Running Refraction Networking for Real

Benjamin VanderSloot, Sergey Frolov, **Jack Wampler**,
Sze Chuen Tan, Irv Simpson, Michalis Kallitsis,
J. Alex Halderman, Nikita Borisov, and Eric Wustrow

Internet Censorship

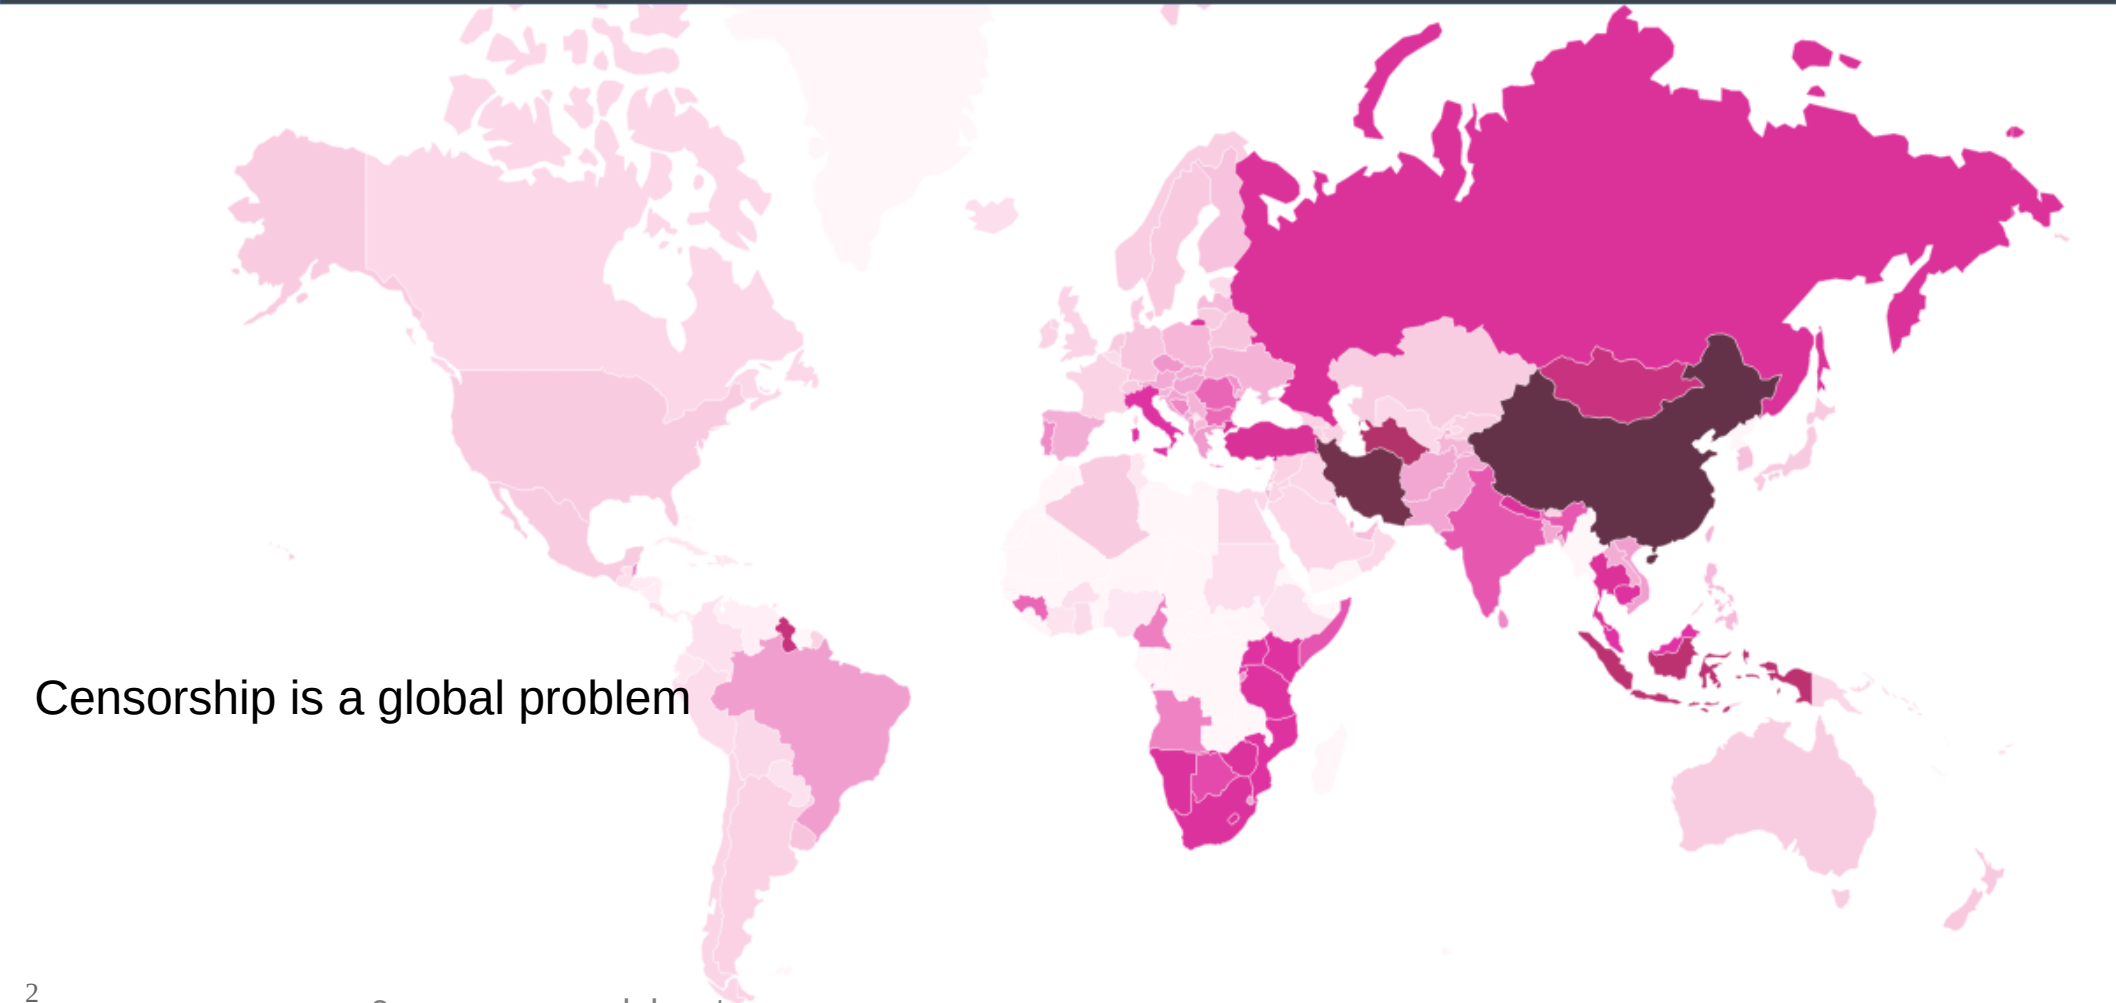

Censorship is a global problem

Source: censoredplanet.org

Proxies

Name some proxies

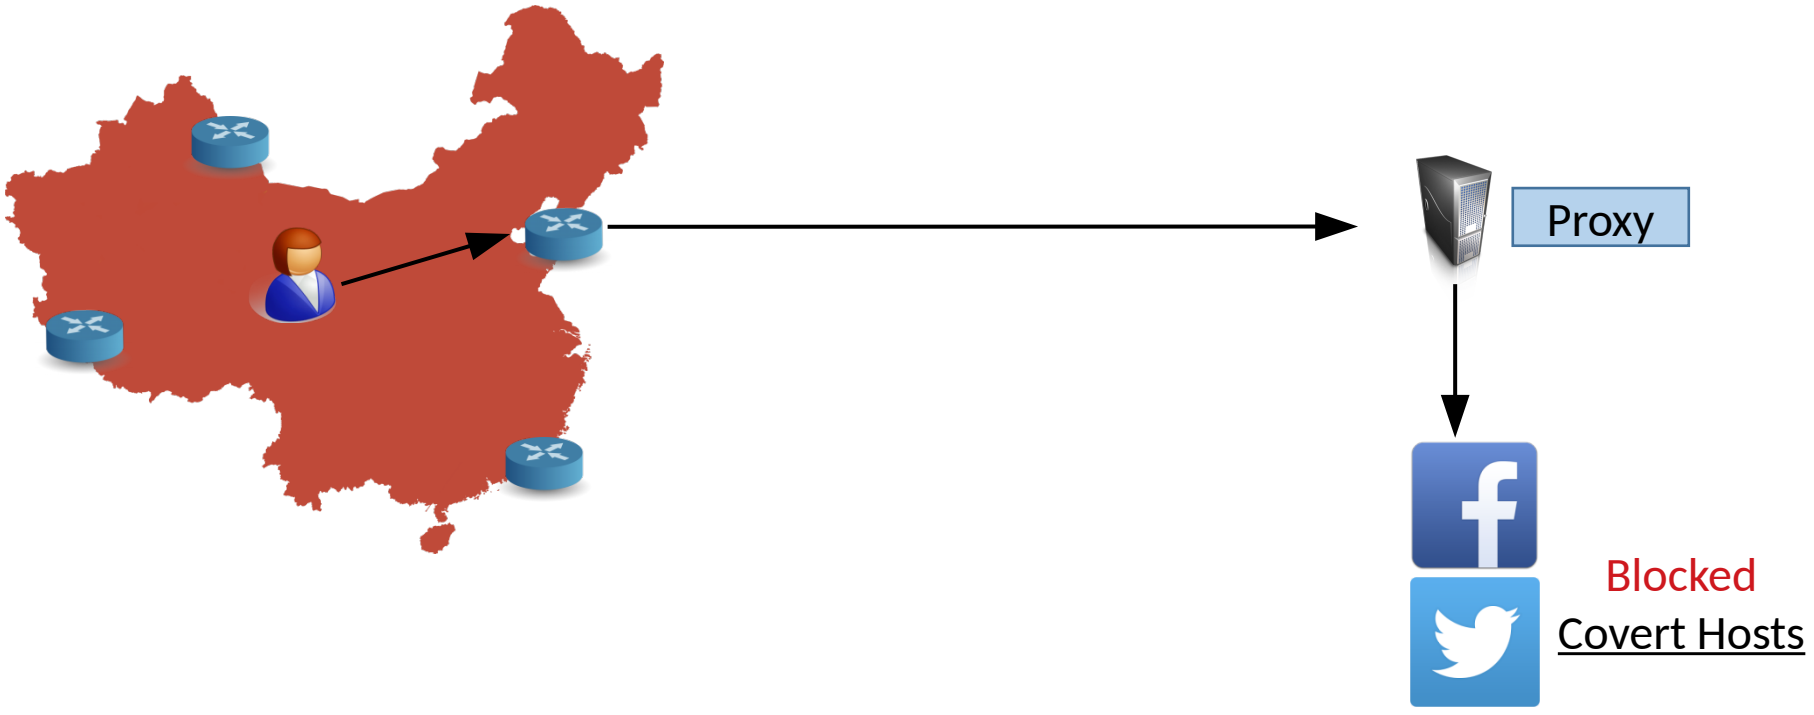

Blocking Proxies

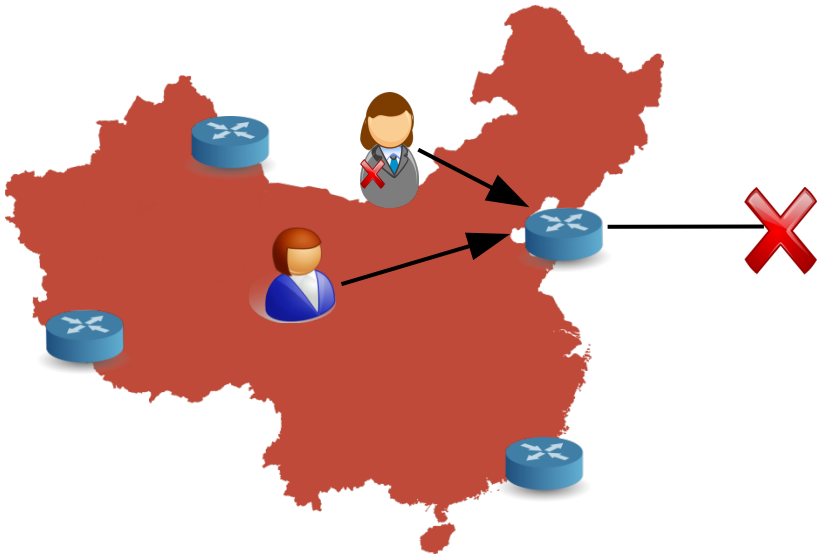

Censors try to discover proxies
by connecting to them as clients

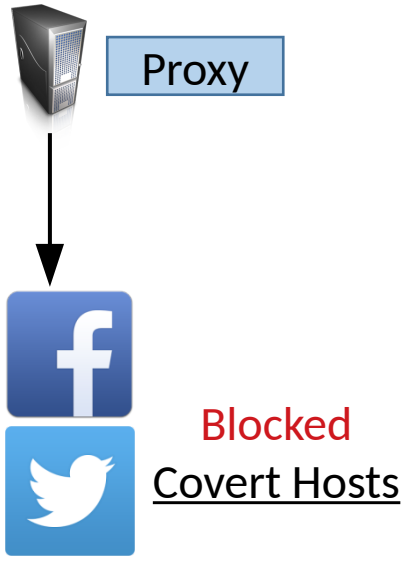

Refraction Networking

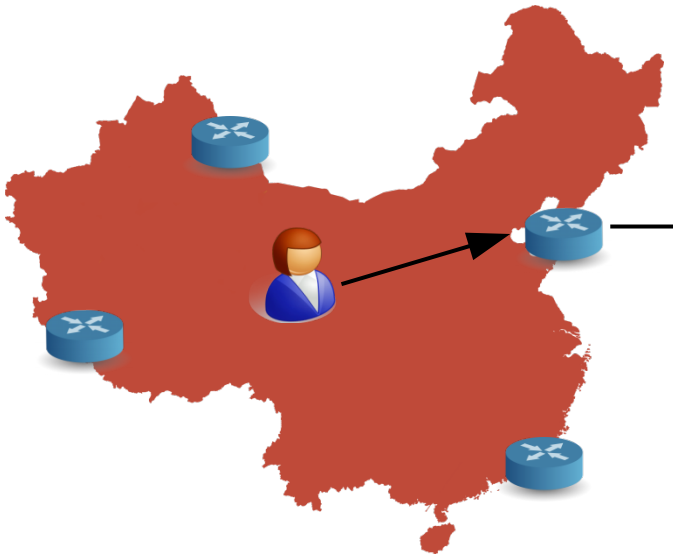

Not Blocked
Decoy Host

Blocked
Covert Hosts

Refraction Networking

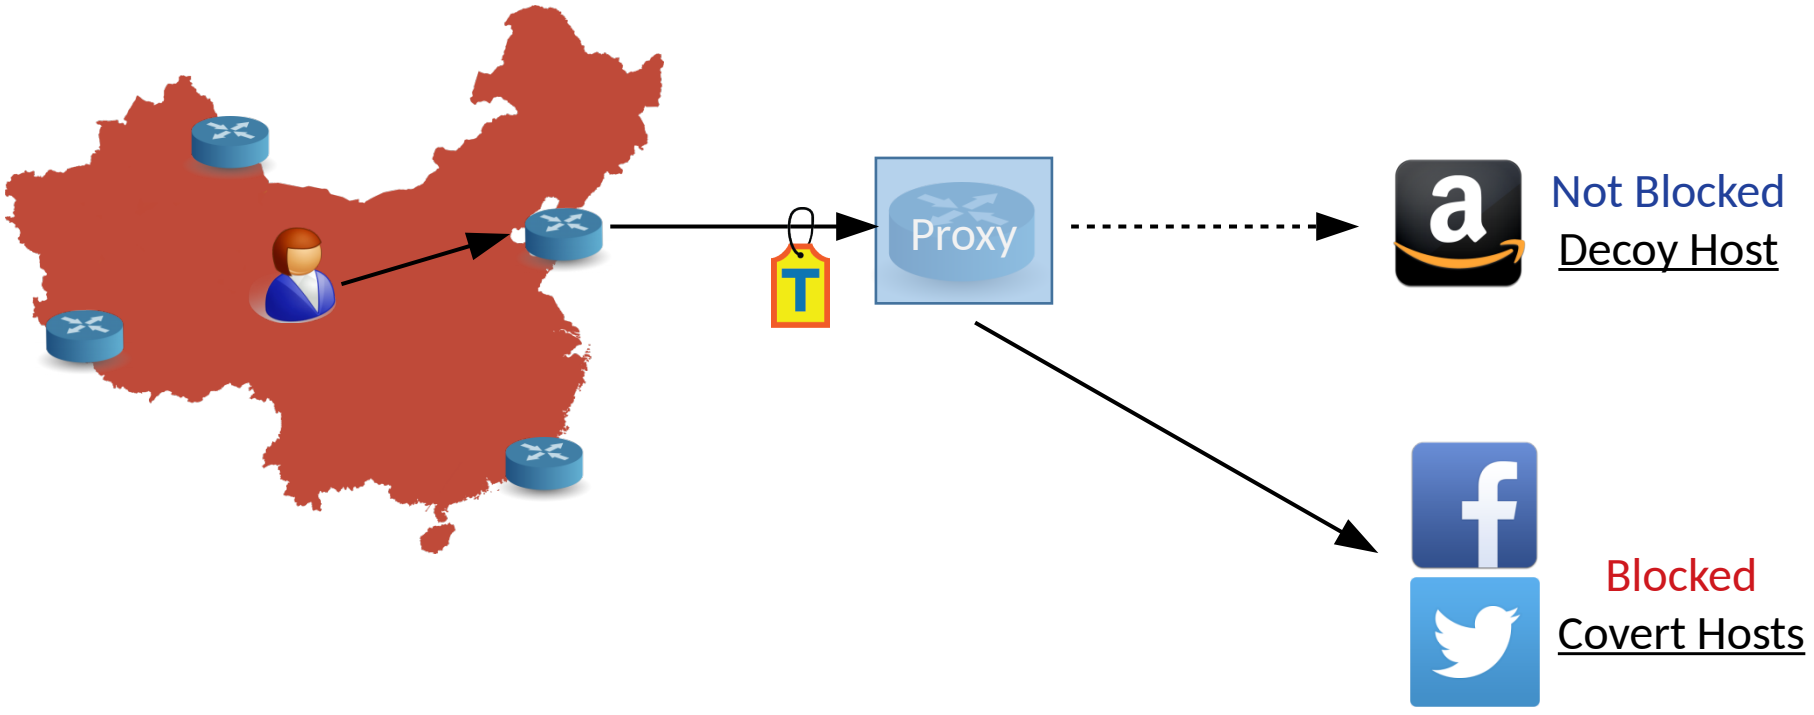

Refraction Networking

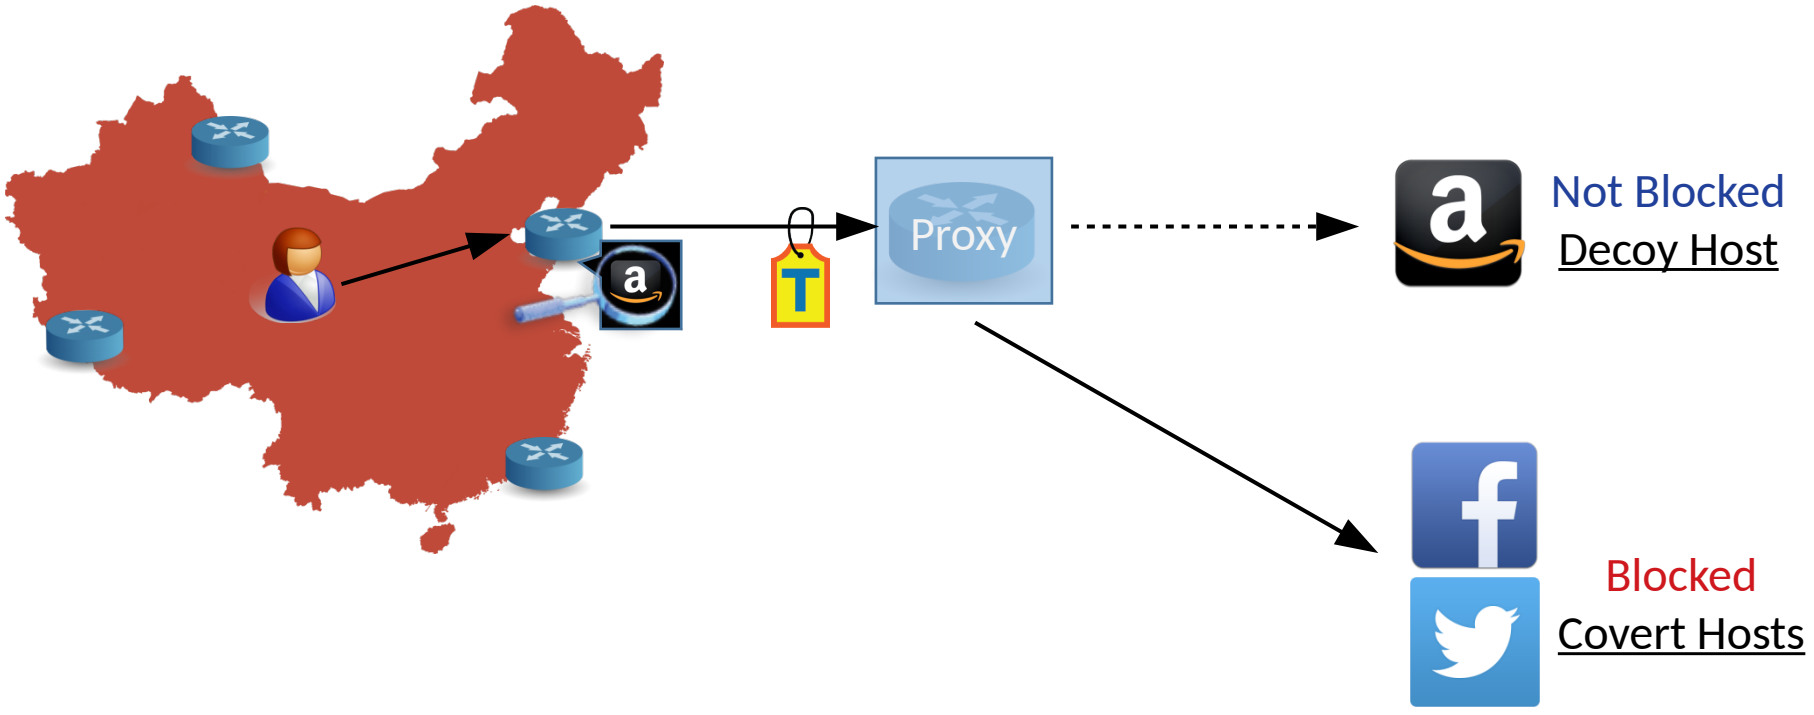

Refraction Networking

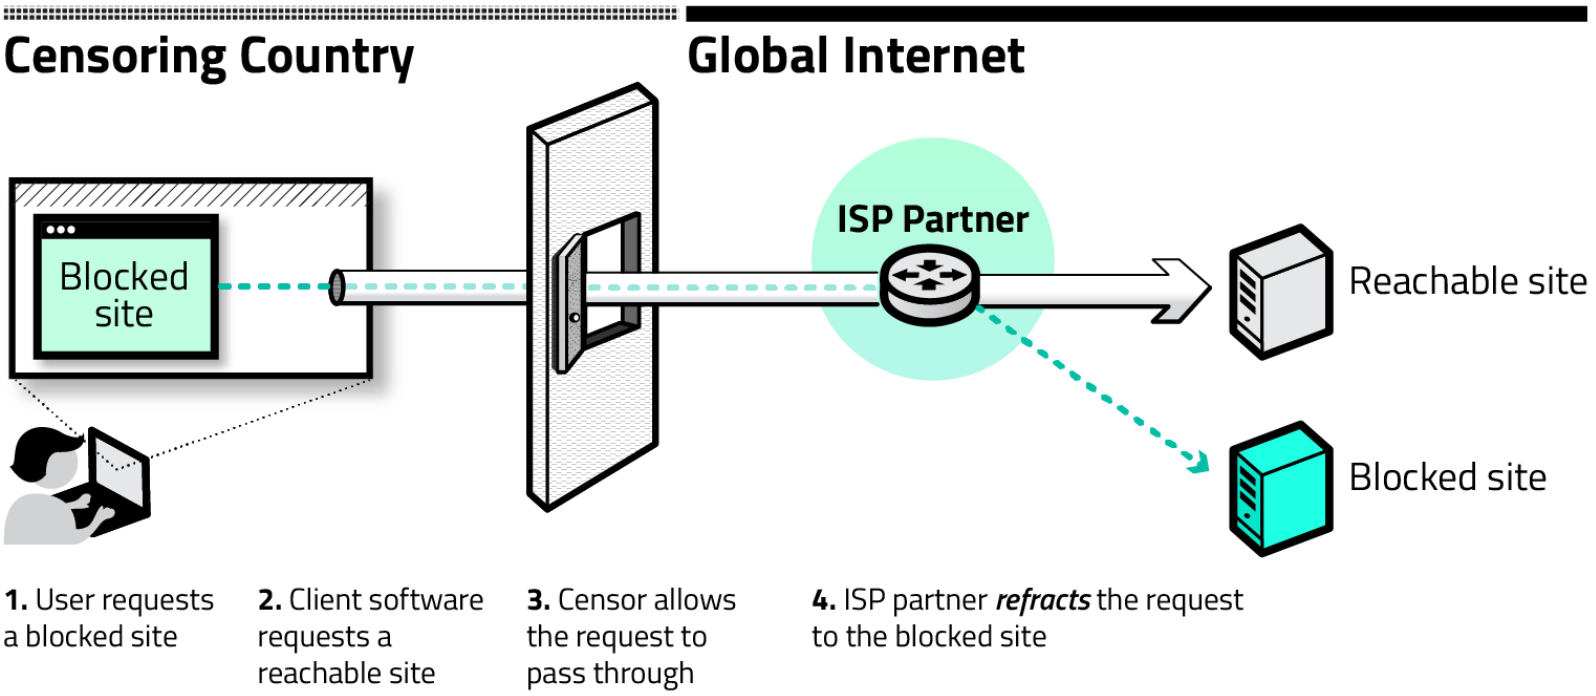

Refraction Networking

FORMERLY *DECOY ROUTING*

Telex: Anticensorship in the Network Infrastructure
Eric Wustrow, Scott Wolchok, Ian Goldberg, J. Alex Halderman [USENIX 2011]

