CENSORSHIP RESISTANCE

GRAD SEC

NOV 14 2017

TODAY'S PAPERS

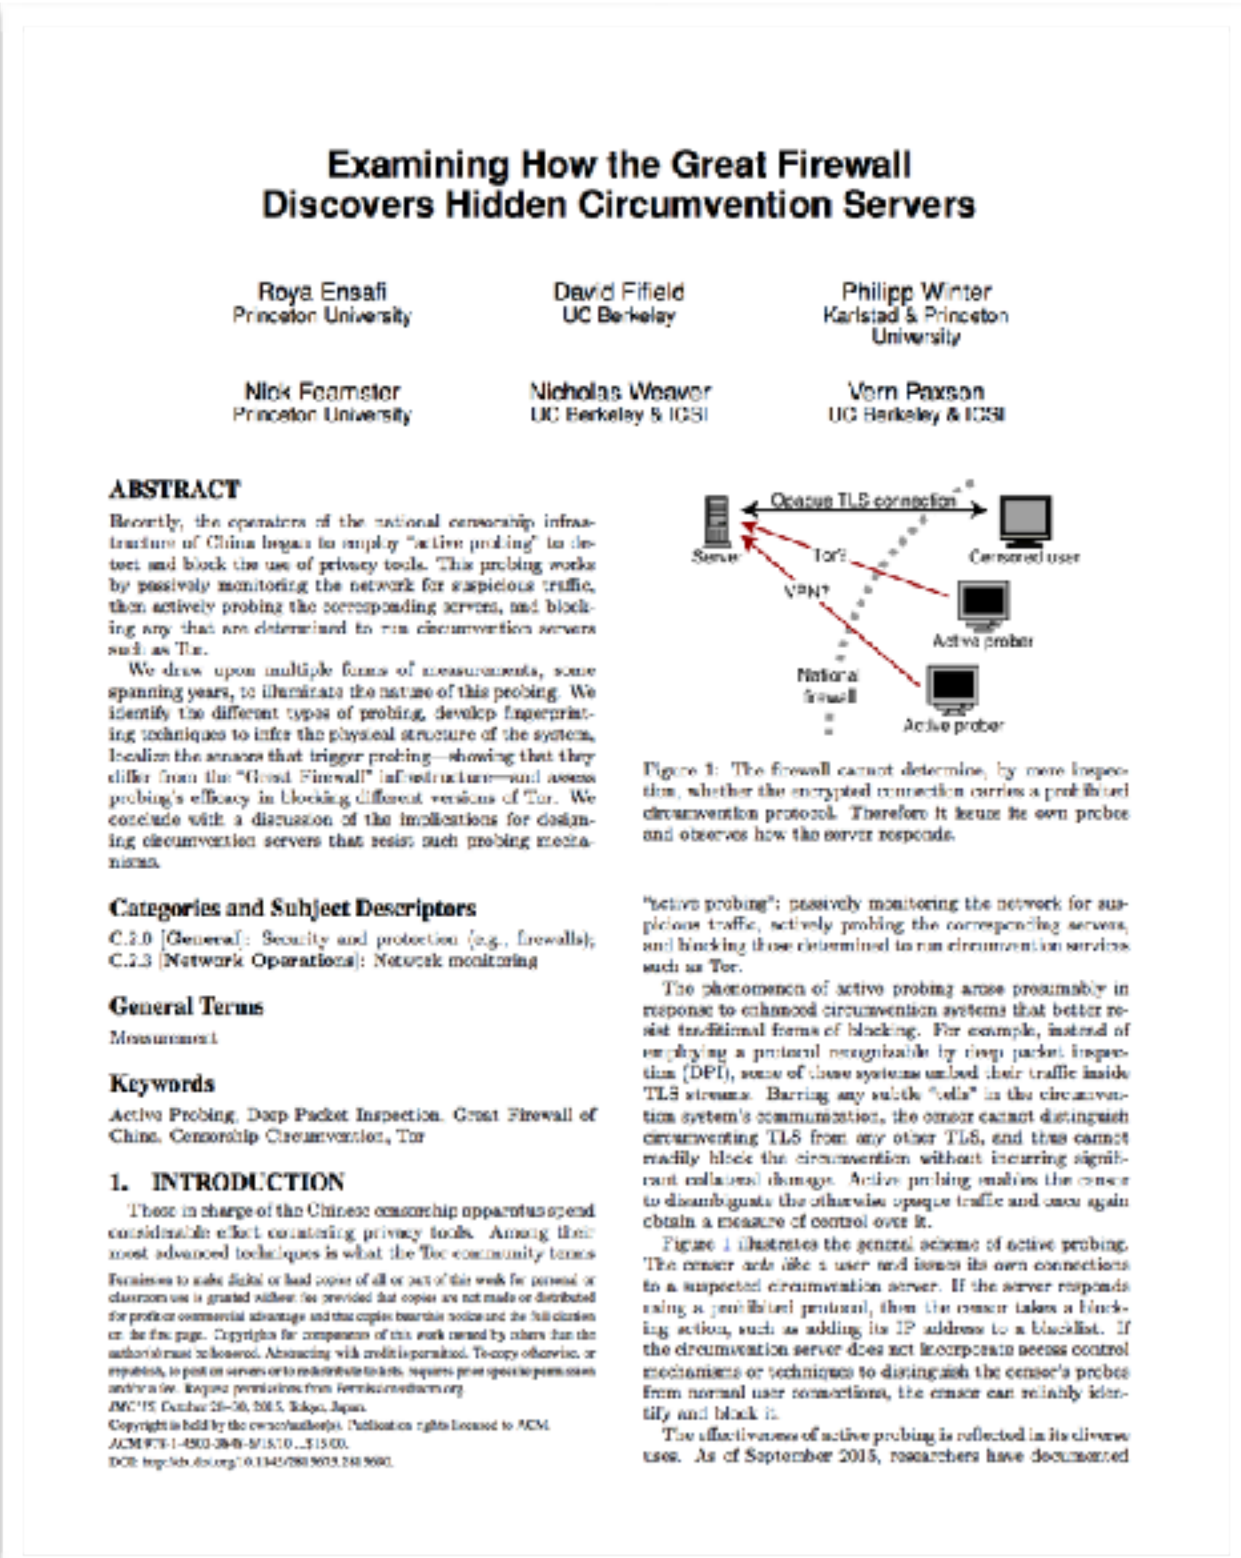

Examining How the Great Firewall Discovers Hidden Circumvention Servers

Roya Ensafi, Princeton University
David Fifield, UC Berkeley
Philipp Winter, Karlstad & Princeton University
Nick Feamster, Princeton University
Nicholas Weaver, UC Berkeley & ICSI
Vern Paxson, UC Berkeley & ICSI

ABSTRACT

Categories and Subject Descriptors

General Terms

Keywords

1. INTRODUCTION

Telex: Anticensorship in the Network Infrastructure

Eric Wustrow* Scott Wolchok* Ian Goldberg† J. Alex Halderman*
*The University of Michigan
†University of Waterloo

Abstract

1 Introduction

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

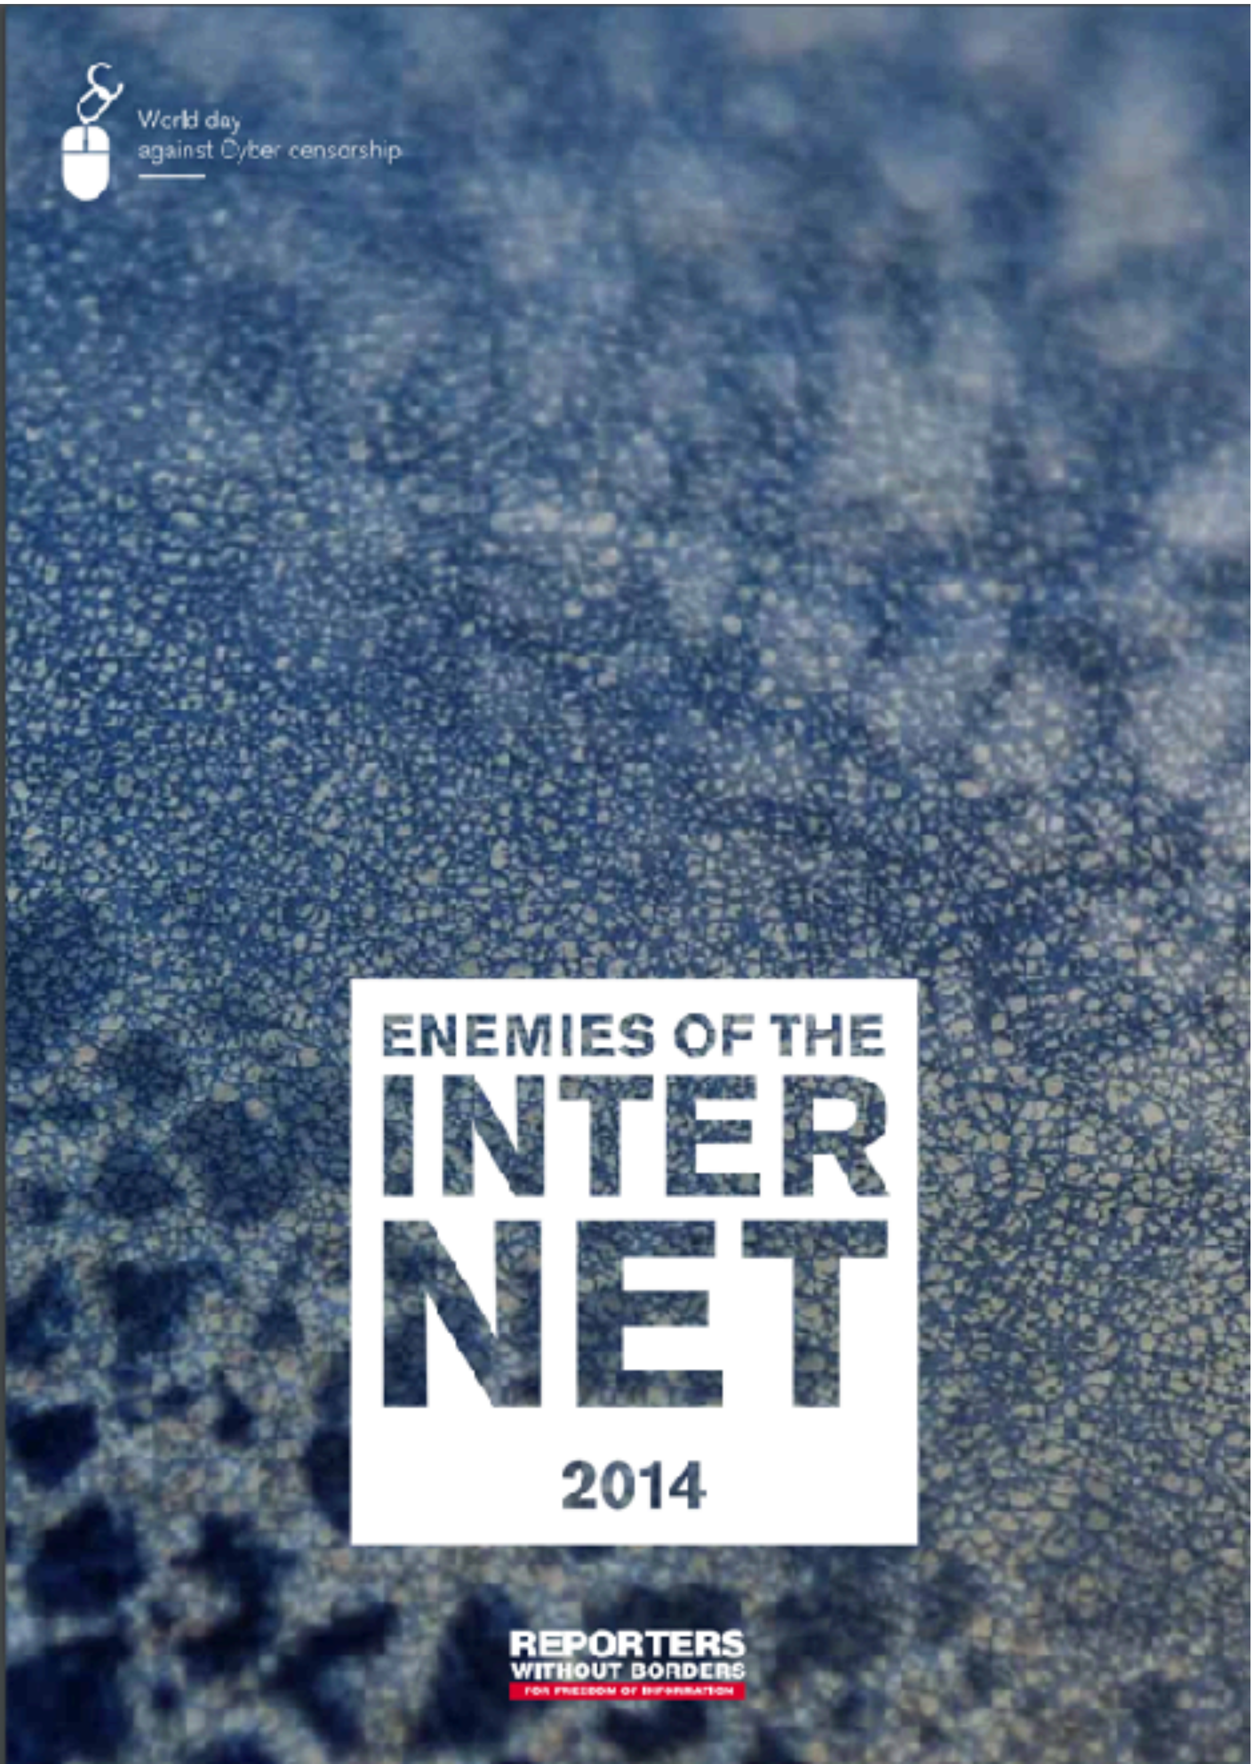

~Annual report by
Reporters without Borders

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

Enemies of the Internet

REPORTERS WITHOUT BORDERS
FOR FREEDOM OF INFORMATION

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

USA: NSA symbolises intelligence services' abuses

In June 2013, computer specialist Edward Snowden disclosed the extent of the surveillance practices of the U.S. and British intelligence services. Snowden, who worked for a government sub-contractor and had access to confidential documents, later exposed more targeted surveillance, focusing on the telecommunications of world leaders and diplomats of allied countries. Activists, governments and international bodies have taken issue with the Obama administration, as the newspapers *The Guardian* and *The Washington Post* have revealed the extent of the surveillance. The main player in this vast surveillance operation is the highly secretive National Security Agency (NSA) which, in the light of Snowden's revelations, has come to symbolize the abuses by the world's 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 publish information from Edward Snowden, while the British edition is not, but the country of the First Amendment has undermined confidence in the Internet and its own standards of security. U.S. surveillance practices and decryption activities are a direct threat to investigative journalists, especially those who work with sensitive sources for whom confidentiality is paramount and who are already under 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.

ENEMIES OF THE INTERNET

Français
Español

World day
against Cyber censorship

Enemies of the Internet

REPORTERS WITHOUT BORDERS
FOR FREEDOM OF INFORMATION

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 w
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 publish information from Edward Snowden, while the British edition is not, but the country of the First Amendment has undermined confidence in the Internet and its own standards of security. U.S. surveillance practices and decryption activities are a direct threat to investigative journalists, especially those who work with sensitive sources for whom confidentiality is paramount and who are already under 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.

ENEMIES OF THE INTERNET

Français
Español

World day
against Cyber censorship

Enemies of the Internet

REPORTERS WITHOUT BORDERS
FOR FREEDOM OF INFORMATION

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 w
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 decryptic
especially those who work with sensitive sources
pressure.

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.

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.

COLLATERAL DAMAGE OF INTERNET CENSORSHIP

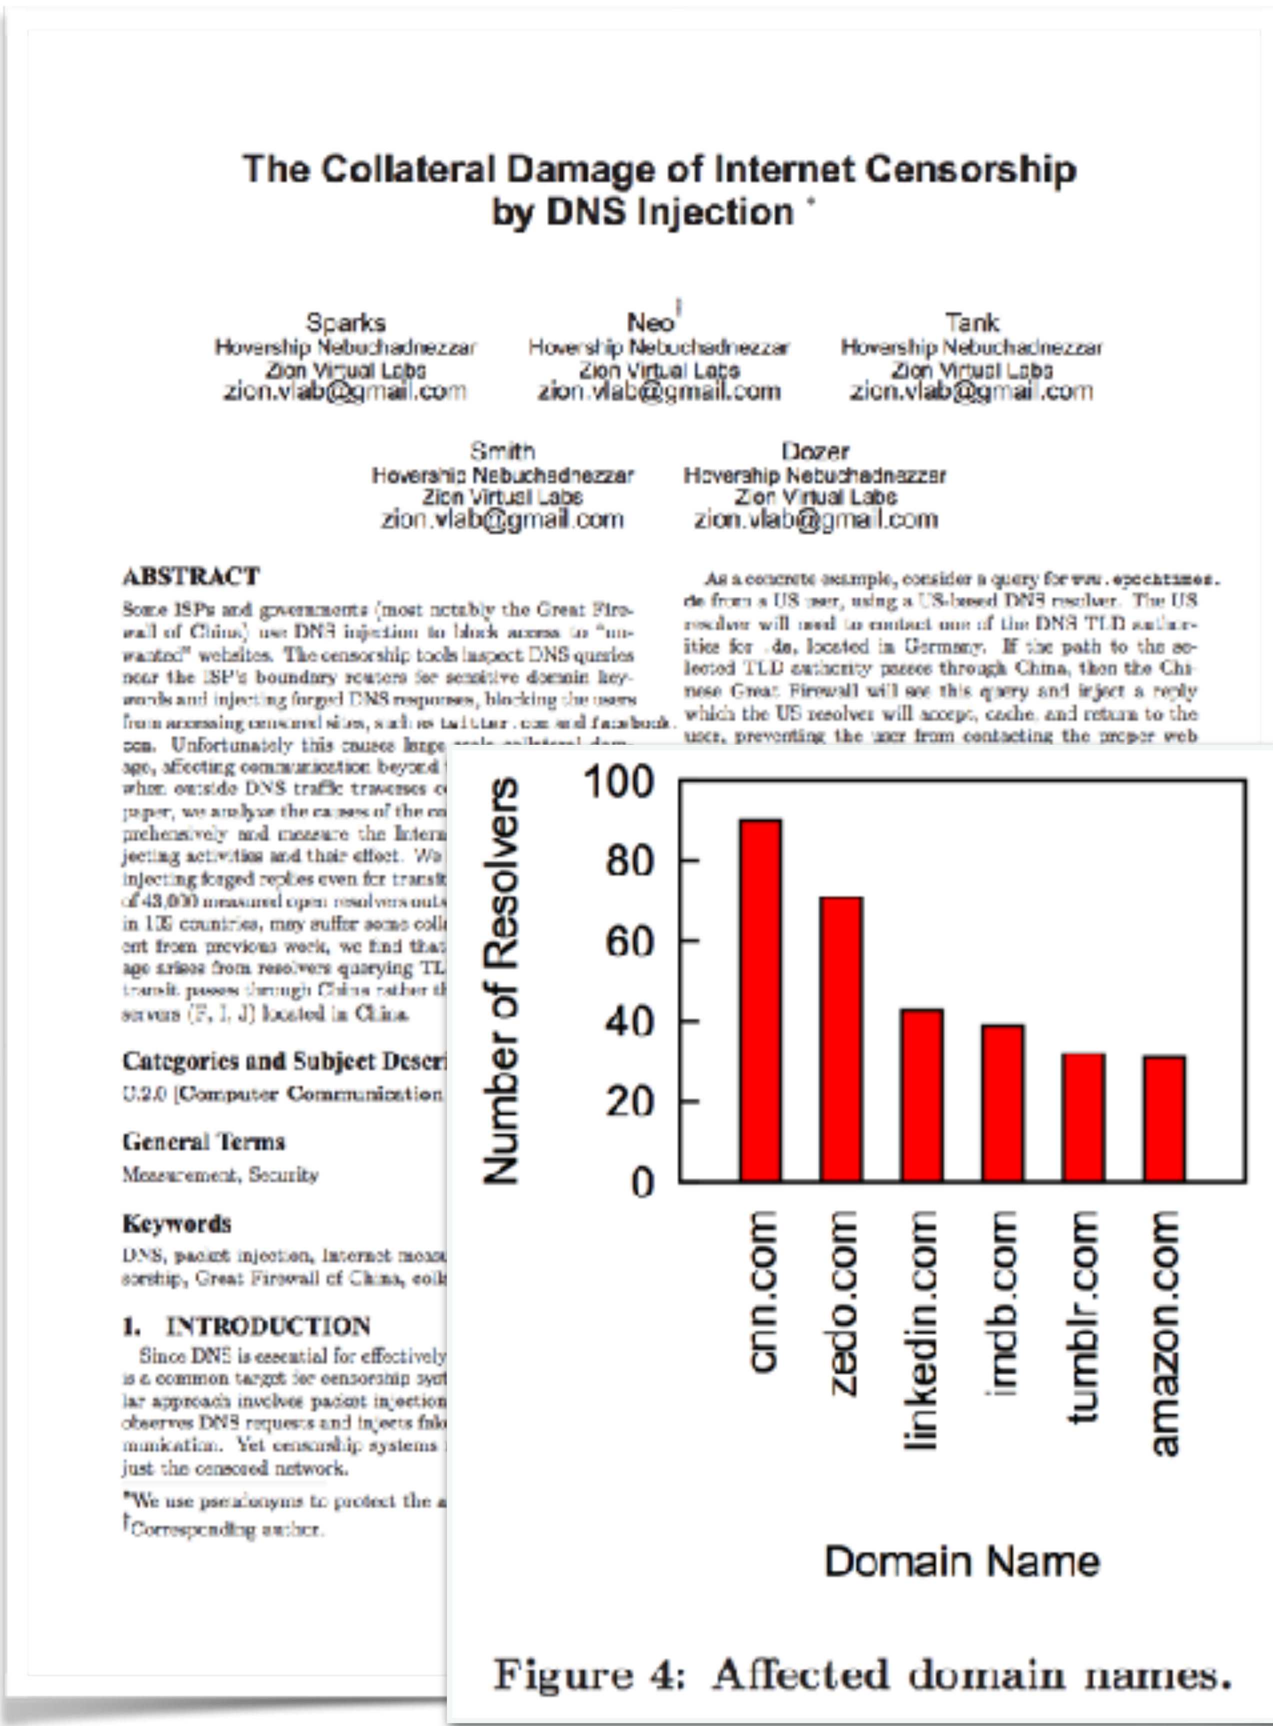

Figure 4: Affected domain names.

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.

317

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

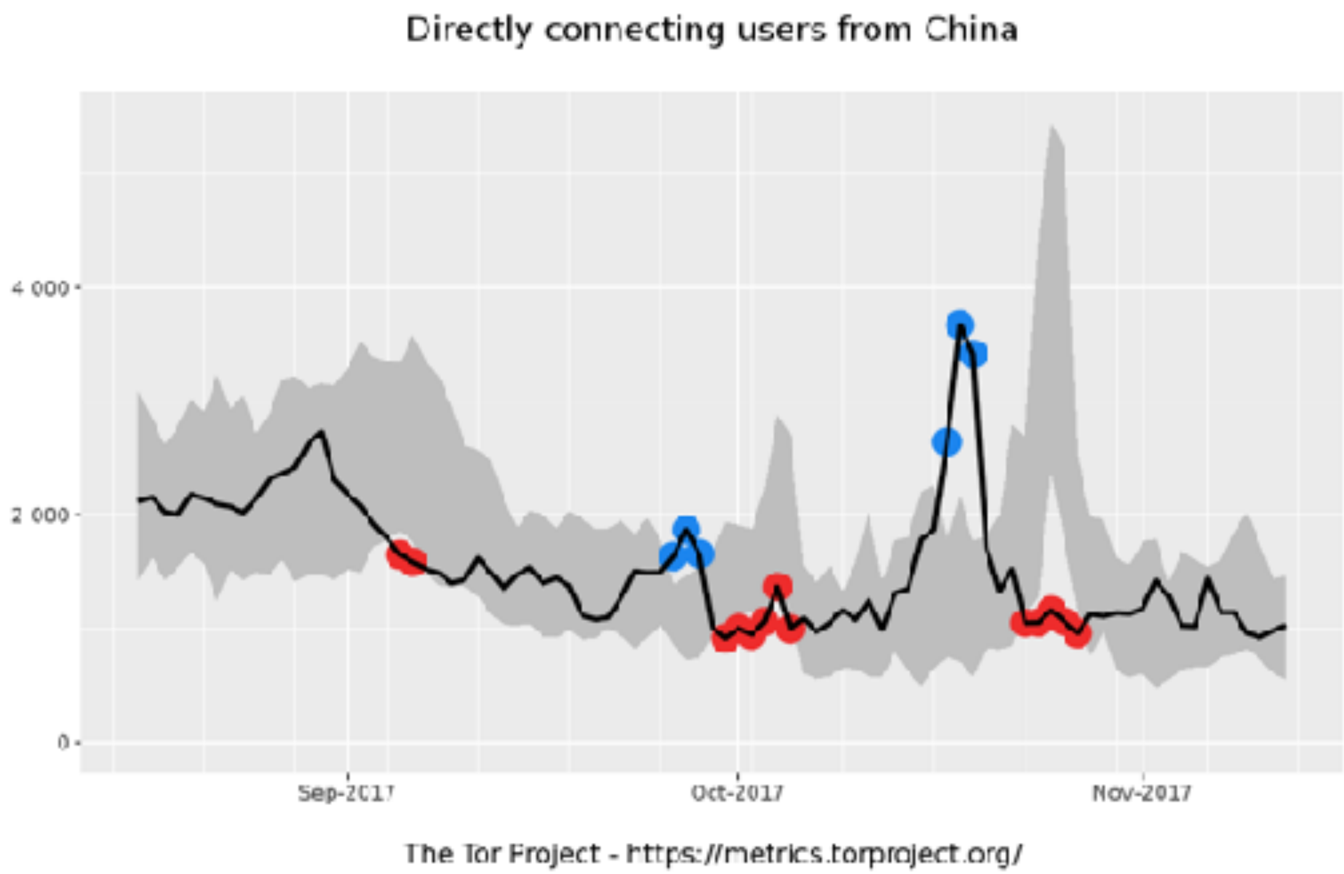

Estimate the number of users on day *i* based on previous days' users

Gray area: Range of estimated users; Usage naturally fluctuates

Downturn event: Drops below
Possibly indicates censorship

Upturn event: Rises above "normal"
Possibly indicates circumvention

BLOCKING TOR

Directly connecting users from China

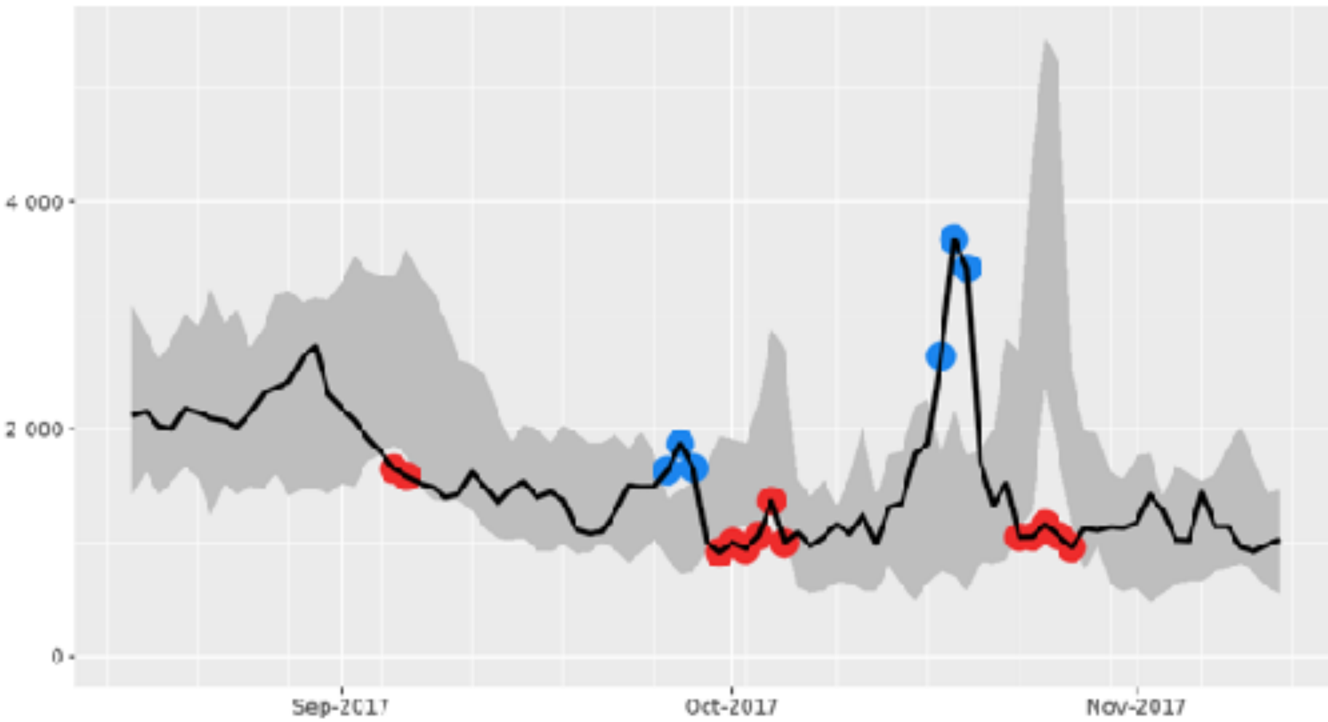

Directly connecting users from Puerto Rico

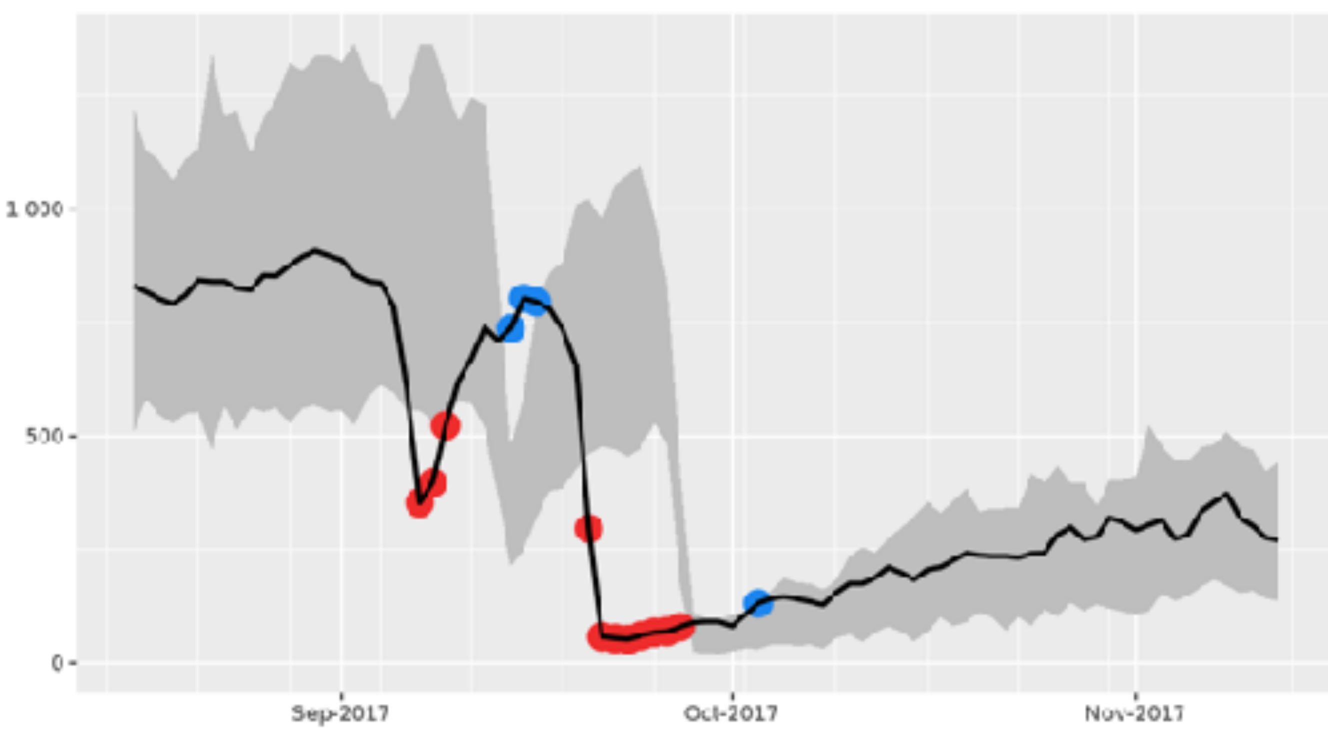

The Tor Project - https://metrics.torproject.org/

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

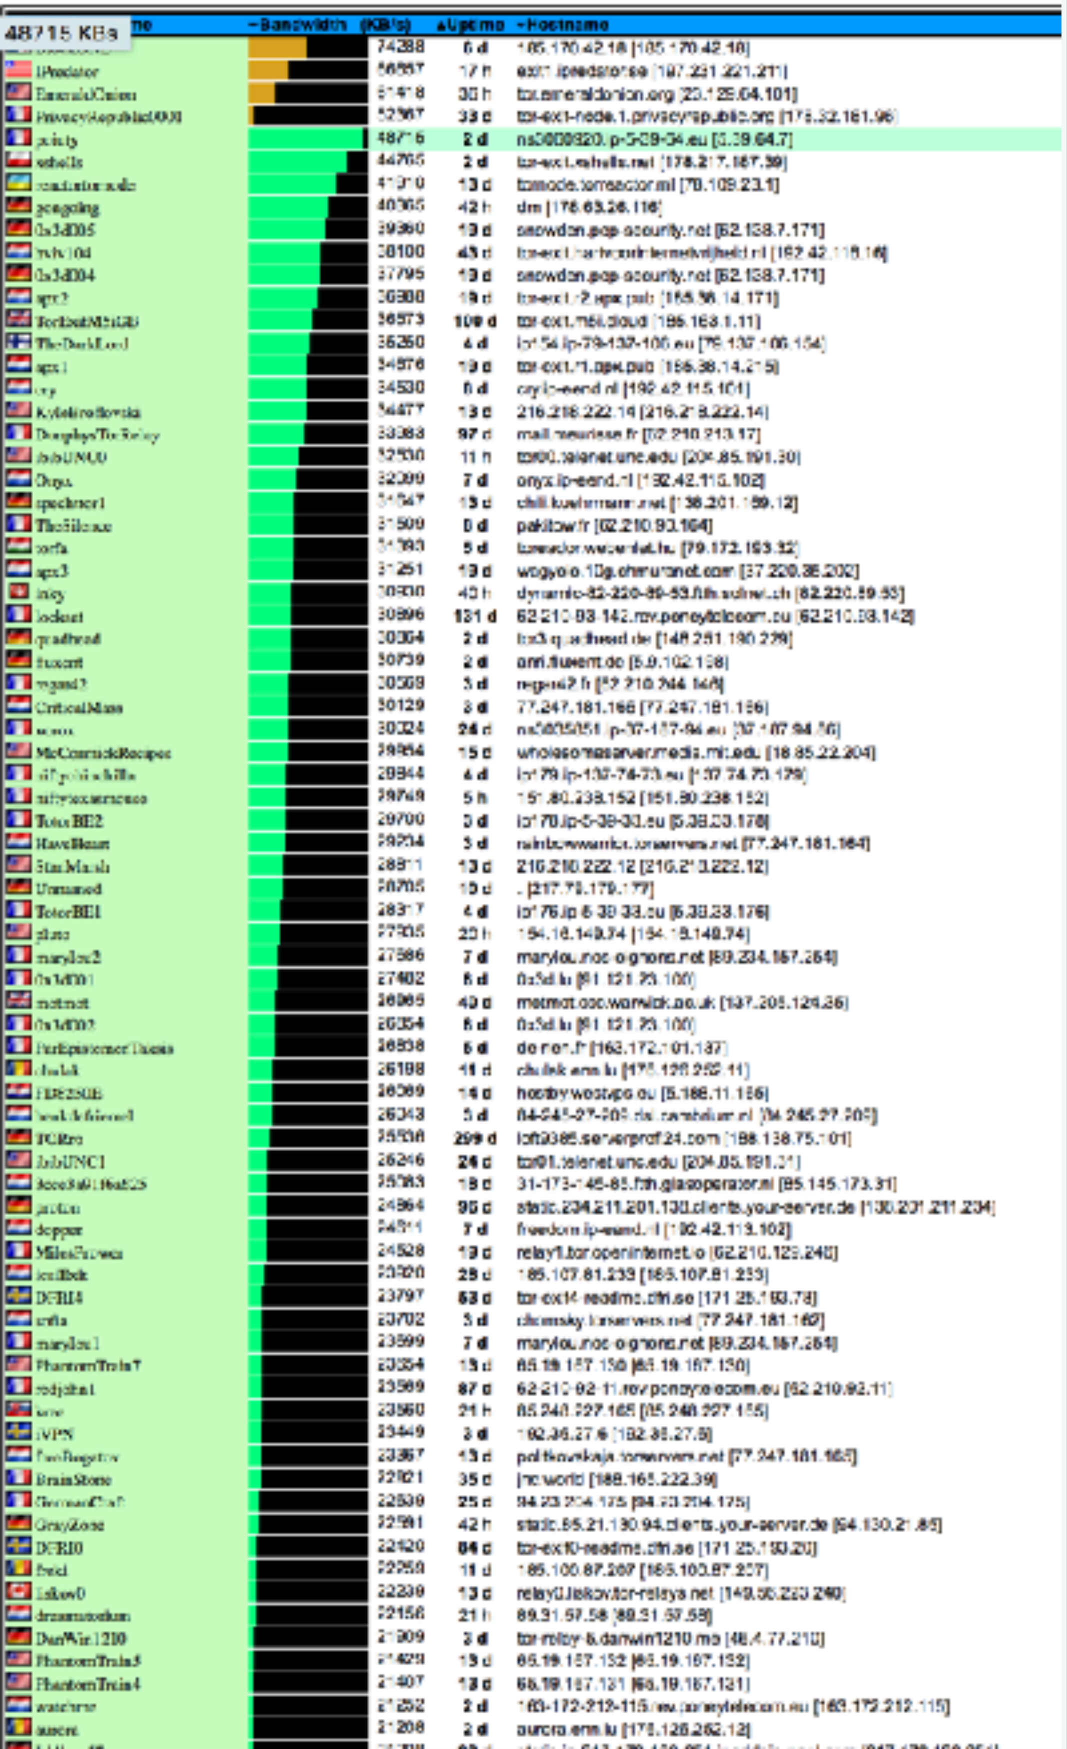

Option 1: Get a list of all Tor nodes

Insert them as firewall rules

Bridge nodes: Tor does not list some nodes;

Users must learn them out of band

This week's paper: Censors 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

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

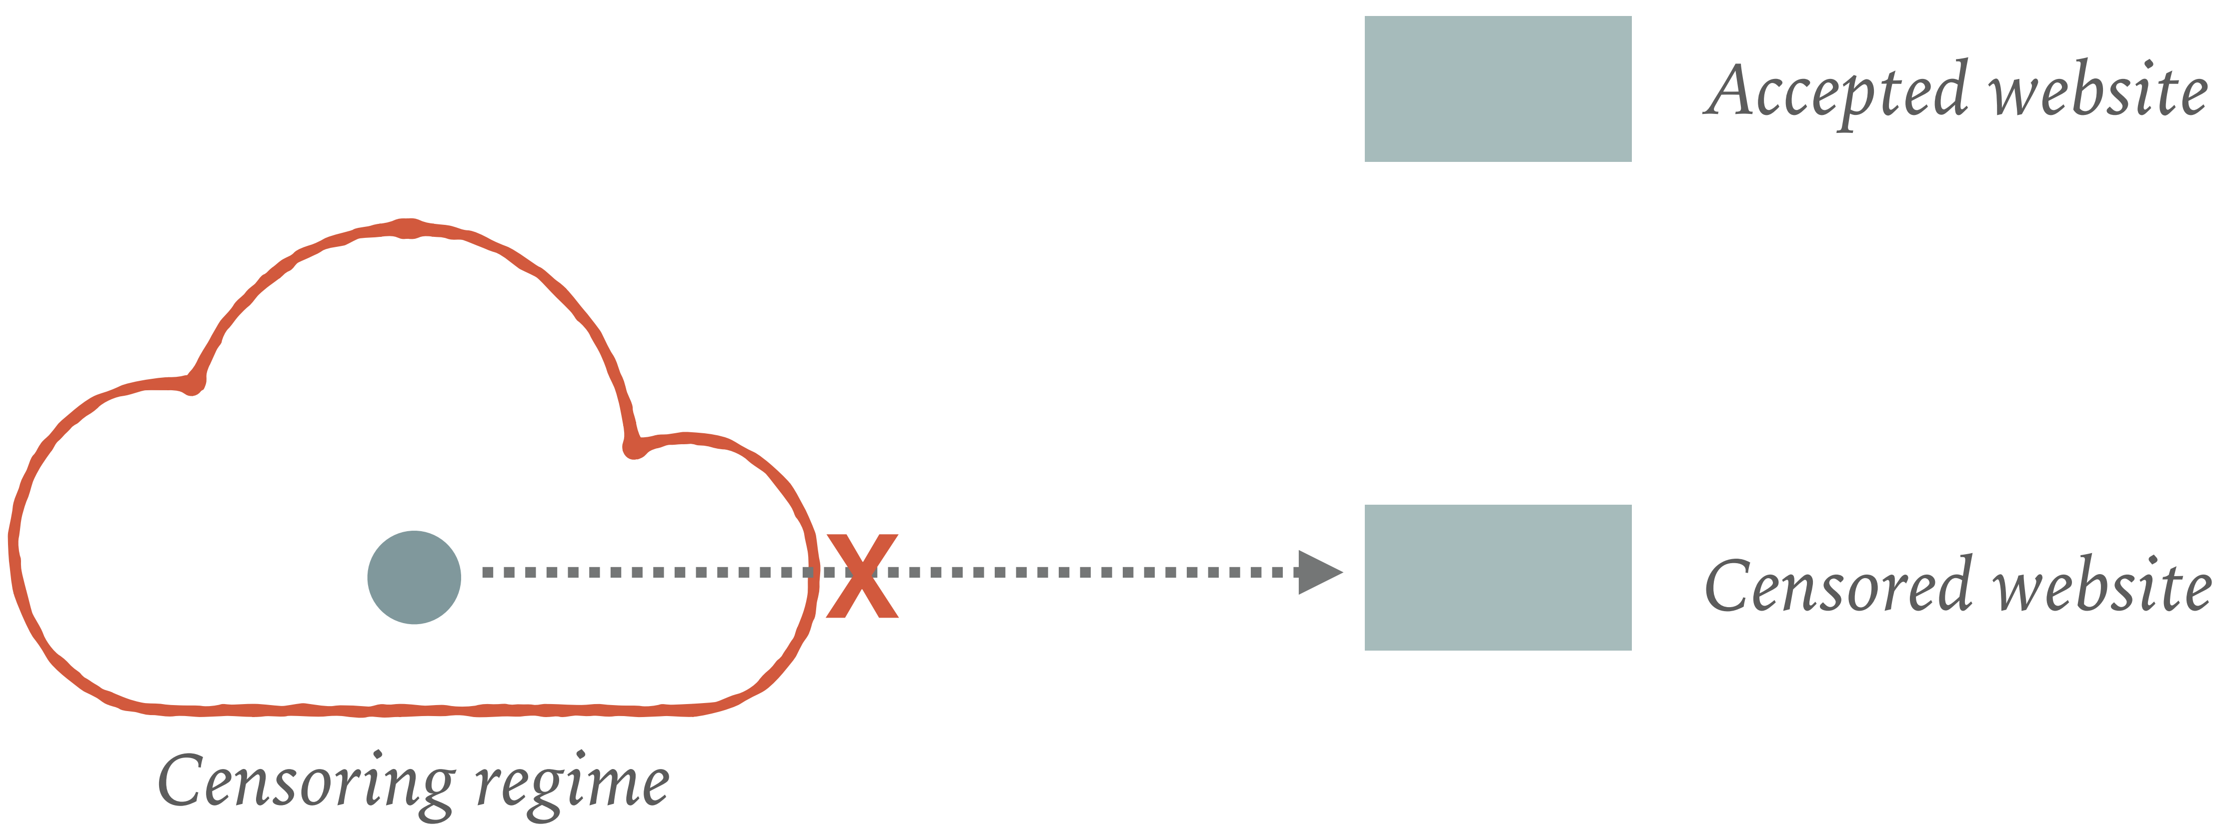

DECOY ROUTING

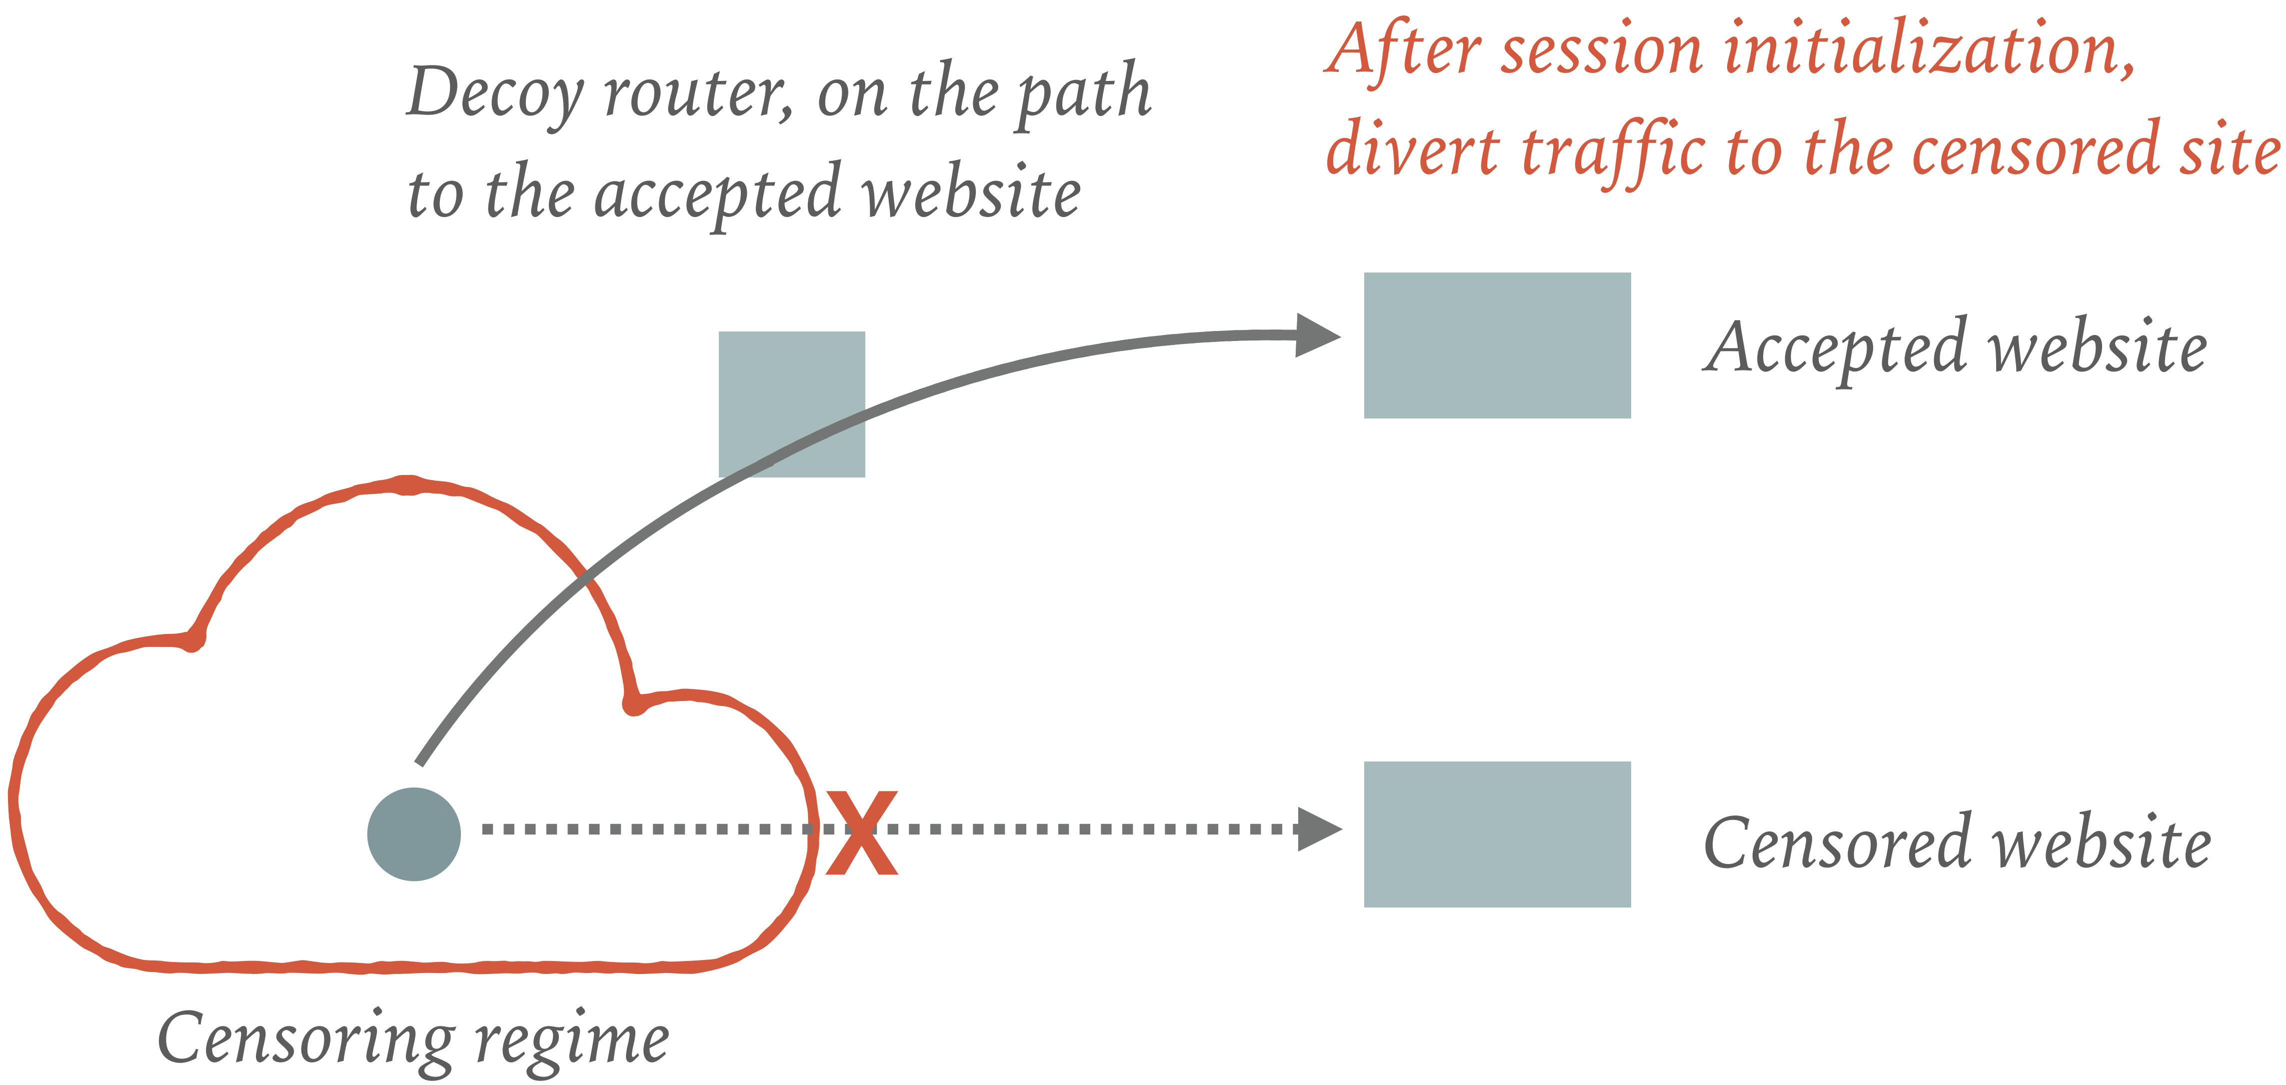

How does the decoy router know the true destination but the censor doesn't?

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

DECOY ROUTING

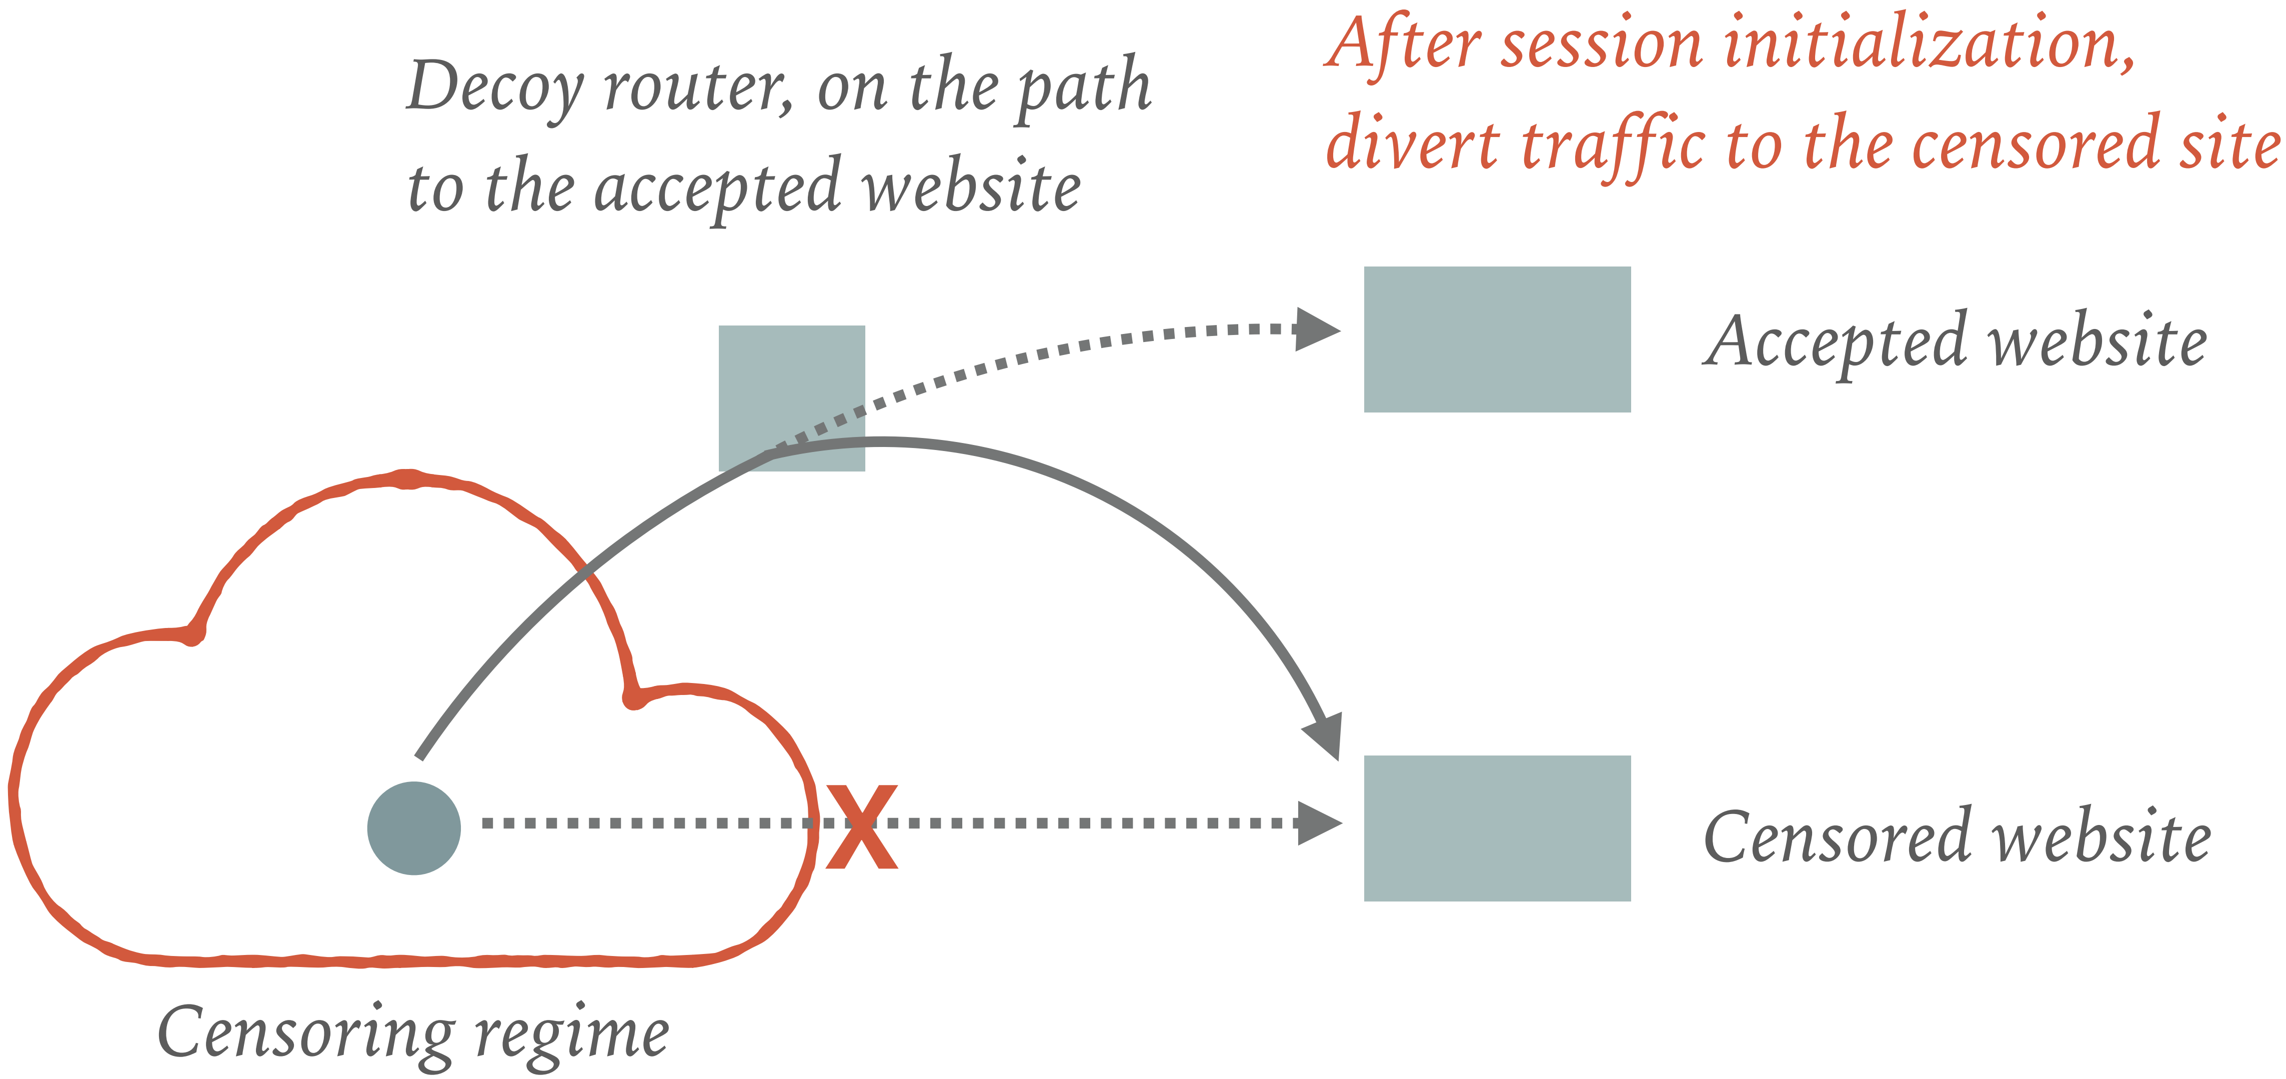

How does the decoy router know the true destination but the censor doesn't?

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

DECOY ROUTING TAGS

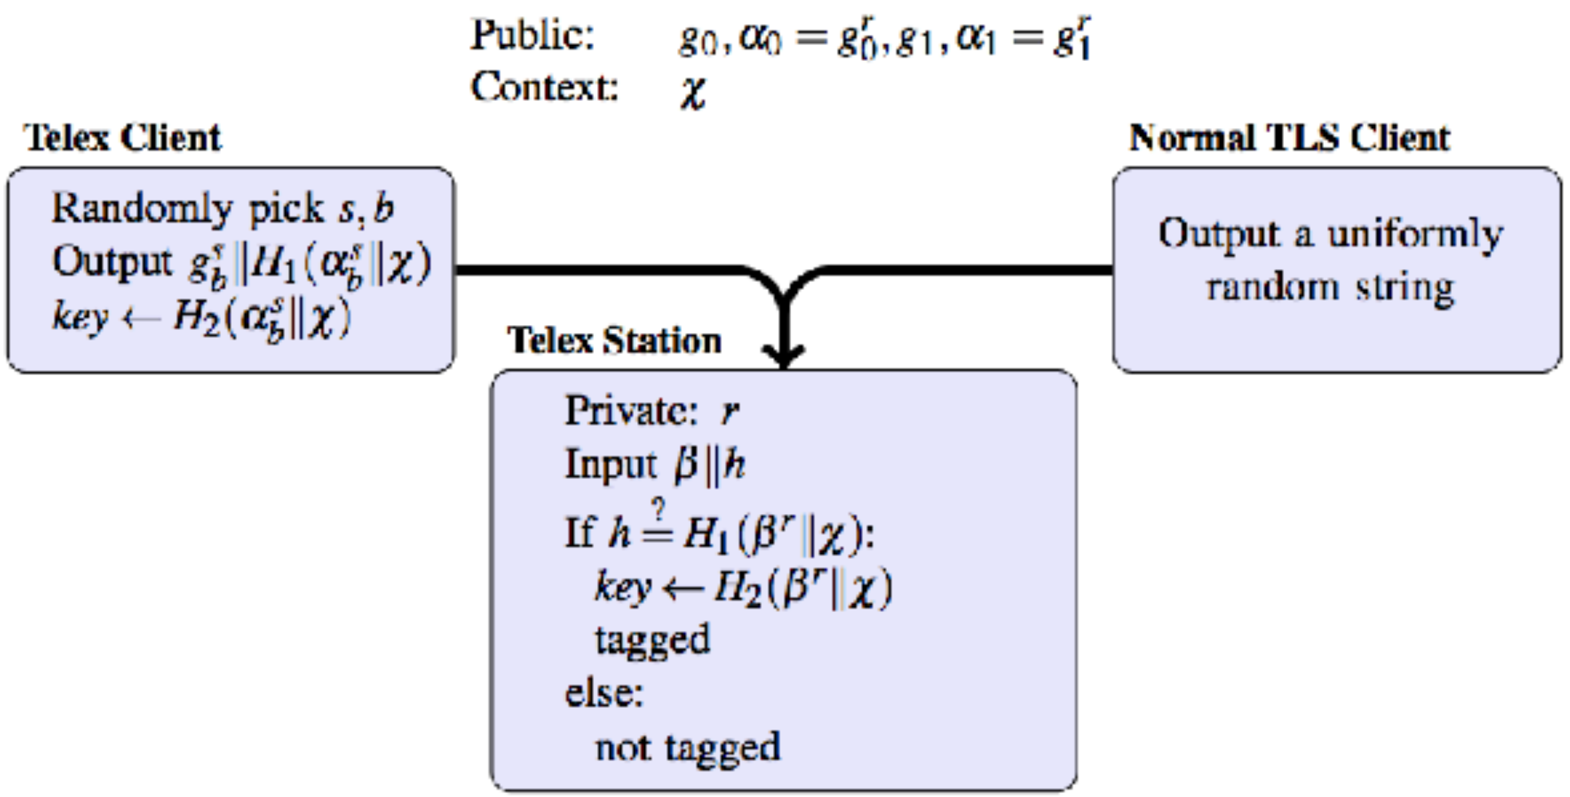

Figure 2: **Tag creation and detection** — Telex intercepts TLS connections that contain a steganographic tag in the ClientHello message's nonce field (normally a uniformly random string). The Telex client generates the tag using public parameters (shown above), but it can only be recognized by using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

AVOIDING CENSORS

One approach

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

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

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*?

ALIBI ROUTING

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?

ALIBI ROUTING

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?

LESSONS LEARNED

Give your papers a more descriptive title; nobody knows what it's about!

Alibi routing
D Levin, Y Lee, L Valenta, Z Li, V Lai… - ACM SIGCOMM …, 2015 - dl.acm.org
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

ALIBI ROUTING

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
© 2015 Copyright held by the owner/author(s). Publication rights licensed to ACM. 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

511

QUESTION

Can we provably avoid countries known to censor/attack?

LESSONS LEARNED

Give your papers a more descriptive title; nobody knows what it's about!

