CENSORSHIP RESISTANCE

CMSC 414

APR 17 2018

CENSORSHIP COMES IN MANY FORMS

DROPPING PACKETS

Network operators: Block traffic in their own networks/countries

Off-path attackers: Inject TCP RST packets (next week)

Routing-capable adversaries: Can influence routes on the Internet

Black-holing: Announce a low-cost path, drop traffic
https://www.youtube.com/watch?v=IzLPKuAOe50

MONITORING TRAFFIC

Boomerang routing: Source/destination close, but route goes through a country known to eavesdrop

DEANONYMIZATION

Identifying and going after **whistleblowers**

MISDIRECTING TRAFFIC

DNS injection: Send back false DNS responses

ENEMIES OF THE INTERNET

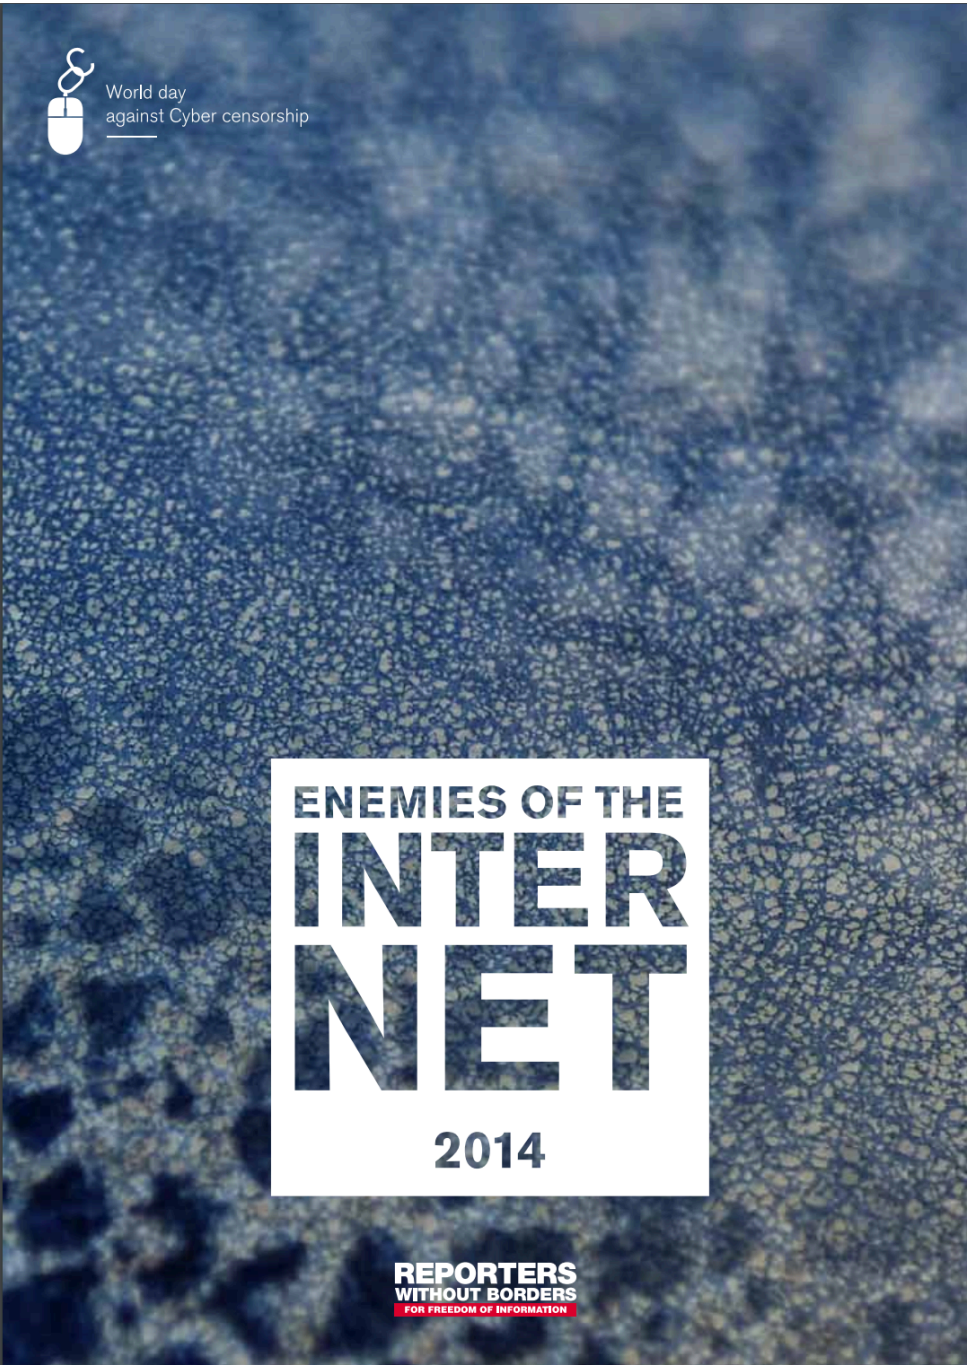

~Annual report by
Reporters without Borders

2014

- *Syria*
- *Russia*
- *Saudia Arabia*
- *UAE*
- *Cuba*
- *Belarus*
- *Pakistan*
- *Vietnam*
- *Turkmenistan*
- *Sudan*
- *Iran*
- *Bahrain*
- *USA*
- *UK*
- *Uzbekistan*
- *India*
- *China*
- *North Korea*
- *Ethiopia*
- *Surveillance dealers*

ENEMIES OF THE INTERNET

Français
Español

World day
against Cyber censorship

REPORTERS WITHOUT BORDERS
FOR FREEDOM OF INFORMATION

Enemies of the Internet

Home | Enemies of the Internet | The Map | Recommendations | Take Action ! | Archives

USA: NSA symbolises intelligence services' abuses

In June 2013, computer specialist Edward Snowde
and British intelligence services. Snowden, who wo
confidential documents, later exposed more target
leaders and diplomats of allied countries. Activist
the Obama administration, as the newspapers *The*
the surveillance. The main player in this vast surve
Agency (NSA) which, in the light of Snowden's reve
intelligence agencies. Against this background, those involved in reporting on security issues have found their sources under increasing pressure.

The U.S. edition of *The Guardian* is still able to pu
is not, but the country of the First Amendment ha
security. U.S. surveillance practices and decryptio
especially those who work with sensitive sources
pressure.

The NSA

Based in Fort Meade, Virginia, the NSA has always operated behind a wall of secrecy. According to legend, its acronym was jokingly said to mean "No Such Agency" because its work took place far from the eyes of U.S.

Pressure on journalists, sources and whistleblowers

The Obama administration has shown itself to be willing to interpret the protection of national security in a broad and abusive manner, at the expense of freedom of information. A witch-hunt was launched against journalists' sources who disclosed confidential information about the powers of the state.

The NSA has been helped in its determined pursuit of WikiLeaks by GCHQ, since all visitors to the website have been monitored by the British agency's TEMPORA surveillance system. Their IP addresses and the terms entered in search engines to access the site are intercepted and recorded.

