Ethical, Scale, and Continuity Concerns for Censorship Measurement

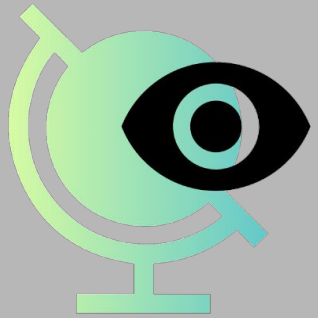

Roya Ensafi
CensoredPlanet.com

In my lab, we...

develop frameworks to
detect network interference,

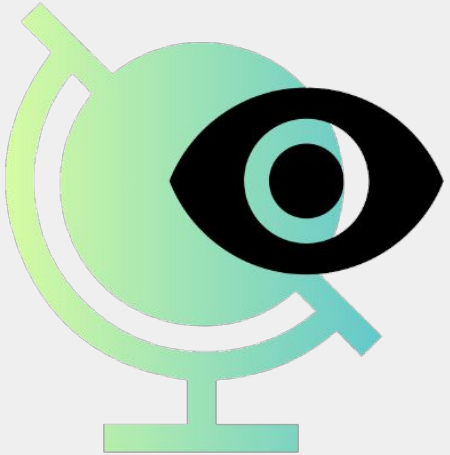

apply these frameworks to **understand the behavior** of network intermediaries,

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

My Group Draws on Diverse Intellectual Methods

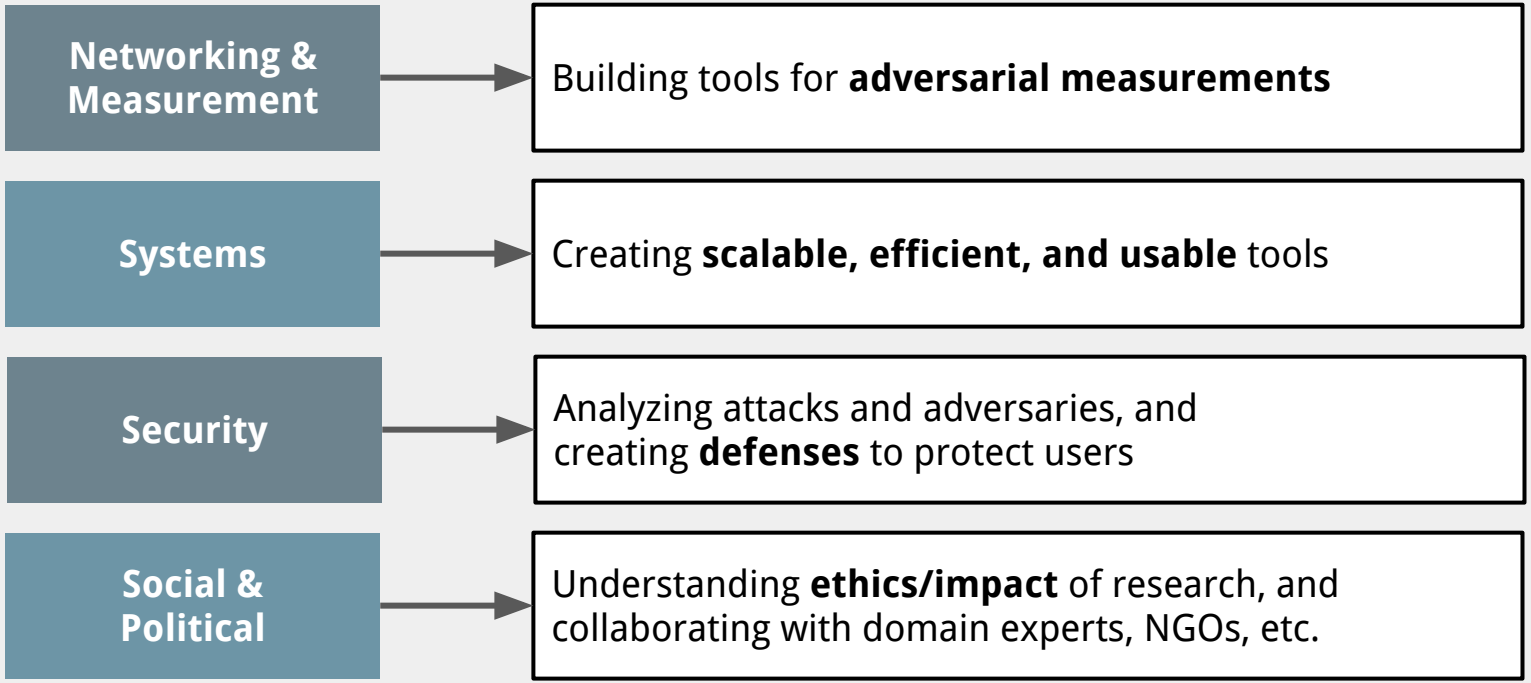

Reports suggest

Internet censorship practices

are diverse in their methods, targets, timing,

differing by regions, as well as across time.

Reports suggest

Internet censorship practices

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

Iraq govt downs Internet in response to massive anti-corruption protests, July, 2018

From Internet Intelligence Map

Russia attempts to block millions of IP addresses in battle against Telegram app

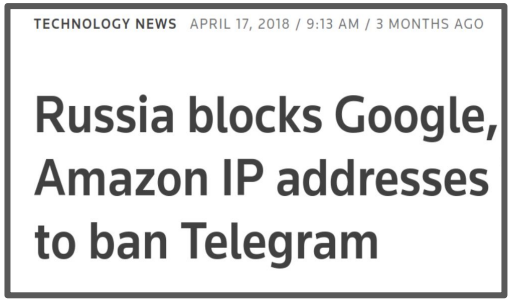

Internet Censorship: A Simplified View

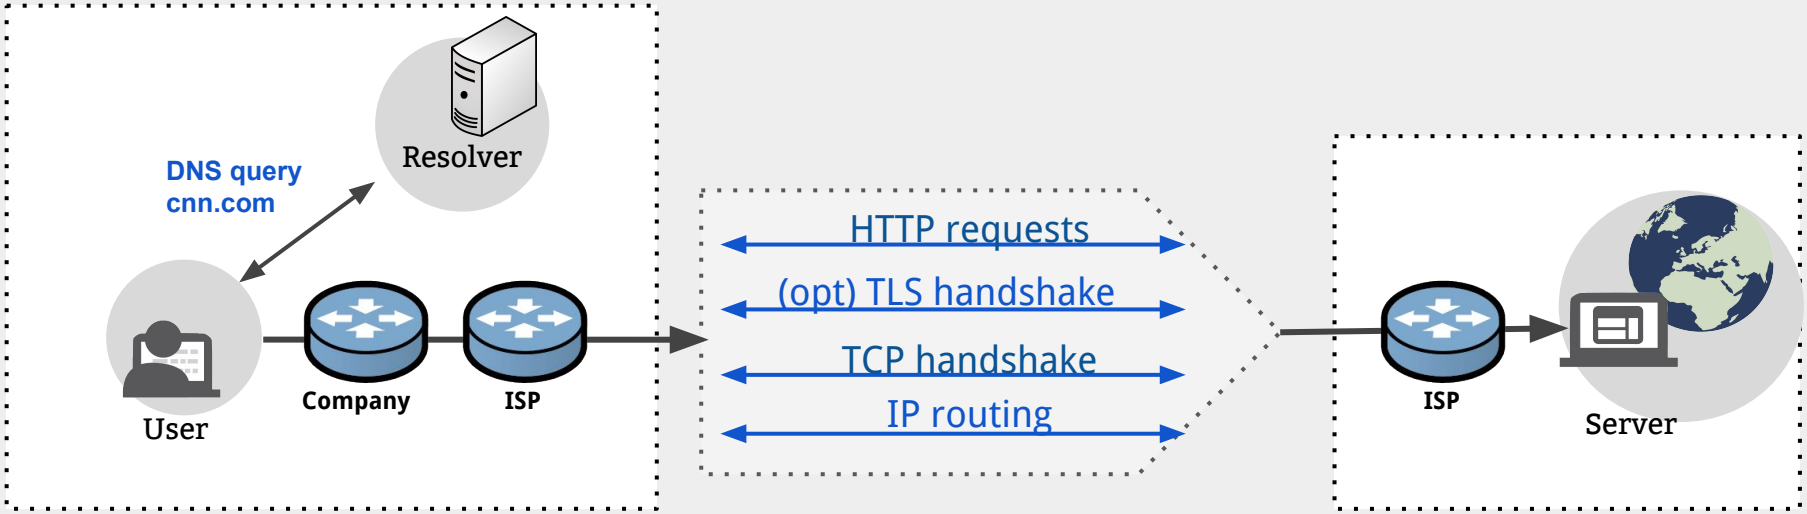

Techniques for disruptions:

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

Why Measure Internet Censorship?

- What is censored, when, for which users, by who
- Advocacy and Transparency are important
  - Inform users about what they are missing
  - Help diplomats and others who make policy decisions
- What technical mechanisms and tools (DPIs) are used
  - Can help to improve defense technology
  - GFW can cause harm → Great Cannon [*]
- Why and how this blocking affects societies

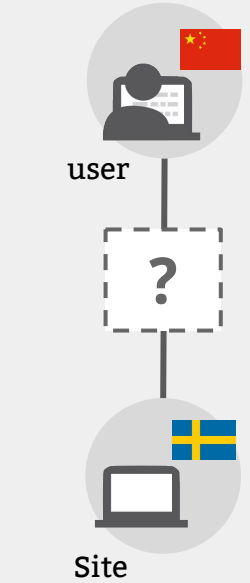

[*] **Analysis of China's "Great Cannon"**
by Marczak, Weaver, Dalek, **Ensafi**, Fifield, McKune, Rey, Scott-Railton, Deibert, and Paxson (In: *USENIX FOCI'15*)

How To Measure Internet Censorship?

PROBLEM:

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

How To Measure Internet Censorship?

PROBLEM:

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

STATE OF THE ART:

- Deploy hardware or software at hosts (RIPE Atlas, OONI probes)
- Ask people on the ground, or use VPNs, or research networks (PlanetLab)

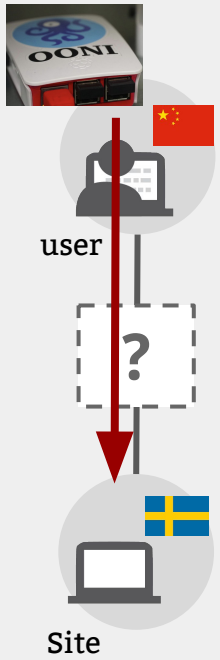

THREE KEY CHALLENGES:

Coverage, ethics, and continuity

OONI: Open Observatory of Network Interference

OONI network:

- Volunteer can install OONI Probe on iOS, Android, Linux & MacOS (web UI)

INTERPRET the DATA:

- If **Control != Experiment → Possible censorship**
- Confirm a case of censorship when they have detected a block page. [*]

Key Challenges for Longitudinal Measurement:

- **Ethics, coverage, and continuity**

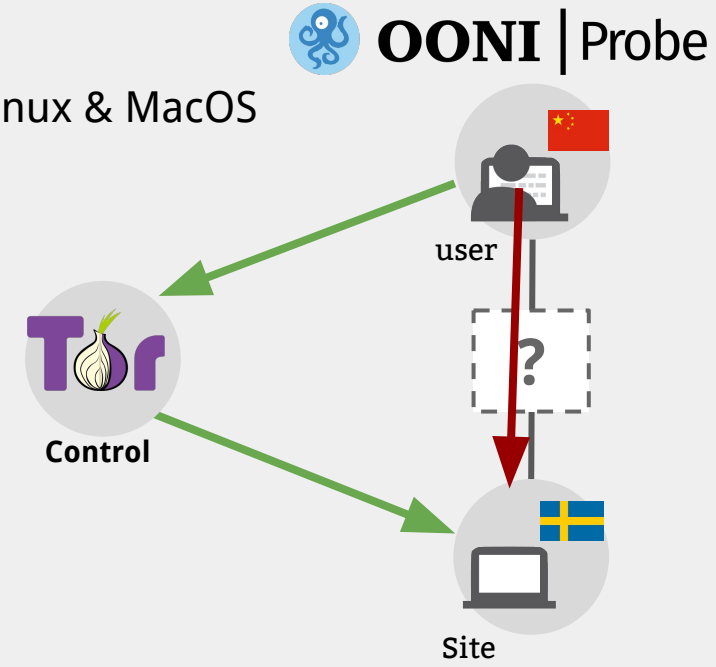

[*] from OONI overview slides

OONI: Ethical, Coverage and Continuity Challenges

EFFORT 1: detailed and honest consent form

- They explicitly say: “Anyone monitoring your internet activity (e.g. ISP) will know that you are running OONI Probe.”

The Open Observatory of Network Interference (OONI) is
collects and processes network measurements with the ai
and traffic manipulation.

