

A Survey and Taxonomy of Internet Censorship Methods

ALEXANDER MASTER, United States Military Academy, USA

EUGENE H. SPAFFORD, Purdue University, USA

Despite the significance of the Internet in global information exchange, authorities in many jurisdictions apply censorship to online communications within their sphere of influence. Some people consider censorship a violation of human rights, such as the freedom of expression or the right to access information. Others view censorship methods as a means of social control, protecting national security or social harmony. Regardless of the motivation, researchers, policymakers, and technology practitioners should have a thorough understanding of the capabilities of Internet censors. We present a taxonomy of Internet censorship methods derived from a systematic review of the literature and analysis of the technology that has enabled the Internet's progression. We describe the technical means that censors implement, contextualized by the legal and social circumstances in which they have occurred over the past three decades. We assert that the taxonomy has pedagogical value and serves as a research framework for further studies. Our study illuminates the dynamic challenges faced in interdisciplinary research on Internet censorship, offering direction for future work on Internet measurement, censorship circumvention, and protocol standards design.

Keywords: Internet censorship methods, taxonomy, nation-state censors, networking, circumvention

1 Introduction

Internet censorship is a relatively new phenomenon in the human experience as a means of controlling information. With the transition from the Industrial Age to the Information Age, information technology and computing systems have become central to society's functions and progress. With the advent of the World Wide Web in the 1990s and decreasing costs of technology, Internet access became available to average citizens in many nations. Internet-connected devices have continued to proliferate, allowing more people to communicate in near real-time over various mediums and protocols. At the same time, governments of nation-states around the world began to realize the power and influence of the "information superhighway." Authorities proceeded to explore methods for controlling information access and exchange through this new medium. Internet censorship has increased in size, scale, and capability over the past three decades, and now impacts billions of people globally [144, 197].

Aceto and Pescapé define Internet censorship as "the intentional impairing or blocking of access to online resources and services," regardless of the goal, scope, or legitimacy of the actions [4]. The technical means by which Internet censors implement censorship vary widely but are generally constrained by standardized protocols that enable the Internet to function as a network of networks. Some Internet censors block web pages with "objectionable"¹ content and inform users that their request was denied. Others simply drop the connection or manipulate the traffic so that the connection appears to fail. Some censors degrade the connectivity of a user when undesirable activity is detected, rendering the communication unusable. Applications that allow for unmoderated communications may be blocklisted (some previous literature refers to this practice as "blacklisting"), denied based on ports used by the application's software or signatures of its data packets. Censors can be nation-state actors or private entities. In addition to blocking communication, Internet surveillance is closely related to censorship activities. Detection of objectionable content is generally the first and most important step toward implementation of a blocking action. More aggressive nation-state censors completely disconnect Internet connectivity [32] during perceived critical events, such as elections or periods of civil unrest [130]. In some extreme cases, censorship may lead to physical confrontation with the content authors by local authorities.

¹Objectionable is locally defined by policy or statute, and appropriateness is thus subjective. In the following, we will refer to objectionable content without judgment as to actual subject matter.

Internet censorship occurs using a broad spectrum of methods and activities. In some cases, the effects of communication disconnection are overt and obvious to the user population, while other countries seek to minimize economic collateral damage from their censorship and use more targeted techniques. Some authorities wish to conceal that they are censoring content and resort to techniques such as the use of artificial Transmission Control Protocol (TCP) Reset packets [99, 306, 310] or bandwidth throttling [98, 208, 334]. In these cases, users may believe that they are experiencing poor network service, which increases the difficulty for Internet measurement researchers to detect censor activity and changes over time. Furthermore, as end-to-end encrypted (E2EE) messaging services gain popularity for their security and privacy properties, censoring nation-states have shown an increased willingness to target these protocols. The proliferation of encrypted traffic analysis (ETA) software or next-generation firewalls (NGFWs) that can block applications such as Signal [237, 336, 338] and Tor Browser [91] may represent new threats to private communications.

Given censorship's wide-ranging implications, in this study we sought to systematize the knowledge of Internet censorship methods by answering the following research question: "What are the technical threats to Internet-based communications within the context of censorship?" We present the results in §5 as a taxonomy of Internet censorship methods. We demonstrate the taxonomy's validity using a systematic review of several decades of multidisciplinary literature surrounding Internet-based censorship and how it has evolved. We assert that the taxonomy has pedagogical value and serves as a research framework for further studies. Our study illuminates challenges faced in interdisciplinary research on Internet censorship, offering direction for future work on Internet measurement, censorship circumvention, and protocol standards design.

2 Background

The authorities of some countries go to great lengths to deny their citizens free and open access to Internet resources. Nation-state Internet censorship is generally characterized as either centralized or decentralized in nature. Centralized censorship often occurs on government-controlled infrastructure. In some nations, there are few (or only one) Internet Service Providers (ISP) or cellular carriers for users to choose from. When the state owns the infrastructure and controls Internet routing, filtering "objectionable" material or limiting access is more straightforward. The People's Republic of China (PRC), whose censorship apparatus is known as the "Great Firewall of China", is considered by many to be the most sophisticated and aggressive centralized state censor [35, 53, 97, 231, 297, 343]. Other examples include small countries with limited access to transnational fiber switching. Syria, which only has one government-controlled autonomous system (AS) [51], can uniformly implement technical censorship measures across its population.

In contrast, decentralized censorship tends to result in fragmented implementation. Websites available in one region may be denied in another. Examples of decentralized censorship regimes are the Russian Federation [244, 332] and India [124, 271, 337]. Authorities in these nations legally compel private-sector service providers to perform web filtering, throttling, or shutdowns. Technical implementations may vary widely among corporations, resulting in a patchwork of censorship. China has even shown signs of uneven implementation of censorship across regions in recent years [326].

Internet censorship also intersects with human conceptions of governance. The proliferation of Internet access has enabled the rapid and widespread exchange of information around the world, accelerating scientific discovery, enabling technological and economic development, and democratizing access to knowledge. The Internet has surpassed television, print media, and radio in terms of media influence and has become one of "the irreplaceable elements of our lives since the 2000s" [88]. While this has been largely a global phenomenon, leaders of some countries wish to assert control over the information environment within their geopolitical borders. In some cases, globalization has been met with the creation of "splinternets," attempts to create separate and independent online resources that reduce or eliminate reliance on global Internet resources. For many years, the Russian Federation has attempted to sever its ties with Western and international web resources and, on several occasions, has tested its ability to operate when

disconnected from transnational fiber connectivity [186, 194, 274]. The Islamic Republic of Iran provides an intranet of government-provided information, maintained by the Islamic Revolutionary Guard Corps (IRGC), for when the government shuts down external Internet access for its population (see §4.9). The Democratic People's Republic of Korea (DPRK, commonly known as North Korea) has long maintained a policy of denying external Internet connectivity to most of its citizens [146]. As part of China's strategic goal of promoting and maintaining social harmony, the Communist Party of China (CCP) has pursued a strategy of "cyber sovereignty," which presupposes that the state has the right to control cyberspace within its sphere of control [198]. Aside from its Social Credit System [54] and the Great Firewall, China offers replacement online platforms for Western websites it blocks, creating a media and information environment in which it can control messages, themes, and narratives through surveillance or coercion [250].² Concerns about "splinternets" and erosion of the global Internet have also emerged in Western contexts, when consumer privacy legislation in the European Union has begun to demand data sovereignty (non-reliance on foreign technology firms) and rights granted to EU citizens concerning their personal data [132, 188]. Among these global developments, Internet censorship remains one of the most direct and pervasive means of limiting individual freedom of expression.

Many instances of Internet censorship have occurred in the last three decades. Social movements, such as the "Arab Spring" in the early 2010s, were aided by the use of technology to organize protests and thus became targets for censorship. Authoritarian leaders have shown willingness to selectively censor information they deem a threat to their maintenance of power. Some governments have gone as far as completely disconnecting Internet connectivity during turbulent periods [2, 130].

Examples of state-sponsored censorship have been documented around the world. From March 2018 to July 2019, the Republic of Chad in Africa blocked access to all major social media platforms, including WhatsApp, Twitter, Instagram, YouTube, and Facebook, for "security reasons in the context of terrorist attacks" [72]. In August 2019, the longest-running Internet shutdown imposed by a democratic government occurred in the region of Jammu and Kashmir after the Indian government revoked the region's special autonomous status. The resulting communications blackout and Internet restrictions lasted approximately 18 months, with cellular mobile Internet access not fully restored until February 2021 [153]. In February 2021, after the military forces of Myanmar (formerly Burma) overthrew the elected government in a coup d'état, the Burmese people experienced cellular network shutdowns, nightly Internet curfews, and blocking of social networks (Facebook, Twitter, and Instagram) [233]. In July 2021, Cubans took to the streets to protest their government's response to the COVID-19 pandemic. Government officials restricted access to Facebook and WhatsApp during demonstrations. Many Cubans used the anti-censorship tool Psiphon for open access to the Internet [78].

In February 2022, metrics of total TOR network usage showed an increase of more than 300% in connections originating from Ukraine immediately after the invasion by Russia [285]. Ukrainians continued to rely on virtual private networks (VPNs) and anti-censorship technologies as the Russian Federation diverted Ukrainian Internet traffic through Russian service providers [258] to censor information sources and conduct surveillance operations. Domestically, Roskomnadzor — the federal agency responsible for monitoring, controlling, and censoring Russian mass media — ordered the blocking of popular apps such as Telegram and filtering traffic from media sources such as the British Broadcasting Company (BBC), Deutsche Welle, Radio Free Europe, and Voice of America, among others [248]. During the United States' 2026 conflict with Iran, Iranians experienced months of Internet blackout, denying them access to information from outside of their country [47].

Censorship is a broad issue that encompasses numerous aspects of modern societies. Some censorship is mandated by law in particular countries, for reasons ranging from morality to political manipulation and control of the information environment. This study does not assert value judgments about the reasons for Internet-based censorship. Instead,

²Examples of Chinese platforms include: Sina Weibo, which provides microblogging services; WeChat, which enables ride-sharing, instant messaging, ticket sales, and more; Douyin, which provides short-form video social media (similar to TikTok); Zhihu, which serves as a question-and-answer website (similar to Quora); and Baidu, China's most popular search engine.

we recognize that international expectations and norms are violated in some locations. Article 19 of the Universal Declaration of Human Rights states that freedom of expression is an inherent right, including the "freedom to hold opinions without interference and to seek, receive and impart information and ideas through any media and regardless of frontiers" [211]. Internet censorship, without broad agreement from those subjected to censorship about what materials are objectionable, limits their individual expression as well as the propagation of knowledge.

2.1 Related Work

Deibert et al. published the seminal *Access Denied* report in 2008, which was the first attempt at a global study on Internet censorship. Their work covered practices in 40 countries and divided technical censorship techniques into four categories: IP blocking, DNS tampering, blockpage, and keyword [82]. The Internet has since become increasingly complex. Global Internet user penetration rates continue to increase, and changing geopolitical landscapes impact the resources that censors are willing to implement to deny access to information. Many Internet measurement organizations and platforms are now providing data that were not available a decade ago.

Aceto and Pescapé wrote a survey of censorship detection systems in 2015 [4]. That work covered academic detection architectures as well as deployed Internet measurement platforms. It relied on the detection systems' design goals (as stated by their authors) for its characterizations, whereas this work focuses on evidence of censorship and characterizes it using the technical methods implemented by nation-state censors, within the historical context of how it occurred over the past three decades. While the section three of their paper lists many of the examples found in this survey, our taxonomy provides the comprehensive abstraction of censor methods, as well as benefits from ten additional years of research to build on.

Hall et al. have a draft informational request for comments (RFC) online, which serves to document ways in which regimes around the world block or impair Internet traffic [134]. The page is a helpful reference for designers, implementers, and users of Internet protocols to understand the properties of Internet technologies that are being exploited to censor end-user access to information. The RFC suggests that it is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress." Our work bolsters these efforts by surveying the academic literature to systemize our knowledge of censor methods, rather than relying on anecdotes.

Master and Garman conducted a worldwide study of nation-state Internet censorship [197]. Their study used a globally representative sample of 70 countries and found that 62 of these countries over a 20-year study period used at least one censorship method against their citizens. They used a preliminary version of our Internet Censorship Methods Taxonomy [196] as a basis for their study.

Khattak et al. performed a systematization of knowledge study on systems they termed "Censorship Resistance Systems" (CRSs). Their subjects included deployed software and theoretical systems from academic papers and "access-centric" and "publication-centric" schemes. The paper surveyed 73 such systems [167] and offered an abstract censor attack model. The authors focused on the goals of CRS designers, not necessarily on censor abilities directly.

Wendzel et al. also conducted a survey of Internet censorship and its measurement [314], similar to Aceto and Pescapé [4]. While their summary of censorship methods is robust, it is a precursor to their primary goal of characterizing censorship measurement methodology. We focus on censors themselves and the historical context surrounding those methods, as well as abstracting censor methods into a comprehensive taxonomy.

In a partially related study, Tschantz et al. performed an extensive survey of evaluation criteria for CRSs [289]. The authors compare the evaluation criteria elements from the literature to observed "real-world" censor behavior derived from field reports and bug tickets, revealing that they are often incongruent. Their goals were to enumerate the evaluation criteria and identify trends rather than to systematize the censor behavior.

To our knowledge, our study is the first to offer a taxonomy of Internet censorship methods. We focus on the technical means by which censors deny access to Internet communications over the past three decades. We provide a synthesis of our systematic literature review in the form of a taxonomy, which has pedagogical value for educators and new researchers to the field, as well as providing a research framework for future studies.

3 Methodology

The research question we address in this work is "What are the technical threats to Internet-based communications within the context of censorship?" To comprehensively address the topic, we used a systematization of knowledge (SoK) approach. SoK studies provide a comprehensive overview of the literature for a particular topic. These works save researchers significant time and effort in discovering the edges of the state-of-the-art.

Fink defines a systematic literature review as "a systematic, explicit, and reproducible method for identifying, evaluating, and synthesizing the existing body of a completed and recorded work by researchers, scholars, and practitioners" [108]. Systematic reviews and meta-analyses found early success in medical research [206], and other communities have since benefited from their insights and approaches. Research must move quickly as new technologies are developed, a challenge in rapidly advancing fields. In 2010, Okoli and Schabram published "A Guide to Conducting a Systematic Literature Review of Information Systems Research" to address the challenges IS researchers face [230]. We utilized Okoli and Schabram's approach to systematize the literature pertaining to Internet censorship.

In our study, Internet censorship methods are technical threats to Internet-based communications that a governmental authority or Internet Service Provider (ISP) can impose. The technical means by which Internet censors implement censorship vary widely but are generally constrained by the standardized protocols that allow the Internet to function: Internet Protocol (IP), Border Gateway Protocol (BGP), Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and QUIC³. Several application-layer protocols are covered in depth because of their importance to the World Wide Web. Thus, literature focused on technology was the primary source of data collection. Communications, political science, and social science journals were also referenced when the domains intersected.

Peer-reviewed journals and conferences were our primary data sources in the research protocol. In cases of emerging technology where scholarly work was not yet available, we utilized government reports, technical reports, and technology blogs. Data sets from Internet measurement communities were often central to these reports. Finally, if appropriate, we cited qualitative data from advocacy organizations and media reporting.

CensorBib⁴ was our starting point for data collection. CensorBib is a curated online archive of research papers in the field of Internet censorship [317]. Its focus is on papers that approach Internet censorship from a technical perspective. Given their predominance in the dataset on CensorBib, the following conferences and workshops were thoroughly surveyed: IEEE Security & Privacy, USENIX Security Symposium, ACM Computer and Communications Security, Network and Distributed System Security (NDSS) Symposium, the Privacy Enhancing Technologies Symposium (PETS), the Workshop on Free and Open Communications on the Internet (FOCI), and the ACM Internet Measurement Conference (IMC). All are venues for high-quality, technology-focused research. Papers were initially screened by title and abstract for relevance. We noted techniques and methods as we worked through the review. We also pursued citations within the references or bibliography section of identified papers, as applicable.

After reviewing the research publications, we constructed the initial elements of the taxonomy. Internet censorship happens on computing systems; using the OSI conceptual model [80] that describes computing system interconnection assisted in ensuring the comprehensiveness of the taxonomy. Each censorship method was grouped and abstracted based on the experience of the authors and informal peer review of researchers at our institutions. A second round

³According to RFC 9000 [187], QUIC is the name of the transport protocol and is not an acronym.

⁴Available at <https://censorbib.nymity.ch/>

of data collection then filled gaps and provided additional supporting sources for the taxonomy elements. We used keywords derived from the first round with Google Scholar, our institutions' digital libraries of subscriptions, and two search engines (Google and DuckDuckGo). This allowed coverage of scholarly sources not included in CensorBib or the preselected venues, as well as government reports and blog posts.

Limitations and Constraints. This study offers a global view of the issues related to Internet censorship. We acknowledge that this work is influenced by Western philosophy given that it originates in the United States and is also limited to sources available in English. In this paper, censorship will specifically refer to that experienced by Internet users, unless otherwise specified. Our research does not quantitatively measure the prevalence of any given censorship method; it uses a qualitative approach, with case studies and a comprehensive survey of research papers, to provide a synthesis of the literature (see §5). Research, in and of itself, can reflect biases of the researchers conducting it. Researchers may perceive relevance or novelty of a particular topic area, and more citations pointing to a particular censorship method (see Table 2) may not reflect widespread use or imply its potential for significant impact to online communications.

4 Survey of Internet Censorship Methods

4.1 HTTP/URL Filtering

Uniform Resource Locators (URLs) were one of the early targets for filtering Internet content. URLs are defined in RFC 1738 [195] and consist of a protocol (e.g., http, https, ftp), subdomain (optional), domain name, top-level domain, and potentially a directory path and file name (e.g., https://testdummy.com/uploads/index.html). URLs direct web browsers and other applications to specific Internet resources. URLs can be restricted using URL blocklists that define restricted sites. An administrator can also use URL regular expressions (regex) to filter based on sets of characters or words.

URL filtering is done for many reasons. Companies may limit the types of websites that employees can access while working: some materials may be deemed inappropriate or pose a distraction in a workplace setting. Filtering may also be performed for security if the company subscribes to cyber threat intelligence (CTI) [92] feeds that provide blocklists of domains associated with malware. Public resources, such as libraries and Internet cafes, may limit the websites patrons can visit to comply with the laws of their jurisdiction [79, 101, 115–117]. Depending on the country and the situation, filtering online content may not be considered censorship — but from a technical perspective, the software and tactics for filtering content are the same. Sambaluk noted that "Students of technology history have long understood that technologies do not have sympathies and do not play sides or favorites... and they seldom voice a preference about who uses them" [257]. This applies not only to privacy-preserving communication applications but also to cybersecurity controls. Controversially, multiple commercial products originating from the USA intended for commercial Internet filtering have been observed in use in nation-state censorship. For example, products such as Blue Coat have been identified in use in Burma, Egypt, Kuwait, Qatar, Saudi Arabia, and United Arab Emirates (UAE); McAfee SmartFilter has been used by Bahrain, Iran, Kuwait, Oman, Saudi Arabia, Tunisia, and UAE; and Netsweeper in Qatar, UAE, Yemen, and Pakistan [74, 75].

URL filtering can be performed locally or as traffic traverses a network. In the context of nation-state censorship, client-side filtering is rare, as managing configurations across numerous platforms (and gaining access to user devices) is difficult. Usually, a proxy, firewall, or middlebox performs network-based filtering upstream of many clients.

URL filtering relies on the ability of the censor to match the destination URL against a blocklist. Transport Layer Security (TLS) encrypts HTTP request and response contents, mitigating some basic URL-filtering mechanisms. Firefox Telemetry reported in October 2021 that 82% of global web traffic from the Firefox browser was encrypted using TLS [95]. Results for Japan and the United States were higher, at 86% and 92%, respectively. In October 2017, Google published a transparency report with an analysis of the top 100 most-accessed (non-Google) web pages globally, showing that

all 100 had a modern TLS configuration allowing use of HTTPS, and 95 of the 100 sites directed users to HTTPS by default without user interaction [122]. This change improves the security and privacy of all Internet users and presents challenges to censors that rely on simple URL-filtering mechanisms to block unencrypted websites.

Despite these developments, some important web pages are still served as unencrypted HTTP [147, 149]. The ease with which these connections can be manipulated or denied is seen as a problem by Internet freedom advocates. Efforts such as the "HTTPS Everywhere" campaign encouraged the use of TLS by default. Major web browsers (Google Chrome, Microsoft Edge, Mozilla Firefox, and Apple Safari) now have built-in functionality to upgrade connections where available [135]. In 2018, Troy Hunt launched a website to perform analysis on global web pages that do not automatically perform TLS redirection by default and offer resources to server administrators to help implement HTTPS at low or no cost [149]. Notably, projects such as *Let's Encrypt*, a non-profit certificate authority, provide X.509 certificates at no cost, regardless of which registrar maintains the domain [131]. Let's Encrypt was founded by the Internet Security Research Group and is credited with the widespread adoption of TLS over the past five years.

More advanced URL filtering approaches have also been observed. One way censors may disallow undesired HTTPS requests is by pattern matching against fields on the TLS certificate that are passed to the client. These will be transmitted in plaintext (unencrypted), and fields such as subjectAltName (SAN) or Common Name (CN) in the IP packet may contain portions of the URL or domain name [215]. However, this blocking method is prone to error. Certificate authorities differ in implementation as to what certificate fields contain, and shared domains could lead to over-blocking.

In addition to TLS, users in censored countries or filtered corporate environments have used proxies to evade filtering. Research has shown that at least tens of thousands of open HTTP/SOCKS proxies are available on the Internet to which any user may connect [249]. Many web browsers and Internet applications allow the manual entry of a proxy server. When users can initiate a connection to a proxy outside of the censor's influence, they may be able to bypass URL filtering; users may chain multiple proxies together to further obfuscate their traffic endpoints. However, these circumvention techniques can easily be thwarted if the censor requires that all connections to external Internet resources be negotiated through the censor's proxies.

A more robust censorship technique involves blocking based on Server Name Indication (SNI). SNI is an extension within the TLS protocol that helps a client to ensure that they are connecting to a web resource using the correct TLS certificate. This is necessary when a client reaches the IP address of a particular server but needs to ensure that it resolves to the correct host, especially on a server that hosts multiple domains. This scenario is commonplace on the modern Internet, with the rise of content delivery networks (CDNs) designed to place content topographically close to users to reduce latency [125]. SNI is transmitted in plaintext, and a censor with the capability to detect SNI using Deep Packet Inspection (DPI) can perform filtering based on it [197].