COLLATERAL DAMAGE OF INTERNET CENSORSHIP

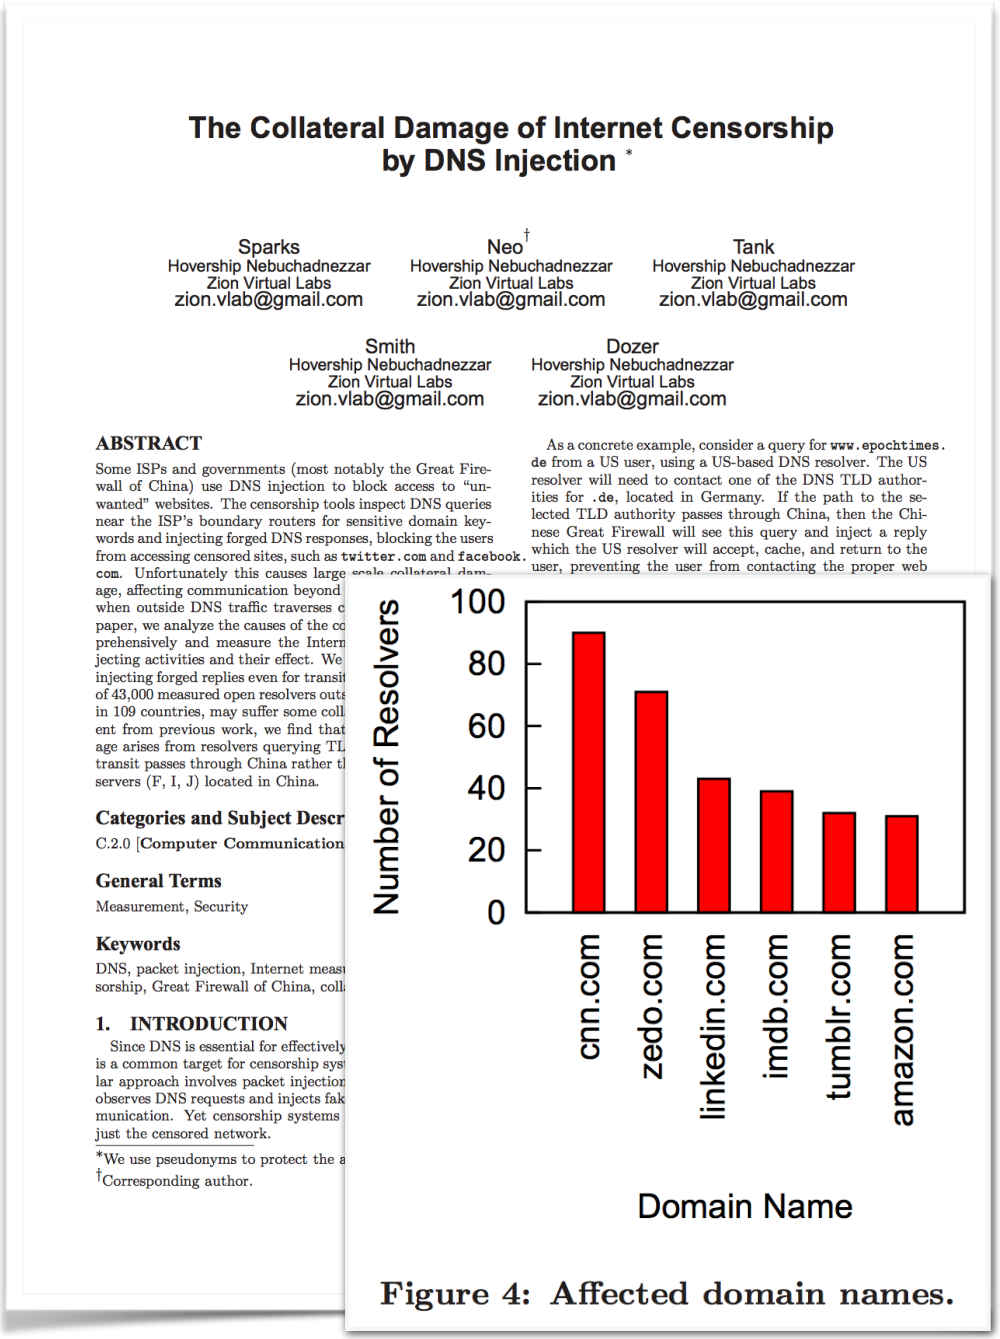

The Collateral Damage of Internet Censorship by DNS Injection *

Sparks
Hovership Nebuchadnezzar
Zion Virtual Labs
zion.vlab@gmail.com

Neo†
Hovership Nebuchadnezzar
Zion Virtual Labs
zion.vlab@gmail.com

Tank
Hovership Nebuchadnezzar
Zion Virtual Labs
zion.vlab@gmail.com

Smith
Hovership Nebuchadnezzar
Zion Virtual Labs
zion.vlab@gmail.com

Dozer
Hovership Nebuchadnezzar
Zion Virtual Labs
zion.vlab@gmail.com

ABSTRACT

Some ISPs and governments (most notably the Great Firewall of China) use DNS injection to block access to "unwanted" websites. The censorship tools inspect DNS queries near the ISP's boundary routers for sensitive domain keywords and injecting forged DNS responses, blocking the users from accessing censored sites, such as `twitter.com` and `facebook.com`. Unfortunately this causes large scale collateral dam-
age, affecting communication beyond
when outside DNS traffic traverses c
paper, we analyze the causes of the co
prehensively and measure the Intern
jecting activities and their effect. We
injecting forged replies even for transit
of 43,000 measured open resolvers outs
in 109 countries, may suffer some coll
ent from previous work, we find that
age arises from resolvers querying TL
transit passes through China rather t
servers (F, I, J) located in China.

As a concrete example, consider a query for `www.epochtimes.de` from a US user, using a US-based DNS resolver. The US resolver will need to contact one of the DNS TLD authorities for `.de`, located in Germany. If the path to the selected TLD authority passes through China, then the Chinese Great Firewall will see this query and inject a reply which the US resolver will accept, cache, and return to the user, preventing the user from contacting the proper web

Categories and Subject Descr

C.2.0 [**Computer Communication**

General Terms

Measurement, Security

Keywords

DNS, packet injection, Internet meas
sorship, Great Firewall of China, coll

1. INTRODUCTION

Since DNS is essential for effectively
is a common target for censorship sys
lar approach involves packet injectio
observes DNS requests and injects fak
munication. Yet censorship systems
just the censored network.

*We use pseudonyms to protect the a
†Corresponding author.

Figure 4: Affected domain names.

China censors the traffic to or from those within its borders *Known*

They do this via DNS injection *Known / expected*

They do this to *any traffic* that traverses its borders *Not known*

More traffic traverses China's borders than we realized *Oh geez..*

CIRCUMVENTING THE CONSTITUTION

LOOPHOLES FOR CIRCUMVENTING THE CONSTITUTION: UNRESTRAINED BULK SURVEILLANCE ON AMERICANS BY COLLECTING NETWORK TRAFFIC ABROAD

*Axel Arnbak and Sharon Goldberg**

Cite as: Axel Arnbak and Sharon Goldberg,
Loopholes for Circumventing the Constitution: Unrestrained Bulk Surveillance on Americans by Collecting Network Traffic Abroad,
21 MICH. TELECOMM. & TECH. L. REV. 317 (2015).
This manuscript may be accessed online at repository.law.umich.edu.

