

GFWeb: Measuring the Great Firewall’s Web Censorship at Scale

Nguyen Phong Hoang^{*†} Jakub Dalek[‡] Masashi Crete-Nishihata[‡]
Nicolas Christin[¶] Vinod Yegneswaran[§] Michalis Polychronakis^{*} Nick Feamster[†]

^{*} University of British Columbia [†] University of Chicago [‡] Citizen Lab - University of Toronto
[¶] Carnegie Mellon University [§] SRI International ^{*} Stony Brook University

Abstract

Censorship systems such as the Great Firewall (GFW) have been continuously refined to enhance their filtering capabilities. However, most prior studies, and in particular the GFW, have been limited in scope and conducted over short time periods, leading to gaps in our understanding of the GFW’s evolving Web censorship mechanisms over time. We introduce GFWeb, a novel system designed to discover domain blocklists used by the GFW for censoring Web access. GFWeb exploits GFW’s bidirectional and loss-tolerant blocking behavior to enable testing hundreds of millions of domains on a monthly basis, thereby facilitating large-scale longitudinal measurement of HTTP and HTTPS blocking mechanisms.

Over the course of 20 months, GFWeb has tested a total of 1.02 billion domains, and detected 943K and 55K pay-level domains censored by the GFW’s HTTP and HTTPS filters, respectively. To the best of our knowledge, our study represents the most extensive set of domains censored by the GFW ever discovered to date, many of which have never been detected by prior systems. Analyzing the longitudinal dataset collected by GFWeb, we observe that the GFW has been upgraded to mitigate several issues previously identified by the research community, including overblocking and failure in reassembling fragmented packets. More importantly, we discover that the GFW’s bidirectional blocking is not symmetric as previously thought, i.e., it can only be triggered by certain domains when probed from inside the country. We discuss the implications of our work on existing censorship measurement and circumvention efforts. We hope insights gained from our study can help inform future research, especially in monitoring censorship and developing new evasion tools.

1 Introduction

The Great Firewall of China (GFW) is one of most sophisticated, well studied, yet opaque censorship apparatuses to date [39]. It has been reported to employ a complex array of filtering mechanisms, including DNS poisoning [11, 12, 32, 37,

47, 54, 71], keyword-based filtering [25, 40, 66, 83], TCP/IP blocking [25, 44, 60, 76, 79], and active probing against network relays [9, 33, 36, 78]. As the GFW continues to evolve, using increasingly advanced techniques [13, 22, 29, 31, 79], prior research efforts have not been able to fully characterize the capabilities and evolution of this dynamic censorship platform. In particular, most measurement efforts to date have been limited in scope, time period, or methodology.

To monitor censorship around the world, global measurement platforms like OONI [38], ICLab [56], and Censored Planet [64] rely on volunteers, commercial VPNs, or public servers to conduct network measurements. Each platform, however, has its own blind spots due to the trade-offs between cost, detail, and breadth of coverage (§2.2). Among prior studies focusing on Internet censorship in China, GFWatch [47] stands out as one of the largest longitudinal measurements of the GFW. However, it is singularly focused on DNS blocking, leaving the GFW’s HTTP and HTTPS filtering policies largely unexplored. These two protocols, together with DNS, form the foundation of all Web communications. As a result, the collective filtering of these protocols plays a critical role in the overall operation of the GFW, and needs to be jointly studied to gain a more comprehensive understanding of the GFW’s filtering capabilities. Unfortunately, no prior work has been able to conduct large-scale, longitudinal measurement of the GFW’s HTTP and HTTPS filtering mechanisms.

To address these gaps, we have developed GFWeb, a large-scale system designed to uncover the GFW’s different blocklists used for Web censorship (§5). GFWeb is capable of continuously testing hundreds of millions of domains on a monthly basis, profiling the censorship behavior of both HTTP and HTTPS filters. To achieve this scale, we have to overcome several challenges, that would otherwise prevent us from scaling up our measurement and obtaining a holistic view of the GFW’s Web censorship. Specifically, our primary technical contributions lie in the design of a novel approach that strategically exploits the loss-tolerant and bidirectional filtering behavior of the GFW to overcome its stateful blocking, residual censorship, and asymmetric interference (§2.2).

Over a period of 20 months, from February 2022 to September 2023, GFWeb tested over one billion fully qualified domains (FQDNs), detecting 943K and 55K pay-level domains (PLDs) censored by the GFW’s HTTP and HTTPS filters, respectively (§5.1). To the best of our knowledge, our findings represent the largest set of censored domains ever discovered and no prior study of the GFW matches this scale of active measurement over an extended period of time.

Our study reveals new insights into the GFW’s evolving Web censorship policies. We discover that the GFW has fixed deficiencies identified by past research, including overblocking [47] and failures to reassemble out-of-order fragmented packets [20]. More significantly, we find its bidirectional blocking is not symmetric as previously thought [47, 72]. Many domains trigger blocking only when probed from within China but not externally (§5.3). We also observe localized network interference which is not caused by the GFW but individual ISPs and cloud providers (§6). These observations underscore the importance of conducting large-scale, longitudinal measurement from both sides of the GFW to better understand China’s complex censorship environment.

Our findings have several implications for Internet freedom efforts (§7), and we are thus committed to keep operating GFWeb and will continue to update the community with GFW’s domain blocklists and censorship behaviors. By sharing our findings, we hope to inform future research that advances the state of censorship measurement and inspires the development of novel censorship circumvention systems.

2 Background and Motivation

We first review the GFW’s Web filtering techniques. We then delve into the challenges encountered by previous efforts and how they have motivated us to design GFWeb.

2.1 Web Filtering Mechanisms

China’s Internet filtering framework, conceptualized in the late 1990s as part of the Golden Shield project [74, 82], serves as the government’s tool for controlling the flow of online information. Often dubbed “the Great Firewall” (GFW) [40], this system comprises middleboxes distributed across border autonomous systems (ASes) [12, 27, 78] and managed in a centralized manner [29, 40, 83]. While the GFW employs several filtering mechanisms (§1), here we focus our discussion on DNS, HTTP, and HTTPS filtering, since these protocols play a foundational role in Web access.

2.1.1 DNS filtering

The unencrypted and unauthenticated nature of traditional DNS resolution is often targeted by censors around the world [61, 69]. The GFW operates as an on-path system and exploits the race condition of UDP-based DNS resolution

by injecting false responses when it detects DNS queries for censored domains. As shown in Figure 1(a), the GFW injects a forged response towards the client upon detecting a DNS query for a restricted domain (e.g., `hrw.org`). Operating as an on-path system, the GFW does not discard the genuine response from the DNS resolver. However, since the GFW is typically closer to the client in terms of network distance, its fake responses often reach the client before the legitimate one [32, 47], effectively tampering with the DNS resolution process. Since almost every Web browsing session begins with a DNS query, this DNS blocking mechanism is the first filter of the GFW’s multi-stage Web censorship operation.

2.1.2 HTTP and HTTPS filtering

When there is no DNS blocking or if a client manages to circumvent it [32, 42, 47, 50], the GFW’s HTTP and HTTPS filters are the next layers of blocking. The GFW’s HTTP and HTTPS filtering mechanisms involve stateful inspection of TCP connections. The GFW begins keeping track of a TCP connection’s state when it sees the first SYN packet from the client sent to initiate the TCP three-way handshake. Upon detecting a censored domain in the HTTP Host header or the Server Name Indication (SNI) extension in the TLS Client Hello, the GFW then tears down the connection by injecting three RST/ACK packets to both the client and the server, as shown in Figures 1(b) and 1(c). Note that URLs explicitly prefixed with `https://` will not trigger the HTTP filter, though most users do not often type this protocol prefix.

2.1.3 Bidirectional filtering

A unique behavior of the GFW is its bidirectional blocking. Its filtering middleboxes are capable of inspecting and injecting traffic in both directions. As a result, not only clients inside the country experience censorship, but traffic sent from outside the country can also trigger censorship. This bidirectional behavior has become a common characteristic often used by the research community to measure the GFW [11, 12, 33, 35, 47, 55, 78] and other similar censorship regimes [18, 57, 58, 75]. We also leverage this behavior to design GFWeb and conduct large-scale measurement of the GFW’s Web censorship against both HTTP and HTTPS connections (§4).

2.2 Existing Efforts and Challenges

Due to its reputation as being one of the most sophisticated censorship systems, impacting not just hundreds of millions of Internet users inside China but also the normal operation of the global Internet [47, 55, 71], the GFW has been the subject of numerous studies over the past two decades. Unfortunately, many of them have been limited in scope or conducted in a one-off manner, leaving gaps in our understanding of the GFW’s Web filtering mechanisms over time.

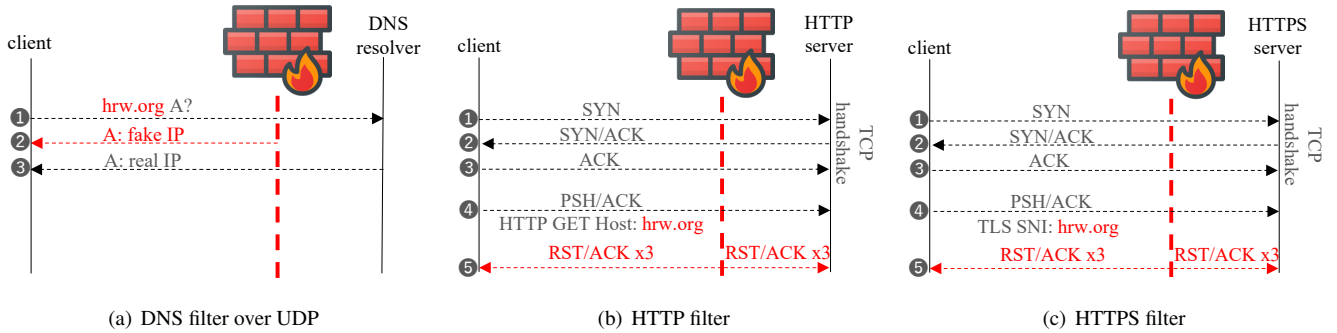


Figure 1: The DNS, HTTP, and HTTPS filtering mechanisms of the Great Firewall. SYN, ACK, PSH, and RST denote TCP synchronization, acknowledgement, push, and reset flags. A packet with the RST flag set is meant to terminate a TCP connection.

