Censored Planet:
Global Censorship Observatory

Roya Ensafi
University of Michigan
Dec 27,2018

In my research lab, we …

develop frameworks to **detect** network interference,

apply these frameworks to **understand the behavior** of network intermediaries,

and use this understanding to **defend against interference** by building tools that safeguard users.

Reports suggest

Internet censorship practices

are at rise!

Network Interference Can Happen on Any Layer

1. A user types www.cnn.com into the browser
2. OS sends a DNS query to learn the IP address
3. Browser fetches the website
4. Browser loads third-party resources

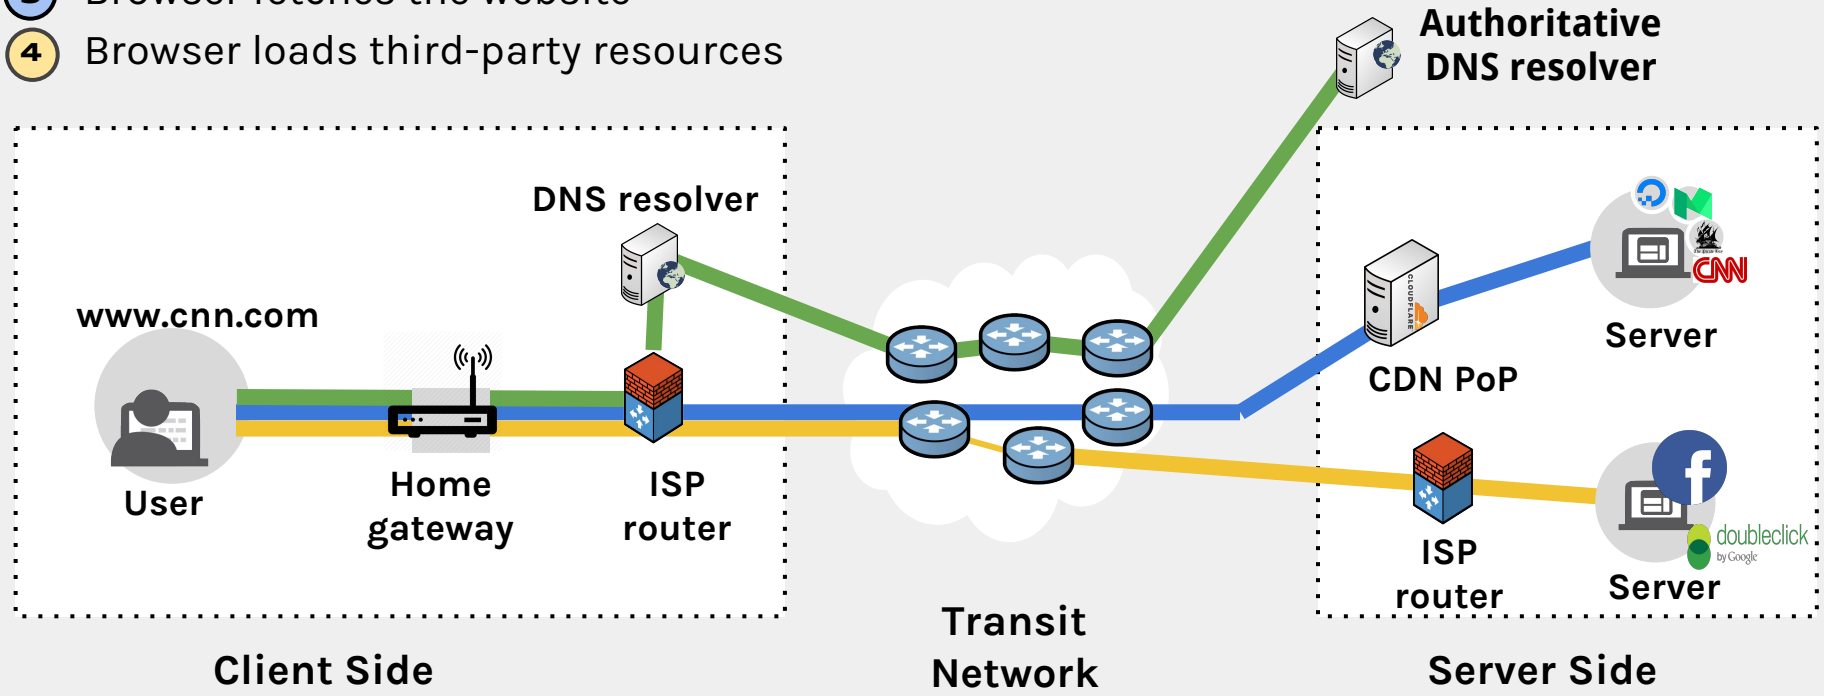

Network Interference Can Happen on Any Layer

1. A user types www.cnn.com into the browser
2. OS sends a DNS query to learn the IP address
3. Browser fetches the website
4. Browser loads third-party resources

403 Forbidden: A Global View of CDN Geoblocking

Allison McDonald
University of Michigan
amcdon@umich.edu

Matthew Bernhard
University of Michigan
matber@umich.edu

Luke Valenta
University of Pennsylvania
lukev@seas.upenn.edu

Benjamin VanderSloot
University of Michigan
benvds@umich.edu

Will Scott
University of Michigan
willrs@umich.edu

Nick Sullivan
Cloudflare
nick@cloudflare.com

J. Alex Halderman
University of Michigan
jhalderm@umich.edu

Roya Ensafi
University of Michigan
ensafi@umich.edu

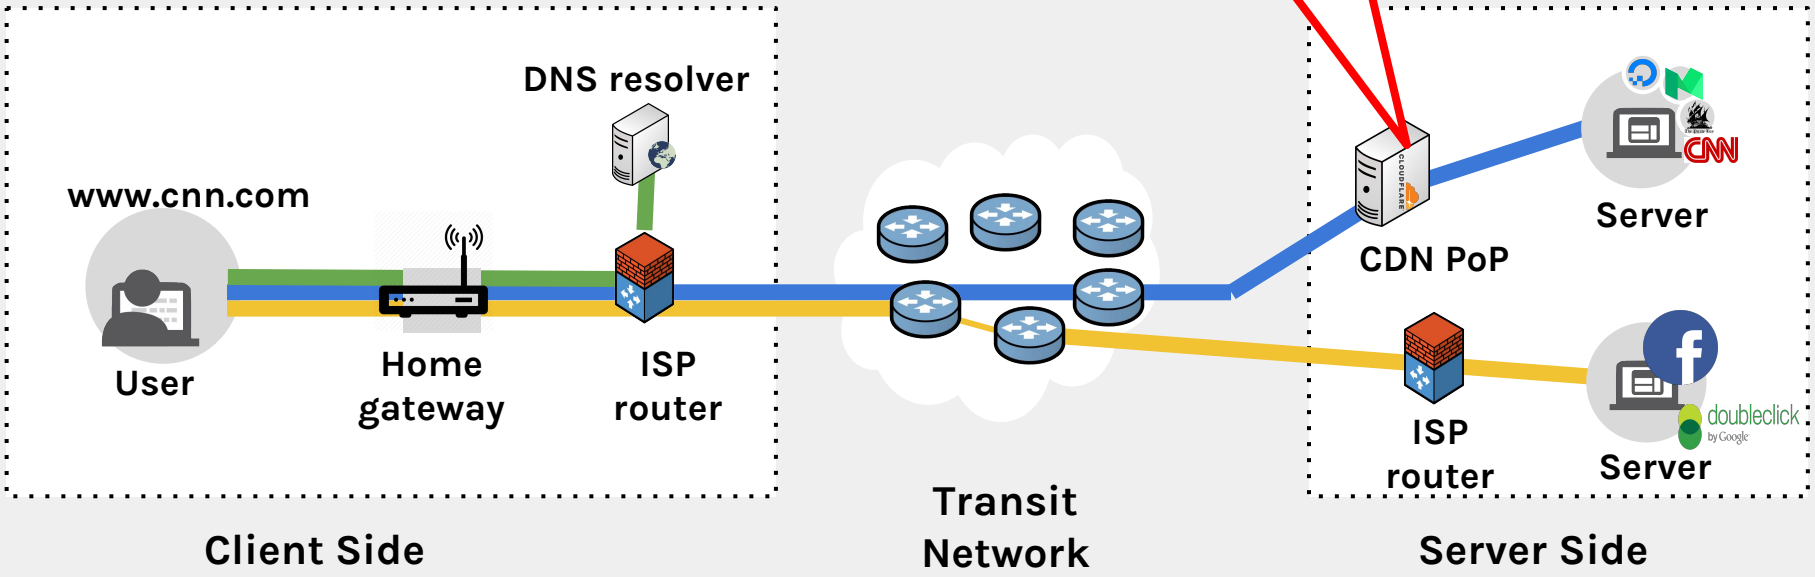

Measuring Censorship is a Complex Problem!

Internet censorship practices are diverse in their methods, targets, timing, differing by regions, as well as across time.

Why Measure Censorship?

NETWORK CENSORSHIP IS ON THE RISE

- Information controls harm citizens
- Spreading beyond the large powers
- Frequently opaque in topic & technique

WE NEED DATA TO:

- Support transparency & accountability
- Improve technological defenses
- Inform users & public policy

Why Measure Censorship?

NETWORK CENSORSHIP IS ON THE RISE

- Information controls harm citizens
- Spreading beyond the large powers
- Frequently opaque in topic & technique

WE NEED DATA TO:

- Support transparency & accountability
- Improve technological defenses
- Inform users & public policy

Freedom on the Net 2018:

"...When users become more aware of censorship, they often take actions that enhance [I]nternet freedom and protect fellow users"

The Vision

“Censorship weather map” to continually monitor Internet censorship around the world

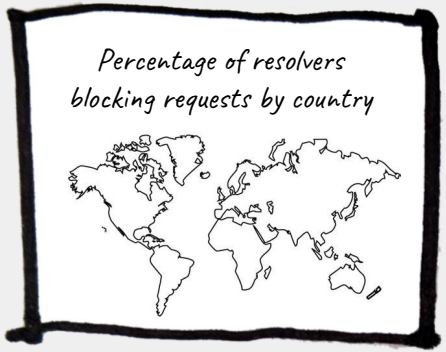

How Have We Collected Data on Censorship?

Common approach:

- Deploy hardware or software in censored region (e.g. RIPE Atlas, OONI probe)
- Ask people on the ground, or use VPNs, or research networks (e.g., FreedomHouse, PlanetLab)

THREE KEY CHALLENGES:

Coverage, continuity, and ethics

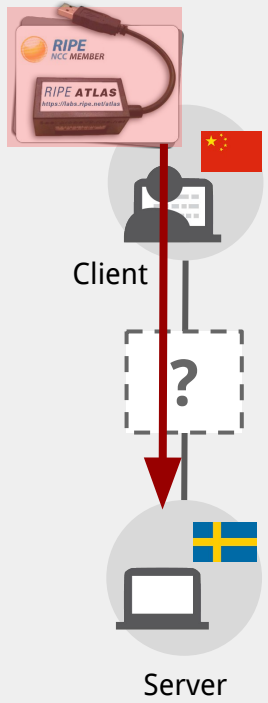

Collecting consistent, continuous, and global data requires a different approach.

From: mic.com

From: CNN.com

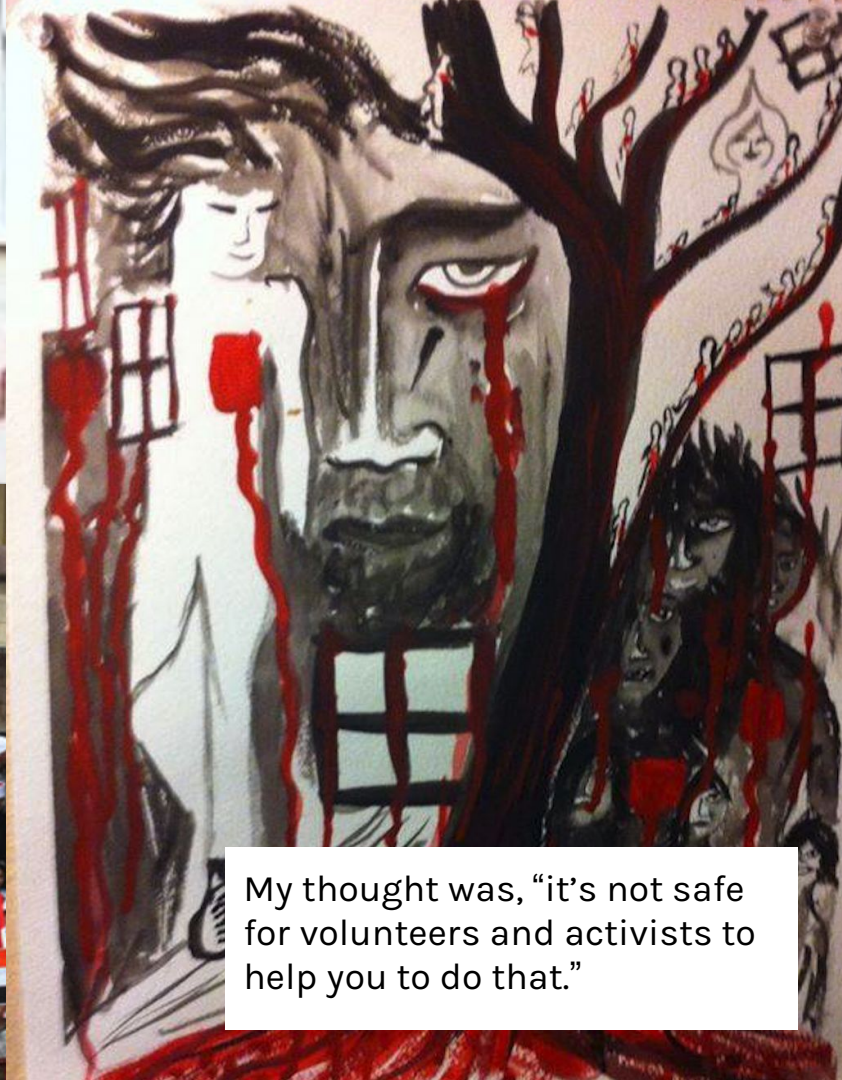

My thought was, “it’s not safe for volunteers and activists to help you to do that.”

