR}_1 \rangle$ Step ❷: 8: $sq \leftarrow \{r.d, r.d \neq d_B\}_{r \in S}$ 9: $\mathcal{FL}_4 \leftarrow \text{PickLayer}(sq, \theta) \cup \{d_B\}$ 10: $\mathcal{FL}_3 \leftarrow \text{PickLayer}(\mathcal{V}_{\text{RPS}}, \theta)$ 11: $\mathcal{FR}_1 \leftarrow [\mathcal{FL}_3, \mathcal{FL}_4]$ 12: send $\langle \mathcal{FR}_1 \rangle$ to d_A </pre>
<pre> 12: on receive $\langle \mathcal{FR}_1 \rangle$ Step ❸: 13: $\mathcal{FL}_1 \leftarrow \text{PickLayer}(\mathcal{V}_{\text{RPS}}, \theta)$ 14: $\mathcal{FL}_2 \leftarrow \text{PickLayer}(\mathcal{V}_{\text{RPS}}, \theta)$ 15: $\mathcal{FR} \leftarrow [\mathcal{FL}_1, \mathcal{FL}_2] \oplus \mathcal{FR}_1$ 16: $r \leftarrow (\text{Now}(), d_A, \text{UL}, fd.\text{ID})$ 17: $S \leftarrow S \cup \{r\}$ ▶ Shared to e-squad 18: Start sending f </pre>	$\leftarrow \leftarrow \leftarrow \leftarrow$	<pre> 13: $\mathcal{BL}_1 \leftarrow \text{PickLayer}(\mathcal{V}_{\text{RPS}}, \theta)$ 14: $\mathcal{BL}_2 \leftarrow \text{PickLayer}(\mathcal{V}_{\text{RPS}}, \theta)$ 15: $\mathcal{BR} \leftarrow [\mathcal{BL}_1, \mathcal{BL}_2] \oplus \mathcal{BR}_1$ 16: $r \leftarrow (\text{Now}(), d_B, \text{DL}, fd.\text{ID})$ 17: $S \leftarrow S \cup \{r\}$ ▶ Shared to e-squad 18: Start receiving f </pre>




Figure 3. For Alice to send a file to Bob through SPoRES, they need to agree upon two routes: a forward route to send file chunks, and a backward one to send acknowledgements. Each user contributes layers to each route, so as to maximise the diversity of the involved relays.

that will transport Bob’s acknowledgements of the chunks. Furthermore, to maximize the peer diversity in the route (for security reasons), both parties compute a portion of each route. In the rest of the paper, we settle with a number of layers of $\#_{\mathcal{L}} = 4$, which is required to have decent security properties, while adding more layers would not make routes significantly more secure (as agreed upon by most onion routing approaches).

We now detail the operations depicted in Fig. 3, and detailed in algorithm 4:

- At ❶, Alice’s sending device d_A crafts the inner part of the header for the route to herself, \mathcal{BR}_1 , at lines 3-6. The final layer \mathcal{BL}_4 is only constituted of Alice’s devices: d_A picks candidates from her e-squad sequence S (line 3), and lastly adds its own descriptor to \mathcal{BL}_4 (l. 4). The third layer \mathcal{BL}_3 is selected using d_A ’s RPS view \mathcal{V}_{RPS} , which contains a pool of global descriptors. On line 7, d_A sends \mathcal{BR}_1 to Bob, along with the file metadata fd .
- At ❷, Bob’s receiving device d_B builds its half of the forward route (\mathcal{FR}_1) just like d_A did at ❶, see lines 8-11. Again, \mathcal{FL}_4 is only made of Bob’s e-squad, while \mathcal{FL}_3 samples devices from the global overlay. d_B sends \mathcal{FR}_1 back to d_A on line 12.
- Finally, at ❸, both devices bootstrapping before starting the file exchange. They first finish the route they will use to reach the other end (lines 13-15), then inform their e-squad that they started sharing a file by adding an interaction r to their sequence S (lines 16-17), and finally start exchanging f (line 18).

Relays selection We now detail the $\text{PickLayer}(\mathcal{V}, \theta)$ function, that takes care of intelligently selecting a layer’s devices. It takes two parameters: an input set of candidate nodes \mathcal{V} , and the *unavailability threshold* $\theta \in]0, 1]$, a configuration parameter that represents the desired maximum probability that all of the layer’s nodes fall offline at the same time (i.e. the probability that the layer be unavailable).

PickLayer iteratively picks a *random* node from \mathcal{V} without replacement, adds it to the output layer \mathcal{L} , and computes the *probability that all of the layer’s nodes fall*

offline at once, $P_{\mathcal{L}}^{\text{off}}$:

$$P_{\mathcal{L}}^{\text{off}} = \prod_{d \in \mathcal{L}} 1 - P_i(d) \quad (4)$$

$P_i(d)$ being the probability that device d remains online (cf. Eq. 2). The function returns either when the offline probability $P_{\mathcal{L}}^{\text{off}}$ falls below the threshold θ , or when the input view \mathcal{V} is emptied. As a baseline that will be used in the evaluation, `PickLayer` randomly chooses a predetermined number of nodes from the input view, without caring for the layer’s probability of becoming unavailable.

Non-e-squad layers are built with \mathcal{V}_{RPS} as input: it comprises a random pool of global `SPORES` participants that were online not long ago at least (cf. Sec. 2.2). The RPS view size $l_{\mathcal{V}}$ should be chosen big enough for `PickLayer` to reach the configured θ , but small enough that the view’s stale descriptors get evicted in a reasonable amount of time. E-squad layers, on the other hand (that is, \mathcal{FL}_4 and \mathcal{BL}_4), only have the less numerous e-squad members as input, such that `PickLayer` might not be able to reach the threshold before emptying the candidate list.

The smaller the threshold θ , the more nodes per layer, the better the route’s availability, but also the bigger the header. There is a trade-off between the readiness of routes and the message transit overhead.

Finally note that randomly picking descriptor from one’s view avoids biasing the relay selection in favour of supposedly highly connected nodes. Indeed, the devices’ availability estimate is published by themselves, and should not be trusted. Our approach gives no interest for attackers to lie on this value, while it encourages everyone to provide good estimates, for the sake of the routes’ reliability.

Exchanging a file As already told, `POR` provides anonymous UDP-like channels: order and integrity of the messages are not guaranteed by the protocol. These features must be supplied by `SPORES` on top of `POR`.

