Quack:
Scalable Remote Measurement of Application-Layer Censorship

Benjamin VanderSloot, Allison McDonald,
Will Scott, J. Alex Halderman, and Roya Ensafi

UNIVERSITY OF MICHIGAN

Censorship

Policy of information control that harms citizens

Spreading beyond the large powers

Frequently opaque in topic and technique

Censorship Measurement

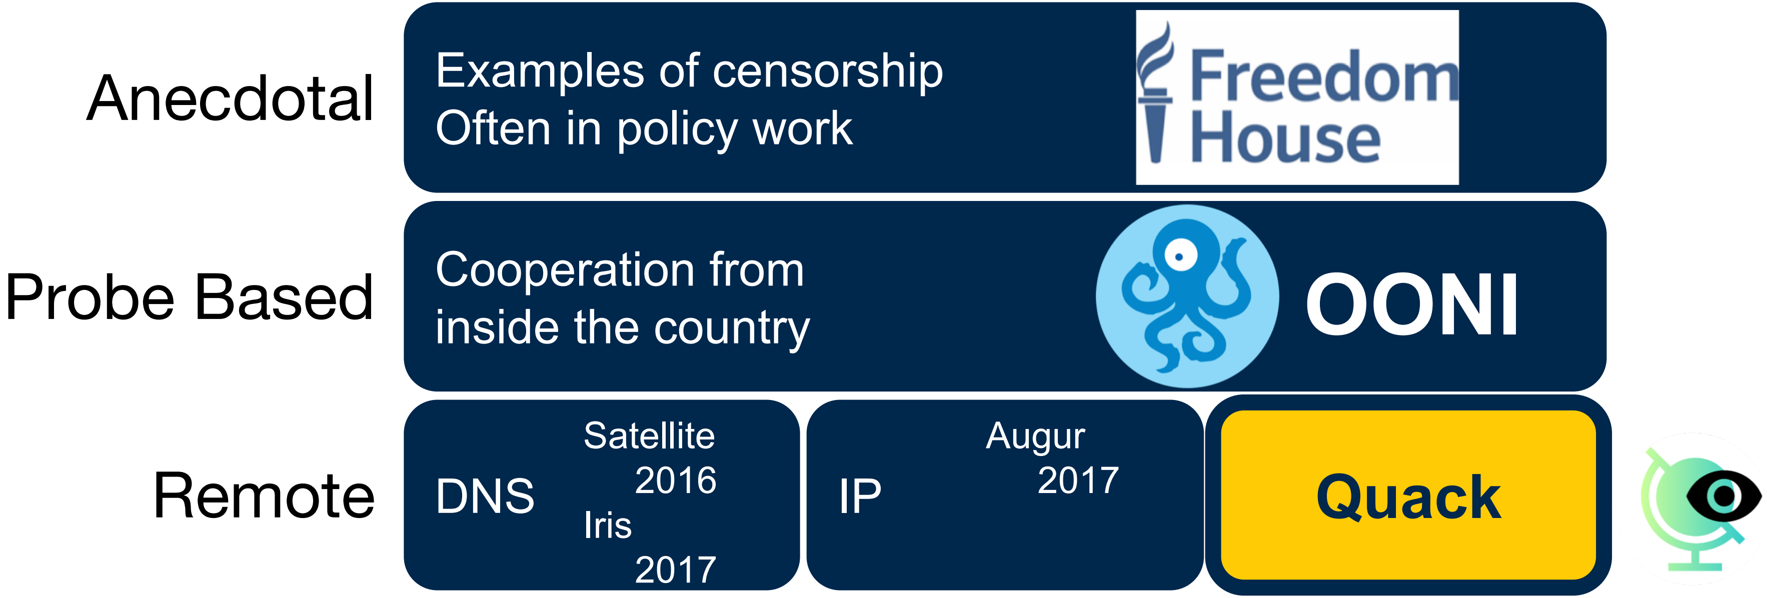

Anecdotal

Examples of censorship
Often in policy work

Freedom House

Probe Based

Cooperation from
inside the country

OONI

Remote

DNS

Satellite
2016
Iris
2017

IP

Augur
2017

Quack

Application-Layer Censorship

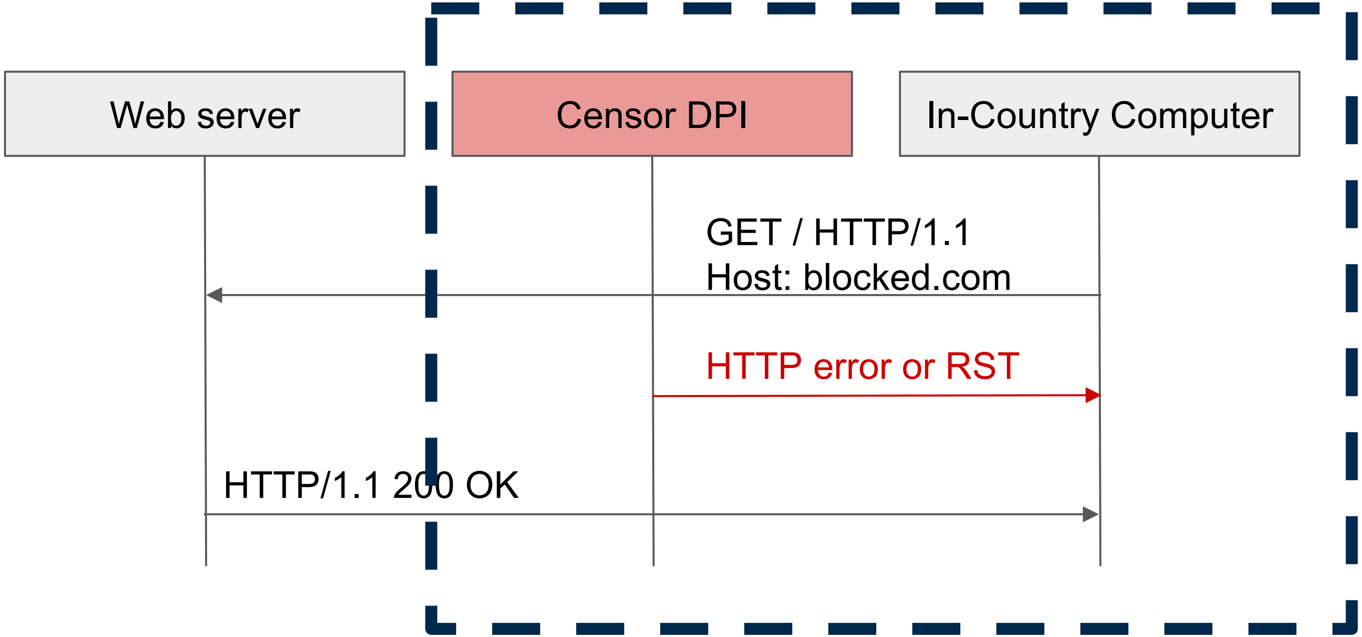

Web server

Censor DPI

In-Country Computer

GET / HTTP/1.1
Host: blocked.com

