

Accountability without accountability

A censorship measurement case study
@willscott

Bio

- Grad school @ University of Washington
- Postdoc @ University of Michigan
- NextGen Scholar @ CSIS



Network Measurement

ACM IMC 2018

Oct 31 - Nov 2, 2018

Boston, MA, USA

27TH USENIX
SECURITY SYMPOSIUM

AUGUST 15–17, 2018
BALTIMORE, MD, USA



SIGCOMM 2018
BUDAPEST

FOCI '18

8th USENIX Workshop on Free and Open Communications on the Internet

AUGUST 14, 2018
BALTIMORE, MD, USA



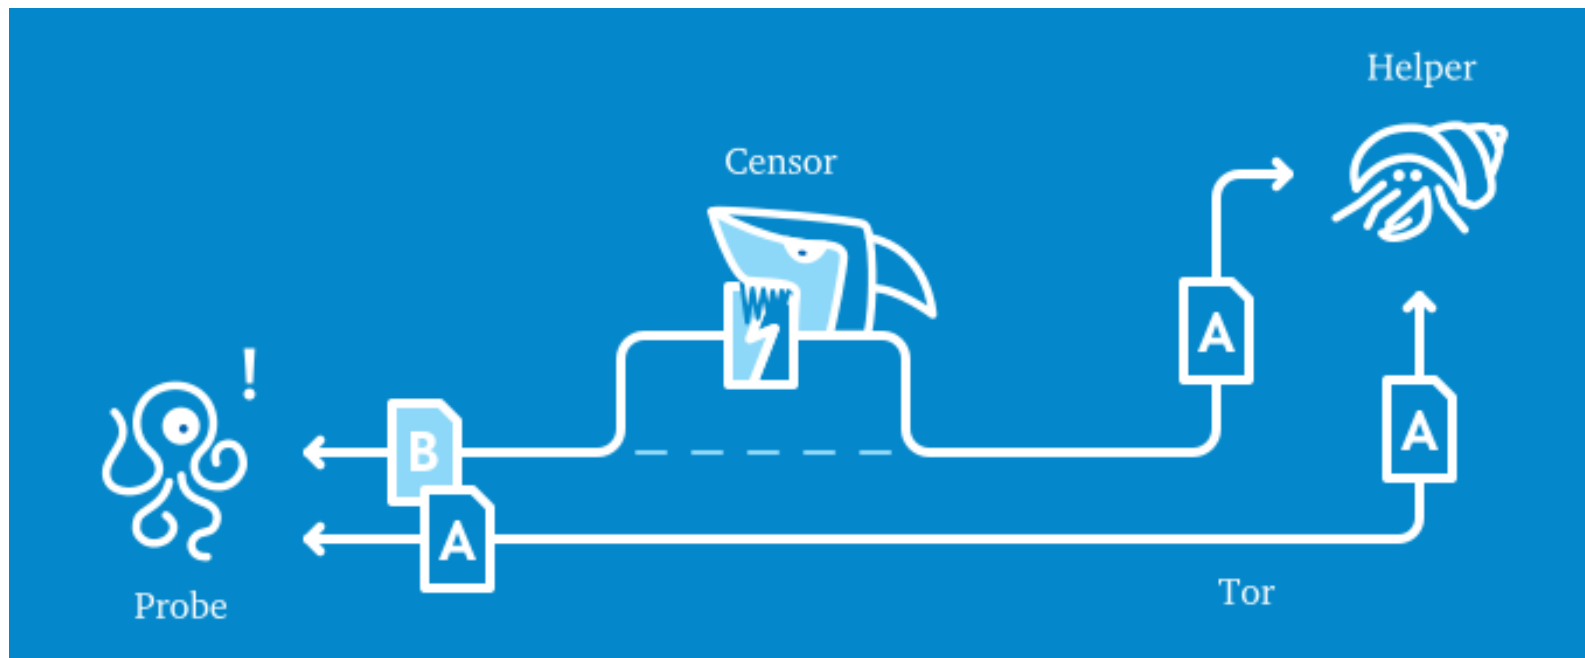
Casa Batlló Overview Barcelona Spain, by ChristianSchd [CC BY-SA 3.0], via Wikimedia Commons

