Understanding the Practices of Global Censorship through Accurate, End-to-End Measurements

Lin Jin
University of Delaware
Newark, Delaware, USA
linjin@udel.edu

Shuai Hao
Old Dominion University
Norfolk, Virginia, USA
shao@odu.edu

Haining Wang
Virginia Tech
Arlington, Virginia, USA
hnw@vt.edu

Chase Cotton
University of Delaware
Newark, Delaware, USA
ccotton@udel.edu

ABSTRACT

It is challenging to conduct a large scale Internet censorship measurement, as it involves triggering censors through artificial requests and identifying abnormalities from corresponding responses. Due to the lack of ground truth on the expected responses from legitimate services, previous studies typically require a heavy, unscalable manual inspection to identify false positives while still leaving false negatives undetected. In this paper, we propose Disguiser, a novel framework that enables end-to-end measurement to accurately detect the censorship activities and reveal the censor deployment without manual efforts. The core of Disguiser is a control server that replies with a static payload to provide the ground truth of server responses. As such, we send requests from various types of vantage points across the world to our control server, and the censorship activities can be recognized if a vantage point receives a different response. In particular, we design and conduct a cache test to pre-exclude the vantage points that could be interfered by cache proxies along the network path. Then we perform application traceroute towards our control server to explore censors' behaviors and their deployment. With Disguiser, we conduct 58 million measurements from vantage points in 177 countries. We observe 292 thousand censorship activities that block DNS, HTTP, or HTTPS requests inside 122 countries, achieving a 10^{-6} false positive rate and zero false negative rate. Furthermore, Disguiser reveals the censor deployment in 13 countries.

ACM Reference Format:
Lin Jin, Shuai Hao, Haining Wang, and Chase Cotton. 2022. Understanding the Practices of Global Censorship through Accurate, End-to-End Measurements. In *Abstract Proceedings of the 2022 ACM SIGMETRICS/IFIP PERFORMANCE Joint International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS/PERFORMANCE '22 Abstracts), June 6–10, 2022, Mumbai, India.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3489048.3522640

1 INTRODUCTION

Internet censorship controls what can be viewed by a certain group of Internet users. Such information control, typically placed by authority entities such as governments, ISPs, or organizations, can be successfully achieved by various techniques ranging from IP-layer censorship (*e.g.*, blocking IP addresses) to application-layer censorship (*e.g.*, DNS manipulation). Internet censorship has been widely witnessed and its severity varies from country to country.

To detect censorship activities in a region, the basic idea is to send a request from a vantage point within the region and then compare the response with a valid response from a legitimate server. However, the dilemma here is that if the request is blocked, the vantage point has no ground truth to identify the valid response. To tackle this issue, existing studies [1–4] collect valid responses from nodes deployed in multiple countries. However, this approach inevitably reduces the detection reliability due to the diversity and flexibility of Internet services. For example, clients at diverse locations may obtain different but valid IP addresses for the same domain. Thus, manual inspection is usually needed, causing the analysis to be unscalable and inefficient. More importantly, manual analysis can only identify false positives (*i.e.*, misclassified censorship) but false negatives (*i.e.*, undetected censorship) remain uncountable due to the lack of ground truth on what should have been received as mentioned above. As a result, the accuracy and reliability of the detection are still questionable after manual analysis.

One recent technique, Quack [5], addresses such a dilemma with servers running the Echo service that reflects back any bytes sent to it. Thus, each request sent to an echo server inside the censored region is reflected and then the outgoing traffic will encounter the censor. In the meantime, the request itself would also be the expected response if no censorship presents. However, the requests sent to and received from the echo servers are not on standard HTTP/HTTPS ports, and such requests cannot trigger a censor if it only examines requests on standard ports. We observe that many censors in 32 countries, including those enforce severe censorship policies such as Saudi Arabia and UAE, only block the requests sent to standard HTTP/HTTPS ports. To this end, it is imperative to explore an accurate and efficient methodology for understanding the censorship practices on a global scale.

In this paper, we propose *Disguiser*, a novel framework that accurately detects censorship activities and explores the deployment of censors. Disguiser introduces a control server as the destination of all probing requests to provide the ground truth of server responses. In particular, we send requests from vantage points, located in the tested regions, to our control server placed outside the tested region, which will return *static* responses we crafted. As such, by comparing the response obtained from the vantage point with the static response, we can accurately identify the censorship activities. Furthermore, we investigate the censor deployment using application traceroute, by which each vantage point makes a three-way handshake with our control server and then sends the requests with incremented Time-to-Live (TTL) values to identify censor behaviors and determine the location of censors.