Freedom on the Net 2018

“Many governments are enforcing criminal penalties for the publication of what they deem false news”

How OONI Deals with Potential Risks?

OONI is a global community of volunteers collecting data on Internet censorship

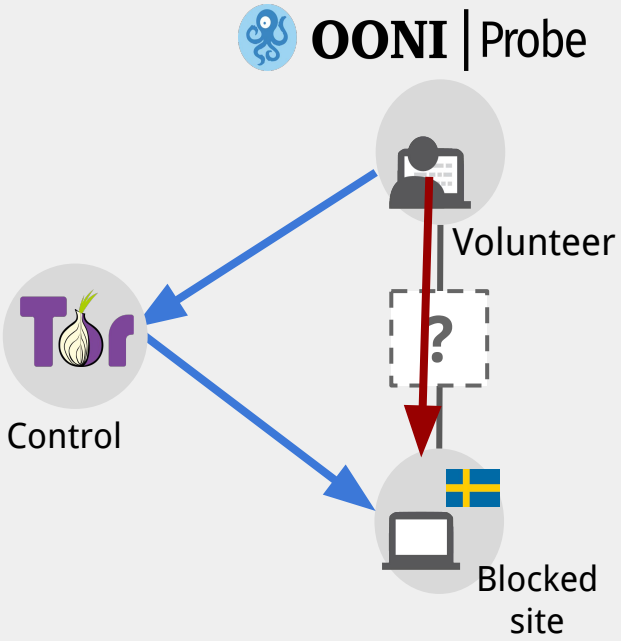

How OONI Deals with Potential Risks?

To minimize potential risk, OONI:

- “Provide as much informed choice to the user as possible => being able to choose which websites to test, whether to upload measurements or not, what type of data to submit, etc.”
- Establish relationships with locals & civil society
- Keep the community of volunteer involved

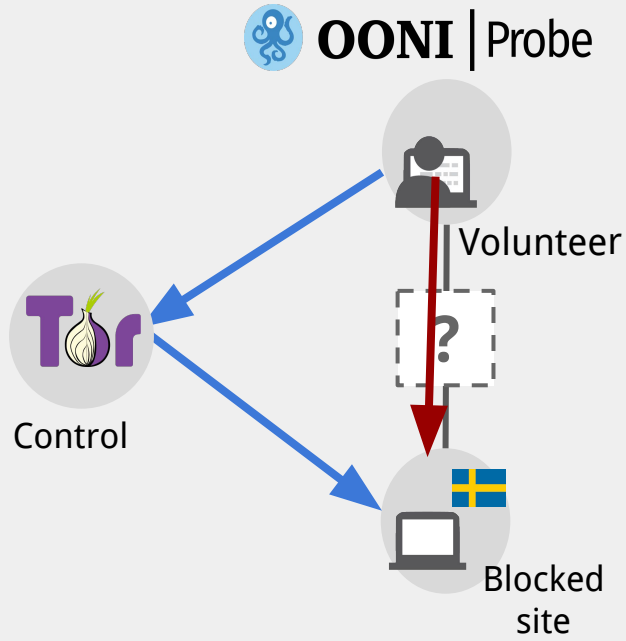

Measuring Internet Censorship Globally... Remotely!

REFRAMING THE PROBLEM:

How can we detect whether pairs of hosts around the world can talk to each other?

... without volunteer participation?

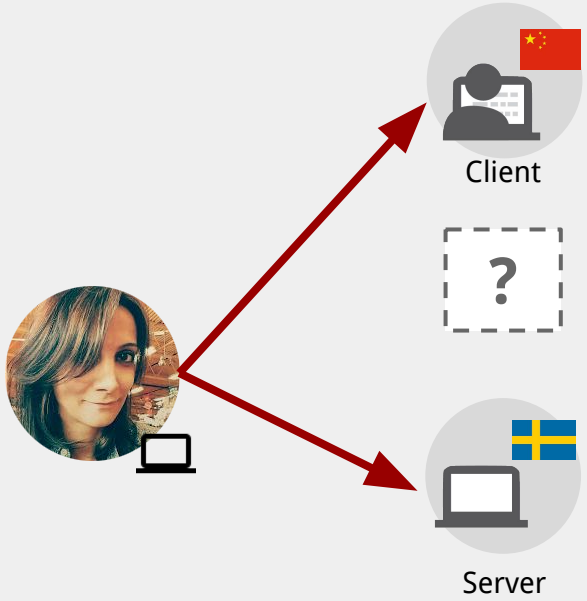

Leveraging Existing Hosts as Vantage Points

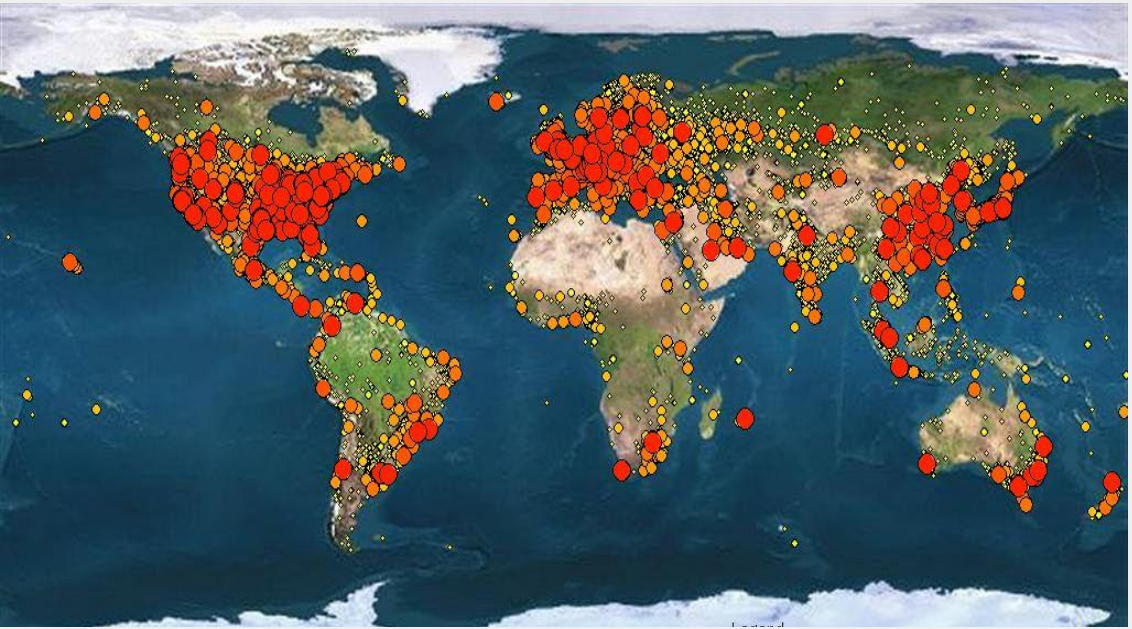

140 million IPv4 hosts that respond to TCP SYNs

These machines speak to the world, and they follow TCP/IP, the basic communication protocol of the Internet.

How can we leverage subtle TCP behavior to detect whether two distant hosts can communicate?

Spooky Scan

Spooky Scan uses TCP/IP side-channels to detect whether a client and server can communicate (and in which direction packets are blocked)

Goal: Detect blocking from off-path

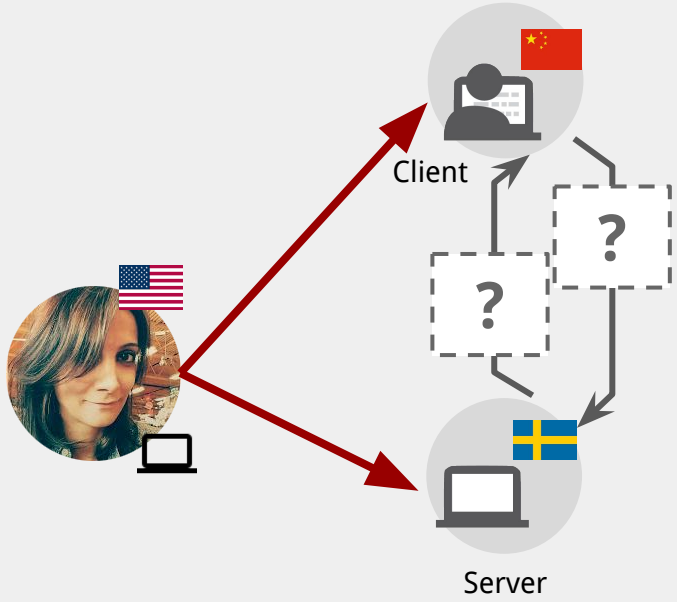

* **Detecting Intentional Packet Drops on the Internet via TCP/IP Side Channels**
  Roya Ensafi, Knockel, Alexander, and Crandall (PAM '14)
* **Idle Port Scanning and Non-interference Analysis of Network Protocol Stacks Using Model Checking**
  Roya Ensafi, Park, Kapur, and Crandall (Usenix Security 2010)
* **TCP Idle Scan** Antirez (Bugtraq 1998)

Background: TCP/IP Protocol

TCP Handshake:

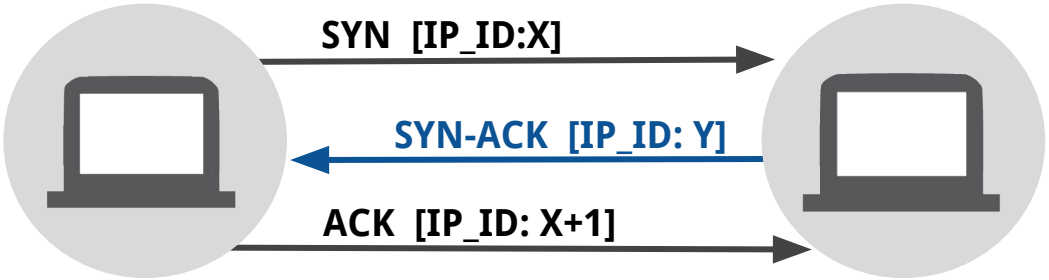

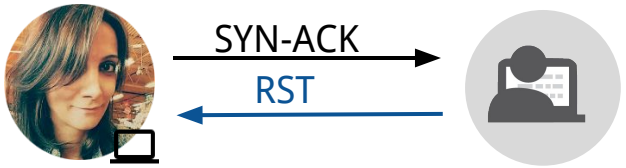

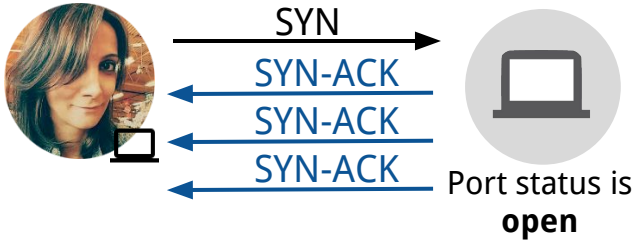

Spooky Scan Requirements

Client

Must maintain a global value for IP_ID

Server

Open port and retransmitting SYN-ACKs

Measurement Machine

Must be able to spoof packets

Spooky Scan

Measurement machine

Client

Client IP_ID

Server

Spooky Scan

No direction blocked

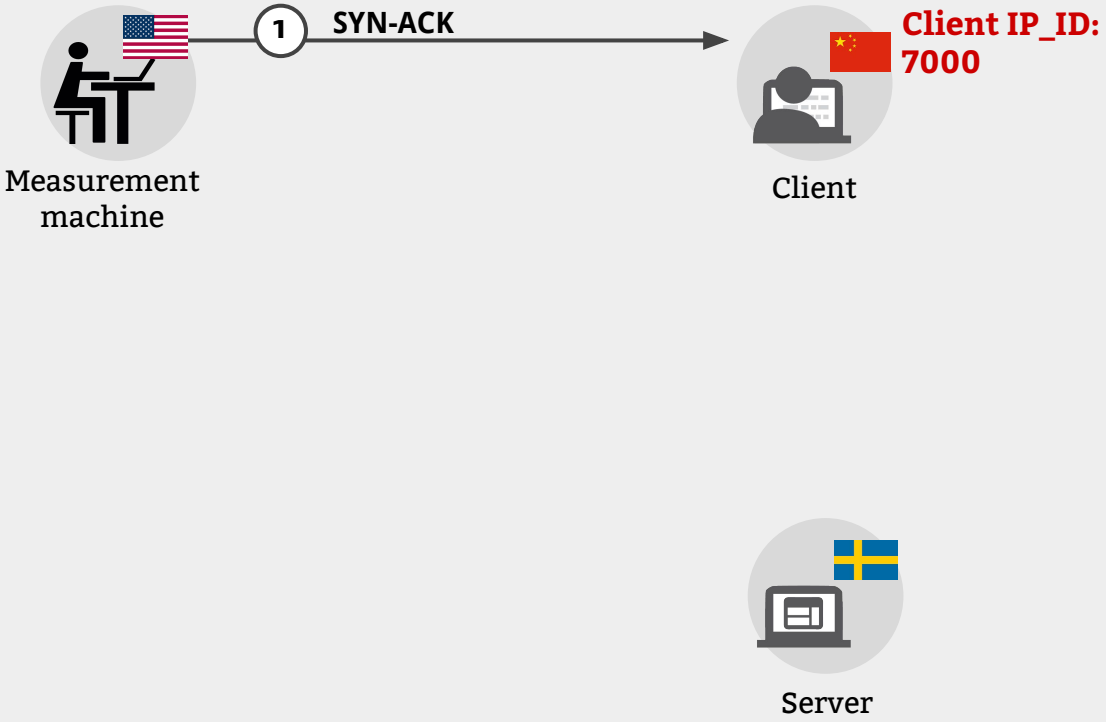

Spooky Scan

No direction blocked

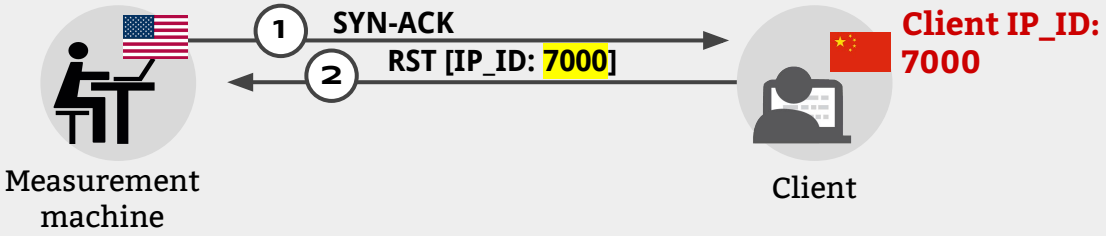

Spooky Scan

No direction blocked

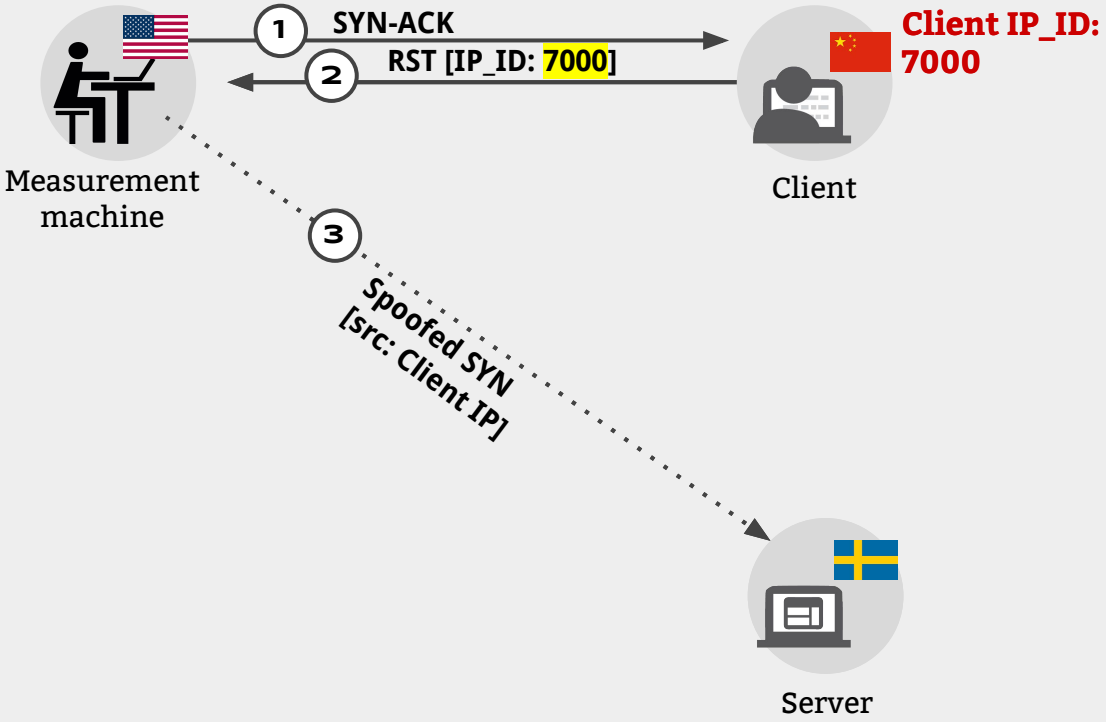

Spooky Scan

No direction blocked

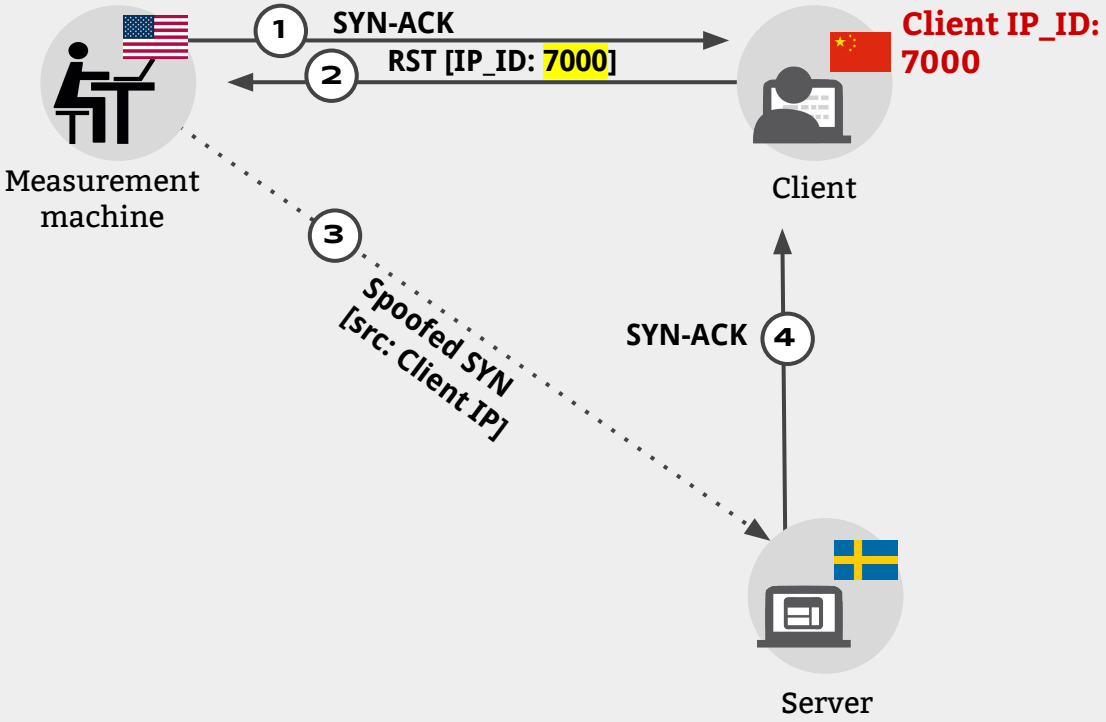

Spooky Scan

No direction blocked

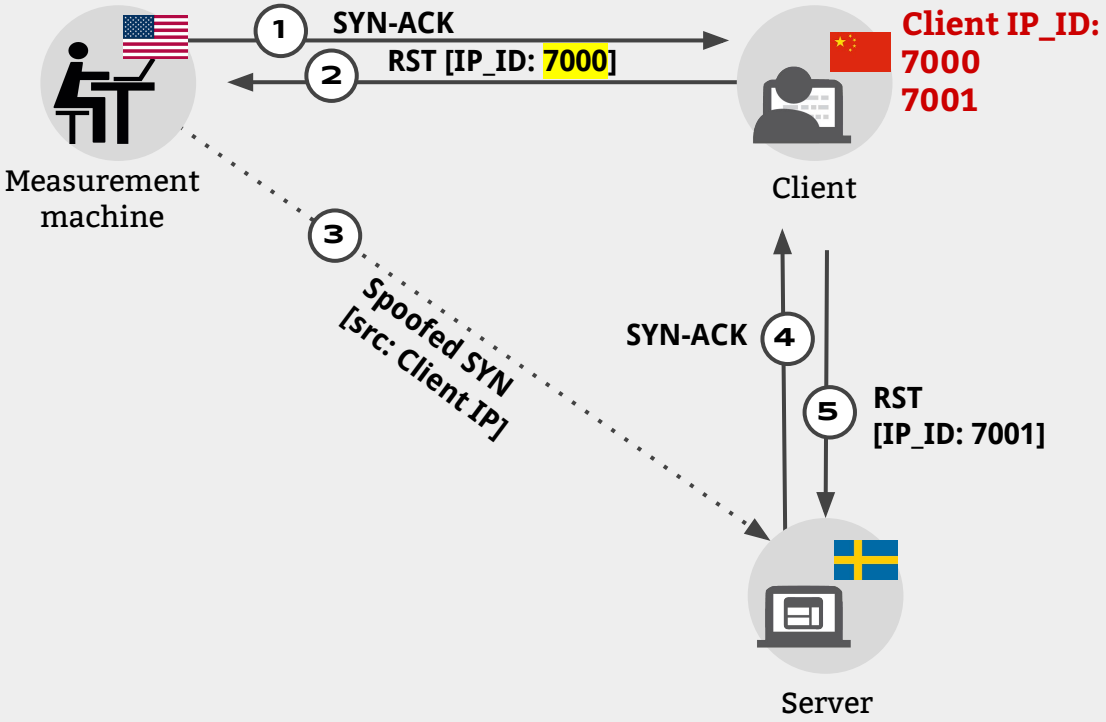

Spooky Scan

No direction blocked

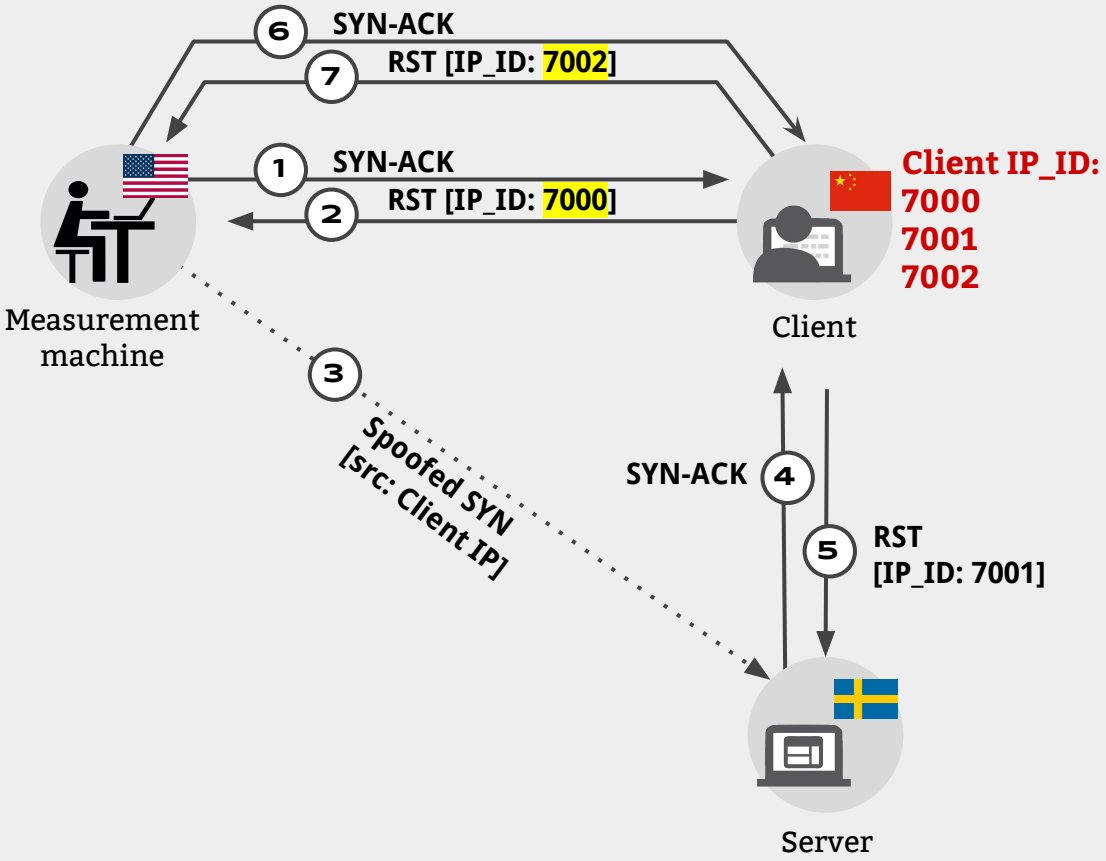

Spooky Scan

No direction blocked

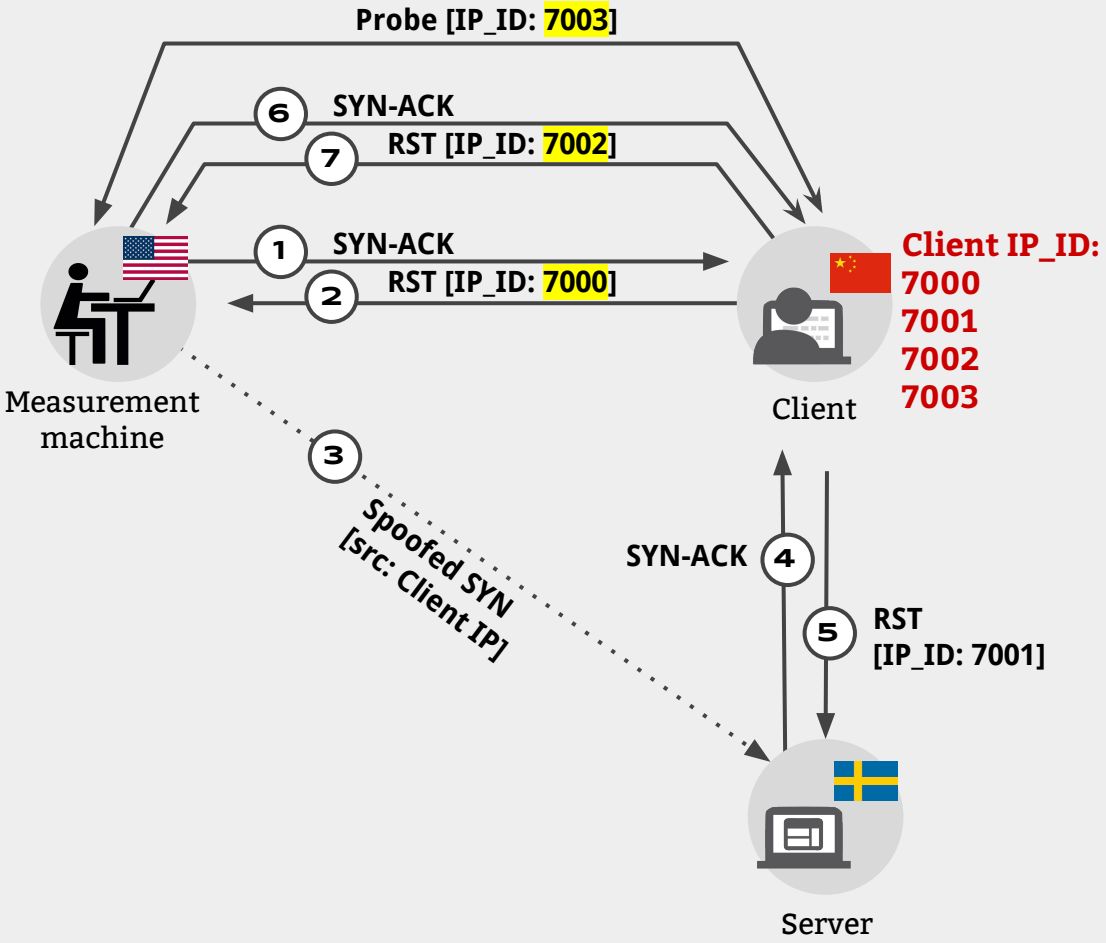

Spooky Scan

Server-to-Client blocked

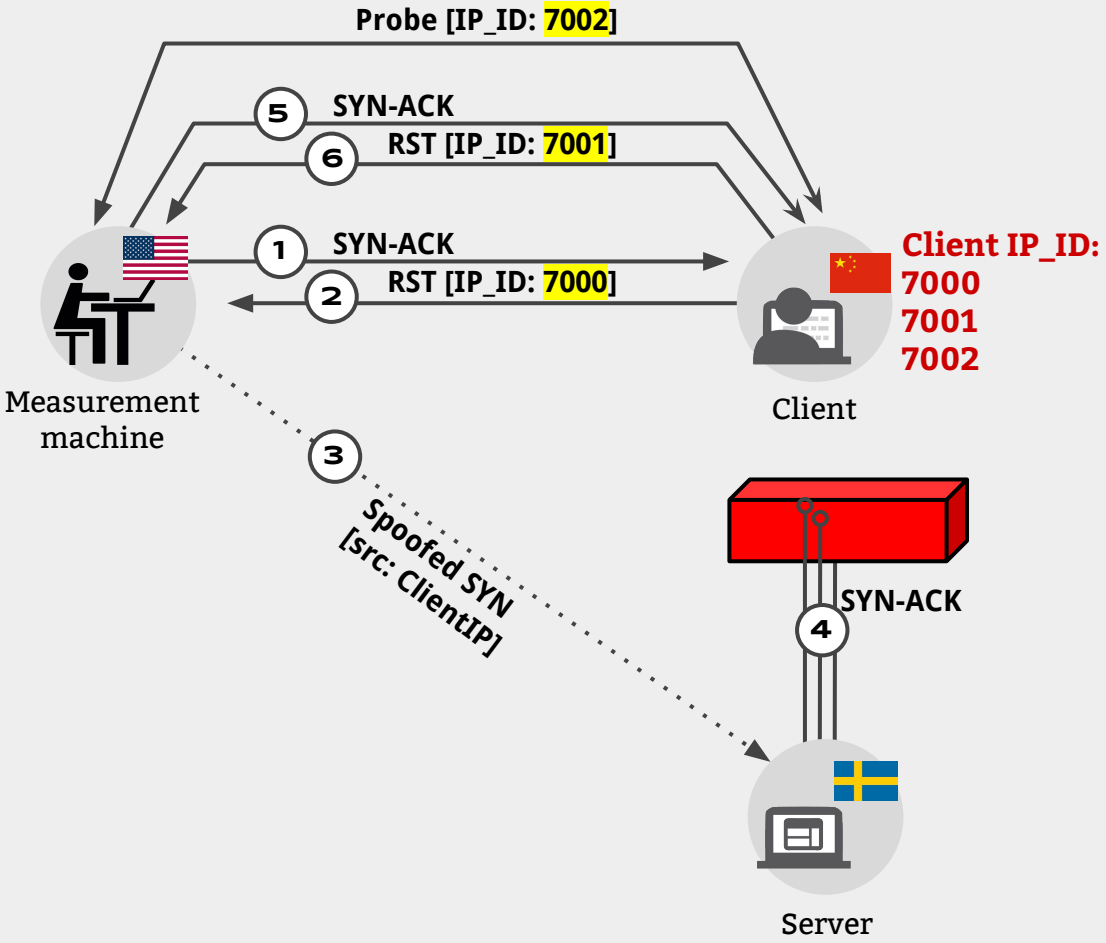

Spooky Scan

Client-to-Server blocked

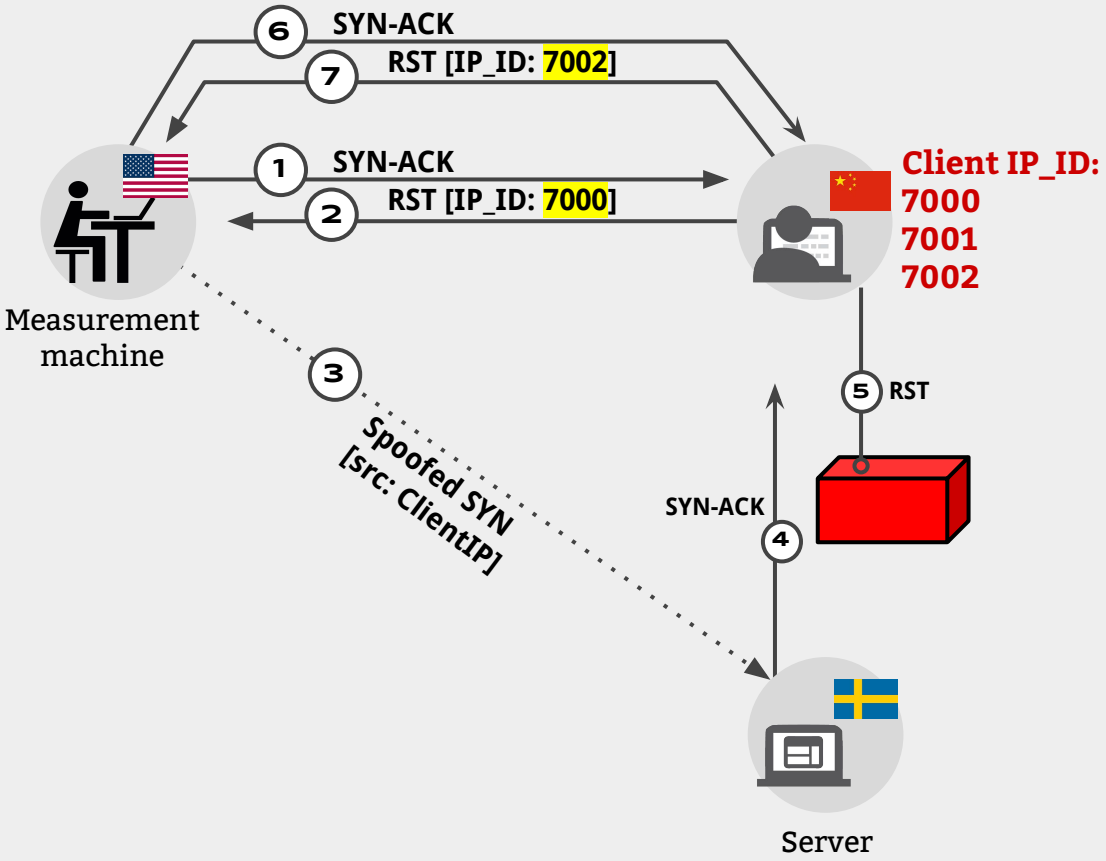

Spooky Scan

Client-to-Server blocked

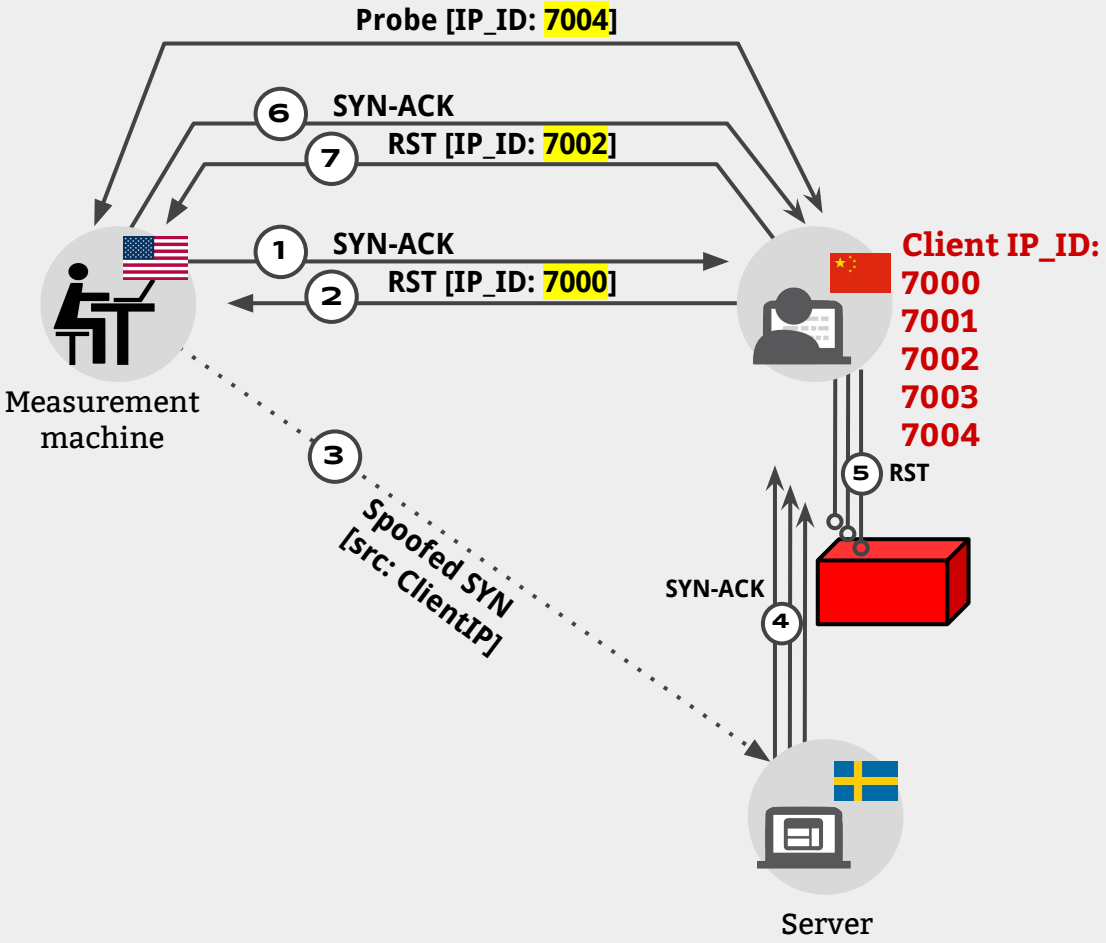

Spooky Scan

Server-to-Client Blocked

Δ IP_ID1 = 1
Δ IP_ID2 = 1

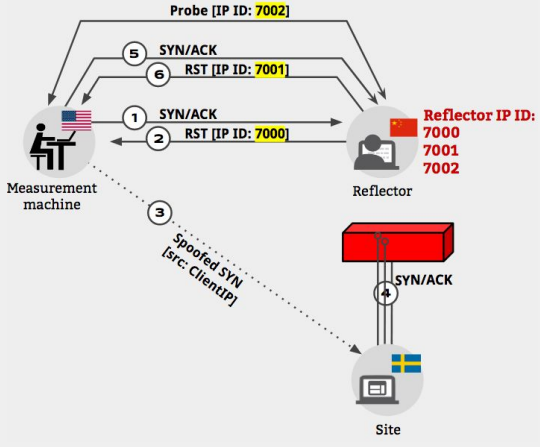

No Direction Blocked

Δ IP_ID1 = 2
Δ IP_ID2 = 1

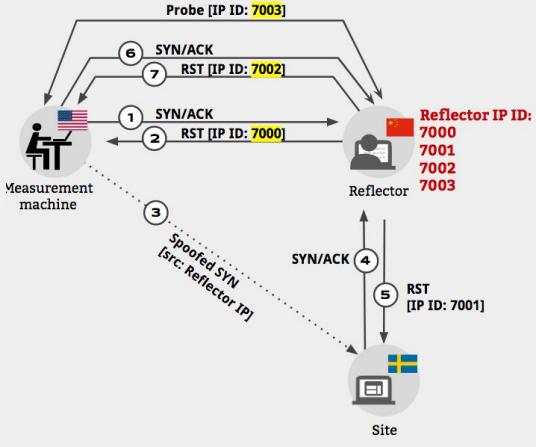

Client-to-Server Blocked

Δ IP_ID1 = 2
Δ IP_ID2 = 2

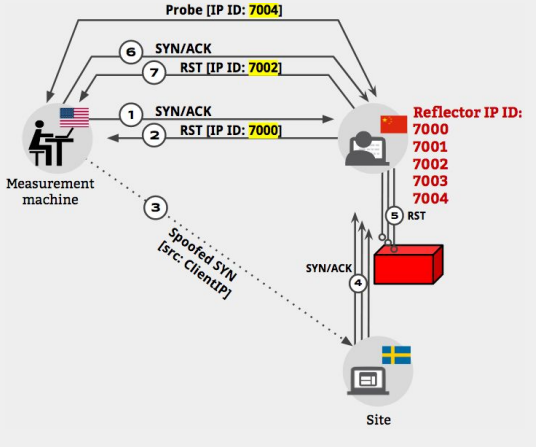

Client IP_ID Noise

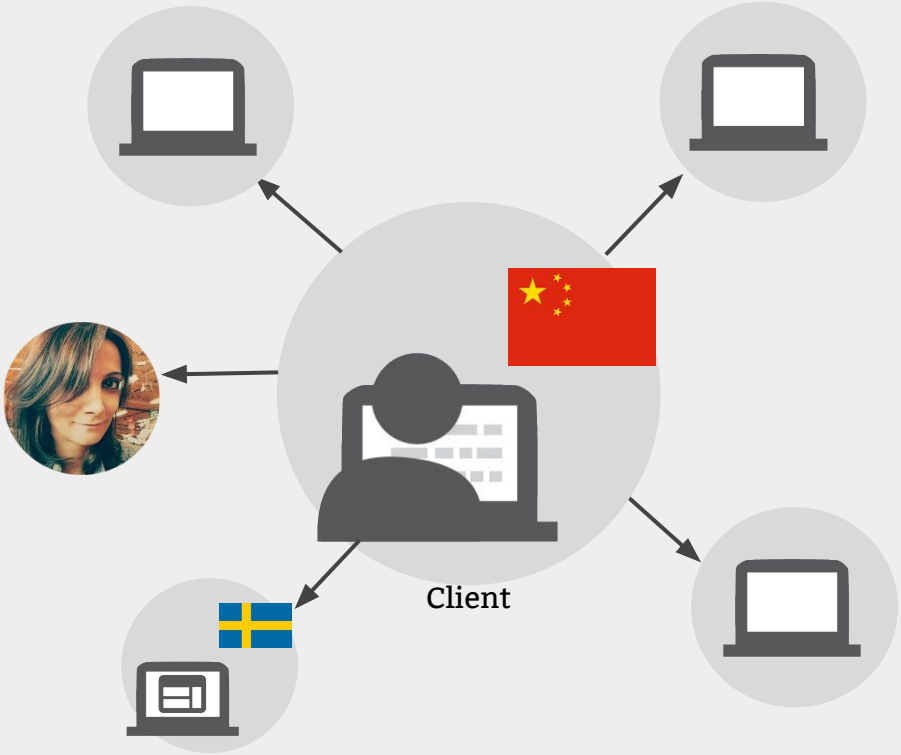

Coping with Client IP_ID Noise

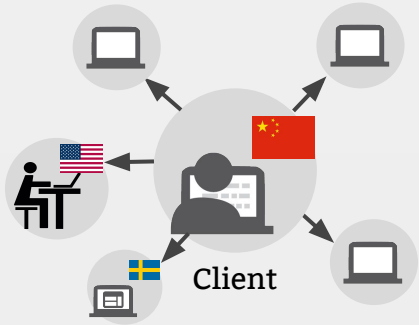

Amplifying the signal

Effect of sending *N* spoofed SYNs:

Server-to-Client Blocked	No Direction Blocked	Client-to-Server Blocked
Δ IP_ID1 = (1 + noise) Δ IP_ID2 = noise	Δ IP_ID1 = (1 + N + noise) Δ IP_ID2 = noise	Δ IP_ID1 = (1 + N + noise) Δ IP_ID2 = (1 + N + noise)

Coping with Client IP_ID Noise

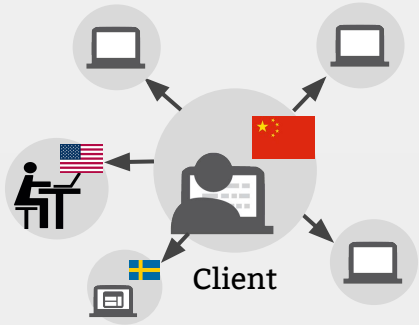

Amplifying the signal

Effect of sending *N* spoofed SYNs:

Server-to-Client Blocked	No Direction Blocked	Client-to-Server Blocked
Δ IP_ID1 = (1 + noise) Δ IP_ID2 = noise	Δ IP_ID1 = (1 + N + noise) Δ IP_ID2 = noise	Δ IP_ID1 = (1 + N + noise) Δ IP_ID2 = (1 + N + noise)

Repeating the experiment

To eliminate the effects of packet loss, sudden bursts of packets, ...

Spooky Scan with Noise: Visualization

Probing method

For first 30s, query IP_IDs. Then, for another 30s

- Send 5 spoofed SYNs per second
- Query IP_ID once per second

Server-to-Client Blocked

Tor relay (SE) to client (CN) blocked

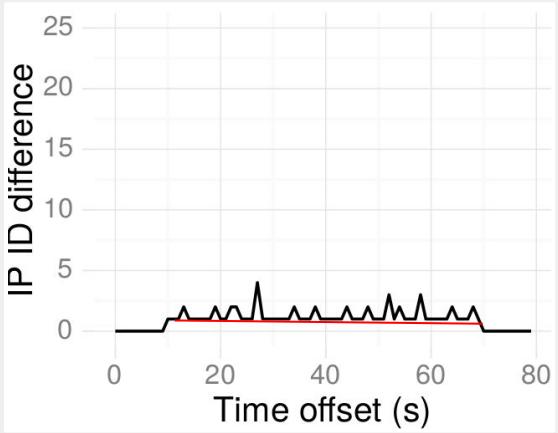

No Direction Blocked

No block btw client (US) and Tor relay (SE)

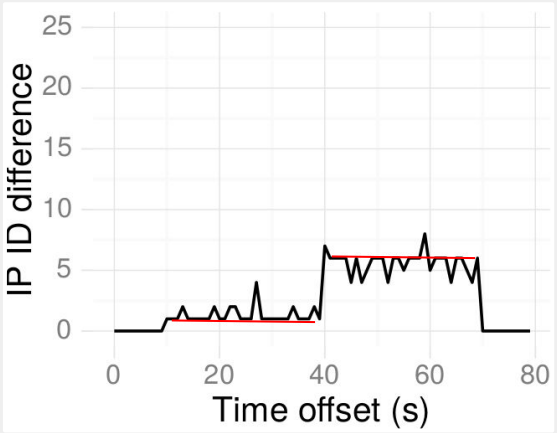

Client-to-Server Blocked

Client (AZ) to Tor relay (SE) blocked

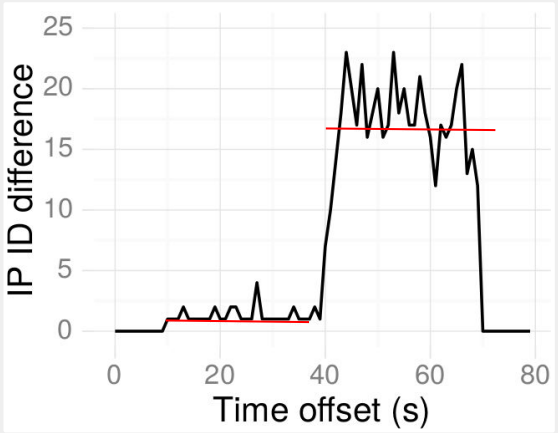

Augur: Spooky for Continuous Scanning

Problem: Want to optimize Spooky to probe many hosts, all the time.

Insight: Some measurements are much noisier than others.

* **Internet-Wide Detection of Connectivity Disruptions**
P. Pearce*, R. Ensafi*, F. Li, N. Feamster, V. Paxson *joint first authors
IEEE S&P (“Oakland”) 2017

Augur: Spooky for Continuous Scanning

Problem: Want to optimize Spooky to probe many hosts, all the time.

Insight: Some measurements are much noisier than others.

