

ScrambleSuit: A Polymorph Network Protocol to Circumvent Censorship

Philipp Winter
Karlstad University

Tobias Pulls
Karlstad University

Juergen Fuss
Upper Austria University of
Applied Sciences

ABSTRACT

Deep packet inspection technology became a cornerstone of Internet censorship by facilitating cheap and effective filtering of what censors consider undesired information. Moreover, filtering is not limited to simple pattern matching but makes use of sophisticated techniques such as active probing and protocol classification to block access to popular circumvention tools such as Tor.

In this paper, we propose *ScrambleSuit*; a thin protocol layer above TCP whose purpose is to obfuscate the transported application data. By using morphing techniques and a secret exchanged out-of-band, we show that *ScrambleSuit* can defend against active probing and other fingerprinting techniques such as protocol classification and regular expressions.

We finally demonstrate that our prototype exhibits little overhead and enables effective and lightweight obfuscation for application layer protocols.

Keywords

Tor, bridge, pluggable transport, active probing, censorship, circumvention

1. INTRODUCTION

We consider deep packet inspection (DPI) harmful. While originally meant to detect attack signatures in packet payload, it is ineffective in practice due to the ease of evasion [1, 2, 3]. At the same time, DPI technology is increasingly used by censoring countries to filter the free flow of information or violate network neutrality [4]. We argue that what makes DPI particularly harmful is the *asymmetry of blocking effectiveness*, i.e., it is hard to stop motivated and skilled network intruders but very easy to censor ordinary user's Internet access. DPI technology ultimately fails to protect critical targets but succeeds in filtering the information flow of entire countries.

Numerous well-documented cases illustrate how DPI technology is used by censoring countries. Amongst others, China is using it to filter HTTP [5] and rewrite DNS responses [6]. Iran is known to use DPI technology to conduct surveillance [7]. In Syria, DPI technology is used for the same purpose [8]. Even more worrying, SSL interception proxies, an increasingly common feature of DPI boxes, are used to transparently decrypt and inspect SSL sessions which effectively breaks SSL's confidentiality and given the rise of opaque Internet traffic [9], there is no reason to believe that this trend will decrease.

The rise of Internet censorship led to the creation of numerous circumvention tools which engage in a rapidly developing arms race with the maintainers of censorship systems. Of particular interest to censoring countries is the Tor network [10]. While originally designed as a low-latency anonymity network, it turned out to be an effective tool to circumvent censorship. The growing success of Tor as circumvention tool did not remain unnoticed, though. Tor is or was documented to be blocked in many countries including Iran [11], China [12] and Ethiopia [13], just to name a few. We argue that many circumvention tools—Tor included—suffer from two shortcomings which can easily be exploited by a censor.

First and most importantly, they are vulnerable to *active probing* as pioneered by the Great Firewall of China (GFW) [12]: the GFW is able to block Tor by first identifying potential Tor connections based on the TLS client cipher list. If such a signature is found on the wire, the GFW reconnects to the suspected Tor bridge and tries to “speak” the Tor protocol with it. If this succeeds, the GFW blacklists the respective bridge. Active probing is not only used to discover Tor but—as we will discuss—also VPN [14] and obsf2 [15], which is a censorship-resistant protocol. From a censor's point of view, active probing is a promising strategy which greatly reduces collateral damage caused by inaccurate signatures. Also, active probing is non-trivial to defend against because censors can easily emulate real computer users.

Second, circumvention tools tend to exhibit a certain “flow signature” which typically remains static. An example is Tor's characteristic 586-byte signature (cf. §5.1). If a censor is able to deploy high-accuracy classifiers trained to recognise these very flow signatures, the respective protocol is blocked. Censorship-resistant protocols are unable to evade these filters by changing their flow signature.

In this work, we present *ScrambleSuit*; a blocking-resistant transport protocol which tackles the two above mentioned problems. *ScrambleSuit* defines a thin protocol layer on top of TCP which provides lightweight obfuscation for the transported application layer protocol. As shown in Figure 1, *ScrambleSuit* is independent of its application layer protocol and works with any application supporting SOCKS. As a result, we envision *ScrambleSuit* to be used by, among other

Tor	VPN	...
ScrambleSuit		
TCP		
IP		

Figure 1: *ScrambleSuit*'s protocol stack.

protocols, Tor and VPN to tackle the GFW’s most recent censorship upgrades.

In particular, ScrambleSuit exhibits the following four features:

Pseudo-random payload: To an observer, ScrambleSuit’s entire traffic is computationally indistinguishable from randomness. As a result, there are no predictable patterns which would otherwise form suitable DPI fingerprints. This renders regular expressions for the purpose of identifying ScrambleSuit useless.

Polymorph: Despite the pseudo-random traffic, a censor could still block our protocol based on flow characteristics such as the packet length distribution. ScrambleSuit is, however, able to change its shape to make it harder for classifiers to exploit flow characteristics.

Shared secret: We defend against active probing by making use of a secret which is shared between client and server and exchanged out-of-band. The server only answers to requests if knowledge of the secret is proven by the client.

Usable: We seek to maximise ScrambleSuit’s usability. Our protocol easily integrates in Tor’s existing ecosystem and does not require architectural changes. Furthermore, the moderate protocol overhead, as shown in §5, facilitates comfortable web surfing.

Blocking-resistant protocols can be split into two groups. While the first group strives to *mimic typically whitelisted protocols* such as HTTP [16] and Skype [17], the second group aims to *look like randomness* [18, 19, 20]. Randomised protocols have the shortcoming of not being able to survive a whitelisting censor. Nevertheless, we decided in favour of randomising because mimicing comes at the cost of high overhead, is difficult to do correctly [21] and we consider whitelisting on a nation scale—at least for most countries—unlikely even though it is often done in corporate networks. So, instead of maximising obfuscation while maintaining an acceptable level of usability, we *maximise usability* while keeping an *acceptable level of obfuscation*.

The contributions of this paper are as follows.

- We propose ScrambleSuit, a blocking-resistant transport protocol.
- We propose two authentication mechanisms based on shared secrets and polymorphism as a practical defence against active probing and protocol classifiers.
- We implement and evaluate a fully functional prototype of our protocol.

We finally point out that unblockable network protocols do not exist. After all, censors could always “pull the plug” as it was already done in Egypt [22] and Syria [23]. By proposing ScrambleSuit, we do not claim to end the arms race in our favour but rather to raise the bar once again.

The remainder of this paper is structured as follows. In §2 we discuss related work which is then followed by an architectural overview in §3. §4 then discusses ScrambleSuit’s design in detail. The protocol is then evaluated in §5 and the results discussed in §6. We finally conclude the paper in §7.

2. RELATED WORK

2.1 Protocol Identification

The identification of protocols is typically motivated by quality of service, traffic shaping and accounting – but also censorship. In order to block protocols, they have to be identified first. Many protocol identification techniques fail in the face of protocols which make an active effort to remain undetected. This led to the research community finding ways to, e.g., detect protocol tunneling in HTTP and SSH [24], the Skype protocol [25] or encrypted traffic [26].

Of particular relevance is the work of Hjelmvik and John [27]. The authors investigated to which extent supposedly obfuscated protocols such as Skype, BitTorrent’s message stream encryption and Spotify can be identified. Based on their findings, Hjelmvik and John suggest evasion techniques for protocol designers which should make it harder to identify obfuscated protocols. Some of our design decisions were motivated by their suggestions. Similar to Hjelmvik and John, Wiley proposed a framework to dynamically classify network protocols based on Bayesian models [28]. This is an important first step towards the ability to compare and evaluate blocking-resistant transport protocols.

