

Stephan Heuser, Bradley Reaves, Praveen Kumar Pendyala, Henry Carter, Alexandra Dmitrienko, William Enck, Negar Kiyavash, Ahmad-Reza Sadeghi, and Patrick Traynor

Phonion: Practical Protection of Metadata in Telephony Networks

Abstract: The majority of people across the globe rely on telephony networks as their primary means of communication. As such, many of the most sensitive personal, corporate and government related communications pass through these systems every day. Unsurprisingly, such connections are subject to a wide range of attacks. Of increasing concern is the use of metadata contained in Call Detail Records (CDRs), which contain source, destination, start time and duration of a call. This information is potentially dangerous as the very act of two parties communicating can reveal significant details about their relationship and put them in the focus of targeted observation or surveillance, which is highly critical especially for journalists and activists.

To address this problem, we develop the Phonion architecture to frustrate such attacks by separating call setup functions from call delivery. Specifically, Phonion allows users to preemptively establish call circuits across multiple providers and technologies before dialing into the circuit and does not require constant Internet connectivity. Since no single carrier can determine the ultimate destination of the call, it provides unlinkability for its users and helps them to avoid passive surveillance. We define and discuss a range of adversary classes and analyze why current obfuscation technologies fail to protect users against such metadata attacks. In our extensive evaluation we further analyze advanced anonymity technologies (e.g., VoIP over Tor), which do not preserve our functional requirements for high voice quality in the *absence* of constant broadband Internet connectivity and compatibility with landline and feature phones. Phonion is the first practical system to provide guarantees of unlinkable communication against a range of practical adversaries in telephony systems.

Keywords: Metadata protection, anonymous telephony, privacy-preserving communications

Received 2017-05-31; revised 2017-08-01; accepted 2017-08-01.

Stephan Heuser: Intel CRI-SC and TU Darmstadt, E-mail: stephan.heuser@trust.tu-darmstadt.de

Bradley Reaves: University of Florida, E-mail: reaves@ufl.edu

1 Introduction

Telecommunication companies record network use by individual customers via Call Data Records (CDRs). CDRs contain important metadata ranging from call source and destination to duration of the connection and the route through the telephony network. Such metadata have most recently been associated with large-scale collection campaigns by intelligence agencies [24]. While these organizations often assert that such programs are necessary to prevent crime and terrorism, privacy advocates argue that the complete cataloging of telephony metadata erodes civil liberties. For example, oppressive regimes can use CDR analysis to identify and harass freedom fighters in civil war zones, such as Syria. However, what researchers and policy makers have failed to consider is that a range of other adversaries may also use CDRs to violate the privacy of targeted individuals. In 2006, for example, detectives hired by executives at Hewlett-Packard were able to use social engineering to acquire phone records and determine the identity of an anonymous corporate board member who leaked sensitive information to journalists [33]. Such attacks are not limited to private detectives, but have also been executed by jealous spouses [2], curious neighbors [58], companies paying for employee cell phones [11] and rogue employees of cellular network providers [22, 53]. In 2011, major security flaws in Vodafone’s data system were re-

Praveen Kumar Pendyala: TU Darmstadt, E-mail: praveen.pendyala@trust.tu-darmstadt.de

Henry Carter: Villanova University, E-mail: henry.carter@villanova.edu

Alexandra Dmitrienko: ETH Zurich, E-mail: alexandra.dmitrienko@inf.ethz.ch

William Enck: North Carolina State University, E-mail: enck@cs.ncsu.edu

Negar Kiyavash: University of Illinois, E-mail: kiyavash@illinois.edu

Ahmad-Reza Sadeghi: Intel CRI-SC and TU Darmstadt E-mail: ahmad.sadeghi@trust.tu-darmstadt.de

Patrick Traynor: University of Florida, E-mail: traynor@cise.ufl.edu

ported, which resulted in CDRs of millions of customers being available on the Internet [44].

These threats have motivated academic research to develop anonymous voice communication systems. The most common and well-studied approach is the use of Voice over IP (VoIP) telephony in combination with low-latency anonymization networks, such as Tor [21]. While Tor is widely believed to provide reasonable anonymity guarantees, there are situations where its shortcomings prohibit adoption. First, compared to the telephony network, Tor relays are more susceptible to congestion, which negatively impacts voice quality. Second, VoIP over Tor mandates constant high-bandwidth Internet connectivity for both caller and callee, which is not always reasonable, for example in rural areas or in developing countries. Finally, in many cases, sensitive communications must take place over traditional telephone systems. This is especially true in journalism, where sources often dictate the use of phone calls [41].

In this paper, we present Phonion, an alternative solution which addresses shortcomings of existing anonymization networks regarding voice communication. Phonion routes calls over the telephony infrastructure and achieves high quality of calls while obfuscating call data records. Our architecture generates alias telephony numbers for its users and does not require Internet connectivity during calls. Phonion is compatible with a wide variety of end user devices – ranging from rotary phones to VoIP clients – and resilient to compromise of (a number of) telephony network operators.

In particular, we make the following contributions:

- *Design of Phonion:* We define the spectrum of adversaries (cf. Section 2) and design the Phonion system of loosely cooperating telephony services to establish and relay calls so that the source and destination of a call are unlinkable using CDR analysis (cf. Section 3). Our contribution lies in combining *existing* technologies in a novel way to create an out-of-band signaling overlay network and phone call forwarding infrastructure.
- *Implementation and extensive evaluation:* We provide a full implementation of Phonion which supports various telephony technologies and is compatible with a vast diversity of end-user devices, ranging from rotary phones to VoIP clients (cf. Section 4). We intend to make our implementation available to the research community. We evaluate our implementation using professional industry-standard voice quality analysis tools and metrics to demonstrate that Phonion maintains call fidelity when compared to standard phone calls (cf. Section 5).

- *Security analysis and comparison against a range of proposed alternatives:* We analyze privacy guarantees provided by Phonion (cf. Section 6) and discuss important deployment considerations (cf. Section 7). We further compare our solution to alternative approaches ranging from Caller ID suppression to “burner” phones that are only used a small number of times (cf. Section 8). To the best of our knowledge, our analysis for the first time shows that the current state of the art fails to address all but the simplest adversaries and fails to scale.

Note: We stress that our main contribution lies in protecting users against CDR analysis by routing calls across a network of multiple independent telephony relays. We use smart engineering to design and implement the Phonion architecture, which is a viable alternative to VoIP over established anonymization networks, that can provide better call fidelity while not relying on constant broadband Internet connectivity on the caller or callee side. We, however, neither attempt to replace well-studied low-latency anonymization networks, such as Tor, nor claim stronger anonymity properties. Indeed, our solution leverages Tor *once* during the initial call circuit setup, but *never* routes actual call contents via Tor.

2 System and Security Model

The primary objective of Phonion is to provide call unlinkability. That is, Phonion prevents an adversary from identifying that specific pairs of users communicated with each other via the Phonion architecture.

2.1 Overview

We first start with a simple use case example of how the Phonion network can be used. We briefly describe the components of Phonion network and will present its detailed design and implementation in Sections 3 and 4. Suppose Alice is a police officer, and she has discovered that a few “bad apples” in her department are routinely violating the civil rights of innocent citizens. She wants to inform someone, but she fears that she could lose her job or face severe harassment if she talks to superiors. By serendipity, at a police department fundraiser she meets Bob, a journalist. That night, she pulls Bob aside and asks if she could call him anonymously with

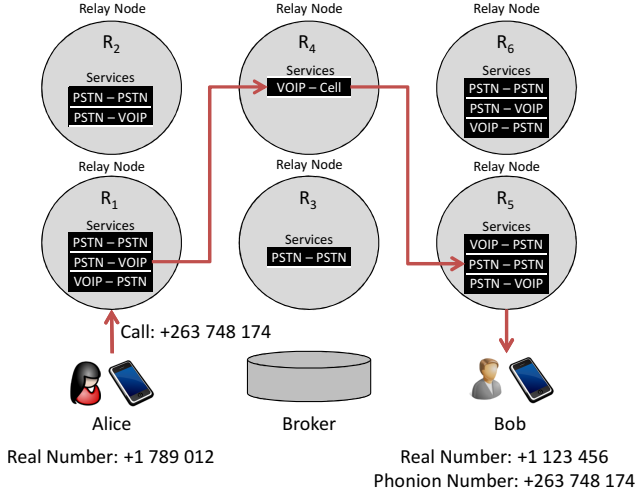


Fig. 1. Phonion Network High-Level Overview

a story. Bob suggests using Phonion to communicate anonymously.

Figure 1 shows a simplified overview of the Phonion network. It consists of a set of users who communicate via a call circuit established across multiple Relay Nodes (denoted N_i in Figure 1). Relay Nodes forward calls from the Caller to the Callee using Relay Services, which connect different telephony technologies (e.g., public switched telephone network (PSTN) to VoIP). The Broker serves as a central directory of available Relay Nodes and Relay Services. In general, any user can setup a Phonion call circuit, but to support our example, we describe the procedure from the perspective of the Callee (Bob).

