

Snowflake, a censorship circumvention system using temporary WebRTC proxies

(Draft February 7, 2024)

Cecylia Bocovich

Arlo Breault

David Fifield

Serene

Xiaokang Wang

Authors are listed alphabetically.

Abstract

Snowflake is a system for circumventing Internet censorship. Its blocking resistance comes from the use of numerous, ultra-light, temporary proxies (“snowflakes”), which accept traffic from censored clients using peer-to-peer WebRTC protocols and forward it to a centralized bridge. The temporary proxies are simple enough to be implemented in JavaScript, in a web page or browser extension, making them vastly cheaper to run than a traditional proxy or VPN server. The large and constantly changing pool of proxy addresses resists enumeration and blocking by a censor. The system is built on the assumption that proxies may appear or disappear at any time: clients discover live proxies dynamically using a secure rendezvous protocol; when an in-use proxy goes offline, its client switches to another on the fly, invisibly to upper network layers.

Snowflake has been deployed with success in Tor Browser and Orbot for several years. It has been a significant circumvention tool during high-profile network disruptions, including in Russia in 2021 and Iran in 2022. In this paper, we explain the composition of Snowflake’s many parts, give a history of deployment and blocking attempts, and reflect on implications for circumvention generally.

1 Introduction

Censorship circumvention systems—systems to enable network communication despite interference by a censor—may be characterized on multiple axes. Some systems imitate a common network protocol; others try not to look like any protocol in particular. Some distribute connections over numerous proxy servers; others concentrate on a single proxy that is, for one reason or another, difficult for a censor to block. What all circumvention systems have in common is that they strive to increase the *cost* to the censor of blocking them—whether that cost be in research and development, human resources,

and hardware; or in the inevitable overblocking that results when a censor tries to selectively block some connections but not others. Snowflake, the subject of this paper, is a circumvention system that uses thousands of temporary proxies and makes switching between them easy and fast. On the spectrum of imitation to randomization, Snowflake falls on the side of imitation; on the scale of diffuse to concentrated, it is diffuse. What characterizes Snowflake the most is that it pushes the idea of distributed, disposable proxies to an extreme: its proxies can run in a web browser, and censored clients communicate with them using WebRTC.

WebRTC is a suite of protocols intended for real-time communication applications on the web [1]. Video and voice chat are typical examples of WebRTC applications. Snowflake exchanges WebRTC data formats in the course of establishing a connection, and uses WebRTC protocols for traversal of NAT (network address translation) and communication between clients and proxies. Crucially for Snowflake, WebRTC APIs are available to JavaScript code in web browsers, meaning it is possible to implement a proxy in a web page or browser extension. WebRTC is also usable outside a browser, which is how we implement the Snowflake client program and alternative, command line-based proxies.

As is usual in circumvention research, we assume a threat model in which *clients* reside in a network controlled by a *censor*. The censor has the power to inspect and interfere with traffic that crosses the border of its network; typical real-world censor behaviors include inspecting IP addresses and hostnames, checking packet contents for keywords, blocking IP addresses, and injecting false DNS responses or TCP RST packets. The client wants to communicate with some *destination* outside the censor’s network, possibly with the aid of third-party *proxies*. The censor is motivated to block the contents of the client’s communication, or even the destination itself. The censor is aware of the possibility of circumvention, and therefore seeks to block not only direct communication, but also indirect communication by way of a proxy or circumven-

tion system. Circumvention is accomplished when the client can reliably reach any proxy, because a proxy, being outside the censor’s control, can then forward the client’s communication to any destination. (In Snowflake, we separate the roles of temporary *proxies* and a stable long-term *bridge*, but the idea is the same.) The censor is presumed to derive benefit from permitting some forms of network access: the censor cannot trivially “win” by shutting down all communication, but must be selective in its blocking decisions, in order to optimize some objective of its own. The art of censorship circumvention is forcing the censor into a dilemma of overblocking or underblocking, by making circumvention traffic difficult to distinguish from traffic that the censor prefers not to block.

Snowflake originates in two earlier projects: flash proxy and uProxy. Flash proxy [10], like Snowflake, used a model of untrusted, temporary JavaScript proxies in web browsers, but the link between client and proxy used WebSocket rather than WebRTC. (WebSocket still finds use in Snowflake, but on the proxy–bridge link, not the client–proxy link.) Flash proxy was deployed in Tor Browser from 2013 to 2016, but never saw much use, probably because the reliance on WebSocket, which lacks the built-in NAT traversal of WebRTC, required client users to do their own port forwarding. WebRTC was then an emerging technology, and while it had been considered as a transport protocol for flash proxy, we decided to start Snowflake as an independent project. uProxy [38], in one of its early incarnations, pioneered the use of WebRTC proxies for circumvention. uProxy’s proxies were browser-based, but its trust and deployment models were different from flash proxy’s and Snowflake’s. Each censored client would arrange, out of band, for an acquaintance outside the censor’s network to run a proxy in their browser [39]. A personal trust relationship was necessary to prevent misuse, since browser proxies fetched destination content directly—meaning the client’s activity would be attributed to the proxy, and the proxy could inspect the client’s traffic. Clients did not change proxies on the fly. uProxy supported protocol obfuscation: the communications protocol was fundamentally WebRTC, but packets could be transformed to resemble something else. This obfuscation was possible because of uProxy’s implementation as a privileged browser extension, with access to real sockets. Because Snowflake uses ordinary unprivileged browser APIs, its WebRTC can only look like WebRTC; on the other hand, for the same reason, Snowflake proxies are easier to deploy. Like flash proxy, uProxy was active in the years 2013–2016.

Among existing circumvention systems, the one that is most similar to Snowflake is MassBrowser [25], which offers proxying through volunteer proxies, called buddies. MassBrowser’s architecture is similar to Snowflake’s: there is a centralized component that coordinates connections between clients and buddies, corresponding to a piece in Snowflake called the broker; buddies play the same role as our proxies. The trust model is intermediate between Snowflake’s and uProxy’s. Buddies preferentially operate as one-hop proxies, as in uProxy, but

are not limited to proxying only for trusted friends. To deter misuse, buddies specify a policy of what categories of content they are willing to proxy. An innovation in MassBrowser not present in Snowflake is client-to-client proxying: clients may act as buddies for other clients, the logic being that what is censored for one client may not be censored for another. The buddy software is not constrained by a web browser environment, and can, like uProxy, use protocol obfuscation on the client–buddy link.

Protozoa [2] and Stegozoa [12] show ways of building a point-to-point covert tunnel over WebRTC, the former by directly replacing encoded media with its own ciphertexts, the latter using video steganography. Designs like these might serve as alternatives for the link between client and proxy in Snowflake. Significantly, where Snowflake now uses WebRTC data channels, Protozoa and Stegozoa use WebRTC media streams, which may be an advantage in blocking resistance. We will say more on this point in Section 3. TorKameleon [40] is a WebRTC-based transport with the dual goals of resisting blocking (circumvention) and complicating traffic correlation attacks (anonymity). It has the notable technical innovation of using a draft API called WebRTC Encoded Transforms to support efficient Protozoa-like embedding of data within media streams, without requiring invasive modifications in a browser.

Our goal in this paper is not to exaggerate the advantages of Snowflake, nor disproportionately emphasize the limitations of other circumvention systems. Circumvention research is a cooperative enterprise, and we recognize and support our colleagues who pursue and maintain their own designs. While challenges remain, today’s circumvention systems by and large accomplish their intended purpose, and are a vital element of day-to-day Internet access for many people. With Snowflake, we have explored a different point in the design space—a fruitful one to be sure—but one with its own tradeoffs. We acknowledge that Snowflake will be a better choice in some censorship environments and worse in others; indeed, one of the ideas we hope to convey is that blocking resistance can be meaningfully understood only in relation to particular censor and its resources, costs, and motivations.

In this paper we present the design of Snowflake, discuss various challenges and considerations, and reflect on over three years of deployment. As of February 2024, Snowflake supports an estimated 42,000 average concurrent users at an average total transfer rate of 3.5 Gbit/s, which works out to around 38 TB of circumvention traffic per day.

2 How it works

A Snowflake proxy connection proceeds in three phases. First, there is rendezvous, in which a client indicates its need for circumvention service and is matched with a temporary proxy. Rendezvous is facilitated by a central server called the broker. Then, there is connection establishment, where the client and its




Figure 1: Architecture of Snowflake. The client contacts the broker through a special rendezvous channel with high blocking resistance. The broker matches the client with one of the proxies that are currently polling. The client and proxy connect to one another using WebRTC. The proxy connects to the bridge, then begins copying traffic in both directions. If the proxy disappears, the client does another rendezvous and resumes its session with a new proxy.

Make this graphic depict STUN servers and indirect rendezvous.

proxy connect to each other with WebRTC, using information exchanged during rendezvous. Finally, there is data transfer, where the proxy transports data between the client and the bridge. The bridge is responsible for directing the client’s traffic to its eventual destination (in our case, by feeding it into the Tor network). Figure 1 illustrates the process.

These phases repeat as needed, as temporary proxies come and go. Proxy failure is not an abnormal condition—it happens whenever a proxy is running in a browser that is closed, for example. A client builds a circumvention session over a sequence of proxies, switching to a new one whenever the current one stops working. State variables stored at the client and the bridge let the session pick up where it left off. The change of proxies is invisible to the applications using Snowflake (except for a brief delay while rendezvous happens): the Snowflake client presents an abstraction of a single, uninterrupted connection.

It does not avail a censor to block the broker or bridge, because Snowflake clients never contact either one directly. Clients reach the broker over an indirect rendezvous channel. Access to the bridge is always mediated by a temporary proxy.

2.1 Rendezvous

A session begins with a client sending a rendezvous message to the broker. There is an ambient population of proxies constantly polling the broker to check for clients in need of service. The broker matches the client with an available proxy, taking into consideration factors like NAT compatibility.

The client’s rendezvous message is a bundle of data that the broker will need to match the client with a proxy, and the proxy will need to connect to the client. The primary element is a Session Description Protocol (SDP) *offer* [28], which contains the information necessary for a WebRTC connection, including the client’s external IP addresses and cryptographic data to

secure a later key exchange. The broker forwards the client’s SDP offer to the proxy, and the proxy sends back an SDP *answer* with its share of connection details. The broker forwards the proxy’s SDP answer to the client. The client and proxy then connect to each other directly. In WebRTC terms, this offer/answer exchange is called “signaling,” and here the broker acts as a signaling server. To gather the information for an SDP offer or answer, clients and proxies communicate with third-party servers, called STUN servers, before contacting the broker. We will say more about how this information is used in Section 2.2. Communication with STUN servers is a normal and expected part of WebRTC, though there are fingerprinting considerations that we discuss in Section 3.