2.2 Protocol Obfuscation

The Tor project developed a blocking-resistant protocol called obfs2 [18]. The protocol implements an obfuscation layer on top of TCP and transports Tor traffic. A passive man-in-the-middle (MITM), however, can decrypt obfs2 traffic. The successor, obfs3 [19], uses a customised Diffie-Hellman handshake to solve this problem. However, both, obfs2 and obfs3 can be actively probed and do not disguise flow properties. In fact, the GFW is already blocking obfs2 bridges by actively probing them [15]. Later in this paper, we extend obfs3’s handshake to be resistant against active probing.

Wiley’s Dust protocol [20] compares to obfs2 and obfs3 in that Dust payload looks like random data. The key exchange is handled out-of-band. Dust also employs packet padding to camouflage packet lengths. However, unlike ScrambleSuit, Dust does not consider inter arrival times.

Weinberg et al. presented StegoTorus [16], a framework for obfuscation modules similar to obfsproxy which is developed by the Tor project [29]. StegoTorus can complicate protocol identification on the application layer as well as on the transport layer. Tor connections can be multiplexed over multiple TCP connections and the application layer is camouflaged by mimicing a cover protocol such as HTTP.

SkypeMorph, as presented by Moghaddam et al. [17] compares to StegoTorus in that it disguises Tor traffic by mimicing an existing protocol; in this particular case Skype video traffic. As long as the censor does not decide to block the cover protocols, SkypeMorph and StegoTorus are able to survive a whitelisting censor. ScrambleSuit differs from SkypeMorph and StegoTorus since it does not mimic a cover protocol. In fact, Houmansadr et al. claim that protocol mimicing—as opposed to tunneling—is a flawed approach due to the immense difficulty of mimicing a protocol correctly [21]. The authors showed that SkypeMorph and StegoTorus differ from their respective cover protocols in numerous ways.

Many blocking-resistant tools blindly employ different obfuscation strategies in the hope to stay under the radar.

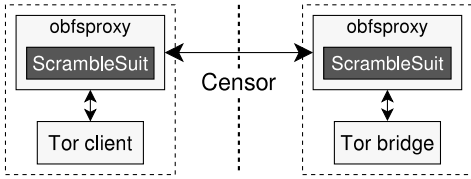


Figure 2: ScrambleSuit is a module for obfsproxy which provides a SOCKS interface for local applications. The traffic between two obfsproxy instances is disguised by ScrambleSuit.

Dyer et al. suggest the opposite [30]. The authors propose to actively learn the regular expressions used by DPI boxes. This knowledge is then used to map cipher text to regular expressions which are guaranteed to pass the filters. This requires the regular expressions of DPI boxes to be known which is typically difficult in practice.

Lincoln et al. proposed DEFIANCE [31]: an architecture to protect Tor bridges from being probed and their respective descriptors¹ from being harvested by crawlers. The authors accomplish these goals by developing a novel rendezvous protocol as well as a technique called address-change signaling.

A solution to the problem of IP address blocking is provided by Fifield et al. [32]. Instead of relying on long-lived static bridge IP addresses, the authors propose to use short-lived proxies which are run by web users visiting special cooperating web sites. A practical problem remains to be solved, however: clients making use of these so-called flash proxies must be able to accept incoming TCP connections. This is not always possible with censored users behind NAT boxes.

2.3 Undetectable Authentication

Vasserman et al. proposed an undetectable authentication system based on port knocking [33]. Their system, SilentKnock, does however have operating system dependencies and does not protect against connection hijacking.

Smits et al. adapted SilentKnock to better work with Tor bridges [34]. The result is called BridgeSPA. When using BridgeSPA, clients can authenticate themselves towards a bridge with just a TCP SYN segment. If the authentication does not succeed, the bridge does not respond with a SYN/ACK segment and the bridge appears to be offline. Just like SilentKnock, BridgeSPA does not protect against connection hijacking and faces a number of practical problems such as the inability to cope with NAT and the dependence on Linux kernels. While ScrambleSuit can not hide its “aliveness”, it is not hindered by NAT or connection hijacking.

3. ARCHITECTURAL OVERVIEW

ScrambleSuit is a module for obfsproxy which is an obfuscation framework developed by the Tor project [29]. As long as obfsproxy is running on the censored client as well as on the server, network traffic in between both communication points can be obfuscated as dictated by the obfuscation modules. As illustrated in Figure 2, obfsproxy acts as a

¹A bridge descriptor is essentially a tuple containing the bridge’s IP address, port and fingerprint.

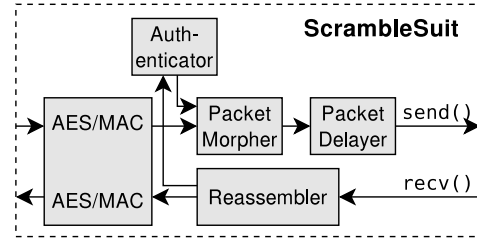


Figure 3: Internally, ScrambleSuit handles authenticated encryption of application data, client authentication as well as flow reshaping using a packet morpher and delayer.

proxy between the Tor client and the Tor bridge. While specifically designed for Tor, obfsproxy can be used by any application as long as it supports the SOCKS protocol.

Internally, ScrambleSuit is composed of several components which are depicted in Figure 3. Outgoing network data is first encrypted and then chopped into pieces (or padded) by the packet morpher. Before these pieces are then sent over the wire, the packet delayer uses small sleep calls to disguise inter arrival times. Finally, incoming network data is first reassembled to complete ScrambleSuit protocol messages and after decryption finally passed on to the local application.

We aim to conceal several aspects of Tor’s communication characteristics. We chose these characteristics based on the work by Hjelmvik and John [27, 35] as well as current DPI capabilities.

Payload By encrypting all ScrambleSuit traffic, we eliminate all payload fingerprints such as Tor’s TLS cipher list [12].

Packet length distribution Among other things, we seek to get rid of Tor’s characteristic 586-byte packets [16, 36]. We do so by morphing Tor’s packet length distribution to a randomly chosen distribution.

Inter arrival times Similar to the packet length obfuscation, we camouflage the inter arrival times by employing small and random sleep intervals before writing data on the wire.

3.1 Threat Model

Our adversary is a *nation-state censor* who desires to block unwanted network protocols and services which would otherwise allow users within the censoring regime the retrieval of unfiltered information or to evade the national filtering system. The censor is making use of payload analysis, flow analysis as well as active probing to identify and then block undesired protocols.

The censor further has full *active and passive* control over the national network. The censor can passively monitor all traffic entering and leaving its networks in line rate. We further expect the censor to actively tamper with traffic; namely to inject, drop and modify traffic as well as hijack TCP sessions. We further expect the censor to select a subset of suspicious traffic for further inspection on the slow path². This could involve *active probing* as done by the

²We define the *slow path* as the minute analysis of a small

GFW in order to block the Tor network [12]. We model our censor to also conduct active MITM attacks. While we believe that passive analysis and active probing are significantly easier to deploy, there is evidence that censors are starting to—or at least have the ability to—conduct active MITM attacks as well [37].

Our adversary is also training and deploying *statistical classifiers* to identify and block protocols. While computationally expensive, it would be imaginable that a censor uses this strategy at least on the slow path and perhaps even on the fast path when using inexpensive flow features.

3.1.1 Adversary Limitations

We expect the censor to be subject to economical constraints. In particular, we assume that the censor is not using a whitelisting approach meaning that only well-defined protocols pass the national filter. Whitelisting implies significant *over-blocking* and we expect this approach to collide with the censor’s economical incentives. We also expect the censor to not block protocols when there is only weak evidence for the protocol being blacklisted. This is a direct consequence of avoiding over-blocking by minimising collateral damage.

Finally, we assume that the censor does not have access to or can otherwise influence censored users’ computers. Once again, we believe that such a scenario is likely to occur in corporate networks but not on a national scale.

4. PROTOCOL DESIGN

This section will discuss ScrambleSuit’s defence against active probing, its encryption, encoding and header format as well as how we achieve polymorphism.