To establish a call circuit, Bob starts the Phonion client app on his mobile phone. He first reserves a set of Relay Services to establish a call circuit which routes to his office phone, for instance over Relay Nodes R_1 , R_4 and R_5 . This *one-time* reservation phase takes place over secure Internet links using Tor, and is the only time Phonion requires Internet access. The client app uses Tor when Bob sets up his preferred call circuit to provide authentication of Phonion infrastructure components, user anonymization and to preserve data confidentiality. After that in the calling phase no Internet connection is needed. Finally, Relay Node R_1 generates a corresponding Phonion telephony number which routes calls via the established call circuit. He gives this number to Alice and tells her to call anytime the next day. With this number, Alice can call from any phone without the need for Internet connectivity and be confident that her employer will never be able to determine that she called Bob.

A significant advantage of Phonion over purely VoIP based solutions, such as VoIP over Tor, is that Phonion supports a wide variety of telephony endpoints, such as cellphones, smartphones, landline phones, or even a desktop or mobile computing platform using VoIP software. This feature will especially be appreciated by users who are comfortable with standard telephony, but less familiar with VoIP technology.

2.2 Adversary Model

We first specify assumptions on adversary capabilities and then define four distinct adversary classes with increasing capabilities. We will use these classes in Section 6 and Section 8 to elucidate the security of Phonion and alternative solutions.

Assumptions. We assume internal and passive adversaries who can obtain Call Data Records (CDRs) for one or more (but not all) entities of the Phonion network. We further assume that adversaries are able to monitor IP network traffic for one or more (but not all) Phonion infrastructure components as well as to deploy malicious Relay Nodes. In addition to CDRs compromised Relay Nodes provide call circuit setup records (CSRs), which contain metadata about Relay Services reserved by a user. However, adversaries are neither capable of compromising voice communication channels nor of monitoring or actively tampering with call content on the wire. This also includes malware on users’ devices, i.e., we assume that the Phonion client on user’s device is trusted and immune, since if the adversary controls the client no privacy guarantee can be given. Moreover, we assume that adversaries are able to map endpoint phone numbers to user identities – there are many means to infer this information directly or indirectly even for the weakest adversaries.

Finally, we exclude Denial of Service (DoS) attacks on Phonion from our security model. Related work has identified that large-scale DoS attacks can potentially have an impact on the security guarantees of distributed anonymization architectures [10], such as Tor. While Phonion is no exception, addressing such attacks is a problem well beyond the scope of this paper.

Adversary Class 1: Associates. A class 1 adversary is capable of compromising telephony metadata *exclusively* at one of the call endpoints, i.e., Alice’s or Bob’s (see Figure 1). More precisely, this adversary models “associates” who know the caller and callee personally and are interested in the activities of that person. For

example, a nosy neighbor may steal a phone bill to see a user’s call records. Another example is found in the Hewlett-Packard case from the introduction (cf. Section 1): a caller’s employer may want to ensure that she is not talking to regulators, journalists, or competitors. Associates have limited interest and face significant challenges in obtaining call data records.

Adversary Class 2: Communications Providers.

A class 2 adversary is capable of compromising metadata passing through a *single* carrier. This models the capabilities of individual communication providers. These may be mobile, landline, or VoIP carriers, resellers (such as MVNOs¹), or value-added service providers (such as Twilio, see Section 4.2). A class 2 adversary has significant statistical insight into typical customer behavior as well as full access to the calling habits of its subscribers. Furthermore, a class 2 adversary may even have records for calls it routes (but does not originate or terminate). Accordingly, we conservatively assume that class 2 adversaries can completely compromise Relay Nodes which establish or route calls via a compromised carrier network². Rogue employees of communications providers may take an unethical interest in the activities of certain customers. Examples include the activities of public figures like celebrities, politicians, and heads of multinational corporations.

Adversary Class 3: Law Enforcement Agencies.

A class 3 adversary is capable of compromising metadata at *multiple* (but not all) carriers (i.e. multiple Relay Services at different Relay Nodes in Figure 1), though we assume that not all carriers are compromised. This models the legitimate behaviors of law enforcement agencies (LEAs). Law enforcement agencies typically only invoke CDR analysis *after* a suspicion of unlawful activity. Upon that suspicion, they will pursue records of a particular caller or callee and their associates. Law enforcement agencies have considerable reach, although that reach is constrained by jurisdiction. LEAs in one country face greater hurdles to obtaining data outside of their jurisdiction.

Adversary Class 4: Intelligence Agencies. A class 4 adversary is capable of compromising metadata at *all* carriers, allowing the adversary to reconstruct a global

view of all concurrent calls. Our fourth class models the capabilities of the most capable and sophisticated intelligence agencies, which routinely engage in bulk metadata collection from domestic and international carriers [24]. We conservatively assume that these agencies, explicitly given license for clandestine activities, are unbound by jurisdictional concerns or legal limits to surveillance. Accordingly, intelligence-class adversaries have the greatest visibility and reach — spanning the customers of many carriers in many countries.

Note that Phonion cannot provide any security guarantees in the presence of global class 4 adversaries who can compromise *all* carriers within the Phonion network. However, this limitation is not specific to Phonion and no other anonymization system can protect against attacks of a class 4 adversary. Nevertheless, Phonion and other advanced anonymization networks, such as Tor, significantly increase costs of such attacks and, hence, frustrate metadata collection and analysis even for such powerful adversaries. Further, our adversary model generally allows powerful adversaries to identify which users are establishing or receiving calls via the Phonion network. Similar to established anonymization networks, such as Tor [21], providing unobservability is generally out of scope of our work.

2.3 Requirements

Security Requirements. Given our main goal and adversary model we formulate the following security requirements for Phonion.

1. **Call unlinkability.** We require that Phonion prevents class 1–3 adversaries from learning that a caller and callee communicated.
2. **Phone number obfuscation.** We require that Phonion prevents class 1–3 adversaries from linking a user’s real phone number to his Phonion number.

Note that we only address metadata analysis and treat voice channel inspection as an orthogonal problem beyond the scope of this paper, because confidentiality of phone calls can be achieved by using orthogonal solutions like encrypting headsets [1] or spoken-word codes. Further, while there are many cases where call audio privacy is necessary, there are two important reasons why protecting call metadata is actually a more critical concern.

First, call records are more easily obtained by any party than real-time audio. For example, in the United States, law enforcement can obtain call records with-

¹ A Mobile Virtual Network Operator is a mobile telephone carrier that does not own its own infrastructure (e.g., towers) and instead leases it from another carrier. Well-known examples are LycaMobile (worldwide) or MetroPCS (USA).

² Note that individual Relay Nodes may offer multiple Relay Services associated with different carriers.

out a court order [42], while those agencies must meet a high burden of proof to obtain court permission for a voice tap. Both nosy spouses and employees of telephony providers have limited ability to actively intercept voice data, but both can easily gain access to telephony bills and the accompanying call history. Further, some law enforcement agencies have begun to compile large internal databases of call metadata [55] that are unencumbered by any oversight. These records could similarly be used to harass whistleblowers, protesters and activists, or to blackmail individuals. Second, even if call audio is encrypted or unavailable for active surveillance, the fact that two parties communicated is sensitive enough that it bears protection. For example, making calls to a known dissident in an oppressive regime could make one a target for more scrutiny, regardless of call content.

Functional Requirements. We specify the following additional functional requirements for Phonion.

1. **Offline calls.** We require that Phonion users can make and receive Phonion calls while they are not connected to the Internet.
2. **Quality of service.** We require that Phonion provides sufficient quality of service. In particular, Phonion must provide acceptable audio latency and minimize voice quality degradation.
3. **Compatibility with legacy telephony devices.** We require that users can make and receive Phonion calls using different types of telephony devices, including cell phones, smartphones, feature phones, and landline phones.

We emphasize that no alternative solution which can be used for anonymous calls can fulfill these requirements (cf. Section 8 for comparison).

3 Phonion Detailed Design

3.1 Phonion Network

As mentioned in Section 2.1 the Phonion network consists of the following components: **Broker**, **Relay Nodes** and **Relay Services**. Furthermore, Phonion users, whom we denote as **Callers** and **Callees**, use **Clients** to interact with the Phonion network and **Phones** to make or receive calls. In the following, we describe the Phonion components in more detail.

Callers and Callees. Callers and Callees are Phonion users who establish outbound and receive inbound Phonion calls, respectively. To evade metadata analysis, ei-

ther Caller or Callee (or both) generate and use a Phonion phone number.

Clients. A Client is a piece of software which is used to obtain a Phonion phone number and to setup a Phonion call circuit upon request of the user. It is executed on a user-controlled platform, such as a PC, laptop or a smartphone, and serves as an interface between the Phonion user and the Phonion network.

Relay Services. Relay Services perform the actual call forwarding in Phonion. A Relay Service connects an incoming call to a Relay Service number to a new outgoing call via another number. It also serves as a gateway between different types of networks, such as public switched telephone networks (PSTNs), cellular and VoIP networks. Each Relay Service has an associated capacity.