A file f exchanged through `SPORES` is chunked into fixed-size pieces, that are transmitted in order by the sender, along with their position (or ID). To ensure chunks integrity, we borrow from `BitTorrent` [11]: the file descriptor fd that is computed with `BuildFileDescriptor(f)` and provided to the receiver on bootstrap notably contains a SHA1 hash per chunk. The receiver verifies that the expected and computed hashes match every time they receive a chunk. The function creates the following descriptor:

$fd = (\text{ID}, \text{size}, \text{chunkSize}, \#\text{Chunks}, \text{chunksHash}, \text{hash})$

Each file is given a unique, random ID, picked by the uploader. The file descriptor also provides the file size, number of chunks and chunk size. The `chunksHash` is the concatenation of each chunk’s SHA1 hash, used by the receiver to verify the integrity of each chunk. Finally, `hash` is the SHA1 hash of `chunksHash`, to verify its own integrity. Using SHA1 hashes,

we ensure the file integrity. The order is guaranteed by the following sliding-window protocol.

To accelerate the file exchange, `SPORES` implements the Selective Repeat Automatic Repeat-reQuest (ARQ) [29, 39, 53] algorithm, a sliding-window protocol that lets the sender send several chunks at once, and allows the receiver to accept them out of order. The sender provides the chunk ID of each piece sent on the forward route, while the receiver sends back an acknowledgement (ACK) with the same ID for each received piece, using the backward route. When the sender does not receive an ACK after sending a chunk, it retries sending after a timeout of several seconds. The file exchange completes once each file chunk has been ACKed.

Finally, as can be seen in lines 4 and 9 of algorithm 4, any e-squad member can receive chunks/ACKs in spite of the proper message recipient. When they do, they can unravel the payload, and forward it to its proper recipient, by finding the recipient’s address in their e-squad sequence S . If the receiver is currently offline, they forward the message to any online e-squad member, until the recipient comes back online and is able to finally receive the message. In essence, the whole e-squad acts as a cache for received messages while the actual recipient is offline.

With these building blocks, we have proposed an entirely decentralised anonymous file exchange service for e-squads. It is specifically tailored for networks with high churn, and, thanks to its gossip components, it can scale to a theoretically unbounded number of users. We now analyse the security properties of `SPORES`, before evaluating its prototype.

3 Security analysis

We claim that using `SPORES` for exchanging files is more anonymous than using traditional onion routing (e.g. `OnionShare` on `Tor` [13]). To this aim, we statistically compare `SPORES`’ and `Tor`’s resilience to de-anonymisation attacks.

3.1 Assumptions and threat model

We take interest in an attacker owning a portion of the network relays (at the very least, their e-squad), and that can tamper with the protocol’s specification (they can notably break the random selection while forwarding messages in alg. 3). Their goal is to link two file exchange participants.

We do not consider the infamous Global Passive Adversary (GPA) attack model, where an attacker listens on all communication pipes. Using people’s devices drastically increases the number of Autonomous Systems (AS) involved in the protocol (mobile carriers, household connections...), rendering the GPA unlikely. In any case, GPA circumvention almost always involves generation of cover traffic [15, 40, 41, 50], which we cannot afford on constrained user end-devices.

Because we leverage secure peer sampling [24], we assume that the global overlay cannot be tampered with, and does return a uniform sample of online peers in the system.

Although headers are of variable size, we further assume that relays cannot guess their position on a route, as they do not know the number of relays per layer.

We already stated that the number of online relays in Tor (~ 6000) was small relative to the number of connected users (~ 2 million). SPORES seeks to involve each user device as a relay. Hence, we make the assumption that the number N_{Sp} of relays in SPORES is a multiple of those of Tor: $N_{\text{Sp}} = C \times N_{\text{TOR}}$, with $C \geq 1$. We further assume that there are N_{adv} colluding devices trying to de-anonymise Alice and Bob as they exchange a file. We write $p_{\text{Sp}} = N_{\text{adv}}/N_{\text{Sp}}$ the proportion in attackers in SPORES, and $p_{\text{TOR}} = N_{\text{adv}}/N_{\text{TOR}}$ the one in Tor. Finally, we consider $\#_{\mathcal{L}} = 3$ hops per route (excluding the final layer composed only of the recipient’s e-squad), and we assume a constant amount of $S_{\mathcal{L}}$ relays per layer.

3.2 Likelihood of the traffic correlation attack

It is well established that onion routing and Tor in particular are not resilient to end-to-end traffic correlation attacks [13, 25, 44, 47]. An attacker listening to each end of an onion route (by owning both end relays or observing traffic) can easily link sender and receiver, and thus de-anonymise the connection. In SPORES, due to the several relays per hop, all messages do not follow the same path. Considering also the increased number of relays in SPORES, we claim that traffic correlation attacks are more difficult than in Tor.

We do not model the operation of the traffic correlation attack. Instead, we study the probability that an adversary successfully positions themselves on a route’s first and last hops, and receives transmitted messages on both ends. We call this overall probability P [analyse mess.].

In Tor, the probability of having an adversary observe the same message on the first and last hops—knowing that they own these relays—is 1, since messages all go through the same relays once the route is built. In other words, the probability of seizing messages in Tor, P^{TOR} [analyse mess.], is simply the probability that the adversary successfully positions themselves on the first and last hop. We assume that the probability P^{TOR} [pick adv.] of selecting an adversary is the same for each layer (an overestimation of Tor’s actual security), and that relays are selected with replacement (which eases the computation, while only having a negligible impact on the outcome). Under these terms:

$$P^{\text{TOR}} [\text{analyse mess.}] = P^{\text{TOR}} [\text{pick adv.}]^2 = \left(\frac{N_{\text{adv}}}{N_{\text{TOR}}} \right)^2$$

In SPORES, it gets more complex. On the first and last layers, the situation is the same: the adversary has to own a number k of relays in the layer of size $S_{\mathcal{L}}$; then, the previous layer has to forward the message to the adversarial nodes. We consider

both cases to have the same independent probability:

$$P^{\text{Sp}} [\text{adv. reads mess.} \in \mathcal{L}] = \sum_{k=1}^{S_{\mathcal{L}}} P^{\text{Sp}} [k \text{ adv.} \in \mathcal{L}] \times P^{\text{Sp}} [\text{adv. reads mess.} \in \mathcal{L} \mid k \text{ adv.} \in \mathcal{L}]$$