Running OONI may be against the terms of service of you
OONI you will connect to web services which may be banr
as Tor. The OONI project will publish data submitted by pr
identifying information. In addition, your use of OONI will
and to anybody who can monitor your internet connection

By running ooniprobe, you are participating as a volunteer
should be aware of and consent to *prior* to running ooniprobe

OONI software tests

The OONI project has developed multiple free software te

- Detect the blocking of websites
- Detect systems responsible for censorship and traffic
- Evaluate the reachability of Tor bridges, proxies, VPN

Below we provide brief descriptions of how these tests wo

Test descriptions

The recommended set of tests that users run through the

Web connectivity: This test examines whether websites a
whether access to them is blocked through DNS tamperin
transparent HTTP proxy. It does so by identifying the reso
establish a TCP session and by sending HTTP GET reques

HTTP invalid request line: This test tries to detect the pre
be responsible for censorship and/or traffic manipulation.
sends an invalid HTTP request line - containing an invalid
request method – to an echo service listening on the stan
network, the invalid HTTP request line will be intercepted
which can help identify the proxy technologies.

HTTP header field manipulation: This test tries to detect
which could be responsible for censorship and/or traffic n
include valid, but non-canonical HTTP headers to a backe
we receive the HTTP headers exactly as we sent them, th
network. If, however, such software is present in the netw
headers that we are sending or add extra headers.

Another test which attempts to detect traffic manipulatior
packets in such a way that they perform a traceroute fron
include **Tor bridge reachability**, **Psiphon**, **Lantern**, **OpenV**
these services work within a tested network by attempting

Further test descriptions can be found here.

Choices

We provide you with choices in regards to which t
you would like to send your measurements to our

Tests

You can *opt-out* from running all of the tests inclu
and by running it/them manually. You can view hov
option.

You can run each test included in `ooniprobe` `deckgen` se

- **Web connectivity test:** `ooniprobe blocking/w`
- **HTTP header field manipulation test:** `ooniprc`
- **HTTP invalid request line test:** `ooniprobe man`

Data collection and publication

OONI software users can *opt-out* from sending OO
ooniprobe configuration](https://github.com/TheT
`~/.ooni/ooniprobe.conf` . Through this file, users
information:

- Country code
- Autonomous System Number (ASN)

By default, OONI does *not* collect users' IP addres
information) through the above configuration file.

Users can also choose to *opt-out* from sending OO
with the `-n` command line option. This option is q
measurements published, due to potential risks th

Learn more about how we handle data through ou

Consent

My consent means the following:

I understand the requirements and the risks of run

I understand that, unless I opt-out (as explained ir
default be sent to the OONI project and published

PRESS q to leave this page

Risks

Many countries have a lengthy history of subjecting digital rights activists to various forms of abuse that could make it dangerous for individuals in these countries to run OONI. The use of OONI might therefore subject users to severe civil, criminal, or extra-judicial penalties, and such sanctions can potentially include:

- Imprisonment
- Physical assaults
- Large fines
- Receiving threats
- Being placed on go
- Targeted for surveil

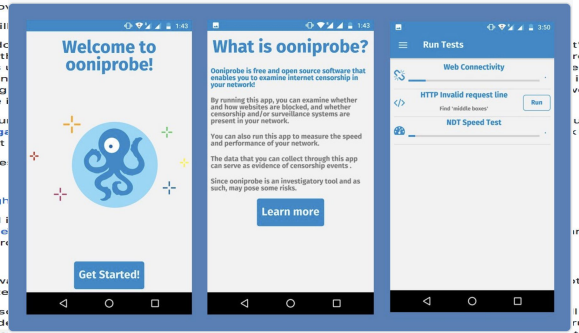

While most countries do
important to note that t
laws if, for example, its
of being criminalized on
viewed as "jeopardizing
measurement could be i

We therefore strongly u
specific inquiries at lege
advice for you or to put

Some relevant resource

- Tor Legal FAQ
- EFF Know Your Righ

Note: The use of OONI i
(https://github.com/The
organization, the Tor Pr

Installing ooniprobe

As with any other softw
access to your compute

The installation of Tor s
measurements are by d
through the `ooniprobe` command line (web connectivity test) is designed to compare HTTP requests over the network of the user and over the Tor network. Similarly, OONI's Psiphon, Lantern and OpenVPN tests require the installation of circumvention software.

We therefore encourage you to consult with a lawyer on the legality of anonymity software (such as Tor, a VPN or a proxy) *prior* to installing ooniprobe.

To remove traces of software usage, you can re-install your operating system or wipe your computer and remove everything (operating system, programs and files) from your hard drive.

Running ooniprobe

Third parties (such as your government, ISP and/or employer) monitoring your internet activity will be able to see all web traffic generated by OONI, including your IP address, and might be able to link it to you personally.

OONI: Ethical, Coverage and Continuity Challenges

EFFORT 1: detailed and honest consent form

EFFORT 2: established close relationships with locals and civil society

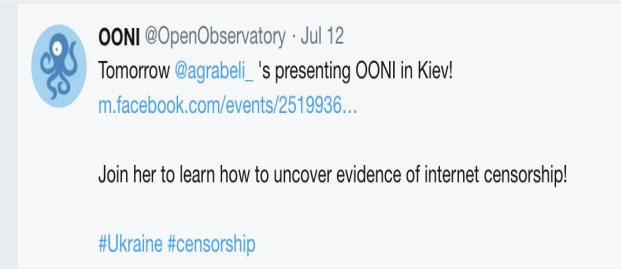

OONI: Ethical, Coverage and Continuity Challenges

OONI

EFFORT 3: Keep the community of volunteer involved

EFFORT 4: Dedicated Focus and Open Source Pledge

EFFORT 5: Greate Funders including Open Technology Fund (OTF), M-lab, FREE PRESS,...

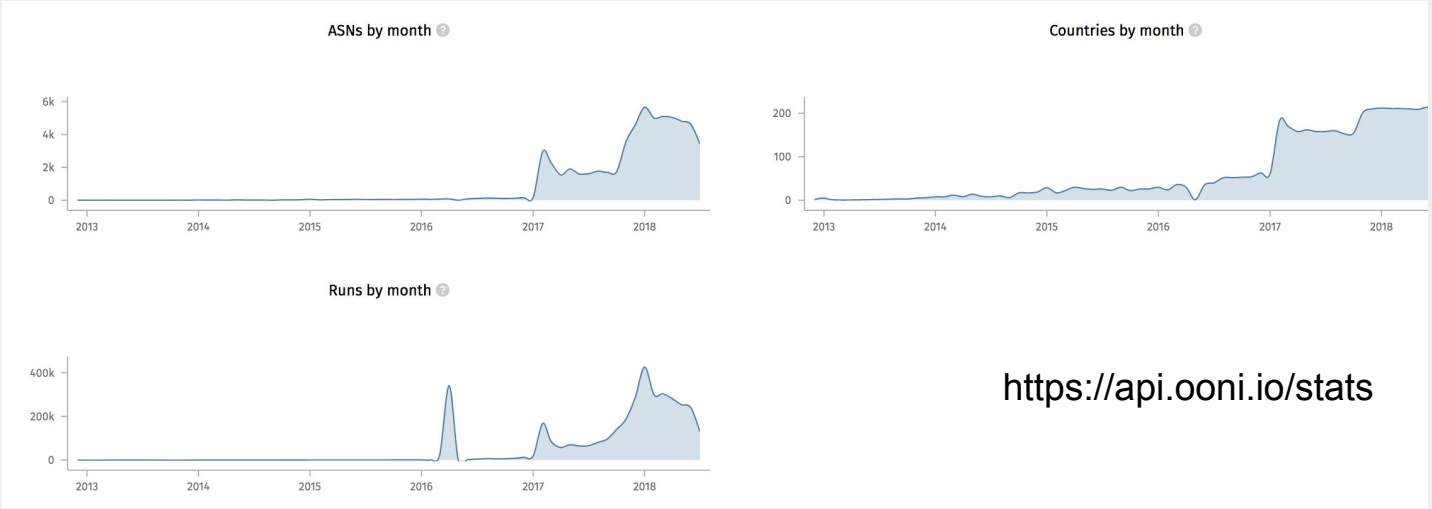

https://api.ooni.io/stats

Thinking Like an “Attacker”…

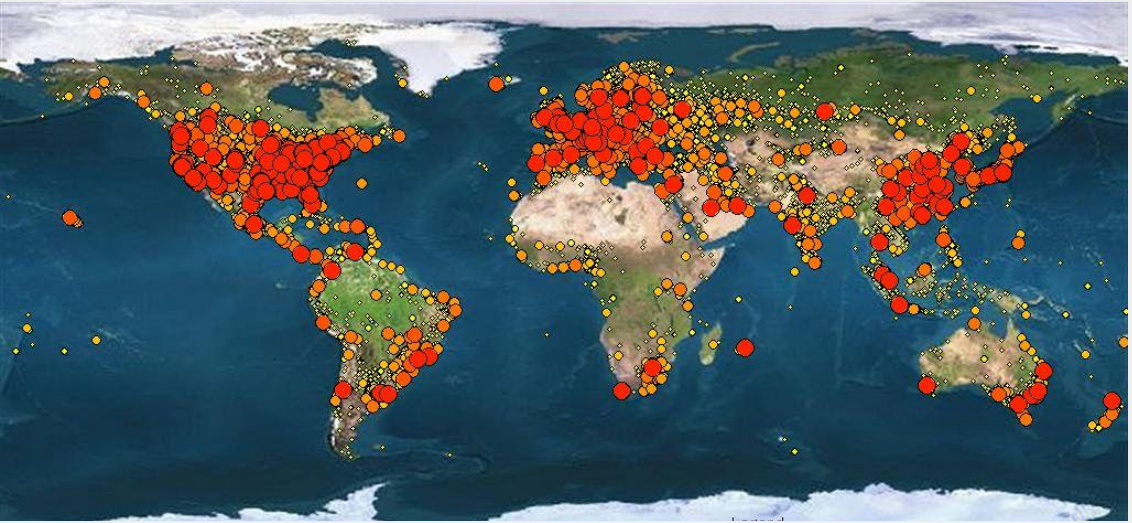

140 million public live IPv4 addresses

These machines blindly follow Internet protocol rules such as TCP/IP.

How can we leverage standard protocol behaviors to detect whether two distant hosts can communicate?

Measuring Internet Censorship Globally… Remotely!

PROBLEM:

- How can we detect whether pairs of hosts around the world can talk to each other from somewhere else in the world

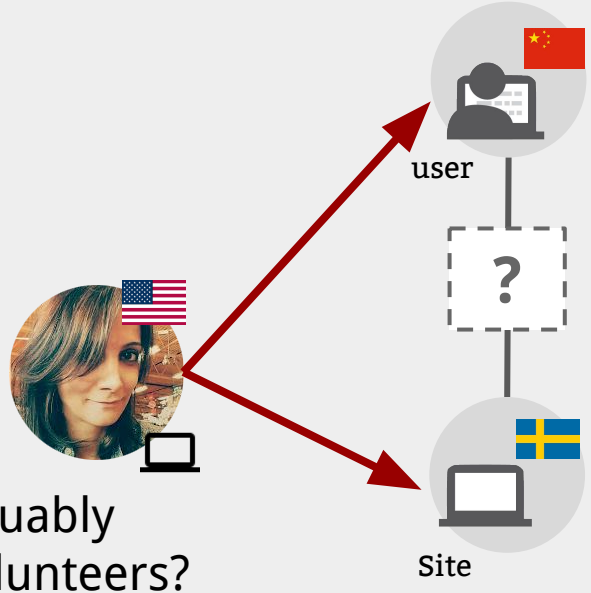

- Can we identify intermediary machines that arguably constitute “infrastructure” to reduce risk for volunteers?

Spooky Scan

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

Goal: **Detect blocking from off-path**

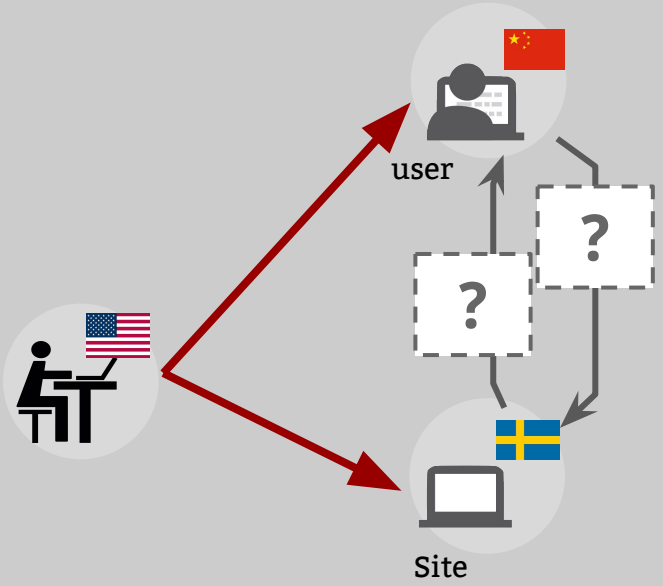

* **TCP Idle Scan** Antirez, (Bugtraq 1998)
* **Detecting Intentional Packet Drops on the Internet via TCP/IP Side Channels**
  Roya Ensafi, Knockel, Alexander, and Crandall (PAM '14)
* **Idle Port Scanning and Non-interference Analysis of Network Protocol Stacks Using Model Checking**
  Roya Ensafi,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

Augur

Augur is a follow up system that uses the same TCP/IP side channels to detect blocking from off-path.