To defend against timing attacks Phonion Relay Services align events to time slots: When a call is placed, each Relay Service introduces a “guard time” to enforce that the Caller waits an amount of time corresponding to when the next “virtual timeslot” begins. Similarly, when one party ends a call by hanging up, the Relay Service on the other end of the call circuit asks the other party to remain on the line until the timeslot ends. This strategy effectively emulates a fixed time slot regime and is designed to protect against CDR analysis based on beginning and end of calls, as we discuss in Section 6.

As a further protection against analysis of CDRs based on the duration of calls we propose that Relay Services introduce asymmetry into call duration. In particular, when one party hangs up, the other party is asked by the last Relay Service on the call circuit to remain longer on the line, and not just until the time slot ends, but even for one or more additional time slots longer. Depending on the Relay Service technology some Relay Services further keep intermediate connections alive for a random period of time to make analysis even harder.

Relay Nodes. Relay Nodes are responsible for the setup of Phonion call circuits. They are associated with (a number of) Relay Services which forward calls between networks of similar or different types (e.g., PSTN-to-VoIP, or PSTN-to-PSTN)³. Upon request of the Client via a secure Internet link, the Relay Node selects a suitable and available Relay Service for the Phonion call circuit (e.g., the one which serves as a gateway between specified types of networks) and configures it to forward

³ Relay Nodes do not necessarily reside on the same platforms as Relay Services, but generally joint deployment is also possible.

the received call to the next Relay Service or to a final destination.

Broker. The Broker serves as a public directory which maintains information about the Phonion network. It describes where Relay Nodes are hosted and which types of Relay Services they provide. For instance, it lists in which jurisdiction the Relay Nodes and their currently available Relay Services are located and between which network types these Relay Services can forward calls.

3.2 Network management

Whenever a new Relay Node enters the network, it registers with the Broker and provides its network address, area of jurisdiction (e.g., a country) and currently available Relay Services. At runtime, Relay Nodes indicate to the Broker whether specific Relay Service have capacity available to forward calls available or not.

3.3 Call circuit establishment

Figure 2 shows the three-step process used by a Callee to establish an inbound call circuit. To support our example from Section 2, we describe the procedure from the perspective of the Callee (“Bob”). However, if the Caller (“Alice”) knows Bob’s number, it is also possible for her to setup an outbound call circuit to his number.

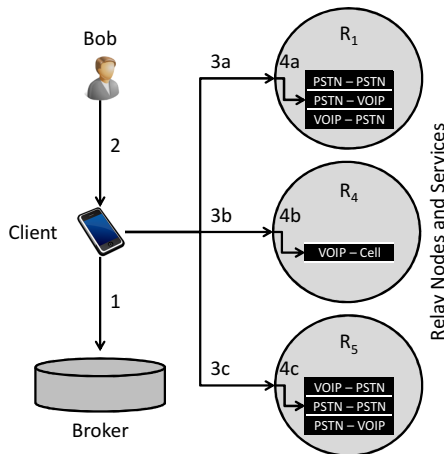


Fig. 2. Call circuit establishment. The Phonion Client queries Broker to locate available Relay Nodes and their Relay Services. It proceeds to reserve Relay Services of the selected Relay Nodes R_1 , R_4 and R_5 to establish a call circuit.

To establish a call circuit, Bob uses his Client software (e.g., a mobile app), which upon launch queries the

Broker for a list of currently available Relay Nodes and their Relay Services (step 1). The list is then displayed to Bob, who selects Relay Nodes and Relay Services for the call circuit and indicates for how long he would like to reserve them (step 2). Here, Bob can choose the number of Relay Nodes and their locations depending on his individual obfuscation requirements. For instance, he could opt for multiple Relay Nodes and Relay Services in different jurisdictions in order to tolerate more powerful adversaries. In contrast, if Bob considers only adversaries with limited capabilities, he would most likely opt for fewer relays and would choose them in a domestic area, which would most likely result in a better voice quality and less costs. In Figure 2 three relays are selected for the call circuit, denoted as R_1 , R_4 , and R_5 .

After receiving input from Bob, the Client contacts the selected Relay Nodes and sends them forwarding rules (steps 3a, 3b, and 3c). These Relay Nodes configure their corresponding Relay Services (steps 4a, 4b, and 4c) to forward calls between them, thereby constructing the final call circuit. The phone number of the first Relay Service in the call circuit (R_1) is the Phonion number to be shared with Alice. Note that by having a user set up the call circuit over Tor node-by-node, Phonion ensures that only the user (and his or her trusted Client) knows the full call circuit, while individual Relay Nodes or Relay Services only know the next and previous hop.

3.4 Placing a Phonion call

Once Bob obtains a Phonion number, he shares it with Alice to receive calls from her. The number is valid for a limited period of time, as requested by Bob during call circuit establishment. When Alice dials the Phonion number, the Relay Services of Relay Nodes R_1 , R_4 , and R_5 forward the call to Bob’s real phone number, potentially over different telephony technologies and over networks located in different jurisdictions.

4 Implementation

We implemented a prototype of our design to evaluate the practicality and performance of a Phonion network. This section discusses implementation details of the Broker, Relay Nodes, Relay Services, and the user Client. We pay special attention to the practical issues of developing Relay Services and our user application.

Initialization. As noted in Section 2, in our current implementation **Clients**, **Relay Nodes** and the **Broker** communicate over Tor [21] using hidden services, which provide resistance to takedown attempts, confidential communication as well as authentication of **Relay Nodes** and the **Broker**. Most importantly, Tor obscures the user identity towards **Relay Nodes** and the **Broker**. Hence, as initialization steps, **Relay Nodes** and the **Broker** setup their services over the Tor network as Tor hidden services, which makes them visible within the Tor network, but does not expose them on the public Internet.

4.1 Broker

The Phonion **Broker** is currently implemented in the Python programming language and based on the Flask web development framework [52]. It is meant to be deployed on a server permanently connected to the Internet. Whenever a **Relay Node** registers its services with the **Broker**, it first sets a password, which allows it to later authenticate itself towards the **Broker**. This simple authentication scheme prohibits an adversary from modifying individual **Relay Nodes**' information, such as currently available capacity, or from deregistering the **Relay Node**.

To achieve high availability of the **Broker** and immutability to password-related attacks, we are currently working on an alternative implementation using blockchain technology. In particular, we are developing the **Broker** as a smart contract on the Ethereum [12] blockchain using the Solidity language [25]. The contract can receive *registration*, *de-registration* and *change configuration* requests from **Relay Nodes** in the form of transactions, which are used to trigger addition, removal or change of records about **Relay Nodes** in a smart contract database. Notably, *change configuration* and *de-registration* requests are accepted by the smart contract only if they arrive from the same origin from which the *registration* request was previously received, thus preventing unauthorized removal or modification of **Relay Node** information in the database. To obtain information about current topology of the Phonion network, **Clients** query a smart contract which responds with data from its database. Note that querying a smart contract – an operation which may happen frequently – does not impose transaction fees, while registration, de-registration and configuration change operations are performed only by **Relay Nodes** (not by **Clients**) and take place only when network topology changes.

4.2 Relay Services

Our implementation currently supports three **Relay Service** variants, which use Twilio, Google Voice, and an Asterisk PBX to construct a call circuit.

Twilio. Twilio [63] is a cloud-based telephony service that offers rental of phone numbers via a HTTPS API and allows the dynamic forwarding of calls to other numbers. Our implementation uses this feature to receive a callback when a call has been received on one of the Twilio phone numbers offered by the **Relay Node**. This callback is issued by Twilio via HTTPS to the corresponding **Relay Node** and contains the Twilio phone number. The **Relay Node** looks up the forwarding rule for this number, and if the number has been reserved by a Phonion **Callee** it forwards the call according to the corresponding forwarding rule.

Google Voice. Google Voice [30] is a cloud-based telephony service which allows users to register a telephone number. Using a web interface, users can configure call forwarding rules for this number. Our Google Voice **Relay Service** interacts with the Google Voice homepage via Selenium WebDriver [57] to establish two outbound calls — for example, one to the **Callee** and another one to the next **Relay Service**— and connects these calls internally. Google Voice also requires user interaction (in the form of a verification call or SMS) to establish forwarding rules.

Asterisk. Asterisk [20] is an open-source PBX software that supports a number of different telecommunication technologies including VoIP, landline, and cellular networks. We use Asterisk **Relay Services** to connect incoming VoIP/SIP calls to landline and cellular numbers. To this end, the **Relay Node** generates a random SIP address for each **Callee** during the reservation process. Our Asterisk **Relay Service** forwards all incoming calls to this number according to the forwarding rules supplied by the **Callee's Client**. Our implementation supports commercial VoIP trunks as well as GSM Voice Modems to establish outbound calls. However, Asterisk's modular architecture allows us to integrate additional telephony technologies, such as ISDN primary rate interfaces, merely by changing the Asterisk configuration.

4.3 Relay Nodes

As one of the primary functions, **Relay Nodes** serve to abstract away complexity of diverse **Relay Services** from **Clients** – they provide a unified communication interface

to **Clients** and integrate various **Relay Services** via an extensible plugin architecture to account for their different properties. Our **Relay Nodes** communicate with Asterisk, Twilio and Google Voice **Relay Services** via HTTPS. Registration, de-registration and configuration changes are signalled to the **Broker** via secure Internet links provided by the Tor software.