4.1 Thwarting Active Probing

We defend against active probing by proposing two mutual *authentication mechanisms* which rely on a secret which is shared *out-of-band*. A ScrambleSuit connection can only be established when both parties can prove knowledge of this very secret. While our first authentication mechanism (see §4.1.2) is designed to work well in Tor’s ecosystem, our second mechanism (see §4.1.3) provides additional security and efficiency if ScrambleSuit is used by other application protocols such as VPN.

With respect to Tor, there *already exists* an out-of-band communication channel which is used to distribute bridge descriptors to censored users. Naturally, we make use of this channel. If, however, ScrambleSuit is used to tunnel protocols other than Tor, users have to handle out-of-band communication themselves.

4.1.1 Proof-of-Work (Again) Proves Not to Work

Before deciding in favour of using a secret exchanged out-of-band, we investigated the suitability of client puzzles. Puzzles—a variant of proof-of-work schemes—could be used by a server to time-lock a secret. This secret can then only be unlocked by clients by spending a moderate amount of computational resources on the problem. One particular puzzle construction, namely time-lock puzzles as proposed by Rivest et al. [38], provides appealing properties such

traffic subset as opposed to the *fast path* which covers the majority of all network traffic and, as a result, has to be processed quickly.

as deterministic unlocking time, asymmetric work load and inherently sequential computation which means that adversaries in the possession of highly parallel architectures have no significant advantage over a client with a single CPU.

While a *single* client puzzle can not be solved in parallel, a censor is able to solve *multiple* puzzles in parallel by assigning all puzzles to the available CPU cores. This is problematic because our threat model includes censors with powerful and parallel architectures. After estimating the Tor bridge churn rate, we came to the conclusion that client puzzles would probably not be able to increase a well-equipped censor’s work load beyond the point of becoming *impractical*; at least not without becoming impractical for *clients as well*. This balancing problem is analogous to why proof-of-work schemes are believed to be unpractical for the spam problem as well [39].

In summary, proof-of-work schemes would not require a shared secret but we believe that this small usability improvement would come at the cost of greatly reduced censorship resistance. A censor in the possession of powerful computational resources would certainly be slowed down but could ultimately not be stopped. Active probing would simply become a matter of investing more resources.

4.1.2 Session Tickets

We now discuss the first of our two authentication mechanisms. A client can authenticate herself towards a ScrambleSuit server by redeeming a *session ticket*. A session ticket needs to be obtained only once out-of-band. Subsequent connections are then bootstrapped using tickets issued by the server during the respective previous connection. A real world analogy would be a person redeeming a ticket in order to gain access to a football stadium. Upon entering the stadium (i.e., successful authentication), the guards give the person a new ticket so that she is able to return for the next match. The same procedure then happens for the next match.

Session tickets are standardised in RFC 5077 [40] and part of TLS since version 1.0. We employ only a subset of the standard since we do not need its full functionality.

The basic idea is illustrated in Figure 4. ScrambleSuit servers issue new session tickets \mathcal{T}_{t+1} which contain a future shared master key k_{t+1} and an issue date d indicating the ticket’s creation time. Session tickets are encrypted and authenticated with secret keys k_S ³ only known to the server, i.e., $\mathcal{T}_{t+1} = \text{Enc}_{k_S}(k_{t+1} \parallel d)$. As a result, a ticket \mathcal{T} is opaque to the client. Note that a client, when obtaining a ticket, also has to learn the master key k_{t+1} in order to be able to derive the same session keys as the server; so clients always obtain the tuple $(k_{t+1} \parallel \mathcal{T}_{t+1})$. Session tickets have the advantage that the server *does not have to keep track of* issued tickets. Instead, the server’s state is outsourced and stored by clients which greatly reduces a server’s load.

Whenever a client successfully connects to a ScrambleSuit server, the server issues a new ticket concatenated to the according master key $(k_{t+1} \parallel \mathcal{T}_{t+1})$ for the client. The tuple is placed in a special ScrambleSuit control message (see §4.2). The new ticket is sent immediately after successful bootstrapping.

We mentioned earlier that a ScrambleSuit server man-

³For simplicity, we refer to these two symmetric keys as just k_S while they are in fact two keys: one for encryption and one for authentication.

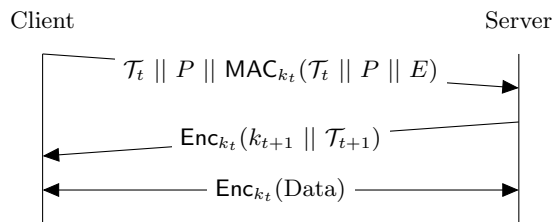


Figure 4: The client redeems a valid session ticket \mathcal{T}_t containing the master key k_t . The server responds by issuing a new ticket \mathcal{T}_{t+1} for future use. Both parties then exchange application data.

ages secret keys k_S which are used to encrypt and authenticate session tickets. This prevents clients from tampering with tickets and the server can verify that a newly received and authenticated ticket was, in fact, issued by the server. Servers rotate their k_S keys after a period of seven days. After the generation of new k_S keys, the superseded keys are kept for another seven days in order to decrypt and verify (but not to issue!) tickets which were issued by the superseded keys. As a result, tickets are always valid and redeemable for a period of *exactly seven days*; no matter when they were issued. As a result, as long as a user keeps reconnecting to a ScrambleSuit server at least once a week, *key continuity* is ensured and there is no need for additional out-of-band communication.

A censor could now conduct traffic analysis by looking for TCP connections which always begin with the client sending $|\mathcal{T}|$ bytes to the server. To obfuscate the ticket length, we introduce random padding P and authenticate the ticket \mathcal{T} as well as the padding P by computing $\text{MAC}_{k_t}(\mathcal{T} || P || E)$ with k_t being the shared master secret which the client obtained together with the ticket and E discussed in the following paragraph. Both parties will use k_t to derive session keys as discussed in §4.2. The server knows that all bytes of the handshake were successfully received when the last bytes form a valid MAC over the previous bytes. The exact amount of random padding is determined by the packet morpher discussed in §4.3.1. We use HMAC-SHA256-128 for the MAC.

At this point, a censor could still intercept tickets and replay them. This would make the server issue a new ticket for the censor. While the censor would not be able to read the resulting ScrambleSuit control message—the shared master key k_t would be unknown—the fact that a replay attack triggers a response can be suspicious. We prevent replay attacks, or in other words ticket double spending, by caching master keys k_t . If a server encounters a cached k_t , it does not reply to prevent the censor from learning the server’s state. We begin to cache a k_t after a new session ticket was issued and the client acknowledged that she correctly received the new ticket by using a special ScrambleSuit message type (see §4.2). To reduce the amount of keys to cache, we add the value E to the MAC. It refers to the Unix epoch divided by 3600, i.e., the current time with a granularity of one hour. While this requires client and server to have loosely synchronised clocks, the server has to cache redeemed keys for a period of only one hour instead of seven days.

Session tickets already provide a strong level of protection. Active probing and replay attacks are foiled while forward

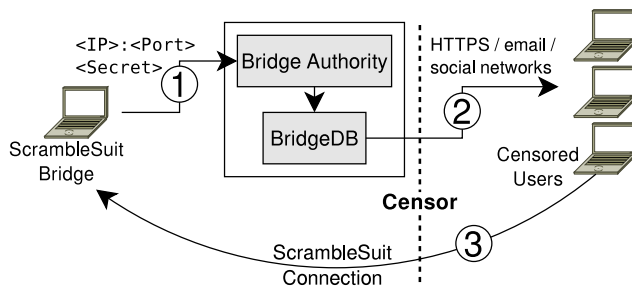


Figure 5: ScrambleSuit bridges send their descriptor to the bridge authority ①. From there, it is distributed to censored users who learn about IP address, port and the secret *out-of-band* ②. Finally, direct connections can be established ③.

secrecy is provided. Therefore, we envision session tickets to be satisfactory for most application protocols to be tunneled over ScrambleSuit.