Goal: Scalable, ethical, and statistically robust system to continuously detect blocking.

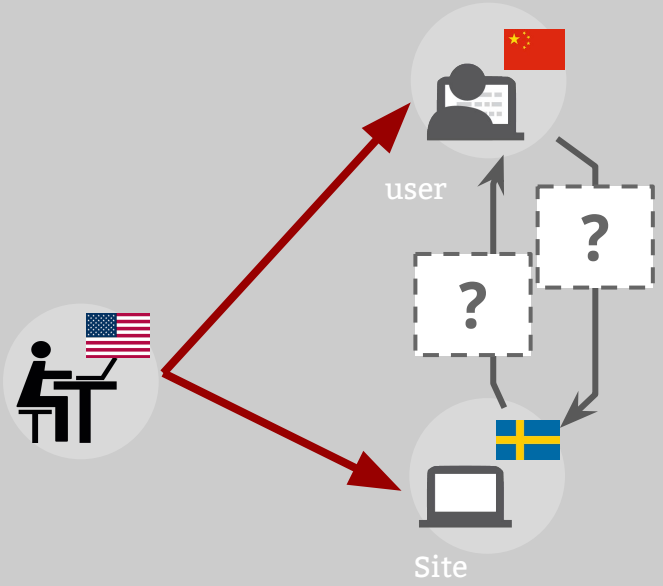

* Augur: Internet-Wide Detection of Connectivity Disruption
P. Pearce*, R. Ensafi*, F. Li, N. Feamster, V. Paxson
(* joint first authors)

TCP/IP

TCP Handshake:

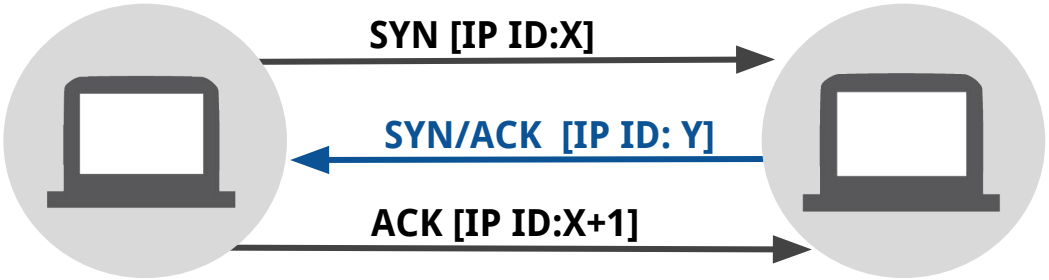

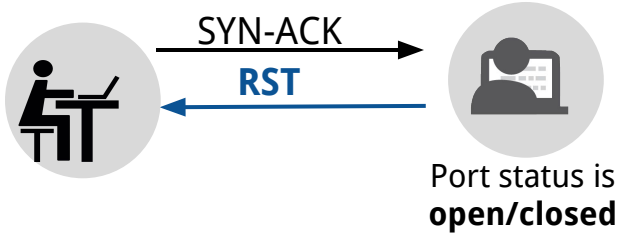

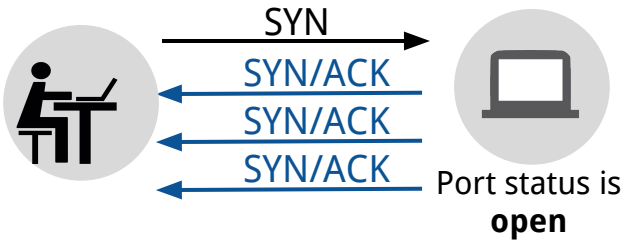

Spooky Scan Requirements

“User” (Reflector)