Interaction with the broker uses a “long-polling” model. An example is shown in Figure 2. Proxies poll the broker periodically, making an HTTPS request to a designated URL path. The broker does not respond immediately to a proxy poll, but instead holds the connection idle for a few seconds to await the possible arrival of a client rendezvous message. If none arrives, the broker sends a response saying “no clients” and the proxy goes to sleep until its next poll. When a client does arrive, the broker sends the SDP offer in response to the proxy’s poll request. The proxy sends its SDP answer to the broker in a separate HTTPS request. The broker responds to the client’s pending request with the proxy’s SDP answer, at the same time sending an acknowledgement to the proxy. At this point rendezvous is finished, and the client and the proxy may connect to one another.

The client must use an indirect, blocking-resistant channel when communicating with the broker. What is needed, essentially, is a miniature circumvention system to bootstrap the full system. What makes rendezvous different from general circumvention are its different (generally more lenient) requirements, which permit a larger solution space. Because rendezvous




Figure 2: The long-polling communication model of Snowflake rendezvous. Proxies poll periodically to check for new clients. When the broker makes a match, the proxy gets the client’s SDP offer, then immediately re-connects to send back its SDP answer. It all happens during one round trip from the client’s perspective. Not shown here is the indirect channel used by the client to access the broker through the censor’s zone of control (shaded background).

is only a small fraction of total communication volume, and it happens relatively infrequently, it may use techniques that would be too slow, expensive, or complicated for real-time or bulk data transfer. Rendezvous is separable and modular: more than one method can be used, and the methods do not necessarily need to bear any relation to the circumvention techniques of the main system. While the assumption of WebRTC permeates Snowflake’s design, its rendezvous modules are independent. We currently support two rendezvous methods in Snowflake:

Domain fronting In this method, the client does an HTTPS exchange with the broker through an intermediary web service such as a content delivery network (CDN), setting the externally visible hostname (the TLS Server Name Indication, or SNI [6 §3]) to a “front domain” different from the broker’s. The CDN routes the HTTPS request to the broker not according to the TLS SNI but rather the HTTP Host header, which, under TLS encryption, reflects the broker’s true hostname [11]. A censor cannot easily block domain-fronted rendezvous without also blocking unrelated connections to the front domain, which should be selected to have high value to the censor. (But see Section 3 for features other than the hostname that a censor might try to use.) The well-known drawback of domain fronting is the high cost of CDN bandwidth. Because we use it only for rendezvous, the cost is much less than if we were to use it for all data transfer.

AMP cache AMP is a framework for web pages written in a restricted dialect of HTML. Part of the framework is a free-to-use cache server [27]. The cache fetches AMP-conformant web pages on demand, which means that it is,

effectively, a restricted sort of HTTP proxy. We have a module that encodes rendezvous messages to AMP specifications, allowing them to be exchanged with the broker via the AMP cache. Rendezvous through the AMP cache is not easily blocked without blocking the cache server as a whole.¹ This rendezvous method still technically requires domain fronting, because the AMP cache protocol would otherwise expose the broker’s hostname in the TLS SNI, but it increases the number of usable intermediaries and front domains.

Amazon SQS Amazon’s Simple Queue Service (SQS) is a message queuing service designed for communication between microservices. Services may create queues, send messages with up to 256KB payloads, and retrieve messages from the queues. For the Snowflake rendezvous, we create a persistent, public broker queue that any client may send to. The broker processes retrieved messages and responds to the client by creating a new single-use queue with the client’s unique ID in the queue name.² Blocking SQS rendezvous requires, at the very least, blocking access to Amazon’s SQS service by region.

Anything that can be persuaded to convey a message of about 1500 bytes indirectly to the broker, and return a response of about the same size, can work as a rendezvous module. For example, encrypted DNS or a chat bot would serve. Though some systems (flash proxy was one) may need only a single, outgoing rendezvous message, Snowflake needs a two-way exchange, to support the SDP offer and answer.

Rendezvous is not unique to Snowflake. Other examples of rendezvous in circumvention include the DEFIANCE Rendezvous Protocol [20 §3], the facilitator interaction in flash proxy [10 §3], and the registration proxy in Conjure [13 §4.1]. A key property of Snowflake and the mentioned systems is that they do not rely on preshared secret information. The client needs only to acquire the necessary software; whatever additional information is required to establish a circumvention session is exchanged dynamically, at runtime. This stands in contrast to another class of systems in which, prior to making a connection, a client must acquire some secret, such as an IP address or password, through an out-of-band channel presumed to be unavailable to the censor—and the system’s blocking resistance depends on keeping that information hidden from the censor. A corollary of the no-secret-information property is that an adversary—the censor—is at no special disadvantage in attacking the system. The censor may download the client software, run it, study its network connections—and the system must maintain its blocking resistance despite this. The disadvantage of a separate rendezvous step is that it is one more thing to get right. Not only the main circumvention channel

¹https://gitlab.torproject.org/tpo/anti-censorship/pluggable-transport/snowflake/-/merge_requests/50

²https://gitlab.torproject.org/tpo/anti-censorship/pluggable-transport/snowflake/-/merge_requests/214

but also the rendezvous must resist blocking: the system is only as strong as the weaker of the two.

2.2 Peer-to-peer connection establishment

Now the client and the proxy connect to each other directly. Even in the absence of censorship, making a direct connection between two Internet peers is not always easy, because of NAT (network address translation) and firewalls. Snowflake clients and proxies alike run in diverse networks with varying NATs and ingress policies. Fortunately for us, WebRTC is designed with this use case in mind, and has built-in support for traversing NAT, in the form of ICE (Interactive Connectivity Establishment) [19], a procedure for testing candidate pairs of peer network addresses to find one that works. ICE makes use of third-party STUN (Session Traversal Utilities for NAT) [29] servers that, among other things, enable a host to learn its external IP addresses. The first part of ICE took place at the beginning of rendezvous, when the client and proxy contacted STUN servers to gather external address candidates and included them in their respective SDP offer and answer.

There is no guarantee that two hosts will be able to make a connection using the facilities of STUN alone. Some address mapping and filtering setups are simply incompatible. In such a case, ICE would normally fall back to using TURN (Traversal Using Relays around NAT) [31], a kind of UDP proxy. Such a fallback would be problematic for Snowflake, because the TURN relays themselves would become a target of blocking by the censor. But Snowflake has an advantage most WebRTC applications do not. Most WebRTC applications want to connect *a particular* pair of peers, whereas we are satisfied when a client can connect to *any* proxy. Snowflake clients and proxies self-measure their NAT type and report it to the broker, which takes NAT compatibility into account and avoids cases that would require a fallback to TURN.

We condense the possible combinations of NAT and firewall features that impact a Snowflake client or proxy’s ability to make a peer-to-peer connection into the following well-known variations:

Full cone The same internal IP–port pair always maps to the same external port. Any remote host may send a packet to an internal IP address and port by sending a packet to the mapped external port.

Restricted cone Like full cone, but incoming packets are allowed only if there has recently been an outgoing packet to the same remote IP address.

Port-restricted cone Like restricted cone, but incoming packets are allowed only if there has recently been an outgoing packet to the same remote IP–port pair.

Symmetric The external port depends on both the internal IP–port pair and the remote IP–port pair. Incoming packets

	No NAT	Full cone	Restricted cone	Port-restricted cone	Symmetric
No NAT	✓	✓	✓	✓	✓
Full cone	✓	✓	✓	✓	✓
Restricted cone	✓	✓	✓	✓	✓
Port-restricted cone	✓	✓	✓	✓	–
Symmetric	✓	✓	✓	–	–
unrestricted client			restricted client		

Table 1: Pairwise compatibility of NAT variants, using the facilities of STUN alone (no fallback to TURN). The incompatible cases are when one peer’s NAT is symmetric and the other’s is symmetric or port-restricted cone. Note the asymmetry in what NAT variants we consider “restricted” in client and proxy.

are allowed only if there has recently been an outgoing packet to the same remote address.

Table 1 shows the pairwise compatibility of NAT variations. As the incompatible cases always involve a symmetric NAT, we further simplify matching by categorizing the variations into the two types *unrestricted* (works with most other NATs) and *restricted* (works only with more permissive NATs). Unrestricted proxies may be matched with any client; restricted proxies may be matched only with unrestricted clients. The broker prefers to match unrestricted clients with restricted proxies, in order to conserve unrestricted proxies for the clients that need them. Symmetric NAT is always considered restricted, but port-restricted cone NAT differs depending on the peer: for proxies it is restricted, but for clients it is unrestricted. The asymmetric categorization is an approximation to help conserve unrestricted proxies for clients with symmetric NATs. Though it creates the potential for an incompatible match, we believe this to be uncommon in practice. In case of a connection failure, clients re-rendezvous and try again.

To self-assess their NAT type, clients use the NAT behavior discovery feature of STUN [22]. Proxies cannot use the same technique, because the necessary STUN features are not exposed to JavaScript. Instead, we adapt a technique from Mass-Browser [25 §V-A]: we run a centralized, always-on WebRTC testing peer behind a simulated symmetric NAT.³ Proxies try connecting to this peer: if the connection succeeds, the proxy’s type is unrestricted; otherwise it is restricted. Clients and proxies retest their NAT type periodically, to account for potential changes in their local networking environment. If a client or proxy is unable to determine its NAT type for some reason, it reports the type “unknown,” which the broker conservatively

³<https://bugs.torproject.org/tpo/anti-censorship/pluggable-transports/snowflake/40013>




Figure 3: Proxy NAT types, in unique IP addresses per day. The places in 2021 and 2022 where there is an increase in the “unknown” NAT type and a decrease in the other types were the result of operational problems with NAT type testing.

treats as if it were restricted.