Probing Methodology:

Until we have high enough confidence (or up to):

Run
- For first 4s, query IP_ID every sec
-
  - Send 10 spoofed SYNs
  - Query IP_ID
- Query IP_ID

Augur: Spooky for Continuous Scanning

Problem: Want to optimize Spooky to probe many hosts, all the time.

Insight: Some measurements are much noisier than others.

Probing Methodology:

Until we have high enough confidence (or up to):

Run
- For first 4s, query IP_ID every sec
- Send 10 spoofed SYNs, Query IP_ID
- Query IP_ID

Repeat runs and use Sequential Hypothesis Testing to gradually build confidence.

Sequential Hypothesis Testing in Augur

Defining a random variable:

$$Y_n(S_i, R_j) = \begin{cases} 1 & \text{if no IP\_ID acceleration occurs} \\ 0 & \text{if IP\_ID acceleration occurs} \end{cases}$$

Sequential Hypothesis Testing in Augur

Defining a random variable:

$$Y_n(S_i, R_j) = \begin{cases} 1 & \text{if no IP\_ID acceleration occurs*} \\ 0 & \text{if IP\_ID acceleration occurs*} \end{cases}$$

*measurement window following injection

Calculate known outcome probabilities (priors):

Prior 1: Prob. of no IP_ID acceleration when there is blocking

Prior 2: Prob. of IP_ID acceleration when there is no blocking

IP_ID evolution in control measurement phase, ~0.5

IP_ID evolution in injection period over all clients, ~1

Sequential Hypothesis Testing in Augur

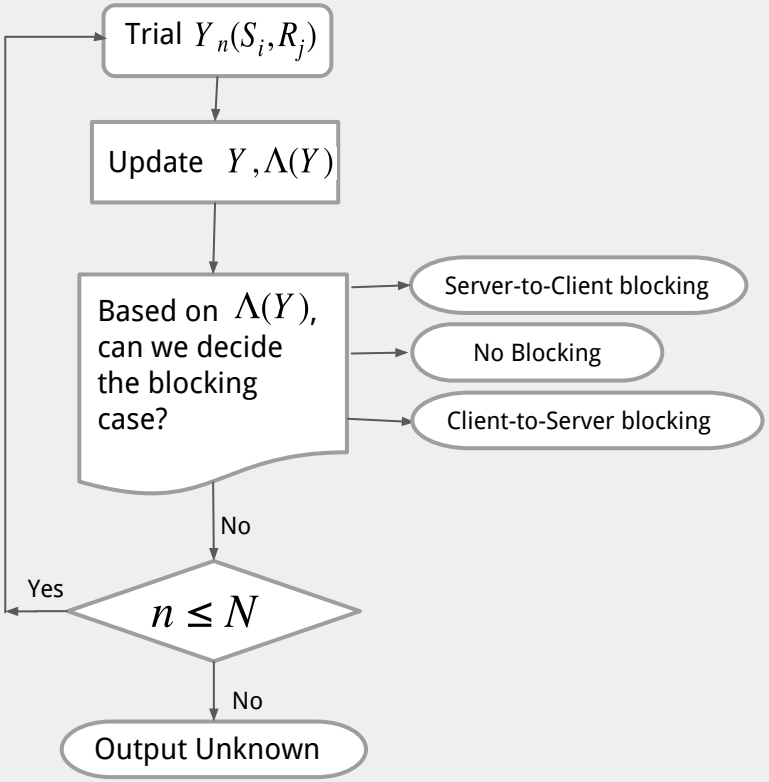

Maximum Likelihood Ratio

$$\Lambda(Y) \equiv \prod_{n=1}^{N} \frac{Pr[Y_n|Blocking]}{Pr[Y_n|No\ Blocking]}$$

Augur Framework

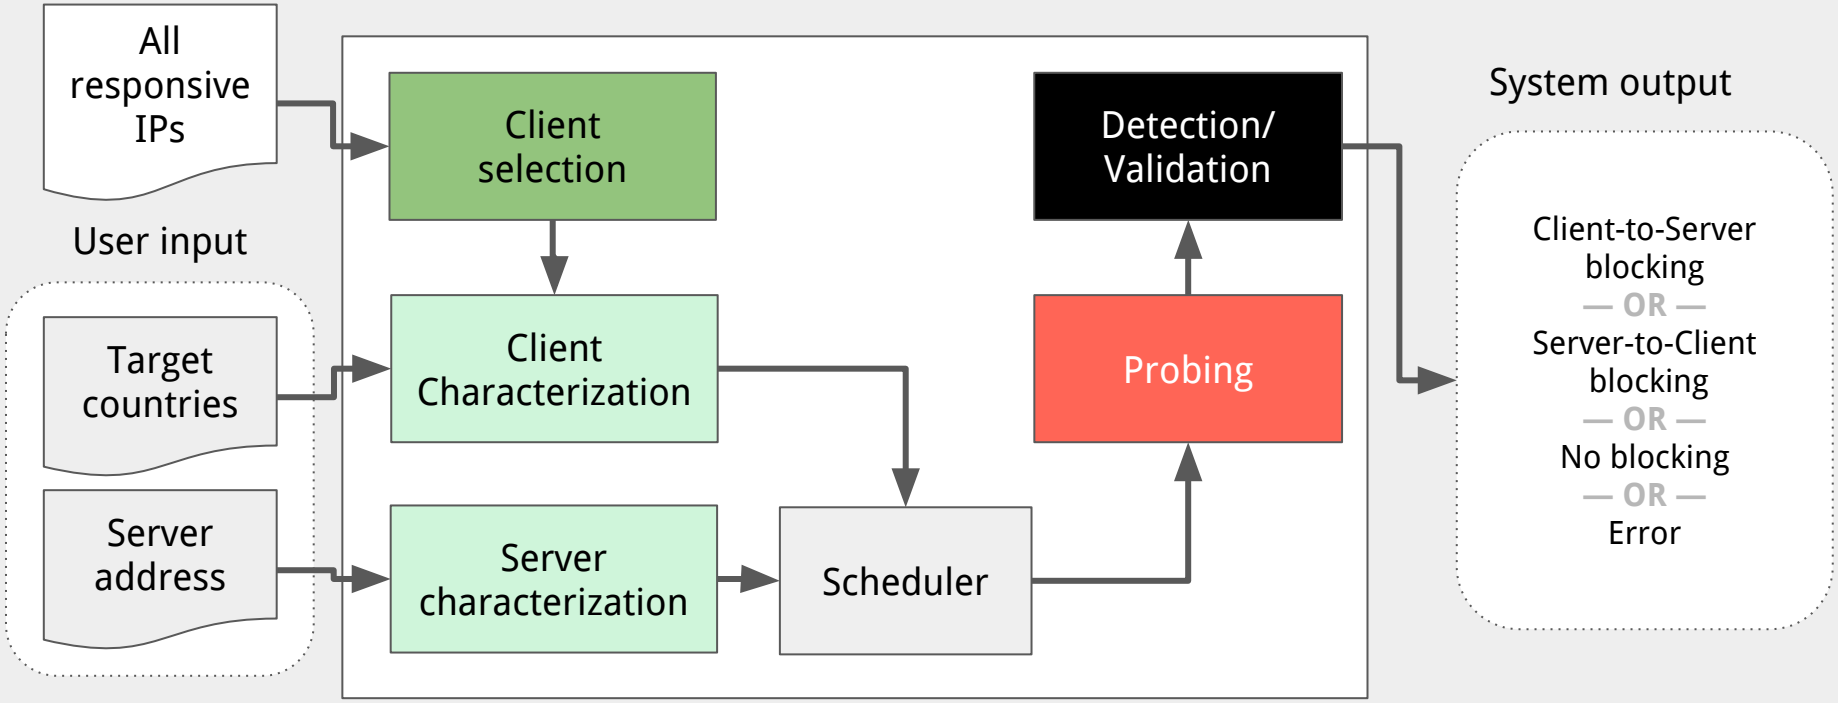

Coverage

CHALLENGE:

Need global vantage points from which to measure

THREE KEY CHALLENGES:
Coverage, continuity, and ethics

Scanning IPv4 on port 80:

22.7 million potential clients (with global IP_ID)

Compare: 10,000 in prior work (RIPE Atlas)

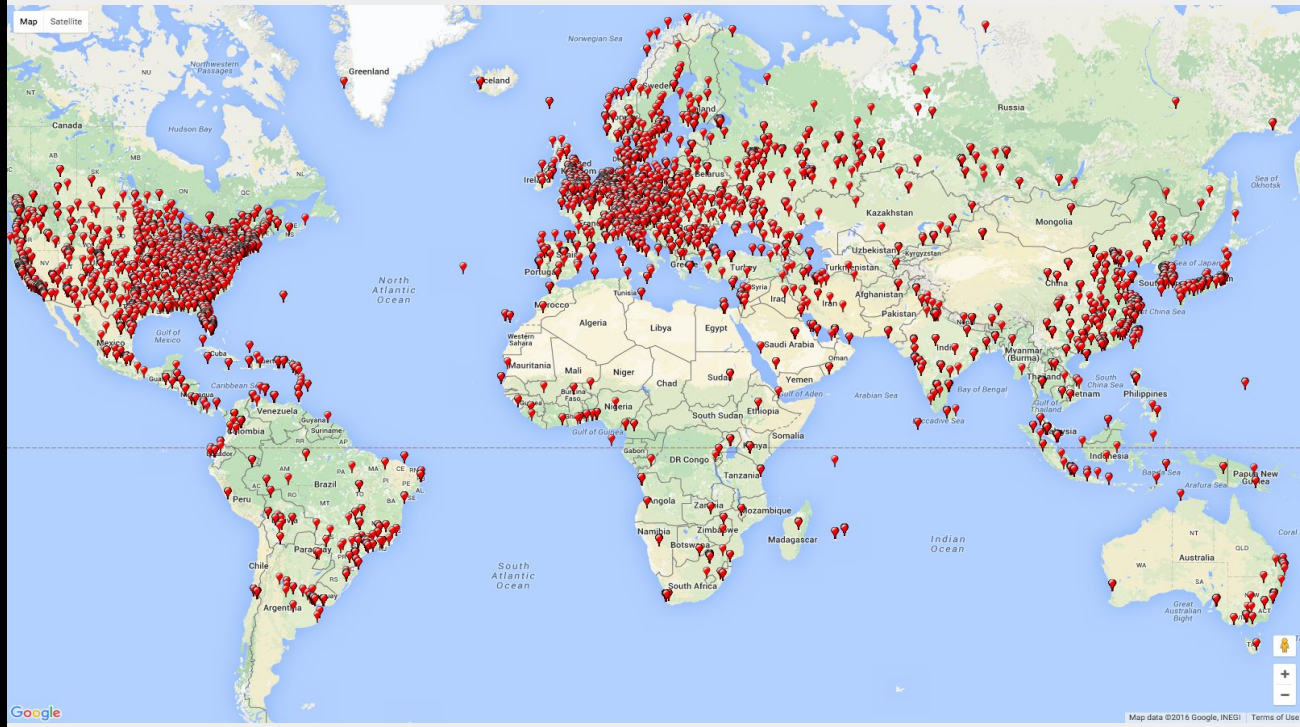

THREE KEY CHALLENGES:
Coverage, continuity, and ethics

Continuity

CHALLENGE:

Need to repeat measurements over time

Augur doesn't depend on end users' participation, allowing us to collect measurements continuously.

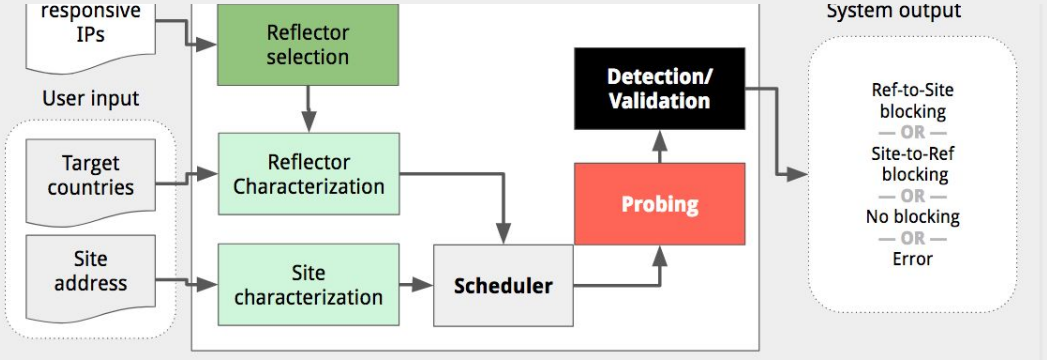

Ethics

CHALLENGE:

Probing banned sites from users' machines creates risk

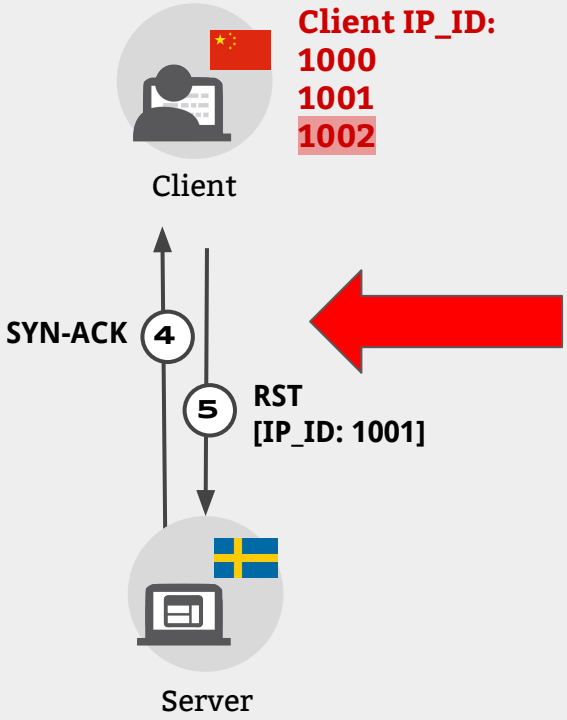

Ethics

CHALLENGE:

Probing banned sites from users' machines creates risk

THREE KEY CHALLENGES:
Coverage, continuity, and ethics

Use only **infrastructure devices** to source probes

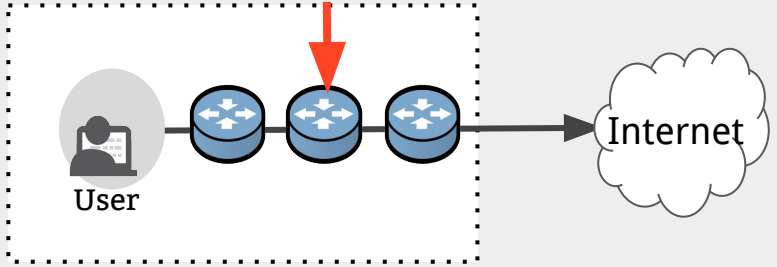

Global IP_ID	22.7 million	236 countries (and dependent territories)
Two hops back from end user	53,000	**180 countries**

Running Augur in the Wild

CHALLENGE:

There is not a good input list of domains, only crowdsource of potentially blocked ones.

Clients: 2,050

Servers: 2,134 (Citizen Lab list + Alexa Top-10K)

Mix of sensitive and popular Sites

Duration: 17 days

Measurements per Client-Server: 47

Overall # of measurements: 207.6 million

Validating Augur

CHALLENGE:

There is **no** ground truth, only anecdotes and reports

✓ **Basic checks based on intuition:**

One Client shouldn't show all sites blocked
99% of clients experience disruption only for 20 or fewer sites

Sites shouldn't be blocked across bulk of Clients
Over 99% of sites exhibit blocking by 100 clients (5%) or less

There should be bias of blocking towards sensitive sites

✓ **Replicating previous findings:**

We should observe countries known to censor heavily

We should observe the same pattern of blocking that

Tor bridges are subject to blocking in China

Augur

Augur is a system that uses infrastructure devices and Spooky's TCP/IP side channel to detect blocking from off-path.

Goal: Scalable, ethical, and statistically robust system to continuously detect TCP/IP disruption

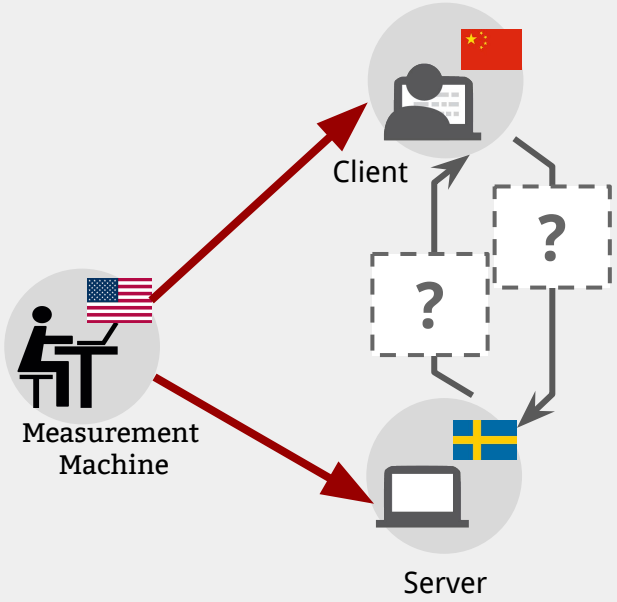

* **Internet-Wide Detection of Connectivity Disruptions**
P. Pearce*, R. Ensafi*, F. Li, N. Feamster, V. Paxson *joint first authors
IEEE S&P ("Oakland") 2017

Censorship Can Happen on Any Layer

CHALLENGE: Design methods to detect interference remotely at all network layers, without end-user participation.

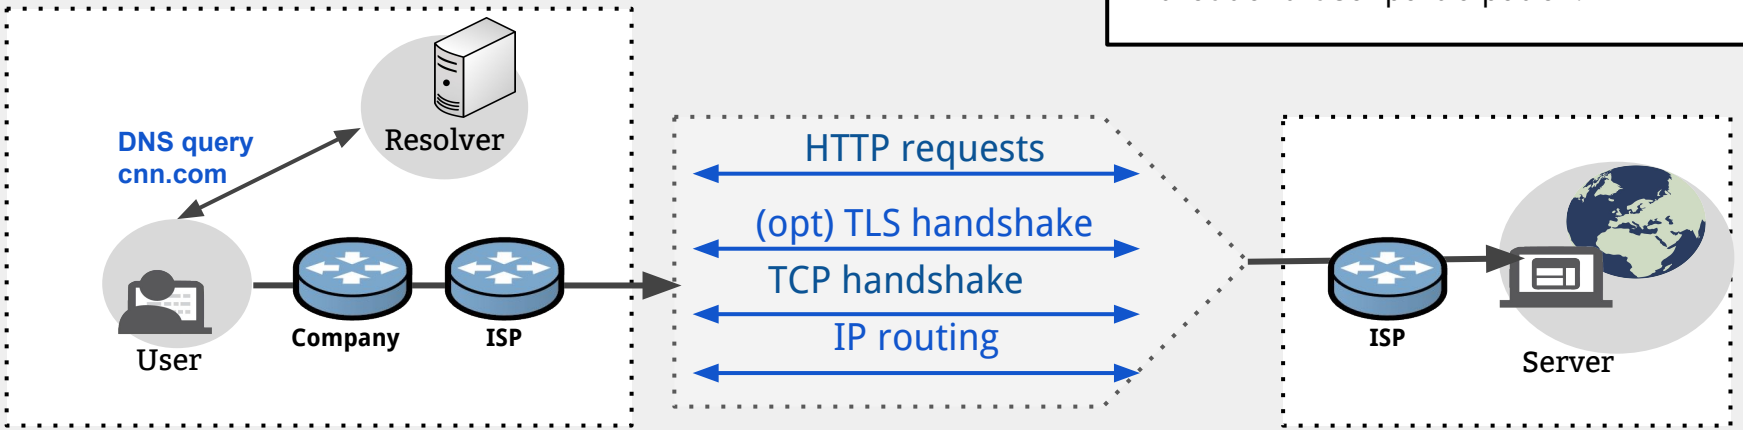

Techniques for disruptions:

- Internet shutdown (IODA)
- IP address blacklisting
- RST injection
- SNI blocking
- HTTP keyword filtering

Remote Way to Detect DNS-Layer Manipulation

PROBLEM:

How can we detect whether DNS queries are being modified anywhere around the world?

