Making Sense of Internet Censorship: A New Frontier for Internet Measurement

Sam Burnett
Georgia Tech
sburnett@cc.gatech.edu

Nick Feamster
Georgia Tech
feamster@cc.gatech.edu

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.
The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

Free and open access to information on the Internet is at risk: more than 60 countries around the world practice some form of Internet censorship, and both the number of countries practicing censorship and the proportion of Internet users who are subject to it are likely to increase. We posit that, although it may not always be feasible to guarantee free and open access to information, citizens have the right to know when their access has been obstructed, restricted, or tampered with, so that they can make informed decisions on information access. We motivate the need for a system that provides accurate, verifiable reports of censorship and discuss the challenges involved in designing such a system. We place these challenges in context by studying their applicability to OONI, a new censorship measurement platform.

Categories and Subject Descriptors

C.2.3 [**Computer-Communication Networks**]: Network Operations—*network monitoring,public networks*

General Terms

Measurements

Keywords

Censorship measurement,OONI

1. INTRODUCTION

Internet censorship is becoming increasingly pervasive: The Open Network Initiative reports that nearly 60 countries around the world restrict Internet communications in some way, and it is likely that far more entities manipulate content or communications in some fashion. The number of countries that interfere with Internet communications is also likely to increase as more countries with historically repressive governments become better connected. Organizations or countries may restrict access to information using techniques that range from blocking communications entirely to degrading performance (sometimes to the point where a service might be unusable) to manipulating content to spread misinformation. Alternatively, an organization might not block or throttle a service, but instead use it to spread false information or otherwise influence public opinion.

Censors are continually developing new techniques to block or manipulate communications, while citizens and activists continue to devise mechanisms for circumventing these technologies. Organizations that wish to censor content or otherwise manipulate information have a variety of advanced techniques and technologies at their disposal, and new circumvention techniques are continually met with attempts to thwart them.

Internet users should know when their Internet access has been obstructed, manipulated, or tampered with; they need an independent, third-party service that helps them determine whether their Internet service provider is restricting access to content or specific protocols, or otherwise degrading Internet service. Despite our attempts to design a variety of censorship circumvention tools, we have very little understanding of the techniques that censors use to restrict or otherwise manipulate access to online information. Citizens need a tool that continually (1) monitors the extent of censorship and manipulation in countries around the world; and (2) evaluates the effectiveness of various technologies that attempt to circumvent censorship in real-world settings.

Monitoring censorship (and the effectiveness of censorship circumvention tools) presents several nuanced challenges beyond simply monitoring Internet performance or availability, due to the inherently adversarial environment in which the measurements are performed. The measurement problem introduces several new challenges. The first set of challenges involve determining *what* to measure in the first place:

- *Internet censorship is not well-defined, and is shifting towards manipulation and propaganda.* Defining censorship clearly is rather nuanced in practice, since forms of "soft censorship" ranging from partial obstruction to distortion of search results to intentional performance degradation to content manipulation, are also emerging.
- *There is no ground truth against which to validate measurement tools and techniques.* Unlike conventional Internet measurement tools, which can be validated against a notion of ground truth, it is difficult to determine the accuracy of censorship monitoring without knowing what the uncensored and untampered content is in the first place.

Even if we understood *what* to measure, determining *how* to measure censorship also presents a unique set of challenges, since a censor could attempt to disrupt or confuse the measurements. First, a censor might block or tamper with measurements from a monitoring tool, or disrupt communication to prevent the tool from reporting measurements. Second,

running a tool that monitors censorship practices could be incriminating. Finally, users may be hesitant to provide reports about sites they visit. Users of such a tool may also wish to determine the cause of unreachability (*e.g.*, whether the unreachability is due to keyword filtering, blocking of a domain name, etc.), and attributing the unreachability to a particular cause may introduce additional challenges.

This paper explores these challenges in more detail and advocates a new direction for Internet measurement research: wide-area monitoring of Internet censorship and control. In Section 2, we explain why monitoring Internet censorship entails an additional set of concerns from conventional Internet measurement. Section 3 outlines several design principles that may help censorship monitoring systems surmount these additional challenges and Section 4 studies how these principles apply to OONI, a nascent censorship measurement framework. Section 5 surveys related work, and Section 6 presents a research agenda for this emerging area.