Figure 3 shows that unrestricted proxies form a relatively small fraction of the proxy population. In absolute terms, there are enough, thanks in large part to the volunteers who run the command-line version of the Snowflake proxy on networks unencumbered by NAT. Though stable, long-term proxies go somewhat against the ethos of Snowflake, it has proved useful, as a matter of practicality, to sacrifice a measure of address diversity for better NAT compatibility in a common case. We can estimate how many tries it takes a client to be matched with a proxy, on average, by counting failed and successful rendezvous attempts at the broker, under the assumption that clients repeat rendezvous attempts until getting a match. In July 2023, unrestricted clients almost always got a match on the first attempt, while restricted clients needed an average of 1.07 attempts (standard deviation 0.05).

While the proxy is connecting to its client, it also connects to the bridge. This connection uses WebSocket [24], which offers a TCP-like, client–server connection layered on HTTPS. The choice of protocol for the proxy–bridge link is arbitrary, and could be changed without affecting the rest of the system. It does not need to be resistant to blocking, it just needs to be available to JavaScript code in web browsers. WebRTC, for example, would work for this link too.

2.3 Data transfer

No complicated processing takes place at the proxy. The main value of a Snowflake proxy is its IP address: it gives the client a peer to connect to that is not on the censor’s address blocklist. Having provided that, the proxy assumes a role of pure data transfer.

Snowflake uses a stack of nested protocol layers. We will walk through the layers and describe the purpose of each.



This is the stack for the client–proxy link, which is the place where WebRTC is used, and which is exposed to observation by the censor (Figure 1). The stack for the proxy–bridge link is the same, but with WebSocket in place of the WebRTC data channel at the top. The layers marked “ephemeral” are skinned off and replaced as proxies come and go. The layers marked “persistent” are instantiated once in each circumvention session, hold long-term state, and are end-to-end between client and bridge.

The connection between a client and its proxy is a WebRTC data channel [18], which provides a way to send arbitrary binary messages between peers. A data channel is its own stack of three protocols: UDP for network transport, DTLS (Datagram TLS) for confidentiality and integrity, and SCTP (Stream Control Transmission Protocol) for delimiting message boundaries and other features like congestion control. Working UDP port numbers will have been discovered using ICE in the previous phase. The peers authenticate one another at the DTLS layer using certificate fingerprints that were exchanged during rendezvous [17 §5.1].

Data channels are well-suited to Snowflake’s needs. (The specification even lists circumvention as a use case [18 §3.2].) But data channels are not the only option: WebRTC also offers *media streams* for unreliable transport of real-time audio and video. Which of these is used may be a fingerprinting vector. We will take up this topic in Section 3.

If clients only ever used one proxy, a WebRTC data channel alone would be sufficient. But a Snowflake proxy might disappear at any moment, and when that happens, its data channel goes with it. If the client was in the middle of a long download, for example, it should be possible to resume the download without interruption after rendezvousing with a new proxy. For this we need a shared notion of session state that exists at the client and the bridge, not tied to any temporary proxy. A lack of session continuity across proxy failures had been an unsolved problem in flash proxy [10 §5.2].

We adopt the Turbo Tunnel design pattern [8] and insert a userspace session and reliability protocol between the ephemeral proxy data channels and the client’s own application streams.⁴ This part of the protocol stack outlives any single proxy; it belongs to the client and the bridge. Its primary function is to attach sequence numbers and acknowledgements to packets of data, so that both ends know what parts of the data stream need to be retransmitted after a temporary loss of proxy connectivity. The client tags its traffic with a random session

⁴<https://lists.torproject.org/pipermail/anti-censorship-team/2020-February/000059.html>

identifier string that remains consistent throughout a session, which the bridge uses to index a map of session variables. For the inner session layer we use a combination of KCP [34] and smux [43]. KCP provides reliability, and smux detects the end of idle sessions and terminates them. KCP and smux have shown their worth in other deployments, and are easy to program, but there is nothing about them on which we depend essentially. Any other transport protocol that provides the necessary features and can be implemented in userspace would do, such as QUIC, TCP, or (another layer of) SCTP. We prototyped successfully with QUIC before deciding on KCP/smux.

Snowflake can be seen as an instance of the “untrusted messengers” model of Feamster et al. [7 §3]: our *proxies* and *bridge* correspond to their *messengers* and *portal*. Proxies are tasked with delivering the client’s data to the bridge, but are not permitted to tamper with or inspect it, which necessitates an inner, end-to-end secure protocol between the client and the bridge. In our deployment, this is Tor protocol. After removing the WebSocket and Turbo Tunnel layers, the Snowflake bridge feeds the client’s Tor streams into a Tor bridge running on the same host. The use of Tor is an implementation choice, not a requirement—many other protocols would work in its place. Tor has the nice quality that not even the bridge sees the plaintext of client streams. But Tor also has certain drawbacks, which we will comment on in Section 4.4 and Section 6.

3 Protocol fingerprinting

Snowflake leans heavily into the “address blocking” side of circumvention, but the “content blocking” part matters too. The goal, as always, is to make circumvention traffic difficult to distinguish from traffic the censor prefers not to block. Snowflake is tied to WebRTC, and can only be effective against a censor that is not willing to block WebRTC protocols wholesale. But even within that scope, there are many variations in how WebRTC is implemented and used, which, if not carefully considered, might enable a censor to selectively block only Snowflake, while leaving other uses of WebRTC undisturbed. Unfortunately for the circumvention developer, the richness of WebRTC protocols creates a large attack surface for fingerprinting. Not only that, WebRTC leaves the details of signaling—the process by which peers exchange the information needed to set up a connection, corresponding to Snowflake rendezvous—unspecified [1 §3], leaving every application to invent its own mechanism.

As WebRTC is designed for the web, most implementations of WebRTC are embedded in web browsers, and are not easily removed from that context. Snowflake originally used a WebRTC library extracted from Chromium, but that eventually proved unworkable for cross-platform deployment. Since 2019, Snowflake has used Pion [30], an independent implementation of WebRTC.⁵ It is not tied to any browser, which

is both good and bad. The good is less development friction, better memory safety (Pion is written in Go, while Chromium WebRTC is C++), and a working relationship with upstream developers that enables us to get fingerprinting-related changes made. The bad is that the protocol fingerprints of Pion do not automatically match the mostly browser-originated WebRTC that Snowflake aims to blend in with.

The following is a list of fingerprinting concerns that bear on Snowflake, and how we have tried to address them. The existence of a fingerprinting vulnerability does not automatically invalidate a circumvention system: the question is whether the vulnerability is reparable. Even among demonstrable vulnerabilities, some are more and some are less practical for a censor to take advantage of. The important thing is to build on a solid foundation; minor flaws may be patched up as necessary.

Selection of STUN servers It is not unusual for a WebRTC application to use STUN, but the choice of what STUN servers to use is up to the application. Running dedicated STUN servers just for Snowflake would not work, because a censor would experience no collateral harm in blocking them. Our deployment uses a pool of public STUN servers that are used for applications other than circumvention, filtered for those that support the NAT behavior discovery feature of Section 2.2. The client chooses a random subset of servers from the pool when it makes a connection; this is because not every STUN server is accessible under every censor.

Format of STUN messages STUN is most often deployed over plaintext UDP, which leaves the formatting of messages open to inspection and potential fingerprinting. STUN messages consist of a fixed header followed by a variable-length list of ordered attributes [29 §5]. What attributes appear, and their order, depends on the STUN implementation and how the application uses it.

We have not done anything in particular to disguise STUN messages. Though plaintext UDP is the most common, STUN specifies other transports, including encrypted ones like DTLS. These may be options for Snowflake in the future—of course, only if they are common enough that their use does not stick out on its own.

Rendezvous Because the rendezvous methods of Section 2.1 are modular, each one needs its own justification as to why it should be difficult to block. In addition, they must be implemented in a way that does not expose accidental distinguishers. For example, the domain fronting and AMP cache rendezvous methods use HTTPS, which is TLS, which means that TLS fingerprinting is a concern [11 §5.1]. Snowflake, like many other circumvention systems, uses the uTLS package [14 §VII] for a client TLS fingerprint that is randomized or that imitates common browsers. See Section 5.2 for an account of when

⁵<https://bugs.torproject.org/tpo/anti-censorship/pluggable-transports/snowflake/28942>

State that “untrusted messenger” protects proxies as well.

domain fronting rendezvous was briefly blocked in Iran, because we were slow in activating uTLS.

Though each rendezvous method may be difficult to block in itself, a censor might combine a low-confidence detection of rendezvous with features from other phases of Snowflake data exchange to strengthen its guess.

DTLS The outermost layer of a WebRTC data connection, directly exposed to a censor, is DTLS (Datagram TLS) over UDP. DTLS is an adaptation of TLS [33 §1] to the datagram setting, and therefore inherits the fingerprinting concerns of TLS [14]. TLS/DTLS fingerprinting may involve, for example, inspecting the ciphersuites and extensions of Client Hello messages, and their order. If a combination is specific to a particular implementation of a circumvention system, it may be blocked at low cost.

Due to practical considerations, Snowflake’s defenses to DTLS fingerprinting are not very robust, and are reactive rather than proactive. In the realm of TLS one may use uTLS, but there is as yet no equivalent for DTLS. The present way of altering DTLS fingerprints in Snowflake is to submit a pull request upstream to Pion when a fingerprint feature used for blocking is identified. Section 5.1 documents how this has happened twice already, in response to blocking in Russia.

Data channel or media stream Besides data channels, WebRTC offers *media streams*, serving the purpose of real-time audio and video communication. Though both are encrypted, data channels and media streams are externally distinguishable because they use different containers. Data channels use DTLS, while media streams use DTLS-SRTP; that is, the Secure Real-Time Transport Protocol with a DTLS key exchange [32 §4.3].

Data channels are a closer match to Snowflake’s communication model: media streams are meant to contain encoded audio and video, not arbitrary binary data. But the use of DTLS rather than DTLS-SRTP could become a significant feature if other WebRTC applications mainly use media streams. Although it would be less convenient, it would be possible to adapt the WebRTC link between client and proxy to use a media stream rather than a data channel, either by modulating binary data into a well-formed encoded audio or video signal in the manner of, say, Stegozoa [12 §3.3], or by replacing encoded media content within SRTP packets, as in Protozoa [2 §4.4] or TorKameleon [40 §III-D].