PETS 2018

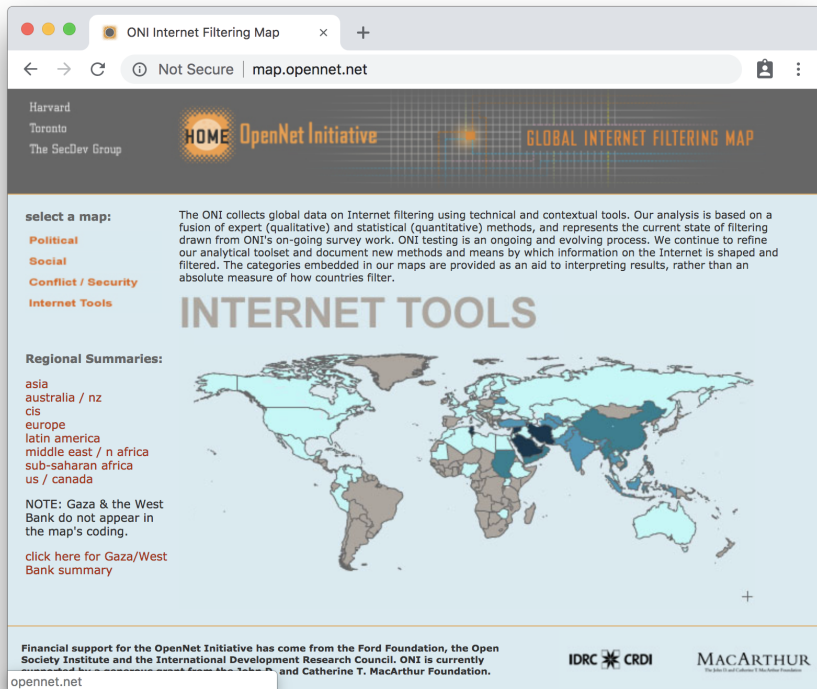
The 18th Privacy Enhancing Technologies Symposium
July 24 – 27, 2018
Barcelona, Spain



Measuring Censorship



Measuring Censorship



Platform	Flexibility	Coverage	Blocking resistance	Main use
PlanetLab [16]	High	Low/Medium	Medium	Network measurements
Atlas [18]	Low	Medium/High	Medium	Network measurements
M-Lab [6]	Low	High	Medium	Network measurements
Tor [5]	Medium	Medium	Low	Low-latency anonymity
OONI [10]	High	Low	Medium	Interference analysis
Herdict [11]	Low	Low/Medium	Low	Interference analysis
OpenNet [14]	Low	Medium	High	Interference analysis

Table 1: Comparison between several popular filtering analysis platforms.

AFTERTHOUGHT | By [Joseph Cox](#) | Apr 9 2018, 9:00am

Chinese Government Forces Residents To Install Surveillance App With Awful Security

Last year, authorities told residents of a Muslim-populated part of China to install JingWang, an app that scans for certain files. Now, researchers have found it transfers the collected data with no encryption.

Encore

Sigcomm 2015

Statement from the SIGCOMM 2015 Program Committee: The SIGCOMM 2015 PC appreciated the technical contributions made in this paper, but found the paper controversial because some of the experiments the authors conducted raise ethical concerns. The controversy arose in large part because the networking research community does not yet have widely accepted guidelines or rules for the ethics of experiments that measure online censorship. In accordance with the published submission guidelines for SIGCOMM 2015, had the authors not engaged with their Institutional Review Boards (IRBs) or had their IRBs determined that their research was unethical, the PC would have rejected the paper without review. But the authors did engage with their IRBs, which did not flag the research as unethical. The PC hopes that discussion of the ethical concerns these experiments raise will advance the development of ethical guidelines in this area. It is the PC's view that future guidelines should include as a core principle that researchers should not engage in experiments that subject users to an appreciable risk of substantial harm absent informed consent. The PC endorses neither the use of the experimental techniques this paper describes nor the experiments the authors conducted.

