Building Performant, Privacy-Enhancing, and Blocking-Resistant Communication Systems

By

Piyush Kumar

Under the supervision of Sambuddho Chakravarty and Mukulika Maity

Computer Science and Engineering

Indraprastha Institute of Information Technology Delhi

New Delhi– 110020

August, 2021

Building Performant, Privacy-Enhancing, and Blocking-Resistant Communication Systems

By

Piyush Kumar

A Thesis

submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Computer Science and Engineering

Indraprastha Institute of Information Technology Delhi

New Delhi– 110020

August, 2021

Thesis Evaluation Committee

- Dr. Amir Houmansadr, Associate Professor, University of Massachusetts
- Dr. Michalis Polychronakis, Associate Professor, Stony Brook University
- Dr. Kent Seamons, Professor, Brigham Young University

Certificate

This is to certify that the thesis titled *Building Performant, Privacy-Enhancing, and Blocking-Resistant Communication Systems* being submitted by *Piyush Kumar* to the Indraprastha Institute of Information Technology Delhi, for the award of the degree of Doctor of Philosophy, is an original research work carried out by him under our supervision. In our opinion, the thesis has reached the standard fulfilling the requirements of the regulations relating to the degree.

The results contained in this thesis have not been submitted in part or full to any other university or institute for the award of any degree or diploma.

August, 2021

Dr. Sambuddho Chakravarty and Dr. Mukulika Maity

Indraprastha Institute of Information Technology Delhi

New Delhi 110020

Abstract

Free and open communication over the Internet is essential for the overall advancement of modern societies. However, there are numerous ways with which malevolent adversaries try to control the flow of information, by either selectively restricting access to content or in extreme scenarios completely shutting down the Internet. Hence, we aim to augment existing research as well as build novel systems to facilitate open communication over the Internet. Thus, in this thesis, we attempt to answer the question — Can we facilitate open access to Internet by building privacy-enhancing technologies that can also provide blocking resistance, good QoS and deployability? To that end, we propose and build solutions for three broad categories of applications. **(i)** applications for which there does not exist any functional and deployed solutions (*e.g.,* anonymous calling over Internet) **(ii)** applications (*e.g.,* accessing censored websites) for which there are functional solutions, but those have great scope for improvement (*e.g.,* blocking resistance, performance) **(iii)** accessing lightweight Internet applications during events of Internet shutdowns (*e.g.,* posting tweet and retrieving news snippets *etc.*).

Firstly, we explore the feasibility of anonymous voice communication over the Internet as it is of great interest to whistleblowers, privacy practitioners *etc.* Current literature sidesteps existing anonymity systems such as Tor for voice calling, citing the performance issues that would render real-time voice calling unusable over Tor. Thus, novel systems tailored for voice have been proposed (*e.g.,* Herd). However, the newly proposed systems are not functional, and thus there does not currently exist any usable anonymous voice calling system. Hence, we revisit Tor and perform a comprehensive analysis of performing voice calling over it. With the help of nearly half a million voice calls performed over the real Tor network over a year under various setups, we establish that it is indeed possible to perform anonymous voice calls using Tor.

Secondly, we look at accessing censored content (*e.g.,* websites, files *etc.*) over the Internet. There already exist solutions such as VPNs, Tor *etc.*, that help accessing restricted content by routing it through single (or multiple) proxy nodes. The censor blocks such proxies as soon as they are discovered, while the researchers attempt to build systems to evade such blocking, effectively leading to an arms race. Decoy Routing (DR) is an approach to potentially break this arms race, where a network router (known as the decoy router) doubles up as a proxy, instead of the end hosts. This makes it extremely difficult for the adversary to censor, as it can no longer block based on end hosts IP, but require blocking routers which carry a significant amount of innocuous traffic.

However, existing DR solutions use commodity servers to act as decoy routers, as traditional routers cannot be programmed to perform the complex operations required for proxying connections. This leads to (1) poor performance for end users due to handling of ISP scale flows by commodity servers, and (2) privacy violations of oblivious users (non-DR clients) due

to the analysis of ISP flows by third party DR maintainers. We thus propose *SiegeBreaker*, which overcomes the aforementioned problems (while also providing other salient features) by building a DR protocol with the help of software defined network (SDN) devices. We show that using SDN provides the essential modularity required in DR designs, eventually leading to good performance along with flexibility in trust relationships for providing better privacy guarantees. We implemented and deployed SiegeBreaker using hardware SDN switches and show with the help of extensive evaluation that SiegeBreaker provides performance comparable to direct TCP downloads. However, despite being a promising solution, DR poses deployment challenges as it relies on the ISPs support to help place the DR infrastructure. Thus, on one hand we have solutions that are readily available (VPNs, Proxies *etc.*) but are easy to block, and on the other hand we have solutions that are hard to block but face deployment challenges (DR).

Thus, next, we attempt to build a solution which is readily available to end users and at the same time provides effective blocking resistance without compromising the performance. To that end we propose a new system *Camoufler* that utilizes the Instant Messaging (IM) apps to transfer censored content. Using IM makes the system immediately usable as these apps are an integral part of the Internet ecosystem. Moreover, using IM apps as-is to transfer censored content makes it difficult for the adversary to distinguish it from normal IM traffic. We implement Camoufler on five popular IM apps (Signal, Telegram, Slack, Skype and Whatsapp) and demonstrate that a user can access censored websites in a few seconds (3.6s on an average for Alexa top-1K websites).

Lastly, we attempt to address a recent and an extreme form of censorship *i.e.,* Internet shutdowns. Such an extreme step of complete Internet disconnectivity leads to various problems for users in such regions (*e.g.,* reporting of power failures, immediate access to medical emergency *etc.*). To the best of our knowledge there do not exist systems that can provide Internet access during shutdowns. Thus, we build a novel system *Dolphin* that can provide access to basic Internet services such as Twitter, email, accessing news *etc.* Dolphin utilizes the cellular voice channel to encode the data bits into audio and transfer it via the cellular call. However, the cellular voice channel is unreliable, inherently lossy, insecure and highly bandwidth constrained. We overcome all the aforementioned challenges and build a complete end to end system. We experimentally demonstrate that Dolphin can be used to access emails, tweets and news all within a few minutes (an email of 500 characters was delivered securely and reliably within 3 minutes).

Acknowledgements

I would like to thank my advisor, Dr. Sambuddho Chakravarty who has immensely supported and guided me throughout my PhD. I would also like to thank my co-advisor Dr. Mukulika Maity for her insightful suggestions. Additionally, I would like to thank Dr. H.B. Acharya for his involvement and guidance in my PhD. I am also thankful to Dr. Vinayak Naik for providing his valuable feedback and insights.

I would also like to take this moment to thank my PhD colleague Dr. Devashish Gosain, who has been a great support to me throughout my PhD journey. Moreover, I would like to thank all the students, faculties and other technical and non-technical people with whom I have worked for the research projects. Special thanks to the IIIT Delhi administration, IT and facility management staff for making work at IIIT hassle free.

Lastly, I would like to thank my parents, who have been extremely supportive throughout my journey.

Research Papers From This Work

1. The Road Not Taken: Re-thinking The Feasibility of Voice Calling over Tor.
   Proceedings on Privacy Enhancing Technologies Symposium (PETS), 2020.
2. SiegeBreaker: An SDN Based Practical Decoy Routing System.
   Proceedings on Privacy Enhancing Technologies Symposium (PETS), 2020.
3. Camoufler: Accessing The Censored Web By Utilizing Instant Messaging Channels.
   Proceedings on Asia Conference on Computer and Communication Security (AsiaCCS) 2021.
4. Dolphin: A Cellular Voice Based Internet Shutdown Resistance System.
   Under submission to a reputed venue.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The Internet has grown well beyond its original purpose of providing a resilient communications infrastructure; in fact, it has become so essential to the free circulation of speech, information, and ideas that Internet access has been considered as a fundamental right by many countries [1]. However, in recent times, adversaries have attempted to contain the Internet with activities like network surveillance, IP address blocking, DNS injection, etc. It is becoming common that Internet Service Providers, working under the orders of malevolent adversaries (e.g., repressive regimes themselves), block, filter, intercept, or even modify traffic between clients on their networks and "controversial" Internet-based services [2, 3]. This has further led to a great deal of tension between civil liberty activists and repressive regimes. Thus, to facilitate freedom over the Internet, researchers have proposed various circumvention schemes. However, existing schemes witness multiple challenges. We broadly divide them into two categories:

Category I: Applications (such as anonymous voice calling) for which there are proposed schemes but none that are functional or immediately usable. This is because, the recently proposed systems (e.g., Herd [4]) are difficult to deploy in practice due to dependence on the recruitment of volunteer operators. Moreover, existing anonymization infrastructure (e.g., Tor [5]) is believed to not provide adequate QoS guarantees particularly for real-time applications such as VoIP [6, 4]. Thus, it becomes of paramount importance to take strides in providing such a solution.

Category II: Applications (e.g., those that allow accessing blocked web content) for which the proposed schemes have functional and usable implementations, but suffer from other drawbacks Arguably, such systems do not simultaneously provide blocking resistance and good performance to the users, while also keeping up with the privacy guarantees.

But all existing schemes that solve the above challenges would only function if the censor is not willing to completely cut-off from the Internet. Sadly, in past few years, we have now witnessed an increase in such extreme censorship cases, i.e., Internet Shutdowns. Thus, we aim to devise a novel scheme to combat such extreme measures.

Category III: Providing basic Internet access even during total Internet shutdowns (e.g., accessing email and news, posting a tweet, etc.). Currently, there does not exist any scheme that can be used to access the Internet during such shutdowns.

We thus aim to improve upon the existing research by proposing and building novel systems for the above three categories. Thus, in this thesis, we aim to answer the questions — "Can we build privacy-preserving communication systems that also provide real-time performance guarantees to the end-users, when connected to the Internet?" "In extreme scenarios of Internet shutdown, can we still provide minimal access to the Internet in a privacy-preserving manner?"

We commenced our research by focusing on the first category of problem i.e., achieving anonymous calling over the Internet. Anonymity has always been of great interest to whistle-blowers, free speech activists, journalists, political activists, etc. There exist solutions that provide anonymity over the Internet for accessing the web, but there does not exist any functional system that supports anonymous voice calling. Hence, we attempt to answer questions like — "Is anonymous voice calling possible over the Internet?", "Does providing anonymity have any quantifiable impact on perceived call quality?", "Do we need a new anonymization network to support VoIP traffic, or can existing anonymous networks (like Tor) be directly used?"

To assess the feasibility of anonymous VoIP calling we considered Tor — a pervasive overlay network with volunteer-operated nodes that are built to support anonymous communication. However, existing research deems popular anonymization systems like Tor unsuitable for providing the requisite performance guarantees that real-time applications like VoIP need. Their claims

are backed by studies [6, 7] that may no longer be valid due to constant advancements in Tor. Moreover, these studies lacked the requisite diversity and comprehensiveness. Thus, conclusions from these studies led them to propose novel and tailored solutions (for achieving anonymous VoIP calling), eventually rejecting Tor as a possible solution.

However, no such system is available for immediate use. Additionally, operating such new systems would incur signi cant costs for recruiting users and volunteered relays, to provide the necessary anonymity guarantees. It thus becomes imperative that the exact performance of VoIP over Tor be quanti ed and analyzed so that the potential performance bottlenecks can be amended. We thus conducted an extensive empirical study across various in-lab and real-world scenarios to shed light on VoIP performance over Tor. In over half a million calls spanning 12 months, across seven countries and covering about 6650 Tor relays, we observed that Tor supports good voice quality (Perceptual Evaluation of Speech Quality (PESQ) [8] >3 and one-way delay <400 ms) in more than 85% of cases. Further analysis indicates that in general for most Tor relays, the contentions due to cross-traf c were low enough to support VoIP calls, that are anyways transmitted at low rates (<120 Kbps). Our ndings are supported by concordant measurements using iperf that show more than the adequate available bandwidth for most cases. Hence, unlike prior efforts, our research reveals that Tor is suitable for supporting anonymous VoIP calls. Thus, in parts of the globe, where Tor is not blocked, it proves to be a readily available solution for achieving anonymous VoIP calling, in addition to accessing the web.

However, determined adversaries like China put constant efforts to block circumvention systems such as Tor, proxies, VPNs, etc. Conversely, researchers keep developing systems that counter the measures taken by such adversaries. This has led to an unsaid arms race between the Internet free speech activists and the censor. However, to put an end to this arms race, researchers proposed a novel solution, viz., Decoy Routing (DR) [9]. It is a promising approach to censorship circumvention, that uses routers (rather than end hosts) as proxy servers. Users of censored networks, who wish to use DR, send specially crafted packets, nominally addressed to an uncensored website. Once safely out of the censored network, the packets encounter a special router (the Decoy Router) which identi es them using a secret handshake, decrypts their content, and proxies them to their true destination (a censored site).

However, DR has implementation problems: it is infeasible to reprogram routers for the complex operations required. Existing DR solutions fall back on using commodity servers as a Decoy Router. But as servers are not efficient at routing, most web applications show poor performance when accessed over DR. A further concern is that the DR has to inspect all flows in order to identify the ones that need DR. This may itself be a breach of privacy for other users (who neither require Decoy Routing nor want to be monitored).

Thus, in the second part of this thesis, we attempt to answer questions like "Can we build a practical privacy-preserving DR solution?", "What performance guarantees can be assured with such a system?" and "How hard would it be to block the proposed DR system?"

To answer such questions, we present a novel DR system, SiegeBreaker, which solves the aforementioned problems using an SDN [10] based architecture. Previous proposals involve a single unit that performs all major operations (inspecting all flows, identifying the Decoy Routing requests, and proxying them). In contrast, SiegeBreaker distributes the tasks for Decoy Routing among three independent modules. (1) The SDN controller identifies Decoy Routing requests via a covert, privacy-preserving scheme, and does not need to inspect all flows. (2) The reconfigurable SDN switch intercepts packets and forwards them to a secret proxy efficiently. (3) The secret proxy server proxies the client's traffic to the censored site. Our modular, lightweight design achieves performance comparable to direct TCP downloads, for both in-lab setups, and Internet-based tests with commercial SDN switches. Thus, we experimentally demonstrate that privacy-preserving highly effective blocking resistant solutions can be built that could further end the arms race between free speech activists and authoritarian regimes.

It must be noted that on one hand, we have systems that are readily available but are relatively easy to block (e.g., VPNs, proxies, etc.). While on the other hand, the systems that provide highly effective blocking resistance (e.g., decoy routing) incur deployment costs (require collaboration with the ISPs to install decoy routers, etc.), rendering them currently unavailable for users and preventing their wide-scale adoption.

Thus, to further aid anti-censorship research, we answer questions like "Is it possible to build a system that is effective in resisting adversaries and at the same time is readily available

to end-users?", "What would be the deployment costs and performance overheads (if any), in building such a system?"

To find a possible answer, we attempt to use the Instant Messaging (IM) channels as a medium to tunnel blocked content. Our motivation to use the IM channel is two-fold. Firstly, IM apps are readily available to end-users, as they are pervasive. Thus, a system built on top of it could be easily accessed by the end-users. Secondly, completely blocking IM apps would result in collateral damage to the censor; an adversary would need to block all IM platforms to completely render the system ineffective, as IM apps are an integral part of today's economic as well as social activities.

Thus, we propose a new anti-censorship system, called Camoufler, which overcomes challenges in existing systems, while attempting to maintain similar blocking resistance available in existing schemes. Camoufler leverages Instant Messaging (IM) platforms to tunnel the client's censored content. This content (encapsulated inside IM traffic) is transported to the Camoufler server (hosted in a free country), which proxies it to the censored website. As a consequence, the censor would observe regular IM traffic being exchanged between the IM peers. Thus, utilizing IM channels as-is for transporting traffic provides unobservability, and also provides good QoS, due to its inherent properties, such as low-latency message transports. Moreover, it does not pose new deployment challenges or incur additional costs.

We implemented Camoufler on five popular IM apps (Telegram, Signal, Whatsapp, Slack and Skype) and it provides adequate QoS for web browsing. E.g., the median time to render the homepages (regular webpages) of Alexa top-1k websites was recorded to be about 3.6 s when using Camoufler implemented over Signal IM application. This is comparable to what one may observe when browsing these pages directly (via browser), in which case it takes around 2.7 s for the same.

Notably, traditional censorship revolves around blocking access to some websites (or services) over the Internet. However, recently there has been a rise in the events of an extreme form of censorship viz., deliberate Internet shutdowns [11] that lead to complete Internet disconnectivity, severely impacting the people residing in such regions. Moreover, these shutdowns render all

existing circumvention schemes unusable. Thus, in the final part of this thesis we aim to answer questions such as – Is it possible to build systems that can provide Internet access during such shutdowns? If so, what type of applications can we access using it? What are the security guarantees we can provide in such a system?

To that end, we present Dolphin, a first of its kind system that can provide access to lightweight and delay tolerant Internet applications (email, tweets, news snippets, etc.) during Internet shutdowns. Dolphin uses the cellular voice channel to transmit data bits. A user in the shutdown region (who wishes to access these applications) requires a trusted peer (end-point) in non-shutdown region to send and retrieve content on its behalf. The data bits between the peers are sent by first encoding them into audio and then transmitting them over a cellular voice call.

We overcome multiple challenges while designing and implementing Dolphin. E.g., the cellular voice channel is inherently lossy and unreliable. But the Internet applications need reliable transfers. Thus, in Dolphin we develop a TCP style reliability layer to overcome the losses. This layer works atop any underlying encoding/modulation technique. Further, to evade eavesdroppers over the insecure voice channel, we use this layer to derive keys for encrypting data. Also, Dolphin can work even without human intervention, by using cellular voice automation services. We experimentally show that Dolphin works for Internet applications, by testing it for sending email, tweets and accessing news snippets. All these applications take a few minutes to be accessed (e.g., a 500 character email was securely received in under 3 minutes).

To summarize, following are the major contributions of this thesis:

- To solve the problems of category I, we demonstrate for the first time that a functional and usable anonymous voice calling system is possible using an existing anonymity system Tor.
- Subsequently for category II, we present the design and implementation of a performant and privacy-preserving decoy routing system in the form of SiegeBreaker.
  Moreover, since DR systems pose deployment challenges, we designed another system, Cloudfer, that is easy to deploy and at the same time provides good performance and

blocking resistance.

- In case of category III, i.e., for extreme scenarios of Internet shutdowns, we develop a novel system Dolphin, that provides access to lightweight Internet applications such as twitter, news etc.

Chapter 2

Background

In this thesis we attempt to build privacy-enhancing systems to facilitate free ow of information over the Internet while also conducting measurement studies to assess their performance. First, we study the feasibility of performing VoIP calls over a popular anonymous network Tor (in Chapter 3). However, in parts of the globe where Tor is censored, we present a new hard to block and performant decoy routing system (SiegeBreaker) built with the help of software de ned networks as an underlying technology (presented in Chapter 4). But, SiegeBreaker may face deployability challenges as it requires collaboration from ISPs. Thus, we show the design and implementation of a new and easily deployable circumvention system (Camou er), that utilizes IM apps to tunnel restricted content. IM apps are ubiquitous and may cause collateral damage to the censor if recklessly blocked (explained in Chapter 5). Lastly, we consider the scenario of Internet shutdowns where all existing circumvention solutions cease to function. Therefore, in Chapter 6 we present a system (Dolphin) that can provide access to basic Internet apps during the events of Internet blackouts. To conclude the thesis, we present our overall takeaways and the potential future work in Chapter 7. The overall thesis road map is depicted in Fig. 2.1.

To better understand the contributions in the individual chapters, that include varied technologies like VoIP, Tor, SDN, and IM applications, we now describe the relevant background.

Figure 2.1: Overall thesis road map depicting different projects covered in each chapter.

2.1 On anonymous VoIP calls

In this section, we begin by describing the basic concepts behind VoIP calling and its evaluation metrics. Next, we brie y describe Tor, followed by a discussion on different types of anonymous calling scenarios.

2.1.1 Basics of VoIP

Voice over IP enables real-time voice communication over IP networks. Any VoIP based system comprises of two primary channels: one for control and signaling, and the other for transporting the actual encoded voice traf c. Session Initiation Protocol (SIP) [12], is an example of a popular VoIP signaling protocol. It includes functionality like authenticating users, establishing and terminating calls, etc. SIP, along with Session Description Protocol (SDP) [13], allows for negotiating the call related parameters like codecs. Once the call is set-up, Realtime Transport Protocol (RTP) [14], a UDP based protocol, solely manages voice traf c. It adds sequence numbers to packets for in-order delivery, and buffers them to minimize the impact of jitter.

2.1.2 QoS Metrics for Voice Calling

The following metrics are often used for measuring the quality of voice calls.

1. Perceptual Evaluation of Speech Quality: PESQ [8] is the ITU speci ed and standardized metric for voice call quality evaluation. It estimates the user perceived call quality. PESQ computation requires both the source and recorded audio for comparison. The PESQ metric generates an objective score using its own algorithm which is then mapped to a

subjective Mean Opinion Score (MOS). It is demonstrated that the score generated by PESQ highly correlates with the MOS score reported by actual users. The metric largely re ects distortions in recorded audio, which in-turn indicates the impact of network losses and jitters.

It returns values between 0 and 5, 5 being excellent, and 0 being poor and unusable. However, in practice, values between 1 and 4.5 are observed. Following ITU speci cations, in our tests, we considered calls with PESQ greater than three as acceptable.

2. Jitter and Packet Loss: Jitter represents variations between subsequent packet arrivals. Such variations may arise due to packet reordering and losses. While VoIP users can endure minor losses, jitter may dramatically hamper the perceived call quality. Even non-permissible variations in either of these two metrics can sharply perturb PESQ, which incorporates the impact of both the metrics.

3. One Way Delay (OWD): OWD is the time duration between when voice packets are encoded at the sender and when they are successfully decoded at the receiver. According to the ITU speci cation [15], OWD should ideally be less than 150 ms. Further, OWD below 400 ms is considered permissible for international calls [16]. Hence, we chose the permissible limit of 400 ms to evaluate the performance of voice calls.

Interestingly PESQ score, other than losses, only captures the impacts of jitters, which represents variations in OWD. It would suffer no perturbations if all the voice packets were uniformly delayed, without jitter or losses. However, for humans, such delays lead to time-shifted audio, eventually leading to poor perceptual quality. Thus, in our evaluation study we considered both PESQ and OWD as metrics for estimating call quality.