Several attempts have been made to increase the privacy of TLS. For example, in Fall 2018, Cloudflare proposed an extension to TLS named Encrypted Server Name Indication (ESNI). This extension would enable TLS 1.3 sessions to encrypt the SNI with a server's public key before sending it to the target [119]. The approach caught the attention of some censors, as research revealed that the People's Republic of China had begun blocking all TLS 1.3 connections utilizing ESNI crossing the Great Firewall [35]. ESNI was not implemented on many production servers, and Firefox was the only major web browser to support it. In a Mozilla blog post in January 2021, Kevin Jacobs announced that Firefox would no longer support ESNI in favor of a new draft RFC in development for Encrypted ClientHello (ECH) [154]. ECH would enable clients to encrypt the entire ClientHello of a TLS handshake (not only SNI), thus concealing the destination from an eavesdropping adversary. In 2023, Mozilla released a version of ECH in Firefox 118, and configured it to be on by default in Firefox 119 [207]. Mozilla's version of ECH relies on DNS over HTTPS (see §4.2). This design decision shifts trust of name resolution to third-party providers (e.g., Mozilla, Cloudflare, NextDNS). In 2024, Firefox had less than 5% market share for web browsers globally [276]; however, Google Chrome had the largest share in the




Fig. 1. Example plaintext SNI for cloudflare-dns.com; domain name shown in SNI extension within packet capture of TLS connection

user market in 2024.⁵ Google developers have implemented their own version of ECH into Chromium (the codebase of Chrome) as of version 117 and subsequently turned it on by default for non-enterprise installations [123]. Despite these advancements, until ECH is ratified as an open standard by the Internet Engineering Task Force (IETF) and adopted by web servers around the world, censors will still be able to use SNI destinations against some users for censorship and allow for mass surveillance of web browsing activities.


4.2 DNS Detection, Manipulation, and Denial

Another method widely used by censors is manipulation of the Domain Name System (DNS). As the "phonebook" of the Internet, the DNS standard (initially defined in RFC 882 and 883, superseded by RFC 1034 and 1035) matches human-readable domain names to IP addresses of web resources. DNS is an application layer protocol. DNS was not originally designed with security or privacy protections; the protocol is the topic of an entire body of literature on the DNS threat landscape, DNS research methods, and entities involved in DNS transactions [169]. Siby et al. have stated that "... virtually every connection to an Internet service is preceded by a DNS lookup" [270].

DNS requests pose several security and privacy considerations. Internet Service Providers (ISPs) can easily surveil users' Internet activity based on DNS. By default, most DNS requests are transmitted unencrypted. ISPs often provide DNS servers for customers to use and can monitor all requests. Even in cases with manually configured external DNS servers, ISPs can "hijack" [145] unencrypted DNS requests and return an IP address of their choosing to the client [181].

Suppose that a censor is only concerned with monitoring activity. In that case, they may transparently proxy the results of manually configured DNS servers, logging the requests, and passing the results back to the client unaltered. A user may believe they are resolving DNS queries using a server of their choosing when the censor has transparently proxied the request and forwarded it along on the user's behalf. These tactics can facilitate surveillance or censorship based on domain queries. If a censor wants to implement a blocking action based on lists of domains (and the observed DNS requests are not encrypted), the censor has many options available. Requests that match selected material can simply be dropped, resulting in timeouts for the client. The censor can also tamper with the DNS result, inserting an IP

⁵<https://gs.statcounter.com/browser-market-share#monthly-202401-202412>



※ 차단안내페이지(warning.or.kr)를 도용한 파밍사이트가 발견되어 각별한 주의가 필요합니다.
 (차단안내페이지는 개인정보를 요구하거나 프로그램 설치를 유도하지 않습니다.)

사이트분야	담당기관	전화번호
불법 도박	사이버 경찰청	1566-0112
	사형산업통합감독위원회	1855-0112
	사형산업통합감독위원회	1855-0112
불법 체육진흥부표권 판매	국민체육진흥공단 클린스포츠 통합콜센터	1899-1119
불법 승차부표권 구매대행	국민체육진흥공단 경찰 경찰총괄본부	1899-0707
불법 마권 구매대행	한국마사회	080-8282-112
불법 의약품 판매	식품의약품안전처 사이버조사팀	(043)719-1913,1921
불법 의약품 판매 및 허위과대광고	식품의약품안전처 사이버조사팀	(043)719-1921,1943
불법 식품 판매 및 허위과대광고	식품의약품안전처 사이버조사팀	(043)719-1912,1914,1919,1944
불법 건강기능식품 판매 및 허위과대광고	식품의약품안전처 사이버조사팀	(043)719-1912,1914,1919,1928
불법 화장품 판매 및 허위과대광고	식품의약품안전처 사이버조사팀	(043)719-1926,1943
불법 의뢰기기 판매	식품의약품안전처 사이버조사팀	(043)719-1903,1905
불법 마약류 판매	식품의약품안전처 마약정책과	(043)719-2806
성매매 및 음란	방송통신심의위원회	(02)3219-5848
성폭력(위조성폭력 유류)	한국지식재산보호원	(02)3183-5825
저작권(불법복제물 유류)	한국저작권보호원 온라인보호부	(02)3153-2460
안보위협행위	사이버 경찰청	1566-0112
명예훼손, 조상권 침해	방송통신심의위원회	(02)3219-5341-4
간혹 혐오, 차별 비하 등	방송통신심의위원회	(02)3219-5164
디지털성범죄정보	방송통신심의위원회	(02)3219-5834

◎ 운영자 아의상정 안내

사이트 운영자는 시청요구를 받은 날로부터 15일 이내에 방송통신심의위원회에 이의신청 할 수 있고, 「행정심판법」, 「행정소송법」 등 관련 법규에 따라 행정심판 및 행정소송을 제기할 수 있습니다.

(a) Korean



※ Farming site, which stole the interception page (warning.or.kr), is found and requires special attention.
 (Blockup does not require personal information or induce program installation.)

Site	Officer	Phone number
Illegal gambling	Cyber Police Agency	1566-0112
	Executive Committee	1855-0112
	Executive Committee	1855-0112
Sale of illegal physical education campaigns	National Gymnaest Clinicport integrated call center	1899-1119
Purchasing illegal winners	National Assembly Agency, Gyeonggi-gyeong, Gyeonggi-gyeong	1899-0707
Illegal ticket purchase	Korean Board of Directors	080-8282-112
Illegal drug sales	Food and Drug Administration Cyber Research Team	(043) 719-1913,1921
Sales of illegal medicinal products and false and large advertising	Food and Drug Administration Cyber Research Team	(043) 719-1913,1921
Illegal food sales and false and high advertising	Food and Drug Administration Cyber Research Team	(043) 719-1912,1914,1919,1944
Illegal health functional food sales and falsehoods and high advertising	Food and Drug Administration Cyber Research Team	(043) 719-1912,1914,1919,1928
Illegal cosmetic sales and false and high gloss	Food and Drug Administration Cyber Research Team	(043) 719-1926,1943
Sale of illegal medical devices	Food and Drug Administration Cyber Research Team	(043) 719-1903,1905
Illegal drug trafficking	Food and Drug Administration Department	(043) 719-2806
Prostitution and pornography	Broadcast Communication Committee	(02)3219-5848
Trademark (distribution of counterfeit goods)	Korea Ji-gyeong Jai-san-san-Kan	(02)3183-5825
Copyright (distribution of illegal drugs)	Korea Ji-Kyung-ho-won online	(02)3153-2460
Security	Cyber Police Agency	1566-0112
Honorary Yikweh Son, Portrait Violation	Broadcast Communication Committee	(02)3219-5341-4
Crusalty, hate, discrimination, etc.	Broadcast Communication Committee	(02)3219-5164
Digital crime	Broadcast Communication Committee	(02)3219-5834

◎ Operator Appeal Guide

(b) English translation

Fig. 2. Example blockpage from the Republic of Korea (Retrieved April 01, 2024)

address of their choosing and redirecting the client to it [183]. This could be used to display a "blocked content warning page" (referred to as a blockpage from this point forward), as some countries have demonstrated [160, 298]. See Figure 2 for an example blockpage. A more malicious adversary could use this tactic to redirect the user to download malware or present the user with disinformation or propaganda materials.

One countermeasure to DNS tampering is to manually define the DNS server for domain resolution. Naive censors will only filter requests they receive, and the user will be able to resolve IP addresses successfully. However, censor and filtering technologies have grown in sophistication — manually defined DNS servers are often not enough to avoid censorship. Even low-resource censors can block all DNS traffic except for their resolver, defeating this countermeasure.

Another countermeasure to DNS tampering is tunneling. Using an encrypted proxy (HTTPS or SOCKS) will obfuscate the request while it traverses the censor boundary, and the unencrypted DNS resolution will take place elsewhere on the Internet outside the censor's influence. Virtual Private Networks (VPNs) for personal use have increased in popularity to circumvent censorship and geographic area filtering for media content [166, 243]. Despite their initial purpose of providing an encrypted connection between two networks across the Internet, commercial providers are increasingly offering paid services to route traffic through their infrastructure. Several censors have taken notice. In particular, the People's Republic of China (PRC) banned non-state-sanctioned VPNs in 2018 [275], and the Russian Federation banned the use of several VPN services in 2021 [286]. VPN connections encrypt traffic between endpoints, allowing users to tunnel their traffic outside the censor's influence. Alternatively, they may use split-tunneling to selectively route only

some of their traffic, such as DNS requests. Although VPN users may be able to connect to content they otherwise could not access, they are shifting trust from their ISP to the VPN provider.

Efforts are underway to improve the security and privacy of DNS as a whole. As early as 1997, work began on DNSSEC (Domain Name System Security Extensions) as a means to add authentication to DNS. DNSSEC adds public-key cryptography to the DNS lookup process. DNSSEC-compliant servers digitally sign DNS records and rely on a chain of trust similar to TLS certificate authorities.⁶ DNSSEC ensures DNS replies have not been manipulated, which could serve as a countermeasure to DNS cache poisoning. However, it does not provide confidentiality, as the responses are still transmitted in plaintext and are subject to other censorship techniques. In 2013, researchers identified that only 0.15% of the top-level domains of ".com" used DNSSEC [180]. Rijswijk-Deij et al. demonstrated how sites using DNSSEC could also be susceptible to packet fragmentation or amplification-based DoS attacks on their DNS traffic [294].

Encryption of DNS can address some of the privacy and authenticity problems of the original protocol. Many public DNS services offer DNS over TLS (DoT) and DNS over HTTPS (DoH). Both have been used as countermeasures to DNS tampering. DoT and DoH encrypt DNS requests in transit and bypass the DNS resolvers of the service provider gateways [270]. Several web browsers support DoH natively. However, there are challenges with implementation. Many Internet of Things (IoT) devices do not support encrypted DNS methods. The backward compatibility issue with many devices, firewalls, and routing equipment currently in production environments will also limit the adoption of encryption for DNS requests. Encrypting DNS also limits the visibility (and potential functionality) of cybersecurity tools that organizations may use to secure their enterprise. Additionally, DoT/DoH server endpoints have been targeted by censors in China, Iran, and Kazakhstan [24]. When used in isolation, DoT and DoH do not fully address Internet privacy concerns, as censors can target other vulnerable aspects of a DNS request. However, DoT/DoH may be helpful if a nation-state only implements its censorship regime via domain name blocklisting.

4.3 IP Address and Port Matching, Blocking

A common form of Internet censorship is blocking based on Internet Protocol (IP) addresses. Censors can maintain blocklists of IP addresses associated with objectionable endpoints. Censor systems monitor traffic and drop connections when a blocked address is observed crossing the censor's boundary. Third parties can provide blocklists, or the censor can maintain their lists based on other content-matching efforts. IP address blocking can be resource intensive, as blocklists may grow large and become outdated as a result of the ephemeral nature of IP addresses [328]. With the growth in the adoption of IPv6, these limitations will continue to increase. Despite these constraints, some nation-states choose to censor Internet traffic using IP blocklists — likely because of ease of implementation.

IP address spoofing can serve as a countermeasure to bypass censorship efforts. The IPv4 protocol lacks authenticity mechanisms in that senders (and network address translation nodes along a traffic path) assign the packet source address field [26]. If a censor implements blocking based on source IP address, spoofing the address of an allowed endpoint may allow traffic that would otherwise be denied. However, source blocking is seldom used for large-scale censorship; censors often focus on the content to which they wish to deny access and apply blocking against all users within the censor's influence. Accordingly, few deployed anti-censorship tools attempt to implement IP address spoofing.

Censors may also choose to block Internet traffic by ports. Port blocking can be paired with IP address blocking or is done indiscriminately across all traffic. Many applications openly advertise their software's ports and transport protocols, and censors may wish to deny access to those specific applications, such as when the Burmese government blocked access to Skype in 2011 [201]. A more advanced blocking regime in Iran involving port blocking was revealed in 2020 by Bock et al. and the Open Observatory of Network Interference (OONI). Four Iranian ISPs were shown to have blocked port 853, associated with DNS over TLS, to many popular DNS providers [23] (in addition to tampering

⁶DNSSEC extensions are defined in RFCs 4033, 4034, and 4035, and summarized in RFC 9364, available at <https://datatracker.ietf.org/doc/html/rfc9364>

with TLS handshakes to those providers). Iran had attempted allowlisting particular ports (80, 443, 53 — HTTP, HTTPS, and DNS, respectively) in 2013 while blocking all other ports [18] before abandoning the effort in the same year. In 2020, Bock et al. demonstrated that Iran had reimplemented what they called a "protocol filter" prior to their censorship apparatus. It filtered out all traffic except HTTP, HTTPS, and DNS before filtering based on content or destination [37].

Traditional port blocking is an approach relying on software following common conventions. This is effective against applications that explicitly follow RFC specifications or their own conventions. However, transport layer protocols can easily use non-standard port mappings, and many applications do so. Most censors are unwilling to rely on allowlist filtering, given the large amount of collateral damage that would result in overblocking (Iran being a notable exception).

A 2021 study [36] highlighted "residual censorship," where censors detect an objectionable connection using one censorship method, then proceed to deny all connections between the two endpoints for a short duration using a 3-tuple (client IP + server IP + port) or 4-tuple (client IP + port + server IP + port). Bock et al. observed this time-based approach to IP and port blocking in China, Iran, and Kazakhstan. More research is needed to determine if other nation-states are implementing similar functionality in their censorship systems.

Using high-number ports is a simple yet detectable countermeasure against censorship. By altering the listening port of a web application or service, censors that only filter based on static port numbers may be bypassed. This technique does not bypass censors that block/reject all traffic except for particular common ports and those that use allowlisting.

Censors may also use hybrid methods to deny undesirable traffic. By injecting an artificial TCP Reset (TCP-RST), network operators can cause an endpoint to shut down communication upon receipt [224, 310], as defined in RFC 793 and RFC 1122 [41, 77]. It is a normal element of TCP/IP functionality for connection timeouts and non-existent TCP endpoints, among other uses. Censors can pair identification techniques (such as those discussed in §4.1) with a TCP Reset to obscure the fact that censorship is occurring. Users may attribute a web browser message or software error to poor network service rather than deliberate censorship. TCP Resets have also been misused to orchestrate distributed denial of service (DDoS) attacks⁷; we discuss these in detail in §4.8.

The complexity of the Internet has complicated the job of administrators ensuring the availability of systems as well as censors hoping to deny access to particular content. An example is content delivery networks (CDNs), designed to provide access to web resources as close to the user as possible to reduce latency [125]. Censors have caused extensive collateral damage in overblocking by denying access to the IP address space of large CDNs [200]. A particularly salient example of inadvertent censorship is the PRC's blocking of Google Scholar. Despite it being considered a valid academic resource and legal service by CCP governmental regulators, the Great Firewall filters traffic destined for Google.com based in the U.S. [184]. Legal frameworks and technical implementations change over time and often conflict.

Tunneling has repeatedly demonstrated its effectiveness in circumventing IP and port blocking. Because the encrypted tunnel obfuscates the metadata of the IP packets, many blocking efforts are thwarted. However, encrypted proxies, VPNs, and anti-censorship tunnels still suffer from the bootstrapping problem [316] — motivated censors will use their resources to prevent the tunnel's initial establishment. These techniques often include simple endpoint IP or port blocking, DNS filtering, and deep packet inspection approaches.

4.4 Deep Packet Inspection and Protocol Fingerprinting

In response to the success of early circumvention tactics, some censors sought advanced capabilities to detect undesirable Internet traffic. As the name suggests, deep packet inspection (DPI) allows searching of payload content in data packets [269]. In an IP packet, the IP header is followed by the packet's payload. A payload consists of the transport, session, presentation, and application (layers 4-7 of the OSI model) data of the communication. DPI is frequently used in cybersecurity contexts by system administrators. DPI-enabled network devices can "be programmed to detect certain

⁷<https://www.akamai.com/glossary/what-is-a-tcp-reset-flood-ddos-attack>




Fig. 3. DPI Diagram from Bendrath and Mueller [25]

bit sequences associated with known malicious code" [127] from large lists of indicators of compromise (IOCs) provided by security tool vendors, threat intelligence sharing feeds, or freely available online. However, any tool that assists with identifying or classifying traffic can also be used to identify content for censorship.

In the context of censorship, DPI allows for content matching, such as with keywords of prohibited materials. While still applicable for Internet censorship, the large-scale proliferation of encryption across many Internet resources [95] has rendered simple content matching obsolete in many protocol exchange scenarios. Additionally, many anti-censorship solutions implement encryption at multiple layers to protect traffic confidentiality and thwart simple content filtering.

DPI is often considered expensive in terms of computing costs, as well as challenging to implement on high-speed network systems [293]. However, the potential of DPI based on heuristic or machine learning techniques remains. Open-source DPI tools are increasing in sophistication and accuracy. The project ntop maintains a software fork of OpenDPI named nDPI, which includes the flow signatures of hundreds of protocols and applications [226]. Some commercial firewalls (e.g., Palo Alto, Fortinet FortiGate, Forcepoint) already claim to be able to filter data from applications on a network [110, 111]. Palo Alto has a blog post guide for their customers on all the steps to block Tor using their product [216]. The most aggressive nation-state censors have shown an increased willingness to implement DPI [197].

Some anti-censorship software developers and researchers have implemented countermeasures against deep packet inspection. One approach is to completely randomize the payload to make the traffic "look like nothing" [87]. This is the approach taken by "obfs" tools such as Dust [315] and ScrambleSuit [319]. These pluggable transports were initially created for integration with Tor or the Tor Browser, allowing some censored users to overcome censor blocks on their first hop. This approach makes inherent assumptions about the adversary in that they only operate against blocklists. If a censor uses an allowlist to define "known good," randomization methods fail.

Other approaches involve mimicking the payloads of other data protocols (dubbed "mimicry" approaches), such as StegoTorus [313] for HTTP, SkypeMorph [205] for Skype, and CensorSpoofer [304] for SIP-based Voice over IP. This approach has been criticized by Houmansadr et al. as "fundamentally flawed," as any slight variation not imitated by the spoofing software can allow a determined censor to detect the anomaly and block the traffic [143]. Balboa works as middleware and attempts to overcome the shortfalls of mimicry approaches by using standard outputs of existing software and injecting traffic into allowed TLS streams [252]. Newer approaches, such as format transformation encryption (FTE), claim that observation-based FTE steganography can be used as an undetectable covert channel [227].

Overall, pluggable transports have become important tools supporting bootstrapping of anti-censorship tool connections. Projects other than Tor have begun to integrate pluggable transport compatibility in their software [228].

Advanced censors have taken additional steps to deny access to anti-censorship, VPN, and other encrypted communication methods. Active probing by the Great Firewall of China has been shown to routinely target Tor bridges [318], as well as SSH [222] and VPN endpoints [223]. Active probing involves the censor masquerading as a user and attempting to connect to the circumvention service. The probe observes how the service responds. If the censor's probe determines

that the server is running an "undesirable" service, the censor can take a blocking action, such as adding the server to an IP blocklist. The "cat-and-mouse" game has continued, as researchers have studied active probing methods and offered solutions to anti-censorship developers to evade active probing [112, 114].

Tor Browser uses obfs4 as its default pluggable transport. Previous versions of obfs were vulnerable to active probing by the GFW of China. Every obfs4 server has a per-bridge secret, and a client must prove knowledge of this secret before gaining access to the protocol. This defeats automated methods of active probing as the censor would need the same out-of-band information as a legitimate client to detect a hidden Tor bridge [106]. Another pluggable transport approach, Meek, uses a routing method known as domain fronting [105] to direct first-hop traffic under the cover of HTTPS to a CDN provider domain that is unlikely to be filtered. Psiphon and Tor include Meek implementations. However, domain fronting has faced significant challenges; developers who relied on Google or Amazon CDNs for cover traffic no longer function after those companies reconfigured their cloud implementations to prevent domain fronting [199].

Some newer anti-censorship tools leverage techniques to bypass less aggressive censors that only use DPI blocking (and no other combinations of network filtering). Software such as GoodbyeDPI [121] and a related fork named PowerTunnel [239] allow users to access DPI-blocked content without other privacy protection mechanisms. These tools are sometimes referred to as "serverless" because they do not rely on an external entity for their functionality. The software creates a local proxy server on the client machine and funnels/manipulates HTTP(S) traffic before it is transmitted over a network interface. Several clever manipulations, such as TCP-level fragmentation of first packets, replacing "Host" header values, adding additional spaces in URIs, mixing letter cases of header values, and fake TTL/sequence/acknowledgment numbers have been demonstrated to fool some DPI implementations.

Encryption does not solve all DPI-based censorship problems; many open research problems related to censorship assume an adversary with advanced DPI capabilities. Some advanced DPI systems characterize traffic from particular applications or protocols, even when the traffic is encrypted. These "fingerprints" or signatures can then be used to determine when users are attempting to use a specific piece of software, despite encryption. One identification technique involves examining the initial unencrypted initiation of an encrypted network stream. Protocol header metadata can often reveal the types of encryption implemented or versions of software applications. A comparison of the available DPI systems that perform this analysis is available at [45].

More sophisticated methods of ETA involve identifying particular byte sequences or patterns in overall encrypted traffic streams to allow for the identification of an application or protocol. The methods may rely on statistical models to identify behavioral changes in bitstreams. Other approaches involve training a machine learning (ML) classifier. The algorithm uses a known "baseline" compared to a bitstream containing a targeted protocol's traffic to observe patterns or identify anomalies. A survey of methods for encrypted traffic classification by Velan et al. is available in [296].