2. CHALLENGES IN MONITORING

This section details some of the new challenges for designing a censorship monitoring system. We believe these challenges highlight an important set of concerns that are distinct from conventional Internet measurement.

2.1 Overview and Context

A *censor* might inhibit communication in several ways; we consider the following attacks, in increasing levels of sophistication:

- *Blocking (complete or partial).* Completely blocking access to a destination (or blocking the use of a protocol) is certainly the most widely practiced form of censorship. Existing tools (*e.g.*, Herdict) essentially report complete blocking. A censor may also practice "partial" blocking, which may be significantly more difficult to detect—for example, a Web page might be reachable, but a censor might still block certain objects referenced within that page.
- *Performance degradation.* Although a censor might not block a service or site outright, we have personally experienced and measured cases where a censor (or an ISP that is acting on behalf of a censor) degrades performance (*e.g.*, by introducing packet loss) to such an extent as to make a service unusable.
- *Content manipulation.* A censor might not block content, but instead manipulate information as it propagates among different Web sites or between users. One example of this type of "soft censorship" might be to remove search results from a user's set of results. Another example might be to propagate versions of a news story that omit or modify key portions of the story. A third example of manipulation might be to use "sock puppet" techniques to create the false appearance of independent opinions in an attempt to persuade or influence public opinion.

2.2 Challenge: What to Measure

We now enumerate and describe various challenges associated with censorship monitoring systems that are distinct from conventional Internet measurement.

Censorship is inherently ill-defined One of the most difficult aspects of measuring Internet censorship is defining its metrics. Because censorship is rooted in the agendas and actions of political actors, it is difficult to formalize the threats or express them quantitatively. Additionally, unlike many conventional network measurement problems, censorship has no single axis or metric along which it might be monitored or measured; as we described in the previous section, censorship can take many forms, and technical and political threats are continually evolving. Finally, there is little to no "ground truth" information available about the existence or extent of censorship practices in any given environment, so determining the effectiveness of various censorship monitoring tools can be difficult. Rather than constructing tools for measuring censorship in general, we narrowly scope specific types of censorship and monitor censorship behavior in the context of each more-limited domain.

Content manipulation may be difficult to disambiguate from personalization or regionalization Ultimately, we aim to monitor not only cases of outright blocking, but also instances where an organization has manipulated content to suit its interests. We believe that this type of "soft censorship" may be the next wave of Internet control. Governments that are savvy about information control may ultimately not block services such as Twitter, Facebook, or popular news sites, but rather manipulate the content from these services to better advance their own agendas and serve their own interests. Similarly, recent cases of "filter bubbles" may also emerge, where a certain user or group of users sees different information from another group. Whether the source of a filter bubble is personalization or the result of content manipulation by a government may be difficult to disambiguate, but, for a monitoring system, this distinction may not matter: ultimately, the system could report differences in content retrieved from various locations (and as various users) and allow the users to draw their own conclusions about the cause of these differences.

Preserving user privacy is difficult when detecting more nuanced forms of censorship Some forms of soft censorship might involve intentional performance degradation or content manipulation. Detecting this type of behavior would require comparing performance or content across groups of users, but such a comparison also implies that each user must reveal their browsing history. This strong requirement and concern for user privacy also constitutes a significant difference from conventional network analysis, which typically takes place within a single organization.

Accurate detection may require a lot of data Although nearly all network measurements encounter this problem, it holds heightened importance in an adversarial environment. Several factors influence the amount of data required for accurate censorship detection. Detecting more granular censorship requires more data; for example, targeted content manipulation of a few unpopular Web pages requires more data then wholesale country-wide blocking. Faulty or malicious monitors, spurious access link failures, and momentary congestion all add noise, which increases the amount of data required. Finally, more sophisticated censorship schemes require more data per measurement; for example, detecting content manipulation on a URL might require recording entire documents, but detecting blocking of the same document would only need a single boolean value.