Protocol fingerprinting is where most research on detecting Snowflake has focused. Fifield and Gil Epner [9] studied the network traffic of WebRTC applications, with the goal of finding fingerprinting pitfalls that might affect Snowflake, which was then in early development. Frolov et al. [14 §V-C] observed that the undisguised TLS fingerprint of domain fronting

rendezvous was distinctive, and introduced the uTLS package that Snowflake now uses to protect it.

MacMillan et al. [23] focused on the DTLS handshake, comparing Snowflake to three other WebRTC applications. They correctly anticipated features of the Pion DTLS handshake that would later be used to block Snowflake in Russia; see details in Section 5.1. Holland et al. [16 §5.3], using the bits of UDP datagrams directly as features, demonstrated approximately equal performance on the same DTLS handshake data set. Their automatically derived classifier assigned high feature importance to length fields in packets, and in fact did well even when deprived of DTLS payload features.

Chen et al. [4] combined features of rendezvous and DTLS in order to reduce false positives. Their classifier begins by looking for DNS queries for STUN servers and front domains typically used by Snowflake clients. They then apply a machine learning classifier to features of a subsequent DTLS handshake. The authors acknowledge that DTLS fingerprinting is fragile, as the DTLS fingerprint is, in principle, controllable by the application. The DNS prefilter may perhaps be mitigated by alternative rendezvous methods (Section 2.1), or by smarter selection of STUN servers. Xie et al. [42] trained a decision tree on packet size, direction, latency, and bandwidth features, with the aim of distinguishing Snowflake’s domain fronting rendezvous from other forms of HTTPS. A challenge in classifying rendezvous flows is that they do not consist of many packets, which limits the features a classifier has to work with. Their lowest reported false positive rate of 0.25% is, in our opinion, unworkably high, given the low base rate of Snowflake rendezvous connections.

Wails et al. [41] criticize past research on detecting circumvention systems, saying that accuracy claims do not hold up with the low base rates of circumvention traffic in practice. They develop classifiers for Snowflake and other circumvention protocols that improve on the state of the art, but find them still prohibitively imprecise at realistic base rates. They propose to reduce false positives by combining multiple observations per IP address—classifying hosts rather than flows—and suggest that Snowflake’s lack of fixed proxies mitigates against this enhancement.

Talk
about
traffic
shaping,
“Ground-
ing in
Empiri-
cism” [36]?

4 Experience

Snowflake has now been in operation for a few years. In lieu of a forward-looking evaluation, here we take a look back at the history of our deployment and reflect on the experience.

4.1 Client counts and bandwidth

Snowflake became available to end users gradually, reflecting a long development process. Development began in late 2015, and deployment in 2017, but the system only really became usable in 2020. It began to attract large numbers of users (enough




Figure 4: Estimated average simultaneous Snowflake users and bandwidth by day. The values at the far left end of the graph, in early July 2021, are about 200 users and 2.7 Mbit/s.

to merit a censor’s attention) in 2022, following network blocking events in Russia and Iran.

Snowflake shipped in the alpha release series of Tor Browser before graduating to the stable series. The first releases of Snowflake were for GNU/Linux in Tor Browser 7.0a1 on 2017-01-24⁶ and for macOS in Tor Browser 7.5a4 on 2017-08-08⁷. But we hit a roadblock in attempting to prepare releases for other platforms: the Chromium-derived WebRTC library we had used to that point presented major difficulties in Tor Browser’s cross-compiling, reproducible build environment. What let us resume making progress was a switch to Pion WebRTC [30] in 2019. With it, we were able to release Snowflake for Windows in Tor Browser 9.0a7 on 2019-10-01⁸, and for Android in Tor Browser 10.0a1 on 2020-06-02⁹.

While at this point Snowflake was available on every platform supported by Tor Browser, it was not yet comfortably usable. Two important parts were missing: no NAT type matching (Section 2.2) meant that a client could not always connect to its assigned proxy; and a lack of persistent session state (Section 2.3) meant that even if a proxy connection was successful, the client’s session would end once that proxy disappeared. For these reasons, by early 2020, the average number of concurrent users had not risen above 40. The Turbo Tunnel session persistence feature became available to users in Tor Browser 9.5a13 on 2020-05-22.¹⁰ The client part of NAT behavior detection was released with Tor Browser 10.0a5 on 2020-08-19¹¹, and proxy support was added on 2020-11-17¹². With these changes, Snowflake became practical for daily browsing, and the number of users began to grow into 2021.

This brings us to Figure 4, which shows the Snowflake users

and daily bandwidth since July 2021. Be aware: the chart does not show a count of unique clients, but rather the *average number of concurrent clients* per day [21]. For example, the value of 12,000 on 2022-05-01 means that, on average, 12,000 clients were using the service at any point in time on that day. The contribution of a client depends on how long it uses the system each day, not how many temporary proxies it uses. The average concurrent client count is estimated from the number of directory requests that are published in the descriptors sent by Tor bridges to the bridge authority and archived by CollecTor.¹³ This method of estimating usage metrics was developed specifically to preserve user anonymity. We discuss the techniques and challenges of obtaining country-specific usage counts more in Section 5 where we provide measurements of Snowflake usage in response to censorship events.

Snowflake’s growth began in earnest when it became part of default installations. Orbot, a mobile app that provides a VPN-like Tor proxy, added a Snowflake client in version 16.4.0 on 2021-01-12.¹⁴ Snowflake graduated to Tor Browser’s stable series in Tor Browser 10.5 on 2021-07-06¹⁵, becoming a third built-in circumvention option alongside meek and obfs4. Being part of a stable release meant that it was easily available to all Tor users, not just a self-selected group of alpha testers. The number of users steadily increased over the next five months, reaching almost 2,000 by December 2021.

A network censorship event may have the effect of either increasing or decreasing the number of users of a circumvention system. The user count decreases when the system is not robust enough and falls to blocking; but increases when it remains one of a diminished number of ways to reach the outside world. Two such censorship events, one in Russia and one in Iran, had the effect of increasing the number of Snowflake users by multiples.

⁶<https://bugs.torproject.org/tpo/applications/tor-browser/20735>

⁷<https://bugs.torproject.org/tpo/applications/tor-browser/22831>

⁸<https://bugs.torproject.org/tpo/anti-censorship/pluggable-transports/snowflake/25483>

⁹<https://bugs.torproject.org/tpo/applications/tor-browser/30318>

¹⁰<https://bugs.torproject.org/tpo/anti-censorship/pluggable-transports/snowflake/33745>

¹¹<https://bugs.torproject.org/tpo/anti-censorship/pluggable-transports/snowflake/34129>

¹²<https://bugs.torproject.org/tpo/anti-censorship/pluggable-transports/snowflake/40013>

¹³<https://metrics.torproject.org/collector.html#type-extra-info>

¹⁴<https://github.com/guardianproject/orbot/releases/tag/16.4.0-RC-1-tor-0.4.4.62021-01-12>

¹⁵<https://blog.torproject.org/new-release-tor-browser-105>

On 2021-12-01, some ISPs in Russia deployed measures to block most forms of access to Tor, including Snowflake [44]. The measures varied in their effectiveness; in the case of Snowflake, blocking was triggered by a particular feature of the DTLS handshake which we were able to mitigate in new releases within a few weeks.¹⁶ Over the next two months the total number of Snowflake users quadrupled. By May 2022, about 70% of Snowflake users were in Russia. The user count in Russia got an additional small boost, visible in the graph, starting on 2022-07-14, when Tor Browser 11.5 added the Connection Assist feature, which automatically enables circumvention options when needed.¹⁷ We will present more details of blocking actions in Russia and their effect on usage in Section 5.1.

The next event to have a major effect on Snowflake usage was the nationwide protests that started in Iran on 2022-09-16. The government imposed network shutdowns and additional network blocking, severe even by the standards of a country already notorious for censorship [3]. Users turned to the few circumvention systems that continued working in the face of the new restrictions, one of which was Snowflake. Adoption was rapid: on 2022-09-20, Iran accounted for only 1% of Snowflake users; by 2022-09-24 it was 67%. The influx of users had us scrambling for a few days to implement performance improvements. Two weeks later, on 2022-10-04, usage dropped almost as quickly as it had risen—the cause was the blocking of a TLS fingerprint used by the Snowflake client.¹⁸ After we released fixes for the TLS fingerprinting issue, the user count began to recover going into 2023. But in our haste to deploy optimizations in September, we had introduced a bug that harmed performance, getting worse with more users¹⁹, which dragged the count down again, until the bug was fixed in mid-March. Umayya et al. happened to do performance tests of Snowflake during this time [37 §4.6]—their results bear out the lessened reliability of connections before the performance bug was fixed²⁰. More details on blocking actions in Iran will appear in Section 5.2.

For most of this history, we ran the backend bridge on a single server, upgrading and optimizing it as needed. But as the bridge reached its hardware capacity, and performance improvements got harder to achieve, we deployed a second bridge to share the load. We discuss the challenges and design considerations of doing so in Section 4.4. The new bridge was made available in Tor Browser 12.0 on 2022-12-07. By July, it supported about 18% of users.

The drop in users by about half on 2023-09-20 was not caused by censor action: rather, it was an unexpected change in the cloud infrastructure we used for domain fronting rendezvous. The front domain we had been using changed its hosting to a different CDN, which caused client rendezvous



Figure 5: Unique proxy IP addresses per day, by proxy type. The two steps in the graph correspond to the invasion of Ukraine by Russia in February 2022, and network restrictions in Iran beginning September 2022, at which times there were campaigns to encourage running Snowflake proxies. Unknown proxy types (amounting to fewer than 50) are not shown.

messages to fail to reach the broker.²¹ The user count began to recover after we made releases with alternative front domains.²²

As of 2024-02-01, Snowflake had transferred 13.9 PB of circumvention data. We are referring to goodput: Tor TLS traffic inside the tunnel, ignoring WebRTC, WebSocket, and KCP/smux overhead. At that time, about 0.7% of all Tor users (24% of bridge users) used Snowflake to connect to Tor.