GFWatch [47] is the largest active measurement system to date focusing exclusively on the GFW’s DNS blocking. No study has longitudinally and systematically investigated the GFW’s HTTP and HTTPS filters, which are equally important—in conjunction with DNS blocking, they form the multiple layers of the GFW’s Web censorship capabilities. Understandably, probing DNS manipulation with UDP packets is much easier compared to measuring HTTP and HTTPS interference. This is because UDP-based DNS resolution is stateless and does not require a handshake process, allowing GFWatch to probe a massive amount of domains in a short period of time. On the other hand, HTTP(S) censorship, as shown in Figures 1(b) and 1(c), is stateful and necessitates several TCP packets to trigger the GFW’s HTTP and HTTPS filters, making it more challenging to measure at scale.

More importantly, unlike the UDP-based DNS censorship which does not trigger residual censorship, once a TCP stream is interrupted, subsequent connections from the same client to the same server will also be blocked for a period of time, preventing continuous probing against HTTP and HTTPS filters. The GFW’s TCP-based censorship has been observed to impose residual censorship on censored connections [17, 25, 66, 81]. This residual censorship is one of the main challenges that we have to tackle to conduct large-scale measurement of the GFW’s HTTP and HTTPS filters.

Although the GFW might not be the main focus, there are platforms that are still actively operating to monitor global censorship, such as OONI [38], ICLab [56], and Censored Planet [64]. OONI [38] relies on volunteers to run accessibility tests. As of July 2023, however, OONI has been blocked in China, hindering regular measurement activities from within the country. More specifically, the platform has been blocked primarily by DNS injection against the main site and HTTPS filtering based on the SNI field of their data collection servers [70].

To minimize risks on volunteers, ICLab [56] uses commercial VPNs as its primary measurement vantage points.

Yet, this method restricts its visibility into China, given the country’s prohibitions on commercial VPNs [15, 23].

Censored Planet [64] employs remote measurement techniques to infer censorship by making use of public servers. Nonetheless, to avoid overburdening these machines the number of domains that can be frequently tested is limited, especially since these remote servers are not owned by the platform. Measuring from outside China also limits the platform’s visibility into censored domains that will only trigger the GFW if tested from inside, due to asymmetric interference (§5.3).

A common design choice of these platforms is the reliance on the Citizen Lab test lists for test domains [4], limiting their measurement scope to only a few thousand domains. Collectively, these challenges and pitfalls of prior works have motivated us to design GFWeb and carry out this study to contribute to the Internet freedom community’s efforts and close the gap in our understanding of the GFW’s Web censorship.

3 Measurement Approach

Considering the challenges faced by prior efforts, we designed GFWeb with the following goals in mind: (1) our system should be able to conduct large-scale measurement of the GFW’s HTTP and HTTPS filtering mechanisms to discover as many censored domains as possible within a reasonable amount of time to provide timely updates to the community; (2) the system should be designed to continuously monitor the GFW’s Web filtering over time to gain a holistic view into its evolving Web censorship; and (3) to avoid posing any potential risks to volunteers and public servers, our system should not rely on these entities for data collection. In the following, we discuss how we overcome the challenges discussed in §2.

3.1 Reducing Cost of TCP Probing

Our first goal is to discover as many censored domains as possible and in a continuous manner. Unlike GFWatch [47],

which measures DNS censorship by simply sending UDP-based queries over port 53 from one side of the GFW to the other, TCP-based HTTP(S) probes are more expensive in terms of the required resources for several reasons, including the need for more than just one packet to establish a connection with an endpoint on the other side of the GFW.

We exploit the GFW’s loss-tolerant blocking behavior to tackle this challenge. This loss-tolerant design is likely needed to cope with (1) the potential high packet loss rate due to the sheer amount of egress and ingress traffic from/to China the filtering middleboxes need to inspect, and (2) asymmetric routing that can cause packets to be routed through different network paths, i.e., packets ② and ③ may not be routed through the same path as ① in Figures 1(b) and 1(c). Consequently, as long as the GFW can observe the initial SYN of a TCP connection (packet ①), it will consider that connection as established and will inject three RST/ACK (packets ⑤) upon detecting a censored domain in the HTTP Host header or the SNI extension of the TLS Client Hello (packet ④).

Taking advantage of this loss-tolerant behavior, the cost of probing HTTP(S) middleboxes can be significantly reduced, as we can send a SYN packet (①) followed by a PSH/ACK packet (④) encapsulating the HTTP GET request or the TLS Client Hello for a domain of interest. Crucially, the PSH/ACK packet’s sequence number *must be* equal to the SYN packet’s sequence number plus one, to properly mimic the behavior of a real completed TCP connection. Without satisfying this condition, the GFW will not consider the connection as established, and will not inject RST/ACK packets even if the request payload contains a censored domain. This approach allows us to efficiently probe many domains by continuously sending SYN and PSH/ACK packet pairs to the other side of the GFW to trigger its HTTP(S) filters without having to wait for the three-way handshake to complete.

3.2 Bypassing Residual Censorship

Despite the above optimization, we still need to tackle another challenge, namely the GFW’s residual censorship, a phenomenon that has not been fully characterized by the research community (we provide more details in §5.4). Previous reports show that once an HTTP connection is interrupted, subsequent packets with the same three-tuple will also be tampered with for a period of 90 seconds [17, 25, 66, 76, 81]. Any SYN packets sent during this period will trigger a forged SYN/ACK while other packets will be responded with RST or RST/ACK [76]. Consequently, this behavior can cause false positives in our measurement because subsequent benign probes would still trigger the filtering middleboxes.

A key mechanism of the GFW’s residual censorship is that it is based on the three-tuple of the banned TCP connection, i.e., the source IP address, the destination IP address, and the destination port number. This means that if we can change any of these three fields, the residual censorship can be bypassed.

Different from the DNS filter, which only inspects packets sent to the default DNS port 53, the HTTP(S) filters are flexible and inspect packets destined to any TCP port number, instead of just the standard HTTP(S) ports 80 and 443. From a censor’s perspective, this policy is necessary for coping with websites hosted on non-standard ports or censored users deliberately changing the destination port to bypass censorship. However, this flexibility also opens up the possibility for us to bypass the residual censorship by changing the destination port number of the TCP connection. In particular, we can send new HTTP(S) probes with a different destination port number than those previously used, while ensuring that GFWeb does not reuse the same port number if 90 seconds have not elapsed. Using this approach, we can bypass the residual censorship and continuously probe the GFW’s HTTP(S) filters.

3.3 Accounting for the Impact of Ephemeral Ports on Censorship Behavior

Bhashkar et al. [14] find that ephemeral source ports can affect the packet routing, which in turn can affect the GFW’s injection behavior. Together with the utilization of different destination port numbers to sidestep residual censorship, we also rotate the ephemeral source port of each probe to account for this potential impact. Naturally, this also becomes a tagging mechanism for us to keep track and later identify which domains trigger the GFW’s HTTP(S) filters. In particular, every probe is uniquely tagged using a different combination of source and destination port numbers, virtually creating a one-to-one mapping between a tested domain and each pair of source and destination ports. We can then easily identify which exact domains are censored by checking the source and destination ports of the respective RST/ACK packets.

3.4 Avoiding Risks Posed on Volunteers and Public Servers

For our study, we opt to use a set of dedicated servers we control, located at both sides of the GFW, to avoid posing any potential risks to end users. Our machines outside China are located in an educational network where we have confirmed that there is no censorship. For machines inside China, we use servers located in two different ASes and at two distant geographical locations, allowing us to probe the GFW via different network paths. These machines are in the data centers of two major cloud providers: Aliyun and QCloud.

Although previous studies have shown the centralized blocking policy of the GFW [29, 40, 83], we still decided to use multiple servers to ensure that our findings are not limited and biased to a specific location. In fact, this decision has been proven to be beneficial, as we show later in §6, as GFWeb was also able to capture some localized network interference events that were not thoroughly investigated by prior works, but can still be of interest to the community.

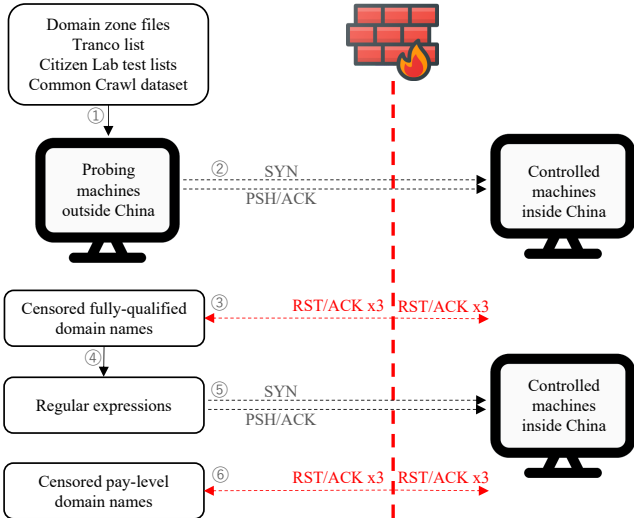


Figure 2: The design of GFWeb’s probing approach, which does not rely on the completion of the three-way handshake for triggering the GFW’s HTTP(S) filters.

Table 1: The number of test domains from each source.

Source	# Domains
ICANN’s TLD Zone Files [2]	389.5M SLDs
Common Crawl’s Web Crawl Data [1]	903.4M FQDNs
Tranco List [62]	33.4M FQDNs
The Citizen Lab Test Lists [4]	25.9K FQDNs

4 GFWeb Measurement Pipeline

We next describe the measurement pipeline of GFWeb, the test lists used, and the rationale behind our design choice.