ABSTRACT

This Article reveals interdependent legal and technical loopholes that the US intelligence community could use to circumvent constitutional and statutory safeguards for Americans. These loopholes involve the collection of Internet traffic on foreign territory, and leave Americans as unprotected as foreigners by current United States (US) surveillance laws. This Article will also describe how modern Internet protocols can be manipulated to deliberately divert American's traffic abroad, where traffic can then be collected under a more permissive legal regime (Executive Order 12333) that is overseen solely by the executive branch of the US government. Although the media has reported on some of the techniques we describe, we cannot establish the extent to which these loopholes are exploited in practice.

An actionable short-term remedy to these loopholes involves updating the antiquated legal definition of "electronic surveillance" in the Foreign Intelligence Surveillance Act (FISA), that has remained largely intact since 1978. In the long term, however, a fundamental reconsideration of established principles in US surveillance law is required, since

* Axel Arnbak is a Faculty Researcher at the Institute for Information Law, University of Amsterdam and a Research Affiliate at the Berkman Center for Internet & Society, Harvard University. Sharon Goldberg is Associate Professor of Computer Science, Boston University and a Research Fellow, Sloan Foundation. She gratefully acknowledges the support of the Sloan Foundation. Both authors thank Timothy H. Edgar, Ethan Heilman, Susan Landau, Alex Marthews, Bruce Schneier, Haya Shulman, Marcy Wheeler and various attendees of the PETS'14 and TPRC'14 conferences for discussions and advice that have greatly aided this work. Alexander Abdo, David Choffnes, Nico van Eijk, Edward Felten, Daniel K. Gillmore, Jennifer Rexford, Julian Sanchez and the anonymous reviewers for HotPETS'14 each provided insightful comments on drafts of this Article. Views and errors expressed in this Article remain the sole responsibility of the authors. This Article was submitted on September 1, 2014 and a brief update was concluded on December 26, 2014. All URLs have been checked on this date. An earlier version of this Article was first posted online on June 27, 2014.

LEGAL REGIMES

Patriot Act

Foreign Intelligence Surveillance Act (FISA)

EO 12333

WHAT CAN BE MONITORED?

Communication with foreign entities

DO ROUTERS COUNT?

What if the US routed traffic out of its borders, then back in — would this count as communication with a foreign entity?