Decoy Routing: Toward Unblockable Internet Communication
Josh Karlin, Daniel Ellard, Alden W. Jackson, Christine E. Jones, Greg Lauer, David P. Mankins, W. Timothy Strayer [FOCI 2011]

Cirripede: Circumvention Infrastructure using Router Redirection with Plausible Deniability
Amir Houmansadr, Giang T. K. Nguyen, Matthew Caesar, Nikita Borisov [CCS 2011]

***TapDance**: End-to-Middle Anticensorship without Flow Blocking*
Eric Wustrow, Colleen M. Swanson, J. Alex Halderman [USENIX 2014]

***Rebound:** Decoy Routing on Asymmetric Routes Via Error Messages*
Daniel Ellard, Alden Jackson, Christine Jones, Victoria Manfredi, W. Timothy Strayer, Bishal Thapa, Megan Van Welie [IEEE LCM 2015]

***Slitheen**: Perfectly Imitated Decoy Routing through Traffic Replacement*
Cecylia Bocovich, Ian Goldberg [CCS 2016]

***The Waterfall of Liberty**: Decoy Routing Circumvention that Resists Routing Attacks*
Milad Nasr, Hadi Zolfaghari, Amir Housmansadr [ACM 2017]

***MultiFlow**: Cross-Connection Decoy Routing using {TLS} 1.3 Session Resumption*
Victoria Manfredi, and Pi Songkuntham [FOCI 2018]