4.1 Domain Test Lists

Similar to GFWatch [47], we use the same domain test lists used by this prior work to conduct our measurement, which also makes our results more comparable. These domain test lists are collected from various sources, including top-level domains (TLD) zone files [2], the Citizen Lab test lists [4], the Tranco list [62], and the Common Crawl project [1]. For every new probing batch, we download the latest version of each list to make sure that GFWeb tests up-to-date domains.

The rationale behind this choice of domain lists is their direct impact on the set of censored domains that will be discovered, as well as on subsequent analyses. Different domain test lists have different characteristics, and the censored domains discovered will be a subset of the domains tested. To avoid any potential bias introduced by individual test lists, we use multiple domain sources in our study to ensure that (1) we

can cover as many domains as possible, and (2) our findings are holistic and not limited to a specific set of domains.

From February 2022 to September 2023, GFWeb has tested an average of 600M FQDNs/month, totaling 1.02 billion FQDNs over the course of 20 months considered in this paper. These domains span 1.5K TLDs¹ and 694M PLDs.² Table 1 summarizes the number of domains from each source.

4.2 Overall Architecture

Figure 2 illustrates the overall architecture of GFWeb and our measurement pipeline. We start with the curation of the test lists (step ①). Then, our probing machines located outside China send SYN and PSH/ACK packet pairs to the other side of the GFW (step ②), with each probe uniquely tagged by a different pair of source and destination port numbers. The payload of the PSH/ACK packet contains the HTTP GET request or the TLS Client Hello for the domain being tested.

Our machines inside China are configured to not respond to any packet sent from the probing machines. This ensures that the observed injected packets are indeed from middleboxes, allowing us to infer which domains are censored (step ③). This configuration is essential for avoiding false positives, because a typical server will respond with a RST or RST/ACK packet to any non-RST TCP packet sent to an open port that does not have a listening service. This is another benefit of using our own servers to conduct the measurement, as we can easily configure them to suit the purpose of our study, eliminating risks posed on volunteers and public servers, while also preventing false positives from tainting our results.

In our previous work [47], we showed evidence of the GFW’s DNS blocking based on regular expressions. Inspired by this finding, for every censored domain discovered in step ③, we also attempt to reverse engineer its blocking pattern by creating eight different permutations by splitting the domain into different substrings and concatenating them with a random string (step ④). This set of permutations is then also probed against the GFW (step ⑤). Finally, by analyzing injections from the GFW’s HTTP(S) filters (step ⑤), we can infer the blocking pattern of the censored domains (step ⑥).

This overall measurement approach is repeated in the other direction from our machines inside China, while discovered censored domains are also re-tested every day to keep track of their blocking status over time. To account for uncontrollable factors, such as packet loss or temporary failures of the GFW when it is overloaded [11, 36], each domain is probed at three different times during the day using different source and destination port number pairs.

¹This number is larger than the number of TLDs currently reported by ICANN [2] because some TLDs that used to be active at the time of our measurement are no longer in use as of this writing.

²We determine pay-level domains (PLDs) using the public suffix list [3].




Figure 3: Number of HTTP(S) censored domains (in log scale) detected over time by GFWeb compared to GFWatch [47].

5 Measurement Results

In this section, we present the results of our measurements and discuss in detail the differences between the observed censorship behaviors of the GFW across the three considered protocols (DNS, HTTP, and HTTPS).

5.1 Censored Domains Over Time

As shown in Figure 3(a), on a monthly basis, GFWeb detects an average of 2.4M and 243K FQDNs censored by the GFW’s HTTP and HTTPS filters, respectively, whereas GFWatch [47] detects about 1.06M FQDNs censored by the DNS filter.³

There is a spike in the number of FQDNs censored by the HTTP filter in May 2022, peaking at more than 6M censored FQDNs. Investigating the cause of this spike, we found that it is due to the addition of many new subdomains of domains that are already censored by the GFW (e.g., *.blog.jp). Our decision to include domains from the Common Crawl dataset [1] is the primary reason behind this spike, as the dataset is regularly updated with crawls containing many new subdomains of PLDs that are already censored. The same spike can also be seen in the GFWatch dataset.

To more accurately quantify the GFW’s censored domains while making our results comparable to GFWatch’s, we analyze the blocking rules of the censored FQDNs by reverse-engineering the regular expressions used by the GFW to identify the base censored domains. A base censored domain is the shortest domain name that matches the blocking rules of the GFW. For example, the base censored domain of en.wikipedia.org and zh.wikipedia.org is wikipedia.org. As shown in Figure 3(b), the number of base censored domains is indeed much smaller than the number of censored FQDNs. In particular, we find an average of 528K and 24K base censored domains per month for HTTP

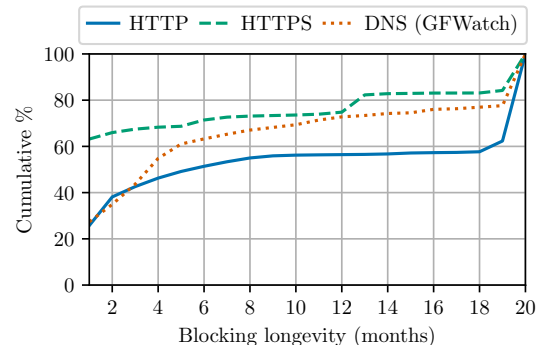


Figure 4: CDF of the longevity of domains censored by the GFW’s DNS, HTTP, and HTTPS filters.

and HTTPS filters, respectively. GFWatch discovers an average of 247K base censored domains per month.

While the number of base domains censored by the HTTP filter is relatively stable, with a slight decrease in 2023, the patterns of the number of base domains blocked by the DNS and HTTPS filters are similar over time, with a slight increase in 2023. As more Web traffic gets centralized and encrypted [34, 45, 46, 48, 49] and major browsers are deprecating HTTP in favor of HTTPS (e.g., Chromium announced in August 2023 that it will automatically upgrade all HTTP connections to HTTPS [24] even if a user types in an HTTP URL), we anticipate that the HTTPS filter will block a constantly increasing number of domains in the near future.

Furthermore, our analysis of blocking longevity reveals a distinct persistence in censorship across protocols. Different from the HTTPS filter, both HTTP and DNS filters demonstrate a propensity to maintain blocks over extended periods, with over 50% of their censored domains experiencing restrictions for at least three months, as shown in Figure 4. This endurance in blocking strategies signifies a nuanced approach to Web censorship, emphasizing the need for continuous monitoring systems like GFWeb.

³We obtained the set of censored domains due to DNS filtering for the same time period from GFWatch’s public dashboard [43].

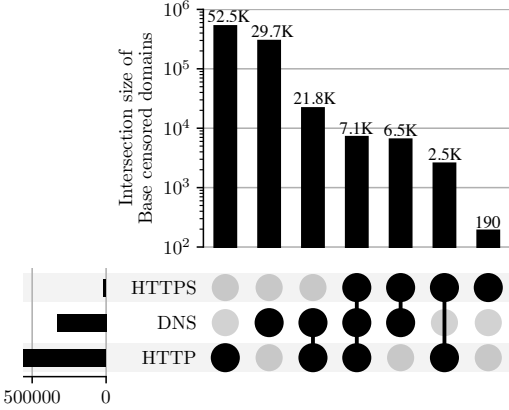


Figure 5: Intersections between the blocklists of the GFW’s DNS, HTTP, and HTTPS filters, with y-axis in log scale.

Overblocking In our previous work [47] we observed that the GFW’s DNS filter overblocks many domains that are totally unrelated to the base censored domains, causing collateral damage to tens of thousands of innocuous websites that happen to be blocked in China. For instance, the Tor Project’s website (`torproject.org`) is blocked by the GFW’s DNS filter using the regular expression `*torproject.org`, which also matches domains like `mentorproject.org` or `ventilatorproject.org`.

We observe that this overblocking behavior is now less prevalent across all three protocols. In particular, we find that the GFW has fixed the overblocking issue of the DNS filter. Nine out of the ten most overblocked rules reported in [47] are now correctly implemented with an additional dot (`\.`) character at the beginning of the regular expressions. Measurement results from GFWeb also indicate that the vast majority of regular expressions used by the GFW’s HTTP and HTTPS filters are also more specific and properly implemented to block only subdomains and country code top-level domains (ccTLDs) of the base censored domains. While a few domains are still being overblocked by the HTTP filter due to overly broad regular expressions (e.g., `*archive.today`), overall our measurement shows that the previous issue with overblocking has been largely mitigated, significantly reducing the collateral damage caused to unrelated domain names.

5.2 Cross-Protocol Comparison

Comparing base censored domains across all three protocols, we find that the HTTP filter has the largest blocklist, followed by the DNS and HTTPS filters. We illustrate the intersections between the blocklists of base censored domains across the three protocols in the UpSet diagram [53] of Figure 5. To prevent short-lived censorship events in which a domain is blocked for only a couple of days from skewing our results, for

this analysis we consider a domain to be blocked by a given filter if it is blocked by that filter for at least three months.

Figure 5 shows that there are around 7.1K base censored domains blocked across all three protocols. There are more than 21.8K domains that are blocked by the HTTP and DNS filters but not by the HTTPS filter, while 2.5K domains are blocked by the HTTP and HTTPS filters but not by the DNS filter. Less than a couple of hundred domains are blocked solely by the HTTPS filter. This result shows that the three filters operate on different blocklists that are not mutually exclusive and complement each other to form the core of the GFW’s Web censorship. A regular Web browsing session of a user from inside China visiting a foreign website, without any usage of circumvention tools, will be inspected and interfered with by either one or a combination of these filters if the visited domain is censored.

While an official report of the GFW’s design has never been published, the differences in the DNS, HTTP, and HTTPS blocklists can be attributed to the technical nature of each protocol, implementation cost, and strategic implications.

HTTP traffic is unencrypted, allowing for granular content-based filtering using Deep Packet Inspection (DPI). The GFW can inspect the entire content of HTTP requests and responses, allowing it to block individual pages even based on specific keywords from the URL [66]. This fine-grained control means that there could be many specific rules and block criteria, leading to a larger blocklist. Another potential explanation is that the HTTP filter has been around for a longer time, when HTTPS was not as prevalent as it is today [8]. The HTTP filter, thus, has had more time to accumulate a larger blocklist.