Must maintain a
<u>global</u> value for IP ID

Site

Open port and
retransmitting SYN-ACKs

Measurement Machine

Must be able to spoof packets

Spooky Scan

Measurement machine

Reflector

Reflector IP ID

Site

Spooky Scan

No direction blocked

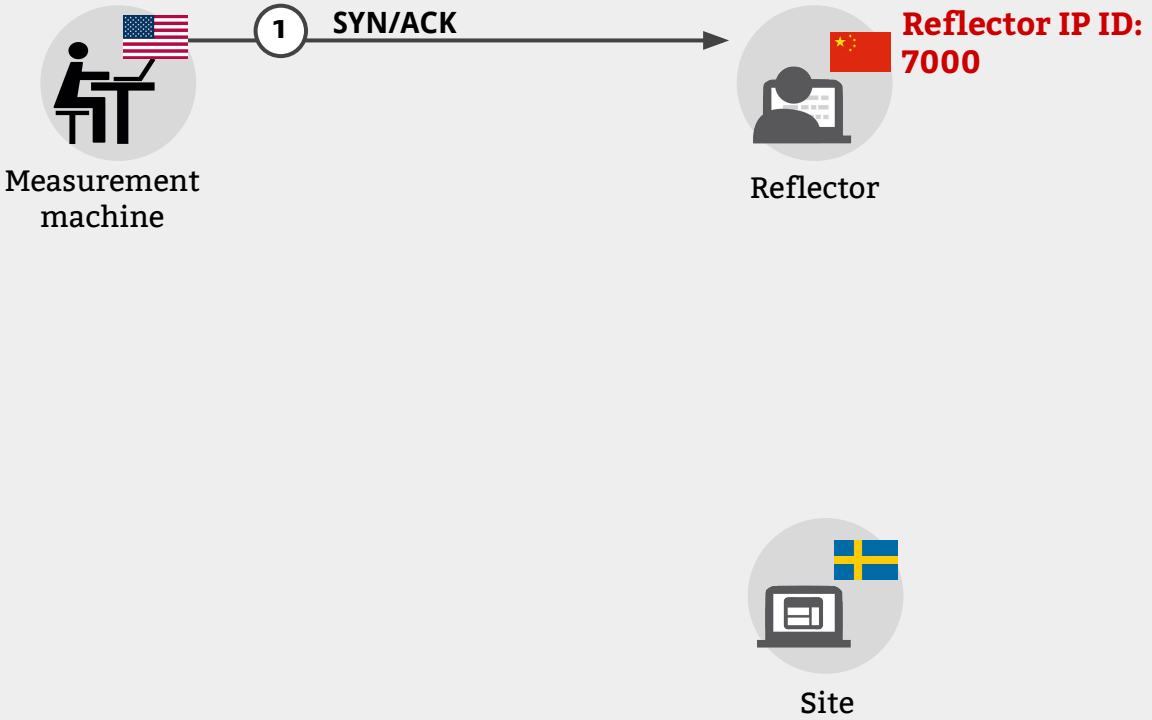

Spooky Scan

No direction blocked

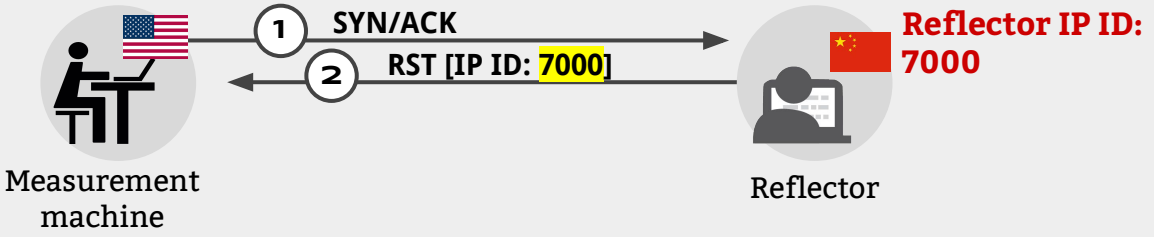

Spooky Scan

No direction blocked

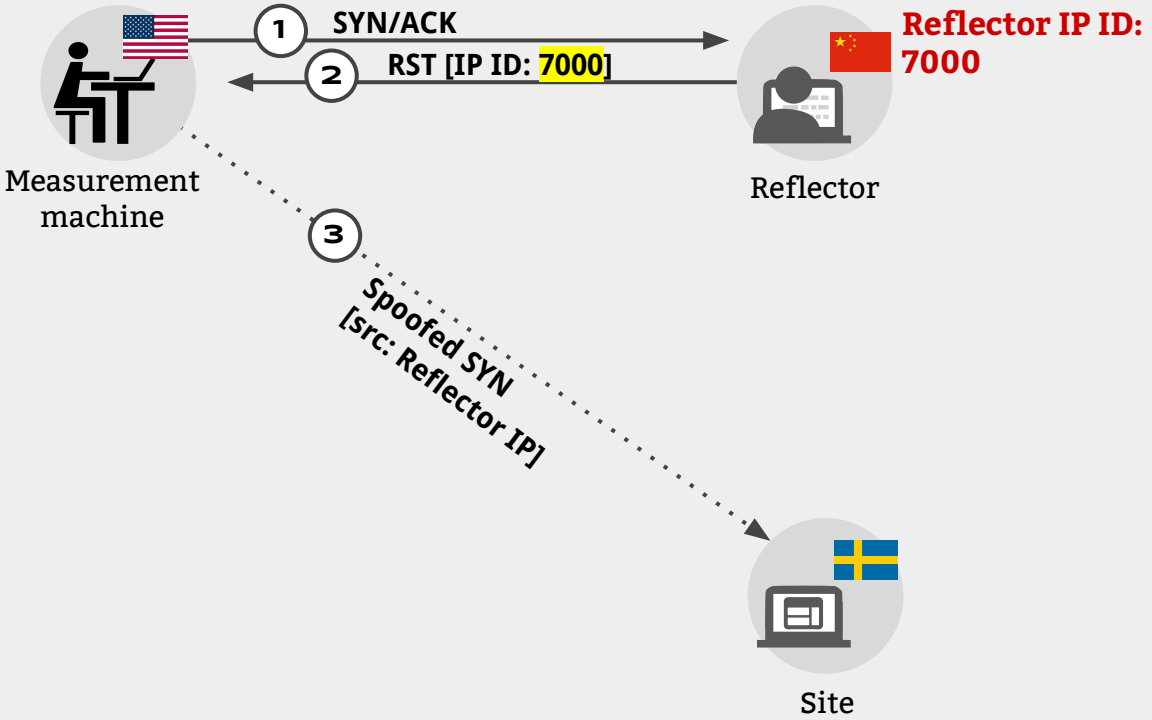

Spooky Scan

No direction blocked

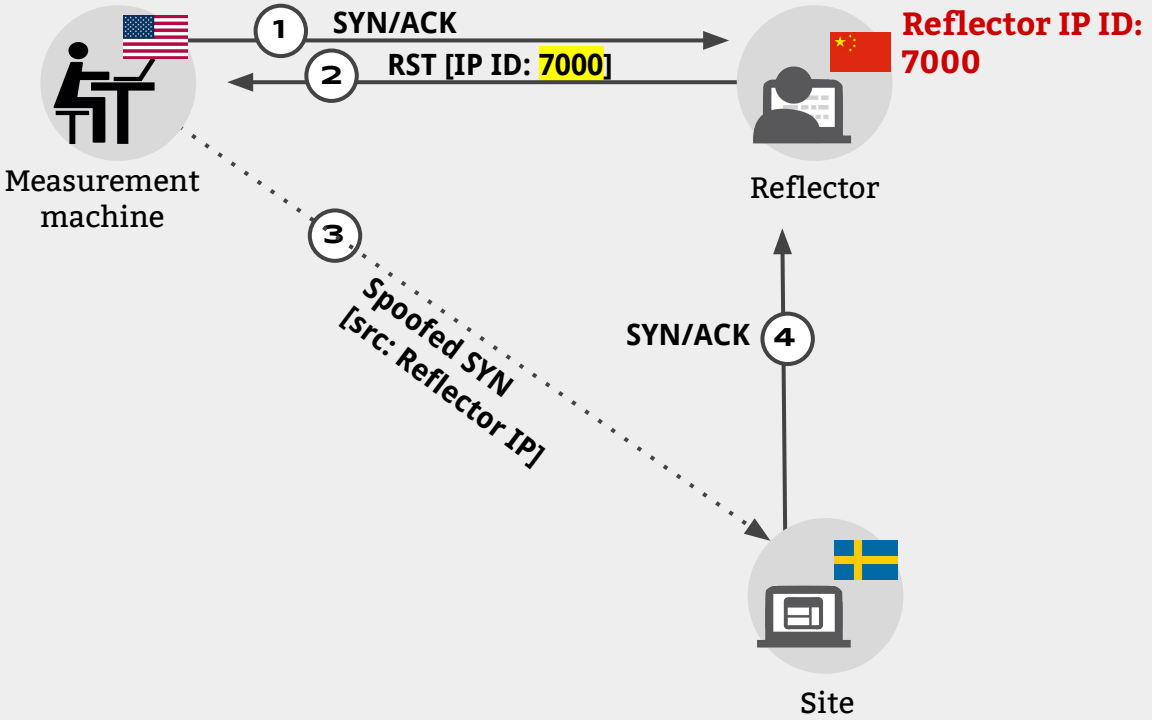

Spooky Scan

No direction blocked

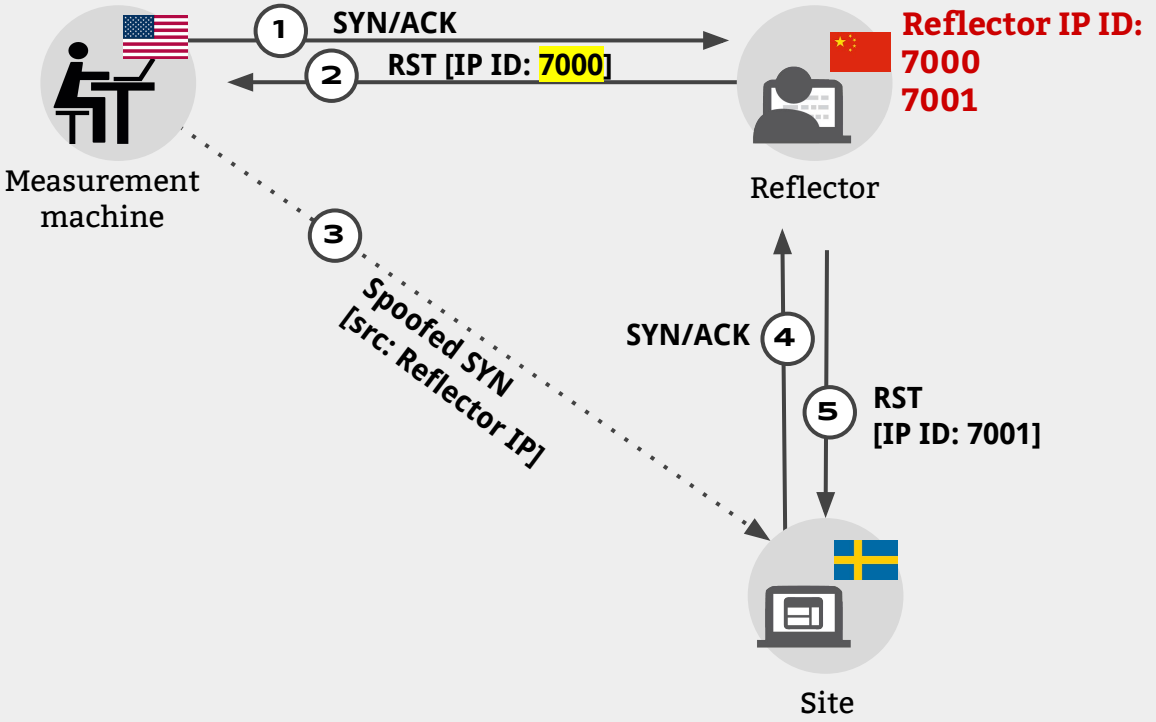

Spooky Scan

No direction blocked

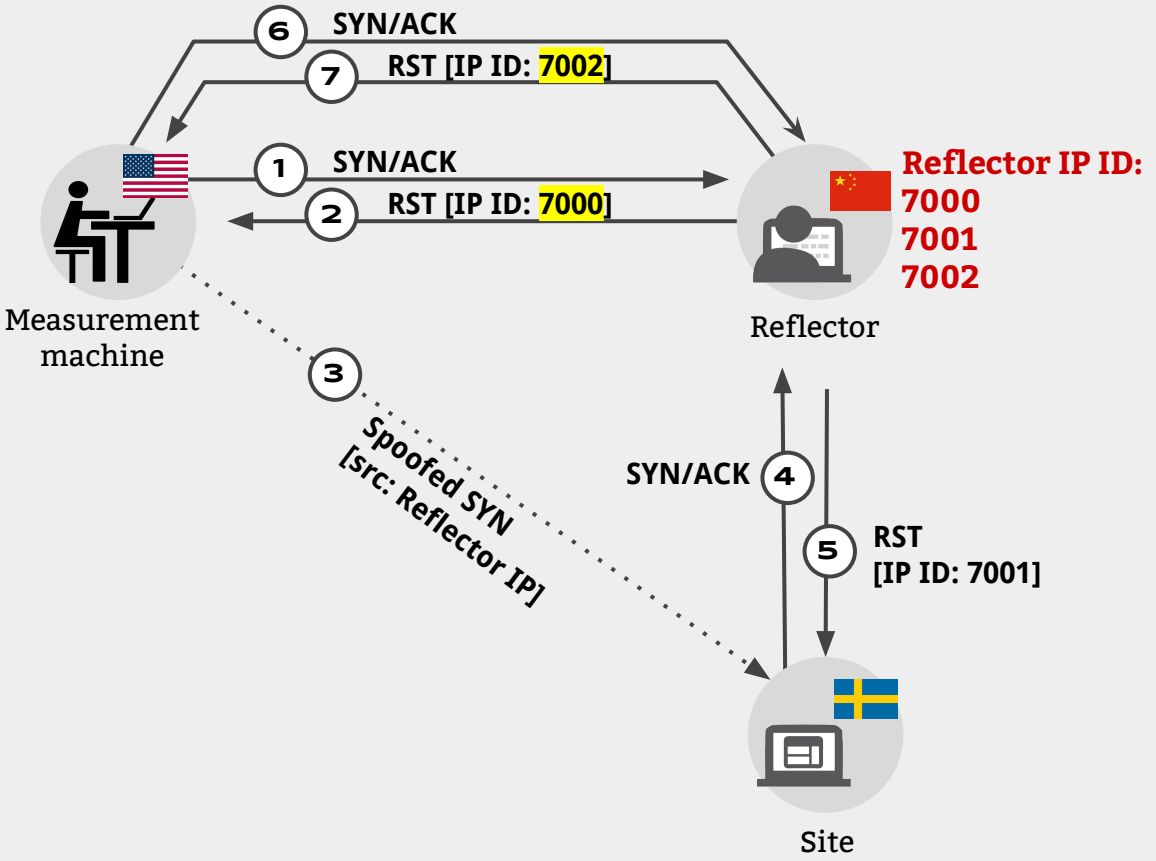

Spooky Scan

No direction blocked

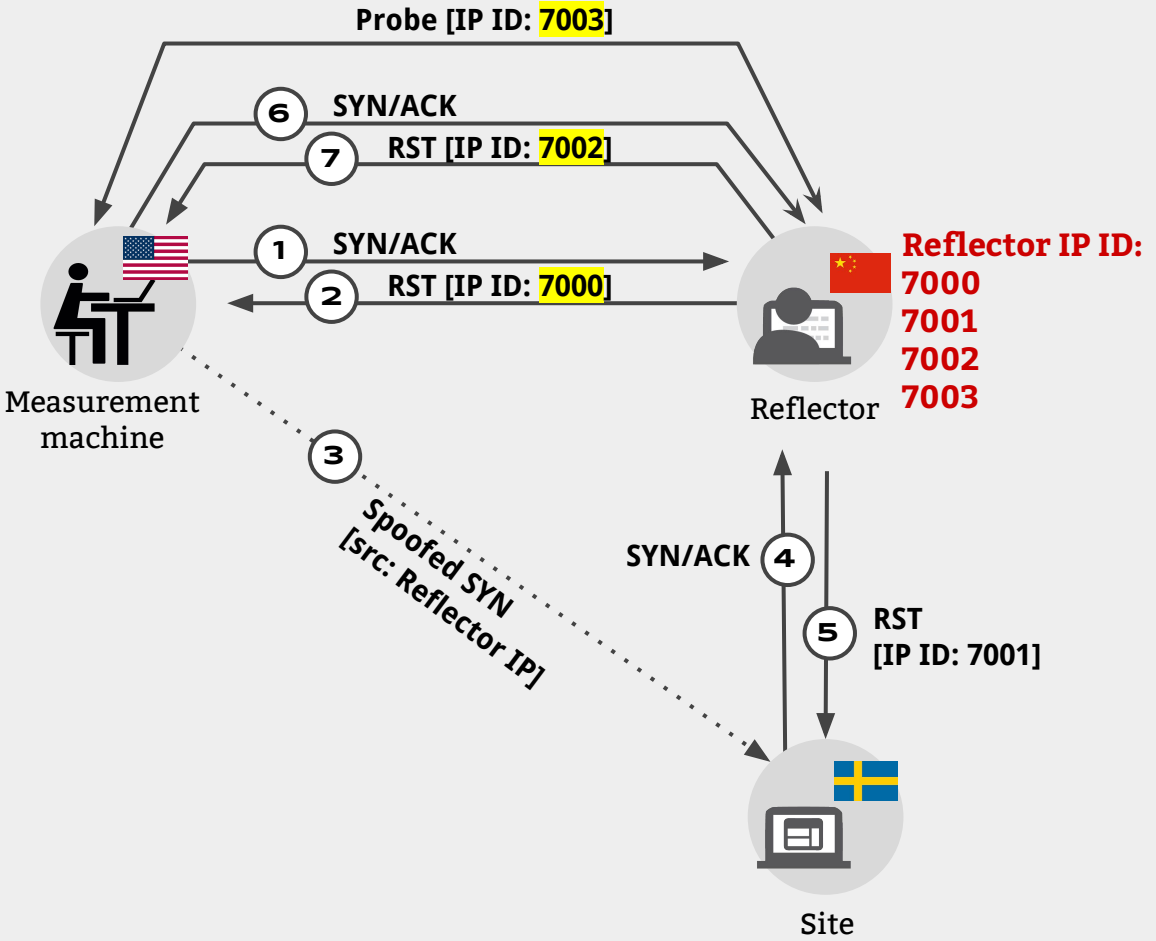

Spooky Scan

Site-to-Reflector
Blocked

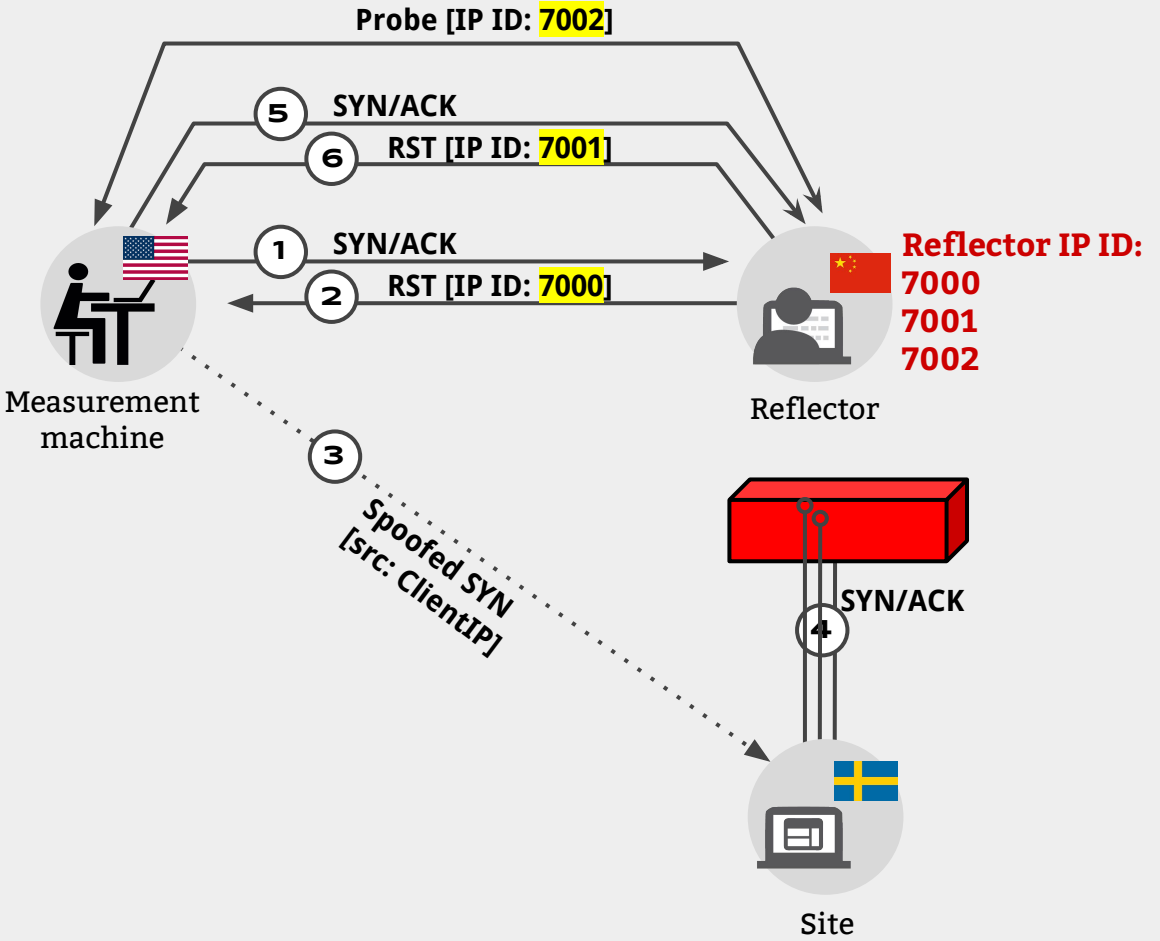

Spooky Scan

Reflector-to-Site
Blocked

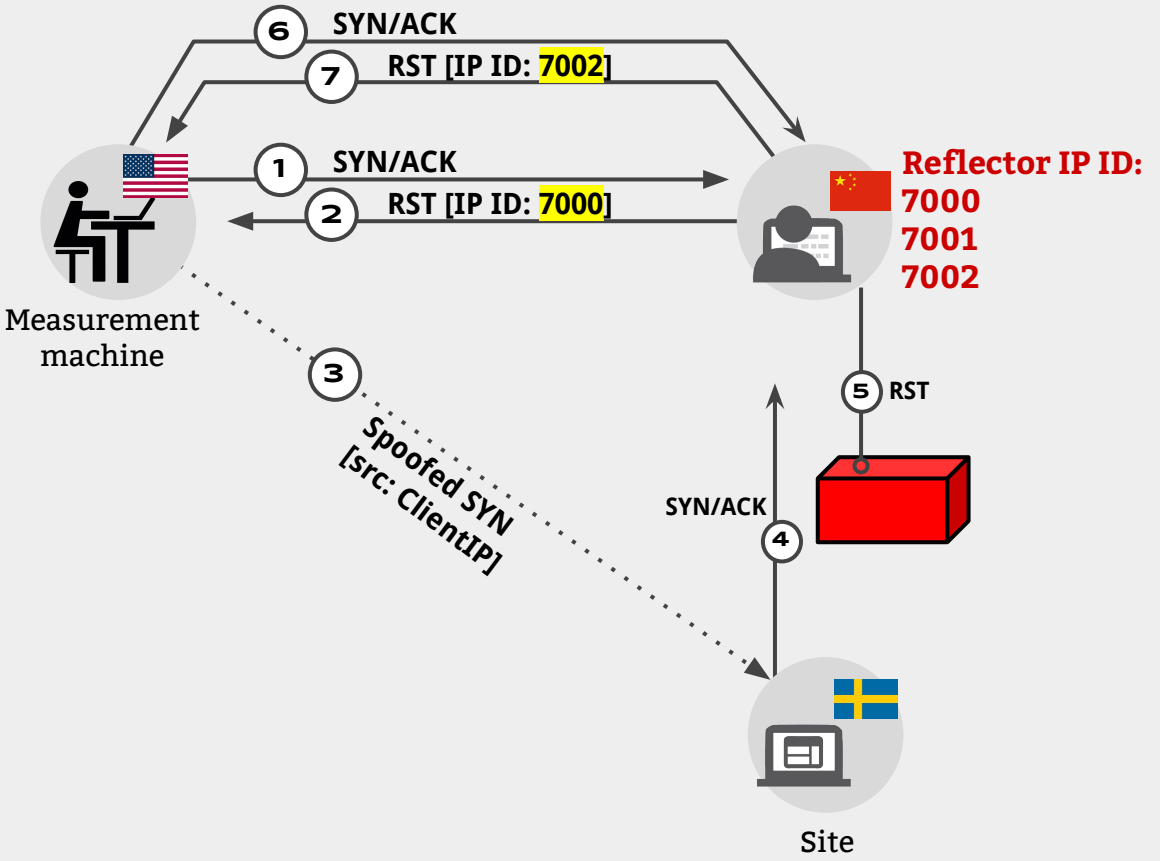