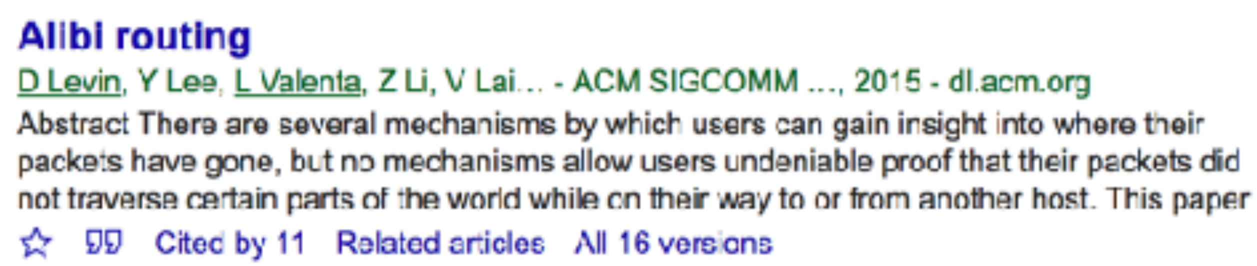

Also, yes, it's possible to get provable avoidance without even knowing where exactly packets went

Users lack control over routing

Mostly relegated to destination-based routing

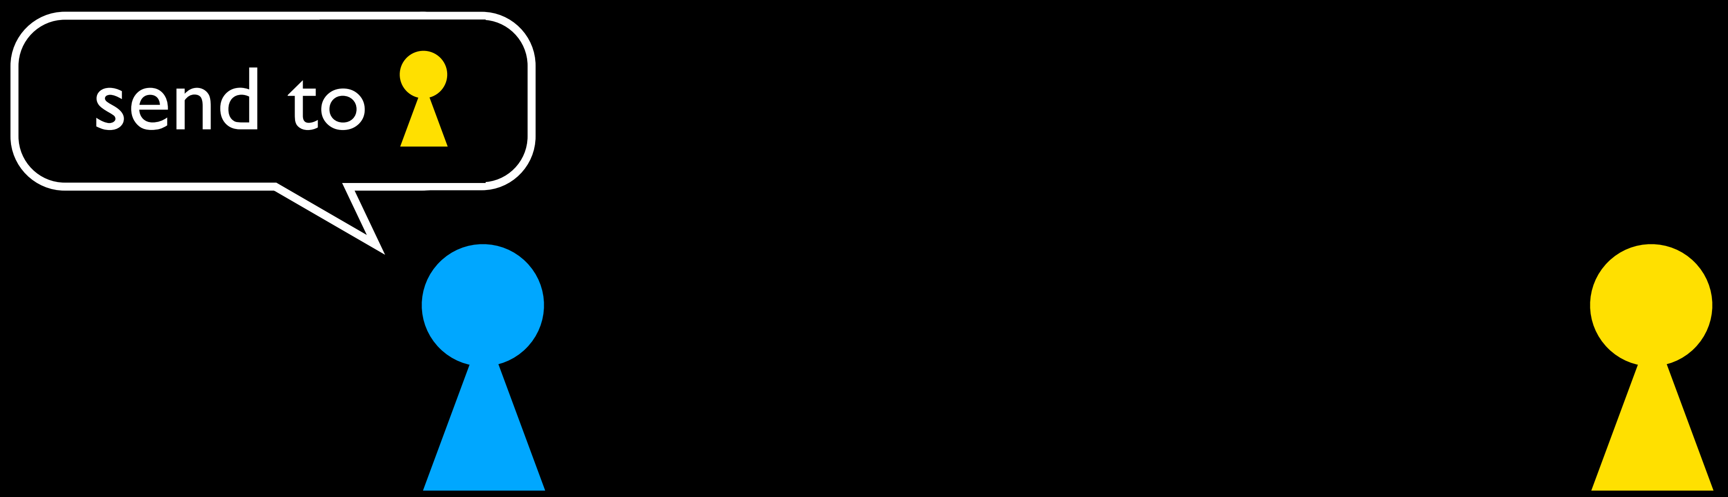

Users lack control over routing

Mostly relegated to destination-based routing

Users lack control over routing

Collateral damage of censorship

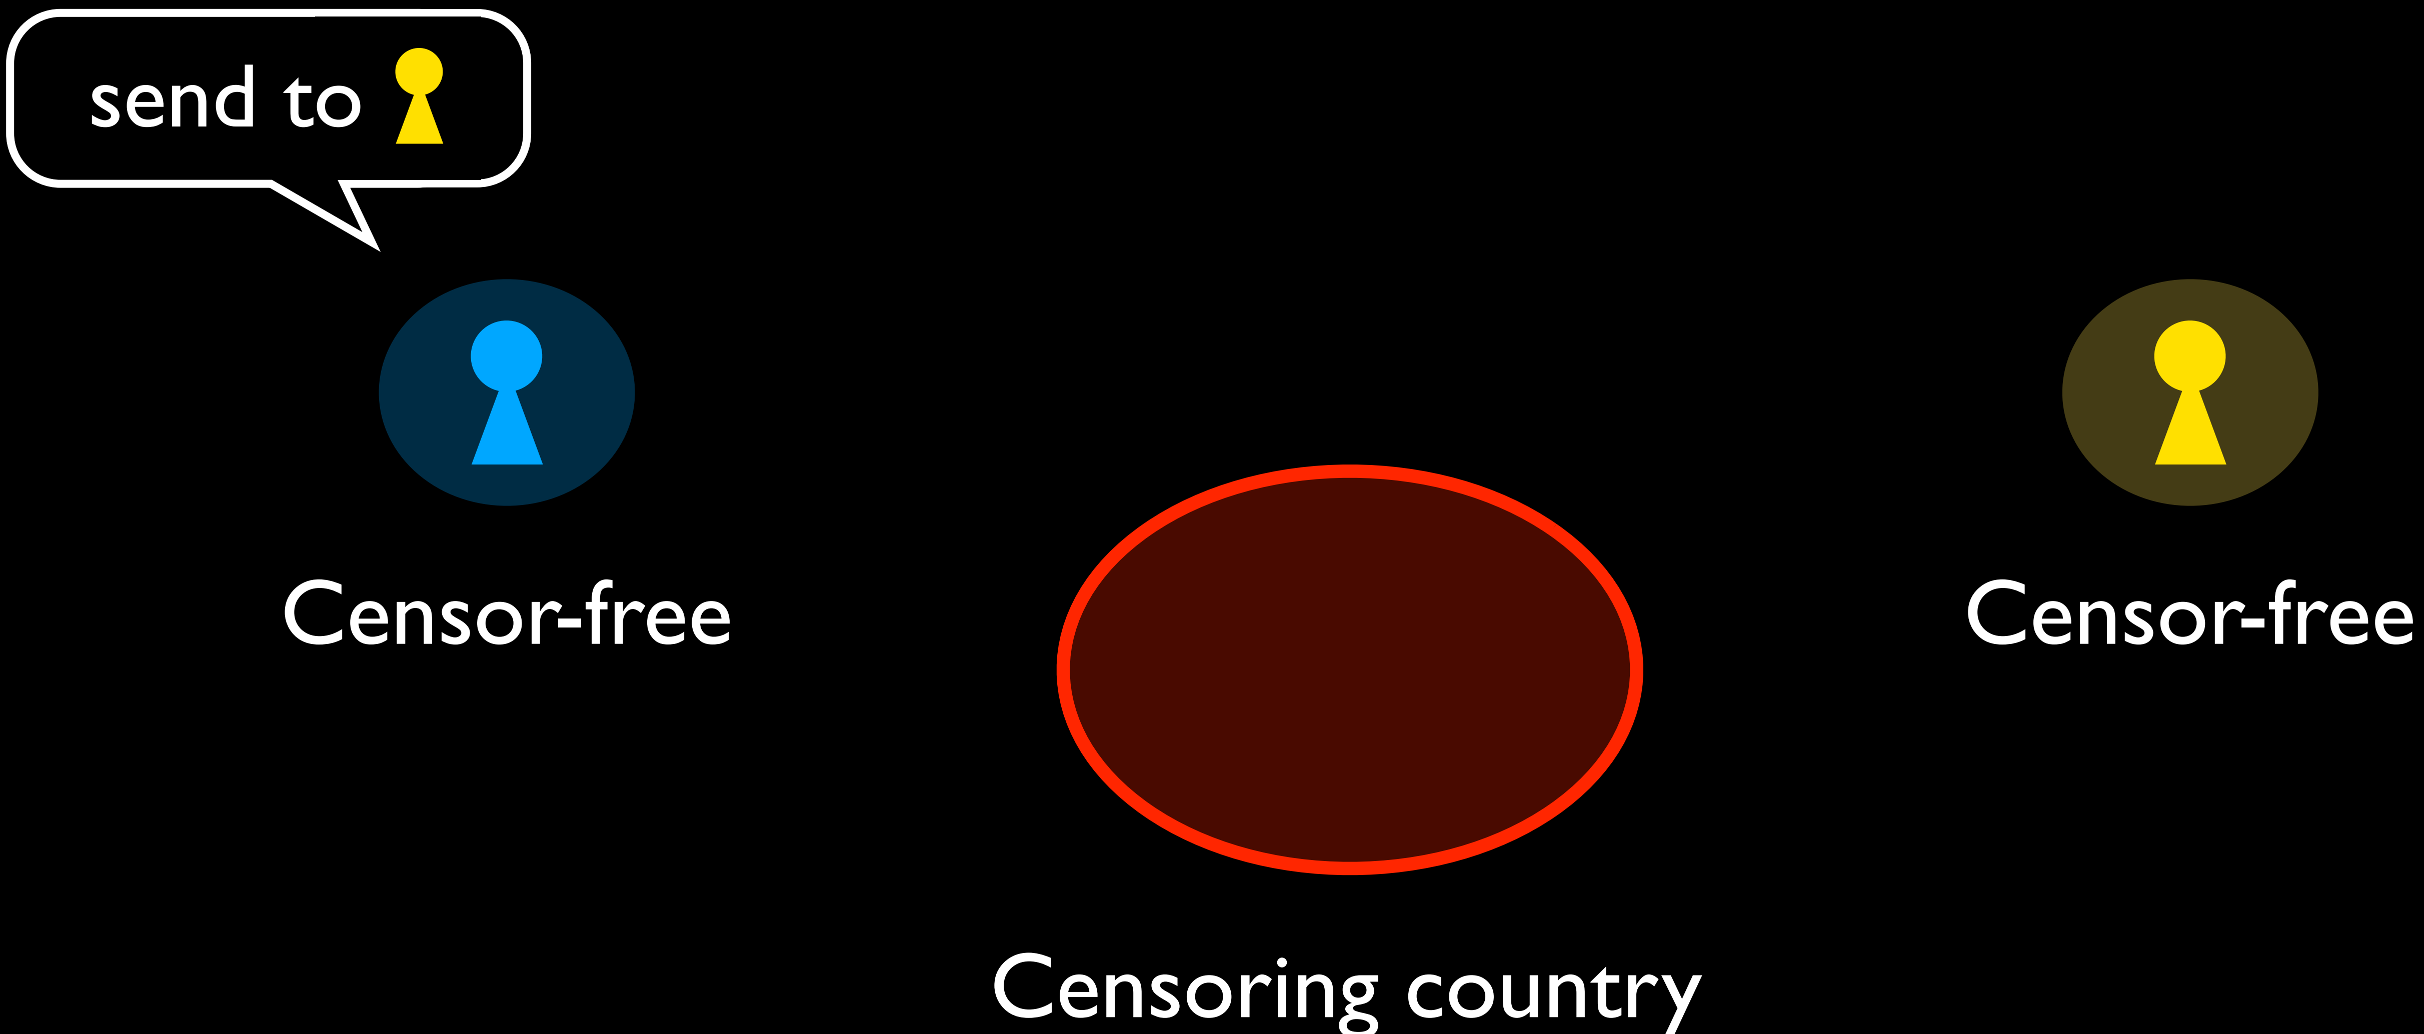

Users lack control over routing

Collateral damage of censorship

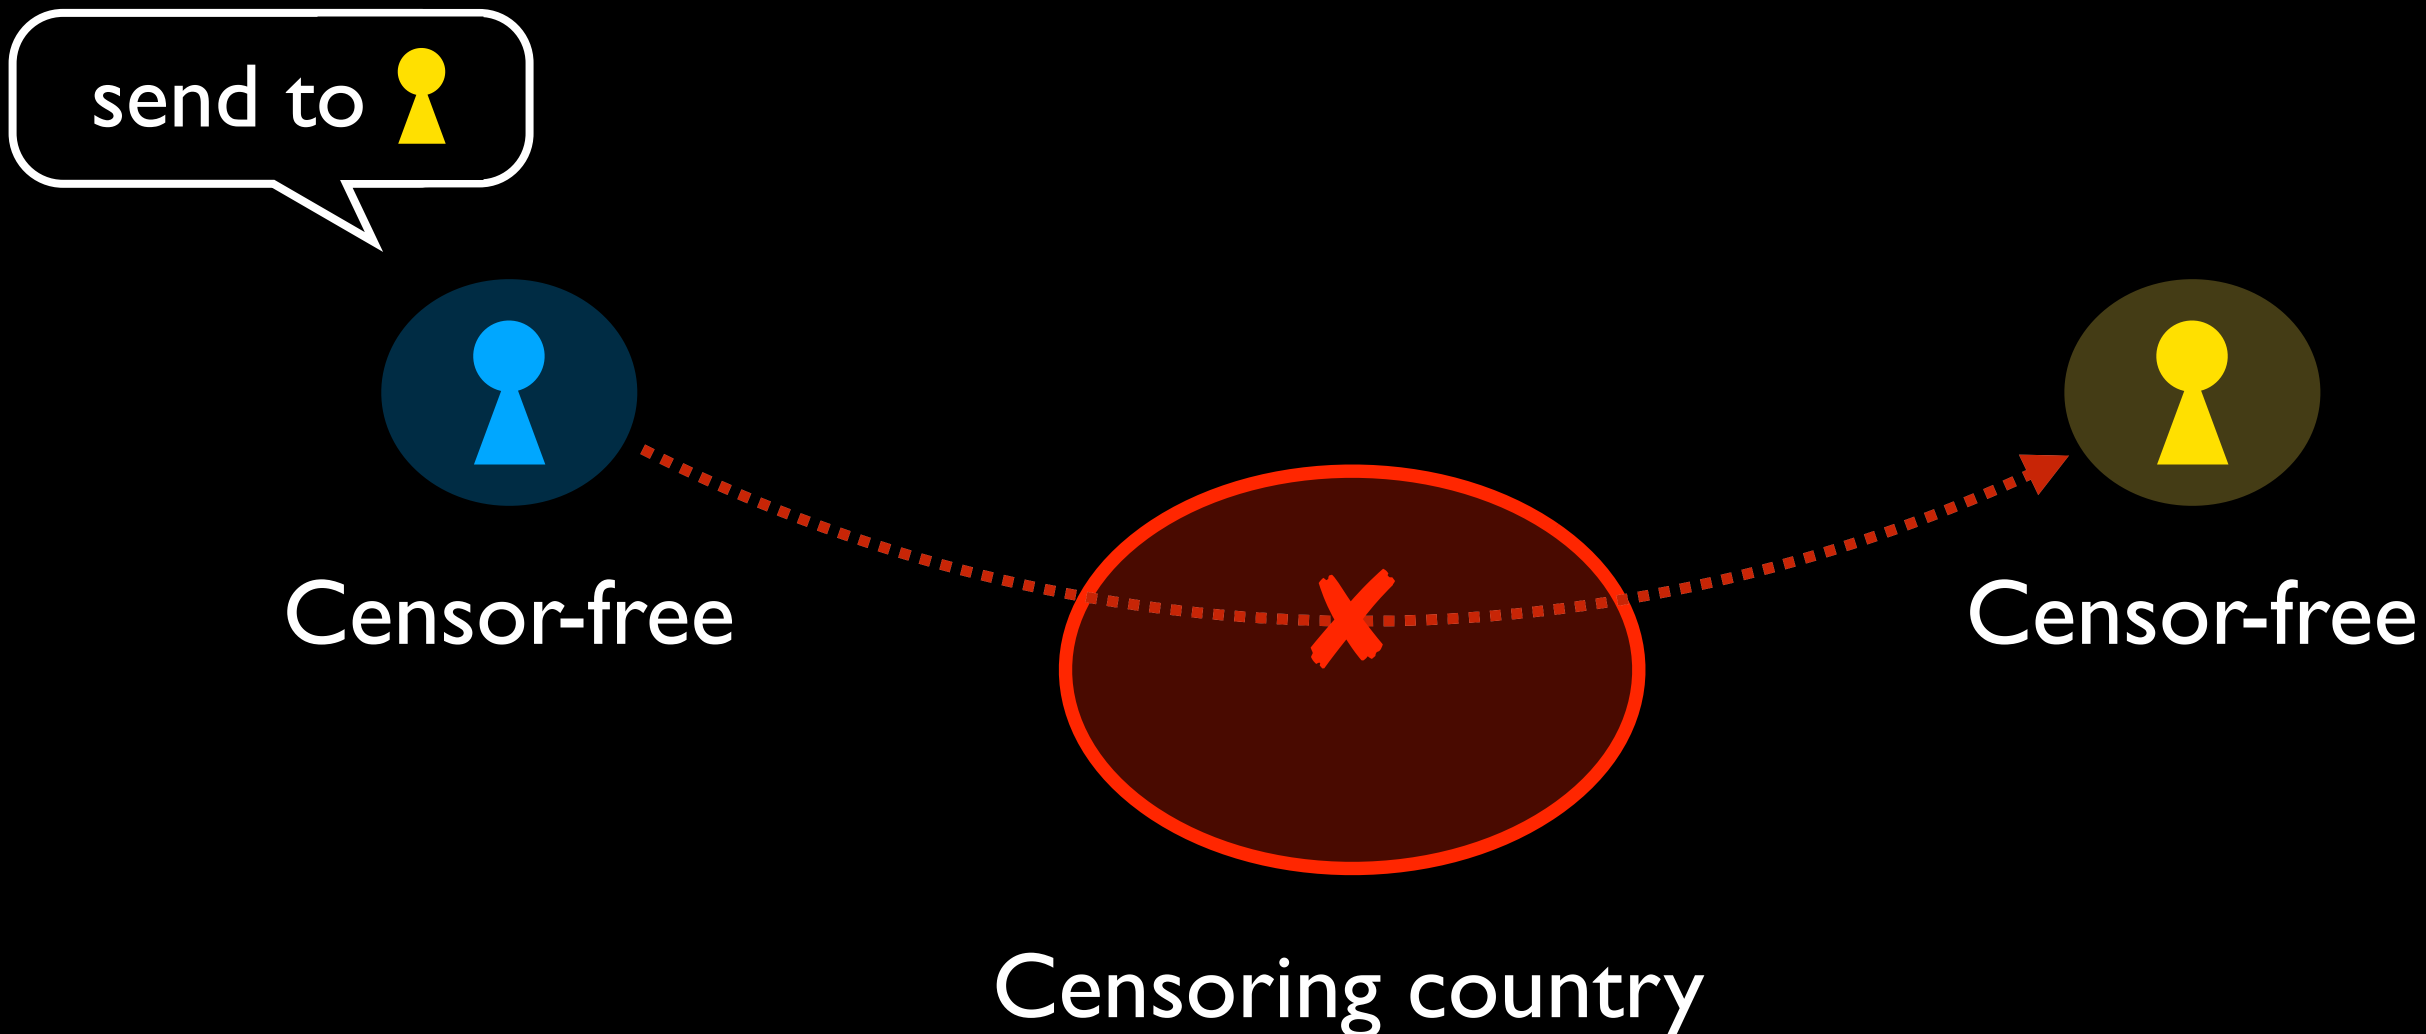

Users lack control over routing

Collateral damage of censorship

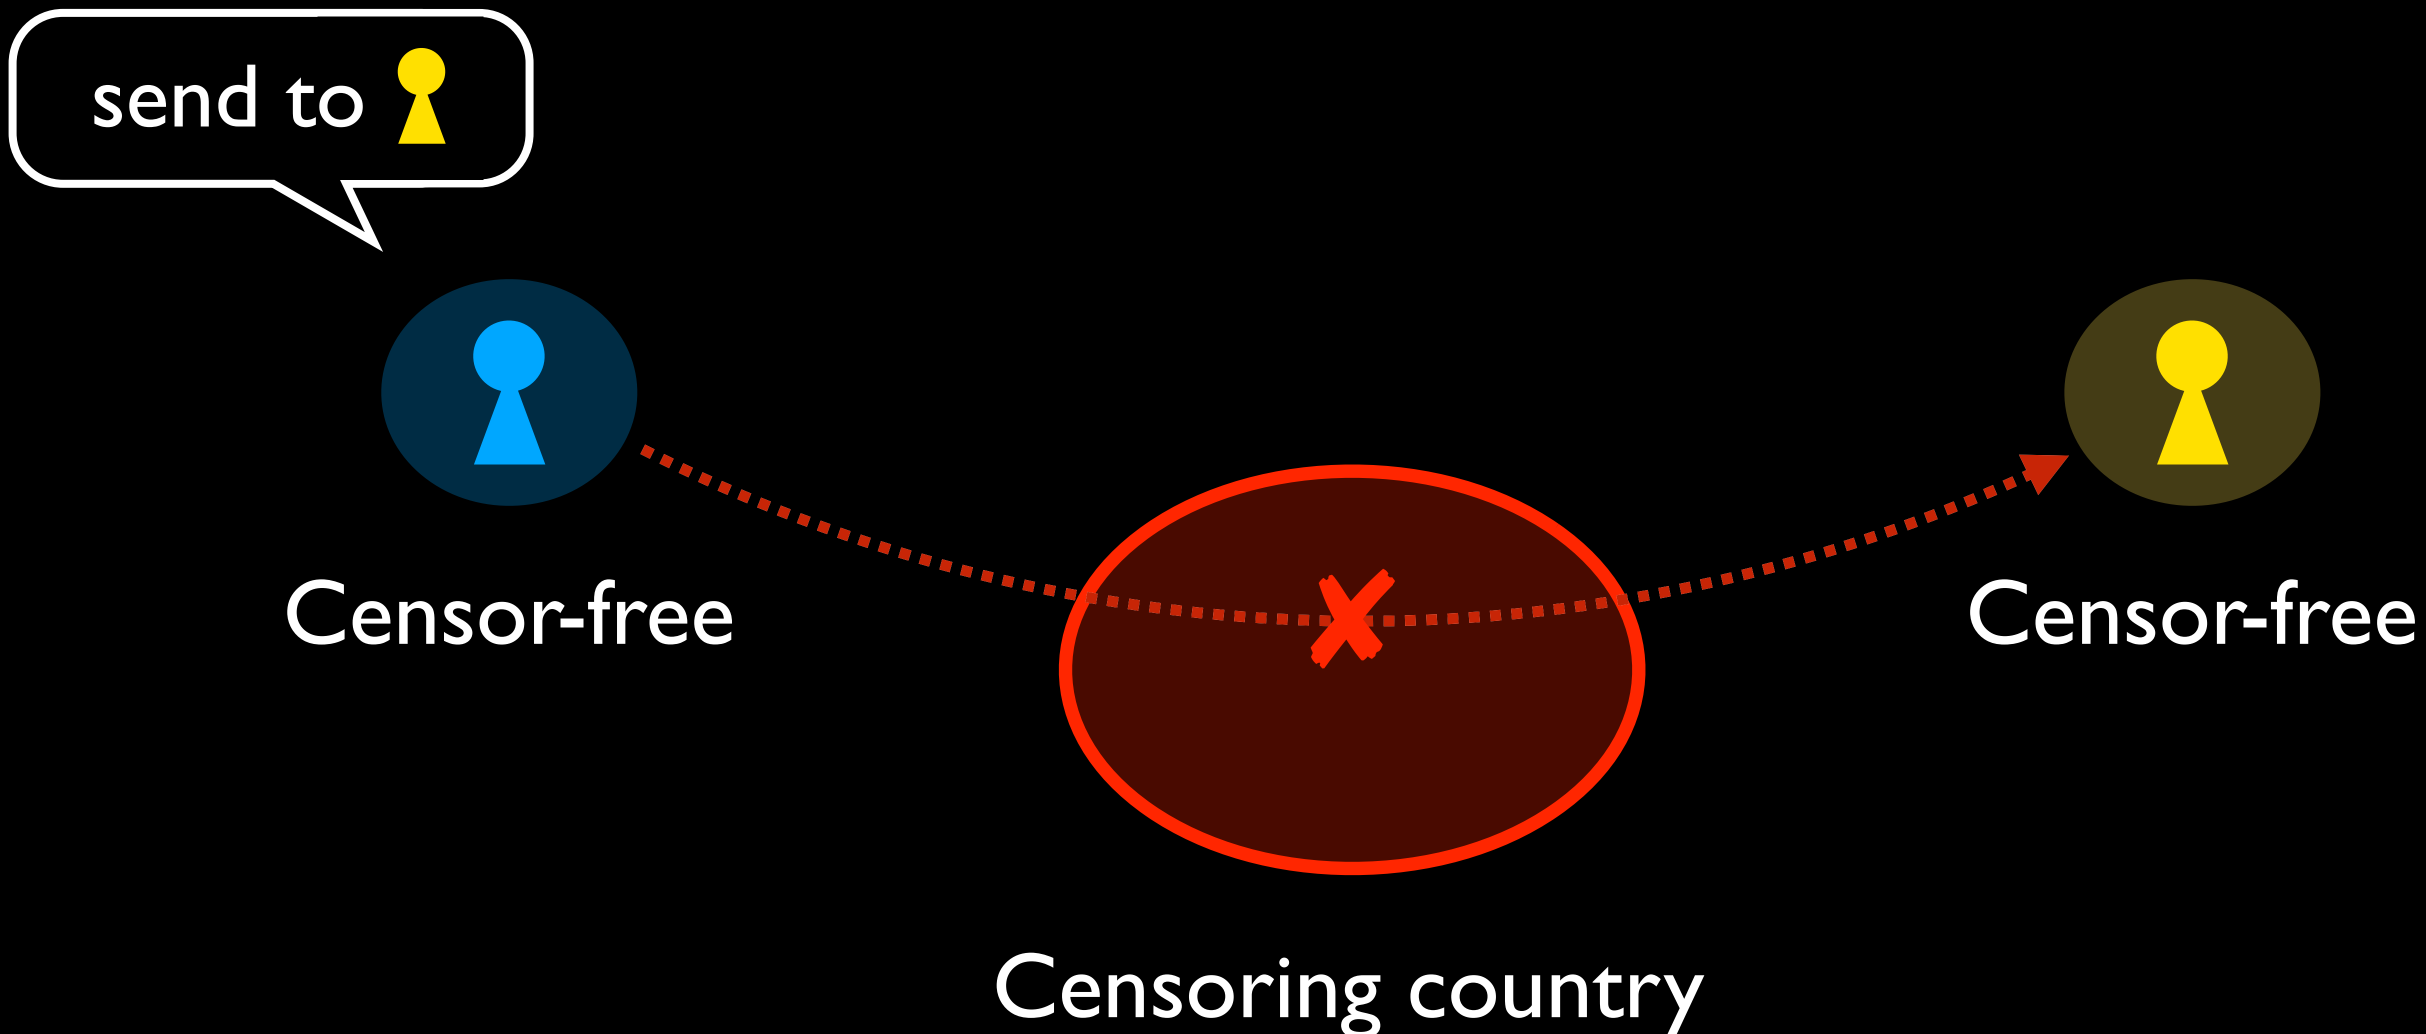

Encryption (HTTPS)

Anonymity (Tor)