An example of pattern or behavioral-based blocking is the denial of secure messaging apps such as Signal [161]. Blocking of Signal has been observed in Iran, China, Cuba, and Uzbekistan [336]. More commonly, these nations use techniques such as DNS manipulation (Iran, China) or DPI targeting SNI fields in a TLS handshake (Cuba, Uzbekistan), likely because they require fewer resources. In 2021, the Russian Federation demonstrated hybrid censorship approaches using SNI targeting with DPI and bandwidth throttling to achieve censorship goals against Twitter, discussed in detail in §4.6.

4.5 BGP Attacks and Disruption

Border Gateway Protocol (BGP) is one of the critical components that allows the Internet to function as a network of networks. Several protocols operate at layer 3 of the OSI model, allowing connectivity among devices worldwide. At the highest levels, the Internet Assigned Numbers Authority (IANA) assigns blocks of IP addresses to regional Internet

registries (RIRs). RIRs allocate blocks to organizations such as Internet Service Providers (ISPs), technology companies, universities, government agencies, and scientific institutions. These large organizations often serve as an Autonomous System (AS) for the purposes of routing Internet traffic. Each AS is assigned an autonomous system number (ASN), which is unique and used by BGP to determine how ASes communicate with one another. For example, AS17 is assigned to Purdue University, which manages the following sets of IPv4 addresses: 128.10.0.0/16, 128.210.0.0/16, 128.211.0.0/16, 128.46.0.0/16, 149.164.0.0/16, 163.245.0.0/16, 192.31.0.0/24, 204.52.32.0/20, 204.52.48.0/20, 205.215.64.0/18, 69.51.160.0/19, and also IPv6 blocks 2001:18e8:800::/44 and 2607:ac80::/32 [281].

There are approximately 64,000 ASNs in use worldwide, with continuous growth [64], and BGP (versions 4 and 6) are the protocols that allow them to communicate with each other. "DNS tells you where you're going, and BGP tells you how to get there" [59]. Although BGP is intended to provide the most efficient routing options available between ASes, there are many associated factors beyond counting the lowest number of router hops between one source IP address and its destination. Factors such as cost of use and physical distance of transport media complicate routing decisions.

In addition to routing complexity, the design choices made when creating BGP have implications for the security and availability of the global Internet. The protocol assumes that advertisements from interconnected networks are truthful about the IP addresses they own. "BGP hijacking" attacks are nearly impossible to prevent [63]. An attacker may choose to advertise a route for a smaller range of IP addresses than other ASes have previously announced. Another attack involves offering a shorter route to particular blocks of IP addresses. Both allow the attacker to divert large swaths of Internet traffic in a direction of their choosing. To conduct a BGP hijack, the attacker must be an AS administrator or take control of an AS system. This limits the scope of who might perform such attacks but does not rule out nation-states or large technology corporations from using these methods to deny access to information.

There have been several BGP-related incidents. In 2004, a Turkish ISP named TTNNet (AS9121) inadvertently modified its configuration to include a full table of entries (>100k). This modification effectively advertised AS9121 as the shortest route to most of the Internet's addresses [291]. For several hours, users worldwide could not access some or all websites because of routing errors. In February 2012, a network operator named Dodo announced internal BGP routes to Telstra, a major ISP in Australia. Telstra incorrectly accepted the routes and caused routing bottlenecks throughout the country. This error resulted in most Australians losing online connectivity for 30 minutes. A similar outage event occurred in August 2012 in Canada, known as the "Bell-Dery" routing leakage incident, after Dery Telecom leaked a full BGP table of routes to Canadian ISP Bell [27]. In 2021, Facebook (now Meta) suffered a global outage partially caused by BGP route errors. During routine maintenance, an engineer accidentally disconnected Facebook's backbone data centers. Failover systems, unable to communicate with backbone infrastructure, withdrew IP address blocks from their BGP advertisements. As BGP routes propagated and DNS server caches worldwide expired, all of Facebook's subsidiary services — Instagram, WhatsApp, Messenger, Mapillary, and Oculus — were inaccessible [155]. The company restored services in six to seven hours but at enormous business and advertising costs.

Malicious BGP-related incidents have also occurred. In 2008, Pakistan Telecom (AS17557) allegedly published an unauthorized BGP announcement of the prefix 208.65.153.0/24 to attempt to censor YouTube from its citizens [214]. This effort temporarily hijacked all YouTube traffic on a global scale, denying access to all users for part of the day on 24 February 2008, until appropriate routes were restored. In 2011, the government of Libya used BGP routing to implement an Internet shutdown on most of its country's users in response to a series of uprisings in the region. A single state-operated AS routed most traffic in Libya at the time, allowing it to withdraw routes and disrupt network connectivity [73]. On February 5, 2021, after the Burmese military overthrew the government of Myanmar in a coup d'état, the junta compelled Campana Mythic (AS136168) to announce the 104.244.42.0/24 prefix, which belonged to Twitter. The junta may have intended to blackhole all Twitter traffic within their borders. The route leaked to the global Internet, making AS136168 appear to own the IP address space, disrupting traffic in Singapore and other

Southeast Asian nations [233]. Other examples have resulted in monetary losses due to criminal activity. In 2018, Russian AS48693 and AS41995 advertised IP blocks that belonged to Amazon Route53 DNS services from the provider eNet (AS10297) of Columbus, Ohio, USA [238]. Criminals forged a fake version of a web-based Ethereum wallet service⁸ and stole approximately \$152,000 worth of cryptocurrency. In 2022, a sophisticated supply chain attack on a Korean cryptocurrency exchange platform (involving BGP hijacking perpetrated by AS9457) also resulted in the theft of approximately \$1.9 million (₩2.37 trillion KRW) [56].

The threat of abuse with BGP manipulation has motivated proposals for security enhancements [204]. However, BGP attacks have not been used in practice as an intentional technique for routine censorship. Instead, other methods with less collateral damage are used. There are few countermeasures against BGP-related disruptions, given the design of the BGP protocol. International economic and political pressure dissuades malicious actors from abusing BGP routing.

4.6 Bandwidth throttling

Bandwidth throttling involves a deliberate slowing of communication speed. In the context of computer networks, throttling can occur at the application software or network management levels. Throttling Internet traffic has been considered for use during periods of high usage. During the global COVID-19 pandemic, governments in Europe and technology companies in the U.S. recommended using bandwidth throttling and traffic shaping techniques to manage the large-scale throughput increases brought on by stay-at-home orders and remote work requirements [162]. However, throttling often arises in debates surrounding "net neutrality," a network design principle that states that all data packets should be treated equally regardless of their content, sites, and platforms [327]. Network neutrality is a contentious topic. It has been interpreted and implemented differently and has faced legal challenges in multiple countries. Proponents of network neutrality claim that traffic on the open Internet should be freely accessible, while opponents claim that prioritizing particular services creates a better user experience.

In the context of censorship, bandwidth throttling has been used to suppress access to information. Throttling provides some deniability on the part of the censor, differentiating it from other methods of Internet censorship. In Iran, the government slowed Internet speeds nationwide in the lead-up to a presidential election in 2009 [12] and 2013 [98], as well as other periods of political unrest [49]. The Russian Federation has deployed selective throttling [334]. In March 2021, the Russian government started limiting the speed of connections to Twitter. Russia had previously requested that Twitter remove content it deemed illegal or objectionable, and until then Twitter had refused to comply. Roskomnadzor issued a statement confirming that it ordered the throttling of Twitter traffic until the company complied with its approximately 28,000 removal orders [253]. Xue et al. discovered that Russian ASes used deep packet inspection to detect SNI within TLS connections specifically for Twitter and its subdomains (twitter.com, t.co, *.twimg.com). If detected, connection speeds were reduced to between 130–150 kilobits per second (kbps). Twitter eventually removed 91% of the requested content, and some of the throttling measures were lifted [334].

Taye described several techniques that governments use to throttle Internet connectivity [208]. First, traditional bandwidth management, such as traffic shaping and policing, can slow objectionable traffic. This can be done based on source or destination IP ranges, virtual local area networks (VLANs), or media access control (MAC) addresses. Second, censors can implement quality of service (QoS) restrictions in network traffic. QoS is typically used to prioritize low-latency traffic, such as voice calls. For censorship, QoS can deprioritize particular kinds of undesirable traffic. Third, inline DPI systems can be used to artificially add latency. Finally, censors can manually alter routing paths at their border gateways to limit traffic. They can change routing paths to be longer, adding latency. Or censors can direct specific traffic through lower capacity links, serving as a chokepoint to delay or halt traffic flows.

⁸<https://www.myetherwallet.com/>

There is currently little anti-censorship research dedicated to the study of bandwidth throttling. There are no purpose-built tools to counteract this kind of censorship method. Anecdotes online suggest that commercial VPNs may bypass ISP throttling, depending on how the throttling is implemented [109]. Xue et al. offered several countermeasures against Russia's throttling of Twitter, such as TCP-level fragmentation and TLS packet stuffing — or more traditional anti-censorship methods, such as encrypted proxies or VPNs — which defeated Russia's detection method [334]. More research and development are needed in this emerging area of censorship studies.

4.7 Internet Shutdowns

Even as societies around the world increasingly rely on Internet connectivity for economic, financial, and social functions, some governments have increased efforts to deny Internet access to their citizens. Feldstein defines Internet shutdowns as "activities undertaken by states to intentionally restrict, constrain, or disrupt internet or electronic communications within a given geographic area or affecting a specific population in order to exert control over the spread of information, within a timebound period" [102]. Authorities have used network disruptions to "quell mass protests, forestall election losses, reinforce military coups, or cut off conflict areas from the outside world. Data from the past few years documented incidents of global shutdowns: 196 documented incidents in 2018, 213 incidents in 2019, and 155 in 2020... [and] Government-instigated internet shutdowns largely took place in relation to five event types: mass demonstrations, military operations and coups, elections, communal violence and religious holidays, and school exams" [103]. Shutdowns may be implemented for minutes, hours, or more. Authorities may assert that their sovereignty grants them the right to control Internet access as they see fit; others cite the need to counter threats to public order, challenges to national security, or moral degradation of their populace. These justifications are often incongruent with international norms. The United Nations Human Rights Committee has stated that "The right to access and use Internet and other digital technologies for the purposes of peaceful assembly is protected by Article 20 of the Universal Declaration of Human Rights and Article 21 of the International Covenant on Civil and Political Rights" [68, 69]. Protection of freedoms of expression and association is often cited as obligations within international human rights law. A plurality of nations across the ideological spectrum agree, but not always in practice. Internet shutdowns often lead to unintended consequences for Internet censors and society at large. From an economic perspective, researchers estimated \$5.62 billion dollars of financial impact on the global economy in 2021 as a result of Internet shutdowns [322].

One of the most cited incidents of Internet shutdown efforts occurred from December 2010 through the end of 2012 in western Asia and northern Africa. Known as the "Arab Spring," these events were a series of anti-government protests, uprisings, and armed rebellions in Tunisia, Egypt, Libya, Yemen, Syria, and Bahrain. The overthrow of several heads of state and some governmental reforms occurred as a result. Demonstrations took place in the streets of Algeria, Iraq, Lebanon, Jordan, Kuwait, Morocco, Oman, and Sudan, while minor protests also occurred in Djibouti, Mauritania, Palestine, and Saudi Arabia [255]. Citizens (often youth or union groups) protested the corruption of political leaders, governmental abuse of power, economic stagnation, extreme poverty, and unemployment. The more authoritarian regimes among these nations took multi-pronged approaches to quell rebellion, including police and military mobilizations against their populations. Social media platforms were cited as a means of mobilizing collective action during Arab Spring protests, and circumventing traditional media and means of communication typically controlled by censoring states [19, 324]. According to some, the impact of Facebook and Twitter during the uprisings has been overstated by popular media [44]. However, several regimes deemed Internet connectivity important enough to the protesters' causes to selectively deny online access to prevent communication or the spread of information [321].

In 2013, Shavitt and Zilberman published an analysis of Internet shutdowns during the Arab Spring, specifically in Egypt, Libya, and Syria [266]. Using combinations of BGP advertisements and altering default IP routing, each country (at different times and to varying degrees) effectively stopped meaningful resolution of TCP/IP connections throughout

regions of their country. Attempts to deny all access to Internet endpoints were relatively straightforward, especially in a small nation such as Syria, which had only one autonomous system controlled by the government [51]. The denial of Internet access may have hampered some protestor mobilization efforts while simultaneously fueling the discontent of Arab youth by disrupting their access to Internet content.

Internet shutdowns have occurred in many countries since the Arab Spring events. Censoring nation-states have diverse geopolitical and social systems and widely varying legal systems. However, controlling access to information is increasingly essential to shaping political narratives. India, the largest self-proclaimed democracy, had 132 regional Internet shutdowns in 2020 [50]. Politicians in India use shutdowns as a tool of governance under the auspices of countering misinformation and disinformation (sometimes termed "fake news"). These leaders cite national campaigns such as "Digital India," intended to use connectivity to promote prosperity and wealth accrual. Shah argued that their efforts do not stop fake news stating "...as infrastructural tools, they enable state (dis)information and propaganda to spread without resistance and thus become potent tools in curbing protests and rightful critique of authoritarian practices" [265]. Shah also argued that given the patchwork nature of ISP infrastructure and implementation, "shutdowns as a way of regulation and governance are both ineffectual and counterproductive" [265].

In 2019, in response to large anti-government demonstrations against rising gasoline prices, Iranians experienced a several-day Internet blackout [217]. In November 2020, Myanmar held national elections in which the candidate for the National League for Democracy Party, Win Myint, won the presidency. The Tatmadaw (Burmese military forces) staged a coup d'état on 01 February 2021, rejecting the election outcome and installing a military junta in place of Myanmar's elected government [256]. In addition to violence, search and seizures, arrests of dissidents, and legal changes to quell dissent, the government blocked social media platforms (Facebook, Twitter, and Instagram) to limit access to information and assembly. Restrictions escalated as protests increased, and nightly shutdowns occurred throughout February and March. Censorship efforts culminated in a near-total Internet shutdown on 02 April 2021 [156]. The junta ordered telecommunications companies to disconnect connectivity to mobile and wireless Internet customers, the only services available in the country to access the Internet. Although Internet penetration rates in Myanmar were only 35% in 2020 [22], the censorship measures implemented are some of the most draconian observed in this century and have spurred debate about free expression, freedom of assembly, and free access to information [212].

Researchers have noted that despite the recent increase in Internet shutdown events, many governments have opted to selectively block (or throttle) individual content or platforms [103]. Rather than total shutdowns, authorities have attempted to avoid a more potent political backlash from their citizens or the international community with selective filtering. More sophisticated censorship methods — as described in previous sections — sometimes offer ambiguity or minimize the likelihood that users will attribute their lack of access to censorship. Network throttling may be blamed on slow or inconsistent ISP connectivity, targeting of a social media platform in short bursts may be viewed as a platform's technical issue, and URL/DNS filtering (that does not present a blockpage but an HTTP error or TCP timeout) might be shrugged off as poor cellular service. Time-based and persistent Internet censorship are important trends for researchers and activists to monitor. As nation-states demonstrate a willingness to use hybrid methods and timing to achieve their goals, anti-censorship software developers [86, 150, 178, 236, 302, 316, 331] remain adaptive in their approaches to circumvent censorship where it occurs.

4.8 Computer Network Attacks and Resource Exhaustion

The study of computer network attacks (CNA) and exploitation has been the subject of cybersecurity research over the years. Subtle methods have been used for cyber espionage, in which a nation-state or non-state actor infiltrates an adversary's network to obtain information for economic, military, or political gain. Subtle and overt methods of cyber attack are also used in the conduct of cyberwarfare, in which one actor destroys, denies, disrupts, deceives, degrades, or

manipulates [66] their targets' computing systems in pursuit of their objectives. Private sector actors often focus their efforts on the World Wide Web to address denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks. Both are forms of resource exhaustion that deny access to services by overwhelming target servers.

DDoS attacks have dominated the computer security literature over the past several decades [182, 189, 190, 203, 287, 341], given the scope and scale of the attacks and their economic impact. The U.S. Cybersecurity and Infrastructure Security Agency (CISA) defines a DoS as an attack that "occurs when legitimate users are unable to access information systems, devices, or other network resources due to the actions of a malicious cyber threat actor. Services affected may include email, websites, online accounts (e.g., banking), or other services that rely on the affected computer or network. A denial-of-service condition is achieved by flooding the target host or network with traffic until the target cannot respond or crashes, preventing access for legitimate users" [57]. DDoS attacks are similar in intent but utilize a distributed system of many computers as the source of the superfluous traffic. This allows the attacker greater anonymity and larger amounts of attack traffic, making a defender's job of mitigating the attack more difficult. Participants in DDoS are often unwitting, as their devices may have been infected with malware that surreptitiously joined the attacker's botnet.

From a technical perspective, DDoS attacks can be broadly classified into network-layer attacks and application-layer attacks. Network-layer attacks generally involve flooding a victim's network bandwidth to exhaust its resources; traffic could be UDP streams, ICMP flooding, DNS traffic flooding, VoIP flooding [11], or other allowed protocols. Application-layer attacks, higher in the OSI model, are protocol-specific and exploit a bug or particular feature of a protocol implementation to consume resources; TCP SYN flooding, TCP SYN-ACK flooding, ACK & PUSH ACK flooding, or RST-FIN flooding [341] are examples. Reflection-based attacks involve spoofing the source of a message, so not only does the victim interact with the flood of traffic, but also the spoofed source, which is now involved when the victim replies. Amplification-based attacks rely on external services to help create additional traffic, "amplifying the message" [287]. They may exploit an IP broadcast feature of a network device or service to reduce the effort the attacker must expend to attack their targets. Reflection and amplification attacks are present in the application layer but occur using protocols such as HTTP or database queries (e.g., SQL). An analysis of application-layer DDoS attacks is available in [287], as is a more general taxonomy of DDoS attacks and defense mitigations in [203].

DDoS attacks have not been observed as a routine method of Internet censorship by nation-states against users. This is likely because of the nature of the attack techniques used. DDoS attacks are effective against large, centralized services and may deny access to the network for all users. Other methods, such as Internet shutdowns, would be used instead for censorship. Zargar et al. characterize the motivations of DDoS attackers into five categories: financial/economic gain, revenge, ideological belief, intellectual challenge, and cyberwarfare [341]. Attribution of blame for the conduct of DDoS attacks is also a difficult problem. Nation-states, governmental proxy groups, criminals, and Internet activists (or "hacktivists") [260] have used DDoS attack techniques to achieve their goals.

There are dozens of examples of DDoS attacks linked to cybercrime or cyberwarfare since the early days of the world wide web. In 2000, a fifteen-year-old using the online handle "Mafiaboy" (he was later revealed to be a Canadian named Michael Calce) compromised computer systems at several universities and used them to conduct DoS attacks. The attack brought down the websites of Yahoo!, CNN, eBay, Dell, Amazon, and E*Trade, and the litigation that followed inspired first-of-their-kind cybercrime laws across several countries [42, 139].

In April 2007, the national government of Estonia was the victim of a campaign of DDoS attacks targeting government services, banking systems, and media outlets. The attacks roughly coincided with the relocation of a memorial statue, the "Bronze Soldier of Tallinn," which was perceived by ethnic Russians living in Estonia and Russia as an affront to their sacrifices during World War II [60]. The economic damage resulting from system downtime and societal disruption was estimated to cost tens of millions of Euros [309]. The campaign against Estonia has been cited as the second

state-sponsored act of cyber warfare (the first being a series of coordinated cyberspace operations, dubbed Titan Rain, against the U.K. Defense Ministry and U.S. Defense Intelligence Agency (DIA) beginning in 2003 [70]).

In 2008, coinciding with an armed attack against Georgia by the Russian Federation, large-scale DDoS attacks began against Georgian web pages [283]. The targets included the Georgian president's website, government ministry pages, and news agencies. While the Russian government may have benefited from the disruption, evidence suggests that the attack resources were crowdsourced from underground hacker communities sympathetic to the Russian cause [283]. There is no published evidence of whether the attackers were encouraged, endorsed, coordinated, or resourced by the state. Other large-scale DDoS incidents were documented in 2013 against Spamhaus (an email spam tracking organization), in 2015 and 2018 against GitHub (an open-source software distribution platform), and in 2016 against Dyn (a major DNS provider) [65]. The Mirai botnet targeted Dyn using compromised Internet of Things (IoT) devices to attack its victims, denying Internet access to many users on the east coast of the United States on 12 October 2016.

In 2017, the U.S. government released a report documenting malicious cyber activity performed by the government of North Korea. The report outlined the activities and capabilities of the Lazarus Group (aka HIDDEN COBRA, or APT38⁹ by Mandiant). The organization was the same group responsible for global computer disruption with their release of the WannaCry malware in 2016 [17, 128]. Among the malware analysis and IOCs in the report, a tool called DeltaCharlie was revealed to have been used for DNS, NTP, and carrier-grade NAT DDoS attacks [58]. Lazarus group likely perpetrated DDoS and offensive cyberspace operations campaigns against South Korean media, financial institutions, and critical infrastructure in 2011 and 2013 [342].

In 2020, Google disclosed that in 2017 its cloud services infrastructure was the target of a 2.54 terabyte per second DDoS attack originating from four Chinese ISPs [65, 202]. A series of attacks, including the "Mantis" attacks, were targeted at Google infrastructure [171, 339]. Another attack against Google cloud services in August 2023 reached a peak of 398 million requests per second, nearly 7.5 times larger than the 2020 campaign [170]. Cloudflare reported mitigation of the largest DDoS attack recorded in history in 2024, with a peak of 666 million packets per second [65]; the targets of the broader campaign were primarily banking and financial services, most of which were located in China.

The literature has documented several DDoS attacks that appear to represent Internet censorship. These attacks are attributed to governmental entities or proxies of that nation-state and tend to target a particular web page or Internet resource they oppose. Notably, DDoS attacks restrict access to Internet resources for *everyone*, not only users within the censor's sphere of influence. After the Snowden classified document leaks of 2013, media outlets reported that the United Kingdom Government Communications Headquarters (GCHQ) intelligence service had conducted DDoS attacks against servers hosting IRC chat services used by members of the hacktivist group Anonymous in 2011 [129]. Concern among Western nations arose from the incident as the use of the TCP-SYN flooding capability (dubbed "Rolling Thunder") may have "chilled" the speech of uninvolved users of the same servers.