THIS PAPER: YES, PROBABLY

So any traffic could be easily monitored

BLOCKING TOR

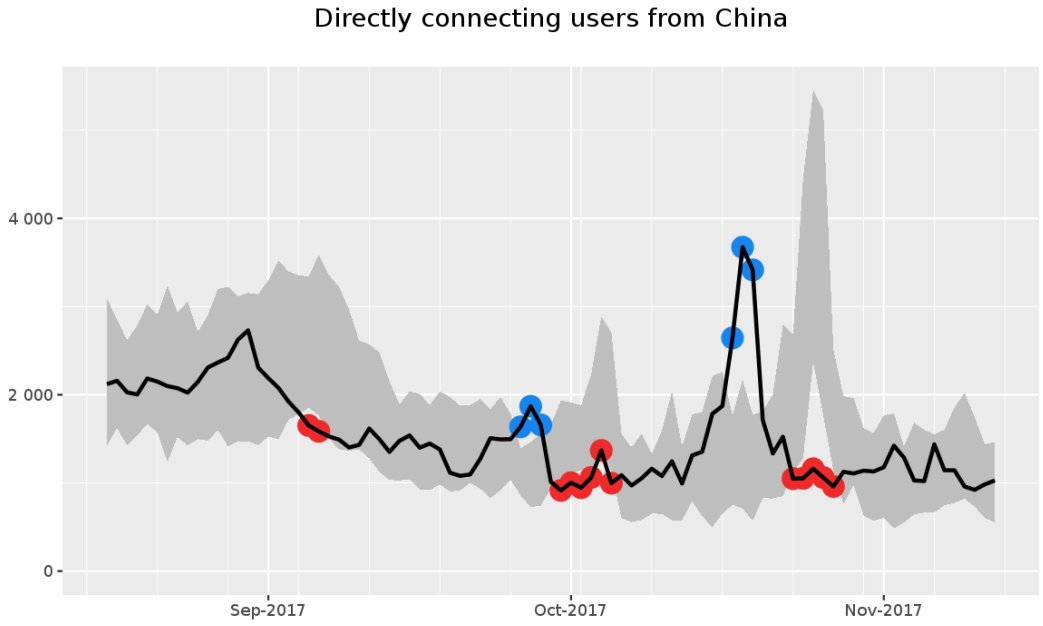

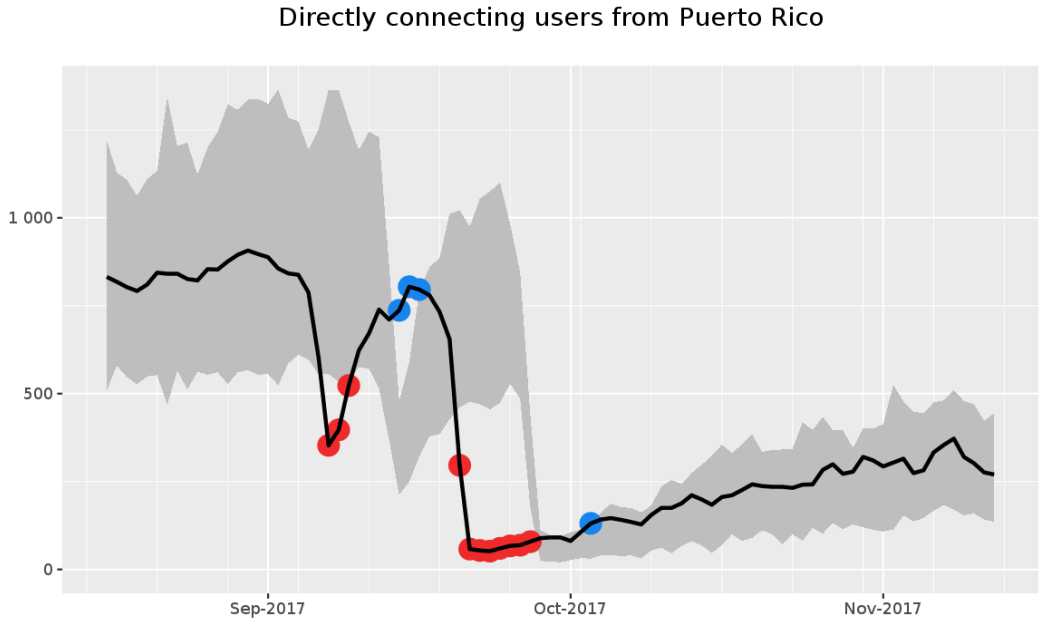

The Tor Project - https://metrics.torproject.org/

Tor | Metrics

Estimate the number of users on day *i* based on previous days' users

Gray area: Range of estimated users; Usage naturally fluctuates

Downturn event: Drops below
Possibly indicates censorship

Upturn event: Rises above "normal"
Possibly indicates circumvention

HOW TO BLOCK TOR

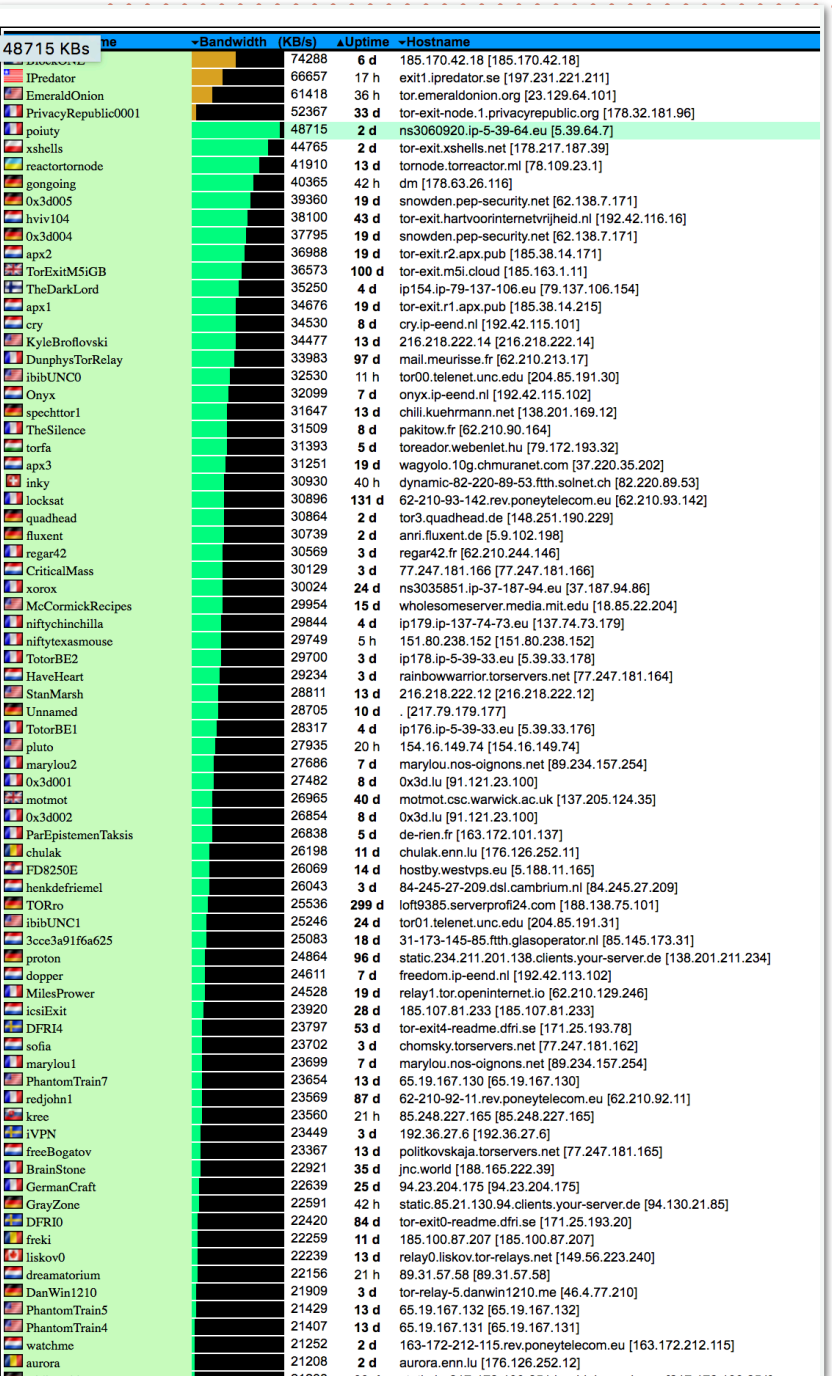

Option 1: Get a list of all Tor nodes
Insert them as firewall rules

Bridge nodes: Tor does not list some nodes;
Users must learn them out of band

Censors ca discover them by actively probing

Scan IP addresses, sending protocol-specific messages: handshake (TLS, obfs), Versions (Tor), HTTPS Post (SoftEther), HTTP GET (AppSpot)

HOW TO BLOCK TOR

Option 2: IP-based reputation schemes;
Will eventually block exit nodes because attackers **launder** their attack traffic thru Tor