Applying the standard equation for sampling with replacement, and considering a uniform probability of picking an adversary when they are k among $S_{\mathcal{L}}$, we obtain:

$$\begin{aligned} P^{\text{Sp}} [\text{adv. reads mess.} \in \mathcal{L}] &= \sum_{k=1}^{S_{\mathcal{L}}} \binom{S_{\mathcal{L}}}{k} p_{\text{Sp}}^k (1 - p_{\text{Sp}})^{S_{\mathcal{L}}-k} \times \frac{k}{S_{\mathcal{L}}} \\ &= p_{\text{Sp}} \sum_{j=0}^M \binom{M}{j} p_{\text{Sp}}^j (1 - p_{\text{Sp}})^{M-j} = p_{\text{Sp}} (p_{\text{Sp}} + (1 - p_{\text{Sp}}))^M = p_{\text{Sp}} \\ &\implies P^{\text{Sp}} [\text{analyse mess.}] = P^{\text{Sp}} [\text{adv. reads mess.} \in \mathcal{L}]^2 = p_{\text{Sp}}^2 \end{aligned}$$

We see that, in SPORES, the lesser security of selecting more nodes per layer (thus augmenting the probability to pick an adversary per hop) is strictly compensated by the probability to send a message to the adversaries. The probability that an adversary reads a message on a layer is equivalent to the probability of picking an adversary.

Comparing the approaches We see that SPORES’ probability of traffic analysis is better than Tor’s as long as C is superior to one (that is, as long as there are more relays in SPORES than in Tor):

$$\begin{aligned} P^{\text{Sp}} [\text{analyse mess.}] &< P^{\text{TOR}} [\text{analyse mess.}] \\ \iff \left(\frac{N_{\text{adv}}}{C \times N_{\text{TOR}}} \right)^2 &< \left(\frac{N_{\text{adv}}}{N_{\text{TOR}}} \right)^2 \iff C > 1 \end{aligned}$$

Since SPORES is specifically tailored to enable low-end client devices to participate in the network, which would increase the number of relays in the network, deploying probabilistic onion routes on a legacy network like Tor would indeed improve security.

3.3 Likelihood of having adversaries on each hop

Tor does not take much interest in the probability that adversaries own all relays on a circuit (trivially de-anonymising the route), as it is negligible with regards to the probability that they perform traffic correlation attacks. Still, because SPORES selects several nodes per layer, and because adversaries could break the random selection of relays while forwarding messages (and intently pick their accomplices in the next layer until destination), this attack vector needs to be studied in our case. We note this attack’s probability P [$\forall i, \text{adv} \in \mathcal{L}_i$].

In Tor’s case, still considering that each layer’s probability of picking an adversary is independent and equal, the probability that adversaries own the whole route is simply: $P^{\text{TOR}} [\forall i, \text{adv} \in \mathcal{L}_i] = (p_{\text{TOR}})^{\#_{\mathcal{L}}}$.

SPORES' case is again more complex. The adversary must first receive a message on the first layer \mathcal{L}_1 , i.e. with a probability of $P^{\text{Sp}}[\text{adv. reads mess.} \in \mathcal{L}] = p_{\text{Sp}}$. Then, they must own at least one relay on each of the next layers \mathcal{L}_2 and \mathcal{L}_3 , which for each layer has the probability:

$$\begin{aligned} P^{\text{Sp}}[\text{adv} \in \mathcal{L}] &= 1 - P^{\text{Sp}}[\text{adv} \notin \mathcal{L}] = 1 - (1 - p_{\text{Sp}})^{S_{\mathcal{L}}} \\ &\stackrel{q=1-p_{\text{Sp}}}{=} (1 - q) \times \sum_{k=0}^{S_{\mathcal{L}}-1} q^k = p_{\text{Sp}} \times \sum_{k=0}^{S_{\mathcal{L}}-1} (1 - p_{\text{Sp}})^k \end{aligned}$$

The probability that the adversary owns relays on each layers and successfully forwards it from source to destination is then:

$$\begin{aligned} P^{\text{Sp}}[\forall i, \text{adv} \in \mathcal{L}_i] &= P^{\text{Sp}}[\text{adv. reads mess.} \in \mathcal{L}_1] \times P^{\text{Sp}}[\text{adv} \in \mathcal{L}]^2 \\ &= p_{\text{Sp}}^3 \times \left(\sum_{k=0}^{S_{\mathcal{L}}-1} (1 - p_{\text{Sp}})^k \right)^2 \end{aligned}$$

Comparing the approaches If we overrate $(1 - p_{\text{Sp}})^k \approx 1$, we get the following inequality:

$$\begin{aligned} P^{\text{Sp}}[\forall i, \text{adv} \in \mathcal{L}_i] &< P^{\text{Tor}}[\forall i, \text{adv} \in \mathcal{L}_i] \\ &\approx \left(\frac{N_{\text{adv}}}{C \times N_{\text{TOR}}} \right)^3 \times S_{\mathcal{L}}^2 < \left(\frac{N_{\text{adv}}}{N_{\text{TOR}}} \right)^3 \implies C > S_{\mathcal{L}}^{2/3} \end{aligned}$$

Considering that increasing the layer size $S_{\mathcal{L}}$ past a certain threshold yields no performance gain (see section 4.2.3), and is costly in terms of message size and transmission time, we recommend values strictly inferior to $S_{\mathcal{L}} = 20$. With this upper bound, $C = 7.4$. We do expect a deployment of SPORES, with the same userbase as Tor, to reach a much bigger number of relays than $7.8 \times N_{\text{TOR}} = 46800$.

In this section, we have seen that SPORES' churn-resilient onion routing approach—the multi-path PORS—was not detrimental to its security under the two attack scenarii that we covered. In fact, assuming a bigger amount of relays than in Tor (i.e. assuming that Tor implements PORS), probabilistic onion routing even yields a security improvement.

4 Evaluation

In this section, we evaluate SPORES in terms of privacy and performance, depending on the user behaviour, and comparing to existing proposals. We first describe our evaluation protocol, before presenting our results in section 4.2.

4.1 Testbed

Let us first present how we simulated user behaviours, before going through our experimental setup.

4.1.1 User behavioural models

To the best of our knowledge, there exists no dataset that we could use to represent the behaviour of an e-squad owner. For this reason, we propose several models for simulating

users, with the objective of generating devices' connection and disconnection patterns encompassing the complexity of human behaviour.

We employ a discrete-time Hidden Markov Model (HMM) of order one [42] to represent a user going to different places (hidden process), and their device usage patterns depending on their location (observable processes). We thus assume that users switch location with a fixed period of T seconds, and that their next location only depends on the previous one (the Markovian hypothesis). Each device is modelled by a independent process: their availability only depends on the user's location (and not on other devices). The concept of availability encapsulates both the power state and transient connectivity of devices.