2.3 Challenge: How to Measure

Measurements must contend with an adversary Unlike conventional Internet measurement tools, those that measure Internet censorship may be subject to censorship themselves, particularly if the protocols or monitoring authority are easily identifiable. To remain robust to blocking, such a monitoring system might thus rely on a system that provides a robust and *covert* channel for circumventing censorship, such as Infranet [9] or Collage [2]. Both of these tools require both communicating parties to exchange secret keys that form the basis of the communication channel; these keys could be distributed with the software itself. (We discuss challenges of software distribution below.) Other circumvention tools that focus on anonymity but not deniability (*e.g.*, Tor without Pluggable Transports [8]) may not be suitable for this purpose, as the presence of encrypted communication might be incriminating and could also be blocked (as a case in point, Iran blocked all SSL communications earlier this year.)

Monitoring agents may themselves be adversarial. Because each agent performs independent measurements from a unique vantage point, it may be difficult to verify the accuracy of these reports. To surmount this challenge, the system must correlate data from multiple independent sources, as we discuss in Section 3.

Even if measurements themselves are accurate, users' access to reports may be blocked. Users should have a robust and deniable means for accessing reports about the censorship of various Internet destinations and services. If the system requires users to query information about censors' attempts to disrupt communication, then such queries from the user should be both robust (*i.e.*, it should be difficult for the censor to block these requests) and deniable (*i.e.*, it should be difficult for a censor to determine that a user is inquiring about censorship). Similar concerns apply if this information is instead passed to users via a continuous feed, although the sheer volume of such a feed might be prohibitive or arouse suspicion. In light of these requirements, systems that provide deniability to the participants who exchange messages, such as Collage or Infranet, may be appropriate for servicing these types of queries or periodically disseminating updates to the user about sites that have been blocked.

Users must acquire and install monitoring agents Because accurate, robust reporting of reachability information entails reachability measurements from many independent vantage points, users must be able to easily acquire and install monitoring agent software, and they must be willing to do so without fear of incrimination. Dual use, which would design the monitoring software for purposes other than censorship monitoring, might mitigate this challenge. One possibility for dual use in this case is to have users deploy monitoring agents that perform reachability tests for network management or troubleshooting purposes. Censorship is just one of many possible causes for unreachability that such a tool might observe (*i.e.*, if the tool helped pinpoint other sources of problems, such as congestion, DNS failure, or network configuration problems), and such a tool might never even explicitly mention that a site was "blocked". A user might draw his or her own conclusions about whether censorship is a likely cause for the observed reachability problems or other artifacts such as performance degradation or differences in content.

Operating the measurement tool or performing certain measurements may be incriminating Gathering information about censorship requires making observations from user vantage points, either by means of software installed directly on client machines or perhaps from a network router or switch positioned at the network edge (*e.g.*, a BISMark router [1]). In either case, performing these measurements might serve to incriminate a user, particularly if the software performs active measurements of reachability from a "target list" that includes destinations that a user might not ordinarily visit. Concealing these measurements presents a challenge, because any attempt to conceal the measurements may alter the network traffic in a way that affects the censor's decision about whether that traffic should be blocked, disrupted, or tampered with.

3. MEASUREMENT PRINCIPLES

We introduce and motivate several principles that can help guarantee robust, accurate monitoring of censorship.

PRINCIPLE 1. *Correlate independent data sources.*

Accurate, robust monitoring of censorship requires verifiable reports about the availability of various destinations and services. In a typical setting where the identity of a sender is trusted, the information that an agent reports could be verified by checking the message's signature. Unfortunately, these monitoring agents may be run by a variety of untrusted sources (possibly even by the censors themselves); this situation creates the possibility whereby any individual monitoring agent might produce inaccurate or intentionally false reports. To account for this possibility, we advocate a monitoring system that incorporates information from a variety of independent sources, rather than using only one type of data to infer censorship. For example, reachability information might come from the following sources:

- *Web server access logs* at various destinations can provide information about which clients are able to access the site, and from where (*i.e.*, geography, ISP, etc.). A sudden decrease in accesses from any particular region of IP address space, ISP, or country might suggest malfeasance.
- *Client-side logging about user downloads* (*e.g.*, from the Alex toolbar) can observe the reachability of various destinations from clients, and, when measurements fail, the likely cause of those failures (*e.g.*, TCP connection reset, DNS lookup failure).
- *Home routers* are always on and can directly measure connectivity issues, since they can measure reachability directly from an access link, independently of the host or local network configuration.
- *DNS lookups* visible from recursive resolvers (*e.g.*, at the access ISP) or at the top-level domain (*e.g.*, at an authority such as Verisign) can provide information about whether DNS requests are being successfully issued from hosts or downstream recursive resolvers. Similarly to monitoring Web access logs, monitoring DNS lookup behavior can expose anomalies, such as a

sudden drop in lookups to a particular domain, or a sudden drop in lookups from a particular geographic region or ISP.

PRINCIPLE 2. *Separate measurements from analysis.*

This principle suggests that measurements be conducted separately from the analysis that determines the likely cause of problems. The primary reason for separating these two components is to allow the analysis to incorporate measurements from a variety of sources, including sources that are outside the reach of the censor. This separation may also help preserve the *dual use* property: monitoring agents who perform measurements about availability can plausibly claim they are performing the measurements for the purposes of network troubleshooting or management. Measuring the availability of services and destinations is an innocuous activity that can be attributed to everyday network management or performance monitoring, thus making the process of performing measurements less suspicious. For example, a user might run a monitoring agent on a home network gateway or as a browser extension to provide information about high packet loss rates or latencies; that monitoring agent would likely not be able to attribute *cause* to the poor performance, but rather would report those measurements to a coordinating authority, which might then perform correlations to better attribute the cause of the performance degradation.

PRINCIPLE 3. *Separate information producers from information consumers.*

Measuring the reachability of various Internet destinations and services can be separated from *receiving* aggregated information or statistics about the reachability of different sites. Users that measure reachability of various destinations need not be the same set of users who access that information. Separating the channels that measure reachability from those that access that information makes it more difficult for a censor to disrupt these channels. It also facilitates gathering information from a variety of sources (*e.g.*, home routers, browser toolbars, user reports, and other measurement services), since the process of gathering information is decoupled from the process of reporting it to users. Finally, this separation may also improve deniability for users.

PRINCIPLE 4. *Employ "dual-use" scenarios wherever possible, to improve robustness.*

A tool that exclusively monitors censorship may be more difficult to deploy than one which serves another purpose, such as helping network operators perform network troubleshooting. As an example, a tool might run periodic throughput tests to different destinations, or execute downloads to popular Web sites, to monitor download or page render times and alert a user or network operator when these times exceed a certain threshold; certain network monitoring platforms (*e.g.*, SamKnows) already perform regular download tests to top Web sites. Such tests naturally serve to monitor access network performance, but anomalies in reachability to these sites could also suggest malfeasance. If censorship monitoring is a secondary purpose (or, better yet, a side effect) of the monitoring agent's functions, users may be more willing to deploy it and censors may find it more difficult to justify blocking the tool or its reports. Therefore, we should, to the extent possible, collect monitoring data from deployed performance monitoring and troubleshooting tools.

PRINCIPLE 5. *Adopt existing techniques to provide robust channels for measuring, reporting, and accessing reports.*

A censor may still attempt to inhibit a censorship monitoring system, in spite of the fact that the tool serves other useful, innocuous purposes. As such, such a system may need to use existing circumvention tools for communication, particularly when disseminating information about availability problems and their possible causes to users. Many of these channels (*e.g.*, Collage) might be extremely low bandwidth, so the system could periodically push information about sites that have been blocked or tampered with to users rather than having the users request that information on a site-by-site basis. On the other hand, designing a communication channel to send monitoring reports or analysis results through a censorship firewall may be easier than circumventing censorship in the general case, since we can make more assumptions about the format, size, or timeliness of the data.