2.1.3 The Onion Router (Tor)

Tor [5] is a widely-used low-latency anonymization network. It allows it's users to communicate without revealing their IP addresses. It consists of globally distributed volunteered hosts acting as relays. Clients communicate to servers by proxying their traf c via a cascade of three such

relays, viz., the entry, the middle, and the exit nodes. The client encrypts the traf c using a three-layered encryption scheme, each corresponding to the three relays, using keys negotiated with each of them, respectively. These encrypted packets are then forwarded via the three chosen relays. Each of these relays de-crypts one layer of encryption and forwards it to the next one in the cascade. Thus, no one ever, other than the client itself, knows the IP address of all the relays and the server. Each relay only knows about the previous and next one in the cascade. The server only sees connections arriving from the exit node, but knows nothing about the client. By design Tor only supports TCP streams.

2.1.4 Types or Use Case of Anonymous Calling

As already mentioned, anonymous calls are of great use to whistleblowers, activists, undercover reporters, etc. However, it may be inquired as to what are the current alternatives (used by such groups) to conduct anonymous calls. To the best of our knowledge, there are no publicly available alternatives. This is the reason that in the past decade, prior efforts [4, 13] attempted to design and build such systems. However, as discussed in Subsec. 3.2, none of these systems are functional. Moreover, there are secure messaging and end-to-end encrypted communication apps such as Signal and Telegram which are highly popular among privacy practitioners. Though these provide security against eavesdroppers, the centralized architecture of all such apps allows the central server to know the details such as who is calling to whom. Thus, even though such apps provide secure calls, they may not be t to be used for anonymous calls.

We now describe the anonymous voice calling scenarios tested in our study.

1. Caller (one-way) anonymity: Here, the caller wants to achieve anonymity (by hiding IP address) against an adversary that may monitor and (or) lter its traf c. Such scenarios represent journalists and whistle-blowers who communicate sensitive information to other individuals or groups (e.g., news headquarters), while evading the adversary.
2. Caller-Callee (two-way) anonymity: Here, both parties want to achieve anonymity. Two individuals who both wish to communicate covertly, while remaining anonymous to their respective adversaries, may use such setups.

Further, these setups and their use cases are discussed in detail in Sec. 3.4.

2.2 SDN Based Decoy Routing

2.2.1 Decoy Routing

DR is an anti-censorship mechanism that that uses “friendly” routers in the Internet as proxy servers. Being routers, they are very dif cult to block (ref. Fig. 2.2). The steps for DR are as follows: (1) A user hosted in a censor regime sets up a TLS connection with an un ltered website (the OD), hosted outside the boundary of the regime. (2) The client-OD traf c carries a special cryptographic signature which can be identi ed by DRs en-route. On identifying this signature, the DR hijacks the connection, and engages a proxy (the DR proxy), that nally fetches the content from the CD. (3) The traf c from CD is returned to the client by the SP, setting TCP/IP header elds to appear as if they were part of the initial client–OD TLS connection.

Figure 2.2: Decoy Routing: The user accesses blocked.com , through DR. To the censor it appears that the client is communicating to an un ltered site, allowed.com .

2.2.2 Software De ned Networking (SDN)

Software De ned Networking [10] is an emergent network design paradigm that decouples the control and data plane of the network. An SDN architecture has a centralized controller, which controls the operation of the switches in the network (usually, by sending commands over a

standard protocol named OpenFlow [18]).

In SDNs, complex operations (such as computing network-wide control policies) are the responsibility of the controller. The policies, in the form of flow tables, are installed dynamically on the network switches which are simple dedicated devices, merely forwarding traffic based on these tables. The advantage of such an architecture is that control plane functionality is consolidated in a single, programmable entity – the controller.

Also, the SDN controller is capable of adding (or updating) any flow table rule or to redirect traffic to itself for analysis. By default OpenFlow compliant switches redirect any previously unseen packet, for which it has no matching rule, to the controller for inspection.

Since SiegeBreaker is also an SDN based anti-censorship system, it naturally inherits the salient features of SDN. In our design the controller and the SDN switches maintained by a friendly ISP are trusted; whereas the Secret Proxy managed by the DR operators/volunteers is not. Thus, any traffic analysed by the controller, will not compromise the ISP's regular customers' privacy, as it is maintained by the ISP itself.

2.3 Utilizing IM Platforms For Exchanging Blocked Content

Instant Messaging (IM) applications are one of the most widely used medium of communication over the Internet. They are used for both personal (WhatsApp [19], Telegram [20] etc.) as well as professional communication (Slack [21], Flock [22] etc.). This is reflected in a recent report [23] which depicts that the number of monthly active users of IM platforms, were 2.4 billion in January 2019. These users are expected to go beyond 3 billion in 2022.

IM applications provide variety of features which are similar across all IM platforms. These include real time messages, image, video and file sharing (for individuals as well as for groups). Some IM apps also support integration of payment wallets [24]. Apart from these features, such apps consider providing security (and privacy) to their clients as one of their design principles. To that end, almost all IM apps provide encryption by either requiring a TLS connection from the client to the IM app provider (i.e., end-to-middle (E2M)) or by having an end-to-end (E2E)

encrypted connection between the clients (ref. Fig. 2.3). This provides confidentiality to the IM clients from eavesdroppers who may attempt to snoop on the network traffic. However, it is desirable that the applications use E2E encryption as it ensures that neither the local eavesdropper, and nor the IM provider is able to see any content. Many applications support this feature, including Whatsapp, Telegram, Viber, Signal, Line, Fleep etc. On the other hand applications such as Skype, Wechat etc. rely on E2M encryption.

Figure 2.3: Difference between end-to-middle (E2M) and end-to-end (E2E) encryption in IM applications.

E2E and E2M encryption schemes. We now give a brief overview of these two different encryption schemes and how they are implemented in IM applications. In an IM application supporting E2E encryption, the peers depend on a centralized server which is responsible for distributing public keys of the peers and relaying messages between them. The client utility uploads a key bundle (containing public keys) to the central server. When an IM user Alice initiates communication with another user Bob, her client utility retrieves the key bundle of Bob from the key server, and derives session key (e.g., using Triple Diffie Hellman algorithm [25]). Then, Alice forwards her key bundle to Bob, which also derives the same session key. At this point, they establish an E2E encrypted connection with each other. With E2E encryption scheme, even the key server (i.e., an IM provider) cannot compromise the confidentiality of the messages being exchanged.

On the other hand, applications supporting E2M encryption generally use the TLS protocol to establish an encrypted connection between the users and the provider. E.g., if Alice and Bob wish to communicate with one another, they first establish individual TLS connections to the IM

provider's server[1]. As evident from Fig. 2.3, E2M encryption scheme allows the provider to observe the content (being relayed) in plain-text. However, it protects the con dentiality of the users' messages from any third-party.

2.4 Internet Shutdowns and Outages

Internet shutdowns are the deliberate act of turning off the Internet connectivity by the competent authorities at the behest of the governments. Such shutdowns have been on the rise, with 213 documented cases reported in 2019 alone. There is a steady increase in the number of countries performing Internet shutdowns with the number increasing from 25 in 2018 to 33 in 2019. These shutdowns could last for less than a day to over a year in some cases (472 days in Chad) [11]

Various projects keep track of these shutdowns at country as well as global scale. E.g., accessnow project [11] categorically reports incidents of shutdowns occurring across the globe, presenting detailed statistics of such events. Further, there are country speci c projects such as [16] which maintain a record of all the shutdowns that happen in India (country with the highest number of shutdowns). Some projects even attempt to estimate the economic losses in icted due to Internet shutdowns, e.g., internetsociety [27].

Other projects attempt to identify Internet outages in general. E.g., IODA [28] keeps track of Internet outages by performing active measurements using various probes, as well as using passive measurements by identifying anomalies in publicly available BGP paths and characterizing them as possible cases of outages. There are some proprietary projects such as ThousandEyes (managed by Cisco) [29] which also keep track of Internet outages across the globe in real-time.

Overall, while there are various studies and platforms that report Internet shutdowns and outages, but none provide solutions to circumvent them. Thus, we present Dolphin, a novel system which provides basic Internet connectivity to the users in shutdown and outage regions by using cellular voice (utilizing just a mobile phone and a laptop/desktop).

[1] A server managed by app maintainers which stores and relays messages.

Chapter 3

The Road Not Taken: Re-thinking the Feasibility of Voice Calling Over Tor

3.1 Introduction

Voice-over-IP (VoIP) applications that support traf c encryption are popular among users who are concerned about their communication privacy. However, these applications do not safeguard the anonymity for Internet users residing in regimes which may conduct surveillance (e.g., the ability of NSA to intercept conversations is widely known [30, 31, 32]). Moreover, to the best of our knowledge, there does not exist any functional VoIP based system that ensures communication privacy and anonymity (along with real-time communication guarantees), required by privacy conscious Internet users and whistle-blowers alike.

Popular privacy and anonymity preserving systems like Tor [5] hide the actual IP address of the communication peers by routing their traf c via a cascade of proxies. Since such systems reroute traf c via circuitous paths, it is a widely held belief that they would incur intolerable delays for real-time applications like VoIP. More importantly, while traditionally VoIP relies on UDP traf c to ensure real-time guarantees, popular systems like Tor are designed to transport TCP traf c. Further, complex cryptographic handshakes, essential to the anonymity guarantees provided by such systems, may exacerbate the impact on performance, making them unsuitable

for real-time applications.

Prior efforts on anonymous calling [3, 4, 17, 34] unanimously agree with aforementioned shortcomings of Tor. Some conclude this based on potentially biased results, involving relays only in Europe or only conducted a few hundred calls [13], lacking diversity. Others, such as Le Blond et al. [4], did not conduct any experiments to measure VoIP performance over Tor. Overall, we believe that existing literature does not quantify the actual interplay of network performance attributes (e.g., one-way delay (OWD), available bandwidth, etc.) and how it impacts VoIP call quality over Tor. However, the belief that Tor is not competent enough to transport VoIP traffic is still prevalent [17].

Thus several novel VoIP architectures were proposed [3, 4, 17] to tackle the shortcomings. Some of these, like Phonion [17] and Herd [4], require a new volunteer-run network with millions of active users (like Tor) for providing anonymity guarantees. This requirement can be a major stumbling block. Moreover, in the absence of active users and significant cross-traffic, comparable to that of Tor (that transports over 200 Gbit/s traffic per day [15]), one cannot adjudge future anonymity and performance assurances of these proposals. Importantly, no such system is currently functional.

Hence, in the absence of an existing anonymous voice calling system, we chose to determine the root cause(s) of poor performance over Tor. After ameliorating them, one may expect to achieve adequate voice call quality with anonymity guarantees equivalent to that provided by Tor. To that end, we began by conducting a pilot study where we made VoIP calls over the Tor network. We also captured the network performance attributes, in order to eventually identify how they impact voice call quality. Following ITU guidelines and prior proposals [17], we used one-way delay (OWD) and Perceptual Evaluation of Speech Quality (PESQ) [8] as metrics to judge VoIP call quality. The PESQ ascribes a value to judge the user-perceived audio quality in an automated manner.

We made 1000 consecutive calls through individual Tor circuits[1] with the callee and caller machines under our control. Contrary to the prevalent notion of poor quality of calls via Tor,

[1] Using the standard Tor client utility.

we observed good call quality, with average PESQ 3.8 and average OWD 280 ms. Overall, 85% of the calls were acceptable (PESQ > 3 and OWD < 400 ms) as per ITU [16, 15].

In the absence of substantial evidence of poor quality calls, we went ahead and conducted an extensive longitudinal study involving 0.5 million voice calls over the Tor network, spread across 12 months. These measurements not only involved varied in-lab and real-world scenarios with diverse Tor relays, VoIP applications, caller-callee locations but also a user study. To our surprise, even then more than 85% of voice calls had acceptable perceptual quality.

The PESQ metric varies inversely with distortions in the perceived audio. Packet drop and jitters, which cause such distortions, are an artifact of increased cross-traffic contentions. A high PESQ score, accompanied with low overall OWD, in the majority of the cases thus indicates low cross-traffic contentions, for VoIP calls. Moreover, VoIP calls which are mostly transmitted at low bit-rates (< 120 Kbps), incur less routing costs, and thus may not suffer much distortions.

Other network performance metrics like available bandwidth also varies with network cross-traffic volume. Thus, concomitant available bandwidth during the call (> 1 Mbps in 90% cases[2]), along with data published by Tor Metrics [35], confirms the reason for obtaining good results.

Overall, this first ever long-term study involving extensive evaluations of calls over Tor, bore some interesting results and insights. We summarize them as follows:

1. In the vast majority of our experiments (85%), involving voice calls over individual Tor circuits, we observed acceptable call quality with PESQ above 3 and OWD less than 400 ms. This holds for a diverse set of scenarios:
   (a) Caller and callee spread across 7 countries (in three continents).
   (b) Coverage of a total of 6650 Tor relays, 22 times more than previous studies [6].
   (c) Popular VoIP apps such as Telegram and Skype.
   (d) Both caller and (or) callee using Tor circuits to establish calls with one another.
   (e) Different codecs and call duration.

[2] The bandwidth was measured by iperf .

Figure 3.1: Existing literature on anonymous voice calling highlighting their inconsistencies and incompleteness. X implies metric not considered or evaluated in the study.

(f) User study involving humans rating voice calls.

2. Acceptable performance in the vast majority of cases was due to relatively low contention for VoIP in most Tor relays.

3.2 Related Work

There exists scant literature on performance evaluation of voice calls over Tor (ref Fig. 3.1). The only attempts made were by Rizal et al. [6] and Heuser et al. [17]. Their efforts involved instantiating a few hundred calls through Tor for evaluating their quality.

However, the aforementioned efforts were limited in scope. E.g., rather than using PESQ (an established call quality metric), Rizal et al. relied only on network parameters like OWD, jitter, and packet loss as evaluation criteria. Also, this study involved Tor relays only in Europe (with only 298 out of the available 4453) for their evaluation, and may not be representative of VoIP performance, for the complete Tor network. Based on these preliminary studies, previous researchers ignored Tor and proposed novel architectures [4, 13, 33] for providing anonymous voice calls. Next, we describe all such efforts.

Inspired by Chaum's mixes [5], Pfitzmann et al. [37] proposed ISDN mixes for anonymous voice communication. Authors proposed two simplex Mix channels, one for the sender and the other for the recipient, thereby enabling full-duplex anonymous communication over the telephony network.

Drac [33] by Daneziset al. involves an architecture for anonymous low latency voice communication using social networks as relays to route traf c. The system, while providing anonymity to a particular user, relies on using the identity of other users in its social circle as alibis. However, Drac presents an analytical model with no functional deployment.

Torphone[38], a system designed over Tor, was an extension to send VoIP traf c via Tor. It reported having achieved an OWD of 2–4 s, which is unsuitable for acceptable call quality. However, it is presently non-functional (and was last used on Windows XP).

In 2015, Le Blondet al. [4], proposed a novel architecture,Herd, to prevent global passive or active adversaries attempts to de-anonymize users, based on call metadata (correlating start and end times of a call). Herd relies on a set of dedicated mixes that relay VoIP traf c to other mixes and endpoints while hiding any distinct traf c patterns. The mixes are also known aszones and a user can select available trustworthy zones to route its call. They believed prior results on evaluating the performance of VoIP over Tor (published back in 2008 [70]) and concluded the RTT to be high (2–4 s), deeming Tor unsuitable for VoIP. However, they refrained from conducting a fresh study to gauge the (then) recent performance of Tor. Additionally, they concluded that the calling entities on Tor could be easily correlated, given the start and end times of a call. They validated this claim by analyzing the call records of a service provider. However, obtaining such data for calls conducted via Tor might not be easy, as Tor is a globally distributed system, and it would require gathering data from geographically diverse ISPs. Moreover, they tested their prototype on only four cloud hosts, performing just 12 calls, in the absence of active users and signi cant cross-traf c when compared to Tor.

Phonion [17] is one of the recent anonymous VoIP systems. It is fundamentally similar to Tor, but speci cally developed for voice communication. It uses relays(similar to Tor relay), relay servicesand broker system(similar to Tor directory authorities). Phonion attempts to anonymize call data records (CDRs) against different adversaries. In Phonion, the calls are relayed via various service providers. This prevents a single provider to gather all the call records for a particular call. The advantage of Phonion is that it works across different voice calling technologies — VoIP, cellular or PSTN (public switched telephone network). Additionally, it requires Internet access only for an initial bootstrap phase after which, anonymous calls could be

instantiated over carrier services. However, the authors of Phonion also side stepped using Tor. Through a pilot study of 100 calls, they concluded Tor to be unsuitable for transporting VoIP calls. The study involved comparing one, two and three-hop circuits of their prototype, with six-hop Tor circuits (two-way anonymity). However, we show in Sec. 3.4.2.1 and Sec. 3.4.2.2 that regular three-hop Tor circuits provide suitable performance for VoIP. Usage of this setup and performing 100 calls might have led them to report an unacceptable average OWD of 777 ms for Tor circuits along with an average PESQ of about 2.4. Moreover, the implementation of their system relied on google voice, a service only available in N. America, making the system unusable elsewhere. Lastly, it required the relay operator to pay an operational fee to telcos, which might serve as a stumbling block for recruiting relays.

Overall, a major hindrance in the wide-scale adoption (and availability) of such systems, is the recruitment of new volunteered relays and users (which Tor already has in abundance). Schatz et al. [34] also highlight this issue. Additionally, as previously mentioned, prior evaluations over Tor, provide very few insights as to how the interplay of network performance attributes affects the voice call quality.

Figure 3.2: V-Tor: The caller establishes a VPN tunnel through Tor, so that all the UDP VoIP traffic traverses the Tor network. On reaching the VPN server, the traffic is sent to the SIP server, which initiates a voice call to the callee.

Figure 3.3: M-Tor: The caller configures the mumble client to work in TCP mode. Thereafter, mumble's traffic is sent over Tor eventually reaching the mumble server, which handles the call procedure between the caller and callee.

Comparison with prior VoIP measurements over Tor: As already mentioned, only two studies actually measured the performance of VoIP over Tor — Rizal's study [6] and Phonion [17] by Heuser et al. We now explain how our measurement study is methodologically different from them.

Rizal's study: This study transported VoIP traffic over Tor by tunnelling it via VPN tunnels. The experiments conducted in this study, involved Tor relays hosted only in Europe. Thus, Rizal

reported a reasonably low average OWD of 152 ms likely due to the geographic proximity of relays. On the other hand, our study included relays from all parts of the globe and hence provided better coverage of the entire Tor network[3]. Thus, our study measured an average OWD of 280 ms likely due to the diversity of involved Tor relays. Moreover, unlike Rizal's study, we used PESQ as an evaluation metric, which is an industry standard for measuring user perceived voice quality.

Phonion: It utilized the Mumble VoIP software [39] to transport VoIP traffic via Tor. Authors of Phonion used PESQ as an evaluation metric. However, their study involved measuring performance only for two-way anonymity, and not the one-way anonymity. On the other hand, our study involved different scenarios corresponding to both the cases, viz., one-way and two-way anonymity. Hence, we involved comparatively more diverse and comprehensive set of measurement setups.

Additionally, both the previous studies involved performing only a few hundred calls – 100 in Phonion and 600 in Rizal. In contrast, we performed 0.5 million calls involving a diverse set of geographic locations, media codecs, etc., along with controlled experiments involving private Tor setups.

Overall, these studies lack the requisite comprehensiveness and a detailed analysis of different performance attributes that could provide deeper insights for VoIP performance over Tor. Thus, our study aimed to fill this research gap by conducting an extensive study with an intent to better understand the behavior of VoIP calls over Tor in varied setups and network conditions.

Other Tor performance measurement studies: There exist abundant studies which measure the Tor network's performance in terms of observed bandwidth and latency [40, 41, 42, 43, 44, 45, 46, 47, 48, 49]. Few, like Shadow [42] and Chutney [50], developed simulators to analyze the performance of Tor. Others like Tor flow [47], EigenSpeed [48], and Peer flow [49] etc., focused on measuring relay bandwidth for calculating relay weights by either directly measuring relay bandwidth (Tor flow) or by indirectly inferring it using statistics such as relays reporting the amount of data exchanged between them (Peer flow). Measuring relay bandwidth is a crucial task

[3] Achieved using the standard Tor utility that helped create different circuits based on default Tor circuit selection algorithm.

in Tor, as it is used to assign weights to different relays that govern the selection probability of a relay in a Tor circuit and thus the amount of traf c the relay might serve. Cangialosi et al. [51] speci cally focused on measuring RTT between relay nodes. Tor metrics [52] is another popular (and actively maintained) project that periodically measures various performance attributes of the Tor networke.g.,circuit RTTs, bandwidth for downloading les of various sizes etc. However, none of these existing projects, measure the perceptual VoIP quality and its associated network characteristics. In our study, we primarily measure the VoIP performance over Tor, along with network characteristics (like bandwidth and RTT) of millions of Tor circuits. The metrics, such as available bandwidth are essential in our study as they help to infer the potential reasons for good or bad performance while conducting VoIP calls. We often referred to the Tor metric results to further support our claims.

3.3 Measurement Approach

In this section, we describe our experimental setups and the approach taken for performing experiments to measure the quality of voice calls over Tor.

3.3.1 Experimental Setup

Previous efforts acknowledge that transporting VoIP traf c anonymously over Tor is non-trivial, as VoIP generally uses UDP and Tor only supports TCP. There are two ways in which one can succeed in sending VoIP traf c over Tor. One way is to tunnel UDP packets inside TCP ows. The other way is to directly encode and send VoIP traf c in TCP packets. Thus, some previous studies [6] utilized VPN tunnels to encapsulate and transfer VoIP packets, while others like Phonion [17], relied on Mumble [39] VoIP software to generate TCP packets and transfer them directly to Tor. Similar to these studies, we tested VoIP performance using both the above approaches. These setups are now described in detail below:

1. SIP client via VPN through Tor (V-Tor): The V-Tor setup (ref. Fig. 3.2), involves a SIP client (caller) connecting through a Tor circuit to a VPN server for establishing a TCP

tunnel. A step by step walkthrough of this setup is described below:

Step1The caller runs the VPN and Tor client utilities. VPN client is con gured to establish a VPN tunnel (to the VPN server) over a Tor circuit, by forwarding VPN traf c to the Tor SOCKS interface. This ensured that all the traf c from the caller would reach the VPN server via the Tor network.