Similar to Tor relays we assume **Relay Nodes** to be operated by volunteers. To encourage adoption of Phonion, **Relay Nodes** may optionally demand a fee from the user for the use of their **Relay Services**. These fees might be required to cover charges by telephony carriers, and to provide additional incentives to deploy and operate the Phonion infrastructure (cf. Section 7). Our current implementation uses Bitcoin payments [43], which allow users to obfuscate their identity and which have similarly been proposed for the reimbursement of Tor relay operators [9]. However, Phonion is agnostic to any particular payment scheme and other suitable solutions, such as more privacy-focused digital currencies [8] or even anonymous pre-paid credit card payments, can be integrated as well. We further plan to additionally support payments based on the Ether cryptocurrency used by Ethereum smart contracts [12] in order to match our Ethereum-based design of the **Broker**.

4.4 Phonion Client App

Although Phonion is platform agnostic (and capable of even connecting calls from rotary phones), we implemented our Phonion Client as an Android app. Our current implementation establishes secure communication channels with the Phonion network components via the Orbot Tor software [60]. Figure 6 in Appendix A shows screenshots of our app.

To acquire a Phonion number, the **Client** app first queries the **Broker** for available **Relay Nodes** and their currently available **Relay Services**. The list of available **Relay Nodes** is displayed to the user, who inputs the **Callee's** landline or mobile phone number and selects the number of **Relay Nodes** according to his or her individual privacy requirements. He or she then either selects individual **Relay Nodes** or lets the software choose a call circuit. The latter option simplifies choosing a secure call circuit by automatically selecting compatible **Relay Nodes** and **Relay Services** spread across diverse jurisdictions. Section 6 discusses the security and privacy issues inherent to call circuit selection in more detail.

Once the **Relay Services** have been reserved, the **Client** app arranges the call circuit by sending forward-

ing rules to the selected **Relay Nodes**. It daisy-chains the individual **Relay Services** so that they forward calls between them towards the **Callee's** real phone number (see Figure 1). Finally, the **Client** displays the Phonion number to the user, which is the phone number of the first **Relay Service** of the call circuit. Calls to this number will be forwarded to the real number of the **Callee**.

5 Evaluation

In this section, we evaluate the impact of Phonion on audio quality during calls. We first describe our experiment setup and then present and interpret the results. Further, we evaluate the latency and audio quality of VoIP over Tor and provide a comparison. Finally we briefly discuss usability aspects of the Phonion client.

5.1 Call Quality Evaluation

Our experiment setup for the evaluation of Phonion consists of two Android smartphones and a measurement PC. The PC uses the Android Debug Bridge [5] to control Phonion calls between the two smartphones. It further feeds an audio signal into the first smartphone via the headset jack, which is then recorded via the second smartphone's headphone jack. Each call lasts 100 seconds, which is the average length of a phone call [61], and we performed 100 calls for each Phonion call circuit.

As we discuss in Section 8, tunneling Voice over IP (VoIP) through a low-latency anonymization network, such as Tor, seems to be a reasonable alternative approach to provide anonymous phone calls. To compare our voice quality results with the audio quality achievable by using VoIP technology over Tor we establish 100 calls between two PCs using the Mumble VoIP software [62] via a Mumble server we deployed on a virtual machine. The Mumble server is contacted by the client PCs through Tor socks proxy services running on each client. The measurement PC remotely controls the Mumble VoIP clients to establish and teardown calls and restarts the Tor proxy after each call. It further feeds the audio signal into the first PC's audio interface and records the received signal from the audio interface of the second PC.

To determine the voice quality degradation we use ITU-T standard P.862 (PESQ) [32], an industry standard algorithm for measuring voice quality in modern telephony networks. Traditionally the audio quality of

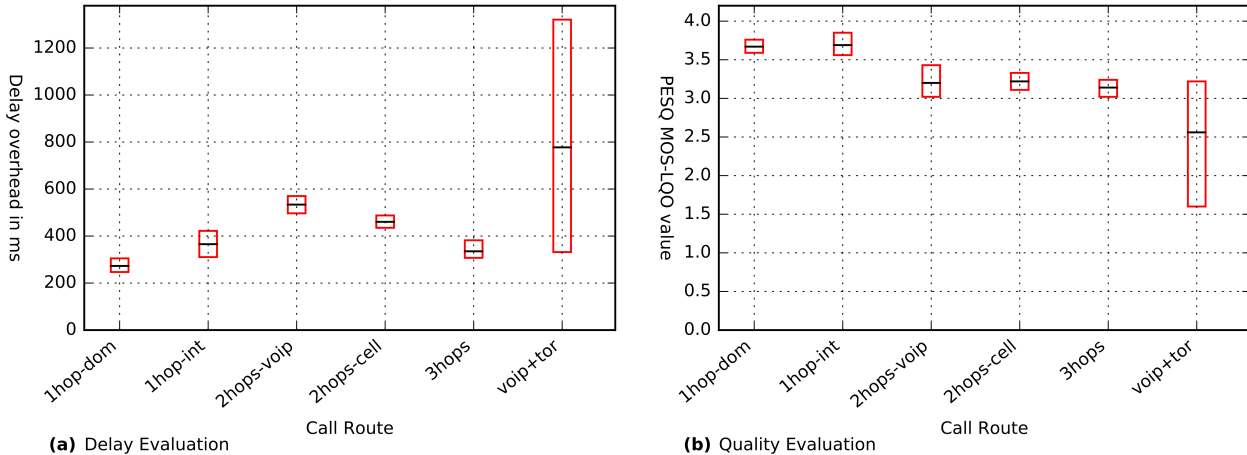


Fig. 3. Evaluation of different call routes over Phonion as well as VoIP using Mumble and Tor. The plots show the average (in black) and 10th to 90th percentiles (in red) for each experiment.

telephony calls has been derived from subjective evaluations performed by humans. In contrast, the PESQ algorithm estimates the quality of recorded voice calls in narrow-band telephony applications by modeling the human perception of audio signals and analyzing common distortions introduced by telephony technology. It calculates the *Mean Opinion Score - Listening Quality Objective* (MOS-LQO) value for recorded audio samples by comparing them to the corresponding reference signals. The MOS-LQO value ranges from 1 (bad audio quality) to 5 (excellent audio quality). A score of 3 is considered “fair”, while a score of 2 is considered “poor” and very annoying. We further measure the latency introduced by Phonion as well as VoIP over Tor by cross-correlating the sent and received audio signals of each call.

5.1.1 Results

The results of our experiments using the call circuits described in Table 1 are presented in Figure 3.

Phonion. Our measurements for Phonion show that latency overhead varies significantly between calls depending on the geographical location of the **Relay Services** and the technologies used on the call circuit (cf. Figure 3a), while the number of **Relay Services** does not necessarily increase latency. For instance, our measurements indicate that calls via VoIP and cellular **Relay Services** introduce more delay than calls via PSTN networks. In additional experiments we found that without even using Phonion, local cellular calls introduce an average mouth-to-ear delay of 385 ms. This aspect also ex-

plains why the longest route in our experiments (*3hops*) introduces less latency than both *2hops* routes: In the *3hops* route only PSTN **Relay Services** are used, whereas the *2hops* routes involve VoIP and cellular telephony.

Table 1. Phonion call circuits used in our evaluation

Call Circuit	Relay Nodes
1hop-dom	National PSTN - PSTN
1hop-int	International PSTN - PSTN
2hops-voip	National PSTN - VoIP, International VoIP - PSTN
2hops-cell	International PSTN - VoIP, International VoIP - Cellular
3hops	National PSTN - PSTN, International PSTN - PSTN, International PSTN - PSTN

Furthermore, we observe that the geographic location of **Relay Services** has an influence on the latency overhead, and it even seems that that geolocations of selected **Relay Services** have a more significant impact than the number of **Relay Services** in the call circuit. When comparing call circuit *1hop-dom*, which uses a relay located in a domestic area, with the call circuit *1hop-int* using an international relay, we see that the former introduces only 273 ms of delay, whereas the latter introduces 365 ms on average. This aspect can at least partially be attributed to differences in the internal telephony network structures of different carriers.

Remarkably, our MOS-LQO measurements indicate that audio quality in Phonion remains stable over the course of multiple calls, as shown in Figure 3b. Surprisingly, unlike latency overhead, audio quality degrada-

tion does not depend on the geolocation of **Relay Nodes**, as one can see when comparing call circuits *1hop-dom* and *1hop-int*. In contrast, the technologies used along the call circuit have significant impact on the MOS-LQO values. The difference stems from the fact that different telephony architectures use different audio codecs for voice data transmission. For example, while the G.711 codec used in high-quality ISDN links and in modern VoIP systems can generally achieve scores higher than 4, the GSM Full Rate voice codec only achieves a maximum score of 3.5 [17, 48]. Further, it is likely that the conversion between different voice codecs has a negative impact on voice quality.