Hide info, but are still subject to censorship

This work

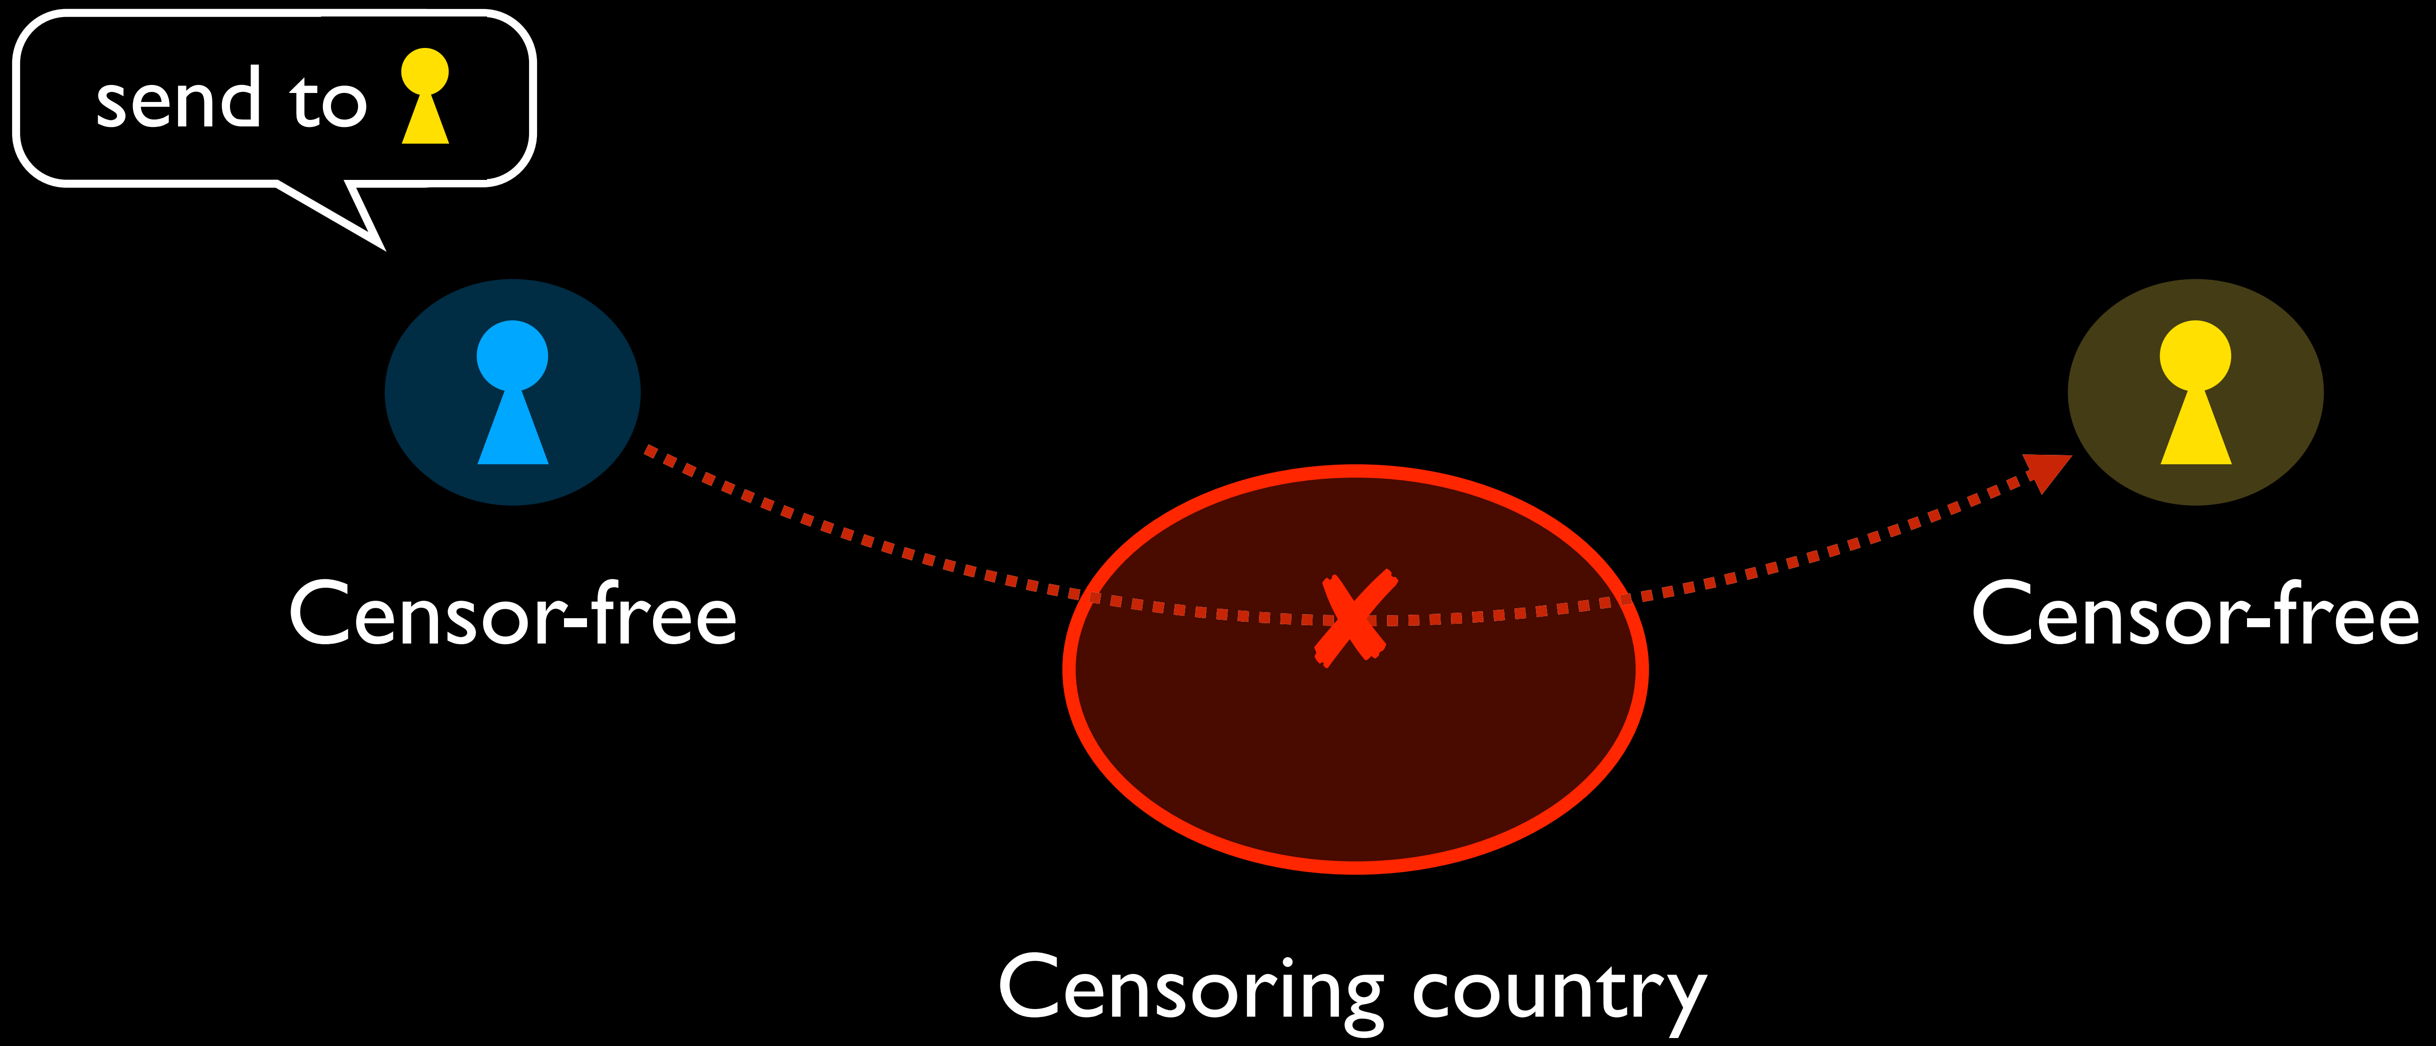

This work

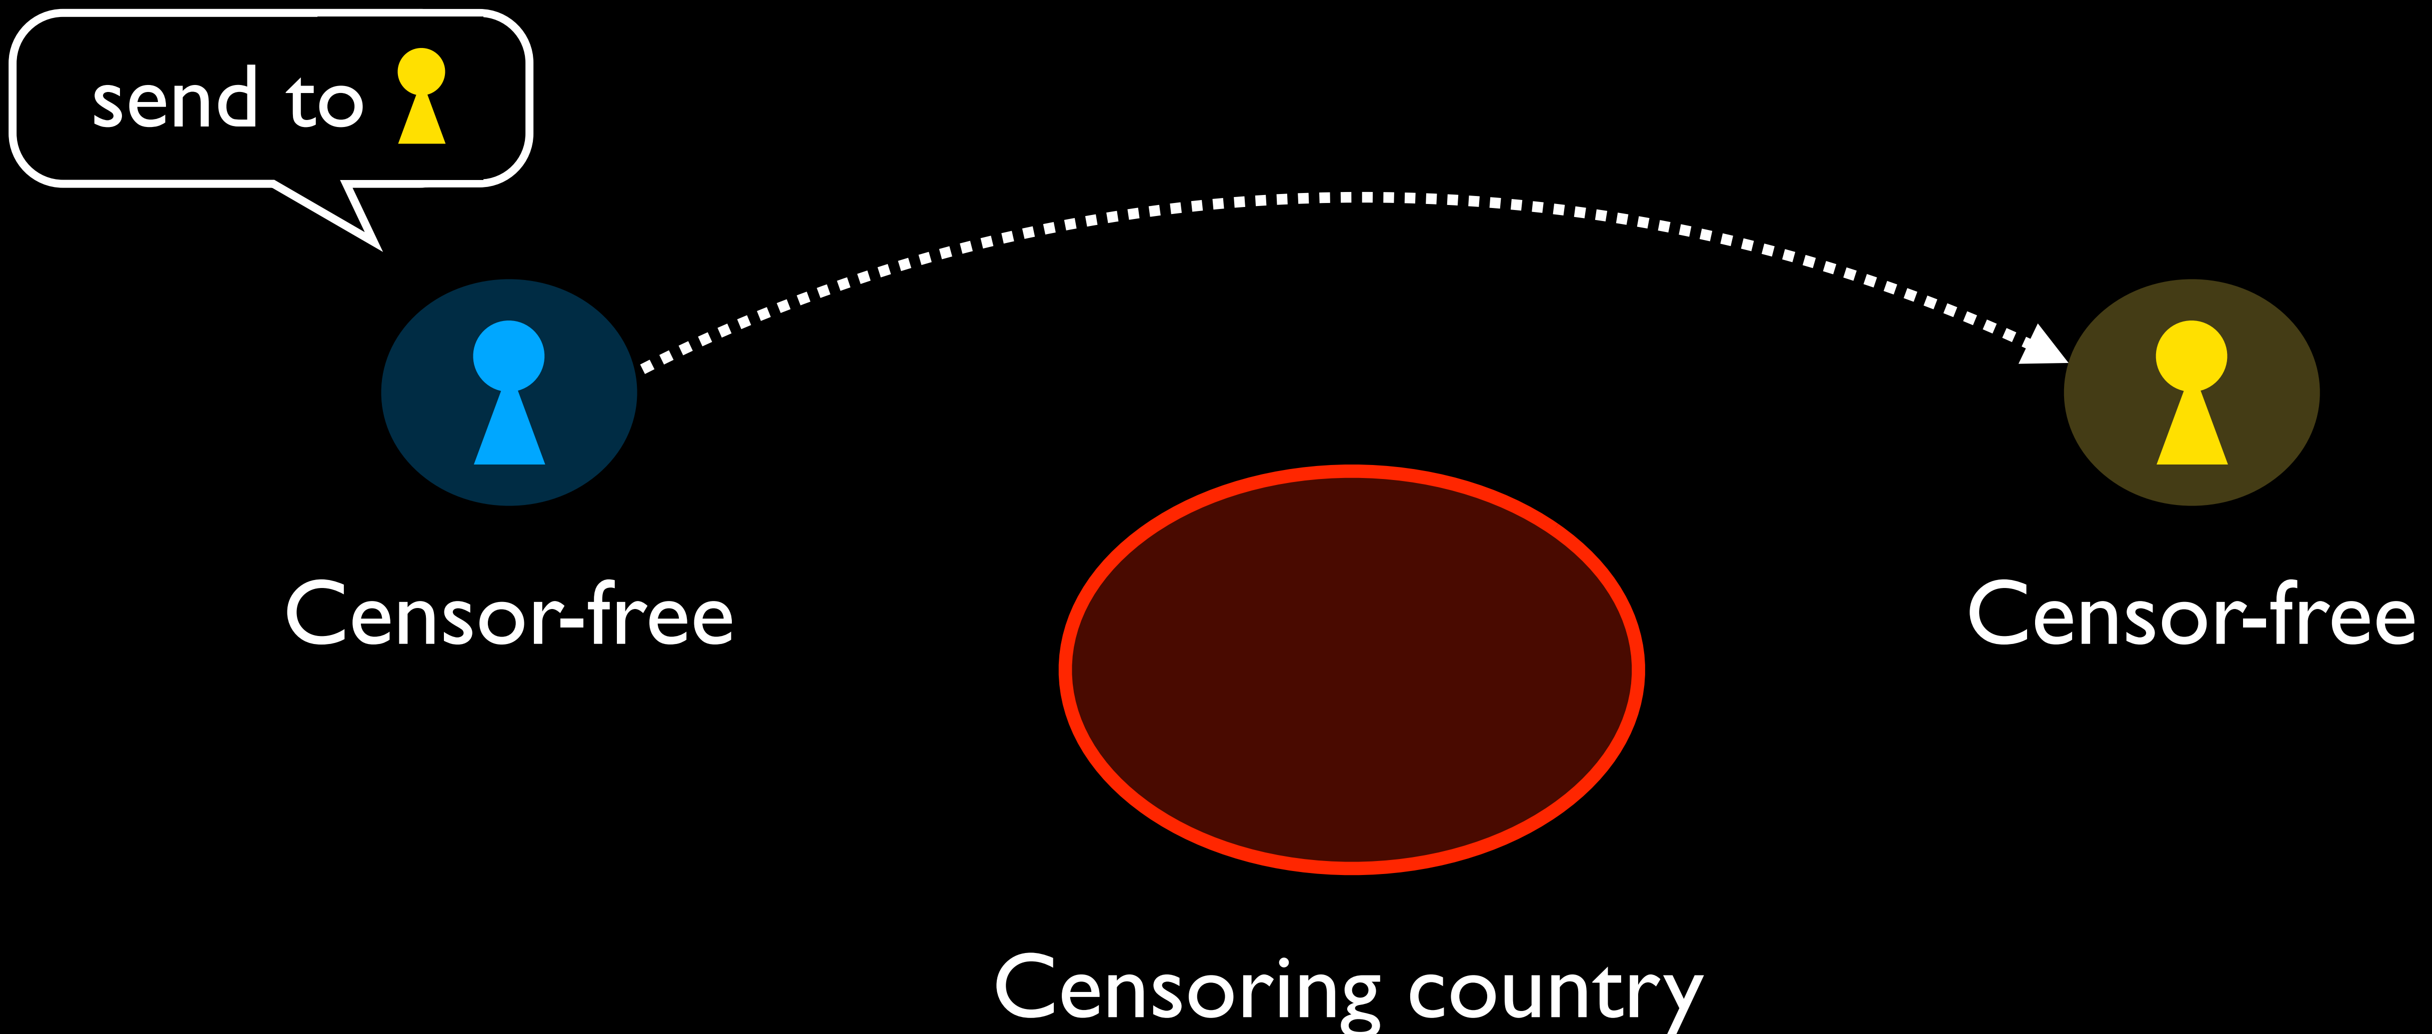

Provable avoidance routing

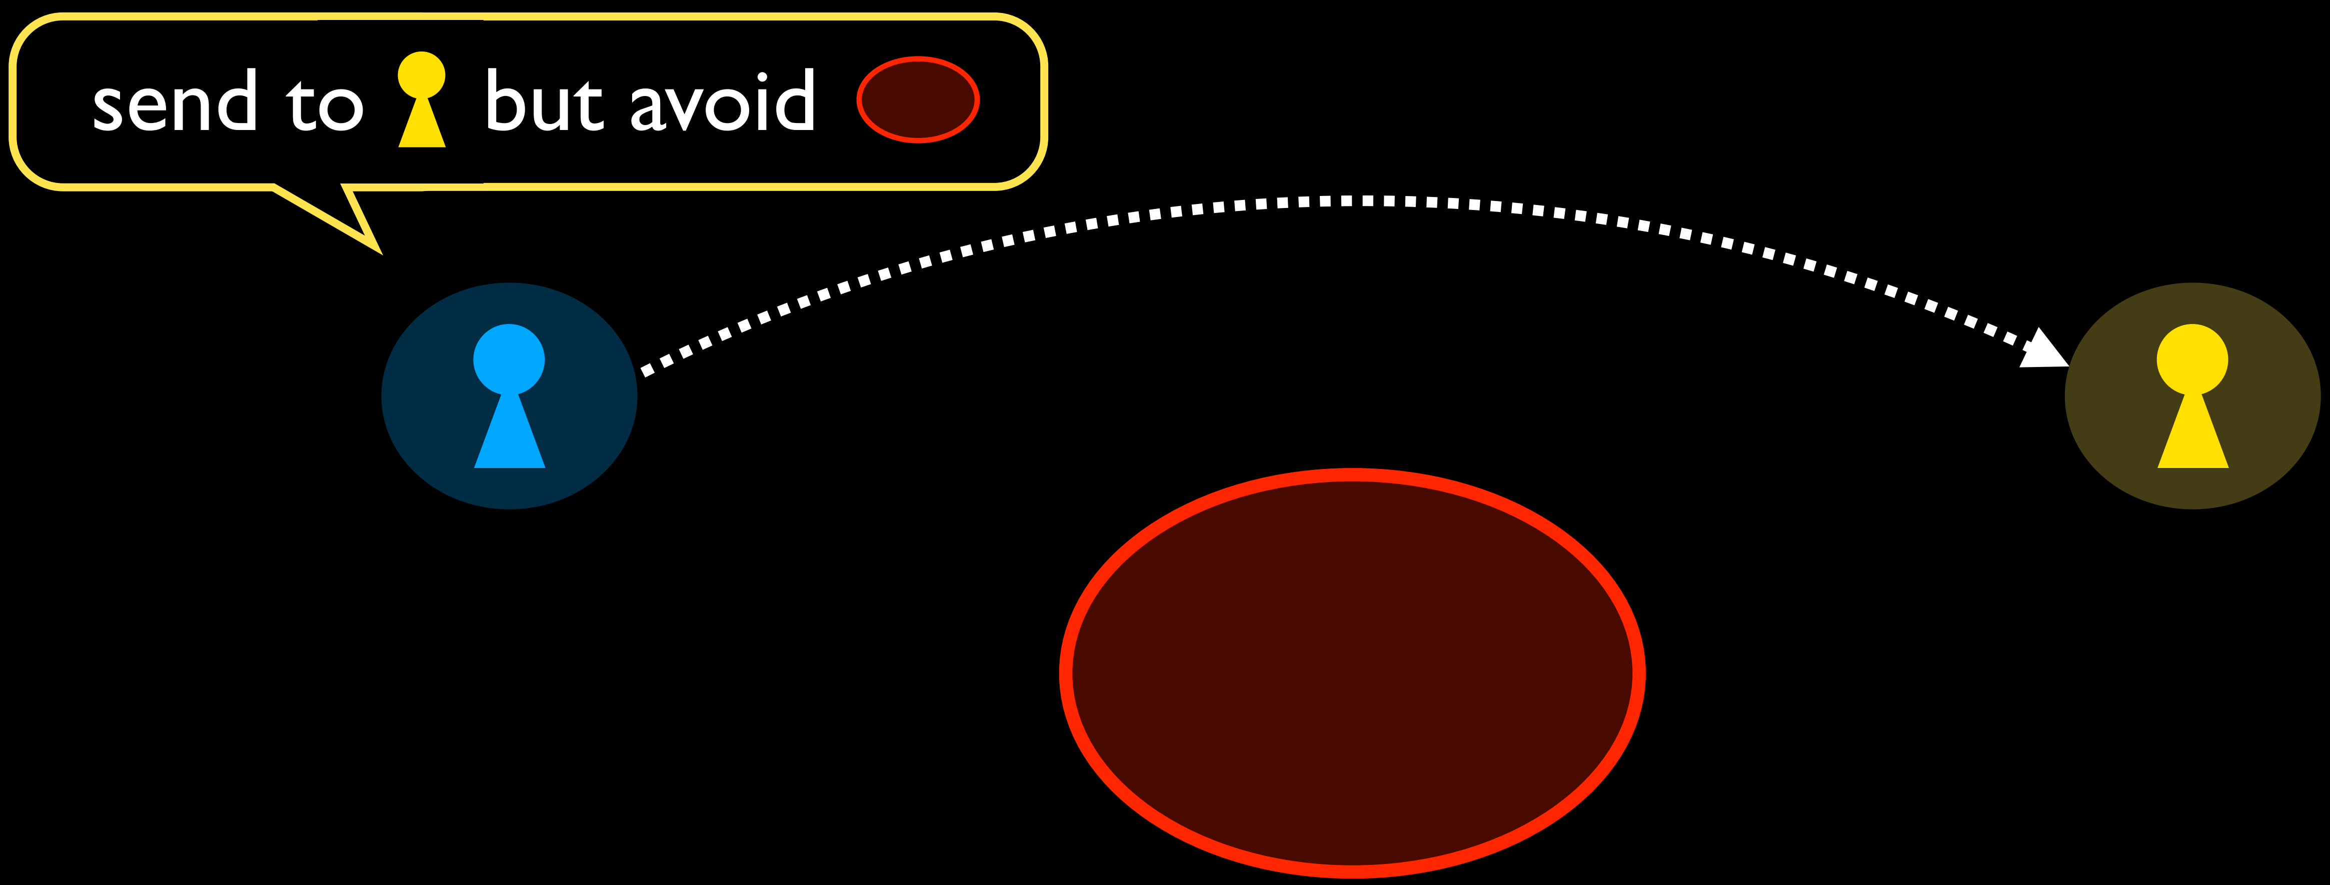

Provable avoidance routing

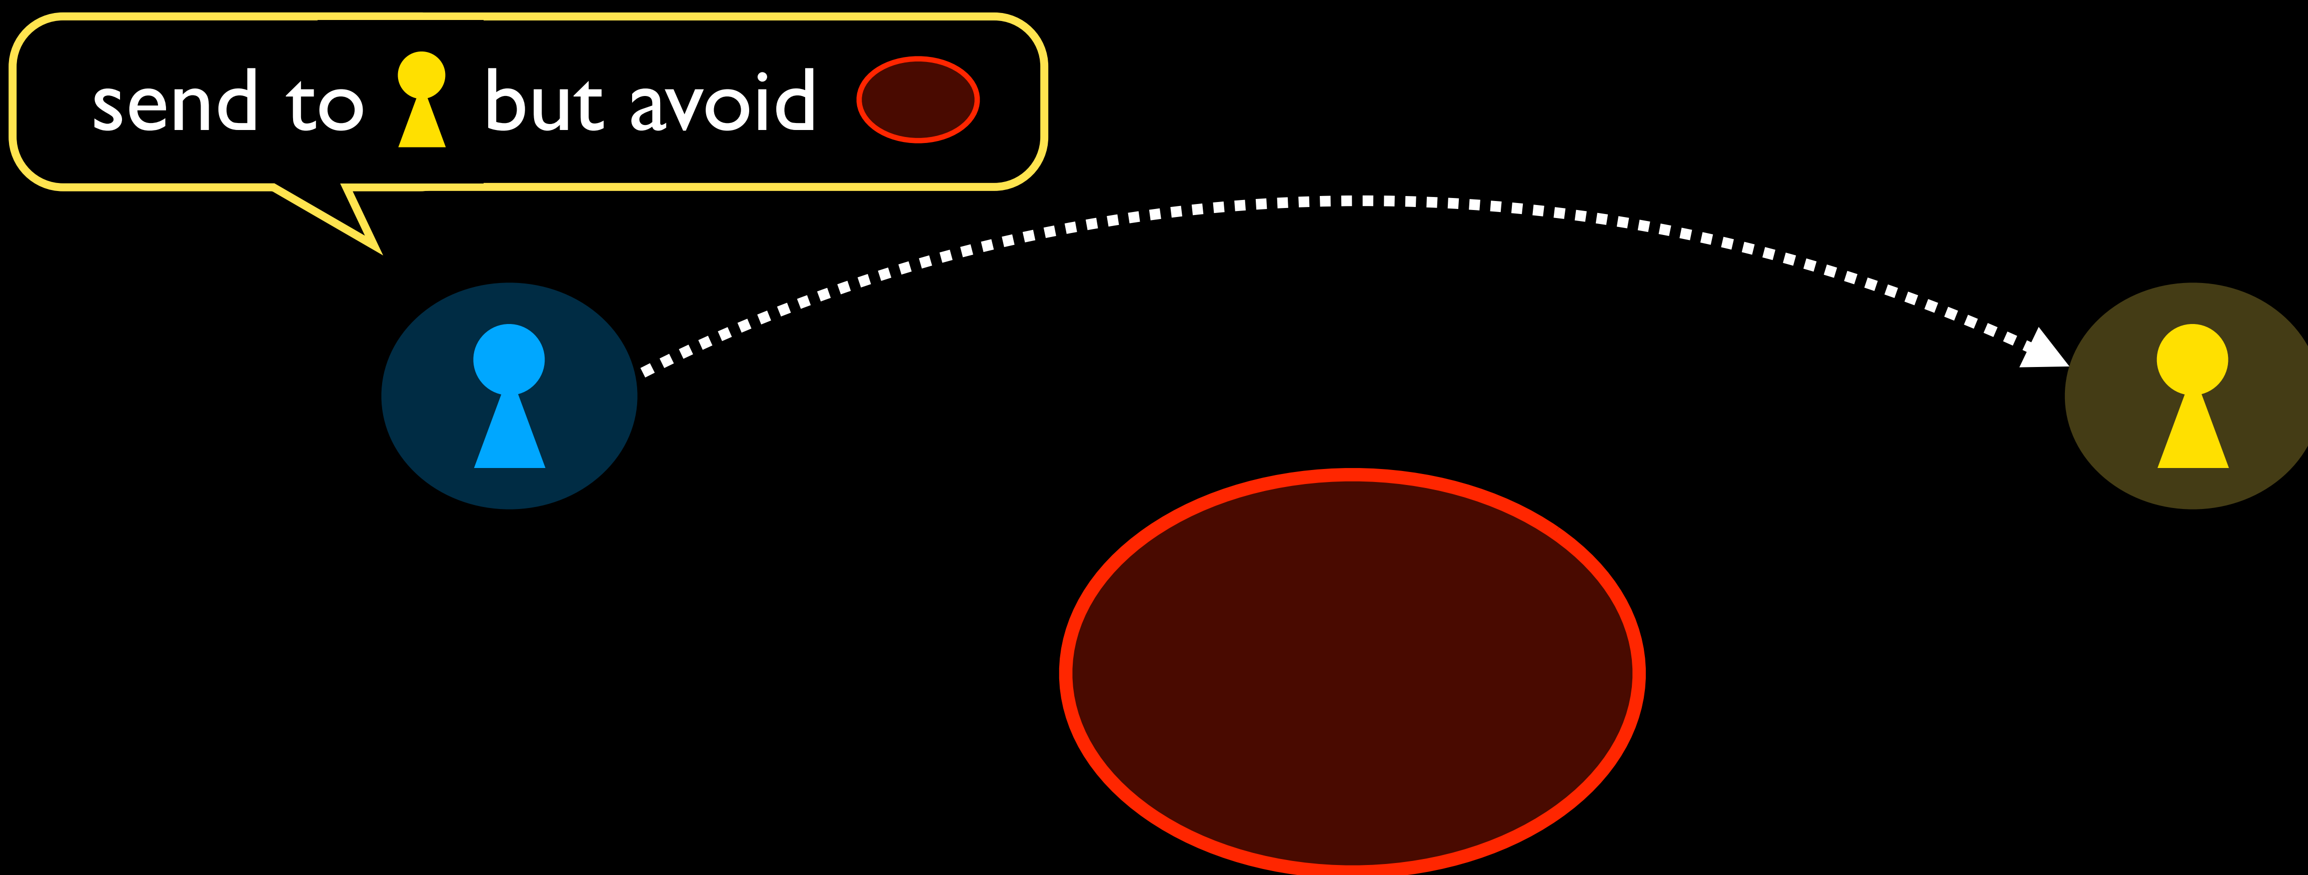

A broadly applicable primitive

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

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

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

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

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

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?*

Proving the impossible

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

A

Proving the impossible

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

A && A ⇒ !X

Mutually exclusive

Proving the impossible

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

A && A ⇒ !X ⇒ !X

Mutually exclusive

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

Achieving provable avoidance

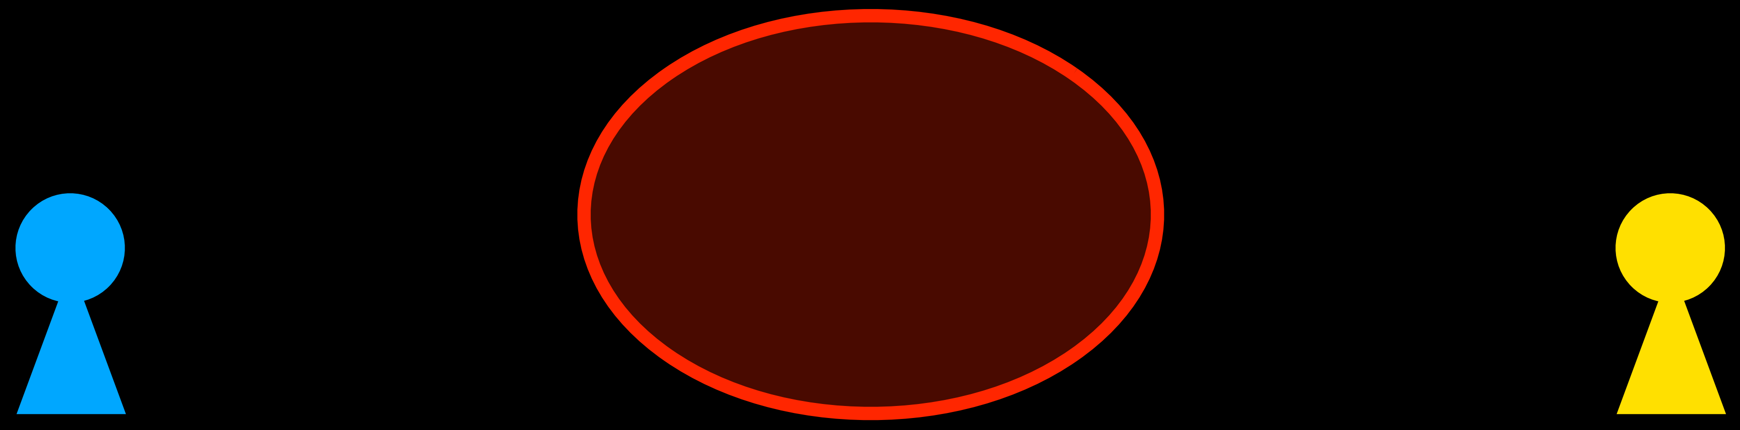

Achieving provable avoidance

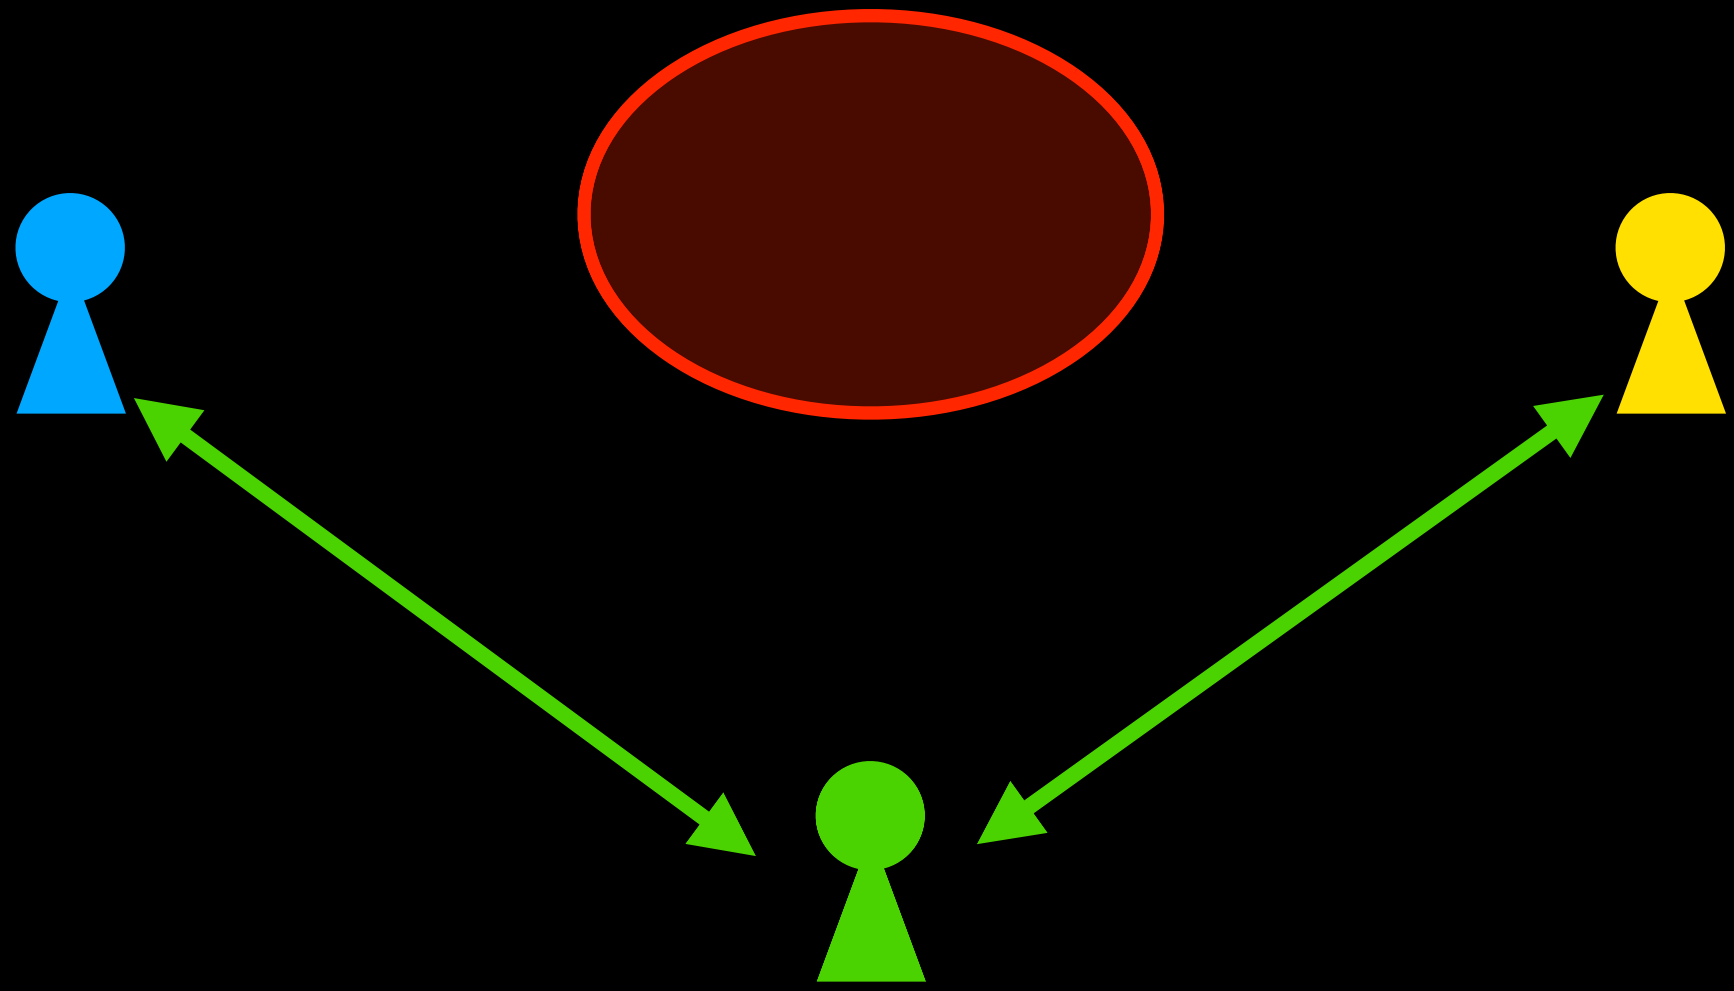

Solicit participation from a relay

Achieving provable avoidance

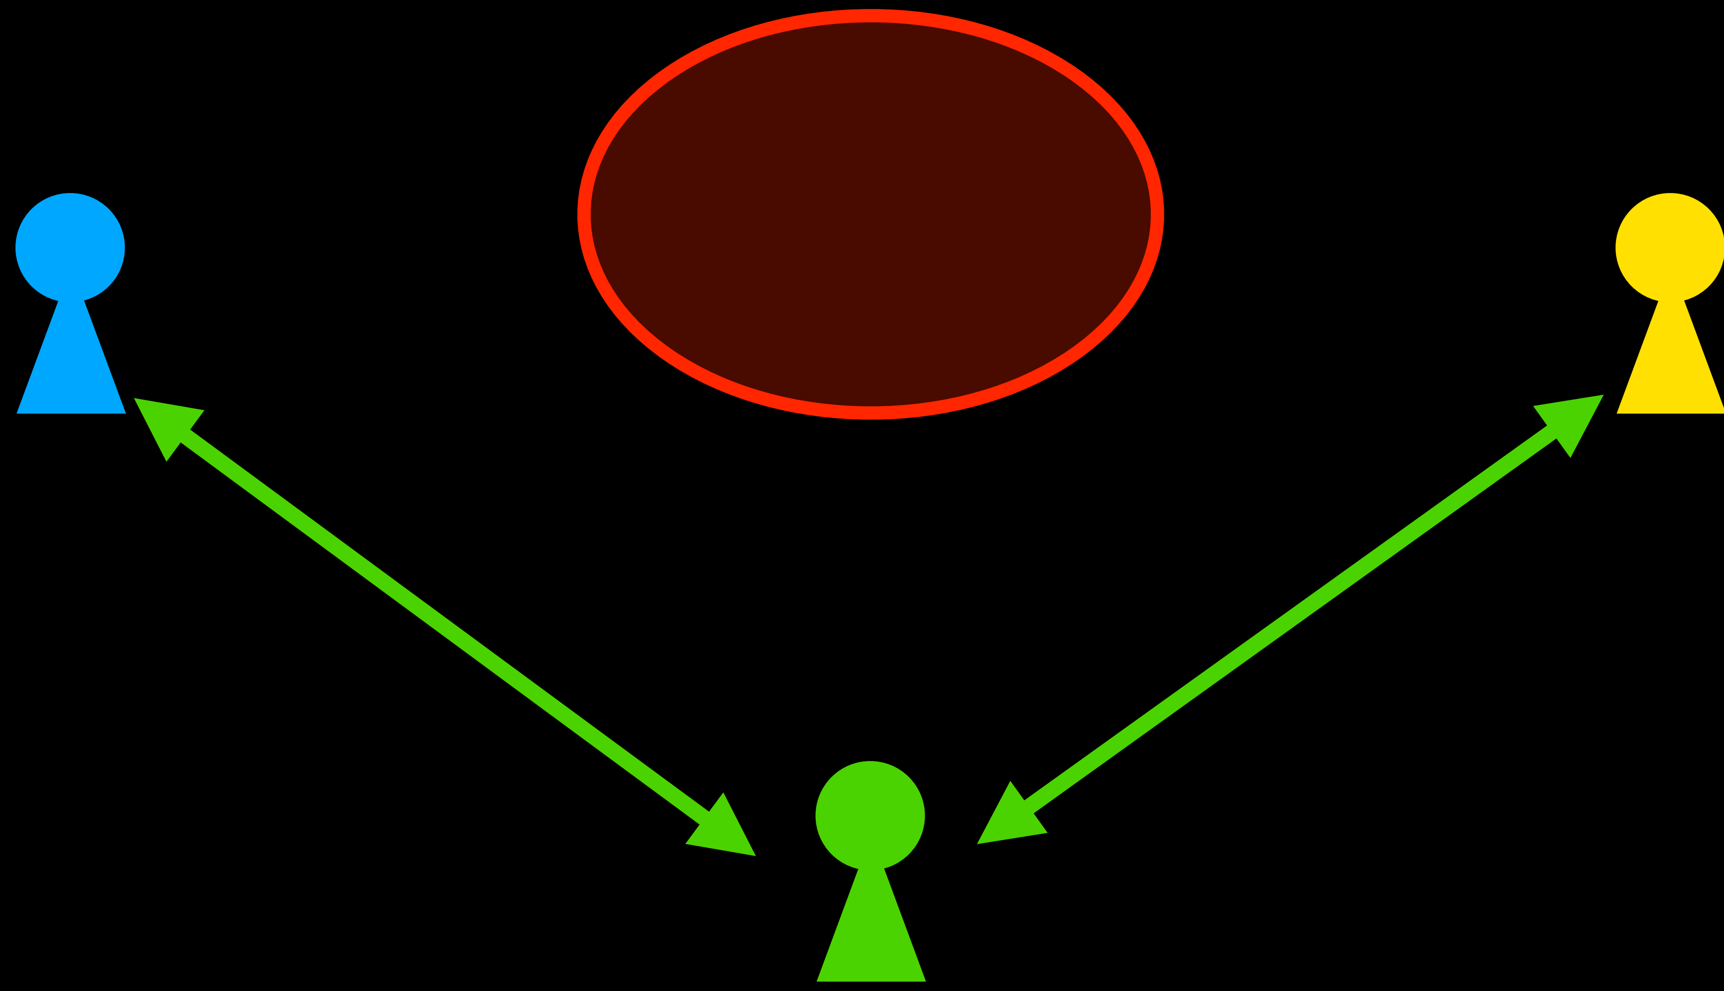

Reply contains a MAC from

⇒ The packet traversed

Achieving provable avoidance

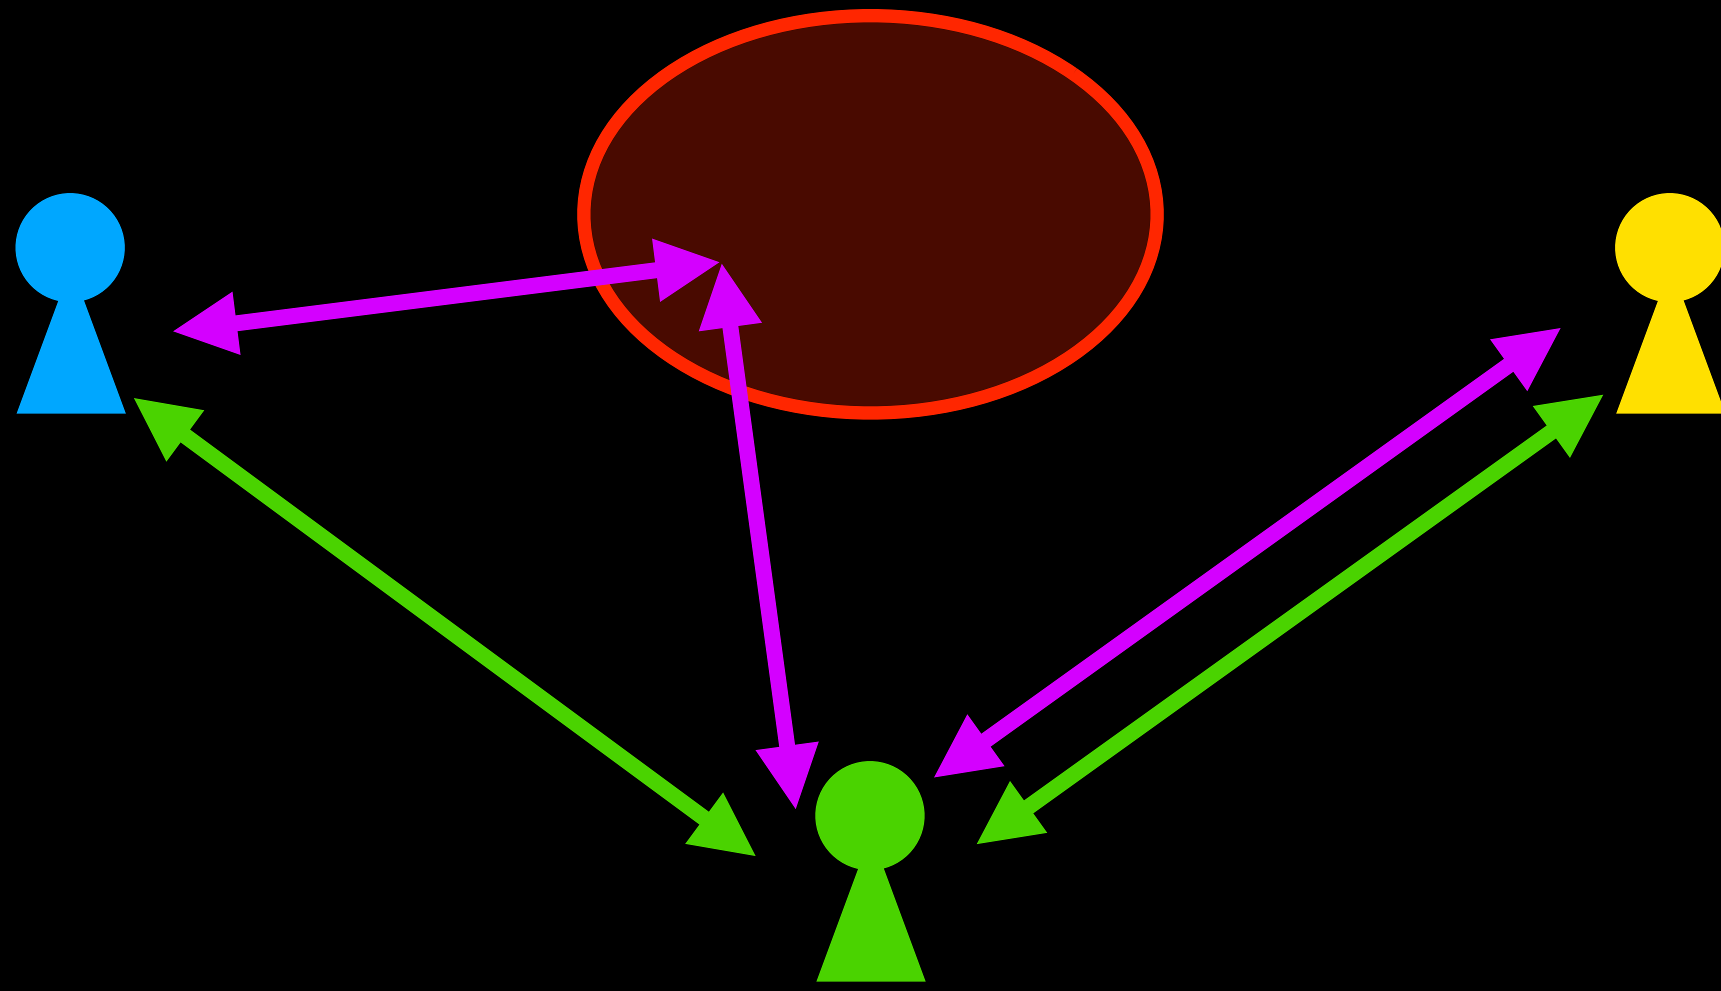

Reply contains a MAC from

⇒ The packet traversed

Achieving provable avoidance

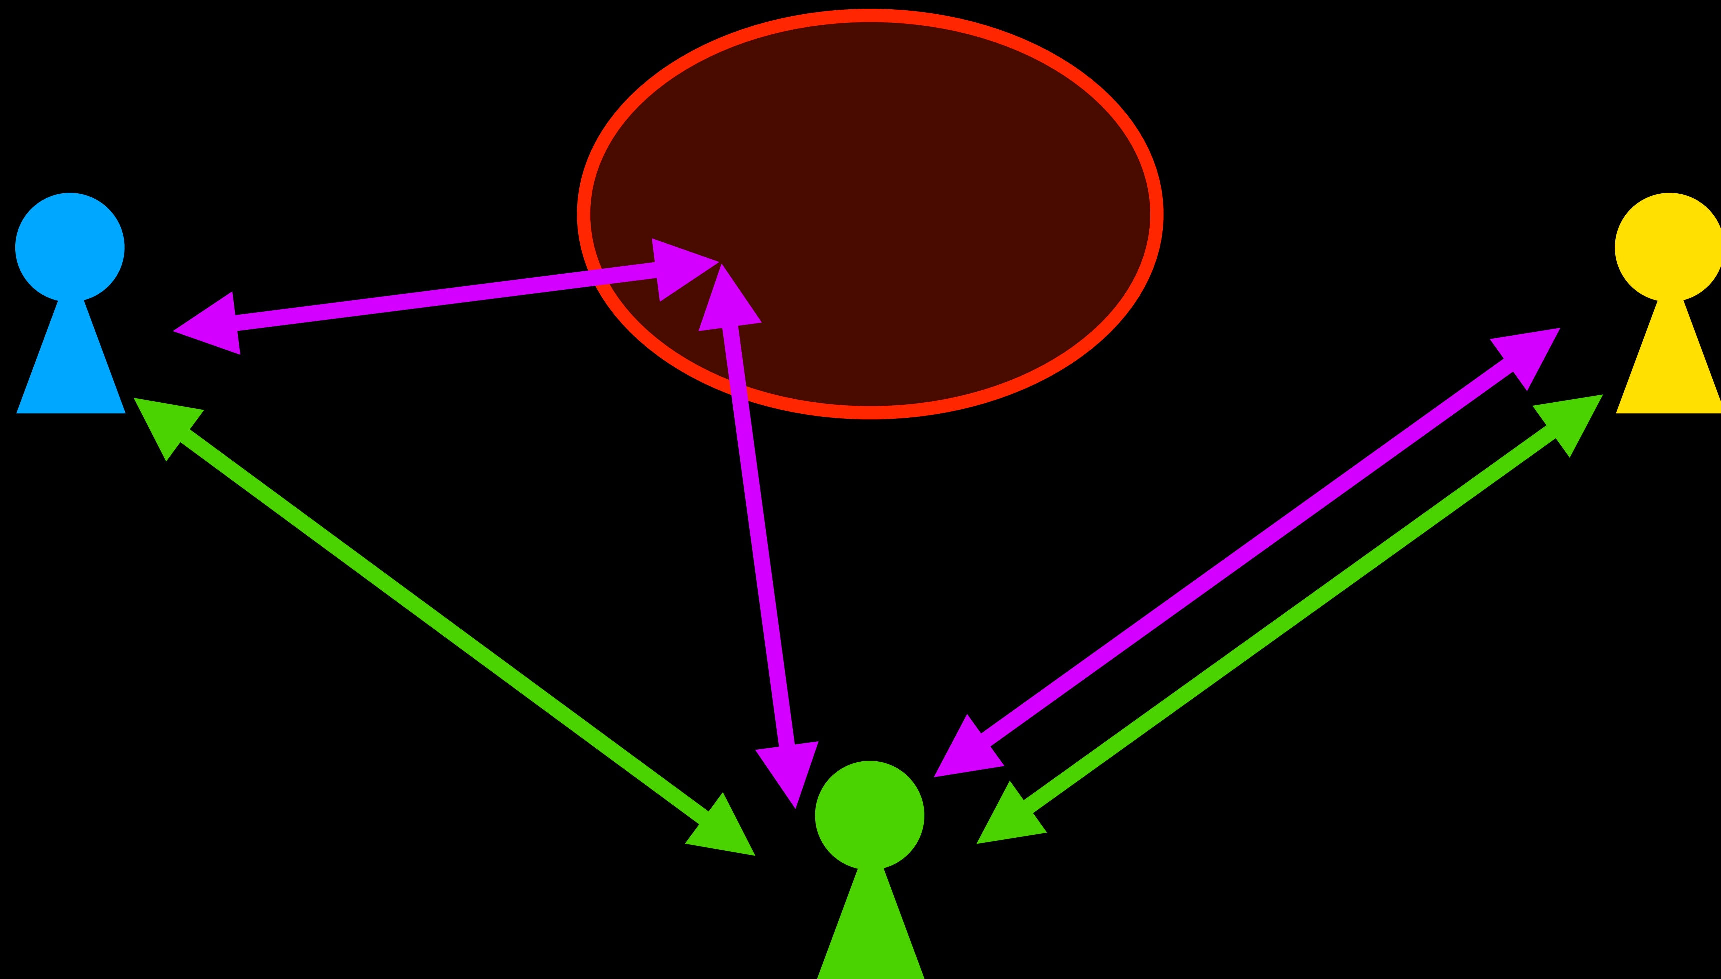

Reply contains a MAC from

⇒ The packet traversed

Achieving provable avoidance

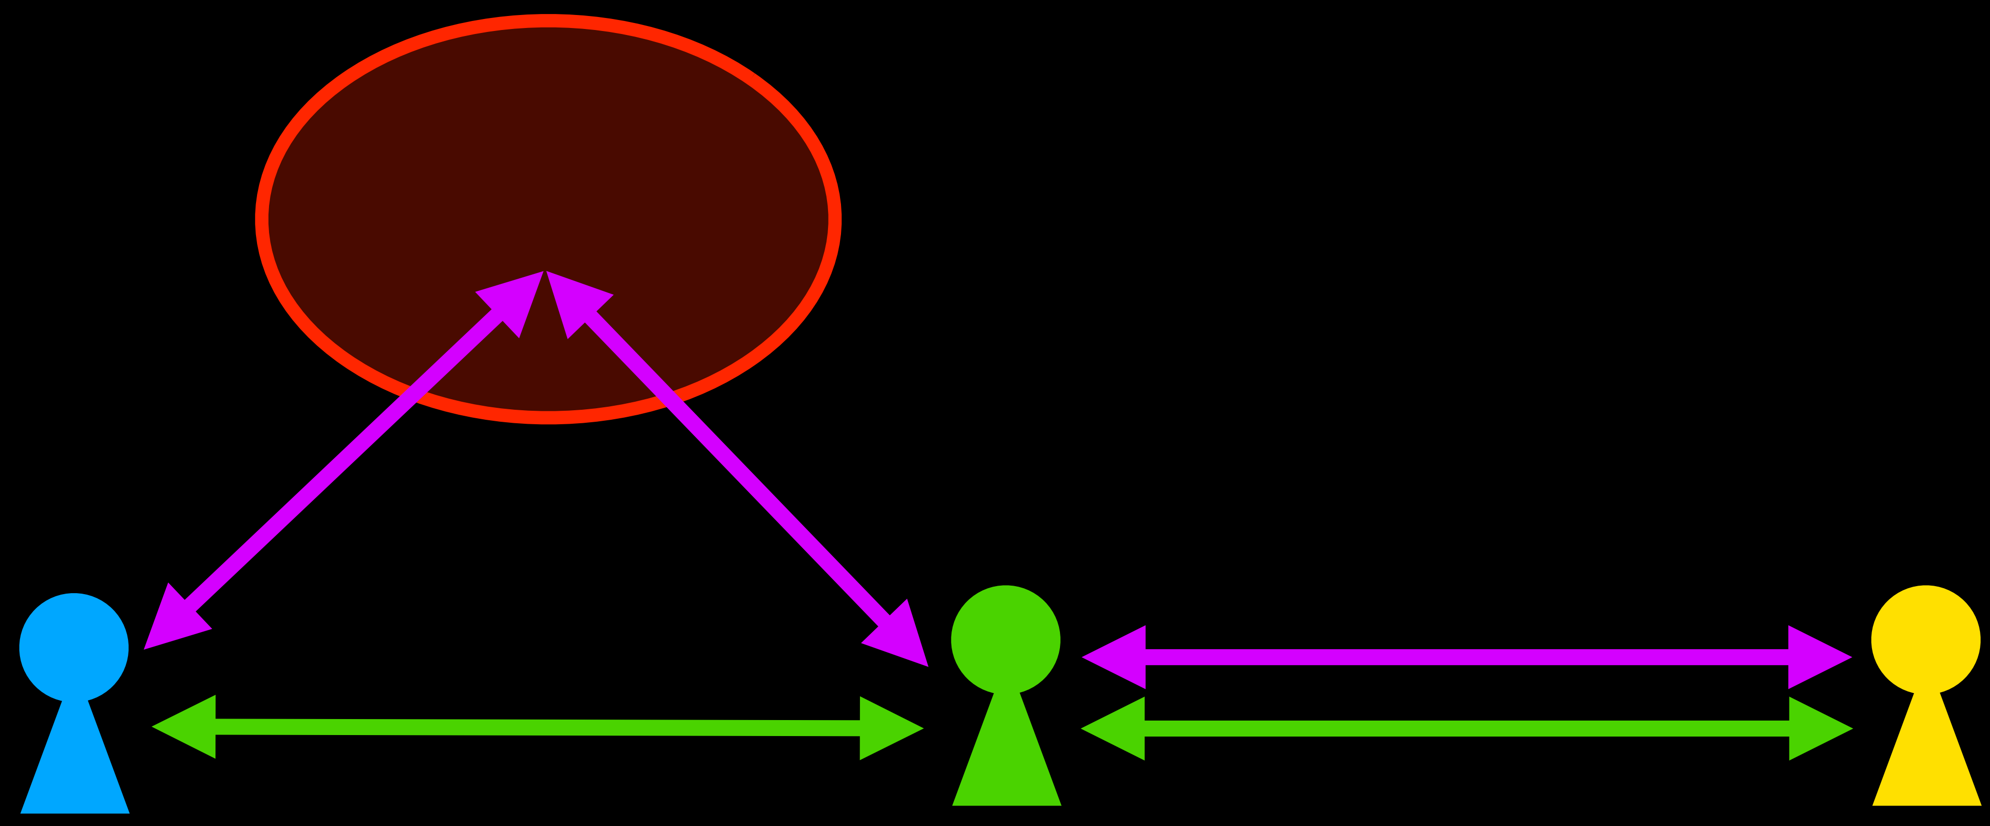

Achieving provable avoidance

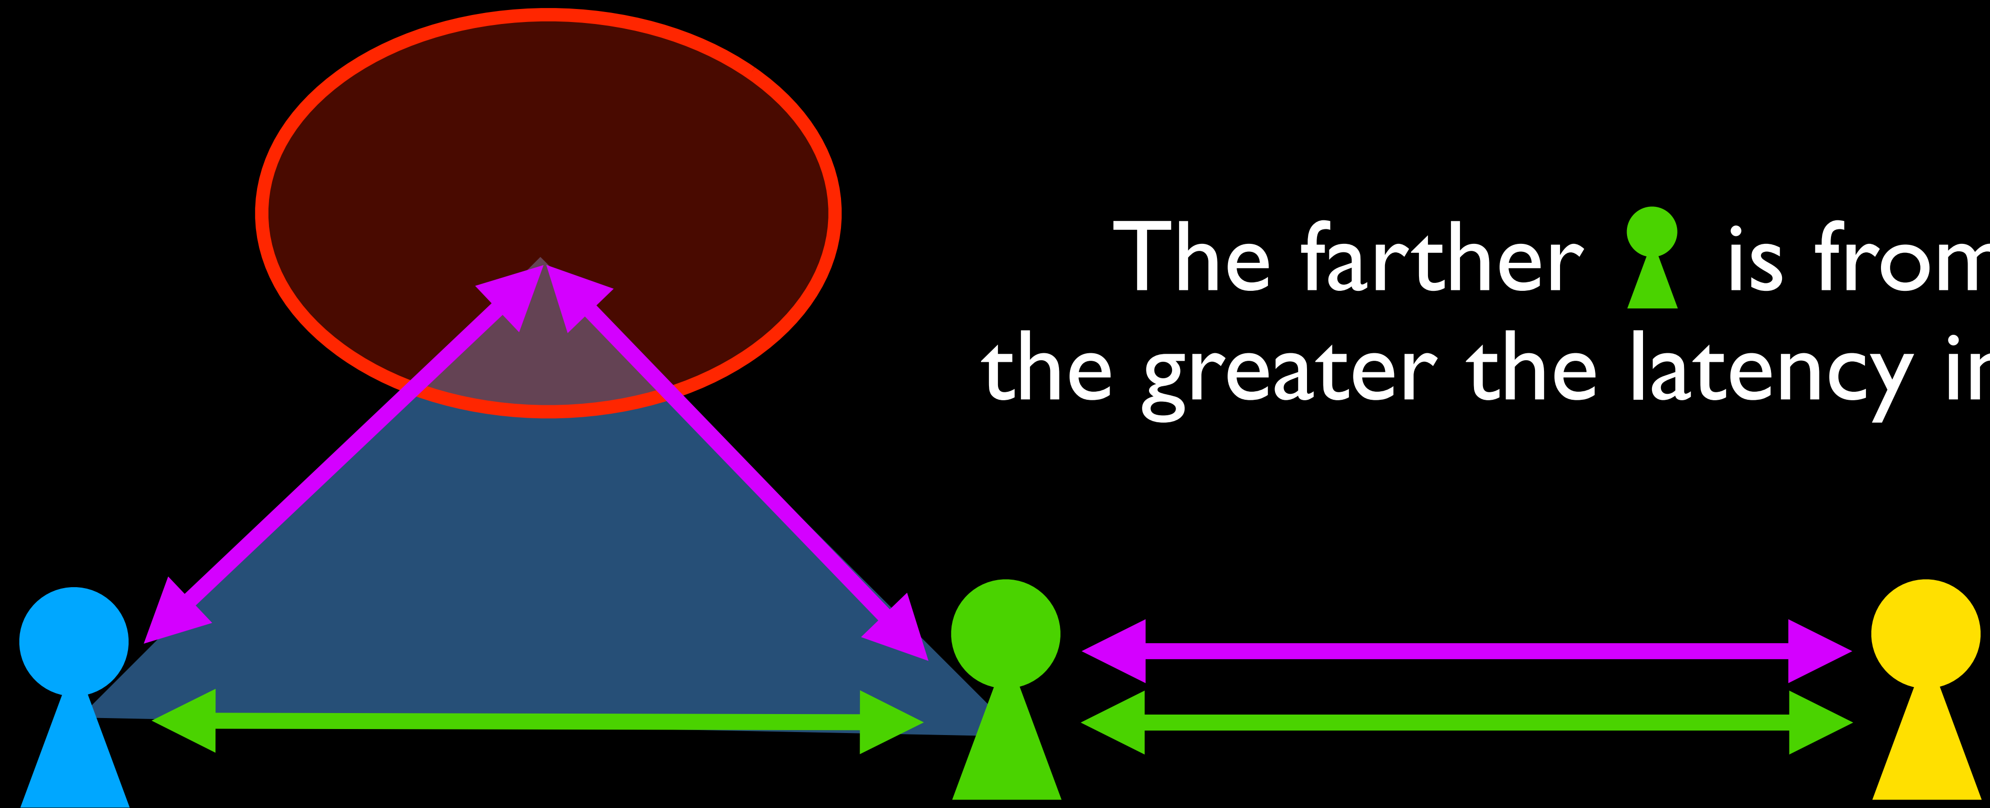

The farther is from
the greater the latency increase

Achieving provable avoidance

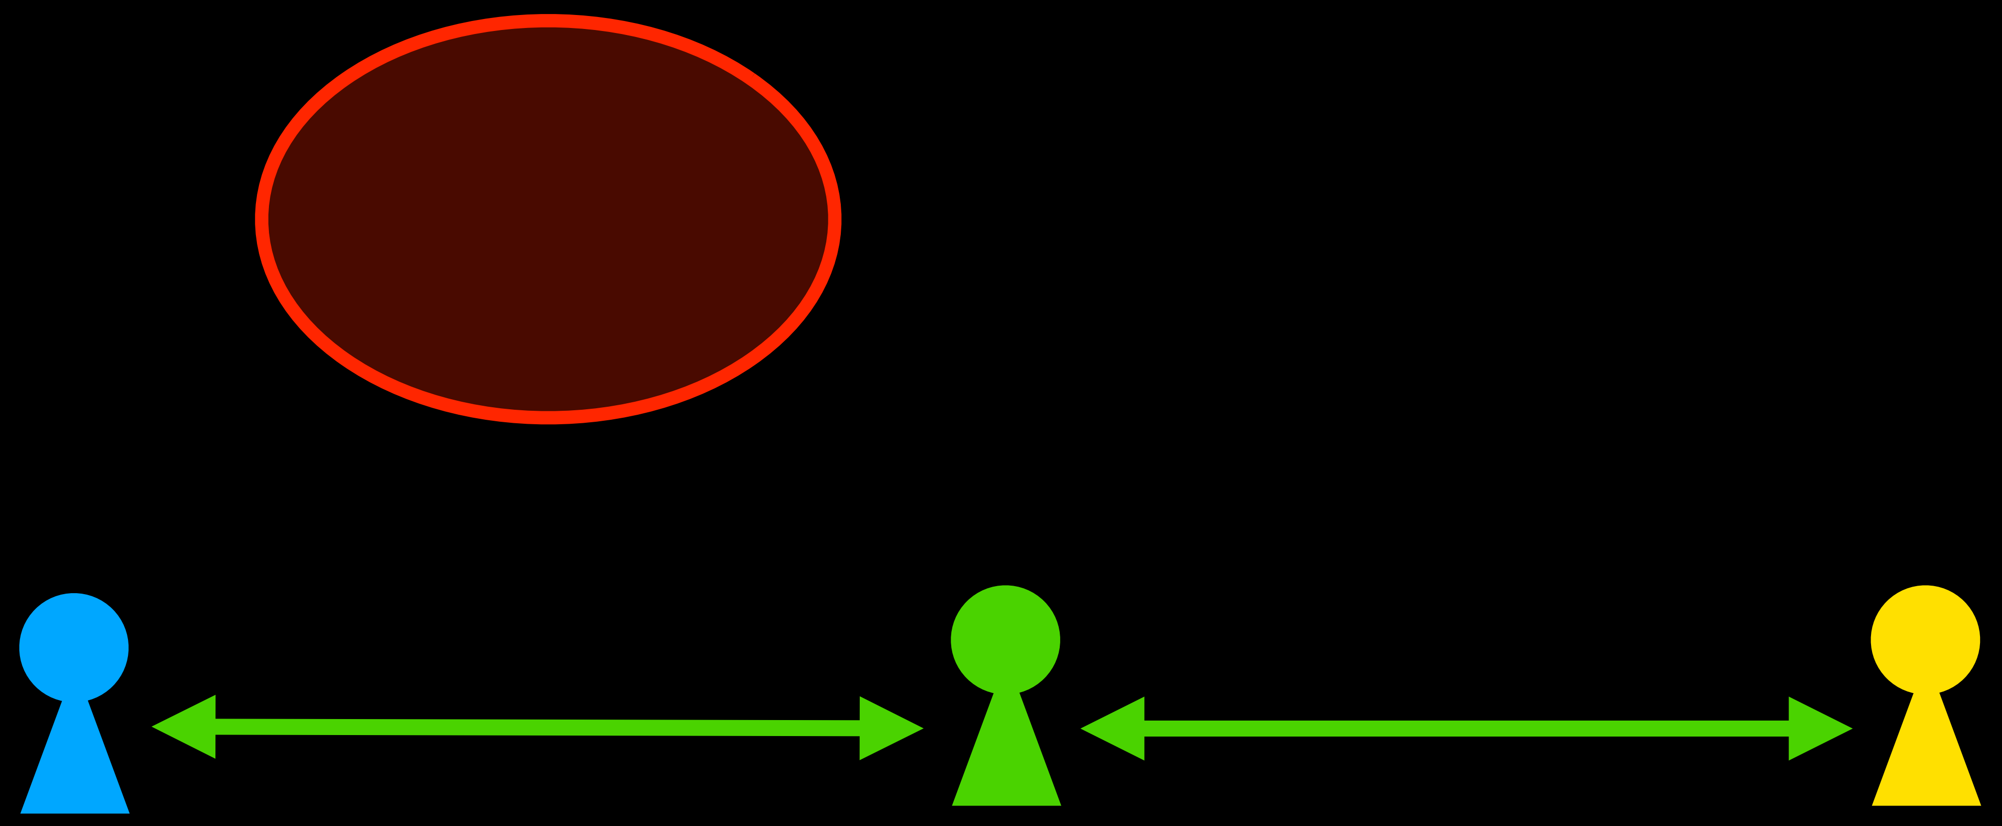

Achieving provable avoidance

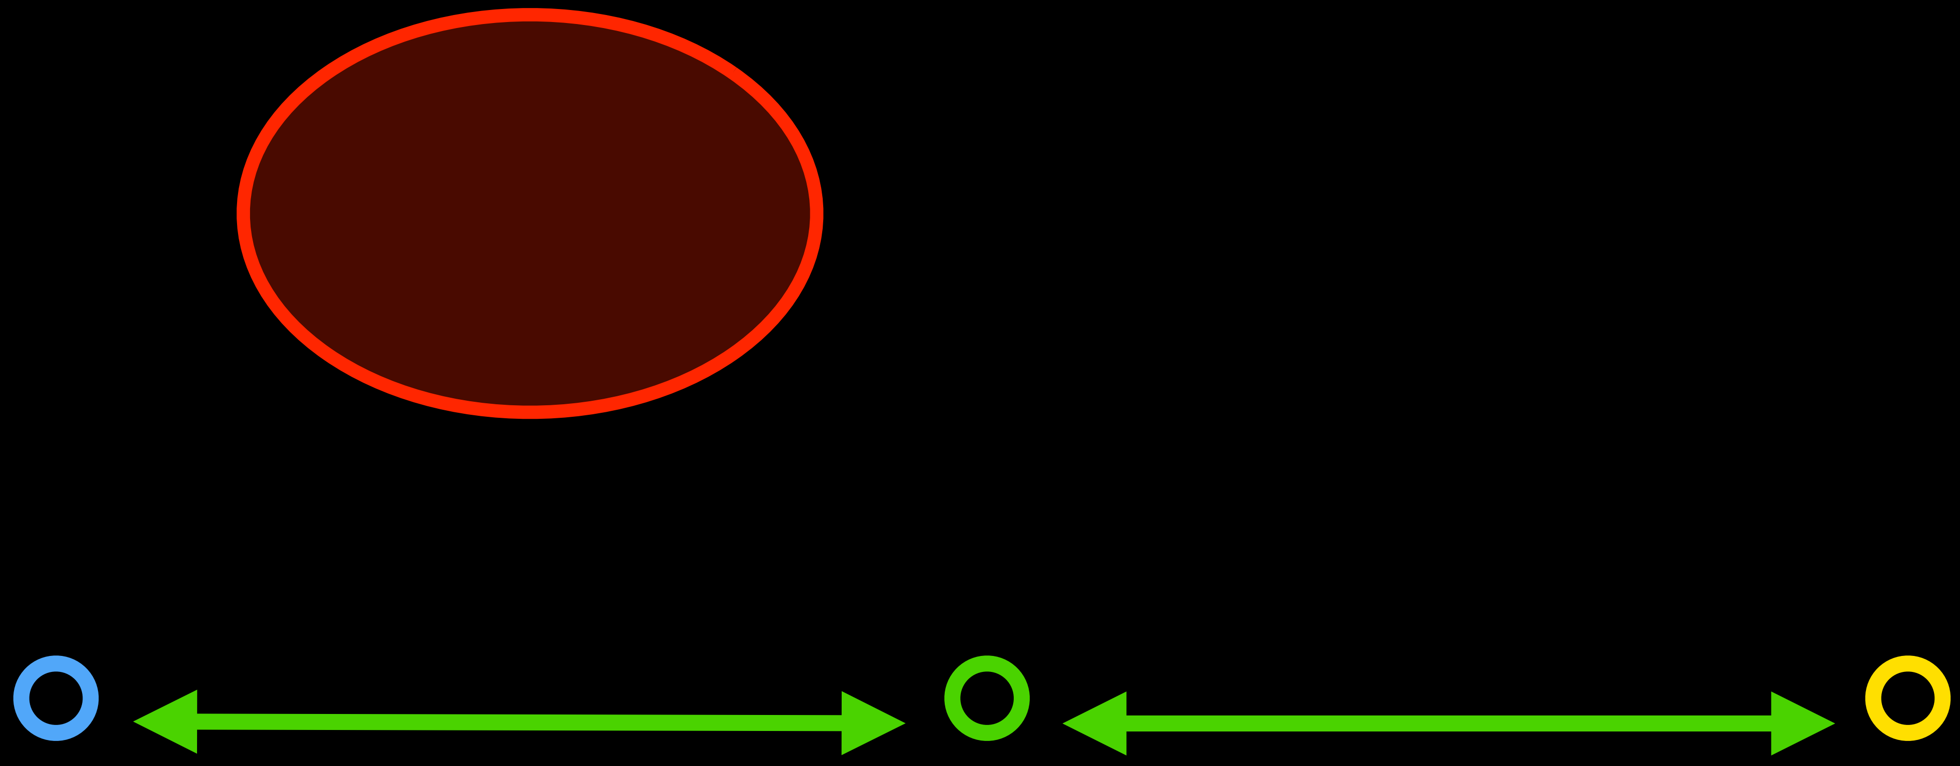

Achieving provable avoidance

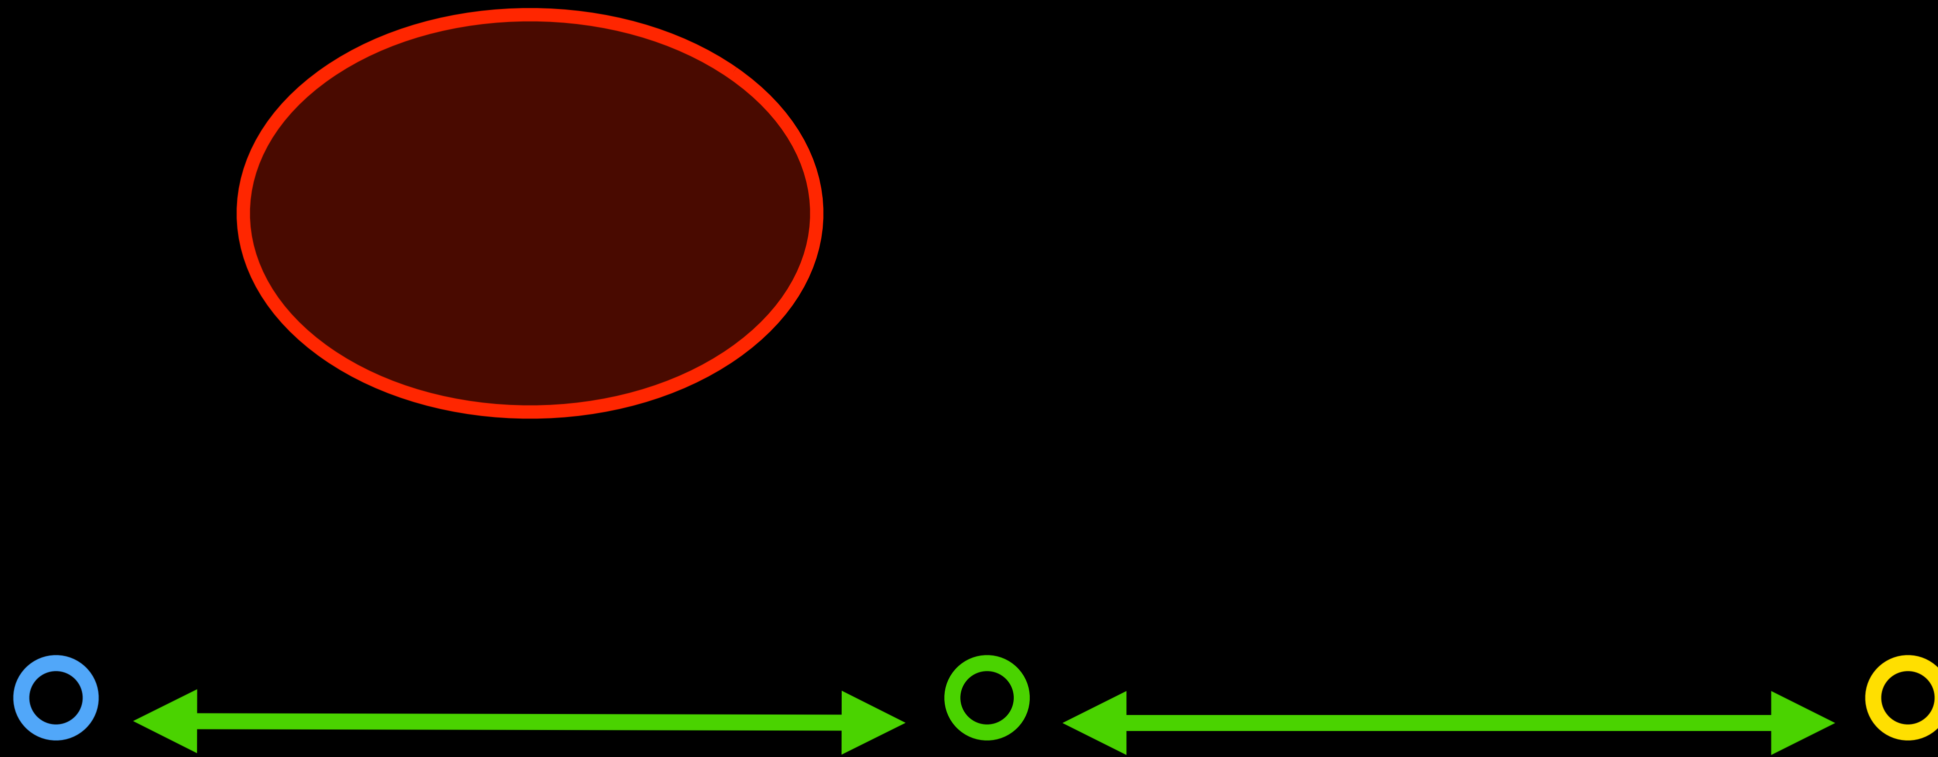

The shortest *possible* distance

thru  and 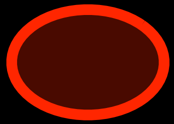

Achieving provable avoidance

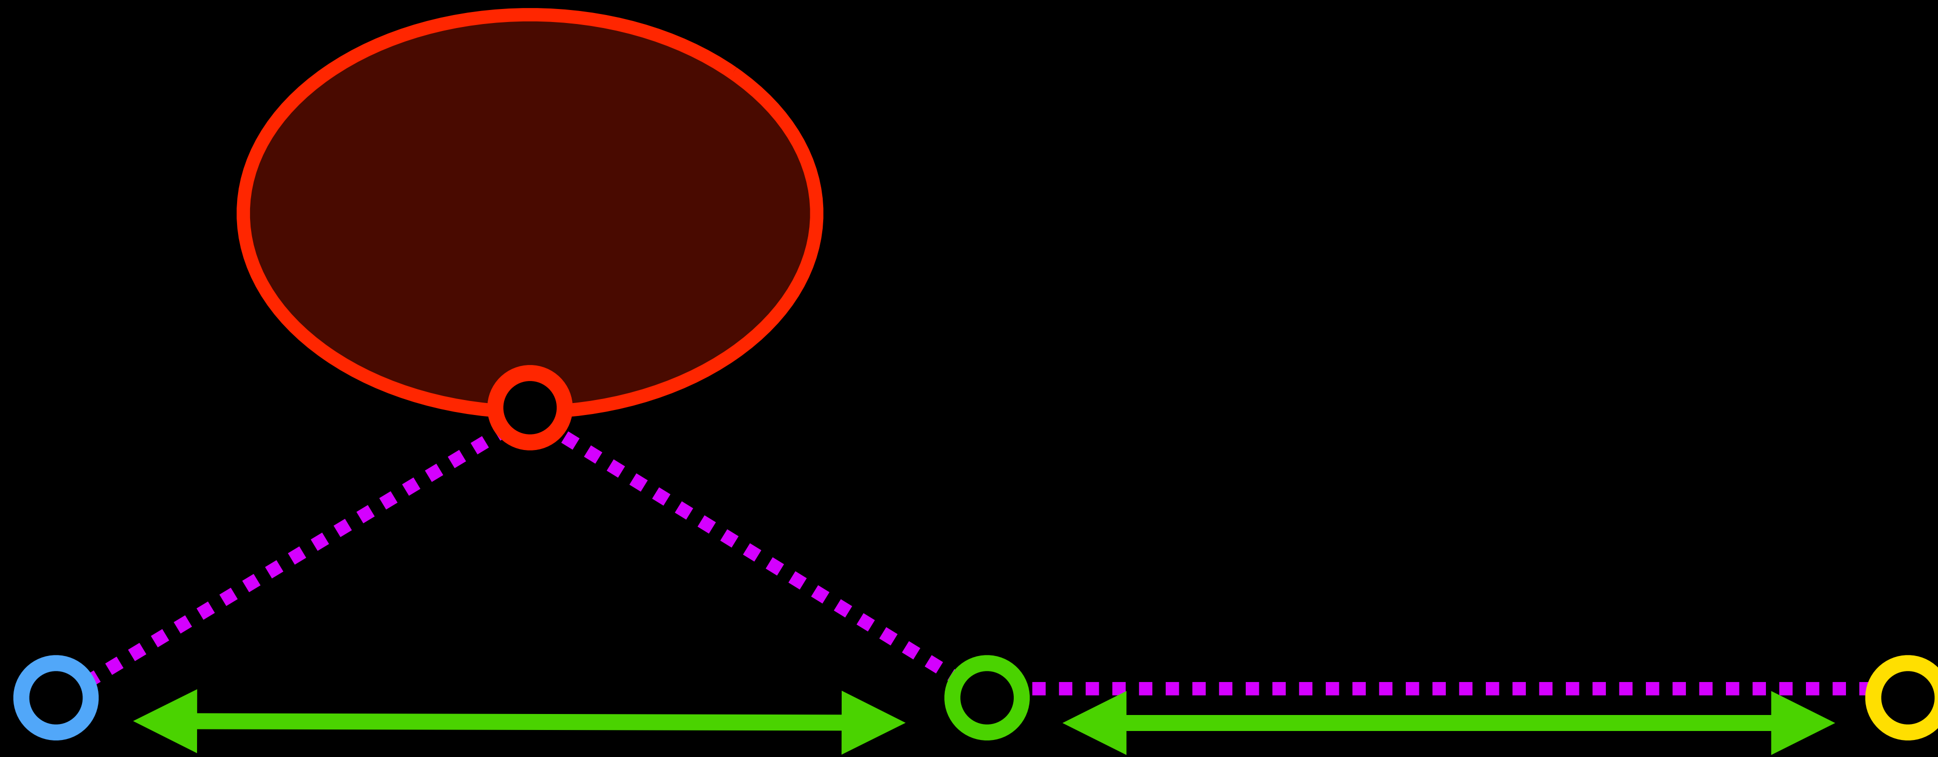

The shortest *possible* distance
thru and

Achieving provable avoidance

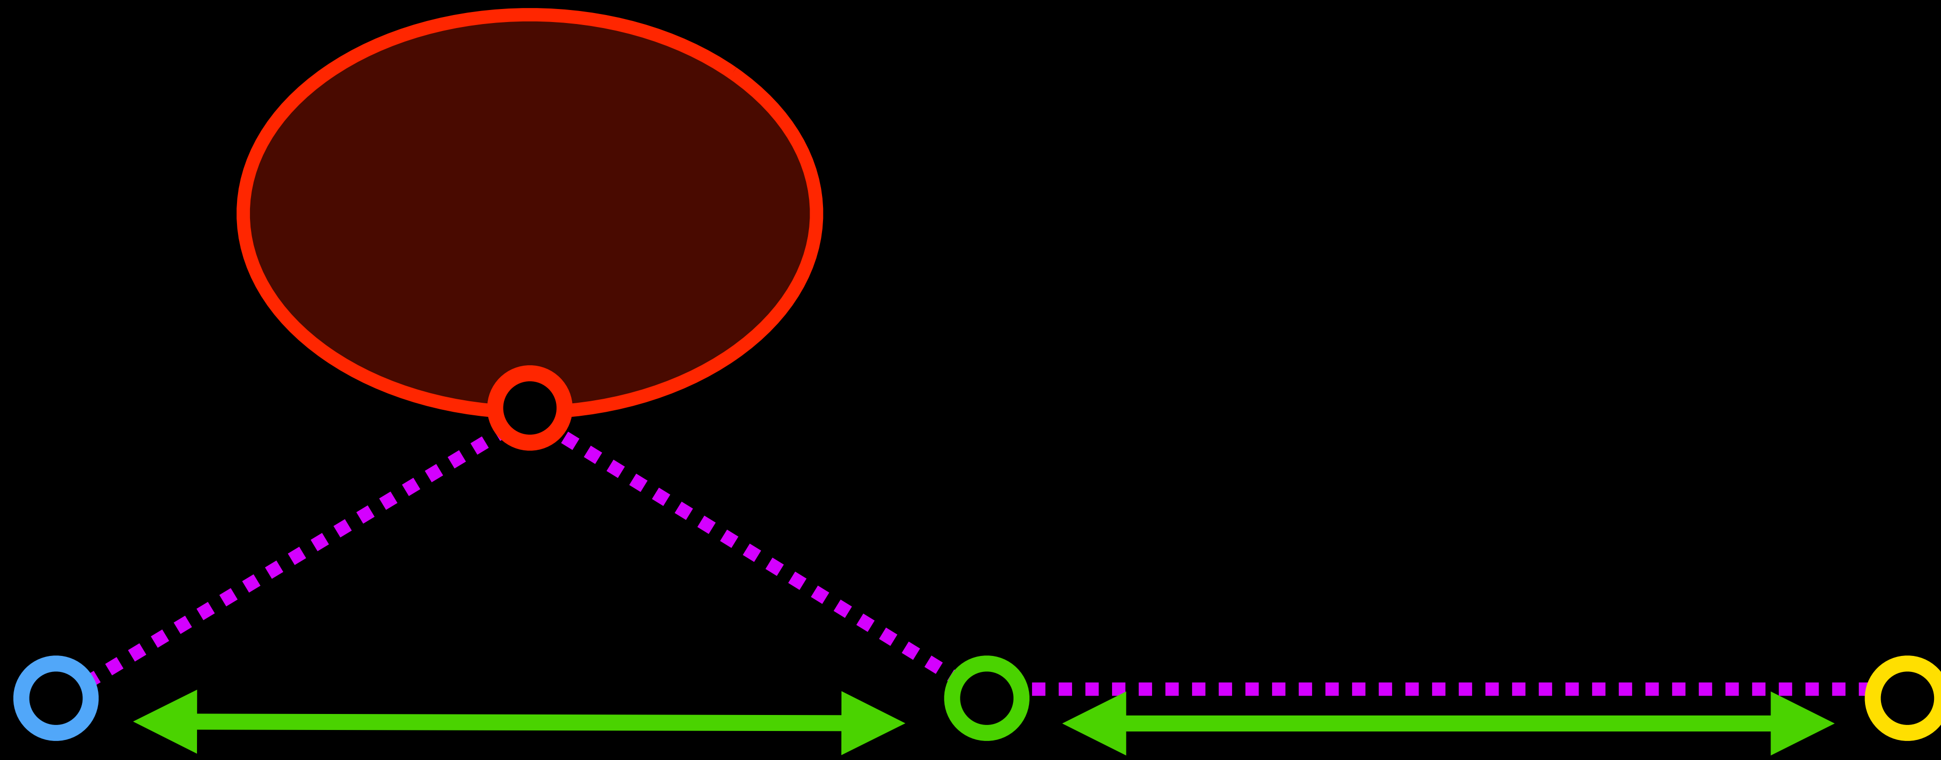

The shortest *possible* distance thru  and 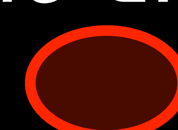 = d