HTTP error or RST

HTTP/1.1 200 OK

Echo Behavior

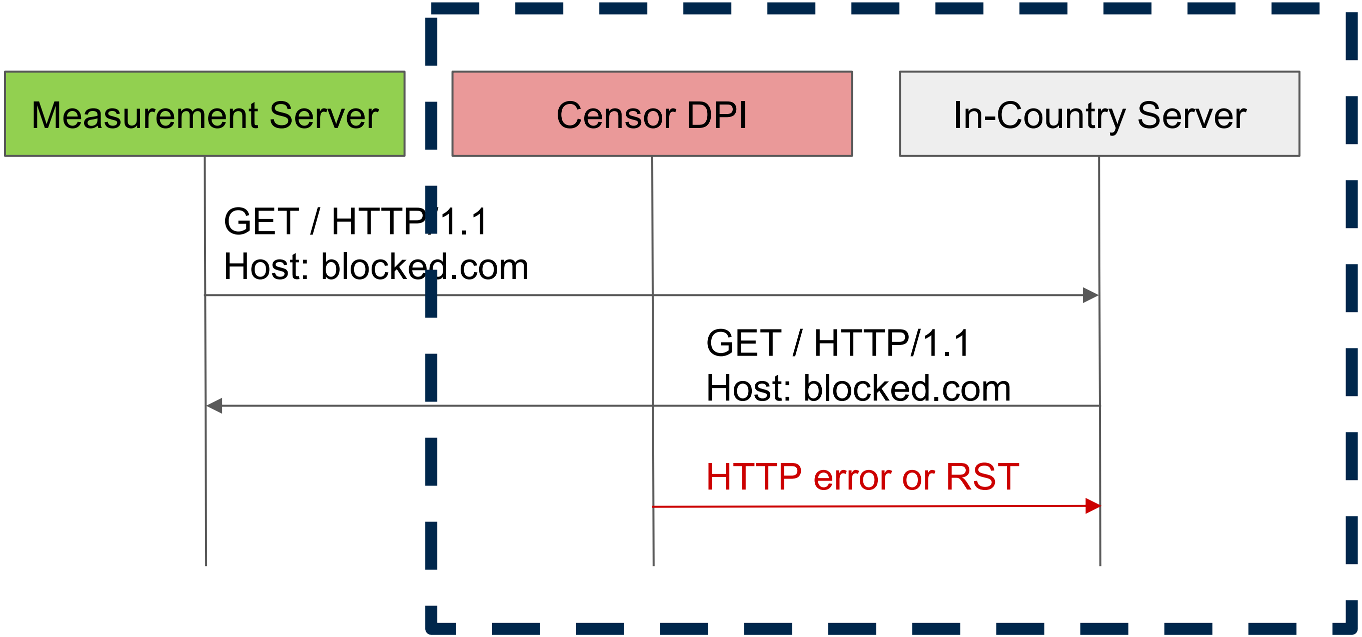

The connection is reset or data has been injected into the protocol

Protocol Selection

Protocols we know can provide echo behavior

TLS

Telnet

FTP

Echo

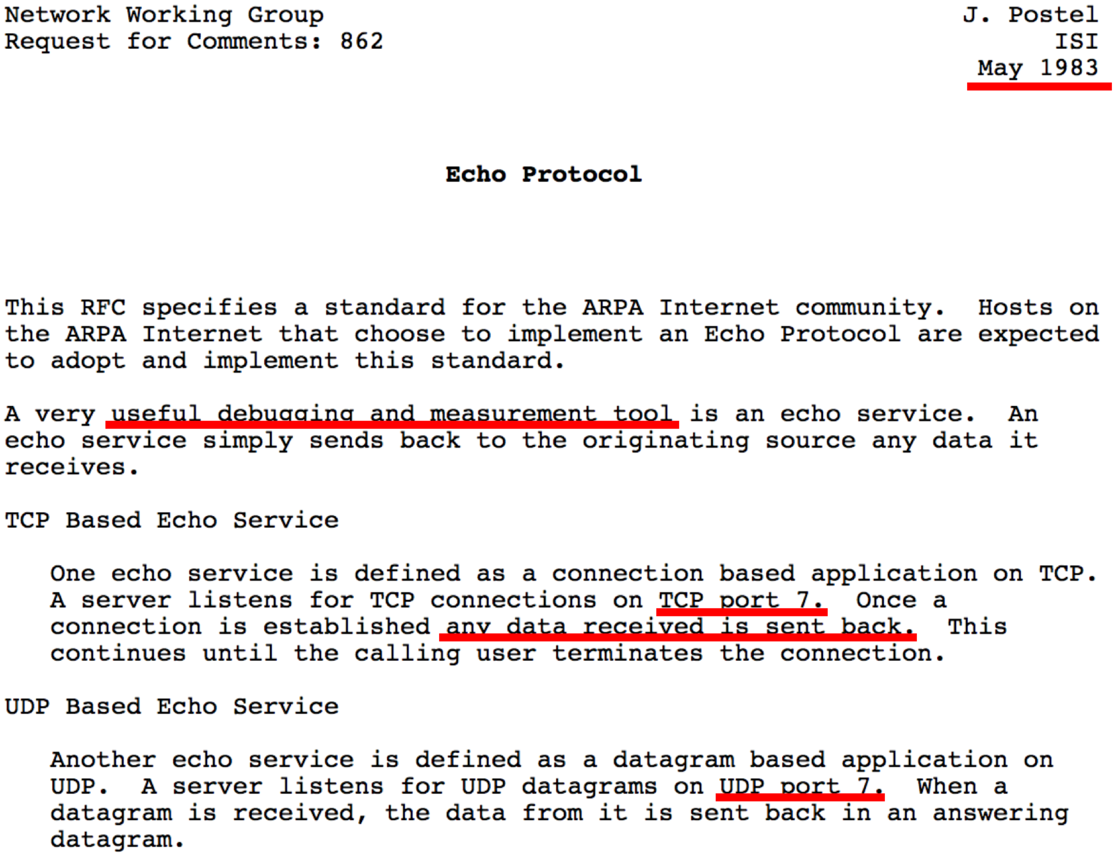

Network Working Group
Request for Comments: 862

J. Postel
ISI
May 1983

Echo Protocol

This RFC specifies a standard for the ARPA Internet community. Hosts on the ARPA Internet that choose to implement an Echo Protocol are expected to adopt and implement this standard.

A very useful debugging and measurement tool is an echo service. An echo service simply sends back to the originating source any data it receives.

TCP Based Echo Service

One echo service is defined as a connection based application on TCP. A server listens for TCP connections on TCP port 7. Once a connection is established any data received is sent back. This continues until the calling user terminates the connection.

UDP Based Echo Service

Another echo service is defined as a datagram based application on UDP. A server listens for UDP datagrams on UDP port 7. When a datagram is received, the data from it is sent back in an answering datagram.

1983

port 7

Design Goals

Detection of keywords are being blocked.

Minimizing risk for the **safety** of people in censored countries.

Robustness in the face of intermittent network failures.

Performing censorship measurement **at scale**.