The DNS filter operates at the forefront of the GFW, allowing it to tamper with a browsing session even before any traffic could be sent to the actual web server if the visited domain is blocked. From an operational perspective, DNS blocking, although easily bypassed [32, 47, 50], is the cheapest to implement, as it only requires an on-path filtering middlebox to inspect DNS queries and forge DNS responses.

Among the three protocols, HTTPS blocking is arguably the most expensive to implement. Due to the encrypted nature of HTTPS, DPI becomes less effective because only the unencrypted portion of the TLS handshake can be inspected. This means that the GFW cannot inspect the content of HTTPS requests and responses, and instead has to rely on the Server Name Indication (SNI) field in the TLS handshake to identify the domain of the visited website. Moreover, a TLS Client Hello packet generally has more fields compared to an HTTP request header or a DNS query, making HTTPS blocking more operationally expensive than the other two protocols.

Collectively, the blocklist sizes reflect a balance between the technical constraints of each protocol, the cost of implementation, and the strategic goals of the GFW. As GFWeb continues to monitor the GFW’s Web censorship, we anticipate that it will continue to ramp up its blocking capabilities against secure protocols like HTTPS in the future. The finding

of different blocklists across the three protocols also has several implications for existing measurement and circumvention efforts, which we discuss in more detail in §7.

5.3 Asymmetric Interference

The GFW has been long believed to be a bidirectional filtering system that symmetrically interferes with both egress and ingress network connections. Our measurements from both sides of the GFW, however, reveal that this is not always the case. Specifically, there are certain domains that will trigger the filtering middleboxes to take injection actions only when the probing direction is from inside China, but not from outside. We refer to this behavior as *asymmetric interference*.

Comparing the sets of domains that trigger the GFW from both sides, we found about 1K domains that only trigger the HTTPS filter to inject RST packets when probed from inside the country. The domains mostly belong to Google (e.g., `google.com.hk`) or are related to circumvention tools (e.g., `torproject.org`, `go-vpn.com`, `dr-wall.com`, `aihuiguo.com`, and `wallvpn.com`). Probing these domains from outside China will not trigger any interference from the GFW’s HTTPS filter. This is also evident by the very low number of network anomalies detected by Censored Planet for these domains. As shown on the platform’s dashboard [7], most `Unexpected Rates` for these domains are below 20%, as opposed to $\sim 100\%$ for domains that are symmetrically interfered with by the GFW’s HTTPS filtering middleboxes.

This finding has very important implications on censorship monitoring systems that perform their measurements under the assumption that a censoring system is always bidirectional. The GFW’s HTTPS filter will not interfere with the probes for these domains that are sent from outside the country, potentially resulting in false negatives, i.e., inferring that they are not filtered while in reality they are.

5.4 Traffic Dropping as Residual Censorship

In addition to the asymmetric interference above, this group of domains also exhibit another interesting behavior: prolonged traffic dropping. In particular, probing these domains will trigger the GFW to *drop* subsequent packets that share the same three-tuple for an extended period of time. This behavior is different from the GFW’s penalty box observed by previous studies, where the GFW will continue to interfere with subsequent connections by injecting RST packets. Instead, the GFW will *keep dropping* any subsequent TCP packets that share the same three-tuple.

Unlike previous reports of the 90-second penalty box [17, 25, 66, 76, 81], which keeps injecting RST packets, we discover that this traffic dropping behavior will continue happening for up to 350 seconds for TCP packets that share the same three-tuple. Although this residual censorship is only

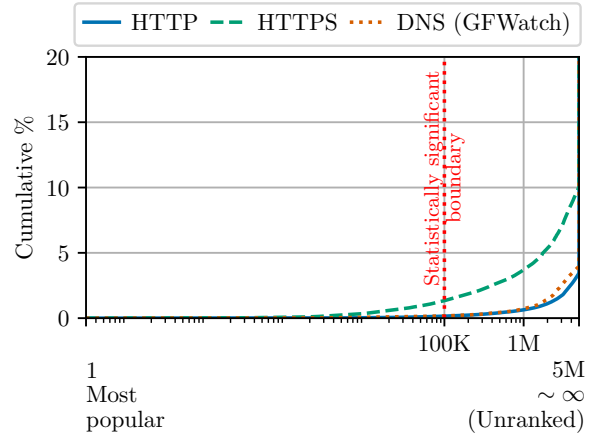


Figure 6: CDF of the popularity ranking of base censored domains (x-axis is in log scale).

imposed on a few domains, it can impact the straightforward use of probing-based evasion tools (§7).

5.5 Popularity Ranking of Censored Domains

Aggregating the ranking information provided by the Tranco list [62] and the Common Crawl dataset [1], we analyze the popularity of the base censored domains discovered by GFWeb for HTTP(S) filters in comparison to the GFWatch (DNS filter) dataset. As shown in Figure 6, the vast majority of the censored domains are not popular, with more than 98% of them having a popularity ranking higher than 100K across all DNS, HTTP, and HTTPS block lists. Less than 0.2% of the censored domains are among the top 100K for both the DNS and HTTPS block lists, and 1.3% for the HTTPS block list. Note the 100K-th position is the statistically significant boundary for website popularity ranking [68].

This result shows how different test lists (or subsets of a single list) can lead to different conclusions about which domains are censored and their popularity, especially when prior studies have often relied on relatively small test lists or just a few thousand domains from website top lists [38, 56, 64]. In contrast, to gain a more holistic view of the GFW’s Web censorship, we have tested as many domains as possible, curated from diverse sources (§4.1).

5.6 A Taxonomy of Censored Domains: Web Categories and Hosting Origins

For categorizing domains, we use VirusTotal’s classification service [5]. Of more than 1M based censored domains discovered, we could only categorize 79.5K domains because many domains no longer exist or do not currently host any content. Figure 7 shows the top ten categories of the base censored domain and their distribution. The categories with the most censored domains include `File sharing/storage`, `adult`




Figure 7: Top ten categories of the base censored domains discovered by GFWeb.

content, malicious, info.tech, and business, comprising 60% of all domains classified.

Ramping Up Censorship against AI-related Domains.

As Artificial Intelligence (AI) is gaining traction in recent years, China has been taking substantial steps in regulating the AI sector. On April 11, 2023, the Cyberspace Administration of China released draft measures for regulating generative AI services [28]. As a result, several popular AI tools and platforms used globally are restricted in China. We observed numerous prominent AI-related domains becoming censored during this time, including `chat.openai.com`, `deepai.org`, and `huggingface.co`. In addition, GFWeb also discovered a new wave of blocking against PLDs containing “GPT” as part of their name. Figure 8 shows the number of censored PLDs containing “GPT” in their domain names that GFWeb and GFWatch [47] detected since the beginning of 2023.

Shortly after the release of the regulations, many new AI services by domestic companies entered China’s market [77]. These events reflect China’s broader strategy to exercise state control over the development of AI technologies, and how monitoring the GFW can provide deeper insights into the country’s policy-making process. This case also highlights the role of China’s Internet censorship in economic protectionism, promoting domestic companies over foreign competitors.

Hosting Origin. We also attempt to identify the origins of the censored domains by looking up their hosting IP address(es) and mapping them to respective Autonomous System Numbers (ASNs) and organizations. The majority of the censored domains are hosted by cloud service providers from Western countries. Table 2 shows the top ten ASes responsible for hosting over 13K base censored domains each. An intriguing aspect of this analysis reveals disparities in the number of base censored domains among various cloud providers, contingent upon the type of censorship applied. More specifically, domains hosted on Google, Cloudflare, BGPNET Global, and




Figure 8: Number of censored PLDs containing “GPT” in their domain names detected by GFWeb and GFWatch [47] since the beginning of 2023.

Table 2: Top ten ASes that host the most base censored domains that are blocked by the GFW’s DNS, HTTP, and HTTPS filters. Note that the numbers are not mutually exclusive as a domain may be censored by more than one filter as discussed earlier in §5.2.

ASN	Organization	Total	DNS	HTTP	HTTPS
15169	Google	92069	81247	12767	2827
13335	Cloudflare	77298	45482	43037	9094
40065	CyberNet Servers	76354	22767	57829	2020
16509	AMAZON-02	65540	29608	41708	5015
396982	Google Cloud Platform	25110	13574	13493	1171
64050	BGPNET Global	22193	17896	5664	770
8075	Microsoft	14235	9702	5685	414
139646	HK Megalayer Tech	13431	7117	6633	206
14618	AMAZON-AES	13125	4820	9298	1272
133618	Trellian Pty. Limited	13039	6917	7960	1251

Microsoft are more targeted by the DNS filter, whereas those on CyberNet Servers, Amazon, and Trellian Pty. Limited experience more blocking by the HTTP filter. This discrepancy underscores the nuanced approach of the GFW towards censorship across different protocols and highlights the critical need for comprehensive measurement tools such as GFWeb to fully understand and document these censorship patterns.

5.7 Detail Comparison with Prior Work

To better illustrate (1) the differences between GFWeb and prior censorship measurement platforms and (2) how our work complements existing efforts, we compare the number of tested and censored PLDs discovered by GFWeb with OONI [38] and Censored Planet [64], two measurement platforms that are actively measuring global Internet censorship, including China’s HTTP and HTTPS filtering.

For the OONI dataset, we obtain Web connectivity tests conducted in China during the same period of time as our measurement from OONI Explorer [59]. For Censored Planet, we obtain the test results from the project’s dashboard [6].

Table 3: The number of tested and censored pay-level domains (PLDs) GFWeb has discovered in comparison to current censorship measurement platforms that are still actively measuring China’s HTTP and HTTPS filtering. GFWatch is not included in this table as it only measures DNS filtering and has been compared in more detail earlier in this section.