Achieving provable avoidance

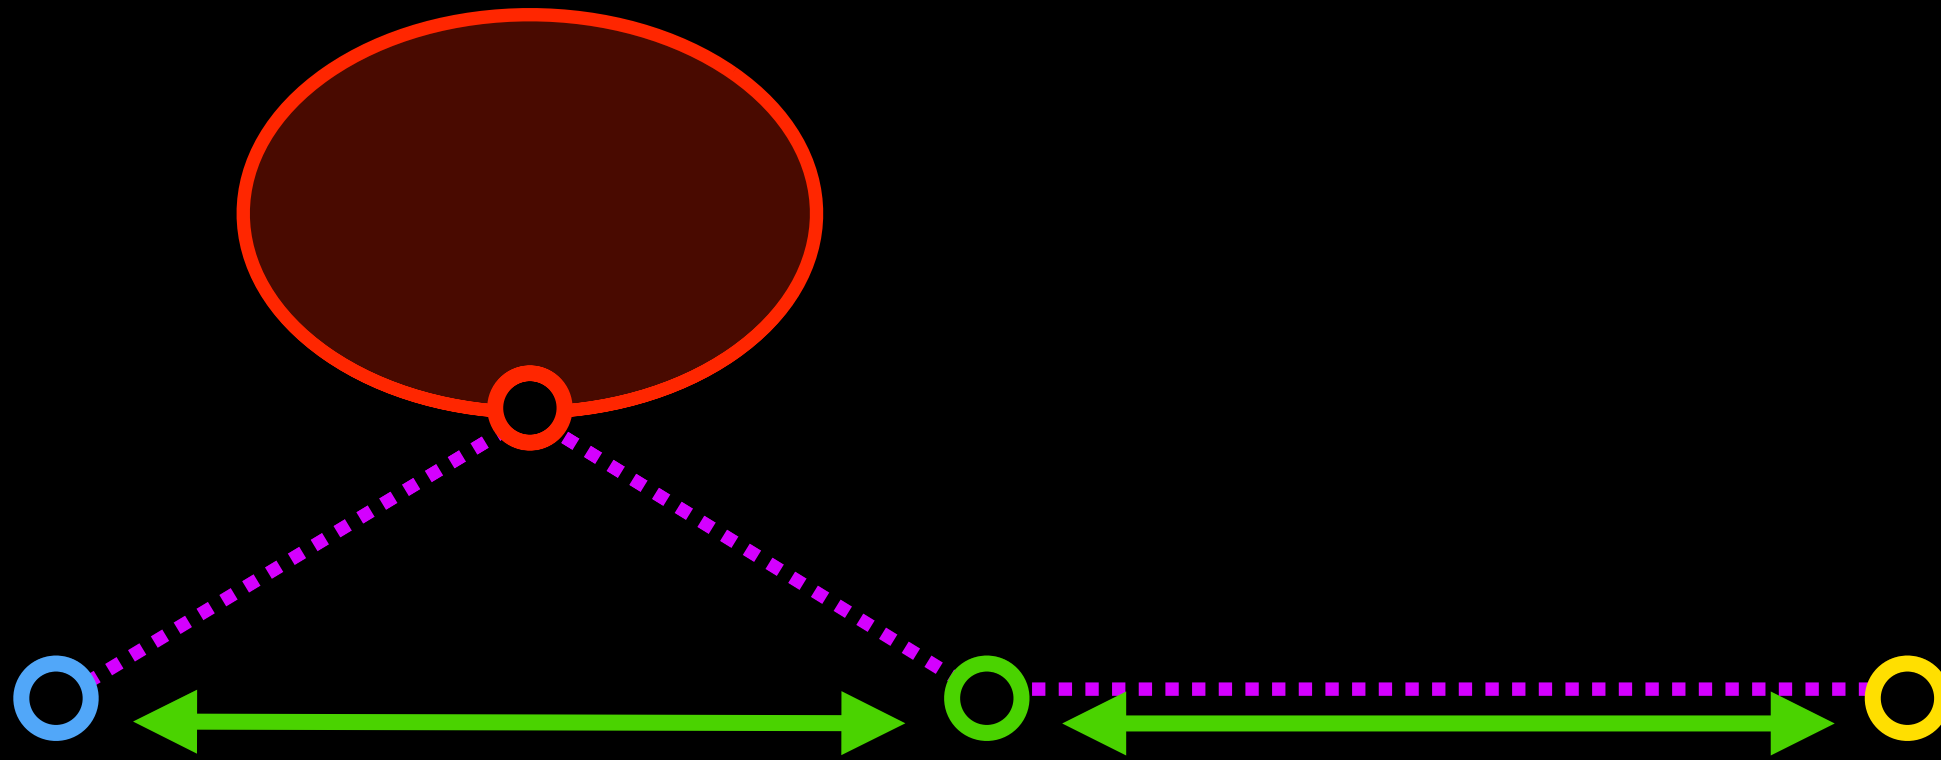

The shortest *possible* RTT
thru and

$= 2\,d\,/\,c$

Achieving provable avoidance

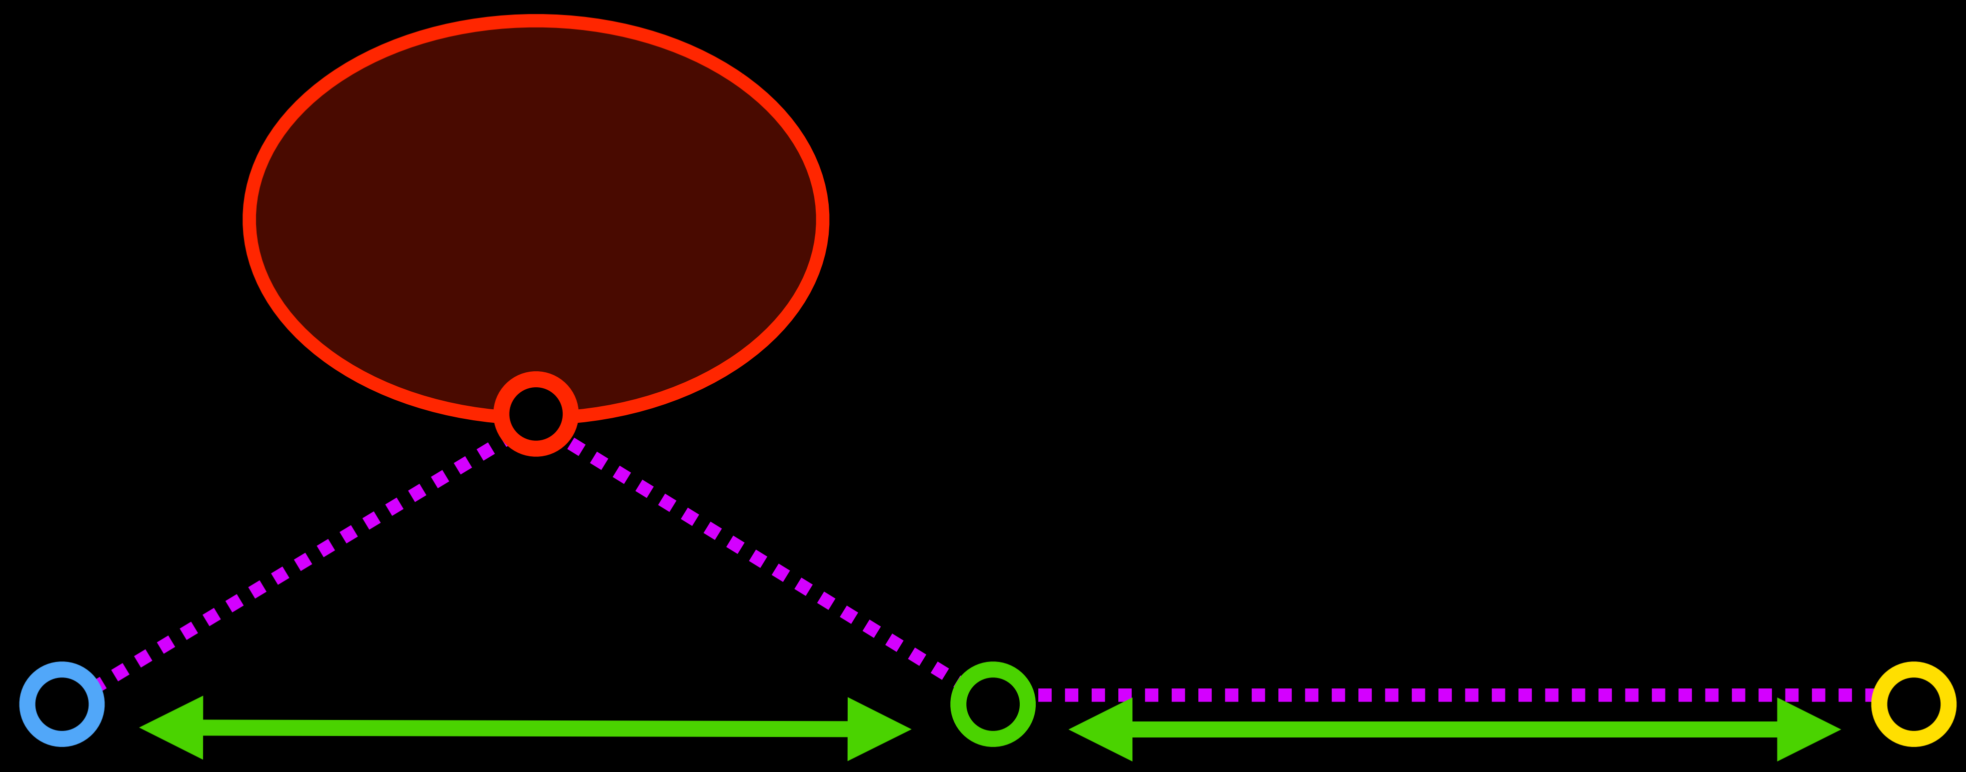

Measured RTT ≪ The shortest *possible* RTT thru and = 2 d / c

Achieving provable avoidance

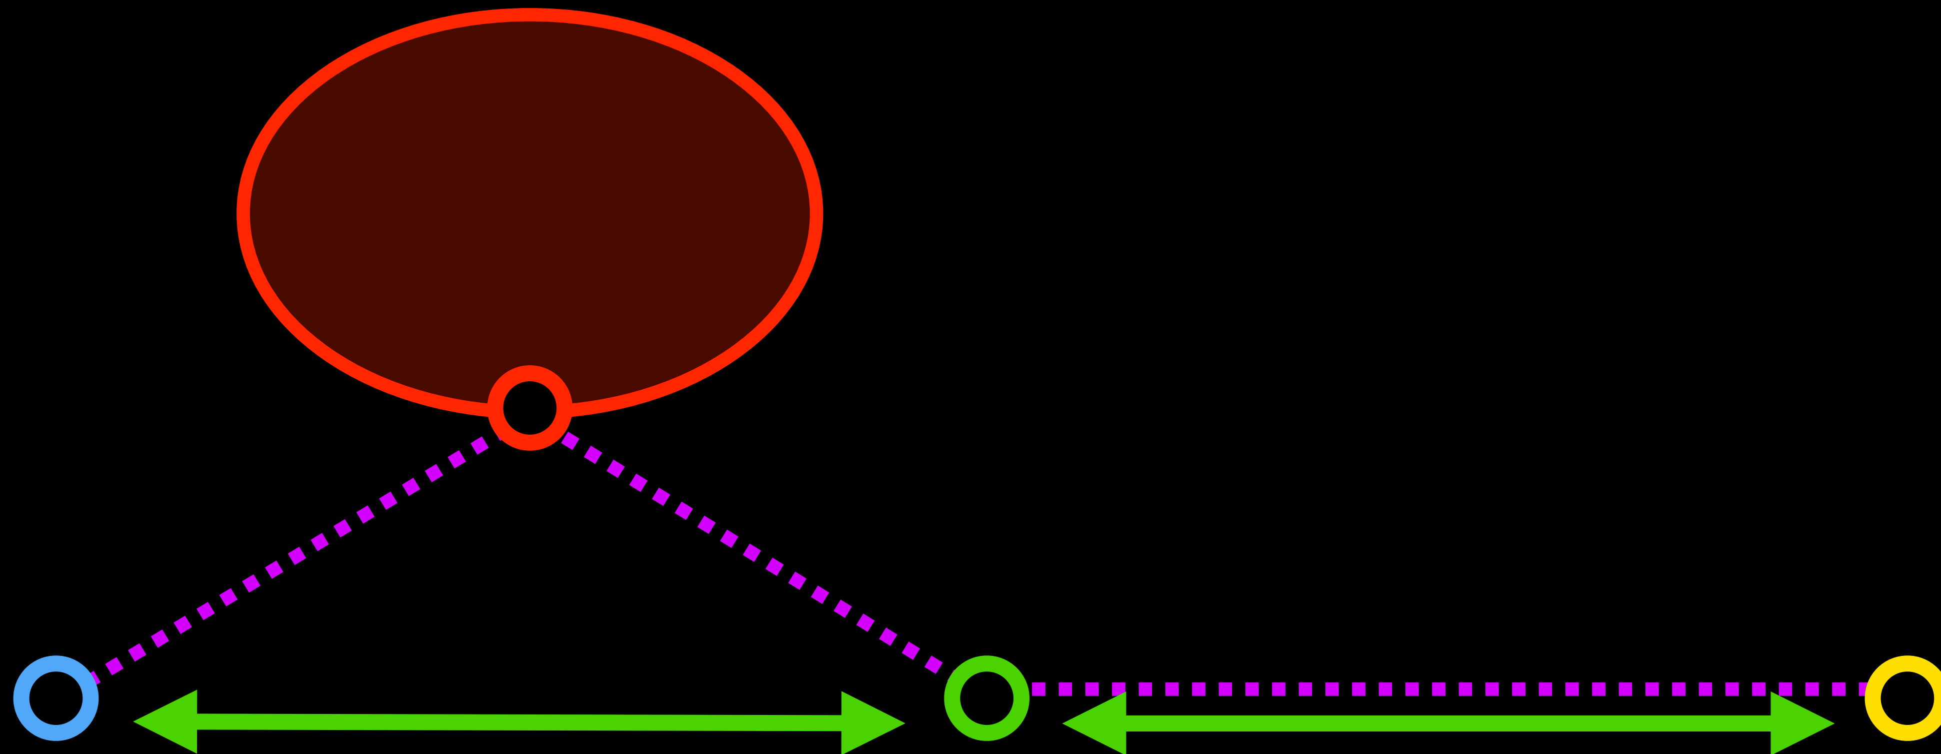

Measured RTT ≪ The shortest *possible* RTT thru and = 2 d / c

⇒ It could not have traversed *and*

Achieving provable avoidance

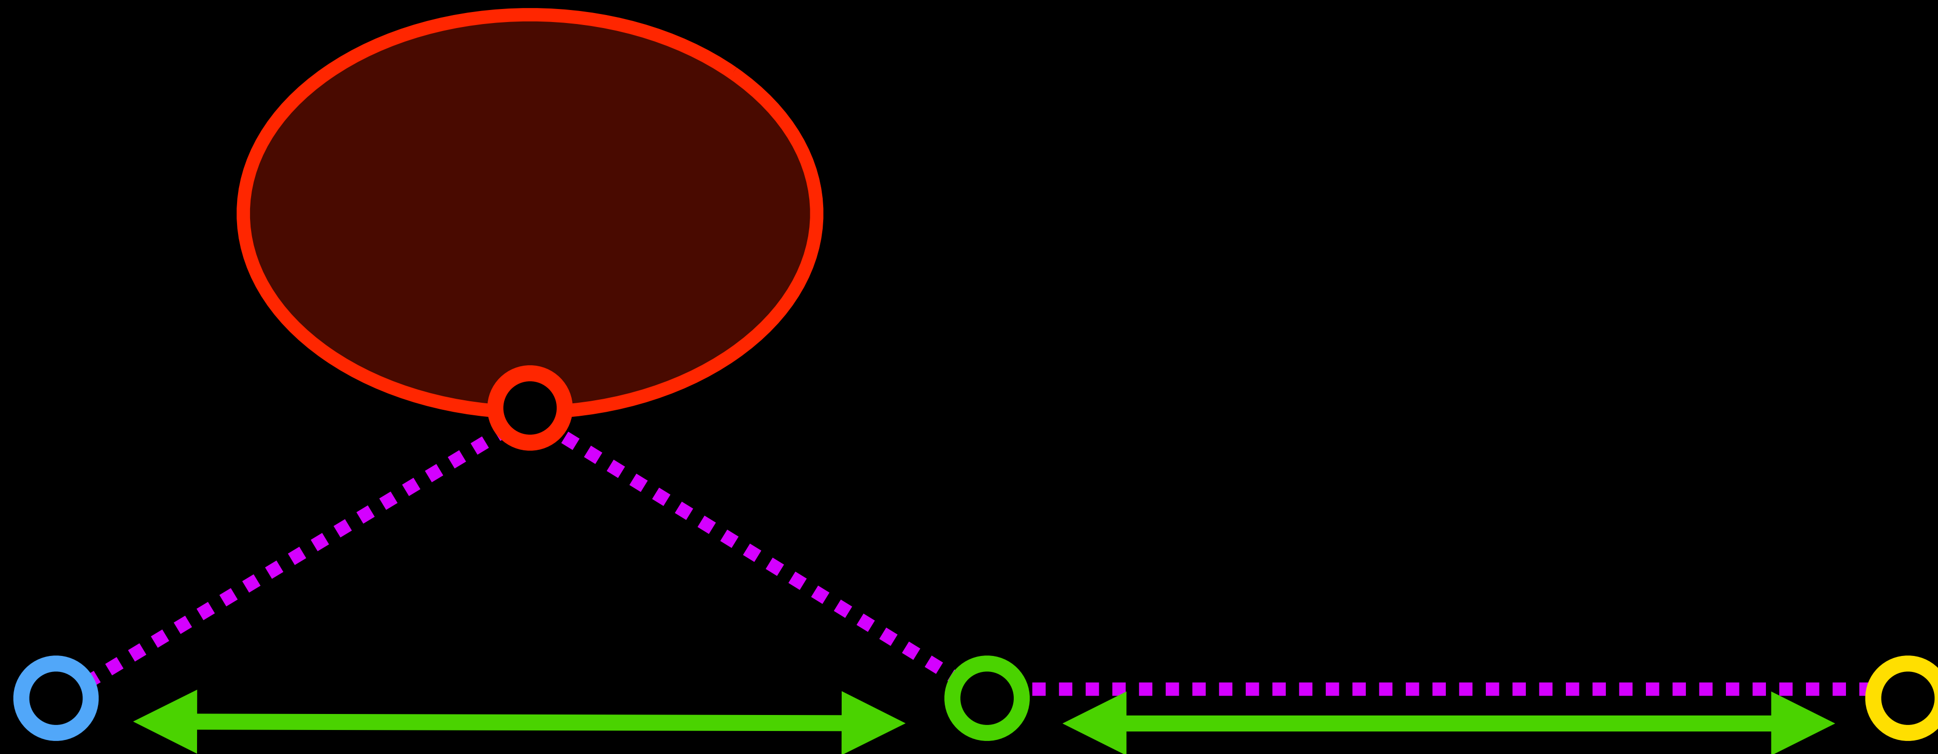

Measured RTT ≪ 2 d / c

⇒ It could not have traversed  *and* 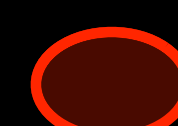

Achieving provable avoidance

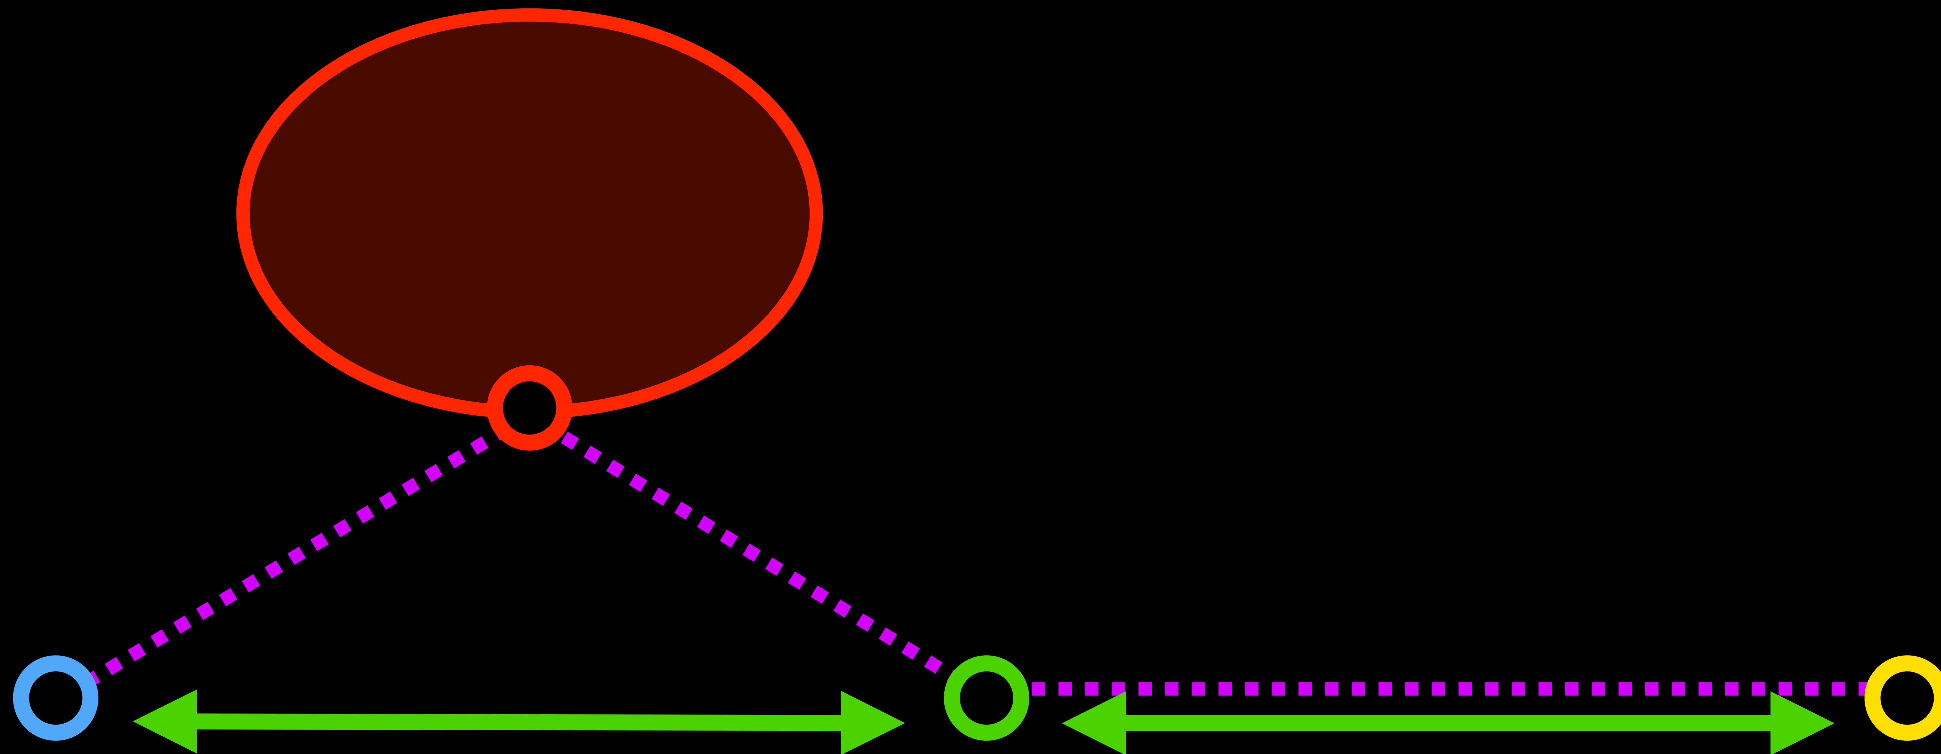

Measured RTT ≪ 3 d / c

⇒ It could not have traversed  *and* 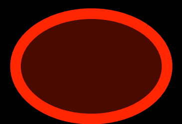

Achieving provable avoidance

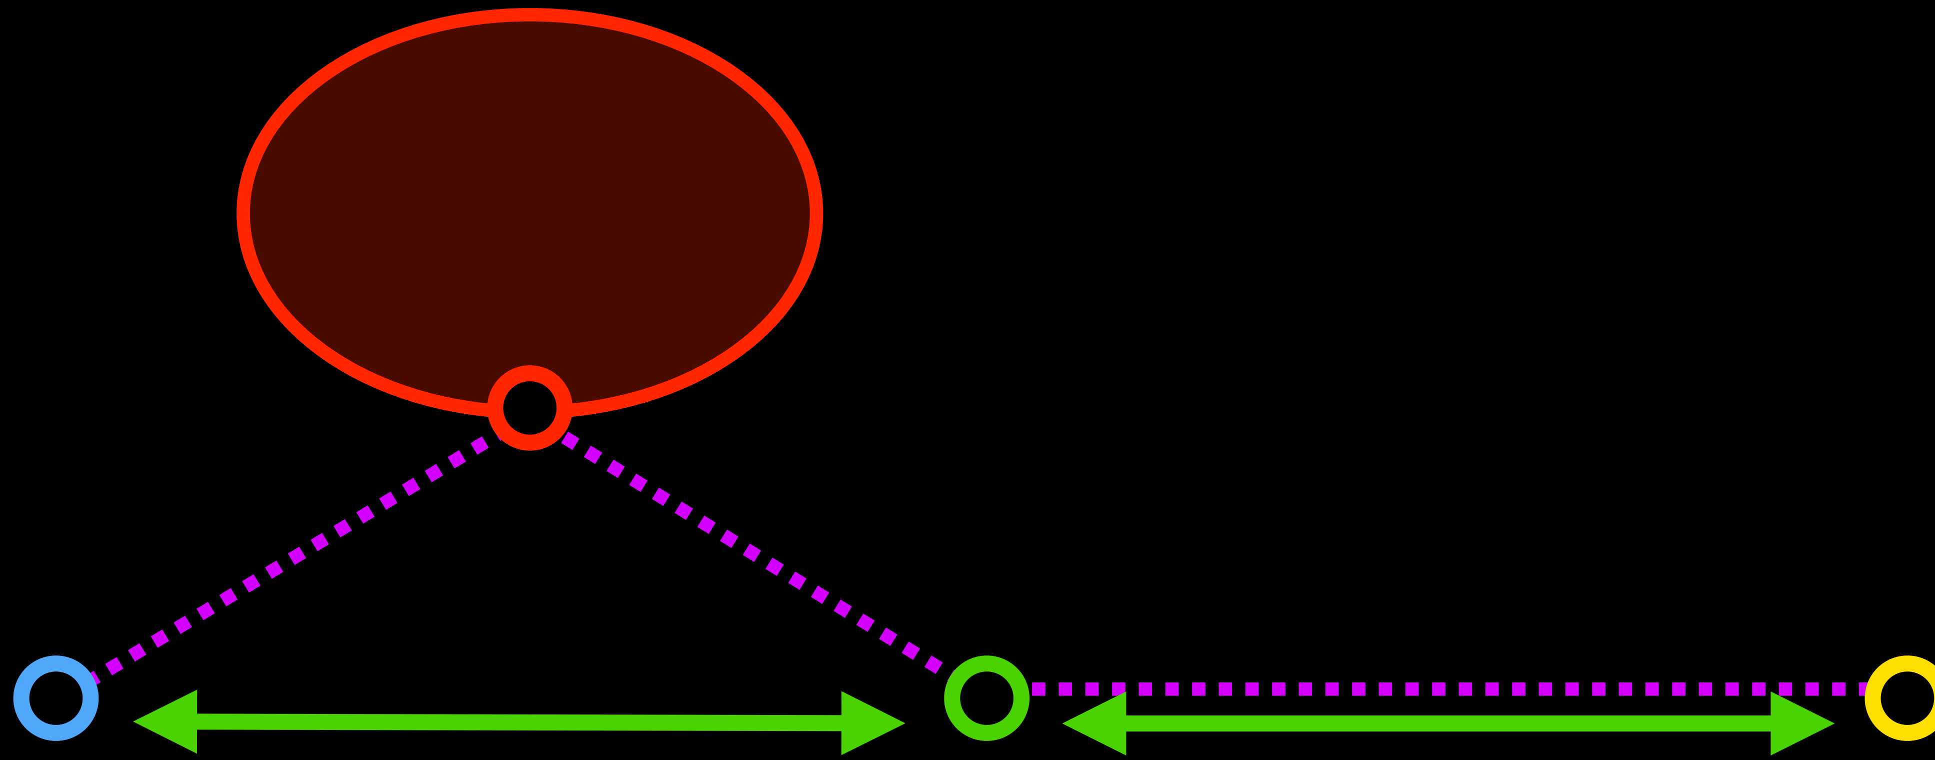

Safety factor

$$(1 + \underbrace{\delta}_{}) * \text{Measured RTT} \ll 3\,d\,/\,c$$

⇒ It could not have traversed *and*

Achieving provable avoidance

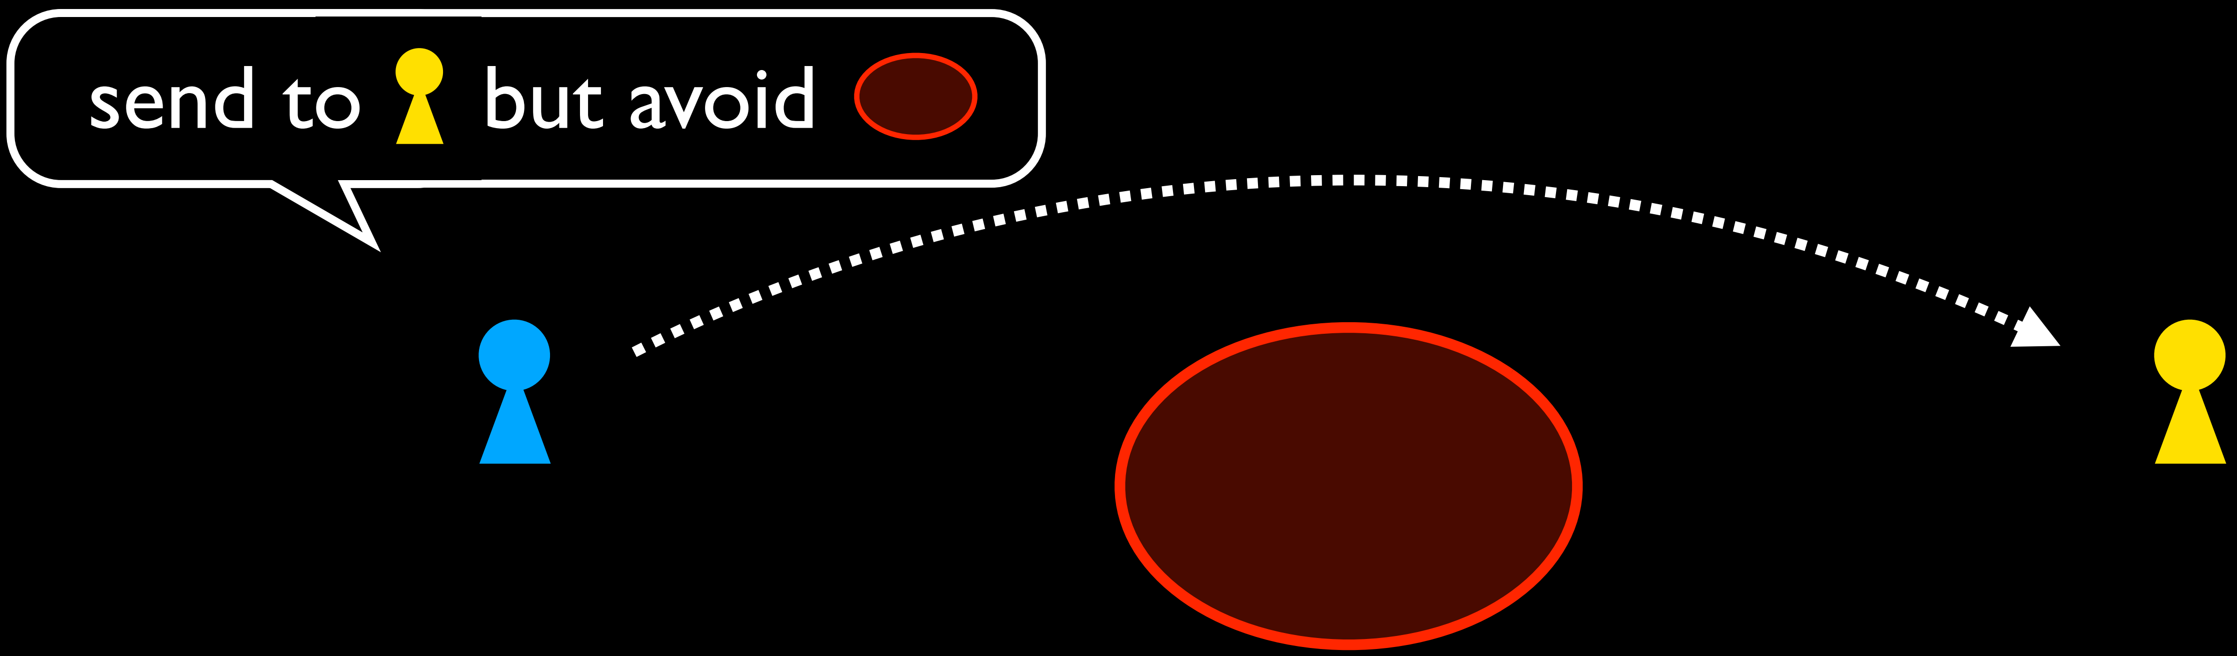

Achieving provable avoidance

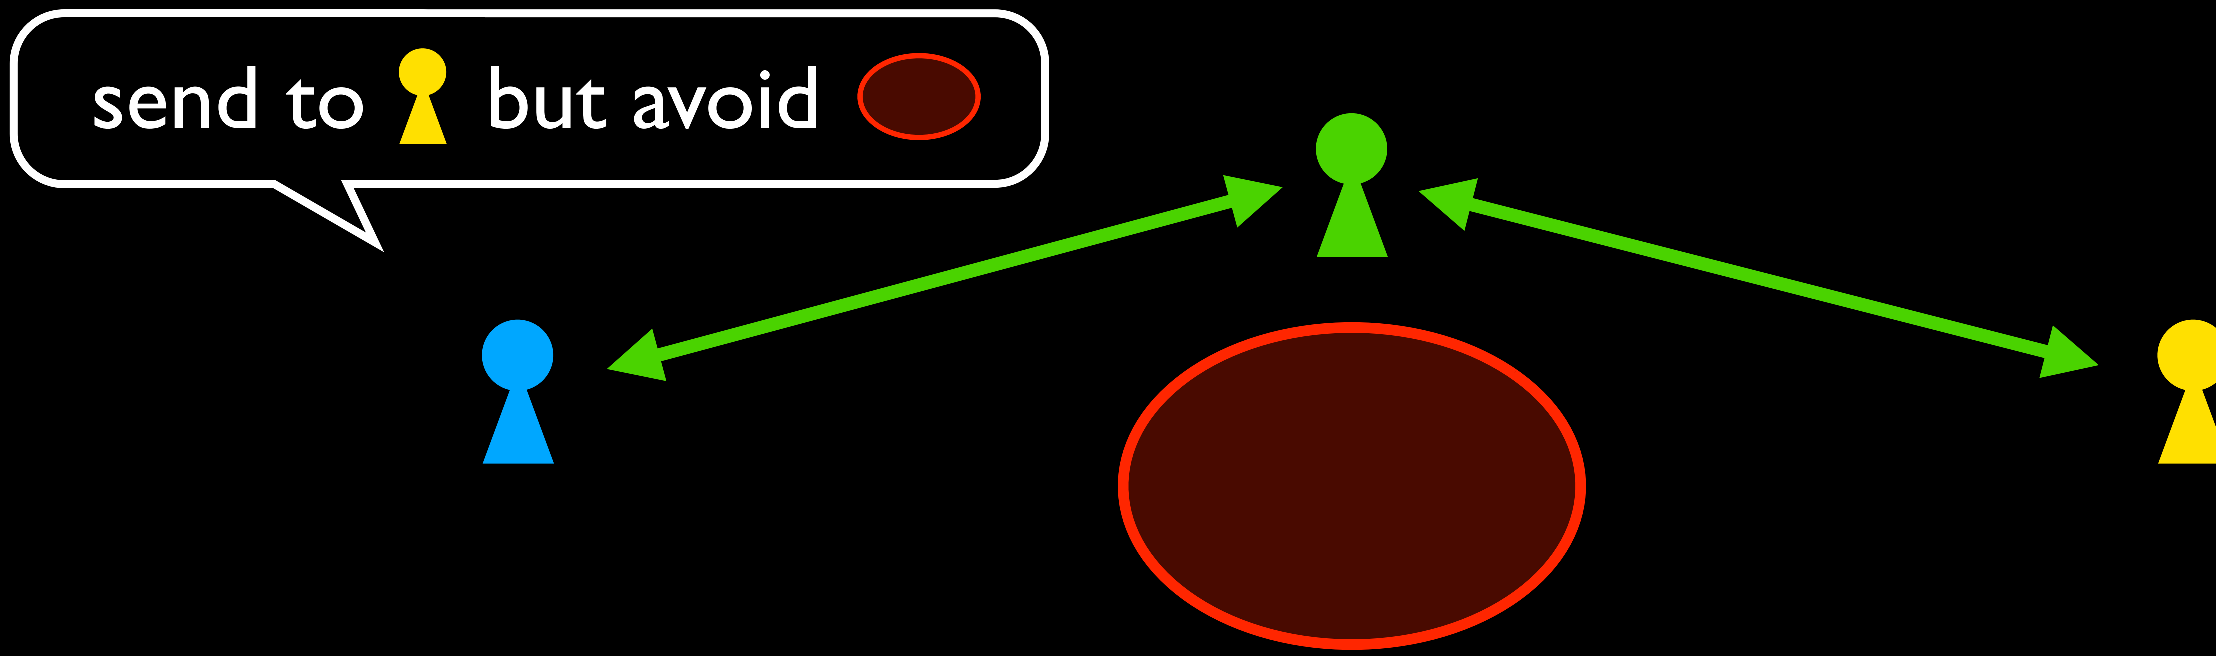

Achieving provable avoidance

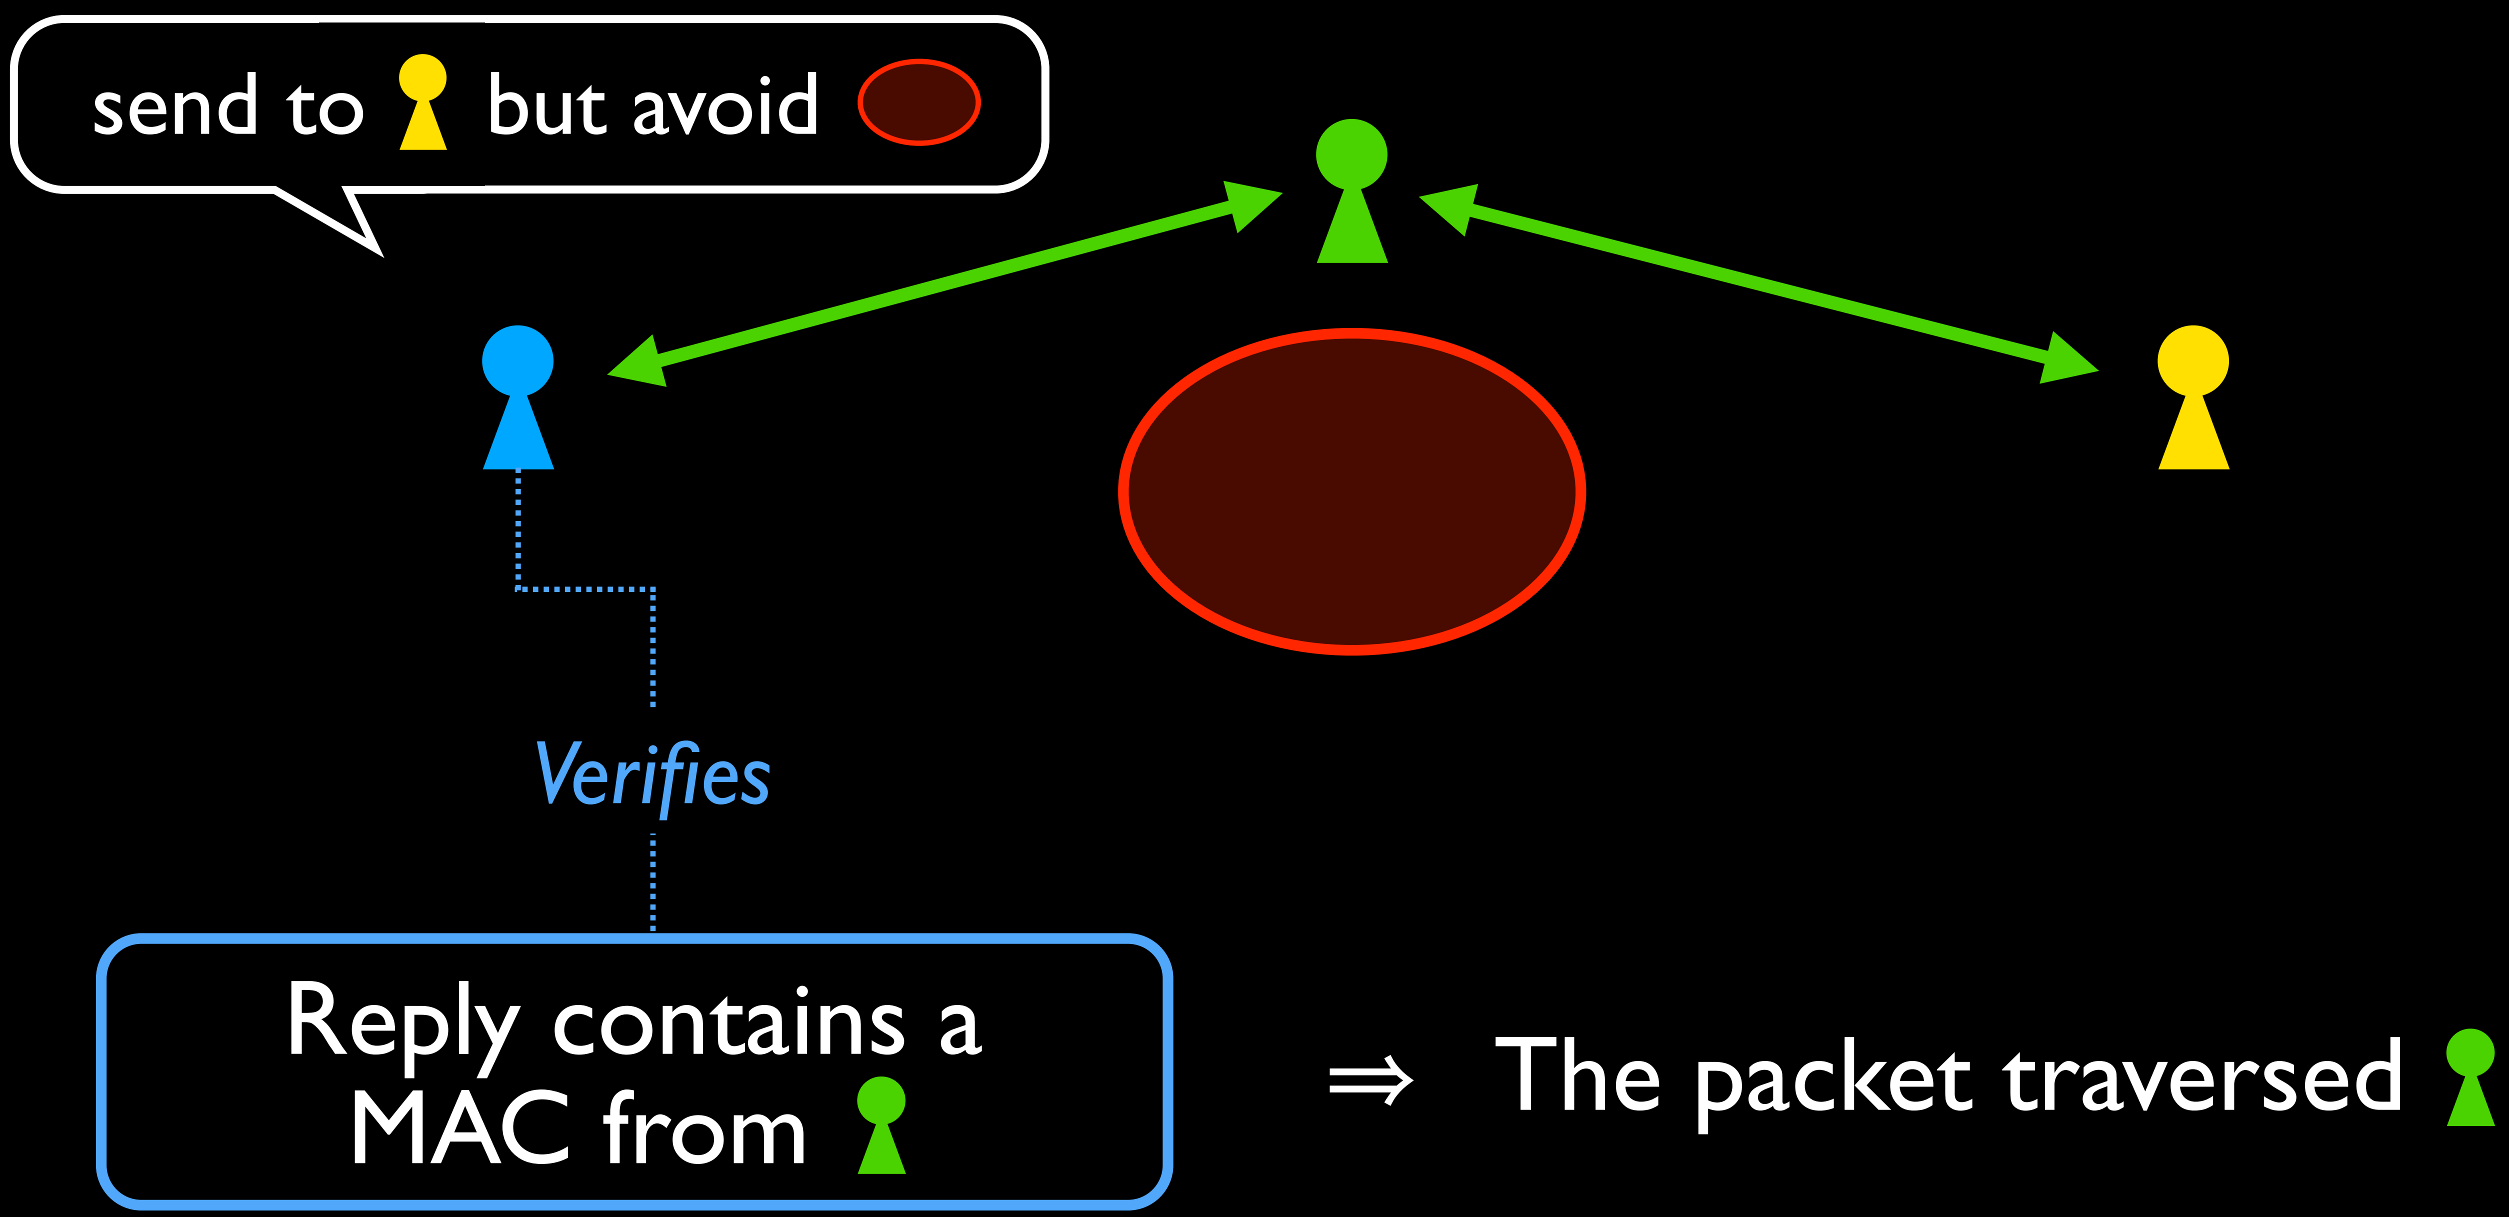

Achieving provable avoidance

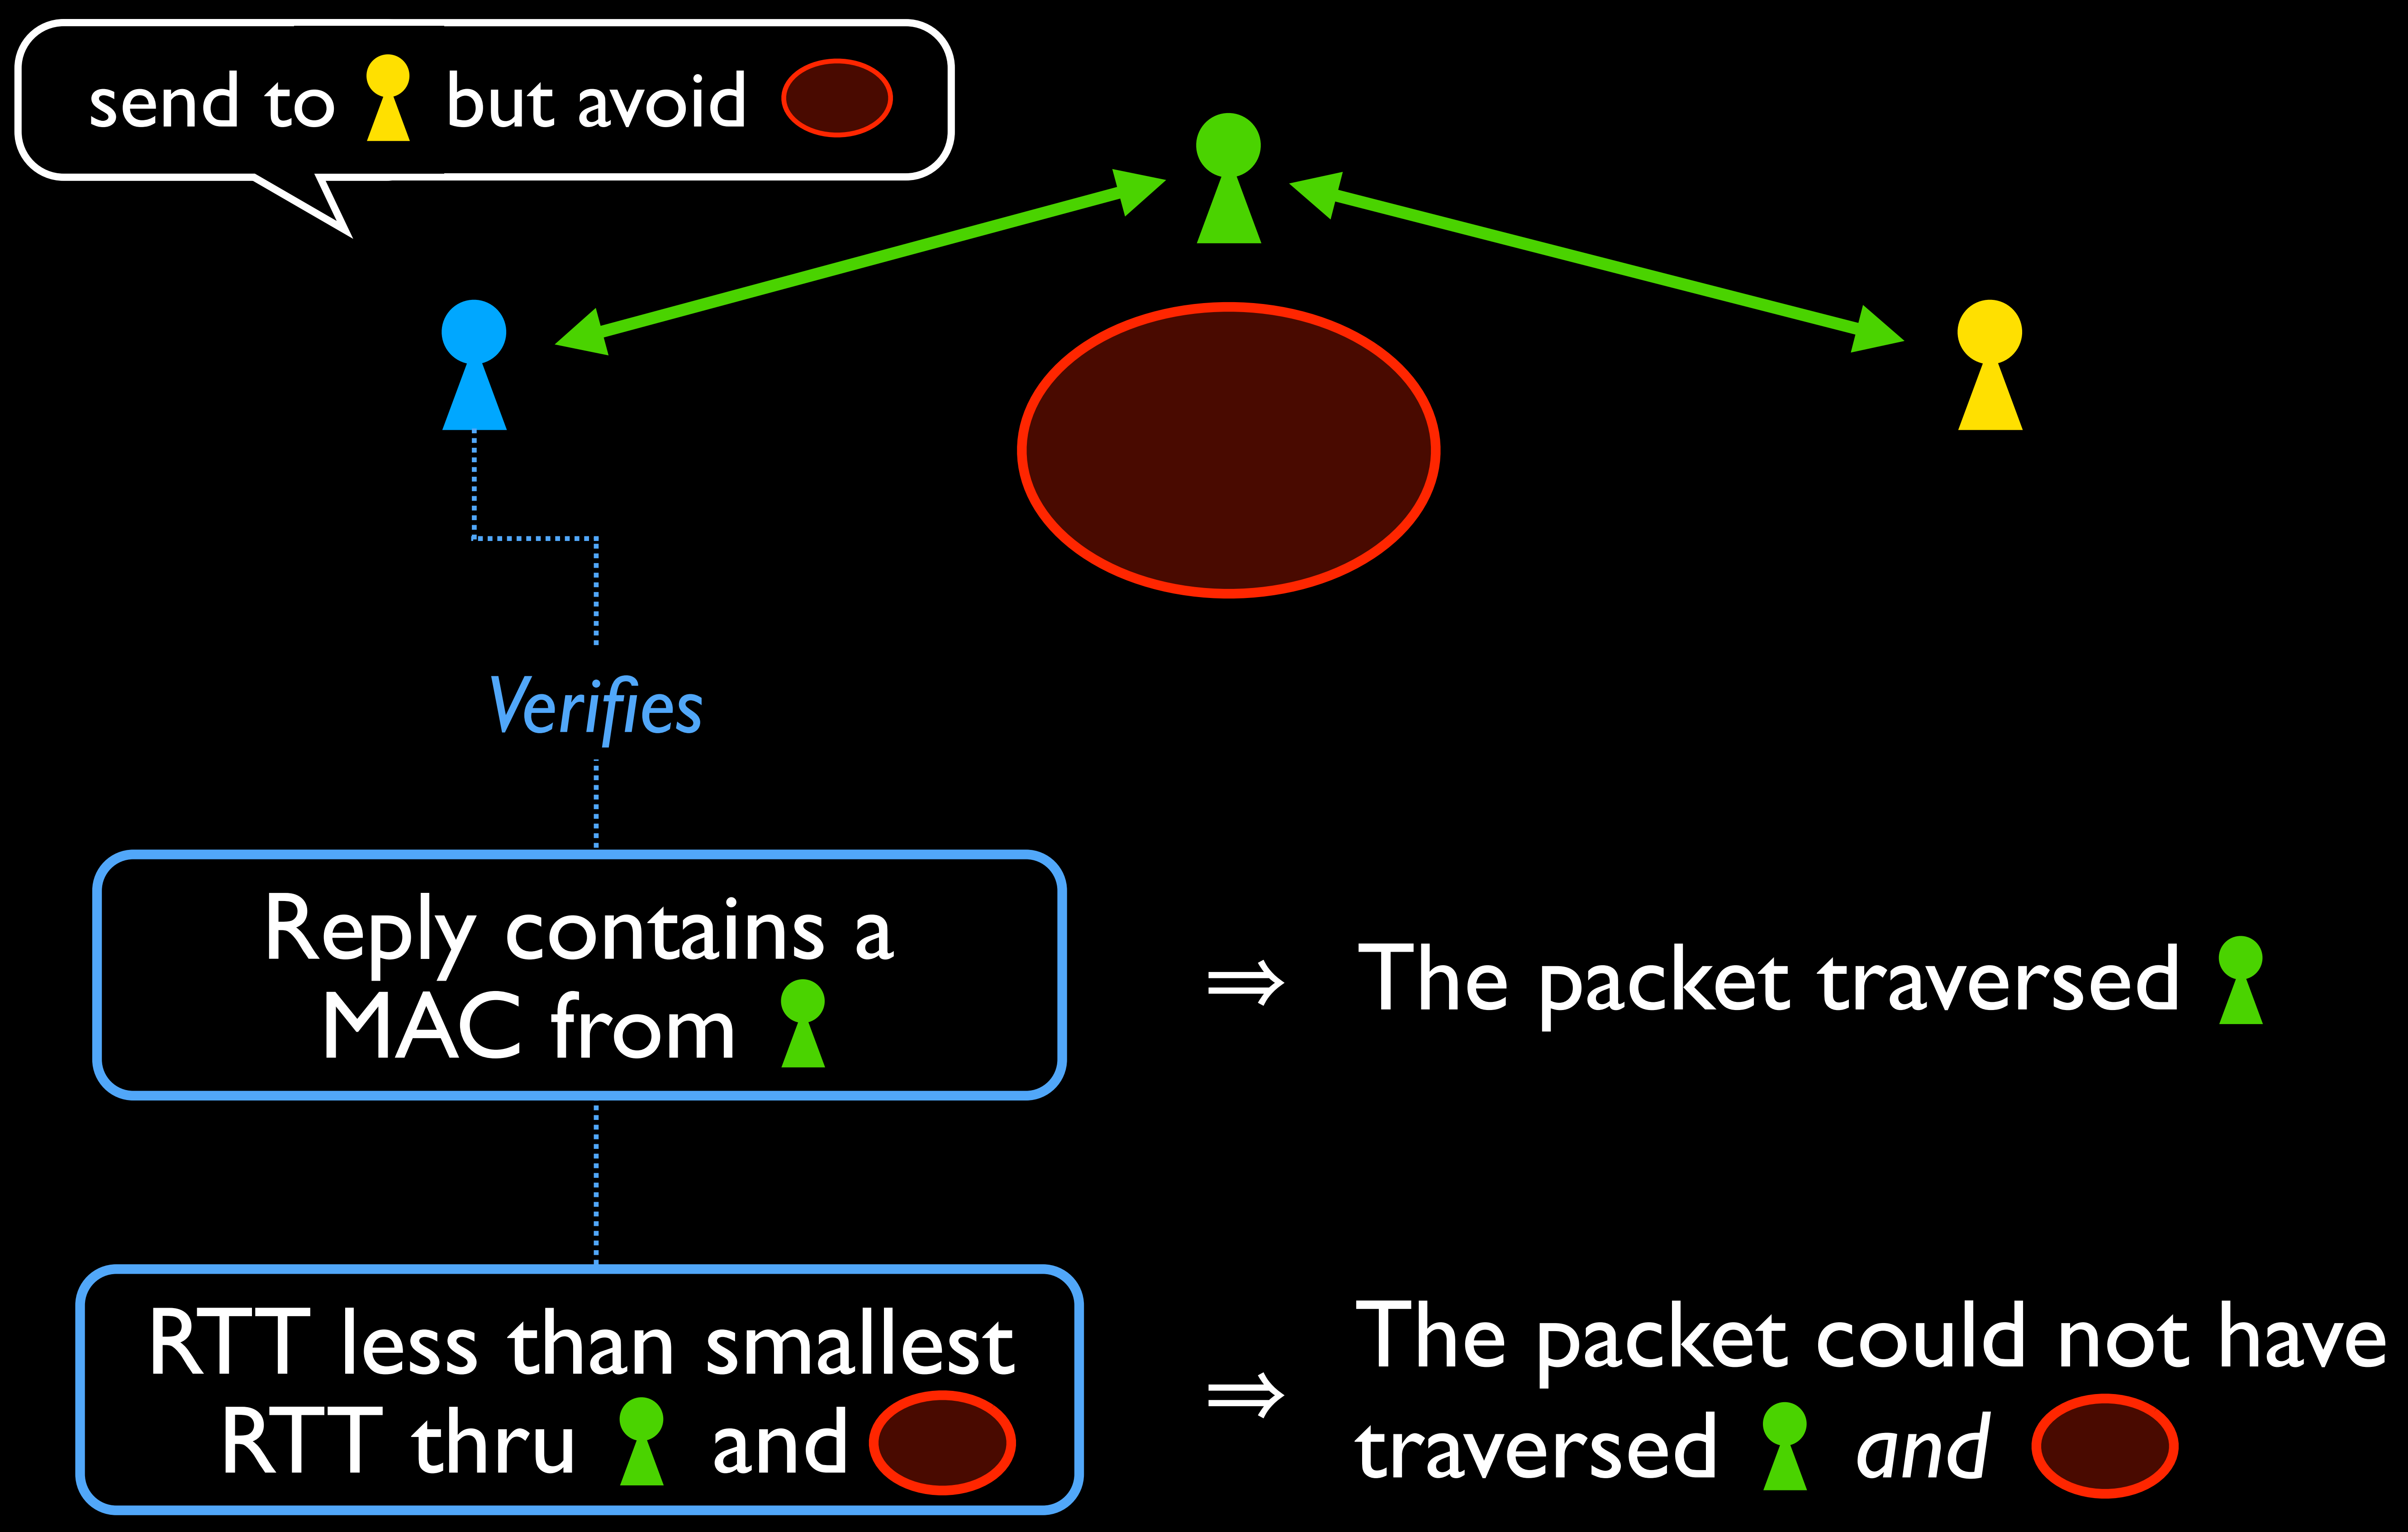

Achieving provable avoidance

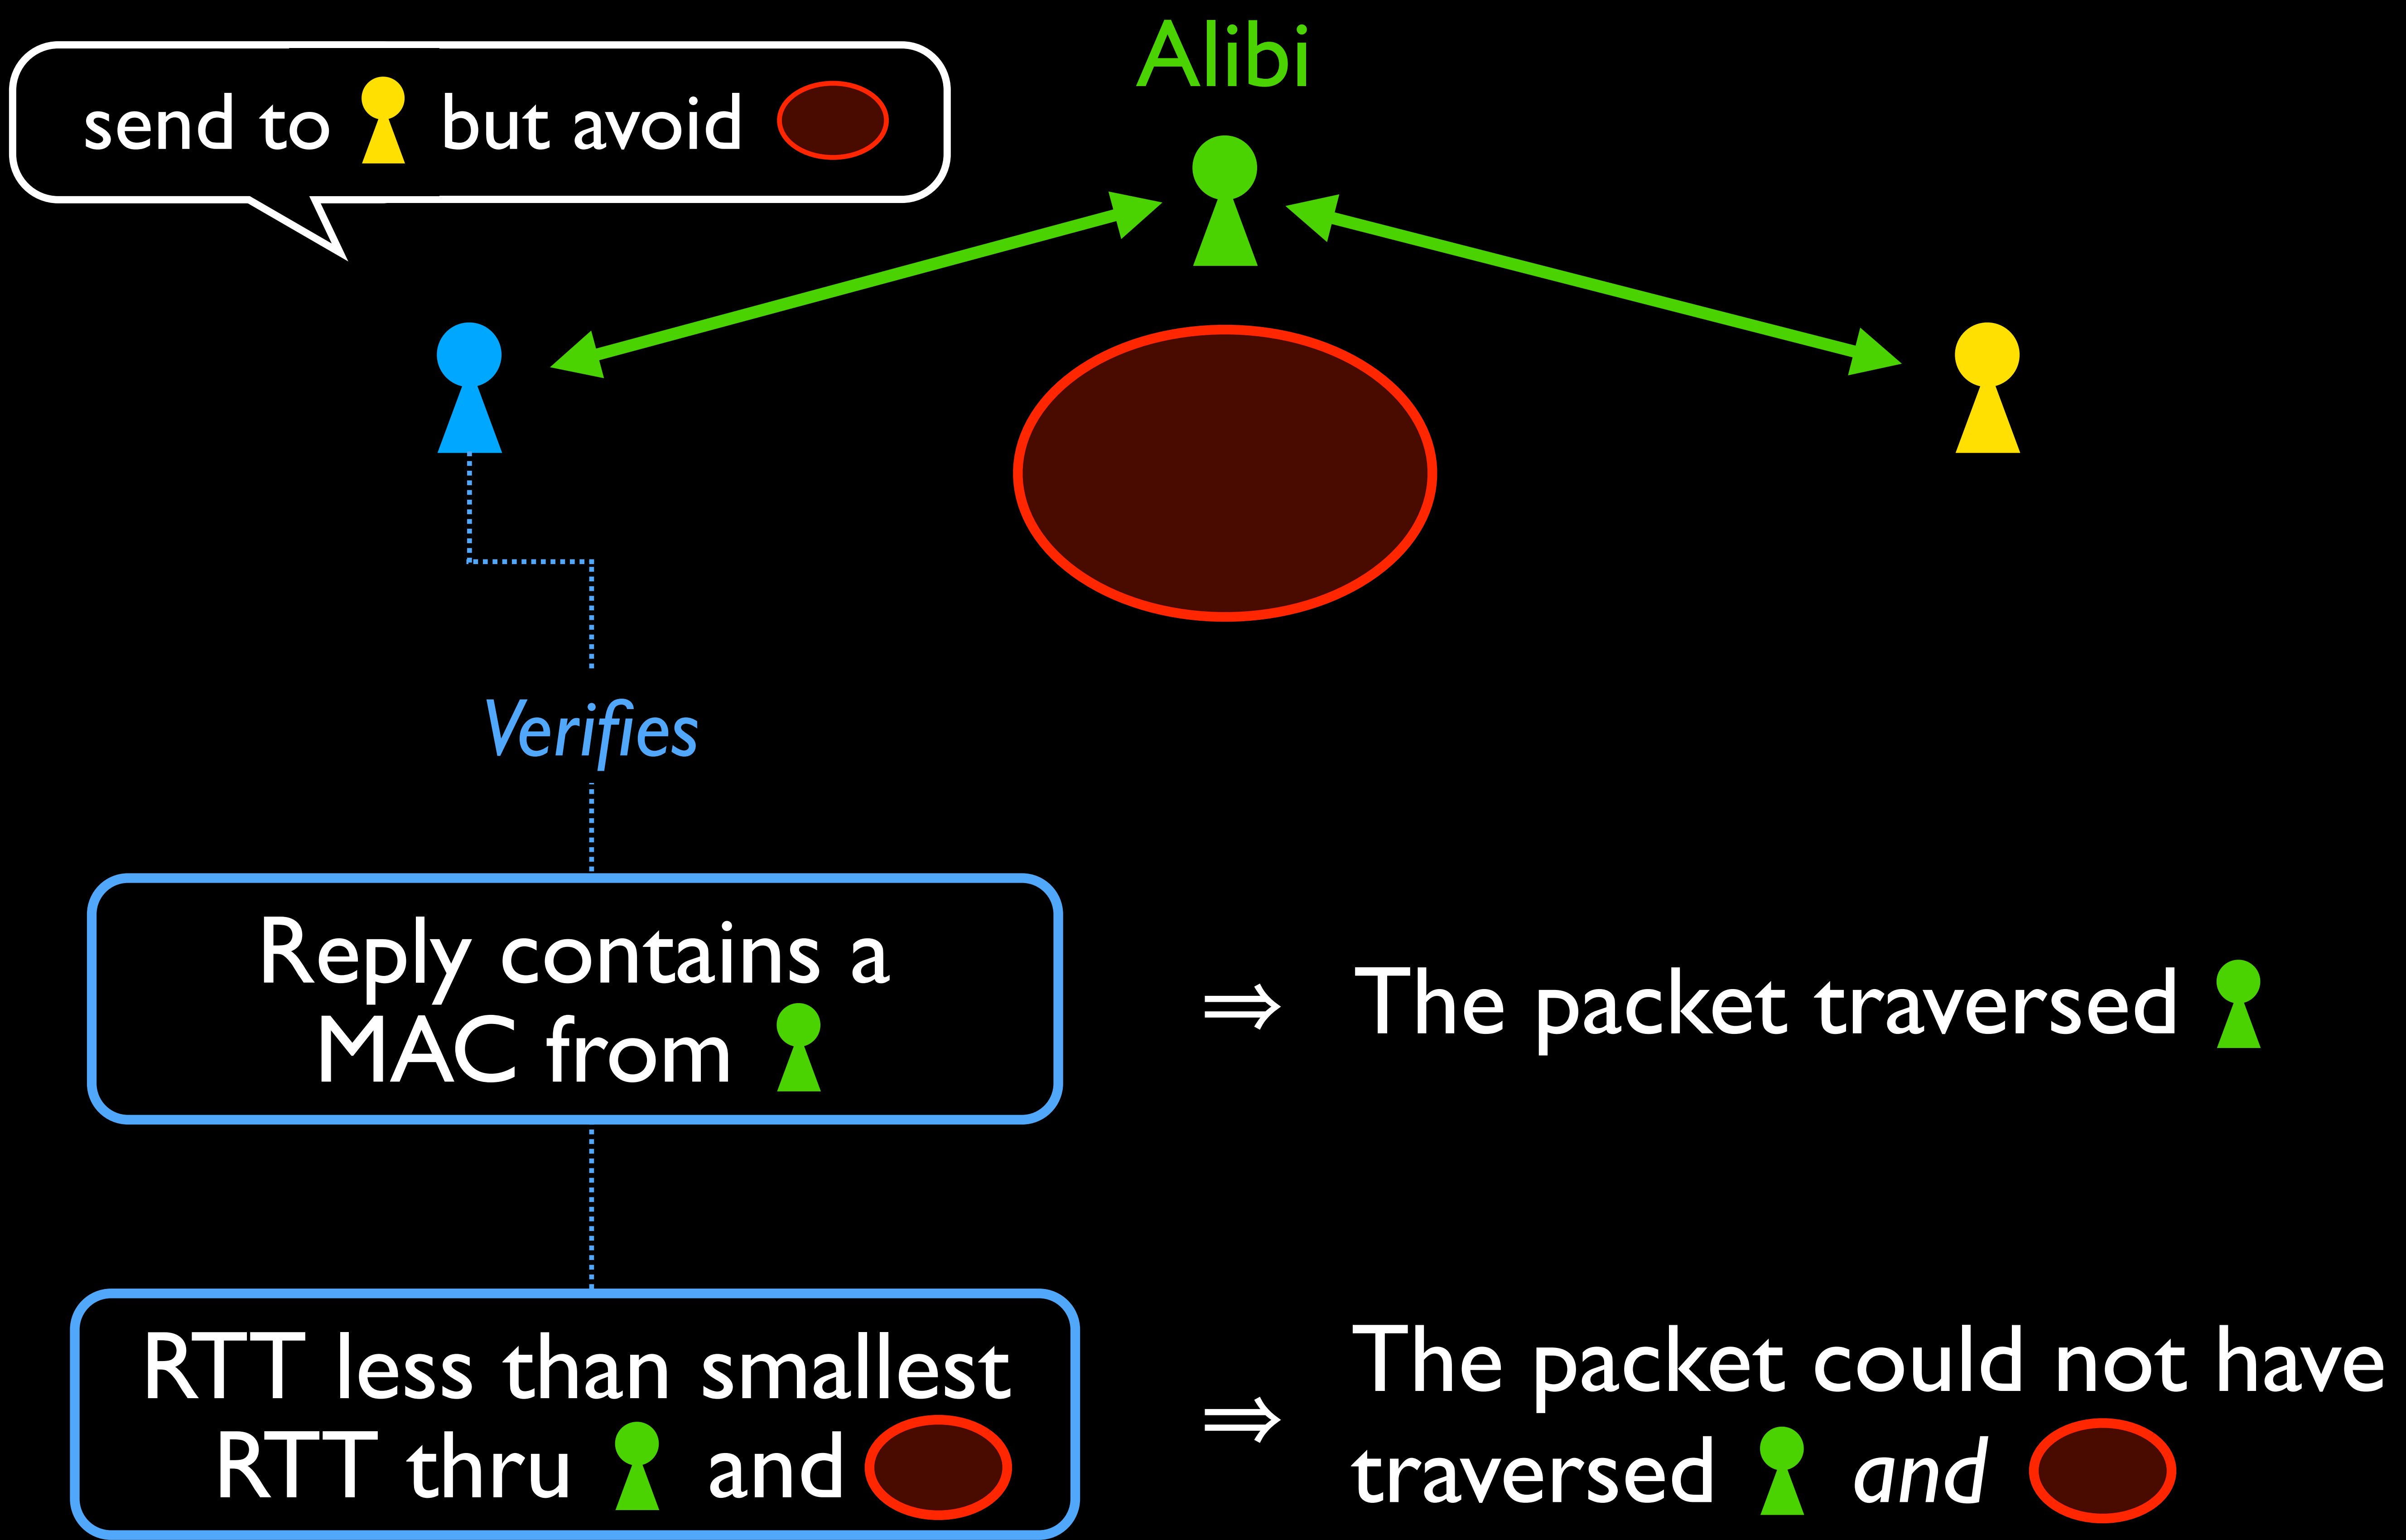

Alibi Routing

Peer-to-peer protocol for finding potential alibis

- Users choose forbidden regions
- Users compute target regions
  - Where alibis *might* be
- Alibi Routing recursively searches for peers within the target regions