Early Refraction Schemes

Refraction networking

- Station listens network router at an ISP

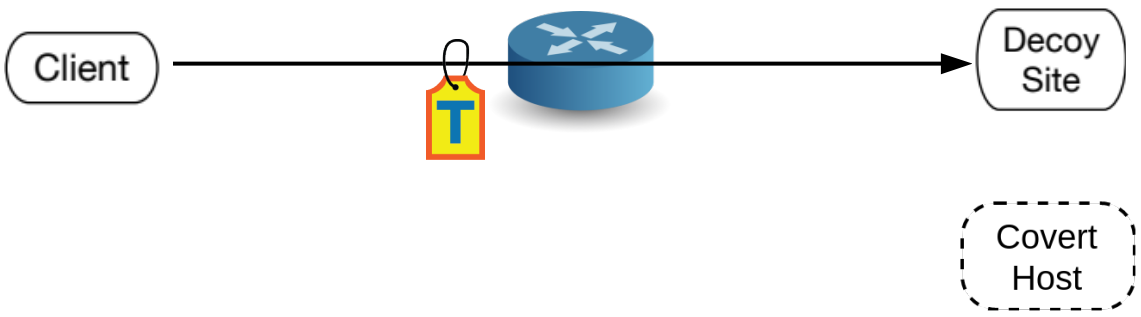

Early Refraction Schemes

Inline Blocking

- Drops connections to decoy sites
- Redirects traffic to covert destination

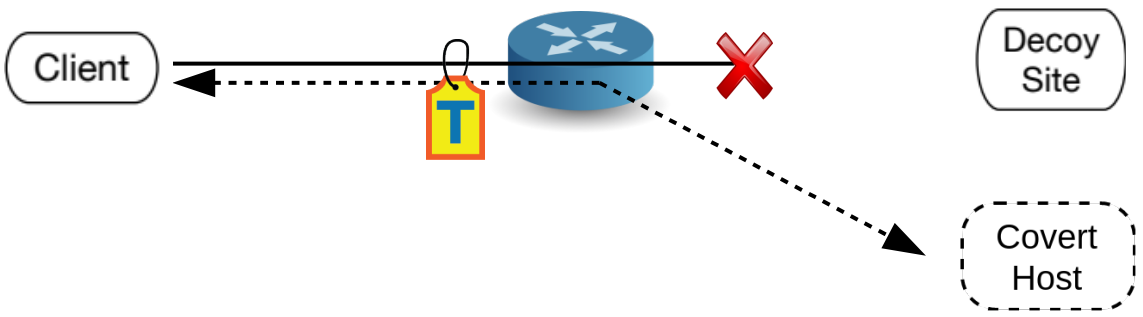

TapDance

TapDance

- Station listens on passive tap at an ISP

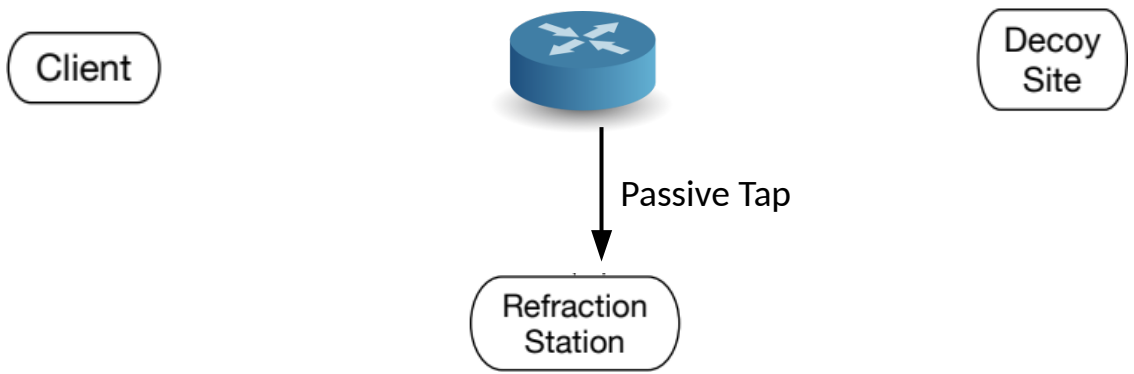

TapDance

TapDance

- Station listens on passive tap at an ISP
- Client connects to the decoy

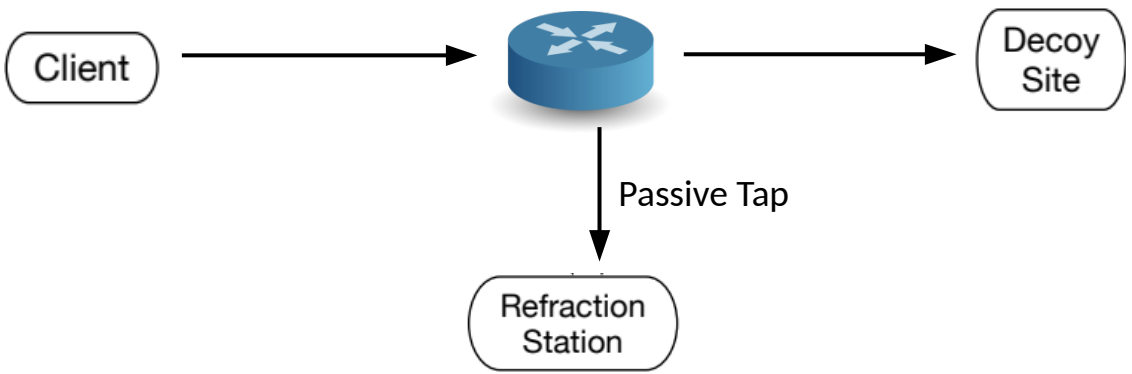

TapDance

TapDance

- Station listens on passive tap at an ISP
- Client connects to the decoy
- Client sends something to silence the decoy

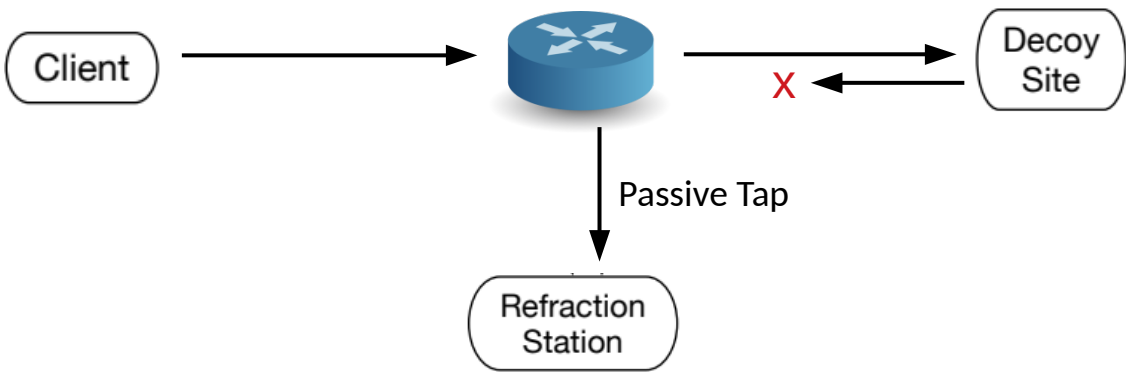

TapDance

TapDance

- Station listens on passive tap at an ISP
- Client connects to the decoy
- Client sends something to silence the decoy
- Station pretends to be the decoy while the connection stays open

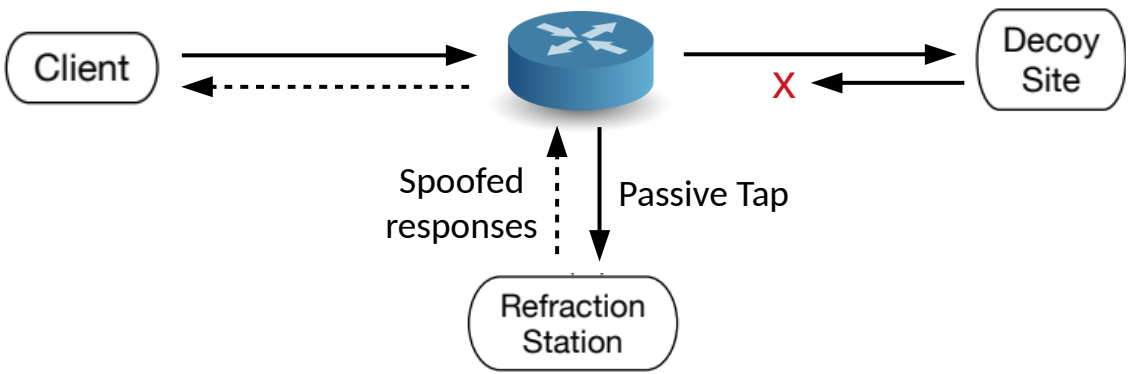

Deployment

Deployment

Trial deployment of Tapdance

We evaluate 4 months of data
from early 2019

Station Placement

- Detectors placed at major ingress points
- Four stations

  3 x (4 x 10Gbps) stations

  +1 x (2 x 10Gbps) station

  ---

  140 Gbps Merit capacity

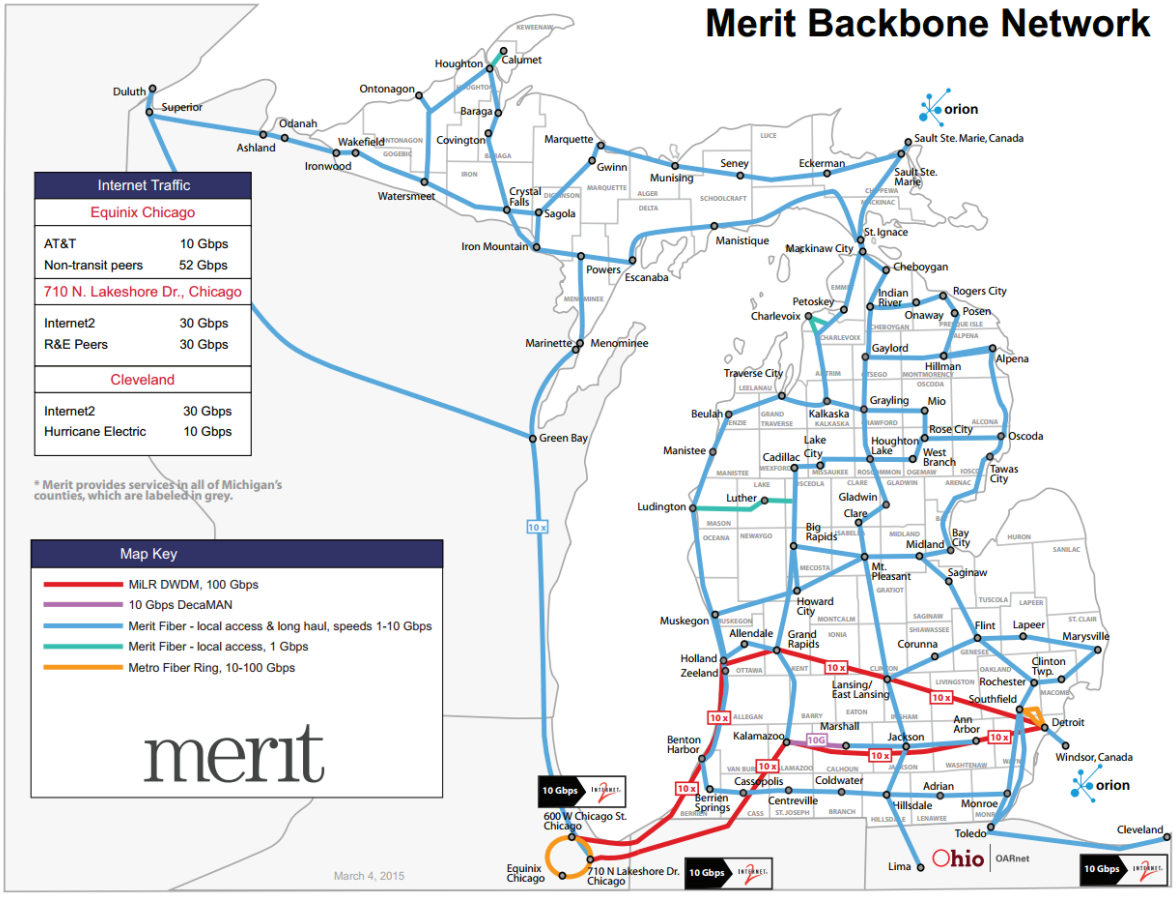

Station Placement

- Detectors placed at major ingress points
- Four stations

  140 Gbps Merit capacity

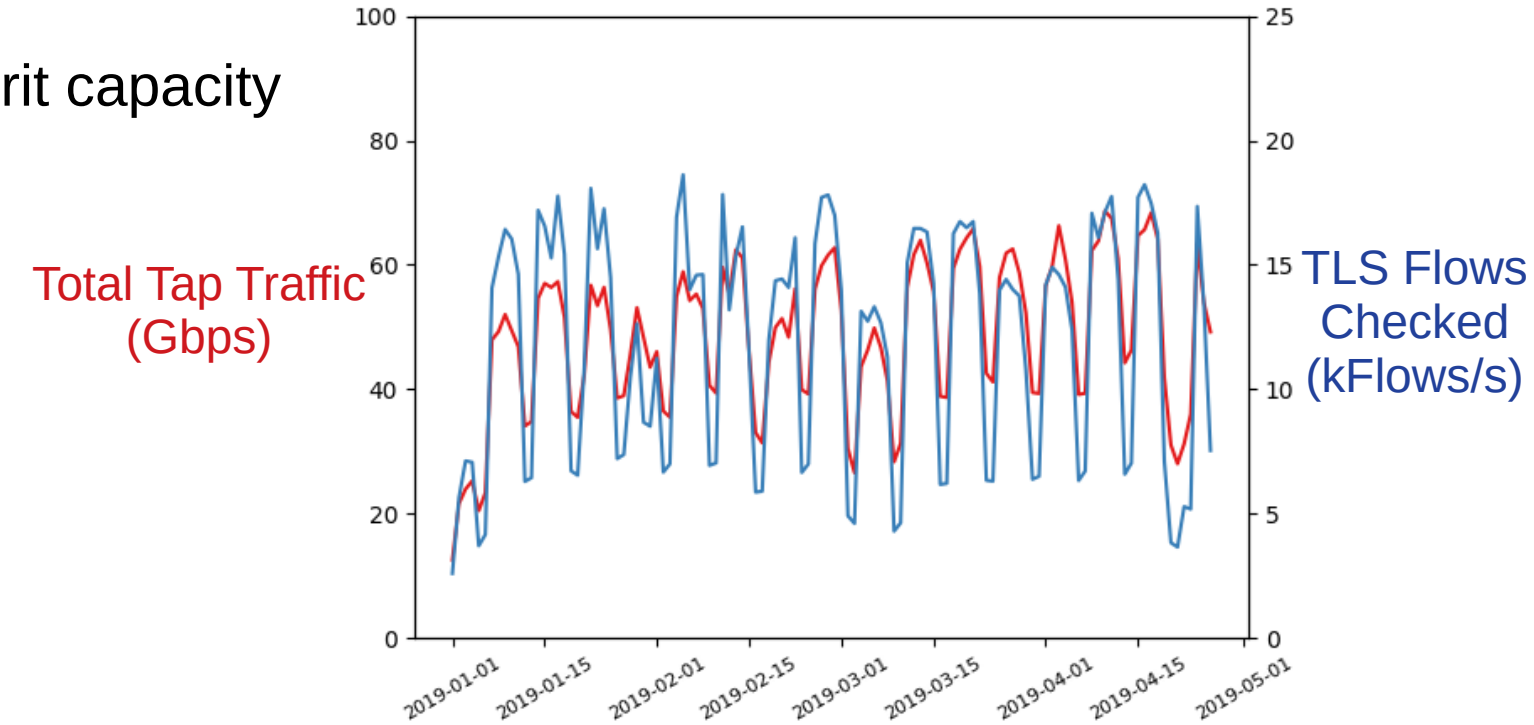

Previous Deployment

Previous TapDance Trial

FOCI ‘17

Tapdance Flows are short, so to support users we multiplex over many short connections