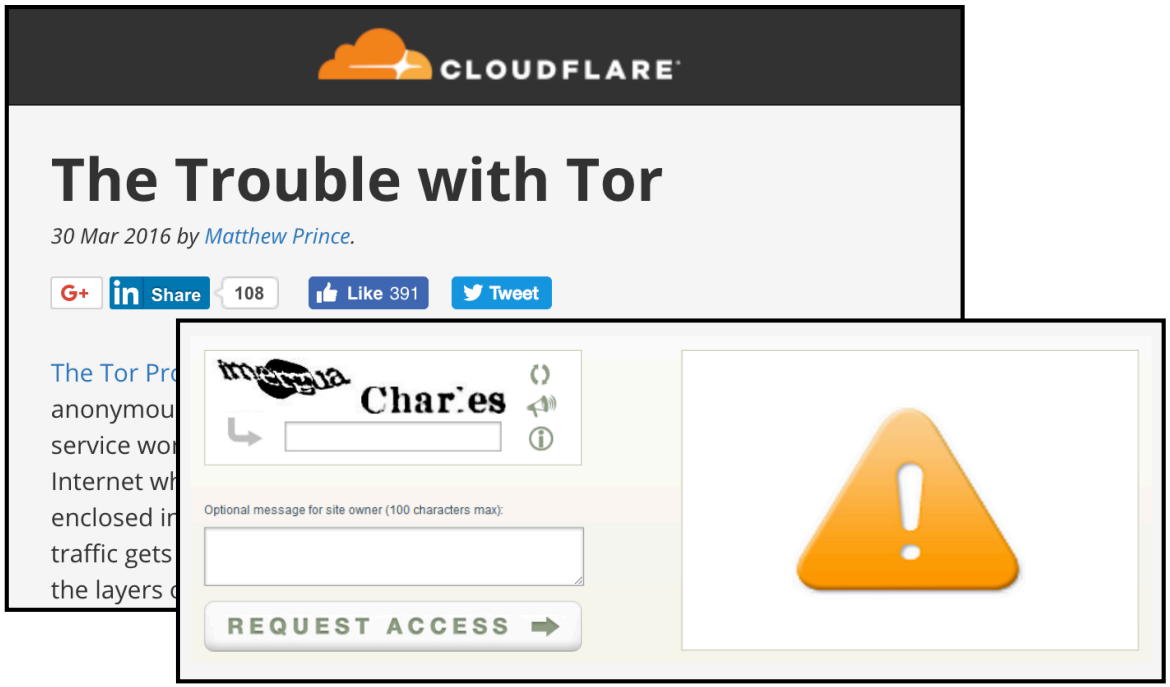

DECOY ROUTING

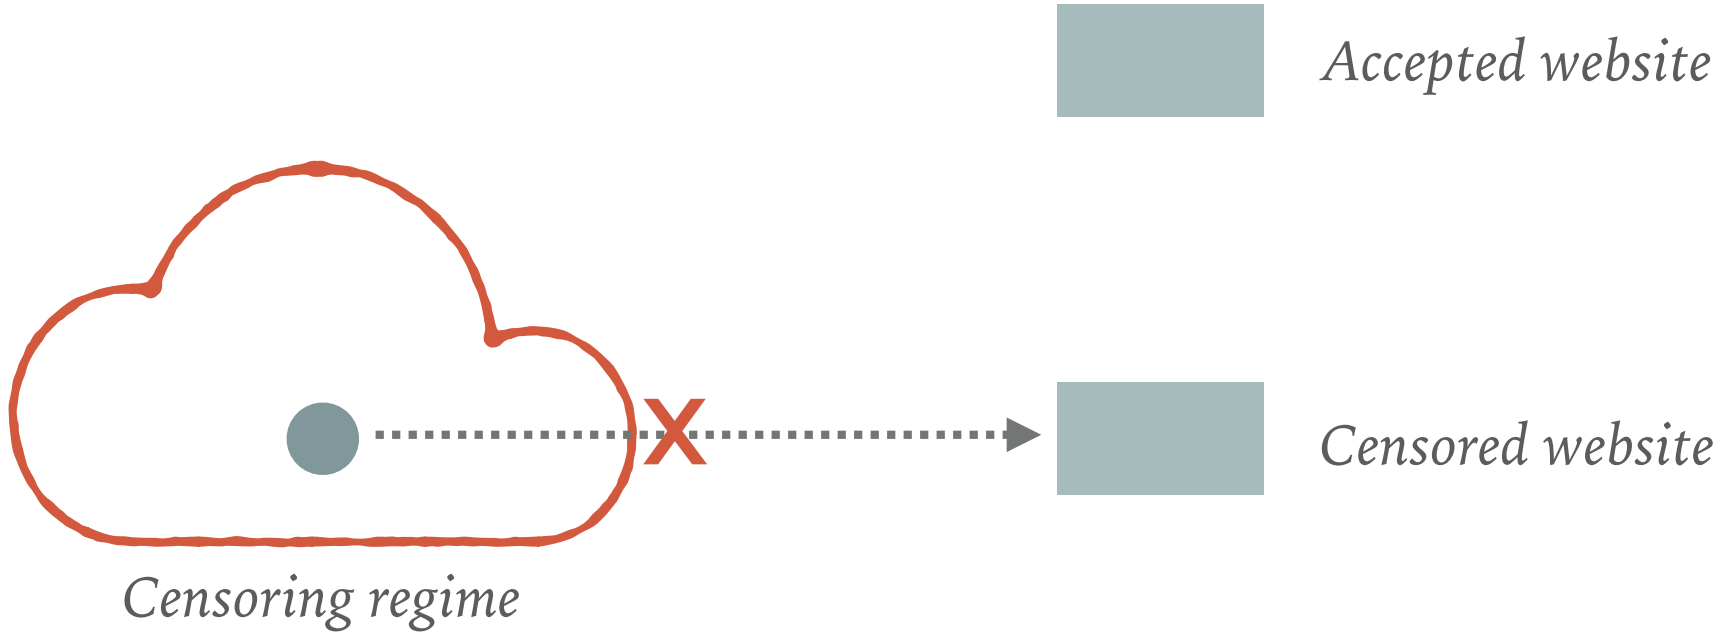

DECOY ROUTING

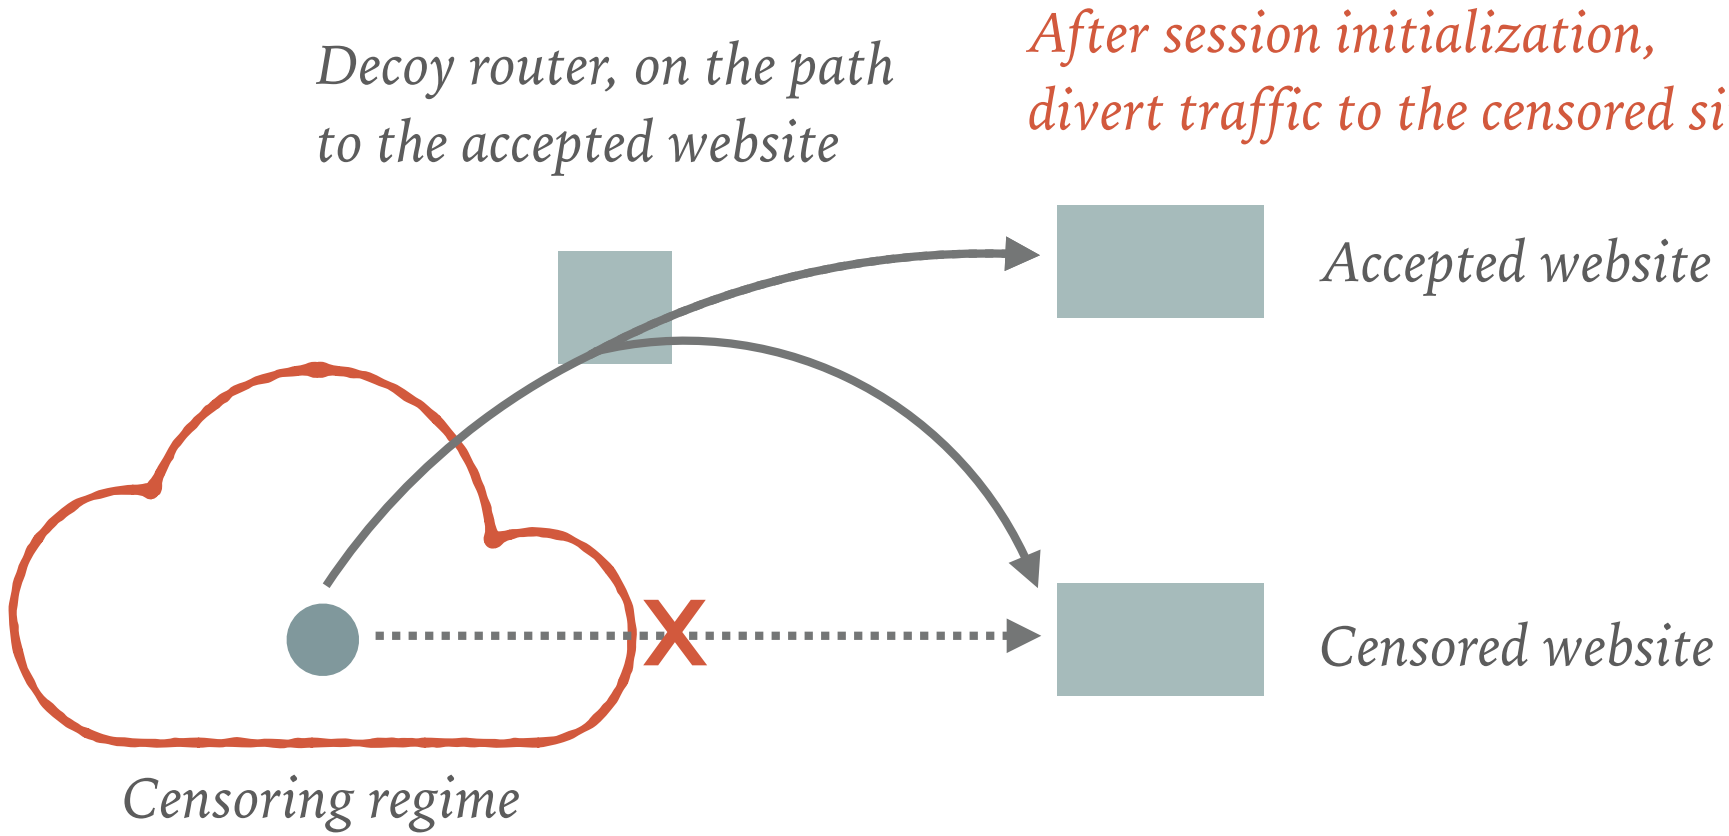

How does the decoy router know the true destination but the censor doesn't?