Session tickets do not, however, integrate well with Tor’s existing ecosystem. The reason lies in how Tor bridges are distributed to users. The process is illustrated in Figure 5. Volunteers will set up ScrambleSuit bridges which then publish their descriptors—including IP address, port and secret—to the bridge authority (1) which feeds this information into the BridgeDB component. In the next step, the gathered descriptors have to be distributed to censored users (2). The two primary distribution channels are email and HTTPS [41]. Users can ask for bridges over email or they can visit the bridge distribution website⁴ and obtain a set of bridges after solving a CAPTCHA. The problem is that *one bridge descriptor* is typically shared by *many users*. All these users would end up with an identical session ticket. This causes two severe problems. First, our replay protection mechanism does not allow reuse of session tickets. Second, session ticket reuse would lead to identical byte strings at the beginning of a ScrambleSuit handshake which would be a strong distinguisher. These problems lead us to our *second authentication mechanism* which is optimised for Tor and can function with a secret which is shared by many users as shown in the scenario in Figure 5.

4.1.3 Uniform Diffie-Hellman

Our second authentication mechanism is an extension of the Uniform Diffie-Hellman (UniformDH) handshake which was proposed in the obfs3 protocol specification [19, §3]. obfs3’s handshake makes use of uniformly distributed public keys which are only negligibly different from random bytes. As a result, UniformDH can be used to agree on a master key k_t without a censor knowing that Diffie-Hellman is used.

UniformDH is based on the 1536-bit modular exponential group defined in RFC 3526 [42]. When initiating a UniformDH handshake, the client first generates a 1536-bit private key x . The least significant bit of x is then unset in order to make the number even. The public key X is defined as $X = g^x \pmod{p}$ where $g = 2$. The server computes its private key y and its public key Y the same way. To prevent a censor from learning that X is a quadratic residue mod p —a clear distinguisher—the client randomly chooses to send either X or $p - X$ to the server. The server can then derive

⁴URL: <https://bridges.torproject.org>.

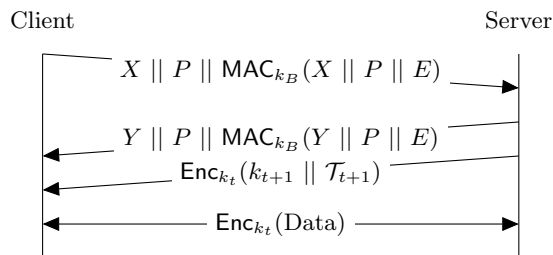


Figure 6: After client and server agreed on the master key k_t using Diffie-Hellman, the server is issuing a new session ticket for the client. Afterwards, both parties exchange application data.

the shared master secret by calculating $k_t = X^y \pmod p$. Since the private keys x and y are even, the exponentiations $X^y \pmod p$ and $(p - X)^y \pmod p$ result in the same shared master secret.

4.1.4 Extending Uniform Diffie-Hellman

In its original form, the UniformDH construction does not protect against active probing. A censor who suspects UniformDH can simply probe the supposable bridge and opportunistically initiate a UniformDH handshake. To prevent that attack, we now turn UniformDH’s anonymous handshake into an authenticated handshake in order to be resistant against active attacks.

We do so quite similar to the session tickets discussed in §4.1.2. As depicted in Figure 6, we concatenate pseudo-random padding P and a MAC to the public keys X and Y . The MAC authenticates the respective public key as well as the padding. The MAC is keyed by a shared secret k_B which is distributed together with the Tor bridge’s IP:port tuple over email or HTTPS (cf. step ① in Figure 5). As with tickets, the server and client know that the handshake message was fully received when the last received bytes form a valid MAC over the previous bytes. Note that k_B can be reused because it is only used to key the MAC. The handshake is conducted using UniformDH with randomly chosen public keys. As a result, two subsequent UniformDH handshakes based on the same k_B will appear to be different to a censor. We defend against replay attacks by adding E , the Unix epoch divided by 3600, to the MAC and cache the MAC for a period of one hour.

A successful UniformDH key agreement is followed by the server issuing a session ticket for the client. The client will then redeem this ticket upon connecting to the server the next time. Accordingly, we expect the UniformDH handshake to be done *only once*, namely when a Tor client connects to a bridge for the first time. From then on, session tickets will be used to connect to the same bridge.

To a censor, the payload of both authentication schemes is computationally indistinguishable from randomness. As a result, a censor who is assuming that a server is running ScrambleSuit is unable to tell whether a client successfully authenticated herself by using UniformDH or by redeeming a session ticket.

We finally stress that bootstrapping ScrambleSuit using UniformDH provides *less security* than when bootstrapped using session tickets. Since the secret key k_B for UniformDH will be used by multiple clients, a malicious client in the pos-

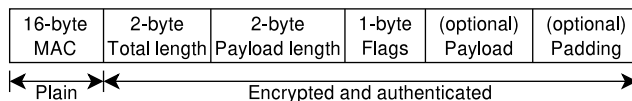


Figure 7: ScrambleSuit’s message header format. The encrypted part is authenticated by a HMAC-SHA256-128. The entire message is computationally indistinguishable from randomness.

session of k_B and who is able to eavesdrop on the connection of another client using the same ScrambleSuit server can conduct active MITM attacks. While Tor does protect against active MITM attacks, this can be problematic for application protocols other than Tor. Therefore, we emphasise that session tickets are the preferred authentication mechanism whereas our UniformDH extension’s sole purpose is to make ScrambleSuit work well in Tor’s infrastructure.

4.1.5 Usability Considerations

In order for a user to successfully connect to a ScrambleSuit server, she needs a *triple*: an IP address, a TCP port and a secret which is either the UniformDH secret k_B or a session ticket tuple $(k_t || T_t)$. We expect these triples to be distributed mostly electronically; over email, instant messaging programs or online social networks. As a result, a user can simply copy and paste the entire triple into her obfsproxy configuration file.

We do, however, also expect limited verbal distribution of ScrambleSuit triples, e.g., over a telephone line. To facilitate this, we define the encoding format of secrets and tickets to be Base32. Base32 strings consist of the letters A–Z, the numbers 2–7⁵ as well as the padding character “=”. Since there is no distinction between uppercase and lowercase letters, we hope to make verbal distribution less confusing and error-prone. After all, a ScrambleSuit bridge descriptor would look like: `Bridge scramblesuit 1.2.3.4:443 password=NCA6I6GZZD42BWUB`. We believe that the prefix `password=` will find more acceptance among users than simply appending the secret.

4.2 Header Format and Confidentiality

Our protocol employs a custom message format whose header is illustrated in Figure 7. ScrambleSuit exchanges variable-sized messages with optional padding which is discarded by the remote machine.

The first 16 bytes of the header are reserved for a HMAC-SHA256-128 which protects the integrity and authenticity of the protocol message. In accordance with the encrypt-then-MAC paradigm, the HMAC is computed over the encrypted remainder of the message. The secret key required by the HMAC is derived from the shared master key.

The HMAC is followed by two bytes which specify the total length of the protocol message. ScrambleSuit’s maximum transmission unit is 1460-byte-sized messages. Together with a minimal IP and TCP header, this adds up to 1500-byte packets which fill an Ethernet frame. In order to be able to distinguish padding from payload, the next two bytes determine the payload length. If no padding is used, the payload length equals the total length.

⁵The numbers 0 and 1 are omitted to prevent confusion with the letters I and O.

To separate application data from protocol signaling, we define a 1-byte message flag field. The first bit signals application data in the message body whereas a message with the second bit set contains a newly issued session ticket. The third bit (which can be set together with the first bit) confirms the receiving of a session ticket. We reserve the remaining bits for future use. In particular, they could be used to negotiate flow properties as discussed in §4.3.