DDoS-related censorship has often been shown to originate from Chinese ASes. On 16 March 2015, greatfire.org observed via their Amazon CloudFront hosting service that their website was under attack by a DDoS and their hosting costs had increased to \$30,000 per day because of the added throughput [273]. Later the same month, two GitHub pages affiliated with their project were also targeted. The organizers of greatfire.org offer resources to Chinese nationals looking to circumvent censorship within PRC. Baidu servers were identified as the source of much of the DDoS traffic, although the company denied involvement. Researchers dubbed the attack tool used during the incident the "great cannon" of China and advised that its overt deployment represents an escalation of state-level information control [192, 193]. The same paper alleges a capability called QUANTUM, maintained by the U.S. National Security Agency (NSA) and U.K. GCHQ, which similarly abuses plaintext HTTP traffic as a "man-on-the-side" attack for exploitation rather than DDoS.

⁹https://www.mandiant.com/sites/default/files/2021-09/rpt-apt38-2018-web_v5-1.pdf

Another example of politically motivated DDoS attacks happened in Southeast Asia in recent years. On 27 February 2022, CNN Philippines journalists were preparing to cover a presidential candidate debate when their website went down because of a DDoS attack [118]. Guest writes, "Since June 2021, opposition politicians, independent media, and fact-checking websites in the Philippines have been hit over and over with brute-force cyberattacks known as distributed denial-of-service, or DDoS, attacks. CNN, major news network ABS-CBN, Rappler (the outlet founded by the 2021 Nobel Peace Prize winner Maria Ressa), and VERA Files, a fact-checking organization, have all been targeted, along with the website of Vice President Leni Robredo, who is a staunch critic of the current president, Rodrigo Duterte" [133]. A Rappler journalistic investigation noted that a local Filipino hacking group called Pinoy Vendetta claimed credit for some of the attacks. A government report from the US-CERT (United States Computer Emergency Readiness Team) identified source IP addresses affiliated with the Philippine Army associated with DDoS attacks on two independent media sites, AlterMidya and Bulatlat [10]. Definitive attribution of DDoS attackers is difficult, and it is unclear whether some attacks denying access to political commentary in the Philippines were carried out by a nation-state actor.

A 2009 book chapter specifically analyzing DoS attacks with political motivations suggests that most DDoS attackers online are non-state actors [213]. A study in 2020 evinced that authoritarian states are increasingly performing attacks on foreign entities during elections and changes of power [185]. However, the ease of access to attack resources coupled with difficulties in attribution often leaves identifying involvement unresolved.

4.9 Additional Censorship Considerations

The Internet censorship methods described above encompass the technical dimensions of *how* censors deny access to Internet resources. These generally involve identifying "objectionable" content as it traverses a network and then implementing a blocking action. Other methods are more blunt instruments, such as Internet shutdowns, denying access to all users. There are additional considerations for Internet censorship from social and legal perspectives.

One element not yet discussed in this work is compelled disconnection or removal of Internet content. Web servers must physically exist and operate on computer hardware, and law enforcement (LE) entities may seize hardware to take a web page offline. LE may also compel a domain registrar to surrender access to a domain so website users can no longer resolve the IP address of the target server that hosts the now-censored content. See Figure 4 for an example of a splash page left behind on the Raid Forums website after a coalition of LE agencies seized the domain and associated web servers. Similar techniques have been used in other cybercrime enforcement measures. LE seizures frequently disrupt the online distribution of child sexual abuse material (CSAM), which is nearly universally denounced as criminal and an exception to the norms of free expression [151]. Web services that host pirated content in violation of the intellectual property laws of certain jurisdictions may also be targeted in a similar manner [246].

Determined state censors may also use substitution as censorship. Redirection of web requests is often used to display a blockpage, warning the user why they were denied access [82]; a censor could also use redirection to display a website that appears legitimate but contains false information. Although specific instances of disinformation through redirection have not been documented in the literature, state-sponsored actors have been shown to spread false information or links to "fake news" websites on social media platforms [136, 263, 340]. These efforts may distract or discourage citizens from seeking the information they want, or reduce overall institutional trust among Internet users [229]. Nation-states may also provide intranets, such as the Iran National Information Network (NIN) built by the IRGC, which aims to normalize among Iranian citizens that they can access the information they want from their government and do not need global Internet access. This normalization allows for creation of "splinternets" with less public outcry or pushback when a censorship apparatus denies access to resources from outside countries.

Another controversial and sometimes ambiguous aspect of Internet censorship deals with content moderation. In addition to heavily censored external connections, China has been shown to delete content on social media networks



Fig. 4. Example law enforcement domain seizure [292] (Retrieved October 2, 2022)

and news sites within its borders [1, 21, 84, 148, 172, 344]. The Chinese government encourages Chinese companies to host alternative services to Western social media (e.g., Weibo as a microblogging service, WeChat as an alternative to messengers such as WhatsApp), which it can compel to remove content as the CCP deems necessary.

In the United States, free speech is protected under the First Amendment to the country's constitution and upheld in the courts under the "state action doctrine." Furthermore, Section 230 of the Communications Decency Act is part of a federal law that protects online services from lawsuits based on user content [176]. While fostering free expression, issues of speech suppression have arisen because U.S. technology companies have wide latitude in determining what content they remove. Drawing distinctions between protected speech (e.g., radical ideas, controversial opinions, religious beliefs) and harmful rhetoric (e.g., hate speech, calls to enact violence, unlawful discrimination) can be difficult and spurs debate. From a legal perspective, U.S. citizens cannot have their speech suppressed by the government, but private companies hosting forums, blogs, and social media platforms are not beholden to the same standard and can moderate content as they see fit. Europe has demonstrated a desire to push further than the U.S.: rather than leaving companies the option of moderation, the European Union (EU) passed the Digital Services Act. The law requires companies that serve users in EU member states to moderate content for items such as "false information, hate speech, and extremism," and levies substantial fines (up to 6% of annual revenue) for noncompliance [261]. It is uncertain how European regulations will affect future global corporation moderation practices and impact Internet users' freedom of expression.

Self-censorship is another adjacent issue in any discussion of censorship. Humans often decide to withhold or selectively disclose their own discourse [67, 142]. Rather than deny access to content, a country may create a social or legal environment that encourages users to keep their views to themselves. A regime may allow access to social media platforms but punish journalists or dissidents who express opinions that portray authorities in a negative light [280]. In extreme examples, citizens may feel compelled to feign positive attitudes toward ruling elites [268].

Another tangential censorship issue is the disruption of information access due to physical attacks. Unlike Internet shutdowns as described above, some nations will attack critical infrastructure that enables Internet connectivity. Instances where European submarine fiber optic lines have been cut have been allegedly attributed to Russia, which "primarily employs deniable disruption and sabotage as part of hybrid warfare, aiming to impose costs on rivals without escalating to open conflict" [33]. During Russia's invasion of Ukraine, Russian forces made concerted efforts to destroy cellular telephone towers, disrupting Ukrainians' ability to communicate [28]. In response, Ukraine leveraged low-earth

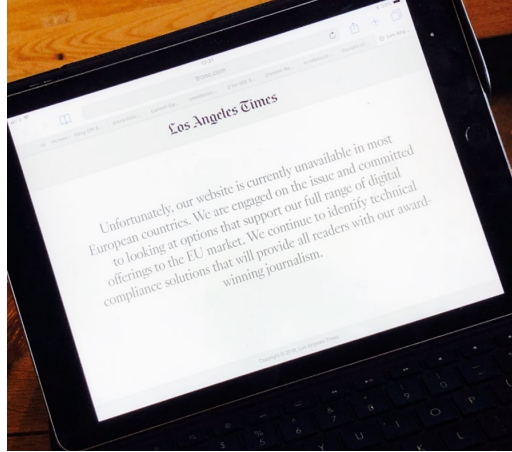


Fig. 5. Website blocking a user with a European IP address after GDPR enacted [284] (May 25, 2018)

orbit satellite constellations from Starlink to enable TCP/IP connectivity [46]. While these physical layer attacks are not intended solely for censorship, their consequences are similar to the network, transport, and application layer attacks described throughout our survey.

Recent technological advances in generative artificial intelligence (AI), such as large language models (LLMs), have become a target of censorship. LLMs allow for access to large datasets of information, which users can have "conversations" with software to analyze, summarize, or create text-based content. Nation-states have begun to demonstrate a willingness to deny their citizens access to web-based LLMs [6, 29]. Similarly, recent research has revealed that automated censorship is applied to Web-based language translation services (three Chinese, one American) [254].

A final consideration for potential Internet censorship is publisher denial of access. It is common practice for web servers to filter inbound traffic. Administrators do this for a variety of security reasons. Subscribing to CTI blocklists or loading indicators of compromise (IOC) lists into an Intrusion Detection System (IDS) or Intrusion Prevention System (IPS) may prevent cyber attacks from occurring against their server. Some servers filter IP addresses that are marked as malicious, send spam emails, or host malware. The Internet measurement community has also identified publisher-side censorship trends [290] in which users who value anonymity are scrutinized. At the time of their study, Khattak et al. found that 3.67% of the top 1000 web pages blocked access to Tor users [167]. Some servers naively block all Tor exit nodes. Others use heuristic methods to detect abusive behavior and punish any user who shares the same exit node as a malicious actor [272]. A Tor user may also receive differential treatment, such as a "paywall" or a CAPTCHA challenge to "prove they are human." Similarly, aggressive filtering is sometimes applied against VPN users [31, 93, 166, 243]. Additionally, publisher-side filtering may occur based on geolocation. Servers in one country may not accept users from another. When the EU enacted the General Data Protection Regulation (GDPR), many web servers worldwide blocked connections from European IP addresses [290]. These decisions likely happened to avoid compliance with GDPR's privacy rules or out of an abundance of caution for liability (see Figure 5 for an example). Mobile device applications have also been shown to discriminate based on the estimated location where a user resides [126]. These kinds of server-side blocking are not necessarily government censorship, but they are relevant to a holistic discussion of what it means to be able to access information freely in the modern, interconnected world.

Table 1. The Internet Censorship Methods Taxonomy

	Censorship Method	Techniques	OSI Model
Shallow Packet Inspection	Internet Shutdowns	Physical network disconnection	Physical Layer
		Logical denial	
		Routing blackhole*	
	Local Network Attacks**	ARP poisoning DoS	Data Link Layer
		MAC address filtering	
	IP Address Blocking	IP address blocklist/allowlist	Network Layer
		IP range blocklist/allowlist	
		Routing manipulation	
	Computer Network Attack	Network DDoS	
Deep Packet Inspection	Port Blocking	Port blocklist/allowlist	Transport Layer
		TCP/UDP/QUIC manipulation	
	BGP Attacks and Disruption	BGP hijacking	
		AS path forgery	
		BGP collusion attack	
	Bandwidth Throttling	Indiscriminate throttling	Session, Application, and Presentation Layers
		Routing manipulation/delays*	
		DPI latency injection	
		Targeted throttling	
		Traffic shaping/policing	
	DNS Tampering	Quality of Service (QoS)	
		DNS blocklist/allowlist	
		DNS cache poisoning	
		DNS protocol manipulation	
		DNS hijacking	
	Protocol/Application Content Filtering	DNS transparent proxy	
		URL blocklist/allowlist	
		Transparent proxy	
		HTTP web content matching	
		Keyword filtering (FTP, SMTP, etc)	
	TLS-based Filtering	SNI blocklist	
		MITM attack	
		Application targeting	
	Computer Network Attack	Application-layer DDoS	
		Offensive cyberspace operations	
	Protocol/Application Fingerprinting	Protocol manipulation	
		Protocol or application blocking	
		Active probing	
		Encrypted traffic analysis	
		Pattern/heuristic filtering	

*Occurs at the Network Layer

**None observed for nation-state censorship

5 A Taxonomy of Internet Censorship Methods

The Internet Censorship Methods Taxonomy, shown in Table 1, is the result of our study. We organized the taxonomy using the OSI model for clarity of presentation. A censorship method is an abstraction of similar techniques used by censors. Examples of techniques that comprise a censorship method are included, but this is not an exhaustive list. Layers 1-4 represent methods that only require "shallow packet inspection" methods (reading of packet headers) and are minimally resource-intensive. Deep packet inspection (DPI) methods analyze contents beyond packet headers. Notably, bandwidth throttling bridges the gap between these methods and can occur network-wide or on an application-layer basis. Distributed denial-of-service attacks also occur at both the network and the application layers.

Moving downwards through the methods categorized as DPI implies increased complexity and use of resources. As an example, consider the observation of patterns (or anomalies) in bitstreams of traffic, including those protected

Table 2. Evidence of Internet Censorship Methods

Censorship Method	References			
	1990-1999	2000-2009	2010-2019	2020-2025
Internet Shutdowns	-	[305]	[73, 138, 266, 289]	[32, 103, 130, 191, 233, 256, 265, 322]
IP Address Blocking	-	[61, 82, 89, 152, 328]	[4, 51, 106, 200, 221, 264, 272, 289, 298, 306]	[7, 30, 94, 218, 233, 244, 278, 332, 336]
Port Blocking	-	-	[4, 282, 289]	[30, 94, 218]
Computer Network Attacks	-	[57, 62, 203]	[4, 20, 73, 81, 83, 182, 192, 289]	[34, 36, 165, 185, 287]
BGP Attacks and Disruption	-	[214, 291]	[4, 27, 204]	[63, 155, 191, 233, 245]
Bandwidth Throttling	-	-	[12, 18, 289]	[334, 343]
DNS Tampering	-	[61, 82, 89, 152, 183, 320]	[3, 5, 13–15, 18, 51, 90, 100, 107, 120, 157, 159, 168, 210, 221, 235, 262, 264, 289, 298–300, 306, 323, 330, 337]	[7, 16, 24, 30, 37, 38, 43, 48, 55, 134, 137, 141, 145, 158, 164, 177, 218, 225, 233, 244, 245, 267, 271, 278, 279, 288, 336]
Protocol/Application Content Filtering	[85]	[61, 62, 71, 82, 89, 320]	[3, 4, 18, 51, 74, 76, 107, 120, 157, 168, 179, 210, 221, 234, 264, 289, 295, 298, 300, 306, 312, 330, 337]	[7, 36–38, 134, 137, 140, 158, 164, 175, 209, 218, 225, 232, 241, 242, 244, 245, 271, 277–279, 288, 301, 307, 307, 311, 329, 346]
TLS-based Filtering	-	-	[4, 18, 52, 113, 163, 216, 221, 289, 295, 318]	[7, 23, 35, 36, 40, 94, 134, 141, 158, 164, 219, 220, 240, 242, 259, 271, 277, 278, 311, 332, 334, 336, 346]
Protocol/Application Fingerprinting	-	-	[18, 91, 96, 97, 106, 143, 161, 215, 289, 303, 306, 318]	[8, 9, 37, 137, 325, 333, 335, 336]

by encryption. Encrypted traffic analysis often uses machine learning or statistical models to create profiles of traffic behavior for particular applications or protocols. Higher layers of OSI (down the taxonomy chart) represent methods of censorship with increasingly complex circumvention analysis and censor implementation. A preliminary version of the taxonomy was presented previously in [196]. In Table 2 we show the bibliography of evidence of Internet censorship methods and techniques, sorted chronologically by method and delineated by decade.

5.1 The Impact of Technological Changes

Foundational improvements in Internet architecture can reduce the ability of censors on a broad scale. For example, a primary means of TLS-based blocking occurs when a censor targets the hostname of a website within the unencrypted SNI extension of a TLS header. When Cloudflare attempted to deploy its own encrypted SNI (ESNI), China blocked all TLS 1.3 traffic that had ESNI [35]. The IETF has a draft RFC for an Encrypted ClientHello (ECH) to be implemented in the TLS standard [247]. If ratified and implemented on the majority of web servers across the Internet, ECH has the potential to eliminate much of an entire class of censor methods. If ECH becomes the global standard for TLS connections, censors will eventually have to abandon many of their TLS-based censorship methods.

Censorship circumvention has been characterized as an arms race or "cat-and-mouse" game: developers of anti-censorship software discover novel ways to evade censor methods, and motivated censors respond with countermeasures to thwart circumvention efforts. Manual discovery may take months or years, while more recent efforts have been made to automate the discovery of evasion strategies [39]. However, many censors have constraints on resources or political will. A country may be willing to implement a censorship regime knowing that it will not deter determined users but be content in knowing that the majority of the population will not have access to denied content. Some researchers

have questioned the assumption of an arms race situation, in favor of developing software that can be fully known by the censor [104], as censors are limited by the fundamental protocols that allow the Internet to operate. The Internet Censorship Methods Taxonomy scopes the problem space for anti-censorship software developers by characterizing what methods and techniques are possible. Future work may involve assigning metrics to a tool's ability to counter particular censor methods. This would allow for an objective comparison of the capabilities of anti-censorship software.

6 Conclusion

The Internet has surpassed television, print media, and radio in terms of media influence and has become one of "the irreplaceable elements of our lives since the 2000s" [88]. Consequently, a growing number of authorities in many jurisdictions apply censorship to online communications. Understanding how censors implement content filtering, throttling of connections, or Internet shutdowns is important for researchers, policymakers, and practitioners.

This work presents a taxonomy of Internet censorship methods derived from a systematic review of the literature and analysis of Internet technologies. The survey systematizes the technical means that censors implement, outlined by the legal and social circumstances in which they have occurred over the past three decades. We contend that this taxonomy has pedagogical and research value. This work lays the groundwork for future research endeavors in Internet design, Internet measurement studies, and online censorship circumvention.

References

- [1] Nicholas Aase, Jedidiah R. Crandall, Álvaro Díaz, Jeffrey Knockel, Jorge Ocaña Molinero, Jared Saia, Dan Wallach, and Tao Zhu. 2012. Whiskey, Weed, and Wukan on the World Wide Web: On Measuring Censors' Resources and Motivations. In *Free and Open Communications on the Internet*. USENIX, Bellevue, WA, 7. <https://www.usenix.org/system/files/conference/foci12/foci12-final17.pdf>
- [2] Access Now. 2026. Internet Shutdowns. <https://www.accessnow.org/issue/internet-shutdowns/>.
- [3] Giuseppe Aceto, Alessio Botta, Antonio Pescapè, M. Faheem Awan, Tahir Ahmad, and Saad Qaisar. 2016. Analyzing Internet Censorship in Pakistan. In *Research and Technologies for Society and Industry*. IEEE, Bologna, Italy, 6. <https://doi.org/10.1109/RTSI.2016.7740626>
- [4] Giuseppe Aceto and Antonio Pescapè. 2015. Internet Censorship Detection: A Survey. *Computer Networks* 83 (June 2015), 381–421. <https://doi.org/10.1016/j.comnet.2015.03.008>
- [5] AFNIC. 2013. *Report of the AFNIC Scientific Council: Consequences of DNS-based Internet Filtering*. Technical Report. AFNIC Scientific Council. <https://www.afnic.fr/wp-media/uploads/2021/01/SC-consequences-of-DNS-based-Internet-filtering.pdf>
- [6] Mohamed Ahmed and Jeffrey Knockel. 2024. Extended Abstract: The Impact of Online Censorship on LLMs. In *Free and Open Communications on the Internet*. Proceedings on Privacy Enhancing Technologies, Bristol, UK, 9. <https://www.petsymposium.org/foci/2024/foci-2024-0006.pdf>
- [7] Fatemah Alharbi, Michalis Faloutsos, and Nael Abu-Ghazaleh. 2020. Opening Digital Borders Cautiously yet Decisively: Digital Filtering in Saudi Arabia. In *FOCI Workshop*. USENIX, Virtual Event, USA, 7. https://www.usenix.org/system/files/foci20-paper-alharbi_0.pdf
- [8] Alice, Bob, Carol, Jan Beznazwy, and Amir Houmansadr. 2020. How China Detects and Blocks Shadowsocks. In *Internet Measurement Conference*. ACM, Virtual Event, USA, 14. <https://doi.org/10.1145/3419394.3423644>
- [9] Sultan Almutairi, Yogev Neumann, and Khaled Harfoush. 2024. Fingerprinting VPNs with Custom Router Firmware: A New Censorship Threat Model. In *Conference Proceedings of the CCNC*. IEEE, Las Vegas, NV, USA, 5. <https://doi.org/10.1109/CCNC51664.2024.10454833>
- [10] AlterMidya. 2021. Joint Statement by Bulatlat and Altermidya: Hold the Philippine Army Accountable for Cyberattacks against PH Media Websites. <https://www.altermidya.net/joint-statement-by-bulatlat-and-altermidya-hold-the-philippine-army-accountable-for-cyberattacks-against-ph-media-websites/>
- [11] Warda Amalou and Merouane Mehdi. 2022. An Approach to Mitigate DDoS Attacks on SIP Based VoIP. In *ICCEIS 2021*. MDPI, Virtual Event, USA, 6. <https://doi.org/10.3390/engproc2022014006>
- [12] Collin Anderson. 2013. *Dimming the Internet: Detecting Throttling as a Mechanism of Censorship in Iran*. Technical Report. University of Pennsylvania, PA. arXiv:1306.4361 [cs] <http://arxiv.org/abs/1306.4361>
- [13] Collin Anderson, Philipp Winter, and Roya. 2014. Global Network Interference Detection over the RIPE Atlas Network. In *Free and Open Communications on the Internet*. USENIX, San Diego, CA, 8. <https://www.usenix.org/system/files/conference/foci14/foci14-anderson.pdf>
- [14] Anonymous. 2012. The Collateral Damage of Internet Censorship by DNS Injection. *ACM SIGCOMM Computer Communication Review* 42, 3 (June 2012), 21–27. <https://doi.org/10.1145/2317307.2317311>
- [15] Anonymous. 2014. Towards a Comprehensive Picture of the Great Firewall's DNS Censorship. In *Free and Open Communications on the Internet*. USENIX, San Diego, CA, 7. <https://www.usenix.org/system/files/conference/foci14/foci14-anonymous.pdf>
- [16] Anonymous, Arian Akhavan Niaki, Nguyen Phong Hoang, Phillipa Gill, and Amir Houmansadr. 2020. Triplet Censors: Demystifying Great Firewall's DNS Censorship Behavior. In *Free and Open Communications on the Internet*. USENIX, Virtual Event, USA, 7. https://www.usenix.org/system/files/foci20-paper-anonymous_0.pdf
- [17] U.S. Army. 2020. *ATP 7-100.2 North Korean Tactics*. Army Techniques Publication. U.S. Army.