System Design, Inside-Out

A Single Trial

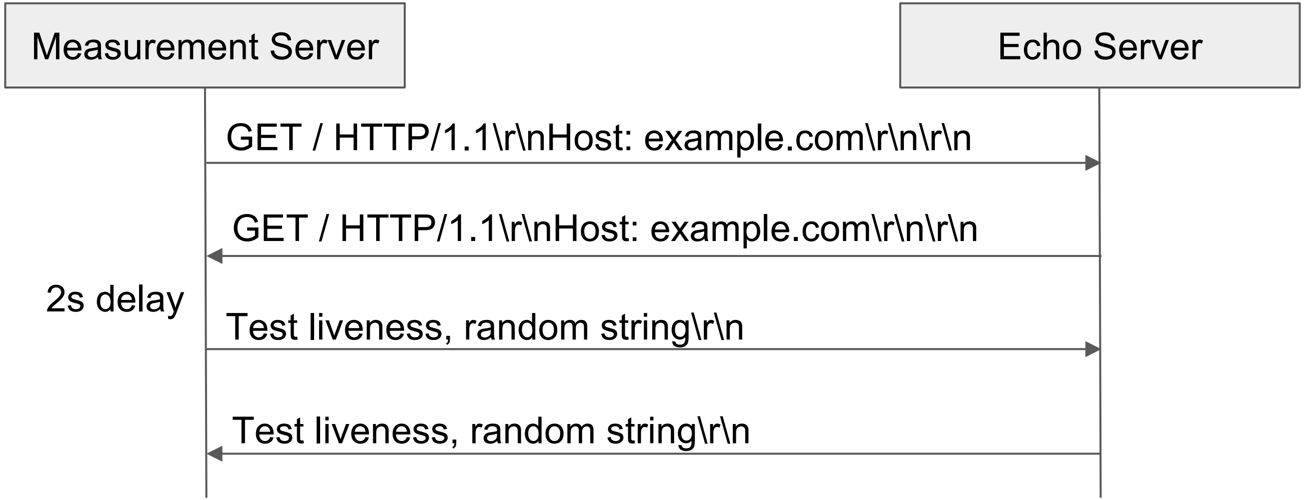
Measurement Server
Echo Server
GET / HTTP/1.1\r\nHost: example.com\r\n\r\n
GET / HTTP/1.1\r\nHost: example.com\r\n\r\n
2s delay
Test liveness, random string\r\n
Test liveness, random string\r\n