PRINCIPLE 6. *Heed and adapt to continually changing political and technical threats.*

Unlike conventional Internet monitoring, the practice of monitoring Internet censorship is a continually moving target, since governments and organizations continually devise new tactics for disrupting communication. Hence, we believe a "plug and play" censorship monitoring system will never exist: any monitoring system will quickly be rendered ineffective by rapidly evolving attempts to block or disrupt it. Maintaining an effective censorship monitoring system will require expertise in a variety of domains, including, but not limited to, public policy.

4. CASE STUDY: OONI

The Tor Project recently proposed the Open Observatory of Network Interference (OONI), a framework for observing blocking, performance degradation, content manipulation, and other forms of network interference [10, 15]. Figure 1 illustrates OONI's conceptual architecture. Inspired by unit testing frameworks, developers use OONI to write *interference tests* that run on many independent *probes* deployed at the network edge. These tests measure information about a variety of *services*. Examples of probes are Web browser extensions, home routers, and standalone programs running on clients. Monitoring can be active, in which case it is triggered by a "target list" of destinations provided by a central authority, other user reports, or even the operator of the probe itself; active tests can use OONI Test Helpers, which are services operated specifically to assist OONI measurements. Measurements can also be passive, where probes collect information about users' normal interactions with services. The OONI test author decides the type and extent of information collected based on the kinds of interference he wants to detect. For example, detecting blocking requires reachability information, detecting performance degradation requires finer grained performance statistics, and detecting content manipulation requires document contents.

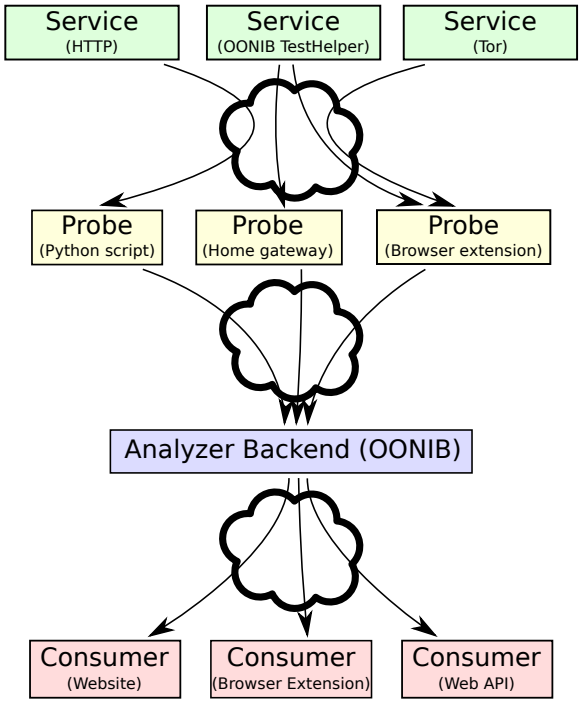

Figure 1: OONI's conceptual design and possible attack points. Both measurements and measurement results must contend with threats from adversaries.

We briefly study OONI's adherence to the principles we introduced in Section 3:

Principle 1 (Correlate independent data sources): OONI developers write tests that collect data from any number of sources, and run on any number of end hosts. Although OONI tests are envisioned to run at the network edge, they could theoretically operate on centralized data like server logs.

Principle 2 (Separate measurements from analysis): OONI Probes send measurements to an OONI backend, which stores all measurements in a public data repository (*e.g.*, M-Lab [14]). Third-party researchers can download, analyze, and interpret these raw measurements.

Principle 3 (Separate producers from consumers): Third party researchers consume OONI measurements from a central repository that is decoupled from data collection.

Principle 4 (Employ "dual-use" scenarios): OONI is a framework for writing network interference tests but specifies nothing about how those tests are presented to users. OONI test developers could employ their own dual-use scenarios to drive adoption of their tests.

Principle 5 (Use existing security techniques): OONI Probes send their measurements to the OONI backend using TLS over Tor. This not only assures confidentiality and integrity of test measurements, but also shields the IP address of the Probe; test developers must still be careful not to compromise the Probe identities in the measurement data itself. If Tor is inaccessible users may use a covert channel to access the Tor network (*e.g.*, StegoTorus [19]). At the backend, measurements are stored in a secure data center and can be accessed with TLS or through Tor or another censorship circumvention tool.