The header is then followed by the message payload which contains the application protocol transported by *ScrambleSuit*. We employ encryption in order to hide the application protocol, the padding as well as *ScrambleSuit*'s header. With regard to Tor, this means that the already encrypted Tor traffic is wrapped inside yet another layer of encryption. For encryption, we use 256-bit AES in counter mode. The counter mode effectively turns AES into a stream cipher. We use two symmetric keys: one for the traffic $C \rightarrow S$ and one for $S \rightarrow C$. Both symmetric keys as well as the respective nonces for the counter mode are derived from the shared master secret using HKDF based on SHA256 [43].

4.3 Changing Shape

So far, we discussed defences against censors who analyse packet payload or conduct active attacks to reveal *ScrambleSuit*'s presence. However, a censor could also make use of *traffic analysis*. In this section, we propose lightweight countermeasures to diminish—but not to defeat!—such attacks. In particular, we will teach *ScrambleSuit* how to change its “protocol shape”⁶.

Our definition of *ScrambleSuit*'s shape is twofold: we consider *packet lengths* and *inter arrival times*. While our transported data is encrypted and exhibits no structure, these flow metrics can still leak information about the application protocol [44, 45, 46]. As a result, we seek to randomise these characteristics in order to decrease the accuracy of protocol classifiers used to detect *ScrambleSuit*.

In general, the kernel's TCP stack is responsible for packet lengths. In order to affect packet lengths in user space, we deactivate Nagle's algorithm which seeks to avoid unnecessarily small TCP segments. This comes, however, at the cost of increased protocol overhead.

The randomisation of packet lengths as well as inter arrival times is based on a randomly generated discrete probability distribution. We generate these distributions by first determining the amount of bins n which is uniformly chosen from the set $\{1..100\}$. In the next step, we assign each bin b_i for $1 \leq i \leq n$ a probability by randomly picking a value in the interval $]0, 1 - \sum_{i=1}^n b_{i-1}[$ for $b_0 = 0$. The following gives an example of four assignments.

$$b_0 \leftarrow 0 \tag{1}$$

$$b_1 \stackrel{R}{\leftarrow}]0, 1 - b_0[\tag{2}$$

$$b_2 \stackrel{R}{\leftarrow}]0, 1 - b_0 - b_1[\tag{3}$$

$$b_n \stackrel{R}{\leftarrow}]0, 1 - b_0 - \dots - b_{n-1}[\tag{4}$$

We will show in §5 that this naive approach turns out to be sufficient to obfuscate Tor's flow characteristics. A specialised algorithm—e.g., to optimise throughput—would be conceivable but is beyond the scope of this paper.

⁶This happens analogous to the *scramble suits* in Philip K. Dick's novel “A Scanner Darkly”.

4.3.1 Packet Length Adaption

It is well known that a network flow's packet length distribution leaks information about the network protocol [27, 44, 47] and even the content [48, 46]. For instance, a large fraction of Tor's traffic contains 568-byte packets which is the result of Tor's internal use of 512-byte cells plus TLS' header (see Figure 9 and 10). These 568-byte packets form a strong distinguisher which can be used to detect Tor by simply capturing a few dozen network packets as shown by Weinberg et al. [16]. To defend against such simple applications of traffic analysis, we modify *ScrambleSuit*'s packet length distribution.

An efficient way to morph a source distribution to a target distribution was proposed by Wright, Coull and Monrose [49]. Their concept, traffic morphing, relies on the computation of a morphing matrix to minimise the overhead when morphing a source distribution to a target distribution. Unfortunately, we cannot make use of traffic morphing because our target distribution is dynamic which would require frequent recomputation of the morphing matrix which is an expensive operation. This would lead to unnecessary CPU load on the client as well as on the bridge. Furthermore, our source distribution is not known since *ScrambleSuit* is designed to be able to handle arbitrary application protocols.

Instead, we adopt *naive sampling* to disguise the application protocol's packet length distribution. Every time *ScrambleSuit* establishes a connection to a server, it randomly generates a fresh discrete probability distribution as discussed earlier. Every bin in the probability distribution is uniformly chosen from the set $\{1..1460\}$. The newly generated distribution is then randomly sampled for every chunk of application data, *ScrambleSuit* is about to send over the wire. After a sample length is obtained, our algorithm either *a*) pads the current packet to fit the sample's length or *b*) splits and sends it and then proceeds to morph the remaining data the same way.

4.3.2 Inter Arrival Time Adaption

Analogous to the packet length distribution, the distribution of inter arrival times between consecutive packets has discriminative power and could be used by censors to identify protocols [50]. In contrast to the packet length distribution, inter arrival times are frequently distorted by network jitter, overloaded middle boxes and the communicating end points. Nevertheless, we believe that it would be no sound strategy to assume the network to be unreliable enough to render measurements useless. Therefore, *ScrambleSuit* is also able to modify its inter arrival times. The mechanism is the same as for the packet length adaption: a random distribution is generated and then random samples are drawn from it. The samples are the parameters for short `sleep()` calls which are invoked prior to sending data to the remote end.

We are only able to increase inter arrival times but not to decrease them. Increased inter arrival times have a direct negative effect on throughput and can easily turn into a nuisance for users when getting too high. As a result, we keep the sleep intervals within the interval of $]0, 100[$ milliseconds. We believe that this interval represents a reasonable trade-off between obfuscation and throughput as we will show in §5.1.

4.3.3 Shortcomings

It is important to note that for a censor armed with a well-chosen set of features, traffic analysis can be a powerful attack and strong defences are believed to be expensive [45, 46]. We made an effort to disguise obvious flow features while keeping throughput high enough to facilitate comfortable web surfing and enable the transportation of low-latency applications over ScrambleSuit.

A censor can still measure derived flow metrics such as “total bytes transferred”, packet directions or the “burstiness” of ScrambleSuit’s behaviour. These metrics would be expensive to disguise and by exploiting them, a censor would at least be able to guess whether ScrambleSuit’s transported application is a bulk file transfer or a request-response protocol. Nevertheless, traffic analysis does not give censors a *certain answer*. False positives are always a problem and can lead to overblocking. As mentioned in our threat model, we believe that the censor might use traffic analysis to select a subset of traffic for closer inspection but not to block flows.

5. EXPERIMENTAL EVALUATION

We implemented a prototype of ScrambleSuit in the form of several Python modules for obfsproxy. Our prototype consists of approximately 2000 lines of code. The measurements below were all conducted using this prototype.

As illustrated in Figure 8, our experimental setup consisted of two Debian GNU/Linux machines which were connected via a router which performed the measurements. All three machines were connected over Fast Ethernet. We expect this setup to be ideal for a censor because it does not cause IP fragmentation or high latency due to overloaded middle boxes. As a result, we believe that a censor would do worse in practice. Both of our machines were running Tor v0.2.4.10-alpha and obfsproxy. The Tor bridge was configured to be private and was only used by our client. The bridge then relayed all traffic into the public Tor network. Note that ScrambleSuit is only “spoken” in between the client and the bridge.

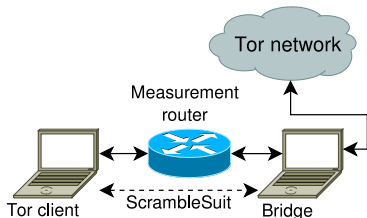


Figure 8: The experimental setup.

5.1 Blocking Resistance

To create network traffic for our measurements, we downloaded the 1 MB Linux kernel v1.0 from kernel.org⁷ on the client. We downloaded the file once over Tor⁸ and 5 times over ScrambleSuit.

⁷<https://www.kernel.org/pub/linux/kernel/v1.0/linux-1.0.tar.bz2>.

⁸In fact, we downloaded the file many times over Tor and found that consecutive runs differed mostly in the ratio between 586-byte and 1448-byte packets. As a result, we only plot one run.