4.2 Number and type of proxies

Snowflake’s effectiveness depends on its proxies, of which there are several types. The primary type is the web browser extension, which, once installed, works in the background while the browser is running. There is also a “web badge” version of the proxy that does not require installation. It uses the same JavaScript code as the extension, but runs in an ordinary web page. Some people leave a browser tab idling on the web badge page, rather than install a browser extension. Apart from the web-based proxies, we provide a standalone, command-line proxy that does not require a browser. This version is convenient to install on a rented VPS, for example. Running a long-term proxy at a fixed IP address is somewhat at odds with Snowflake’s goal of proxy address diversity and agility, but these standalone proxies are valuable because they tend to have less restrictive NATs, making them compatible with more clients. Finally, Orbot, a mobile app for accessing Tor, besides being able to *use* Snowflake for circumvention, can also *provide* Snowflake proxy service to others, a feature called “kindness mode.”

We coordinated with the Tor Project’s network health team to collect privacy preserving metrics at the broker during the client and proxy polls of the Snowflake rendezvous.²³ The resulting metrics are published at the end of every 24 hour

¹⁶<https://bugs.torproject.org/tpo/applications/tor-browser-build/40393>

¹⁷<https://blog.torproject.org/new-release-tor-browser-115/>

¹⁸<https://bugs.torproject.org/tpo/anti-censorship/pluggable-transports/snowflake/40207>

¹⁹<https://bugs.torproject.org/tpo/anti-censorship/pluggable-transports/snowflake/40260>

²⁰<https://bugs.torproject.org/tpo/anti-censorship/pluggable-transports/snowflake/40262>

²¹<https://forum.torproject.org/t/9346>

²²<https://bugs.torproject.org/tpo/applications/tor-browser/42120>

²³<https://gitlab.torproject.org/tpo/network-health/metrics/collector/-/issues/29461>

collection period²⁴ in aggregate and we do not publish or store client or proxy IP addresses. Metrics concerning client polls are rounded up to the nearest multiple of 8 to prevent individual participation patterns from becoming visible in the aggregate counts. The collected metrics allow us to determine daily unique proxy IP counts, along with the associated country codes, proxy types, NAT behavior types (i.e., restricted, unrestricted, or unknown), and how many times a proxy was matched with a client. Figure 5 shows the daily counts of each proxy type. Browser extension proxies predominate, representing about 80% of 140,000 daily IP addresses. For comparison, there were about 1,900 of the more traditional style of Tor bridge at this time. The difference is attributable to the relative ease of running a Snowflake proxy versus a Tor bridge—though the comparison is not quite direct, because Tor bridges have better defenses against enumeration than do Snowflake proxies.

It was not clear, at the outset, that it would even be possible to attract enough proxies to make Snowflake meaningfully blocking resistant and support a reasonable number of users. Lowering the technical barriers to running a proxy was only part of it; getting there also took intentional advocacy and outreach. In the early days, circa 2017, the only round-the-clock proxy support was a few standalone proxies, run by us for the benefit of alpha tester clients. The browser extension became available in mid-2019.²⁵ In the latter half of 2019, additional proxy capacity came when Cupcake, a browser extension for flash proxy with an existing user base, was repurposed for Snowflake.²⁶ Orbot’s Snowflake proxy feature was added in version 16.4.1 in February 2021.²⁷ (In Figure 5, Orbot is counted among the standalone proxies until January 2022, when it got its own proxy type designation.)

It is worth reflecting on the greater popularity of the browser extension compared to the web badge. The latter had been envisioned as the primary source of proxies in flash proxy, the idea being that people’s browsers would automatically become proxies while reading sites that had the flash proxy badge installed, unless they checked an option to prevent it. We decided, early on, that flash proxy’s opt-out permission had been a mistake, and that Snowflake would be opt-in. In order to run a proxy, a person must take a positive action such as installing a browser extension or activating a toggle on a web page. Our initial worry that this policy would reduce the number of proxies turned out to be unfounded. People find an informative, interactive proxy control panel more appealing than a nondescript badge graphic, and install the browser extension in greater numbers than ever used the web badge in flash proxy.

²⁴<https://metrics.torproject.org/collector.html#snowflake-stats>

²⁵https://bugs.torproject.org/tpo/anti-censorship/pluggable-transport/snowflake/30931#note_2593598, https://bugs.torproject.org/tpo/anti-censorship/pluggable-transport/snowflake/30990#note_2593718

²⁶<https://github.com/glamrock/cupcake/issues/24>

²⁷<https://github.com/guardianproject/orbot/releases/tag/16.4.1-BETA-2-tor.0.4.4.6>




Figure 6: Proxy pool churn in January 2023. The dark upper line shows the number of unique proxy IP addresses in a 24-hour window starting at the point indicated. The lighter descending lines show how many of the same IP addresses remain in the pool, at 1-hour intervals up to 40 hours later. It takes about 20 hours for 50% of the proxy pool to turn over.

4.3 Proxy churn

The size of the proxy pool is not the only measure of its quality. Also important is its “churn,” the rate at which it is replenished with fresh proxy IP addresses. Churn determines how hard a censor would have to work to keep a blocklist of proxy IP addresses up to date; or alternatively, how quickly a momentarily complete blocklist would lose effectiveness.

We ran an experiment²⁸ to measure churn. Every hour, the broker logged a record of the proxy IP addresses it had seen in the past hour. To avoid storing real proxy IP addresses, each record was not a transparent list, but a HyperLogLog++ sketch [15], a probabilistic data structure for estimating the number of distinct elements in a multiset. We additionally hashed proxy IP addresses with a secret string before adding them to a sketch, to prevent their recovery from our published data. A sketch supports two basic operations: count and merge. Given a sketch X , we may compute an approximate count $|X|$ of its unique elements, and given two sketches X and Y , we may merge them into a new sketch representing the union $X \cup Y$. The quantity we are interested in, the size of the intersection of two sketches, is computed using the formula $|X| + |Y| - |X \cup Y|$. Such a computation estimates how many IP addresses are shared across two samples of the proxy pool.

Figure 6 visualizes the results of the churn experiment. We merged consecutive sketches over a 24-hour window to serve as a reference, then computed the size of its intersection with other windows of the same size, offset by $+1, +2, \dots, +40$ hours. After 1 hour, the shifted window still has, on average, 97.3% of addresses in common with the reference; after 12 hours the fraction has fallen to 68.8%; by the time 24 hours have elapsed, only 38.2% of proxy IP addresses are ones that had been seen in the previous day.

²⁸<https://bugs.torproject.org/tpo/anti-censorship/pluggable-transport/snowflake/34075>

4.4 Multiple bridges

In the abstract model of Figure 1, the bridge is a single, centralized entity. It *can* be centralized because it is never accessed directly, but only via temporary proxies. Unlike more traditional static proxy systems, Snowflake does not benefit, in terms of blocking resistance, from having multiple bridges. For scalability reasons, though, it is useful for “the” bridge to be realized as multiple servers, each handling a fraction of client traffic.

Our deployment now uses two bridges. Generalizing from one bridge to two required changes to the messages exchanged between clients, proxies, and the broker. Unfortunately, the fact of multiple bridges cannot be made fully transparent to clients, for technical reasons related to Tor. In our design, the client informs the broker of what bridge it wants to use, the broker conveys the choice to the proxy, and the proxy connects to the client’s chosen bridge. This is in contrast to other imaginable designs where the choice of bridge is made by the broker or the proxy. We will discuss design considerations and tradeoffs.

One minor difficulty is distributing the Turbo Tunnel layer. Recall from Section 2.3 that Snowflake has the notion of an end-to-end session between a client and the bridge, independent of temporary proxy connections that carry it. This is made possible by extensive state stored at the bridge: a table of clients, reassembly buffers, transmission queues, timers, and so on. While it is certainly possible to instantiate one such bundle of state variables per bridge, a session begun in one instance must remain with that instance—no other has the context necessary to make the packets of the session meaningful. This difficulty might be resolved by hashing the client’s session identifier string to index a consistent bridge per session, as long as the set of bridges does not change too frequently.

There is another difficulty that is harder to work around. A Tor bridge is identified by a long-term identity public key. If, on connecting to a bridge, the client finds that the bridge’s identity is not the expected one, the client will terminate the connection [5 §4.2]. The Tor client can configure at most one identity per bridge; there is no way to indicate (with a certificate, for example) that multiple identities should be considered equivalent. This constraint leaves two options: either all Snowflake bridges must share the same cryptographic identity, or else it must be the client that makes the choice of what bridge to use. While the former option is possible to do (by synchronizing identity keys across servers), every added bridge would increase the risk of compromising the all-important identity keys. Our vision was that different bridge sites would run in different locations with their own management teams, and that any compromise of a bridge site should affect that site only.

These considerations led us to a multi-bridge design in which clients have awareness of (at least a subset of) all bridges, and it is the client that chooses which bridge will be used for a particular session.²⁹ The client includes a bridge identity string

in its rendezvous message to the broker (Section 2.1); then the broker maps the identity to the WebSocket URL of the corresponding bridge, and conveys that URL to the proxy that’s chosen to serve the client. We rely on clients choosing uniformly to equalize load across bridges. A consequence is that every bridge must meet a minimum performance standard: we cannot, say, centrally assign 20% of clients to one and 80% to another according to their relative capacity. Another drawback is that there is currently no way to instruct Tor to connect to only one of the bridges it knows about (short of rewriting the configuration file): if two bridges are configured, Tor starts two sessions through Snowflake, each doing its own rendezvous, which is wasteful and makes for a more conspicuous network fingerprint. Still, this is the best solution we have found, given the constraints. A deployment not based on Tor would have more flexibility.

A client-chooses design risks misuse by clients, if not handled carefully. Clients should only be able to select from a limited set of known bridges, not cause proxies to connect to arbitrary destinations—otherwise the tens of thousands of Snowflake proxies might be weaponized to attack third parties. The client’s bridge selection in its rendezvous message is represented not as an IP address or hostname, but as a hash of the bridge’s public identity key. The broker maps the identity to a WebSocket URL by consulting its own local database of known bridges, and rejects rendezvous messages that refer to an unknown bridge. After the broker tells the proxy what WebSocket URL to connect to, the proxy does its own check, verifying that the hostname in the URL is a subdomain of a known suffix reserved for Snowflake bridges. So there are two independent safeguards against misuse.

5 Notable blocking attempts

In Section 4.1 we saw how Snowflake’s user counts have at times been affected by the blocking actions of censors. Now we take a closer look at selected censorship events. The effect of censorship has usually been to increase, rather than decrease, the number of Snowflake users. This is no paradox: as censorship intensifies, users are displaced from less resilient to more resilient systems. Snowflake’s blocking resistance has not in every case been a success, though, and here we also reflect on missteps and persistent challenges. The examples are taken from Russia, Iran, China, and Turkmenistan, and are selected for being significant and instructive. Common lessons are that communication with affected users is invaluable in quickly understanding and reacting to blocking; and that blocking resistance is relative to a given censor, because every censor’s cost calculus is different.