An example To illustrate our model, we display a fictitious user model comprising $N_{\text{loc}} = 3$ different locations and $N_{\text{dev}} = 4$ devices. Below are represented the $N_{\text{loc}} \times N_{\text{loc}}$ matrix A , that drives the user's movements, and the $N_{\text{loc}} \times N_{\text{dev}}$ matrix B , that is the concatenation of each device's probability at each state (i.e. nothing sums to one). Note that nothing prevents the user from using several devices at a time (B 's rows do not sum to one):

$$\begin{aligned} A = \begin{matrix} & \begin{matrix} Home & Outside & Work \end{matrix} \\ \begin{matrix} Home \\ Outside \\ Work \end{matrix} & \begin{bmatrix} 0.6 & 0.4 & 0 \\ 0.2 & 0.6 & 0.2 \\ 0 & 0.4 & 0.6 \end{bmatrix} \end{matrix} \\ B = \begin{matrix} & \begin{matrix} Phone & Laptop & Home computer & Workstation \end{matrix} \\ \begin{matrix} Home \\ Outside \\ Work \end{matrix} & \begin{bmatrix} 0.8 & 0.6 & 0.7 & 0 \\ 0.6 & 0.2 & 0 & 0 \\ 0.7 & 0.2 & 0 & 0.7 \end{bmatrix} \end{matrix} \end{aligned}$$

Once a model is built, we perform a random walk of L rounds to generate a sequence, or timeline, of interactions $X \in \{0, 1\}^{N_{\text{dev}} \times L}$. Like in section 2.1.2, $X_i(d) = 1$ means that device d was online at round i , and equals 0 otherwise. Devices, through the e-squad overlay, learn the timeline X to predict their future availability, but not the hidden location sequence of their user.

Diverse user models With this HMM ground, we can build models with variable mean availability and predictability. Consider, for instance, a model with only one location, with every device's probability equal to 0.5. Intuitively, this model is the most unpredictable we could build, as all devices switch state with uniform probability. On the contrary, a model with L locations that are visited in order by the user, and devices probabilities that are either 1 or 0, is very predictable: the device connection timeline is deterministic, and loops every L usage rounds.

To build matrices A and B that display such diversity, we sample their content using the beta distribution, a versatile probability distribution function defined on $[0, 1]$, first studied by Pearson in 1895 [38]. The matrix A is then normalized as needs be. The beta distribution has two shape parameters $(\alpha, \beta) \in \mathbb{R}^{+*}$; we are interested in the function's smoothed binomial shape when both parameters are below

one. Small values of $\alpha = \beta$ lead to samples closer to 0 or 1, while $\alpha = \beta = 1$ is the uniform distribution. We skew the distribution, for a fixed $\beta = 0.6$, by varying the expected value μ by picking $\alpha = \frac{\beta}{\mu^{-1}-1}$.

We measure the predictability of a given model by generating a timeline X of length L . To evaluate the predictability, we first compute the probability $P_i(d)$ that each device d stays online, for each round X_i such that $L_{\text{init}} < i < L$. (Since computing $P_i(d)$ requires initial information, we only compute it for steps past a number of rounds L_{init} .) Then, we compare $P_i(d)$ with the actual outcome $X_{i+1}(d)$ using a logarithmic scoring rule, that is: $sc_i(d) = X_{i+1}(d)\log(P_i(d)) + (1 - X_{i+1}(d))\log(1 - P_i(d))$. Finally, the total predictability of a model is the average of all computed log scores.

We proposed 4 user behavioural models representing diverse predictabilities. For each of them, unless otherwise noted, we set $N_{\text{loc}} = 4$ locations and $N_{\text{dev}} = 6$ devices. L_{init} is always set to 50, while the total sequence length L depends on the experiment duration:

1. Uniform (*Uni.*): $N_{\text{loc}} = 1$; each device’s probability is equal to μ . This model shows no periodicity, and is thus the least predictable.
2. Unpredictable (*Unpred.*): as above, the HMM matrices are sampled from a beta distribution, this time with $\beta = 0.8$, generating transitions and device probabilities closer to 0 or 1.
3. Predictable (*Pred.*): here, $\beta = 0.2$, which brings probabilities even closer to 0 or 1.
4. Deterministic (*Det.*): the user cycles deterministically through the set of N_{loc} locations. Devices probability being always 0 or 1, the timeline X is entirely deterministic.

We evaluate the performance of SPORES with regard to these different models in section 4.2.1.

4.1.2 Methodology

To evaluate SPORES, we built a prototype in 6100 lines of Go, including all core functions except the cryptography. The users’ behaviours, driving the devices churn, were simulated with 1600 lines of Python. Each device runs as a Docker container, participating in a single virtual network. Due to the scale of the experiment, and to generate somewhat realistic network traffic, each user’s devices are scattered over a multi-host Docker Swarm. The experiments were deployed on 6 AWS ‘r5.large’ VMs, plus another one to orchestrate the experiments.

The experimental process is the following: we initialise the experiment by letting each user’s device spawn on a random VM. At this time, booted devices start exchanging descriptors in the global overlay. Once every device is started, we start scheduling each device according to their user’s behavioural model, updating their availability state every T

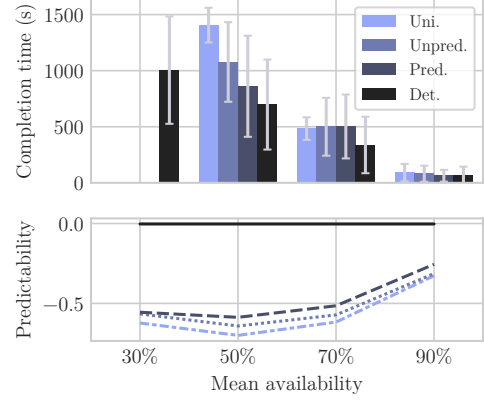


Figure 4. *On top:* File transfer median completion times for each user model, per mean availability μ . Error bars represent the interquartile range. *On the bottom:* The predictability of each user model per mean availability. Higher is better, as this is a log score.

seconds. We then pick two random online devices belonging to different users, and perform the initial route creation through REST calls to each of them. We repeat the operation until we consider enough files were exchanged, leaving a reasonable time between exchanges to avoid saturating the network. We tear down the network and retrieve results after leaving some time for the devices to proceed with their file exchanges. Note that the orchestrating VM is not able to assess whether files finished exchanging, such that some files fail downloading for lack of time.