Encore: Lightweight Measurement of Web Censorship with Cross-Origin Requests

Sam Burnett
School of Computer Science, Georgia Tech
sam.burnett@gatech.edu

Nick Feamster
Department of Computer Science, Princeton
feamster@cs.princeton.edu

Abstract

Despite the pervasiveness of Internet censorship, we have scant data on its extent, mechanisms, and evolution. Measuring censorship is challenging: it requires continual measurement of reachability to many target sites from diverse vantage points. Amassing suitable vantage points for longitudinal measurement is difficult; existing systems have achieved only small, short-lived deployments. We observe, however, that most Internet users access content via Web browsers, and the very nature of Web site design allows browsers to make requests to domains with different origins than the main Web page. We present Encore, a system that harnesses cross-origin requests to measure Web filtering from a diverse set of vantage points without requiring users to install custom software, enabling longitudinal measurements from many vantage points. We explain how Encore induces Web clients to perform cross-origin requests that measure Web filtering, design a distributed platform for scheduling and collecting these measurements, show the feasibility of a global-scale deployment with a pilot study and an analysis of potentially censored Web content, identify several cases of filtering in six

months of measurements, and discuss ethical concerns that would arise with widespread deployment.

Categories and Subject Descriptors

• **Networks** → Network measurement; *Web protocol security* • **Social and professional topics** → **Technology and censorship**

Keywords

Web censorship; Network measurement; Web security

1 Introduction

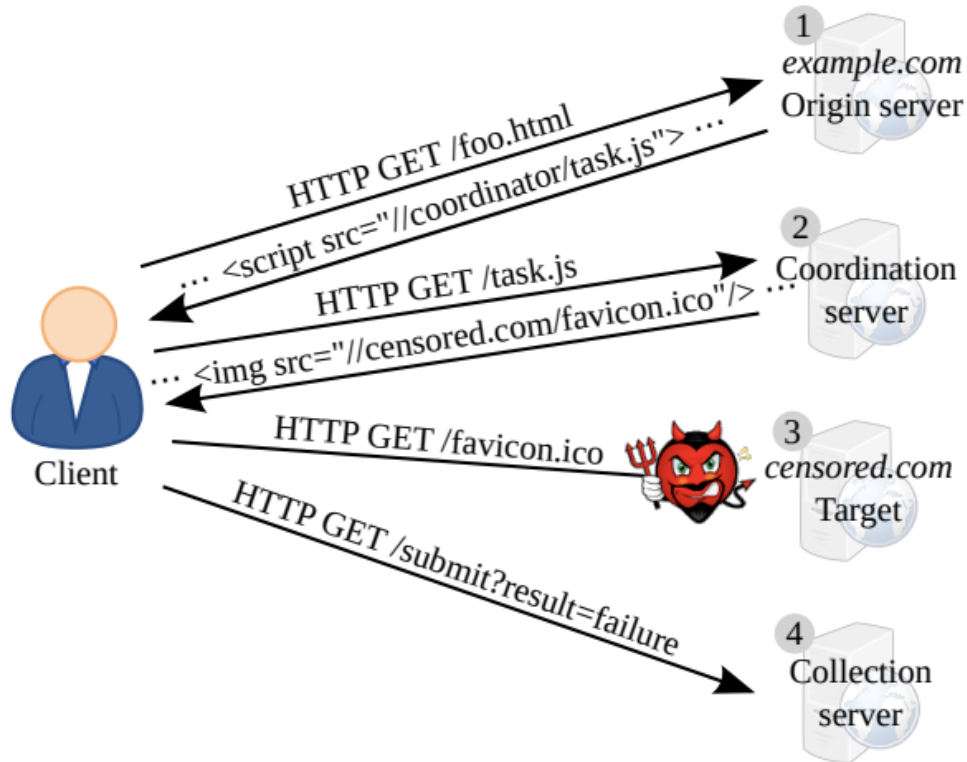
Internet censorship is pervasive: by some estimates, nearly 60 countries restrict Internet communication in some way [35]. As more citizens in countries with historically repressive governments gain Internet access, government controls are likely to increase. Collecting pervasive, longitudinal measurements that capture the evolving nature and extent of Internet censorship is more important than ever.

Researchers, activists, and citizens aim to understand what, where, when, and how governments and organizations implement Internet censorship. This knowledge can shed light on government censorship policies and guide the development of new circumvention techniques. Although drastic actions such as introducing country-wide outages (as has occurred in Libya, Syria, and Egypt) are eminently observable, the most common forms of Internet censorship are more subtle and challenging to measure. Censorship typically targets specific domains, URLs, keywords, or content; varies over time in response to changing social or political conditions (e.g., a national election); and can be indistinguishable from application errors or poor performance (e.g., high delay or packet

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGCOMM '15, August 17–21, 2015, London, United Kingdom
Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-3542-3/15/08 ...\$15.00
DOI: <http://dx.doi.org/10.1145/2785956.2787485>