Snowflake is blockable by a censor that is willing to block WebRTC. We would not argue otherwise. Indeed, we believe this is how a circumvention system should be presented: not by

²⁹<https://bugs.torproject.org/tpo/anti-censorship/pluggable-transports/snowflake/>

arguing its unblockability in absolute terms, but by laying out what actions by a censor would suffice to block it—or more to the point, *what sacrifices a censor would have to make* in order to block it. Advancing the state of the art of censorship circumvention consists in pushing blocking beyond the capabilities of more and more censors.

Tor bridges report aggregate binned counts by country code of connected unique IP addresses per day in the descriptors uploaded to the bridge authority. We use the Tor Metrics method of combining the distribution of counts by country code with the number of directory requests to obtain an estimate of the average number of concurrent clients per day for each location [21]. The mapping of IP addresses to country codes is not without flaws. During the time of the measurements shown here, Tor uses the IPFire location database.³⁰ There is at least one instance where we were able to detect geolocation inaccuracies after noting a significant drop in Snowflake users thought to be located in the US that correlated directly with a blocking event in Iran.³¹

5.1 Blocking in Russia

Snowflake, along with other common ways of accessing Tor, was blocked in a subset of ISPs in Russia on 2021-12-01 [44]. The event was evidently coordinated and targeted, as it happened suddenly and affected many Tor-related protocols at once. Besides Snowflake, a portion of Tor relays and bridges, as well as some servers of the circumvention transports meek and obfs4, were blocked, at least temporarily. The blocking campaign was less than totally successful—one of its effects was to substantially increase the number of users accessing Tor via circumvention transports, Snowflake among them.

We benefited from established relationships with developers and users in Russia, one of whom, through manual testing, found what traffic feature was being used to distinguish Snowflake. It was DTLS fingerprinting, of the kind cautioned about in Section 3.³² Specifically, it was the presence of a supported_groups extension in the DTLS Server Hello message produced by Pion. The extension being present in Server Hello was a bug³³—but one that afforded the censor a feature to distinguish DTLS connections with a Pion implementation in the server role from other forms of DTLS. The process of finding the flaw, fixing it, and shipping new releases of Tor Browser took a few weeks³⁴, after which the user count rose quickly: from the beginning to the end of December 2021, the number of users in Russia grew from about 400 to over 4,000 (Figure 7). Snowflake was to become a significant tool amid the general intensification of censorship in Russia following the invasion of Ukraine in February 2022.




Figure 7: Snowflake users in Russia (average concurrent). Events discussed in the text are marked. The attempted blocking of Tor-related transports in December 2021 led to Snowflake’s first surge in usage. The decrease in September–October 2022 coincided with an even larger influx from Iran.




Figure 8: Snowflake users in Iran. Heightened censorship beginning in September 2022 caused Iran to become the single biggest source of Snowflake users. The drop in October 2022 was the result of TLS fingerprint blocking, which interfered with rendezvous and took some time to mitigate.




Figure 9: Snowflake users in China. Though no sustained blocking is evident, disruption of domain fronting rendezvous for three days in May 2023 briefly depressed user numbers.




Figure 10: Snowflake users in Turkmenistan. This graph shows a different range of dates than the other three. Though there have never been many Snowflake users in Turkmenistan, blocking events are evident on 2021-10-24 and 2022-08-03.

³⁰<https://www.ipfire.org/projects/location/>

³¹https://bugs.torproject.org/tpo/anti-censorship/pluggable-transports/snowflake/40207#note_2844116

³²https://bugs.torproject.org/tpo/anti-censorship/pluggable-transports/snowflake/40014#note_2765074

³³<https://github.com/pion/dtls/issues/409>

³⁴https://gitlab.torproject.org/tor-browser-build/-/merge_requests/375

The Server Hello supported_groups distinguisher had been discovered and documented by MacMillan et al. [23 §3] already in 2020. We might have avoided this blocking event by proactively fixing the known distinguisher—but it was not necessarily the wrong call not to have done so. There is always more to do than time to do it; one must consider the opportunity cost of preempting specific blocking that may not come to pass. In this case, a reactive approach by us was enough: the loss was minor, and we were able to patch the problem quickly. Even in ISPs where the blocking rule was present, it did not block 100% of Snowflake connections, because of the how it targeted a quirk in Pion, and only in Server Hello. When the DTLS server role in the WebRTC data channel was played by a non-Pion peer, such as a web browser proxy, the feature was not present.

In May 2022 we got a report of a new detection rule, this time keying on not just the presence, but the *contents* of the supported_groups extension, at a byte offset suggesting that it targeted the Client Hello message, not Server Hello.³⁵ The presence of a supported_groups extension in Client Hello is not at all unusual, but the specific groups offered by Pion’s implementation differed from those of common browsers. Though we confirmed the existence of the blocking rule, testers reported that Snowflake continued to work—which may have something to do with the fact that the Snowflake client does not always play the client role in DTLS. If the Snowflake client is the DTLS server, and the DTLS client is a browser proxy, then the byte pattern looked for by the blocking rule does not appear. We developed a mitigation, but by the time we prepared a testing release in July 2022, the new rule had apparently been removed and replaced by another. We can only speculate as to reasons, but it may be that the old rule had too many false positives, or was just not effective enough.

The detection rule that replaced supported_groups in Client Hello looked for the presence of a Hello Verify Request message.³⁶ Hello Verify Request is an anti-denial-of-service feature in DTLS, in which the server sends a random cookie to the client, and the client sends a second Client Hello message, this one containing a copy of the cookie [33 §5.1]. It is not an error to send Hello Verify Request (it is a “MAY” in the RFC), but because the Pion implementation in Snowflake sent it, and major browsers did not, it was a reliable indicator of Snowflake connections. (Those, at least, in which the DTLS server role was played by a Snowflake client or standalone proxy.) This distinguisher, too, had been anticipated by MacMillan et al. in 2020 [23 §3]. The first reports of the blocking rule arrived in July 2022; but as you can see in Figure 7, it had no apparent immediate effect. It is hard to say whether the drastic decline in October 2022 was a consequence of this rule, or some other, unidentified one. That decline coincided with an explosion of users from Iran, which temporarily affected the usability of the whole system. We deployed a mitigation to remove the Hello

Verify Request message from Snowflake, regrettably, only in February 2023³⁷, after which the number of users in Russia began to recover.

The case of Snowflake in Russia illustrates some of the complexity of censorship measurement. The answer to a question like “Does Snowflake work in Russia?” is not a simple yes or no. It may depend on the date, the ISP, and even such factors as which endpoint plays the DTLS server role.

5.2 Blocking in Iran

In late September 2022, users from Iran became the majority of Snowflake users almost overnight, only to fall just as quickly two weeks later. See Figure 8. The cause of the rise was extraordinary new network restrictions amid mass protests [3]; the cause of the decline was TLS fingerprint blocking, which stopped Snowflake rendezvous from working. The crypto/tls package of the Go programming language (in which the Snowflake client is written) may produce several slightly different TLS fingerprints, depending on hardware capabilities and how it was compiled.³⁸ It was one of these fingerprints that was blocked. Because the blocking rule was so specific, some users were affected and others were not. Why would a censor block only one (even if the most common) TLS fingerprint? It may have been a simple oversight. On the other hand, it is not certain that the blocking was meant for Snowflake specifically. Go is a popular language for implementing circumvention systems; Snowflake may have been caught up in blocking that was intended for another system.

The fact that simple TLS fingerprinting worked to block Snowflake rendezvous was carelessness on our part. Aware of the possibility, we had already implemented TLS camouflage using uTLS in the Snowflake client, but failed to turn it on by default. Activating the feature required only a small configuration change³⁹, but we had to wait for new releases of Tor Browser and Orbot to get it into the hands of users: see the September–November 2022 interval in Figure 8.

After repairing the TLS fingerprinting flaw, the number of users from Iran gradually recovered to near its former peak. We are aware of only minor disruptions after this time. The default rendezvous front domain was blocked (by TLS SNI) in some ISPs between 2023-01-16 and 2023-01-24⁴⁰, which we confirmed using data from the censorship measurement platform OONI. A reduction in users is visible at this time. AMP cache rendezvous continued to work. OONI measurements in the weeks after the block was lifted showed sporadic failures to connect to the front domain. If these were further attempts at blocking, they did not have much of an effect.

³⁵<https://bugs.torproject.org/tpo/anti-censorship/censorship-analysis/40030>

³⁶https://bugs.torproject.org/tpo/anti-censorship/censorship-analysis/40030#note_2823140

³⁷https://gitlab.torproject.org/tpo/applications/tor-browser-build/-/merge_requests/637

³⁸https://bugs.torproject.org/tpo/anti-censorship/pluggable-transports/snowflake/40207#note_2844163

³⁹https://gitlab.torproject.org/tpo/applications/tor-browser-build/-/merge_requests/540

⁴⁰https://bugs.torproject.org/tpo/anti-censorship/team/115#note_2873040

5.3 Blocking in China

The user count graph from China, Figure 9, does not show any drastic changes like others we have seen so far. There is a modest but respectable number of Snowflake users in China. Though there have been no singular, sustained events, we have seen evidence of short-term or tentative blocking attempts.

In May 2019, when Snowflake was still in alpha release, a user in China reported a failure to connect. Investigation revealed that the cause was IP address blocking of the few proxies that existed at the time.⁴¹ Rendezvous happened, and the STUN exchange worked, but the client and proxy could not establish a connection. We experimented with running a proxy at a previously unused IP address: clients in China could connect when they were assigned that proxy by the broker. This was back before the web browser extension proxy existed, and the only consistent proxy support was a few standalone proxies that we, the developers, ran at a static IP address. It ceased to be an issue as the proxy pool grew in size.