Parameters There are $N_U = 25$ users in the system, each running 6 devices, resulting in a network of 150 relays, randomly scattered over the hosts. Users switch between $N_{\text{loc}} = 4$ states.

We fixed the file size to 50MiB, and the chunk size to 512KiB, resulting in 100 chunks per file transfer. To bootstrap the e-squad overlay, we provide them with an initial user activity sequence of $L_{\text{init}} = 50$ device usages. This way, even the initial availability predictions are backed by a reasonably accurate model.

Each user interaction lasts $T = 6$ seconds; we leave $5 * T = 30$ s between each file exchange; exchange 50 files per experiment, and tear down the experiment $20 * T = 2m$ after the last transfer started.

Unless otherwise noted, the unavailability threshold θ for creating routes equals 0.001. When the user model is not specified, the unpredictable is under study. When the mean availability μ is not specified, it equals 50%.

4.2 Conducted experiments

We evaluated our system on three different regards: its performance under different conditions, its security against colluding attackers, and finally the dynamics of PORs.

4.2.1 Influence of the users' behaviour

We first study the influence of the user models and mean availability of devices on the file transfer performances. Towards this goal, we perform an experiment per model described in section 4.1.1 and per $\mu \in [0.3, 0.5, 0.7, 0.9]$. On figure 4, we display the file transfers median completion times for each of these experiments, along with the predictability of each model. The top error bars represent the interquartile range (that is the range between the 25th and 75th percentiles of the completion time distribution). Empty error bars mean that no file exchange succeeded (at $\mu = 30\%$, all models fail except the unpredictable (*Unpred.*) one). Because the file transfer is handled by devices that suffer from churn just as their fellow peers, the transfer times should not be taken literally: they merely serve as a metric to compare outcomes.

We see that the deterministic model stands out in terms of predictability, while the other models follow a similar pattern with their predictability being minimal at $\mu = 50\%$. They are still ordered as was predicted in sec. 4.1.1. Consequently, the deterministic model always shows better file exchange completion times than the other models. Most importantly, it is the only model where transfers are entirely completed when devices are only available 30% of the time. We also see that the influence of the user model decreases as the network get more available: it is more interesting to perform predictive routing when the risk of dropping messages is big.

4.2.2 Security measurements

Using the logs from the Unpred. model with $\mu = 50\%$, we now study the occurrence of the attacks presented in section 3. Depending on the proportion of colluding users attacking the system, we display on top of figure 5 the percentage of routes that were corrupted—either on each hop, as discussed in section 3.3, or only on the first and last hops, as seen in sec. 3.2. The bottom of the figure reads the amount of seized messages once routes are compromised, showing how SPORES circumvents traffic analysis attacks. We only showed results for a single experiment/mean availability couple, because the probability of such attacks does not depend on either parameter (the output curves were mostly equal).

All these statistics were computed as follows: knowing that there are $N_U = 25$ users in each experiment, there are always 23 potential adversaries per file exchange (excluding the sender and receiver). We consider a number $N_{adv}^U \in \llbracket 1, \dots, 17 \rrbracket$ of evil users, conspiring to de-anonymise the whole network. For each value of N_{adv}^U , we computed up to a thousand combinations of conniving users, and counted

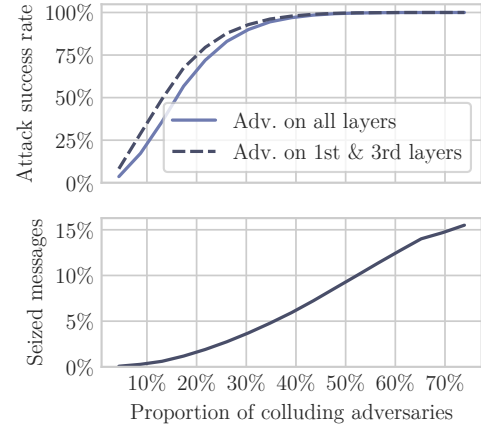


Figure 5. *On top:* Proportion of compromised routes as a function of the percentage of colluding adversaries over the whole network, for both attack types. We see that as long as the percentage of colluding attackers remains low (which is likely to be the case in our context), only few routes are corrupted. *On the bottom:* Proportion of messages observed by the attacker on each end of the route, considering only routes that were compromised.

the number of times they successfully compromised a route, for each route that was created during the experiment. Note that, doing so, we consider that attackers follow the same churn model as other users.

For the top of figure 5, we counted the number of times the set of attackers successfully positioned themselves on a route, normalised by the total number of routes created in the experiment. We plotted two curves for the case where an attacker is on the whole route, and when they are only on the route's ends. The bottom of the figure displays the proportion of messages effectively observed by both ends of a compromised route while attackers perform an end-to-end traffic correlation attack. To compute this statistics, for each compromised route, we counted the number of messages that passed through the adversarial relays, normalised by the total number of messages passing through this route.

We firstly see that it suffices to own around 40% of the network for all created routes to be compromised. 3.6% of the routes would be entirely compromised by an attacker owning 4.3% of the nodes. Although daunting, this observation constitutes a major argument in favour of the multiplication of relays in any onion network. To resist such de-anonymisation attempts, it is crucial for onion networks to let any participating device partake in the routing.

Where SPORES stands out, in terms of security, is on its resilience to traffic analysis attacks: even when an attacker successfully positions himself on both ends of the route, they can hardly observe 15% of traffic, even if they subverted most of the network. Given that most existing end-to-end




Figure 6. Success rate of the messages transit depending on the unavailability threshold θ . With $\theta = 1$, SPORES behaves like traditional onion routing with one node per layer.

traffic correlation attacks assume that the eavesdropper sees all the traffic, the multi-path routing severely hampers this attack vector.

4.2.3 Tuning Probabilistic Onion Routes

We have performed a last experiment using the Unpred. model with $\mu = 0.5$, where we varied the unavailability threshold θ used by the PickLayer route selection function. We took $\theta = \{1, .1, .01, .001, .0001\}$, and exchanged 50 files, resulting in 10 files per value of θ . Remember that, when $\theta = 1$, the PickLayer function is satisfied as long as the layer’s probability of being offline ($P_{\mathcal{L}}^{\text{off}}$) is lower than 1, that is, when there is one device per layer. At $\theta = 1$, the resulting route will resemble traditional onion routing.