A Single Test

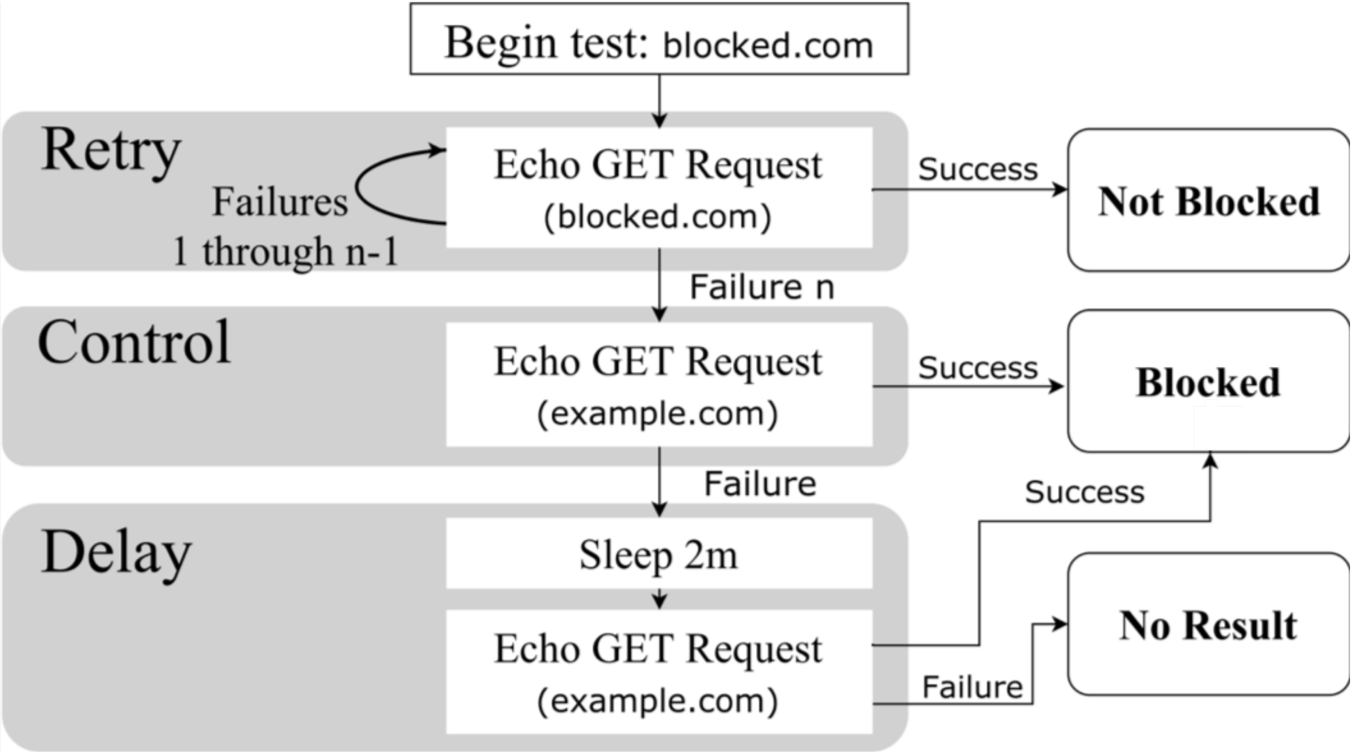

Classification

Adding further redundancy

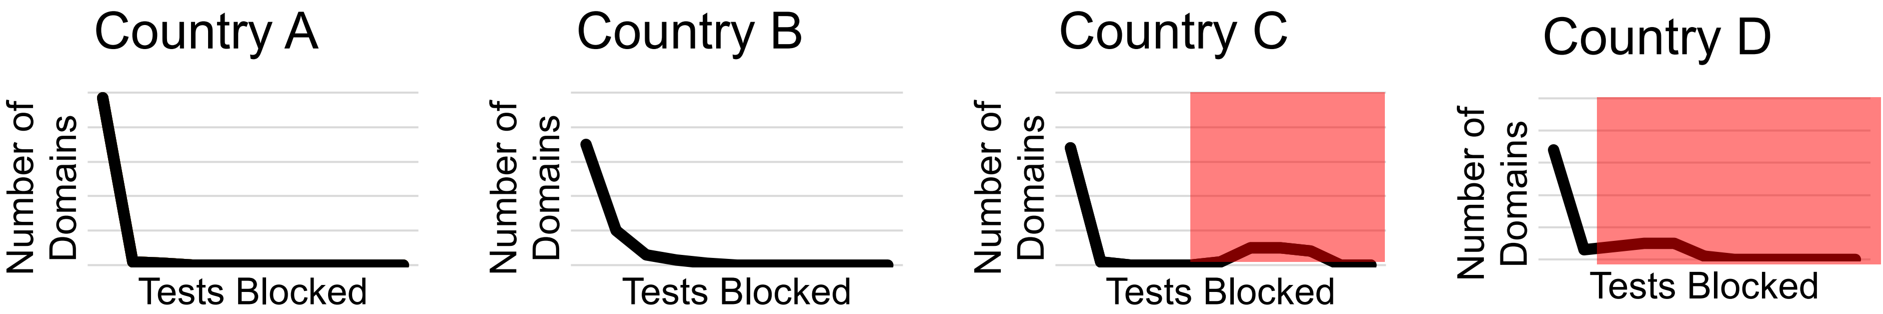

Ignore countries with Blocked results in one autonomous system

We validate countries we classify as being true positives

Ethical Considerations

We can provide useful transparency into censorship

However, we create potential risks for Echo server operators

Is participation voluntary?

Is informed consent feasible?

Do subjects incur no more than minimal risk?

Is this risk reasonable?

Upholding our Principles

We opt not to use informed consent,
but we minimize risk and respect blacklist requests

Our measurements look unlike real traffic on several network layers

Servers don't connect to real sites
Our servers indicate our research use
- DNS PTR
- WHOIS
- Explanatory webserver

Connections are port 7 to ephemeral
Real website data never sent to echo server
Requests have different headers
Echo servers are normally not user devices

Sought IRB approval, but were deemed outside their scope

Discovering Echo Servers

5,000,000 servers reply with a SYNACK on port 7 in 6,900 ASes (198 countries)

Only 57,000 complete a trial, from 3,766 ASes (172 countries)

Only **47,000** are there a day later, from **3,463 ASes (167 countries)**

What are Echo servers?

TCP-level detection could work at scale, but tells kernel level behavior

Take a 1% sample, and perform NMap OS detection

44.7% “server,” “router,” or “switch”
12.5% “Linux” but not the above
17.0% unidentified

4% are non-server windows machines and 2 Android devices.

Citizen Lab List

Hand curated and labeled list of ~1000 domains

Topics and domains that are either censored or interesting

What can we learn from testing with HTTP formatted echo scans?