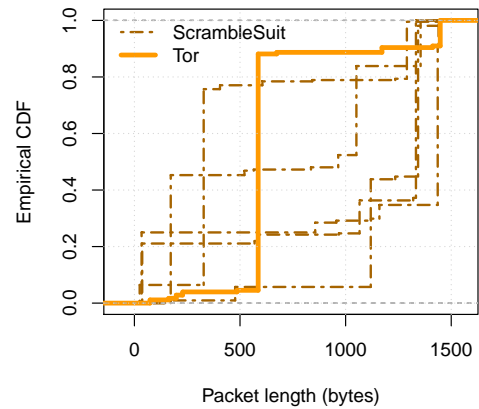


Figure 9: Client-to-server packet lengths.

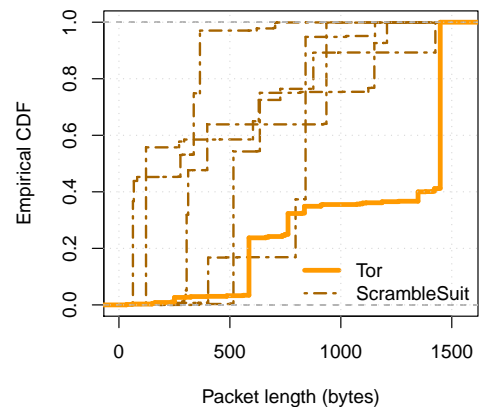


Figure 10: Server-to-client packet lengths.

The packet length distribution for client-to-server traffic is illustrated in Figure 9 and the server-to-client traffic in Figure 10. The solid orange line represents the download over Tor. The prevalence of 586-byte packets is clearly visible; especially for the client-to-server traffic. These packets contain internal Tor cells which handle flow control. All these segments had to be wrapped into a 512-byte Tor cell. In addition, more than 50% of the server-to-client traffic consists of 1448-byte packets. The remaining brown lines represent 5 consecutive downloads over ScrambleSuit. Both empirical cumulative distribution functions visibly deviate from Tor’s.

Figure 11 and 12 depict the inter arrival times of the same data. The delays in Figure 11 tend to be rather high—only 40% of Tor packets had an inter arrival delay under 10 ms—because the client only acknowledged the bulk data coming from the server. For this reason, the delays are much smaller in Figure 12. For both, the packet lengths as well as the inter arrival times, a two-sample Kolmogorov-Smirnov test rejected the hypothesis that the Tor distributions equal any of ScrambleSuit’s distributions.

5.2 Performance

In order to evaluate ScrambleSuit’s overhead and goodput, we created a 1,000,000-byte file consisting of random bytes and placed it on a web server operated by our institution. We then downloaded the file with `wget` 25 times over HTTP, Tor and ScrambleSuit, respectively. For Tor and Scramble-

Table 1: Mean (μ) and standard deviation (σ) of the goodput, transferred KBytes and the total overhead. The data was generated based on the download of a 1,000,000-byte file.

	HTTP		Tor		ScrambleSuit	
	μ	σ	μ	σ	μ	σ
Goodput	6.3 MB/s	3.4 MB/s	279.7 KB/s	293.9 KB/s	89.8 KB/s	41.1 KB/s
C→S KBytes	23.1	1.6	71.9	7.7	132.5	35.5
S→C KBytes	1047	20.7	1121.6	38.8	1242.3	70.2
Total overhead	9.5%	2.2%	22.2%	4.4%	40.7%	10.1%

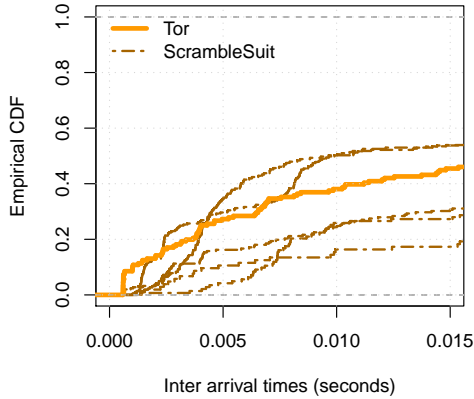


Figure 11: Client-to-server inter arrival times.

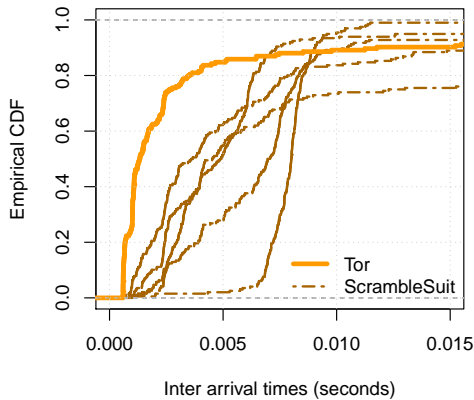


Figure 12: Server-to-client inter arrival times.

Suit, we established a new circuit for every download but we used the same entry guard in order to make the results more comparable. Based on the measured data, we calculated the mean μ and the standard deviation σ for several performance metrics. The results are depicted in Table 1.

The goodput refers to the application layer throughput. We achieved very high values for the HTTP download because the file transfer could be carried out over the LAN. Tor averaged at roughly 280 KB/s and ScrambleSuit achieved approximately one third of that. Just like Tor, ScrambleSuit exhibits a high standard deviation. We believe that this was mostly caused by differences in circuit throughput but ScrambleSuit also fluctuates due to its different shapes. Shapes featuring many large packet lengths lead to higher throughput whereas shapes with small packets tend to harm

throughput.

The next two rows of Table 1 refer to the transferred KBytes from client to server (C→S) and server to client (S→C). Note that this covers all the data which was present on the wire; including IP and TCP header. The consideration of IP and TCP overhead is important because one of the reasons for ScrambleSuit’s overhead is the varying packet lengths which imply an increase in IP and TCP headers. Unsurprisingly, Tor transferred more data than HTTP because of Tor’s and TLS’ protocol overhead. ScrambleSuit transferred the most data because of the additional protocol header (see §4.2) as well as the varying packet lengths. This is emphasised by the high standard deviation of 35 and 70 KBytes, respectively.

The last row illustrates the total protocol overhead. Once again, we also consider IP and TCP headers. HTTP has the lowest overhead followed by Tor and finally ScrambleSuit. Our protocol exhibits 40% overhead with a high standard deviation of 10% which stems from the shape shifts.

6. DISCUSSION

We made an effort to design ScrambleSuit in a way to be resistant against current censorship threats; most notably active probing. In this section we discuss the remaining attack surface and point out emerging problems.

6.1 Attacks on ScrambleSuit

Active Probing: A censor could still actively probe a ScrambleSuit server. Upon establishing a TCP connection, a censor could proceed by sending arbitrary data. However, the server will not respond without prior authentication. In contrast to SilentKnock [33] and BridgeSPA [34], ScrambleSuit does not disguise its “aliveness”. While this approach does leak information⁹, it has the benefit of making ScrambleSuit significantly easier to deploy due to lack of platform dependencies.

Packet Injection, Modification and Dropping: A censor could tamper with an existing ScrambleSuit connection by injecting data, modifying bytes on the wire or dropping packets. Both communicating parties will detect injected or modified data due to the MAC being invalid. Hijacking a ScrambleSuit TCP connection boils down to the same problem; a censor would bypass the authentication but is unable to talk to the ScrambleSuit server because the session keys are unknown. Finally, dropped packets are handled by the application protocol and could trigger TCP re-transmissions.

⁹A censor learns that a server is online but unwilling to talk unless given the “correct” data.

Payload & Flow Analysis: Payload analysis would only yield data which is computationally indistinguishable from randomness. Flow analysis, on the other hand, would yield a certain distribution of packet lengths and inter arrival times which changes from connection to connection. While §4.3.3 showed that defending against traffic analysis can be costly, our main objective is to thwart active probing attacks because they enable *deterministic protocol identification*. Sophisticated traffic analysis attacks will always have a range of *uncertainty*. We believe that the GFW’s very reason to conduct active probing is to obtain certainty and avoid collateral damage. Protocol blocking based on traffic analysis will unavoidably imply false positives.