VoIP over Tor. Our latency results for Mumble over Tor show a rather high delay overhead of 777 ms. Moreover, latency varies significantly between calls. The average PESQ value for calls using Mumble over Tor is 2.5, indicating “fair” to “poor” audio quality. Call quality also varies significantly between calls, as shown in Figure 3b. In a separate experiment we established that Mumble was able to achieve a relatively high MOS-LQO value of 3.6 without using Tor. Overall, these results confirm observations made by Rizal [51], who reported a similar and significant quality impairment affecting VoIP over Tor (cf. Section 8).

5.1.2 Comparison

Our results show that VoIP over Tor introduces more delay compared to Phonion. Further, voice quality degradation is significantly lower in the Phonion network compared to VoIP over Tor, even in the presence of multiple audio codec conversions. For example, even when Phonion uses VoIP (in this case G.711) and low-quality GSM codecs, Phonion can provide adequate call quality, in contrast to VoIP over Tor, which achieves only “poor” to “fair” voice quality.

Note that the recently published Herd architecture [38], which implements a low-latency anonymization network for VoIP, aims to address Tor’s deficiencies regarding telephony applications. Herd shows promising latency and audio quality results during an initial evaluation within an idealized environment where all relays are hosted within Amazon EC2 data centers. The Herd authors, however, used a fundamentally different experiment setup to evaluate their approach: First, the E-Model algorithm [31] used by the authors to evaluate voice quality degradation cannot be applied to Phonion. E-Model merely *estimates* call fidelity based on a

model of the telephony network. But this model does not take different telephony technologies within a call circuit into account. The PESQ algorithm we selected for our evaluation instead compares actual call audio recordings on both the caller- and callee side. Second, it is unclear whether or not the actual mouth-to-ear delay was considered during the latency evaluation of Herd. We contacted the Herd authors, who kindly supplied a version of their implementation, and we are currently working on an extended evaluation environment which will allow us to apply our experiments to Herd.

5.2 Usability Evaluation

The user interface of the Phonion Client is the result of an initial usability study with students and staff members performing user interface walk-through experiments. We incorporated study results into our implementation, which led to the clear step-by-step process to setup Phonion numbers shown in Appendix A. Further, hands-on demonstrations at IT exhibitions revealed not only considerable public interest in our architecture, but also that users were able to operate the Phonion Client without any significant guidance.

6 Privacy Guarantees

In this section, we characterize the anonymity provided by Phonion under the adversary model defined in Section 2. As noted in Section 2.3 the main goals of Phonion are to prevent an adversary from 1) linking pairs of **Callers** and **Callees** who communicate via Phonion and 2) linking a user’s real phone number to his Phonion number. Since Phonion uses out-of-band signaling we consider both call data records (CDRs) as well as call circuit setup records (CSRs).

Call data records are generated by telephony carriers for each call they route across their network. A call data record $CDR = \{t_s, t_e, n_{caller}, n_{callee}\}$ contains start time t_s and end time t_e of a particular call as well as the participants’ phone numbers n_{caller} and n_{callee} . If an adversary is able to *compromise* a **Relay Service** he can obtain CDRs for this particular **Relay Service**.

In addition, as stated in Section 2.2 we assume that class 2, 3 and 4 adversaries (telecommunication carriers, law enforcement agencies (LEAs) and intelligence agencies) are capable of infiltrating the Phonion network by operating malicious **Relay Nodes**. Such a *com-*

promised Relay Node not only provides CDRs for calls routed via its Relay Services to the adversary, but also call circuit setup records generated whenever users reserve one of its Relay Services. A call circuit setup record $CSR = \{t_s, t_e, IP_s, n_{relayservice}, n_{nexthop}\}$ contains start time t_s and end time t_e of a Relay Service reservation, the source IP address used during call circuit setup IP_s , the Relay Service number $n_{relayservice}$ and the number of the next hop on the call circuit $n_{nexthop}$, which is either the Callee's or another Relay Service's phone number.

To analyze metadata attacks we represent the Phonion network as a graph. In this graph, an edge connecting two compromised Relay Services represents either a single CDR or CSR. An edge between two uncompromised Relay Nodes may represent many multiplexed calls or reservations over a single carrier or different carriers (e.g., one call over VoIP and one over PSTN).

The adversary can now use his available CDRs and CSRs obtained from compromised Relay Services and Relay Nodes to construct a graph that represents its view of the Phonion network during a specified time. Intuitively, the adversary accomplishes this by assuming all parties *could* be connected, and then removing connections that are proven impossible based on his prior information.

6.1 Trivial Case: Class 4 Adversaries

A class 4 adversary models the capabilities of globally operating intelligence agencies and thus is assumed to have a global view of the Phonion network, since he is able to compromise all Relay Nodes and Relay Services. Accordingly, he can link any caller and callee pair at any given time by ruling out edges based on his global view of the Phonion network. This is illustrated in Figure 4a, where the adversary can clearly trace connections between users across all Relay Services of Relay Nodes R_1 , R_2 , R_3 and R_4 . Obviously he can also trivially link users to their Phonion numbers by tracing all call circuits to their last Relay Services.

6.2 Class 1 Adversaries

A single class 1 adversary, such as a spouse inspecting a Phonion user's telephony bill, does not have access to the required metadata to link a user's real phone number and his Phonion number. This is because a Relay Service never uses the user's Phonion number as an outgoing caller ID. Instead, the Relay Service picks a number randomly from the pool of registered numbers available

to it, which prevents the Phonion number from ever appearing in a user's call logs or telephony bills. Linking individual Callers and Callees is not possible without additional metadata, as we elaborate in the following sections. Accordingly, a class 1 adversary is never able to link users communicating via Phonion.

6.3 Class 2 Adversaries

Class 2 adversaries (telecommunication carriers) possess CDRs and CSRs for calls they route. In our model, they can accordingly only compromise individual Relay Nodes and (a subset of) their Relay Services. This is shown in Figure 4b, where the adversary possesses CDRs and CSRs for Relay Node R_1 and all its Relay Services. Class 2 adversaries cannot assume that edges connecting uncompromised Relay Nodes (here R_2 , R_3 and R_4) and their Relay Services do not exist. Thus, they cannot reconstruct complete call circuits, which makes it impossible for them to link pairs of communicating Phonion users by tracing calls within the Phonion network and prevents them from linking users' real phone numbers to their Phonion numbers.

Timing Attacks. If the adversary has start and end times of calls or corresponding CSRs of two communicating users, he may use this information to link pairs of communicating Phonion users as well as real numbers and Phonion numbers without tracing them through the Phonion network. Such "end-to-end" timing attacks are applicable to many anonymity systems, including Tor [21], and Phonion is no exception. Furthermore, the adversary can use heuristic information to reduce the anonymity set of users across multiple Phonion calls. Note that class 2 adversaries only have sufficient capabilities to perform such attacks if they route both the Caller's outgoing and Callee's incoming call. This is for example the case if both Caller and Callee use the same telephony provider.

Let N be the number of Phonion users at a given time period $t = [t_1, t_2]$. The adversary aims to link a pair of users $[u_1, u_2]$ which communicate using Phonion within t . Let q denote the probability of any pair of Phonion users establishing a call via two Relay Services of Relay Nodes an adversary is able to compromise (i.e., obtain CDRs or CSRs for) within t . We now calculate the adversary's chance of linking $[u_1, u_2]$ when they communicated via Phonion in t . If i is the *exact* number of additional user pairs communicating via the compromised Relay Services in t , the size of the anonymity set for the user pair $[u_1, u_2]$ is $|\Omega_{[u_1, u_2]}| = \binom{2i+2}{2}$. The

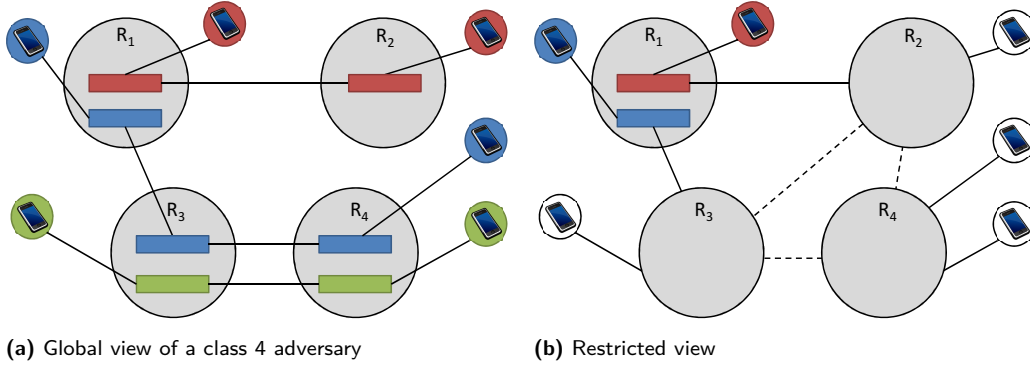


Fig. 4. Adversary views of the Phonion network. In (a), the adversary has compromised all Relay Nodes and Relay Services. As such, all connections between Relay Nodes (grey) and their Relay Services (colored boxes) are absolute, and thus all users can be linked. In contrast, in (b) the adversary, who compromised R_1 , must assume connections exist between Relay Nodes R_2 , R_3 and R_4 , which are depicted as dashed lines.