Encore

Sigcomm 2015



Encore

Sigcomm 2015

Statement from the SIGCOMM 2015 Program Committee: The SIGCOMM 2015 PC appreciated the technical contributions made in this paper, but found the paper controversial because some of the experiments the authors conducted raise ethical concerns. The controversy arose in large part because the networking research community does not yet have widely accepted guidelines or rules for the ethics of experiments that measure online censorship. In accordance with the published submission guidelines for SIGCOMM 2015, had the authors not engaged with their Institutional Review Boards (IRBs) or had their IRBs determined that their research was unethical, the PC would have rejected the paper without review. But the authors did engage with their IRBs, which did not flag the research as unethical. The PC hopes that discussion of the ethical concerns these experiments raise will advance the development of ethical guidelines in this area. It is the PC's view that future guidelines should include as a core principle that researchers should not engage in experiments that subject users to an appreciable risk of substantial harm absent informed consent. The PC endorses neither the use of the experimental techniques this paper describes nor the experiments the authors conducted.

Encore: Lightweight Measurement of Web Censorship with Cross-Origin Requests

Sam Burnett
School of Computer Science, Georgia Tech
sam.burnett@gatech.edu

Nick Feamster
Department of Computer Science, Princeton
feamster@cs.princeton.edu

Abstract

Despite the pervasiveness of Internet censorship, we have scant data on its extent, mechanisms, and evolution. Measuring censorship is challenging: it requires continual measurement of reachability to many target sites from diverse vantage points. Amassing suitable vantage points for longitudinal measurement is difficult; existing systems have achieved only small, short-lived deployments. We observe, however, that most Internet users access content via Web browsers, and the very nature of Web site design allows browsers to

months of measurements, and discuss ethical concerns that would arise with widespread deployment.

Categories and Subject Descriptors

• **Networks** → **Network measurement**; *Web protocol security* • **Social and professional topics** → **Technology and censorship**

Keywords

Web censorship; Network measurement; Web security

Statement from the SIGCOMM 2015 Program Committee: The SIGCOMM 2015 PC appreciated the technical contributions made in this paper, but found the paper controversial because some of the experiments the authors conducted raise ethical concerns. The controversy arose in large part because the networking research community does not yet have widely accepted guidelines or rules for the ethics of experiments that measure online censorship. In accordance with the published submission guidelines for SIGCOMM 2015, had the authors not engaged with their Institutional Review Boards (IRBs) or had their IRBs determined that their research was unethical, the PC would have rejected the paper without review. But the authors did engage with their IRBs, which did not flag the research as unethical. The PC hopes that discussion of the ethical concerns these experiments raise will advance the development of ethical guidelines in this area. It is the PC's view that future guidelines should include as a core principle that researchers should not engage in experiments that subject users to an appreciable risk of substantial harm absent informed consent. The PC endorses neither the use of the experimental techniques this paper describes nor the experiments the authors conducted.

Institutional Review Board

Public Announcement

**WE WILL PAY YOU \$4.00 FOR
ONE HOUR OF YOUR TIME**

Persons Needed for a Study of Memory

*We will pay five hundred New Haven men to help us complete a scientific study of memory and learning. The study is being done at Yale University.

*Each person who participates will be paid \$4.00 (plus 50c carfare) for approximately 1 hour's time. We need you for only one hour: there are no further obligations. You may choose the time you would like to come (evenings, weekdays, or weekends).

*No special training, education, or experience is needed. We want:

Factory workers	Businessmen	Construction workers
City employees	Clerks	Salespeople
Laborers	Professional people	White-collar workers
Barbers	Telephone workers	Others

All persons must be between the ages of 20 and 50. High school and college students cannot be used.

*If you meet these qualifications, fill out the coupon below and mail it now to Professor Stanley Milgram, Department of Psychology, Yale University, New Haven. You will be notified later of the specific time and place of the study. We reserve the right to decline any application.

*You will be paid \$4.00 (plus 50c carfare) as soon as you arrive at the laboratory.

TO:
PROF. STANLEY MILGRAM, DEPARTMENT OF PSYCHOLOGY,
YALE UNIVERSITY, NEW HAVEN, CONN. I want to take part in
this study of memory and learning. I am between the ages of 20 and
50. I will be paid \$4.00 (plus 50c carfare) if I participate.