Platform	Direction	Duration	Tested PLDs		Censored PLDs		Common censored PLDs	
			HTTP	HTTPS	HTTP	HTTPS	HTTP	HTTPS
OONI [38]	inside	2022/02 - 2023/09	13K	85K	275	211	212 (77%)	151 (71.5%)
Censored Planet [64]	outside	2023/03 - 2023/09	2K		1K	220	512 (50%)	220 (100%)
GFWeb	both sides	2022/02 - 2023/09	694M		943K	55K	N/A	

Note that we could only obtain data from March 2023 to September 2023 from the project’s dashboard at the time of conducting this comparison. As shown in Table 3, GFWeb has tested far more domains and also discovered a much larger number of censored PLDs than the other two platforms.

Unlike GFWeb, OONI [38] does not design dedicated tests to probe the GFW’s HTTP and HTTPS filters separately. Instead, OONI’s Web connectivity tests are designed to replicate the normal browsing session of a user by sequentially (1) conducting DNS lookup, (2) connecting to the web server using the IP(s) returned, and (3) sending an HTTP(S) request to the web server depending on whatever the protocol of the URL provided by the Citizen Lab’s test lists [4]. Thus, to determine the number of domains tested against the GFW’s HTTP and HTTPS filters, we utilize the protocol prefix (i.e., `http://` or `https://`) of the URLs tested to count the number of HTTP and HTTPS domains. Over the same period of time, OONI has tested 13K and 85K HTTP and HTTPS PLDs, respectively.

Like GFWeb, Censored Planet [64] conducts independent tests for HTTP and HTTPS. From March to September 2023, Censored Planet has tested about 2K PLDs for both protocols. The discrepancy between the number of domains tested by Censored Planet and OONI despite both platforms relying on the Citizen Lab’s test lists [4] is due to Censored Planet’s design choice of focusing more on testing the China-specific list while OONI’s volunteers are free to choose which domains they would like to test.

Overall, GFWeb largely agrees with the other two platforms in terms of the number of censored domains that are commonly detected. As shown in Table 3, GFWeb and OONI agree on more than 70% of the censored domains detected by both platforms. GFWeb also agrees with Censored Planet on 100% of the domains censored by the HTTPS filter. Analyzing the remaining discrepancies, we find that the majority of them are due to localized network interference events but not the GFW’s HTTP and HTTPS filters, as we will show in §6.

6 Other Cases of Network Interference

While the ultimate goal of our study is to investigate the GFW’s Web censorship, we encounter some interesting cases

where localized network interference events could also be observed from major cloud providers and ISPs in China.

6.1 Hosting Providers’ Redirection

Our measurement machines in China are located in the data centers of two major cloud providers: Aliyun and QCloud. While analyzing network traffic collected by GFWeb, we found that both providers have deployed DPI middleboxes to interfere with HTTP connections and attempt to redirect users to warning pages when certain domains are requested despite the location of the request client (i.e., inside or outside China). Specifically, Aliyun injects a redirection to `batit.aliyun.com/alww.html` and QCloud injects a redirection to `dnspod.qcloud.com/static/webblock.html` as shown in Figure 9. These pages warn that the domain being requested is not registered with the Chinese government. For a domain to be hosted from within China, it needs to be registered with the Chinese Ministry of Industry and Information Technology (MIIT) and obtain an Internet Content Provider (ICP) license.

Both providers’ DPI middleboxes are deployed as on-path devices and intercept packets in a stateless manner. In other words, these redirection injections can be triggered without initiating a complete TCP handshake (i.e., packets ❶-❸ in Figure 1(b)). An HTTP PSH/ACK packet (❹) with the payload of a trigger domain is sufficient to prompt the middleboxes to inject a redirection towards the side of the connection that sends the PSH/ACK packet. Bock et al. [16] has recently found that this behavior could be weaponized for TCP-based amplification attacks since the injected warning pages are much larger compared to the original HTTP request.

Over the course of our study, Aliyun and QCloud middleboxes have interfered with 36.5M and 39.1M FQDNs, respectively. Clustering these domains by their PLDs, we find a total of 7.8M and 6.8M PLDs whose HTTP requests were redirected to Aliyun and QCloud ICP warning pages. It’s worth noting that this interference is not caused by the GFW and should not be considered as nation-state censorship. Instead, they are caused by the hosting providers to enforce the ICP license requirements since a benign domain like `example.com` also triggers the redirection of QCloud’s middleboxes.

温馨提示：该网站暂时无法进行访问

原因一：根据工信部相关法规，您尚未进行备案；

原因二：根据工信部相关法规，您当前的接入商不是阿里云或万网；


原因三：您的网站可能存在不适宜传播的信息，请联系网站管理员。

本页面为默认提示页面，如网站存在以上问题请及时进行处理。

用户备案请登录阿里云代备案管理系统；

谢谢合作！

(a) Aliyun redirection to `batit.aliyun.com/alww.html`



网站暂时无法访问

该网站未根据工信部相关法律法规在腾讯云进行备案。

国务院令292号《互联网信息服务管理办法》和《非经营性互联网信息服务备案管理办法》规定，所有对中国大陆境内提供服务的网站都必须先进行ICP备案。

(b) QCloud redirection to `dnspod.qcloud.com/static/webblock.html`

Figure 9: Redirection injected by Aliyun and QCloud to warn about ICP license requirements.

Figure 10: China’s “Anti-Fraud” redirection.

6.2 China’s “Anti-Fraud” Redirection

In late 2021, there were some reports of new network interference events across major state-owned ISPs [26, 73], including three largest operators: China Telecom, China Unicom, and China Mobile [51]. Many users reported that their browsing sessions were redirected to a warning page showing an “Anti-Fraud” message. As indicated in Figure 10, the warning page advises users that the site they are trying to access is suspected of fraud and asks them to install an app developed by the State Anti-Fraud Center from the Android or the Apple app stores.

GFWeb has also observed this network interference. Applying the limited time-to-live (TTL) probing approach used in prior works [11, 65], we confirm that it is indeed caused by the ISPs. While we did not have direct access to vantage points within these ISPs, one of our measurement machines in China is located in a data center whose upstream provider is China Telecom, allowing us to observe the redirection injected by this ISP’s middleboxes. Similarly to the GFW, they can inspect and inject packets bidirectionally. They are also deployed as on-path devices and have a loss-tolerant design (i.e., can be triggered without a complete TCP handshake).

GFWeb observed a total of 2.3M redirection attempts caused by China Telecom’s middleboxes. In addition to two URLs reported on the Tor Project’s GitLab [73] (i.e., URLs 1 and 2 in Table 4), we have also observed six other URLs that share the same pattern but were injected with lower frequency. Our data also indicates that URLs ending with `parameter1` and `parameter2` were not deployed until February 2023.

Table 4 also lists some injected URLs that we deem as “buggy” because they contain either an invalid or non-routable IP address (i.e., `0.0.0.*`). We believe that these URLs are a result of misconfiguration because the injection of these URLs will not lead to any redirection, and thus they are not effective for the ISP’s intended purpose.

GFWeb observed 478K unique FQDNs that trigger these injections. Still, we could not find any patterns that could explain why these domains were targeted, since they only triggered the redirection for a short period of time. Looking up the IP addresses of these domains, we find that less than half of them are associated with an IP address, while the remaining are either not associated with any IP address or not existing (i.e., NXDOMAIN [21]). Instead of targeting specific domains that are truly associated with fraud, this observation suggests that the redirection could have been randomly triggered in an opportunistic manner to redirect users to the warning page, persuading them to install the government’s “anti-fraud” app. This is evident by the fact that one of the trigger domains was `baidu.com` [73], which is obviously not fraudulent.

The original anecdote [73] also reported that OONI [38] and Censored Planet [64] observed this network interference happened via DNS injection as well. However, we could not reproduce redirection cases caused by DNS injection of `182.43.124.6` anymore. This strengthens our hypothesis that the redirection is performed in an opportunistic manner and does not target any particular domains for a long period of time. At the time of writing, we are still observing this network interference from China Telecom and will continue monitoring for any change in the future.

Table 4: Redirection URLs injected by China Telecom’s middleboxes.

Index	# Triggered	Redirection URLs
1	1.7M (75.2%)	182.43.124.6
2	182K (7.9%)	182.43.124.6/fzyujing?parameter
3	123K (5.3%)	0.0.0.0/fzyujing?parameter1
4	86K (3.7%)	0.0.0.0/fzyujing?parameter
5	74K (3.2%)	182.43.124.6/fzyujing?parameter1
6	67K (2.9%)	182.43.124.6/fzyujing?parameter2
7	26K (1.1%)	0.0.0.0
8	6K (0.3%)	0.0.0.0124.6/fzyujing?parameter

7 Implications

This section discusses the implications of our findings for censorship measurement and evasion, especially in the context of monitoring and circumventing the GFW’s Web censorship.

7.1 External Measurement Based on Bidirectional Filtering

The asymmetric filtering policy discovered in §5.3 impacts platforms that rely on bidirectional censorship. Within the research community, there have been two main hypotheses on why the GFW’s Web censorship is designed to be bidirectional. One is that the GFW is designed to also geo-block some domestic websites hosted inside China, preventing their access from outside the country. The other is that maintaining the direction information of every connection can be operationally expensive for the GFW, given the massive amount of traffic that leaves and enters the country at any given time.

In our previous work [47] we found evidence to support the former by showing that the GFW’s DNS filter indeed injects forged responses to DNS queries for some websites hosted inside China. Our finding of asymmetric interference, however, shows that the GFW can maintain the direction information of network connections, and only takes interference actions when the direction of a triggering probe is from inside China.

While this asymmetric interference is only applied to a small group of domains, this filtering policy can affect systems like Hyperquack [72] which rely on bidirectional censorship to conduct remote measurement as the GFW’s HTTPS filter will not interfere with certain probes from outside China. Indeed, HTTPS analysis from the project’s dashboard shows relatively low `Unexpected Rates` for these domains [7].