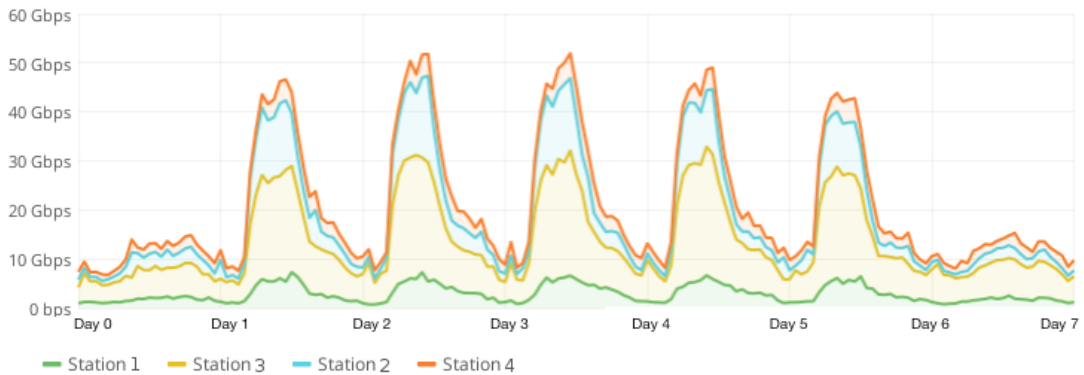

An ISP-Scale Deployment of TapDance

Sergey Frolov[1], Fred Douglas[3], Will Scott[5], Allison McDonald[5], Benjamin VanderSloot[5], Rod Hynes[6], Adam Kruger[6], Michalis Kallitsis[4], David G. Robinson[7], Steve Schultze[2], Nikita Borisov[3], J. Alex Halderman[5], and Eric Wustrow[1]

[1]University of Colorado Boulder [2]Georgetown University Law Center [3]University of Illinois Urbana-Champaign [4]Merit Network [5]University of Michigan [6]Psiphon [7]Upturn

Abstract

We report initial results from the world’s first ISP-scale field trial of a refraction networking system. Refraction networking is a next-generation censorship circumvention approach that locates proxy functionality in the middle of the network, at participating ISPs or other network operators. We built a high-performance implementation of the TapDance refraction networking scheme and deployed it on four ISP uplinks with an aggregate bandwidth of 100 Gbps. Over one week of operation, our deployment served more than 50,000 real users. The experience demonstrates that TapDance can be practically realized at ISP scale with good performance and at a reasonable cost, potentially paving the way for long-term, large-scale deployments of TapDance or other refraction networking schemes in the future.

1 Introduction

Censorship circumvention tools typically operate by connecting users to a proxy server located outside the censoring country [3, 12, 15, 18]. Although existing tools use a variety of techniques to conceal the locations of their proxies [5, 9, 13, 17, 19], governments are deploying increasingly sophisticated and effective means to discover and block the proxies [7, 8, 20].

Refraction networking [16][1] is a next-generation circumvention approach with the potential to escape from this cat-and-mouse game. Rather than running proxies at specific edge-hosts and attempting to hide them from censors, refraction works via Internet service providers (ISPs) or other network operators, who provide censorship circumvention functionality for any connection that *passes through* their networks. To accomplish this, clients make HTTPS connections to sites that they can reach, where such connections traverse a participating network. The participating network operator recognizes a steganographic signal from the client and appends the user’s requested data to the encrypted connection response. From the perspective of the censor, these connections are indistinguishable from normal TLS connections to sites the censor has not blocked. To block the refraction connections, the censor would need to block all connections that traverse a participating network. The more ISPs participate in such a system, the greater the extent of collateral damage that would-be censors would suffer by blocking the refracted connections.

A variety of refraction networking systems have been proposed in recent years [2, 6, 10, 11, 21, 22], representing different trade-offs among practicality, stealthiness, and performance. The basic idea is to watch all of the traffic passing through a router, selecting flows which are steganographically tagged as participating in the protocol, and then modifying that traffic by extracting and making the encapsulated request on behalf of the client. While each of these schemes has been prototyped in the lab, implementing refraction within a real ISP poses significant additional challenges. An ISP-scale deployment must be able to:

- Identify client connections on high-speed backbone links operating at 10–40 Gbps or more. This is at the limits of commodity network hardware.
- Be built within reasonable cost constraints, in terms both of required hardware and of necessary rack space at crowded Internet exchange points.
- Operate reliably without disrupting the ISP’s network or the reachable sites clients connect to.
- Have a mechanism for identifying reachable sites for which connections pass through the ISP, and for disseminating this information to clients.
- Coordinate traffic across multiple Internet uplinks or even multiple ISPs.

To demonstrate that these challenges can be solved, we constructed a large trial deployment of the TapDance refraction scheme [21] and operated a trial deployment in partnership with two mid-sized network operators: a

[1]Previous works used the term *decoy routing*, which confusingly shares the name of a specific refraction scheme. We use refraction networking as an umbrella term to refer to all schemes.

Improved Operation

Multiple independent stations

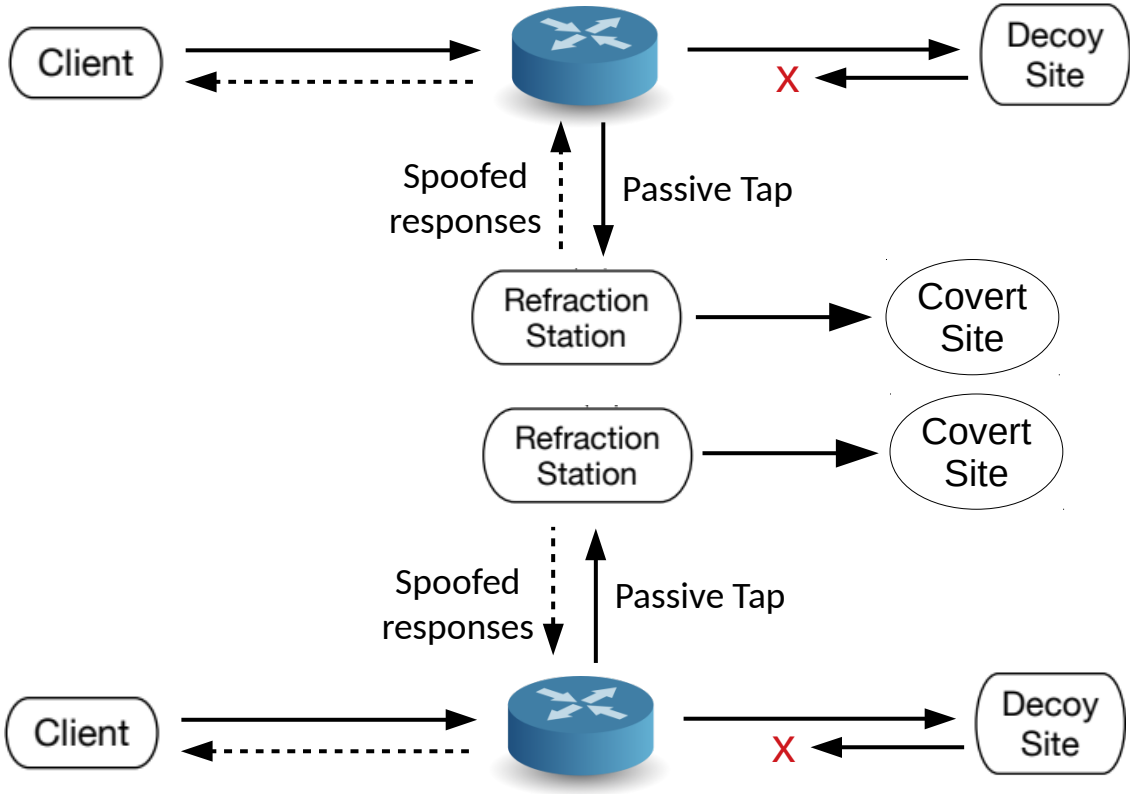

Previous Operation

Multiple Detectors, One Proxy

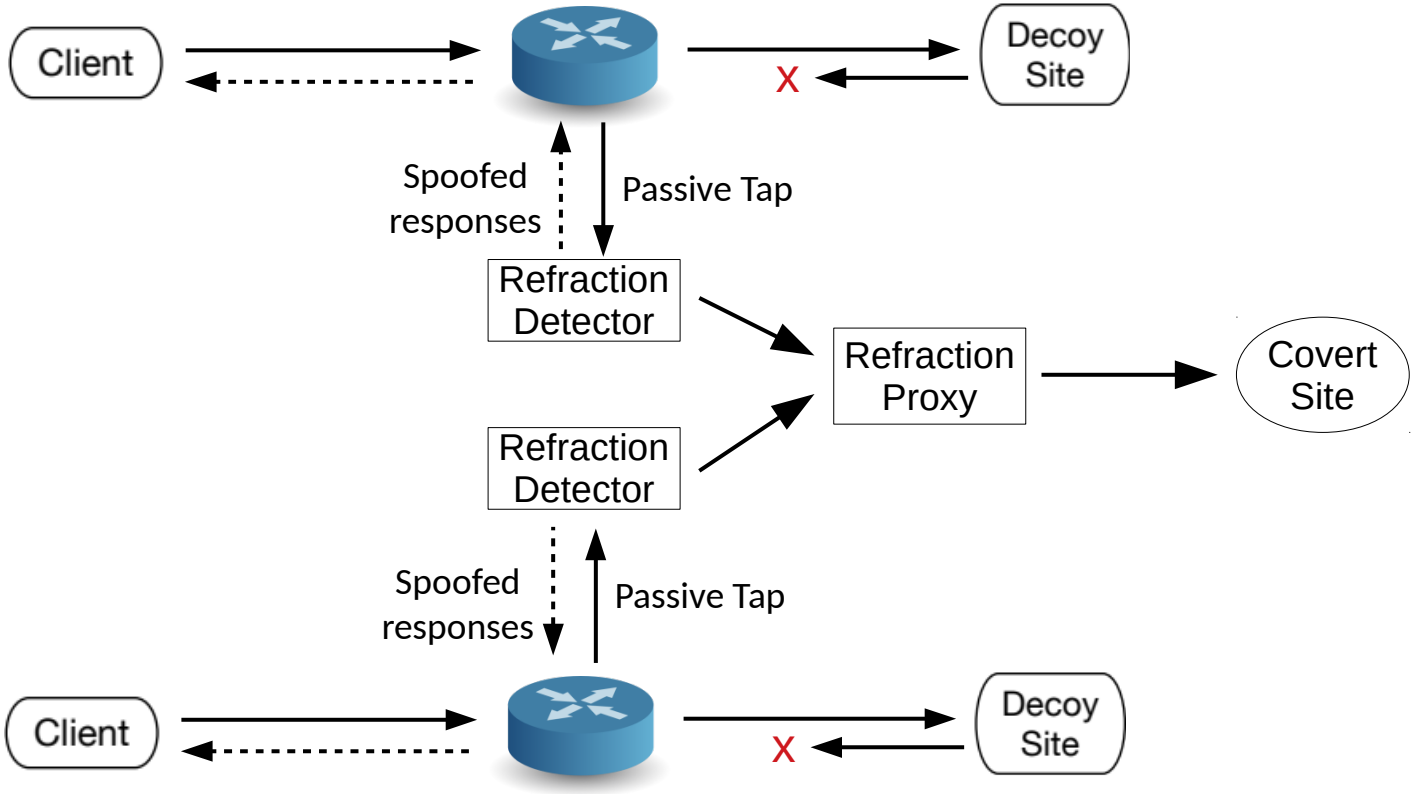

Station Operation & Coordination

- Detectors monitor network taps
- One centralized proxy manager

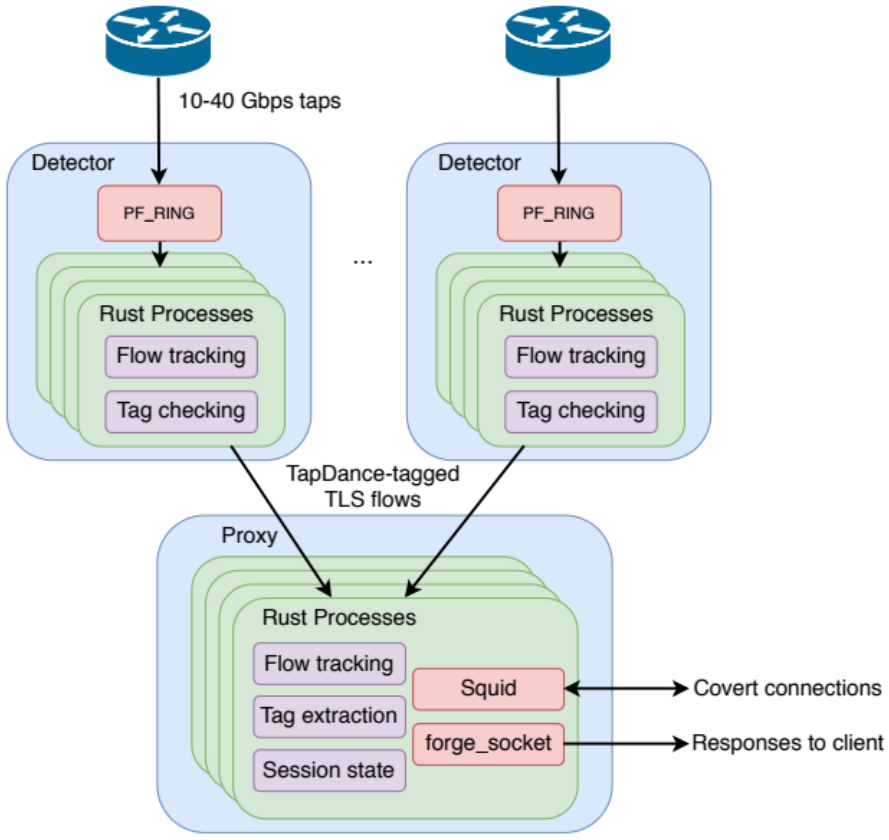

Decoys

- Discovered by scanning port 443 across Merit Address Space
- Filtered to retain only reliable decoys