NAME (Please Print)

ADDRESS

TELEPHONE NO. Best time to call you

AGE OCCUPATION SEX

CAN YOU COME:

WEEKDAYS EVENINGS WEEKENDS

Encore

Sigcomm 2015

in the set of measurements we report on in this paper. The Institutional Review Boards (IRBs) at both Georgia Tech and Princeton declined to formally review Encore because it does not collect or analyze Personally Identifiable Information (PII) and is not human subjects research [9]. Yet, Encore is clearly capable of exposing its users to some level of risk. Because we do not understand the risks that a tool like Encore presents, we have focused most of our research efforts on developing the measurement technology, not on reporting results from the measurements that we gather. Other censorship measurement tools have and will continue to face similar ethical questions, and we believe that our role as researchers is to lead a responsible dialogue in the context of these emerging tools.

Menlo Report



The Menlo Report

Ethical Principles Guiding Information and
Communication Technology Research


August 2012



**Homeland
Security**

Science and Technology

Scanning



-
1. Coordinate closely with local network admins to reduce risks and handle inquiries.
 2. Verify that scans will not overwhelm the local network or upstream provider.
 3. Signal the benign nature of the scans in web pages and DNS entries of the source addresses.
 4. Clearly explain the purpose and scope of the scans in all communications.
 5. Provide a simple means of opting out, and honor requests promptly.
 6. Conduct scans no larger or more frequent than is necessary for research objectives.
 7. Spread scan traffic over time or source addresses when feasible.
-

Table 5: **Recommended Practices** — We offer these suggestions for other researchers conducting fast Internet-wide scans as guidelines for good Internet citizenship.

<https://zmap.io/paper.pdf>

DNS




Figure 1: DNS servers discovered in each Country. We find 169 countries hosting DNS resolvers in more than 20 class-c networks.

Satellite '16 (250,000 servers)

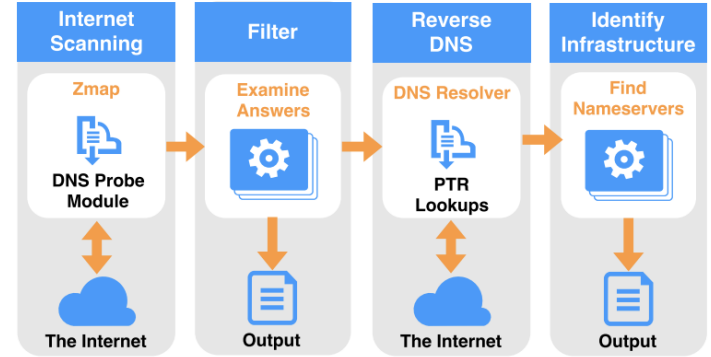


Figure 1: Overview of Iris's DNS resolver identification and selection pipeline. Iris begins with a global scan of the entire IPv4 address space, followed by reverse DNS PTR lookups for all open resolvers, and finally filtering resolvers to only include DNS infrastructure.

Iris '17 (8,000 servers)

Commercial Scanning

- Shodan.io



- Rapid7



- Censys



The Core Ethical Questions

What responsibility comes with our privilege?

How likely is the potential for harm users are exposed to?

How much value is gained by the increased transparency?



Ethics

The in-depth tale of Bylock, the Turkish messenger app whose 1x1 tracking GIF was the basis for tens of thousands of treason accusations




<https://arrestedlawyers.org/2018/01/21/ever-changing-evidence-bylock/>

Today

Given these factors, we believe that the risks of our work to echo server operators are extremely small. We considered seeking informed consent from them anyway, but we rejected this route for several reasons.¹ First, the risk to these users is low, but if we were to contact them to seek consent, this interaction with foreign censorship researchers would *in and of itself* carry a small risk of drawing negative attention from the authorities. Second, if we only used servers for which the operators granted consent, these operators would face a much higher risk of reprisal, since their participation would be easy to observe and would imply knowing complicity. Third, obtaining consent would be infeasible in most cases, due to the difficulty of identifying and contacting the server operators; if we limited our study to echo servers for which we could find owner contact information, this would lead to far fewer usable servers, thus severely reducing the benefit of the study. The communities that stand to benefit most

Echo (USENIX Security '18)

Remedies

- Blending in to the "background noise"
- Limiting Netflow records between end-users and sensitive content
- "Sub-prime" repackaging  **hola!**

Accountability without Accountability

@willscott