probability of *exactly* i additional user pairs communicating in t via the compromised Relay Services is $\binom{\lfloor \frac{N-2}{2} \rfloor}{i} q^i (1-q)^{\lfloor \frac{N-2}{2} \rfloor - i}$. Thus, the probability of linking the pair of interest $P_l([u_1, u_2])$ is as follows:

$$\begin{aligned}
 P_l([u_1, u_2]) &= \sum_{i=0}^{\lfloor \frac{N-2}{2} \rfloor} P_l([u_1, u_2], i) \\
 &= \sum_{i=0}^{\lfloor \frac{N-2}{2} \rfloor} P_l([u_1, u_2] | i) P_l(i) \\
 &= \sum_{i=0}^{\lfloor \frac{N-2}{2} \rfloor} \frac{1}{|\Omega_{[u_1, u_2]}|} \binom{\lfloor \frac{N-2}{2} \rfloor}{i} q^i (1-q)^{\lfloor \frac{N-2}{2} \rfloor - i} \\
 &= \sum_{i=0}^{\lfloor \frac{N-2}{2} \rfloor} \frac{1}{\binom{2i+2}{2}} \binom{\lfloor \frac{N-2}{2} \rfloor}{i} q^i (1-q)^{\lfloor \frac{N-2}{2} \rfloor - i}
 \end{aligned}$$

Obviously, the anonymity of the system depends on the number of users who are communicating at the period of interest. Thus, if the call occurs when few users are active probability of correctly deanonymizing the communicating pair would increase. Figure 5 depicts the probability of correctly linking the users of interest, $P_l([u_1, u_2])$, for selected values of q and $2 \leq N \leq 1000$ users and shows the intuitive result that probability of linking users of interest for each value of q goes to zero as N grows. If the adversary is able to distinguish between the Callers and Callees, he can further improve his analysis as follows: Assume for the pair of interest, user u_1 was the Caller and u_2 was the Callee. In this case, u_1 must have been talking to one of the $i+1$ Callees. Similarly u_2 must have been engaged with one of the $i+1$ Callers. Thus the anonymity set $|\Omega_{[u_1, u_2]}|$ reduces to $\binom{i+1}{2}$ and the rest of the analysis remains unchanged.

Note that in the above calculation for ease of description, we assumed that q is a constant fixed value across all possible pairs of users. In general q could be a function of time and user dependent. For instance, the probability of a pair engaging in a call could be higher during work hours than late in the night and therefore vary for users from various geographical regions. Additionally, some users might be less often engaged in conversations. Further, q can also be used to model the effect of our virtual timeslot regime discussed in Section 3: When Relay Services align calls to virtual timeslots the value of q would increase, since the decreased time resolution would increase the chance of multiple calls ending at the same time. While the above model can easily be extended to capture such cases, we acknowledge that care must be taken when aligning Phonion calls to virtual timeslots to prevent the adversary from extracting the added noise [34]. We refer to future work for a detailed analysis of this aspect.




Fig. 5. Probability of linking two users u_1 and u_2 $P_l([u_1, u_2])$ for selected values of q and N .

Long Term Intersection Attacks. When two parties in a low-latency anonymity system communicate repeatedly over a period of time, an adversary can determine that those parties are more likely to appear in the network than other pairs of users. These “long-term intersection attacks” are an inescapable reality for all low-latency anonymity systems [19, 27, 40], including Tor and Phonion. A simple mitigation strategy against such attacks is to change the user’s real phone number between multiple calls, while keeping these individual numbers unlinkable to the real user identity. However, as we will discuss in more detail in Section 8, phone numbers unlinked to an established identity (e.g., anonymous prepaid SIM cards) are often difficult to obtain in a legal way. Further, depending on the adversary’s capabilities he might be able to use heuristic information to link multiple phone numbers [54].

6.4 Class 3 Adversaries

Class 3 adversaries (law enforcement agencies) are able to compromise multiple Relay Nodes and Relay Services and thus can obtain CDRs and CSRs concerning multiple different telecommunication carriers. However, we assume they are unable to compromise at least one Relay Node and his Relay Services on any given call circuit. Similar to class 2 adversaries they thus cannot trace calls between two Phonion users: A single uncompromised Relay Node and corresponding Relay Service on a call circuit is sufficient to prevent an adversary from restoring the call circuit by merely tracing calls, since he cannot rule out that additional edges connecting uncompromised Relay Nodes and Relay Services exist. However, class 3 adversaries can use both end to end timing attacks as well as long-term intersection attacks to attempt to link Callers and Callees (cf. Section 6.3).

7 Deployment Considerations

We assume Phonion to be operated primarily by volunteers optionally using our payment scheme to cover their costs. These costs comprise fixed (e.g., monthly) costs and variable rates for calls⁴.

Fixed costs constitute costs for Internet and telephony network connectivity. Phonion Relay Nodes re-

quire constant Internet connectivity. The required bandwidth is however generally very low, since Relay Nodes are only involved in call circuit setup and signaling but do not handle any voice data. Bandwidth requirements for Relay Services vary significantly. For example, locally deployed Asterisk Relay Services using cellular and analog voice modems do not require Internet connectivity at all. In case such relays also offer VoIP services the bandwidth depends on the number of calls and the VoIP codecs in use.

Relay Services which receive and establish calls solely via VoIP only require broadband Internet connectivity, while all Relay Services which establish connections to the PSTN or cellular network require corresponding modems and subscriptions. Cloud-based Relay Services, such as Twilio and Google Voice, offer such subscriptions for a fixed monthly price. Google Voice provides a free landline telephone number for all customers, while Twilio’s pricing depends on the country and telephony technology. For example, in most countries Twilio charges more for cellular numbers than for landline numbers. Our Asterisk-based Relay Services operate similar to SIM Boxes, which terminate VoIP calls and establish outgoing calls via cellular modems. Connectivity to PSTN phone lines can be established via affordable ISDN or analog voice modems.

Calls established via Phonion Relay Services potentially generate costs for Relay Node operators. VoIP calls only cause a rise in Internet traffic, whereas calls involving cellular or PSTN Relay Services generate costs depending on type and location of caller and callee, involved telephony carriers and call duration.

8 Related Work

Related work has scrutinized security and privacy issues of VoIP [37], but our paper was more influenced by work on anonymous Internet communication. Chaum mixes [14] were the forerunners of all anonymity systems, but cannot provide low-latency communication, which is paramount for voice applications. In contrast, onion routing [29], used in Tor [21], provides low-latency anonymity for interactive communications. These systems (and others discussed in Edman and Yener’s survey [23]) focus exclusively on IP communications. Other than early work on techniques for implementing mixes in ISDN [46, 47] and cellular [26] networks in cooperation with the network provider, privacy-preserving communication in classical telephony networks has not received

⁴ We focus on Internet and telephony network connectivity and exclude deployment-specific costs, such as electricity, for now.

much attention. In the following, we will classify the existing approaches and compare the most prominent and user-accessible ones to Phonion in Table 2.

Caller ID Spoofing. A naive attempt at call obfuscation might be to use Caller ID spoofing techniques or number blocking services (e.g., *67 in the United States). While this may protect against the Callee recognizing the Caller’s number, the Caller’s carrier has a complete and accurate record of the call. Accordingly, it offers no privacy guarantees against CDR analysis and is not robust against any of our adversaries.

Conference Calls, Google Voice, and Burner App. A second naive attempt might be to use a conference calling service to act as a simple relay between two parties. In fact, services like Google Voice [30] and the Burner App [3], which provide alias numbers to their users, can establish calls in a manner similar to a conference call; when these calls are established, it is actually Google Voice or Burner that calls both Caller and Callee. This means that these services protect against class 1 adversaries (associates) but no others. The services themselves will require registration and maintain CDRs, in addition to providing a single point of surveillance.

Disposable Phones. A third and very popular approach is to use disposable pre-paid phones (popularly known as “burner phones”). Disposable phones unlinked to an established identity are often difficult to obtain and illegal to possess in many places. Further, even phone numbers of disposable phones can be mapped to true identity of users. For example, an adversary can find the disposable phone in the Caller’s possession [56]. Further, location records of the disposable phone can leak the real identity of the user. Adversaries may even use statistics to identify disposable phones [54]. Accordingly, disposable phones provide no true privacy guarantees. Only if an adversary has absolutely no means for obtaining the true identity of the Caller, then disposable phones can provide unlinkable calls.

VoIP based approaches. Two parties may use encrypted VoIP calls to evade metadata analysis, for example using on the Signal [45] or SilentPhone [59] apps, which use the SRTP protocol [7] for voice encryption and the ZRTP protocol [13] for key agreement. While call audio is encrypted and difficult (but not impossible [65]) to recover, IP addresses still clearly identify the endpoints. Any adversary with VoIP flow records of the Caller or Callee can identify the other endpoint of a conversation. This includes class 1 adversaries (associates) who monitor the local network of Caller or Callee

to gain access to VoIP flow records. While gaining no metadata protection, users also have to accept the inconvenience of only being able to make calls when high-speed cellular Internet connectivity is available, which is not always guaranteed in many places.