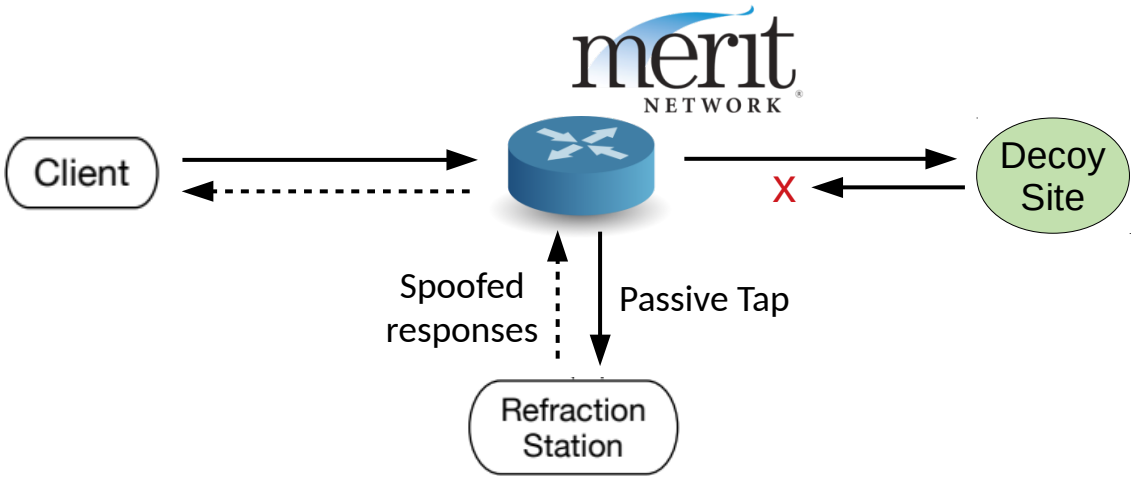

Decoys

- Discovered by scanning port 443 across Merit Address Space
- Filtered to retain only reliable decoys

- Compatible TLS ciphersuite
- Has not requested to be excluded
  - Which decoys actually opted out?

Decoy Collection & Filtering

Decoys

- Discovered by scanning port 443 across Merit Address Space
- Filtered to retain only reliable decoys

Total: 1500 – 2000 Decoys

...

839: www.uofmhosting.net
840: openjericho.com
841: vpn.norcocmh.org
842: afs.msu.edu
843: publicapps.nscl.msu.edu
844: michross-uat.bus.umich.edu
845: www.hillsdale.edu
846: michiganross.umich.edu
847: www.firelab.org
848: kb.lsa.umich.edu
849: charmm-dev.org
850: www.wayne.edu
851: lhfacility.msu.edu
852: umphoto-portals.photos.ns.umich.edu
853: www.umflint.edu
...

TapDance Client

Psiphon Proxy

- Integrated TapDance in Psiphon's Android app
- Deployed to ~560K users in censored countries
- TapDance "Competes" with other proxy protocols transparently to users
  - Meek
  - Tapdance
  - OSSH
  - And other variants

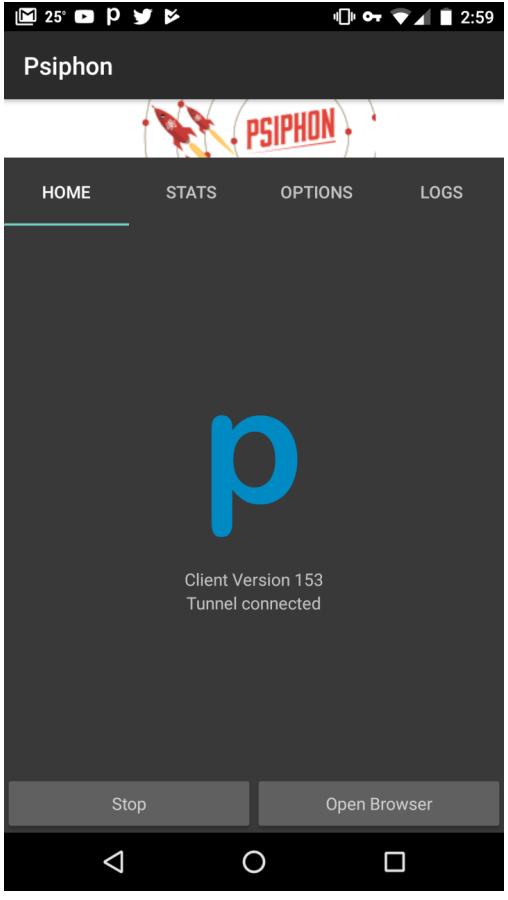

TapDance All-Together

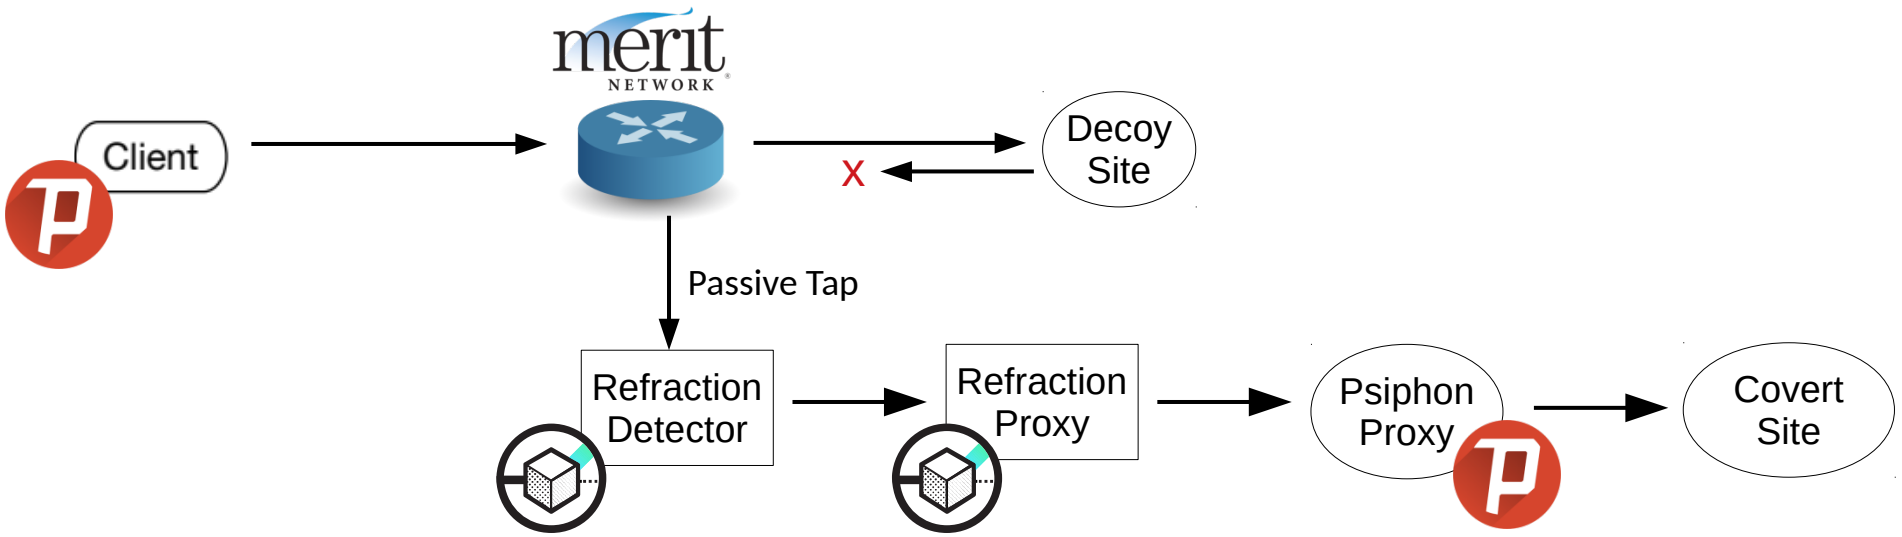

Performance

Operation

Tap Operation

Total System Goodput (Mbps)

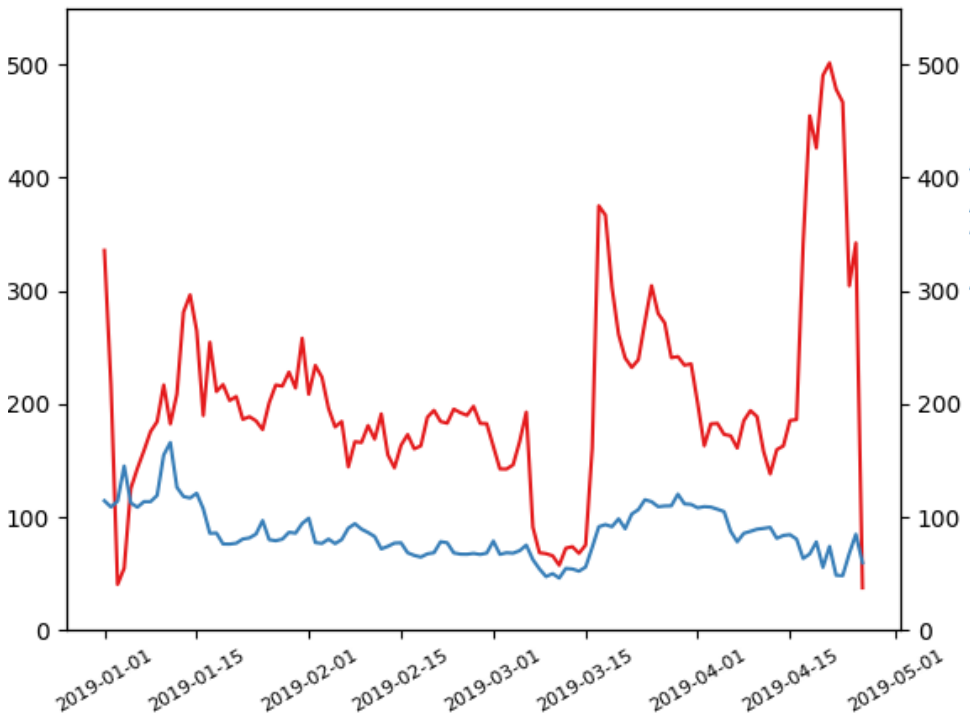

Mean User Goodput (Kbps)

Client Experience

Connection Establishment Latency (s)

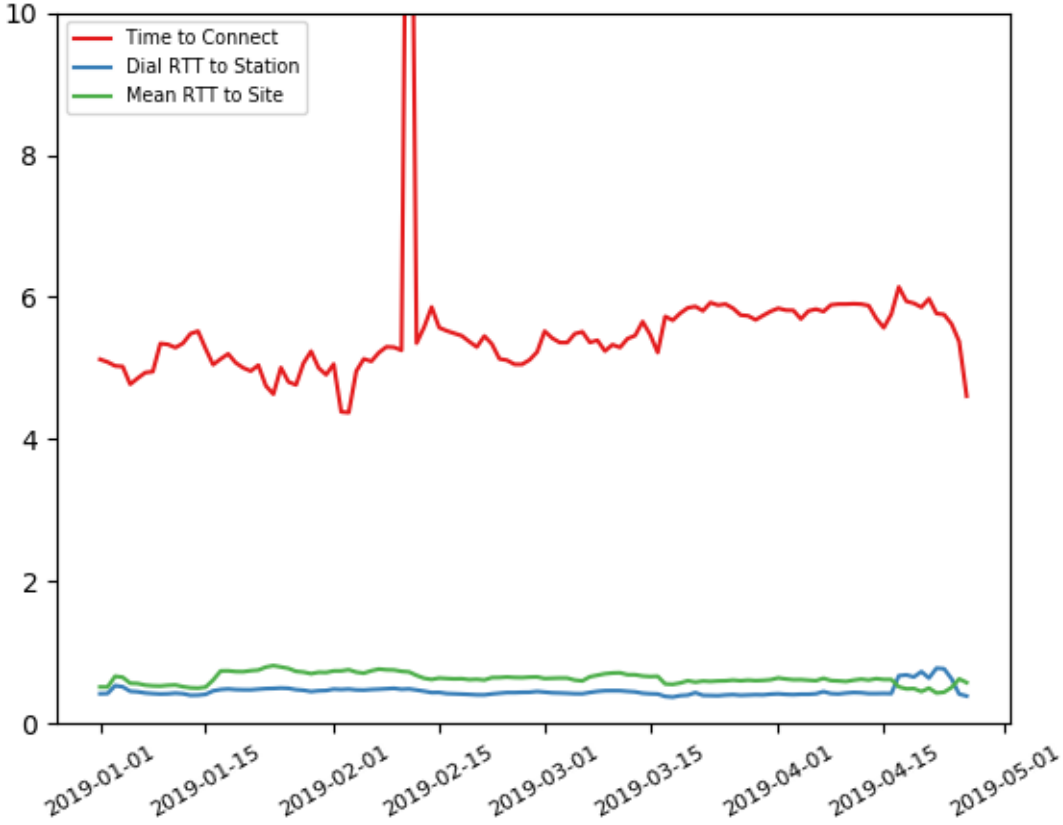

Client Experience

Connection Establishment Latency (s)