Step2Client would initiate a VoIP call using a SIP utility. The VoIP traf c (generated from SIP) would reach the VPN server via Tor (explained in step 1). VPN server decapsulates SIP packets and forwards them to the SIP server.

Step3SIP server then helps negotiate the call between the caller and the callee.

2. Mumble with TCP mode over Tor (M-Tor): This setup relies on using the Mumble client program (in TCP mode) for encoding voice traf c through TCP streams. These streams are transported via Tor to a Mumble server that mediates the voice call between the parties.

It must be noted that we performed experiments for both the V-Tor and M-Tor setups. However, since we obtained similar experimental outcomes from both the setups, we present description of V-Tor experiments in the subsequent sections and sub-sections, and the details of M-Tor experiments in Appendix A.1.

3.3.2 Overview of Experiments

We now describe the experiments performed to measure the VoIP call quality over Tor. Our experiments were primarily conducted to identify the root cause of the hitherto believed poor voice call quality of Tor.

We began with a pilot study that involved conducting 1000 consecutive calls over Tor, with the caller and callee under our control, but in different geo-locations. Our results, using V-Tor (and M-Tor) setup, showed high call quality in a large fraction of the cases. Considering these results to be potential outliers, compared to ndings of previous authors, we went ahead and conducted a comprehensive measurement study spread across 12 months.

We conducted two different sets of experiments — (1) involving in-lab setups and (2) involving circuits through the public Tor relays.

In-lab experiments: In the in-lab setup, our goal was to measure the performance of VoIP over V-Tor (M-Tor) setups with competing cross-traffic entirely under our control. For this, we setup a private network in our lab, consisting of Tor nodes along with client and server (VPN, SIP, etc.) machines. This private network was deployed on real machines, with three of them serving as relays, while one of them was also serving as a directory authority. Other machines acted as clients and servers. In these experiments, we measured performance attributes and established baseline values (e.g., bandwidth requirement) for a VoIP call. These experiments were performed over setups involving different combinations of VPN, Mumble, and Tor in the following manner:

1. Direct VoIP calls: VoIP calls between the caller and callee were conducted without involving Tor and VPN. The caller and callee communicated directly using SIP or Mumble protocol. This helped us in observing the minimum bandwidth and delay requirements when no overhead was introduced (due to Tor or VPN).

2. VoIP calls over VPN: Caller and callee communicated using SIP protocol. However, calls were encapsulated through VPN connections, in order to measure the impact (if any) due to the overhead of running a VPN.

3. VoIP calls over Tor: Encapsulating the calls through VPN connections (or Mumble) and then transporting them via Tor circuits, to evaluate the impact of the overheads due to Tor.

Moreover, we performed some additional experiments to observe the impact of variation in background cross-traffic on VoIP calls. We observed the variation in performance when VoIP call(s) were made in the presence of heterogeneous background traffic (other VoIP calls and web traffic). Such in-lab experiments may potentially present clues regarding the number of VoIP clients that could be simultaneously supported by the real Tor network.

Internet based experiments over public Tor: After performing various in-lab tests, we carried out multiple experiments involving public Tor relays, where we had no control over the background cross-traffic and network conditions. These experiments involved measuring performance

across diverse scenarios (Tor relays, end-points, codecs, etc.), with the intention of studying the variation in performance under real-world conditions. More speci cally, the experiments involved measuring the call quality by varying:

- Tor circuits: Involved measuring quality across a large number of circuits (650 unique relays) created using the regular Tor client program.
- Geo-location of communication peers: Involved measuring quality by instantiating several calls, varying the location of the calling peers.
- Type of anonymity achieved: Involved measuring quality while achieving one-way and two-way anonymity, in accordance with the use cases already described in Sec. 2.1.4.
- Circuit lengths: Involved measuring quality over two-hop circuits. By default, Tor circuits are built using three relays.
- Codecs: Involved measuring quality by varying the call codecs used.
- Call duration: Involved measuring quality when call duration was varied.
- Type of relays used: Involved measuring quality when using bridges instead of public Tor relays.

We also measured the voice call quality for few popular voice calling apps such as Telegram [53] and Skype [54], when used over Tor. Users may choose to rely on using these already popular and familiar apps, instead of having to setup V-Tor and M-Tor. Additionally, we also conducted a user study involving 20 participants who rated calls via Tor. Most of experimental results are presented in the next section (Sec. 3.4). A small fraction, addressing important concerns related to VoIP performance, e.g., impact of call codecs, Tor bridges, etc. are presented in Sec. 3.6.

3.3.3 Implementation Details

Host con gurations: All machines of the in-lab experiments used Intel Core i5 8th gen CPUs with 8 GB of RAM. The hosts used in the experiments involving public Tor relays were hosted on

Digital Ocean's cloud based infrastructure, distributed across seven countries. These machines were equipped with Intel Xeon 2.2 GHz single core CPUs and had 2 GB RAM. Since the latter experiments required initiating only a single call at a time, we did not require machines with higher memory.

Tor con guration: In the basic con guration, all the caller and callee hosts had the latest Tor v0.3.9 installed. We used the Tor stem python library [55] to ensure that (1) only a single circuit was enabled at a time (2) the performance metrics (such as bandwidth measured using iperf , RTT using ping , etc.) were measured for the same circuit through which the call was performed.

Communication peers: For V-Tor experiments, the end hosts were installed with OpenVPN [56] v 2:4:7. The VoIP traf c generated was encapsulated through OpenVPN client and transported via Tor (by specifying the SOCKS port in the OpenVPN con guration). The clients used the python based SIP client, pjsua [57], a command-line softphone, to automate SIP calls. The module supported audio playout and recording.

VPN/SIP server: The V-Tor setup uses a machine con gured to be a VPN as well as a SIP server. OpenVPN [56] v 2:4:7 was con gured to work as the VPN server. It forwarded the encapsulated SIP calls to the SIP server. Freeswitch PBX [58] was used as the SIP server for handling SIP calls. We used different con guration les that come with Freeswitch to make call extensions, route calls, select codecs, etc.

Popular VoIP apps: Among all the popular apps, Telegram is the only one that provides user APIs. Still, we require to overcome several challenges to measure the performance of Telegram calls over Tor.

Firstly, the API provides only instant messaging automation facility. We thus mined the source code and discovered that it relies on libtgvoip [59] for making voice calls. We used this library for automated calls. Secondly, the library uses UDP for transporting VoIP. We thus modi ed it to enforce transporting voice calls via TCP streams. Thirdly, libtgvoip is con gured, by default, to play audio clips endlessly in a loop. We applied appropriate modi cations to control the playout duration, to suit our calls. Fourthly, we also required

modifying the library so as to synchronize the call setup and termination events with starting and stopping of voice recording. This synchronization is required for accurately measuring PESQ. Additionally, pyrogram [60] redirected Telegram traf c to the Tor SOCKS port.

Other popular voice calling apps like Skype poses two challenges, viz., absence of resources like APIs or source codes to aid call automation, and ability to redirect traf c through Tor by con gurable SOCKS interface. To overcome the rst obstacle we initially synced the caller and callee machines. Thereafter once the call was initiated, the caller plays out the audio clip using *mplayer* [61], while the callee captured and recorded the call audio directly from the sound card using *pactl* [62] utility. To overcome the second challenge, we used OpenVPN to encapsulate and redirect Skype call traf c to the Tor SOCKS interface.

3.4 Measurement Results

In this section, we describe the experiments conducted to test the performance of anonymous calls over Tor, and their corresponding outcomes. We begin by enlisting some common steps we followed while conducting the experiments:

- In all our experiments, a caller host played out an audio clip containing 30 s of human speech. It was encoded and transported, via a unique Tor circuit, to the callee. The callee recorded the audio, which is later used for computing PESQ.
- For every call, we recorded the network traf c through *pcap* les, and also measured various network performance attributes of the Tor circuit (through which the call is performed) like available bandwidth and RTT using *iperf* and *ping*, respectively. The *ping* utility was run during a call, as it is not a bandwidth intensive test. On the contrary, *iperf* (bandwidth intensive) tests were conducted after the completion of the call, so that it does not have any impact on the call quality.

  For the in-lab experiments also, we measured the stream bandwidth using *iperf*.
- PESQ score for every call was calculated by comparing the original (one played out at

the caller) and recorded (at the callee) audio clips. Any score above three was considered good [15].

- One way delay was also calculated for the duration of the call. We used ping to calculate OWD. As per ITU guidelines for international calls [16], the upper limit of OWD for acceptable call quality is 400 ms.
- We ensured that for a call, all the performance metrics were measured for the same circuit through which the call was instantiated.

In every experiment, the above steps were repeated for each call. Thereafter, we analyzed the measurements and performance metrics across all these iterations.

3.4.1 In-Lab Experiments

We performed these experiments, with an intent of measuring call performance under different testing conditions, while fully controlling network link capabilities and background cross-traffic. To establish the baselines, we created three test scenarios for the setup. These involved: (1) Direct SIP calls (2) SIP calls over VPN tunnel (3) SIP calls through VPN over Tor (V-Tor). All these experiments followed the setup described in Fig. 3.2 For the scenarios where Tor was not used, the nodes between caller and callee (in Fig. 3.2) merely functioned as routers. This was done to minimize any biases in performing experiments, by ensuring that the packets traverse the same number of hops[4].

Experiments using V-Tor setup: We started by initiating VoIP calls over all the three setups and computed their respective PESQ scores and OWD values. The capacity of the link between the caller and callee was 100 Mbps. The measured PESQ score averaged across individual samples was the same for all the three scenarios i.e. 4.5. Whereas OWD was below 100 ms. This result established that for a single call, with no competing cross-traffic, the overheads introduced by the VPN and Tor had no significant impact on the call quality. We additionally observed that the available bandwidth requirement for a single call in all the three scenarios was no more than

[4] Additionally, removing these hops do not have any observable impact on our results. We kept them just to have uniformity among our experiments.

120 Kbps (ref. Tab. 3.1). As expected, direct calls transmitted at the lowest rates. Additional overheads due to the headers introduced by VPN and Tor progressively increased the bandwidth requirements.

Call category	Bandwidth (Kbps)
Direct SIP call	84
SIP call via VPN	108
V-Tor	120

Table 3.1: Baseline bandwidth (in Kbps) requirement of VoIP in different scenarios.

Next, to understand the impact of cross-traf c on VoIP call quality, we initiated VoIP calls in the presence of cross-traf c. The experiments were carried out for three link bandwidth con guration—2 Mbps, 5 Mbps and 10 Mbps. Studying the performance under cross-traf c, for different link bandwidth would help us understand if the observed behavior is consistent or not. These experiments were speci cally conducted for the VPN via Tor (V-Tor) setup. Further, these lab experiments provided us insights on the number of calls that could potentially be made under varied network conditions on the real-Tor network.

Thus, we gradually introduced the cross-traf c by increasing the number of parallel le downloads (using wget) from another client that shared the link with the caller. We made sure that the cross-traf c was in the direction of call, to ensure adequate cross-traf c contention. We measured the degradation in the call quality, by computing the average PESQ score, for every new parallel connection introduced. The cross-traf c was gradually increased such that it utilized the total link capacity from 5%, to 10%, and then all the way up to the point where the call under consideration received less than 120 Kbps of the total available bandwidth. At this point, we observed a sharp decline in PESQ for the call (2.3). This corresponds to unacceptable call quality. Our ndings are summarized in Tab. 3.2.

Competing Streams	Link Bandwidth	Available Bandwidth Per Stream	Call Requirement	PESQ Score
< 75	10 Mbits	> 133 Kbps	120 Kbits	> 4.2
80	10 Mbits	125 Kbits	120 Kbits	3.4 #
> 85	10 Mbits	< 117 Kbits	120 Kbits	< 2.3 ##

Table 3.2: Analysis of V-Tor under the presence of competing non-VoIP (web or le downloads) cross-traf c.

We then performed experiments, where the background cross-traf c constituted of other

VoIP calls. Similar to the previous experiment, we performed this test for three different link capacities (2, 5 and 10 Mbps) for the V-Tor setup. The results of the 5 Mbps link bandwidth test are summarized in Tab. 3.3. We obtained similar behavior for the other two bandwidth categories.

Competing VoIP Calls	Link Bandwidth	Available Bandwidth Per Call	Call Requirement	PESQ Score
< 35	5 Mbits	> 145 Kbps	120 Kbits	> 4.2
40-43	5 Mbits	128 Kbits	120 Kbits	3.3 #
> 43	5 Mbits	< 120 Kbits	120 Kbits	< 2.3 ##

Table 3.3: Analysis of V-Tor under the presence of competing VoIP cross-traffic.

The results indicate that when the contention on the shared link increases, the PESQ drops. The PESQ metric is very sensitive to the impact of even minor network drops or delays. Even a small increase in contention, e.g., only five additional download streams reduce PESQ from 4.2 to 3.4 (ref. Tab. 3.2). Corresponding to increased contentions, the available bandwidth for every stream (including VoIP) drops. The constant bit-rate voice traffic (120 Kbps) suffers significant distortions that may, however, have little impact on the non-VoIP flows. This held when the cross-traffic was non-VoIP as well as when it was VoIP.

These in-lab measurements indicate that, a client should be able to conduct good quality calls, if the constructed circuit provides a bandwidth of above 120 Kbps. This should hold true on the real Tor network as well. E.g., if a circuit has an available bandwidth of about 1.2 Mbps, then it is capable to simultaneously support a maximum of 10 VoIP clients (with acceptable call quality).

3.4.2 Experiments Involving Public Tor Relays

Having obtained good performance in in-lab tests, we went ahead to evaluate the performance of voice calls over public Tor network. We expected the results to vary significantly due to the dynamic nature of competing cross-traffic and network conditions over the Internet.

We began with our pilot study that involved a client host (caller), positioned in our university, establishing VoIP call to a cloud hosted peer (callee). The call traffic was transported via Tor. We

sequentially started 100 calls, each transported via a different Tor circuit, and measured the call quality by computing PESQ score. To our surprise, for both the setups we observed an average PESQ of 3.8 and an acceptable OWD of 280 ms. Even after 1000 calls, each transported via a freshly created Tor circuit, we observed very similar performance measures (PESQ 3.80 and OWD 273 ms).

However, one may argue that our positive results might have been a small anomalous fraction. These may have been different from the bulk of poor outcomes that may have led others to deem Tor as unfit for VoIP. In order to test that this was not a fluke, we conducted a longitudinal experimental study covering diverse scenarios. The experiments are described below.

3.4.2.1 Caller anonymity: co-located voice server and callee (Scenario I).

We begin with the fundamental scenario where a caller is positioned in a censored network and wishes to call someone who is beyond the censor's control. The caller makes calls to the callee, through the public Tor network, using setups similar to one shown in Fig. 3.2. It is assumed that callee runs a publicly accessible VPN server (for V-Tor) on its host.

In our experiments, we chose seven individual cloud machines as callers, and three other as callees. Each of these was selected from Europe, N. America, and Asia. For every caller–callee pair we made 1000 calls using the V-Tor setup. In each case, we measured the PESQ and OWD. A total of 42000 calls were made. The average PESQ and OWD across all measurement was 3.88 and 217 ms, respectively.

By default, Tor circuits are three hops long. To reduce the potential impact of hop length on performance, we repeated the said experiments by making calls over two-hop circuits. This did not led to any significant impact on the PESQ (3.90). However, as expected, the OWD reduced to 205 ms. The CDF of PESQ scores obtained is depicted in Fig. 3.4. Results clearly show PESQ above 3 in over 93% of the calls.

Figure 3.4: V-Tor: CDF of PESQ for Caller Anonymity when server is co-located with callee (Scenario I).

3.4.2.2 Caller anonymity: separate VPN/voice server (Scenario II).

There are, however, certain limitations of Scenario 1. Firstly, the setup requires a publicly accessible VPN server, which may be infeasible when the callee is behind a NAT. Secondly, there may be cases where multiple whistleblowers or covert reporters (i.e., several callers), communicate to callees working for a common organization. In such cases, having a commonly shared VPN/SIP server, with high availability, supporting features like voicemail, removes the need for the callee to be always online. Thirdly, it reduces the hassle for every callee to port VoIP server to different platforms.

Therefore we considered an alternative setup where the VPN / SIP server and callee were not hosted on the same host. They were distributed among seven different cloud hosts, positioned across Europe, N. America and Asia. This separation may incur higher OWD between the communication peers, due to the intervening network between the VPN/SIP server and the callee, thus impacting call quality.

To test this, the caller made 1000 individual calls (through Tor) to every callee (seven locations), via VPN/SIP servers (three locations). Similar to the previous experiment, a total of 42000 calls were conducted. We observed acceptable quality with average PESQ 3.51 and average OWD 270 ms, slightly higher than the previous scenario.

Here again, we further tried to optimize the performance using shorter 2-hop circuits. We

Figure 3.5: V-Tor: CDF of PESQ for Caller Anonymity when server is separately hosted (Scenario II).

saw the average PESQ increased by 1, and the average OWD reduced by 210ms. The results are presented in Fig. 3.5.

The CDF of OWD for V-Tor is depicted in Fig. 3.6. Evident from the results, we observed PESQ above 3 and OWD less than 400ms in over 92% of the calls.

Figure 3.6: V-Tor: CDF of OWD variation for Caller Anonymity in both Scenario I and II.

3.4.2.3 Caller and Callee (two-way) anonymity (Scenario III).

There may be cases where both the caller and callee are positioned in censored networks. In such cases, they may connect via Tor to a VPN/SIP server placed outside their respective censors' jurisdictions. Their calls would be routed via their individual Tor circuits. We thus tried to observe the impact of such scenarios (traf c traversing two circuits) on the overall call quality as the additional network hops may increase OWD.

Similar to previous experiments, we varied the caller and callee locations across seven

countries, while the VPN/SIP server was distributed across three. Each caller initiated 1000 voice calls to a callee, resulting in a total of 2000 calls. For V-Tor we observed an average PESQ of about 3.2 with 81% calls above PESQ 3. However, the average OWD, as expected due to the increased network hops, was about 458 ms. This is slightly above the acceptable limit.

We thus tried to optimize performance by using shorter two-hop circuits. We hence repeated the above tests using two-hop circuits and observed a reduction of the average OWD to 396 ms. The results are shown in Fig. 3.7. In general, for such scenarios, regular three-hop circuits incur higher OWD, compared to two-hop ones. Hence, two-hop circuits seem a better choice for such cases. However, the results of our user study (ref Sec. 3.4.4) shows that users did not have any noticeable performance impact (due to the delay introduced) when both the users conversed via Tor for two-way anonymity.

Figure 3.7: CDF of delay for V-Tor setup when two-way Anonymity was achieved (Scenario III).

To summarize, in all the three scenarios over the public Tor network we observed good call quality (PESQ >3 and OWD <400 ms) in about 85% cases.

We use 400 ms as a threshold for good call quality following the ITU recommendation for international calls [15]. However, these recommendations also reported some user dissatisfaction even when OWD was between 300 ms and 400 ms. On analyzing the results, we found that more than 80% of calls had OWD below 300 ms. This has been further analyzed in detail in the Appendix. A.2.

3.4.2.4 Experiments involving popular apps

Next, we evaluate the performance of two popularly used VoIP apps, Telegram, and Skype, when running over Tor. Evaluating these apps would be bene cial from the usability point of view as most users generally use these apps for their day to day tasks. Hence using them for anonymous calls would be relatively easy, as they would not require installing V-Tor or M-Tor setups. However, the users might not be completely secure or anonymous, when using these apps as the app maintainers would know the calling parties.

In this experiment, we instantiated 100 consecutive calls using each of these apps (for both the setups) and computed their call quality. The average PESQ score was 3.8 and 3.54 for Telegram and Skype, respectively. Skype had more than 80% calls with PESQ score more than 3, whereas Telegram had approx 85% above 3. Overall, there was not much difference in terms of call quality between the two applications.

In our results, we observe that popular voice calling apps perform acceptably well over Tor.

3.4.3 Direct Calls

In the previous subsections, we have established that users in the majority of the cases would obtain good call quality when calls are performed over Tor. However, it would be interesting to compare the performance when calls are instantiated with and without Tor to understand the relative change in call quality.

Figure 3.8: Comparison of call quality in terms of PESQ obtained with and without Tor.

Hence, we performed direct calls involving both the setups. In the V-Tor setup, the calls were carried via VPN, and for M-Tor they were carried out directly using mumble. We instantiated 1000calls for each scenario described in earlier subsections. The results are shown in Fig. 5.6. As can be seen from the results, there is an expected relative drop in performance when moving from non-Tor to Tor setups. This is obvious as there are expected to be much fewer distortions on non-Tor setups. However, the performance of calls via Tor is good enough for the thresholds of performance metrics used in our study (OWD 400 ms and PESQ 3).

3.4.4 Users' Perspective

In our research, we conducted an extensive experimental study to adjudge the quality of voice calls via Tor. The study relied on two standard metrics, i.e., PESQ and OWD. These metrics are robust in judging the call quality. In fact, PESQ was established as a substitute for subjective tests (as already discussed in Sec. 2.1.2).

However, in all our experiments, we conducted calls with a maximum duration of 30 s. This is because, PESQ does not support the evaluation of calls whose duration is longer than 30 s [63]. Moreover, the PESQ score evaluation is a non-linear function, with respect to call duration. Hence, the mean of multiple samples of 30 s will not correspond to the PESQ of the actual combined duration call, as clearly stated by ITU [63]. Also, as per ITU recommendations [63] and PESQ [8] draft, a duration of 8 s - 12 s is sufficient to judge the call quality of the channel under test. However, one might argue that the real users' experience might be different for calls above 30 s duration, as these are not tested directly with PESQ. Therefore, we conducted a user study that involved human subjects evaluating the call quality with prior approval from ethical review committee (ref. Appendix A.2).

The user study involved 20 participants[5], which were randomly divided into five groups of four participants each. The groups were given individual sets of calls, each containing three different recorded calls. These calls were conducted over the Tor network, using setups described in Subsection 3.3.1

[5]The location of participants does not have much effect as Tor clients by default select relays in varied locations. However, our participants were spread across three countries.