- [18] Simurgh Aryan, Homa Aryan, and J Alex Halderman. 2013. Internet Censorship in Iran: A First Look. In *Free and Open Communications on the Internet*. USENIX, Washington, D.C., 8. <https://www.usenix.org/system/files/conference/foci13/foci13-aryan.pdf>
- [19] Albany Associates. 2012. The Arab Spring and the Impact of Social Media. <https://web.archive.org/web/20130512021954/http://www.albanyassociates.com/notebook/2012/03/the-arab-spring-and-the-impact-of-social-media/>
- [20] Michael Bailey and Craig Labovitz. 2011. *Censorship and Co-option of the Internet Infrastructure*. Technical Report. University of Michigan, Ann Arbor, MI. 7 pages. <https://faculty.cc.gatech.edu/~mbailey/publications/CSE-TR-572-11.pdf>
- [21] David Bamman, Brendan O'Connor, and Noah Smith. 2012. Censorship and Deletion Practices in Chinese Social Media. *First Monday* 17, 3 (March 2012), 16. <https://doi.org/10.5210/fm.v17i3.3943>
- [22] The World Bank. 2022. % of Population Using the Internet - Myanmar. <https://data.worldbank.org/indicator/IT.NET.USER.ZS?locations=MM>
- [23] Simone Basso. 2020. DNS over TLS Blocked in Iran. <https://ooni.org/post/2020-iran-dot/>
- [24] Simone Basso. 2021. Measuring DoT/DoH Blocking Using OONI Probe: A Preliminary Study. In *Network and Distributed System Security (NDSS) Symposium*. NDSS, Virtual Event, USA, 10. <https://www.ndss-symposium.org/wp-content/uploads/dnspriv21-02-paper.pdf>
- [25] Ralf Bendrath and Milton Mueller. 2011. The End of the Net as We Know It? Deep Packet Inspection and Internet Governance. *New Media & Society* 13, 7 (Nov. 2011), 1142–1160. <https://doi.org/10.1177/1461444811398031>
- [26] Karyn Benson. 2016. *Leveraging Internet Background Radiation for Opportunistic Network Analysis*. Ph.D. Dissertation. University of California, San Diego. <https://www.proquest.com/dissertations-theses/leveraging-internet-background-radiation/docview/1828386151/se-2>
- [27] Karyn Benson, Alberto Dainotti, K. C. Claffy, and Emile Aben. 2013. Gaining Insight into AS-level Outages through Analysis of Internet Background Radiation. In *IEEE Conference on Computer Communications Workshops*. IEEE, Turin, Italy, 6. <https://cseweb.ucsd.edu/~kbenson/papers/tma13.pdf>
- [28] Vera Bergengruen. 2022. *The Battle for Control Over Ukraine's Internet*. TIME. <https://time.com/6222111/ukraine-internet-russia-reclaimed-territory/>
- [29] Harel Berger and Yuval Shavitt. 2024. *Measuring DNS Censorship of Generative AI Platforms*. Technical Report. ArXiv. 14 pages. <https://doi.org/10.48550/arXiv.2412.14286>
- [30] Abhishek Bhaskar and Paul Pearce. 2022. Many Roads Lead To Rome: How Packet Headers Influence DNS Censorship Measurement. In *USENIX Security Symposium*. USENIX, Boston, MA, USA, 17. <https://www.usenix.org/system/files/sec22-bhaskar.pdf>
- [31] Francine Biazin do Nascimento. 2021. *Measuring the Accessibility of Popular Websites While Using Mullvad VPN*. Ph.D. Dissertation. Delft University of Technology. <http://resolver.tudelft.nl/uuid:190ac398-50a0-4a47-9ec4-5514d5b22e25>
- [32] Zachary S. Bischof, Kennedy Pitcher, Esteban Carisimo, Amanda Meng, Rafael Bezerra Nunes, Ramakrishna Padmanabhan, Margaret E. Roberts, Alex C. Snoeren, and Alberto Dainotti. 2023. Destination Unreachable: Characterizing Internet Outages and Shutdowns. In *Proceedings of the ACM SIGCOMM 2023 Conference*. ACM, New York NY USA, 14. <https://doi.org/10.1145/3603269.3604883>
- [33] Bergur Bjartmarsson. 2026. *Beneath the Surface: Understanding Russian and Chinese Actions in Submarine Cable Context: A Comparative Analysis*. Master's thesis. University of Iceland, Reykjavik, Iceland. <https://hdl.handle.net/1946/51989>
- [34] Kevin Bock, Abdulrahman Alaraj, Yair Fax, Kyle Hurley, Eric Wustrow, and Dave Levin. 2021. Weaponizing Middleboxes for TCP Reflected Amplification. In *USENIX Security Symposium*. USENIX, Virtual Event, USA, 18. <https://www.usenix.org/system/files/sec21-bock.pdf>
- [35] Kevin Bock, Anonymous, Louis-Henri Merino, David Fifield, Amir Houmansadr, and Dave Levin. 2020. Exposing and Circumventing China's Censorship of ESNI. https://gfw.report/blog/gfw_esni_blocking/en/
- [36] Kevin Bock, Pranav Bharadwaj, Jasraj Singh, and Dave Levin. 2021. Your Censor Is My Censor: Weaponizing Censorship Infrastructure for Availability Attacks. In *2021 IEEE Security and Privacy Workshops (SPW)*. IEEE, San Francisco, CA, USA, 398–409. <https://doi.org/10.1109/SPW53761.2021.00059>
- [37] Kevin Bock, Yair Fax, Kyle Reese, Jasraj Singh, and Dave Levin. 2020. Detecting and Evading Censorship-in-Depth: A Case Study of Iran's Protocol Filter. In *Free and Open Communications on the Internet*. USENIX, Virtual Event, USA, 7. <https://www.usenix.org/system/files/foci20-paper-bock.pdf>
- [38] Kevin Bock, George Hughey, Louis-Henri Merino, Tania Arya, Daniel Liscinsky, Regina Pogolian, and Dave Levin. 2020. Come as You Are: Helping Unmodified Clients Bypass Censorship with Server-side Evasion. In *ACM SIGCOMM*. ACM, Virtual Event, USA, 13.
- [39] Kevin Bock and Dave Levin. 2020. Automating the Censorship Arms Race. *XRDS: Crossroads, The ACM Magazine for Students* 27, 2 (Dec. 2020), 30–35. <https://dl.acm.org/doi/10.1145/3437410>
- [40] Kevin Bock, Gabriel Naval, Kyle Reese, and Dave Levin. 2021. Even Censors Have a Backup: Examining China's Double HTTPS Censorship Middleboxes. In *Free and Open Communications on the Internet*. ACM, Virtual Event, USA, 7. <https://doi.org/10.1145/3473604.3474559>
- [41] Robert Braden. 1989. RFC 1122 - Requirements for Internet Hosts - Communication Layers. <https://datatracker.ietf.org/doc/html/rfc1122>
- [42] Susan W Brenner. 2007. At Light Speed: Attribution and Response to Cybercrime/Terrorism/Warfare. *The journal of criminal law & criminology* 97 (2007), 379–475. <https://www.proquest.com/docview/218442098>
- [43] Jacob Brown, Xi Jiang, Van Tran, Arjun Nitin Bhagoji, Nguyen Phong Hoang, Nick Feamster, Prateek Mittal, and Vinod Yegneswaran. 2023. 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*. ACM, Long Beach CA USA, 13. <https://doi.org/10.1145/3580305.3599775>
- [44] Axel Bruns, Tim Highfield, and Jean Burgess. 2013. The Arab Spring and Social Media Audiences: English and Arabic Twitter Users and Their Networks. *American Behavioral Scientist* 57, 7 (July 2013), 871–898. <https://doi.org/10.1177/0002764213479374>
- [45] Tomasz Bujlow, Valentin Carela-Español, and Pere Barlet-Ros. 2015. Independent Comparison of Popular DPI Tools for Traffic Classification. *Computer Networks* 76 (Jan. 2015), 75–89. <https://doi.org/10.1016/j.comnet.2014.11.001>
- [46] Matt Burgess. 2022. *How Starlink Scrambled to Keep Ukraine Online*. WIRED. <https://www.wired.com/story/starlink-ukraine-internet/>
- [47] Matt Burgess and Lily Hay Newman. 2026. Internet Starts to Return in Iran After 3-Month Blackout. <https://www.wired.com/story/internet-in-iran-starts-to-return-after-3-month-blackout/>
- [48] Paola Calle, Larissa Savitsky, Arjun Nitin Bhagoji, Nguyen Phong Hoang, and Shinyoung Cho. 2024. Toward Automated DNS Tampering Detection Using Machine Learning. In *Free and Open Communications on the Internet*. Proceedings on Privacy Enhancing Technologies, Bristol, UK, 9.

- <https://www.petsymposium.org/foci/2024/foci-2024-0008.pdf>
- [49] Center for Human Rights in Iran. 2018. Iran’s Severely Disrupted Internet During Protests: “Websites Hardly Open”. <https://iranhumanrights.org/2018/01/irans-severely-disrupted-internet-during-protests-websites-hardly-open/>
 - [50] Software Freedom Law Centre. 2020. Internet Shutdowns. <https://internetsutdowns.in/>
 - [51] Abdelberi Chaabane, Terence Chen, Mathieu Cunche, Emiliano De Cristofaro, Arik Friedman, and Mohamed Ali Kaafar. 2014. Censorship in the Wild: Analyzing Internet Filtering in Syria. In *Internet Measurement Conference*. ACM, Vancouver, BC, Canada, 14. <https://doi.org/10.1145/2663716.2663720>
 - [52] Zimo Chai, Amirhossein Ghafari, and Amir Houmansadr. 2019. On the Importance of Encrypted-SNI (ESNI) to Censorship Circumvention. In *Free and Open Communications on the Internet*. USENIX, Santa Clara, CA, USA, 8. https://www.usenix.org/system/files/foci19-paper_chai_update.pdf
 - [53] Sonali Chandel, Zang Jingji, Yu Yunnan, Sun Jingyao, and Zhang Zhipeng. 2019. The Golden Shield Project of China: A Decade Later—An in-Depth Study of the Great Firewall. In *2019 International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery (CyberC)*. IEEE, Guilin, China, 111–119. <https://doi.org/10.1109/CyberC.2019.00027>
 - [54] Mo Chen and Jens Grossklags. 2022. Social Control in the Digital Transformation of Society: A Case Study of the Chinese Social Credit System. *Social Sciences* 11, 6 (2022), 23. <https://doi.org/10.3390/socsci11060229>
 - [55] Yanan Cheng, Yali Liu, Chao Li, Zhaoxin Zhang, Ning Li, and Yuejin Du. 2022. In-Depth Evaluation of the Impact of National-Level DNS Filtering on DNS Resolvers over Space and Time. *Electronics* 11, 8 (April 2022), 1276. <https://doi.org/10.3390/electronics11081276>
 - [56] Catalin Cimpanu. 2022. KlaySwap Crypto Users Lose Funds after BGP Hijack. <https://therecord.media/klayswap-crypto-users-lose-funds-after-bgp-hijack/>
 - [57] CISA. 2009. *Understanding Denial-of-Service Attacks*. Technical Report ST04-015. U.S. Cybersecurity and Infrastructure Agency. <https://www.cisa.gov/uscirt/ncas/tips/ST04-015>
 - [58] CISA. 2017. *HIDDEN COBRA – North Korea’s DDoS Botnet Infrastructure*. Technical Report TA17-164A. U.S. Cybersecurity and Infrastructure Agency. <https://www.cisa.gov/uscirt/ncas/alerts/TA17-164A>
 - [59] Mitchell Clark. 2021. What Is BGP, and What Role Did It Play in Facebook’s Massive Outage. <https://www.theverge.com/2021/10/4/22709260/what-is-bgp-border-gateway-protocol-explainer-internet-facebook-outage>
 - [60] Richard A. Clarke and Robert K. Knake. 2012. *Cyber War: The next Threat to National Security and What to Do about It* (1st ed.). Ecco, New York.
 - [61] Richard Clayton. 2006. Failures in a Hybrid Content Blocking System. In *Privacy Enhancing Technologies*. Springer, Cambridge, UK, 78–92. <https://www.cl.cam.ac.uk/~rnc1/cleanfeed.pdf>
 - [62] Richard Clayton, Steven J. Murdoch, and Robert N. M. Watson. 2006. Ignoring the Great Firewall of China. In *Privacy Enhancing Technologies*. Vol. 4258. Springer Berlin Heidelberg, Berlin, Heidelberg, 20–35. https://doi.org/10.1007/11957454_2
 - [63] Cloudflare. 2018. BGP Hijacking. <https://www.cloudflare.com/learning/security/glossary/bgp-hijacking/>
 - [64] Cloudflare. 2018. What Is BGP? | BGP Routing Explained. <https://www.cloudflare.com/learning/security/glossary/what-is-bgp/>
 - [65] Cloudflare. 2024. Famous DDoS Attacks | The Largest DDoS Attacks of All Time. <https://www.cloudflare.com/learning/ddos/famous-ddos-attacks/>
 - [66] US Strategic Command. 2009. *The Cyber Warfare Lexicon*. Technical Report. USSTRATCOM. <https://info.publicintelligence.net/USSTRATCOM-CyberWarfareLexicon.pdf>
 - [67] Philip Cook and Conrad Heilmann. 2013. Two Types of Self-Censorship: Public and Private. *Political Studies* 61, 1 (March 2013), 178–196. <https://doi.org/10.1111/j.1467-9248.2012.00957.x>
 - [68] Human Rights Council. 2020. *General Comment No. 37 (2020) on the Right of Peaceful Assembly (Article 21)*. Technical Report CCPR/C/GC/37. United Nations, United Nations. <https://digitallibrary.un.org/record/3884725?ln=en>
 - [69] Human Rights Council. 2021. *Ending Internet Shutdowns: A Path Forward*. Technical Report A/HRC/47/24/Add.2. United Nations, United Nations. <https://documents-dds-ny.un.org/doc/UNDOC/GEN/G21/149/66/PDF/G2114966.pdf?OpenElement>
 - [70] Council on Foreign Relations. 2005. Titan Rain. <https://www.cfr.org/cyber-operations/titan-rain>
 - [71] Jedidiah R. Crandall, Daniel Zinn, Michael Byrd, Earl Barr, and Rich East. 2007. ConceptDoppler: A Weather Tracker for Internet Censorship. In *Computer and Communications Security*. ACM, Alexandria, VA, USA, 14. <http://www.csd.uoc.gr/~hy558/papers/conceptdoppler.pdf>
 - [72] Abdi Latif Dahir. 2019. After a Record 16-Month Ban, This President Has Unblocked Social Media Access. <https://qz.com/africa/1667263/chads-idriss-deby-unblocks-social-media-after-record-shutdown>
 - [73] Alberto Dainotti, Claudio Squarcella, Emile Aben, Kimberly Claffy, Marco Chiesa, Michele Russo, and Antonio Pescapè. 2011. Analysis of Country-wide Internet Outages Caused by Censorship. *ACM IMC 2011* 1 (2011), 18. <https://doi.org/10.1145/2068816.2068818>
 - [74] Jakub Dalek, Bennett Haselton, Helmi Noman, Adam Senfit, Masashi Crete-Nishihata, Phillipa Gill, and Ronald Deibert. 2013. A Method for Identifying and Confirming the Use of URL Filtering Products for Censorship. In *Proceedings of the ACM IMC 2013*. ACM, Barcelona Spain, 23–30. <https://doi.org/10.1145/2504730.2504763>
 - [75] Jakub Dalek and Adam Senfit. 2011. Behind Blue Coat: Investigations of Commercial Filtering in Syria and Burma. <https://citizenlab.ca/2011/11/behind-blue-coat/>
 - [76] Alexander Darer, Oliver Farnan, and Joss Wright. 2017. FilteredWeb: A Framework for the Automated Search-Based Discovery of Blocked URLs. In *Network Traffic Measurement and Analysis*. IFIP, Dublin, Ireland, 9. http://tma.ifip.org/wordpress/wp-content/uploads/2017/06/tma2017_paper32.pdf
 - [77] DARPA. 1981. RFC 793 - Transmission Control Protocol (TCP). <https://datatracker.ietf.org/doc/html/rfc793>
 - [78] Shepardson David. 2021. Censorship Circumvention Tool Helps 1.4 Million Cubans Get Internet Access. <https://www.reuters.com/world/americas/censorship-circumvention-tool-helps-14-million-cubans-get-internet-access-2021-07-16/>
 - [79] Dawn. 2010. 5,000 Internet Porn Users Held in China Last Year. <https://www.dawn.com/news/825971/5-000-internet-porn-users-held-in-china-last-year>
 - [80] J.D. Day and H. Zimmermann. 1983. The OSI Reference Model. *Proc. IEEE* 71, 12 (1983), 16. <https://doi.org/10.1109/PROC.1983.12775>

- [81] Ronald Deibert, John Palfrey, Rafal Rohozinski, Jonathan Zittrain, and OpenNet Initiative (Eds.). 2010. *Access Controlled: The Shaping of Power, Rights, and Rule in Cyberspace*. MIT Press, Cambridge, MA.
- [82] Ronald Deibert, John Palfrey, Rafal Rohozinski, Jonathan Zittrain, and Janice Stein. 2008. *Access Denied: The Practice and Policy of Global Internet Filtering*. The MIT Press, Cambridge, MA. <https://muse.jhu.edu/book/60844>
- [83] Ronald Deibert, Jonathan Zittrain, Rafal Rohozinski, and John Palfrey (Eds.). 2012. *Access Contested: Security, Identity, and Resistance in Asian Cyberspace Information Revolution and Global Politics*. MIT Press, Cambridge, MA.
- [84] Kevin Michael Deluca. 2016. Weibo, WeChat, and the Transformative Events of Environmental Activism on China's Wild Public Screens. *International Journal of Communications* 10 (2016), 19.
- [85] Chen Ding, Chi-Hung Chi, Jing Deng, and Chun-Lei Dong. 1999. Centralized Content-Based Web Filtering and Blocking: How Far Can It Go?. In *Conference Proceedings*, Vol. 2. IEEE, Tokyo, Japan, 115–119. <https://doi.org/10.1109/ICSMC.1999.825218>
- [86] Roger Dingledine, Nick Mathewson, and Paul Syverson. 2004. Tor: The Second-Generation Onion Router. *13th USENIX Security Symposium* 1 (2004), 18. http://usenix.org/publications/library/proceedings/sec04/tech/full_papers/dingledine/dingledine.pdf
- [87] Lucas Dixon, Thomas Ristenpart, and Thomas Shrimpton. 2016. Network Traffic Obfuscation and Automated Internet Censorship. *IEEE Security & Privacy* 14, 6 (Nov. 2016), 43–53. <https://doi.org/10.1109/MSP.2016.121>
- [88] Ayhan Dolunay, Fevzi Kasap, and Gökçe Keçeci. 2017. Freedom of Mass Communication in the Digital Age in the Case of the Internet: "Freedom House" and the USA Example. *Sustainability* 9, 10 (Oct. 2017), 1739. <https://doi.org/10.3390/su9101739>
- [89] Maximilian Dornseif. 2003. Government Mandated Blocking of Foreign Web Content. In *DFN-Arbeitsstagung Über Kommunikationsnetze*. Gesellschaft für Informatik, Trier, Germany, 617–647. <https://censorbib.nymity.ch/pdf/Dornseif2003a.pdf>
- [90] Haixin Duan, Nicholas Weaver, Zongxu Zhao, Meng Hu, Jinjin Liang, Jian Jiang, Kang Li, and Vern Paxson. 2012. Hold-On: Protecting Against On-Path DNS Poisoning. In *Securing and Trusting Internet Names*. National Physical Laboratory, UK, 7. <http://www.cs.albany.edu/~mariya/courses/csi516F19/papers/holdon.pdf>
- [91] Arun Dunna, Ciarán O'Brien, and Phillipa Gill. 2018. Analyzing China's Blocking of Unpublished Tor Bridges. In *Free and Open Communications on the Internet*. USENIX, Baltimore, MD, USA, 7. <https://www.usenix.org/system/files/conference/foci18/foci18-paper-dunna.pdf>
- [92] Josiah Dykstra, Lawrence A. Gordon, Martin P. Loeb, and Lei Zhou. 2022. The Economics of Sharing Unclassified Cyber Threat Intelligence by Government Agencies and Departments. *Journal of Information Security* 13, 03 (2022), 85–100. <https://doi.org/10.4236/jis.2022.133006>
- [93] Sabrina Earle. 2016. The Battle against Geo-Blocking: The Consumer Strikes Back. *Richmond journal of global law and business* 15 (2016), 20. <https://scholarship.richmond.edu/global/vol15/iss1/2>
- [94] Kathrin Elmenhorst, Bertram Schütz, Nils Aschenbruck, and Simone Basso. 2021. Web Censorship Measurements of HTTP/3 over QUIC. In *Proceedings of the 21st ACM Internet Measurement Conference*. ACM, Virtual Event, 276–282. <https://doi.org/10.1145/3487552.3487836>
- [95] Let's Encrypt. 2021. Let's Encrypt Statistics. <https://letsencrypt.org/stats/>
- [96] Roya Ensafi, David Fifield, Philipp Winter, Nick Feamster, Nicholas Weaver, and Vern Paxson. 2015. Examining How the Great Firewall Discovers Hidden Circumvention Servers. In *Proceedings of the 2015 Internet Measurement Conference*. ACM, Tokyo Japan, 445–458. <https://doi.org/10.1145/2815675.2815690>
- [97] Roya Ensafi, Philipp Winter, Abdullah Mueen, and Jedidiah R. Crandall. 2015. Analyzing the Great Firewall of China Over Space and Time. *Proceedings on Privacy Enhancing Technologies* 2015, 1 (April 2015), 61–76. <https://doi.org/10.1515/popets-2015-0005>
- [98] Golnaz Esfandiari. 2013. Iran Admits Throttling Internet To 'Preserve Calm' During Election. <https://www.rferl.org/a/iran-internet-disruptions-election/25028696.html>
- [99] James Fallows. 2008. The Connection Has Been Reset. <https://www.theatlantic.com/magazine/archive/2008/03/the-connection-has-been-reset/306650/>
- [100] Oliver Farnan, Alexander Darer, and Joss Wright. 2016. Poisoning the Well – Exploring the Great Firewall's Poisoned DNS Responses. In *Workshop on Privacy in the Electronic Society*. ACM, Vienna, Austria, 4. <https://doi.org/10.1145/2994620.2994636>
- [101] FCC. 2000. Children's Internet Protection Act (CIPA). <https://www.fcc.gov/consumers/guides/childrens-internet-protection-act>
- [102] Steven Feldstein. 2021. *The Rise of Digital Repression: How Technology Is Reshaping Power, Politics, and Resistance*. Oxford Press, New York, NY.
- [103] Steven Feldstein. 2022. Internet Shutdowns Are Changing. How Should Citizens and Democracies Respond? <https://carnegieendowment.org/2022/03/31/government-internet-shutdowns-are-changing-how-should-citizens-and-democracies-respond-pub-86687>
- [104] David Fifield. 2024. Against the "Arms Race". <https://bamssoftware.com/talks/arms-race-foci-2024/>
- [105] David Fifield, Chang Lan, Rod Hynes, Percy Wegmann, and Vern Paxson. 2015. Blocking-Resistant Communication through Domain Fronting. *Proceedings on Privacy Enhancing Technologies* 2015, 2 (June 2015), 46–64. <https://doi.org/10.1515/popets-2015-0009>
- [106] David Fifield and Lynn Tsai. 2016. Censors' Delay in Blocking Circumvention Proxies. In *Free and Open Communications on the Internet*. USENIX, Austin, TX, 7. <https://www.usenix.org/system/files/conference/foci16/foci16-paper-fifield.pdf>
- [107] Arturo Filastò and Jacob Appelbaum. 2012. OONI: Open Observatory of Network Interference. In *Free and Open Communications on the Internet*. USENIX, Bellevue, WA, 8. <https://www.usenix.org/system/files/conference/foci12/foci12-final12.pdf>
- [108] Arlene Fink. 2005. *Conducting Research Literature Reviews: From the Internet to Paper* (2nd ed ed.). Sage Publications, Thousand Oaks, CA.
- [109] FireninjaDD. 2017. Will Using a VPN Prevent ISP Throttling Due to the Removal of Net Neutrality? https://www.reddit.com/r/technology/comments/6c1kck/will_using_a_vpn_prevent_isp_throttling_due_to/
- [110] Forcepoint. 2023. Protocol Agents on NGFW Engines. <https://help.forcepoint.com/ngfw/en-us/7.0.0/GUID-8A8FD672-D113-4D47-945A-46298F96C0F9.html>
- [111] Fortinet. 2021. FortiOS 7.0.1 Administration Guide. <https://docs.fortinet.com/document/fortigate/7.0.1/administration-guide/233445/blocking-applications-with-custom-signatures>