6.2 Future Work

Improving Tor’s censorship resistance is a two-sided problem. On the one hand, bridge descriptors need to be distributed to users while not falling into the hands of censors and on the other hand, the subsequent Tor connection should be hard to identify for censors. While we focused on the latter, the former remains an open problem as well. Recent work focused on reputation-based models to maximise bridge aliveness [51, 52].

The arms race with circumvention tools might pressure censors into introducing whitelisting. While we are not aware of country-wide whitelists, Russia is experimenting with a “Clean Internet” [53]. Should this approach turn out to be successful for censors, the arms race will shift towards tunneling circumvention traffic through whitelisted protocols.

We finally point out that no protocol is “fingerprintless”. In our design, we tried to avoid obvious distinguishers and minimised the interaction surface for attackers who lack the shared secret. But since ScrambleSuit does not mimic a cover protocol, it hides within the set of *unknown rather than known* protocols. As long as a censor’s network features a high diversity of network protocols, the censor is unable to fully control or model, we expect our approach to provide a decent level of protection.

7. CONCLUSION

We presented ScrambleSuit; a transport protocol which provides lightweight obfuscation for application protocols such as Tor and VPN. The two major contributions of our protocol are the ability to defend against *active probing* and *protocol classifiers*. We achieve the former by proposing two authentication mechanisms—one general-purpose and the other specifically for Tor—and the latter by proposing morphing techniques to disguise packet lengths and inter arrival times.

We further developed a prototype of ScrambleSuit and used it to conduct an experimental evaluation. In particular, we discussed the effectiveness of our obfuscation techniques as well as ScrambleSuit’s overhead. Our evaluation suggests that ScrambleSuit can provide strong protection against censors who do not overblock significantly. Finally, we believe that the low protocol overhead makes ScrambleSuit comfortable to use for web surfing and other low-latency applications.

Acknowledgements

We want to thank George Kadianakis, Harald Lampesberger, Stefan Lindskog, and Michael Rogers who all provided valuable feedback which improved this paper. We further want to thank Internetfonden of the Swedish Internet Infrastructure Foundation for supporting the main author’s work with a research grant.

Our code is publicly available at <http://www.cs.kau.se/philwint/scramblesuit/>. Finally, we point out that all of the references listed below contain a link to an open access version of the respective resource. Please consider doing the same in your papers.

References

- [1] Thomas H. Ptacek and Timothy N. Newsham. *Insertion, Evasion, and Denial of Service: Eluding Network Intrusion Detection*. Tech. rep. Secure Networks, Inc., 1998. URL: <http://cs.unc.edu/~fabian/course.papers/PtacekNewsham98.pdf>.
- [2] Olli-Pekka Niemi, Antti Levomäki, and Jukka Manner. “Dismantling Intrusion Prevention Systems (Demo)”. In: *SIGCOMM*. ACM, 2012. URL: <http://conferences.sigcomm.org/sigcomm/2012/paper/sigcomm/p285.pdf>.
- [3] Mark Handley, Vern Paxson, and Christian Kreibich. “Network Intrusion Detection: Evasion, Traffic Normalization, and End-to-End Protocol Semantics”. In: *USENIX Security*. USENIX Association, 2001. URL: http://static.usenix.org/events/sec01/full_papers/handley/handley.pdf.
- [4] Marcel Dischinger, Alan Mislove, Andreas Haeberlen, and Krishna P. Gummadi. “Detecting BitTorrent Blocking”. In: *IMC*. ACM, 2008. URL: http://www.mpi-sws.org/~mdischin/papers/08_imc_blocking.pdf.
- [5] Richard Clayton, Steven J. Murdoch, and Robert N. M. Watson. “Ignoring the Great Firewall of China”. In: *PETS*. Springer, 2006. URL: <http://www.cl.cam.ac.uk/~rnc1/ignoring.pdf>.
- [6] Sparks et al. “The Collateral Damage of Internet Censorship by DNS Injection”. In: *SIGCOMM Computer Communication Review* 42.3 (2012). URL: <http://conferences.sigcomm.org/sigcomm/2012/paper/ccr-paper266.pdf>.
- [7] Christopher Rhoads and Loretta Chao. *Iran’s Web Spying Aided By Western Technology*. 2009. URL: <http://online.wsj.com/article/SB124562668777335653.html>.
- [8] Jillian C. York. *Government Internet Surveillance Starts With Eyes Built in the West*. 2011. URL: <https://www.eff.org/deeplinks/2011/09/government-internet-surveillance-starts-eyes-built>.
- [9] Andrew M. White et al. “Clear and Present Data: Opaque Traffic and its Security Implications for the Future”. In: *NDSS*. The Internet Society, 2013. URL: <http://cs.unc.edu/~amw/resources/opaque.pdf>.