To assess the influence of θ , we counted the number of messages that successfully traversed their POR for each parameter value, resulting in the message transit success rates displayed on figure 6. Given the high churn of the experiment, traditional onion routes ($\theta = 1$) only allowed 9.2% of their traffic to go through. We see that adding any number of devices per layer allows to reach 75% of reliability at least. The most reliable route occurs when $\theta = .01$, where 91% of the messages go through. Adding more devices per layer only weakens the route despite the added redundancy: this is explained by the bigger network cost (due to the increased header size) and latency (due to the sequential attempts at reaching the next layer’s relays) of adding alternatives. It is interesting to note that, as θ shrinks exponentially, the number of relays per layer seems to grow linearly: it is 2.5 for $\theta = 0.1$, 4, when $\theta = 0.01$, 6 when $\theta = 0.001$, and 8 when $\theta = 0.0001$.

We see that multi-path routing is a very promising prospect for onion routing over an unreliable network, and that a reasonable threshold is enough to provide maximum route efficiency. The header size linearly grows (per increments of a symmetric block size) per the number of relays per hop, while the added cryptographic cost of deciphering a handful of envelopes per hop is minimal. We hope to improve on our header format in future works to

make PORs even more powerful, *via* the use of Sphinx [12] to enable more compact and secure header formats.

5 Related works

Anonymous file sharing between people is not an easy problem, for technical and political reasons [21, 33]. In 2014, OnionShare (<https://onionshare.org>) solved the issue by proposing a solution over Tor: one of the two persons (the ‘server’) willing to exchange a file creates an onion service over Tor [13], and provides the other person (the ‘client’) with an onion link (a random hash URL finishing with ‘.onion’) pointing to the service. The client then visits that site through the Tor browser, and can either download from or upload to the server (depending on the configuration mode). Once the transfer is completed, the server tears down the service, leaving no further trace of the file exchange. OnionShare requires that the server be created *prior* to the file exchange, while SPORES requires no such bootstrap, simplifying the exchange. Furthermore, OnionShare leverages Tor, consequently it is inherently susceptible to a variety of traffic analysis attacks [7, 36, 45]. A variety of proposals attempt to circumvent the attacks by enhancing the route selection process [3, 49, 52]. Our work, orthogonal to these, takes another approach: we promote a multiplication of relays while being churn tolerant to effectively improve anonymity. Supporting the same claim, HORNET [9] proposes a new onion routing strategy aiming better performance and resistance to mass surveillance programs by rendering relays stateless. However, SPORES is still more resilient to the aforementioned attacks, though, thanks to our multi-path routing approach.

Some academic proposals, such as Tarzan [15], Vuvuzela [50] or Loopix [40] do tackle also traffic analysis, and even the Global Passive Adversary (GPA) model where an attacker would listen on all communication pipes. All of them achieve this feat by generating dummy cover traffic, which we consider undesirable due to the important footprint of such approach, when we target mobile appliances with constrained resources.

We take the most interest in efforts to decentralise anonymity networks, which would allow them to scale and be more resilient. I2P, being more than a decade years old, has to be cited as a fully P2P anonymity system, comprising 34k daily users [22]. Although, the lack of coverage of its security properties does not allow to compare it to other systems. Recent prospects to allow decentralisation of networks take interest in leveraging blockchain technologies, or trusted computer zones, to realise critical building blocks of decentralised systems. NextLeap [19], for instance, offers to solve the problem of identifying peers using blockchain technologies. SGX-Tor [26] proposes to make onion relays more secure by running them in encrypted enclaves; and ConsenSGX suggests that Tor’s centralised Directory Authorities consensus could scale to more relay servers

using the same components. SPORES does not need trusted computing to warrant its security properties. Although, blockchains or trusted enclaves would be an interesting avenue for future developments in user authentication, for instance.

6 Conclusion

With SPORES, we have proposed an anonymous P2P file transfer protocol by revisiting traditional onion routing and leveraging on the people's own devices. Overall, we have seen that SPORES was a sound approach to onion routing in challenging network conditions. Through its predictive component, it can successfully accomplish file transfers even in the worst connectivity scenarios. Its security properties are novel, as it is one of the first onion routing approaches to finally hinder traffic correlation attacks. Its design makes it fit for large scale deployments on commodity hardware, which would lower the risk of de-anonymised routes. Finally, the multi-path routing approach proves its worth when compared to the legacy on an unstable network: the fact of proposing just one alternative node per layer already increases the routes' reliability by 65%.

In future works, we hope to improve on SPORES on several aspects. Our predictive model could be enhanced; it would firstly require some field studies about people's usage of their e-squads. We also look upon contributions like Sphinx [12] to have more compact and secure header formats.