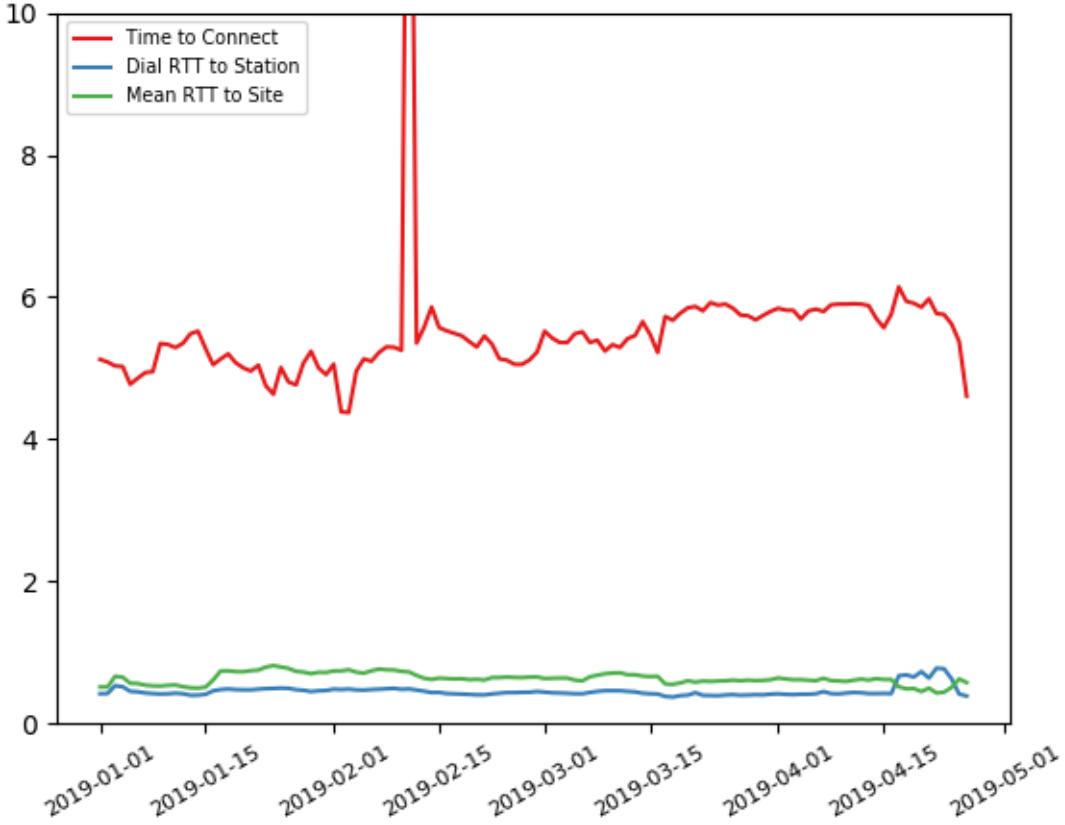

Checkpoint CDF

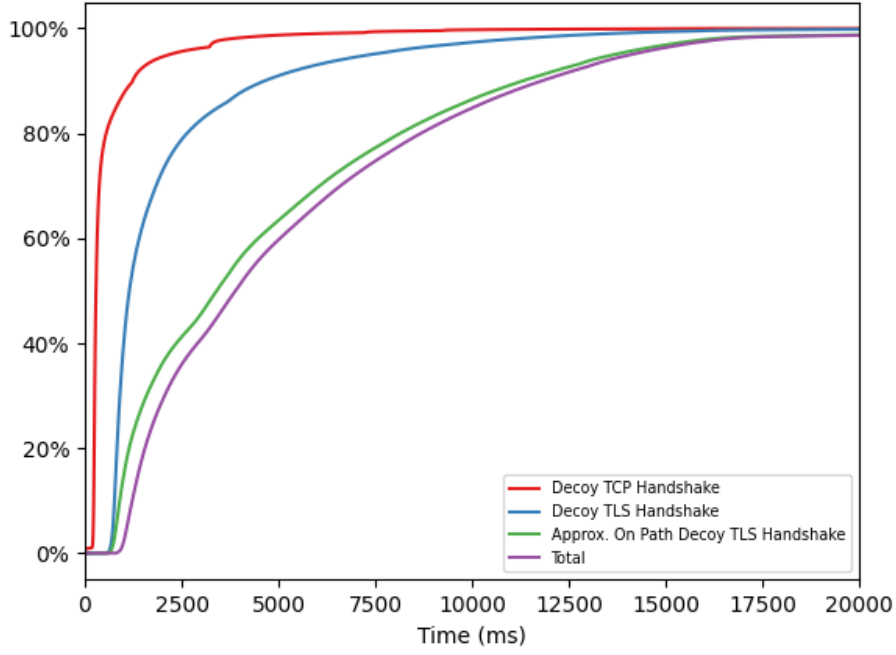

Client Experience

Connection Establishment Latency (s)

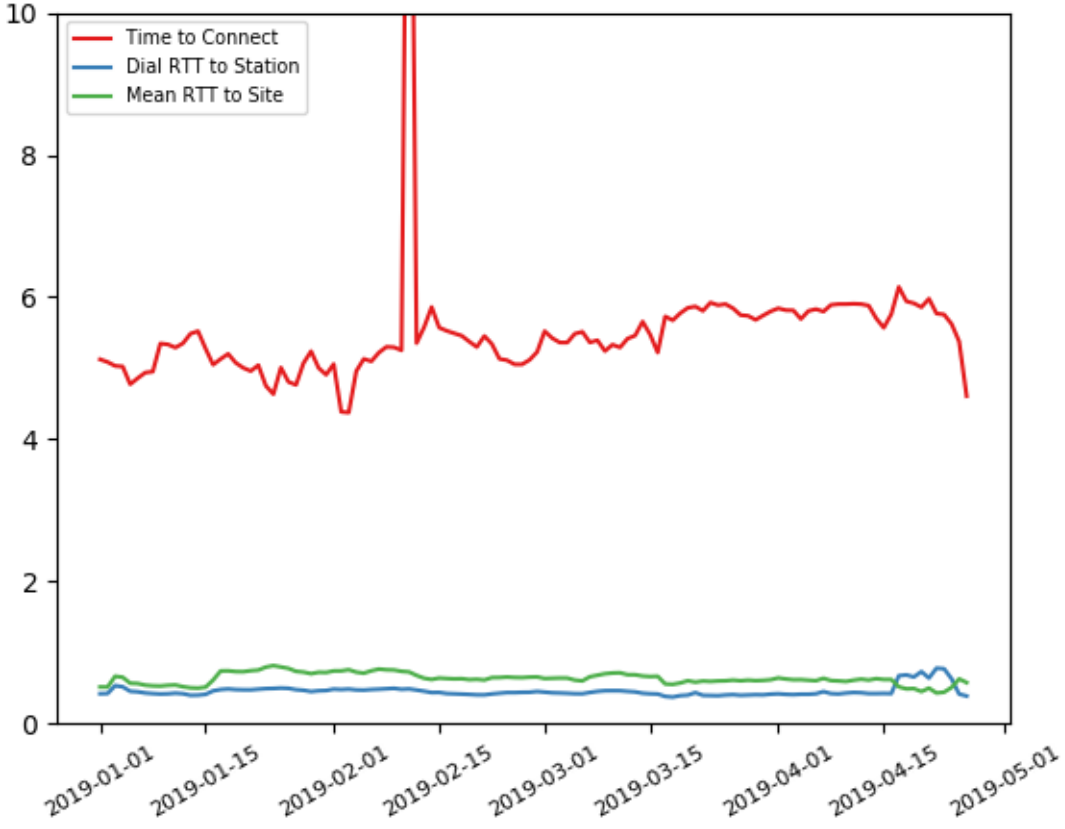

Checkpoint CDF

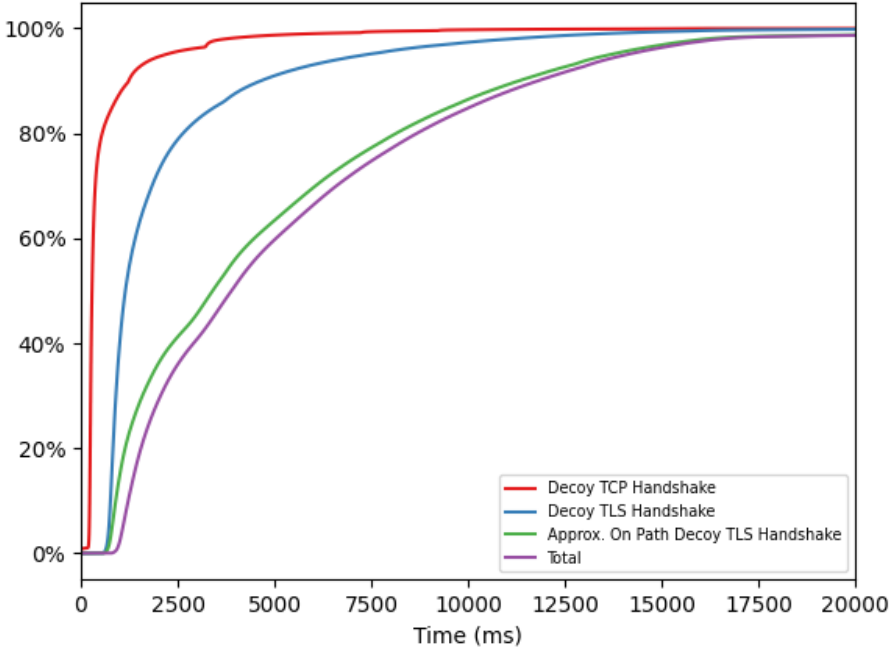

Clients that fail first connection must retry the entire handshake process incurring high latency penalty

Decoys

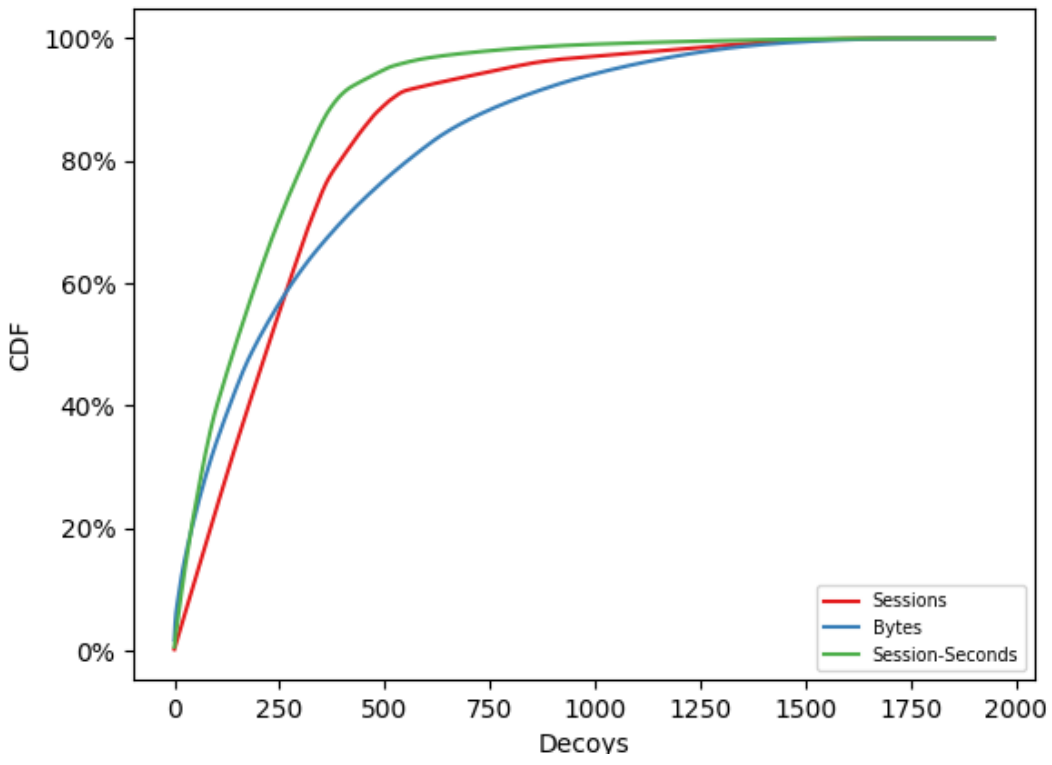

Are client sessions distributed evenly across decoys?

- By number
- By bytes
- By duration

Decoys

Are client sessions evenly distributed across decoys?

- Some worked harder than others

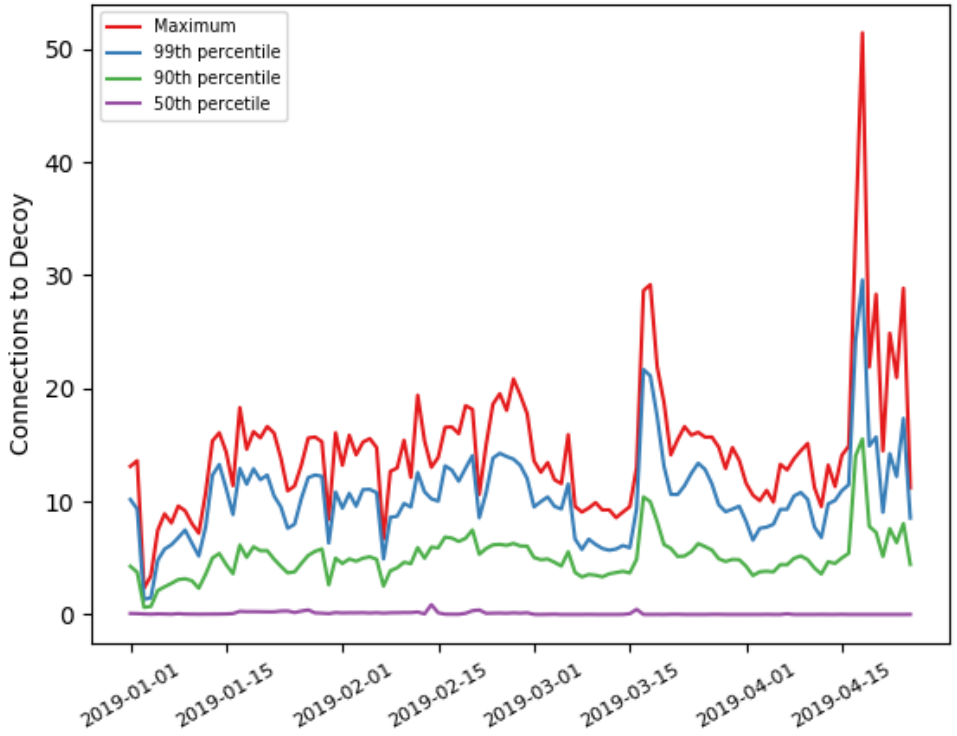

Decoys

Even the decoys that work hardest are not extremely heavily burdened

Rank	Hostname	Mean Concurrent Connections	Connections	Average Transfer Rate (bps)
1		13.24	163,991	1140.74
2		12.76	167,277	994.44
3		12.00	167,144	990.74
4		10.75	167,507	846.30
5		10.70	128,691	1230.55
6		10.68	151,699	744.44
7		10.48	127980	1193.52
8		10.42	161,146	847.22
9		10.41	127,971	1240.74
10		10.34	127,948	1173.15

Proxy Partner

Psiphon Proxy

- Meek
- Tapdance
- OSSH
- etc.