7.2 Measurement Based on a Single Protocol

Different blocklists across DNS, HTTP, and HTTPS filters discussed in §5.2 could lead to incomplete visibility into the GFW’s Web censorship of measurement systems that rely on a single protocol to detect censorship. OONI [38], for instance, is designed to capture censorship in multiple protocols for

Web connectivity, but its results can be biased towards the first protocol being interfered with due to the sequential nature of its testing. For example, `hrw.org` is blocked across all three protocols, but many OONI’s test results will indicate only DNS blocking because DNS is the first layer of censorship triggered. Subsequent connections may not be indicative and conclusive of HTTP(S) filtering since the resolved IPs have already been falsified by the DNS filter. Unless the testing client could evade the DNS filter [32, 42], or the GFW sometimes fails the race condition [47], OONI may not be able to observe interference caused by the HTTP(S) filters. Moreover, OONI will also not observe interference by the HTTP filter if the input URL from the Citizen Lab test list [4] is already prefixed with `https://`.

Similarly, DNS-based systems like GFWatch [47] will miss domain names that are only blocked by the HTTP(S) filters. On the other hand, TCP-based measurement systems like Cloudflare Radar [63] will miss domains that are blocked by the DNS filter. Cloudflare Radar relies on passive TCP connections observed at their CDN edge servers to detect connection tampering. However, since many censored domains may have already been blocked by the GFW’s DNS filter even before the TCP connection could be established, Cloudflare will never see the TCP connections initiated for such domains. As a result, what Cloudflare Radar sees is a subset of the GFW’s Web censorship imposed on domains that the DNS filter does not or fails to block.

Even outside the context of censorship measurement, Xie et al. [80] recently proposed Secrank, a China-specific domain top list built from passive DNS traffic collected by the largest public DNS resolver in China. Checking the censored domains detected by GFWeb against Secrank top SLDs, we find more than 73K common SLDs that are censored by the GFW’s HTTP(S) filters, including many popular known censored domains (e.g., `google.com`, `facebook.com`, and `hrw.org`). This observation suggests that DNS queries alone may not be sufficient to reflect the actual situation of Web access in China as the GFW’s HTTP(S) filters can still block a significant number of domains after the DNS resolution process.

7.3 Probing-based Evasion Strategy

The discovery of the GFW’s traffic dropping as a form of residual censorship against certain domains (§5.4) also has an important implication for evasion tools that rely on probing the GFW to automate strategy searching.

Geneva [19], for example, is the state-of-the-art evasion strategy generator that functions based on genetic algorithms to automate the generation of evasion strategies. The process involves continuously probing a target censor to search for evasion strategies, and requires prior knowledge of the censor’s blocking behavior to guide the search process. In particular, a consistent censoring signal (i.e., injected packets

or traffic dropping) has to be known in advance to determine whether an evasion strategy is effective or not.

Thus, the coexistence of both injected packets and traffic dropping in the GFW’s blocking behavior against the same protocol as shown in §5.4 can be problematic for Geneva-like tools for two reasons. First, the GFW’s traffic dropping behavior can be a significant obstacle for the automation process because once activated it can stay effective for at least 350 seconds, preventing continuous probing. Second, the absence of packets injected by the GFW may trick Geneva-like tools into concluding that an effective evasion strategy has been found if the initial assumption is that the GFW will inject packets if a strategy is unsuccessful, and vice versa.

With the increasing popularity of automated evasion tools [19, 42], Amich et al. [10] recently examined the possibility of censors using machine learning to detect probing traffic and impose traffic dropping, to prevent the discovery of straightforward automated probing-based strategies. To the best of our knowledge, this is the first time prolonged traffic dropping triggered by probing censored domains against the GFW has been observed in the wild. All traffic dropping cases observed by prior studies are in the context of blocking fully encrypted connections of network relays [9, 33, 36, 78, 79] or Encrypted SNI [52], but not Web censorship.

This newly discovered prolonged traffic dropping behavior is, perhaps, the GFW’s response to recent developments in automated evasion discovery to hinder the straightforward use of such tools. In fact, we have also noticed that the GFW’s HTTPS filter already fixed an issue that prevented it from reassembling out-of-order fragmented TLS Client Hello packets [20]. At the time of this writing, it is now capable of correctly reassembling these fragmented packets and takes interference actions based on the reassembled domain name detected in the SNI extension.

8 Discussion

We next discuss the ethical considerations and limitations of our work, and make suggestions for future censorship monitoring and circumvention efforts.

8.1 Ethical Considerations

Internet censorship, particularly when implemented by nation-states, is often driven by political motives [30, 41, 67], making it a delicate subject of study. Consequently, assessing network interference resulting from nation-state censorship demands careful execution that does not impose risk on any end users. In designing GFWeb (§4), we send probes between machines under our control for data collection, thus eliminating the need for participation from end users. This core strategy enables GFWeb to fully operate in an autonomous manner and allows for straightforward system redeployment from different vantage points if needed. To further reduce potential impact on

the hosting providers where our measurement machines are deployed in China, we will apply the limited TTL-based probing strategy so that our external probing packets will reach the GFW but not the hosting providers’ network.

We have also considered the broader implications of our research on different communities including academics, policymakers, and the general public. Our goal through this empirical measurement study is to furnish a factual, holistic, and unbiased analysis that contributes to the understanding of the GFW’s Web censorship and its impact on Internet freedom of hundreds of millions of users in China. We thus design GFWeb in a scalable and longitudinal fashion such that the system can continuously test an extensive range of domain names from multiple sources instead of relying on any single test list of domain names (§4.1), which could potentially introduce bias into our measurement results.

By adhering to these considerations, our study contributes valuable knowledge to the field of Internet censorship research while upholding our commitment to responsible and ethical research practices.

8.2 Limitations

Localized Censorship. Each measurement system is designed with specific goals and limitations, and our system, GFWeb, is no different. In order to measure a wide array of domains, we are faced with the necessity of balancing the depth and breadth of our coverage. This means that certain localized censorship events, which might occur in specific residential network areas, could go undetected by our system but may be observable by others with more extensive vantage points, such as OONI [38] or Censored Planet [64]. For example, our system was able to identify network interference linked to China Telecom’s “Anti-fraud” campaign, yet there have been reports of additional interference events at various network locations outside our coverage [73].

Detectability and Blockability. Similar to the blocking case of OONI [70], our measurement infrastructure is also susceptible to being identified and blocked by the GFW. Throughout our study, we did not observe any targeted blocking against our machines. However, the blocking of OONI’s main site and its data collection servers significantly impeded their data collection efforts within China [70]. This incident highlights that host-based traffic analysis poses a considerable risk to the functionality of our system, particularly because GFWeb is designed to send a substantial volume of probes to elicit responses from the GFW’s censoring middleboxes.

In an effort to achieve comprehensive measurement coverage while minimizing risk to end users, we accept the trade-off that GFWeb’s probing traffic may be detectable by the GFW. Consequently, there exists a possibility that our machines could be blocked in the future. Crucially, however, GFWeb’s operations are not confined to any particular machine and can

be readily migrated to alternative machines as necessary. Furthermore, other strategies such as splitting the probing traffic across multiple machines and/or imposing delays and rate limits on the probing traffic could be employed to mitigate the risk of detection and blocking [10]. We are actively exploring these strategies to ensure the continuous operation of GFWeb.

Moreover, our measurement strategy could become ineffective if the GFW decides to become truly stateful. While this is a valid concern, such a change is unlikely to happen since the loss-tolerant nature of the GFW, rather than a bug, is a design choice to handle common situations when packets are lost or routed through different network paths [14]. Should such a change occur, it may only slow us down but will not prevent us from running our measurement, as we can just configure our machines at the other side of the GFW to actually complete the TCP handshake, satisfying the statefulness requirement.

8.3 Suggestions

8.3.1 Censorship Measurement

The discovery of asymmetric interference (§5.3) underscores the importance of conducting measurements from both sides of the GFW, since filtering policies can be different depending on the probing direction and the domain being tested.

Furthermore, measurement systems that function based on continuous probing against remote servers using the same destination ports will need to be aware of the two different types of residual censorship and take appropriate actions to avoid incorrect inferences. More specifically, the residual censorship that “keeps injecting” packets may cause false positives as subsequent benign probes would still trigger the GFW to emit forged packets, whereas the residual censorship that “keeps dropping” subsequent traffic may cause false negatives due to the absence of forged packets that are usually anticipated. The co-existence of these two types of residual censorship in HTTPS filtering also suggests that it is non-trivial to determine whether a domain is blocked or not based solely on the presence or absence of forged packets if residual censorship is not taken into account. To that end, it is important to design measurement approaches that can sidestep the residual censorship to avoid both false positives and false negatives.

8.3.2 Censorship Circumvention

The GFW’s Web censorship is composed of multiple layers of filtering based on different blocklists and protocols. While various efforts have attempted to circumvent the GFW’s Web censorship at different layers [20, 25, 32, 42, 47, 76], an effective circumvention solution will need to tackle the GFW’s multi-layered filtering architecture. Otherwise, circumvention solutions that only target a single filtering layer may not be sufficient to achieve the desired result.

Similarly to the suggestion for censorship measurement, probing-based evasion techniques like Geneva [19] also need

to be aware of the two different types of residual censorship to avoid being tricked into thinking that the censorship has been successfully evaded when it is actually not, especially when the residual censorship is of the “keeps dropping” type [10].

8.3.3 Using Measurement Data

Internet censorship measurement is a challenging task. Each measurement system is designed with different resources and constraints. Consumers of censorship measurement data (e.g., journalists, researchers, and policy makers) should be aware of the strengths and drawbacks of each system, and consider multiple measurement results from different protocols and data sources to obtain a more complete picture of the censorship landscape. When it comes to determining the censorship status of a domain, it is important to gather results from multiple systems and protocols to obtain a more conclusive result.

9 Conclusion

DNS, HTTP, and HTTPS filtering middleboxes together form the primary pillars of the GFW’s Web censorship. In this work, we present GFWeb, a longitudinal measurement system designed to discover domain blocklists used by the GFW for censoring Web access. Over the course of 20 months, GFWeb has tested over a billion fully qualified domains, and detected 943K and 55K pay-level domains censored by the GFW’s HTTP and HTTPS filtering middleboxes, respectively. Our study not only complements prior efforts by providing a more comprehensive view into the GFW’s Web censorship over time, but also reveals several new findings, including the GFW’s asymmetric blocking behavior, and patches of overblocking and failure in reassembling fragmented packets.