The first call was 30 s long, and the remaining were two and four minutes long. We calculated PESQ for the 30 s call, and recorded OWD and jitter for longer duration calls (as PESQ cannot be calculated for calls longer than 30 s). The users listened to these recorded files and gave an Absolute Category Rating (ACR)[6] in the range of 1 – 5. We then calculated the Mean Opinion Score (MOS) by averaging the score given by all four participants of each group (results of which are summarized in Tab. 3.4). The users reported an average MOS of 4.25, 4.5 and 4.0 for the 30 s, 2 mins, and 4 mins calls, respectively. These results show that, in general, users reported good call quality over Tor.

Group No.	MOS 30s	2min	4min
I	4.5	4.5	4
II	4.5	4.25	4.5
III	3.5	4.75	4.25
IV	4.25	4.5	3
V	4.75	4.5	4.25

Table 3.4: MOS by different user groups for varied call length.

Comparing MOS with network attributes: We confirmed the aforementioned MOS values (obtained from users) against recorded performance metrics. For the 30 s call, we compared the MOS values with the corresponding PESQ scores. We observed an average MOS of 4.25, and a correlated average PESQ of 4.2, thus supporting our observations. Similarly, for the remaining two calls, we found that OWD and jitter were well within bounds for "good quality" (OWD < 400 ms and jitter < 30 ms). The average OWD was 278 ms, and jitter was about 24 ms.

All the above experiments involved the users listening to recorded calls. Thereafter, we went a step ahead, and asked ten users to converse daily via Tor for usual conversations. It must be noted that both the users connected via Tor, simulating Scenario III (two-way anonymity). Users reported a score for each of the calls. Call length varied from 1 min to a max of half an hour. This experiment was conducted for about 15 days, and users in the majority of the cases reported good call quality comparable to that achieved via popular VoIP applications rating an average MOS of 4.1.

The results in this user study further strengthen the claim about obtaining adequate call

[6] This is in accordance with ITU guidelines [64] for rating calls in subjective tests.

quality for anonymous calls over Tor.

3.5 Insights from Measurements

We now present explanation of our experimental results, along with other interesting insights we observed from these results.

3.5.1 Overall Performance Analysis

Computing PESQ involves the comparison of the original audio clip, as played out by the caller, with what is recorded at the callee. Effects of network delays, jitters, and losses, re ected in the recorded audio, are captured by this metric. Variations in network conditions, like increase in contentious cross-traf c, leads to an increase in the drops and delays, and thus negatively impacts the perceived quality (and thus PESQ). Besides PESQ, such contention also impacts other network performance metrics like OWD, RTT and available bandwidth.

Our overwhelmingly positive results, with PESQ above 3 and OWD under 400 ms in 85% cases are indicative of relatively low network contentions that can impact VoIP call quality. VoIP calls are encoded at low sending rates (< 128 Kbps) and thus require low available bandwidth.

Further, in 90.6% of our Tor circuits, we observed adequate available bandwidth (> 1 Mbps), as reported by iperf . About 95% of these 90.6% circuits supported calls with acceptable performance. This indicates that circuits with suf cient available bandwidth improves the chances of call obtaining good perceptual quality. This can be further understood by analyzing the 90.6% circuits where we measured over 1 Mbps bandwidth. We tabulate the frequency of these circuits, along with their corresponding PESQ scores. As evident from Tab. 3.5, for calls where we obtained good perceived quality ($PESQ > 3$), the frequency of occurrence of circuits where available bandwidth was more than 1 Mbps was also very high. Similarly, we also observed that for bad quality calls ($PESQ < 3$), the instances of circuits obtaining a bandwidth above 1 Mbps were relatively low.

PESQ	1-2	2-3	3-4	4-5
Frequency	6K	22K	162K	259K

Table 3.5: Variation of frequency of Tor circuits (>1 Mbps bandwidth) with PESQ of calls via them.

In general, we observe that with an increase in network contention, both call quality and available bandwidth decrease. This is evident from Fig 3.9. As incidences of high PESQ coincide with cases when the recorded available bandwidth is high, and vice versa.

Figure 3.9: Fraction of PESQ scores at different available bandwidths.

Impact on performance over time: Additionally, we also analyzed whether the performance of VoIP calls changed over time. For this, we randomly sampled 100 calls from our measurements and analyzed the distribution of calls based on quality. We repeated this experiment several times, and in each iteration, we observed that more than 85% calls always observed good perceived quality. This indicates that the overall distribution is uniform, and there was no observable change over time.

To ascertain the above observation, we plotted a graph incorporating results from our complete dataset. We draw a box plot consisting of variation in PESQ scores for different months. The whiskers in the plot represent 10 and 90 percentile values. The `x' represents the 99 and 1 percentile values, with `-' representing the min and max values respectively[7]. As evident from Fig. 3.10, we did not obtain any signi cant change in performance over time. Thus depicting that there was no observable change in performance over time.

[7]All subsequent box plots follow the same style as described.

Figure 3.10: Change in performance over time.

3.5.2 Performance Dependence on Types of Relays

We next analyzed whether different type of relays (viz., Guard, Middle or Exit) along with their frequencies of occurrence in circuits, had any observable impact on the VoIP call quality over Tor. To begin with, we distributed the relay frequency into three bins — low (< 100), moderate (150 – 200) and high (> 500). Then, from each group, we randomly selected a few entry, middle, and exit relays (around ten each) and manually inspected the PESQ scores of the calls involving them. For each type of relays (in all frequency bins), we observed PESQ scores ranging from 1 – 4.5, with the majority of them being over 3. Overall, there was no obvious difference in the distribution of calls with different PESQ scores. This observation indicates that PESQ neither depends on any specific type of relay, and nor on their corresponding frequency of occurrence.

To further ascertain our claims and to obtain a comprehensive picture for all our measurements, we plotted the distribution of PESQ scores corresponding to all guard, middle, and exit nodes. The box plot for exit nodes is shown in Fig. 3.11. As evident, the PESQ values show no dependence on the frequency of occurrence of Tor relays. Corresponding to both less (< 150) and more frequently appearing relays (> 500) we observed PESQ above 3 for a large fraction of calls. The trend was similar for guard and middle relays.

Tor churn analysis : Tor relay churn is defined as the rate of relays joining or leaving the network from one consensus to the other (according to Tor Metrics [52] and Winteret al. [65]). To that end, we determined whether Tor relay churn had any impact on our results. We calculated

Figure 3.11: PESQ of individual calls vs frequency of exit relays.

the monthly relay churn[8] for the entire duration of our study (ref Fig. 3.12). It is evident from Fig. 3.12 that relay churn was low (an avg. of 0.2%). Further, we also observed that more than 85% of our measured VoIP calls had acceptable quality. Our overall results thus indicate that such a low value of churn has an insignificant impact on call quality.

Figure 3.12: Tor relay churn for the entire duration of our study, i.e., 12 months.

However, one may argue that a large proportion of exit relays might not support the default ports used by VoIP applications (SIP or Mumble) involved in our study. This might lead to bias in the churn analysis, as we may be considering only a small fraction of all the available exit relays that support VoIP calls. Thus, in all our experiments, we configured our VoIP servers to listen on port 80 and 443 as these are generally allowed on the majority of the exit relays.

Further, we studied the prevalence of exit nodes supporting VoIP applications by default and

[8] The Tor consensus is updated every hour; in one month it is updated around 720 times. The individual box plot corresponds to the change in consensus of these 720 values. Thus, there are 12 box plots each corresponding to a different month.

compared them with those who support port 80 and 443. For this, we analyzed the exit relays in the Tor consensus files for the duration of our study. First, we identified the number of exit relays allowing the default VoIP applications ports (64738 for M-Tor and 1194 for V-Tor), and also the ports used for our study (80 and 443). Once we obtained the number of exit relays allowing these ports, we added their corresponding bandwidth weights and divided it with the cumulative bandwidth weight of all the exit relays. This gave us a normalized value of the bandwidth weight of these exit relays (that allow the aforementioned ports). The results are depicted in Fig. 3.13.

Figure 3.13: Analysis of the bandwidth coverage of exit relays when using default VoIP application ports (64738 for Mumble and 1194 for VPN), compared to when port 80 and 443 were used.

As evident from the results, almost all the exit relays supported 80 (99.9 %) and 443 (99.6 %). On the other hand, 90.1 % and 83.4 % relays allowed the default VPN and Mumble ports, respectively. Thus, it indicates that our results did not rely on some specific set of exit relays. Moreover, a great fraction of exit relays (> 80%) would support our VoIP applications by default. These results, along with the churn analysis conducted previously, depict that Tor relay churn did not have any significant impact on our results.

Overall, the analysis in this subsection clearly indicates that the performance was not dependent on any specific type of relay. The apparent independence is because, for a call to be of good quality, we require less cross traffic contention, and bandwidth of about 200 Kbps for the entire call duration. The prevalence of such an ecosystem naturally on Tor has already been argued in Sec. 3.5.

3.6 Discussion

In this section, we address concerns like how VoIP performs over Tor when using bridges, using different codecs, etc.

Call quality over Tor bridges: Tor bridges [66], are unadvertised entry (guard) relays, whose information is closely guarded. They are conservatively distributed either through BridgeDB [67], or covertly via out-of-band means (e.g., emails). Censors may identify and filter Tor traffic using entry node IP addresses and (or) port numbers. User residing in such networks may use bridges to access other Tor relays and set up the circuits.

Hence, to study the impact (if any) when using a bridge, we performed 1000 measurements, using the V-Tor setup. In this experiment, we used a bridge to connect to a new Tor circuit each time. The average PESQ score was about 3.7. In about 85% calls, we observed good performance (PESQ $\geq$ 3.0 and OWD $<$ 400 ms). It must be noted that, we restricted our measurements to a single bridge node as they are scarce.

Impact of codecs: Codecs define the way audio is encoded and decoded to transmit them as packets. Hence we measured the impact of different codecs on call quality. The calls for this experiment were conducted over the real Tor network. We used some popular codecs for our

Figure 3.14: CDF of PESQ scores when different codecs were used.

evaluation, viz., Opus, G.711, Speex, and GSM. The codecs were selected with the help of configuration changes in the Freeswitch SIP server. We conducted 1000 calls corresponding to each of these codecs and measured their PESQ scores (ref. Fig. 3.14). We observed that the

fraction of calls with PESQ above 3 were roughly equal (85%) in all the cases. Thus any of the tested popular codecs could have been used for initiating good quality calls. However, GSM and Speex are lossy codecs, compared to lossless codecs like G.711, and thus may not provide very high quality calls (PESQ 4) [68]. Besides, such lossy codecs encode at lower bitrates compared to the lossless ones.[9] Thus, they may be used for Tor circuits with low available bandwidth to receive adequate call quality.

Thresholds for PESQ: The MOS scale proposed by ITU delineates perceived performance through sharp integral differences (e.g., two corresponds to "annoying but usable" and three corresponds to "fair"). However, it does not extend such annotations for values in-between. Further, upon listening to a few calls manually where PESQ was between 2.8 and 3.2, we observed no audible differences. Moreover, recent studies [68] indicate that humans may be unable to differentiate the quality of samples, when the difference between their MOS scores is less than 0.4.

In our study, we conservatively selected a PESQ over 3 as "acceptable", and anything below it as not (in accordance with ITU). We also observed a significant cases which may be categorized as "annoying but usable." Finally, we also observed about 8% with PESQ under 2. ITU classifies these as "totally unacceptable." Thus, we believe that actual user experience may be even better than what our study reports (as evident from the conducted user study in Subsec. 3.4.4).

Coverage of Tor relays: We recorded and analyzed Tor circuit information for all our experiments. We now present some interesting insights we observed from this analysis. A total of about 600K Tor circuits were created during our study.[10] These circuits involved a total of 650 unique Tor relays. Prior research (Rizal et al. [6]) reportedly used only about 298 relays that too restricted to Europe.

Selection of caller-callee endpoints: Throughout our experiments, we used endpoints either as cloud hosts or in-lab machines. All of these were sufficiently provisioned for a voice call

[9] GSM and Speex encode at bitrates <40 Kbps whereas G.711 encodes at 84 Kbps.

[10] The number of Tor circuits are slightly higher than the total number of experiments as the two way anonymity experiments involve creating two Tor circuits for a single call.

with more than adequate bandwidth at the caller-callee end. However, one might argue that our results might be biased as all our endpoints may have suf cient bandwidth for VoIP calls. But in general, the performance bottleneck was introduced at the Tor relays. As already described in Sec. 3.5, there were a signi cant number of Tor circuits which observed a low bandwidth (1 Mbps). Hence, even if the endpoints are well provisioned, it does not bias our results, as mostly the Tor relays were the bottlenecks.

3.7 Conclusion

Real-time anonymous VoIP calls are of interest to privacy and anonymity conscious citizens, whistle-blowers and covert reporters etc. However, existing research on performance evaluation of VoIP calls over Tor is not comprehensive. They contraindicate transporting VoIP packets over Tor and thus favored novel architectures to support anonymous calling. Moreover, there does not exist a functional system to achieve the same. Additionally, the costs involved in recruiting volunteer operated relays (like Tor) along with users, and managing such a system, might outweigh the bene ts.

Thus, it was essential to identify the causes of poor voice call quality over Tor by observing how the interplay of various network attributes (RTT, available bandwidth etc.) impacts VoIP quality. To that end, we conducted a longitudinal study (spread across 12 months). It involved extensive testing of about half a million voice calls over Tor, including a user study, using various Tor circuits, peer locations, popular apps etc. To our surprise, in over 85% cases, we observed good performance (PESQ 3 and OWD 400ms), with only under 8% cases which were totally unacceptable (PESQ 2). The results of the user study also corroborate our ndings. Our study is the rst to demonstrate that anonymous VoIP calls are indeed possible using Tor.

Although, Tor can be used to perform anonymous VoIP calls along with its traditional usage of accessing web content, but various adversaries attempt to block Tor and in certain instances succeed in doing so. Thus, there is a need to build systems that are dif cult to block in practice. Hence, in the next chapter we attempt to build a decoy routing system as it provides effective blocking resistance.

Chapter 4

SiegeBreaker: An SDN Based Practical Decoy Routing System

4.1 Introduction

Free speech on the Internet is a political battleground. On one hand, several governments claim that their sovereignty extends to cyberspace, and that they should control the content seen (or written) by their citizens. On the other hand, as censorship is frequently used to silence the opposition [70], access to free speech online is considered a human right by the United Nations [1]. This tension between censors and free speech advocates has led to an "arms race": users use proxy based services (e.g., VPNs and Tor [71]) to access content on the Internet, and in response, censors design stronger counter measures [72]. Even the best anti-censorship solutions, such as relay addresses [73] and protocol obfuscation [74, 75], are temporary measures [76].

Decoy Routing [9, 77, 78] is an attempt to break out of this "arms race", through the use of "smart routers" that double as proxies. DR clients – residents of censor countries, who need to access a blocked site – start by sending packets addressed to allowed sites (the overt destination (OD)). These packets appear innocuous, but actually carry a covert cryptographic message, carefully chosen to make it hard for the censor to distinguish the packets from regular TLS messages. Once safely past the network boundaries of the censor, these packets are identified by

DRs that hijack the client's connection and redirect that flow to the covert destination (CD), the actual censored site the client wishes to contact. Even if a DR is discovered, it is hard for the adversary to "Route Around Decoys" [79, 80] without losing connectivity from a major part of the Internet.

Interest in Decoy Routing has led to a rich body of prior proposals [76, 77, 81, 82, 83, 84], but not led to the development of practical and efficient DR systems. There are several challenges that must be addressed before DR becomes widely adopted in practice.

Firstly, Internet applications perform poorly when using existing DR systems. As reported in previous research, DR clients experienced overall low throughput [85] and high latency [84]. This is because existing solutions are implemented on commodity servers (rather than real routers); they cannot serve traffic at line rates while at the same time matching network flows using cryptographic operations. On the other hand, it is virtually impossible to simply port existing solutions to conventional routers (which are not designed to be reprogrammed).

Secondly, existing proposals assume (unfounded) mutual trust between third-party DR operators (analogous to Tor maintainers), and the "friendly" ISPs (who simply host the DR). These proposals involve DR operators inspecting all flows traversing the ISP. Thus, DR operators may compromise the privacy of non-DR flows passing through "friendly" ISPs (e.g., by recording their packet headers, metadata, or even content), breaching the trust extended to them.

A third issue is identifying the networks where DRs may be positioned (to intercept requests from several users), and providing incentives for these network operators to deploy DR. Finally, a fourth issue is defending against various attacks (traffic analysis) from different adversaries.

The DR community has begun research efforts to address the third [79, 80] and the fourth issue [83]. However, the first two problems – assuring performance for regular DR flows, without compromising privacy of non-DR flows – remain open challenges to practical DR.

We thus propose the first practical system to answer such concerns: SilentBreaker (SB), a Software-Defined Network (SDN) [86] based, efficient, and privacy preserving DR system. It works as follows:

- The users of SB signal the Decoy Router through an email, addressed to an email ID associated with the SDN controller (that manages the SDN switches). The email bears encrypted information, that can only be decrypted by the SDN controller; at a later step, this information authenticates a legitimate DR client to the controller.
- After deciphering a DR request, the controller re-con gures the switch (the Decoy Router), to redirect all subsequent TCP packets, only the SB client – OD connection, to itself.
- These TCP packets carry covert information (indicated in the previous email message) that authenticate a legitimate DR request, without requiring the inspection of all other ows traversing the DR (as required in all previous proposals). This drastically reduces the inspection load on the controller.
- After authenticating the client, the controller re-con gures the switch to divert subsequent SB client – OD packets (i.e. only the DR ows) to a third-party secret proxy (SP).
- These subsequent packets help covertly establish a fresh session key with the SP, and expose to it the URL of the CD. The SP then communicates with the CD and sends the responses back to the client, ensuring reliable and ef cient delivery of packets.

Overall, the protocol ensures that the third-party DR operator managing SP, observes only the DR requests, without having any knowledge of the non-DR ows, thereby preserving their privacy.

SB proves to be a practical and privacy-preserving DR system. It builds upon the idea that DR is not a monolithic task. Thus, it distributes the tasks of DR over three loosely-coupled, synergistic modules: (1) The SDN controller, that orchestrates the switch(es), independently identi es DR requests and re-con gures SDN switches on-the- y. (2) The switch, which primarily forwards packets between the client and the SP, bypassing engaging the controller. (3) The SP, which focuses on communicating with CD, along with other tasks to ensure secrecy and reliability. Such a design not only leads to a massive improvement in performance, but also ensures that the DR system (and its operator) only sees those few ows which are associated with DR requests.

Our major contributions can be summarized as follows:

1. The prototype of SB, a practical and efficient DR solution, which achieves multiple goals:
    - Modularity. DR requests are handled by three loosely coupled entities – the SDN controller, the switch and the SP – reducing computational load on each.
    - Privacy preserving signalling. Using the above modular design, SB inspects a small fraction of all traffic and serves only the DR requests.
    - Dedicated hardware. SB uses SDN switches (instead of commodity servers proposed in previous designs) to divert flows efficiently at line rates.
    - Reliable communication. SP provides reliable and efficient client–CD communication, by preserving the characteristics of TCP like behaviour (when there is no actual TCP connection between them).
2. Comprehensive performance analysis of SB, on real testbeds (involving a commercial SDN switch), shows near-native (TCP) performance for both web surfing and bulk downloads, even in the presence of 50K parallel flows or 15Gbps network-traffic on the SDN switch. In ideal conditions, SB is capable of serving individual client requests at line rates of 10Gbps.
3. Thorough security analysis of the protocol, showing how SB is resilient to a variety of advanced attacks (such as forced asymmetry [87]).

Finally, we note that existing efforts to deploy DRs [78, 79] could be very costly — even modest proposals [80] involve replacing about 11,000 routers across 30 ASes with DRs. However, in an SDN-based AS, any switch could double as a DR on-demand, without disrupting regular routing. This solves the problem of DR placement within an AS. We further demonstrate how the client could use inconspicuous protocol messages to covertly signal the controller, and help identify the best switch for installing flow rules, in Subsec. 4.4.4.

4.2 Related Work

Existing DR systems follow the general pattern described in the previous section, but vary in their choice of secret handshake, side channel, and whether they take additional measures to suppress unwanted messages from the OD to the client (which might make the censor suspicious). We

now provide a brief survey of existing DR systems.

Karlin et al. [9] proposed an initial DR architecture, called Curveball. However, the first-generation DR implementations include Cirripede [76] and Telex [77]. Cirripede uses the TCP Initial Sequence Number (ISN) field as a covert channel for registering the client to the Registration Server. This server provides a fixed set of sequence numbers that the client may use in subsequent connections. In Telex, the client embeds a cryptographic tag (created using the Telex station's public key) in the TLS nonce field of the Client Hello message to OD. Seeing a tag, the Telex station establishes a shared secret with the client (using Diffie-Hellman key exchange), which is subsequently used to covertly proxy traffic to the filtered sites. Telex requires both sides of the client–OD communication be intercepted; it is, thus, fragile where protocol messages (and acknowledgements) between the client and OD traverse asymmetric network paths. The recent systems by Bocovich et al. [83, 88] and by Nasr et al. [82] build upon the signaling architecture of Telex and Cirrepede respectively.

Tapdance [81], an improvement over Telex [77], uses chosen-ciphertext steganography to signal the DR in an incomplete HTTPS request destined to the OD. Further, this request eliminates the need for inline blocking of the OD [77] and the requirement to intercept both request and response, as the OD never responds to such requests. Recently, Frolov et al. [85] partnered with a small ISP (and a university) to deploy Tapdance for practical usage. Their experiments reveal that incomplete HTTPS requests, often terminate in roughly 30 seconds, requiring frequent re-negotiation. This issue, coupled with the average Tapdance client throughput of 5 KBps, makes it unsuitable for accessing much of the web content. This reveals how existing DR solutions struggle with real-life deployment. Even on a very small ISP, with three uplink routers (compared to thousands, for a backbone ISP like Level-3 or Cogent [86]), there are serious performance problems. Deploying at scale would require a much larger infrastructure, capable of monitoring millions of flows, identifying the DR requests, and providing them proxy service.