Client includes "tags" in TLS handshakes that only the decoy router can identify

DECOY ROUTING TAGS

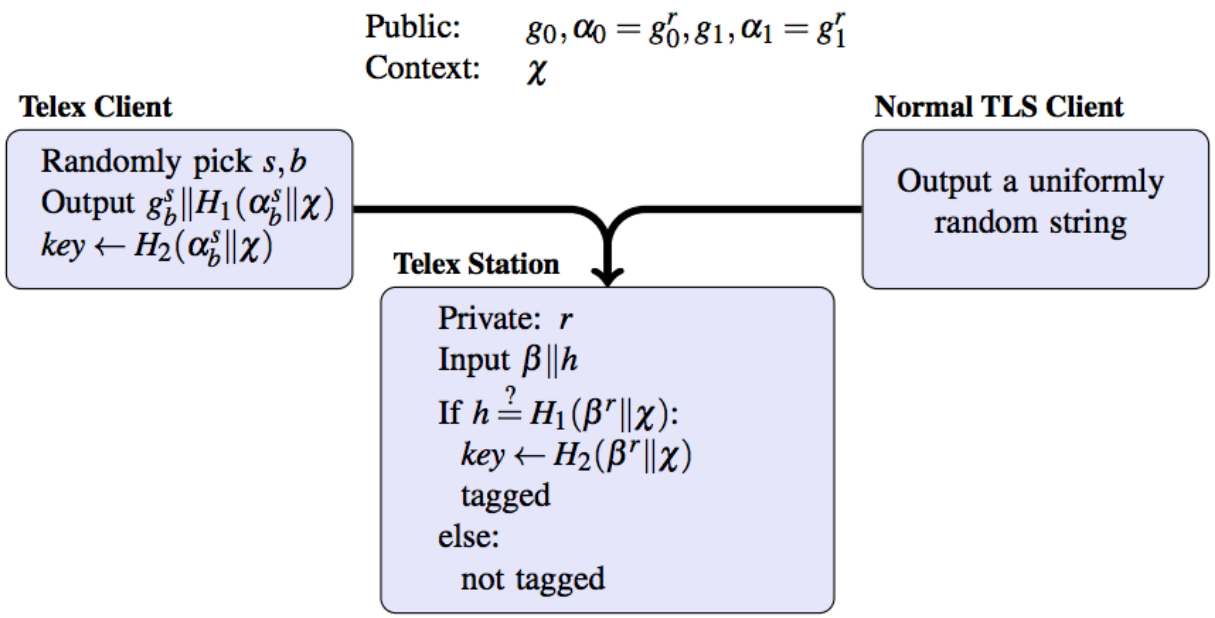

Figure 2: **Tag creation and detection** — Telex intercepts TLS connections that contain a steganographic tag in the ClientHello message's nonce field (normally a uniformly random string). The Telex client generates the tag using public parameters (shown above), but it can only be recognized by using the private key r embedded in the Telex station.

AVOIDING CENSORS

One approach

1. Map the Internet

Incredibly difficult research problem unto itself!

2. Choose paths that do not go through the attackers' countries

Is it possible to get *provable avoidance*?

SOME RESEARCH HERE AT UMD

Alibi Routing

Dave Levin* Youndo Lee* Luke Valenta† Zhihao Li* Victoria Lai*
Cristian Lumezanu‡ Neil Spring* Bobby Bhattacharjee*

*University of Maryland †University of Pennsylvania ‡NEC Labs

ABSTRACT

There are several mechanisms by which users can gain insight into where their packets have gone, but no mechanisms allow users undeniable proof that their packets did *not* traverse certain parts of the world while on their way to or from another host. This paper introduces the problem of finding "proofs of avoidance": evidence that the paths taken by a packet and its response avoided a user-specified set of "forbidden" geographic regions. Proving that something did *not* happen is often intractable, but we demonstrate a low-overhead proof structure built around the idea of what we call "alibis": relays with particular timing constraints that, when upheld, would make it impossible to traverse both the relay and the forbidden regions.

We present *Alibi Routing*, a peer-to-peer overlay routing system for finding alibis securely and efficiently. One of the primary distinguishing characteristics of Alibi Routing is that it does not require knowledge of—or modifications to—the Internet's routing hardware or policies. Rather, Alibi Routing is able to derive its proofs of avoidance from user-provided GPS coordinates and speed of light propagation delays. Using a PlanetLab deployment and larger-scale simulations, we evaluate Alibi Routing to demonstrate that many source-destination pairs can avoid countries of their choosing with little latency inflation. We also identify when Alibi Routing does not work: it has difficulty avoiding regions that users are very close to (or, of course, inside of).

Categories and Subject Descriptors

C.2.2 [**Computer-Communication Networks**]: Network Protocols; C.2.0 [**Computer-Communication Networks**]: General—*Security and protection*

SIGCOMM '15, August 17–21, 2015, London, United Kingdom

 ISBN 978-1-4503-3542-3/15/08…$15.00

DOI: http://dx.doi.org/10.1145/2785956.2787509

Keywords

Alibi Routing; Provable route avoidance; Censorship avoidance; Peer-to-peer; Overlay routing

1. INTRODUCTION

Users have little control over where in the world their packets travel en route to their destinations. Some mechanisms exist to provide insight into where packets traveled, such as the record-route IP option, overlay routing systems (§7), or to a lesser extent source-routing. While these approaches expose a subset of the path the user's packets took, they do not allow a user to determine or provably influence where their packets do *not* go.

This paper introduces a new primitive we call *provable avoidance routing*. With provable avoidance routing, a user specifies arbitrary geographic regions—such as countries or UN voting blocs—to be avoided while communicating with a destination. If successful, the primitive returns *proof* that the user's packets did not traverse the forbidden regions. If it is unsuccessful, it concludes only that the packets *may have* traversed them.

The goal of provable avoidance routing is *detection*, as opposed to *prevention*. In other words, alone, it is unable to ensure a user's packets *will not* traverse a region of the world—we do not require modifications to the underlying routing protocols or hardware, and so we are subject to all of today's uncertainties as to where packets will travel. Rather, what we are able to provide is assurance that the user's packets and their respective responses took paths that *did not* traverse regions of the world. Our proofs of avoidance are provided on a per-packet basis, and are *a posteriori*: only after sending the packet and getting a reply can we ascertain whether or not the round-trip communication avoided the forbidden region.

While outright prevention would be ideal, detection can be a powerful tool, as well. For example, consider one of the greatest threats to open communication on the Internet: censorship. Beyond just dropping [34] or logging [29] users' traffic, censorship can take many forms, including *injecting* packets with false information [4]. Recent results indicate that many users may be censored not by their (or their destination's) countries, but by regimes through which their packets transit; a group of anonymous researchers demonstrated that DNS queries that merely traverse China's borders are

QUESTION

Can we provably avoid countries known to censor/attack?