... without volunteer participation?

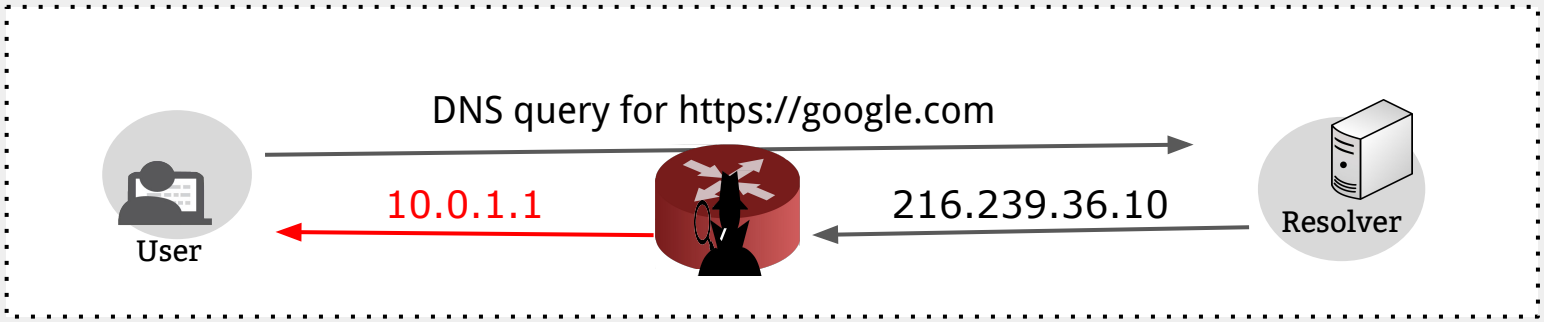

Satellite

Satellite* is a system that uses organizational open DNS resolvers to detect whether a user can resolve a domain correctly

Goal: Scalable, ethical, and statistically robust system to continuously detect DNS level manipulation

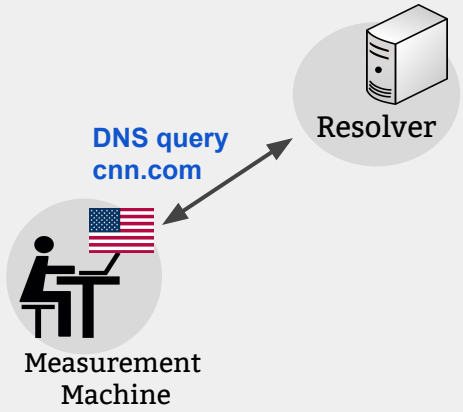

* **Satellite: Joint Analysis of CDNs and Network-Level Interference,**
  W. Scott, T. Anderson, Y. Kohno, and A. Krishnamurthy.
  In USENIX ATC, 2016.
* **Global Measurement of DNS Manipulation,**
  P. Pearce, B. Jones, F. Li, R. Ensafi, N. Feamster, V. Paxson
  USENIX Security, August 2017

* NOTE: Our deployed system benefits from both research papers, for simplicity, we use Satellite because of it seniority

Deploying Satellite

CHALLENGE:

Identify “wrong” DNS responses

THREE KEY CHALLENGES:
Coverage, continuity, and ethics

Coverage:

- Scan IPv4 for open resolvers: 4.2 M, 232 countries
- Heavy rate limit queries to resolvers and domains

Continuity:

- Satellite doesn’t depend on end users’ availability, and resolvers have less downtime

Ethics:

- Using resolvers reasonably attributed to Internet naming infrastructures: they can be resolvers with a valid PTR record beginning with the subdomain ns[0-9]+ or nameserver[0-9]*-->14k
-

Censorship Can Happen on Any Layer

CHALLENGE: Design methods to detect interference remotely at all network layers, without end-user participation.

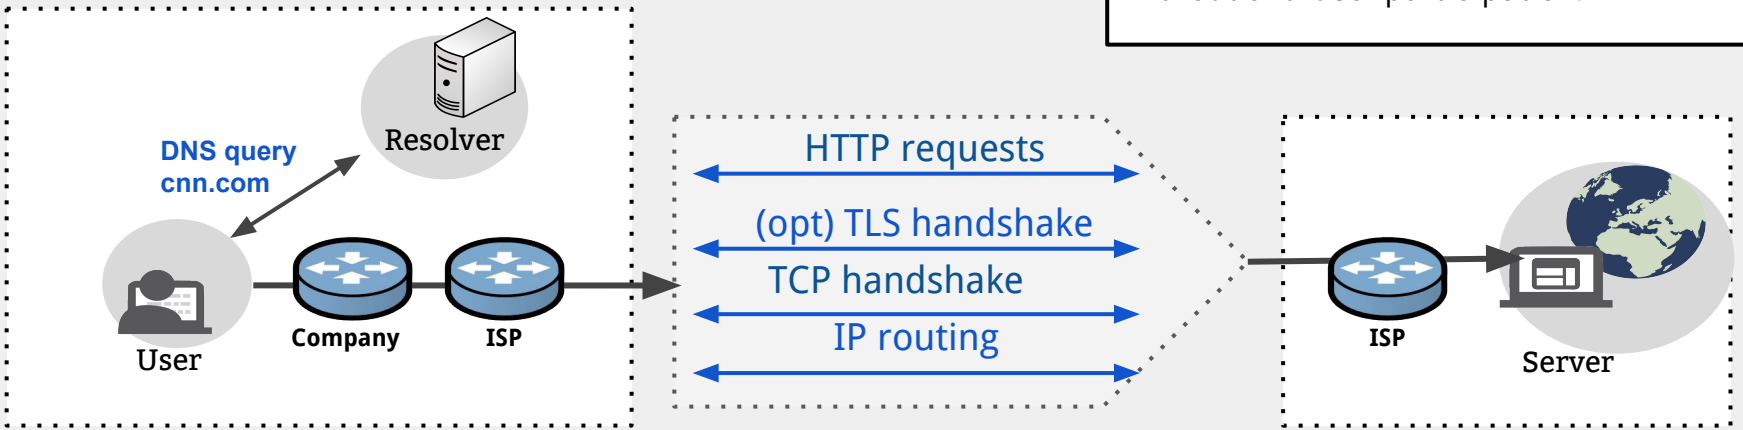

Techniques for disruptions:

- Internet shutdown (IODA)
- IP address blacklisting
- RST injection
- SNI blocking
- HTTP keyword filtering

Side Channel to Detect Application-Layer Blocking

PROBLEM:

How can we detect **keywords/URLs** are blocked?

… without volunteer participation?

Echo Protocol to the Rescue!

Using the Echo Protocol:

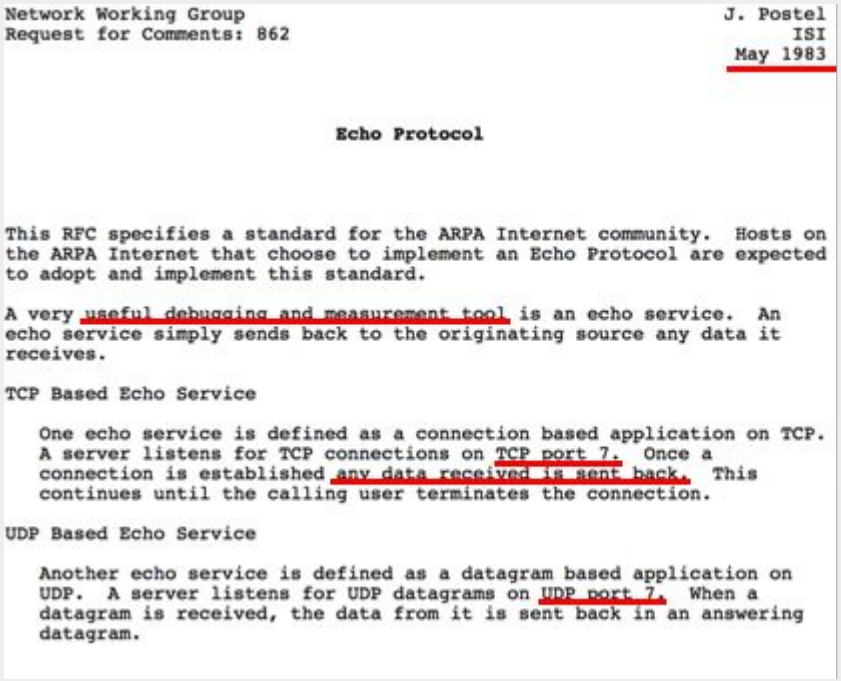

Network Working Group
Request for Comments: 862

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

Echo Protocol to the Rescue!

Using the Echo Protocol:

- An Echo service simply sends back to the originating source any data it receives.

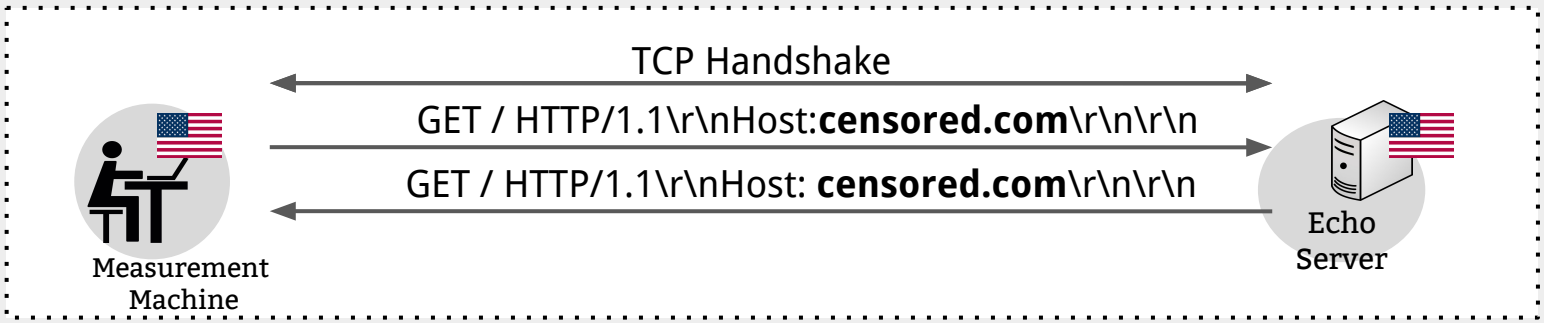

Echo Protocol to the Rescue!

Using the Echo Protocol:

- An Echo service simply sends back to the originating source any data it receives.

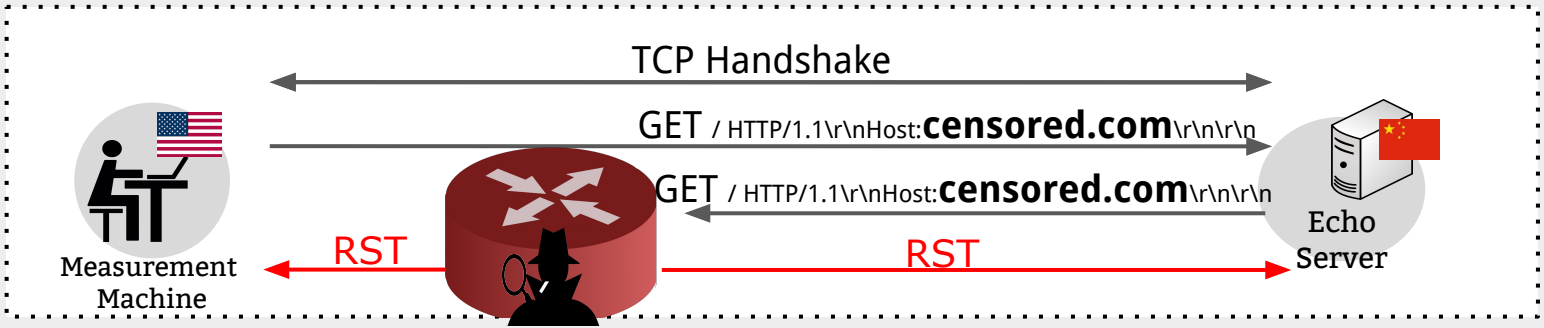

Quack

Quack is a system that uses Echo servers to detect whether keywords/URLs are blocked

Goal: Scalable, ethical, and statistically robust system to continuously detect application-layer blocking

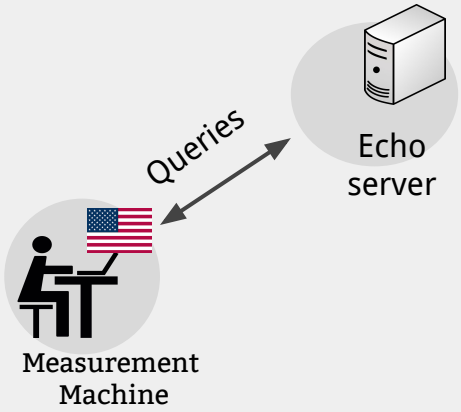

* **Quack: Scalable Remote Measurement of Application-Layer Censorship,**
VanderSloot, McDonald, Scott, Halderman, Ensafi.
USENIX Security, August 2018

Deploying Quack

CHALLENGE:

Attributing Echo servers to Internet infrastructures is tricky!