Bocovich et al. [83] took a different approach where they address latency-based and website-fingerprinting attacks. They achieved resistance to such attacks by sending the covert content in the leaf elements (images etc.) of the webpage served by OD. The client maintained an active connection with OD, and only the leaf content in the response of the OD was replaced with covert

content. Further, as Slitheen still required analysing content in both the reverse and forward path, the same authors in their subsequent paper [38] proposed Gossip, a new protocol which allows all existing routing symmetry-dependent DR systems to work despite path asymmetry. Another approach was proposed by Nasr et al. [82], in which the downstream content, i.e. replies from OD, are used to detect DR, rather than the (upstream) requests. This helps in the mitigation of strong RAD [37] attacks. However, both Slitheen and Waterfall were developed and deployed on commodity servers, which are not practical to scale in an ISP.

In brief, existing DR solutions have not succeeded in utilizing commodity routers to match/decrypt messages and serve as proxies on-demand at an ISP scale.

4.3 SiegeBreaker vs Prior Research

Properties	Siege-Breaker	Curve-ball	Telex	Cirripede	Tap-Dance	Rebound	Slitheen	Waterfall of Liberty	Secure Asymmetry
Implementation on real routers	●	○	○	○	○	○	○	○	○
Real Internet Workload Performance	●	○	○	○	○	○	○	○	○
Privacy preservation of non-DR clients	●	○	○	●†	○	○	○	●†	○
Handling asymmetric routing	●	○	○	●	●	●	○	●	●
Resistant to replay attacks	●	●	●	●	○	●	●	●	●
Resistant to latency analysis	○	○	○	○	○	●	●	●	●
Resistant to website fingerprinting	○	○	○	○	○	○	●	●	●

Table 4.1: A comparison of different features of existing DR schemes. ● - feature supported, ○ - feature unsupported & † - feature supported with some assumptions.

We now describe how SB significantly differs from prior attempts at DR (summarized in Tab. 4.1).

- Implementation On Real Routers: Existing DR proposals use commodity servers as Decoy Routers – either by using the server as a router [78], or by attaching the server to one (through a wiretap) [31]. A server running a non-realtime OS, that has to inspect all the network connections to identify DR requests, and also proxy such requests to the intended CDs, is a clear bottleneck. SB uses hardware (SDN) switches instead of commodity servers, allowing it to process ISP-scale traffic at line rates. Although, Tapdance was deployed in an ISP in

2017 [85], it did all the processing on a server attached to a router. It is important to note that, unlike SB, their implementation was not on the router itself.

- Real Internet Workload Performance: Among existing DR prototypes, Frolov [85] demonstrates the practical use of Tapdance, inside a university and within an ISP. However, they reported overall low client throughput (5 KB/s), and precludes testing with large files. Other approaches have restricted themselves to downloads of a few (less than ten) small (<1 MB [78]) webpages, in a controlled laboratory environment [77, 81], and even so, report over half a minute to download the home pages of popular sites (1.5 MB in size) [84].

  SB's prototype is extensively tested for large files (up to 1GB) and high client link bandwidths (up to 1Gbps). Even under heavy cross-traffic conditions, and with CD and OD over the Internet, SB's client requests can be served at line rates of 1 Gbps. SB's performance was comparable to direct TCP downloads (near-native) in every scenario we tested.

- Privacy Preservation of non-DR clients: All prior approaches require the DR to inspect both the DR and non-DR flows, either by analysing TLS or TCP SYN packets.[1] This allows the DR volunteer hosts (e.g., Telex/Tapdance/Slitheen station) to also see what sites other non-DR clients are visiting. This may be considered an unwelcome intrusion on their privacy. We introduce a novel signaling scheme that allows the controller to easily isolate DR requests (i.e. re-configure the SDN switch to selectively forward DR packets to the SP). Not only is this efficient (reduces load on the system) but also privacy-preserving (i.e. removes the need to inspect non-DR flows.

- Handling Asymmetric Routing: Routing of traffic in the Internet is not always symmetric, and most DR systems need to intercept client–OD traffic in both directions [2, 83, 77]. Some systems survive route asymmetry through special schemes (such as a gossip protocol) [68]. However, a few DR systems [81, 84] have no trouble with asymmetric routes; SB is one of these robust systems.

- Reliability and Efficiency: In all DR systems, the DR hijacks client–OD TCP connections.

[1] In Cirripede and Waterfall, if we assume that the registration server is maintained by the ISP's operator, then they could inherit this property. Additionally, how SDN implementation can be used to complement existing DR schemes (with this feature) is elaborated in B.2.

DR traf c between the client and the CD lacks the reliability and ow-control guarantees of regular TCP connections, and is thus easily impacted by background Internet cross-traf c. SB emulates TCP's message transmission mechanism, to ensure reliable, in-order and ef cient delivery of packets to the client. This has signi cant positive impact on clients' performance.

- Limitations of SiegeBreaker: SB is not resilient to latency based [87] and website nger-printing attacks [89]. Moreover, we have not tested SB in an ISP. (Such testing has only been demonstrated for Tapdance.) One may ask that in an SDN-based ISP, on which switch should the redirection rules be installed? As a solution, we have proposed a novel scheme (in Subsec. 4.4.4) using which a client can covertly signal the controller about the appropriate switch. Our proposed scheme, however, may be susceptible to ngerprint attacks by a determined adversary. In such scenarios, the controller would install rules on all the switches. But, due to the associated timeout, these rules would be eventually purged out from all the switches, except the one that would actually intercept the DR ows. This is explained in detail in Subsec. 4.4.4. However, we acknowledge, it still increases overhead on the system in terms of installing the unwanted rules whenever a user requests DR service.

4.4 System Design

4.4.1 Threat Model

We consider two types of adversaries for our system. Our primary adversary is a powerful nation-state that censors its citizens' Internet access. Not only can adversary censor regular Internet traf c using any known technique (DNS injection, IP address/URL ltering, DPI, etc.), it may also covertly monitor network packets. Further, in an attempt to detect/disrupt circumvention systems alone, it may also subtly manipulate network packets that cross its boundaries, without disturbing regular Internet users. Also, for the functioning of SB, we assume that users have access to some form of regular email service (either webmail or SMTPS).

We also consider a secondary (Byzantine) adversary — a motivated individual or group, who intends to disrupt normal network routing. This adversary sends forged messages, aiming to

Figure 4.1: SB Protocol: The numbered arrows correspond to the various protocol messages, exchanged between the client, the SDN switch (acting as DR), the OD, the SP, and CD, as described in Subsec. 4.4.2.

install unwanted rules in SDN switches, or in any other way disrupt network routing through attacking the SB infrastructure.

4.4.2 SiegeBreaker Protocol

Pre-requisites

The SB architecture assumes a network of centrally controlled SDN switches, con gured to function as regular routers (i.e. they forward packets based on their destination IP address). It is assumed that the client *a priori* has access to the following entities:

1. Public key (PUB_{CO}) of the controller's 2048-bit RSA public-private key-pair.
2. The 256-bit DH public exponent g^y of SP.
3. The controller's email ID.

Protocol Description

SB's DR protocol comprises roughly of three phases, viz. bootstrapping, hijacking and proxying. We now present a step-by-step walk-through of a DR session, with reference to Fig. 4.1.

Bootstrapping (Phase I):

1. The client initiates DR by sending an email to the controller's email address. The payload of this email contains the DR request. It is encrypted with the public key of the controller, and thus not understandable to the client's email provider or the adversary.

2. The controller, on receiving an email, attempts to decrypt the content using its private key, and searches for four fields embedded in it: (1) The magic word "SIEGE", signifying DR request. (2) A TCP Initial Sequence Number (ISN) that the client will eventually use while handshaking with the OD. (3) IP addresses of client and OD. (4) A DH exponent public value (g^x), the client derived using privately chosen number x.

3. On successfully identifying a DR request, the controller installs a rule on the SDN switch, so that all client–OD packets are redirected to it (the controller). This inspection rule will expire after a hard timeout.

4. Subsequently, the client initiates a fresh TCP handshake with the OD. The SYN packet of the handshake bears the same ISN, as the one sent in the initial email.

   This packet is redirected to the controller (as per the rule installed in step 3), who checks that its ISN matches the ISN specified in the initial email and forwards it to OD. Thereafter the client completes the TCP connection, and initiates a TLS handshake with the OD, negotiating the session key (Client-ODSessionKey).

   This initial ISN match is the vital step that authenticates the client and identifies the DR requests, without the need to inspect any other flow.

   In case the client's SYN packet does not have the expected ISN, the controller continues to poll the packets of the client–OD flow until the timeout period of inspection rule lapses, causing it to expire.

Hijacking (Phase II):

5. After completing the TLS handshake, the client sends a TLS data packet, carrying the HTTP GET request (with OD as the host field), encrypted with the session key that was negotiated with the OD.

   The inspection rule is still in place, so this packet goes to the controller which replaces the payload of the TLS data packet with the client's DH exponent g^x (which it received in step 2) and the OD IP. It then forwards this packet to the SP.

   Further, the controller also updates the redirection rule on the switch, to divert subsequent packets of the flow to the SP. This rule matches the client-OD flow based on the client's source port, client IP and OD IP as seen in the TCP SYN packet. We denote this rule the *SP redirection rule*, and associate an idle timeout with it. Once a rule is installed for a specific client-OD pair, no other rule, for the same pair, would be installed until the said rule times out. Thus, clients behind a NAT firewall who try to use the DR service, would require selecting different ODs. This ensures that all such clients can simultaneously access the DR services. This is explained in detail in Subsec. 4.6.1.

6. The SP, on receiving the first TLS data packet, assumes that the client intends to use DR.

   SP terminates the existing client–OD connection by sending a RST packet (spoofed as client, with correct sequence and acknowledgement numbers) to the OD[2] (obtained in step 5).

   Further, the SP derives a pre-master key (PMK) (g^{xy}) using client's g^x (received via the first TLS data packet) and its own DH private number y. The SP (spoofing as OD) then crafts a TLS data packet carrying a random nonce, Nc_s, along with its HMAC computed using PMK.

   To the censor, this appears to be a regular TLS data packet, carrying random bits. The client, however, treats these bits as a nonce Nc_s and calculates its HMAC. It derives the PMK (g^{xy}) using the private DH number x, and the publicly known g^y. Successful verification of the HMAC allows the client to confirm that the DR request was successful.

Master-key derivation: Following a successful validation of the nonce Nc_s, the client selects its

[2] The RST causes the client–OD connection to be forcefully terminated, without any responses being sent to the client that may raise the censor's suspicion.

own nonce N_C. Using these nonces and the PMK, the client derives a 6-tuple key (K_C, H_C, IV_{SP}, K_{SP}, H_{SP}). IV_C, K_C and H_C are the IV, cipher and HMAC keys for the client, while IV_{SP}, K_{SP} and H_{SP} are those of the SP. These keys are generated using a Psuedo Random Function (PRF) involving SHA-256 hash function (Referring to TLS v1.2 [90]).

Proxying (Phase III):

7. Thereafter the client crafts a TLS data packet carrying the CD URL (encrypted with key K_C using 256-bit AES in CBC mode), and a HMAC of the URL and N_C (using HMAC key H_C).

   The client sends this packet, addressed to the OD. En-route the packet encounters the switch that redirects it to the SP.

8. SP, on successful reception of the aforementioned TLS data packet, extracts the nonce N_C. Using N_C, N_S and the PMK, SP also computes the same 6-tuple as the client. Upon successful HMAC validation, the SP decrypts the URL (using key). Finally, SP connects to the CD and requests data. The SP focuses entirely on serving DR requests, without the burden of identifying them from among all ows via costly cryptography operations

9. CD serves the requested content to SP.

10. The SP encrypts these responses with key and signs them with HMAC key H_{SP}. It then sends them back to the client, spoo ng the source IP address of the OD (maintaining the state of Client–OD connection). This keeps up the pretense, to the client's censor ISP, that the client is communicating with the OD. The same session, can also be used for requesting content from various other CDs (before the idle timeout expires).

This walk-through raises several questions, such as why email was used as a covert channel? Among multiple switches, exactly on which switch(es) the controller would install the inspection/redirection rules? What rule timeout values should be selected in our system? Given that there is no true TCP session between client and CD, how do they compensate for dropped packets etc. We now explain the design decision of using email along with how SB would select switch(es) to install the rules and discuss the remaining concerns in Sec. 4.6.

4.4.3 Improved Covert Signalling

Existing DR implementations use TCP initial sequence numbers (ISNs) [78], the ClientRandom field inside TLS client hello [77], and the encrypted body of an HTTPS request [81] as covert signaling fields. These schemes require analyzing either all SYN packets [78, 82] or TLS flows [81, 77], as the covert patterns are embedded in the packets sent from the DR client to the OD. Inspecting large volume of traffic for identifying the embedded covert patterns, in innocuous looking packets, make it difficult for the censor to detect DR requests. However, these well thought signalling schemes pose new performance challenges for the DR users.

All existing DR systems rely on commodity servers to analyze the large volume of traffic [81]. This may pose as a performance bottleneck as NICs of these servers are not built to handle such high speed traffic. Further, these systems analyze all flows, including the non-DR ones, to detect DR requests. This may inadvertently compromise the privacy of non-DR users. In principle, SB can easily adopt any of the existing signalling schemes. However, SB aims to reduce the burden of analyzing all flows, while still retaining the same standard of unobservability (from the censor), compared to existing architectures. Additionally it also aims to preserve the privacy of non-DR users.

Thus, our covert signaling indicates the controller about the upcoming DR flows, such that it inspects only the potential DR traffic, reducing the amount of traffic to be analyzed. For this, a client can rely on any out-of-band channel (e.g., IMs, SMS and email etc.). In our present implementation, SB clients use an email to covertly signal the controller. This email contains client's source IP, OD's IP and the TCP ISN that would be used by the client in the subsequent DR request. Thereafter, the controller installs inspection rule for only flows with indicated client IP and OD IP address (i.e. a potential DR flow), disregarding the rest of flows. For potential DR flows, the controller matches the ISN in the SYN packet, with the one indicated in the email. A successful match, confirms a DR flow, and the SP redirection rule is installed on the switch.

Thereafter, only the DR flows are diverted to the SP, which has no knowledge of non-DR flows. This final step helps preserves the privacy of non-DR users from the DR volunteers maintaining the SP.

It must be noted that, as an alternative to ISN, the TLS ClientRandom may be sent in the email. The controller would thus con rm DR ows by matching the content of TLS ClientRandom in ClientHello (of potential DR ows), to the one received in the email.

On the use of email: In the past, email has been successfully used as a covert channel for transporting censored content [92, 93, 94, 95]. Similar to such efforts, we also assume that users have access to some form of TLS supported email (either webmail or regular SMTPS). However, in our design email is not directly used as the data channel; it is a control channel, used for signalling the SDN controller.

A determined adversary may attempt to snoop and block all emails destined to controller's mail ID. Since the communication is TLS encrypted, our approach is not vulnerable to wire snif ng adversaries. They are only vulnerable to a very powerful adversary who can assume control over the clients' email services. However, even when the email provider is controlled by the adversary, one can issue unique email IDs to each individual client, as proposed by Houmansadr et al. [94]. This makes it impractical for the censor to learn and block all such email IDs.

Moreover, an adversary may cripple the system by randomly delaying suspicious emails, thereby delaying the installation of the inspection rule. Oblivious DR clients' packets would thus reach the switch before this rule is installed. Hence, they would go undetected by the controller, failing to avail DR service. Further, our email body contains four elds and thus may have a xed length, thereby raising suspicion. However, we make it dif cult for the adversary to identify such emails E.g., we pad the email body with random nonces.

However, a determined adversary may nd some other identi able patterns in the encrypted (registration) email. In such cases, we can resort to using other carriers, such as images, PDFs etc., to steganographically encode the covert signal. as already used in existing anti-censorship systems [96, 97]. However, our current implementation does not incorporate this feature.

In the extremity, an adversary may attempt to delay every email (both DR and non-DR). In such cases, SB may be tailored such that the controller sends an acknowledgment email back to the client. Reception of this email, con rms the availability of DR services to the client. As a

drawback, this may incur a delay in the bootstrapping phase. However, in the presence of such disruptive adversaries, the email con rmation assures availability, albeit at cost of such delays.

4.4.4 Auxiliary Signalling Scheme

In SB, the controller needs to install OpenFlow rules to redirect the packets of a client–OD ow to itself or the SP. However, as a network will have multiple SDN switches, SB poses a new problem: how can the controller know which switch to install these rules on?

Figure 4.2:Ping packet assisting the controller in selecting the appropriate switch to install redirection rule.

The simplest solution would be to install switches in the SDN with the redirection rules. But intuitively, this approach maybe expensive considering rule installation time and the nite memory of the switch. (Note that a rule is added for every client–OD ow. They cannot be aggregated.) We therefore propose a secondary signal to achieve the same. The client sends a crafted ping packet as a covert signal, to help the controller identify the switch it should install the DR rule on (ref. Fig. 4.2). The overall approach is explained as following:

1. Switches are bootstrapped with a rule that forwards all ping packets to the controller.
2. While crafting the email to the controller, the client adds three more elds: (i) IP-Identi er (from IP header), (ii) ping sequence number and (iii) ping identi er (from ICMP header).
3. Next, the client sends a crafted yet innocuous looking ping packet to the OD (with three elds described in previous step). En route to OD, decoy AS's SDN controller receives the

ping packet and compares the aforementioned three fields to the one sent in the initial email. A successful match confirms that the ping packet is indeed from the authenticated client.

4. At this step, the controller selects the switch from which it received (and validated) the ping packet. Thus the redirection rules are installed on this switch, assuming the subsequent client-OD packets shall also arrive at the same switch.

Selection of three fields: The goal of the controller is to correctly associate the ping packet to the legitimate DR client that initially sent the email. An adversary may try to brute force this association by sending fake ping packets (spoofed as DR client), guessing the three header fields. In case the adversary succeeds, and the route taken by packets of the adversary differ from that of the client, the controller would end up installing a rule on a switch which might not appear on the client–OD route. This renders DR service unusable for intended DR clients. With our architecture, adversary needs to correctly identify 48 bits (IP-ID, ping seq. no. and ping ID, each being 16-bit field) making the brute force attack highly impractical (requiring 2^{48} 281 trillion ping packets per client–OD pair.)

Indistinguishable Pings: We only rely on the IP-ID, sequence number and identifier fields of the ping packet for covert signalling (which are by default random numbers). We keep rest of the fields (and payload) identical to pings generated by standard OSes. Thus our ping packet appear indistinguishable from regular ping packets generated by popular OSes.

However, we acknowledge that sending pings for covert signalling, before accessing DR, may be classified as a pattern (i.e. a fingerprint attack). We try to make it difficult for the adversary to detect such patterns by making these ping packets innocuous. Thereafter, we introduce random delays between ping packets and the corresponding DR requests.

Our auxiliary signalling mechanism is prone to fingerprint attacks. However, SB can also function without this scheme. In the absence of this scheme, the rule would be installed on all the switches. But, these would be automatically purged, after a short duration (due to timeouts) from all the switches, except the one which later identifies the actual SB clients packets. However, the installation of rules on all the switches may incur an additional memory overhead. This may eventually impact the total number of clients that can be supported by the switches. We thus

analyzed the number of clients that can be supported simultaneously by the SDN switch. The HP3500yl SDN switch, that we used for evaluating our system (described ahead in Sec. 5.4), supports 74K table entries [98]. More advanced SDN switches [99] e.g., HP10500 series can support 1; 152; 000 table entries (i.e. 1; 152; 000 open ow clients). Such switches can transport a total of about 8:8 Tbps [100] traf c. Moreover, the SDN switches that we used, do not impede performance when the number of clients (or entries in the table) increases. These are built for commercial deployment within large ISPs, and can easily handle a large volume of traf c without any slowdown. Thus, we believe that our proposed system will be suitable for deployment.

4.5 Experimental Evaluation

In this section, we describe the experiments conducted to evaluate the performance of SB. First, to test the ef cacy of the protocol, we tested SB on DETER [101] testbed. We describe this in detail in Appendix B.1. Next, to comprehensively evaluate the performance, we tested SB on a physical testbed (using a commercial SDN switch, HP3500YL). Our performance tests are broadly divided into two categories:

1. Controlled experiments, which were designed to test the robustness of the complete system including the client, proxy, and the SDN switch in a lab setup.
2. Internet experiments, which were designed to show that our system works well in realistic scenarios — i.e. downloading content from Internet websites that were otherwise blocked.

We begin by describing the setup used in the experiments, followed by the tests performed, and their respective results.

Experimental Topology:

SiegeBeaker was tested using the setup in Fig. 4.3, for controlled experiments (where all the entities were hosted inside our lab), and the setup in Fig. 4.7 for Internet experiments (when OD and CD were websites hosted over the Internet).

Figure 4.3: The topology used for evaluating SB; 1 hardware SDN switch (S1), 2 routers (R1,R2), an SDN controller and 6 host nodes. R1 blocks traffic to CD.

4.5.1 Controlled Experiments

Our first experiment involved the SB client downloading files of various sizes from the CD, and comparing the download times to that achieved when downloading over `wget`. The results of this test (ref. Fig. 4.4) depicts that SB performs considerably well in comparison to `wget`.

Figure 4.4: Comparison of SB and `wget` on a controlled setup, in terms of download time for file sizes up to 1GB.

Our next experiment was to evaluate the performance of SB's SP–client congestion and flow control (explained in Sec. 4.6) against native web (TCP) traffic. Both SB and `wget` clients simultaneously downloaded large files (varying between 100 MB and 1GB) from different CDs, over 1 Gbps shared network link (between R1 and S1, see Fig. 4.3). The throughput achieved over this shared link, for both SB and `wget` is shown in Fig. 4.5. As evident from the results, SB and the `wget` client achieve roughly the same throughput (sharing the link capacity almost uniformly). The result demonstrates that our system effectively emulates TCP's congestion and

ow control mechanisms, evenly sharing the link capacity with background TCP traf c.

Figure 4.5: TCP performance: SB and wget simultaneously downloading a le on a common link. They share the available bandwidth almost equally.

To determine the impact of cross-traf c on SB's performance (by increasing the load on SDN switch), we connected several hosts to the switch and exchanged increasing volume of traf c between them, via P2P connections. At the same time, we used a SB client to download a 1GB le from the CD. The results of this test are shown in Fig. 4.6. As evident, increasing the cross-traf c load had no impact on the client's throughput. This is because, once con gured, the switch merely forwards selective DR ows to the SP. This effectively isolates them from the rest of the traf c, avoiding contention.

Figure 4.6: SB client's performance with increasing load on the SDN switch.