While CDR analysis conceptually cannot be addressed using content encryption we acknowledge that the protection of call content confidentiality is a desirable goal for any architecture improving caller and callee privacy. Due to our explicit requirement for offline calls simply integrating VoIP encryption is, however, not a viable option. Nonetheless external solutions, such as encrypting headsets [1], are available and can be used in combination with Phonion. Similarly, cryptographic authentication of end users can be provided by systems like AuthLoop [49].

VoIP calls routed through the Tor [21] low-latency anonymization network can provide protection against adversaries of classes 1-3, similarly to Phonion. When using hidden services for VoIP calls Tor additionally provides call confidentiality, but as we stressed our goal is to frustrate metadata linkage. However, both related work [39, 51] as well as our analysis (cf. Section 5) have shown that Tor cannot reliably provide adequate quality for VoIP communication. Recent advances in low-latency anonymization networks designed specifically for VoIP telephony [38] may to some extent be able to address this concern. Nonetheless, all purely VoIP-based solutions require constant high-quality broadband Internet connectivity and are inoperable with standard PSTN/cellular networks and telephony devices.

We stress that Phonion is not meant to replace Tor, which is well understood and widely used in practice. However, challenges in adopting fundamental architectural changes to *significantly* improve the performance of established low-latency anonymization networks [15] motivate the need for Phonion, which specifically addresses shortcomings of Tor and comparable approaches regarding voice communication.

Tor-based anonymization systems have inspired a number of alternative developments which target security and privacy concerns of VoIP communication. VoIP Session Initiation Protocol messages contain URIs that can be used to unambiguously identify users, which PrivateSIP [36] addressed by encrypting user URIs using public keys of final destination proxies; in later work, the authors combine this technique with Tor to provide complete path anonymity [35]. Pr2-P2PSIP [28] proposes changes to the P2PSIP IETF standard draft to protect users’ locations and patterns of communication from unauthorized disclosure.

Table 2. Comparison of Call Obfuscation Strategies

	CallerID Spoofing	Conference Calls, Google Voice, Burner App	Disposable Phone	Encrypted VoIP	VoIP over Tor	Phonion
Unlinkable for Associates	×	✓	×	×	✓	✓
Unlinkable for Providers	×	×	×	×	✓	✓
Unlinkable for Law Enforcement	×	×	×	×	✓	✓
Unlinkable for Intelligence Agencies	×	×	×	×	×	×
Usable with any carrier worldwide	×	✓	✓	×	×	✓
Supports offline calls	✓	✓	✓	×	×	✓
Supports legacy PSTN phones	✓	✓	✓	×	×	✓
Supports legacy cellular phones	✓	✓	✓	×	×	✓
Supports legacy VoIP clients	×	×	×	✓	✓	✓
Adequate call quality	✓	✓	✓	✓	×	✓

Like VoIP over Tor, all purely VoIP-based approaches require constant broadband Internet connectivity. In contrast, Phonion only requires Internet connectivity during call circuit setup, but not for actual calls. Furthermore, since call routing in Phonion is independent of IP anonymization networks, it can generally achieve a high quality of service (see Section 5).

Several proposals have addressed anonymous routing for VoIP. Danezis et al. develop a P2P architecture that uses a social network to provide anonymous call routing by having “friends” act as VoIP proxies [18]. Similarly, Phonion could shift the functionality of the Brokers to a social network, where Relay Nodes would be operated by a user’s trusted friends. Aguilar Melchor et al. outline traffic analysis-resistant VoIP mix systems that rely on a single central server to provide cover traffic and anonymous routing [4]. In contrast, Phonion achieves its security goals by adopting a hopping scheme where calls are routed over distributed telephony relays.

Even if endpoint identifiers in VoIP protocols, such as session initiation protocol (SIP), are blinded, and IP addresses anonymized, VoIP users can still be linked to calls using traffic analysis attacks based on watermarking [16], conversational dynamics [64], or audio artifacts [6, 50]. Furthermore, prior work has shown that an adversary can derive call content when encrypted VoIP calls use variable bit-rate codecs [65]. These traffic analysis attacks also potentially apply to Phonion call circuits, even when adopting voice encryption. However, we note that Phonion focuses on obfuscating CDRs and not on protecting call content confidentiality, authenticity and integrity. Thus, we consider these attacks to be out of scope of this paper, since they require monitoring and/or manipulation of call content.

9 Conclusion

In this paper, we presented Phonion, a system to provide phone call unlinkability in the face of metadata analysis attacks. Phonion uses a series of relays to obscure the true source and destination of a call. We show that where existing technologies fail to provide meaningful security guarantees or adequate call quality, Phonion provides call unlinkability against several strong adversary classes, including associates, compromised carriers, and powerful law enforcement agencies. Our prototype shows that such systems are practical to build and provide adequate call quality. With Phonion, users have a new tool to protect themselves from the privacy risks of metadata collection and analysis.

Acknowledgments

We would like to thank Prof. Stefan Katzenbeisser, Emiliano De Cristofaro, and Micah Sherr, who provided feedback that improved the quality of this manuscript.

This work was supported in part by the German Science Foundation (project S2, CRC 1119 CROSSING), the European Union’s Seventh Framework Programme (609611, PRACTICE), and the German Federal Ministry of Education and Research within CRISP.

This work was also supported in part by the US National Science Foundation under grant numbers CNS-1318167 and CNS-1464088. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- [1] JackPair: secure your voice phone calls against wiretapping. https://www.kickstarter.com/projects/620001568/jackpair-safeguard-your-phone-conversation/video_share.
- [2] C. Action and the National Consumers League. Protect your phone records. <http://www.consumer-action.org/downloads/english/Pretexting.pdf>, 2007.
- [3] Ad Hoc Labs. Burner: Free phone number, temporary disposable numbers. <http://www.burnerapp.com/>.
- [4] C. Aguilar Melchor, Y. Deswarte, and J. Iguchi-Cartigny. Closed-circuit unobservable Voice over IP. In *Annual Computer Security Applications Conference*. IEEE, 2007.
- [5] Android Developers Documentation. Android Debug Bridge. <http://developer.android.com/tools/help/adb.html>.
- [6] V. A. Balasubramaniyan, A. Poonawalla, M. Ahamad, M. T. Hunter, and P. Traynor. PinDr0P: using single-ended audio features to determine call provenance. In *ACM Conference on Computer and Communications Security*, 2010.
- [7] M. Baugher. The Secure Real-time Transport Protocol (SRTP). IETF RFC 3711, Mar. 2013. <https://rfc-editor.org/rfc/rfc3711.txt>.
- [8] E. Ben Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza. Zerocash: Decentralized anonymous payments from Bitcoin. In *2014 IEEE Symposium on Security and Privacy (SP)*, 2014.
- [9] A. Biryukov and I. Pustogarov. Proof-of-work as anonymous micropayment: Rewarding a Tor relay. In *Financial Cryptography and Data Security 2015*, 2015.
- [10] N. Borisov, G. Danezis, P. Mittal, and P. Tabriz. Denial of service or denial of security? In *ACM Conference on Computer and Communications Security*, 2007.
- [11] M. Bowman. Employers can snoop through your cell phone. <http://blogs.lawyers.com/2013/01/employers-snoop-through-cell-phone/>, 2013.
- [12] V. Buterin. A next-generation smart contract and decentralized application platform, 2014. <https://github.com/ethereum/wiki/wiki/White-Paper>.
- [13] J. Callas, A. Johnston, and P. Zimmermann. ZRTP: Media path key agreement for unicast secure RTP. IETF RFC 6189, Oct. 2015. <https://rfc-editor.org/rfc/rfc6189.txt>.
- [14] D. L. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2), 1981.
- [15] C. Chen, D. E. Asoni, D. Barrera, G. Danezis, and A. Perrig. HORNET: high-speed onion routing at the network layer. In *ACM Conference on Computer and Communications Security*, 2015.
- [16] S. Chen, X. Wang, and S. Jajodia. On the anonymity and traceability of peer-to-peer VoIP calls. *IEEE Network*, 20(5), 2006.
- [17] Cisco Systems. Understanding codecs: Complexity, hardware support, MOS, and negotiation. <http://www.cisco.com/c/en/us/support/docs/voice/h323/14069-codec-complexity.html#mos>.
- [18] G. Danezis, C. Diaz, C. Troncoso, and B. Laurie. Drac: An architecture for anonymous low-volume communications. In *Privacy Enhancing Technologies*, 2010.
- [19] G. Danezis and A. Serjantov. Statistical disclosure or intersection attacks on anonymity systems. In *International Workshop on Information Hiding*, 2004.
- [20] Digium Inc. Asterisk. <http://www.asterisk.org>.
- [21] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The second-generation onion router. In *USENIX Security Symposium*, 2004.
- [22] B. Doherty. Vodafone australia admits hacking Fairfax journalist's phone. <http://www.theguardian.com/business/2015/sep/13/vodafone-australia-admits-hacking-fairfax-journalists-phone>, 2015.
- [23] M. Edman and B. Yener. On anonymity in an electronic society: A survey of anonymous communication systems. *ACM Computing Survey*, 42(1), Dec. 2009.
- [24] Electronic Frontier Foundation. NSA spying. <https://www.eff.org/nsa-spying>.
- [25] Ethereum. Solidity 0.2.0 documentation. <http://solidity.readthedocs.io/en/latest/index.html>.
- [26] H. Federrath, A. Jerichow, D. Kesdogan, and A. Pfitzmann. Security in public mobile communication networks. In *IFIP TC6 International Workshop on Personal Wireless Communications*, 1995.
- [27] J. Feigenbaum, A. Johnson, and P. Syverson. Probabilistic analysis of onion routing in a black-box model. *ACM Transactions on Information and Systems Security*, 15(3), 2012.
- [28] A. Fessi, N. Evans, H. Niedermayer, and R. Holz. Pr2-P2PSIP: privacy preserving P2P signaling for VoIP and IM. In *Principles, Systems and Applications of IP Telecommunications*, 2010.
- [29] D. M. Goldschlag, M. G. Reed, and P. F. Syverson. Hiding routing information. In *International Workshop on Information Hiding*, 1996.
- [30] Google Inc. Google Voice. <https://www.google.com/voice>.
- [31] ITU-T. The E-model: A computational model for use in transmission planning. <https://www.itu.int/rec/T-REC-G.107>.
- [32] ITU-T. P.862: Perceptual evaluation of speech quality (PESQ): An objective method for end-to-end speech quality assessment of narrow-band telephone networks and speech codecs. <http://www.itu.int/rec/T-REC-P.862>.
- [33] D. Kaplan. Suspicions and spies in Silicon Valley. <http://www.newsweek.com/suspicions-and-spies-silicon-valley-109827>, 2006.
- [34] H. Kargupta, S. Datta, Q. Wang, and K. Sivakumar. On the privacy preserving properties of random data perturbation techniques. In *IEEE International Conference on Data Mining*, 2003.
- [35] G. Karopoulos, G. Kambourakis, and S. Gritzalis. PrivaSIP: Ad-hoc identity privacy in SIP. *Computer Standards & Interfaces*, 33(3), 2011.
- [36] G. Karopoulos, G. Kambourakis, S. Gritzalis, and E. Konstantinou. A framework for identity privacy in SIP. *Journal of Network and Computer Applications*, 33(1), Jan. 2010.
- [37] A. D. Keromytis. A comprehensive survey of voice over IP security research. *IEEE Communications Surveys & Tutorials*, 14(2), 2012.
- [38] S. Le Blond, D. Choffnes, W. Caldwell, P. Druschel, and N. Merritt. Herd: A scalable, traffic analysis resistant anonymity network for VoIP systems. *SIGCOMM Comput.*