Principle 6 (Adapt to evolving threats): OONI places the onus of adapting to new threats on test developers by giving them significant flexibility in test authorship. Ultimately, humans must monitor and respond to the global technical and political landscape.

5. RELATED WORK

Along with OONI, The Tor Project hosts the Tor Censorship Detector, which monitors censorship of the Tor network itself and often observes nation-wide censorship as a side effect [5]. Data from the detector could corroborate OONI's measurements of the Tor network. OONI grew from a desire to expand on earlier studies of censorship [11, 6, 7].

Several projects attempt to measure the extent of various forms of Internet censorship. Perhaps the most well-known project is Google's Transparency Report [12], which shows access patterns to various Google services (*e.g.*, search, mail, news) from around the world. Unfortunately, the raw information regarding reachability to these services is not publicly available, and the information is limited to Google services. The Report also lacks an automated mechanism for inferring outages from the data it reports. Should any information from the Transparency Report become publicly available, it would still need to be combined with information from other independent data sources, for the reasons discussed in Section 3.

The Berkman Center's Herdict project gathers reports from users about the reachability of various sites using manual reports from users [13]. The manual nature of these user reports and the lack of independent sources of reachability information requires the operators of the Herdict service to manually verify each report; hence, the monitoring effort cannot be particularly large-scale.

Several efforts focus specifically on China. Crandall *et al.* performed Internet measurements to study the keyword filtering implemented by the Great Firewall of China, and found that effective censorship does not require the blocking of every illicit word, but must simply block enough to promote self-censorship [4]; they also proposed an architecture for maintaining a "weather report" about the keywords that are filtered over time. Clayton *et al.* studied the Chinese firewall's implementation of filtering based on TCP connection resets using measurements to a single Web server for each border AS [3]; such a tool could be a useful input to our proposed system if applied across a wider variety of destinations and clients.

Several projects study censorship around the world from PlanetLab vantage points. CensMon probes for DNS, IP, and HTTP filtering from PlanetLab and presents compelling preliminary evidence of censorship in China [16]. Verkamp and Gupta observe censorship outside China through a combination of PlanetLab nodes and case-by-case edge measurements [18]. PlanetLab measurements are convenient for researchers, but its vantage points are hosted on university networks that may not be affected by censorship in the same ways as home users.

Truthy analyzes and visualizes how information spreads across Twitter and tries to identify bursts of activity around memes [17]; its meme detection algorithms could serve as a

useful starting point for helping track the emergence of sock puppets or astroturfing campaigns.

Although our proposed monitoring system focuses on monitoring, not circumvention, a robust, accurate monitoring system could help determine the effectiveness of existing circumvention tools [8, 9, 2, 20].

6. SUMMARY AND RESEARCH AGENDA

The challenges and opportunities we outlined in this position paper make the case for a new class of Internet measurement, focused on measuring Internet availability in the face of adversaries. Perhaps one of the most significant differences of censorship monitoring from conventional Internet measurement is that the technical aspects of performing the measurements are not the most difficult challenge. Specifically, determining that a particular site has been blocked—or even what technology a censor is using to execute censorship—is doable; determining which part of a site's content has been tampered with is hard but still doable given sufficient measurement data; what is more difficult is determining *why* a particular site is being blocked, and, specifically, what content (*e.g.*, keyword, topic, domain, author) triggered the censor to block it. The research community should achieve this level of understanding through deeper inference, and by incorporating information from a wide variety of sources, including users, ISPs, Web site operators, companies, universities, and advertisers.

An immediate next step in the research agenda for censorship monitoring is to explore different information sources and determine the type of information they might provide for helping users infer availability problems and their causes. For example, in this paper, we have roughly outlined how server logs, client request results, home routers, DNS lookups, and wide-area measurements might each contribute to a better understanding of censorship practices, but the lion's share of the work will be to analyze each of these sources (and possibly other data sources) more thoroughly to understand what information they can provide independently, how reliable that information might be, and how these information sources might be coalesced to produce more robust, reliable reports.