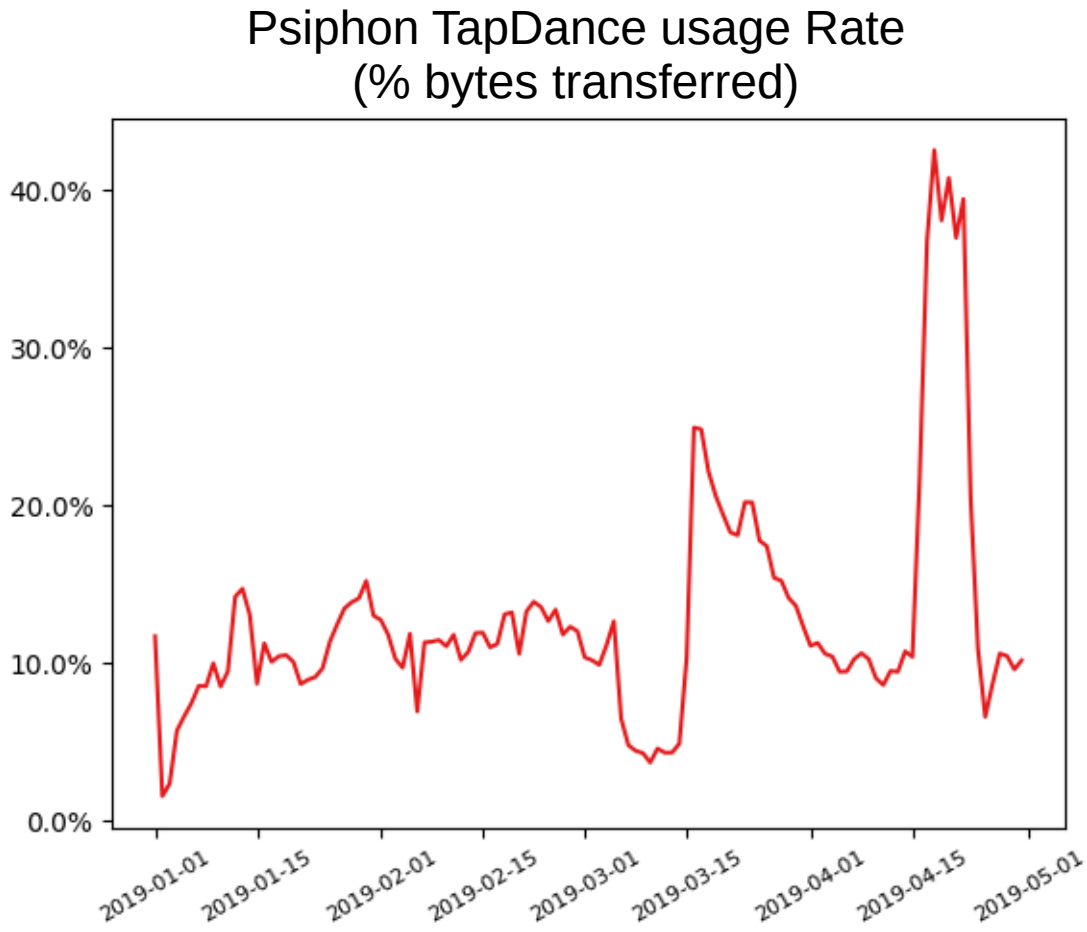

Proxy Partner

Psiphon Proxy

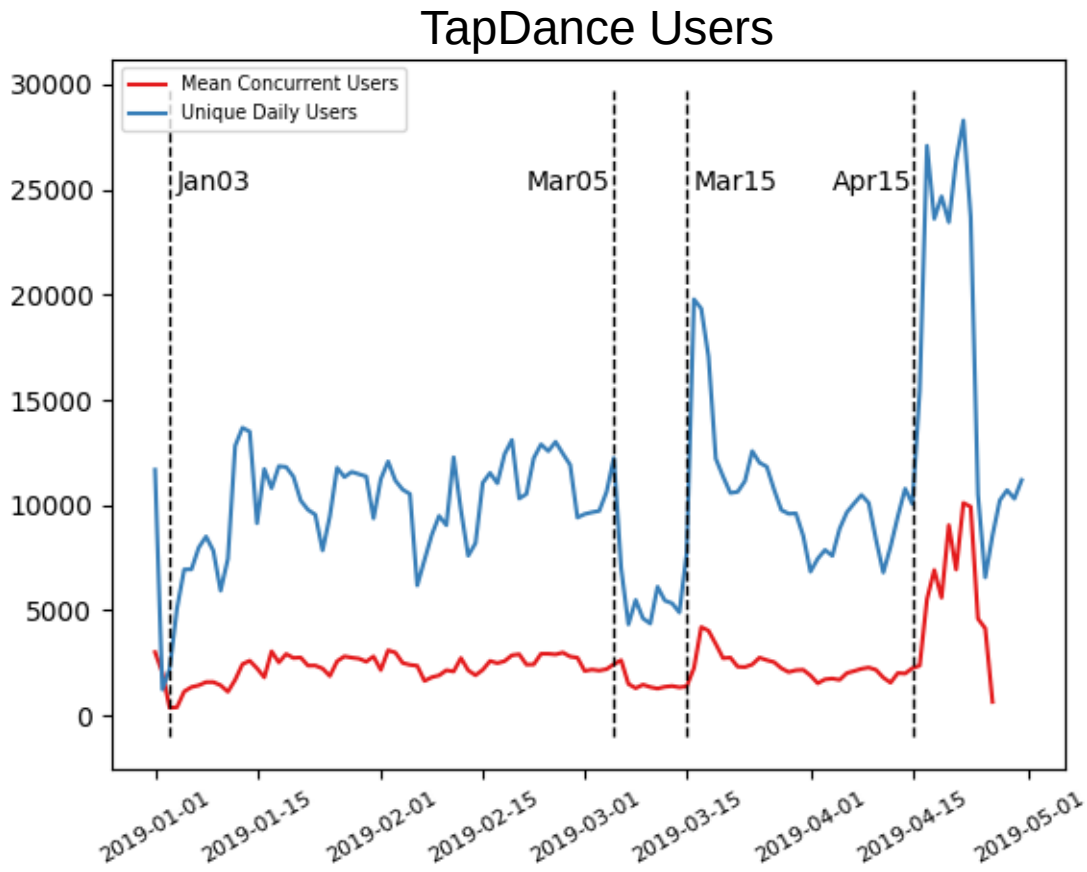

Censorship events

TapDance supplements other proxies under censorship

Jan 03 - Domain fronting methods are unblocked for a short period of time.

Mar 05 - Direct proxy methods are unblocked favoring alternative Psiphon transports