References

- [1] 2019. *Cisco Visual Networking Index: Forecast and Trends, 2017–2022*. Technical Report. Cisco. <https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white-paper-c11-741490.html>
- [2] A. Bakker and M. v Steen. 2008. PuppetCast: A Secure Peer Sampling Protocol. In *2008 European Conference on Computer Network Defense*. 3–10. DOI : <http://dx.doi.org/10.1109/EC2ND.2008.7>
- [3] Armon Barton, Mohsen Imani, Jiang Ming, and Matthew Wright. 2018. Towards Predicting Efficient and Anonymous Tor Circuits. *arXiv:1805.01977 [cs]* (May 2018). arXiv:cs/1805.01977
- [4] Edward Bortnikov, Maxim Gurevich, Idit Keidar, Gabriel Kliot, and Alexander Shraer. 2009. Brahms: Byzantine Resilient Random Membership Sampling. *Computer Networks* 53, 13 (Aug. 2009), 2340–2359. DOI : <http://dx.doi.org/10.1016/j.comnet.2009.03.008>
- [5] Yérom-David Bromberg, Adrien Luxey, and François Taïani. 2018. CASCADE: Reliable Distributed Session Handoff for Continuous Interaction Across Devices. In *2018 IEEE 38th International Conference on Distributed Computing Systems (ICDCS)*. 244–254. DOI : <http://dx.doi.org/10.1109/ICDCS.2018.00033>
- [6] Carole Cadwalladr. 2018. ‘I Made Steve Bannon’s Psychological Warfare Tool’: Meet the Data War Whistleblower. *The Guardian* (March 2018). <https://www.theguardian.com/news/2018/mar/17/data-war-whistleblower-christopher-wylie-faceook-nix-bannon-trump>
- [7] Xiang Cai, Xin Cheng Zhang, Brijesh Joshi, and Rob Johnson. 2012. Touching from a Distance: Website Fingerprinting Attacks and Defenses. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security - CCS ’12*. ACM Press, Raleigh, North Carolina, USA, 605. DOI : <http://dx.doi.org/10.1145/2382196.2382260>
- [8] David L. Chaum. 1981. Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms. *Commun. ACM* 24, 2 (2 1981), 84–90.
- [9] Chen Chen, Daniele Enrico Asoni, David Barrera, George Danezis, and Adrian Perrig. 2015. HORNET: High-Speed Onion Routing at the Network Layer. *arXiv:1507.05724 [cs]* (July 2015). arXiv:cs/1507.05724
- [10] Ian Clarke, Oskar Sandberg, Brandon Wiley, and Theodore W. Hong. 2001. Freenet: A Distributed Anonymous Information Storage and Retrieval System. In *International Workshop on Designing Privacy Enhancing Technologies: Design Issues in Anonymity and Unobservability*. Springer-Verlag New York, Inc., 46–66.
- [11] Bram Cohen. 2008. *The BitTorrent Protocol Specification*. Technical Report. BitTorrent Inc. http://www.bittorrent.org/beps/bep_0003.html
- [12] G. Danezis and I. Goldberg. 2009. Sphinx: A Compact and Provably Secure Mix Format. In *2009 30th IEEE Symposium on Security and Privacy*. 269–282. DOI : <http://dx.doi.org/10.1109/SP.2009.15>
- [13] Roger Dingledine, Nick Mathewson, and Paul Syverson. 2004. *Tor: The Second-Generation Onion Router*. Technical Report. Naval Research Lab Washington DC.
- [14] Amos Fiat and Moni Naor. 1993. Broadcast Encryption. In *Advances in Cryptology — CRYPTO ’93 (Lecture Notes in Computer Science)*, Douglas R. Stinson (Ed.). Springer Berlin Heidelberg, 480–491.
- [15] Michael J. Freedman and Robert Morris. 2002. Tarzan: A Peer-to-peer Anonymizing Network Layer. In *Proceedings of the 9th ACM Conference on Computer and Communications Security (CCS ’02)*.
- [16] David M. Goldschlag, Michael G. Reed, and Paul F. Syverson. 1996. Hiding Routing Information. In *Information Hiding*, Gerhard Goos, Juris Hartmanis, Jan Leeuwen, and Ross Anderson (Eds.). Vol. 1174. Springer Berlin Heidelberg, Berlin, Heidelberg, 137–150. DOI : http://dx.doi.org/10.1007/3-540-61996-8_37
- [17] Jennifer Stisa Granick. 2017. *We Kill People Based on Metadata*. Cambridge University Press, 53–66. DOI : <http://dx.doi.org/10.1017/9781316216088.006>
- [18] Christian Grothoff. 2017. *The GNUUnet System*. PhD Thesis. Université de Rennes 1.
- [19] Harry Halpin. 2017. NEXTLEAP: Decentralizing Identity with Privacy for Secure Messaging. In *Proceedings of the 12th International Conference on Availability, Reliability and Security - ARES ’17*. ACM Press, Reggio Calabria, Italy, 1–10. DOI : <http://dx.doi.org/10.1145/3098954.3104056>
- [20] Danny Harnik, Benny Pinkas, and Alexandra Shulman-Peleg. 2010. Side Channels in Cloud Services: Deduplication in Cloud Storage. *IEEE Security & Privacy* 8, 6 (2010), 40–47.
- [21] Parker Higgins. 2014. The Troubling Truth of Why It’s Still So Hard to Share Files Directly. <https://www.wired.com/2014/06/the-troubling-truth-of-why-its-still-so-hard-to-share-files-directly/>, *Wired* (June 2014). <https://www.wired.com/2014/06/the-troubling-truth-of-why-its-still-so-hard-to-share-files-directly/>
- [22] Nguyen Phong Hoang, Panagiotis Kintis, Manos Antonakakis, and Michalis Polychronakis. 2018. An Empirical Study of the I2P Anonymity Network and Its Censorship Resistance. In *Proceedings of the Internet Measurement Conference 2018 on - IMC ’18*. ACM Press, Boston, MA, USA, 379–392. DOI : <http://dx.doi.org/10.1145/3278532.3278565>
- [23] Márk Jelasity, Spyros Voulgaris, Rachid Guerraoui, Anne-Marie Kermarrec, and Maarten Van Steen. 2007. Gossip-based peer sampling. *ACM ToCS* 25, 3 (2007), 8.
- [24] Gian Paolo Jesi, Alberto Montresor, and Maarten van Steen. 2010. Secure Peer Sampling. *Computer Networks* 54, 12 (Aug. 2010), 2086–2098. DOI : <http://dx.doi.org/10.1016/j.comnet.2010.03.020>
- [25] Aaron Johnson, Chris Wacek, Rob Jansen, Micah Sherr, and Paul Syverson. 2013. Users Get Routed: Traffic Correlation on Tor by Realistic Adversaries. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security - CCS ’13*. ACM Press, Berlin, Germany, 337–348. DOI : <http://dx.doi.org/10.1145/2508859.2516651>
- [26] Seongmin Kim, Juhyeng Han, Jaehyeong Ha, Taesoo Kim, and Dongsu Han. 2018. SGX-Tor: A Secure and Practical Tor Anonymity Network With SGX Enclaves. *IEEE/ACM Transactions on Networking* 26, 5 (Oct. 2018), 2174–2187. DOI : <http://dx.doi.org/10.1109/TNET.2018.2868054>
- [27] Michal Kosinski, David Stillwell, and Thore Graepel. 2013. Private Traits and Attributes Are Predictable from Digital Records of Human Behavior. *Proceedings of the National Academy of Sciences* 110, 15 (April 2013), 5802–5805. DOI : <http://dx.doi.org/10.1073/pnas.1218721110>
- [28] Stevens Le Blond, David Choffnes, Wenxuan Zhou, Peter Druschel, Hitesh Ballani, and Paul Francis. 2013. Towards Efficient Traffic-Analysis Resistant Anonymity Networks. In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM - SIGCOMM ’13*. ACM Press, Hong Kong, China, 303–314. DOI : <http://dx.doi.org/10.1145/2486001.2486002>
- [29] J. A. Lockitt, A. G. Gatfield, and T. R. Dobyns. 1975. A Selective Repeat ARQ System. In *3rd International Conference on Digital Satellite Communications*. 189–195.
- [30] Adrien Luxey, Yérom-David Bromberg, Fábio M. Costa, Vinícius Lima, Ricardo C. A. da Rocha, and François Taïani. 2018. Sprinkler: A Probabilistic Dissemination Protocol to Provide Fluid User Interaction in Multi-Device Ecosystems. In *2018 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. 1–10. DOI : <http://dx.doi.org/10.1109/PERCOM.2018.8444577>
- [31] Ewen MacAskill, Gabriel Dance, Feilding Cage, Greg Chen, and Nadja Popovich. 2013. NSA Files Decoded: Edward Snowden’s Surveillance Revelations Explained. *The Guardian* (Nov. 2013).
- [32] S. C. Matz, M. Kosinski, G. Nave, and D. J. Stillwell. 2017. Psychological Targeting as an Effective Approach to Digital Mass Persuasion. *Proceedings of the National Academy of Sciences* 114, 48 (Nov. 2017), 12714–12719. DOI : <http://dx.doi.org/10.1073/pnas.1710966114>
- [33] Randall Munroe. 2011. File Transfer. (Sept. 2011).