Spooky Scan

Reflector-to-Site
Blocked

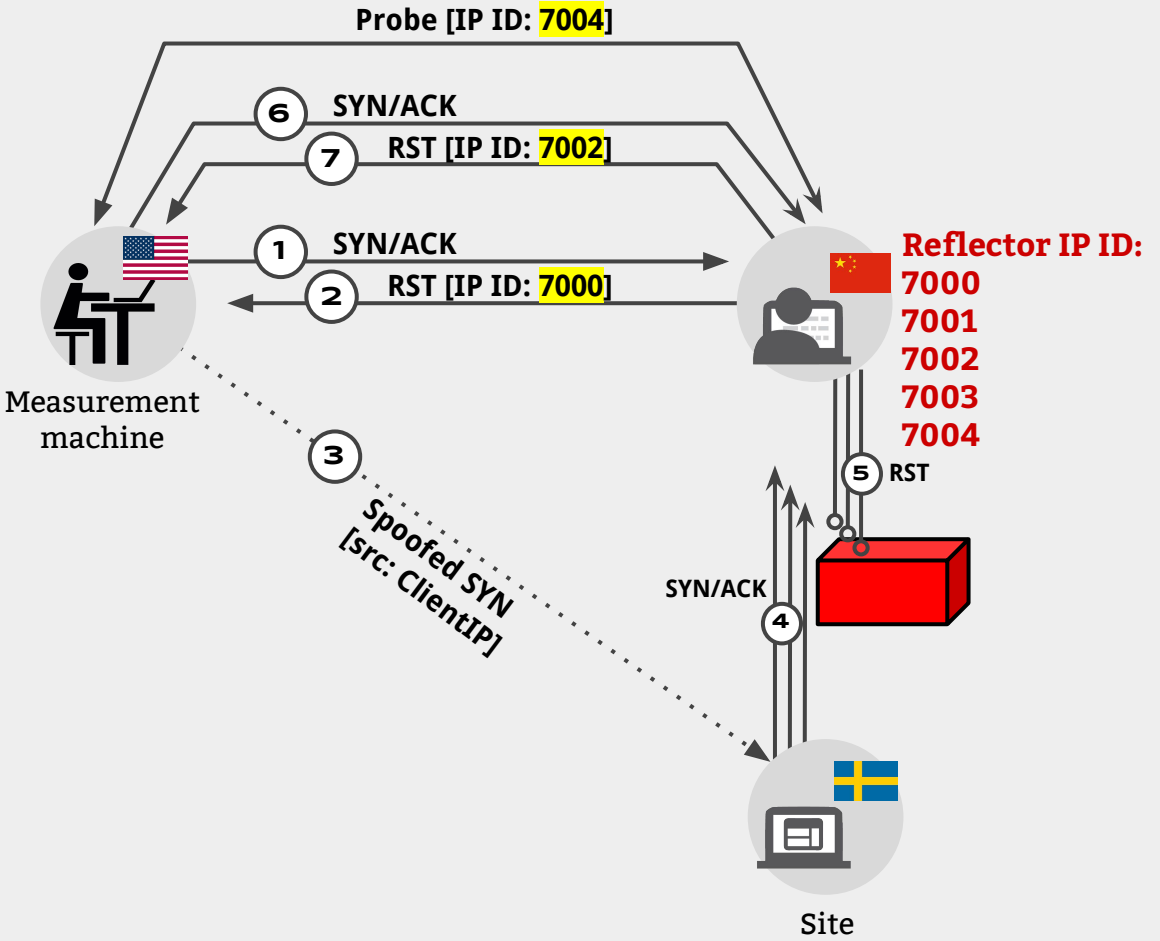

Spooky Scan

Site-to-Reflector Blocked

Δ IP ID1 = 1
Δ IP ID2 = 1

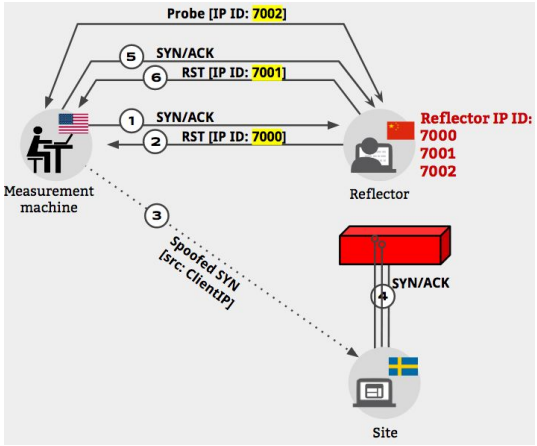

No Direction Blocked

Δ IP ID1 = 2
Δ IP ID2 = 1

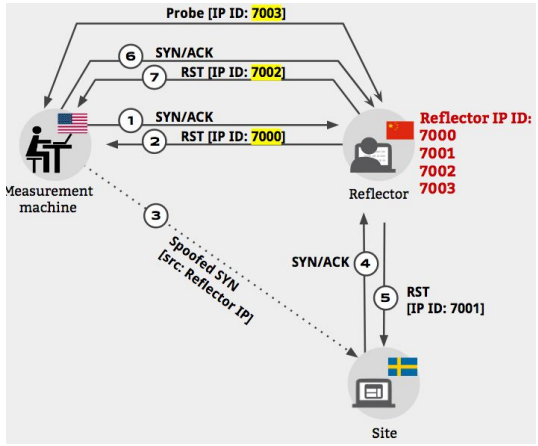

Reflector-to-Site Blocked

Δ IP ID1 = 2
Δ IP ID2 = 2

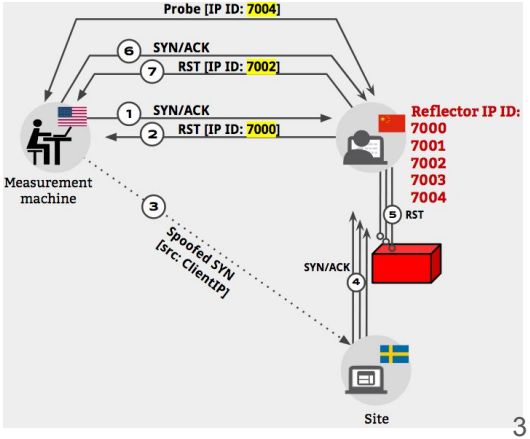

Coping with Reflector IP ID Noise

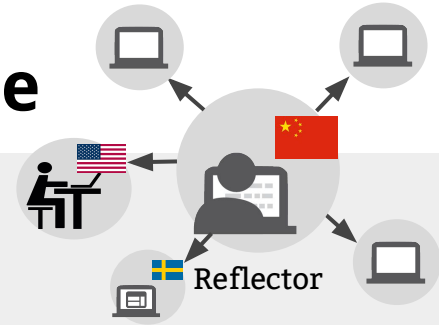

Amplifying the signal

Effect of sending *N* spoofed SYNs:

Site-to-Reflector Blocked	No Direction Blocked	Reflector-to-Site 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 Reflector IP ID Noise

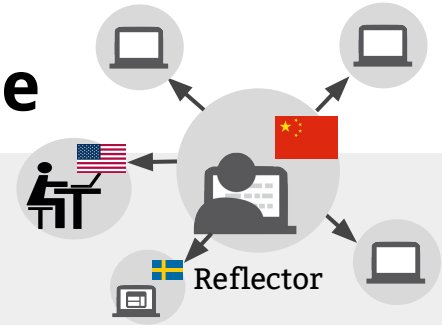

Amplifying the signal

Effect of sending *N* spoofed SYNs:

Site-to-Reflector Blocked	No Direction Blocked	Reflector-to-Site 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, ...