The implications of our investigation extend far beyond academic circles, touching on the fabric of global Internet governance and the ongoing struggle for digital freedom. The adaptive nature of the GFW signals a future where Internet censorship will become more nuanced and technically complex, posing significant challenges for circumvention technologies and international policy efforts.

In light of these insights, our work underscores the need for a reinvigorated approach to understanding and combating Internet censorship. The dynamic between censorship and circumvention is not static; it evolves as part of a larger geopolitical and technological landscape, with implications for global Internet freedom, the free flow of information, and the resistance against digital authoritarianism.

As we keep operating GFWeb, we hope that our data (publicly available at <https://gfweb.ca>) will not only provide fresh insights into technical observations, but also promptly update the public regarding changes in the GFW’s blocking policies and support other initiatives, especially those focusing on censorship detection and circumvention.

Acknowledgments

We would like to thank the anonymous reviewers and shepherd for their thorough feedback on this paper. We also thank the team at GreatFire.org and others who preferred to remain anonymous for helpful discussions and suggestions.

This research was supported in part by the Open Technology Fund's Internet Freedom program, by DARPA under award HR00112190126, and by the National Science Foundation under award CNS-1814817. The opinions in this paper are those of the authors and do not necessarily reflect the opinions of the sponsors.

References

- [1] Common Crawl Project. <https://commoncrawl.org>.
- [2] ICANN Centralized Zone Data Service. <https://czds.icann.org>.
- [3] Public Suffix List. <https://publicsuffix.org/>.
- [4] The Citizen Lab Test Lists. <https://github.com/citizenlab/test-lists>.
- [5] VirusTotal: URL Scanning Service. <https://www.virustotal.com/gui/home/url>.
- [6] Censored Planet Dashboard, Accessed: 2023-10-01. <https://dashboard.censoredplanet.org/access.html>.
- [7] Censored Planet Dashboard - Evidence of Asymmetric Interference of the Great Firewall's HTTPS filter, Accessed: 2023-10-11. <https://archive.ph/1PXS6>.
- [8] Josh Aas, Richard Barnes, Benton Case, Zakir Durumeric, Peter Eckersley, Alan Flores-López, J. Alex Halderman, Jacob Hoffman-Andrews, James Kasten, Eric Rescorla, Seth D. Schoen, and Brad Warren. Let's Encrypt: An automated certificate authority to encrypt the entire web. *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, 2019.
- [9] Alice, Bob, Carol, Jan Beznazwy, and Amir Houmansadr. How China detects and blocks shadowsocks. *Proceedings of the ACM Internet Measurement Conference*, 2020.
- [10] Abderrahmen Amich, Birhanu Eshete, Vinod Yegneswaran, and Nguyen Phong Hoang. Deresistor: Toward detection-resistant probing for evasion of internet censorship. In *Proceedings of the USENIX Security Symposium*, 2023.
- [11] Anonymous, Arian Akhavan Niaki, Nguyen Phong Hoang, Phillipa Gill, and Amir Houmansadr. Triplet censors: Demystifying Great Firewall's DNS censorship behavior. In *Proceedings of the USENIX Workshop on Free and Open Communications on the Internet (FOCI)*, 2020.
- [12] Anonymous Author(s). Towards a comprehensive picture of the Great Firewall's DNS censorship. In *Proceedings of the USENIX Workshop on Free and Open Communications on the Internet (FOCI)*, 2014.
- [13] Derek E. Bambauer, Ronald J. Deibert, J. Palfrey, Rafal Rohozinski, N. Villeneuve, and J. Zittrain. Internet Filtering in China in 2004-2005: A Country Study. 2005.
- [14] Abhishek Bhaskar and Paul Pearce. Many roads lead to Rome: How packet headers influence DNS censorship measurement. In *Proceedings of the USENIX Security Symposium*, 2022.
- [15] Bloomberg. China Tells Carriers to Block Access to Personal VPNs by February, 2017-07-10. <https://www.bloomberg.com/news/articles/2017-07-10/china-is-said-to-order-carriers-to-bar-personal-vpns-by-february>.
- [16] Kevin Bock, Abdulrahman Alaraj, Yair Fax, Kyle Hurley, Eric Wustrow, and Dave Levin. Weaponizing middleboxes for TCP reflected amplification. In *Proceedings of the USENIX Security Symposium*, 2021.
- [17] Kevin Bock, Pranav Bharadwaj, Jasraj Singh, and Dave Levin. Your censor is my censor: Weaponizing censorship infrastructure for availability attacks. In *Proceedings of the IEEE Workshop on Offensive Technologies (WOOT)*, 2021.
- [18] Kevin Bock, Yair Fax, Kyle Reese, Jasraj Singh, and Dave Levin. Detecting and evading censorship-in-depth: A case study of Iran's protocol filter. In *Proceedings of the USENIX Workshop on Free and Open Communications on the Internet (FOCI)*, 2020.
- [19] Kevin Bock, George Hughey, Xiao Qiang, and Dave Levin. Geneva: Evolving censorship evasion strategies. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2019.
- [20] Kevin Bock, Gabriel Naval, Kyle Reese, and Dave Levin. Even censors have a backup: Examining China's double HTTPS censorship middleboxes. In *Proceedings of the ACM SIGCOMM Workshop on Free and Open Communications on the Internet*, 2021.
- [21] S. Bortzmeyer and S. Huque. NXDOMAIN: There Really Is Nothing Underneath. RFC 8020, IETF, November 2016.
- [22] Jacob Brown, Xi Jiang, Van Tran, Arjun Nitin Bhagoji, Nguyen Phong Hoang, Nick Feamster, Prateek Mittal, and Vinod Yegneswaran. Augmenting rule-based DNS censorship detection at scale with machine learning. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2023.
- [23] Cate Cadell. Apple says it is removing VPN services from China App Store. Reuters, 2017-07-29. <https://www.reuters.com/article/us-china-apple-vpn/apple-says-it-is-removing-vpn-services-from-china-app-store-idUSKBN1AE0BQ>.
- [24] Chromium Blog. Towards HTTPS by default, 2023-08-16. <https://blog.chromium.org/2023/08/towards-https-by-default.html>.
- [25] Richard Clayton, Steven J. Murdoch, and Robert N. M. Watson. Ignoring the Great Firewall of China. In *Proceedings of the 6th Workshop on Privacy Enhancing Technologies (PET)*, pages 20–35. Springer, 2006.
- [26] Cloudflare Community. China Telecom HTTP Redirect of 1.1.1.1, 2022-04-22. <https://community.cloudflare.com/t/china-telecom-http-redirect-of-1-1-1-1/378187>.
- [27] Jedidiah R. Crandall, Daniel Zinn, Michael Byrd, Earl Barr, and Rich East. ConceptDoppler: A weather tracker for Internet censorship. In *Proceedings of the 14th ACM conference on Computer and communications security (CCS)*, pages 352–365, 2007.
- [28] Cyberspace Administration of China. Administrative Measures for Generative Artificial Intelligence Services (in Chinese), 2023-04. http://www.cac.gov.cn/2023-04/11/c_1682854275475410.htm.
- [29] Ronald Deibert. China's Cyberspace Control Strategy: An Overview and Consideration of Issues for Canadian Policy, 2010.
- [30] Ronald Deibert. Reset: Reclaiming The Internet for Civil Society, 2020.