Further, we also included provisions at the proxy to handle multiple clients. For a sample run of 9 clients simultaneously downloading a 100 MB le, we observed a nearly even distribution of bandwidthi.e. 11%. Similar results were observed when this experiment was repeated for larger les (such as 200 MB).

4.5.2 Internet Experiments

To further evaluate our prototype, we conducted experiments where the OD and CD were chosen to be websites hosted on the Internet. We tested our system's performance against two kinds of network workloads — regular web browsing and large file downloads respectively.

Figure 4.7: The topology used for evaluating SB over Internet with clients inside a university campus.

Assessing the performance for regular web browsing. This experiment involved recording the download times for the home pages of several blocked websites. To simulate web censorship, we configured our university firewall to filter access to Alexa top-500 sites, for a particular client under our control. The SDN switch and the SP were installed outside the firewall, and had unhindered access to the Internet.

In our test we spawned 500 concurrent SB client instances, using blocked Alexa top-500 sites as CDs, and several random unfiltered sites as ODs. In all cases, we were successful in downloading the home pages of the blocked sites, which varied in size from a few kB to 1.5 MB; we measured their respective download times. Next, we reconfigured the firewall to disable filtering, and again accessed the same web pages using Tor. As Fig. 4.8 shows, the average download time for SB (1.8 s) was very close to that for Tor (1.7 s).

Assessing the performance for bulk downloads. To test SB with bulk file downloads, we set up web servers on six geographically distributed machines, each serving files of various sizes. The firewall was set to block direct access to these machines from our clients, but we could download the files using SB, using a random unfiltered website as OD. Regardless of the background network conditions, the download times for SB do not significantly differ from our baseline,

Figure 4.8: Download times for Alexa top-500 websites — accessed by 500 parallel clients of SB and wget

Figure 4.9: Comparing time taken by SB and wget for downloading files, with OD and CD hosted on Internet.

download times using wget . Fig. 4.9 represents one such case, corresponding to a cloud server, comparing the download times w.r.t various file sizes (varying from 100 MB to 1 GB).

To gauge the performance of SB under heavy cross-traffic loads, we initially planned to divert the entire university's traffic through the switch. However, we chose against doing so, for two reasons. Firstly, it would force users uninvolved in the study to send their traffic through our switch, raising ethical concerns. Secondly, the cumulative volume of our university's traffic rarely exceeded 500 Mbps, and would not saturate the switch.

We therefore continued with our original setup, but connected 15 hosts to the switch so as to generate varying background loads. These hosts generated cross-traffic by establishing P2P connections with one another and exchanging large volumes of data (1.5 Gbps). At the same time, we downloaded large files (100 MB and 1 GB) from the cloud servers, using both SB and wget . As Fig. 4.10 shows, SB's download times are comparable to those of wget , regardless of variations in cross-traffic.

Figure 4.10: SB vs. wget when increasing the cross-traffic on SDN switch, with OD and CD hosted on Internet.

Similar tests were repeated by varying the number of parallel flows traversing the switch. The hosts, instead of exchanging P2P data, ran Apache Benchmark Version 2.4 and connected to a web server (also connected to the switch), spawning as many as 50k concurrent web connections[3]. Fig. 4.11 presents the outcomes of these tests. Increasing the number of flows, had no impact on the download times achieved using SB. This was comparable to what is achieved using wget. Thus, even in the presence of high cross-traffic, SB proves to be practical for downloading files of all sizes.

Additionally, SB has no measurable impact on non-DR traffic due to increasing DR traffic. Non-DR traffic is forwarded directly by the SDN switch without any intervention from the controller. Since SDN switch is a specialized hardware designed to transmit traffic at line rates, the non-DR traffic is thus not impacted.

Overall, the promising performance of SB can be attributed to our modular design (which makes use of hardware SDN switches), and avoids the unnecessary cryptographic inspection of non-DR flows.

Measuring the setup time. A SB client is required to covertly signal the controller to install a DR redirection rule, before it begins the actual download from the CD. The time elapsed between initial email and the installation of final DR redirection rule (i.e. between step 1 and step 5 of protocol) is called *setup time*.

[3] According to the NOC, this is much larger than the total number of flows at the university's edge router, that rarely exceeds 20k flows.

Figure 4.11: SB vs wget when increasing the number of flows on SDN switch with OD and CD hosted on Internet.

In 100 experimental trials, we observed an average setup time of 3-4 seconds. Similar to Sweet [94], we observed that most of this delay comes from email handling modules (selenium scripts composing mail, SMTP connection time at controller etc.) rather than network latency (which is of the order of milliseconds). [Note: The download time figures in previous subsections do not include setup time.]

Browsing experience: We used wget for benchmarking purposes. However, SB is also integrated in a browser. It can easily load static as well as dynamic websites without any significant performance degradation. Similar to Telex [77], we also used SB in our lab to access normal day to day websites for months without any problems.

4.5.3 Implementation Details

We now describe the implementation details of each individual component of SB. All in-lab machines, had Intel(R) Core(TM) i5-7400 CPUs @ 3.00GHz, provisioned with 8GB RAM, fitted with Gigabit Ethernet adapters.

Client: The client code is written in C (≈ 1100 LoC). It uses the OpenSSL library v1.1.0f for handling TLS connections. We use raw sockets API for crafting packets and libpcap for receiving and processing packets.

The email component is implemented for both SMTP and webmail using Python. The webmail version uses Selenium 3.14 [102] and SMTP version uses smtplib [103] in Python.

The client program listens on a local port on the client machine for new websites requests and serves the content as and when requested.

Secret Proxy: The code for SP is also written in C (1500 LoC). The client – SP TLS connections are handled using the OpenSSL library. We modify the header fields via raw sockets and use libpcap for packet processing.

SDN switch: Our implementation uses the HP 3500yl, a 24-port Gigabit SDN switch. This switch works in hybrid mode, i.e. both OpenFlow and non-OpenFlow environments simultaneously. It has three different types of immutable tables: Default (0), Hardware (100) and Software (200). Our redirection and inspection rules all make use of the Hardware table, as this allows for the best performance.

Controller: We designed and implemented our controller application using Ryu v4.15, which is written in Python (2068 LoC), and communicates with the switch using OpenFlow 1.3. Email reception and processing is implemented for both IMAP and webmail in Python. The webmail version uses Selenium and the IMAP version uses Python imaplib [104]. In our experiments we used IMAP for fetching Emails, as it was easier to integrate with the controller.

OD and CD: We use Linux nodes running Apache 2.4.18 configured as HTTPS server, to serve as the OD and CD.
The source code of SB can be found at at [105].

4.6 Discussion

In this section, we discuss the various design aspects along with the security analysis of the SB protocol.

4.6.1 System Design

1) Selection of timeouts: In our design, the controller needs to install redirection rules at steps 3 and 5. As the switch has limited memory, these rules need to be purged eventually.

The inspection rule (step 3 of protocol) match packets based on client–OD IP pair and has a hard timeout associated with it. We cannot use an idle-timeout for it because if the DR and non-DR clients are behind a common NAT device (i.e. share same source IP), then the aforementioned rule would also capture packets for the non-DR clients accessing the same OD. The rule may remain active even when there are no DR flows that match the rule. Enough of these (NAT based) inspection rules may unnecessarily remain in the switch exhausting its limited memory. Since, in all our experiments, we always observed that the inspection rule (in step 3) was installed in less than 3s, so we set the value of hard time out to be 4s.

For SP redirection rule (step 5 of protocol), controller installs the rule based on client IP, source port and OD IP. Here again, non-DR flows (sharing the same source IP, accessing the same OD as of DR client), shall not match the rule, because different flows have different ports associated with them. Thus, we use an idle timeout with this rule which expires once the flow is inactive for the said duration. Following [106], we set the value of idle timeout to 5s.

2) Proxy–Client traffic reliability, efficiency, and congestion control: The connection between SP and CD is a standard TCP connection, but the “connection” between client and SP merely appears to be so. Client–SP packets carry spoofed headers, so they can appear as client–OD traffic to the censor. However, they cannot rely on the kernel's TCP stack (due to spoofing). Hence, in order to ensure packets between client and SP are reliably and efficiently delivered (despite the absence of a true connection), we emulate TCP's congestion and flow control mechanisms at the application layer The SP–client traffic, not only bear sequence and acknowledgement numbers corresponding to the initial client–OD traffic, but also mimic TCP sending patterns (varying the sender window sizes, based on acknowledgements and timeouts, as in TCP).

3) Handling clients behind NAT: It must be noted that attempts by multiple clients behind a NAT to use the same OD for DR may lead to service unavailability. This can be explained with an example. Let us assume that for a client C_1 that is behind a NAT, an SP redirection rule (bearing NAT IP IP_{P1}, OD IP IP_{OD1}, and client's source port) is already installed. At the same time, another client C_2 attempts using the DR service with same OD (IP address: IP_{OD1}). This would lead to installation of a new inspection rule for the client C_2 for the same IP_{P1} and IP_{OD1} pair. Due to this new rule, the DR packets of C_1 would also be redirected to the controller, rather than

to the SP. Thus, disrupting the TCP DR session. As a preventive measure, we do not install any inspection rule for a client IP–OD IP pair, if there is already a SP redirection rule in place for the same. Thus, the clients behind a NAT device would require trying out different ODs to access DR service.

4) Location of Decoy Station (SP): Previous proposals (except Cirripede and Waterfall) require Decoy Stations to be located close to the Decoy Router. In our design, the Decoy Router is simply an SDN switch, which forwards packets based on flow rules; it can transport DR flows to the SP regardless of its network location. This helps defend against Forced Asymmetry attack discussed ahead in Subsec. 4.6.2.

4.6.2 Possible Attacks and Countermeasures

We now consider attack scenarios that may be launched by attackers within our threat model, and describe how SB addresses these threats.

1) Fingerprinting Attacks

a) TLS handshake fingerprinting: To prevent an attacker from characterizing the TLS handshake signature (e.g., the list of ciphers supported), we use the cipher suite of a popular browser Mozilla Firefox[4].

b) TCP/IP protocol fingerprinting: An adversary that records TCP/IP header values for traffic to ODs, may use it to distinguish the DR flows from non-DR flows. The DR flows may have different TTL values, window sizes etc. (due to packets originating from SP) compared to non-DR flows, making them easily distinguishable. In such a scenario the SP can generate identical TCP/IP characteristics as that of OD, making them relatively indistinguishable from the regular flows.

2) SB Confirmation Attacks: The aim of the adversary is to distinguish DR flows from the non-DR flows. To achieve this, an adversary may modify, drop or replay some packets of a TLS connection. However, tampering (or replaying) of regular TLS packets in a normal connection

[4]We can also use a recent library uTLS [107] to achieve the same.

results in the receiver replying with TLS error messages. Since SB hijacks an existing TLS connection, it does not natively respond with such error messages. The adversary might use this distinctive feature to con rm a DR connection. We now describe the kind of attacks such an adversary may launch, along with the preventive measures employed in SB.

a.) Record and Replay (without disrupting regular Internet clients): To con rm usage of DR, censor can record all traf c of a suspected DR user and can try to replay a part of (or complete) connection, at a later time. The rst interaction of DR client with SP occurs at step 7. All prior steps corresponding to the original client–OD TLS connection. Thus replaying packets of all the prior steps would not result in con rmation of DR. Hence, the censor may initiate a new DR connection and try to replay only the second TLS data packet, carrying the CD URL (step 7 of Fig. 4.1). Any response from SP, other than an TLS alert, would con rm the use of DR. However, our design is resilient to such an attack as the said packet contains a HMAC derived from the session's HMAC key. Hence, replaying it will lead to failure in HMAC veri cation. Additionally, none of the subsequent packets can be replayed because of the usage of AES in CBC mode.

b.) Modi cation and/or Replay (disrupting regular Internet clients): An adversary can attempt to con rm if an arbitrary TLS ow is intended for using DR service. Such an adversary can actively modify and/or drop packets of an ongoing ow. We believe that modifying the packets of such arbitrary TLS ow may prove costly to the adversary. It will disrupt regular Internet communication, eventually degrading the user experience. Such attackers are outside the scope of our threat model (ref. Subsec. 4.4.1). However, even if the adversary carries out such attacks, we enumerate the possible solutions to mitigate them in Appendix B.3.

3) Forced Asymmetry: This attack [87] assumes that traf c from anywhere in the Internet, destined to IP pre xes of censored regime can be (1) classi ed into traf c originated (or transited), from Decoy ASes (ASes which host DR infrastructure) and non-decoy ASes and (2) will enter the censorious regime at different locations (e.g., BGP routers) in the network using a load balancing technique called as hole punching [108]. E.g., if a client accesses OD without DR, its responses will enter at one particular network location. Whereas, if it pretends to use OD while using DR, traf c will enter at different location, con rming the use of DR to the censor.

In the existing DR systems (e.g., TapDance [81], Slitheen [83] etc.) the decoy router and the decoy proxy are placed together within a same decoy AS. The decoy router intercepts the DR requests of the client and forwards it to decoy proxy, which on behalf of client fetches content from the CD and sends it to the client (spoofing as OD). Thus, response packets will always come from decoy AS to the client. This makes the Forced Asymmetry attack successful against such systems.

However, due to SB's modularity, the DR infrastructure is not confined to an AS, because the SP (decoy proxy) can be hosted anywhere on the Internet, even in non-DR ASes. As a consequence, response packets originate from non-DR ASes rather than DR ASes. This would result in DR and non-DR traffic entering the censor's network through the same location, making the attack ineffective.

4) Routing Capable Attacks: A powerful adversary may disrupt any DR system (including SB), by preventing traffic originated from its networks to reach the decoy AS. Schuchard et al. [87], demonstrated few such attacks, which rely on the censor's ability to control the route taken by packets, either through a tainted path (which has DR between client and OD) or a clean path (devoid of DR):

a) Crazy Ivan attack: For ongoing client–OD connections (either through a tainted or a clean path), the censor can enforce changes in routing policies such that these connections (DR and non-DR) follow only the clean path. The non-DR flows would not be affected by this change whereas, DR flows would be disrupted.

b) Packet Injection: For a suspected ongoing connection, an adversary can send packets (spoofing as client) along a clean path in order to reveal the actual connection state between client and OD. E.g., an adversary could send a crafted TCP packet, with sequence number corresponding to the initial client–OD connection, to the OD. Responses like TCP ACK or Duplicate ACK would indicate absence of DR session, whereas a TCP RST response from the OD would reveal the presence of an ongoing DR connection.

c) Routing around decoys: An adversary might change its routing policies permanently, in order to avoid the DR ASes altogether (by always selecting a clean path), thereby denying DR

service.

All the aforementioned attacks rely on the assumption that clean paths can be easily obtained by the adversary. However, recent DR placement strategies (local [74], as well as global [80]) suggest that, obtaining clean paths would be difficult for the adversary, considering the massive topological changes (with the associated costs) required to obtain the same. On the contrary, Schuchard in [109] argued that, if the censor manages to obtain a small (yet significant) fraction of clean paths — rerouting its traffic via clean paths would incur heavy economic losses to decoy ASes (which are likely transit ASes that earn revenue by transiting the Internet traffic of their customer ASes). This might build pressure on decoy ASes to remove DRs by inflicting economic losses.

However, to successfully launch former [87] and latter RAD attacks [109], there are some practical challenges. Firstly, a censor needs to find all possible tainted paths from ISPs under its jurisdiction to different Internet destinations, by active probing [87] — a non-trivial exercise for the censor. Secondly, after obtaining tainted and clean paths, censor would have to introduce nationwide BGP policy changes, including changes to routing business relationships, and may lead to network downtimes etc.

5) Delayed SYN: The client is expected to initiate a TCP connection before the redirection rule expires on the SDN switch (step 3 of the protocol). If client's connection request is delayed and misses this window, intentionally by an adversary (or due to congestion), it would arrive at the OD (without being observed by the controller). The regular TCP and TLS handshakes would ensue. Packets sent by the client would reach OD, not the controller or SP. This does not pose a problem: after the handshakes (TCP and TLS), the subsequent GET request by the client fetches a regular response from the OD. The client immediately realizes that it missed the window, and initiates a new connection request. The censor sees nothing suspicious.

6) DoS Attacks: We discuss different ways in which an adversary may abuse the system to deny service to DR and/or non-DR clients. Like others, SB is also vulnerable to such attacks. We also discuss the strategies that are in place, or can be adopted to mitigate them.

a) Forced Decoy Routing: We consider the Byzantine attacker that tries to disrupt normal

routing for a non-DR user. To do so, the attacker simulates the client side of the SB protocol, while spoofing the source address of the victim (the non-DR user). Eventually, the controller installs redirection rules for the victim's IP address (assuming it to be a DR client). Thus when the victim tries to access OD, its traffic is redirected to the SP which fails to decrypt these packets, and drops them. Thus, by abusing the DR, non-DR traffic is prevented from reaching its intended destination.

However, in our design, the SP redirection rule matches packets based on source IP, destination IP and source port number. Thus, to successfully launch the attack, the adversary must anticipate the port number the victim will use while connecting to a particular OD; the probability of success of this anticipation is 0.00006 (assuming ephemeral port range = 16383).

b) Memory Exhaustion Attacks: As already mentioned in previous attacks (like Forced Decoy Routing), a powerful adversary can force installation of enough flow rules to exhaust the TCAM memory of SDN switch by pretending to be many legitimate DR clients. To minimize this threat, we incorporate timeouts for different flow rules (step 3 and 5 of protocol). Methods to minimize TCAM memory usage in SDN [110, 111, 112, 113, 114, 115] can be further used to mitigate this attack.

c) Spamming the Controller: A simple attack would be to flood the controller's email with spam. The adversary could send random emails, or more subtly, emails requesting DR service, from thousands of email addresses. This noise might prevent the controller from detecting legitimate requests. In such scenarios, we resort to issuing a unique email ID for the controller to individual (or groups of) clients. Thus, even if one email ID is spammed, clients can use other email IDs to access DR service. Additionally, similar to Mailet [95], we can also enforce usage limitations on clients or can use Captcha [116] and/or puzzles.

d) Fake Sessions Attack: As a more sophisticated attack, the adversary could send the controller "legitimate-looking" emails, which set up decoy routing sessions with random source IP addresses and ISNs. The controller, receiving an email, would install an inspection rule and begin to analyze the corresponding (irrelevant) flows. To minimize the impact of this attack, we include a hard timeout with the inspection rule. A fake email can make the controller inspect unwanted traffic

only for the timeout duration, i.e. 3s (ref. Subsec. 4.6.1).

e) Fake DR Request Registration: With an intent to deny some specific users of DR services, an adversary can send an email containing their source IP, some OD-IP and some random ISN. When the actual DR user sends an email, with the same source and OD IP, but a different ISN, the controller stores both the ISNs for matching. Hence, when the DR user sends a TCP SYN to the aforementioned OD IP, it's ISN would match the one indicated in the latter email, and the SP redirection rule would be installed.

An adversary could send multiple such emails with different ISNs. Now the controller would require maintaining a set of ISNs for every client and OD IP pair. Searching through such sets to match ISNs of the incoming SYN packets could marginally increase the overhead at the controller. However, DR services would not be hindered.

Moreover, the presence of the hard timeout of the inspection rule, would force the adversary to keep sending such emails regularly. This could be very expensive for the adversary if the attack has to be carried out for all possible clients. Thus, SB is more resilient to such attacks, unlike other registration based systems (such as Cirrepede and Waterfall of Liberty) where the adversary just needs to register once for a client.

4.6.3 SDNs and SiegeBreaker

It can be argued that SDNs are primarily deployed in small networks (like data centers), and are not suitable for use on the Internet. However, SDNs have been shown to scale to an entire AS [117, 118, 119], and practical ISP-scale SDN has been demonstrated by Rexford et al [120]. Moreover, there is a recent research that suggests how ISPs can (and should) migrate to SDNs [121]. Further, SB can work even for a non-SDN AS. Friendly non-SDN ISPs could position open flow switches such that they intercept traffic to critical network entities (e.g. routers which transport a large fraction of traffic [80]), without replacing the existing infrastructure

Scalability of Controller: Would the SB controller become a bottleneck at ISP scale? Recent work [122, 123, 124, 125, 126] suggests that SDN controllers can scale to very large use cases –

for instance, Cuttlefish [124] proposes a hierarchy of local and global controllers, which can offload tasks to each other, providing higher control plane throughput and better scalability. We suggest that, if needed, such approaches can be used when SB is deployed at scale.

4.7 Conclusion

We present SiegeBreaker, a practical and efficient Decoy Routing system. Using an SDN architecture, SiegeBeaker divides the tasks of Decoy Routing among three loosely-coupled modules. The SDN controller focuses on detecting packets of interest, using an efficient privacy-preservingsignaling scheme, by reconfiguring SDN switches on-the-fly. The switch then redirects all packets of the Decoy Routing flow to a secret proxy server. Finally, the proxy server communicates with the covert website on behalf of the client, and transmits back the responses reliably and efficiently.

In extensive tests involving commercial SDN switches, our prototype shows promising performance — nearly equal to that of direct TCP connections. Additionally, SiegeBreaker's flows uniformly share the available link bandwidth with other non-DR connections. Along with privacy preserving signaling, such performance results show promise for future implementation and deployment by SDN-based networks.

It must be noted that DR systems provide great blocking resistance. However, such systems pose deployment challenges as they require collaboration from ISPs to deploy DR infrastructure. Thus, DR provides a promising but futuristic solution for circumventing censors. However, in the meantime we require systems that can be immediately used without posing deployment challenges. Thus, in the next chapter we aim to build such a system by utilizing instant messaging applications to transfer restricted content.

Chapter 5

Camou er

5.1 Introduction

In the last decade the Internet has become an integral part in various aspects of our lives, ranging from education, healthcare to social interactions, technological advancements in different spheres of sciencœtc. Free ow of ideas and unrestricted access to information has become a necessity not only for personal development but also for the advancement of the society. However, repressive regimes continuously attempt to surveil and censor this ow of information by restricting the content, users can share and access.

In opposition, free speech activists and researchers developed systems [31, 127, 128] which aim at providing unhindered access to information for clients in repressive regimes. This strife of ideologies between the censors and the free speech advocates has led to the evolution and development of effective censorship as well as anti-censorship technologies [128, 127]. This has led to an arms race between the censors and free speech activists. As censors advance their craft, researchers try to stay a step ahead by developing hard to block anti-censorship systems [130].