Augur Framework

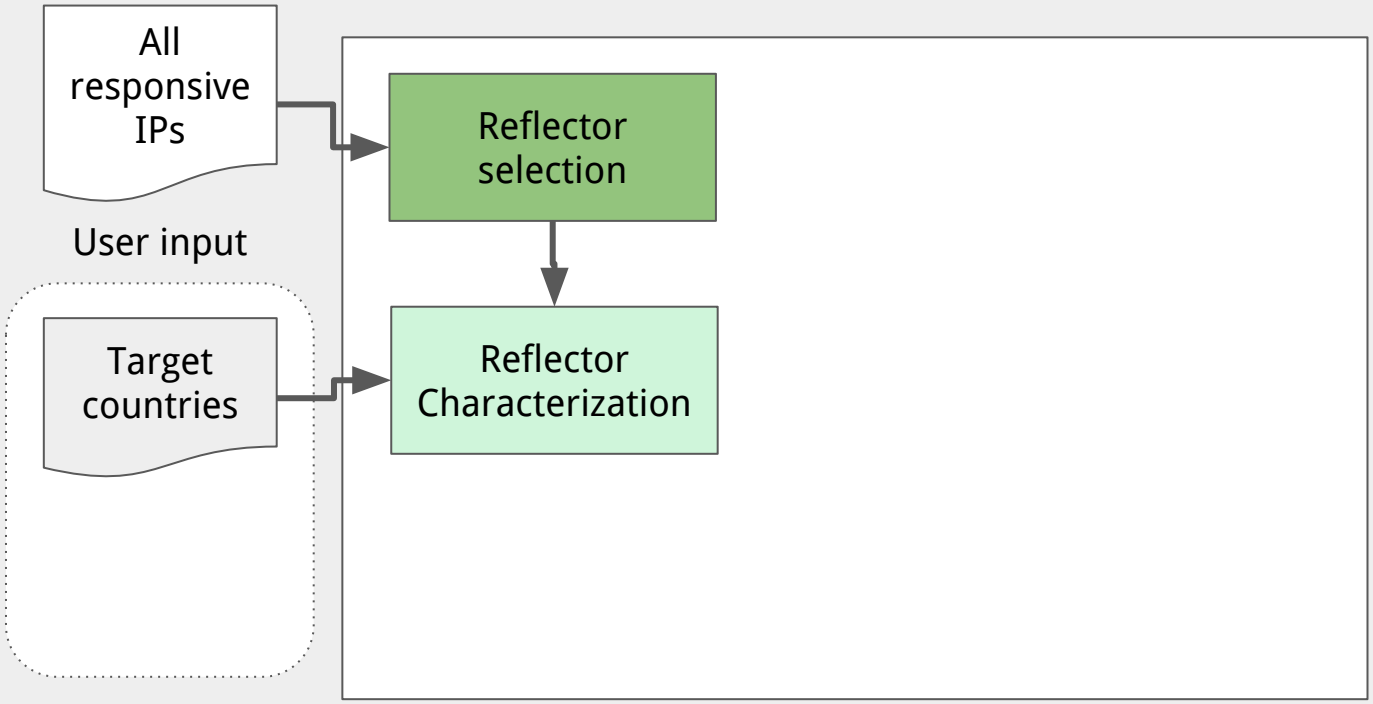

Augur Framework

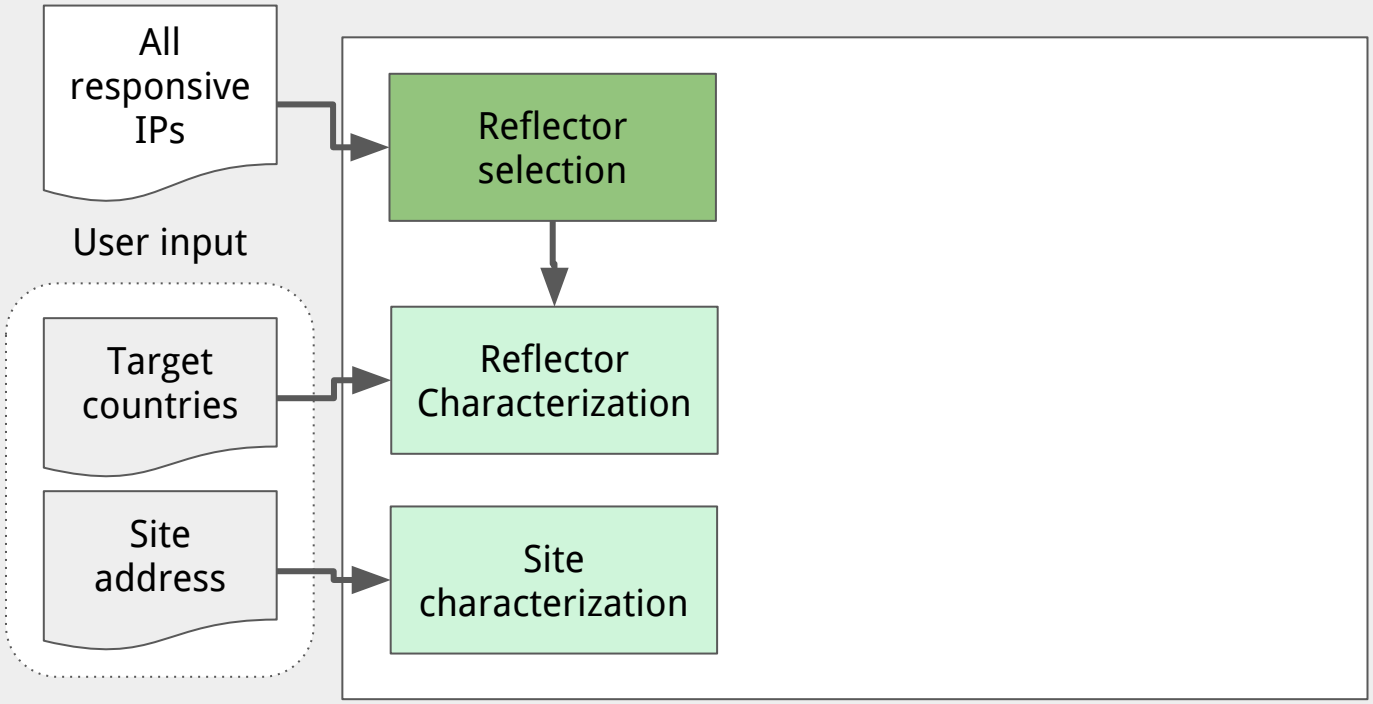

Augur Framework

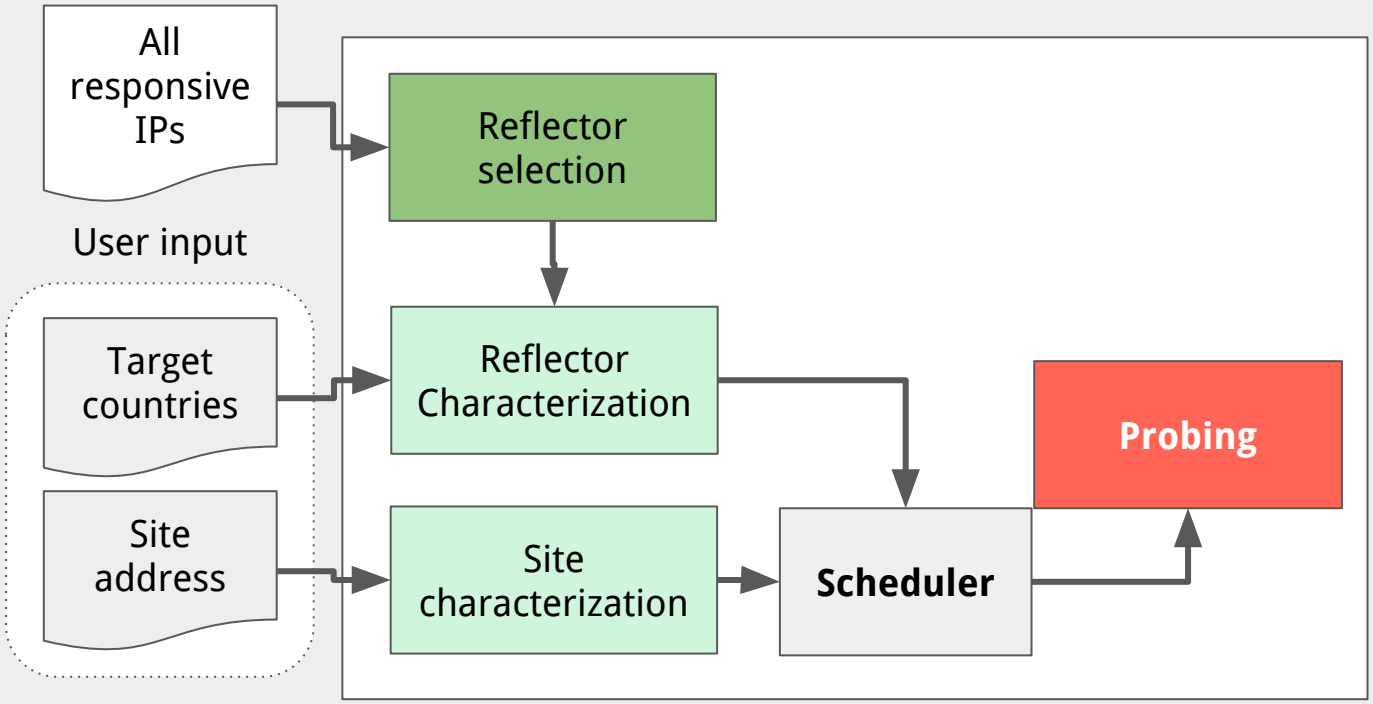

Augur Framework

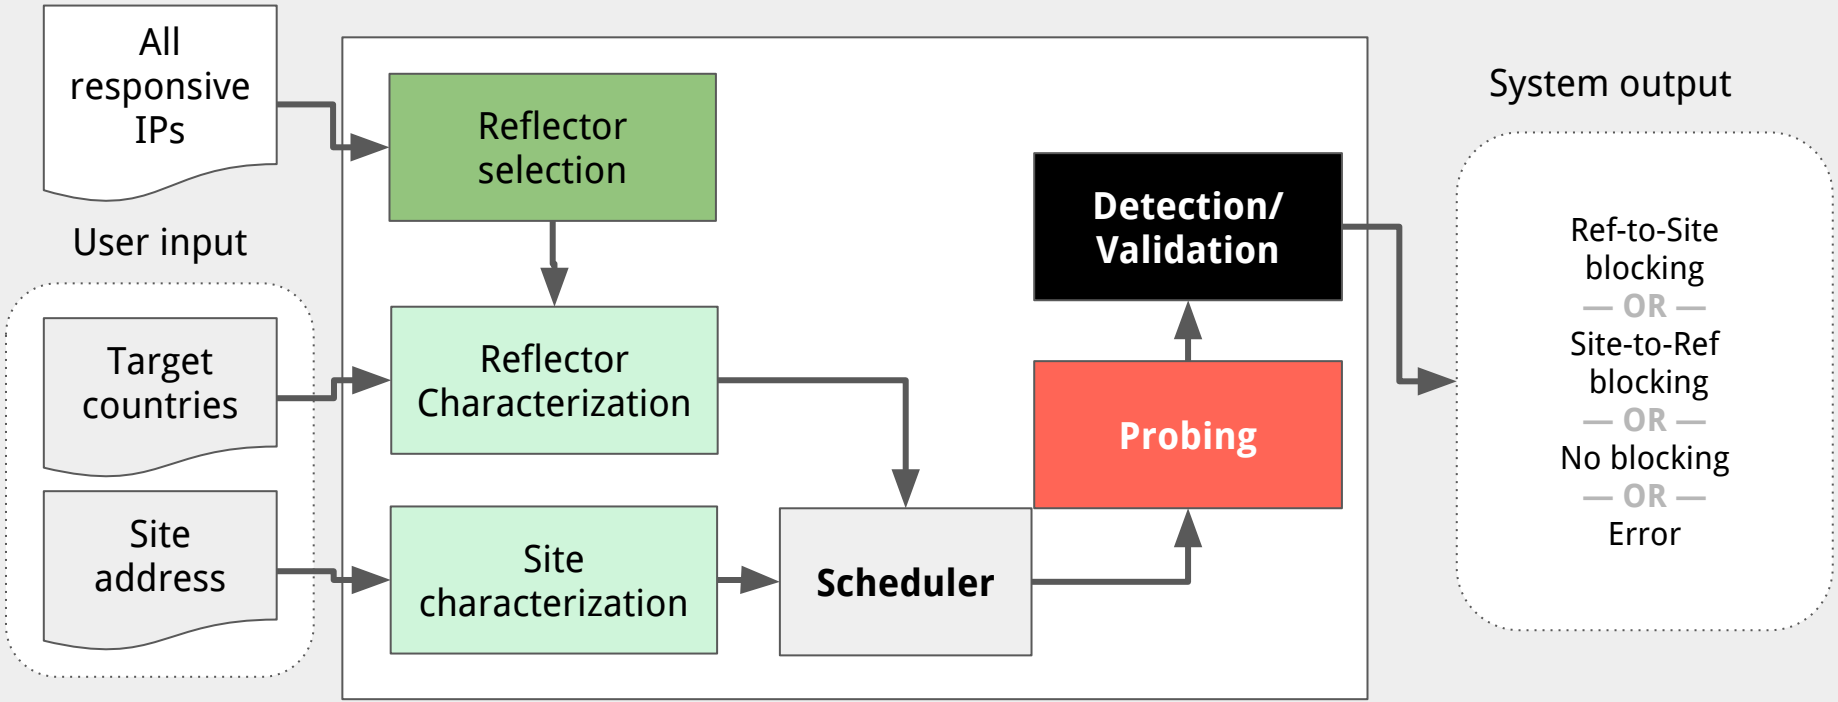

Coverage

Challenge: Need global vantage points from which to measure

Scanning IPv4 on port 80:

- 22.7 million potential reflectors!