That same month, we noticed blocking of the default STUN server, of which there was only one at the time.⁴² The solution was to add more STUN servers⁴³, and select a subset of them on each rendezvous attempt⁴⁴. Curiously, it seems that when the STUN server was blocked, the standalone proxies that had been blocked earlier in the month became unblocked.⁴⁵

The next incidents we are aware of did not occur until 2023, recent enough to appear in Figure 9. On May 12, 13, and 14, a few users reported problems with domain fronting rendezvous.⁴⁶ We could not get systematic measurements, but it appeared that censorship was triggered by observing multiple (two or three) HTTPS connections with the same TLS SNI to certain IP addresses within a short time. It is possible that Snowflake was not the target of this blocking behavior, and was affected only as a side effect. If it indeed had to do with Snowflake, our best guess is that it was aimed at the multiple rendezvous mentioned in Section 4.4—though such a policy would certainly also affect a large number of non-Snowflake connections. The user count from China was about halved during those three days. On May 15, the blocking went away and user counts returned to normal.

Also in May 2023, one user reported apparent throttling (artificial reduction in speed by packet dropping) of TLS-in-DTLS connections, based on packet size and timing features.⁴⁷ Such a policy would affect Snowflake, because it transports Tor TLS inside DTLS data channels. Reportedly, adding padding to the first few packets to disrupt the size and timing signature was enough to prevent throttling. Our own speed tests run at the

⁴¹https://bugs.torproject.org/tpo/anti-censorship/pluggable-transport/snowflake/30350#note_2593274

⁴²https://bugs.torproject.org/tpo/anti-censorship/pluggable-transport/snowflake/30368#note_2593357

⁴³<https://bugs.torproject.org/tpo/anti-censorship/pluggable-transport/snowflake/30579>

⁴⁴https://gitlab.torproject.org/tpo/anti-censorship/pluggable-transport/snowflake/-/merge_requests/7

⁴⁵https://bugs.torproject.org/tpo/anti-censorship/pluggable-transport/snowflake/30368#note_2593360

⁴⁶<https://bugs.torproject.org/tpo/anti-censorship/censorship-analysis/40038>

⁴⁷<https://github.com/net4people/bbs/issues/255>

time did not show evidence of throttling, with or without added padding.⁴⁸ There was no obvious reduction in the number of users. It may have been a localized, ISP-specific phenomenon.

5.4 Blocking in Turkmenistan

There have never been more than a few tens of Snowflake users in Turkmenistan. Even so, it has happened at least twice that the number of users dropped suddenly to zero, as shown in Figure 10. We found a variety of causes: domain name blocking by DNS and TCP RST injection; and blocking of certain UDP port numbers commonly used for STUN.

Turkmenistan is a particularly challenging environment for circumvention. Though relatively unsophisticated, censorship there is more severe and indiscriminate than in the other places we have discussed. Only a small fraction of the population has access to the Internet at all, which makes it hard to communicate with volunteer testers and lengthens testing cycles. We have been able to mitigate Snowflake blocking in Turkmenistan, but only partially, and after protracted effort.

The drop on 2021-10-24 was caused by blocking of the default broker front domain.⁴⁹ We determined this by taking advantage of the bidirectionality of the Turkmenistan firewall. Nourin et al. [26 §2] provide more details; we will state just the essential information here. Among the censorship techniques used in Turkmenistan are DNS response injection and TCP RST injection. DNS queries for filtered hostnames receive an injected response containing a false IP address; TLS handshakes with a filtered SNI receive an injected TCP RST packet that tears down the connection. Conveniently for analysis, it works in both directions: packets that *enter* the country are subject to injection just as those that *exit* it are. By sending probes into the country from outside, we found that the default broker front domain was blocked at both the DNS and TLS layers. It was some time—not until August 2022—before we got confirmation from testers that an alternative front domain worked to get around the block of the broker.

The increase in the number of users from May to August 2022 was caused by a partial unblocking of the broker front domain on 2023-05-03. We realized this only in retrospect, from examination of data from Censored Planet [35], a censorship measurement platform that had continuous measurements of the domain at that time, in one autonomous system in Turkmenistan. There was a shift from RST responses to successful TLS connections on that date. DNS measurements did not catch the moment of the shift, but they also showed no signs of blocking after that date. Evidently, some users were then able to connect. But the unblocking must not have been everywhere, because as late as 2022-08-18, users reported that RST injection was still in place for them (though DNS injection had stopped).

⁴⁸https://bugs.torproject.org/tpo/anti-censorship/pluggable-transport/snowflake/40251#note_2906723

⁴⁹<https://bugs.torproject.org/tpo/anti-censorship/censorship-analysis/40024>

There was yet another layer to the blocking. Even if they could contact the broker (at the default or an alternative front domain), clients could not then establish a connection with a proxy. Testing revealed blocking of the default STUN port, UDP 3478. A client that cannot communicate with a STUN server cannot find its ICE candidate addresses (Section 2.2), without which most WebRTC proxy connections will fail. (The exceptions are proxies without NAT or ingress filtering. While there are some such proxies, censorship in Turkmenistan also outright blocks large parts of IP address space, including data center address ranges where those proxies tend to run.) As chance would have it, the NAT discovery feature we rely on for testing the NAT type of clients requires STUN servers to open a second, functionally equivalent listener on a different port [22 §6], commonly 3479. Changing to those alternative port numbers enabled some users to connect to Snowflake again. Specifically, STUN servers on port 3479 worked in AGTS, one of two major affected ISPs. The workaround did not work in Turkmentelecom, the other ISP, where port 3479 was blocked. Though we do not have continuous measurements to be sure, we suspect that the STUN port blocking began on 2022-08-03 and precipitated the drop seen on that date in Figure 10.

The blocking techniques described in this section are crude, and surely result in significant overblocking—but they nevertheless offer greater challenges to circumvention than the more considered blocking of Russia and Iran. We highlight this to make the point that blocking resistance cannot be defined in absolute terms, but only relative to a particular censor. Censors differ not only in resources (time, money, equipment, personnel), but also in tolerance for the social and economic harms of overblocking. Circumvention can only respond to and act within these constraints. The government of Turkmenistan has evidently chosen to prioritize political control over a functioning network, to an extreme degree. To paraphrase one of our collaborators: “What they have in Turkmenistan can hardly be called an Internet.”⁵⁰ In a network already damaged by oppressive policy, the additional harm caused by the clumsy blocking of this or that circumvention system is comparatively small. This shows the sense in which a resource-poor censor can “afford” certain blocking actions that a richer, more capable censor cannot.

6 Future work

A natural extension of Snowflake would be to have it access systems other than Tor—ordinary VPNs, for example. Tor has its benefits: an existing user base, a standard (pluggable transports) for integrating circumvention modules, and exit nodes separate from entry nodes, which relieve the circumvention developer of the concerns associated with actually exiting traffic to its destination. But Tor has drawbacks as well, notably

its lower speed and lack of support for UDP and other non-TCP protocols. Nothing inherently ties Snowflake to Tor, and it might easily be adapted to other systems. One question is whether every Snowflake-like deployment should manage its own pool of proxies, or if proxies can somehow be shared. Building Snowflake’s population of proxies has been a substantial undertaking in itself—for every project to have to repeat the process from scratch would be a regrettable duplication of effort. There is no reason why one proxy might not serve multiple projects, the client expressing its preference in the same way it now signals which Tor bridge to use (Section 4.4). But there would be design issues to work out. While some proxy operators may be happy to donate bandwidth to a free-to-use project like Tor, they may need more incentive than altruism to help a commercial VPN. A shared deployment would impose additional friction on development (making it harder to alter the proxy protocol, for example). Rather than retrofit the current Tor-based proxies with support for other systems, a next-generation proxy pool might be designed from the ground up with multiple cooperating projects in mind. If it proved successful, the Tor deployment could migrate to it.

The Turbo Tunnel reliability layer of Section 2.3 was necessary for providing a continuous session abstraction over a sequence of unreliable proxies. But it might do even more: in particular, it should be possible for a client to multiplex its traffic over multiple proxies not just sequentially, but in parallel. (Something like multipath TCP.) Sequence numbers in the inner reliability layer would ensure a reliable stream, even when proxies have different lifetimes and performance characteristics. Multiplexing could increase performance by using the sum of the bandwidths of the individual proxies, and reduce variability by hedging against the client being assigned one very slow proxy. Using two or more proxies at once would eliminate the brief pause for re-rendezvous between consecutive proxies that now occurs. Our experiments with multiplexing have so far not shown enough benefit to justify the change, though it may be a matter of tuning.⁵¹ And of course, analysis would be required to determine whether simultaneous WebRTC connections form a distinctive network fingerprint.

Present research questions.

Availability

The project web site, <https://snowflake.torproject.org/>, has links to source code and instructions for installing the proxy browser extensions.

Add Git clone URL or similar for the paper itself. Say it shows how to reproduce our figures. Must also include the churn logs of Section 4.3.

Acknowledgements

The Snowflake project has been made possible by the cooperation and support of many people and organizations. We want

⁵⁰https://bugs.torproject.org/tpo/anti-censorship/censorship-analysis/40024#note_2889792

⁵¹https://bugs.torproject.org/tpo/anti-censorship/pluggable-transports/snowflake/25723#note_2718643

to thank particularly: Chris Ball, Griffin Boyce, Roger Dingledine, Sean DuBois, Arthur Edelstein, Mia Gil Epner, gus-tavo gus, J. Alex Halderman, Haz Æ 41, Jordan Holland, Armin Huremagic, Ximin Luo, Kyle MacMillan, Ivan Markin, meskio, Prateek Mittal, Erik Nordberg, Linus Nordberg, Vern Paxson, Sukhbir Singh, Aaron Swartz, ValdikSS, Vort, Philipp Winter, WofWca, Censored Planet, the Counter-Power Lab at UC Berkeley, Greenhost, Guardian Project, Mullvad VPN, the Net4People BBS and NTC forums, OONI, the Open Technology Fund, Pion, the Tor Project, financial donors, and the volunteers who run Snowflake proxies.