To that end, recent anti-censorship systems [9, 131, 93] are designed on a fundamental principlei.e., to incur collateral damage to the censor, if it attempts to disrupt the circumvention schemeThis makes it dif cult for the adversary to completely block these systems. Approaches

like decoy routing [9], domain fronting [131] etc. are examples of such systems. Restricting Decoy Routing requires the censor to update nation-wide routing policies, and in the process sustain heavy collateral damages (e.g., increased performance overheads [132]). Similarly, Domain Fronting requires the adversary to block cloud services (such as Google App engine), which might also be hosting essential services for oblivious users. Other systems such as Conjure [133] and MassBrowser [134] require the censor to block some IPs or IP pre xes, thereby also blocking other innocuous services running behind those. Lastly, tunnelling systems such as SWEET [94], Covercast [96], Freewave [135] etc., transport censored traf c via services and protocols essential for smooth functioning of businesses, and thus a censor's economy. These systems exploit Email [94], VoIP [135], video streams [96, 136] etc., as covert channels to transport content. Since these systems use the underlying protocol as-is, it becomes hard for an adversary to distinguish circumvented traf c from the underlying protocol's messages. Hence, a determined adversary may attempt to disrupt the use of the underlying protocol itself (e.g., emails in case of SWEET). Although, blocking such channels may incur collateral damage to the censor.

Figure 5.1: Camou er basic architecture

However, despite these systems providing ef cient blocking resistance from the adversary, they exhibit other challenges which hinders them from being widely used. These challenges include deployment limitations, high cost of operation and low performance for the users. For instance, Decoy Routing and Conjure requires collaboration from ISPs to install and maintain additional network hardware for them to function. Similarly, Domain Fronting requires hosting a proxy on a fronting service (such as Google App Engine, Amazon Cloudfront etc.) which incur high periodic subscription [131] costs. Moreover, DeltaShaper [136] (a tunneling based system) provides 2.56 Kbps throughput, which is insuf cient for providing web browsing.

Thus, we propose and build a new tunnelling based system, Camou er, that aims at over-

coming the shortcomings of the existing systems while maintaining similar blocking resistance. Camou er utilizes Instant Messaging (IM) platform as a medium to tunnel the censored traf c. IM channel seems to be better suited to act as a tunnelling medium for developing such a system in comparison to other existing counterparts. This is because it has some salient features, that in general, anti-censorship schemes strive to achieve. They are:

1. Minimized latency: IM platforms aim at minimizing the latency (<1s) [137] when user exchange messages with each other (ref. Sec. 5.4). This provides good QoS, unlike other non-realtime channels such as emails.

2. Adequate throughput: IM platforms have suf cient data transport capacity (in the form of attachments), in comparison to channels such as VoIP (which encode data at low bit rate). Thus IM proves to be more suitable for regular web browsing.

3. Reliability: IM is also a reliable channel in comparison to others such as VoIP and video. While IM ensures reliable delivery of messages, real-time VoIP and video generally do not incorporate this feature, as the information lost in the latter channels becomes irrelevant and is thus not recovered.

4. Blocking Resistance:Similar to existing systems, restricting the underlying IM applications may incur collateral damage to the adversary as IM apps are an important part of personal as well as professional spaces [138, 139, 140, 141, 142][1]. At present, businesses utilize IMs as a medium to advertise, communicate and expand, e.g., several airline reservation and movie ticketing services utilize IMs to directly send e-tickets to customers. Further, IM based collaboration platforms such as Slack, Flock, etc. are now widely used as an alternate to email for professional communication [143, 144, 145].

5. Deployment Ease:IM applications are ubiquitously used by netizens. As a consequence, to use Camou er, a user merely needs to install programs at the client and server ends. Apart from that, there are no additional requirements as assumed by previous circumvention proposals (e.g.,collaborating with ISPs [9, 80, 85]).

[1]There were 2.5 billion active IM users till January 2019 and this number is expected to easily cross 3 billion till 2020 [142].

Thus, we developed Camo er with IM channels as its underlying tunneling protocol. The basic architecture of Camou er (depicted in Fig. 5.1) requires the censored user to have an IM ID on any popular IM platform. The user using the Camou er client, tunnels requests to a Camou er server hosted in a free country, which acts as a proxy and serves the censored content to the client. All the censored content is encrypted and exchanged via the underlying IM platform that the user has access to. This would give the pretense to the censor that regular IM clients are communicating, providing unobservability to Camou er.

Camou er has been implemented and tested to work on several popular IM applications including Whatsapp, Signal, Telegram, Slack and Skype, and can be extended to others as well. Camou er clients take an average and median time of 4.1s and 3.6s respectively, to access webpages of Alexa top-1000 sites.

To summarize, following are the major contributions of this work:

- The design of a new anti-censorship system Camou er, which utilizes IM apps as covert media to tunnel censored content.
  - Usage of IM covert channel serves multiple advantages in comparison to existing tunneling channels by ensuring: (i) low latency (ii) reliability (iii) similar blocking resistance and (iv) high data transport capacity.
- A prototype implementation with a detailed performance evaluation of Camou er on ve popular apps including, Signal, Telegram, Slack, Whatsapp, and Skype, depicting the feasibility of our design.
- A detailed security analysis, depicting Camou er's robustness against variety of attacks including traf c analysis.

5.2 Related Work

To promote free speech and unhindered information exchange over the Internet, researchers and free speech activists have proposed several anti-censorship solution [148, 147, 148, 74, 128, 127,

9, 131, 149] over the years. Given the plethora of such systems and proposals available, we try to categorize them and enlist their advantages and disadvantages.

1. Proxy based systems:These systems include proxies, VPNs, Tor [71] [etc. These involve clients relaying their traf c via intermediate proxies, that connect to the censored sites on behalf of the clients. Such systems are easily deployable and thus readily available to end users. However, they can be easily blocked by the adversaries, as generally such systems publicly advertise the proxies' IP address. This, makes it trivial for the adversary to block them as soon as discovered.

2. Decoy routing systems:Decoy routing is a promising approach, that requires client in the censoring regime to connect to un ltered websites. These requests contain covert information, that allows special intermediate routers (decoy routers) en-route to intercept them, and decipher the true censored destination that the client wishes to access. These decoy routers then proxy the requests and responses between the clients and censored sites, while keeping up with the pretense to the adversary that the client is connected to the un ltered website. Examples of such systems include, Telex [150], Cirripede [78], Slitheen [151], Tapdance [149], Waterfall of Liberty [82],SiegeBreaker [152] etc.

   To censor decoy routers, the adversary may require undertaking daunting measures such as changing entire nation's routing policy [87] in order to bypass such systems. Such routing changes are prohibitively expensive to achieve in practice [79, 80]. Thus it becomes very dif cult for the adversaries to block them. However, such systems require collaboration from the ISPs in order to function and thus pose a hurdle in deployment [80].

3. Mimicry based systems:These systems attempt to disguise and transfer censored content as regular applications' protocol messages. E.g., SkypeMorph [74], helps access censored websites by mimicking Skype's communication protocol. Others, like CensorSpoofer [158], obfuscate requests to camou age censored sites as VoIP messages transported over SIP. However, such systems are relatively easier for the adversary to block as it is very dif cult to mimic all the features of the underlying protocol [75]. Moreover, their performance depends on the traf c rate of the cover protocols. VoIP, used commonly, has very low transmission

rates (e.g.,5 - 40 Kbps in Skype), thus leading to low QoS for web browsing.

4. Tunneling based systems:Tunneling based systems rely on encapsulating covert traf c in standard application protocol messages (e.g., email, VoIP calls, streaming video etc. Examples include SWEET [94], Covertcast [96], Delta Shaper [136], Freewave [135], Cloud-Transport [154] etc. These systems are an improvement over mimicry based systems as they do not mimic protocol messages, rather they directly use them as the covert channels. This ensures that the features of the underlying protocol (e.g., packet size) remain unaltered, while the censored content is encapsulated inside the payload. In turn, this makes it dif cult for the censor to disambiguate them from regular (underlying) protocol messages. Thus such systems provide more ef cient blocking resistance against adversaries, as the latter may require blocking the complete underlying applications (such as email, cloud services etc.) which may result in massive collateral damages. However, such systems provide poor QoS for web downloads (e.g.,due to low offered bandwidth in these channels) and are limited in their deployment.

5. Miscellaneous systems:There exist other anti-censorship systems, which do not fall in any of the above categories, like domain fronting [131], CacheBrowser [93], MassBrowser [134] etc. Domain fronting makes use of different popular cloud services such as Google app engine , to access censored content. The request to the proxy server (hidden behind a cloud server), is concealed in a HTTPS request, which is destined to the domain name of an innocuous front-end of the cloud server. This front-end decrypts the HTTPS request and forwards it to the proxy server. To block such services the adversary may require blocking the entire fronting service (e.g., Google app engine), thereby blocking other third party applications that use the platform. However, leveraging fronting services is not cost effective for deployment [131].

   Similarly, CacheBrowser works by utilizing content distribution networks (CDNs) frontends, located outside the censor's boundary. It requires the blocked content to be hosted on the CDN network, and thus also has an associated cost. Moreover, websites not hosted on these CDNs cannot be accessed.

   A recent system MassBrowser by Nasr et al., leverages some of the existing techniques (such

as Domain fronting and CDN browsing) to build client-to-client proxy system. It is similar to proxy based systems (described before), with a fundamental difference that these proxies are not hosted on a public IP but are rather behind public NATs. Thus, blocking the MassBrowser is not as easy as blocking the IP addresses as it would lead to blocking other clients behind those NATed IPs.

Lastly, similar to Decoy Routing a recent approach by Frolov et al. (Conjure [133]) tries to host proxies on an ISP's unused IP addresse space. However, Conjure also requires collaboration from ISPs to install network taps that would allow them to inspect the traffic transiting the ISP.

Camoufler is an example of tunneling based systems. It aims to address the shortcomings of tunneling systems such as low QoS, low-bandwidth channel etc. Additionally, similar to other systems, we attempt to achieve effective blocking resistance.

5.3 Camoufler Architecture

This section describes the threat model we consider for Camoufler, and present its design (ref. Fig. 5.2) as well as a step-by-step walkthrough of the protocol.

5.3.1 Threat Model

We assume the adversary to be a nation state, which can employ any existing censorship techniques e.g., IP/DNS filtering, URL and keyword filtering [129, 155, 156] etc. However, we assume that the adversary is not willing to be disconnected from important Internet services (or from the complete Internet for that matter) and inflict loss to itself, for achieving extensive censorship. Specifically, censor would allow at least some IM based channel(s) to function within its jurisdiction. However, it may attempt to selectively block certain IM platforms. This is because IM platforms have proliferated significantly and are widely used in personal as well as professional space [138, 139, 141]. Hence, it would be difficult for the censor to completely block them without sustaining significant collateral damage.

Figure 5.2: Camou er Detailed Architecture.

Additionally, we also assume that the IM channels are encrypted, either end-to-end or end-to-middle, as is already the case with most of the existing popular IM applications [157]. Thus the adversary can monitor or actively analyze the encrypted traf c. He may also attempt to drop, modify or replay packets for deliberately inducing perturbations (for some suspicious IM connections) to detect the usage of Camou er. However, he would refrain from launching these active attacks at a large scale so as to not disrupt the communication of regular IM clients.

5.3.2 System Design

We now describe the protocol design of our proposed system. Our system tunnels the web requests of IM users (i.e., Camou er clients) residing in censored regime to an IM peer in a free country (acting as Camou er server), using IM applications. The users of Camou er could themselves rent out VPS servers in such countries and run the Camou er server, or may rely on trusted peers (friends in free countries) for running the server on residential or educational hosts.

The Camou er server proxies the received web requests to the censored destination. This enables the clients (in a censored regime) to access blocked content. General working of our end-to-end system can be understood by referring to Fig. 5.2. We begin by describing the individual

components of Camou er and their functioning, followed by a step by step walkthrough of Camou ers' operation.

Camou er client consists of the following components:

1. Local Proxy (LP): It is a standard HTTP proxy that acts as an interface between the user's browser and the client engine (described ahead). More speci cally, it accepts content requests from the browser and passes them to the client engine and vice versa.
2. Client Engine (CE): Client engine acts as an interface between the LP and the underlying IM application. It receives a web request from the LP, processes it (encryption, compression), and then forwards the processed web request to the underlying IM application.

   IM application oblivious to the aforementioned process sends the received content to the other IM peer (Camou er server) as standard IM packets. Similarly, CE receives content from the IM application, decompresses and decrypts it before sending it to the LP.

The IM server consists of the following components:

1. Server Engine (SE): The server engine works similar to CE. When SE receives a web request, it forwards it to the CP. Later, when SE receives the web content from the CP, it compresses and encrypts the web content and sends it back to the Camou er client, using the underlying IM application.
2. Content Proxy (CP): The content proxy retrieves the requested censored content from the blocked website and sends it back to the SE.

Furthermore, similar to other existing systems, we assume that the user has access to the Camou er software. We now describe a complete walkthrough of our system. The steps involved in accessing Camou er are as follows:

1. The user con gures his browser of choice (Firefox, Chrome, Opera, etc.) to forward all the requests of the browser to the LP. Once con gured, the client can use its browser to access the blocked content freely.

2. Next, the user inputs the URL of a censored website in the browser. This is forwarded to the LP, which forwards this connection request to the CE and waits for the retrieved content.

3. On receiving the content request (from LP), CE encrypts it (using a derived shared key as described ahead in §5.3.2.1), and uses the underlying IM application to transport this request to the Camou er server (SE).

4. SE, on receiving content request, decrypts it and then forwards the request to the CP. The CP, retrieves the censored content and sends it back to the SE. The SE then nally, encrypts and compresses this content and tunnels it back to Camou er client using the IM channel.

5. On receiving the website content from Camou er server, the CE decompresses, decrypts and forwards it to the LP, which in turn provides it to the browser for appropriate rendering.

5.3.2.1 On encrypting exchanged content

Most of the IM applications support end-to-end encryption and thus ideally it is not required to additionally encrypt the messages exchanged via them (ref. §2.3). However, there might be scenarios where E2E applications are not allowed within the censor's jurisdiction (only E2M apps are allowed), or the client wants to be extra cautious by additionally encrypting the exchanged content. In such scenarios, the client derives a shared key with the server.

For this, the client program needs the RSA public key (key_{pub}) of the server along with the public key of its DH exponent (g^y). Since the Camou er server is managed by the user himself (or by his trusted peer), these keys are assumed to be with the client program.

Similarly, the client would generate its DH private key x and eventually derive the shared key g^{xy}. When the client utility is run by the user, as a background process, it informs the server that it intends to encrypt its communication. Then the utility sends the public key of its DH exponent g^x encrypted with the RSA public key of the server (key_{pub}).

The server extracts g^x by decrypting it using its RSA private key (key_{priv}) and derives the shared key g^{xy} using its private part of DH exponent y. Once the key is derived, both the parties derive a hash of g^{xy} (using SHA-256), and use these resulting hashed bits to encrypt the

subsequent messages exchanged between them using AES-256. The corresponding DH private keys x and y along with the derived key g^{xy} are then deleted to ensure perfect forward secrecy.

5.4 Evaluation

We now evaluate the performance of Camou er using our prototype implementation on several IM appsviz., Signal, Skype, Slack, Telegram and Whatsapp. It must be noted that Camou er has a general architecture which can be implemented on any IM app.

Time to access Alexa top-1k websites. In the rst experiment, we assessed the time, Camou er takes to access different websites. Thus, we downloaded the default webpages of the Alexa top 1000 websites and recorded the time it took Camou er to complete this operation. The Camou er client and server were geographically apart by a distance of 8300 miles. The Camou er client was running on a machine hosted in our university, whereas the server was running on a cloud hosting service. We performed this experiment for all the ve IM apps. The results (in the form of a box plot for each IM app), are represented in Fig. 5.3. As evident from

Figure 5.3: Download time of Alexa top 1000 websites using different IM apps and its comparison with direct downloads.

the gure, we were able to access most of the websites in a few seconds that is comparable to direct download timeE.g., for Signal, we recorded a median download time of 3.6s with the average being 4.1s. Similarly, using Telegram we were able to access these websites with an average time of 2.7s and median time of 2.3s. Time taken by Whatsapp was higher (average

of 7.6s) compared to other IM apps, as it was automated using Selenium web automation framework. The details can be found in Appendix. C.4. However, it must be noted that the performance obtained by Camoufler using Whatsapp was still better than most of the existing systems such as Covercast [86], Deltashaper [136], etc. which incur an overhead of more than 10s for similar operations.

Time to First byte: This experiment was conducted to test the responsiveness efficiency of Camoufler server. We record the time it took for the first byte of the content to reach the Camoufler client (from the Camoufler server), after it sent the initial request.

For this experiment we accessed the Alexa popular 10 websites (100 times each) and plot the CDF of the results obtained for the Telegram app in Fig. C.1. As evident from the graph, most of the websites (in over 90% of trials) were able to receive their first byte in less than 2s with half of the websites retrieving it under 1s. We obtained similar results for other apps. The details of the individual apps can be referred in Appendix. C.1.

Figure 5.4: CDF of Time To First Byte (TTFB) for 10 popular Alexa websites (each downloaded 100 times).

Transmission time of messages from Camoufler client to server: Next, we evaluated the time spent for sending request from Camoufler client to Camoufler server via IM platforms (analogous to one way delay) using the aforementioned setup. This provides a measure of the end-to-end latency incurred due to the underlying IM platforms in transferring messages. Thus, in this experiment, we sent the same web requests via different IM platforms (100 times each), and recorded the time taken in receiving them at the other end. As evident from Fig. 5.5, in majority of the cases, the apps take less than a second to transfer the content in one direction. The inherent

low overhead of the IM platforms is very helpful in providing good QoS for the clients, and is re ected in our results.

Figure 5.5: CDF of time taken by a message to travel from Camou er Client to Camou er server for different IM apps.

Location diversity: In this set of experiments we varied the location of both the Camou er client and the server to analyze if there was any signi cant change in performance. We varied the client across six locations, and the server across three (each in America, Europe and Asia). Alexa top-1k websites were downloaded for each of the 18 (6x3) client-server pairs. The result for the server in Asia (Singapore) is depicted in Fig. 5.6. It is evident from the box plot that there was not much variation when the client location was varied. The trend remained similar for other server locations as well (ref. Appendix. C.2).

Figure 5.6: Box plot depicting download time (from a server in Singapore) of top Alexa-1000 websites for varying client location.

Bulk downloads: Camou er also supports bulk content transfer by transmitting the large

les as compressed attachments. The Camou er server rst downloads the requested content, compresses it and then transfers it back to the client. In our tests, Camou er client successfully downloaded les of various sizes (10, 20,...,100 MB) that were hosted on cloud servers. A sample of our results depicting download time variation is summarized in Tab. 5.1. Using Telegram as

Downloaded Using	Time (in s)						
	10 MB	20 MB	30 MB	40 MB	50 MB	75 MB	100 MB
Direct (Wget)	7.9	15	23	29	35	51	68
Camou er	13.6	23.5	34.6	45.3	52.1	77.2	93.3

Table 5.1: Large File Downloads: Comparison of download times of Camou er and Wget.

our underlying IM channel, we downloaded each le ve times. Across different measurements we observed that download time with Camou er is higher when compared to download time with wget . This is because, Camou er server rst downloads the complete le (at its end) and then sends it to the Camou er client. However, large le downloads are generally delay tolerant and thus we believe this additional delay could be acceptable by the Camou er clients. Overall, we observed a similar performance with other IM channels as well.

5.4.1 Implementation Details

We now describe the details of our proof of concept implementation of Camou er. As an example we describe the implementation details on the Signal messenger platform as it is very popular among security and privacy practitioners. However, we similarly implemented Camou er on other platforms as well, including Whatsapp, Telegram, Skype and Slack. The details of their implementation can be referred in Appendix C.4.

Camou er Client: The client implementation was performed on a Linux host running Ubuntu 18.04 LTS, consisting of a 4 GHz processor, and provisioned with 8GB of RAM. The client's browser was con gured to forward its requests to local port 8000 where LP listens for Camou er requests. LP is written in python and uses the socket API [158] to manage connections to and from the browser.

CE comprises of scripts written using python and shell-scripting language. It interacts with the underlying Signal messenger to send and receive messages, using the signal-cli [159]

interface. signal-cli helps automate exchanging messages (using CLI commands), over the signal messenger. Thus, CE uses it to craft and send blocked websites request to the SE, using the send -m command of signal-cli . Secondly, CE use the daemon mode and the dbus feature, shipped with the signal-cli interface to listen for incoming messages from SE. The dbus feature allows applications to create a listener which can easily receive and process different events that are generated when the signal app receives messages. Pydbus [160] was used to interact with the dbus interface. On receiving the response, CE decompresses and decrypts it using the gzip [161] and Crypto [162] libraries in python, before forwarding it to the LP. The compression and decompression process are lossless.

Camou er server: Camou er server was also implemented on a Linux machine running Ubuntu 18.04 LTS OS with 4 GHz processor and 8 GB RAM. The source code was also written in python and shell-scripting language, and similar to CE it utilizes pydbus [160], gzip , and Crypto library for performing various tasks.

SE also utilizes the signal-cli interface and accesses the dbus feature for processing incoming messages. This processing enables SE to extract the censored website. The CP connects to the blocked website using python socktes API to retrieve responses. The SE compresses and encrypts the responses using gzip and Crypto library before sending it back to CE (using the signal-cli interface).

5.5 Security Analysis

We now describe various attacks the censor might attempt to block access to Camou er.

5.5.1 Traf c Analysis

Censors may attempt to analyze Camou er client's traf c in order to identify distinguishing features and block them from using Camou er. As already described in our threat model, the censor can inspect the encrypted Camou er client content within its network boundaries.

It must be noted that the functionality of all IM apps (with respect to their traf c characteristics) are very similar [157]. Thus the proposed attacks (and defenses) discussed subsequently are applicable across all IM platforms. We now enlist the possible attacks.