SIGMETRICS/PERFORMANCE '22 Abstracts, June 6–10, 2022, Mumbai, India
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9141-2/22/06.
https://doi.org/10.1145/3489048.3522640

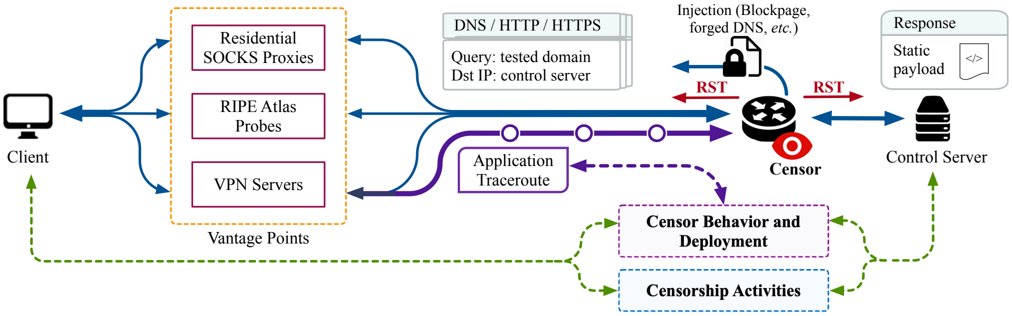

Figure 1: Architecture of Disguiser Framework.

With Disguiser, we conduct comprehensive and large-scale measurements on censorship with three fundamental protocols: DNS, HTTP, and HTTPS. During our study, we conduct 58 million measurements from 177 countries and observe 292 thousand censorship activities from vantage points in 122 countries. Disguiser achieves a 10^{-6} false positive rate and zero false negative rate in detecting censorship activities. Furthermore, we perform application traceroute and reveal censor deployment in 13 countries.

2 SYSTEM DESIGN

Disguiser is an end-to-end measurement framework for *accurately* investigating the practices of global Internet censorship that is based on either DNS or HTTP/HTTPS. Figure 1 illustrates the system design of Disguiser. The objective of Disguiser is to detect censorship activities and explore the censor deployment, while effectively eliminating false negatives and minimizing false positives without manual inspection. The high-level idea is that a client instructs the vantage points to (1) craft DNS/HTTP/HTTPS requests with the test domain names embedded, (2) send the packets to our control server to trigger censorship, and (3) collect the response back for later analysis. On the other side, our control server replies to arbitrary requests with a static payload for each type of protocol. To identify the location of censors and examine their deployment, application traceroute is performed in which the packets with increased TTL values are repeatedly sent for encountering the censors. Importantly, to eliminate the noise data, we carefully design tests to exclude the vantage points which could be potentially affected by the proxies/middleboxes placed in the network path.

Disguiser conducts censorship measurements through three types of vantage points (SOCKS Proxies, RIPE Atlas, and VPN). In particular, SOCKS Proxies conduct TCP-based DNS tests and HTTP/HTTPS tests, while RIPE Atlas conducts UDP-based DNS test. These two types of vantage points are used to study the global censorship activities. We then leverage VPN servers to conduct application traceroute to detect the censor deployment.

3 EXPERIMENTS AND RESULTS

We conducted a global scale measurement study with Disguiser over two six-week periods, one from April 2020 to May 2020, and another from June 2021 to Aug 2021. In total, we conduct 58 million measurements from vantage points in 177 countries, and identify 292 thousand censorship activities in 122 countries. We achieve a 10^{-6} false positive rate and zero false negative rate.

Overall, we find that HTTP-based blocking is the most prevalent censorship behavior. For the UDP-based censorship, we find that Iran and China are ranked top-2 in the percentage of blocked domains, and China's censorship policy is most consistent in blocking domains with UDP- and TCP-based queries. For the HTTP/HTTP censorship, we find that Iran leads the percentage of the blocked domains for the two protocols, and Saudi Arabia, UAE, and China all block more than 15% of the tested domains. Also, the top three categories of the censored domains are almost the same for the two protocols, and the percentages are close as well. By comparing the censorship activities in two years, we observe a significant increase in the number of censored domains in China in 2021, and Russia adopts HTTPS interception in 2021.

Disguiser reveals the deployment of censors in 13 countries with application traceroute. We find that 10 out of 13 countries deploy in-path censors that would directly drop or modify the actual packets. Also, we identify that censor devices tend to be deployed close to the nation's border routers, and they tend to be deployed in the ASes with a relatively high AS rank, indicating that they can monitor a large number of Internet users in that country.

ACKNOWLEDGMENT

We would like to thank our shepherd, Zihui Ge, and the anonymous reviewers for their detailed and insightful comments, which help to improve the quality of this paper. This work was supported in part by the National Science Foundation (NSF) grant DGE-1821744, the Army Research Office (ARO) grant W911NF19-1-0049, and an Internet Freedom Fund from Open Technology Fund (OTF).

REFERENCES

[1] Arturo Filastò and Jacob Appelbaum. 2012. OONI: Open Observatory of Network Interference. In *USENIX Workshop on Free and Open Communications the Internet (FOCI)*.
[2] 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 *ACM Internet Measurement Conference (IMC)*.
[3] 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 *IEEE Symposium on Security and Privacy (S&P)*.
[4] Paul Pearce, Ben Jones, Frank Li, Roya Ensafi, Nick Feamster, Nick Weaver, and Vern Paxson. 2017. Global Measurement of DNS Manipulation. In *USENIX Security Symposium*.
[5] 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*.