Ultimately, as governments and organizations become increasingly savvy, they may *use* various communication media as tools for influence and persuasion, rather than simply blocking them. For example, "sock puppet" and "astroturfing" behavior are well-known, but there has been little attempt to monitor or quantify these effects on a large scale, and today a user has no way of knowing whether the content they receive is part of a sock-puppet or astroturfing campaign. Additionally, citizens may be subject to cases of content or information manipulation whereby content (*e.g.*, a news story) is altered as it is translated, paraphrased, or copied to a different Web site. As censorship and information control become more sophisticated, techniques for monitoring censorship will also need to evolve to expose new forms of "soft censorship".

Acknowledgments

This work was funded in part by NSF award CNS-1111723 and a Google Focused Research Award. We thank Hans Klein for valuable feedback on this work.

7. REFERENCES

[1] BISMark: Broadband Internet Service Benchmark. `http://projectbismark.net/`.

[2] S. Burnett, N. Feamster, and S. Vempala. Chipping away at censorship firewalls with user-generated content. In *Proc. 19th USENIX Security Symposium*, Washington, DC, Aug. 2010.

[3] R. Clayton, S. Murdoch, and R. Watson. Ignoring the Great Firewall of China. In *Privacy Enhancing Technologies*, pages 20–35. Springer, 2006.

[4] J. Crandall, D. Zinn, M. Byrd, E. Barr, and R. East. ConceptDoppler: A Weather Tracker for Internet Censorship. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, Arlington, VA, Oct. 2007.

[5] G. Danezis. An anomaly-based censorship detection system for Tor. `https://metrics.torproject.org/papers/detector-2011-09-09.pdf`, Sept. 2011.

[6] R. Deibert. *Access denied: The practice and policy of global internet filtering*. MIT Press, 2008.

[7] R. Deibert. *Access controlled: The shaping of power, rights, and rule in cyberspace*. The MIT Press, 2010.

[8] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The second-generation onion router. In *Proc. 13th USENIX Security Symposium*, San Diego, CA, Aug. 2004.

[9] N. Feamster, M. Balazinska, G. Harfst, H. Balakrishnan, and D. Karger. Infranet: Circumventing Web censorship and surveillance. In *Proc. 11th USENIX Security Symposium*, San Francisco, CA, Aug. 2002.

[10] A. Filastò and J. Appelbaum. Ooni: Open observatory of network interference. In *USENIX FOCI*, Aug. 2012.

[11] P. Gill. Characterizing global web censorship: Why is it so hard? Presented at the Workshop on Active Internet Measurements, Feb. 2013.

[12] Google Transparency Report. `http://www.google.com/transparencyreport/`.

[13] HerdictWeb: The Verdict of the Herd. `http://herdict.org`.

[14] Measurement Lab. `http://measurementlab.net`, Jan. 2009.

[15] Open observatory of network interference. `https://ooni.torproject.org`.

[16] A. Sfakianakis, E. Athanasopoulos, and S. Ioannidis. Censmon: A web censorship monitor. In *USENIX FOCI*, Aug. 2011.

[17] Truthy. `http://truthy.indiana.edu`.

[18] J.-P. Verkamp and M. Gupta. Inferring mechanics of web censorship around the world. In *Proceedings of the 2nd USENIX Workshop on Free and Open Communications on the Internet*, Aug 2013.

[19] Z. Weinberg, J. Wang, V. Yegneswaran, L. Briesemeister, S. Cheung, F. Wang, and D. Boneh. Stegotorus: a camouflage proxy for the tor anonymity system. In *Proceedings of the 2012 ACM conference on Computer and communications security*, pages 109–120. ACM, 2012.

[20] E. Wustrow, S. Wolchok, I. Goldberg, and J. A. Halderman. Telex: Anticensorship in the Network Infrastructure. In *Proc. 20th USENIX Security Symposium*, San Francisco, CA, Aug. 2011.