- Commun. Rev.*, 45(4), 2015.
- [39] M. Liberatore, B. Gurung, B. N. Levine, and M. Wright. Empirical tests of anonymous voice over IP. *Journal of Network and Computer Applications*, 34(1), 2011.
- [40] N. Mathewson and R. Dingledine. Practical Traffic Analysis: Extending and Resisting Statistical Disclosure. In *Privacy Enhancing Technologies*, 2005.
- [41] S. E. McGregor, P. Charters, T. Holliday, and F. Roesner. Investigating the computer security practices and needs of journalists. In *USENIX Security Symposium*, Aug. 2015.
- [42] T. Meyer. No warrant, no problem: How the government can get your digital data. <https://www.propublica.org/special/no-warrant-no-problem-how-the-government-can-still-get-your-digital-data>, June 2014.
- [43] S. Nakamoto. Bitcoin: A peer-to-peer electronic cash system. <http://bitcoin.org/bitcoin.pdf>, 2008.
- [44] N. O'Brien. Mobile security outrage: Private details accessible on net. <http://www.smh.com.au/technology/security/mobile-security-outrage-private-details-accessible-on-net-20110108-19j9j.html>, 2011.
- [45] Open Whisper Systems. Signal. <https://whispersystems.org>.
- [46] A. Pfizmann, B. Pfizmann, and M. Waidner. ISDN-mixes: Untraceable communication with very small bandwidth overhead. In *Kommunikation in verteilten Systemen*, volume 267, 1991.
- [47] A. Pfizmann and M. Waidner. Networks without user observability — design options. In *Advances in Cryptology - EUROCRYPT. International Conference on the Theory and Applications of Cryptographic Techniques*, 1986.
- [48] A. Ramo. Voice quality evaluation of various codecs. In *IEEE International Conference on Acoustics Speech and Signal Processing*, March 2010.
- [49] B. Reaves, L. Blue, and P. Traynor. AuthLoop: Practical end-to-end cryptographic authentication for telephony over voice channels. In *USENIX Security Symposium*, Aug. 2016.
- [50] B. Reaves, E. Shernan, A. Bates, H. Carter, and P. Traynor. Boxed out: Blocking cellular interconnect bypass fraud at the network edge. In *USENIX Security Symposium*, 2015.
- [51] M. Rizal. *A Study of VoIP performance in anonymous network – The onion routing (Tor)*. PhD thesis, Niedersächsische Staats- und Universitätsbibliothek Göttingen, 2014.
- [52] A. Ronacher. Flask (A Python microframework). <http://flask.pocoo.org/>.
- [53] J. Sanchez. Verizon employees fired after peeping Obama cell records. <http://arstechnica.com/tech-policy/2008/11/verizon-employees-suspended-after-peeping-obama-cell-records/>, 2008.
- [54] J. Sanchez. Other uses of the NSA call records database – Fingerprinting burners? <http://justsecurity.org/1971/nsa-call-records-database-fingerprinting-burners/>, Oct. 2013.
- [55] G. W. Schulz. Virginia police have been secretly stockpiling private phone records. <http://www.wired.com/2014/10/virginia-police-secretively-stockpiling-private-phone-records/>, Oct. 2014.
- [56] M. Schwartz. Lose the burners: Court okays prepaid phone tracking. <http://www.informationweek.com/security/mobile/lose-the-burners-court-okays-prepaid-pho/240005614>, Aug. 2012.
- [57] Selenium HQ. Selenium WebDriver. <http://www.seleniumhq.org>.
- [58] R. Siciliano. Protecting mail from identity theft. <http://robertsiciliano.com/blog/2011/03/22/protecting-mail-from-identity-theft/>, 2011.
- [59] Silent Circle. SilentPhone. <https://www.silentcircle.com/products-and-solutions/software/>.
- [60] The Guardian Project. Orbot: Mobile anonymity + circumvention. <https://guardianproject.info/apps/orbot/>.
- [61] The Register. The death of voice: Mobile phone calls now 50 per cent shorter. http://www.theregister.co.uk/2013/01/30/mobile_phone_calls_shorter/.
- [62] The Tor Project. Mumble – Tor bug tracker & wiki. https://wiki.mumble.info/wiki/Main_Page.
- [63] Twilio. APIs for text messaging, VoIP & voice in the cloud. <https://www.twilio.com>.
- [64] O. Verscheure, M. Vlachos, A. Anagnostopoulos, P. Frossard, E. Bouillet, and P. S. Yu. Finding "who is talking to whom" in VoIP networks via progressive stream clustering. In *International Conference on Data Mining*. IEEE, 2006.
- [65] A. M. White, A. R. Matthews, K. Z. Snow, and F. Monrose. Phonotactic reconstruction of encrypted VoIP conversations. In *IEEE Symposium on Security and Privacy*, 2011.

Appendix

A Screenshots of the Phonion Client Application

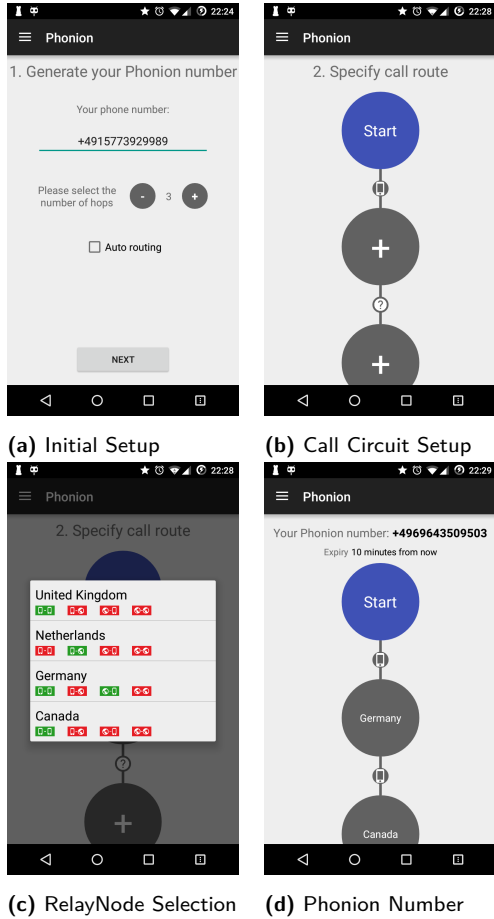


Fig. 6. Screenshots of the Phonion Android Client. The user first enters his real phone number and selects the desired number of Relay Nodes (Step a). The Client contacts Phonion Brokers, and the user selects Relay Nodes according to his or her requirements (Step b and c). Finally, the Client arranges the call circuit and displays the Phonion number to the user (Step d).