THREE KEY CHALLENGES:
Coverage, continuity, and ethics

Coverage:

- Scan IPv4 for Echo servers: 47k , 167 countries

Continuity:

- Quack doesn’t depend on end users’ availability, and Echo servers have less downtime

Ethics:

- Using Echo servers reasonably attributed to Internet infrastructures

Techniques for Remotely Measuring Interference

TCP/IP Layer →

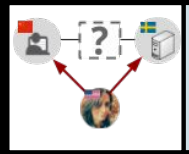

Spooky/Augur (2014-17) → Global IP_ID routers

DNS Layer →

Satellite (2016-2017) → Institutional open resolvers

Application Layer →

Quack (2018) → Services that reflect data (e.g. Echo)

The Vision

“Censorship weather map” to continually monitor Internet censorship around the world

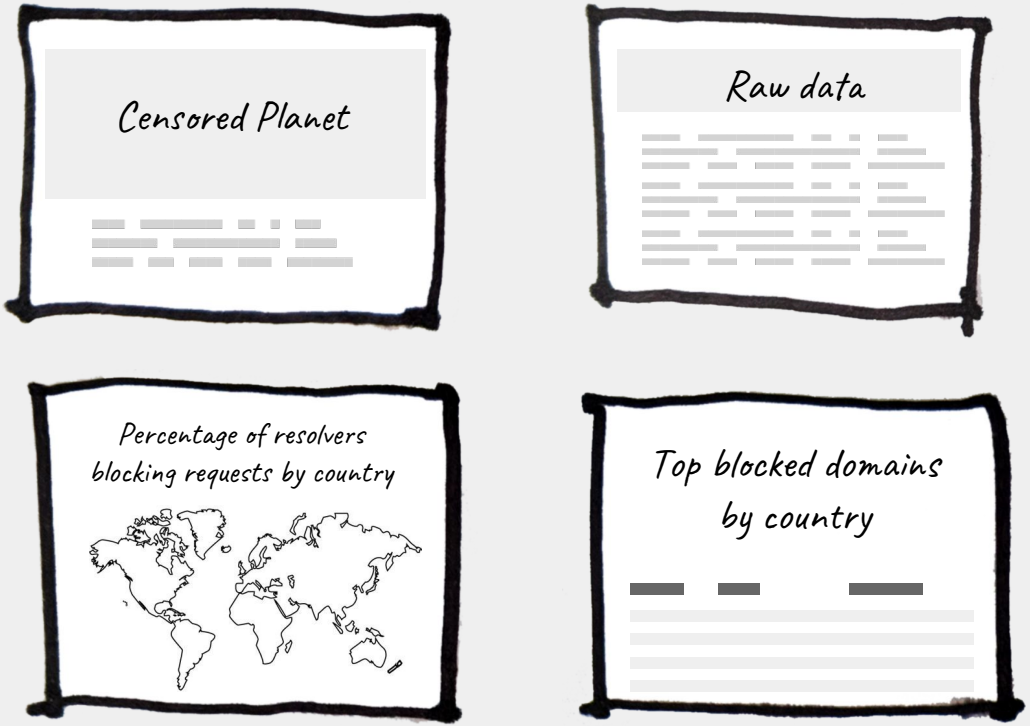

Reality

Censored Planet

A platform for continuously monitoring global Internet censorship

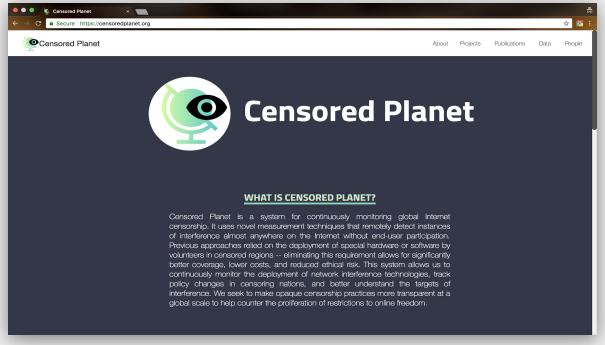

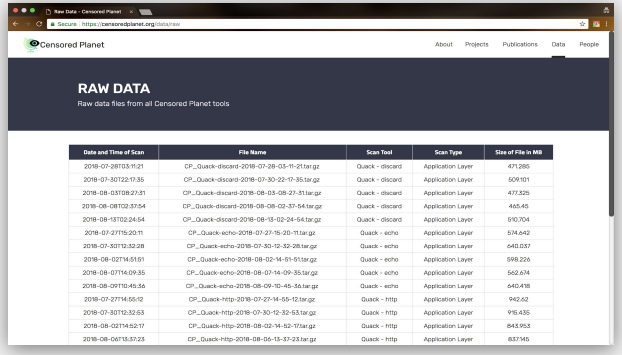

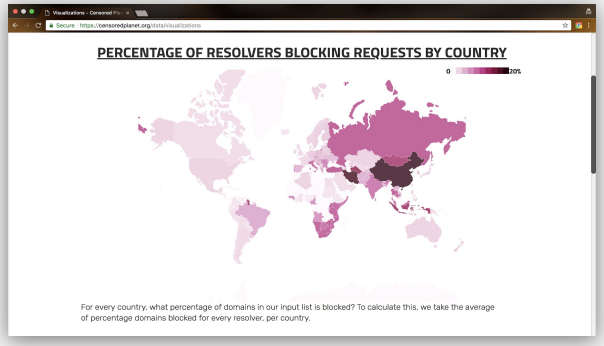

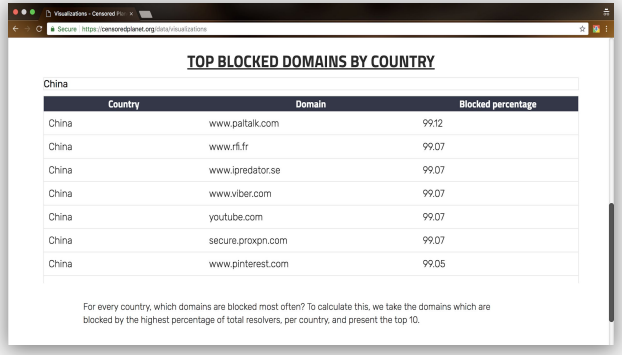

Special thanks to my amazing students and collaborators who worked extremely hard to launch this project in August.

What can Censored Planet Data Reveal?

Global, continuous data lets us **watch how censors react** to major political events

Jamal Khashoggi's disappearance and killing widely reported by world media in October 2018

Censored Planet tests reachability from **214 vantage points in Saudi Arabia** every week

In mid-October, Saudi Arabia began **blocking more than twice as many news sites** we test than prior to Khashoggi's death

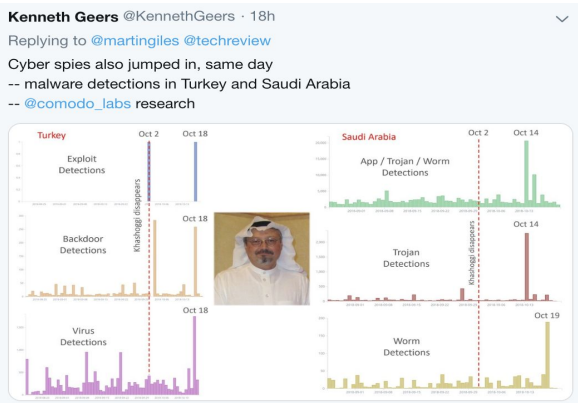

Censored Planet's Future Plan

+ OONI

Side Channels are unable to **replicate the full level of detail** of dedicated local vantage points.

→ **Integrate remote and local measurements to provide the best of both worlds**

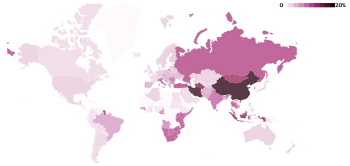

Developing visualization, statistical tools to automate **spotting patterns and trends**.

→ **develop the empirical science of understanding Internet censorship**

Censored Planet is looking for excited and dedicated

engineer & political science researcher,

if you are interested, come talk to me!

Detecting Network Interference with Side Channels

Quack: Scalable Remote Measurement of Application-Layer Censorship
B. VanderSloot, A. McDonald, W. Scott, J. A. Halderman, **R. Ensafi**
USENIX Security 2018

Internet-Wide Detection of Connectivity Disruptions
P. Pearce*, **R. Ensafi***, F. Li, N. Feamster, V. Paxson *joint first authors
IEEE S&P ("Oakland") 2017
Invited to appear in the IEEE Security & Privacy Magazine

Global Measurement of DNS Manipulation
P. Pearce, B. Jones, F. Li, **R. Ensafi**, N. Feamster, V. Paxson
USENIX Security 2017
Invited to appear in USENIX ;login:, Winter 2017 Issue

Analyzing the Great Firewall of China Over Space and Time
R. Ensafi, P. Winter, M. Abdullah, J. Crandall
Privacy Enhancing Technologies Symposium (PETS), 2015

Detecting Intentional Packet Drops on the Internet via TCP/IP Side Channels
R. Ensafi, J. Knockel, G. Alexander, J. Crandall
Passive and Active Measurement (PAM), 2014

Idle Scanning and Non-interference Analysis of Network Protocol Stacks Using Model Checking
R. Ensafi, J. Park, D. Kapur, J. Crandall
USENIX Security 2010

Censored Planet:
Global Censorship Observatory

Roya Ensafi
University of Michigan
Dec 27,2018

Ethics in Censorship Measurement

More generally, censorship research frequently raises ethical considerations.

E.g., under what conditions is it safe enough to use remote vantage points?

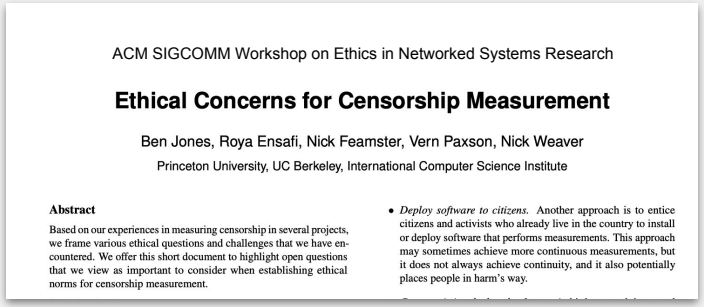

ACM SIGCOMM Workshop on Ethics in Networked Systems Research

Ethical Concerns for Censorship Measurement

Ben Jones, Roya Ensafi, Nick Feamster, Vern Paxson, Nick Weaver

Princeton University, UC Berkeley, International Computer Science Institute

Abstract

Based on our experiences in measuring censorship in several projects, we frame various ethical questions and challenges that we have encountered. We offer this short document to highlight open questions that we view as important to consider when establishing ethical norms for censorship measurement.

- *Deploy software to citizens.* Another approach is to entice citizens and activists who already live in the country to install or deploy software that performs measurements. This approach may sometimes achieve more continuous measurements, but it does not always achieve continuity, and it also potentially places people in harm's way.

IRBs are often not positioned to help.

> Common Rule (45 CFR 46.102(f)) defines a human subject as "a living individual about whom an investigator conducting research obtains (1) data through intervention or interaction with the individual or (2) identifiable private information."

We turn to authorities such as the **Belmont and Menlo Reports** to guide ethical thinking.

Frequently consult with colleagues to check our reasoning and conclusions.

Questions we regularly consider include:

- **What populations of users are affected?**
- **Is informed consent feasible?**
- **Have we considered all anticipatable risks?**
- **Do humans incur no more than minimal risk?**
- **Can we take steps to further reduce risks?**
- **Do benefits accrue to the population that is subjected to the risk?**

My Research Community

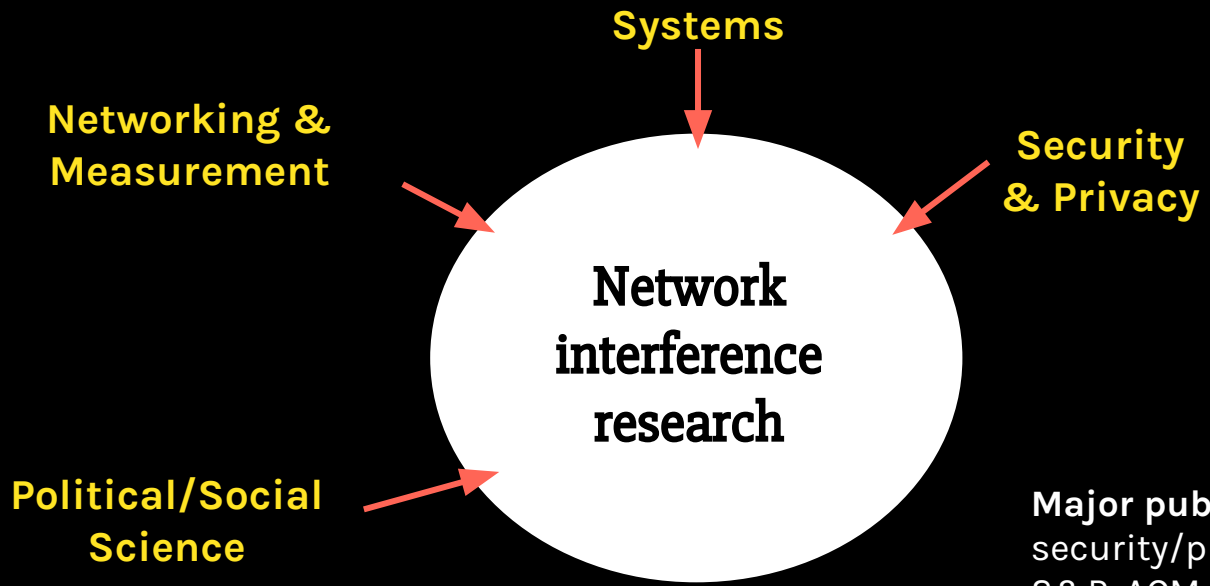

Major publication venues:
security/privacy (USENIX Sec., IEEE S&P, ACM CCS, NDSS, and PETS), Internet measurement (ACM IMC), specialized workshops (USENIX FOCI)