DEMONSTRATES:

It is possible to get "provable avoidance" without even knowing where exactly packets go

Users lack control over routing

Mostly relegated to destination-based routing

Users lack control over routing

Collateral damage of censorship

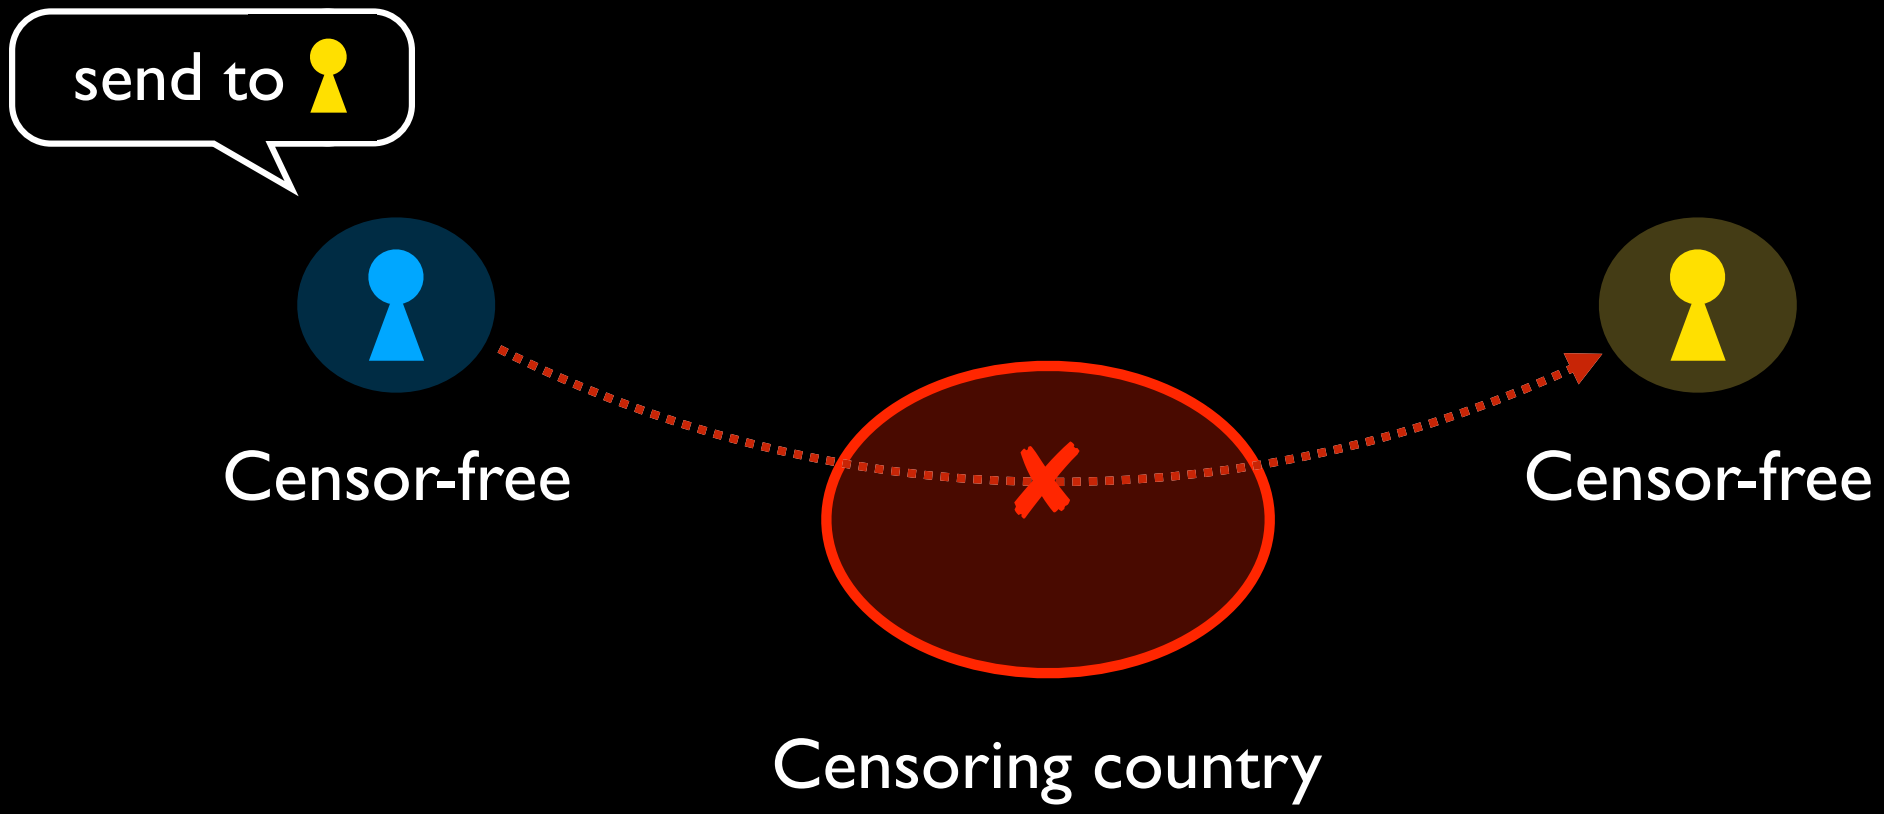

Encryption (HTTPS)

Anonymity (Tor)