- [10] Roger Dingledine, Nick Mathewson, and Paul Syverson. "Tor: The Second-Generation Onion Router". In: *USENIX Security*. USENIX Association, 2004. URL: http://static.usenix.org/event/sec04/tech/full_papers/dingledine/dingledine.pdf.
- [11] The Tor Project. *Iran*. URL: <https://censorshipwiki.torproject.org/CensorshipByCountry/Iran>.
- [12] Philipp Winter and Stefan Lindskog. "How the Great Firewall of China is Blocking Tor". In: *FOCI*. USENIX Association, 2012. URL: <https://www.usenix.org/system/files/conference/foci12/foci12-final2.pdf>.
- [13] The Tor Project. *Ethiopia*. URL: <https://censorshipwiki.torproject.org/CensorshipByCountry/Ethiopia>.
- [14] Charles Arthur. *China tightens 'Great Firewall' internet control with new technology*. 2012. URL: <http://www.guardian.co.uk/technology/2012/dec/14/china-tightens-great-firewall-internet-control>.
- [15] *GFW actively probes obfs2 bridges*. 2013. URL: <https://trac.torproject.org/projects/tor/ticket/8591>.
- [16] Zachary Weinberg et al. "StegoTorus: A Camouflage Proxy for the Tor Anonymity System". In: *CCS*. ACM, 2012. URL: <http://www.owlfolio.org/media/2010/05/stegotorus.pdf>.
- [17] Hooman Mohajeri Moghaddam, Baiyu Li, Mohammad Derakhshani, and Ian Goldberg. "SkypeMorph: Protocol Obfuscation for Tor Bridges". In: *CCS*. ACM, 2012. URL: <http://www.cypherpunks.ca/~iang/pubs/skypemorph-ccs.pdf>.
- [18] The Tor Project. *obfs2 (The Twobfuscator)*. URL: <https://gitweb.torproject.org/pluggable-transport/obfsproxy.git/blob/HEAD:/doc/obfs2/obfs2-protocol-spec.txt>.
- [19] The Tor Project. *obfs3 (The Threebfuscator)*. URL: <https://gitweb.torproject.org/pluggable-transport/obfsproxy.git/blob/HEAD:/doc/obfs3/obfs3-protocol-spec.txt>.
- [20] Brandon Wiley. *Dust: A Blocking-Resistant Internet Transport Protocol*. Tech. rep. University of Texas at Austin, 2011. URL: <http://blanu.net/Dust.pdf>.
- [21] Amir Houmansadr, Chad Brubaker, and Vitaly Shmatikov. "The Parrot is Dead: Observing Unobservable Network Communications". In: *Security & Privacy*. IEEE, 2013. URL: <http://www.cs.utexas.edu/~amir/papers/parrot.pdf>.
- [22] Alberto Dainotti et al. "Analysis of Country-wide Internet Outages Caused by Censorship". In: *IMC*. ACM, 2011. URL: http://www.caida.org/publications/papers/2011/outages_censorship/outages_censorship.pdf.
- [23] Eva Galperin and Jillian C. York. *Syria Goes Dark*. 2012. URL: <https://www.eff.org/deeplinks/2012/11/syria-goes-dark>.
- [24] Maurizio Dusi, Manuel Crotti, Francesco Gringoli, and Luca Salgarelli. "Tunnel Hunter: Detecting Application-Layer Tunnels with Statistical Fingerprinting". In: *Computer Networks* 53.1 (2009). URL: <http://www.ing.unibs.it/~salga/pub/2009-tunnel.pdf>.
- [25] Dario Bonfiglio et al. "Revealing Skype Traffic: When Randomness Plays with You". In: *SIGCOMM*. ACM, 2007. URL: http://www.telematica.polito.it/mellia/papers/Skype_Sigcomm2007.pdf.
- [26] Roni Bar-Yanai, Michael Langberg, David Peleg, and Liam Roditty. "Realtime Classification for Encrypted Traffic". In: *SEA*. Springer, 2010. URL: <http://www.openu.ac.il/home/mikel/papers/60490373%5B1%5D.pdf>.
- [27] Erik Hjelmvik and Wolfgang John. *Breaking and Improving Protocol Obfuscation*. Tech. rep. Chalmers University of Technology, 2010. URL: http://www.iis.se/docs/hjelmvik_breaking.pdf.
- [28] Brandon Wiley. *Blocking-Resistant Protocol Classification Using Bayesian Model Selection*. Tech. rep. University of Texas at Austin, 2011. URL: <http://blanu.net/BayesianClassification.pdf>.
- [29] The Tor Project. *obfsproxy*. URL: <https://www.torproject.org/projects/obfsproxy>.
- [30] Kevin P. Dyer, Scott E. Coull, Thomas Ristenpart, and Thomas Shrimpton. *Format-Transforming Encryption: More than Meets the DPI*. Tech. rep. Portland State University, 2012. URL: <http://eprint.iacr.org/2012/494.pdf>.
- [31] Patrick Lincoln et al. "Bootstrapping Communications into an Anti-Censorship System". In: *FOCI*. USENIX Association, 2012. URL: <https://www.usenix.org/system/files/conference/foci12/foci12-final7.pdf>.
- [32] David Fifield et al. "Evading Censorship with Browser-Based Proxies". In: *PETS*. Springer, 2012. URL: <http://freehaven.net/~arma/flashproxy.pdf>.
- [33] Eugene Y. Vasserman, Nicholas Hopper, John Laxson, and James Tyra. "SilentKnock: Practical, Provably Undetectable Authentication". In: *ESORICS*. Springer, 2007. URL: http://www-users.cs.umn.edu/~hopper/silentknock_esorics.pdf.
- [34] Rob Smits et al. "BridgeSPA: Improving Tor Bridges with Single Packet Authorization". In: *WPES*. ACM, 2011. URL: <http://www.cypherpunks.ca/~iang/pubs/bridgespa-wpes.pdf>.
- [35] Erik Hjelmvik and Wolfgang John. "Statistical Protocol IDentification with SPID: Preliminary Results". In: *SNCNW*. 2009. URL: http://www.cse.chalmers.se/~johnwolf/publications/sncnw09-hjelmvik_john-CR.pdf.
- [36] George Kadianakis. *Packet Size Pluggable Transport and Traffic Morphing*. Tech. rep. The Tor Project, 2012. URL: <https://research.torproject.org/techreports/morpher-2012-03-13.pdf>.
- [37] Martin Johnson. *China, GitHub and the man-in-the-middle*. 2013. URL: <https://en.greatfire.org/blog/2013/jan/china-github-and-man-middle>.

- [38] Ronald L. Rivest, Adi Shamir, and David A. Wagner. *Time-lock Puzzles and Timed-release Crypto*. Tech. rep. Massachusetts Institute of Technology, 1996. URL: <http://people.csail.mit.edu/rivest/RivestShamirWagner-timelock.ps>.
- [39] Ben Laurie and Richard Clayton. ““Proof-of-Work” Proves Not to Work”. In: *WEIS*. 2004. URL: <http://www.cl.cam.ac.uk/~rnc1/proofwork2.pdf>.
- [40] Joseph Salowey, Hao Zhou, Pasi Eronen, and Hannes Tschofenig. *RFC 5077: Transport Layer Security (TLS) Session Resumption without Server-Side State*. 2008. URL: <https://tools.ietf.org/html/rfc5077>.
- [41] Zhen Ling et al. “Extensive Analysis and Large-Scale Empirical Evaluation of Tor Bridge Discovery”. In: *INFOCOM*. IEEE, 2012. URL: <http://www.cs.uml.edu/~xinwenfu/paper/Bridge.pdf>.
- [42] Tero Kivinen and Mika Kojo. *RFC 3526: More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange (IKE)*. 2003. URL: <http://tools.ietf.org/html/rfc3526>.
- [43] Hugo Krawczyk and Pasi Eronen. *RFC 5869: HMAC-based Extract-and-Expand Key Derivation Function (HKDF)*. 2010. URL: <https://tools.ietf.org/html/rfc5869>.
- [44] Manuel Crotti, Maurizio Dusi, Francesco Gringoli, and Luca Salgarelli. “Traffic Classification through Simple Statistical Fingerprinting”. In: *SIGCOMM Computer Communication Review* 37.1 (2007). URL: <http://www.sigcomm.org/sites/default/files/ccr/papers/2007/January/1198255-1198257.pdf>.
- [45] Kevin P. Dyer, Scott E. Coull, Thomas Ristenpart, and Thomas Shrimpton. “Peek-a-Boo, I Still See You: Why Efficient Traffic Analysis Countermeasures Fail”. In: *Security & Privacy*. IEEE, 2012. URL: <http://kpdyer.com/publications/oakland2012-peekaboo.pdf>.
- [46] Xiang Cai, Xin Cheng Zhang, Brijesh Joshi, and Rob Johnson. “Touching from a Distance: Website Fingerprinting Attacks and Defenses”. In: *CCS*. ACM, 2012. URL: <http://www.cs.sunysb.edu/~xcai/fp.pdf>.
- [47] Yeon-sup Lim et al. “Internet Traffic Classification Demystified: On the Sources of the Discriminative Power”. In: *CoNEXT*. ACM, 2010. URL: http://conferences.sigcomm.org/co-next/2010/CoNEXT_papers/09-Lim.pdf.
- [48] Andriy Panchenko, Lukas Niessen, Andreas Zinnen, and Thomas Engel. “Website Fingerprinting in Onion Routing Based Anonymization Networks”. In: *WPES*. ACM, 2011. URL: <http://lorre.uni.lu/~andriy/papers/acmccs-wpes11-fingerprinting.pdf>.
- [49] Charles V. Wright, Scott E. Coull, and Fabian Monrose. “Traffic Morphing: An Efficient Defense Against Statistical Traffic Analysis”. In: *NDSS*. The Internet Society, 2009. URL: <http://mirror.robert-marquardt.com/anonbib/cache/morphing09.pdf>.
- [50] Mohamad Jaber, Roberto G. Cascella, and Chadi Barakat. “Can we trust the inter-packet time for traffic classification?”. In: *ICC*. IEEE, 2011. URL: <http://www-sop.inria.fr/members/Chadi.Barakat/ICC2011.pdf>.
- [51] Damon McCoy, Jose Andre Morales, and Kirill Levchenko. “Proximax: A Measurement Based System for Proxies Dissemination”. In: *Financial Cryptography and Data Security*. Springer, 2011. URL: <http://cseweb.ucsd.edu/~klevchen/mml-fc11.pdf>.
- [52] Qiyang Wang, Zi Lin, Nikita Borisov, and Nicholas J. Hopper. “rBridge: User Reputation based Tor Bridge Distribution with Privacy Preservation”. In: *NDSS*. The Internet Society, 2013. URL: http://www-users.cs.umn.edu/~hopper/rbridge_ndss13.pdf.
- [53] Russian “Clean Internet” experiment gets green light. 2013. URL: <http://rt.com/politics/anti-pedophile-safe-internet-russian-169/>.