- [31] Ronald J. Deibert. Dark Guests and Great Firewalls: The Internet and Chinese Security Policy. *Journal of Social Issues*, 58:143–159, 2002.
- [32] Haixin Duan, Nicholas Weaver, Zongxu Zhao, Meng Hu, Jinjin Liang, Jian Jiang, Kang Li, and Vern Paxson. Hold-On: Protecting against on-path DNS poisoning. In *Proceedings of the Conference on Securing and Trusting Internet Names (SATIN)*, 2012.
- [33] Arun Dunna, Ciarán O’Brien, and Phillipa Gill. Analyzing China’s blocking of unpublished Tor bridges. In *Proceedings of the USENIX Workshop on Free and Open Communications on the Internet (FOCI)*, 2018.
- [34] Let’s Encrypt. Let’s Encrypt Stats, 2019. <https://letsencrypt.org/stats/>.
- [35] Roya Ensafi, David Fifield, Philipp Winter, Nick Feamster, Nicholas Weaver, and Vern Paxson. Examining how the Great Firewall discovers hidden circumvention servers. In *Proceedings of the ACM Internet Measurement Conference (IMC)*. ACM, 2015.
- [36] Roya Ensafi, Philipp Winter, Abdullah Mueen, and Jedidiah R. Crandall. Analyzing the Great Firewall of China over space and time. In *Proceedings of the 15th Privacy Enhancing Technologies Symposium (PoPETs)*, 2015.
- [37] Oliver Farnan, Alexander Darer, and Joss Wright. Poisoning the well – exploring the Great Firewall’s poisoned DNS responses. In *Proceedings of the ACM Workshop on Privacy in the Electronic Society (WPES)*, 2016.
- [38] Arturo Filastò and Jacob Appelbaum. OONI: Open observatory of network interference. In *Proceedings of the USENIX Workshop on Free and Open Communications on the Internet (FOCI)*, 2012.
- [39] Freedom House. Freedom in the World 2023 - China, 2023. <https://freedomhouse.org/country/china/freedom-world/2023>.
- [40] Geremie R. Barme And Sang Ye. The Great Firewall of China, 1997-06-01. <https://www.wired.com/1997/06/china-3/>.
- [41] Phillipa Gill, Masashi Crete-Nishihata, Jakub Dalek, Sharon Goldberg, Adam Senft, and Greg Wiseman. Characterizing web censorship worldwide: Another look at the OpenNet Initiative data. *ACM Transactions on the Web*, 9(1), 2015.
- [42] Michael Harrity, Kevin Bock, Frederick Sell, and Dave Levin. GET /out: Automated discovery of application-layer censorship evasion strategies. In *Proceedings of the USENIX Security Symposium*, 2022.
- [43] Nguyen Phong Hoang. GFWatch Dashboard, 2020. <https://gfwatch.org>.
- [44] Nguyen Phong Hoang, Sadie Doreen, and Michalis Polychronakis. Measuring I2P censorship at a global scale. In *Proceedings of the USENIX Workshop on Free and Open Communications on the Internet (FOCI)*, 2019.
- [45] Nguyen Phong Hoang, Ivan Lin, Seyedhamed Ghavamnia, and Michalis Polychronakis. K-resolver: Towards decentralizing encrypted DNS resolution. In *Proceedings of the 2nd Workshop on Measurements, Attacks, and Defenses for the Web (MADWeb)*, 2020.
- [46] Nguyen Phong Hoang, Arian Akhavan Niaki, Nikita Borisov, Phillipa Gill, and Michalis Polychronakis. Assessing the privacy benefits of domain name encryption. In *Proceedings of the 15th ACM Asia Conference on Computer and Communications Security (ASIACCS)*, 2020.
- [47] Nguyen Phong Hoang, Arian Akhavan Niaki, Jakub Dalek, Jeffrey Knockel, Pellaeon Lin, Bill Marczak, Masashi Crete-Nishihata, Phillipa Gill, and Michalis Polychronakis. How great is the Great Firewall? Measuring China’s DNS censorship. In *Proceedings of the 30th USENIX Security Symposium*, 2021.
- [48] Nguyen Phong Hoang, Arian Akhavan Niaki, Phillipa Gill, and Michalis Polychronakis. Domain name encryption is not enough: Privacy leakage via IP-based website fingerprinting. In *Proceedings on the 21st Privacy Enhancing Technologies Symposium (PoPETs)*, 2021.
- [49] Nguyen Phong Hoang, Arian Akhavan Niaki, Michalis Polychronakis, and Phillipa Gill. The Web is Still Small After More Than a Decade. *ACM SIGCOMM Computer Communication Review (CCR)*, 50(2):24–31, 2020.
- [50] Nguyen Phong Hoang, Michalis Polychronakis, and Phillipa Gill. Measuring the accessibility of domain name encryption and its impact on Internet filtering. In *Proceedings of the Passive and Active Measurement Conference (PAM)*, 2022.
- [51] Hugo Butcher Piat. Navigating the Internet in China: Top Concerns for Foreign Businesses, 2019-03-12. <https://www.china-briefing.com/news/internet-china-top-concerns-foreign-businesses>.
- [52] Kevin Bock and iyouport and Anonymous and Louis-Henri Merino and David Fifield and Amir Houmansadr and Dave Levin. Exposing and Circumventing China’s Censorship of ESNI, 2020-08-07. <https://geneva.cs.umd.edu/posts/china-censors-esni/esni>.
- [53] Alexander Lex, Nils Gehlenborg, Hendrik Strobelt, Romain Vuillemot, and Hanspeter Pfister. Upset: Visualization of intersecting sets. *IEEE Transactions on Visualization and Computer Graphics*, 20, 2014.
- [54] Graham Lowe, Patrick Winters, and Michael L. Marcus. The Great DNS Wall of China. Technical report, New York University, 2007.
- [55] Bill Marczak, Nicholas Weaver, Jakub Dalek, Roya Ensafi, David Fifield, Sarah McKune, Arn Rey, John Scott-Railton, Ron Deibert, and Vern Paxson. An analysis of China’s “Great Cannon”. In *Proceedings of the USENIX Workshop on Free and Open Communications on the Internet (FOCI)*, 2015.
- [56] Arian Akhavan Niaki, Shinyoung Cho, Zachary Weinberg, Nguyen Phong Hoang, Abbas Razaghpanah, Nicolas Christin, and Phillipa Gill. ICLab: A global, longitudinal internet censorship measurement platform. In *Proceedings of the IEEE Symposium on Security & Privacy*, 2020.
- [57] Sadia Nourin, Kevin Bock, Nguyen Phong Hoang, and Dave Levin. Detecting network interference without endpoint participation. In *Proceedings of the USENIX Workshop on Free and Open Communications on the Internet (FOCI)*, 2023.
- [58] Sadia Nourin, Van Tran, Xi Jiang, Kevin Bock, Nick Feamster, Nguyen Phong Hoang, and Dave Levin. Measuring and evading Turkmenistan’s internet censorship. In *Proceedings of the ACM Web Conference (WWW)*, 2023.
- [59] OONI. OONI Explorer: Uncover evidence of internet censorship worldwide, Accessed: 2023-10-01. <https://explorer.ooni.org/>.
- [60] Jong Chun Park and Jedidiah R. Crandall. Empirical study of a national-scale distributed intrusion detection system: Backbone-level filtering of HTML responses in China. In *Proceedings of the 30th International Conference on Distributed Computing Systems (ICDCS)*, pages 315–326, 2010.

- [61] Paul Pearce, Ben Jones, Frank Li, Roya Ensafi, Nick Feamster, Nick Weaver, and Vern Paxson. Global measurement of DNS manipulation. In *Proceedings of the USENIX Security Symposium*, 2017.
- [62] Victor Le Pochat, Tom van Goethem, Samaneh Tajalizadehkhooob, Maciej Korczyński, and Wouter Joosen. Tranco: A research-oriented top sites ranking hardened against manipulation. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2019.
- [63] Ram Sundara Raman, Louis-Henri Merino, Kevin Bock, Marwan M. Fayed, Dave Levin, Nick Sullivan, and Luke Valenta. Global, passive detection of connection tampering. *Proceedings of the ACM SIGCOMM Conference*, 2023.
- [64] Ram Sundara Raman, Prerana Shenoy, Katharina Kohls, and Roya Ensafi. Censored Planet: An Internet-wide, longitudinal censorship observatory. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, 2020.
- [65] Ram Sundara Raman, Mona Wang, Jakub Dalek, Jonathan Mayer, and Roya Ensafi. Network measurement methods for locating and examining censorship devices. In *Proceedings of the 18th International Conference on Emerging Networking Experiments and Technologies (CoNEXT)*, 2022.
- [66] Raymond Rambert, Zachary Weinberg, Diogo Barradas, and Nicolas Christin. Chinese wall or Swiss cheese? keyword filtering in the Great Firewall of China. In *Proceedings of the ACM Web Conference (WWW)*, 2021.
- [67] Lotus Ruan, Masashi Crete-Nishihata, Jeffrey Knockel, Ruohan Xiong, and Jakub Dalek. The intermingling of state and private companies: Analysing censorship of the 19th national communist party congress on wechat. *The China Quarterly*, 246:497 – 526, 2020.
- [68] Walter Rweyemamu, Tobias Lauinger, Christo Wilson, William K. Robertson, and Engin Kirda. Clustering and the weekend effect: Recommendations for the use of top domain lists in security research. In *Proceedings of the Passive and Active Network Measurement Conference (PAM)*, 2019.
- [69] Will Scott, Thomas Anderson, Tadayoshi Kohno, and Arvind Krishnamurthy. Satellite: Joint analysis of CDNs and network-level interference. In *Proceedings of the USENIX Annual Technical Conference (ATC)*, 2016.
- [70] Simone Basso and Maria Xynou and Arturo Filasto. China is blocking OONI, 2023-07-28. <https://ooni.org/post/2023-china-blocks-ooni/>.
- [71] Sparks, Neo, Tank, Smith, and Dozer. The collateral damage of Internet censorship by DNS injection. *SIGCOMM Computer Communication Review*, 42(3):21–27, 2012.
- [72] Ram Sundara Raman, Adrian Stoll, Jakub Dalek, Reethika Ramesh, Will Scott, and Roya Ensafi. Measuring the deployment of network censorship filters at global scale. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2020.
- [73] Tor Project’s Anti-Censorship Team. China Anti-Fraud Webpage Redirection Censorship, 2022-01-12. <https://gitlab.torproject.org/tpo/anti-censorship/censorship-analysis/-/issues/40026>.
- [74] Lokman Tsui. An inadequate metaphor: the Great Firewall and Chinese internet censorship. *Global Dialogue*, 9(1/2):60, 2007.
- [75] Benjamin VanderSloot, Allison McDonald, Will Scott, J. Alex Halderman, and Roya Ensafi. Quack: Scalable remote measurement of application-layer censorship. In *Proceedings of the USENIX Security Symposium*, 2018.
- [76] Zhongjie Wang, Yue Cao, Zhiyun Qian, Chengyu Song, and Srikanth V. Krishnamurthy. Your state is not mine: A closer look at evading stateful Internet censorship. In *Proceedings of the Internet Measurement Conference (IMC)*, 2017.
- [77] Will Henshall. How China’s New AI Rules Could Affect U.S. Companies, 2023-09-19. <https://time.com/6314790/china-ai-regulation-us>.
- [78] Philipp Winter and Stefan Lindskog. How the Great Firewall of China is blocking Tor. In *Proceedings of the USENIX Workshop on Free and Open Communications on the Internet (FOCI)*, 2012.
- [79] Mingshi Wu, Jackson Sippe, Danesh Sivakumar, Jack Burg, Peter Anderson, Xiaokang Wang, Kevin Bock, Amir Houmansadr, Dave Levin, and Eric Wustrow. How the Great Firewall of China detects and blocks fully encrypted traffic. In *Proceedings of the USENIX Security Symposium*, 2023.
- [80] Qinge Xie, Shujun Tang, Xiaofeng Zheng, Qingran Lin, Baojun Liu, Haixin Duan, and Frank H. Li. Building an open, robust, and stable voting-based domain top list. In *Proceedings of the USENIX Security Symposium*, 2022.
- [81] Xueyang Xu, Z. Morley Mao, and J. Alex Halderman. Internet censorship in China: Where does the filtering occur? In *Proceedings of the Passive and Active Measurement Conference (PAM)*, pages 133–142, 2011.
- [82] Young Xu. Deconstructing the Great Firewall of China. Technical report, Thousand Eyes, 2016.
- [83] Jonathan Zittrain and Benjamin Edelman. Internet filtering in china. *IEEE Internet Computing*, 7(2):70–77, mar 2003.