- [112] Sergey Frolov, Jack Wampler, and Eric Wustrow. 2020. Detecting Probe-resistant Proxies. In *Network and Distributed System Security*. The Internet Society, San Diego, California, 17. <https://doi.org/10.14722/ndss.2020.23087>
- [113] Sergey Frolov and Eric Wustrow. 2019. The Use of TLS in Censorship Circumvention. In *Network and Distributed System Security*. The Internet Society, San Diego, CA, 15. https://www.ndss-symposium.org/wp-content/uploads/2019/02/ndss2019_03B-2-1_Frolov_paper.pdf
- [114] Sergey Frolov and Eric Wustrow. 2020. HTTP: A Probe-Resistant Proxy. In *Free and Open Communications on the Internet*, Vol. 1. USENIX, Virtual Event, USA, 8. <https://www.usenix.org/system/files/foci20-paper-frolov.pdf>
- [115] FTC. 1998. Children’s Online Privacy Protection Rule (“COPPA”). <https://www.ftc.gov/enforcement/rules/rulemaking-regulatory-reform-proceedings/childrens-online-privacy-protection-rule>
- [116] Raphael Garcia. 2019. Internet Censorship Is Part of South Korea’s Democracy Package. <https://international.thenewslens.com/article/122579>
- [117] Frank Gardner. 2000. Saudis ‘defeating’ Internet Porn. http://news.bbc.co.uk/2/hi/middle_east/742798.stm
- [118] Jodesz Gavilan. 2022. Cyberattack Hits CNN Philippines on Day of Presidential Debate. <https://www.rappler.com/nation/cyberattack-hits-cnn-philippines-presidential-debate-february-27-2022/>
- [119] Alessandro Ghedini. 2018. Encrypt It or Lose It: How Encrypted SNI Works. <https://blog.cloudflare.com/encrypted-sni/>
- [120] Phillipa Gill, Masashi Crete-Nishihata, Jakub Dalek, Sharon Goldberg, Adam Senft, and Greg Wiseman. 2015. Characterizing Web Censorship Worldwide: Another Look at the OpenNet Initiative Data. *ACM Transactions on the Web* 9, 1 (2015), 29. <https://doi.org/10.1145/2700339>
- [121] GoodbyeDPI. 2017. GoodbyeDPI – Deep Packet Inspection Circumvention Utility. <https://github.com/ValdikSS/GoodbyeDPI>
- [122] Google. 2017. HTTPS Encryption on the Web. <https://transparencyreport.google.com/https/overview>
- [123] Google. 2023. Feature: TLS Encrypted Client Hello (ECH). <https://chromestatus.com/feature/6196703843581952>
- [124] Devashish Gosain, Anshika Agarwal, Sahil Shekhawat, H. B. Acharya, and S. Chakravarty. 2018. Mending Wall: On the Implementation of Censorship in India. *arXiv* 1 (June 2018), 20. <https://doi.org/10.48550/arXiv.1806.06518>
- [125] Devashish Gosain, Mayank Mohindra, and Sambuddho Chakravarty. 2021. Too Close for Comfort: Morasses of (Anti-) Censorship in the Era of CDNs. *Proceedings on Privacy Enhancing Technologies* 2021, 2 (April 2021), 173–193. <https://doi.org/10.2478/popets-2021-0023>
- [126] Devashish Gosain, Kartikey Singh, Rishi Sharma, Jithin Suresh Babu, and Sambuddho Chakravaty. 2024. Out in the Open: On the Implementation of Mobile App Filtering in India. In *Passive and Active Measurement*, Philipp Richter, Vaibhav Bajpai, and Esteban Carisimo (Eds.). Vol. 14538. Springer Nature Switzerland, Cham, 19–36. https://doi.org/10.1007/978-3-031-56252-5_2
- [127] James A. Green (Ed.). 2015. *Cyber Warfare: A Multidisciplinary Analysis*. Routledge, Abingdon, Oxon ; New York, NY.
- [128] Andy Greenberg. 2017. The WannaCry Ransomware Has a Link to Suspected North Korean Hackers. <https://www.wired.com/2017/05/wannacry-ransomware-link-suspected-north-korean-hackers/>
- [129] Glenn Greenwald. 2014. Exclusive: Snowden Docs Show UK Spies Attacked Anonymous, Hackers. <https://www.nbcnews.com/feature/edward-snowdeninterview/exclusive-snowden-docs-show-uk-spies-attackedanonymous-hackers-n21361>
- [130] Giovanni De Gregorio. 2020. Internet Shutdowns and the Limits of Law. *International Journal of Communication* 1 (2020), 20. <https://ijoc.org/index.php/ijoc/article/view/13752/3183>
- [131] Internet Security Research Group. 2023. Let’s Encrypt. <https://letsencrypt.org/>
- [132] Rohan Grover, Kyooeun Jang, and Li Wen Su. 2024. Beyond Digital Protection(ism). *Journal of Information Policy* 14 (May 2024), 32. <https://doi.org/10.5325/jinfopoli.14.2024.0005>
- [133] Peter Guest. 2022. “It’s like Being under Siege”: How DDoS Became a Censorship Tool. <https://restofworld.org/2022/blackouts-ddos/>
- [134] Hall et al. 2022. A Survey of Worldwide Censorship Techniques (Draft in Progress). <https://datatracker.ietf.org/doc/draft-irtf-pearg-censorship/08/>
- [135] Alexis Hancock. 2021. HTTPS Is Actually Everywhere. <https://www.eff.org/deeplinks/2021/09/https-actually-everywhere>
- [136] Nicholas Harrell, Alexander Master, Nicolas Starck, and Daniel Eerhart. 2025. Tactics and Techniques of Information Operations: Gaps in US Response to Counter Malign Influence. *International Conference on Cyber Warfare and Security* 1 (2025), 10. <https://doi.org/10.34190/iccws.20.1.3271>
- [137] Michael Harrity, Kevin Bock, Frederick Sell, and Dave Levin. 2022. GET /out: Automated Discovery of Application-Layer Censorship Evasion Strategies. In *USENIX Security Symposium*. USENIX, Boston, MA, USA, 19. <https://www.usenix.org/system/files/sec22-harrity.pdf>
- [138] Shaddi Hasan, Yahel Ben-David, Giulia Fanti, Eric Brewer, and Scott Shenker. 2013. Building Dissent Networks: Towards Effective Countermeasures against Large-Scale Communications Blackouts. In *Free and Open Communications on the Internet*. USENIX, San Jose, CA, USA, 8. <https://censorbib.nymity.ch/pdf/Hasan2013a.pdf>
- [139] Rebecca Hersher. 2015. Meet Mafiaboy, The ‘Bratty Kid’ Who Took Down The Internet. <https://www.npr.org/sections/alltechconsidered/2015/02/07/384567322/meet-mafiaboy-the-bratty-kid-who-took-down-the-internet>
- [140] Nguyen Phong Hoang, Jakub Dalek, Masashi Crete-Nishihata, Nicolas Christin, Vinod Yegneswaran, Michalis Polychronakis, and Nick Feamster. 2024. GFWeb: Measuring the Great Firewall’s Web Censorship at Scale. In *USENIX Security Symposium*. USENIX, Philadelphia, PA, 17. <https://www.usenix.org/system/files/sec24fall-prepub-310-hoang.pdf>
- [141] Nguyen Phong Hoang, Michalis Polychronakis, and Phillipa Gill. 2022. Measuring the Accessibility of Domain Name Encryption and Its Impact on Internet Filtering. In *Passive and Active Measurement Conference*. Springer, Virtual Event, 17. https://doi.org/10.1007/978-3-030-98785-5_23
- [142] John Horton. 2011. Self-Censorship. *Res Publica* 17, 1 (Feb. 2011), 91–106. <https://doi.org/10.1007/s11158-011-9145-3>
- [143] A. Houmansadr, C. Brubaker, and V. Shmatikov. 2013. The Parrot Is Dead: Observing Unobservable Network Communications. In *2013 IEEE Symposium on Security and Privacy*. IEEE, Berkeley, CA, 65–79. <https://doi.org/10.1109/SP.2013.14>
- [144] Freedom House. 2026. *Freedom on the Net 2025 Report*. Technical Report. Freedom House, Washington, D.C. 43 pages. https://freedomhouse.org/sites/default/files/2025-12/FOTN%202025_final_digital_120525.pdf
- [145] Rebekah Houser, Shuai Hao, Zhou Li, Daiping Liu, Chase Cotton, and Haining Wang. 2021. A Comprehensive Measurement-Based Investigation of DNS Hijacking. In *Symposium on Reliable Distributed Systems*. IEEE, Virtual Event, USA, 11. <https://doi.org/10.1109/SRDS53918.2021.00029>

- [146] Philip N. Howard, Sheetal D. Agarwal, and Muzammil M. Hussain. 2015. The Hermit Kingdom in Cyberspace: Unveiling the North Korean Internet. *Information, Communication & Society* 18, 3 (2015), 252–267. <https://doi.org/10.1080/1369118X.2014.940363>
- [147] HTTP Archive. 2024. HTTP - The Web Almanac 2024. <https://almanac.httparchive.org/en/2024/http>
- [148] Jinling Hua and Rajib Shaw. 2020. Corona Virus (COVID-19) “Infodemic” and Emerging Issues through a Data Lens: The Case of China. *International Journal of Environmental Research and Public Health* 17, 7 (March 2020), 2309. <https://doi.org/10.3390/ijerph17072309>
- [149] Troy Hunt and Scott Helme. 2021. Why No HTTPS? <https://whynohttps.com/>
- [150] Saro Husaini. 2024. *Circumventing Internet Censorship*. Thesis. University of South-Eastern Norway. <https://openarchive.usn.no/usn-xmlui/handle/11250/3135775>
- [151] ICMEC. 2021. The Growing Global Threat of Child Sexual Abuse Material (CSAM). <https://icmec.org.au/blog/the-growing-global-threat-of-child-sexual-abuse-material-csam/>
- [152] OpenNet Initiative. 2009. *Internet Filtering in China*. Technical Report. OpenNet Initiative, Technical Report. http://opennet.net/sites/opennet.net/files/ONI_China_2009.pdf
- [153] Internet Society Pulse. 2021. *Internet Shutdown in Jammu and Kashmir, India*. Internet Society. <https://pulse.internetsociety.org/en/shutdowns/india-110/>
- [154] Kevin Jacobs. 2021. Encrypted Client Hello: The Future of ESNi in Firefox. <https://blog.mozilla.org/security/2021/01/07/encrypted-client-hello-the-future-of-esni-in-firefox/>
- [155] Santosh Janardhan. 2021. More Details about the October 4 Outage. <https://engineering.fb.com/2021/10/05/networking-traffic/outage-details/>
- [156] Andrea Januta and Minami Funakoshi. 2021. Myanmar’s Internet Suppression. <https://www.reuters.com/graphics/MYANMAR-POLITICS/INTERNET-RESTRICTION/rlgpdbreepo/>
- [157] Jill Jermyn and Nicholas Weaver. 2017. Autosonda: Discovering Rules and Triggers of Censorship Devices. In *Free and Open Communications on the Internet*. USENIX, Vancouver, BC, Canada, 9. <https://www.usenix.org/system/files/conference/foci17/foci17-paper-jermyn.pdf>
- [158] Lin Jin, Shuai Hao, Haining Wang, and Chase Cotton. 2021. Understanding the Practices of Global Censorship through Accurate, End-to-End Measurements. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 5, 3 (Dec. 2021), 1–25. <https://doi.org/10.1145/3491055>
- [159] Ben Jones and Nick Feamster. 2015. Can Censorship Measurements Be Safe(r)? In *HotNets ’15*. ACM, Philadelphia, PA, 7. <http://conferences.sigcomm.org/hotnets/2015/papers/jones.pdf>
- [160] Ben Jones, Tzu-Wen Lee, Nick Feamster, and Phillipa Gill. 2014. Automated Detection and Fingerprinting of Censorship Block Pages. In *Proceedings of the Internet Measurement Conference 2014*. ACM, Vancouver BC Canada, 299–304. <https://doi.org/10.1145/2663716.2663722>
- [161] Kadak. 2018. How to Block Traffic Based on Application Filters with an Exception. <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000CIXfCAK>
- [162] Cecilia Kang, Davey Alba, and Adam Satariano. 2020. Surging Traffic Is Slowing Down Our Internet. <https://www.nytimes.com/2020/03/26/business/coronavirus-internet-traffic-speed.html>
- [163] Nikolaos Karapanos and Srdjan Capkun. 2014. On the Effective Prevention of TLS Man-In-The-Middle Attacks in Web Applications. In *USENIX Security Symposium*. USENIX, San Diego, CA, 17. <https://www.usenix.org/system/files/conference/usenixsecurity14/sec14-paper-karapanos.pdf>
- [164] Divyank Katira, Gurshabad Grover, Kushagra Singh, and Varun Bansal. 2023. CensorWatch: On the Implementation of Online Censorship in India. In *Free and Open Communications on the Internet*. Proceedings on Privacy Enhancing Technologies, Lausanne, Switzerland, 12.
- [165] Lukas Kawerau, Nils B. Weidmann, and Alberto Dainotti. 2023. Attack or Block? Repertoires of Digital Censorship in Autocracies. *Journal of Information Technology & Politics* 20, 1 (Jan. 2023), 60–73. <https://doi.org/10.1080/19331681.2022.2037118>
- [166] Mohammad Taha Khan, Joe DeBlasio, Geoffrey M. Voelker, Alex C. Snoeren, Chris Kanich, and Narseo Vallina-Rodriguez. 2018. An Empirical Analysis of the Commercial VPN Ecosystem. In *Proceedings of the Internet Measurement Conference 2018*. ACM, Boston, MA, USA, 443–456.
- [167] Sheharbano Khattak, David Fifield, Sadia Afroz, Mobin Javed, Srikanth Sundaresan, Vern Paxson, Steven J. Murdoch, and Damon McCoy. 2016. Do You See What I See? Differential Treatment of Anonymous Users. In *Proceedings 2016 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA, 15. <https://doi.org/10.14722/ndss.2016.23342>
- [168] Sheharbano Khattak, Mobin Javed, Syed Ali Khayam, Zartash Afzal Uzmi, and Vern Paxson. 2014. A Look at the Consequences of Internet Censorship Through an ISP Lens. In *Internet Measurement Conference*. ACM, Barcelona, Spain, 13. <http://conferences2.sigcomm.org/imc/2014/papers/p271.pdf>
- [169] Aminollah Khormali, Jeman Park, Hisham Alasmay, Afsah Anwar, Muhammad Saad, and David Mohaisen. 2021. Domain Name System Security and Privacy: A Contemporary Survey. *Computer Networks* 185 (Feb. 2021), 107699. <https://doi.org/10.1016/j.comnet.2020.107699>
- [170] Emil Kiner and Tim April. 2023. Google Mitigated the Largest DDoS Attack to Date, Peaking above 398 Million Rps. <https://cloud.google.com/blog/products/identity-security/google-cloud-mitigated-largest-ddos-attack-peaking-above-398-million-rps>
- [171] Emil Kiner and Satya Konduru. 2022. How Google Cloud Blocked the Largest Layer 7 DDoS Attack at 46 Million Rps. <https://cloud.google.com/blog/products/identity-security/how-google-cloud-blocked-largest-layer-7-ddos-attack-at-46-million-rps>
- [172] Gary King, Jennifer Pan, and Margaret E. Roberts. 2013. How Censorship in China Allows Government Criticism but Silences Collective Expression. *American Political Science Review* 1 (2013), 18. <https://gking.harvard.edu/files/censored.pdf>
- [173] Gary King, Jennifer Pan, and Margaret E. Roberts. 2014. Reverse-Engineering Censorship in China: Randomized Experimentation and Participant Observation. *Science* 345, 6199 (2014), 11. <http://cryptome.org/2014/08/reverse-eng-cn-censorship.pdf>
- [174] Jeffrey Knockel, Masashi Crete-Nishihata, Jason Q. Ng, Adam Senft, and Jedidiah R. Crandall. 2015. Every Rose Has Its Thorn: Censorship and Surveillance on Social Video Platforms in China. In *Free and Open Communications on the Internet*. USENIX, Washington, D.C., 10.
- [175] Jeffrey Knockel and Lotus Ruan. 2021. Measuring QQMail’s Automated Email Censorship in China. In *Free and Open Communications on the Internet*. ACM, Virtual Event, 8. <https://doi.org/10.1145/3473604.3474560>
- [176] Jeff Kosseff. 2019. *The Twenty-Six Words That Created the Internet*. Cornell University Press, Ithaca, New York.