- [34] Ellen Nakashima and Joby Warrick. 2013. For NSA Chief, Terrorist Threat Drives Passion to ‘Collect It All’. *Washington Post* (14 July 2013).
- [35] Arjun Nambiar and Matthew Wright. 2006. Salsa: A Structured Approach to Large-scale Anonymity. In *Proceedings of the 13th ACM Conference on Computer and Communications Security (CCS '06)*. 17–26.
- [36] Milad Nasr, Amir Houmansadr, and Arya Mazumdar. 2017. Compressive Traffic Analysis: A New Paradigm for Scalable Traffic Analysis. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security - CCS '17*. ACM Press, Dallas, Texas, USA, 2053–2069. DOI: <http://dx.doi.org/10.1145/3133956.3134074>
- [37] John Naughton. 2016. Death by drone strike, dished out by algorithm. *The Guardian* (2 2016). <https://www.theguardian.com/commentisfree/2016/feb/21/death-from-above-nia-csa-skynet-algorithm-drones-pakistan>
- [38] Karl Pearson and Olaus Magnus Friedrich Erdmann Henrici. 1895. X. Contributions to the Mathematical Theory of Evolution.—II. Skew Variation in Homogeneous Material. *Philosophical Transactions of the Royal Society of London. (A.)* 186 (Jan. 1895), 343–414. DOI: <http://dx.doi.org/10.1098/rsta.1895.0010>
- [39] Larry L. Peterson and Bruce S. Davie. 2003. *Computer Networks: A Systems Approach, 3rd Edition*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA. 97–110 pages.
- [40] Ania M. Piotrowska, Jamie Hayes, Tariq Elahi, Sebastian Meiser, and George Danezis. 2017. The Loopix Anonymity System. In *26th USENIX Security Symposium (USENIX Security 17)*. 1199–1216.
- [41] John P. Podolanko, Revanth Pobala, Hussain Mucklai, George Danezis, and Matthew Wright. 2017. LiLAC: Lightweight Low-Latency Anonymous Chat. *IEEE*, 141–151. DOI: <http://dx.doi.org/10.1109/PAC.2017.14>
- [42] Lawrence R. Rabiner and Biing Hwang Juang. 1986. An Introduction to Hidden Markov Models. *IEEE ASSP Magazine* 3 (1986), 4–16. DOI: <http://dx.doi.org/10.1109/MASSP.1986.1165342>
- [43] Marc Rennhard and Bernhard Plattner. 2002. Introducing MorphMix: Peer-to-peer Based Anonymous Internet Usage with Collusion Detection. In *Proceedings of the 2002 ACM Workshop on Privacy in the Electronic Society (WPES '02)*.
- [44] Florentin Rochet and Olivier Pereira. 2017. Waterfilling: Balancing the Tor Network with Maximum Diversity. *Proceedings on Privacy Enhancing Technologies* 2017, 2 (April 2017), 4–22. DOI: <http://dx.doi.org/10.1515/popets-2017-0013>
- [45] Florentin Rochet and Olivier Pereira. 2018. Dropping on the Edge: Flexibility and Traffic Confirmation in Onion Routing Protocols. *Proceedings on Privacy Enhancing Technologies* 2018, 2 (April 2018), 27–46. DOI: <http://dx.doi.org/10.1515/popets-2018-0011>
- [46] Stuart J. Russell and Peter Norvig. 2003. *Artificial Intelligence, A Modern Approach. Second Edition*. University of Michigan Press. 835 pages.
- [47] Andrei Serjantov and Peter Sewell. 2003. Passive Attack Analysis for Connection-Based Anonymity Systems. In *Computer Security – ESORICS 2003 (Lecture Notes in Computer Science)*, Einar Snekkenes and Dieter Gollmann (Eds.). Springer Berlin Heidelberg, 116–131.
- [48] Douglas R Stinson. 2005. *Cryptography: theory and practice*. Chapman and Hall/CRC.
- [49] Yixin Sun, Anne Edmundson, Nick Feamster, Mung Chiang, and Prateek Mittal. 2017. Counter-RAPTOR: Safeguarding Tor Against Active Routing Attacks. In *2017 IEEE Symposium on Security and Privacy (SP)*. 977–992. DOI: <http://dx.doi.org/10.1109/SP.2017.34>
- [50] Jelle van den Hooff, David Lazar, Matei Zaharia, and Nikolai Zeldovich. 2015. Vuvuzela: Scalable Private Messaging Resistant to Traffic Analysis. In *Proceedings of the 25th Symposium on Operating Systems Principles - SOSP '15*. ACM Press, Monterey, California, 137–152. DOI: <http://dx.doi.org/10.1145/2815400.2815417>
- [51] Spyros Voulgaris, Daniela Gavidia, and Maarten van Steen. CYCLON: Inexpensive Membership Management for Unstructured P2P Overlays. 13, 2 (????), 197–217. DOI: <http://dx.doi.org/10.1007/s10922-005-4441-x>
- [52] Gerry Wan, Aaron Johnson, Ryan Wails, Sameer Wagh, and Prateek Mittal. 2019. Guard Placement Attacks on Path Selection Algorithms for Tor. *Proceedings on Privacy Enhancing Technologies* 2019, 4 (Oct. 2019), 272–291. DOI: <http://dx.doi.org/10.2478/popets-2019-0069>
- [53] E. Weldon. 1982. An Improved Selective-Repeat ARQ Strategy. *IEEE Transactions on Communications* 30, 3 (March 1982), 480–486. DOI: <http://dx.doi.org/10.1109/TCOM.1982.1095497>
- [54] Wu Youyou, Michal Kosinski, and David Stillwell. 2015. Computer-Based Personality Judgments Are More Accurate than Those Made by Humans. *Proceedings of the National Academy of Sciences* 112, 4 (Jan. 2015), 1036–1040. DOI: <http://dx.doi.org/10.1073/pnas.1418680112>