Citizen Lab List HTTP Experiments

Censorship observed in 12 countries:

China, Egypt, Iran, Jordan, Kazakhstan, Saudi Arabia,
South Korea, Thailand, Turkey, UAE, Uzbekistan

Most frequent categories:

News, Anonymization, Pornography

Citizen Lab List HTTP Experiments

Censorship observed in 12 countries:

China, **Egypt**, Iran, **Jordan**, Kazakhstan, **Saudi Arabia**, South Korea, **Thailand**, **Turkey**, **UAE**, Uzbekistan

Most frequent categories:

News, Anonymization, Pornography

Citizen Lab List HTTP Experiments

Censorship observed in 12 countries:

China, **Egypt**, Iran, **Jordan**, Kazakhstan, **Saudi Arabia**,
South Korea, Thailand, **Turkey**, UAE, Uzbekistan

Most frequent categories:

News, **Anonymization**, Pornography

Citizen Lab List HTTP Experiments

Censorship observed in 12 countries:

China, Egypt, **Iran**, Jordan, Kazakhstan, Saudi Arabia, **South Korea**, **Thailand**, Turkey, UAE, Uzbekistan

Most frequent categories:

News, Anonymization, **Pornography**

Validating Detected Disruption

Freedom on the Net by Freedom House and OpenNet Initiative as ground truth

All countries from the previous list are known to have deployed technical means!

Several countries that are “Not Free” that we tested.

e.g. Pakistan deploys DNS based censorship

HTTP vs HTTPS

TLS Server Name Indication allows web-hosts to serve the correct certificate

Of our original countries, the bold blocked HTTPS:

China, **Egypt**, **Iran**, **Jordan**, Kazakhstan, Saudi Arabia,
South Korea, Thailand, **Turkey**, **UAE**, **Uzbekistan**