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

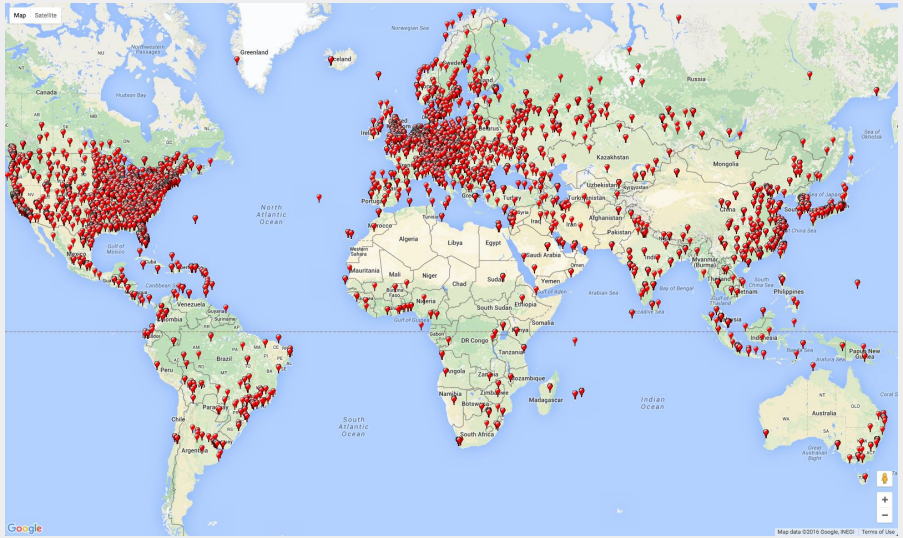

Ethics

Challenge: Probing banned sites from users' machines creates risk

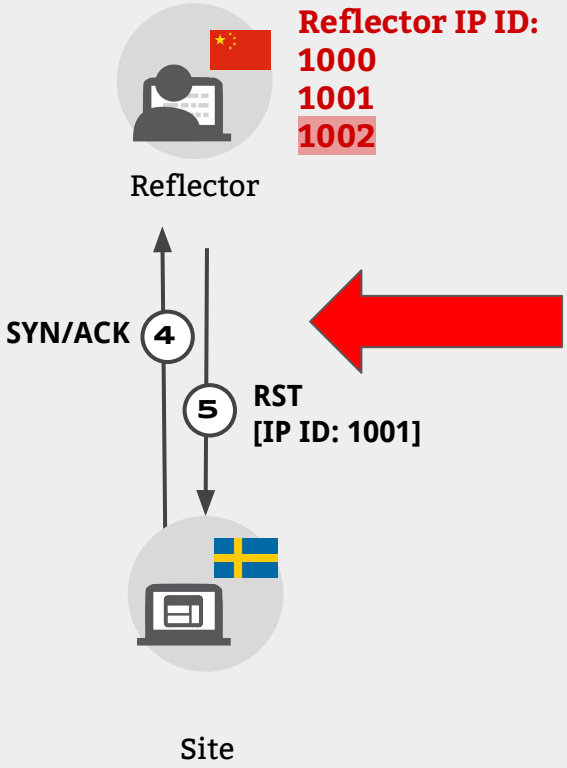

Ethics

Challenge: Probing banned sites from users' machines creates risk

THREE KEY CHALLENGES:
Coverage, ethics, and continuity

Use only **infrastructure devices** to source probes

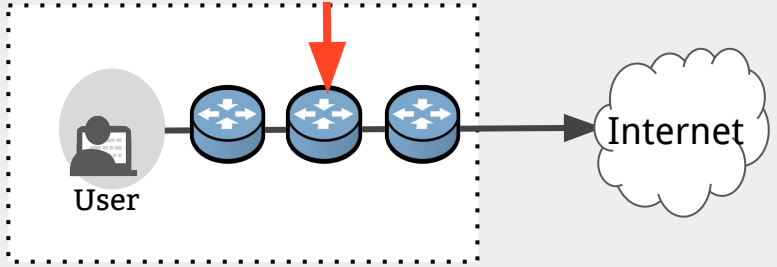

Global IP ID	22.7 million	236 countries (and dependent territories)
Two hops back from end user	**53,000**	**180 countries**

Continuity

Challenge: Need to repeat measurements over time

Augur doesn't depend on end users' availability, and routers have less downtime, allowing us to collect measurements continuously.

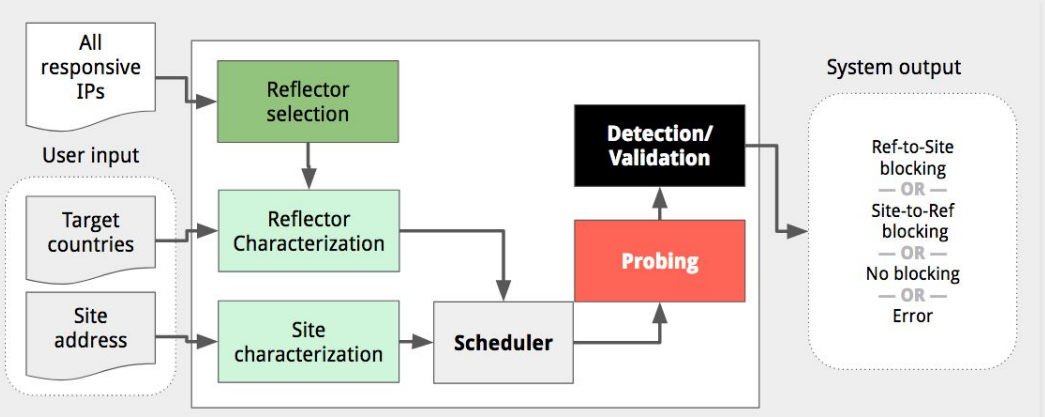

Running Augur In the Wild

Reflectors: 2,050

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

Mix of sensitive and popular sites

Duration: 17 days

Measurements per reflector-site: 47

Overall # of measurements: 207.6 million

Top Blocked Sites

Site-to-Reflector Blocked

Site-to-Reflector blocking

No.	Site	% Refs	% Cnt.	Class
1.	hrcr.org	41.7	83.0	Human Rights
2.	alstrangers.[LJ].com	37.9	78.8	Militants
3.	varlamov.ru	37.7	78.0	Foreign relations
	nordrus-norna.[LJ].com			Hate speech
4.	www.stratcom.mil	37.5	78.6	Foreign relations
5.	www.demonoid.me	21.7	58.5	P2P file sharing
6.	amateurpages.com	21.2	57.9	Adult contents
	voice.yahoo.jajah.com			Voice over IP
	amtrak.com			ALEXA

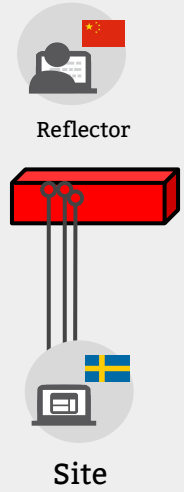

Interesting example:

- amtrak.com was blocked for 21% of reflectors, 57% of countries (ranked 6) → Collateral damage

Top Blocked Sites

Reflector-to-site Blocked

Reflector-to-site blocking

No.	Site	% Refs	% Cnt.	Class
1.	nsa.gov	7.4	23.3	US Gov.
2.	scientology.org	2.2	6.9	Minority faiths
3.	goarch.org	1.9	4.4	Minority faiths
4.	yandex.ru	1.8	3.8	Freedom of Expression
5.	hushmail.com	1.8	4.4	Free email
6.	carnegieendowment.org	1.6	4.4	Political reforms

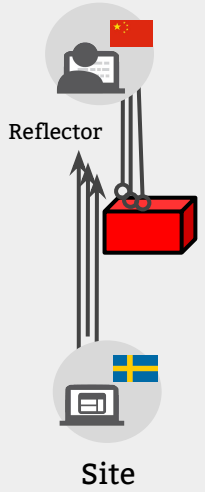

Interesting example:

- nsa.gov was blocked for 7.4% of reflectors, 23% of countries (ranked 1)

Note: Some servers discriminate by providing their services to specific regions

Examples: Dating sites, banking sites, or sites that have to follow embargo rules

Side Channels at Other Network Layers

Network interference happens at all layers

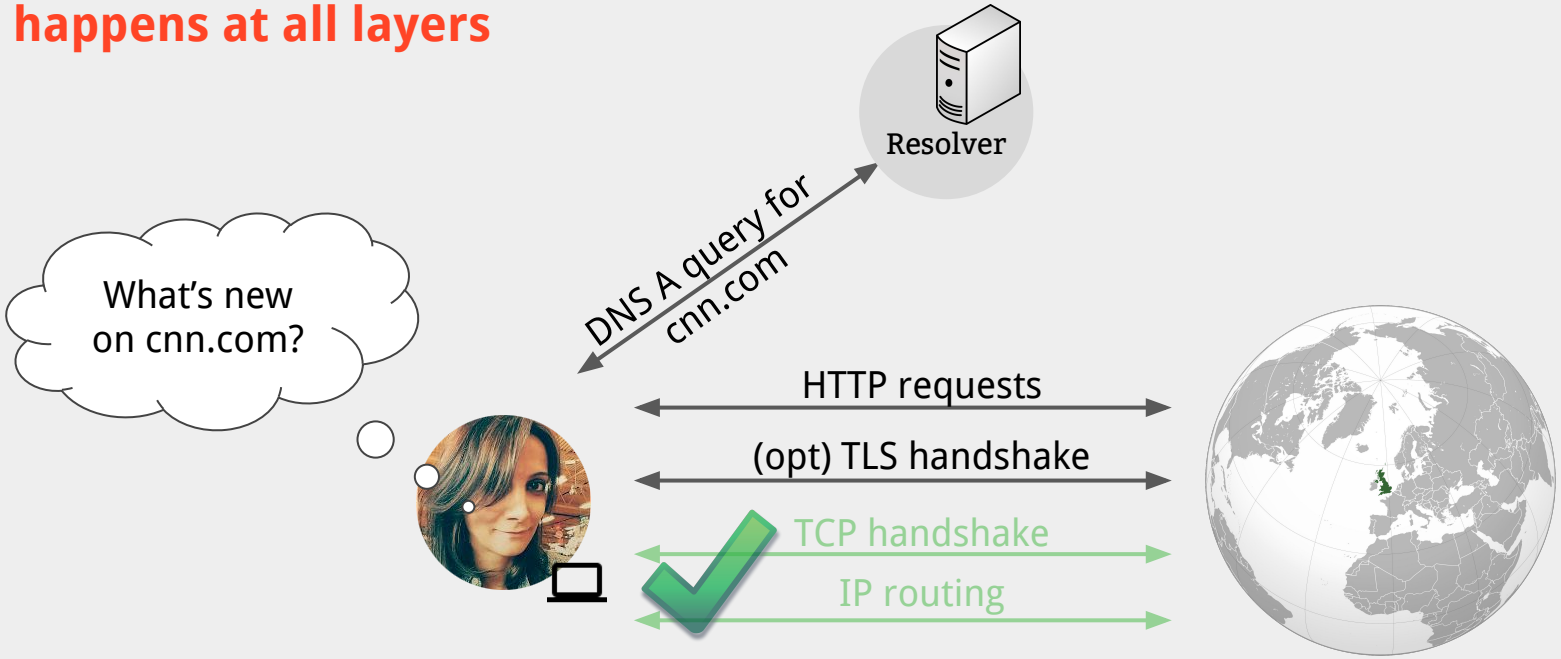

Satellite (Iris)

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

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

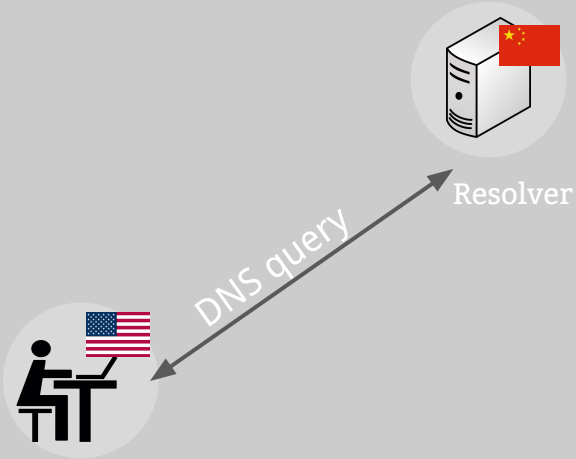

* **Satellite: Joint Analysis of CDNs and Network-Level Interference,Satelite,** Scott, Anderson, Kohno, and Krishnamurthy. In USENIX ATC, 2016.
* **Global Measurement of DNS Manipulation,** Pearce, Jones, Li, Ensafi , Feamster, Paxson, USENIX Security, August 2017

Deploying Satellite

Challenge:
Identify “wrong” DNS responses

THREE KEY CHALLENGES:
Coverage, ethics, and continuity

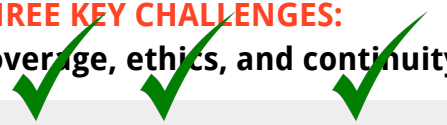

Coverage:

- Scan IPv4 for open resolvers: 11M resolvers, 6M returned correct answer, 166 countries

Ethical:

- Using resolvers reasonably attributed to Internet naming infrastructures: ~ 14k

Continuity:

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

Detecting DNS manipulation:

- Using consistency and independent verifiability heuristics.

Side Channels at Other Network Layers

Network interference happens at all layers

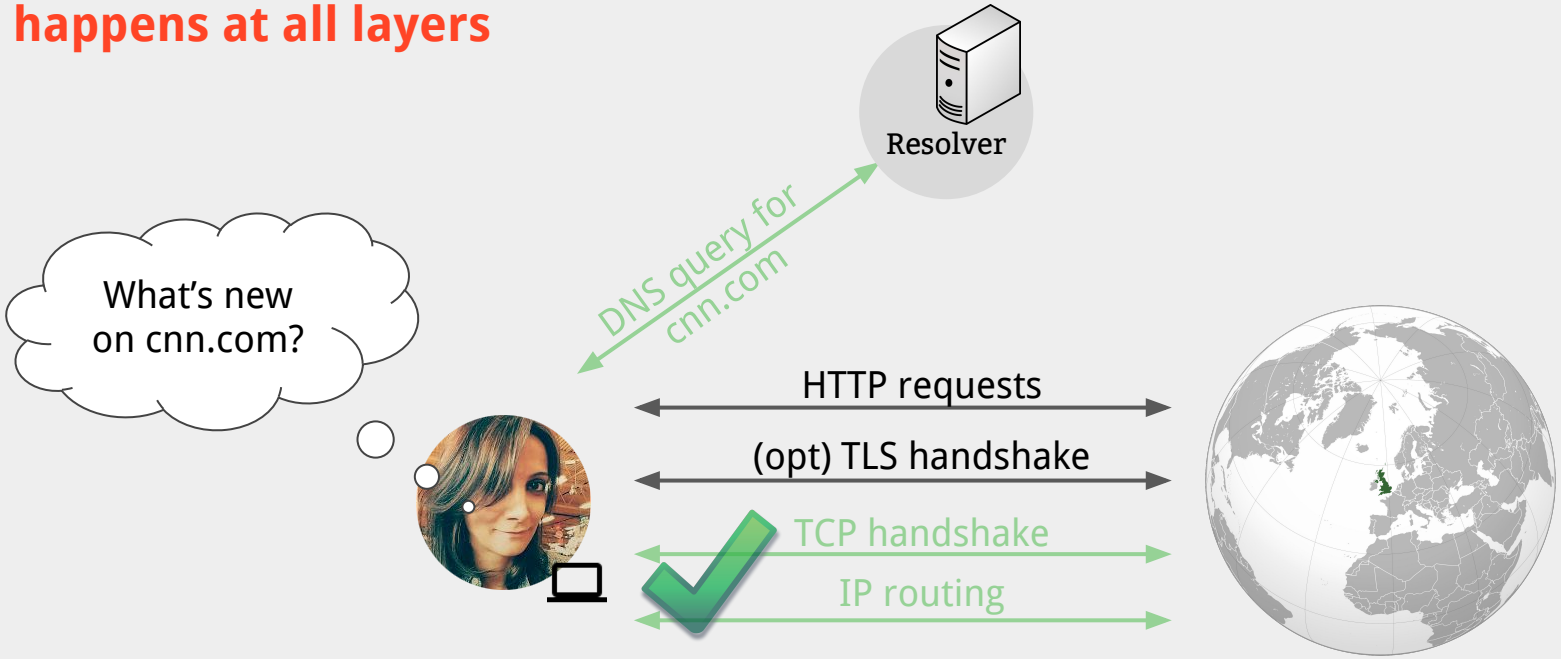

Side Channel to Detect Application-Layer Blocking

PROBLEM:

- How can we detect whether a keyword/URLs are being blocked around the world?