On Traf c patterns: Camou er involves downloading and accessing blocked content by exchanging IM messages. As analyzed previously (by authors of SWEET [94] and Mailet [95] etc), an adversary could attempt to distinguish regular IM traf c from IM ows that transport Camou er traf c. Such attacks work on the premise that the behaviour of the tunneling/encapsulating protocol could be different when used with and without anti-censorship schemes. For instance, if there are differences in the packet exchange rates between a regular IM client and a Camou er client, then it could be used to disambiguate the two. On one hand, it is already known that, IM clients (other than chatting) exchange signi cant amount of multimedia content [163]. In a recent study [164], authors reported that above 50% of the messages exchanged over IM applications constituted multimedia content. However, on the other hand, Camou er clients would mostly fetch blocked content (e.g.,websites). Thus a determined adversary may attempt to differentiate web content downloaded (using Camou er) from multimedia content downloaded (using standard IM apps). But, since the underlying traf c (both of regular IM and Camou er) is encrypted, it is plausible that adversary may opt of for traf c analysis based on differences in packet exchange rate and packet sizes [94].

Thus to observe such differences, we performed tests involving regular IM clients accessing multimedia content and Camou er clients accessing websites. In our experimental setup, we used one machine (located in our lab) for running both regular IM and Camou er clients. This machine communicated with another one (located in a different country) that ran an IM client (the peer) and the Camou er server. We began by measuring the packet exchanged rate when (1) multimedia content was downloaded by a regular IM client and (2) web content was downloaded by using Camou er. Fig. 5.7 depicts the scenario when we downloaded a PDF le, a GIF animation, an image and a video clip using regular IM app. Further, we downloaded the webpages from cnn.com , youtube.com andgithub.com , and a DOC le using Camou er. It is evident from the gure that, there is a sudden rise (spike) in packet exchange rate when multimedia content is shared between IM clients. This is because, multimedia attachments (e.g., large video)

involves a lot of content being transferred. Thus, underlying IM applications send data (packets) at a faster rate resulting in a spike in packet exchange rate.

Figure 5.7: Packets exchange rate of regular IM client accessing multimedia content (images, GIF animation, video etc.) vs a Camou er client accessing websites (, github etc.) and a doc le. As evident, traf c characteristics of both regular IM and Camou er are very similar.

It must be noted that, depending upon the size of the multimedia object being shared over the IM channel, one could expect spikes in packet exchange rate. For a large object (e.g., a video of size 1 MB) the spike would be very high when compared to a smaller object (an image of size 100 KB). This trend hold good for Camou er traf c as well, as it uses the same IM app for transferring the content. For instance, when we downloaded a 1.5 MB video using regular IM client, we observed packet exchange rate peaked at 800 packets per second (ref. Fig. 5.7). Further, on downloading a 1.3 MB document using Camou er, we observed a similar packet exchange rate (i.e., more than 700 packets per second. This trend holds good for smaller size les, websites and multimedia objects as well. Our experiment thus indicates that it is hard to differentiate Camou er traf c from the regular IM traf c.

Further, we also plotted the packet size distribution for the above set of experiments for multimedia content (exchanged by regular IM clients) and websites (accessed by Camou er clients). It is evident from the histogram presented in Fig. 5.8, that the maximum number of packets are clustered in two bins, the rst bin with less than 100 byte packets, and the second bin with more than 1200 byte packets. The former corresponds to mostly the acknowledgement packets generated from the regular IM (or Camou er) client and the latter corresponds to the data packets sent by the other IM client (or the Camou er server). Overall, it is evident that

depending on the size of the content, the number of data packets vary; large content download would result in large number of larger sized data packets (over 1200 bytes) while smaller content download would result in fewer bigger size data packets. For instance, the video downloaded by the regular IM client resulted in about 1200 data packets (over 1200 bytes each). In contrast, the image download resulted in only about 100 data packets of comparable sizes. The trend was very similar for Camou er client as well, a document download resulted in 1000 data packets, whereas accessing website such as cnn.com resulted in only around 150 data packets.

This is further highlighted in Fig. 5.9. Large content download would result in large number of bigger size packets (over 1200 bytes) irrespective of the type of client (regular IM/Camou er) and the content type (video/document etc.). E.g., the box plot for a video download (by a regular IM client) and a document download (by a Camou er client) are very similar — packet size distribution mostly consisting of more than 1200 bytes packets. Similarly, the box plot for smaller objects (like image/GIF) download by a regular IM client looks very similar to the website downloaded using Camou er. However, the box plots in this case have more spread; packet size distribution consist of a fewer bigger size data packets[2].

Additionally, we performed multiple such experiments to further strengthen our claims. We downloaded different multimedia objects using regular IM clients as well as web content using Camou er, and measured the packet exchange rates and packet sizes. Across all our tests, the observations remain consistent, i.e., the traf c characteristics of a web download using Camou er is akin to multimedia download using regular IM apps (ref. App. C.5 for details).

However, it could still be argued that our observations would hold good only if regular IM clients often download multimedia content. Otherwise, the spikes in packet exchange rate, or a packet size distribution consisting of large number of bigger size packets, could be a uniquely identi able characteristics of Camou er. To that end, recent studies [103, 104] suggest that IM clients very often exchange multimedia content. For instance, in [104] authors report that more than 50% of the messages exchanged over IM platforms constitute multimedia objects. These would also result in spikes in packet exchange rate and also the transmission of bigger sized packets, by

[2] Additionally, if the average size of multimedia objects being exchanged over IM platforms differ from the average webpage sizes, the censor may attempt to distinguish them. However, in such a scenario, Camou er can easily add cover traf c (by padding extra bytes) to even out the differences.

regular IM clients. Hence, overall we believe that any attempts by a wire-snif ng adversary to distinguish Camou er clients from regular IM clients could lead to high false positives.

Figure 5.8: Histogram of packet sizes when regular IM client accessing multimedia content and when they use Camou er to access websites.

Figure 5.9: Box plot of packet sizes observed when regular IM client accessing multimedia content and when they use Camou er to access websites.

Inducing traf c perturbations: Next, an adversary may attempt to identify Camou er ows by actively dropping, delaying, modifying packets in some connections so as to see if Camou er and regular IM clients behaved differently to compensate for such perturbations. However, it must be noted that, Camou er is not mimicry based, rather a tunnelling system. It does not "pretend" to use the IM apps, rather it uses the underlying IM channel without modifying its default behavior. Therefore, such analysis would not provide any observable changes in the packet level features of the IM apps being used by Camou er. The IM channel would continue to respond in exactly the same way to perturbations such as drops and modi cations, regardless of whether they are used with the Camou er or not.

5.5.2 Other Attacks

Collusion attacks: The adversary can attempt to coerce, and thus collude, with the IM service provider which would enable it to access much more information than it could normally obtain by analyzing merely the encrypted traf c. This information could further be used for identifying and blocking Camou er clients. Thus, there could be two possibilities:

1. Collusion with an end-to-end encrypted IM provider: In this scenario, the IM provider, and thus the adversary, would not be able to inspect the content of IM messages as they are end-to-end encrypted (ref. §2.3). The lack of plain-text messages would hinder the adversary from obtaining any identifying information, such as the kind of content being transported.

   Further, the censor could observe the metadata of messages from the IM provider, e.g., their IM IDs. However, since the Camou er server's ID is not publicly known (ref. §5.3.2), the censor would not be able to differentiate it from regular IM IDs.

2. Collusion with end-to-middle IM providers: In end-to-middle IM applications, an encrypted channel is established between the client and the IM providers. Thus, if the adversary colludes with the IM provider, it would be able to inspect IM clients' content in plain text. However, Camou er client derives a shared secret with the Camou er server (using the scheme described in Subsec. 5.3.2.1), and uses that to encrypt the messages. This would not allow the censor or the IM provider to inspect the plain-text traf c and thus they could not attempt to identify clients by ltering requests seeking censored URLs.

   However, in extreme cases the adversary could attempt to identify and drop all encrypted IM messages, to disrupt Camou er. In such a scenario, we could use stenographic techniques [153] to hide our content from the adversary in plain sight, as also assumed in other anti-censorship systems [94, 96]. This may reduce the overall QoS and thus could be seen as a trade-off between unobservability and QoS.

Identifying Camou er servers: An adversary may attempt to identify Camou er servers' IM IDs, after which he/she may attempt to censor it. If the adversary owns the IM platform it could simply lter the IDs by itself, otherwise it may coerce the IM provider to block the said

IDs. However, as already mentioned, Camou er servers' IM IDs are not publicly known — either a Camou er client would run its own Camou er server in some hosting service or would request someone trusted to run the Camou er server utility. Thus, it is extremely hard for an adversary to determine the Camou er servers' IM IDs. Additionally, a determined adversary may further attempt to actively probe different IM IDs (on all IM platforms) by pretending to be a Camou er client. Responses to such probes could lead to the detection of the Camou er server. As a mitigation, the Camou er server responds only to the trusted IM IDs. Thus, active reconnaissance by censors would be futile.

Long term user pro ling: An adversary could attempt to longitudinally pro le individual IM clients. Any deviations from the pro led behavior (such as sending messages at odd times of the day etc.) may evoke suspicion of use of a tunneling based system (including Camou er). The success of such attacks would largely depend on accurately pro ling clients, using some advanced machine learning techniques. Studying such attacks is an important part of our future work.

Properties	Camou er	SWEET	CloudTransport	Facet	CovertCast	Mailet	DeltaShaper	Freewave
Blocking Resistance	●	●	●	●	●	●	●	●
Deployment ease	●	●	●	●	●	●	●	●
No Cost of operation	◑	●	○	●	●	●	●	●
Requisite QoS	●	○	○	○	○	○	○	○
Collusion Defense	●	●	●	○	○	○	○	○

Table 5.2: A comparison of different features of existing tunnelling based anti-censorship schemes. ● - feature supported, ○ - feature unsupported, ◑ - feature partially supported.

Properties	Camou er	Proxy Systems	Domain Fronting	Decoy Routing	MassBrowser	Cache Browser	Conjure
Blocking Resistance	●	○	●	●	●	●	●
Deployment ease	●	●	●	○	●	○	○
No Cost of operation	◑	◑	○	○	○	○	○
Requisite QoS	●	◑	○	○	●	●	●
Collusion Defense	●	○	○	○	○	●	○

Table 5.3: A comparison of different features of other existing anti-censorship schemes. ● - feature supported, ○ - feature unsupported, ◑ - feature partially supported.

5.6 Comparison With Prior Systems

We now compare Camou er with existing systems (described in Sec. 5.2) based on different features that circumvention schemes strive to provide. We compared Camou er with existing tunnelling systems (e.g., SWEET, CloudTransport etc.), and also with other anti-censorship systems (e.g., Proxy based system, Decoy Routing etc.). A summarization of this comparison is done in Tab. 5.2 and Tab. 5.3 respectively.

- Blocking Resistance

  Tunnelling Systems: All such systems rely on using some underlying channel to covertly transport censored contents. Attempts to block the channel often incurs collateral damages to the adversary itself. Thus, all these systems, including Camou er provide blocking resistance, against adversaries that attempt to censor the entire communication channel (e.g., blocking all email services, IM services etc.).

  Other Systems: Most of the current promising systems also provide adequate blocking resistance by using the same principle of incurring collateral damage to the adversary. E.g., decoy routing requires the adversary to change nation-wide routing policies in order to prevent users from accessing the system. However, proxy based systems are relatively easier to block as the adversary merely needs to lter traf c destined to their IP addresses, incurring no collateral damages.

- Deployment Ease

  Tunnelling Systems: Camou er, like most tunnelling based systems, require installing programs at the client and server ends. Apart from that, there are no additional requirements.

  Other Systems: Decoy routing systems require collaboration from the ISPs, and thus poses a hurdle for deployment. Similarly, Cache Browser relies on content publishers to host their content on some CDNs, thereby posing deployment challenges[3]. Similar to Decoy Routing, Conjure requires ISPs assistance to install multiple servers with taps having access to all content transiting the ISP and thus pose challenges for deployment. Apart from these, other

[3]Non-CDN (blocked) websites can not be accessed by the CacheBrowser.

systems do not pose much hurdle for deployment.

- No Cost of Operation

  Tunnelling Systems: Except for CloudTransport, tunnelling based systems do not incur any upfront cost to the users, or to the content providers. CloudTransport requires hosting servers (e.g., Amazon s3), incurring monetary operational costs. Camou er does not incur any cost when the client has a peer to run Camou er server. Otherwise, if the client decides to run Camou er server on a cloud host, it would incur hosting charges.

  Other Systems: Domain fronting, Decoy routing and MassBrowser incur a monetary cost for their functioning. While Massbrowser and Domain Fronting requires subscription for services like Google App Engine , Amazon CloudFront etc., most Decoy Routing proposals (including Conjure) ideally require ISPs to change their existing network routing infrastructure. CacheBrowser requires that the censored content is hosted on CDNs, thereby incurring periodic subscription charges. Some proxy based systems such as VPNs, etc. also require subscription costs.

- Requisite QoS

Figure 5.10: Camou er vs Sweet: Download time of Alexa top-1k websites.

Tunnelling Systems: Most tunnelling systems fail to provide requisite QoS due to the limitation of the underlying channel they use for exchanging content. E.g., freewave uses VoIP, which encodes and transports data at low bit rates (19 Kbps), insuf cient for providing requisite QoS for applications like web browsing. Similarly, Facet, CovertCast and DeltaShaper use video streams to encode censored contents. However, factors like lossy video encoding of the

underlying platforms, result in unsuitable QoS for web browsing (e.g., Mcpherson et. al. [96] report that loading BBC news homepage along with three articles takes almost 120s. Further, DeltaShaper only provides an effective bandwidth of 2.56 Kbps [136]. Mailet and SWEET rely on emails which can carry signi cant content but inherently involve substantial delays (when emails transit email servers). However, Camou er due to its usage of IM applications (which involve minimum transit delay and substantial content carrying capacity) provides good QoS with signi cant improvement in comparison to other systems such as SWEET[4] (see Fig. 5.10). CloudTransport traf c also transits multiple hops (cloud provider, cloud bridge, etc.), incurring delays and reducing the overall QoS.

Other Systems: Except for Decoy Routing and Domain Fronting, all other systems provide satisfactory QoS. The recent deployment of Decoy Routing within an ISP do not provide adequate QoS for the end users [85, 133]. On the other hand, Domain Fronting incurs a signi cant amount of delay due to functional overheads (e.g., Domain Fronting introduces a delay of 16s more compared to Psiphon, when downloading a le [134]). Proxy systems such as Tor may sometimes incur substantial delays due to the selection of low bandwidth relays.

- Collusion Defense

Tunnelling Systems: It is generally dif cult to safeguard against a covert channel application provider who could collude with the censor to help identify clients. The CloudTransport architecture ensure that the cloud provider has no information about the destinations the client visits, even when the cloud provider colludes with the censor. SWEET attempts to resists collusion by distributing unique email IDs to individual users. In Camou er, the Camou er server IM IDs are known only to the speci c clients. This makes it dif cult for the adversary (that colludes with the IM provider) to block traf c by observing destination IDs in clients' requests.

Other Systems: Only CacheBrowser has a way to defend against collusion, as it uses frontends located outside the censor's boundary to access censored content. Thus, even if the CDN provider agrees to lter content in the censor's country, it may not do so in foreign countries, due to its own business motivations and regulatory compliance. In proxy systems, the collusion

[4] We could not obtain working codes of CloudTrasport for comparison. However, SWEET's code was publicly available.

of the VPN provider or proxy maintainer with the adversary, makes it relatively easy for the latter to identify clients. Similarly for Domain Fronting, the adversary by colluding with fronting service provider, could identify and block the Domain Fronting servers hosted on these services. Also, in Decoy Routing and Conjure, if the ISP colludes, then the respective systems would cease to function.

5.7 Discussion And Future Work

On blocking of IM apps

It can argued that an adversary may attempt to block IM apps, and in extreme case may block all existing IM apps to disrupt Camou er. In such a scenario, like any other tunneling system, Camou er may cease to functio E.g., if apps like Skype are completely blocked by the censor, then circumvention solutions like DeltaShaper [136], Facet [165], Freewave [135] etc. would be disrupted. But we believe that such a move by the censor could be prohibitively expensive leading to collateral damage as IM apps are extremely popular and have penetrated deeply into businesses and commercial spaces as well. However, it must be noted that, even if the censor allows only a single IM app to function, controlled by itself (ref. §5.5.2), Camou er would continue to function.

A more rational approach that could be opted by the censor is to selectively block only the suspicious IM IDs. Thus, a detailed analysis of such threats is already presented in §5.5.2.

On scalability of Camou er

As and when the popularity of Camou er grows, the system has the potential to scale up for larger deployments. Camou er uses IM apps as an overlay network to tunnel traf c. The ubiquity of IM applications is potentially bene cial for scaling Camou er into a distributed system with a large user base. Similar to Tor, Camou er volunteer could act as the Camou er servers. Thus, an increase in the number of Camou er clients could be handled by these distributed volunteer Camou er servers. We forsee that in future the popularity of Camou er may drive the recruitment of server hosters and maintainer, much like Tor.

Text vs attachment for transporting content

Camou er has a choice of selecting how it transports content using the underlying IM channel viz. as text, or as an attachment. The Camou er server initially used text messages to transport content. However, we discovered that in a few cases the contents of the websites accessed were truncated. Upon investigation, we found that IM applications generally restrict the volume of data that can be sent via a single text message. To overcome such restrictions, the Camou er server compressed the content before sending it to the other end. Thereafter, almost all the websites could be accessed, without data being truncated. Next, we also considered transporting our content as an attachment. This is because, with attachments we can transport more data as compared to text. However, we noticed a slight increase in the download time of Alexa top-1k websites, compared to when data was transported via texts (ref. Fig. 5.11).

Figure 5.11: Download time of Alexa top-1k websites using Camou er, when downloaded as text and as an attachment.

Thus, our design combines the best of both approaches. By default, we used text to transport content due to its obvious performance advantage. However, for the few cases where text could not be used, we relied on using attachments. Alternatively, the content can be segmented into multiple chunks, such that the length of each chunk is within the limits enforced on the length of text messages. However, unlike ours, this approach may incur extra round trips. and increase the overall delay.

On the use of Camou er

In general, tunnelling systems involve sending censored content using some underlying tunnelling channel (email, video streams, IM apps, etc.). It could be argued that when such circumvention

systems would become popular, the channel providers might be coerced by the censor to stop their services. Thus building circumvention solutions that use such channels might prove to be detrimental to the providers' business in future.

However, similar to Camou er, there already exist multiple anti-censorship solutions that hide the censored content in plain sight (e.g., CovertCast [96], DeltaShaper [136], Sweet [94], Protozoa [166], Skypemorph [74], Decoy Routing/Refraction Networking [150, 78, 9]). Notably, anti-censorship research inherently involves a trade-off between unobservability (of censored content) and the danger of inadvertently blocking the underlying channel. Thus to reduce the collateral damage, proper care must be taken while using such systems; to mitigate the inadvertent harm they can be used as a last resort when the communication over the Internet is severely disrupted.

5.8 Conclusion

We presented Camou er, a new anti-censorship system to provide unhindered access to information over the Internet. Camou er utilizes standard IM platforms (such as Whatsapp, Signal), to tunnel and transport censored content, and thus attempts to make it dif cult for the adversary to detect it.Camou er provides satisfactory performance, reliability, blocking resistance and deployment ease. Using the prototype implementation of Camou er on popular IM apps, we experimentally demonstrate that it provides acceptable performance for regular web browsing and bulk downloads. A detailed security analysis of Camou er highlights that it may be hard to be detected by an adversary (e.g., an ISP working at the behest of an authoritative regime).

Notably, all the systems that have been developed in the previous chapters assume that the adversary does not aim to completely disconnect the Internet for its citizens. However, in the recent past there is a rise in the cases of such deliberate Internet blackouts, performed by the adversaries. Thus, in the next chapter we try to develop a system that can provide access to Internet in such blackout regions.

Chapter 6

Dolphin

6.1 Introduction

The original idea of Internet was to provide a platform to facilitate free flow of information across the globe. This unhindered access to information has promulgated the rampant growth in all walks of life (including technology). On one hand the Internet is so vital to the modern world that free speech over it is considered a fundamental human right by the UN [41]. But on the other hand many censoring nation states attempt to disrupt the free flow of information (as per their convenience), opposing the original idea of Internet. As a result, in the past decade, there has been an exponential rise in the events of Internet censorship globally [34, 21, 167]. This has lead to an ongoing arms race between adversaries and free speech activists across the globe; adversaries continue to evolve various censorship techniques [180, 155, 168, 169, 170], whereas civil liberty activists counter them with wide range of novel circumvention systems [171, 133, 128, 172].

Traditional censorship involves restricting access to a particular resource (such as a website) on the Internet. However, in the recent past, an extreme form of Internet censorship, Internet shutdowns, has been on the rise. With such extreme measures, the adversary has gone a step ahead in the arms race by completely disabling Internet connectivity in a particular region. These shutdowns can range from a day to over a year in some cases (like in Myanmar and Chad [10]). The complete cut-off from Internet renders all available circumvention tools non-functional.

Figure 6.1: Overview of Dolphin's architecture.

Moreover, such a step has a severe impact on the people residing in the shutdown region. They are even devoid of accessing essential services over Internet, e.g., access to news, reporting power failures and outages, sending and receiving important emails etc. The recent COVID-19 pandemic further exacerbates the impact — a large population of the globe has moved to working online, both for professional and personal tasks. Thus, the regions with such shutdowns have been adversely impacted in these trying times. E.g., due to Internet shutdown in Myanmar some villages in the country were not even aware of the pandemic for many months [173]. It is even alarming that more and more countries are opting for Internet shutdowns. E.g., the number of countries who performed Internet shutdown increased from 25 in 2018 to 33 in 2019, with overall documented shutdown events increasing from 75 in 2016 to 213 in 2019 [174]. Considering that such trends are becoming common, it is plausible that more nation states opt for such measures [174]. Thus, it becomes an imperative to explore solutions using which people living in Internet shutdown regions could access basic Internet services like email, accessing news articles, tweets etc.

To that end, we introduce *Dolphin*, a novel system that can provide access to the aforementioned "lightweight" delay-tolerant Internet applications even during shutdowns. Dolphin uses only the cellular voice channel to transmit data. This idea rests on the observation that in Internet shutdown regions cellular voice connectivity is maintained (possibly for performing important executive and administrative tasks, by the governments). There are multiple documented evidences to support this observation [175, 176, 177, 178].

However, cellular voice channel is by design not built for running Internet applications. It is an unreliable, lossy, insecure and highly bandwidth constrained channel. We describe ahead how we overcome these challenges. But before that we begin by describing the overall working of