Iran censored 374 domains when tested with HTTPS vs 25 with HTTP

Alexa Top 100,000

Significant improvement in domain test list size

Achieved by restricting to 20 tests per domain per country,

We require 100 servers, so we only test 40 countries

Took less than three days on one test machine

Alexa Top 100,000 Blocking Experiments

China, Egypt, Iran, Saudi Arabia, South Korea, Thailand, Turkey

3293 censored domains, 180 from Citizen Lab List

Number of domains blocked by country does not correlate with our earlier tests

Censored category not seen in Citizen Lab List testing: **Shopping**

Popularity of Blocked Domains

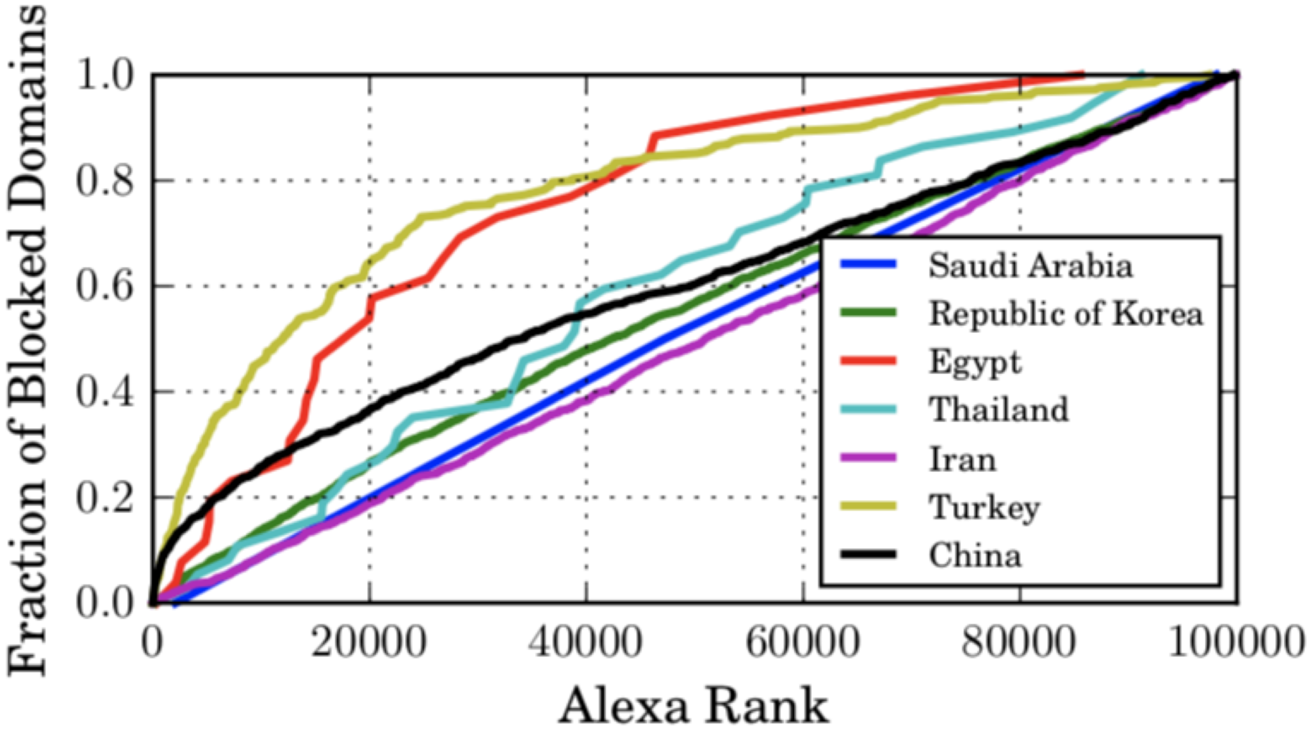

Limitations

Quack is easier to block due to minimal collateral damage

Further research is required to explore further

Countries could be blocking block TCP direction that sent SYN

May not detect heterogenous deployments

Coverage reduces as we require more tests in each country

Conclusions

Application-layer censorship can be measured remotely

We test an order of magnitude more domains than prior work

Future work should combine Quack with other remote measurements

We would like to acknowledge

Bill Marczak
Adam Bates
Our Reviewers

Quack:
Scalable Remote Measurement of Application-Layer Censorship

Benjamin VanderSloot, Allison McDonald,
Will Scott, J. Alex Halderman, and Roya Ensafi

UNIVERSITY OF
MICHIGAN

Measuring Asymmetry in HTTP Blocking

Many Echo servers are also Discard servers

```
$ netcat -l 9 > /dev/null
```

Of our original countries, the bold required echo behavior:

China, Egypt, **Iran**, Jordan, Kazakhstan, **Saudi Arabia**,
South Korea, **Thailand**, Turkey, **UAE**

Churn

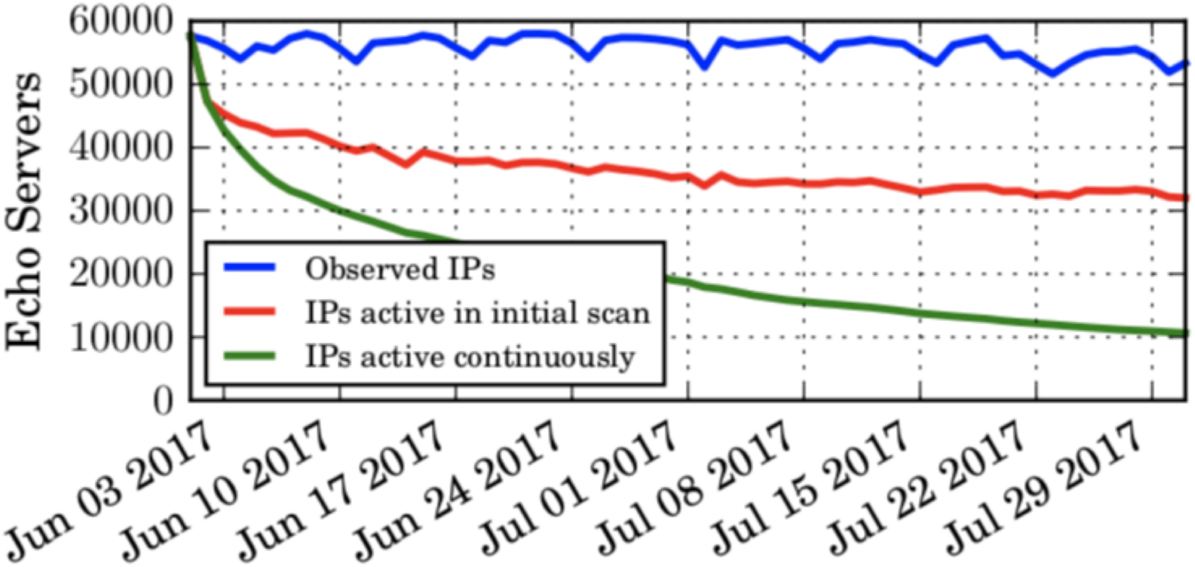
Echo Servers
60000
50000
40000
30000
20000
10000
0
Jun 03 2017
Jun 10 2017
Jun 17 2017
Jun 24 2017
Jul 01 2017
Jul 08 2017
Jul 15 2017
Jul 22 2017
Jul 29 2017
Observed IPs
IPs active in initial scan
IPs active continuously