References

- [1] Harald T. Alvestrand. Overview: Real-time protocols for browser-based applications. RFC 8825, January 2021. <https://www.rfc-editor.org/info/rfc8825>.
- [2] Diogo Barradas, Nuno Santos, Luís Rodrigues, and Vítor Nunes. Poking a hole in the wall: Efficient censorship-resistant Internet communications by parasitizing on WebRTC. In *Computer and Communications Security*. ACM, 2020. https://www.gsd.inesc-id.pt/~nsantos/papers/barradas_ccs20.pdf.
- [3] Simone Basso, Maria Xynou, Arturo Filastò, and Amanda Meng. Iran blocks social media, app stores and encrypted DNS amid Mahsa Amini protests, September 2022. <https://ooni.org/post/2022-iran-blocks-social-media-mahsa-amini-protests/>.
- [4] Junqiang Chen, Guang Cheng, and Hantao Mei. F-ACCUMUL: A protocol fingerprint and accumulative payload length sample-based Tor-Snowflake traffic-identifying framework. *Applied Sciences*, 13(1), 2023. <https://www.mdpi.com/2076-3417/13/1/622>.
- [5] Roger Dingledine and Nick Mathewson. Tor protocol specification, March 2023. <https://spec.torproject.org/tor-spec>.
- [6] Donald E. Eastlake 3rd. Transport Layer Security (TLS) extensions: Extension definitions. RFC 6066, January 2011. <https://www.rfc-editor.org/info/rfc6066>.
- [7] Nick Feamster, Magdalena Balazinska, Winston Wang, Hari Balakrishnan, and David Karger. Thwarting web censorship with untrusted messenger discovery. In *Privacy Enhancing Technologies*. Springer, 2003. <http://nms.csail.mit.edu/papers/disc-pet2003.pdf>.
- [8] David Fifield. Turbo Tunnel, a good way to design censorship circumvention protocols. In *Free and Open Communications on the Internet*. USENIX, 2020. <https://www.bamsoftware.com/papers/turbotunnel/>.
- [9] David Fifield and Mia Gil Epner. Fingerprintability of WebRTC. *CoRR*, abs/1605.08805, 2016. <https://arxiv.org/abs/1605.08805>.
- [10] David Fifield, Nate Hardison, Jonathan Ellithorpe, Emily Stark, Roger Dingledine, Phil Porras, and Dan Boneh. Evading censorship with browser-based proxies. In *Privacy Enhancing Technologies*. Springer, 2012. <https://crypto.stanford.edu/flashproxy/flashproxy.pdf>.
- [11] David Fifield, Chang Lan, Rod Hynes, Percy Wegmann, and Vern Paxson. Blocking-resistant communication through domain fronting. *Privacy Enhancing Technologies*, 2015(2), 2015. <https://www.bamsoftware.com/papers/fronting/>.
- [12] Gabriel Figueira, Diogo Barradas, and Nuno Santos. Stegozoa: Enhancing WebRTC covert channels with video steganography for Internet censorship circumvention. In *Asia CCS*. ACM, 2022. <https://dl.acm.org/doi/10.1145/3488932.3517419>.
- [13] Sergey Frolov, Jack Wampler, Sze Chuen Tan, J. Alex Halderman, Nikita Borisov, and Eric Wustrow. Conjure: Summoning proxies from unused address space. In *Computer and Communications Security*. ACM, 2019. <https://jhalderm.com/pub/papers/conjure-ccs19.pdf>.
- [14] Sergey Frolov and Eric Wustrow. The use of TLS in censorship circumvention. In *Network and Distributed System Security*. The Internet Society, 2019. <https://tlsfingerprint.io/static/frolov2019.pdf>.
- [15] Stefan Heule, Marc Nunkesser, and Alex Hall. HyperLogLog in practice: Algorithmic engineering of a state of the art cardinality estimation algorithm. In *Extending Database Technology*. ACM, 2013. <https://research.google/pubs/pub40671/>.
- [16] Jordan Holland, Paul Schmitt, Nick Feamster, and Prateek Mittal. New directions in automated traffic analysis. In *Computer and Communications Security*. ACM, 2021. <https://dl.acm.org/doi/10.1145/3460120.3484758>.
- [17] Christer Holmberg and Roman Shpount. Session Description Protocol (SDP) offer/answer considerations for Datagram Transport Layer Security (DTLS) and Transport Layer Security (TLS). RFC 8842, January 2021. <https://www.rfc-editor.org/info/rfc8842>.
- [18] Randell Jesup, Salvatore Loreto, and Michael Tüxen. WebRTC data channels. RFC 8831, January 2021. <https://www.rfc-editor.org/info/rfc8831>.
- [19] Ari Keränen, Christer Holmberg, and Jonathan Rosenberg. Interactive Connectivity Establishment (ICE): A protocol for network address translator (NAT)

traversal. RFC 8445, July 2018.
<https://www.rfc-editor.org/info/rfc8445>.

[20] Patrick Lincoln, Ian Mason, Phillip Porras, Vinod Yegneswaran, Zachary Weinberg, Jeroen Massar, William Simpson, Paul Vixie, and Dan Boneh. Bootstrapping communications into an anti-censorship system. In *Free and Open Communications on the Internet*. USENIX, 2012.
<https://www.usenix.org/conference/foci12/workshop-program/presentation/lincoln>.

[21] Karsten Loesing. Counting daily bridge users. Technical Report 2012-10-001, The Tor Project, October 2012.
<https://research.torproject.org/techreports/counting-daily-bridge-users-2012-10-24.pdf>.

[22] Derek MacDonald and Bruce Lowekamp. NAT behavior discovery using session traversal utilities for NAT (STUN). RFC 5780, May 2010.
<https://www.rfc-editor.org/info/rfc5780>.

[23] Kyle MacMillan, Jordan Holland, and Prateek Mittal. Evaluating Snowflake as an indistinguishable censorship circumvention tool. *CoRR*, abs/2008.03254, 2020.
<https://arxiv.org/abs/2008.03254>.

[24] Alexey Melnikov and Ian Fette. The WebSocket protocol. RFC 6455, December 2011.
<https://www.rfc-editor.org/info/rfc6455>.

[25] Milad Nasr, Hadi Zolfaghari, Amir Houmansadr, and Amirhossein Ghafari. MassBrowser: Unblocking the censored web for the masses, by the masses. In *Network and Distributed System Security*. The Internet Society, 2020. <https://www.ndss-symposium.org/ndss-paper/massbrowser-unblocking-the-censored-web-for-the-masses-by-the-masses/>.

[26] Sadia Nourin, Van Tran, Xi Jiang, Kevin Bock, Nick Feamster, Nguyen Phong Hoang, and Dave Levin. Measuring and evading Turkmenistan’s Internet censorship. In *The International World Wide Web Conference*. ACM, 2023.
<https://dl.acm.org/doi/abs/10.1145/3543507.3583189>.

[27] OpenJS Foundation. How AMP pages are cached.
https://amp.dev/documentation/guides-and-tutorials/learn/amp-caches-and-cors/how_amp_pages_are_cached [cited 2023-06-10].

[28] Marc Petit-Huguenin, Suhas Nandakumar, Christer Holmberg, Ari Keränen, and Roman Shpount. Session Description Protocol (SDP) offer/answer procedures for Interactive Connectivity Establishment (ICE). RFC 8839, January 2021.
<https://www.rfc-editor.org/info/rfc8839>.

[29] Marc Petit-Huguenin, Gonzalo Salgueiro, Jonathan Rosenberg, Dan Wing, Rohan Mahy, and Philip Matthews. Session Traversal Utilities for NAT (STUN). RFC 8489, February 2020.
<https://www.rfc-editor.org/info/rfc8489>.

[30] Pion WebRTC. <https://github.com/pion/webrtc>.

[31] Tirumaleswar Reddy, Alan Johnston, Philip Matthews, and Jonathan Rosenberg. Traversal Using Relays around NAT (TURN): Relay extensions to Session Traversal Utilities for NAT (STUN). RFC 8656, February 2020.
<https://www.rfc-editor.org/info/rfc8656>.

[32] Eric Rescorla. WebRTC security architecture. RFC 8827, January 2021.
<https://www.rfc-editor.org/info/rfc8827>.

[33] Eric Rescorla, Hannes Tschofenig, and Nagendra Modadugu. The Datagram Transport Layer Security (DTLS) protocol version 1.3. RFC 9147, April 2022.
<https://www.rfc-editor.org/info/rfc9147>.

[34] skywind3000. KCP - A fast and reliable ARQ protocol, January 2020. <https://github.com/skywind3000/kcp/blob/1.7/README.en.md>.

[35] Ram Sundara Raman, Prerana Shenoy, Katharina Kohls, and Roya Ensafi. Censored Planet: An Internet-wide, longitudinal censorship observatory. In *Computer and Communications Security*. ACM, 2020.
<https://censoredplanet.org/censoredplanet>.

[36] Michael Carl Tschantz, Sadia Afroz, Anonymous, and Vern Paxson. SoK: Towards grounding censorship circumvention in empiricism. In *Symposium on Security & Privacy*. IEEE, 2016. <https://internet-freedom-science.org/circumvention-survey/>.

[37] Zeya Umayya, Dhruv Malik, Devashish Gosain, and Piyush Kumar Sharma. PTPerf: On the performance evaluation of Tor pluggable transports. In *Internet Measurement Conference*. ACM, 2023.
<https://ptperf.github.io/>.

[38] uProxy. <https://www.uproxy.org/>.

[39] uProxy v1.2.5 - design doc. Archived at <https://archive.org/details/uProxy-Design-Doc-v1.2.5>.
https://docs.google.com/document/d/1t_30vX7RcrEGuWwcg0Jub-HiNI0Ko3kB0yqXgrQN3Kw [cited 2023-02-25].

[40] João Afonso Vilalonga, João S. Resende, and Henrique Domingos. TorKameleon: Improving Tor’s censorship resistance with K-anonymization and media-based covert channels. *CoRR*, abs/2303.17544, 2023.
<https://arxiv.org/abs/2303.17544>.

- [41] Ryan Wails, George Arnold Sullivan, Micah Sherr, and Rob Jansen. On precisely detecting censorship circumvention in real-world networks. In *Network and Distributed System Security Symposium*. The Internet Society, 2024. <https://www.robgjansen.com/publications/precisedetect-ndss2024.html>.
- [42] Yibo Xie, Gaopeng Gou, Gang Xiong, Zhen Li, and Mingxin Cui. Covertness analysis of Snowflake proxy request. In *International Conference on Computer Supported Cooperative Work in Design*. IEEE, 2023. <https://ieeexplore.ieee.org/document/10152736>.
- [43] xtaci. smux, February 2023. <https://github.com/xtaci/smux>.
- [44] Maria Xynou and Arturo Filastò. Russia started blocking Tor, December 2021. <https://ooni.org/post/2021-russia-blocks-tor/>.