Mar 15 - Direct and domain fronting are blocked once more

Apr 15 - New techniques for blocking previously reliable proxies rolled out

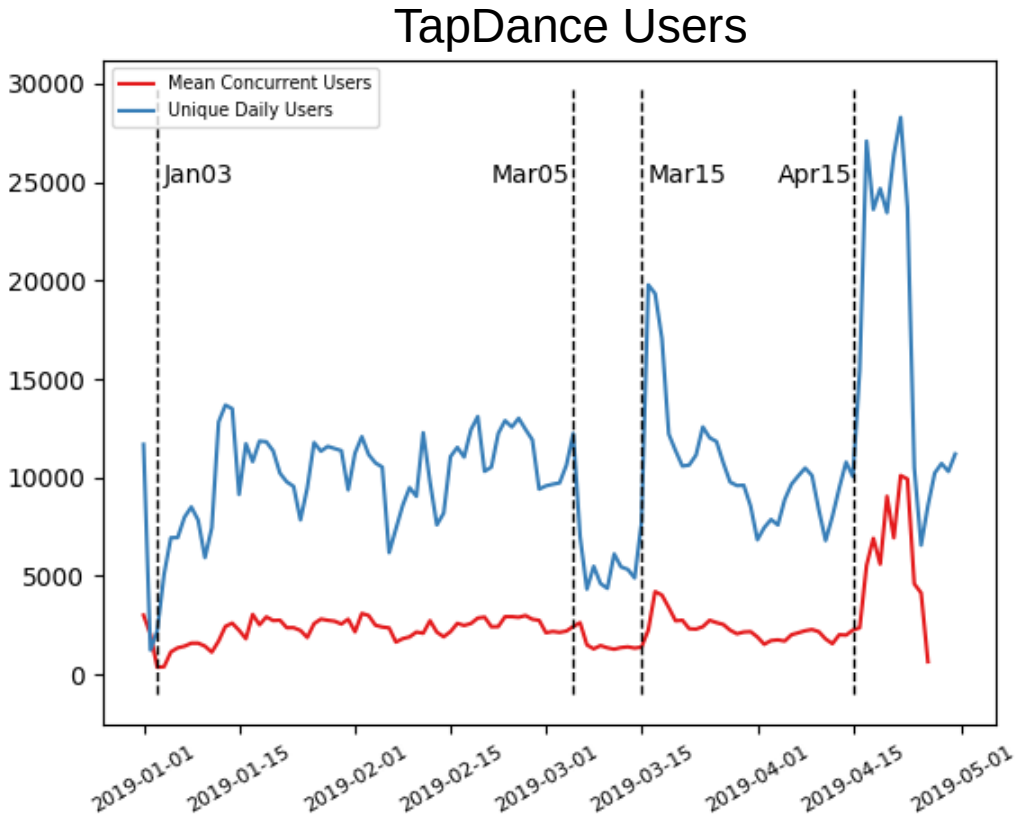

Decoys

Session stats

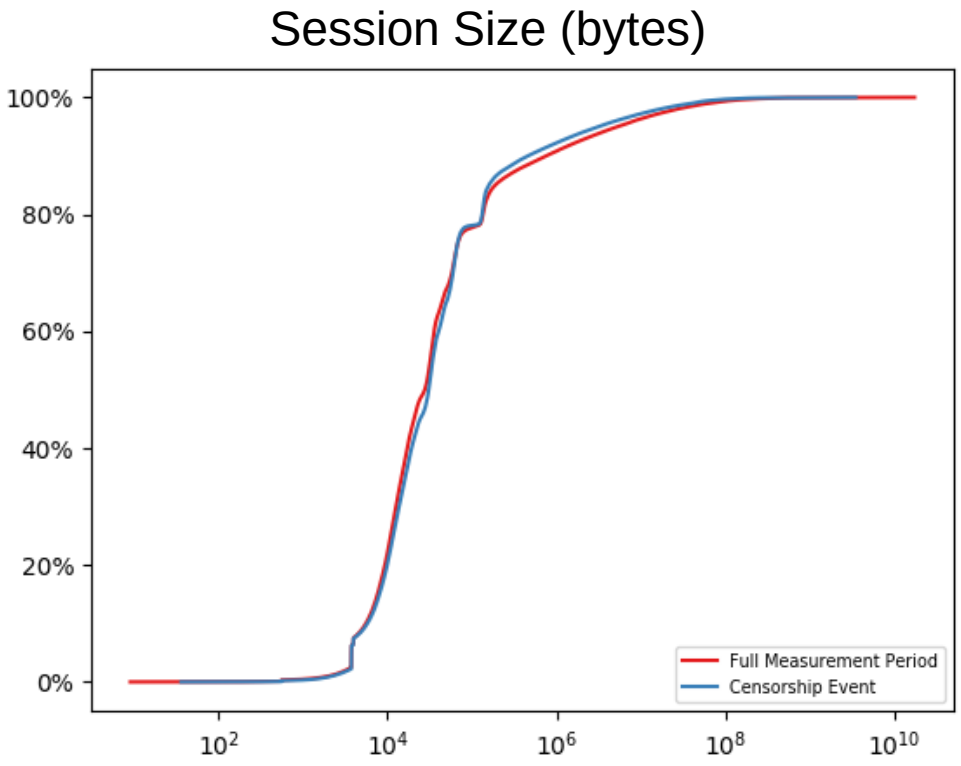

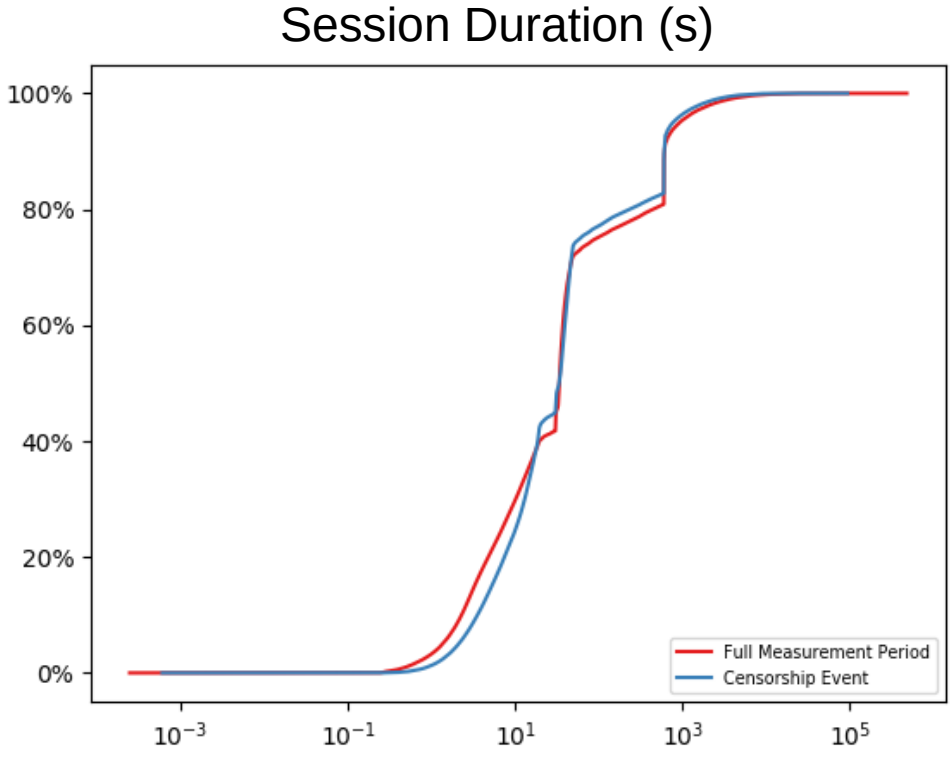

Lessons

Decoy Failure

Selecting decoys is difficult

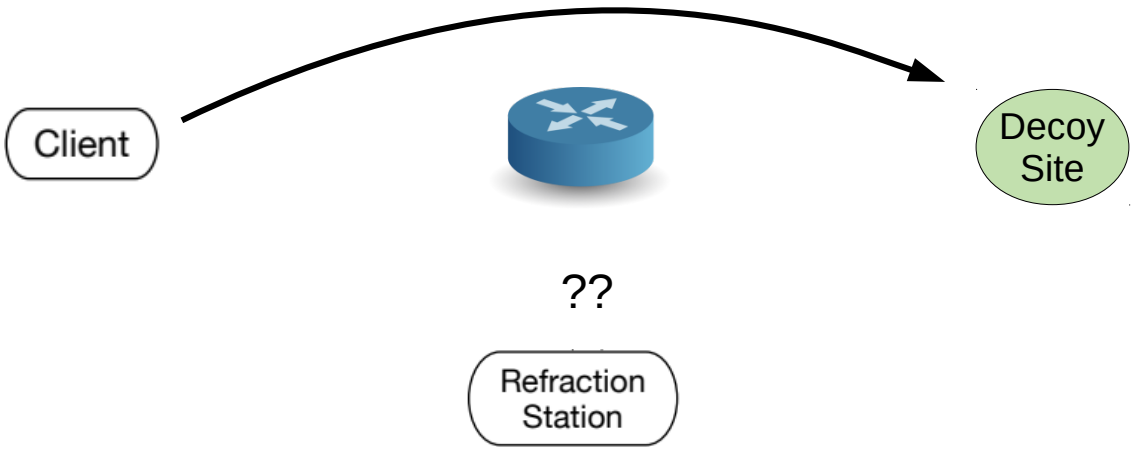

Decoy Failure

Selecting decoys is difficult

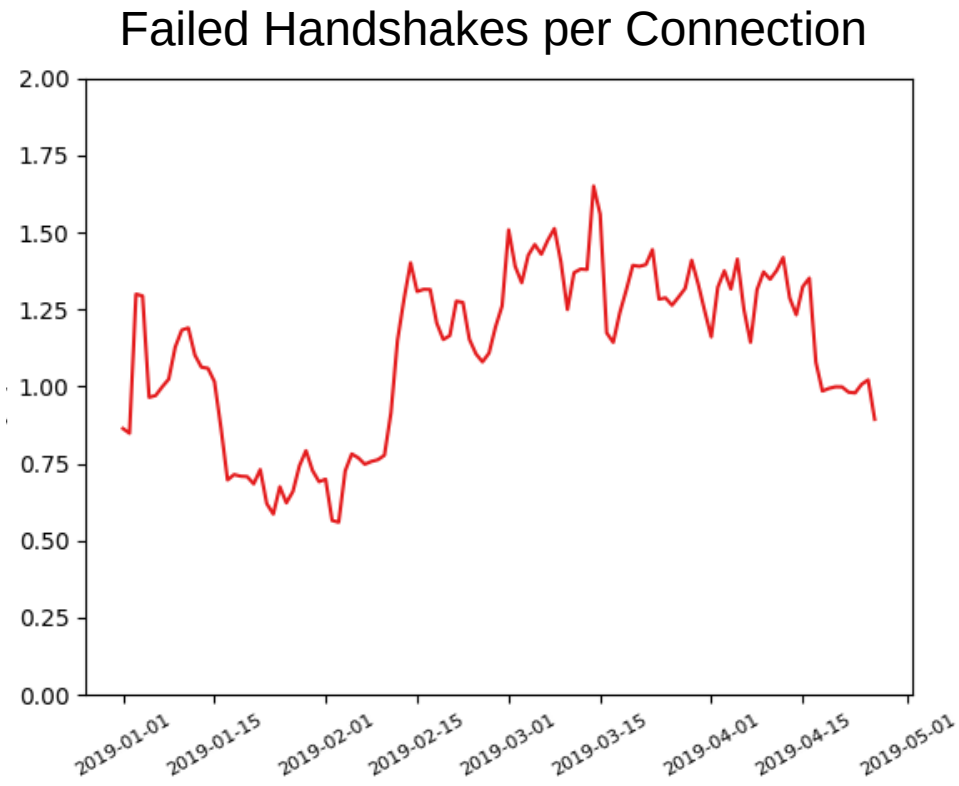

TapDance Limitations

TapDance connection limitations

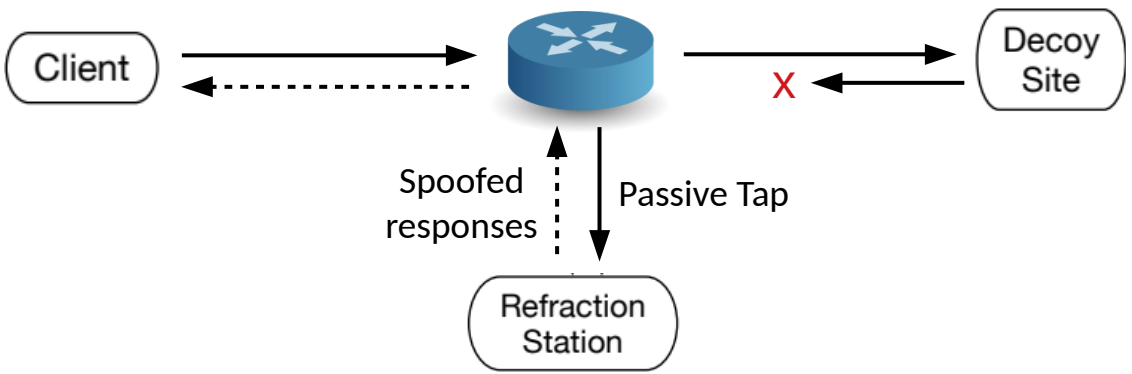

TapDance Limitations

TapDance connection limitations

- Client sends something to silence the decoy
- Station pretends to be the decoy while the connection stays open
  - Connection upload limit
  - Connection duration limit

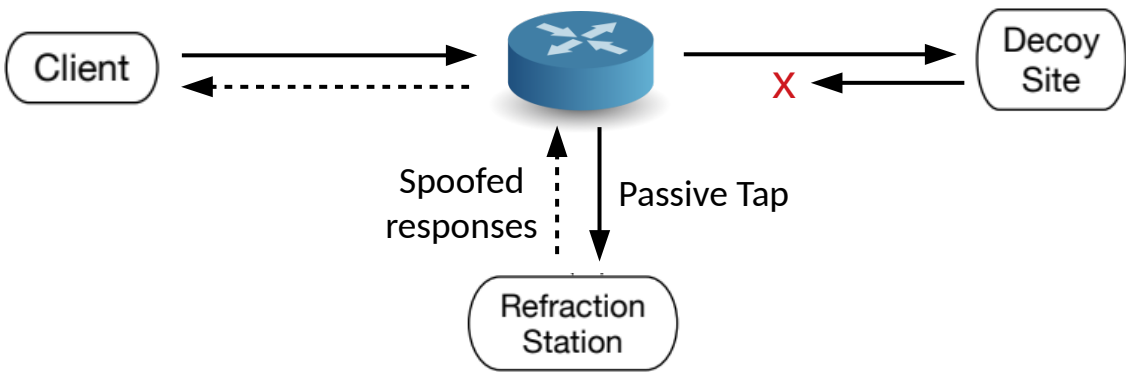

TapDance Limitations

DittoTap – Slitheen + TapDance

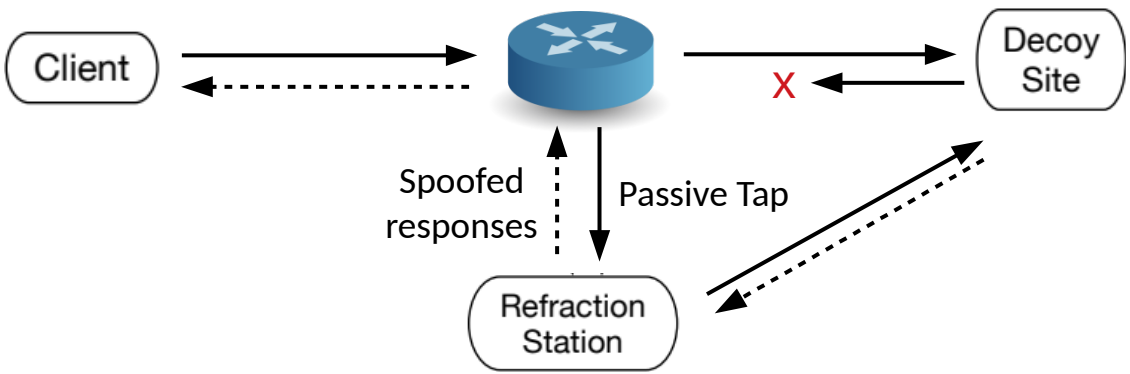

TapDance Limitations

Split Flows – Slitheen + TapDance

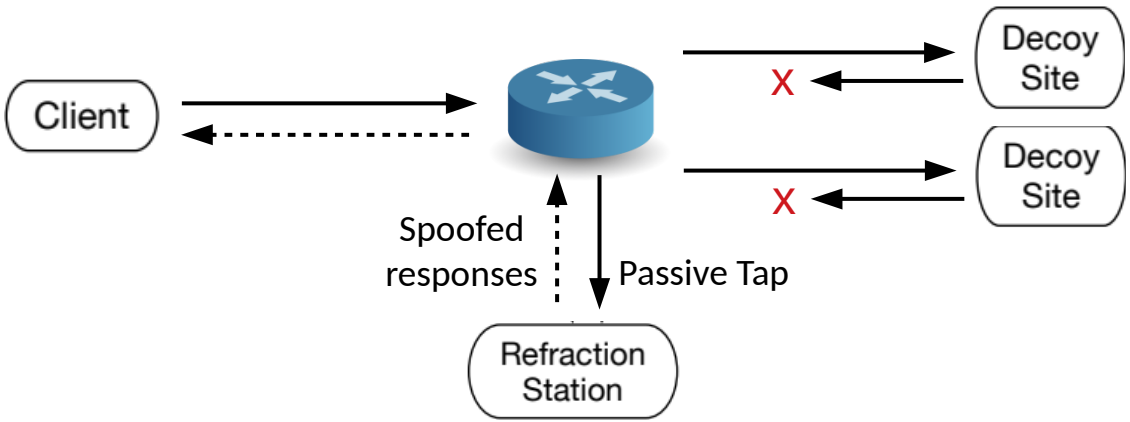

TapDance Limitations

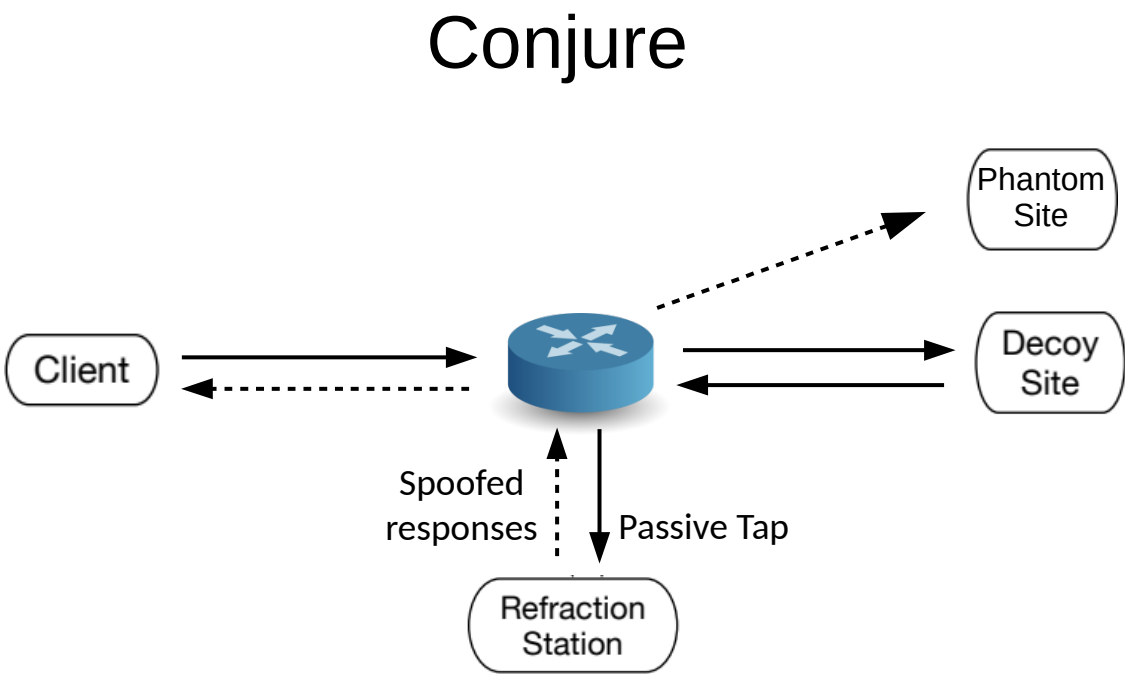

Addressing Partner Concerns

Minimal Production Impacts

Manageable Decoy Loads

No Observed Censor Retaliation

Take Away

TapDance supplements other proxies in the event of censorship events by providing uniquely censorship resistant service.

Running Refraction Networking for Real

Benjamin VanderSloot, Sergey Frolov, Jack Wampler,
Sze Chuen Tan, Irv Simpson, Michalis Kallitsis,
J. Alex Halderman, Nikita Borisov, and Eric Wustrow