Coverage

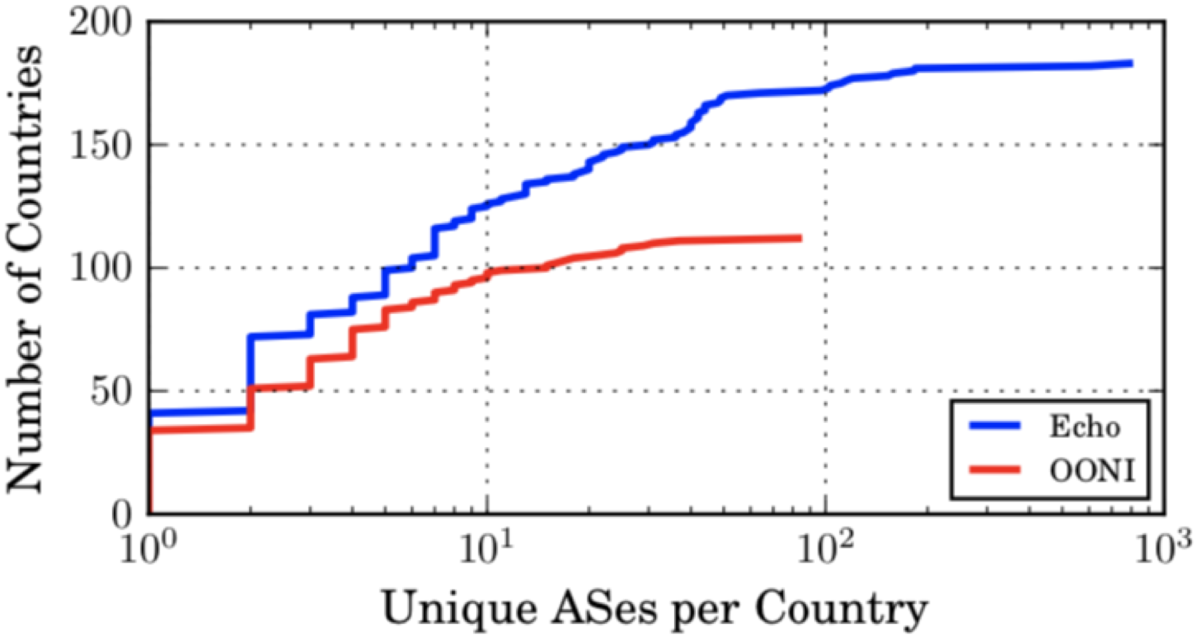
Number of Countries
200
150
100
50
0
10^0
10^1
10^2
10^3
Unique ASes per Country
Echo
OONI

Validation

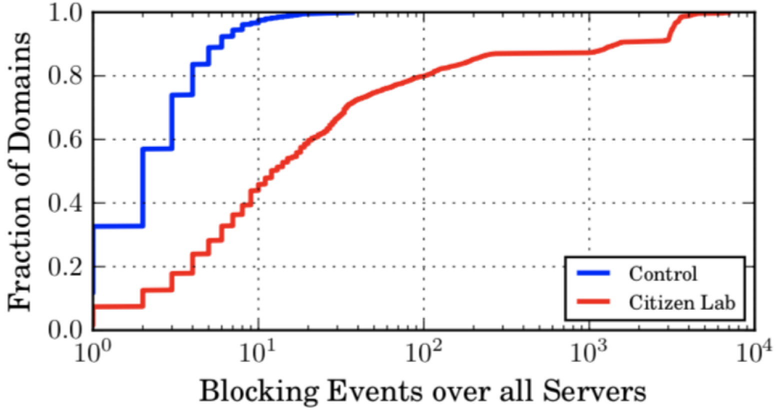
Fraction of Domains
1.0
0.8
0.6
0.4
0.2
0.0
10^0
10^1
10^2
10^3
10^4
Blocking Events over all Servers
Control
Citizen Lab

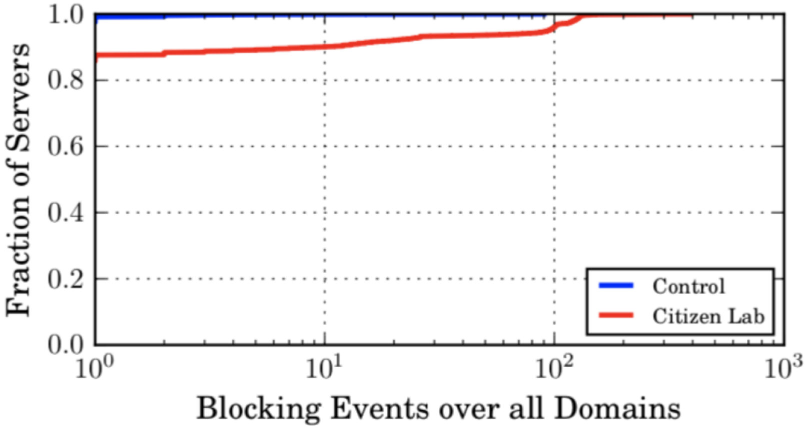
Fraction of Servers
1.0
0.8
0.6
0.4
0.2
0.0
10^0
10^1
10^2
10^3
Blocking Events over all Domains
Control
Citizen Lab