Echo Protocol to the Rescue!

Echo Protocol:

- The Echo Protocol, as defined in RFC862 in 1983 by J. Postel, is a network debugging service, predating ICMP Ping.

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 a keywords/URLs are blocked

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

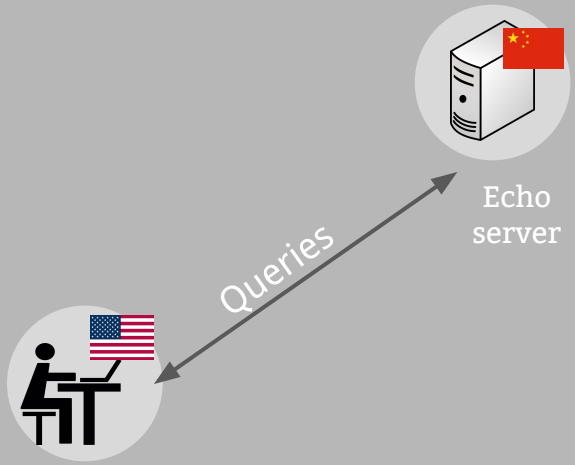

Side Channels at Other Network Layers

Network interference happens at all layers

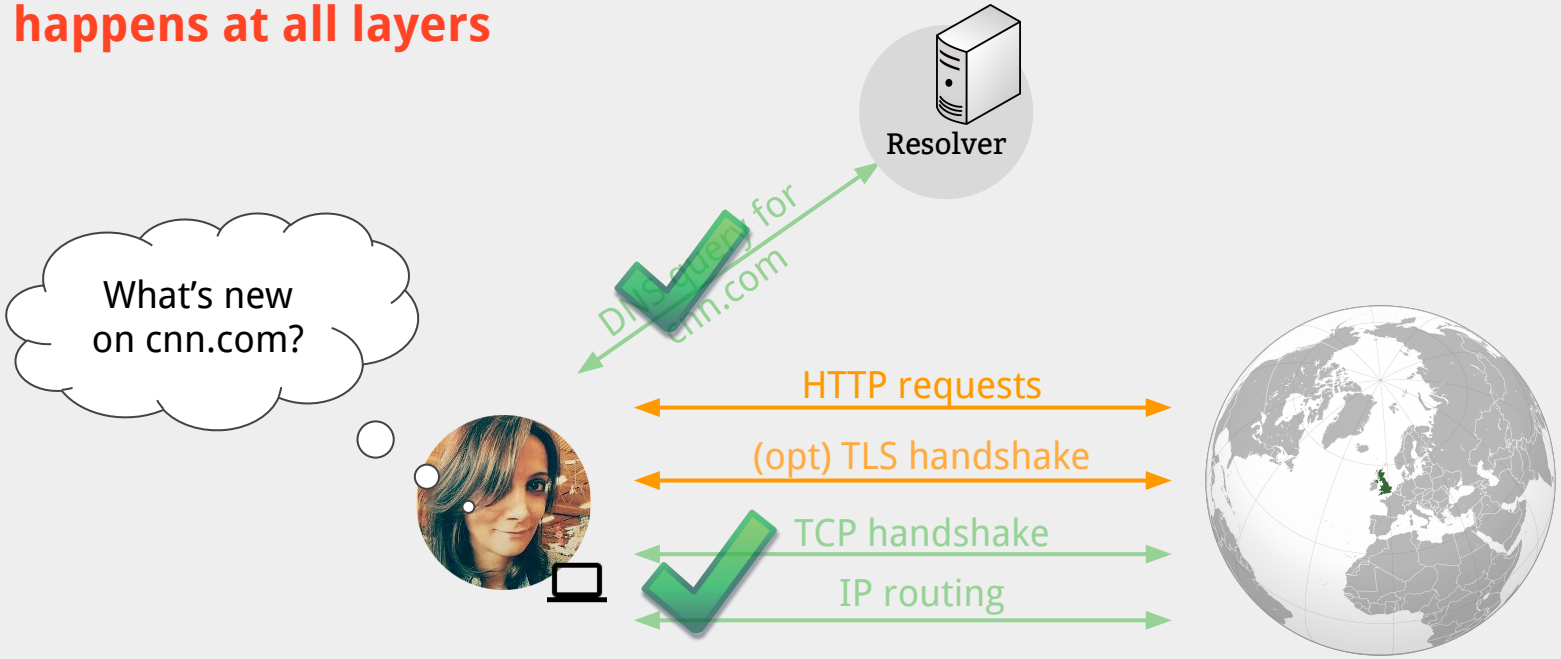

Side Channels at Other Network Layers

Network interference happens at all layers

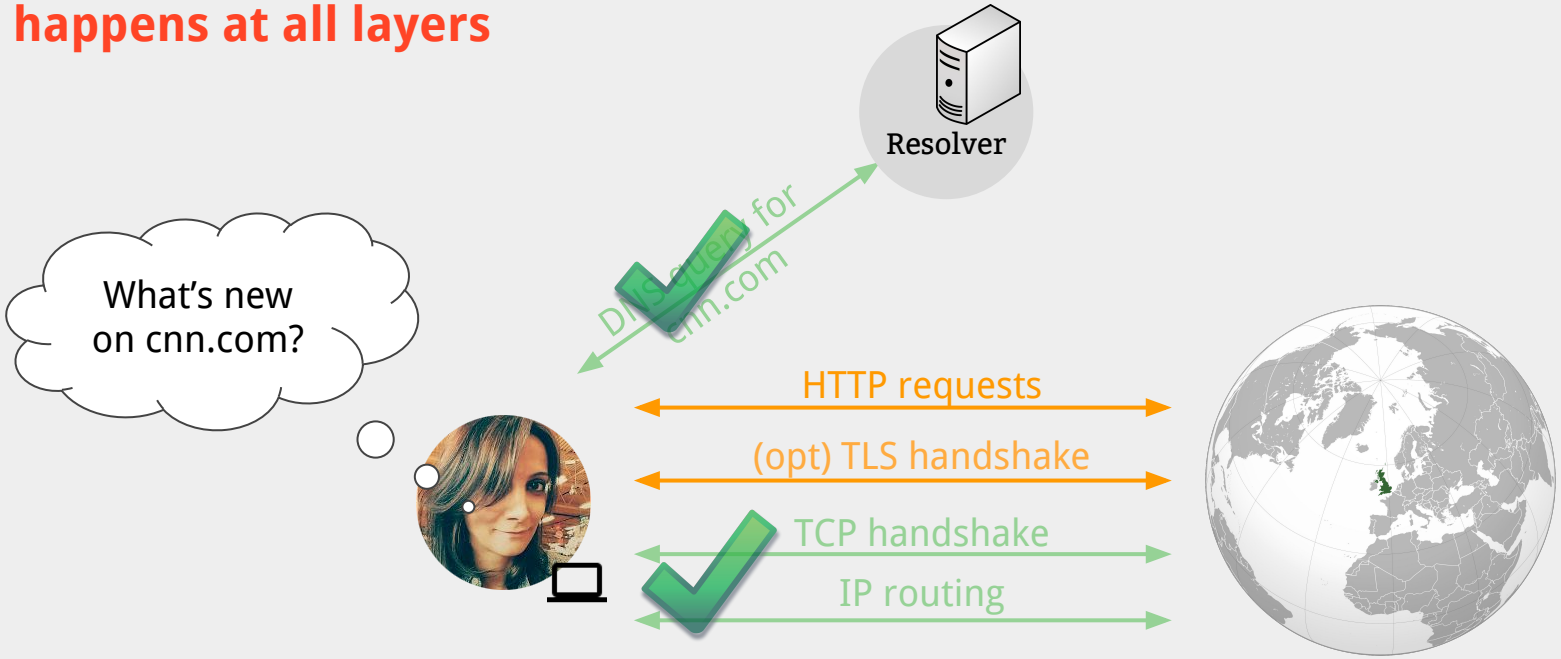

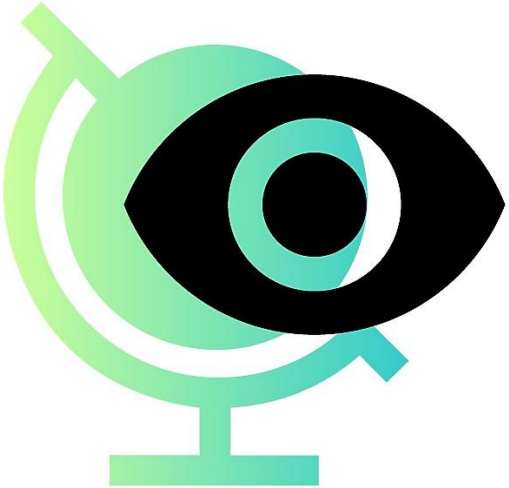

Censored Planet, a system that provides a continual and global view of Internet censorship

- **Daily reachability measurements** for key websites from countries worldwide
- Data collected with Augur, Satellite, and Quack combined with **side channels at other network layers**
- Tools for mapping and **comparative analyses** across locations and time

Backup slides :)

Augur for Continuous Scanning

Insight: Some measurements much noisier than others.

Augur for Continuous Scanning

Insight: Some measurements much noisier than others.

Probing Methodology:

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

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

Augur for Continuous Scanning

Insight: Some measurements much noisier than others.

Probing Methodology:

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

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

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

Augur: Sequential Hypothesis Testing

Defining a random variable:

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

Augur: Sequential Hypothesis Testing

Defining a random variable:

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

Calculate known outcome probabilities (priors):

Prior 1: Prob. of no IPID acceleration when there is blocking
Prior 2: Prob. of IPID acceleration when there is no blocking

Augur: Sequential Hypothesis Testing

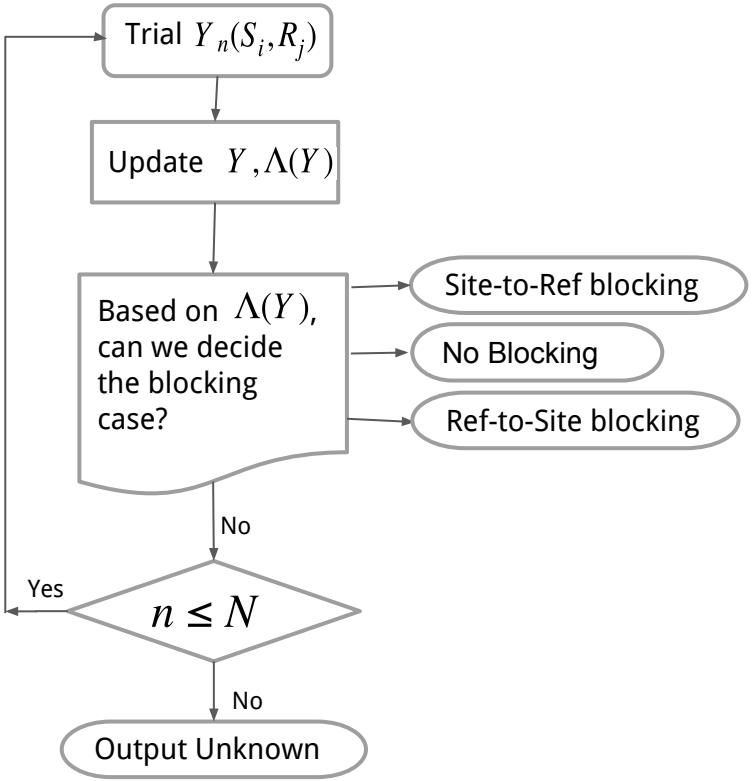

Maximum Likelihood Ratio

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

Augur Framework