- [177] Carmen Kwan, Paul Janiszewski, Shela Qiu, Cathy Wang, and Cecylia Bocovich. 2021. Exploring Simple Detection Techniques for DNS-over-HTTPS Tunnels. In *Free and Open Communications on the Internet*. ACM, Virtual Event, USA, 6. <https://doi.org/10.1145/3473604.3474563>
- [178] Yumi Lamansky. 2024. *TCP Segmentation to Evade NIDS*. MS Thesis. Arizona State University. <https://keep.lib.asu.edu/items/198214>
- [179] Fangfan Li, Abbas Razaghpanah, Arash Molavi Kakhki, Arian Akhavan Niaki, David Choffnes, Phillipa Gill, and Alan Mislove. 2017. Lib•erate, (n): A Library for Exposing (Traffic-Classification) Rules and Avoiding Them Efficiently. In *Proceedings of the 2017 Internet Measurement Conference*. ACM, London, United Kingdom, 128–141. <https://doi.org/10.1145/3131365.3131376>
- [180] Wilson Lian, Eric Rescorla, Hovav Shacham, and Stefan Savage. 2013. Measuring the Practical Impact of DNSSEC Deployment. *USENIX Security Symposium* 1 (Aug. 2013), 16. https://www.usenix.org/system/files/conference/usenixsecurity13/sec13-paper_lian.pdf
- [181] IVPN Limited. 2023. Transparent DNS Proxy, Dnsleaktest.Com. <https://dnsleaktest.com/what-is-transparent-dns-proxy.html>
- [182] G. Loukas and G. Oke. 2010. Protection Against Denial of Service Attacks. *Computer Journal* 53 (2010), 17. <https://doi.org/10.1093/comjnl/bxp078>
- [183] Graham Lowe, Patrick Winters, and Michael L Marcus. 2007. *The Great DNS Wall of China*. Technical Report. New York University, Technical Report, New York University. <https://censorbib.nymity.ch/pdf/Lowe2007a.pdf>
- [184] Zhen Lu, Zhenhua Li, Jian Yang, Tianyin Xu, Ennan Zhai, Yao Liu, and Christo Wilson. 2017. Accessing Google Scholar under Extreme Internet Censorship: A Legal Avenue. In *Middleware*. ACM, Las Vegas, NV, 7. <https://censorbib.nymity.ch/pdf/Lu2017a.pdf>
- [185] Philipp M. Lutscher, Nils B. Weidmann, Margaret E. Roberts, Mattijs Jonker, Alistair King, and Alberto Dainotti. 2020. At Home and Abroad: The Use of Denial-of-service Attacks during Elections in Nondemocratic Regimes. *Journal of Conflict Resolution* 64, 2-3 (Feb. 2020), 373–401. <https://doi.org/10.1177/0022002719861676>
- [186] Matthew Luxmoore and Milan Czerny. 2026. *Russia Shuts Off Internet in Moscow as It Tests Nationwide Censorship System*. The Wall Street Journal. <https://www.wsj.com/world/russia/russia-shuts-off-internet-in-moscow-as-it-tests-nationwide-censorship-system-3b44c0af>
- [187] J. Lyengar and M. Thomson. 2021. RFC 9000: A UDP-Based Multiplexed and Secure Transport. <https://doi.org/10.17487/RFC9000>
- [188] Tambiana Madiega and Stefano De Luca. 2023. Collapse of the Internet. In *Future Shocks 2023: Anticipating and Weathering the Next Storms*. European Parliament, EU, 34–39. <https://doi.org/10.2861/88235>
- [189] Tasnuva Mahjabin, Yang Xiao, Guang Sun, and Wangdong Jiang. 2017. A Survey of Distributed Denial-of-Service Attack, Prevention, and Mitigation Techniques. *International Journal of Distributed Sensor Networks* 13, 12 (Dec. 2017), 155014771774146. <https://doi.org/10.1177/1550147717741463>
- [190] Mousa Taghizadeh Manavi. 2018. Defense Mechanisms against Distributed Denial of Service Attacks : A Survey. *Computers & Electrical Engineering* 72 (Nov. 2018), 26–38. <https://doi.org/10.1016/j.compeleceng.2018.09.001>
- [191] Antonio Mangino and Elias Bou-Harb. 2021. A Multidimensional Network Forensics Investigation of a State-Sanctioned Internet Outage. In *IWCMC 2021*. IEEE, Harbin City, China, 813–818. <https://doi.org/10.1109/IWCMC51323.2021.9498743>
- [192] Bill Marczak, Nicholas Weaver, Jakub Dalek, Roya Ensafi, David Fifield, Sarah McKune, Arn Rey, John Scott-Railton, Ron Deibert, and Vern Paxson. 2015. An Analysis of China’s “Great Cannon”. In *Free and Open Communications on the Internet*. USENIX, Washington, D.C., 11. <https://www.usenix.org/system/files/conference/foci15/foci15-paper-marczak.pdf>
- [193] Bill Marczak, Nicholas Weaver, Jakub Dalek, Roya Ensafi, David Fifield, Sarah McKune, Arn Rey, John Scott-Railton, Ron Deibert, and Vern Paxson. 2015. *China’s Great Cannon*. Citizen Lab, Munk School of Global Affairs, University of Toronto. <https://citizenlab.ca/research/chinas-great-cannon/>
- [194] Alexander Marrow and Dmitry Antonov. 2021. *Russia Disconnected from Global Internet in Tests: RBC Daily*. Reuters. <https://www.reuters.com/technology/russia-disconnected-global-internet-tests-rbc-daily-2021-07-22/>
- [195] L. Masinter and M. McCahill. 1994. RFC 1738: Uniform Resource Locators (URL). <https://datatracker.ietf.org/doc/html/rfc1738>
- [196] Alexander Master. 2023. *Modeling and Characterization of Internet Censorship Technologies*. Ph.D. Dissertation. Purdue University. <https://doi.org/10.25394/PGS.23666784>
- [197] Alexander Master and Christina Garman. 2023. A Worldwide View of Nation-state Internet Censorship. In *Free and Open Communications on the Internet*. Proceedings on Privacy Enhancing Technologies, Lausanne, Switzerland, 22. <https://www.petsymposium.org/foci/2023/foci-2023-0008.pdf>
- [198] Alexander Master and Kendall Gomber. 2026. The PRC and Cyberspace: Internet Censorship, the Social Credit System, and Chinese Cyberspace Operations. In *De Gruyter Handbook of Cyber Conflict and Warfare*. De Gruyter, DE. <https://doi.org/10.1515/978311613031-016>
- [199] Matan Mates. 2022. *Tor, Meek & The Rise And Fall Of Domain Fronting*. Sentinel One. <https://www.sentinelone.com/blog/privacy-2019-tor-meek-rise-fall-domain-fronting/> Updated October 27, 2022.
- [200] Allison McDonald, Matthew Bernhard, Luke Valenta, Benjamin VanderSloot, Will Scott, Nick Sullivan, J. Alex Halderman, and Roya Ensafi. 2018. 403 Forbidden: A Global View of CDN Geoblocking. In *Internet Measurement Conference*, Vol. 1. ACM, Boston, MA, USA, 13.
- [201] Molly McHugh. 2011. Burma Bans Skype, Severing Global Communication. <https://www.digitaltrends.com/computing/burma-bans-skype-severing-global-communication/>
- [202] Damian Menscher. 2020. Exponential Growth in DDoS Attack Volumes. <https://cloud.google.com/blog/products/identity-security/identifying-and-protecting-against-the-largest-ddos-attacks>
- [203] Jelena Mirkovic and Peter Reiher. 2004. A Taxonomy of DDoS Attack and DDoS Defense Mechanisms. *ACM SIGCOMM Computer Communication Review* 34, 2 (April 2004), 39–53. <https://doi.org/10.1145/997150.997156>
- [204] Asya Mitseva, Andriy Panchenko, and Thomas Engel. 2018. The State of Affairs in BGP Security: A Survey of Attacks and Defenses. *Computer Communications* 124 (June 2018), 45–60. <https://doi.org/10.1016/j.comcom.2018.04.013>
- [205] Hooman Mohajeri Moghaddam, Baiyu Li, Mohammad Derakhshani, and Ian Goldberg. 2012. SkypeMorph: Protocol Obfuscation for Tor Bridges. In *ACM CCS 2012*. ACM Press, Raleigh, NC, 97. <https://doi.org/10.1145/2382196.2382210>
- [206] David Moher, Alessandro Liberati, Jennifer Tetzlaff, Douglas G. Altman, and The PRISMA Group. 2009. Preferred Reporting Items for Systematic Reviews and Meta-Analyses: The PRISMA Statement. *PLoS Medicine* 6, 7 (July 2009), e1000097. <https://doi.org/10.1371/journal.pmed.1000097>
- [207] Mozilla. 2023. Encrypted Client Hello (ECH) - Frequently Asked Questions. <https://support.mozilla.org/en-US/kb/faq-encrypted-client-hello>

- [208] Lennart Mühlenmeier. 2020. Jordan Does Not Block, It Throttles Internet Access. <https://netzpolitik.org/2020/jordan-throttles-not-blocks-internet-access-shutdowns-keepit/#netzpolitik-pw>
- [209] Philipp Müller, Niklas Niere, Felix Lange, and Juraj Somorovsky. 2024. Turning Attacks into Advantages: Evading HTTP Censorship with HTTP Request Smuggling. In *Free and Open Communications on the Internet*. Proceedings on Privacy Enhancing Technologies, Bristol, UK, 12. <https://www.petsymposium.org/foci/2024/foci-2024-0012.pdf>
- [210] Zubair Nabi. 2013. The Anatomy of Web Censorship in Pakistan. In *Free and Open Communications on the Internet*. USENIX, Washington, D.C., 7. <https://www.usenix.org/system/files/conference/foci13/foci13-nabi.pdf>
- [211] United Nations. 1948. Universal Declaration of Human Rights: Article 19. <https://www.un.org/en/about-us/universal-declaration-of-human-rights>
- [212] United Nations. 2021. Right to Freedom of Expression in Myanmar Must Be Guaranteed, UN Expert Urges Military Coup Leader. <https://news.un.org/en/story/2021/04/1090742>
- [213] Jose Nazario. 2009. Politically Motivated Denial of Service Attacks. In *Virtual Battlefield: Perspectives on Cyber Warfare*. Vol. 3. IOS Press, Ebook, 20.
- [214] RIPE NCC. 2008. YouTube Hijacking: A RIPE NCC RIS Case Study. <https://web.archive.org/web/20080405030750/http://www.ripe.net/news/study-youtube-hijacking.html>
- [215] Palo Alto Networks. 2018. How to Block a Specific HTTPS Site with URL Filtering. <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000CIIPCA0>
- [216] Palo Alto Networks. 2018. How to Block Tor. <https://knowledgebase.paloaltonetworks.com/KCSArticleDetail?id=kA10g000000CIRtCAK>
- [217] Lily Hay Newman. 2019. How the Iranian Government Shut Off the Internet. <https://www.wired.com/story/iran-internet-shutoff/>
- [218] Arian Akhavan Niaki, Shinyoung Cho, Zachary Weinberg, Nguyen Phong Hoang, Abbas Razaghpanah, Nicolas Christin, and Phillipa Gill. 2020. ICLab: A Global, Longitudinal Internet Censorship Measurement Platform. In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Francisco, CA, USA, 135–151. <https://doi.org/10.1109/SP40000.2020.00014>
- [219] Niklas Niere, Felix Lange, Nico Heitmann, and Juraj Somorovsky. 2025. Encrypted Client Hello (ECH) in Censorship Circumvention. In *Free and Open Communications on the Internet*. FOCI, Washington D.C., 10. <https://petsymposium.org/foci/2025/foci-2025-0016.pdf>
- [220] Niklas Niere, Felix Lange, Robert Merget, and Juraj Somorovsky. 2025. Transport Layer Obscurity: Circumventing SNI Censorship on the TLS Layer. In *Symposium on Security & Privacy*. IEEE, San Francisco, USA, 18. <https://www.computer.org/csdl/pds/api/csdl/proceedings/download-article/26hiUekZ19S/pdf>
- [221] Aqib Nisar, Aqsa Kashaf, Ihsan Ayyub Qazi, and Zartash Afzal Uzmi. 2018. Incentivizing Censorship Measurements via Circumvention. In *SIGCOMM*. ACM, Budapest, Hungary, 14. <https://doi.org/10.1145/3230543.3230568>
- [222] Leif Nixon. 2011. Some Observations on the Great Firewall of China. <https://www.nsc.liu.se/~nixon/sshprobes.html>
- [223] Daiyuu Nobori and Yasushi Shinjo. 2014. VPN Gate: A Volunteer-Organized Public VPN Relay System with Blocking Resistance for Bypassing Government Censorship Firewalls. In *Networked Systems Design and Implementation*. USENIX, Philadelphia, PA, 14. <https://www.usenix.org/system/files/conference/nsdi14/nsdi14-paper-nobori.pdf>
- [224] Sadia Nourin, Kevin Bock, Nguyen Phong Hoang, and Dave Levin. 2023. Detecting Network Interference without Endpoint Participation. In *Free and Open Communications on the Internet*. Proceedings on Privacy Enhancing Technologies, Lausanne, Switzerland, 5. <https://www.petsymposium.org/foci/2023/foci-2023-0010.pdf>
- [225] Sadia Nourin, Van Tran, Xi Jiang, Kevin Bock, Nick Feamster, Nguyen Phong Hoang, and Dave Levin. 2023. Measuring and Evading Turkmenistan's Internet Censorship: A Case Study in Large-Scale Measurements of a Low-Penetration Country. In *Proceedings of the ACM Web Conference 2023*. ACM, Austin TX USA, 11. <https://doi.org/10.1145/3543507.3583189>
- [226] ntop. 2022. nDPI - Open and Extensible LGPLv3 Deep Packet Inspection Library. <https://www.ntop.org/products/deep-packet-inspection/ndpi/>
- [227] Jonathan Oakley, Lu Yu, Xingsi Zhong, Ganesh Kumar Venayagamoorthy, and Richard Brooks. 2020. Protocol Proxy: An FTE-based Covert Channel. *Computers & Security* 92 (2020), 11. <https://censorbib.nymity.ch/pdf/Oakley2020a.pdf>
- [228] obfuscation group. 2015. Pluggable Transports. <https://obfuscation.github.io/>
- [229] Katherine Ognyanova, David Lazer, Ronald Robertson, and Christo Wilson. 2020. Misinformation in Action: Fake News Exposure Is Linked to Lower Trust in Media, Higher Trust in Government When Your Side Is in Power. *The Harvard Kennedy School Misinformation Review* 1, 4 (June 2020), 12. <https://doi.org/10.37016/mr-2020-024>
- [230] Chitu Okoli and Kira Schabram. 2010. A Guide to Conducting a Systematic Literature Review of Information Systems Research. *SSRN Electronic Journal* 1 (2010), 51. <https://doi.org/10.2139/ssrn.1954824>
- [231] Juan Ortiz Freuler. 2022. The Weaponization of Private Corporate Infrastructure: Internet Fragmentation and Coercive Diplomacy in the 21st Century. *Global Media and China* 1 (Nov. 2022), 18. <https://doi.org/10.1177/20594364221139729>
- [232] Aaron Ortwein, Kevin Bock, and Dave Levin. 2023. Towards a Comprehensive Understanding of Russian Transit Censorship. In *Free and Open Communications on the Internet*. PoPETS, Lausanne, Switzerland, 7. <https://www.petsymposium.org/foci/2023/foci-2023-0012.pdf>
- [233] Ramakrishna Padmanabhan, Arturo Filastò, Maria Xynou, Ram Sundara Raman, Kennedy Middleton, Mingwei Zhang, Doug Madory, Molly Roberts, and Alberto Dainotti. 2021. A Multi-Perspective View of Internet Censorship in Myanmar. In *Free and Open Communications on the Internet*. ACM, Virtual Event, USA, 10. <https://doi.org/10.1145/3473604.3474562>
- [234] Jong Chun Park and Jeddiah R. Crandall. 2010. Empirical Study of a National-Scale Distributed Intrusion Detection System: Backbone-Level Filtering of HTML Responses in China. In *Distributed Computing Systems*. IEEE, Genoa, Italy, 12. <https://www.cs.unm.edu/~crandall/icdcs2010.pdf>
- [235] Paul Pearce, Ben Jones, Frank Li, Roya Ensafi, Nick Feamster, Nick Weaver, and Vern Paxson. 2017. Global Measurement of DNS Manipulation. In *USENIX Security*. USENIX, Vancouver, Canada, 19. <https://www.usenix.org/system/files/conference/usenixsecurity17/sec17-pearce.pdf>
- [236] Hugo Gamaliel Dos Santos Pereira. 2024. *MIRACE: Multi-Path Integrated Routing Architecture for Censorship Evasion*. Thesis. Universidade Nova De Lisboa. https://run.unl.pt/bitstream/10362/178441/1/Pereira_2024.pdf

- [237] Jay Peters. 2024. Signal Has Been Blocked by Venezuela and Russia. <https://www.theverge.com/2024/8/9/24217008/signal-blocked-venezuela-russia>
- [238] Louis Poinssignon. 2018. BGP Leaks and Cryptocurrencies. <https://blog.cloudflare.com/bgp-leaks-and-crypto-currencies/>
- [239] PowerTunnel. 2023. PowerTunnel - Powerful and Extensible Proxy Server. <https://github.com/krlvm/PowerTunnel>
- [240] Ram Sundara Raman, Leonid Evdokimov, Eric Wurstrow, J. Alex Halderman, and Roya Ensafi. 2020. Investigating Large Scale HTTPS Interception in Kazakhstan. In *Proceedings of the ACM Internet Measurement Conference*. ACM, Virtual Event, USA, 125–132. <https://doi.org/10.1145/3419394.3423665>
- [241] Ram Sundara Raman, Adrian Stoll, Jakub Dalek, Reethika Ramesh, Will Scott, and Roya Ensafi. 2020. Measuring the Deployment of Network Censorship Filters at Global Scale. In *Proceedings of 2020 NDSS*. NDSS, San Diego, CA, 16. <https://doi.org/10.14722/ndss.2020.23099>
- [242] Ram Sundara Raman, Mona Wang, Jakub Dalek, Jonathan Mayer, and Roya Ensafi. 2022. Network Measurement Methods for Locating and Examining Censorship Devices. In *Emerging Networking Experiments and Technologies*. ACM, Roma, Italy, 17. <https://doi.org/10.1145/3555050.3569133>
- [243] Reethika Ramesh, Leonid Evdokimov, Diwen Xue, and Roya Ensafi. 2022. VPNalyzer: Systematic Investigation of the VPN Ecosystem. In *Proceedings 2022 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA, USA, 16. <https://doi.org/10.14722/ndss.2022.24285>
- [244] Reethika Ramesh, Ram Sundara Raman, Matthew Bernhard, Victor Ongkowitz, Leonid Evdokimov, Anne Edmundson, Steven Sprecher, Muhammad Ikram, and Roya Ensafi. 2020. Decentralized Control: A Case Study of Russia. In *Proceedings 2020 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA, 18. <https://doi.org/10.14722/ndss.2020.23098>
- [245] Reethika Ramesh, Ram Sundara Raman, Apurva Virkud, Alexandra Dirksen, Armin Huremagic, David Fifield, Dirk Rodenburg, Rod Hynes, Doug Madory, and Roya Ensafi. 2023. Network Responses to Russia’s Invasion of Ukraine in 2022: A Cautionary Tale for Internet Freedom. In *USENIX Security Symposium*. USENIX, Anaheim, CA, USA, 19. <https://censoredplanet.org/assets/russia-ukraine-invasion.pdf>
- [246] Amanda Reid. 2019. Considering Fair Use: DMCA’s Take Down & Repeat Infringers Policies. *Communication Law and Policy* 24, 1 (Jan. 2019), 101–141. <https://doi.org/10.1080/10811680.2018.1551036>
- [247] Eric Rescorla, Kazuho Oku, Nick Sullivan, and Christopher Wood. 2022. *TLS Encrypted Client Hello, Internet-Draft RFC*. Draft RFC. Internet Engineering Task Force (IETF). <https://datatracker.ietf.org/doc/draft-ietf-tls-esni/15>
- [248] Reuters. 2022. Russia Blocks Access to BBC and Voice of America Websites. <https://www.reuters.com/business/media-telecom/russia-restricts-access-bbc-russian-service-radio-liberty-ria-2022-03-04/>
- [249] Hal Roberts, Ethan Zuckerman, Jillian York, Robert Faris, and John Palfrey. 2010. 2010 Circumvention Tool Usage Report. *Berkman Center Research Publication No. 2010 1* (2010), 13. https://cyber.harvard.edu/sites/cyber.harvard.edu/files/2010_Circumvention_Tool_Usage_Report.pdf
- [250] Margaret E. Roberts. 2018. *Blocked: How Control of China’s Internet and Social Media Allows the State to Dictate the Narrative*. Cambridge University Press, Cambridge. <https://doi.org/10.1017/9781108551502>
- [251] Margaret E. Roberts. 2020. *Censored: Distraction and Diversion inside China’s Great Firewall*. Princeton University Press, Princeton.
- [252] Marc B. Rosen, James Parker, and Alex J. Malozemoff. 2021. Balboa: Bobbing and Weaving around Network Censorship. In *USENIX Security Symposium*. USENIX, Virtual Event, USA, 16. <https://www.usenix.org/system/files/sec21-rosen.pdf>
- [253] Roskomnadzor. 2021. Roskomnadzor Takes Measures to Protect Russian Citizens from the Influence of Illegal Content. <https://rkn.gov.ru/news/rsoc/news73464.htm>
- [254] Samuel Ruø, Jeffrey Knockel, and Zoë Reichert. 2024. Lost in Translation: Characterizing Automated Censorship in Online Translation Services. In *Free and Open Communications on the Internet*. Proceedings on Privacy Enhancing Technologies, Bristol, UK, 9.
- [255] Malise Ruthven. 2016. How to Understand ISIS. <https://www.nybooks.com/articles/2016/06/23/how-to-understand-isis/>
- [256] Julia Ryng, Guillemette Guicherd, Judy Al Saman, Priyanka Choudhury, and Angharad Kellett. 2022. Internet Shutdowns: A Human Rights Issue. *The RUSI Journal* 167, 4-5 (July 2022), 50–63. <https://doi.org/10.1080/03071847.2022.2156234>
- [257] Nicholas Michael Sambaluk and Eugene Spafford. 2020. *Myths and Realities of Cyber Warfare: Conflict in the Digital Realm*. Praeger Security International, Santa Barbara, CA.
- [258] Adam Satariano. 2022. How Russia Took Over Ukraine’s Internet in Occupied Territories. <https://www.nytimes.com/interactive/2022/08/09/technology/ukraine-internet-russia-censorship.html>
- [259] Sambhav Satija and Rahul Chatterjee. 2021. BlindTLS: Circumventing TLS-Based HTTPS Censorship. In *Free and Open Communications on the Internet*. ACM, Virtual Event, USA, 7. <https://doi.org/10.1145/3473604.3474564>
- [260] Molly Sauter. 2014. *The Coming Swarm: DDoS Actions, Hacktivism, and Civil Disobedience on the Internet* (1 ed.). Bloomsbury Publishing Inc, New York, NY. <https://doi.org/10.5040/9781628926705>
- [261] Jared Schroeder. 2022. Meet the EU Law That Could Reshape Online Speech in the U.S. <https://slate.com/technology/2022/10/digital-services-act-european-union-content-moderation.html>
- [262] Will Scott, Thomas Anderson, Tadayoshi Kohno, and Arvind Krishnamurthy. 2016. Satellite: Joint Analysis of CDNs and Network-Level Interference. In *Proceedings of 2016 USENIX Conference*. USENIX, Austin, TX, 15. https://www.usenix.org/system/files/conference/atc16/atc16_paper-scott.pdf
- [263] Ryan Serabian and Daniel Kapellmann Zafra. 2022. Pro-PRC “HaiEnergy” Information Operations Campaign Leverages Infrastructure from Public Relations Firm to Disseminate Content on Inauthentic News Sites. <https://www.mandiant.com/resources/blog/pro-prc-information-operations-campaign-haienergy>
- [264] Andreas Sfakianakis, Elias Athanasopoulos, and Sotiris Ioannidis. 2011. CensMon: A Web Censorship Monitor. In *Free and Open Communications on the Internet*. USENIX, Bellevue, WA, 6. https://www.usenix.org/legacy/events/foci11/tech/final_files/Sfakianakis.pdf
- [265] Nishant Shah. 2021. (Dis)Information Blackouts: Politics and Practices of Internet Shutdowns. *International Journal of Communication* 15 (2021), 17.
- [266] Y. Shavitt and N. Zilberman. 2012. Arabian Nights: Measuring the Arab Internet during the 2011 Events. *IEEE Network* 26, 6 (Nov. 2012), 75–80. <https://doi.org/10.1109/MNET.2012.6375897>
- [267] Jade Sheffey, Ali Zohaib, Dayeon Kang, Zakir Durumeric, Amir Houmansadr, and Qiang Wu. 2025. Extended Abstract: I’ll Shake Your Hand: What Happens After DNS Poisoning. In *Free and Open Communications on the Internet*. FOCI, Washington D.C., 4. <https://petsymposium.org/foci/2025/foci->