Hide info, but are still subject to censorship

This work

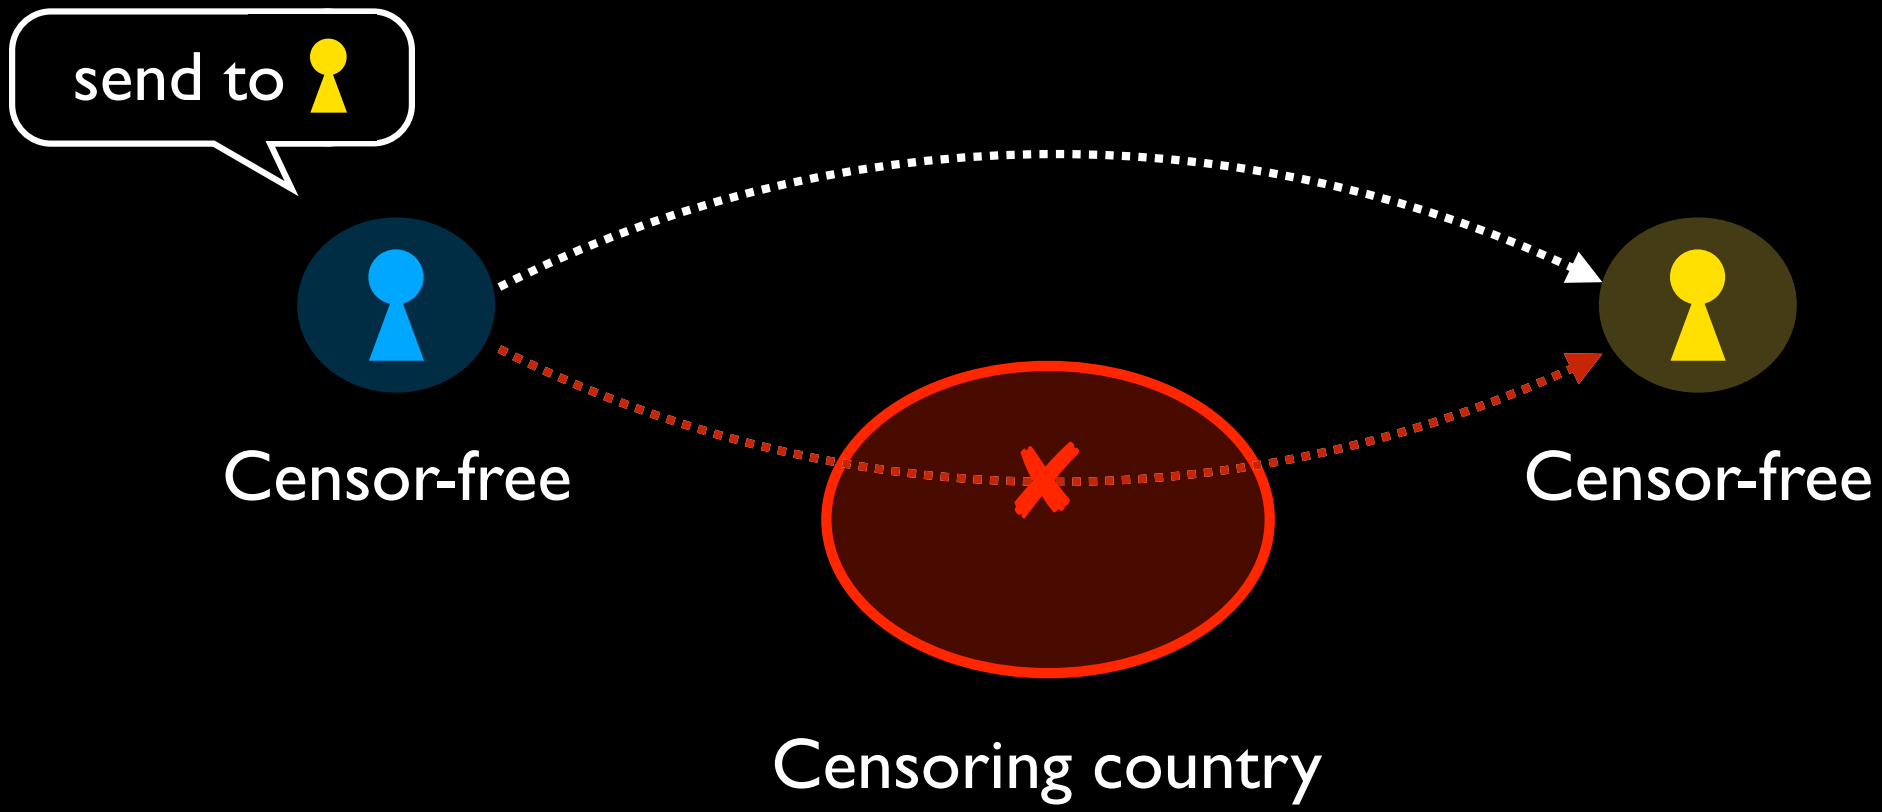

Provable avoidance routing

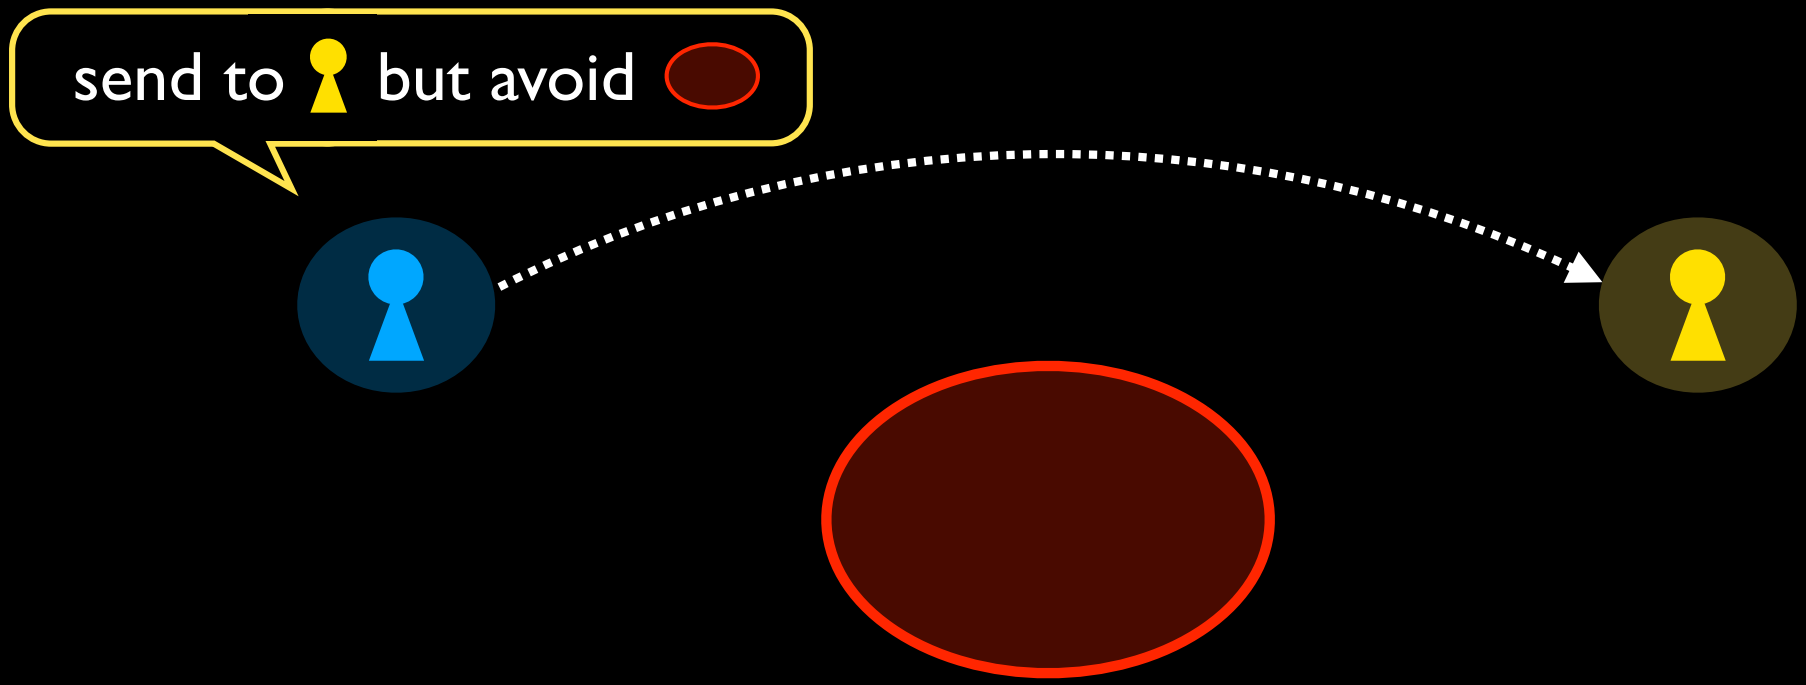

A broadly applicable primitive

Provably disjoint paths

Diffie-Hellman

Avoiding boomerangs

Distinct vantage points

Provable route avoidance goals

Flexibility

Users request their traffic to avoid transiting arbitrary geographic regions

Proof

Provide proofs of avoidance

Provable route avoidance goals

Flexibility

Users request their traffic to avoid transiting arbitrary geographic regions

Without having to know underlying routes

Proof

Provide proofs of avoidance

Provable route avoidance goals

Flexibility

Users request their traffic to avoid transiting arbitrary geographic regions

Proof

Provide proofs of avoidance

Goal: proof that it *did not* traverse

Unadulterated roundtrip of communication

Non-goal: proof that it *cannot* traverse

Provable route avoidance goals

Flexibility — Users request their traffic to avoid transiting arbitrary geographic regions

Proof — Provide proofs of avoidance

How do you prove that something *did not* happen?

Proving the impossible

How do you prove (X) did *not* happen
without enumerating everything that *could have*?

(A) && (A) ⇒ (!X) ⇒ (!X)

Mutually exclusive

(A) is an alibi