2025-0015.pdf

- [268] Xiaoxiao Shen and Rory Truex. 2021. In Search of Self-Censorship. *British Journal of Political Science* 51, 4 (Oct. 2021), 1672–1684. <https://doi.org/10.1017/S0007123419000735>
- [269] Justine Sherry, Chang Lan, Raluca Popa, and Sylvia Ratnasamy. 2015. BlindBox: Deep Packet Inspection over Encrypted Traffic. In *Proceedings of 2015 ACM Conference on Special Interest Group on Data Communication*. ACM, London, UK, 13. <https://doi.org/10.1145/2785956.2787502>
- [270] Sandra Siby, Marc Juarez, Claudia Diaz, Narseo Vallina-Rodriguez, and Carmela Troncoso. 2020. Encrypted DNS -> Privacy? A Traffic Analysis Perspective. In *Proceedings of 2020 NDSS*. Internet Society, San Diego, CA, 19. <https://doi.org/10.14722/ndss.2020.24301>
- [271] Kushagra Singh, Gurshabad Grover, and Varun Bansal. 2020. How India Censors the Web. In *12th ACM Conference on Web Science*. ACM, Southampton United Kingdom, 21–28. <https://doi.org/10.1145/3394231.3397891>
- [272] Rachee Singh, Rishab Nithyanand, Sadia Afroz, Paul Pearce, Michael Tschantz, Phillipa Gill, and Vern Paxson. 2017. Characterizing the Nature and Dynamics of Tor Exit Blocking. In *USENIX Security Symposium*. USENIX, Vancouver, Canada, 19. <https://www.usenix.org/system/files/conference/usenixsecurity17/sec17-singh.pdf>
- [273] Charlie Smith. 2015. We Are under Attack. <https://en.greatfire.org/blog/2015/mar/we-are-under-attack>
- [274] Andrei Soldatov and Irina Borogan. 2015. *The Red Web: The Struggle Between Russia's Digital Dictators and the New Online Revolutionaries*. PublicAffairs, New York, NY.
- [275] Reuters Staff. 2018. Businesses, Consumers Uncertain Ahead of China VPN Ban. <https://www.reuters.com/article/us-china-vpns/businesses-consumers-uncertain-ahead-of-china-vpn-ban-idUSKBN1H612F>
- [276] Statista. 2024. Global Market Share Held by Leading Internet Browsers from January 2012 to August 2024. <https://www.statista.com/statistics/268254/market-share-of-internet-browsers-worldwide-since-2009/>
- [277] Ram Sundara Raman, Louis-Henri Merino, Kevin Bock, Marwan Fayed, Dave Levin, Nick Sullivan, and Luke Valenta. 2023. Global, Passive Detection of Connection Tampering. In *Proceedings of the ACM SIGCOMM*. ACM, New York, NY, 622–636. <https://doi.org/10.1145/3603269.3604875>
- [278] Ram Sundara Raman, Prerana Shenoy, Katharina Kohls, and Roya Ensafi. 2020. Censored Planet: An Internet-wide, Longitudinal Censorship Observatory. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. ACM, Virtual Event, USA, 49–66. <https://doi.org/10.1145/3372297.3417883>
- [279] Jonas Tai, Karthik Nishanth Sengottuvelavan, Peter Whiting, and Nguyen Phong Hoang. 2025. IRBlock: A Large-Scale Measurement Study of the Great Firewall of Iran. In *USENIX Security Symposium*. USENIX, Seattle, USA, 19. <https://www.usenix.org/system/files/usenixsecurity25-tai.pdf>
- [280] Rima Tanash, Zhouhan Chen, Dan Wallach, and Melissa Marschall. 2017. The Decline of Social Media Censorship and the Rise of Self-Censorship after the 2016 Failed Turkish Coup. In *Free and Open Communications on the Internet*. USENIX, Vancouver, BC, Canada, 7. <https://www.usenix.org/system/files/conference/foci17/foci17-paper-tanash.pdf>
- [281] Hacker Target. 2022. WHOIS ASN IP Lookup Tool. <https://hackertarget.com/as-ip-lookup/>
- [282] BITAG Technical Working Group. 2013. *Port Blocking*. Technical Report. BITAG, Virtual. 34 pages. <http://www.ssrn.com/abstract=2701485>
- [283] Eneken Tikk, Kadri Kaska, Kristel Rünnermeri, Mari Kert, Anna-Maria Talihärm, and Liis Vihul. 2008. *Cyber Attacks Against Georgia: Legal Lessons Identified*. Technical Report. Cooperative Cyber Defense Center of Excellence, Tallinn, Estonia. 46 pages. <http://www.carlisle.army.mil/DIME/documents/Georgia%201%200.pdf>
- [284] Adam Tinworth. 2018. Happy GDPR Day! <https://twitter.com/adders/status/999976701089189890/photo/1>
- [285] Tor. 2022. Tor Usage Statistics during Invasion of Ukraine 2022. https://web.archive.org/web/20230105232136/https://web.ics.purdue.edu/~amaster/anticensorship/Tor_usage_Ukraine_invasion.png
- [286] Mark Trevelyan. 2022. Russians' Demand for VPNs Skyrockets after Meta Block. <https://www.reuters.com/technology/russians-demand-vpns-skyrockets-after-meta-block-2022-03-14/>
- [287] Nikhil Tripathi and Neminath Hubballi. 2021. Application Layer Denial-of-Service Attacks and Defense Mechanisms: A Survey. *Comput. Surveys* 54, 4 (April 2021), 1–33. <https://doi.org/10.1145/3448291>
- [288] Elisa Tsai, Ram Sundara Raman, Atul Prakash, and Roya Ensafi. 2024. Modeling and Detecting Internet Censorship Events. In *Network and Distributed System Security*. NDSS, San Diego, CA, USA, 17. <https://www.ndss-symposium.org/wp-content/uploads/2024-409-paper.pdf>
- [289] Michael Carl Tschantz, Sadia Afroz, Anonymous, and Vern Paxson. 2016. SoK: Towards Grounding Censorship Circumvention in Empiricism. In *2016 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Jose, CA, 914–933. <https://doi.org/10.1109/SP.2016.59>
- [290] Michael Carl Tschantz, Sadia Afroz, Shaarif Sajid, Shoaib Asif Qazi, Mobin Javed, and Vern Paxson. 2018. A Bestiary of Blocking: The Motivations and Modes behind Website Unavailability. In *Free and Open Communications on the Internet*. USENIX, Baltimore, MD, USA, 9. <https://www.usenix.org/system/files/conference/foci18/foci18-paper-tschantz.pdf>
- [291] Todd Underwood. 2005. Internet-Wide Catastrophe - Last Year. https://web.archive.org/web/20080228131639/http://www.renesys.com/blog/2005/12/internetwide_nearcatastrophela.shtml
- [292] US Department of Justice. 2022. Press Release: United States Leads Seizure of One of the World's Largest Hacker Forums and Arrests Administrator. <https://www.justice.gov/opa/pr/united-states-leads-seizure-one-world-s-largest-hacker-forums-and-arrests-administrator>
- [293] Brian Van Leeuwen, Jason Gao, Haikuo Yin, Benjamin Anthony, and Vincent Urias. 2019. *Networked-Based Cyber Analysis Using Deep Packet Inspection (DPI) for High-Speed Networks*. Technical Report SAND2019-13774, 1863848, 705224. Sandia National Lab, U.S. Department of Energy. SAND2019-13774, 1863848, 705224 pages. <https://doi.org/10.2172/1863848>
- [294] Roland van Rijswijk-Deij, Anna Sperotto, and Aiko Pras. 2015. Making the Case for Elliptic Curves in DNSSEC. *ACM SIGCOMM Computer Communication Review* 45, 5 (Sept. 2015), 13–19. <https://doi.org/10.1145/2831347.2831350>
- [295] Benjamin VanderSloot, Allison McDonald, Will Scott, J Alex Halderman, and Roya Ensafi. 2018. Quack: Scalable Remote Measurement of Application-Layer Censorship. In *USENIX Security Symposium*. USENIX, Baltimore, MD, USA, 16.

- [296] Petr Velan, Milan Čermák, Pavel Čeleda, and Martin Drašar. 2015. A Survey of Methods for Encrypted Traffic Classification and Analysis. *International Journal of Network Management* 25, 5 (Sept. 2015), 355–374. <https://doi.org/10.1002/nem.1901>
- [297] Pieter Velghe. 2019. “Reading China”: The Internet of Things, Surveillance, and Social Management in the PRC. *China Perspectives* 2019, 1 (March 2019), 85–89. <https://doi.org/10.4000/chinaperspectives.8874>
- [298] John-Paul Verkamp and Minaxi Gupta. 2012. Inferring Mechanics of Web Censorship Around the World. In *Free and Open Communications on the Internet*. USENIX, Bellevue, WA, 7. <https://www.usenix.org/conference/foci12/workshop-program/presentation/Verkamp>
- [299] Vasilis Ververis, Marios Isaakidis, Chrystalleni Loizidou, and Benjamin Fabian. 2017. Internet Censorship Capabilities in Cyprus: An Investigation of Online Gambling Blocklisting. In *E-Democracy – Privacy-Preserving, Secure, Intelligent E-Government Services*, Sokratis K. Katsikas and Vasilios Zorkadis (Eds.). Vol. 792. Springer International Publishing, Cham, 136–149. https://doi.org/10.1007/978-3-319-71117-1_10
- [300] Vasilis Ververis, George Kargiotakis, Arturo Filastò, Benjamin Fabian, and Afentoulis Alexandros. 2015. Understanding Internet Censorship Policy: The Case of Greece. In *Free and Open Communications on the Internet*. USENIX, Washington, D.C., 11.
- [301] Chenxu Wang, Jiangyi Yin, Zhao Li, Hongbo Xu, Zhongyi Zhang, and Qingyun Liu. 2024. Identifying VPN Servers through Graph-Represented Behaviors. In *Proceedings of the ACM Web Conference 2024*. ACM, Singapore Singapore, 10. <https://doi.org/10.1145/3589334.3645552>
- [302] Gaukas Wang, Jackson Sippe, Anonymous, Hai Chi, and Eric Wustrow. 2023. Chasing Shadows: A Security Analysis of the ShadowTLS Proxy. In *Free and Open Communications on the Internet*. Proceedings on Privacy Enhancing Technologies, Lausanne, Switzerland, 7. <https://petsymposium.org/foci/2023/foci-2023-0002.pdf>
- [303] Liang Wang, Kevin P. Dyer, Aditya Akella, Thomas Ristenpart, and Thomas Shrimpton. 2015. Seeing through Network-Protocol Obfuscation. In *Computer and Communications Security*. ACM, Denver, CO, USA, 13. <http://pages.cs.wisc.edu/~liangw/pub/ccsf653-wangA.pdf>
- [304] Qiyan Wang, Xun Gong, Giang T.K. Nguyen, Amir Houmansadr, and Nikita Borisov. 2012. CensorSpoofer: Asymmetric Communication Using IP Spoofing for Censorship-Resistant Web Browsing. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security - CCS '12*. ACM Press, Raleigh, North Carolina, USA, 121. <https://doi.org/10.1145/2382196.2382212>
- [305] Stephanie Wang and Shishir Nagaraja. 2007. *A Technical Review of the Internet Shutdown in Burma*. Technical Report. ONI, OpenNet Initiative. 14 pages. https://web.archive.org/web/20210701131103/https://opennet.net/sites/opennet.net/files/ONI_Bulletin_Burma_2007.pdf
- [306] Zhongjie Wang, Yue Cao, Zhiyun Qian, Chengyu Song, and Srikanth V. Krishnamurthy. 2017. Your State Is Not Mine: A Closer Look at Evading Stateful Internet Censorship. In *Internet Measurement Conference*. ACM, London United Kingdom, 14. <https://doi.org/10.1145/3131365.3131374>
- [307] Zhongjie Wang, Shitong Zhu, Yue Cao, Zhiyun Qian, Chengyu Song, Srikanth V. Krishnamurthy, Kevin S. Chan, and Tracy D. Braun. 2020. SymTCP: Eluding Stateful Deep Packet Inspection with Automated Discrepancy Discovery. In *Proceedings 2020 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA, 17. <https://doi.org/10.14722/ndss.2020.24083>
- [308] Human Rights Watch. 2006. “Race to the Bottom” Corporate Complicity in Chinese Internet Censorship. Technical Report. Human Rights Watch, Technical Report. 149 pages. <https://www.hrw.org/reports/2006/china0806/china0806webwcover.pdf>
- [309] Murdoch Watney. 2014. Determining When Conduct in Cyberspace Constitutes Cyber Warfare in Terms of the International Law and Tallinn Manual on the International Law Applicable to Cyber Warfare: A Synopsis. In *International Conference on Digital Forensics and Cyber Crime*, Vol. 132. Institute for Computer Sciences, Johannesburg, South Africa, 130–143. https://doi.org/10.1007/978-3-319-14289-0_10
- [310] Nicholas Weaver, Robin Sommer, and Vern Paxson. 2009. Detecting Forged TCP Reset Packets. In *NDSS*. NDSS, San Diego, CA, USA, 14. <https://www.ndss-symposium.org/wp-content/uploads/2017/09/weav.pdf>
- [311] Zachary Weinberg, Diogo Barradas, and Nicolas Christin. 2021. Chinese Wall or Swiss Cheese? Keyword Filtering in the Great Firewall of China. In *Proceedings of the Web Conference 2021*. ACM, Ljubljana Slovenia, 472–483. <https://doi.org/10.1145/3442381.3450076>
- [312] Zachary Weinberg, Mahmood Sharif, Janos Szurdi, and Nicolas Christin. 2017. Topics of Controversy: An Empirical Analysis of Web Censorship Lists. *Privacy Enhancing Technologies* 2017, 1 (2017), 42–61. <https://petsymposium.org/2017/papers/issue1/paper06-2017-1-source.pdf>
- [313] Zachary Weinberg, Jeffrey Wang, Vinod Yegneswaran, Linda Briesemeister, Steven Cheung, Frank Wang, and Dan Boneh. 2012. StegoTorus: A Camouflage Proxy for the Tor Anonymity System. In *Proceedings of the 2012 ACM Conference on Computer and Communications Security - CCS '12*. ACM Press, Raleigh, North Carolina, USA, 109. <https://doi.org/10.1145/2382196.2382211>
- [314] Steffen Wendzel, Simon Volpert, Sebastian Zillien, Julia Lenz, Philip Runz, and Luca Caviglione. 2025. A Survey of Internet Censorship and Its Measurement: Methodology, Trends, and Challenges. <https://arxiv.org/abs/2502.14945>
- [315] Brandon Wiley. 2011. *Dust: A Blocking-Resistant Internet Transport Protocol*. Technical Report. School of Information, University of Texas at Austin, University of Texas, Austin. 12 pages. <https://www.freehaven.net/anonbib/cache/wileydust.pdf>
- [316] Philipp Winter. 2014. *Measuring and Circumventing Internet Censorship*. Ph.D. Dissertation. Karlstad University. <https://www.diva-portal.org/smash/get/diva2:758124/FULLTEXT01.pdf>
- [317] Philipp Winter. 2023. NullHypothesis / Censorbib. <https://github.com/NullHypothesis/censorbib>
- [318] Philipp Winter and Stefan Lindskog. 2012. How the Great Firewall of China Is Blocking Tor. In *Free and Open Communications on the Internet*. USENIX, Bellevue, WA, 7. <https://www.usenix.org/system/files/conference/foci12/foci12-final2.pdf>
- [319] Philipp Winter, Tobias Pulls, and Juergen Fuss. 2013. ScrambleSuit: A Polymorphic Network Protocol to Circumvent Censorship. In *Proceedings of the 12th ACM Workshop on Privacy in the Electronic Society*. ACM, Berlin Germany, 213–224. <https://doi.org/10.1145/2517840.2517856>
- [320] Sebastian Wolfgarten. 2006. *Investigating Large-Scale Internet Content Filtering*. Technical Report. Dublin City University, Dublin, Ireland. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.133.5778&rep=rep1&type=pdf>
- [321] Bill Woodcock. 2011. Overview of the Egyptian Internet Shutdown. <https://web.archive.org/web/20120125051846/http://www.pch.net/resources/misc/Egypt-PCH-Overview.pdf>
- [322] Samuel Woodhams and Simon Migliano. 2022. Cost of Internet Shutdowns Report 2021. <https://www.top10vpn.com/research/cost-of-internet-shutdowns/2021/>

- [323] Joss Wright. 2012. *Regional Variation in Chinese Internet Filtering*. Technical Report. University of Oxford. https://papers.ssrn.com/sol3/Delivery.cfm/SSRN_ID2265775_code1448244.pdf?abstractid=2265775&mirid=3
- [324] Robin B. Wright. 2011. *Rock the Casbah: Rage and Rebellion across the Islamic World* (1st ed.). Simon & Schuster, New York.
- [325] Mingshi Wu, Jackson Sippe, Danesh Sivakumar, Jack Burg, Peter Anderson, Xiaokang Wang, Kevin Bock, Amir Houmansadr, Dave Levin, and Eric Wustrow. 2023. How the Great Firewall of China Detects and Blocks Fully Encrypted Traffic. In *USENIX Security Symposium*. USENIX, Anaheim, CA, USA, 12. <https://www.usenix.org/system/files/sec23fall-prepub-234-wu-mingshi.pdf>
- [326] Mingshi Wu, Ali Zohaib, Zakir Durumeric, Amir Houmansadr, and Eric Wustrow. 2025. A Wall Behind A Wall: Emerging Regional Censorship in China. In *Symposium on Security & Privacy*. IEEE, San Francisco, USA, 18. <https://gfw.report/publications/sp25/data/paper/paper.pdf>
- [327] Tim Wu. 2003. Network Neutrality, Broadband Discrimination. *Journal on telecommunications & high technology law* 2, 1 (2003), 36.
- [328] Yinglian Xie, Fang Yu, Kannan Achan, Eliot Gillum, Moises Goldszmidt, and Ted Wobber. 2007. How Dynamic Are IP Addresses?. In *Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*. ACM, Kyoto Japan, 301–312. <https://doi.org/10.1145/1282380.1282415>
- [329] Ruohan Xiong and Jeffrey Knockel. 2019. An Efficient Method to Determine Which Combination of Keywords Triggered Automatic Filtering of a Message. In *FOCI Workshop*. USENIX, Santa Clara, CA, USA, 9. https://www.usenix.org/system/files/foci19-paper_xiong.pdf
- [330] Xueyang Xu, Z. Morley Mao, and J. Halderman. 2011. Internet Censorship in China: Where Does the Filtering Occur? In *Passive and Active Measurement*, Neil Spring and George F. Riley (Eds.). Vol. 6579. Springer, Berlin, Heidelberg, 133–142. https://doi.org/10.1007/978-3-642-19260-9_14
- [331] Diwen Xue, Anna Ablove, Reethika Ramesh, Grace Danciu, and Roya Ensafi. 2024. Bridging Barriers: A Survey of Challenges and Priorities in the Censorship Circumvention Landscape. In *USENIX Security Symposium*. USENIX, Philadelphia, PA, 19. <https://www.usenix.org/system/files/usenixsecurity24-xue-bridging.pdf>
- [332] Diwen Xue, Benjamin Mixon-Baca, Anna Ablove, Beau Kujath, Jedidiah Crandall, and Roya Ensafi. 2022. TSPU: Russia’s Decentralized Censorship System. In *Proceedings of the 22nd ACM Internet Measurement Conference*. ACM, Nice France, 179–194. <https://doi.org/10.1145/3517745.3561461>
- [333] Diwen Xue, Reethika Ramesh, Arham Jain, Michalis Kallitsis, J. Alex Halderman, Jedidiah R. Crandall, and Roya Ensafi. 2022. OpenVPN Is Open to VPN Fingerprinting. In *USENIX Security Symposium*. USENIX, Boston, MA, USA, 19. <https://www.usenix.org/system/files/sec22-xue-diwen.pdf>
- [334] Diwen Xue, Reethika Ramesh, Valdik S S, Leonid Evdokimov, Andrey Viktorov, Arham Jain, Eric Wustrow, Simone Basso, and Roya Ensafi. 2021. Throttling Twitter: An Emerging Censorship Technique in Russia. In *Proceedings of the 21st ACM Internet Measurement Conference*. ACM, Virtual Event, 435–443. <https://doi.org/10.1145/3487552.3487858>
- [335] Diwen Xue, Robert Stanley, Piyush Kumar, and Roya Ensafi. 2025. The Discriminative Power of Cross-layer RTTs in Fingerprinting Proxy Traffic. In *Network and Distributed System Security*. The Internet Society, San Diego, USA, 18. <https://www.ndss-symposium.org/wp-content/uploads/2025-966-paper.pdf>
- [336] Maria Xynou and Arturo Filastò. 2021. How Countries Attempt to Block Signal Private Messenger App around the World. <https://ooni.org/post/2021-how-signal-private-messenger-blocked-around-the-world/>
- [337] Tarun Kumar Yadav, Akshat Sinha, Devashish Gosain, Piyush Sharma, and Sambuddho Chakravarty. 2018. Where The Light Gets In: Analyzing Web Censorship Mechanisms in India. In *Internet Measurement Conference*. ACM, Boston, USA, 13. <https://doi.org/10.1145/3278532.3278555>
- [338] Stephanie Yang. 2021. China Appears to Block Popular Encrypted Messaging App Signal. <https://www.wsj.com/articles/china-appears-to-block-signal-one-of-last-popular-encrypted-messaging-apps-11615883434>
- [339] Omer Yoachimik. 2022. Cloudflare Mitigates 26 Million Request per Second DDoS Attack. <https://blog.cloudflare.com/26m-rps-ddos/>
- [340] Savvas Zannettou, Tristan Caulfield, Emiliano De Cristofaro, Michael Sirivianos, Gianluca Stringhini, and Jeremy Blackburn. 2019. Disinformation Warfare: Understanding State-Sponsored Trolls on Twitter and Their Influence on the Web. In *Companion Proceedings of The 2019 World Wide Web Conference*. ACM, San Francisco USA, 218–226. <https://doi.org/10.1145/3308560.3316495>
- [341] Saman Taghavi Zargar, James Joshi, and David Tipper. 2013. A Survey of Defense Mechanisms Against Distributed Denial of Service (DDoS) Flooding Attacks. *IEEE Communications Surveys & Tutorials* 15, 4 (2013), 2046–2069. <https://doi.org/10.1109/SURV.2013.031413.00127>
- [342] Kim Zetter. 2016. The Sony Hackers Were Causing Mayhem Years Before They Hit the Company. <https://www.wired.com/2016/02/sony-hackers-causing-mayhem-years-hit-company/>
- [343] Pengxiong Zhu, Keyu Man, Zhongjie Wang, Zhiyun Qian, Roya Ensafi, J. Halderman, and Haixin Duan. 2020. Characterizing Transnational Internet Performance and the Great Bottleneck of China. *Measurement and Analysis of Computing Systems* 4 (2020), 23. <https://doi.org/10.1145/3379479>
- [344] Tao Zhu, David Phipps, Adam Pridgen, Jedidiah R. Crandall, and Dan S. Wallach. 2013. The Velocity of Censorship: High-Fidelity Detection of Microblog Post Deletions. In *USENIX Security Symposium*. USENIX, Washington, D.C., 14. <https://www.cs.unm.edu/~crandall/usenix13.pdf>
- [345] Jonathan Zittrain and Benjamin Edelman. 2003. Internet Filtering in China. *IEEE Internet Computing* 7 (2003), 10. <https://doi.org/10.1109/MIC.2003.1189191>
- [346] Ali Zohaib, Qiang Zao, Jackson Sippe, Abdulrahman Alaraj, Amir Houmansadr, Zakir Durumeric, and Eric Wustrow. 2025. Exposing and Circumventing SNI-based QUIC Censorship of the Great Firewall of China. In *USENIX Security Symposium*. USENIX, Seattle, USA, 20. <https://gfw.report/publications/usenixsecurity25/data/paper/quic-sni.pdf>