Choose forbidden regions

User-specified regions to avoid

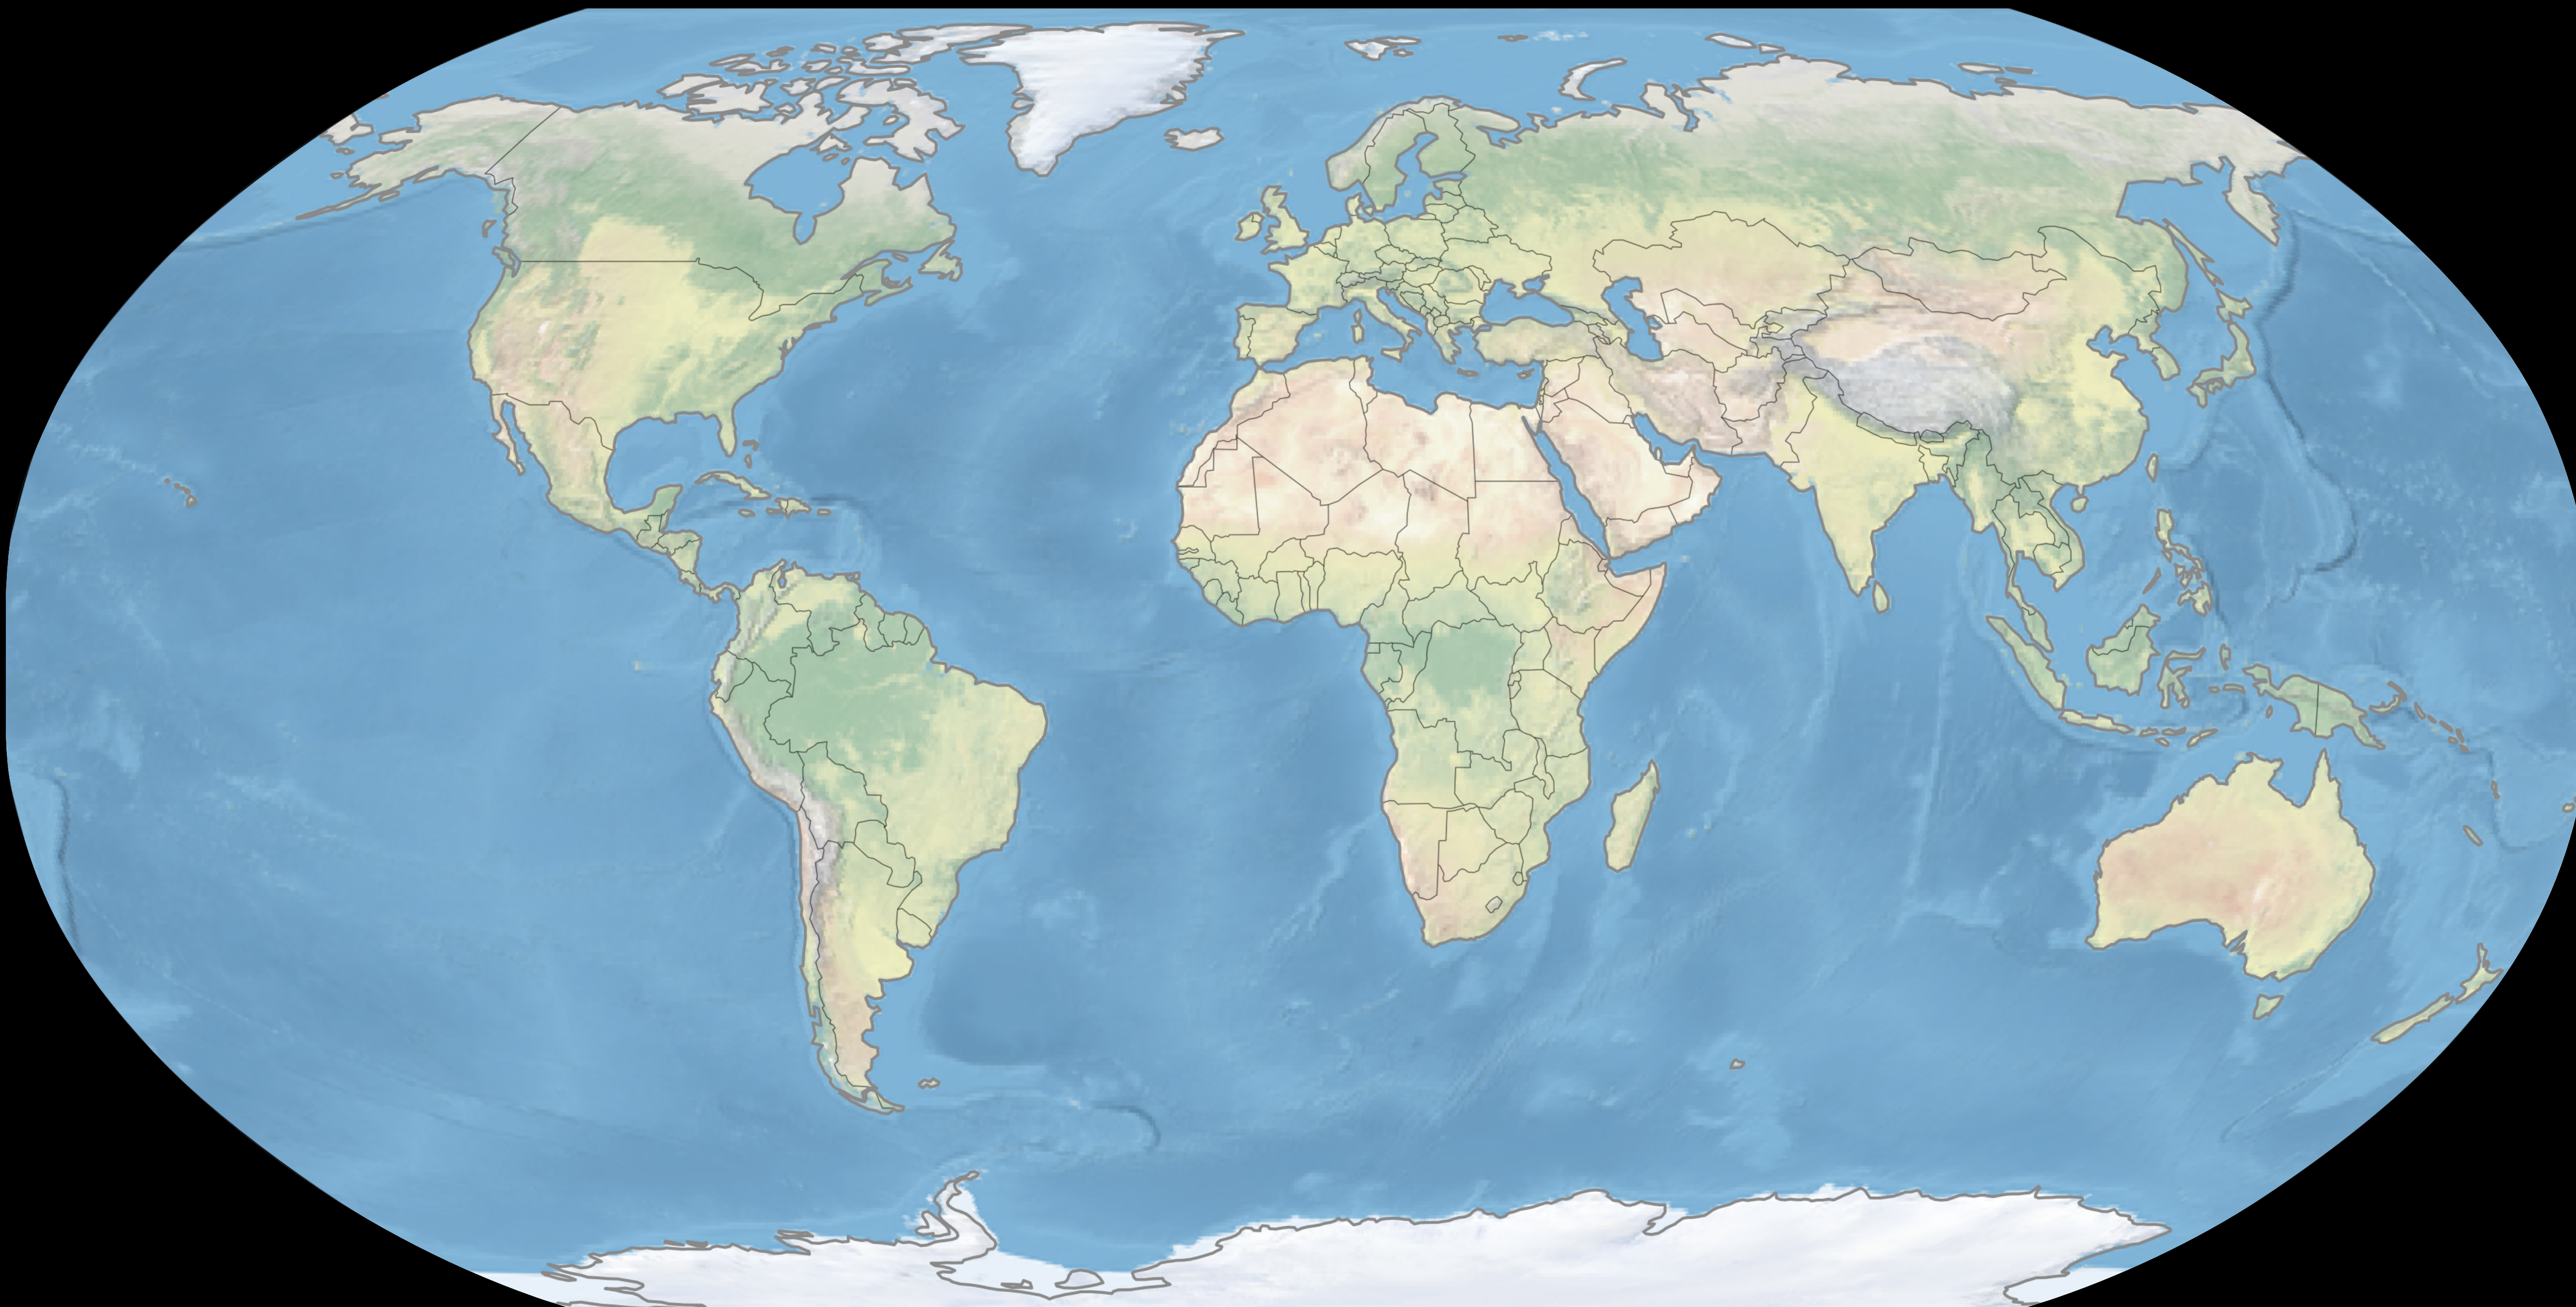

Choose forbidden regions

User-specified regions to avoid

Arbitrary sets of polygons, defined over lat/lon

Choose forbidden regions

User-specified regions to avoid

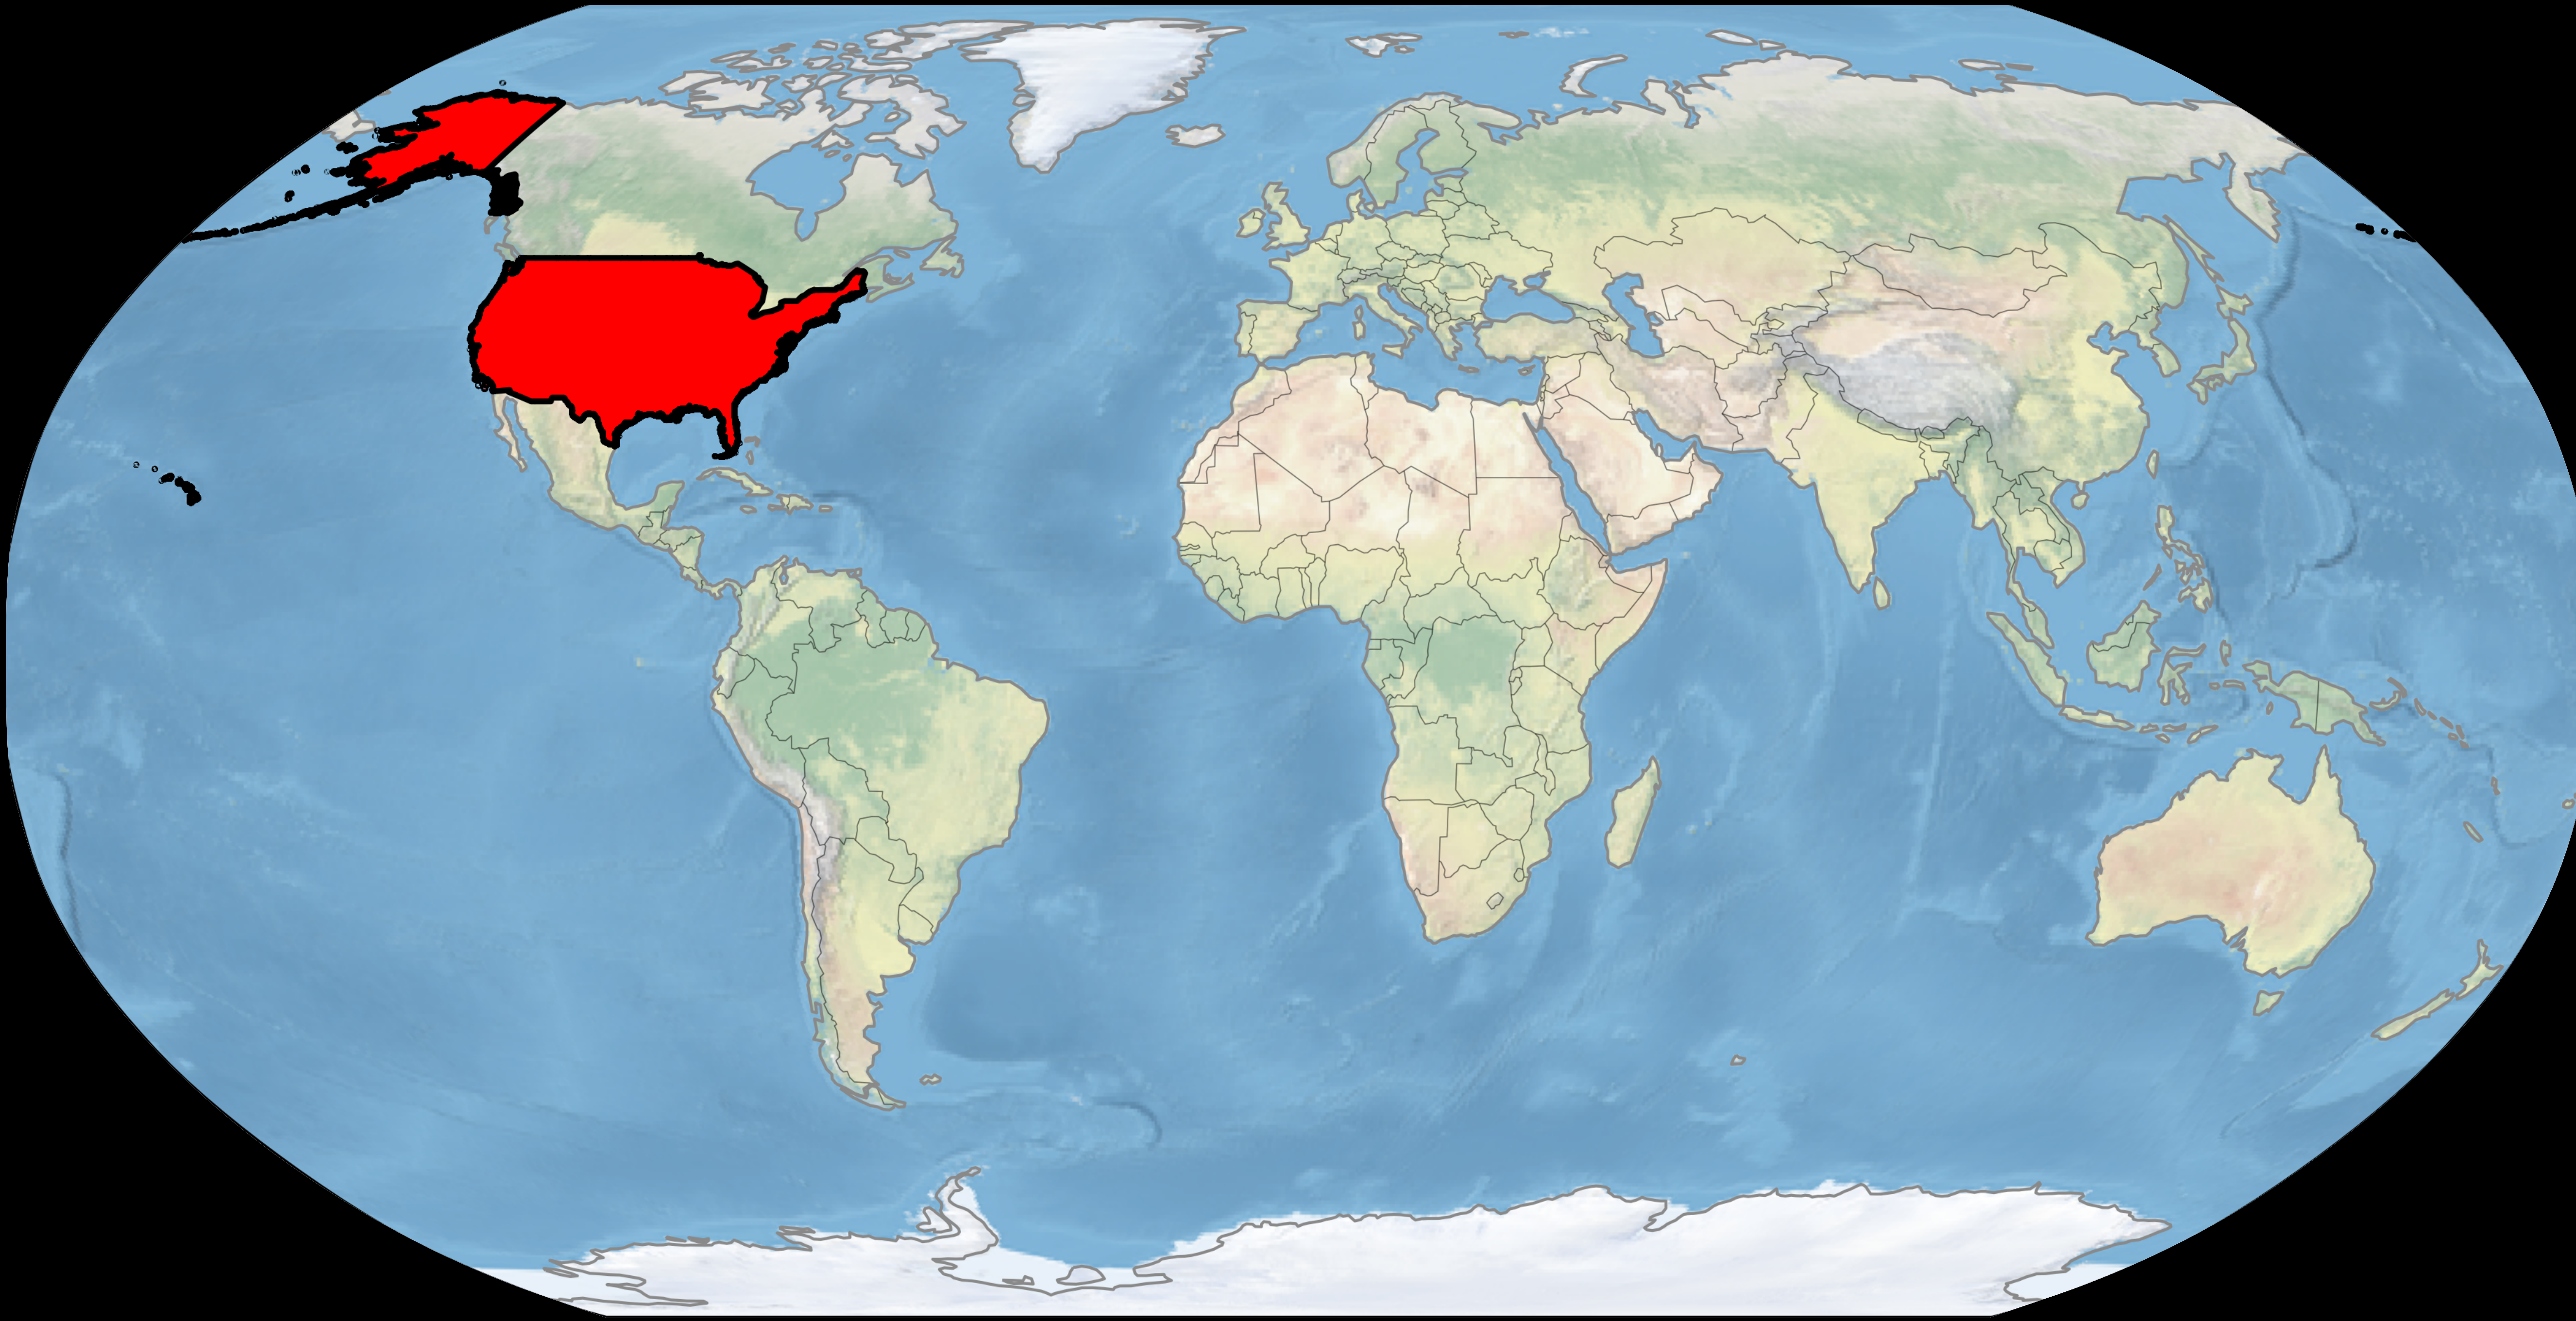

Arbitrary sets of polygons, defined over lat/lon

Choose forbidden regions

User-specified regions to avoid

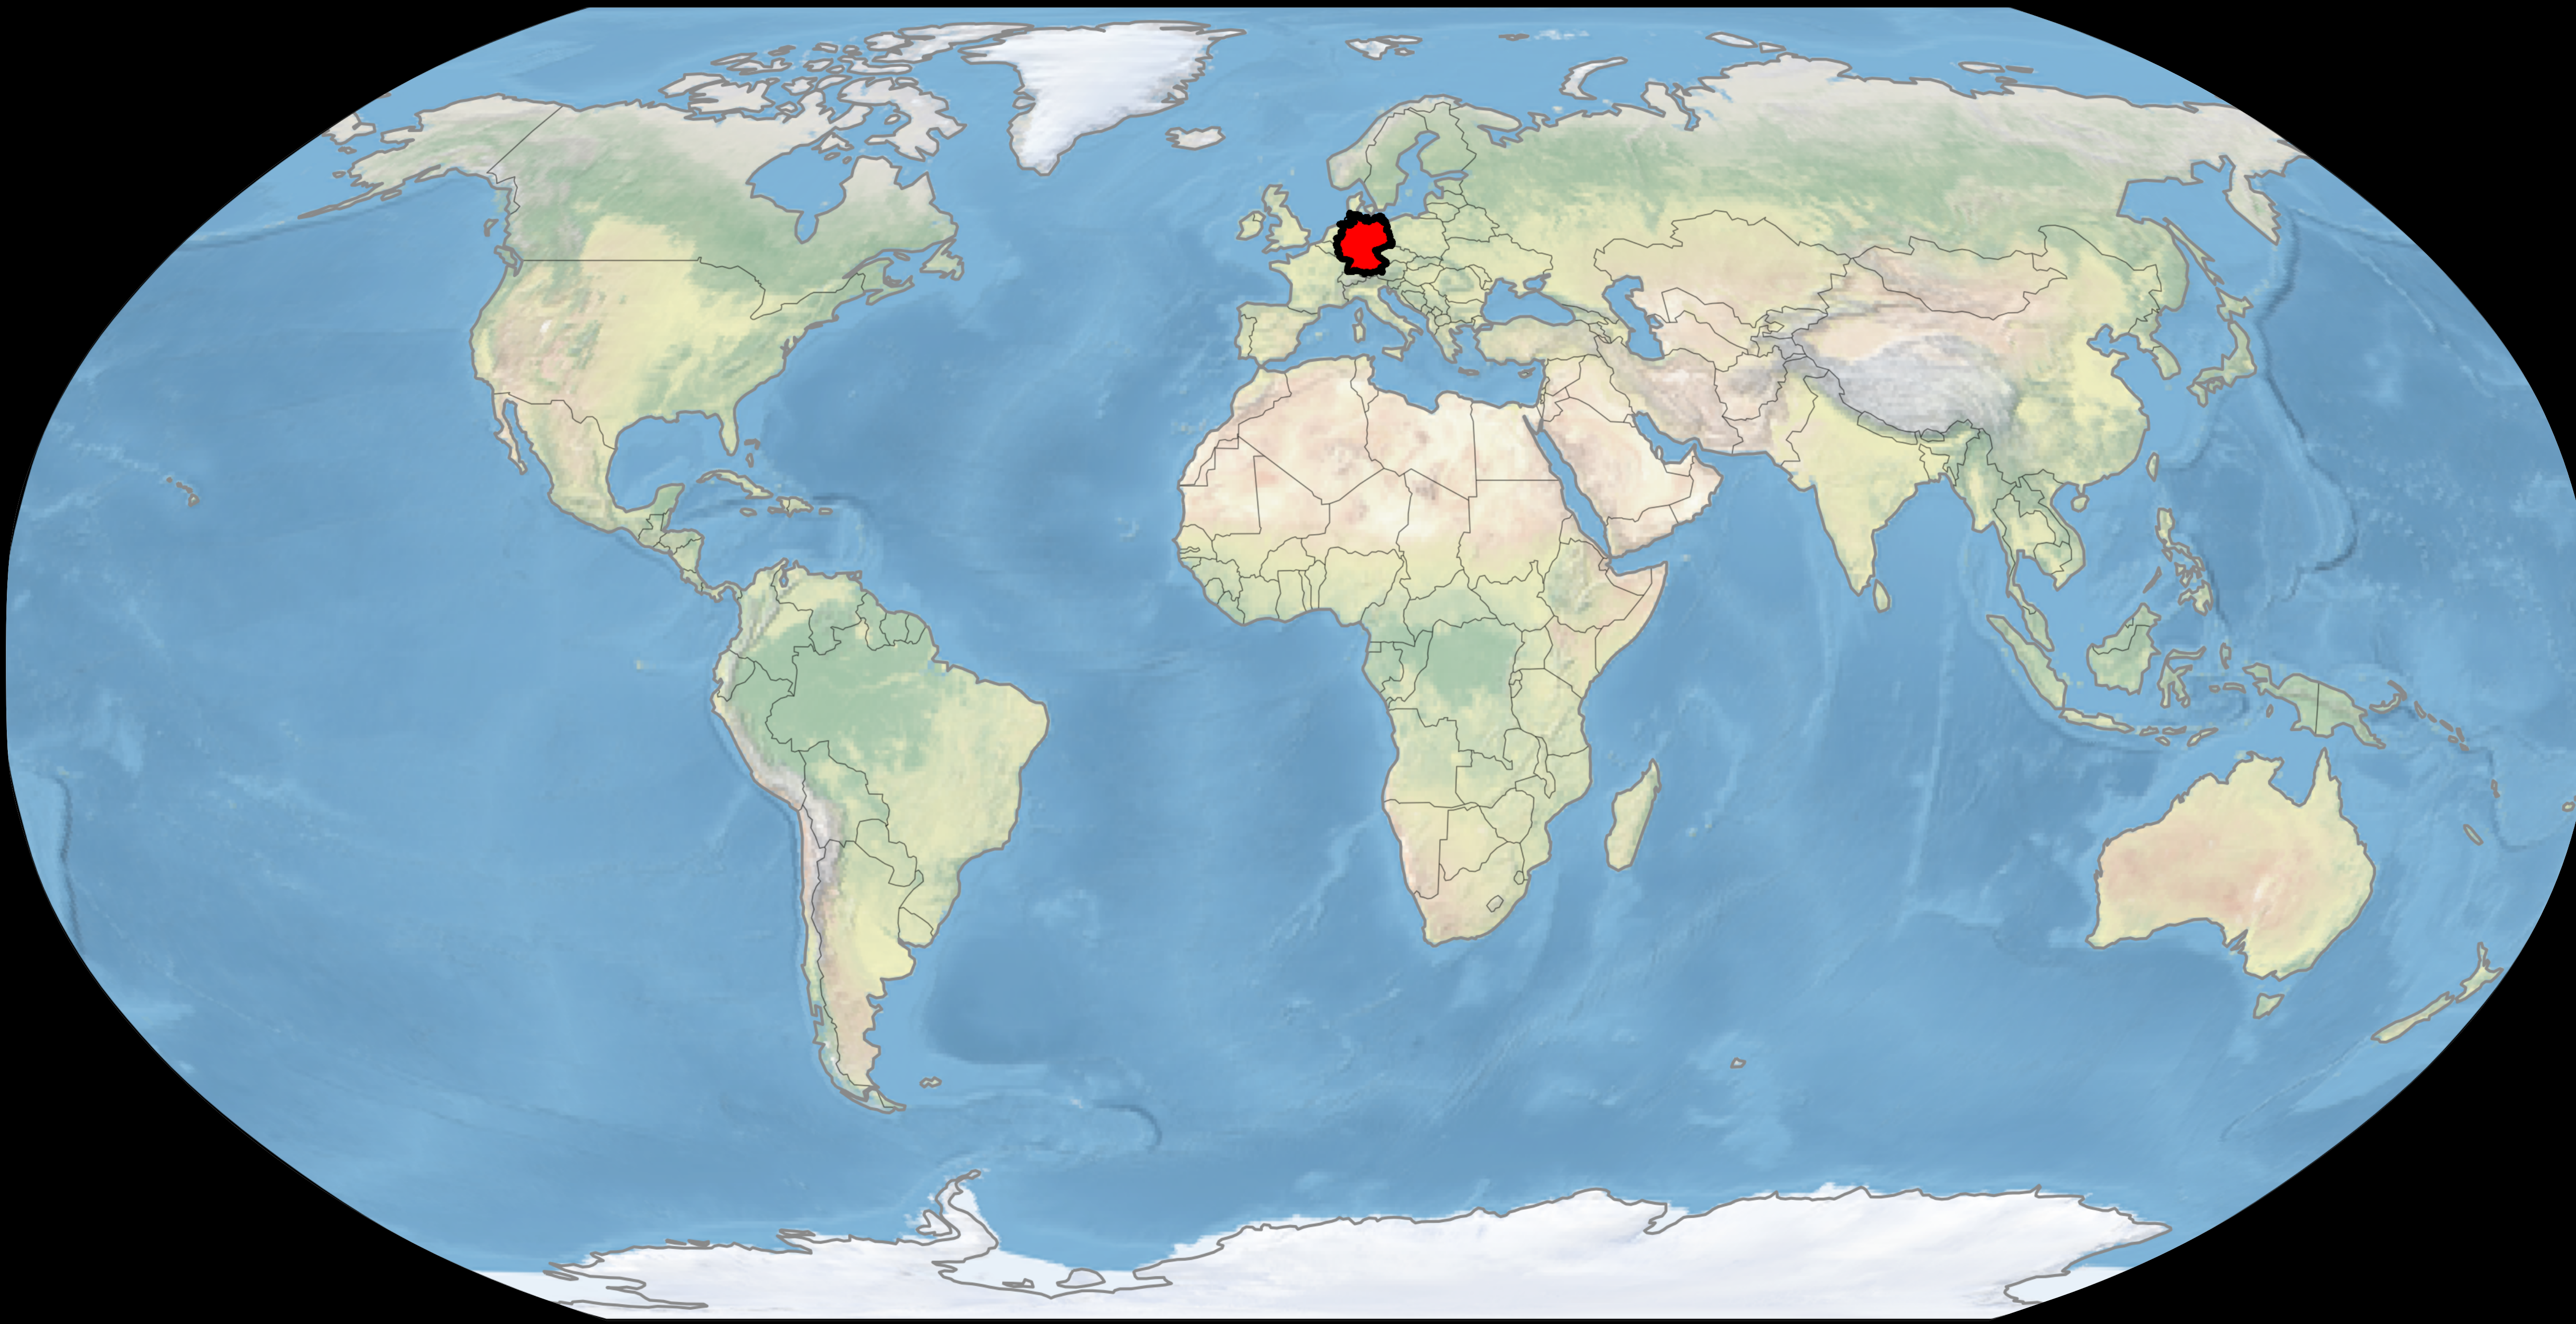

Arbitrary sets of polygons, defined over lat/lon

Compute target regions

Where alibis *might* be

Compute target regions

Where alibis *might* be

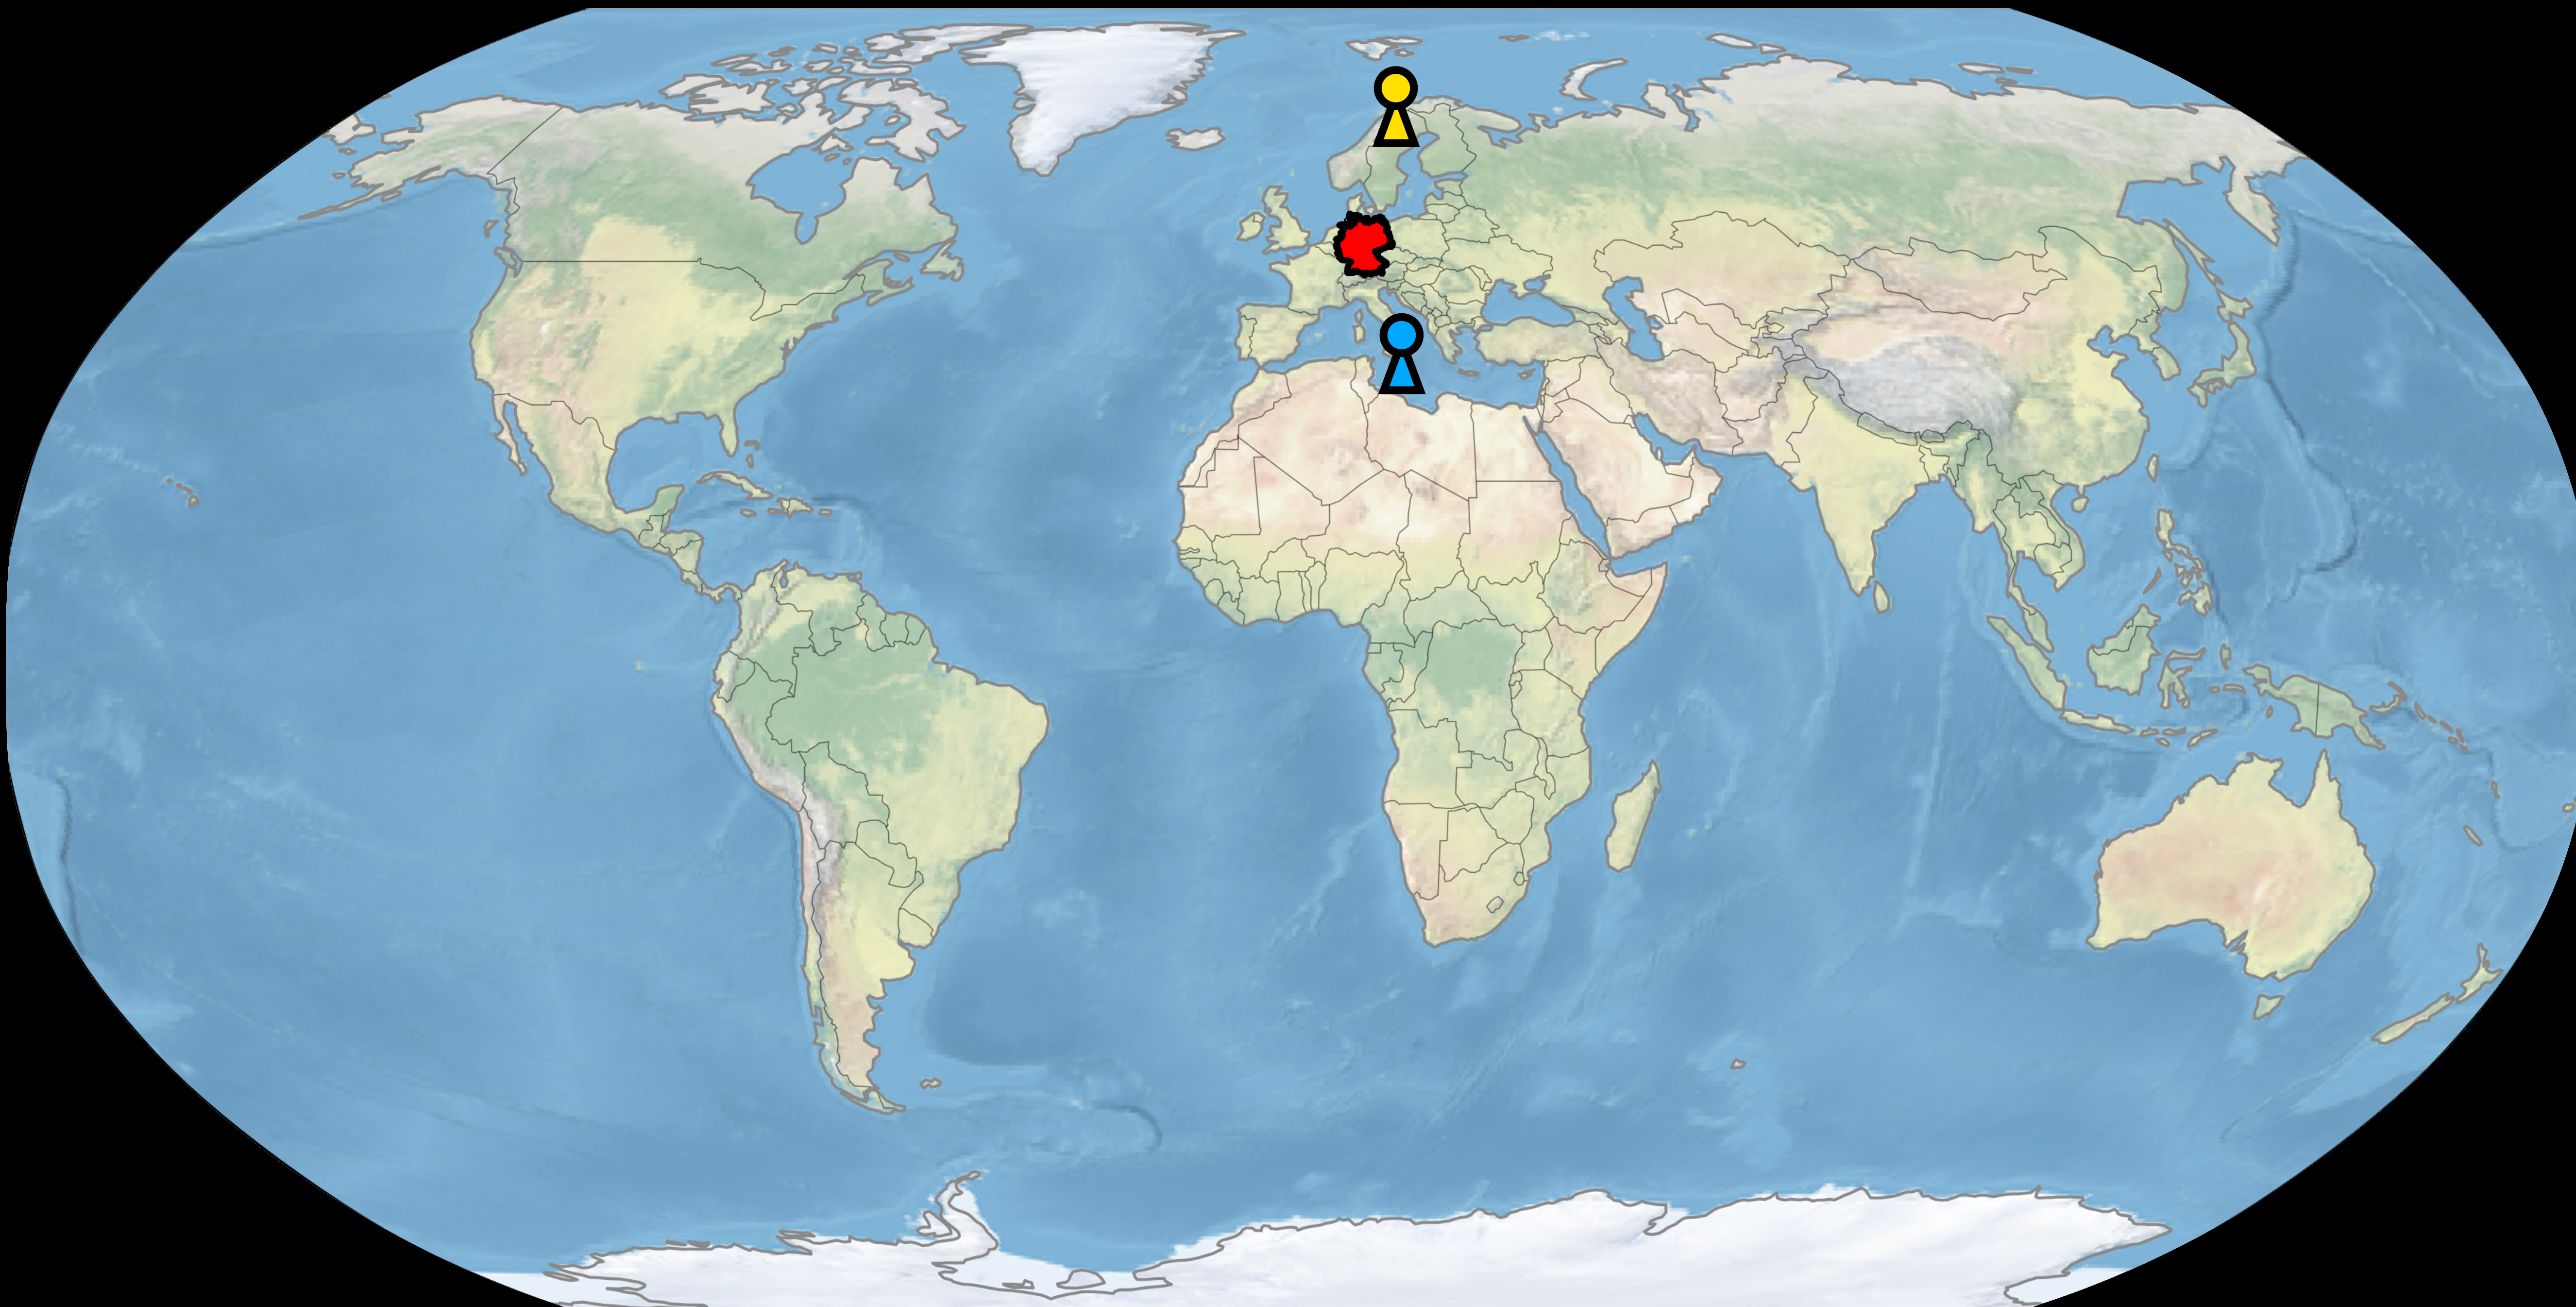

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

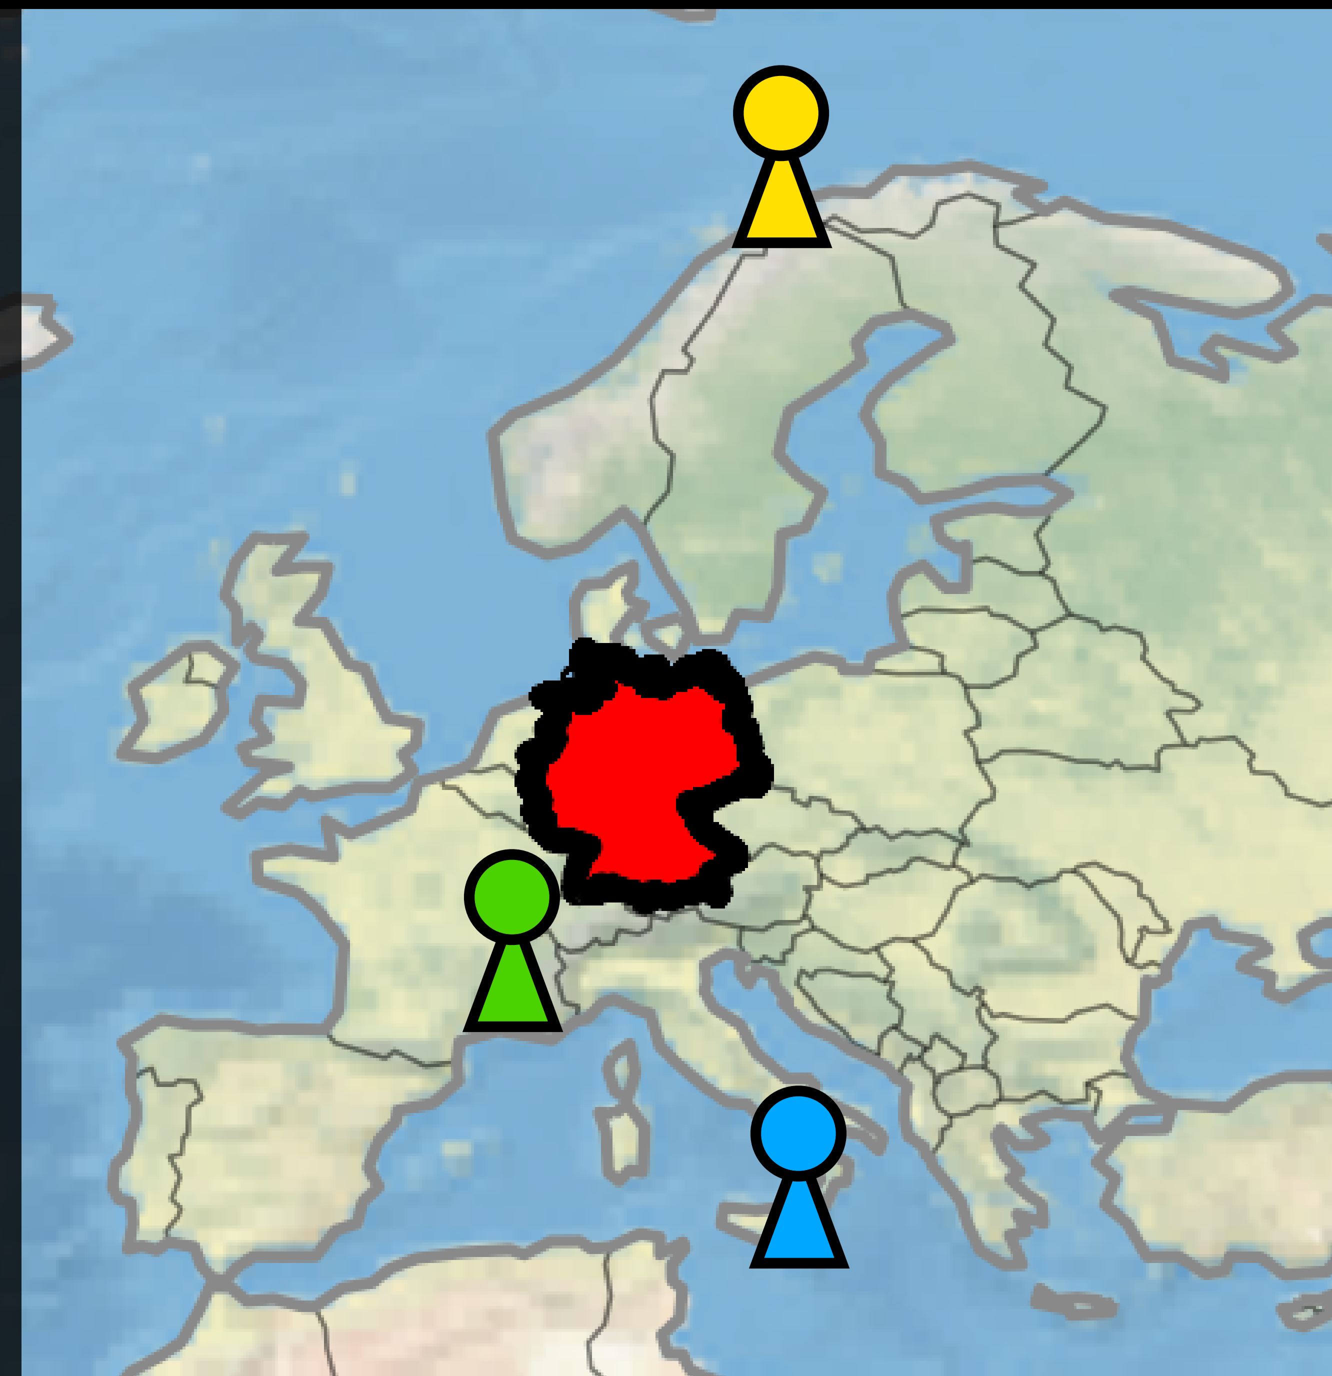

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

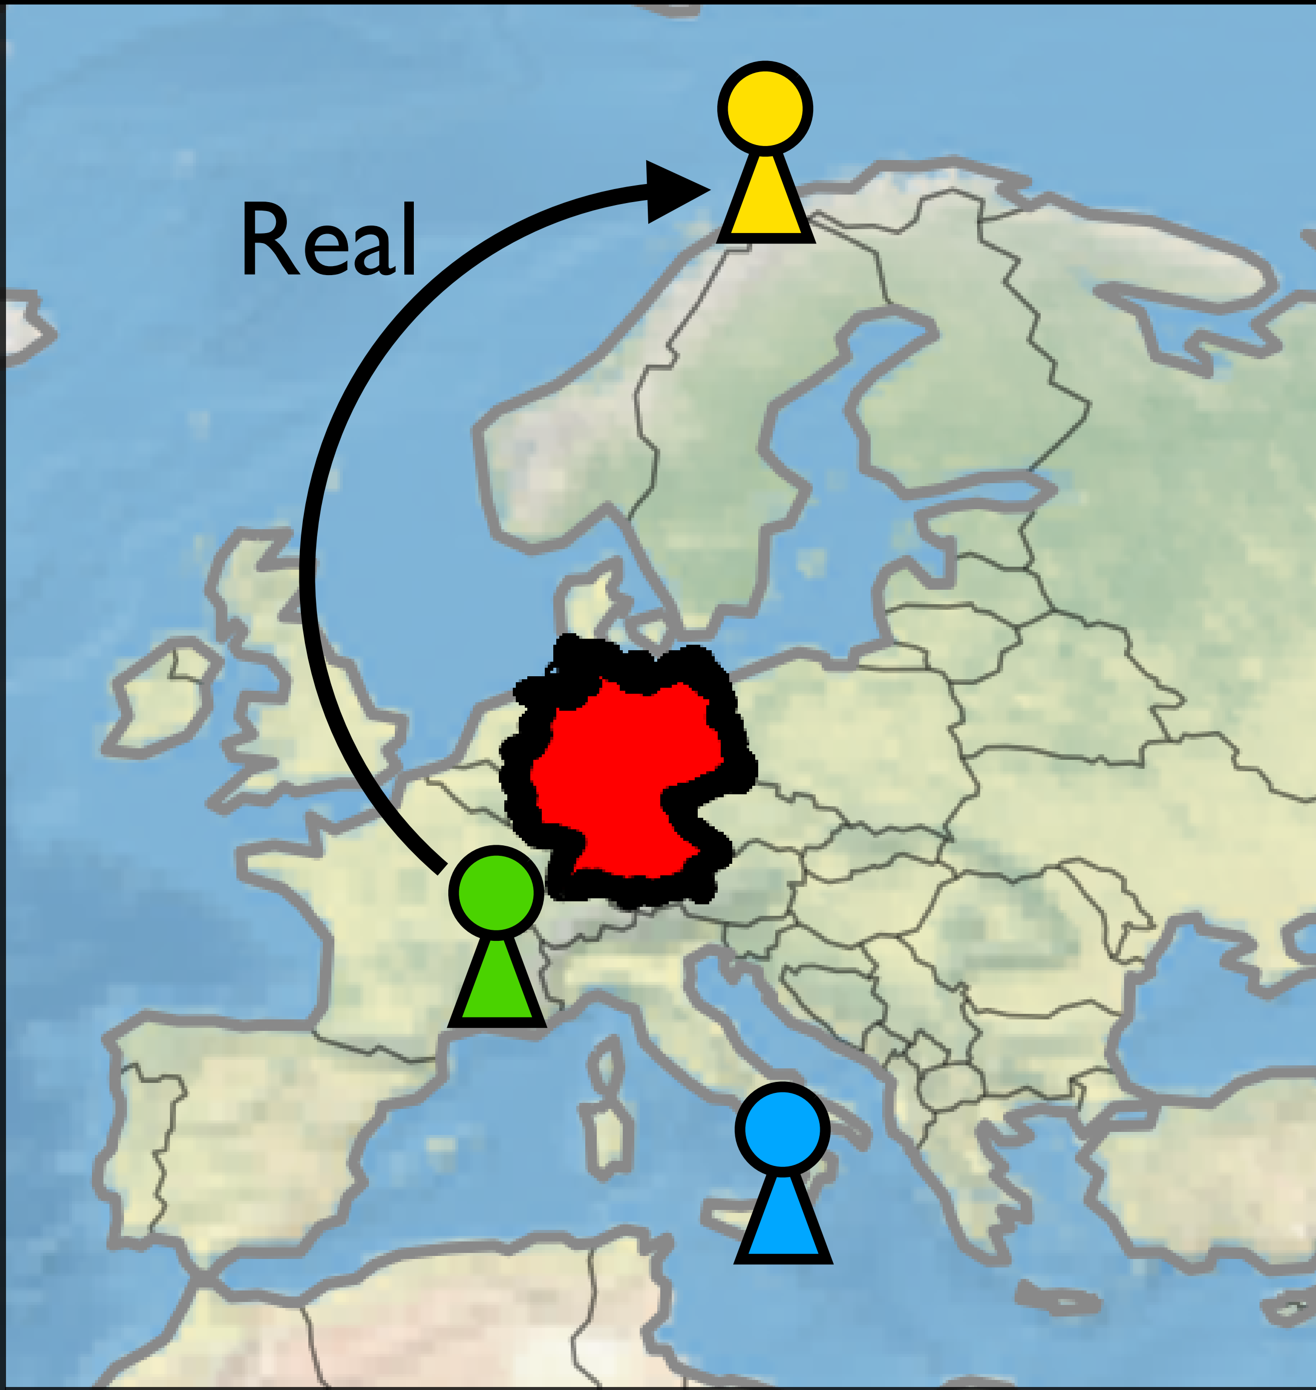

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

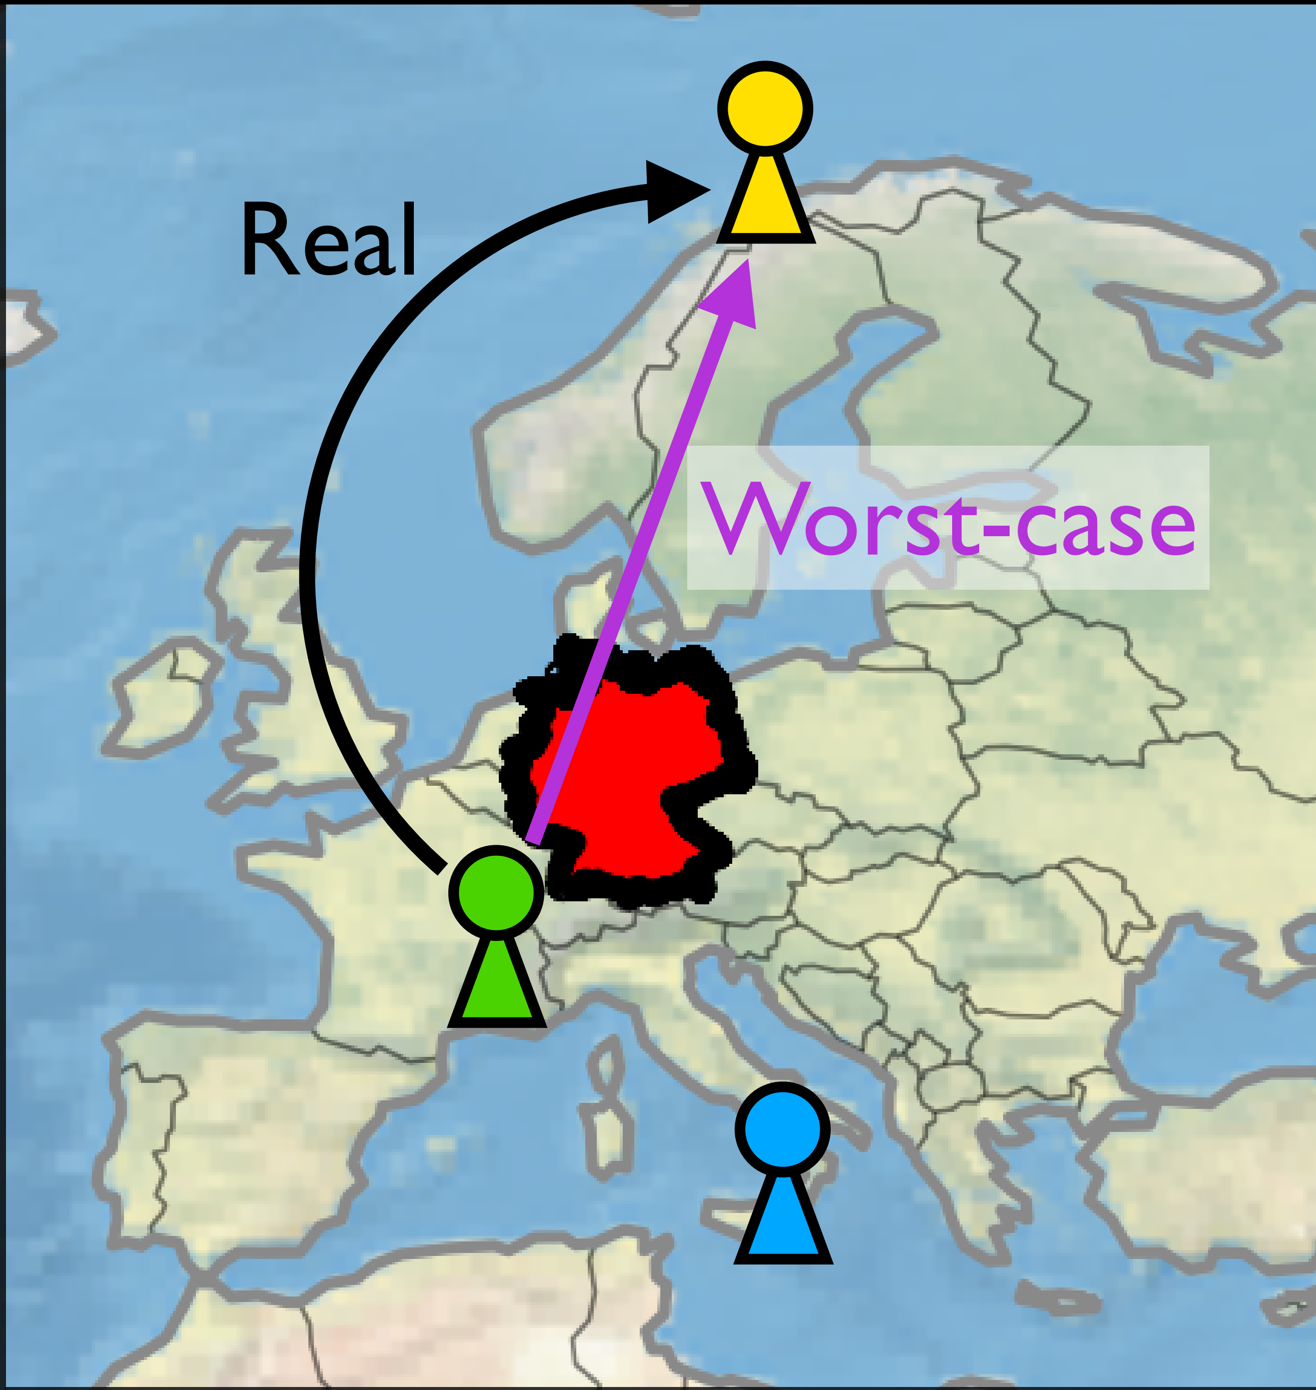

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

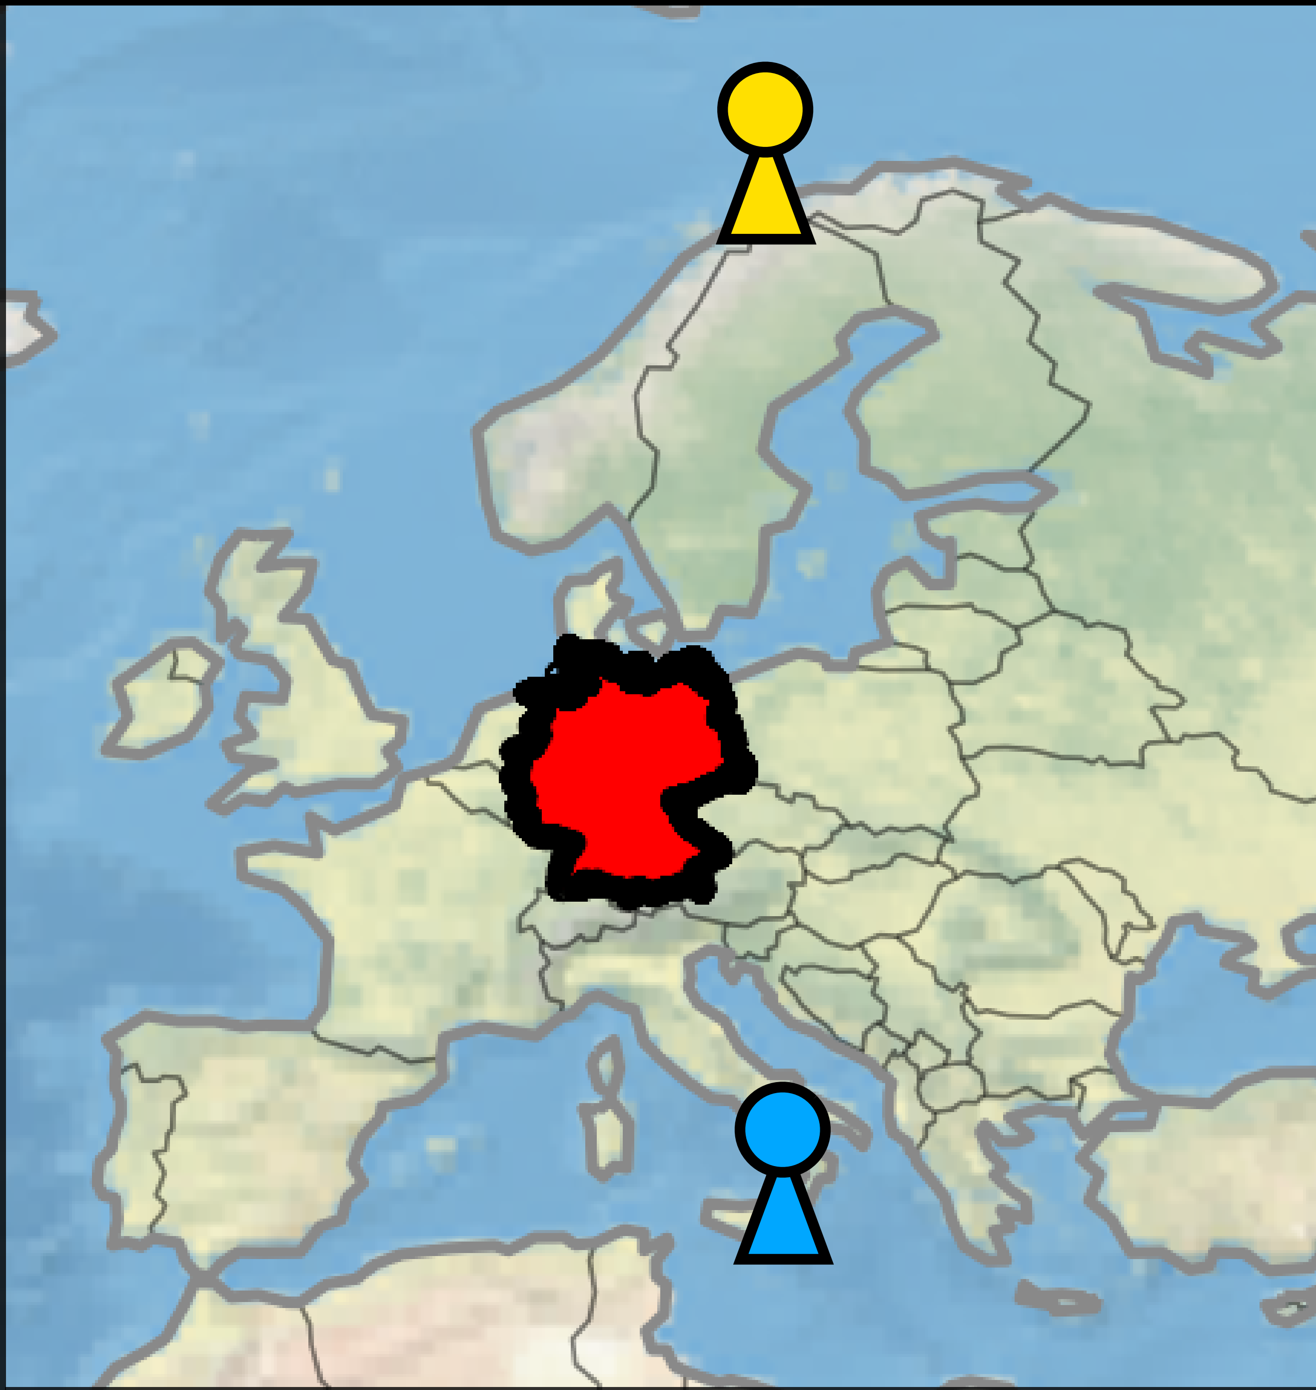

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

Segment the world into a grid

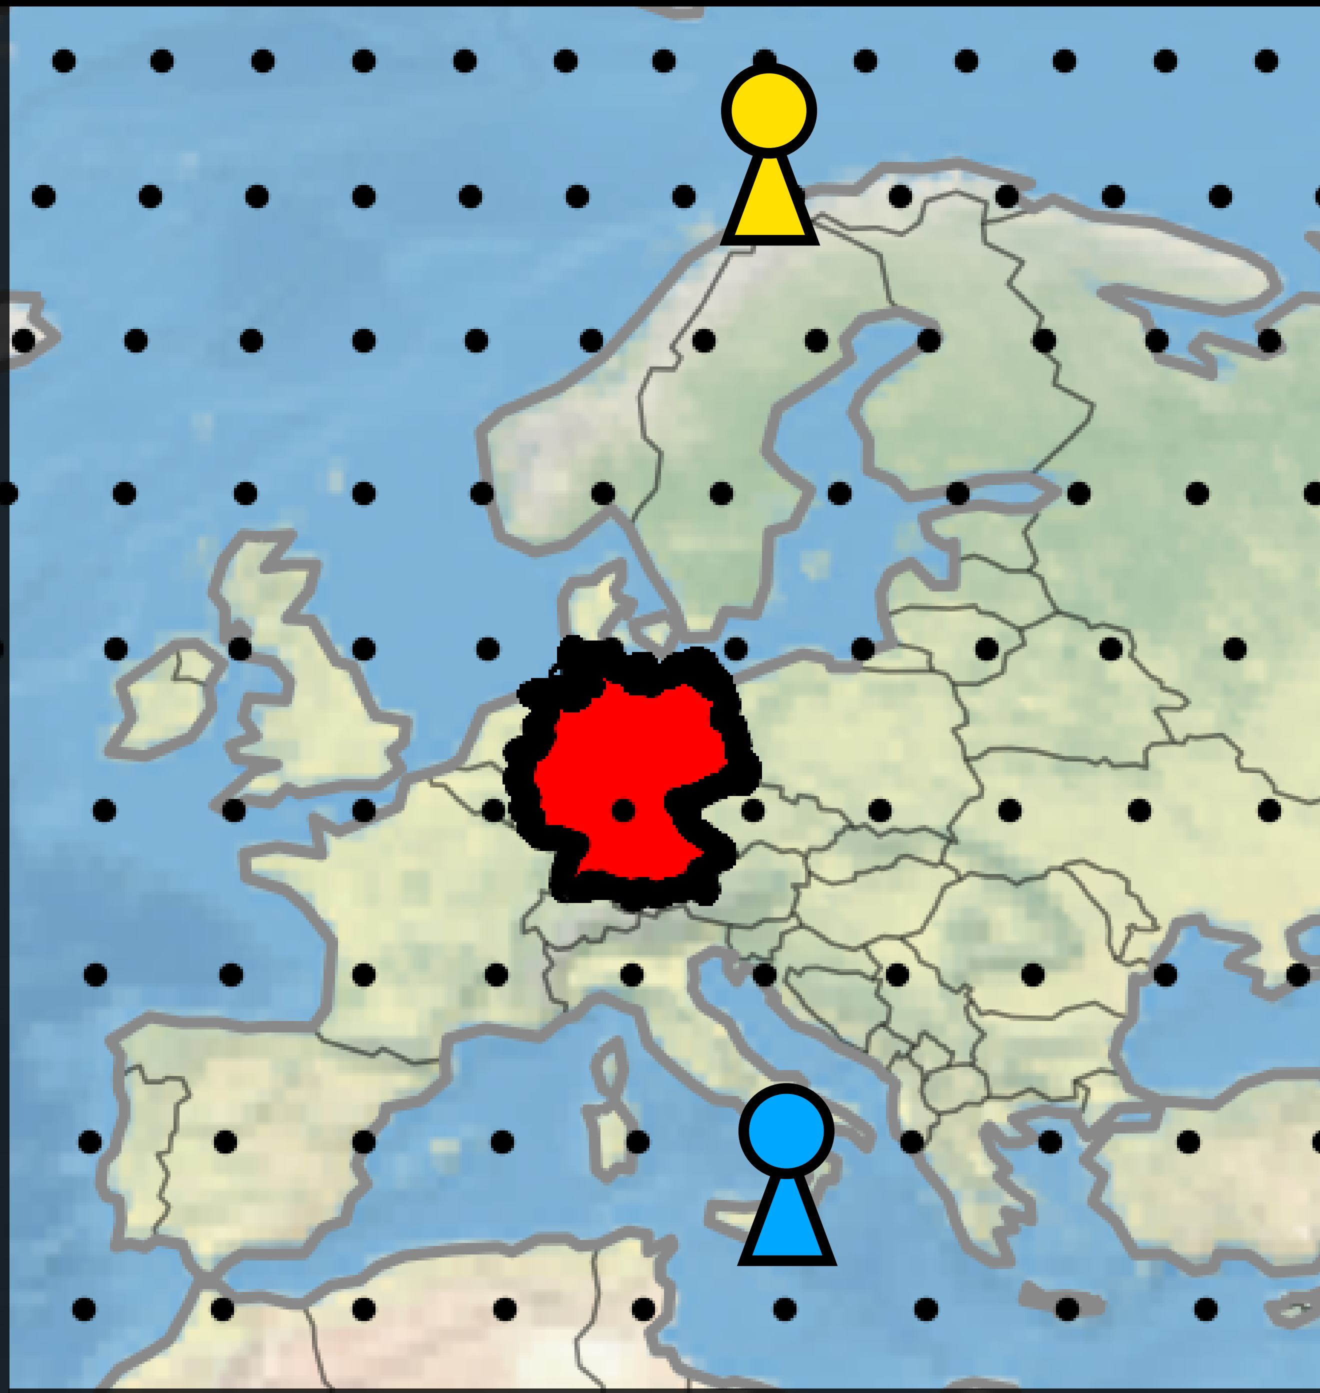

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

Segment the world into a grid

Include a grid point if:

$$(1 + \delta) * \text{Measured RTT} \leq 3\, d \,/\, c$$

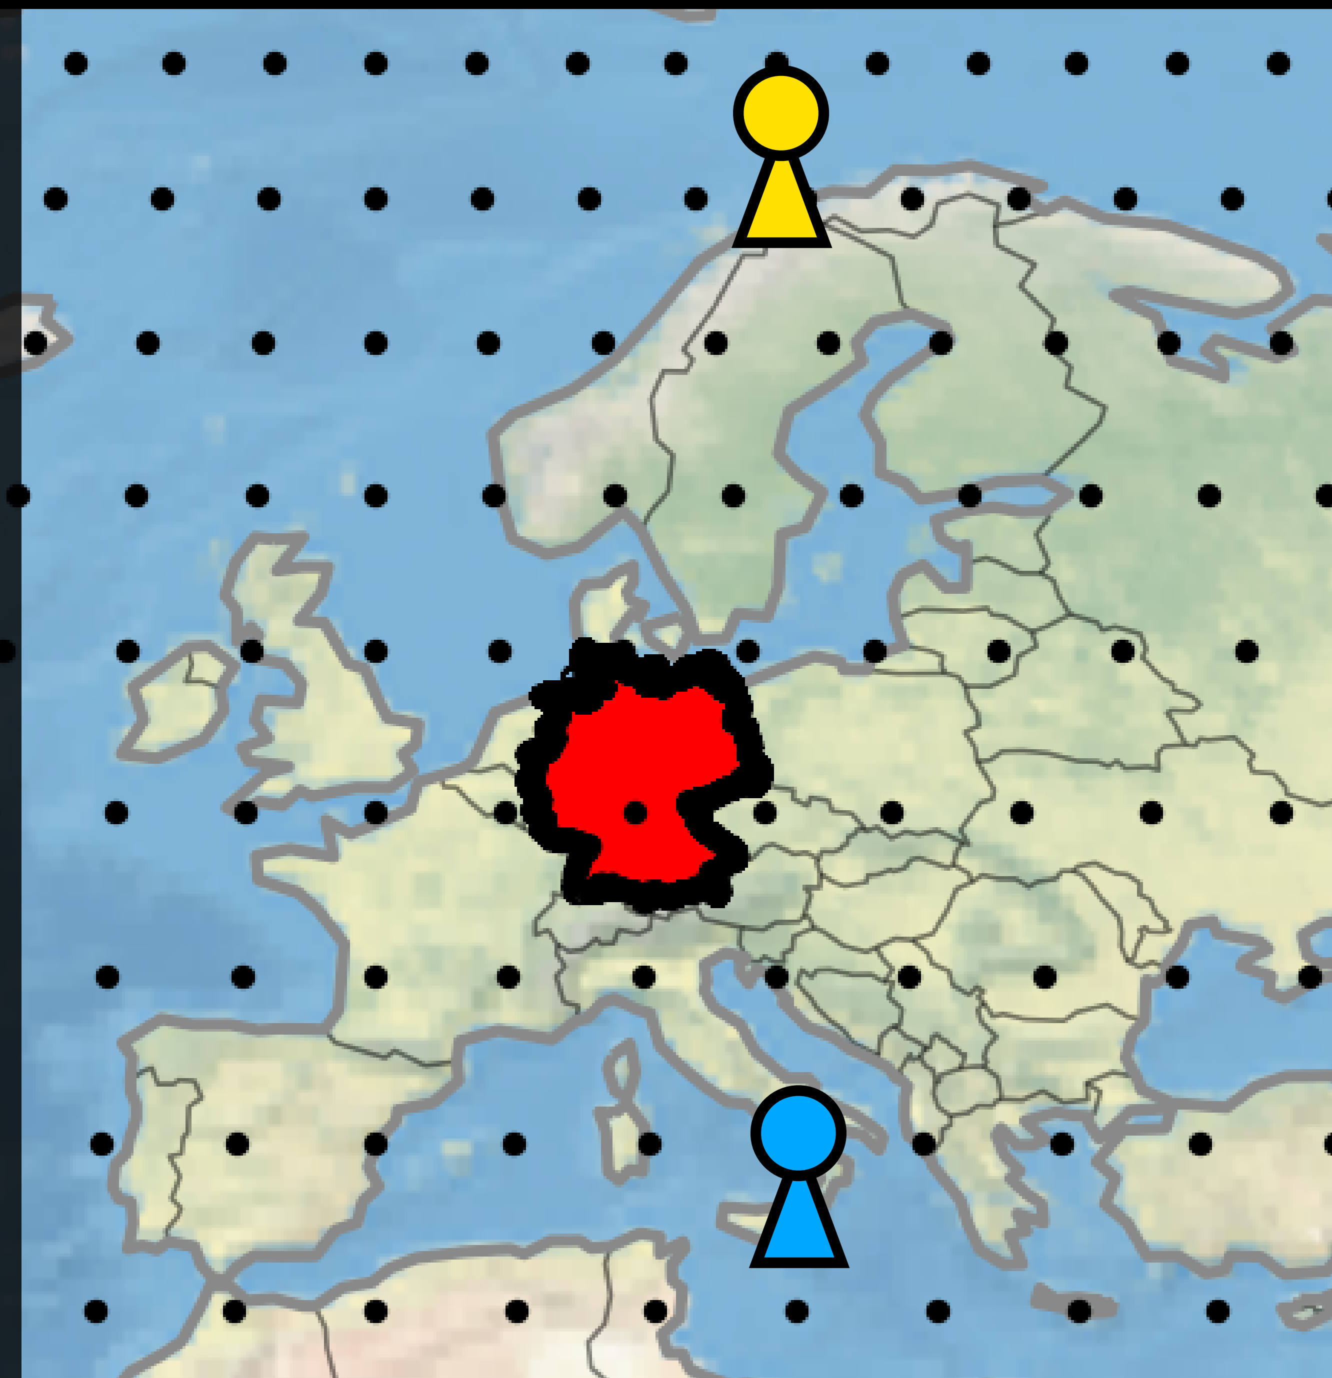

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

Segment the world into a grid

Include a grid point if:

$$(1 + \delta) * \boxed{\text{Measured RTT}} \leq 3\ d\ /\ c$$

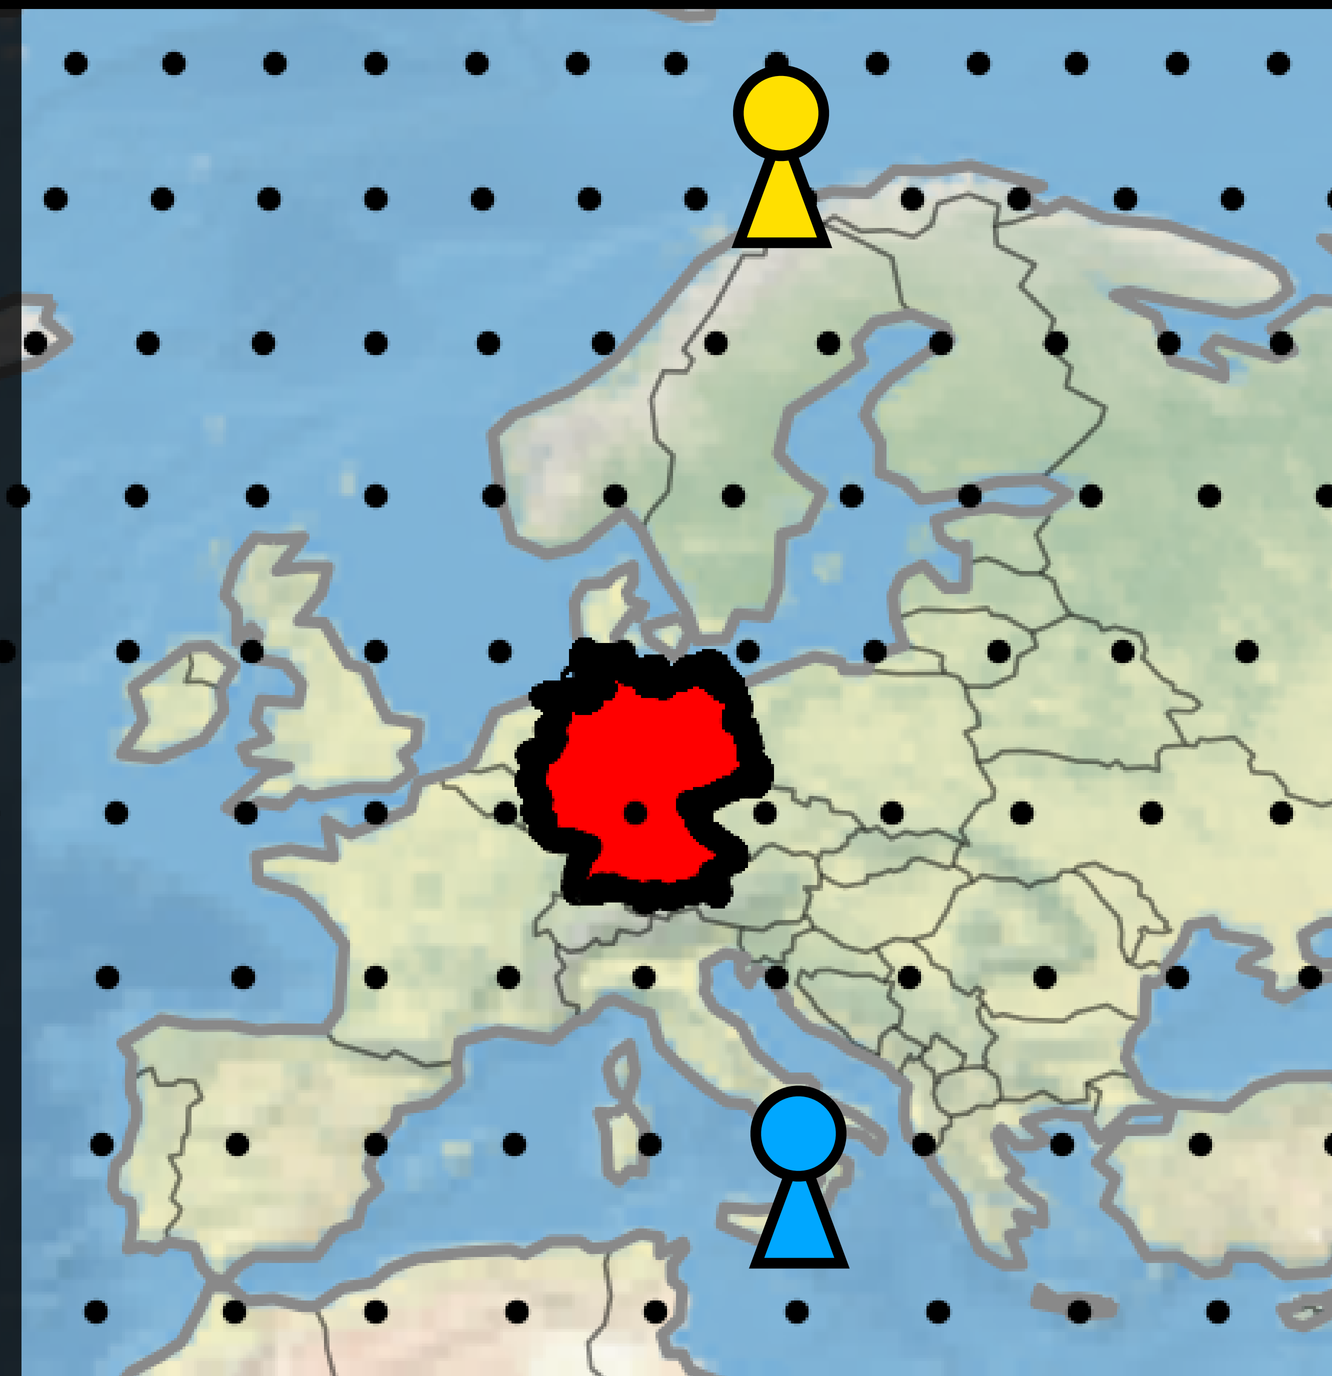

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

Segment the world into a grid

Include a grid point if:

$$(1 + \delta) * \text{Min } \textit{possible} \text{ RTT} \leq 3\, d / c$$

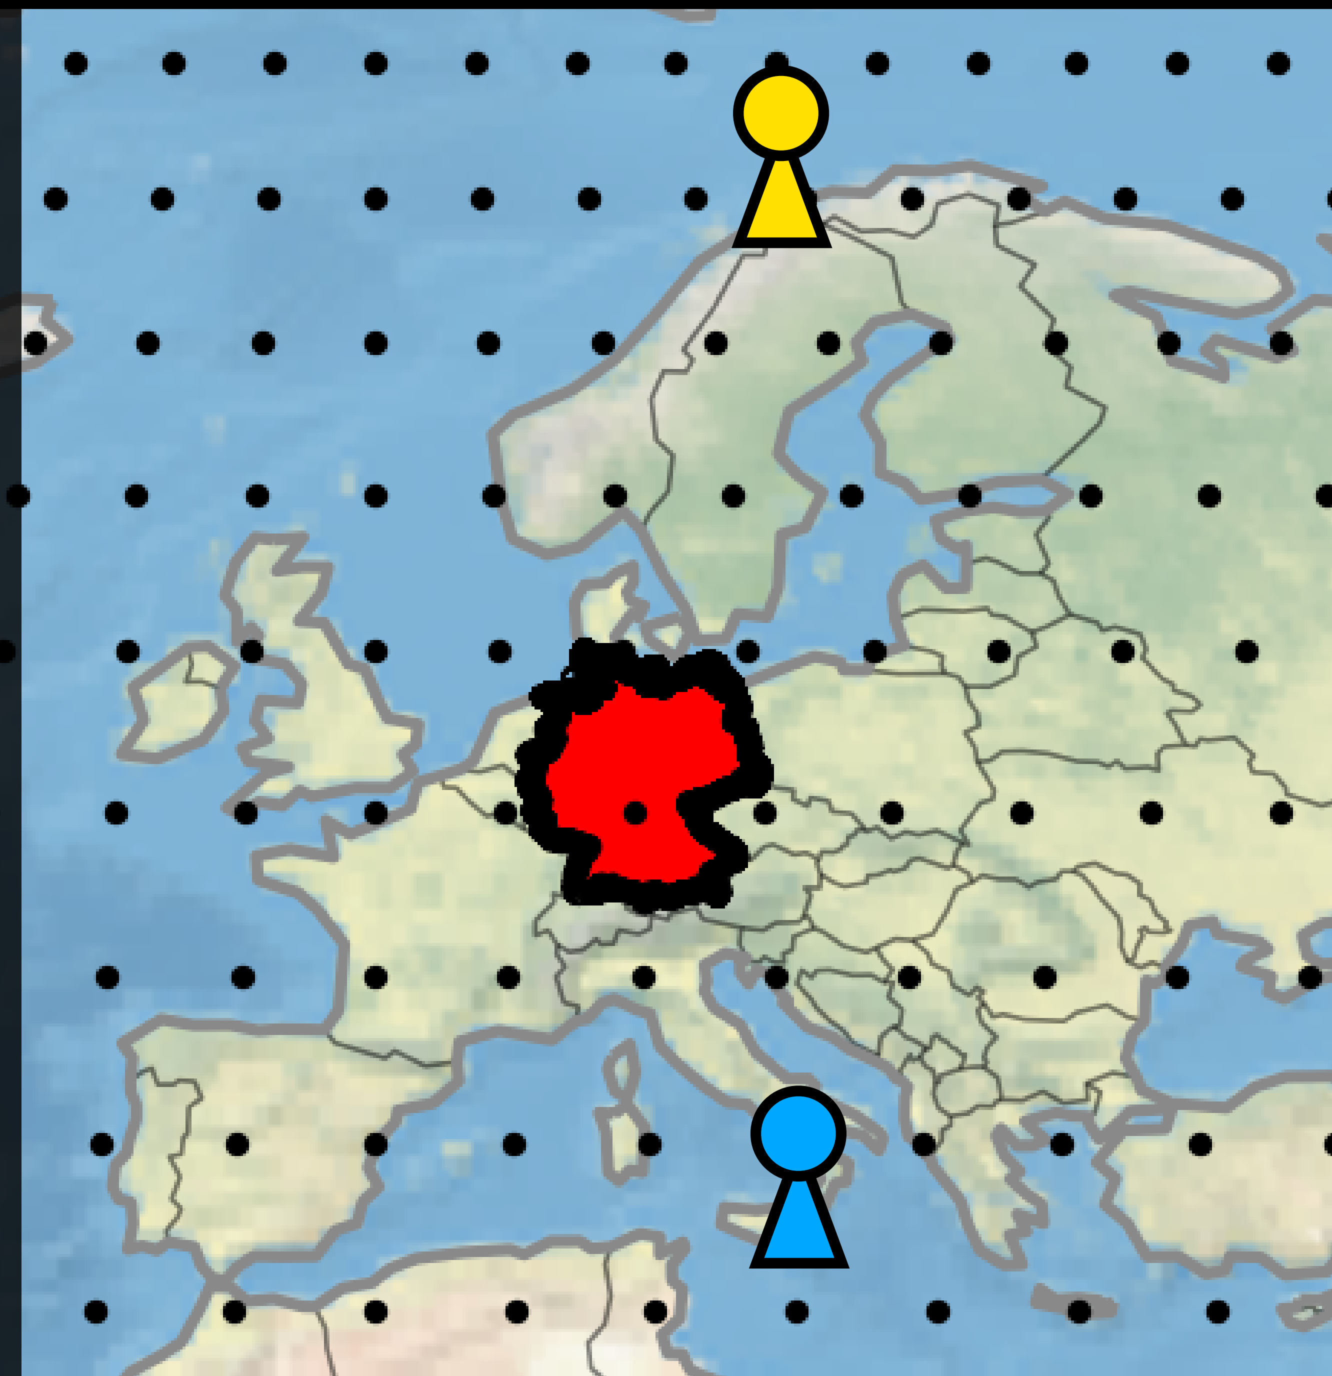

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

Segment the world into a grid

Include a grid point if:

$$(1+\delta) * \underset{\text{RTT}}{\text{Min } possible} \leq 3\, d / c$$

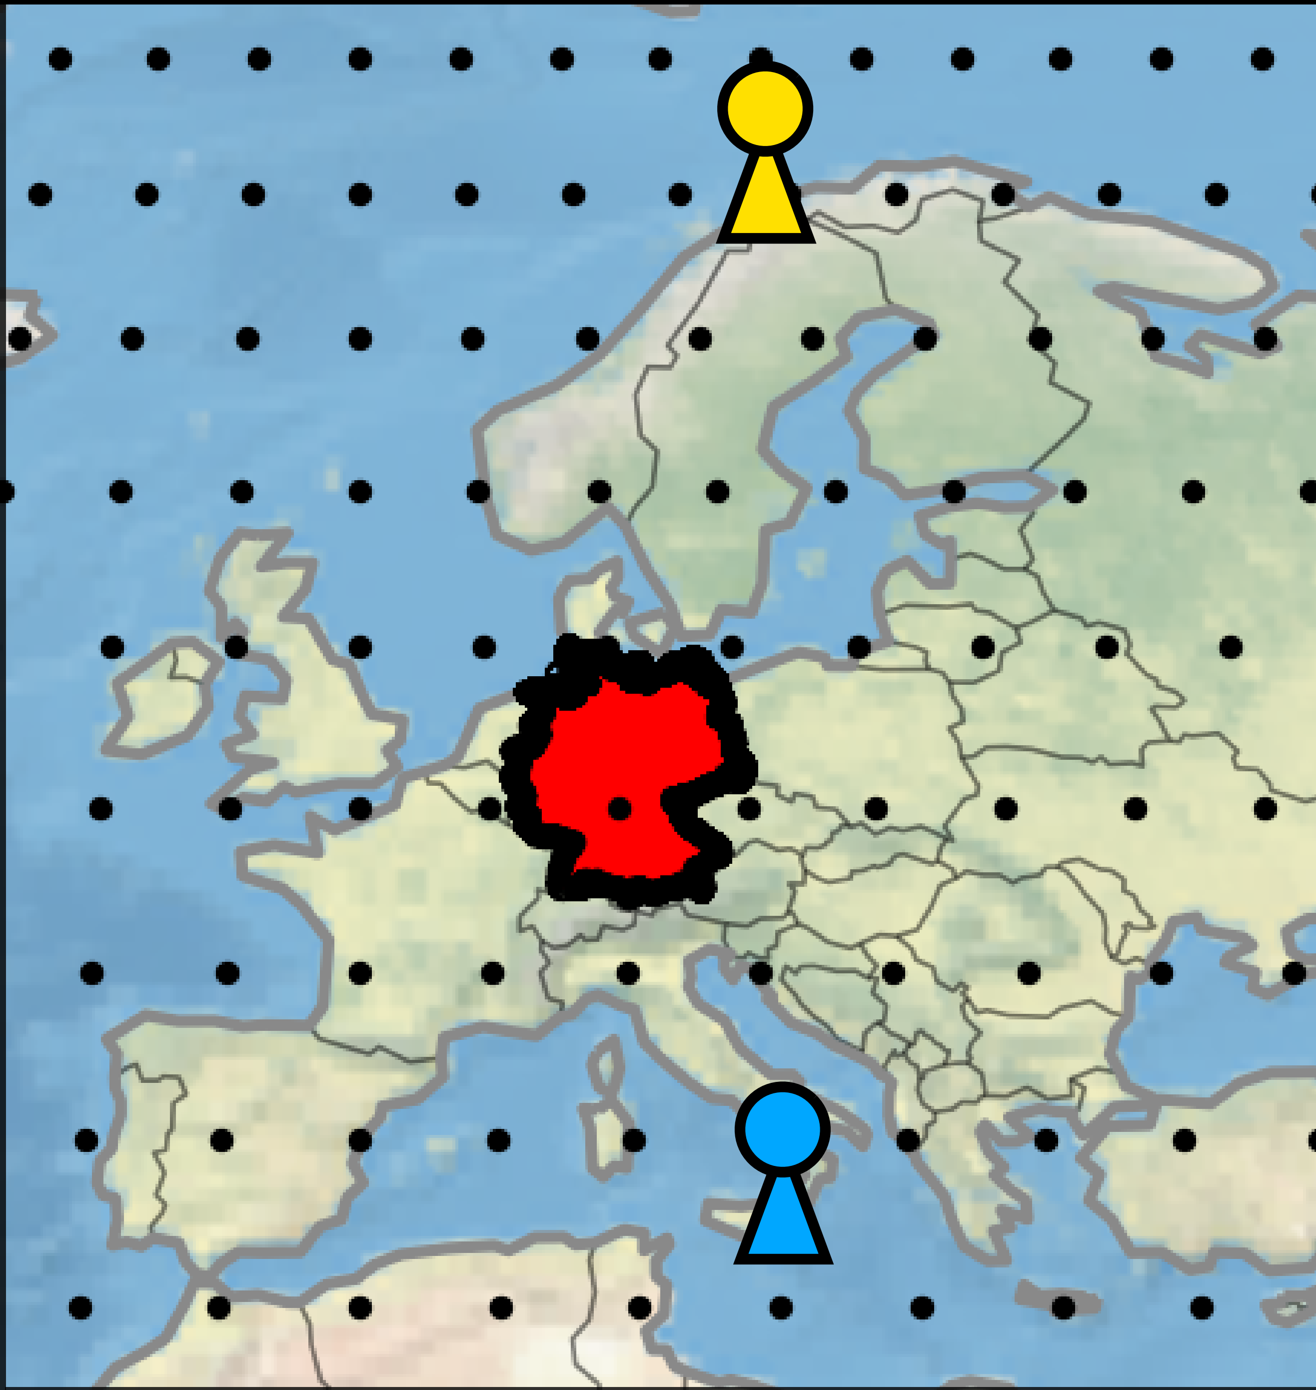

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

Segment the world into a grid

Include a grid point if:

$$(1 + \delta) * \text{Min } \textit{possible} \text{ RTT} \leq 3\, d / c$$

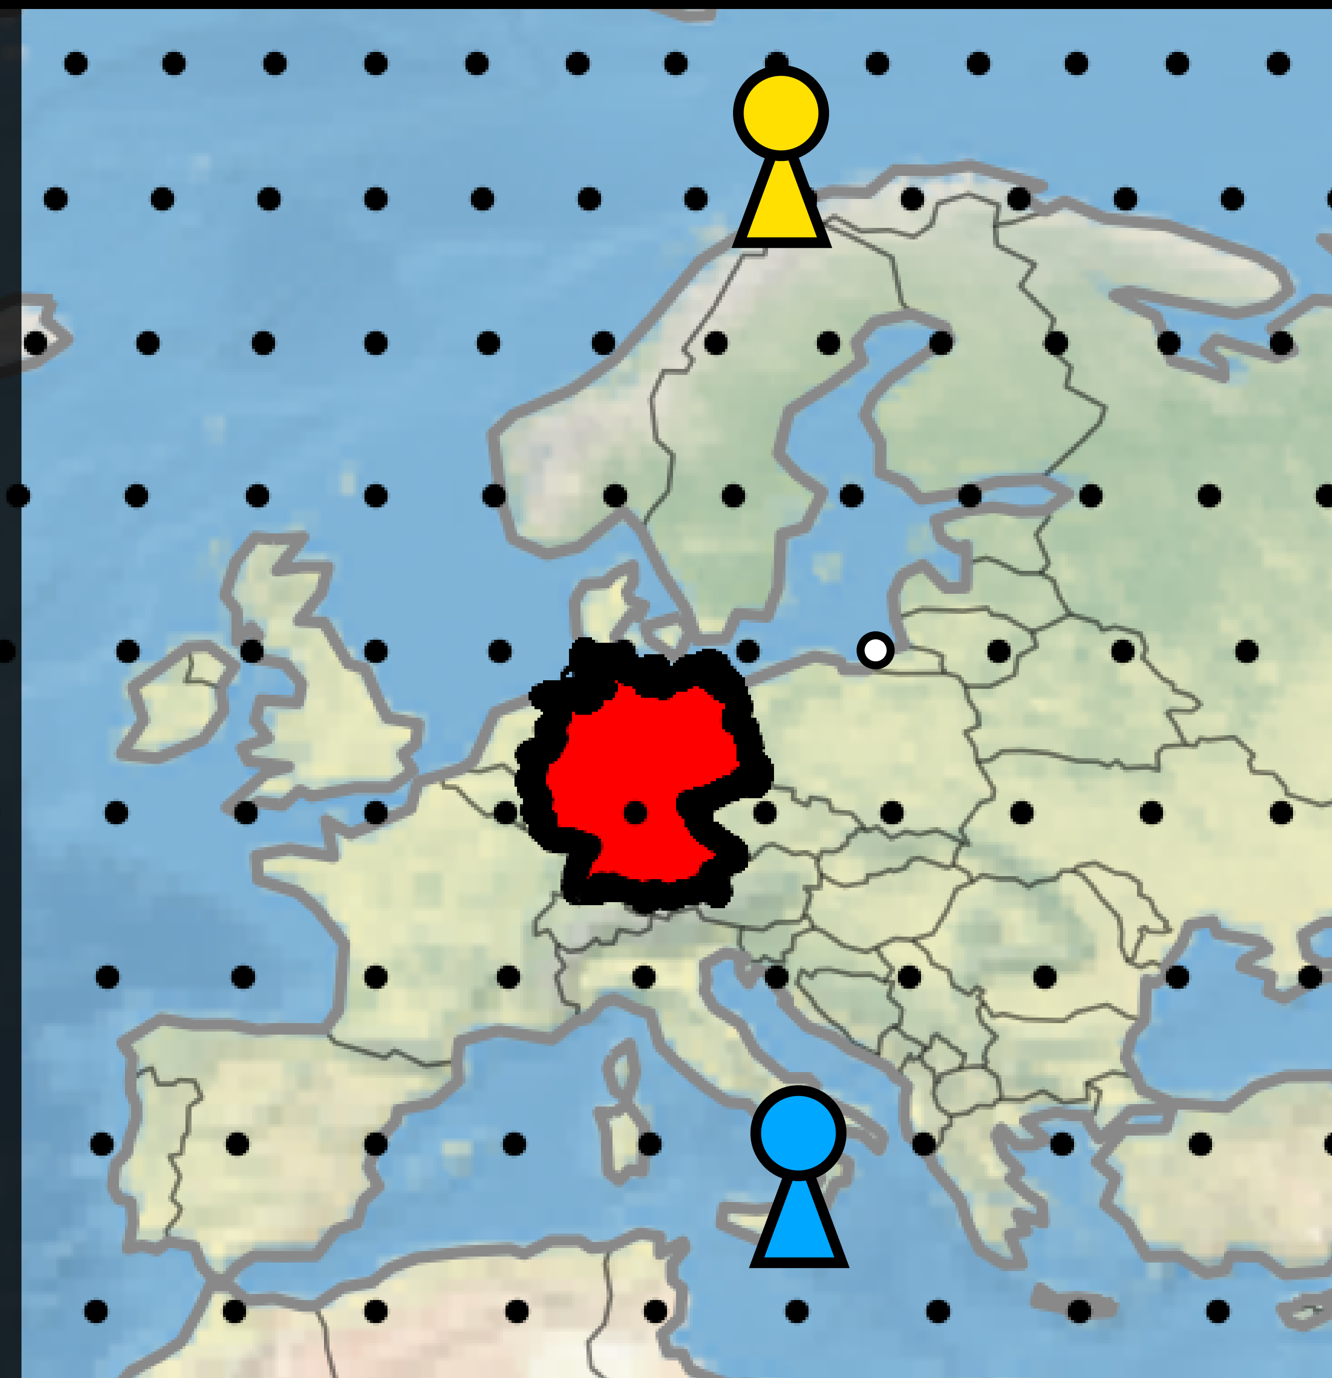

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

Segment the world into a grid

Include a grid point if:

$$(1 + \delta) * \text{Min } \textit{possible} \text{ RTT} \leq 3\, d / c$$

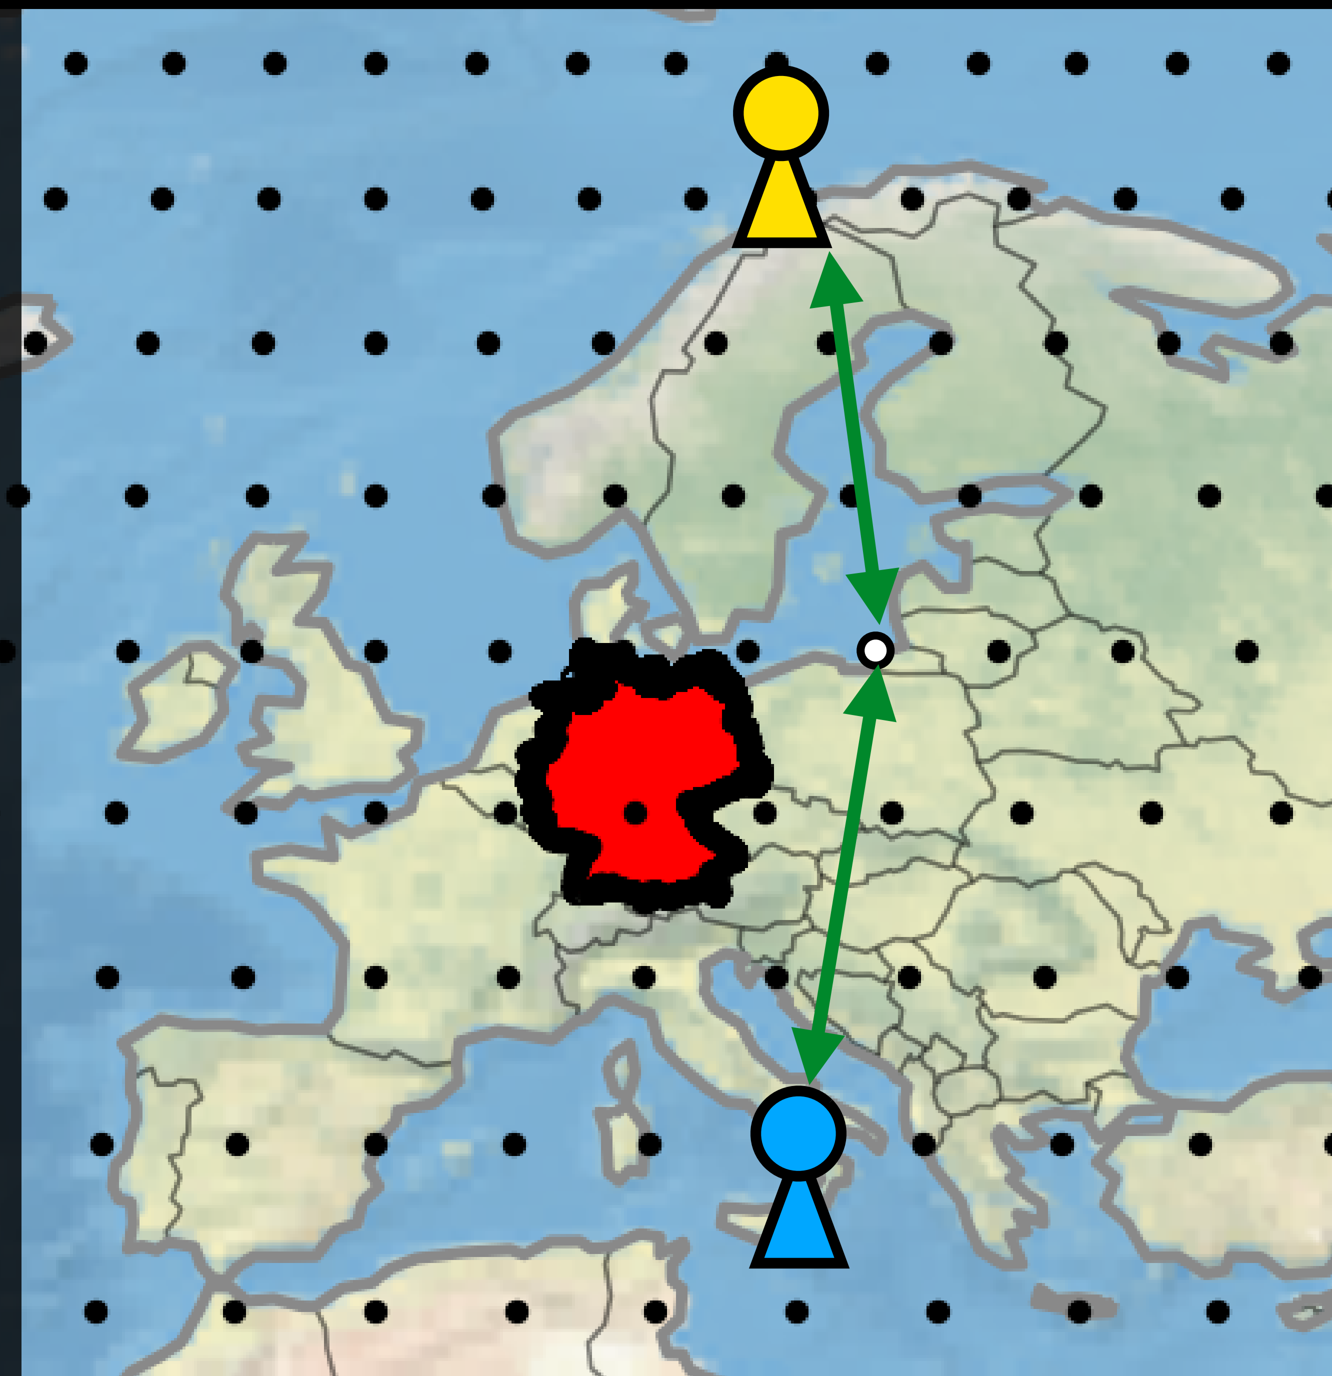

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

Segment the world into a grid

Include a grid point if:

$$(1 + \delta) * \text{Min } \textit{possible} \text{ RTT} \leq 3\, d / c$$

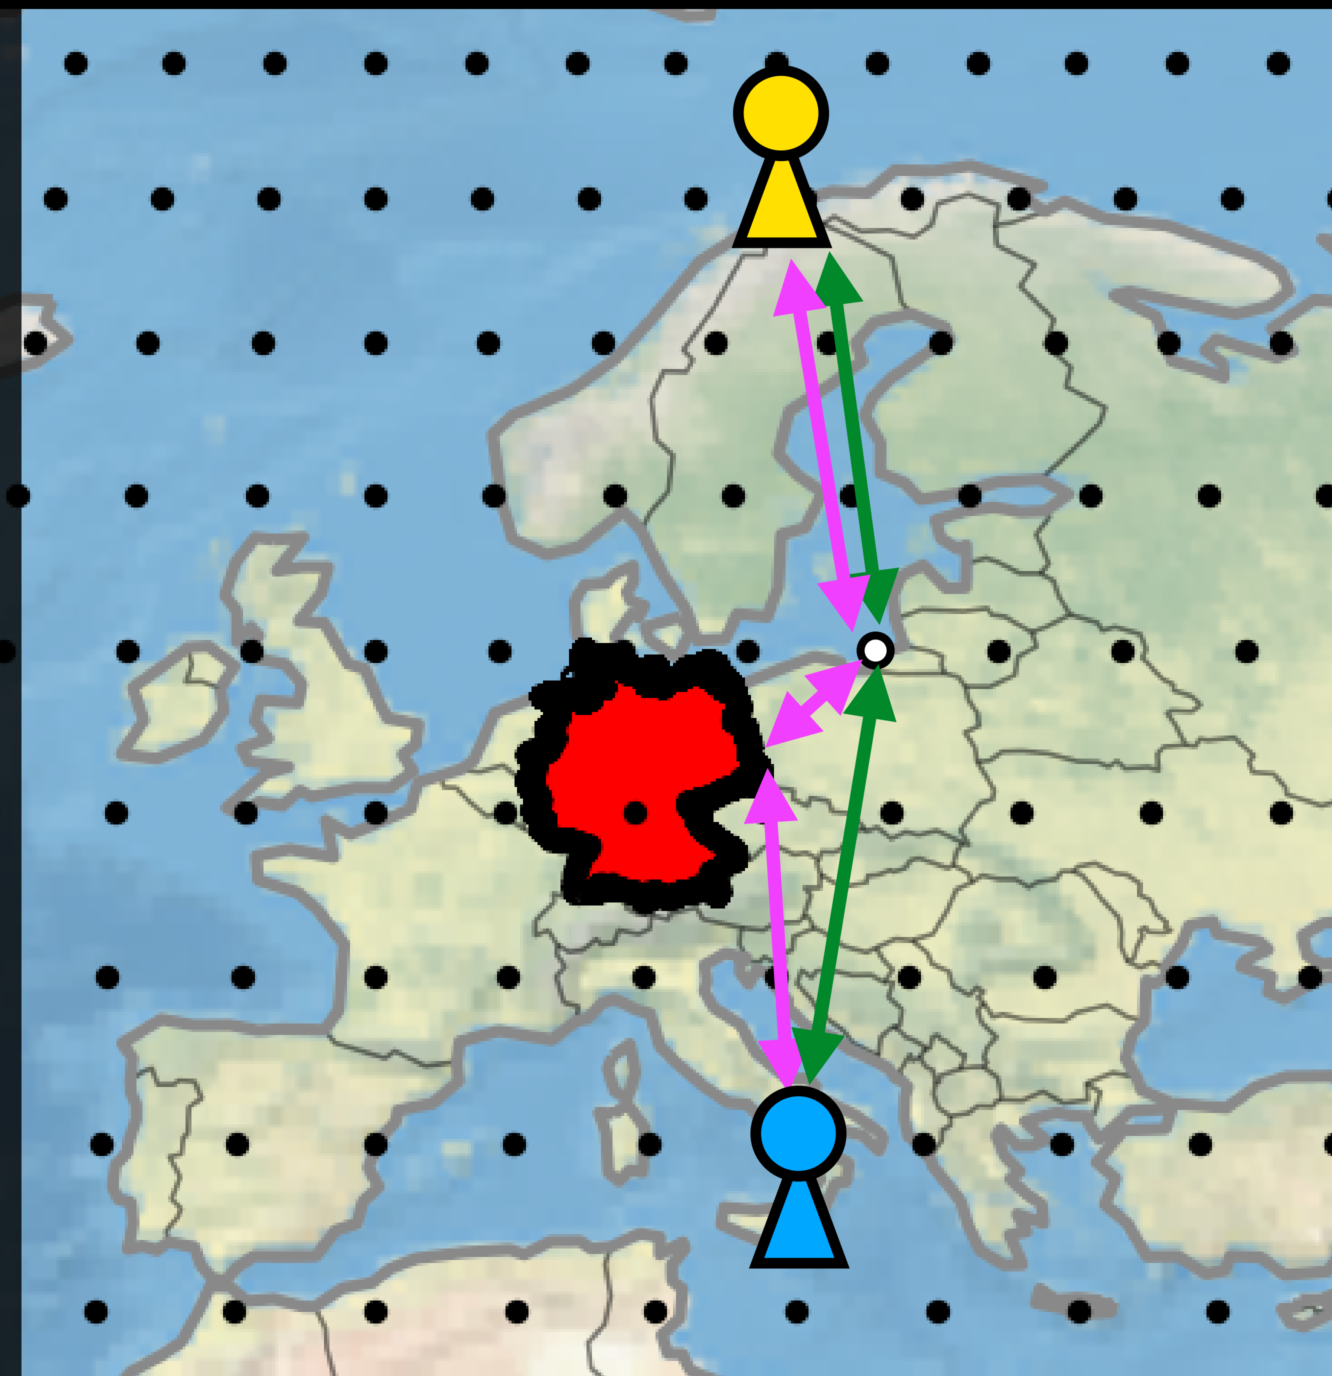

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

Segment the world into a grid

Include a grid point if:

$$(1 + \delta) \; * \; \text{Min } \textit{possible} \text{ RTT} \;\leq\; 3\, d \,/\, c$$

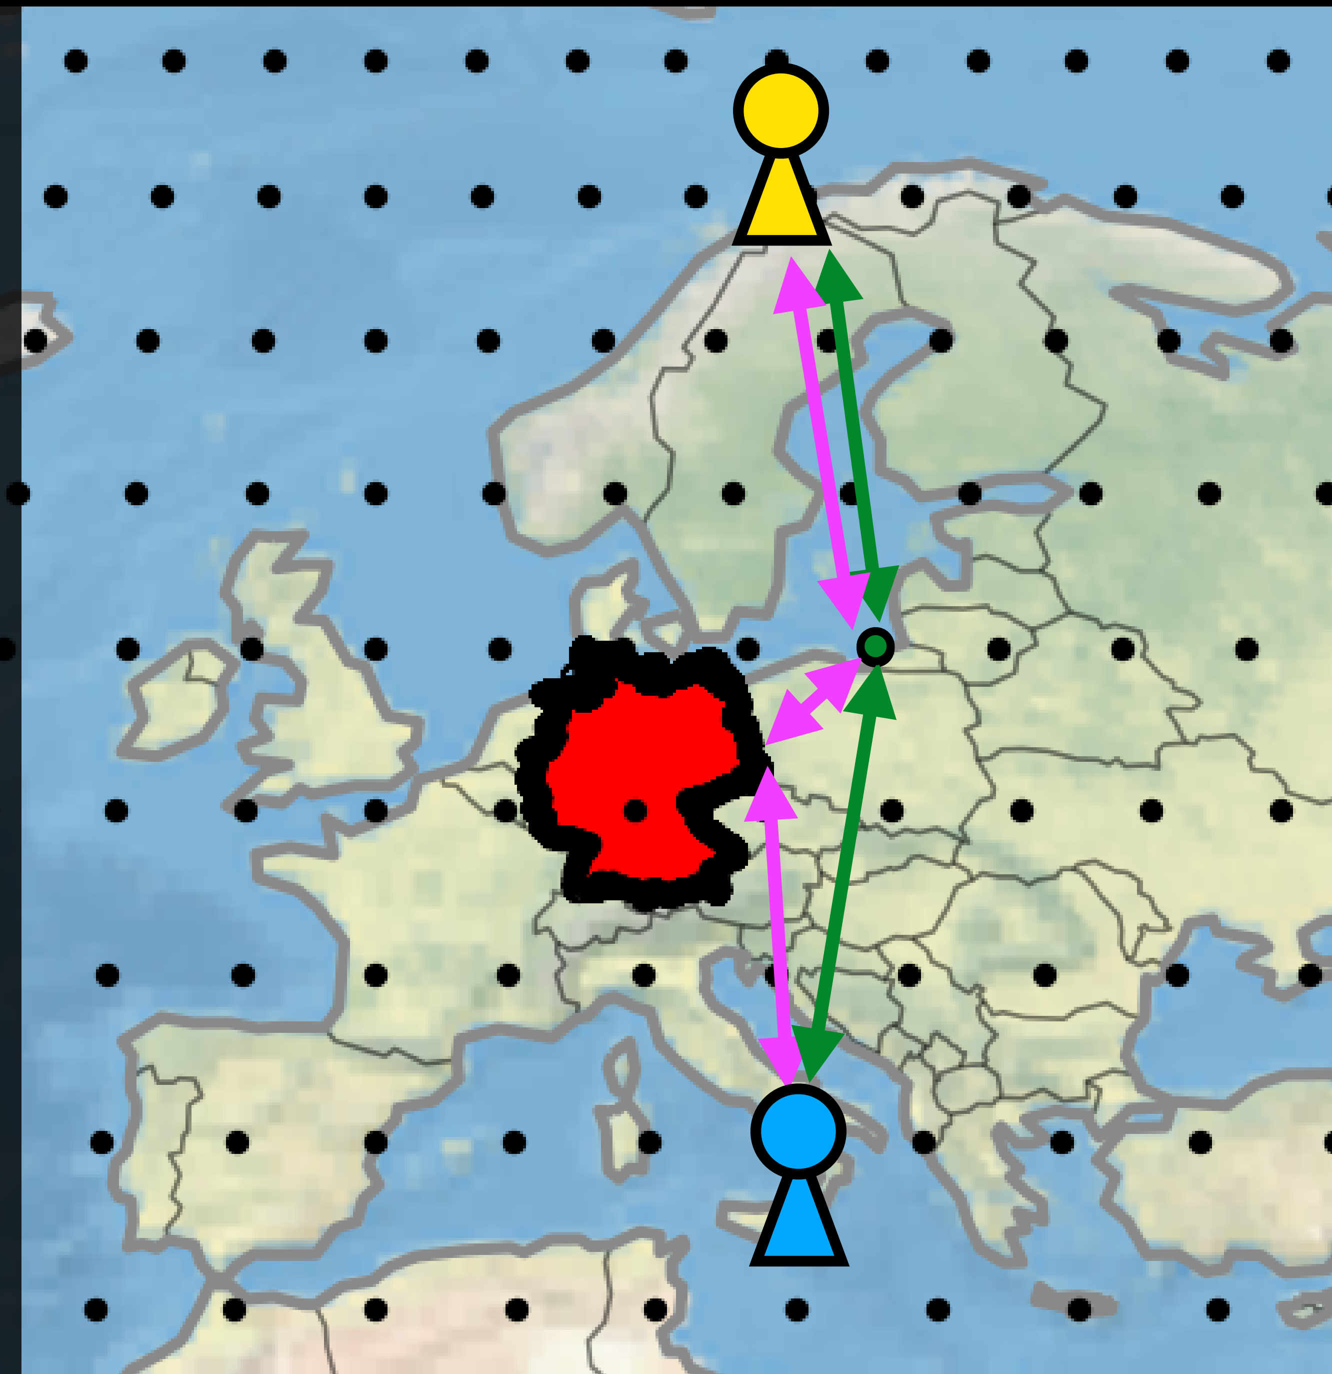

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

Segment the world into a grid

Include a grid point if:

$$(1+\delta) * \text{Min } \textit{possible}\ \text{RTT} \leq 3\ d\ /\ c$$

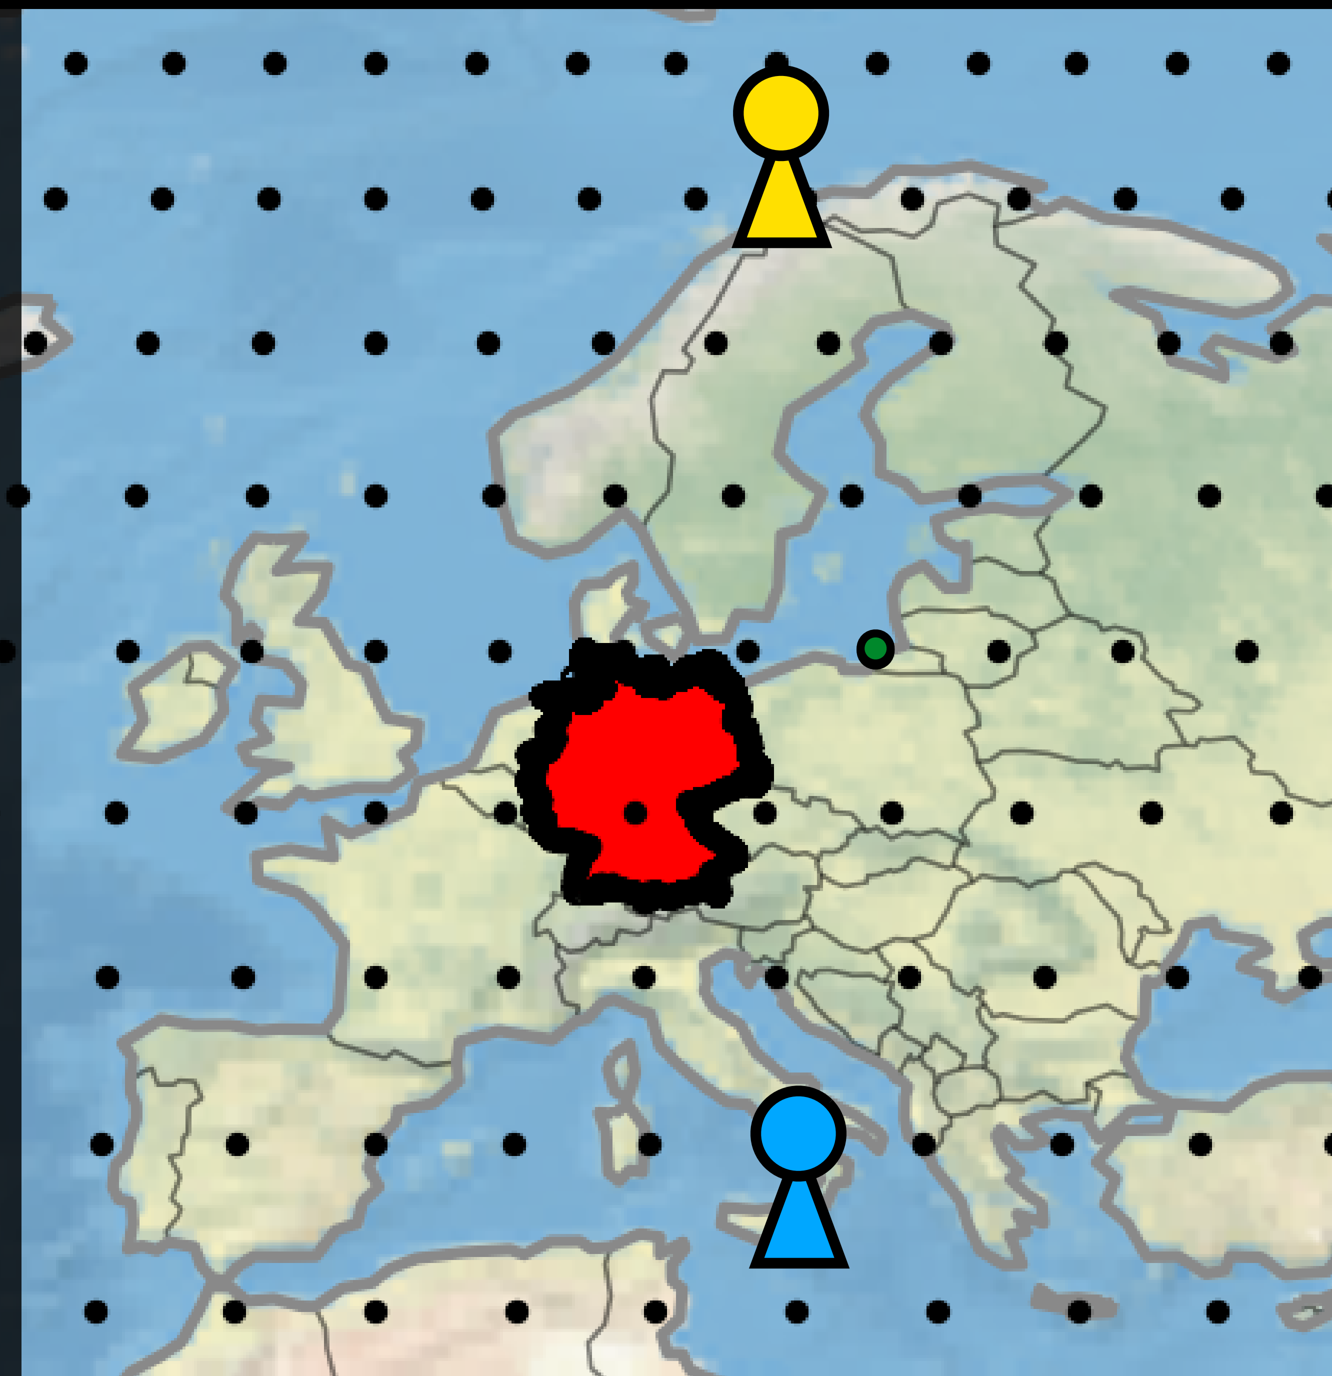

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

Segment the world into a grid

Include a grid point if:

$$(1 + \delta) \ast \text{Min } \textit{possible} \text{ RTT} \leq 3\, d / c$$

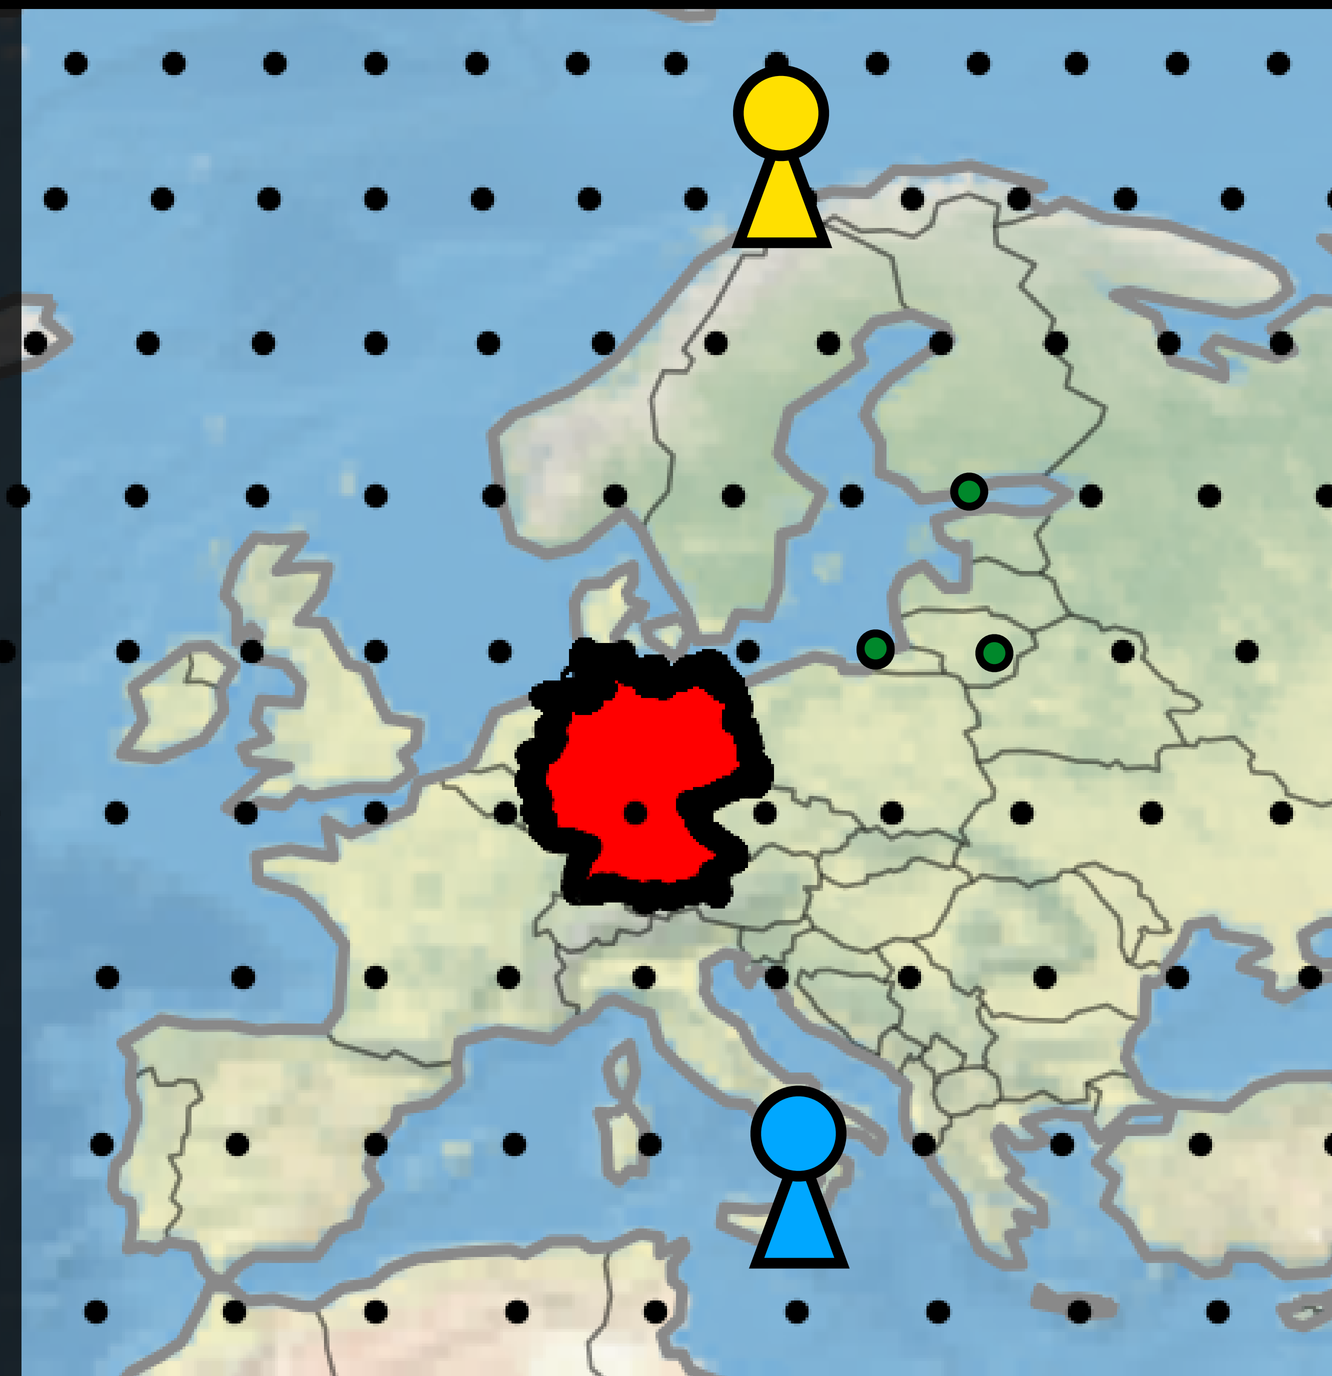

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

Segment the world into a grid

Include a grid point if:

$$(1 + \delta) * \text{Min } \textit{possible} \text{ RTT} \leq 3\, d\, /\, c$$

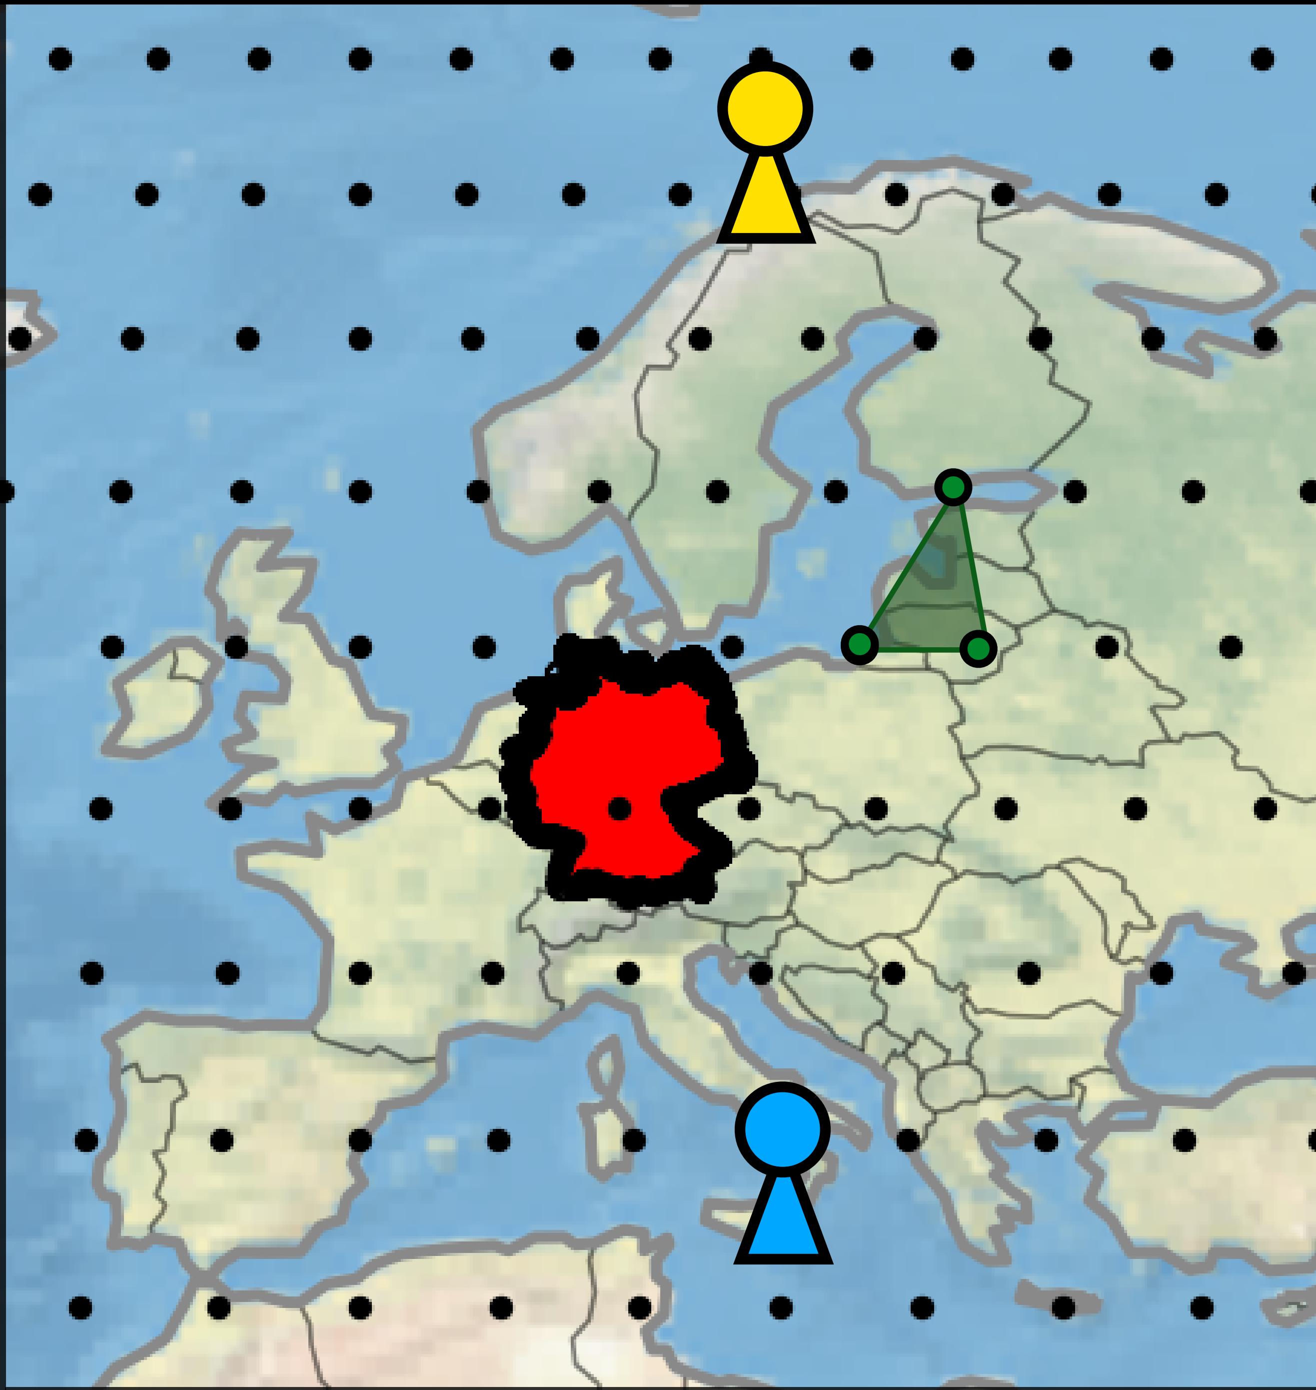

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

Segment the world into a grid

Include a grid point if:

$$(1 + \delta) * \text{Min } \textit{possible} \text{ RTT} \leq 3\, d \,/\, c$$

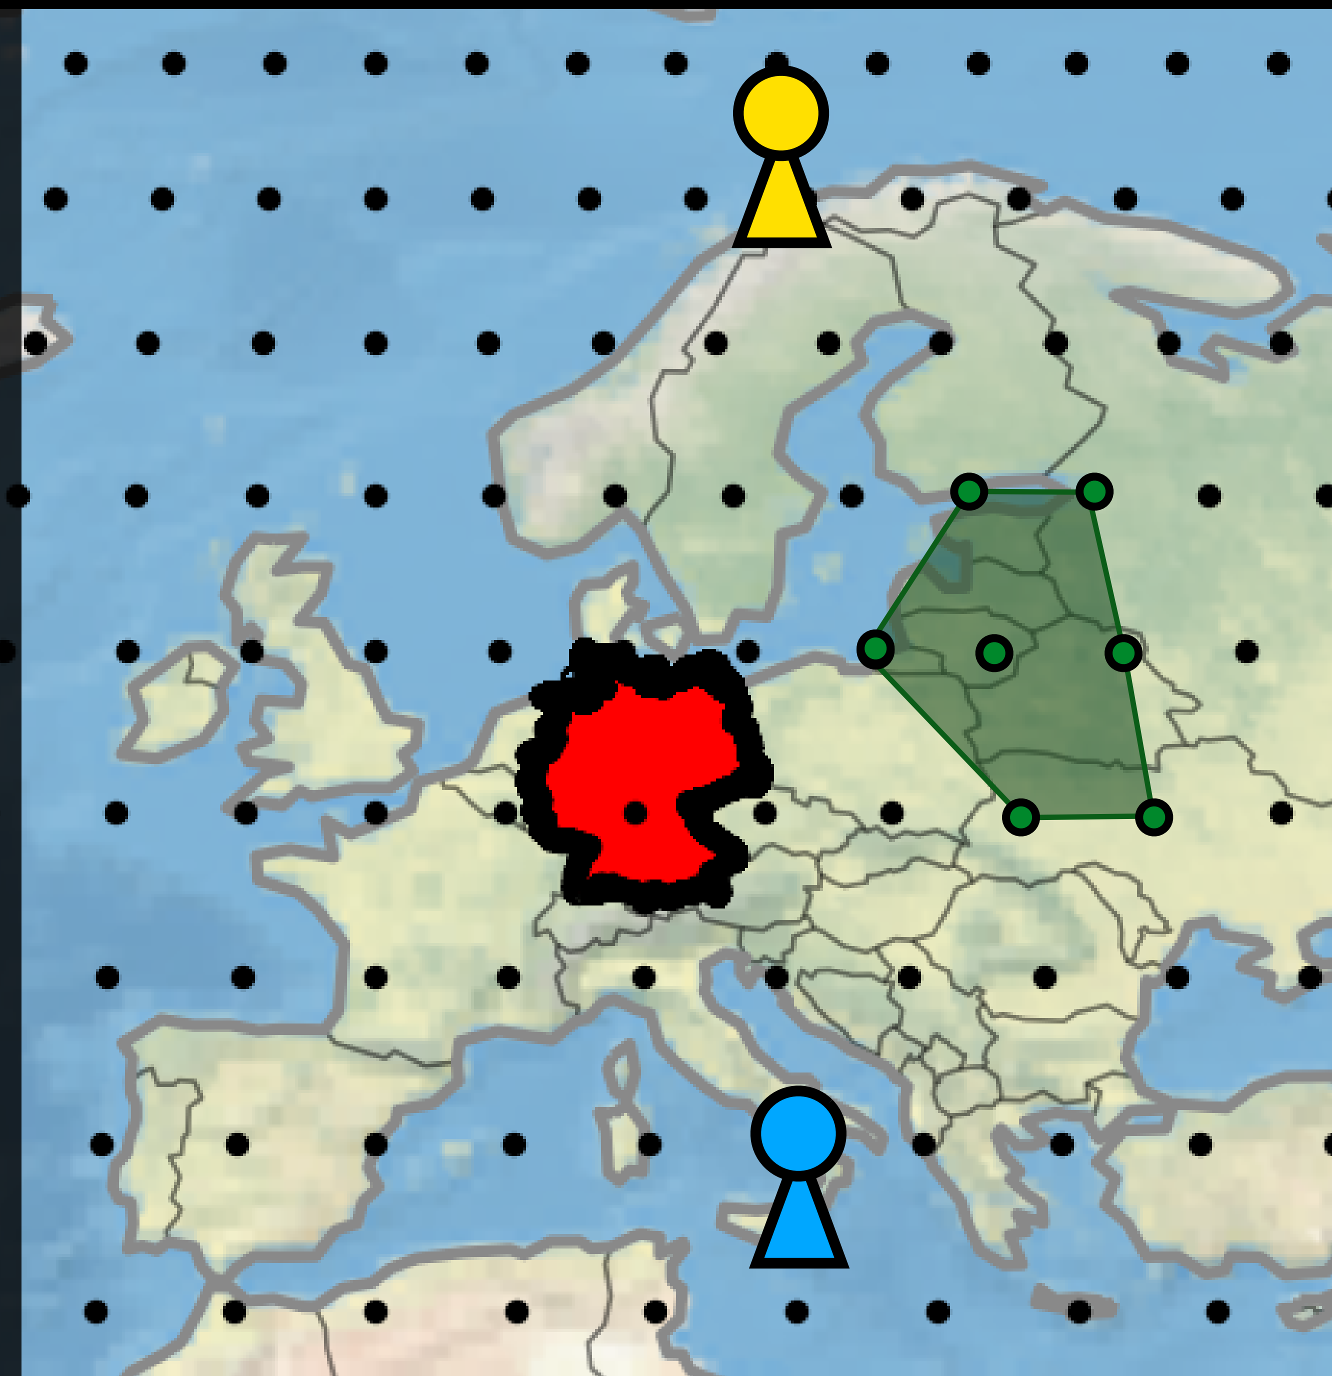

Compute target regions

Where alibis *might* be

Exclude locations where alibis cannot exist

Segment the world into a grid

Include a grid point if:

$$(1 + \delta) * \text{Min } \textit{possible} \text{ RTT} \leq 3\,d\,/\,c$$

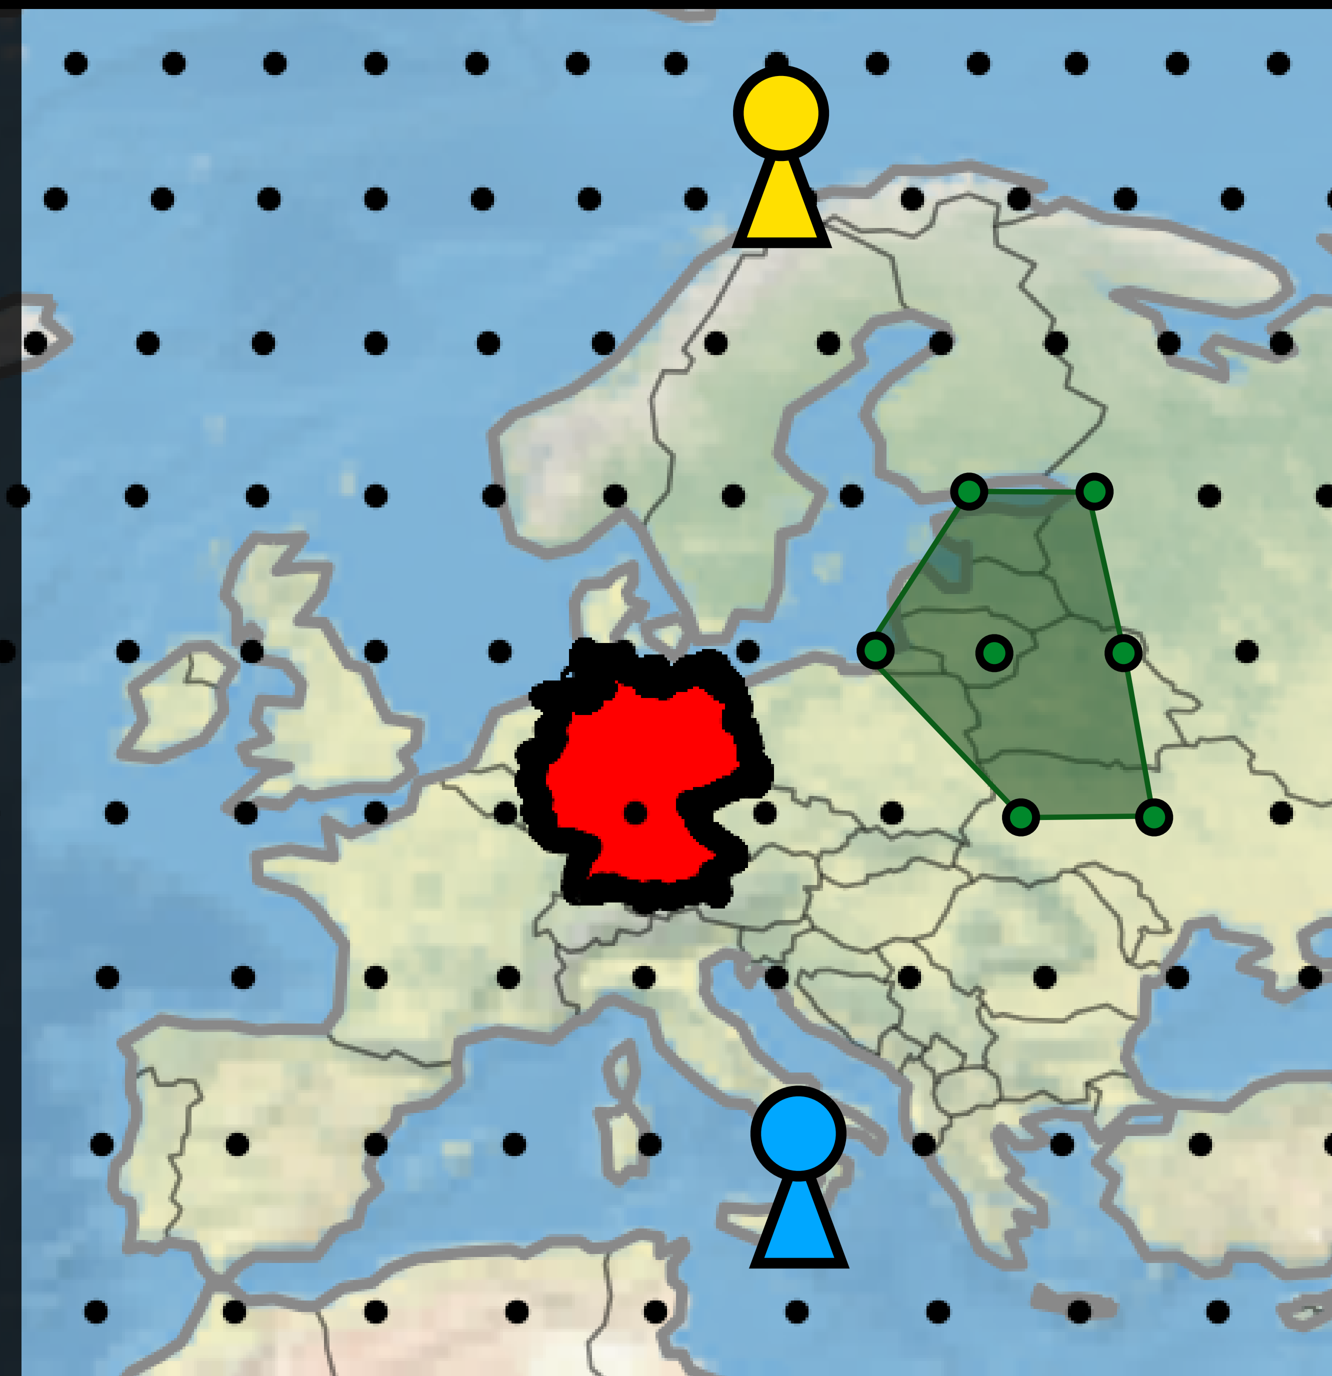

Compute target regions

Where alibis *might* be

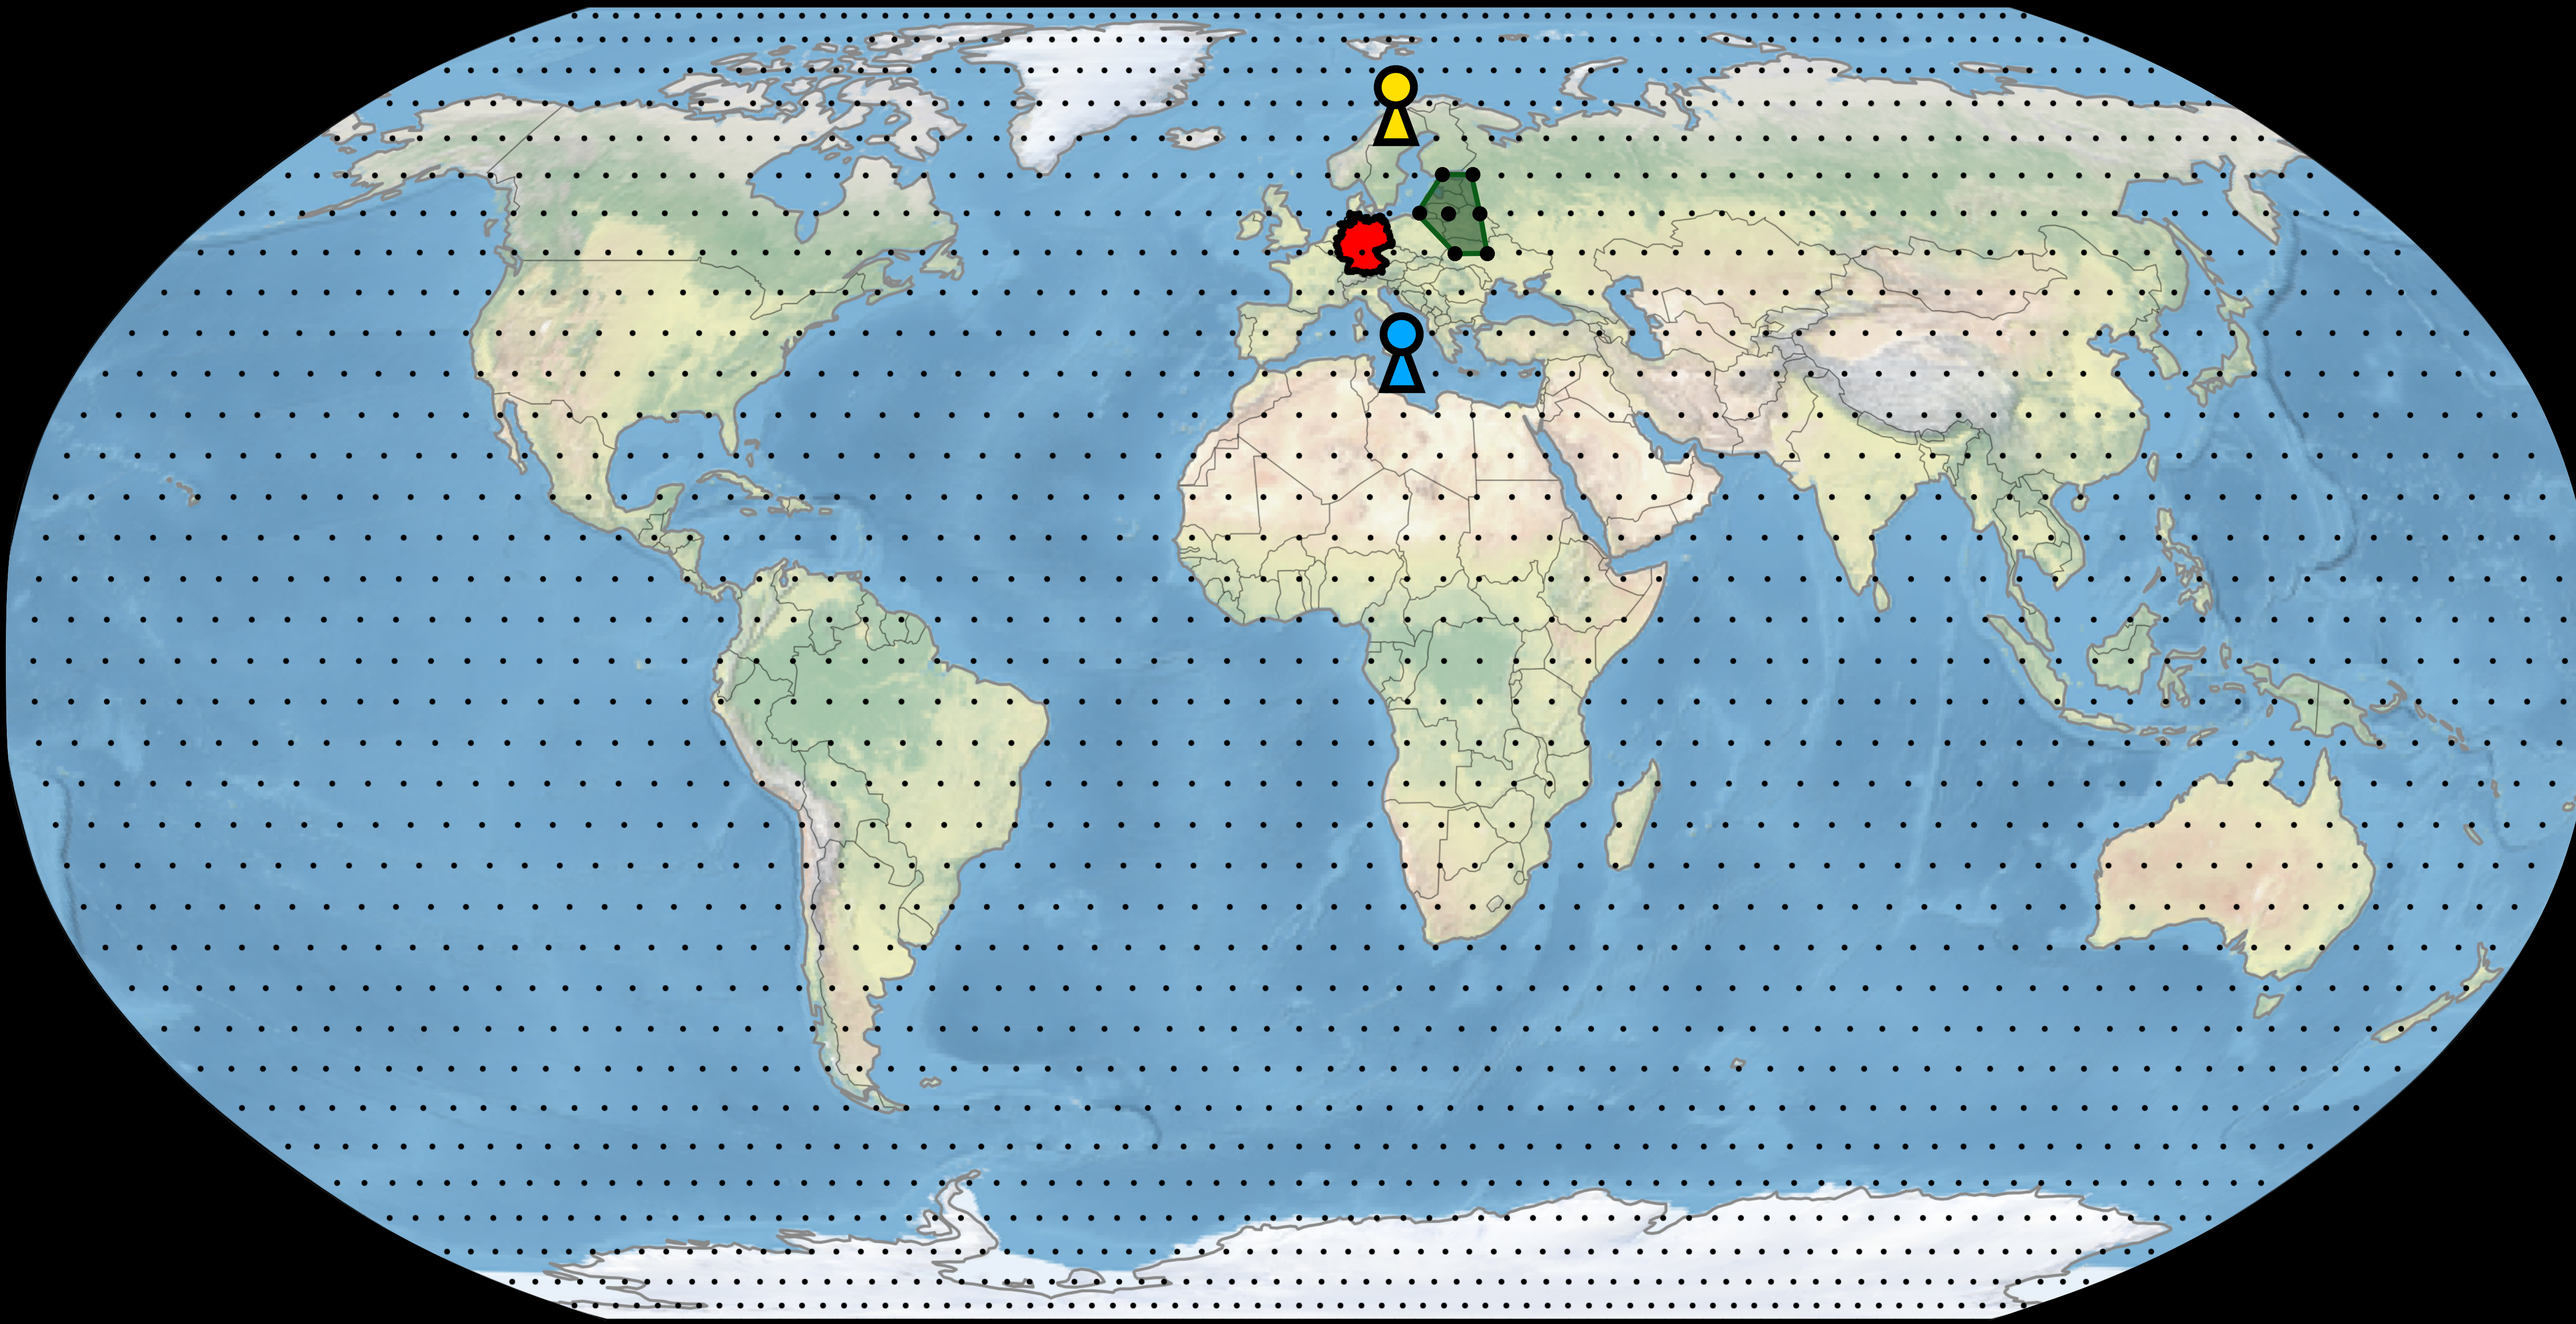

Compute target regions

Where alibis *might* be

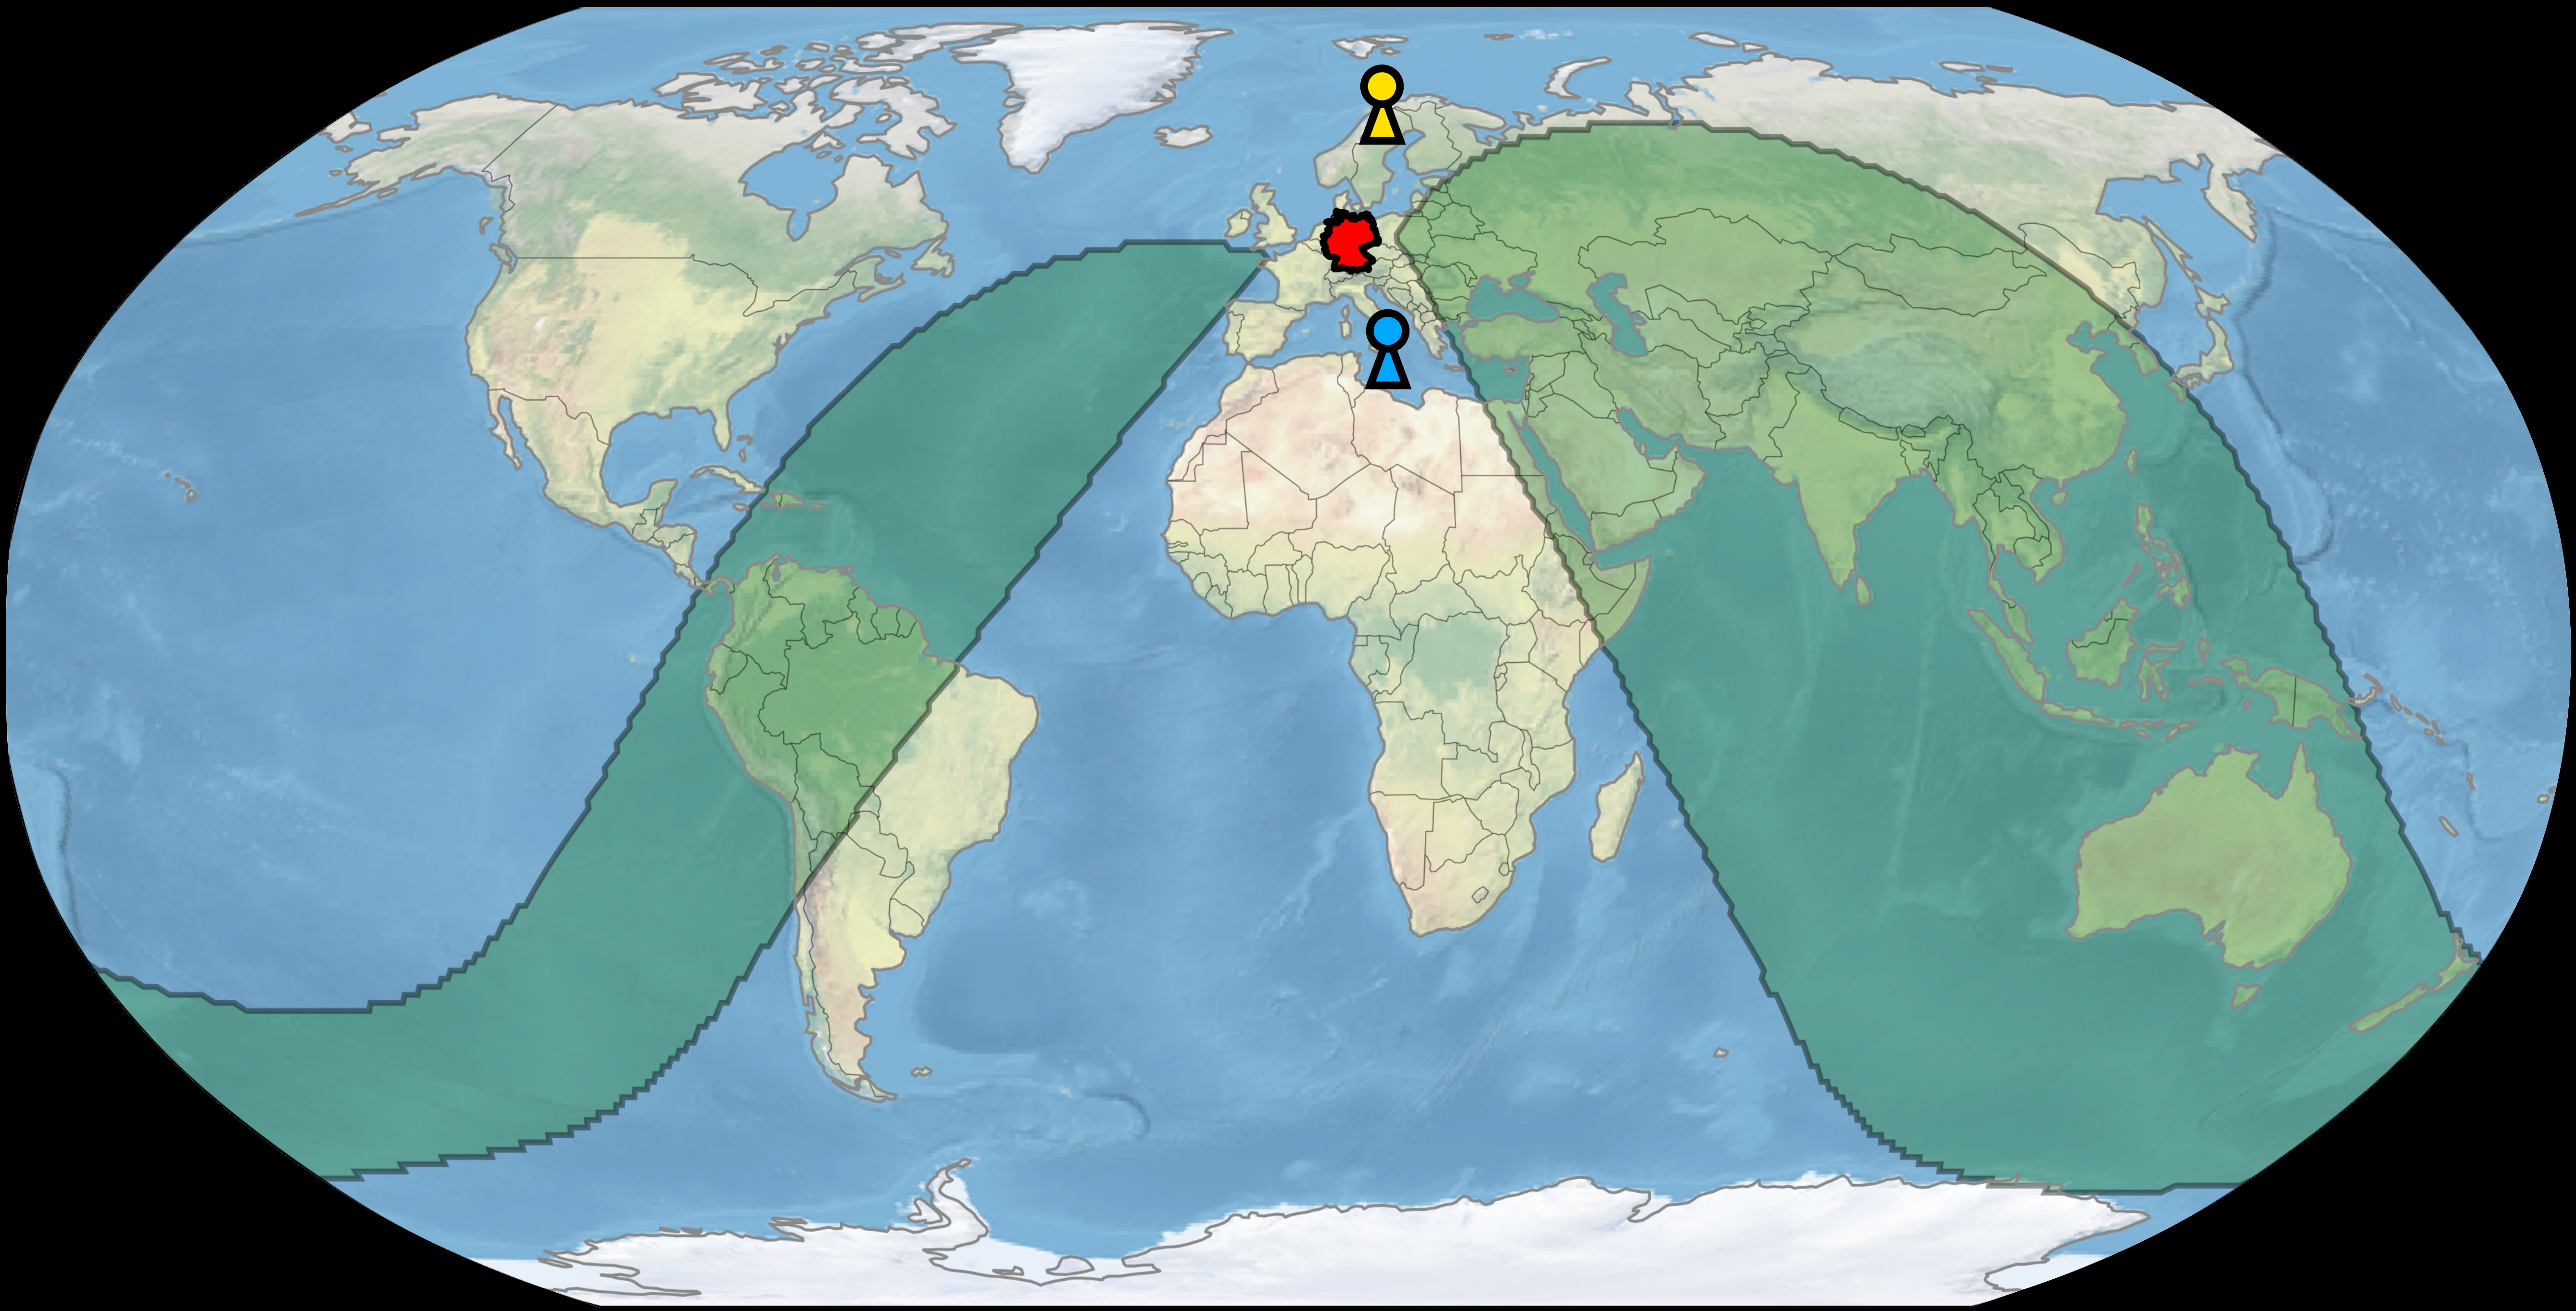

Compute target regions

Where alibis *might* be

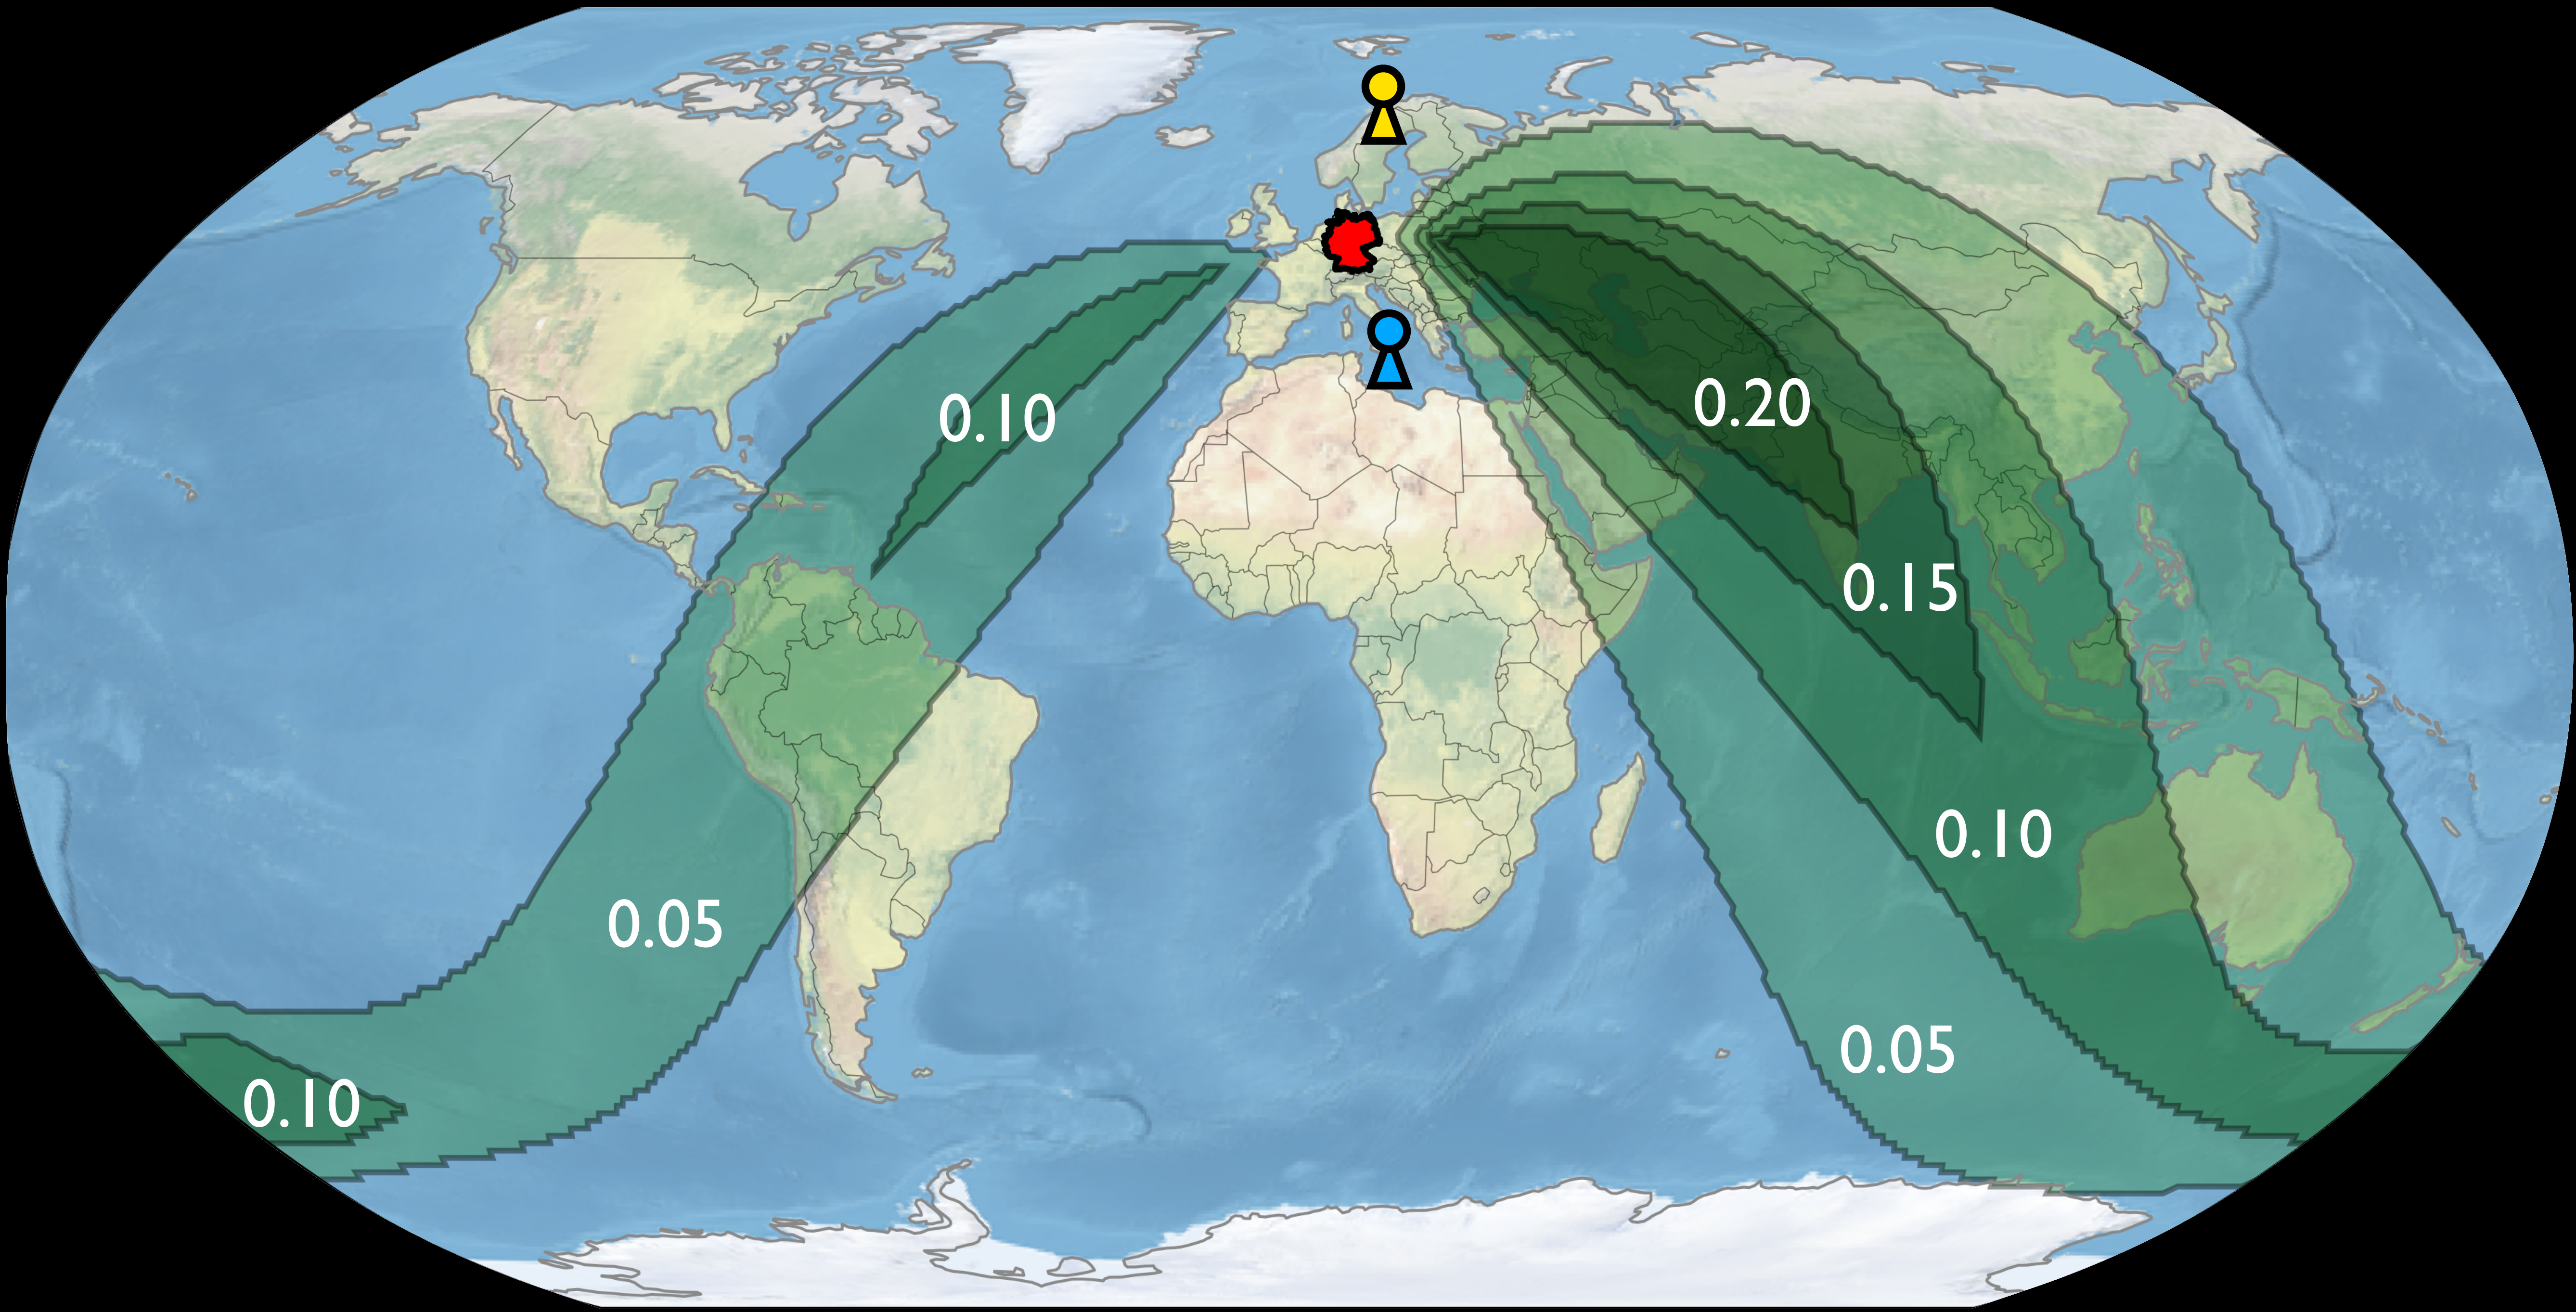

Compute target regions

Where alibis *might* be

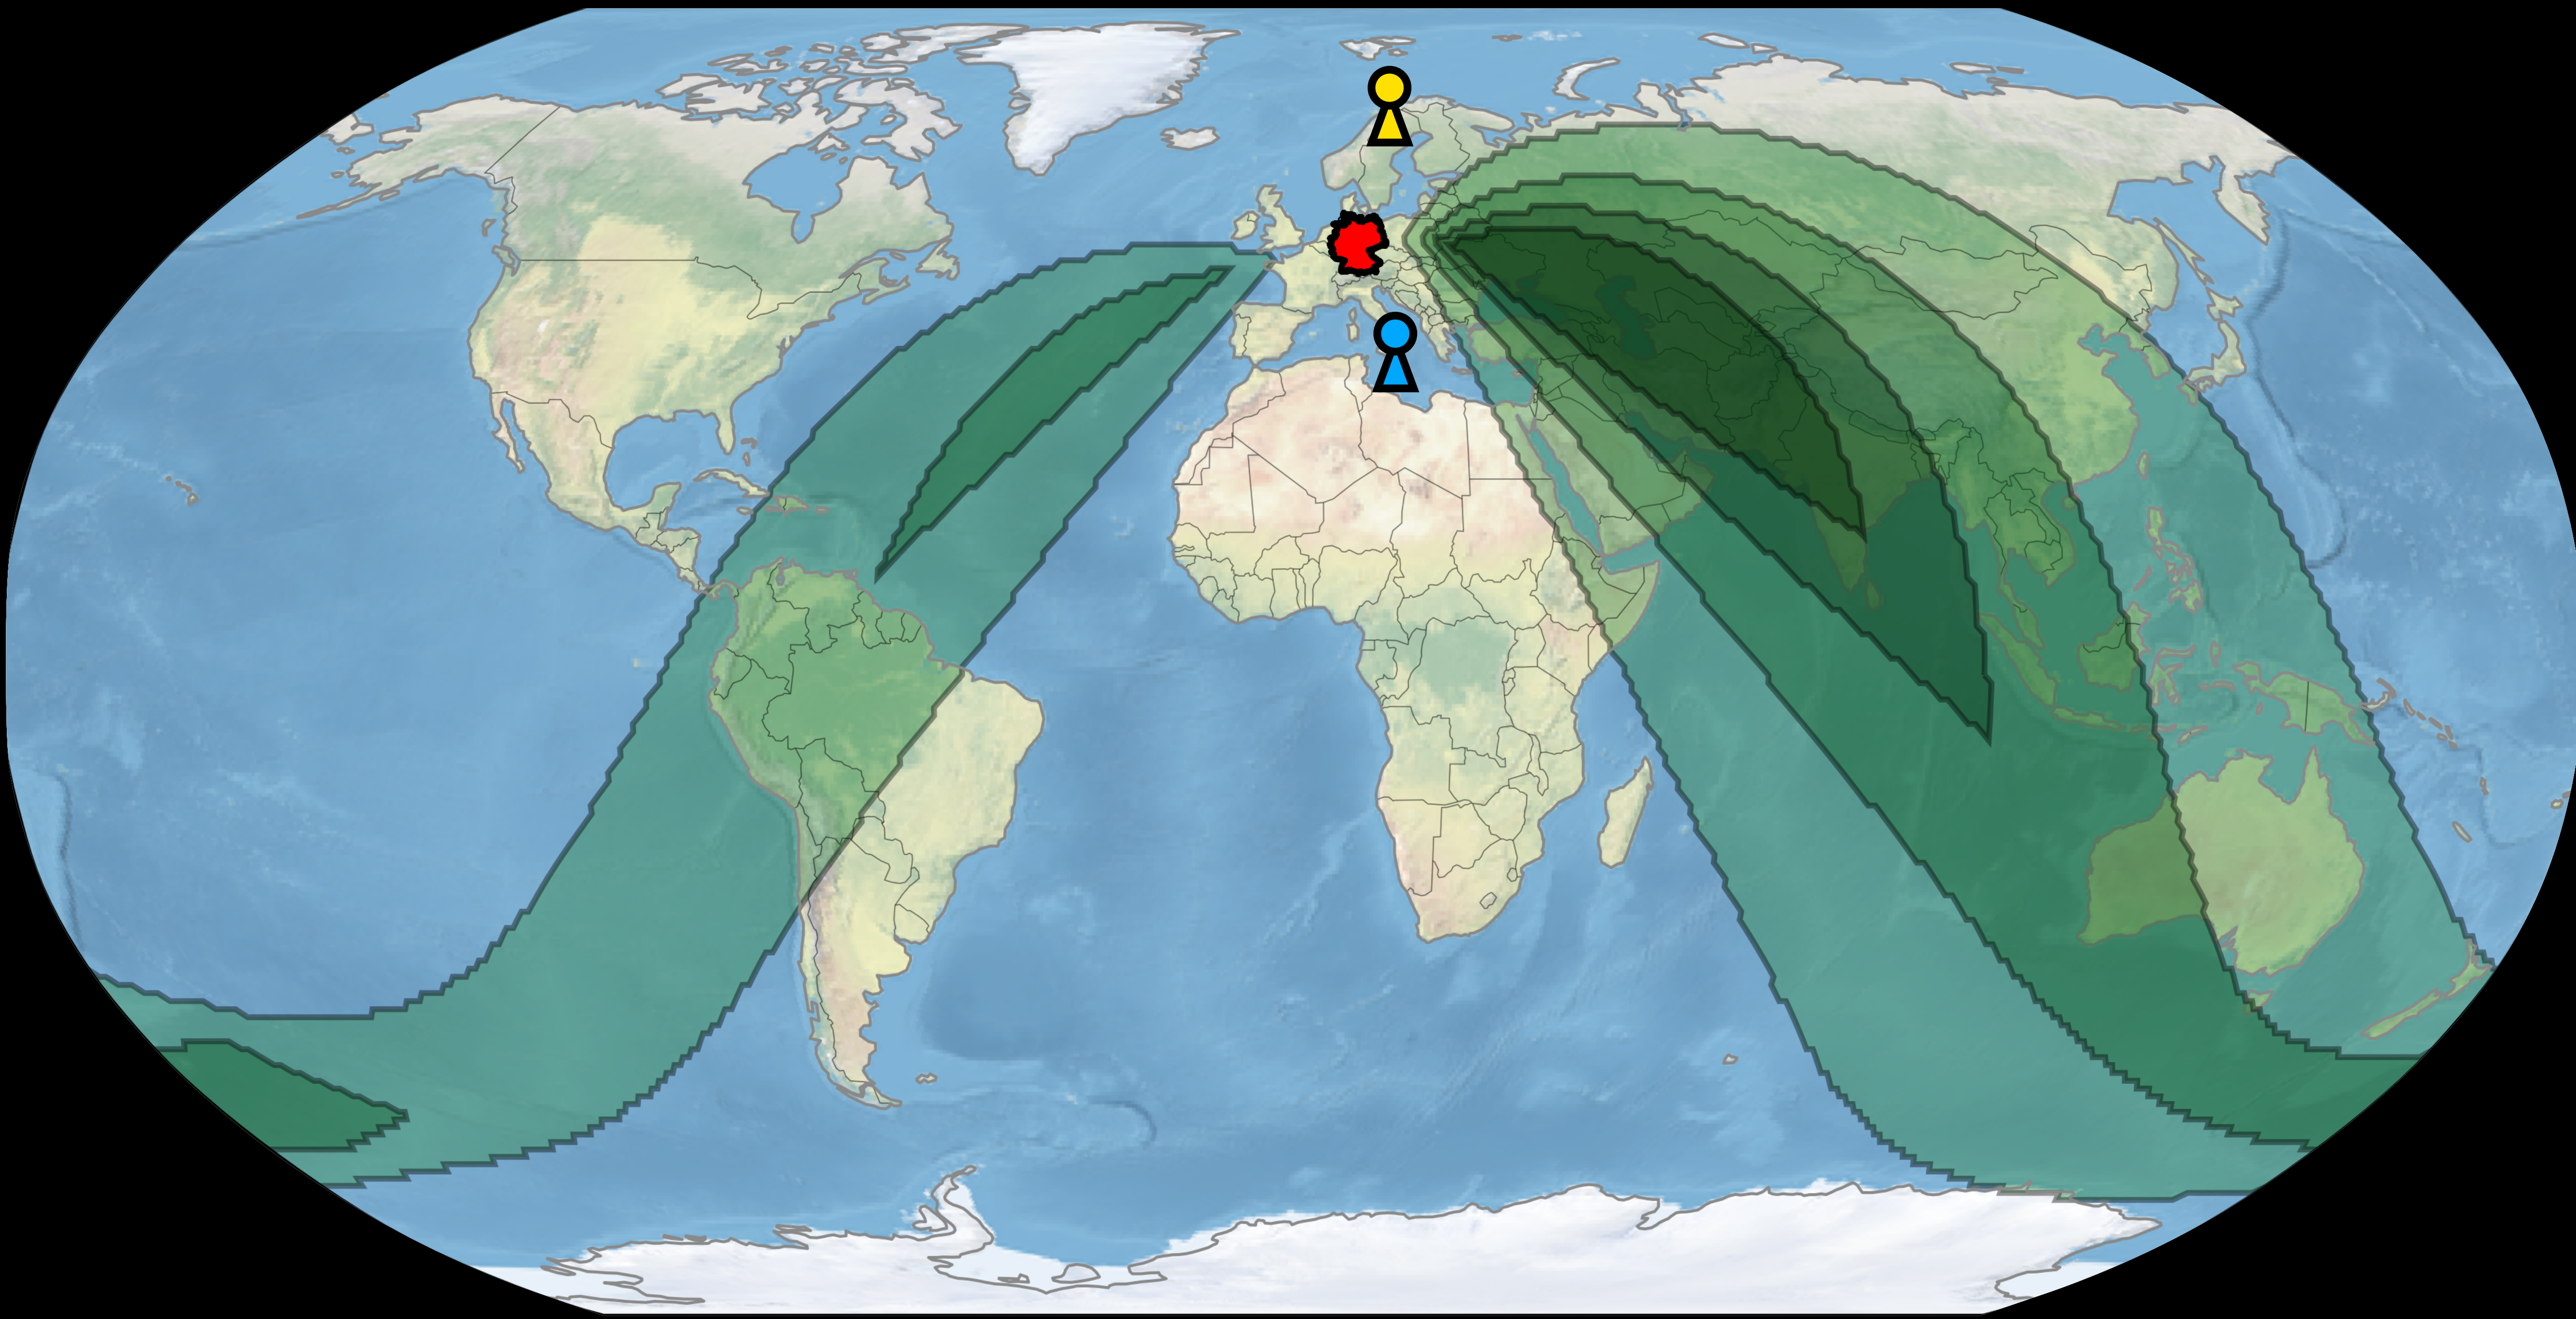

Being in a target region is a
necessary but not sufficient condition of an alibi

Compute target regions

Where alibis *might* be

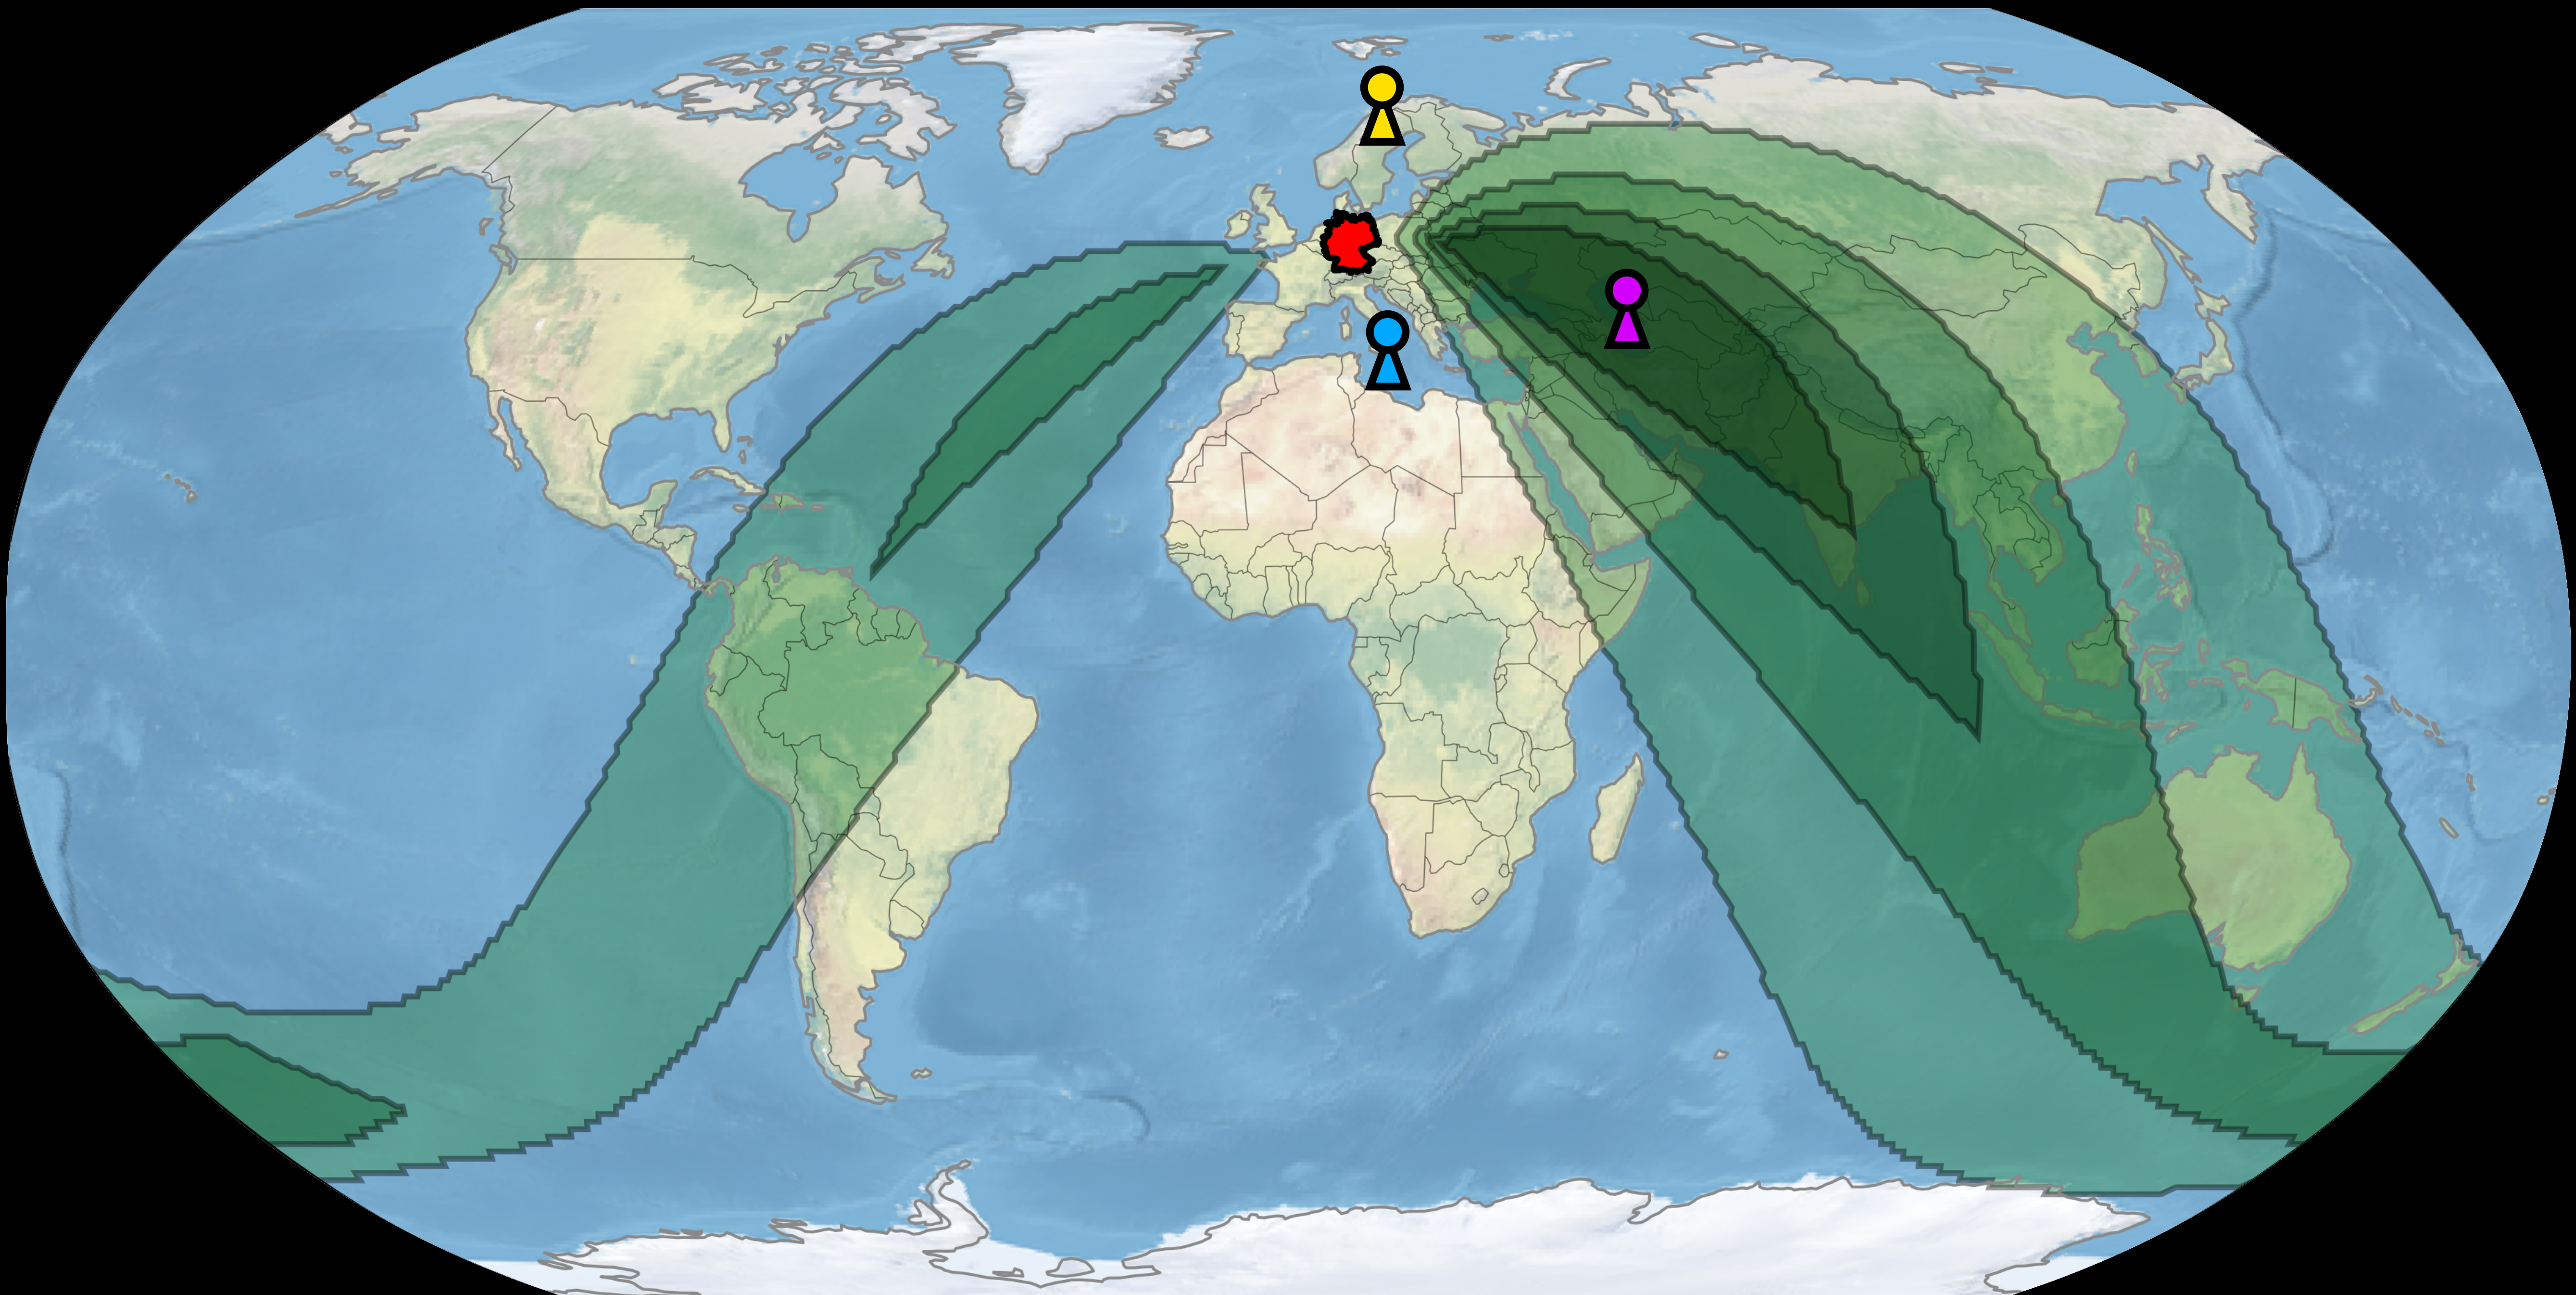

Being in a target region is a
necessary but not sufficient condition of an alibi

Compute target regions

Where alibis *might* be

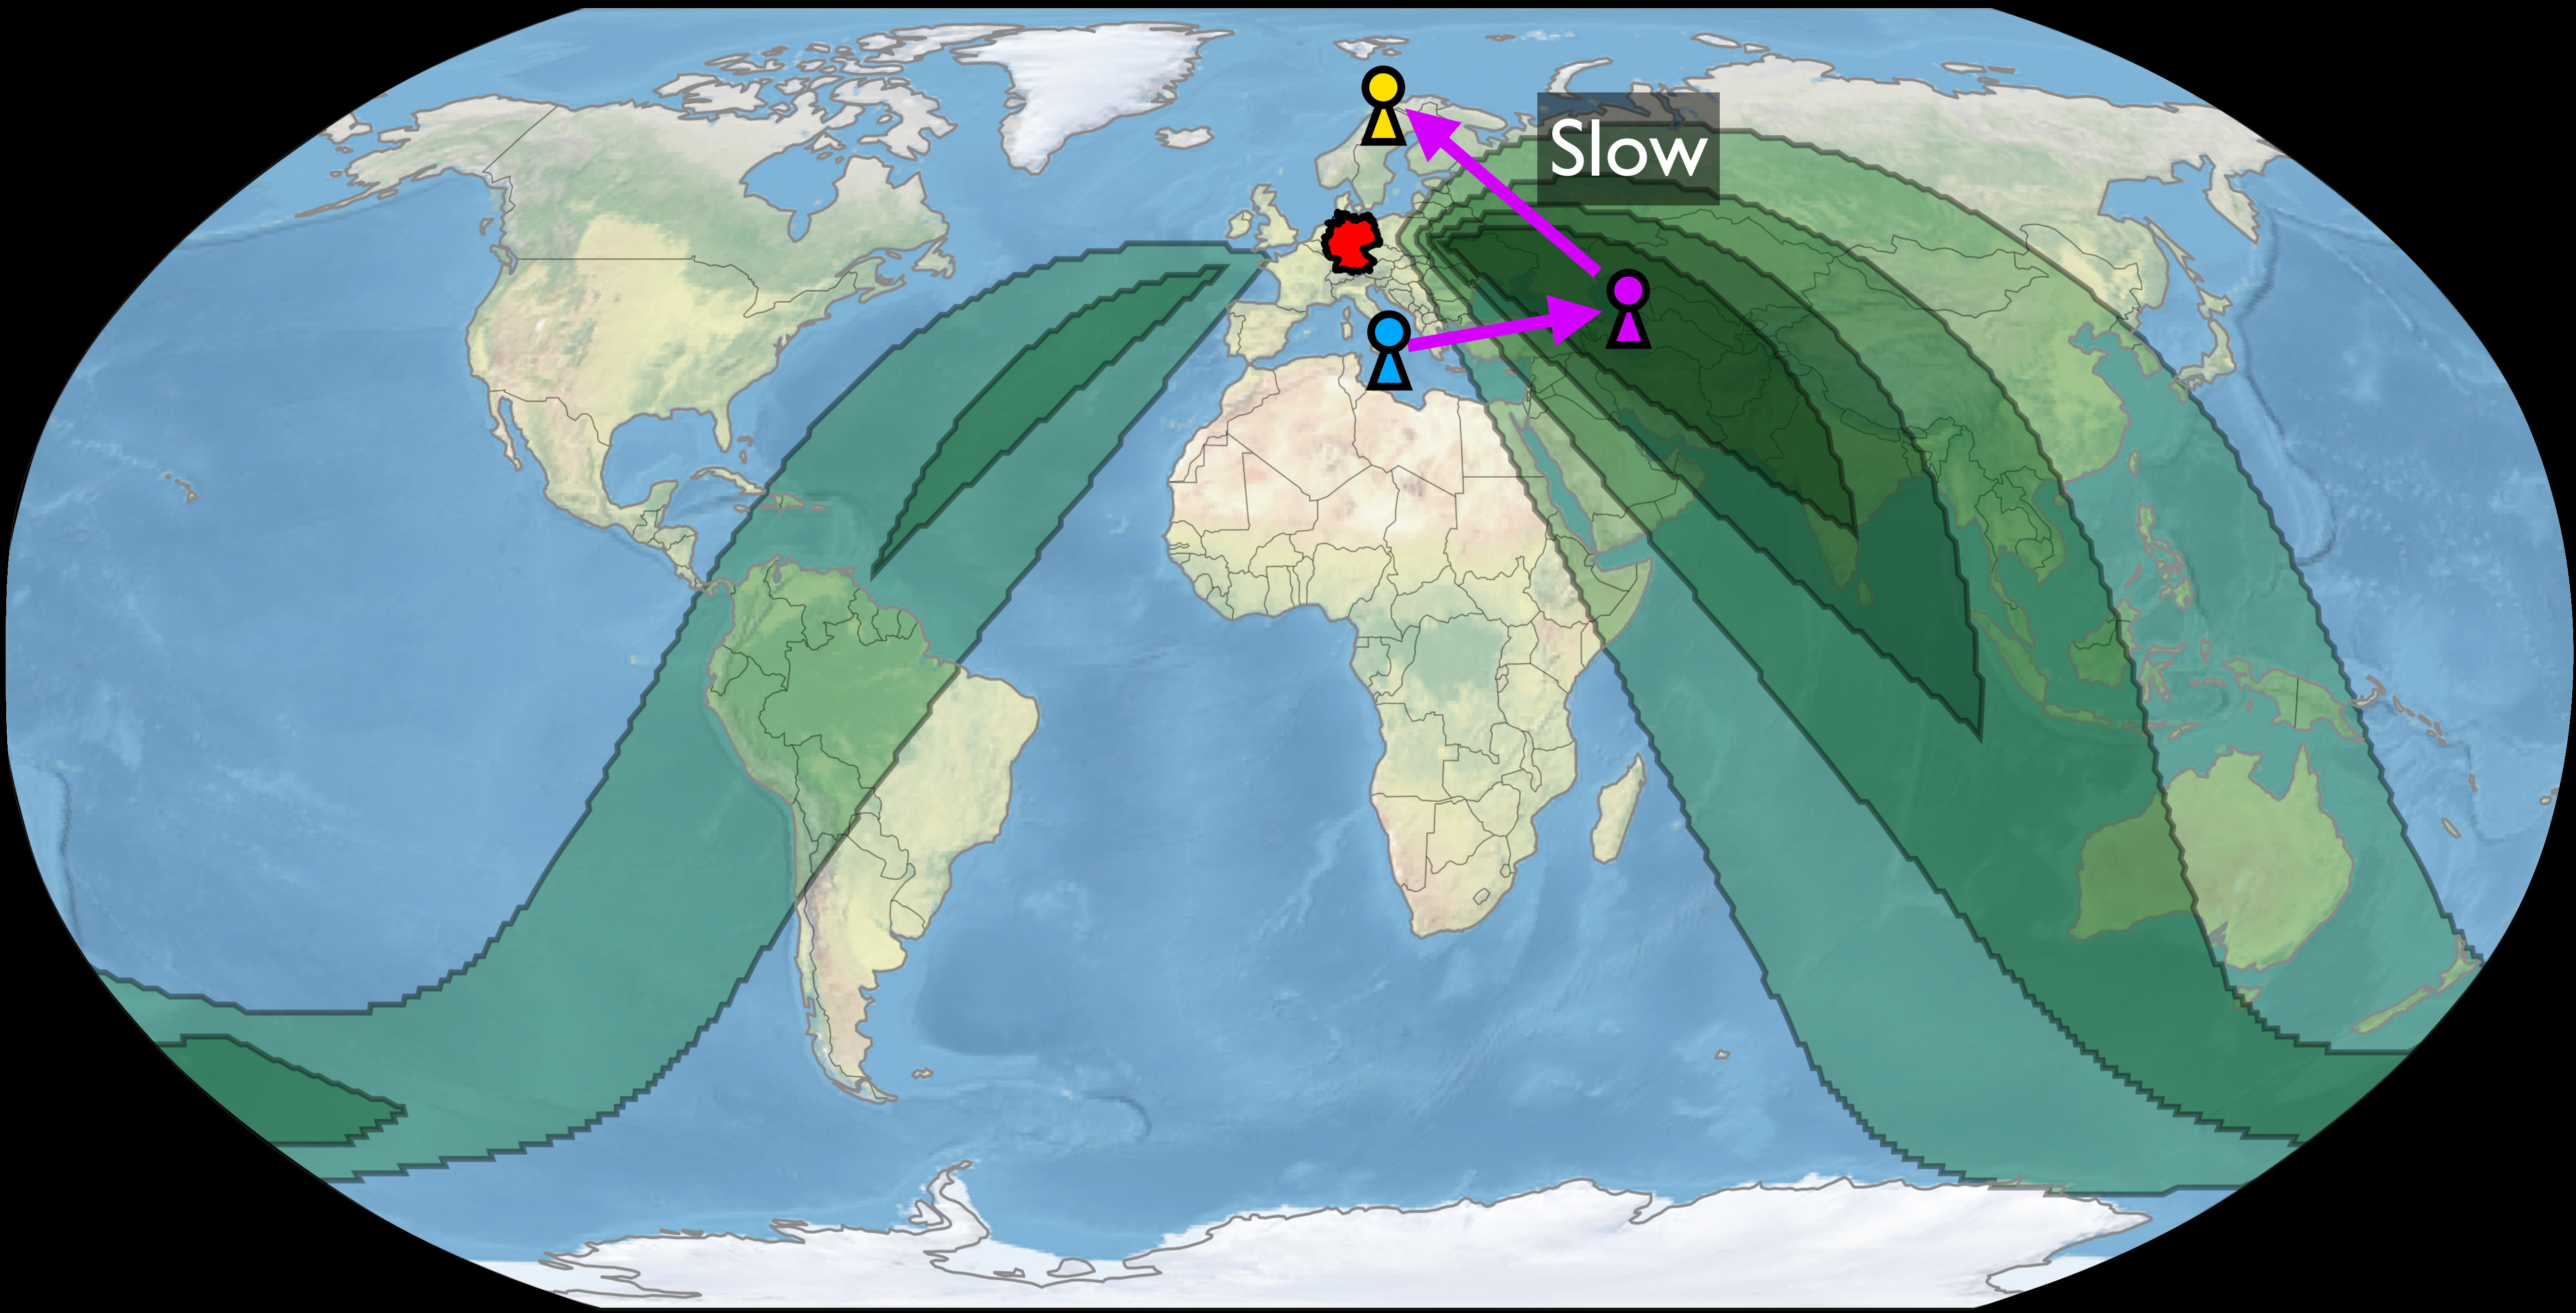

Being in a target region is a
necessary but not sufficient condition of an alibi

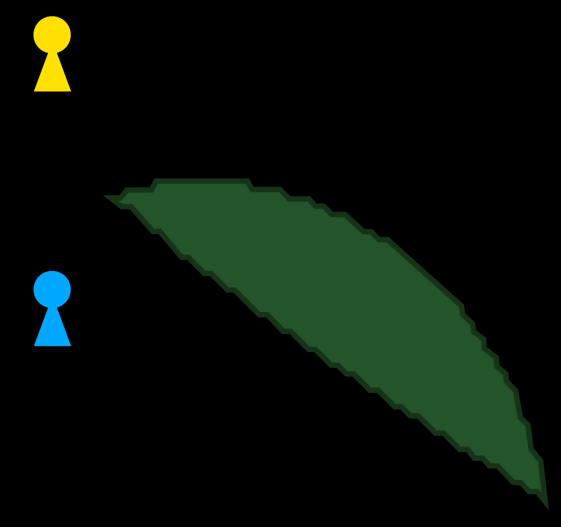

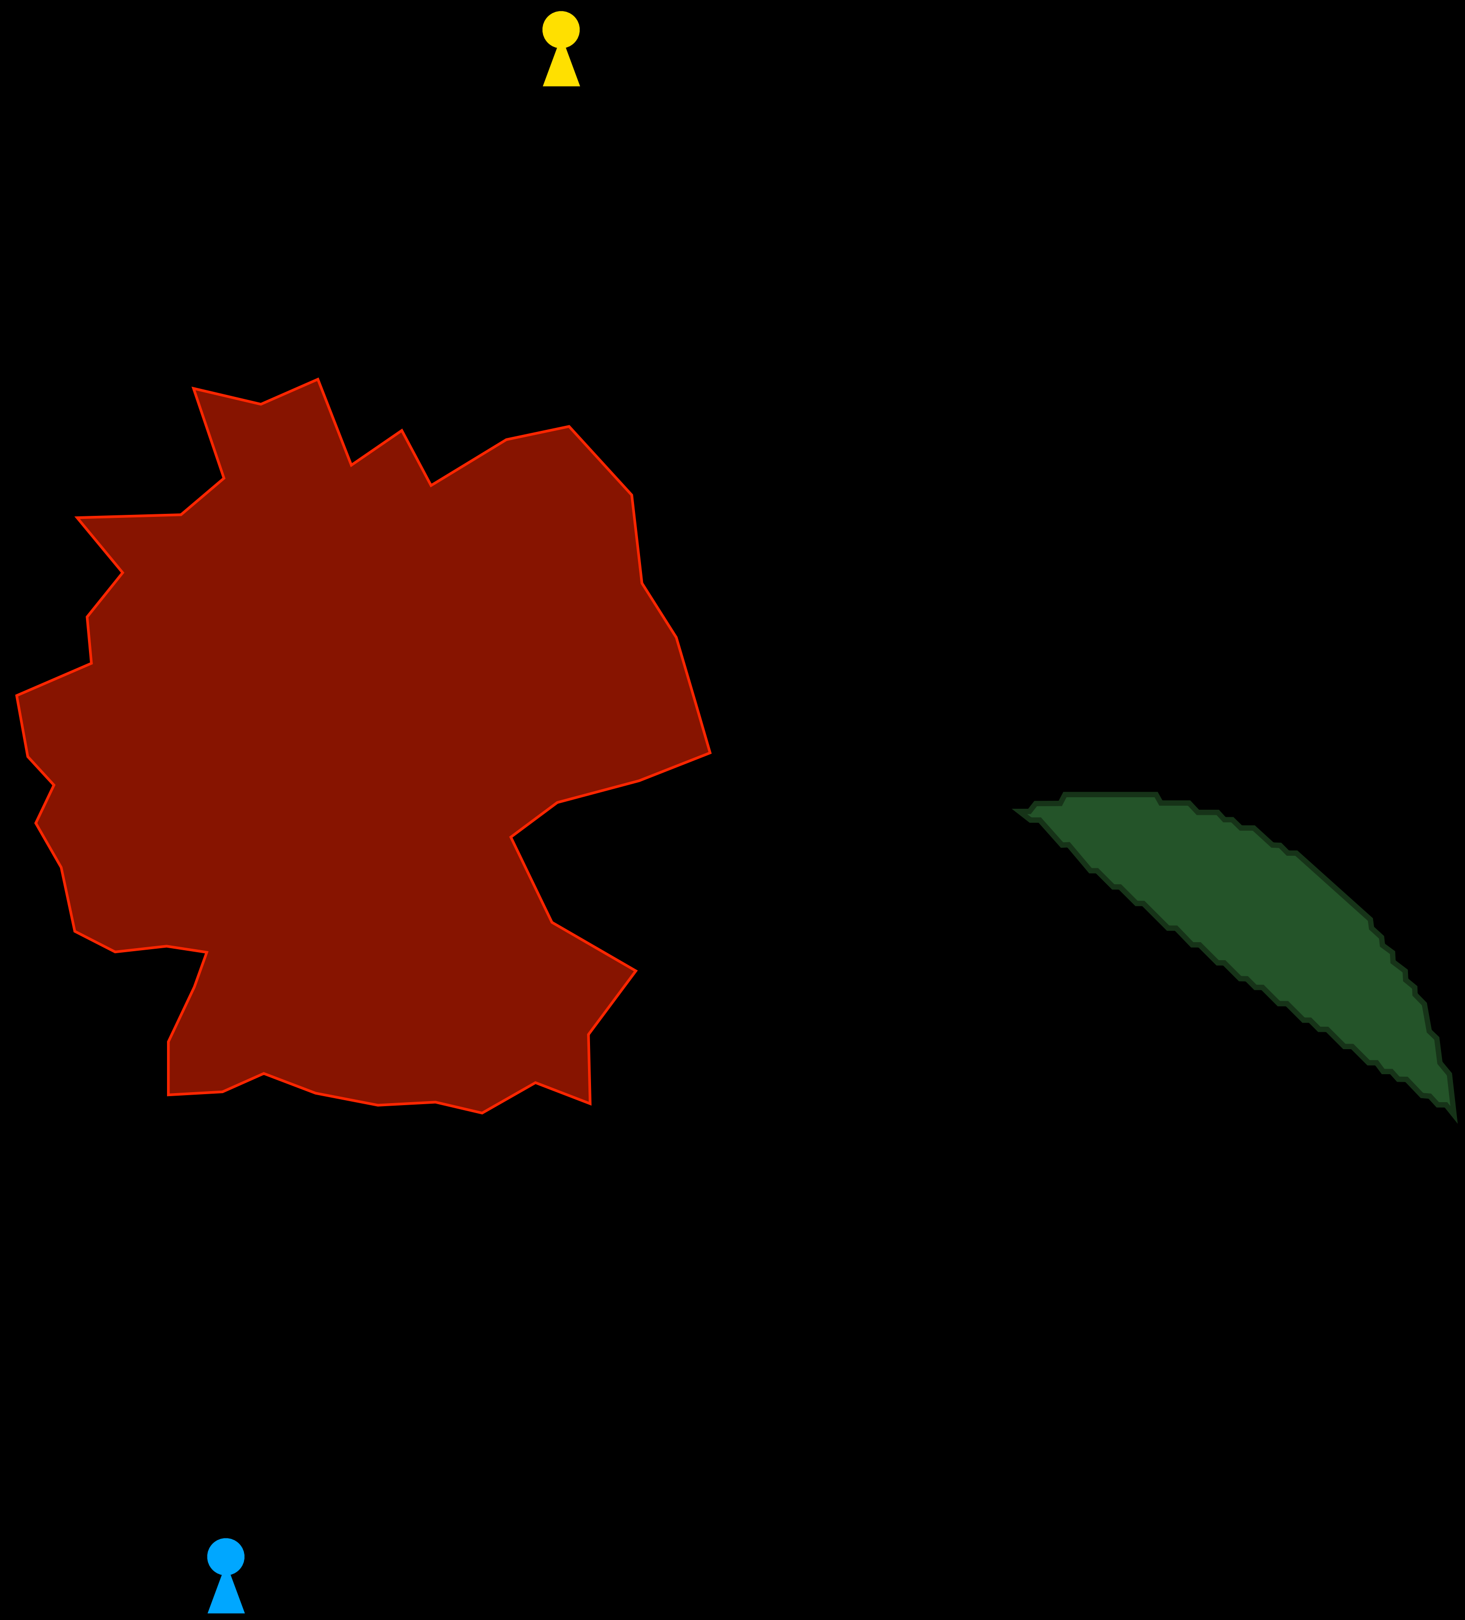

Peer-to-peer:

Every participant has a set of "neighbor" peers

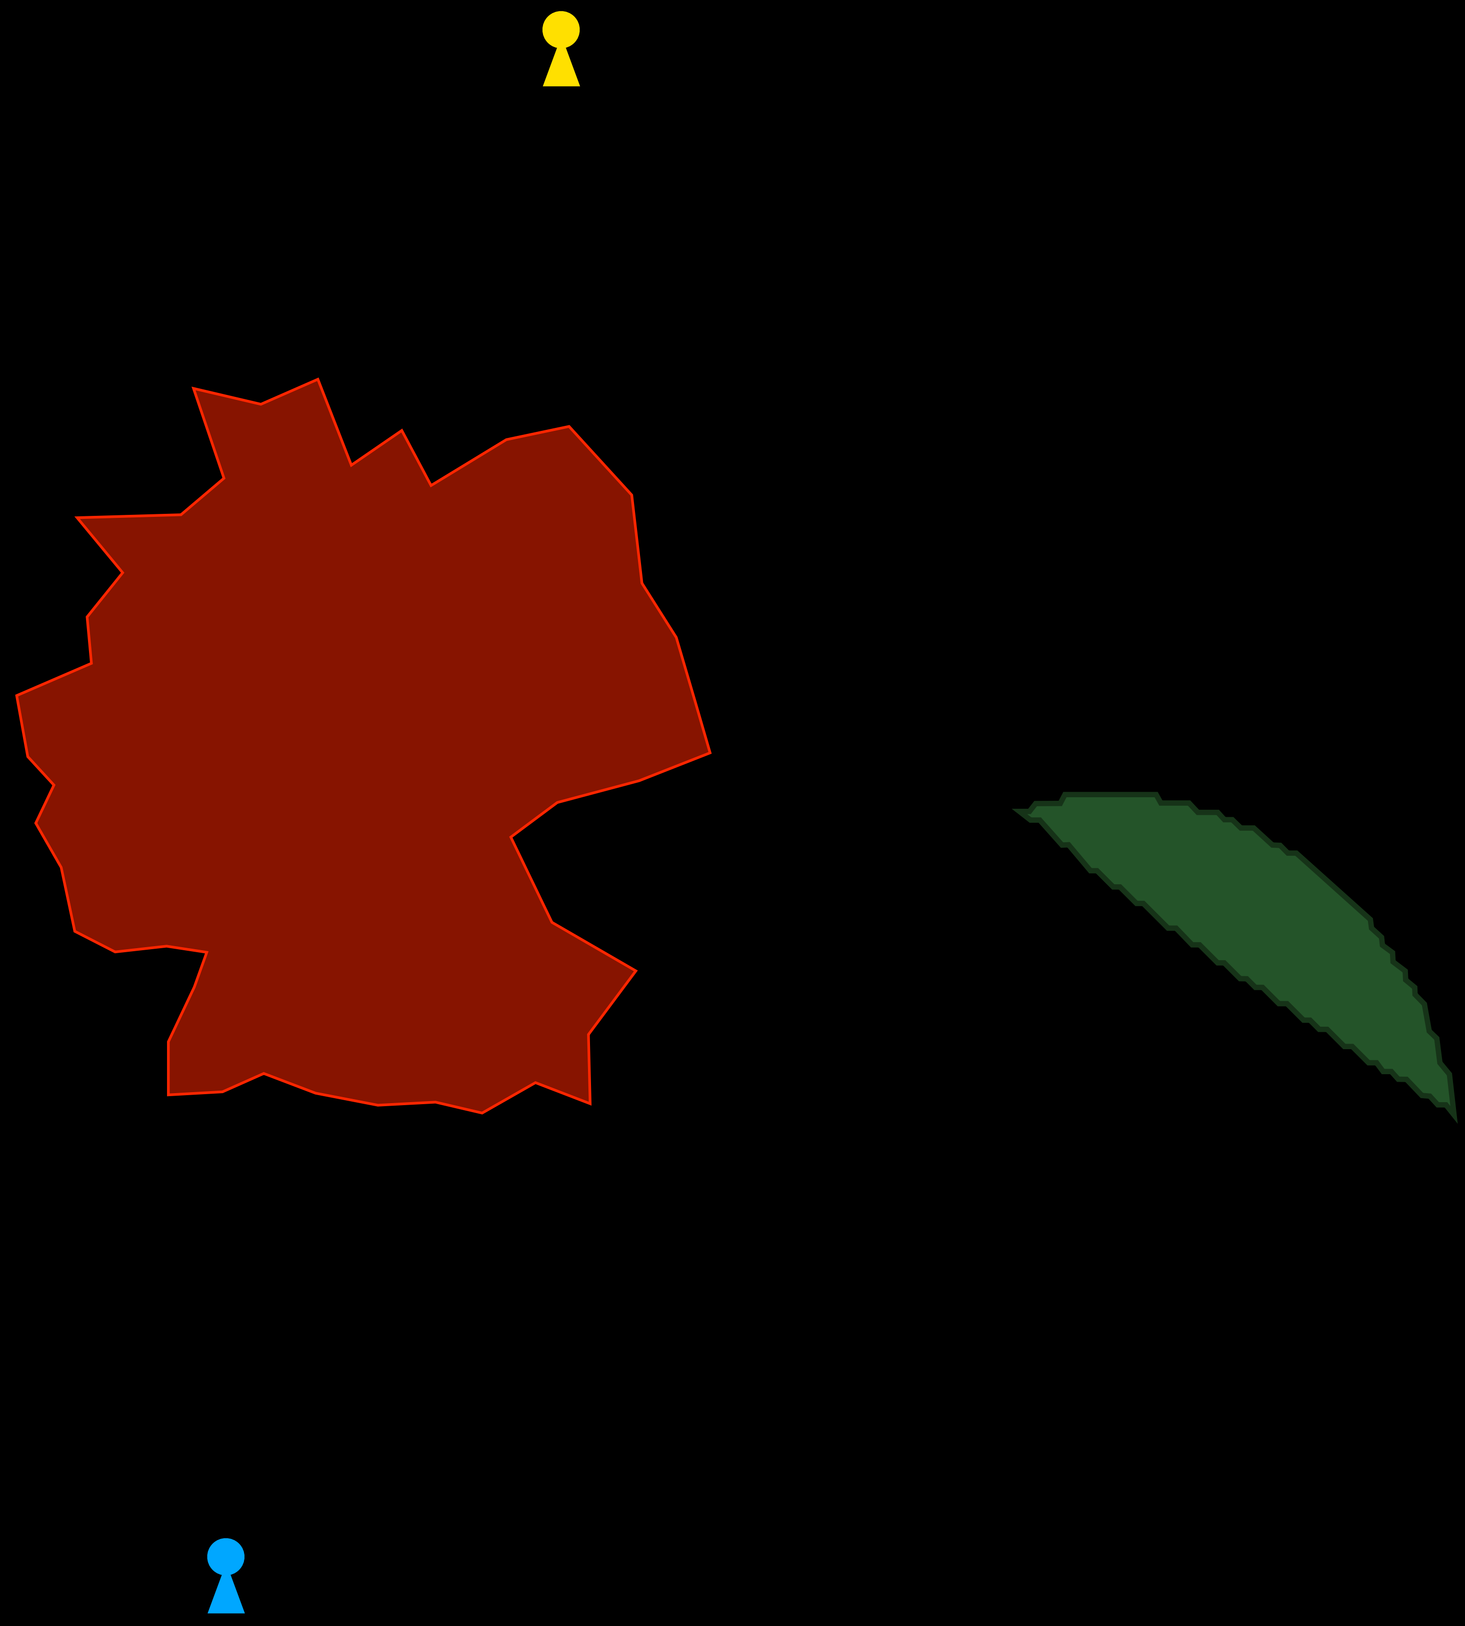

Safety:

Only forward to neighbors whom you *know* aren't in *F*

Safety:

Only forward to neighbors whom you *know* aren't in *F*

Safety:

Only forward to neighbors whom you *know* aren't in *F*

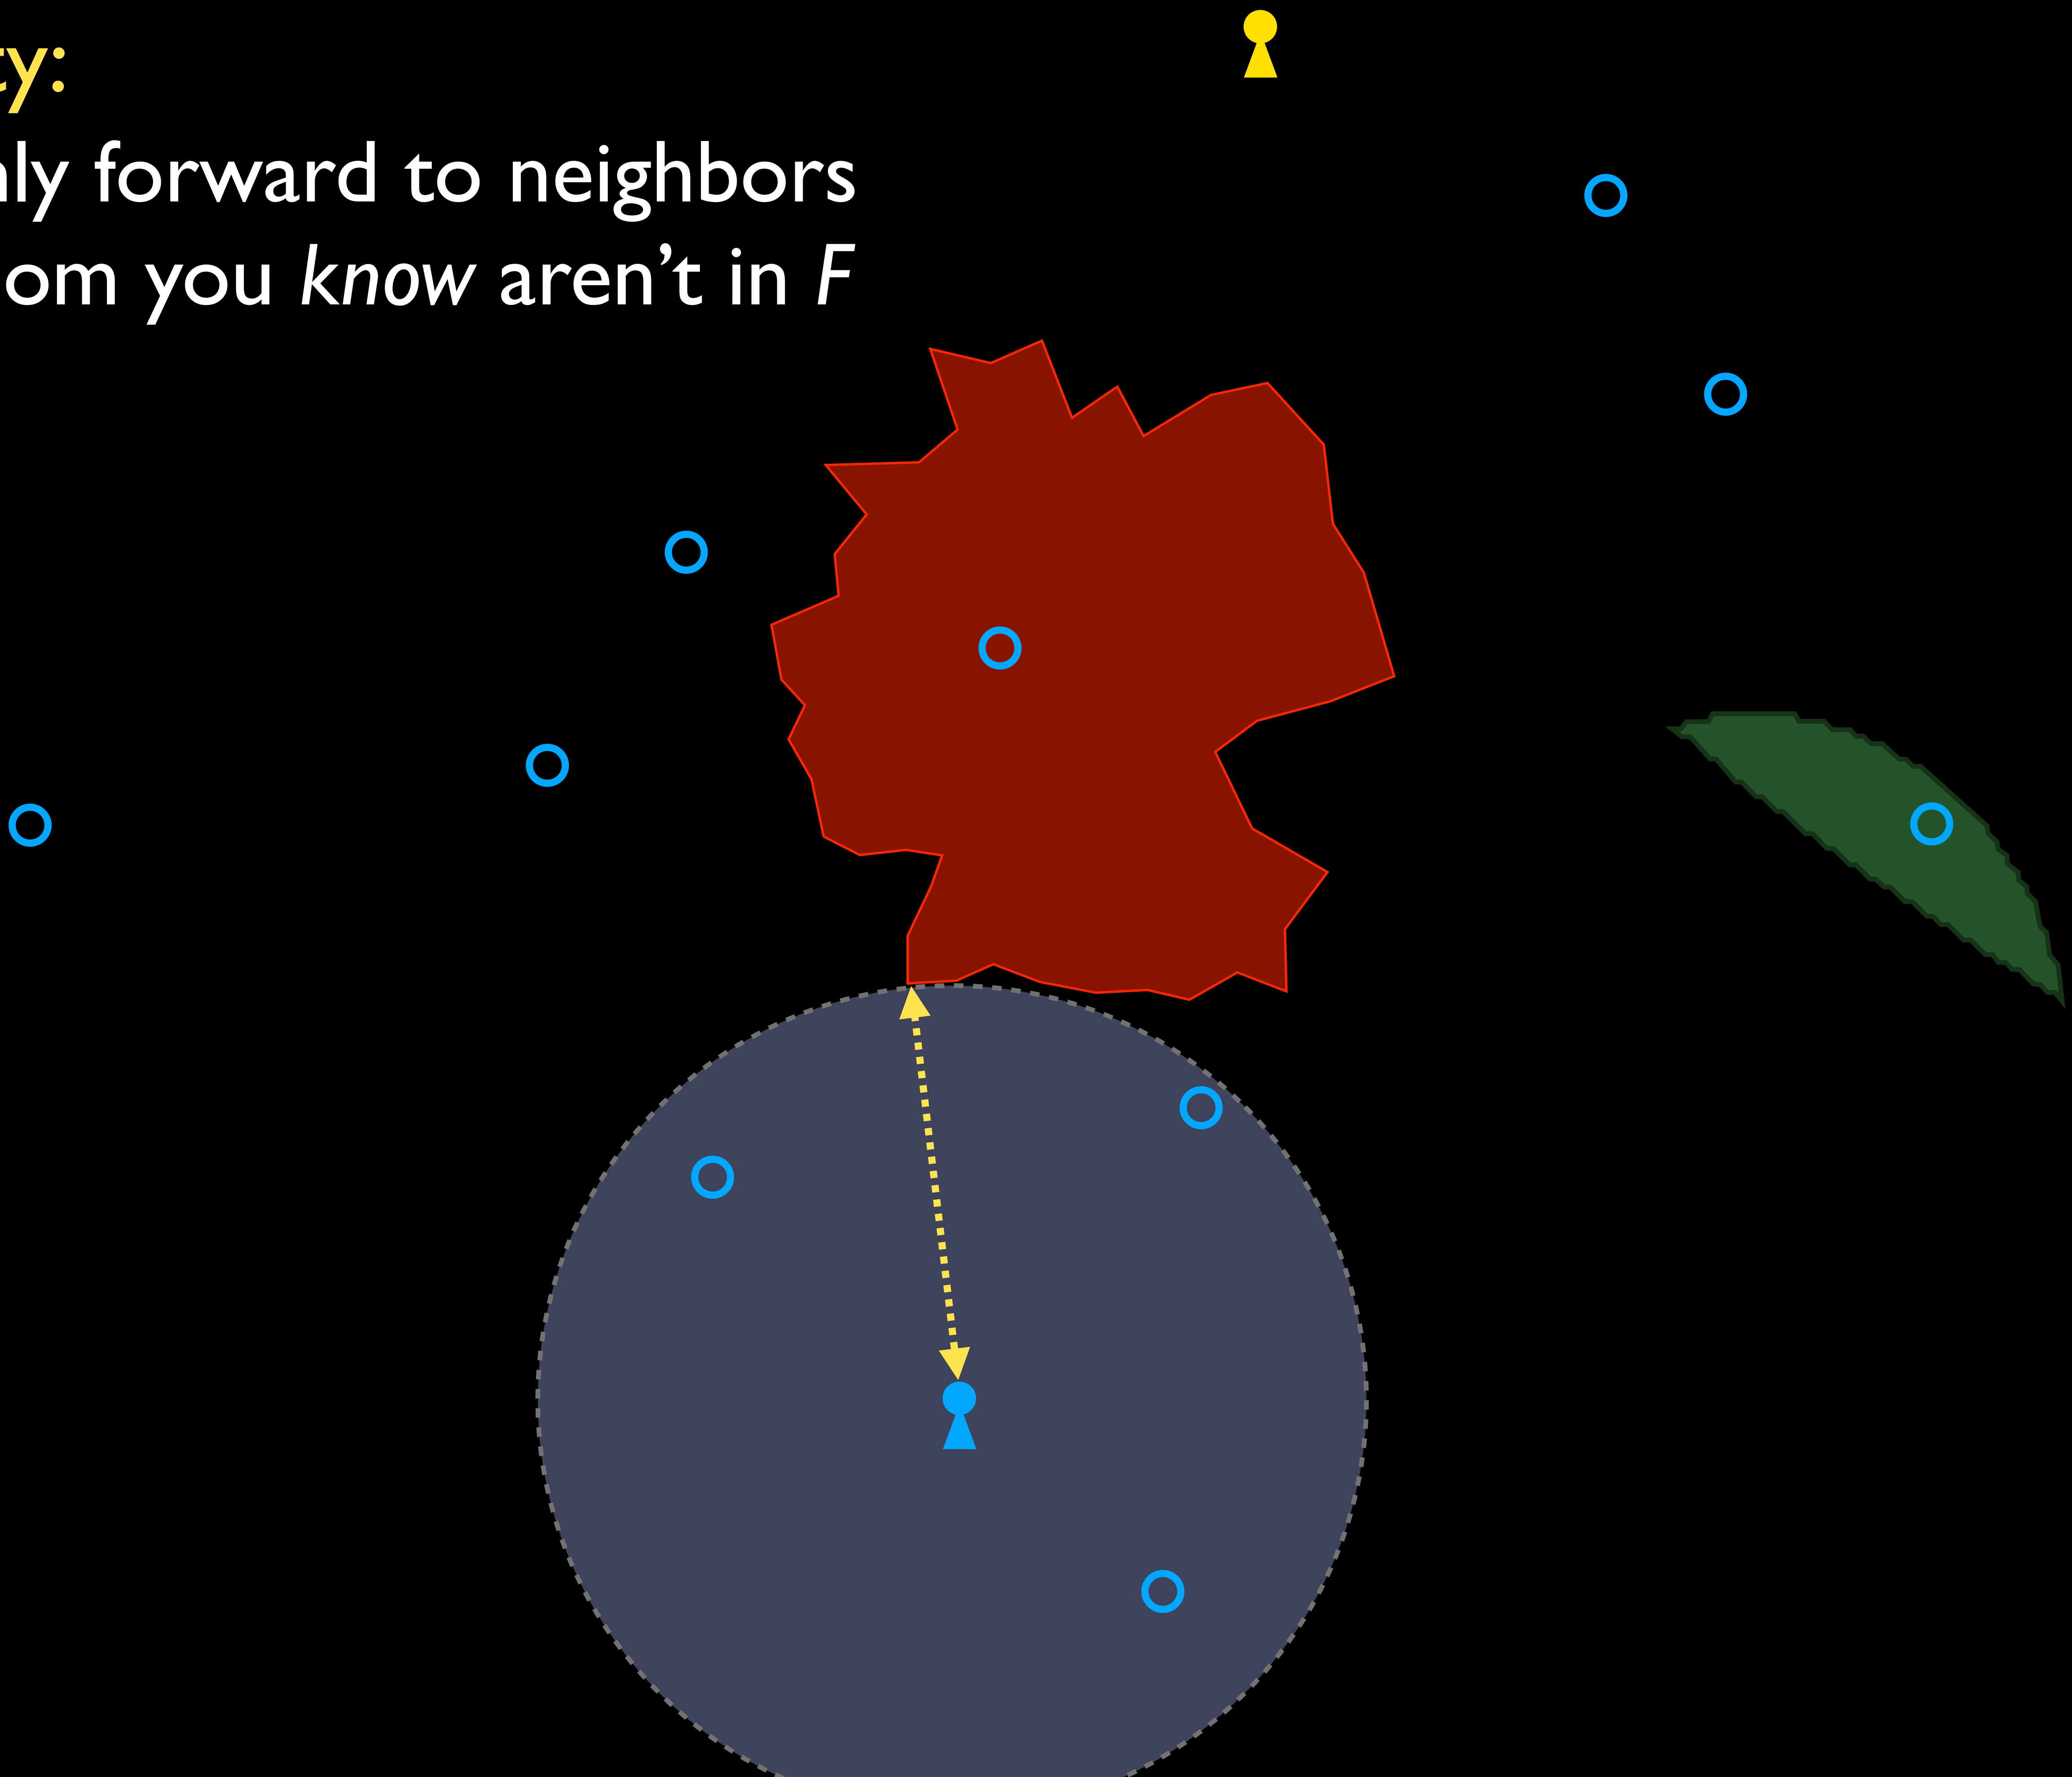

Safety:

Only forward to neighbors whom you *know* aren't in *F*

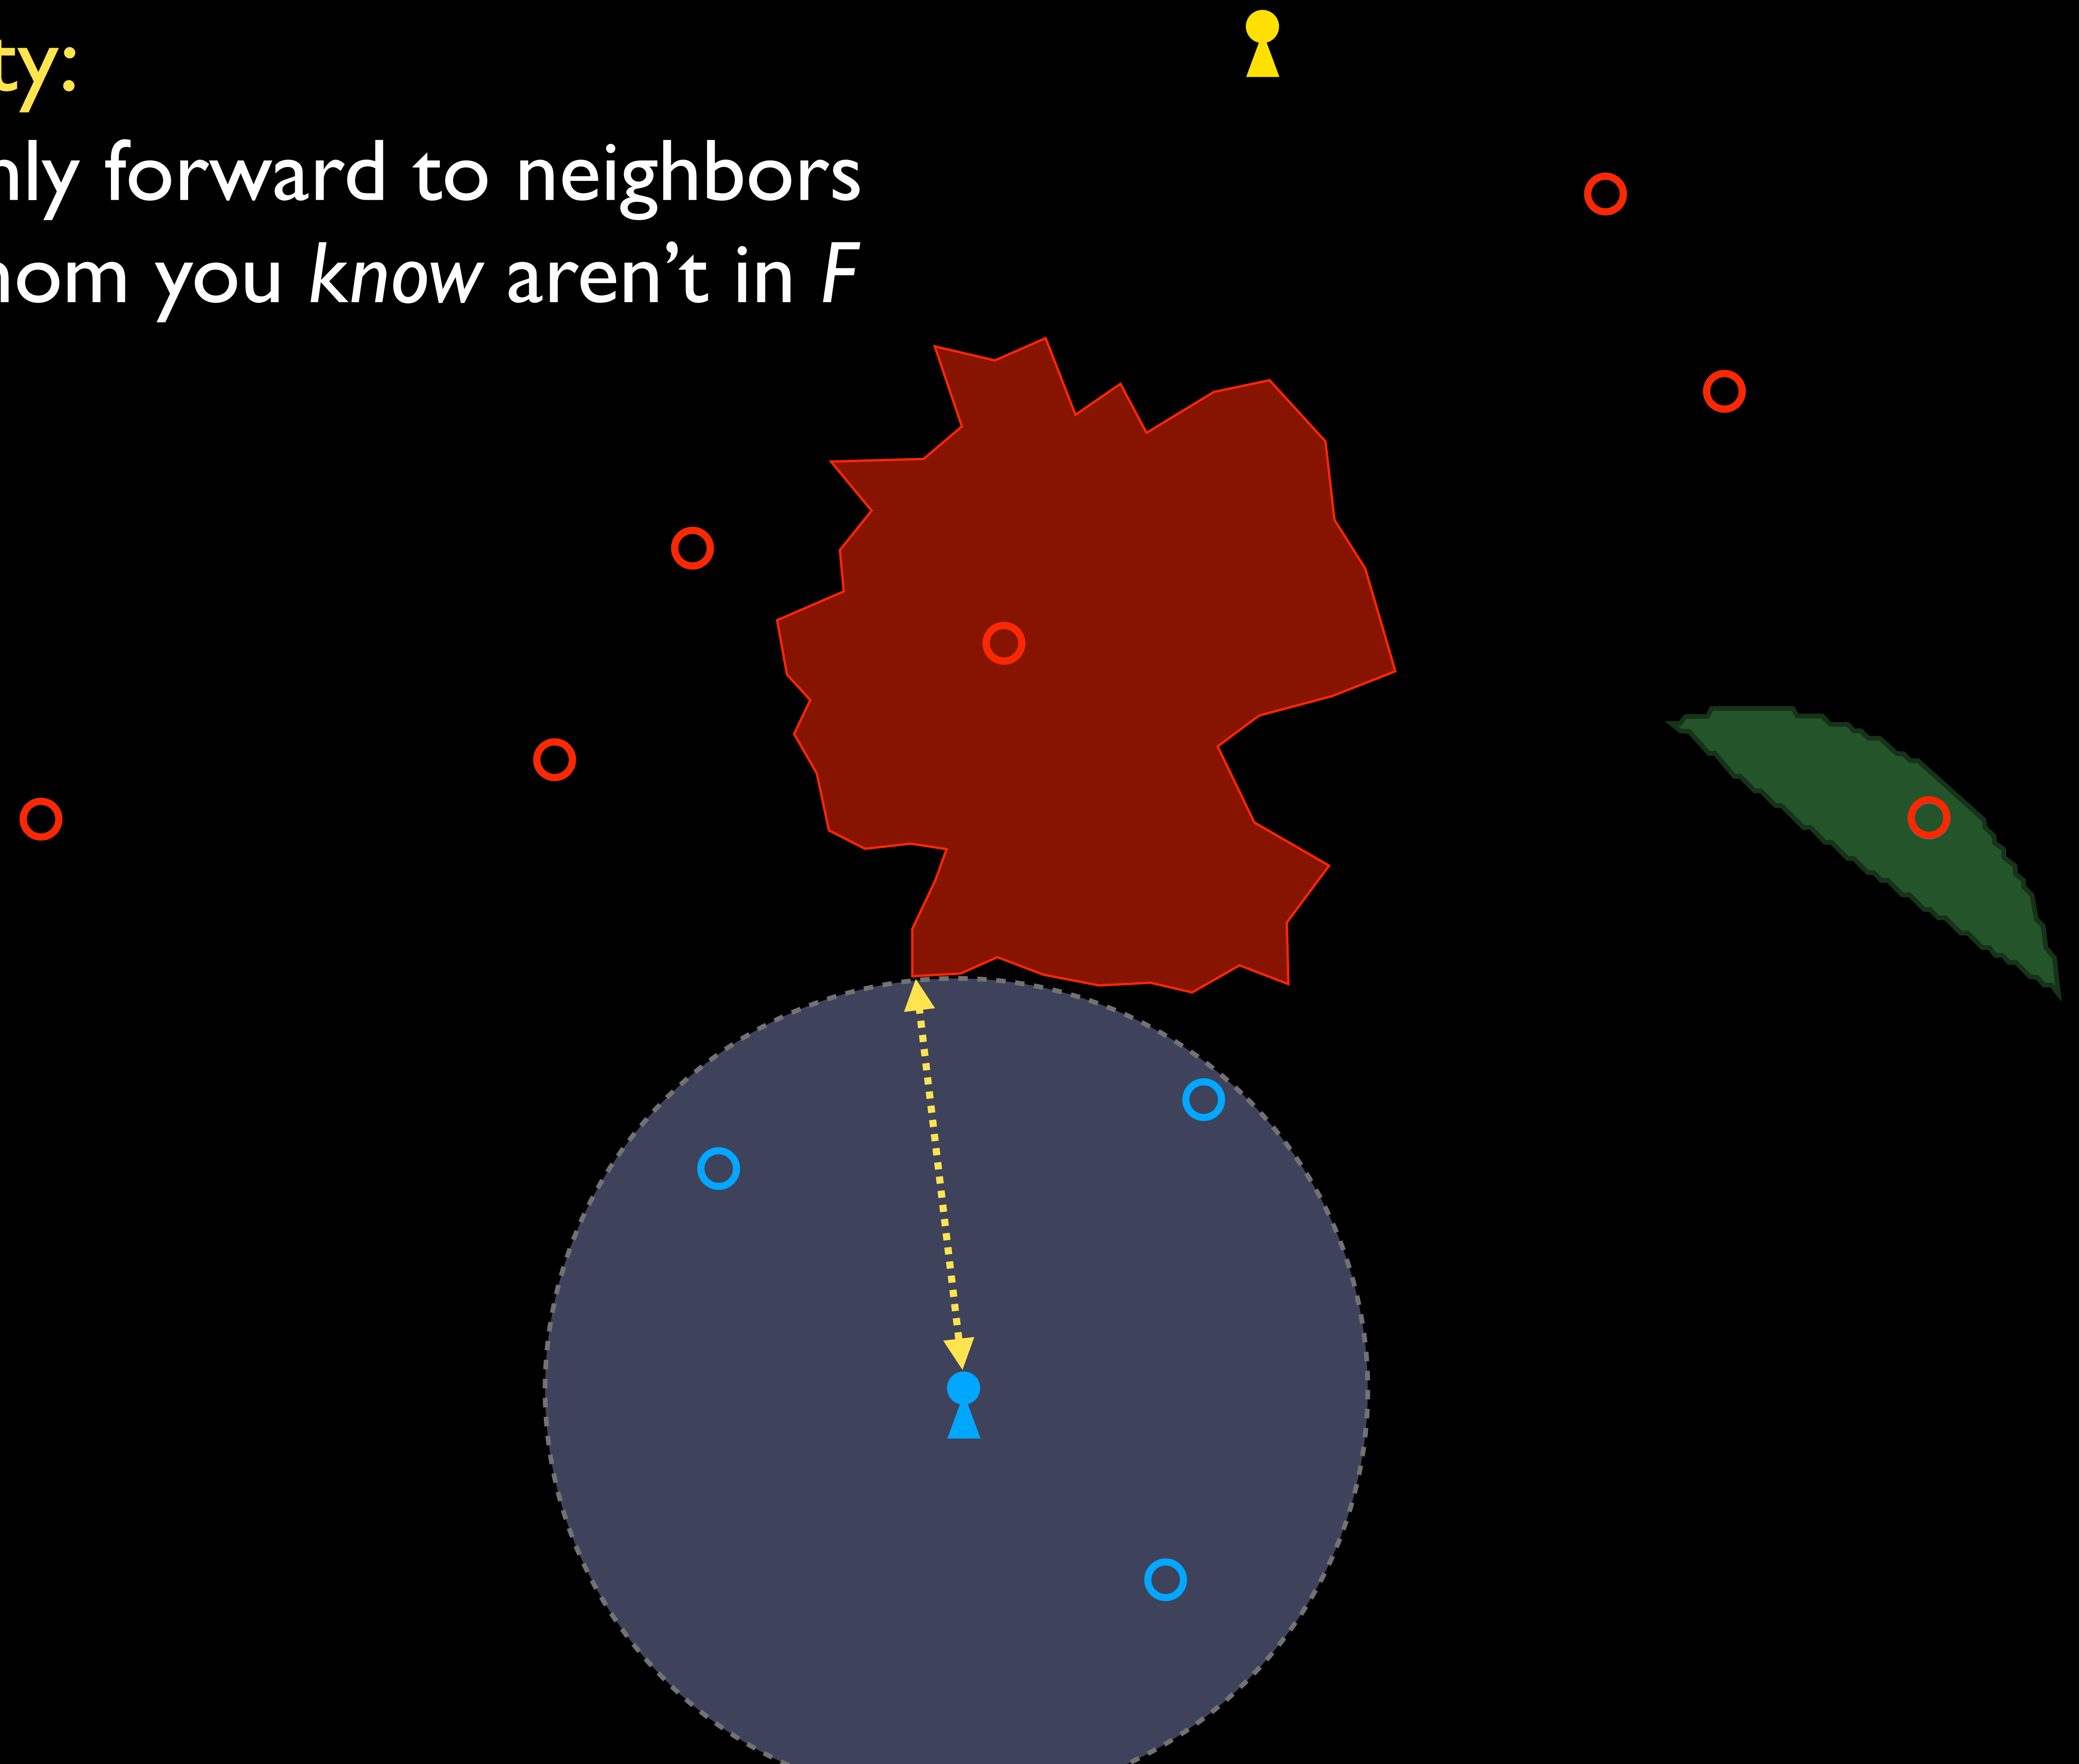

Safety:

Only forward to neighbors whom you *know* aren't in *F*

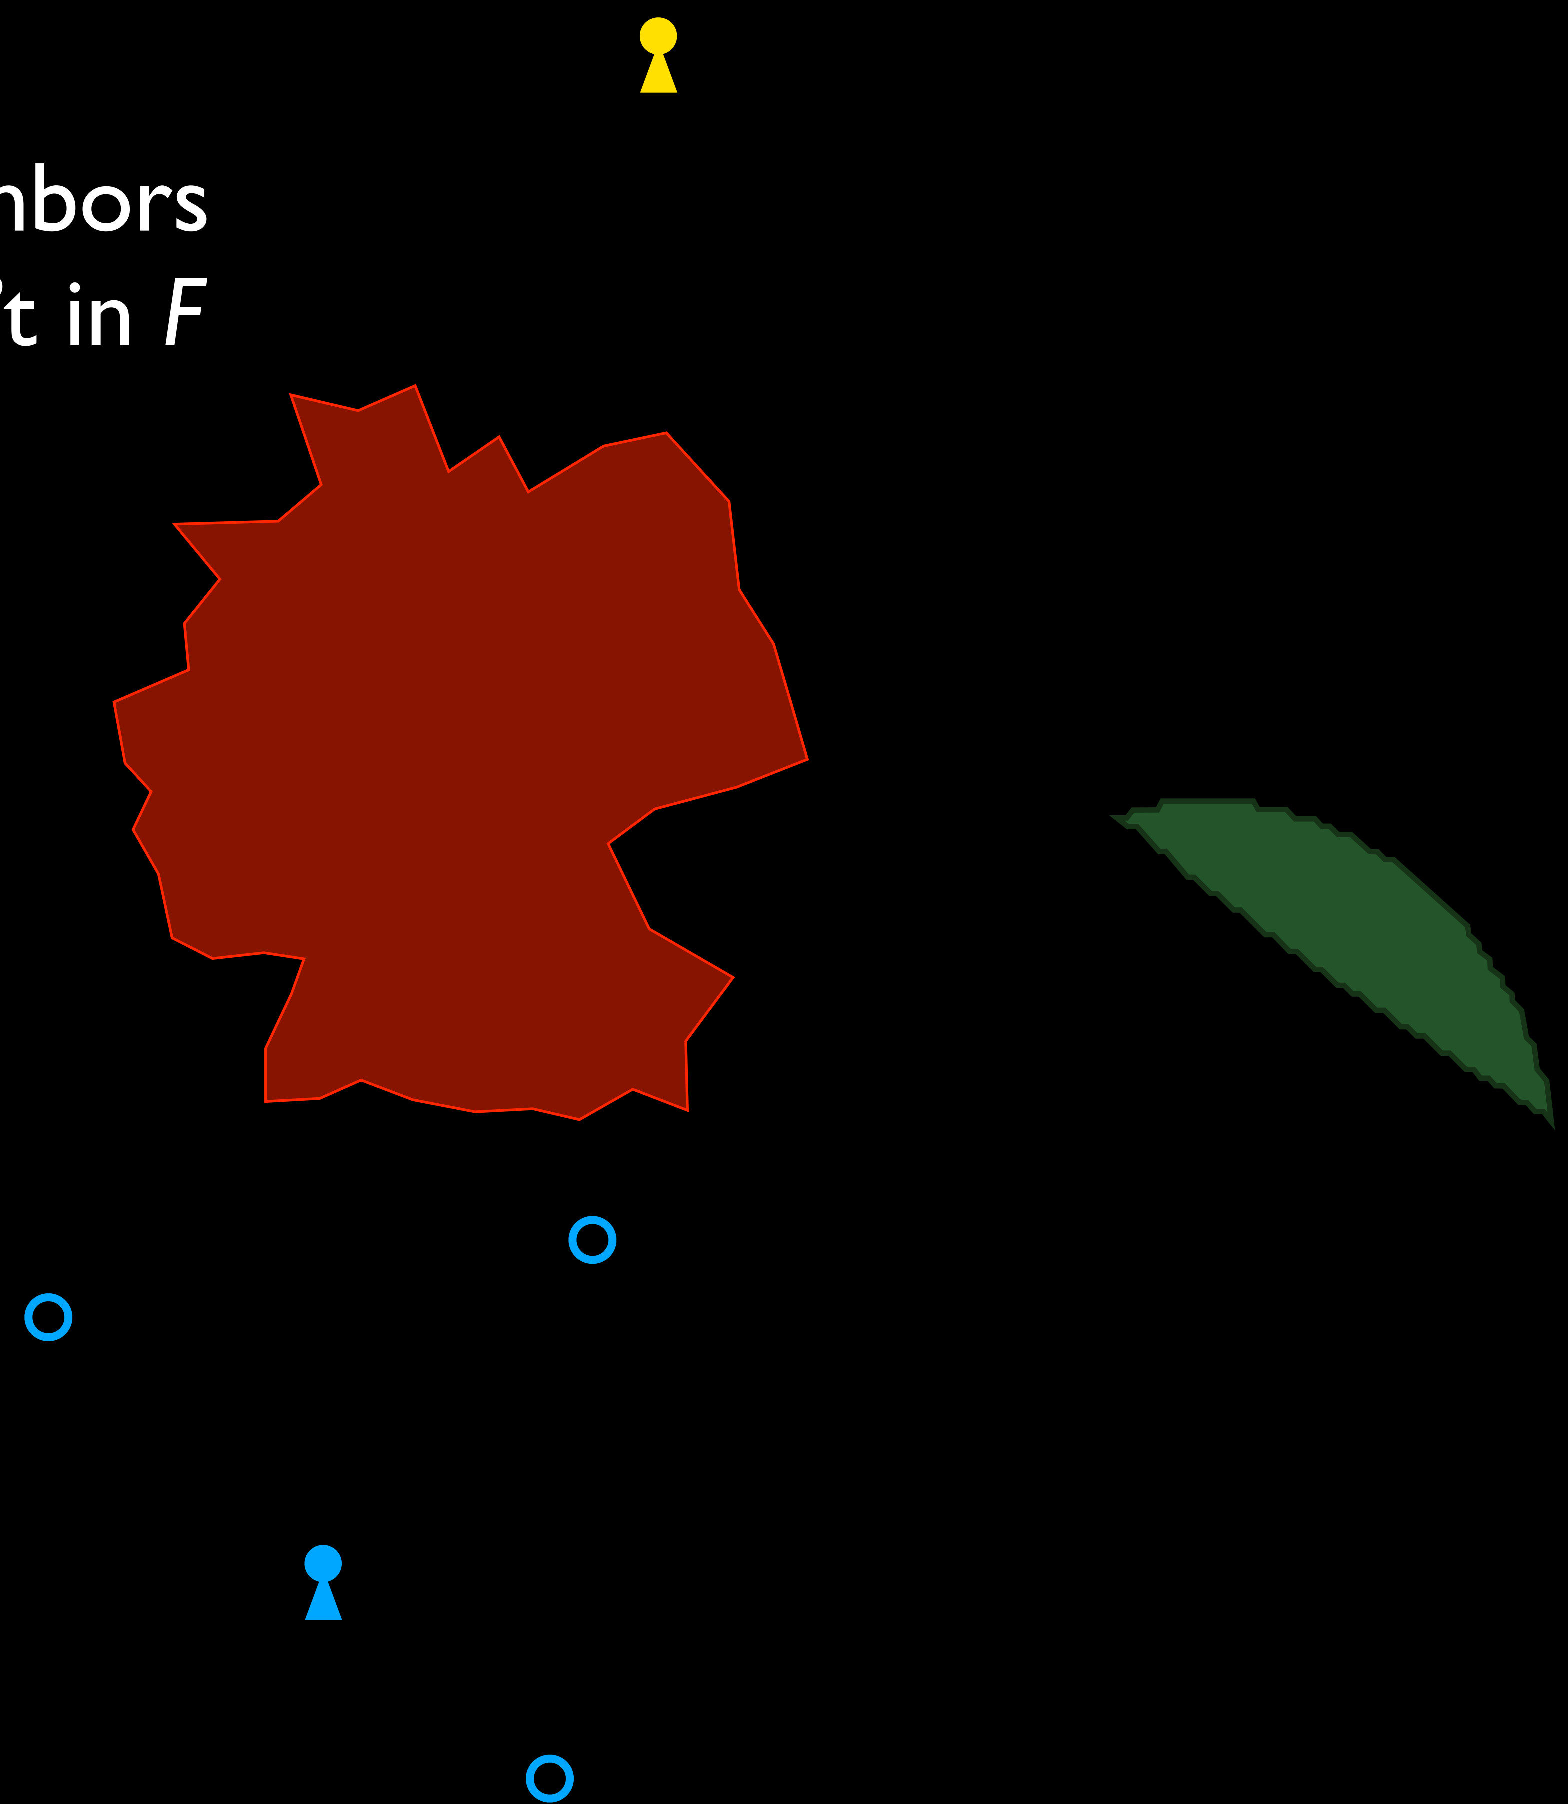

Safety:

Only forward to neighbors whom you *know* aren't in *F*

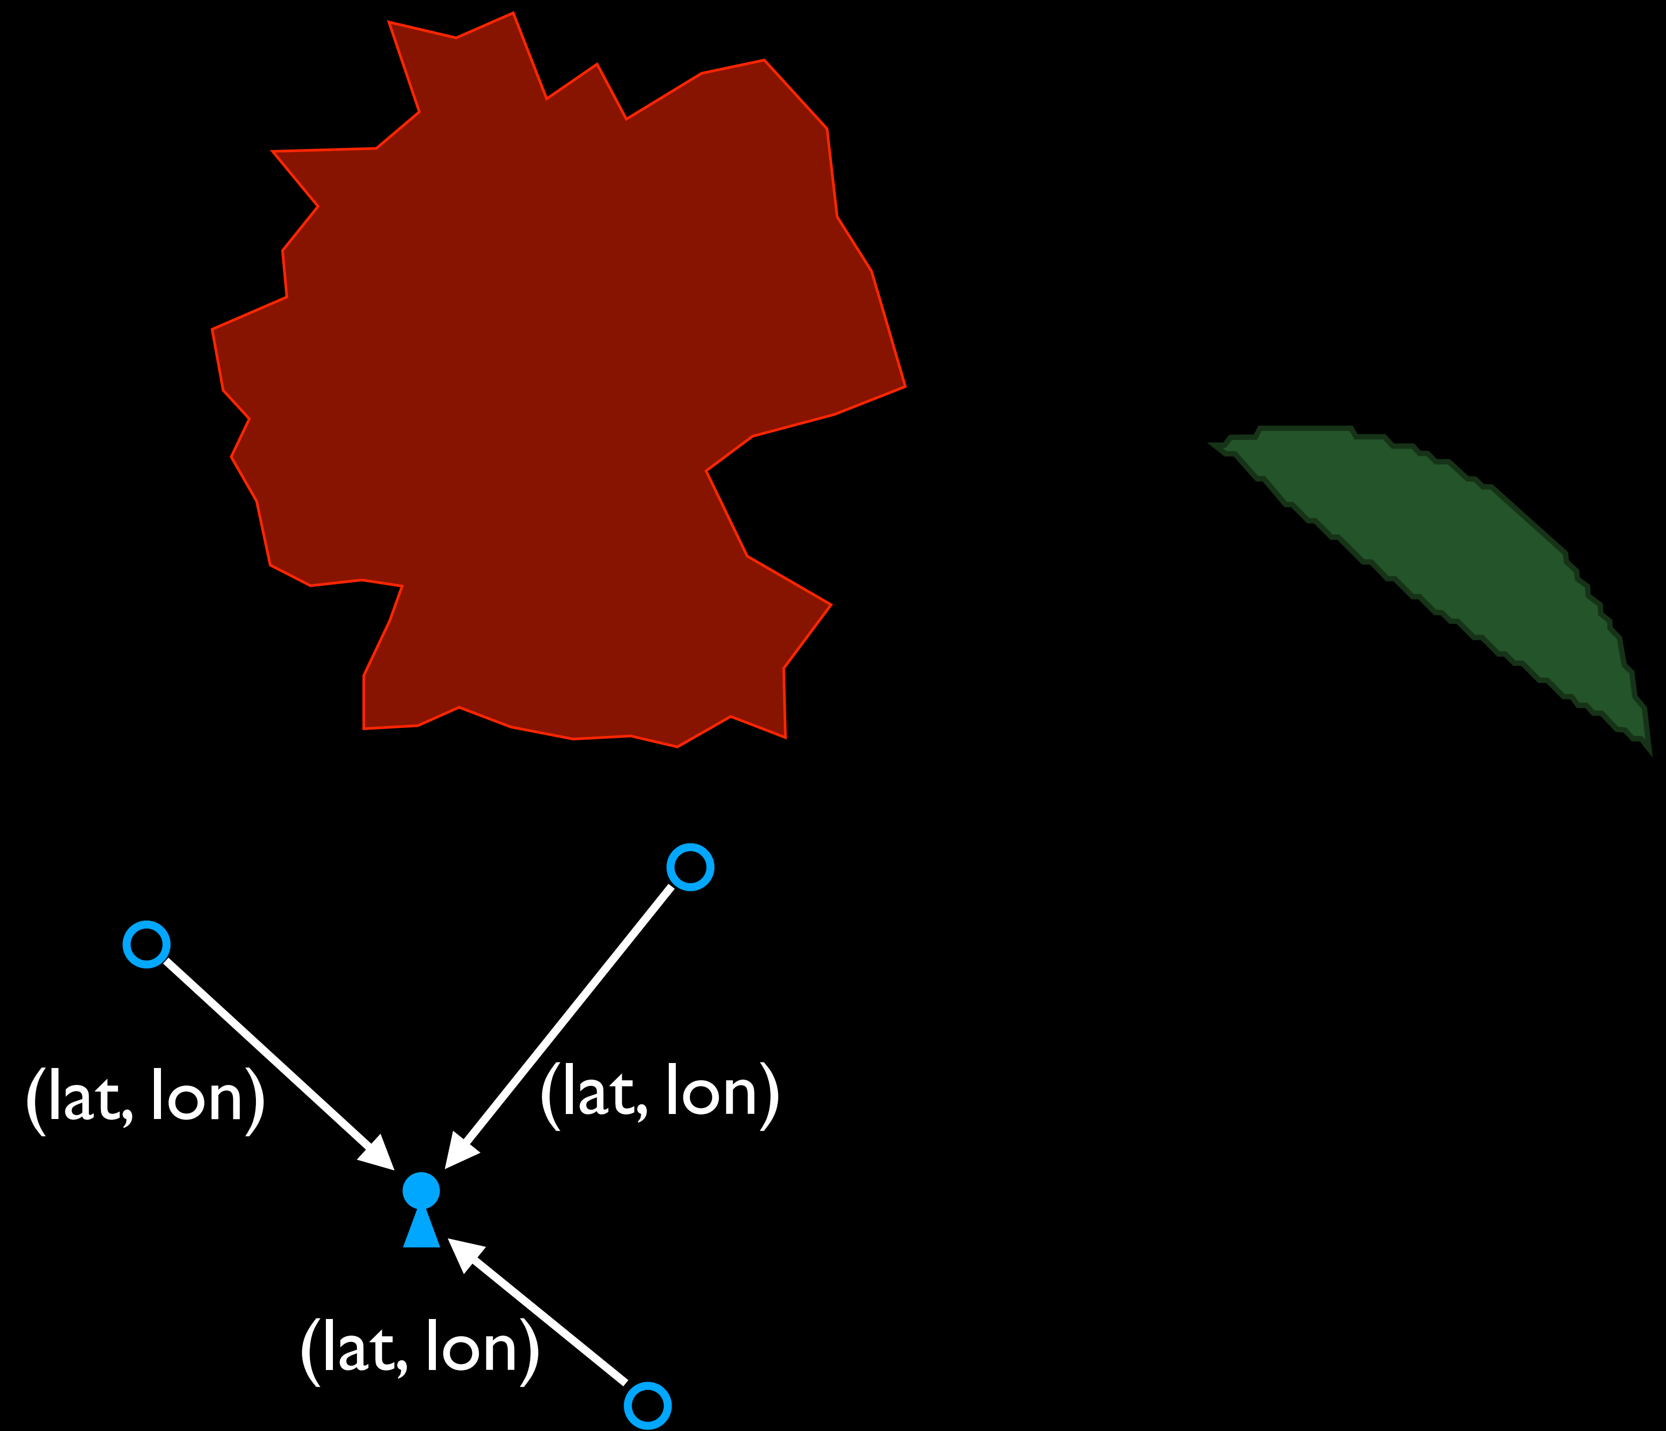

Progress:

Forward to the (safe) neighbor whose safe zone is closest to T

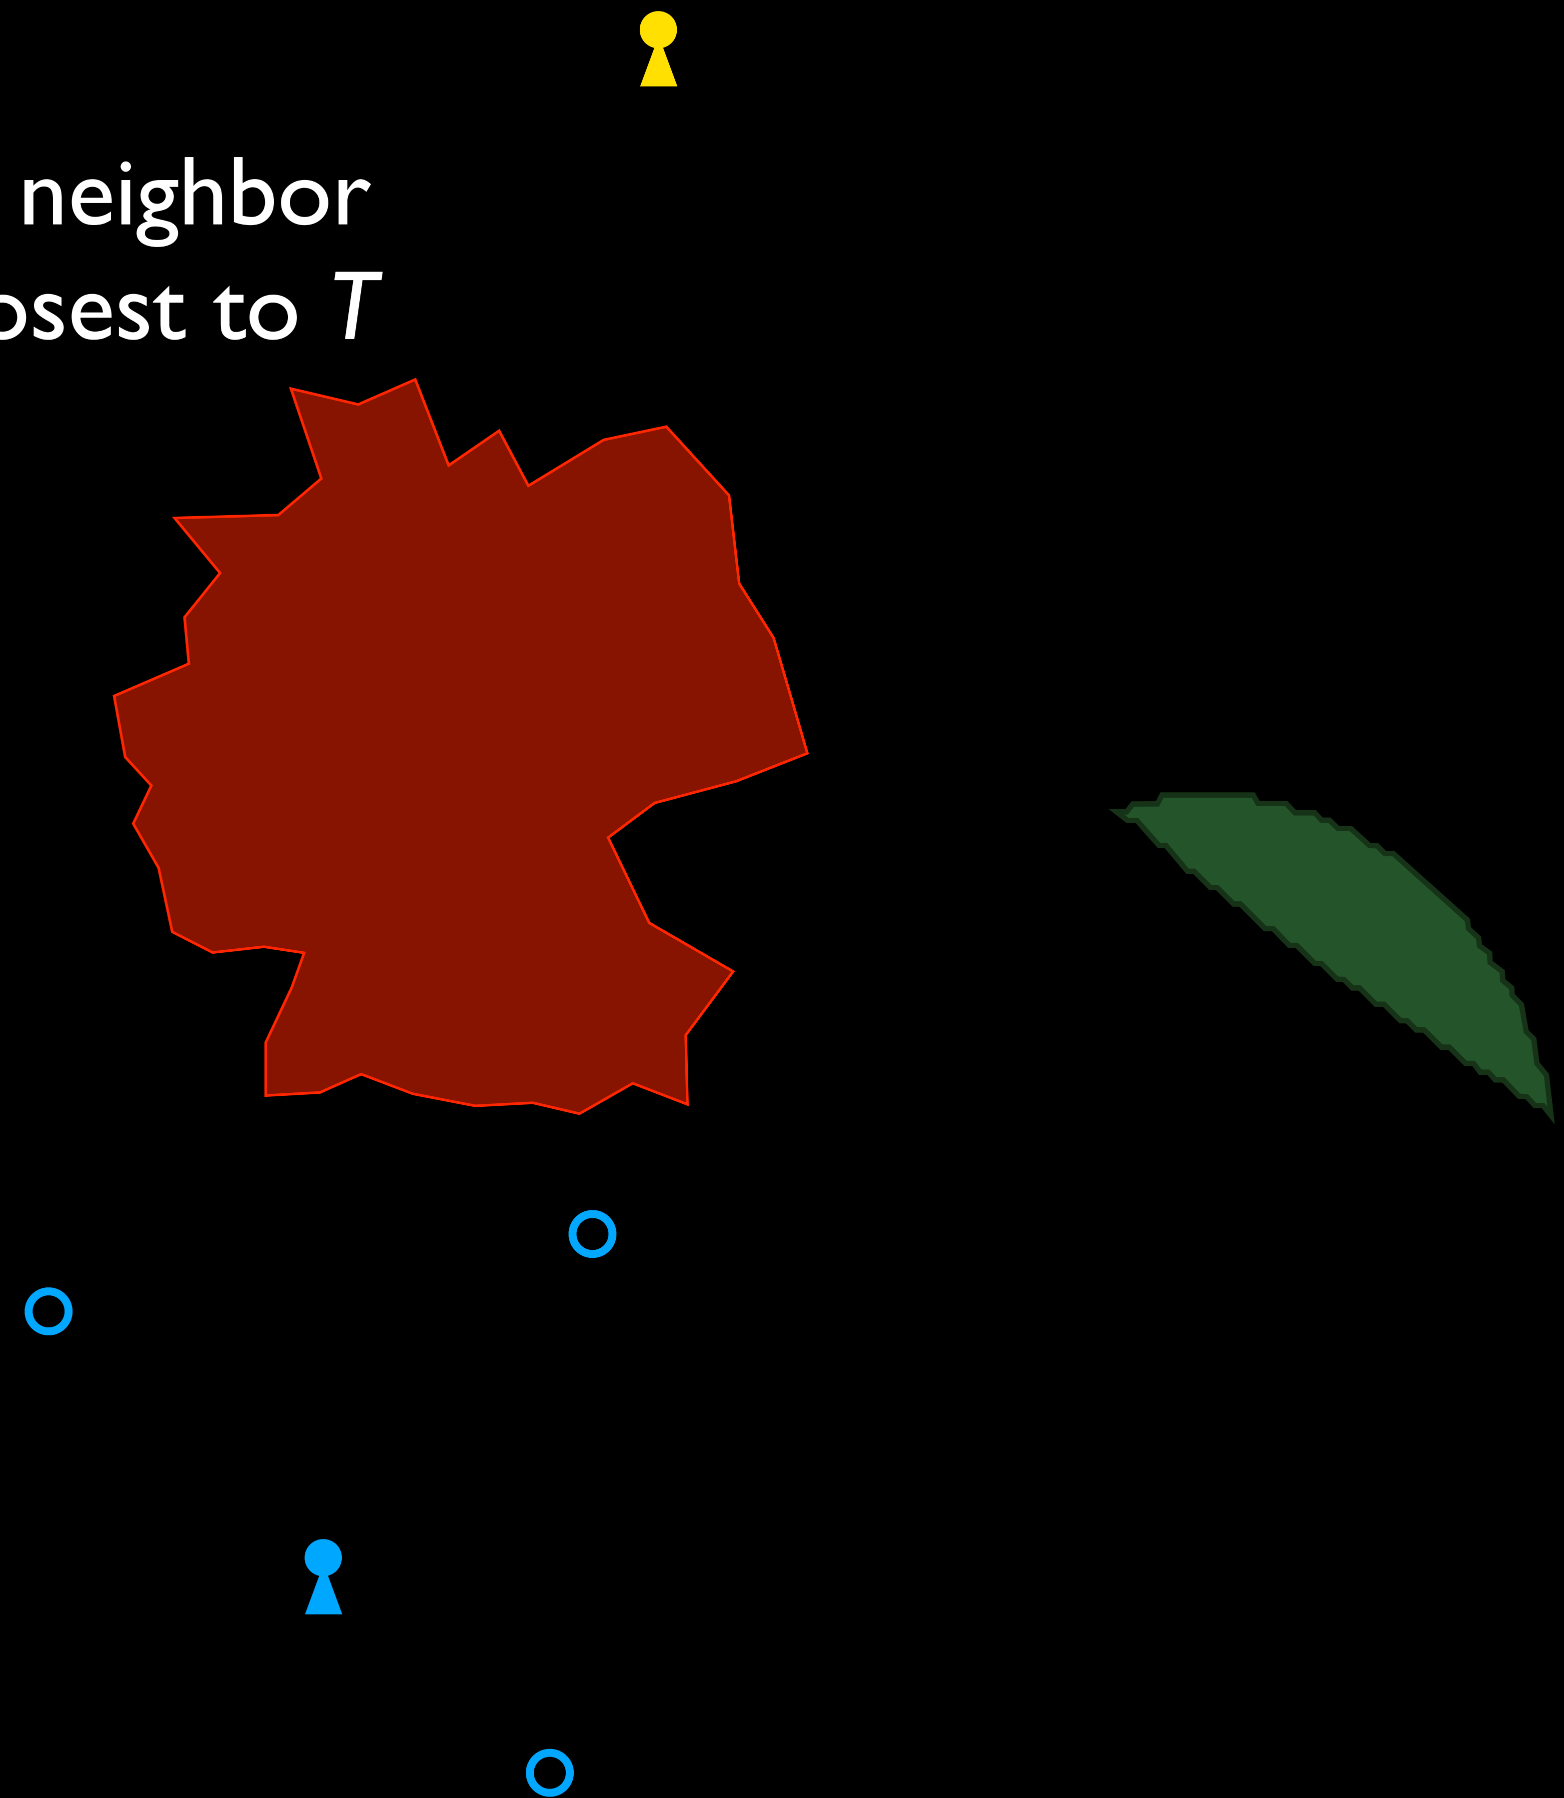

Progress:

Forward to the (safe) neighbor whose safe zone is closest to T

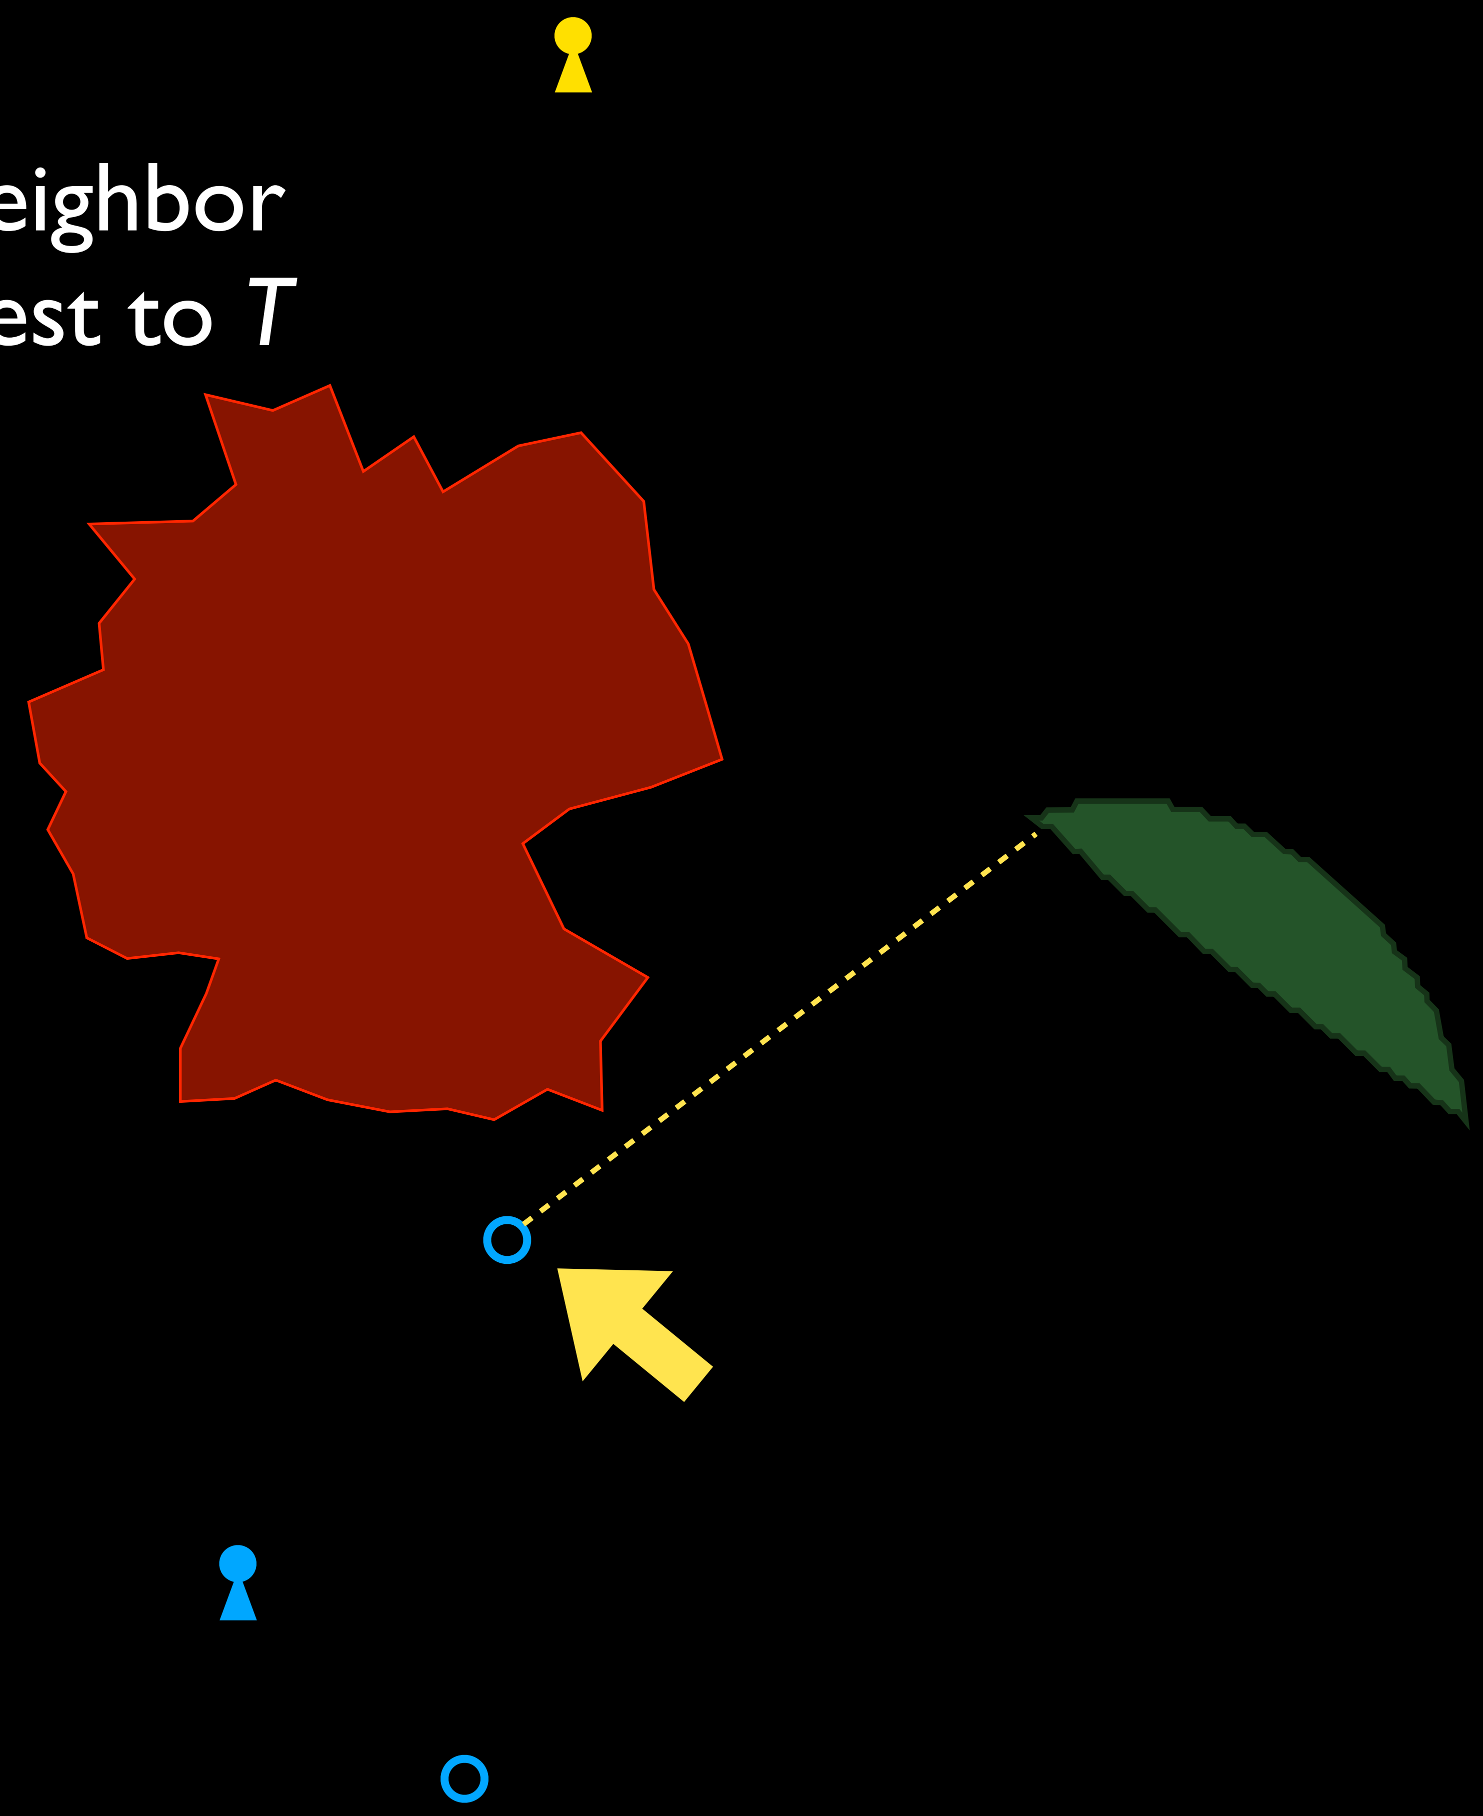

Progress:

Forward to the (safe) neighbor whose safe zone is closest to T

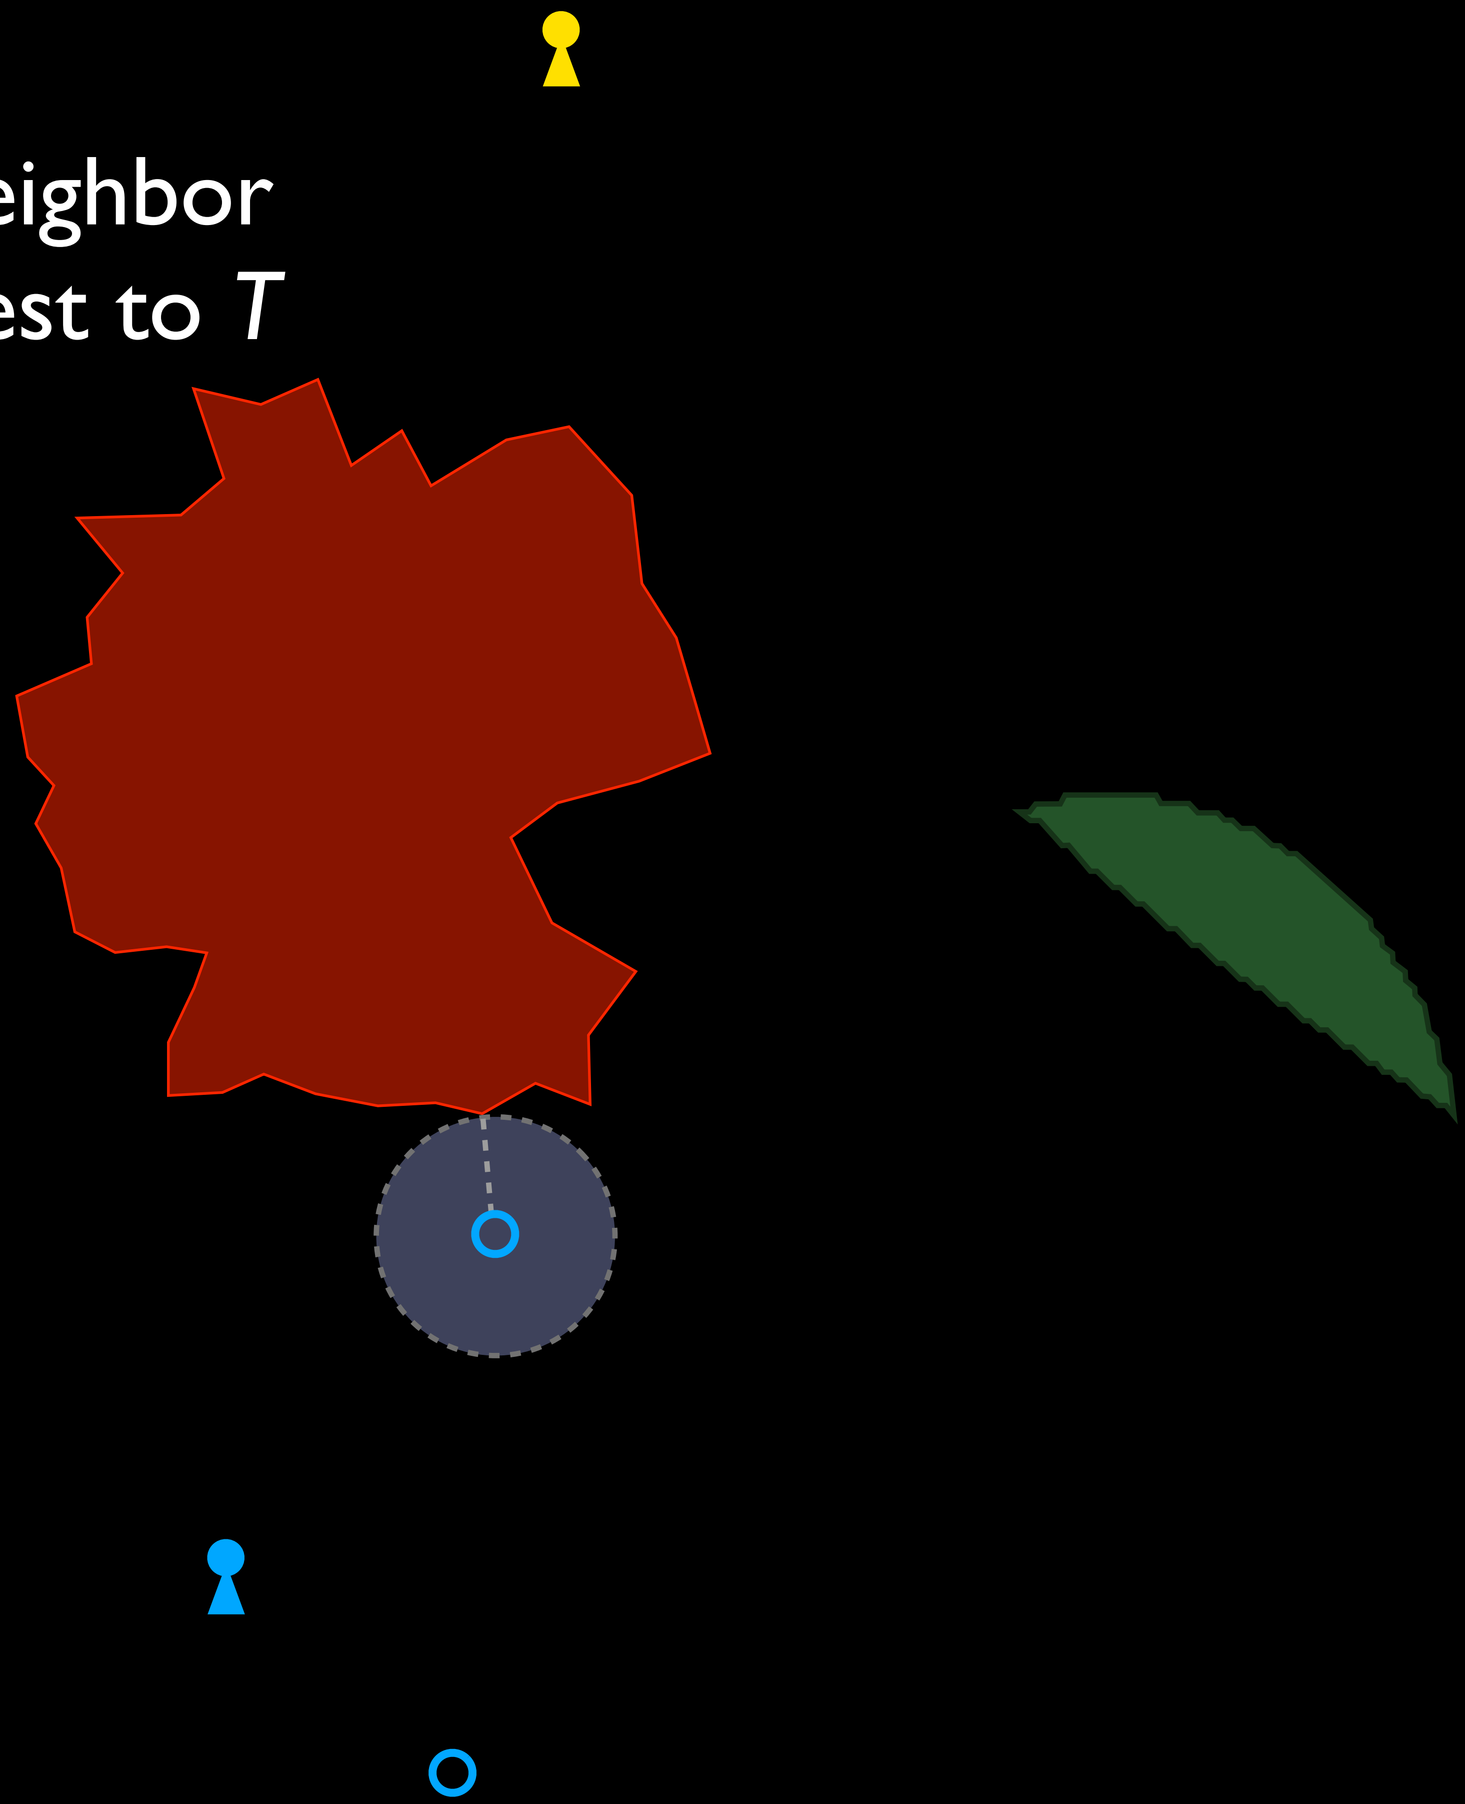

Progress:

Forward to the (safe) neighbor whose safe zone is closest to *T*

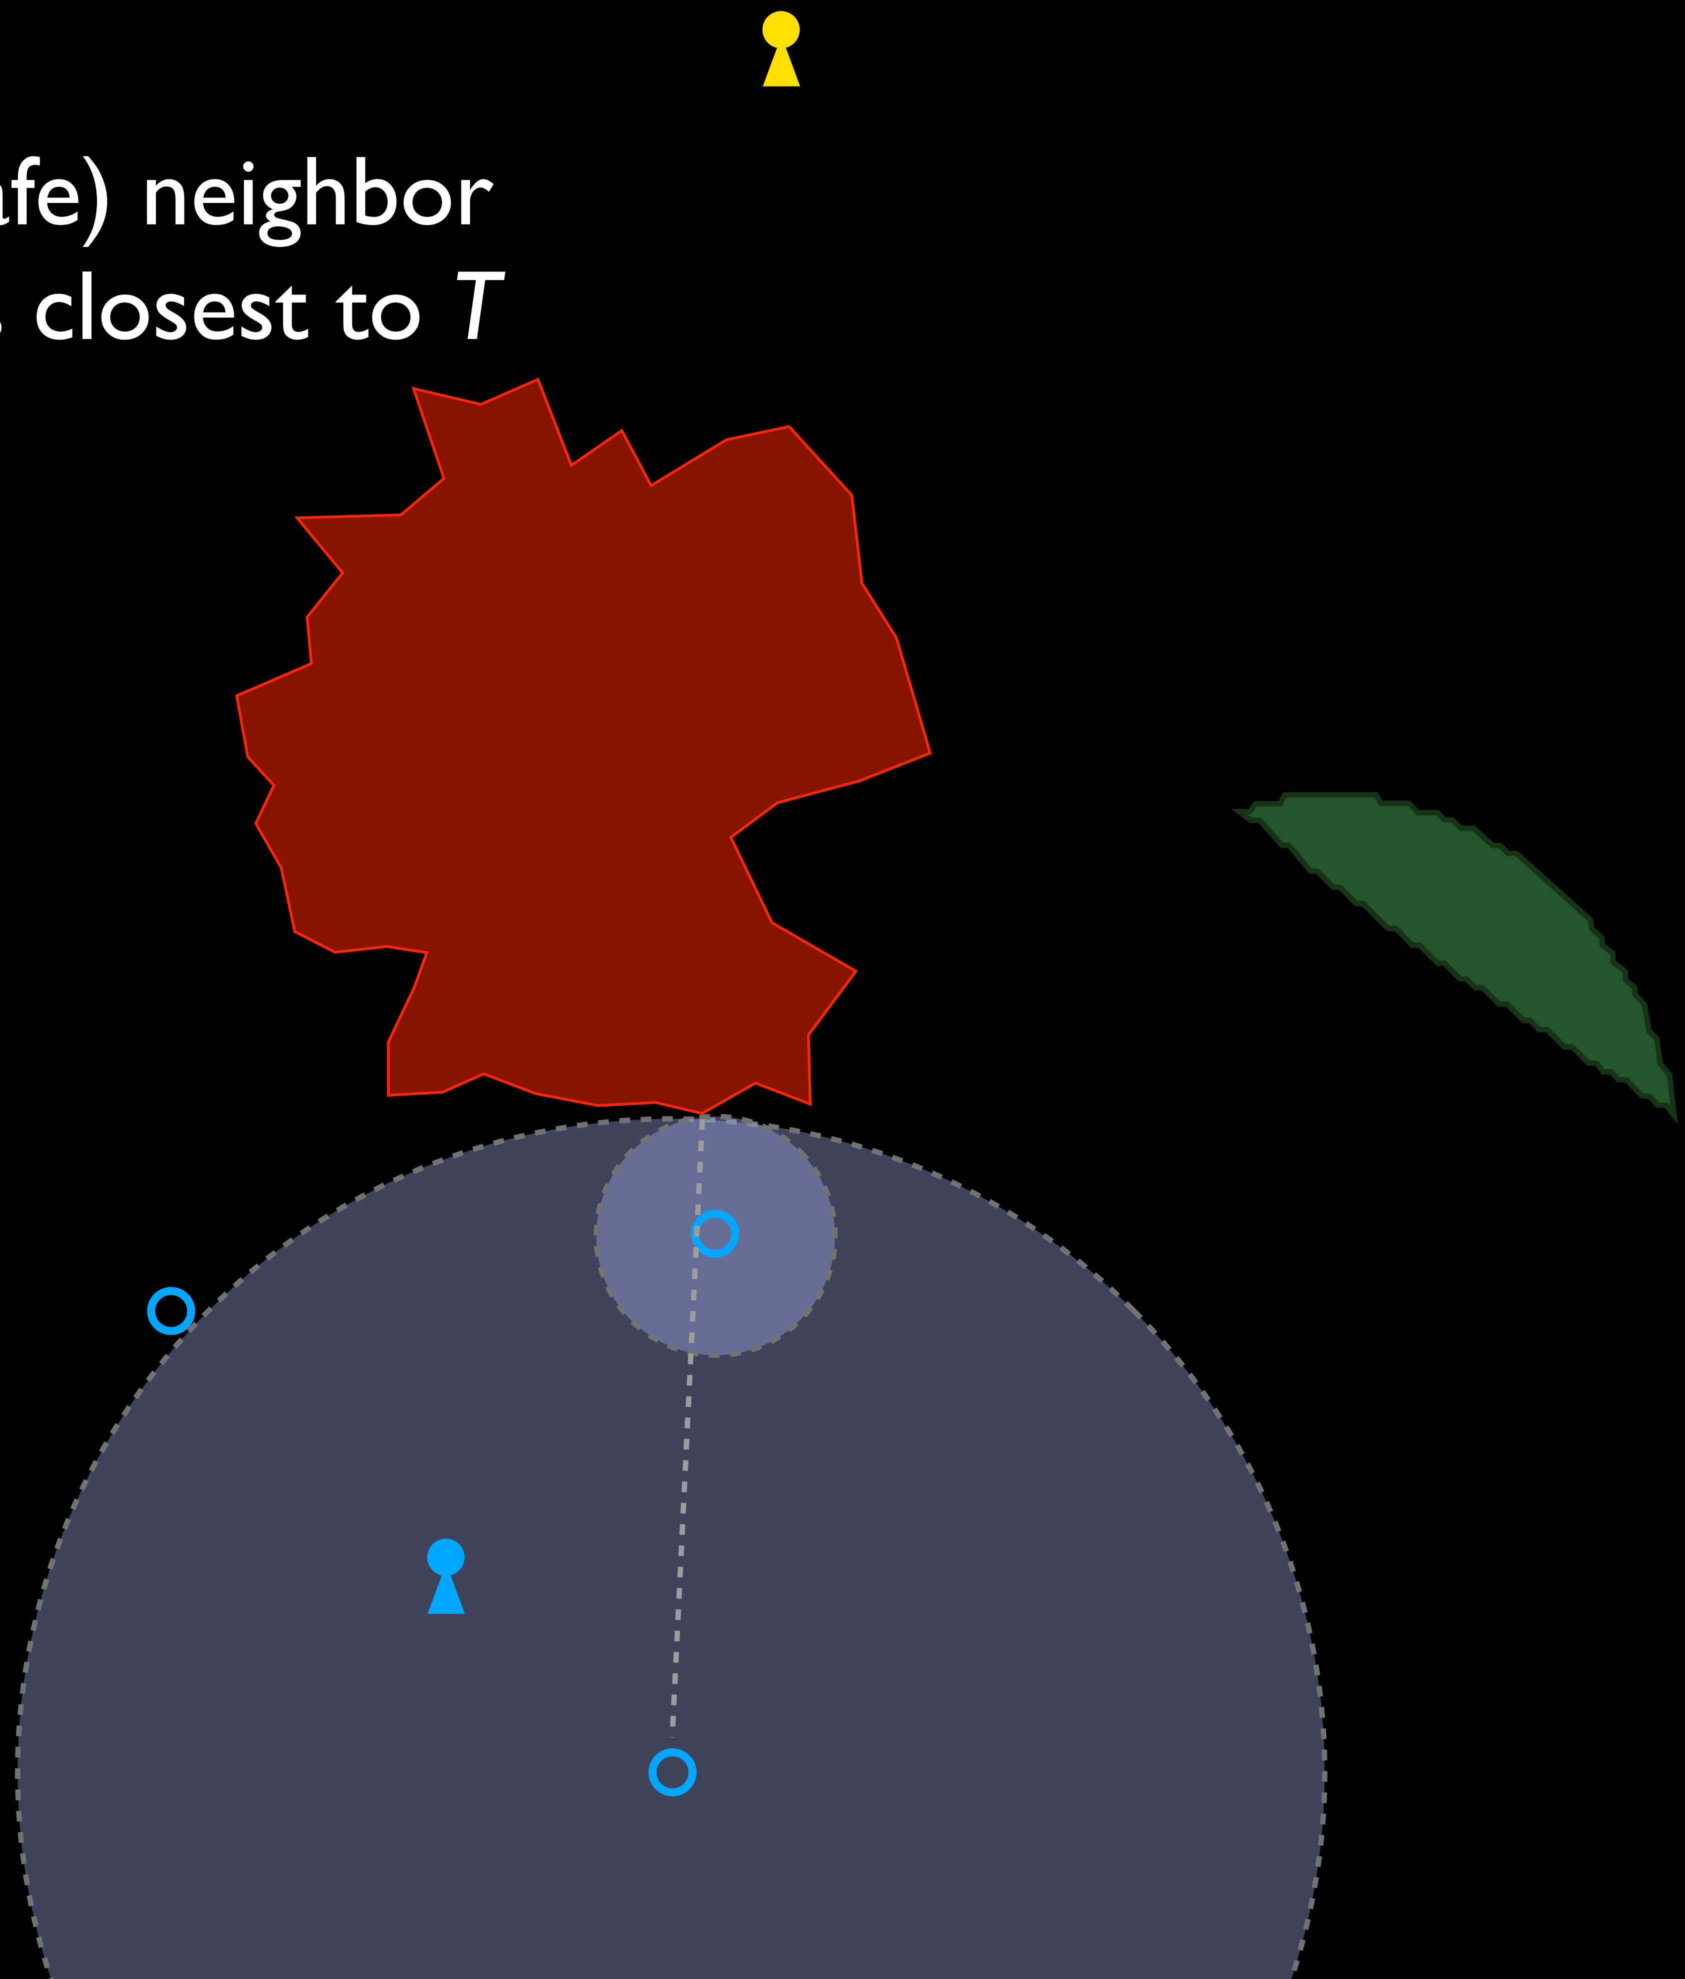

Progress:

Forward to the (safe) neighbor whose safe zone is closest to *T*

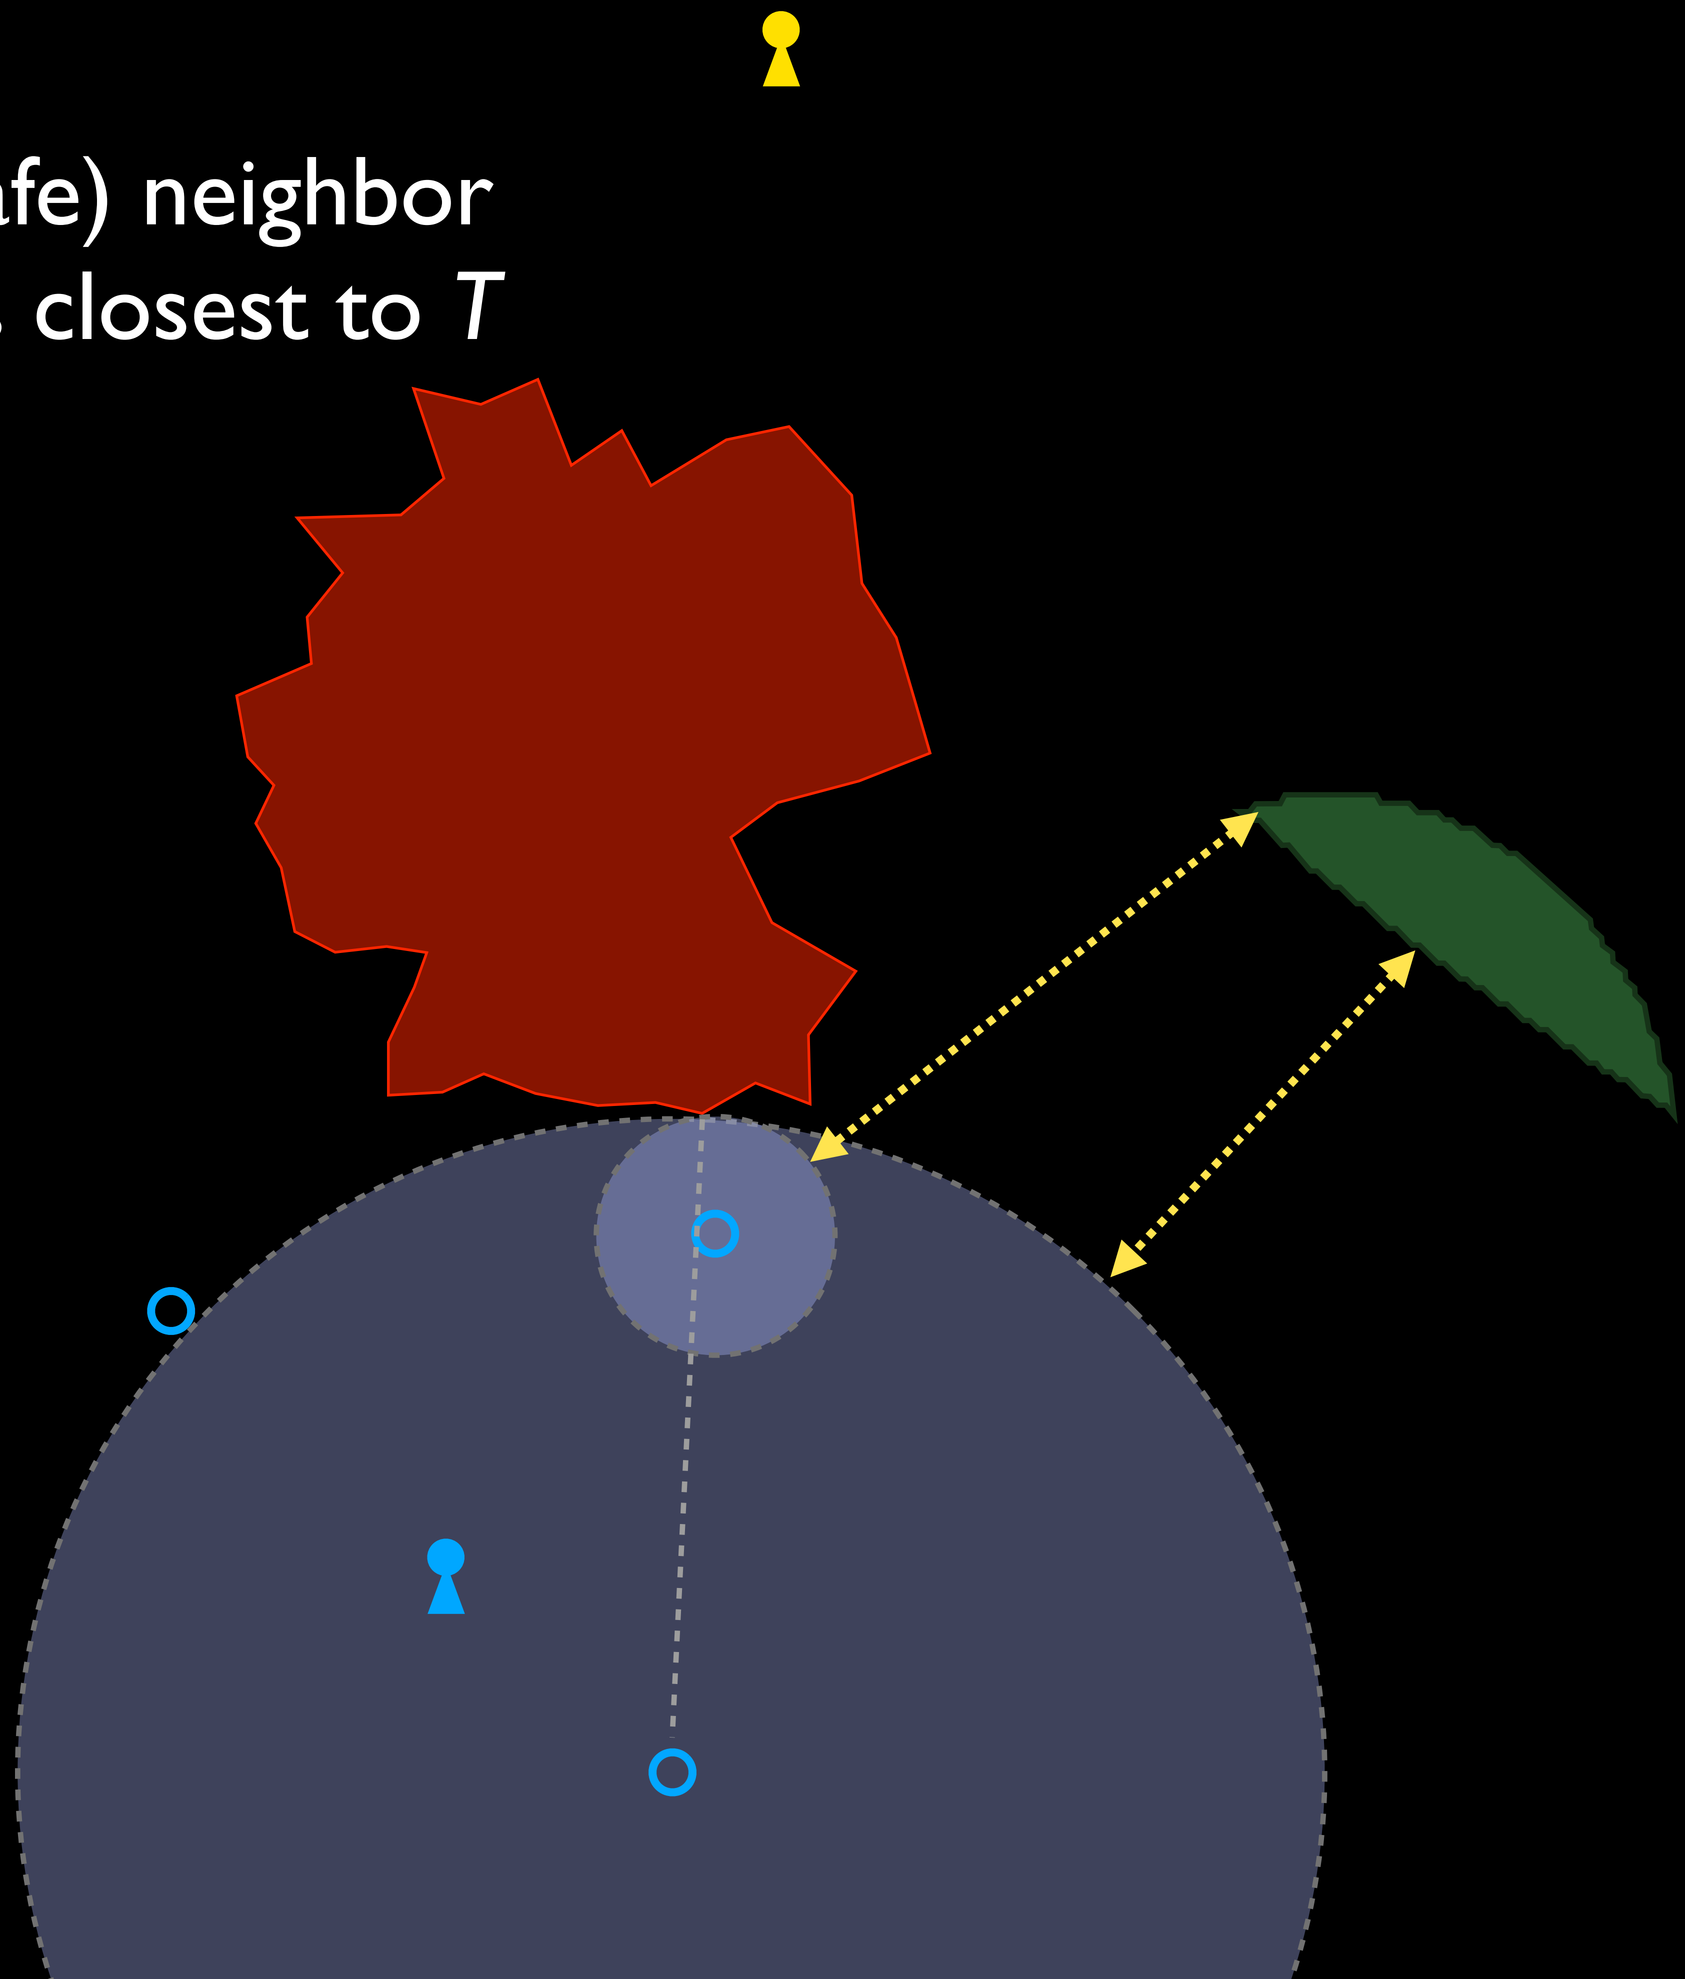

Recursive forwarding

Forward *F* and *T*

Continue until progress stops

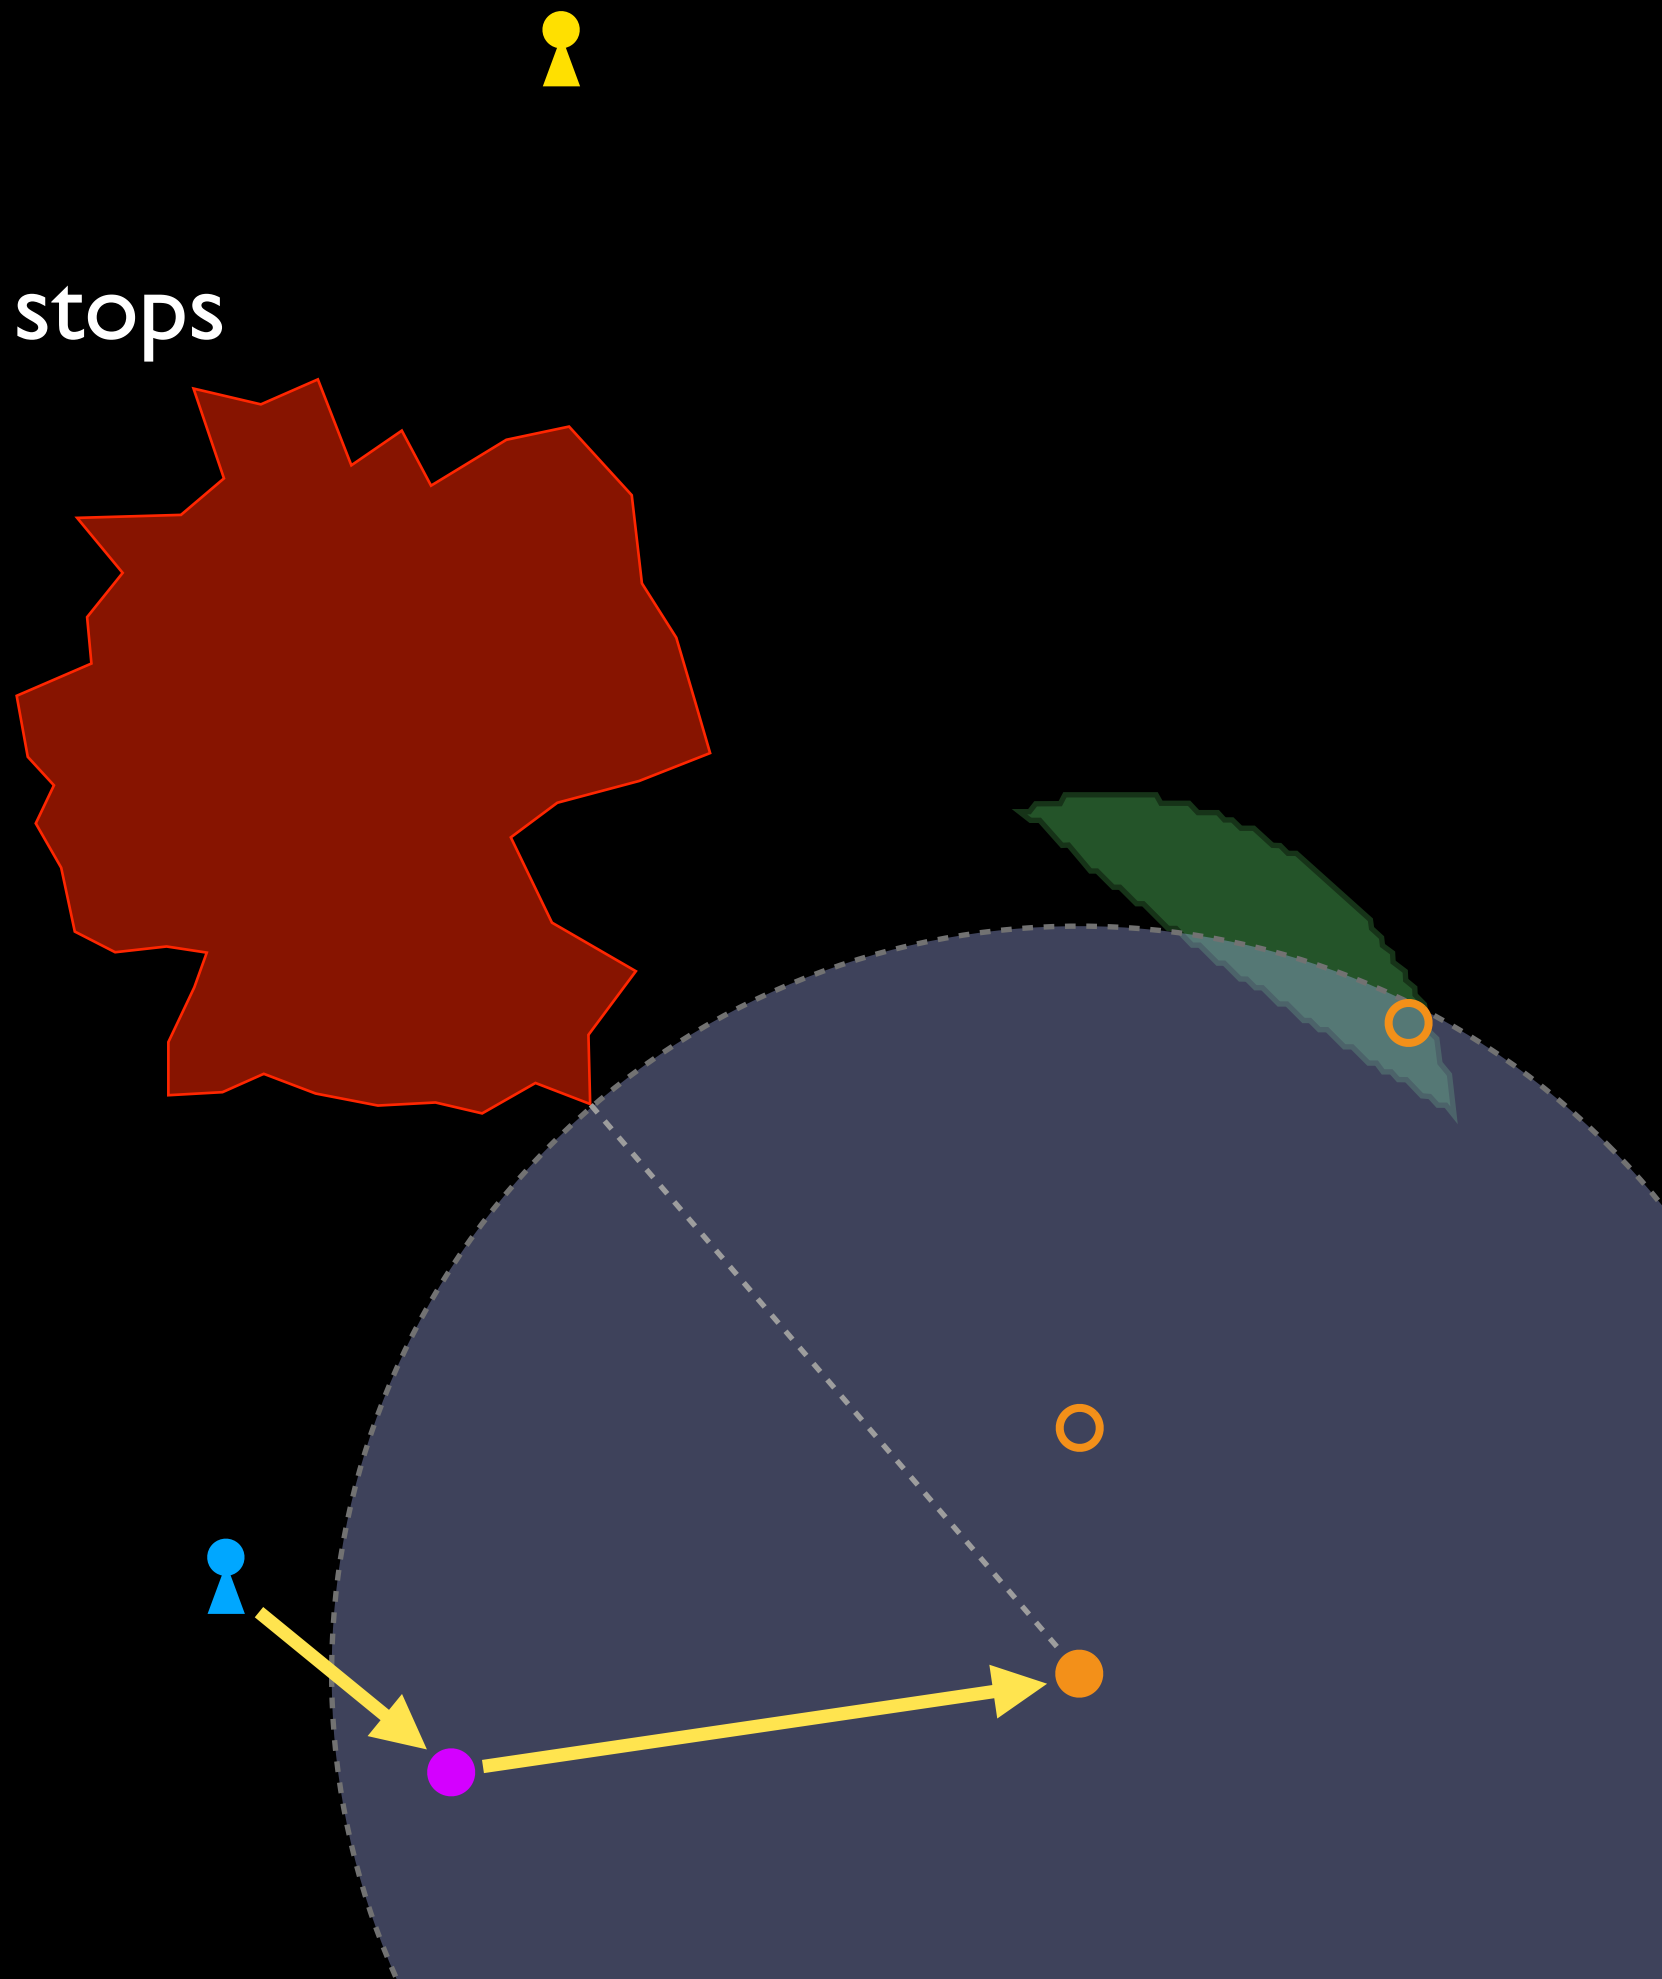

Recursive forwarding

- Forward *F* and *T*
- Continue until progress stops

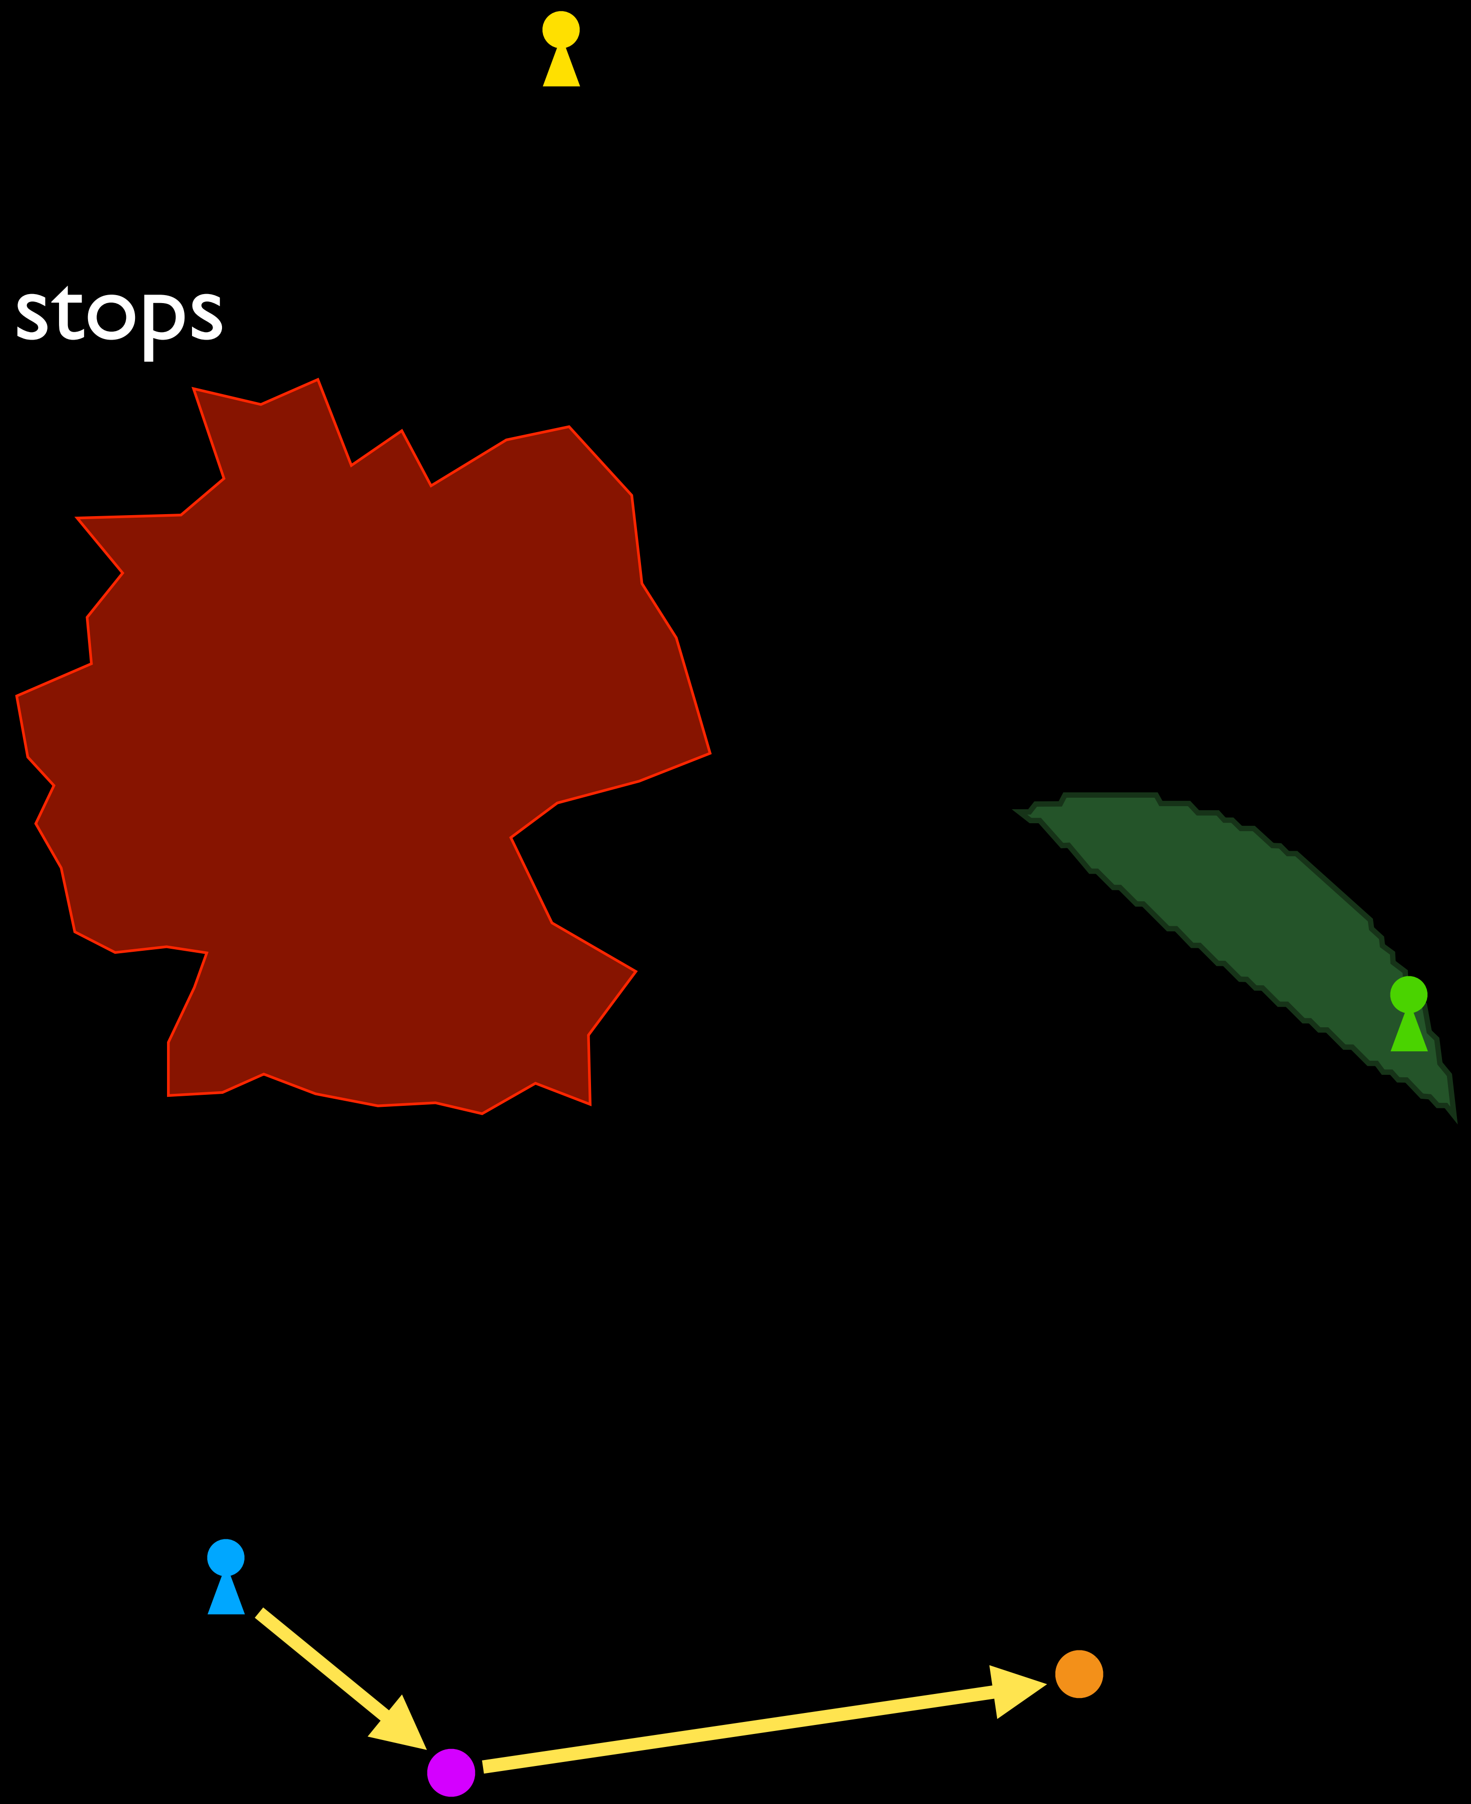

Recursive forwarding

- Forward *F* and *T*
- Continue until progress stops

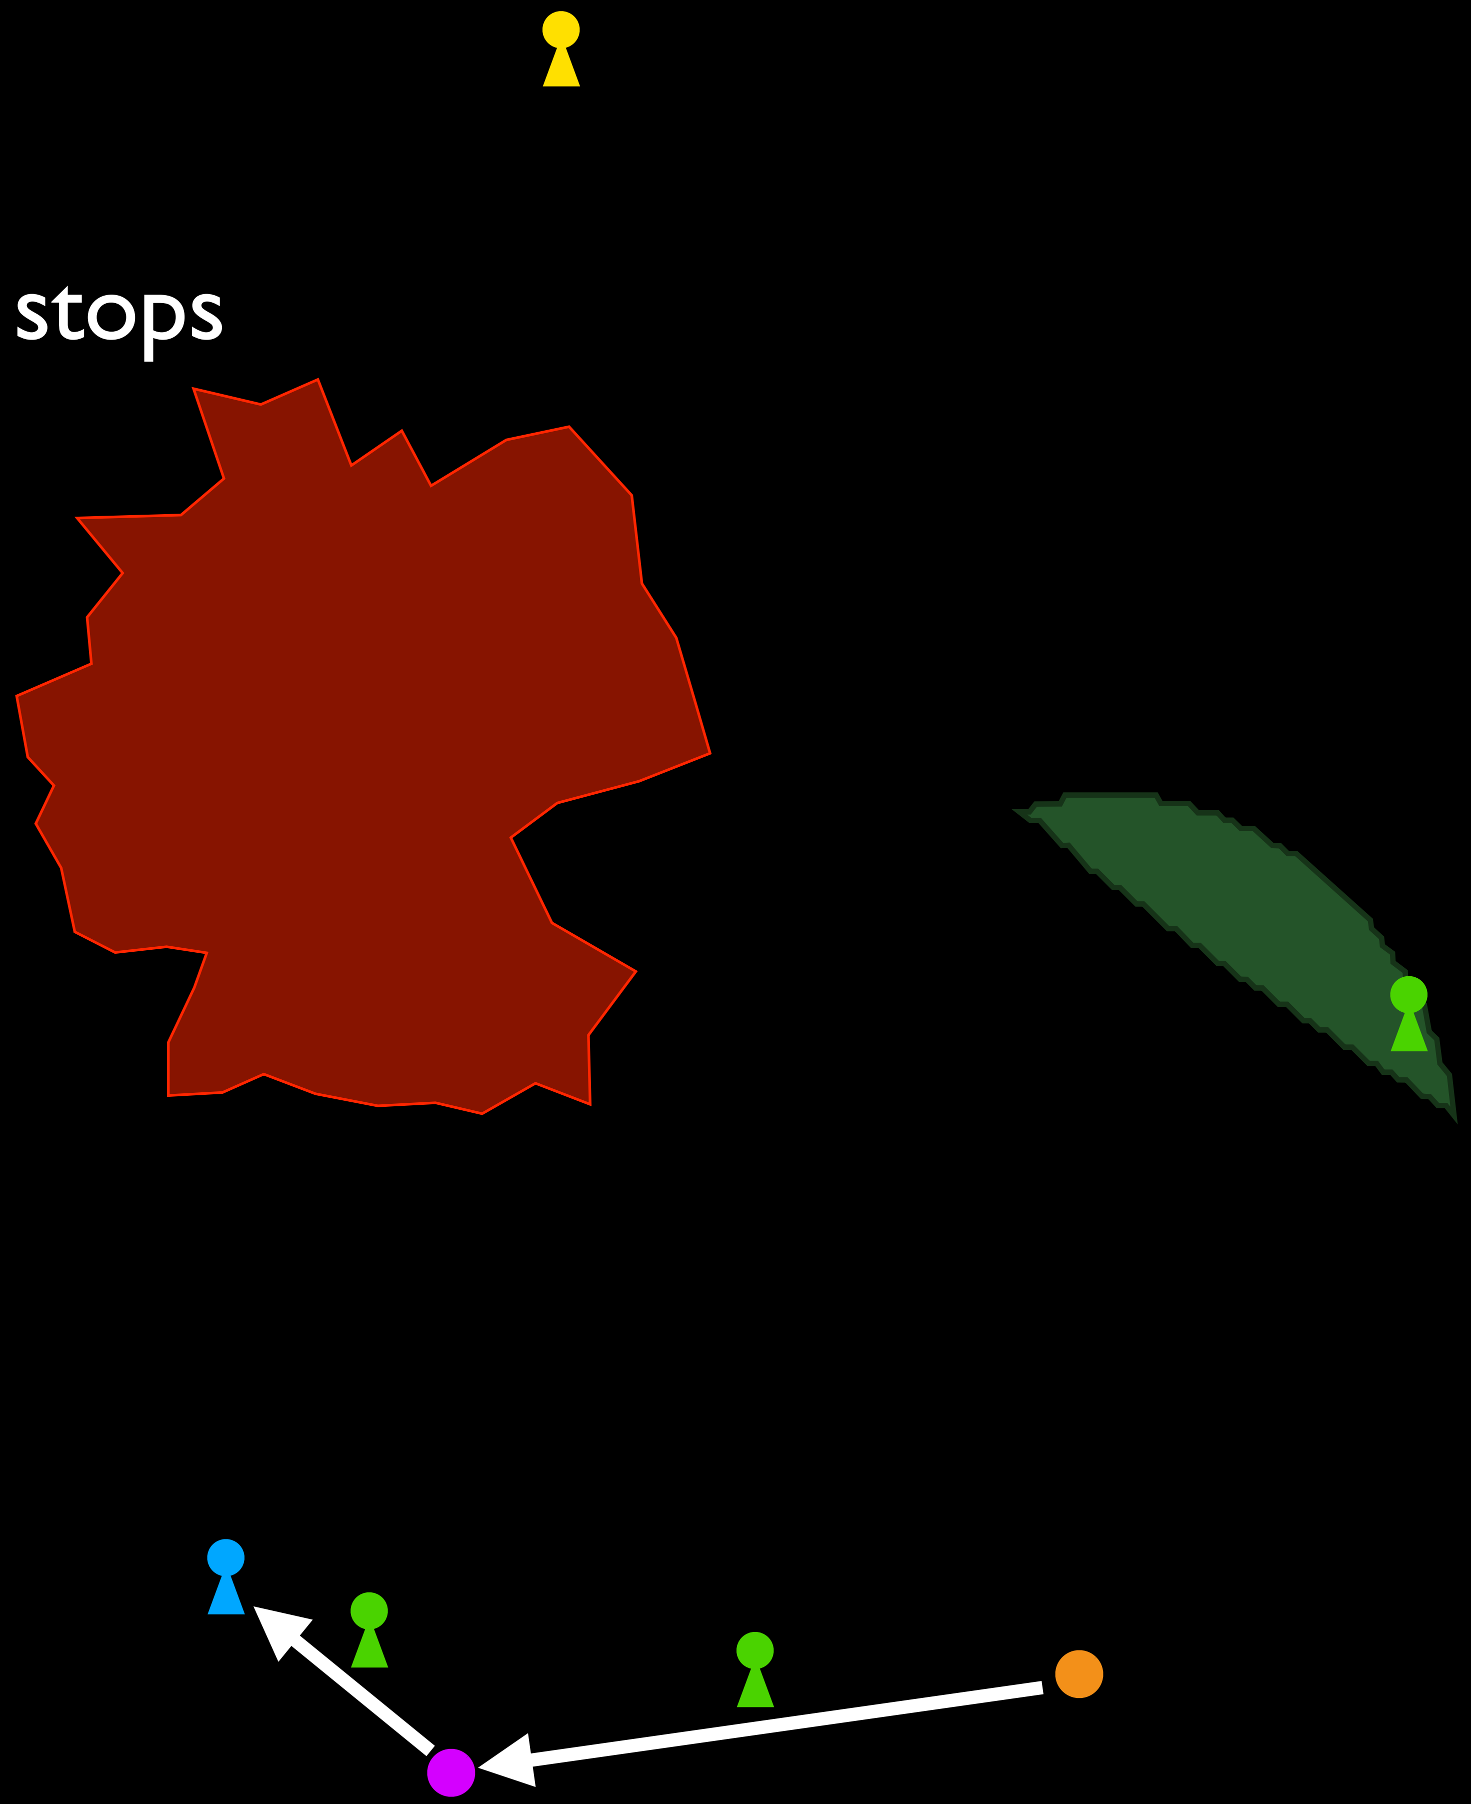

Recursive forwarding

- Forward *F* and *T*
- Continue until progress stops

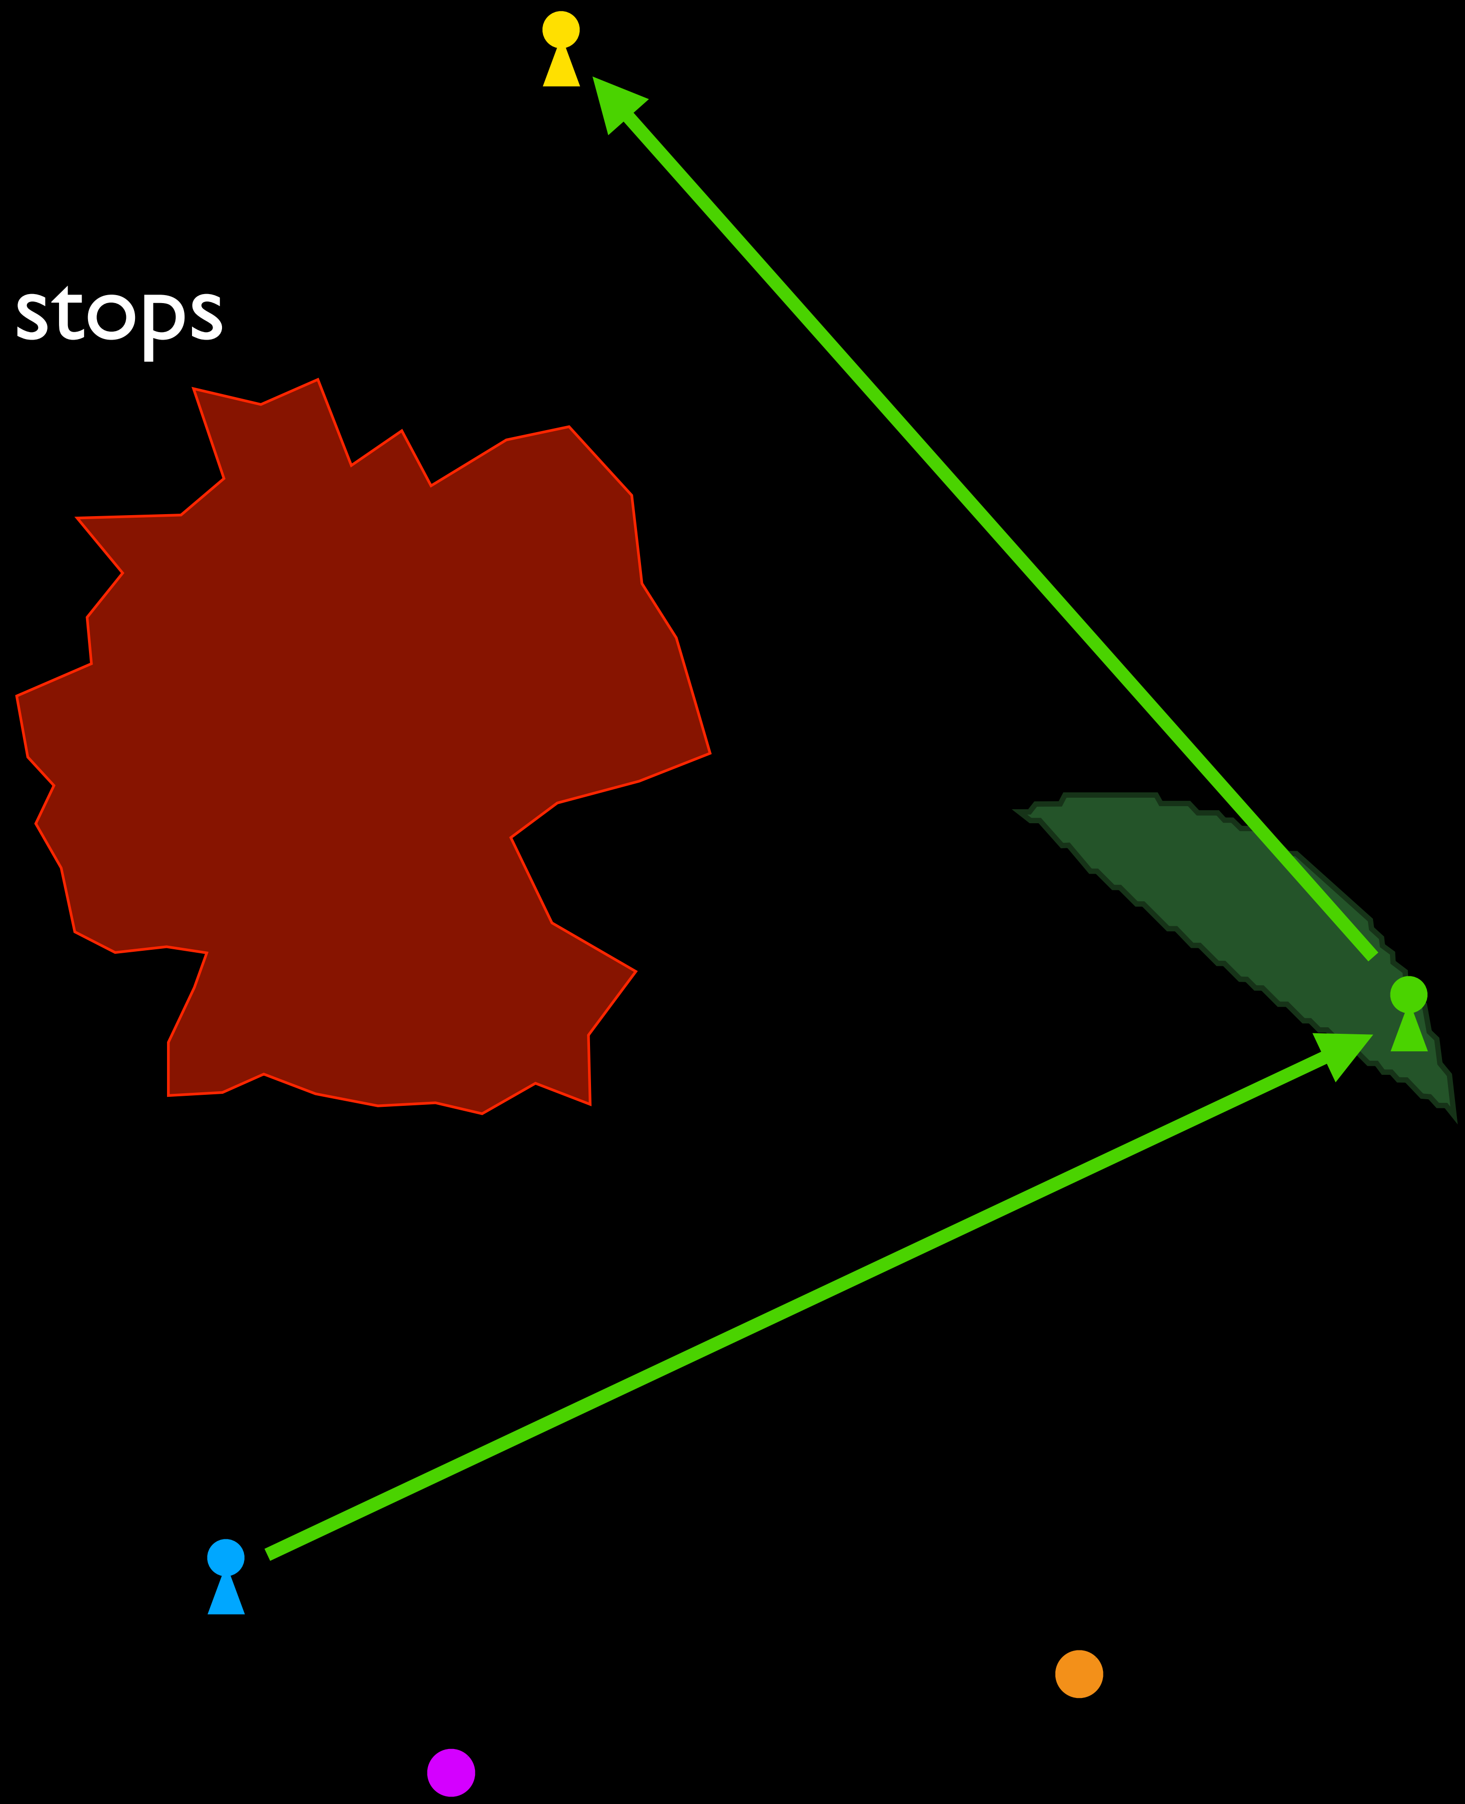

Recursive forwarding

Forward *F* and *T*

Continue until progress stops

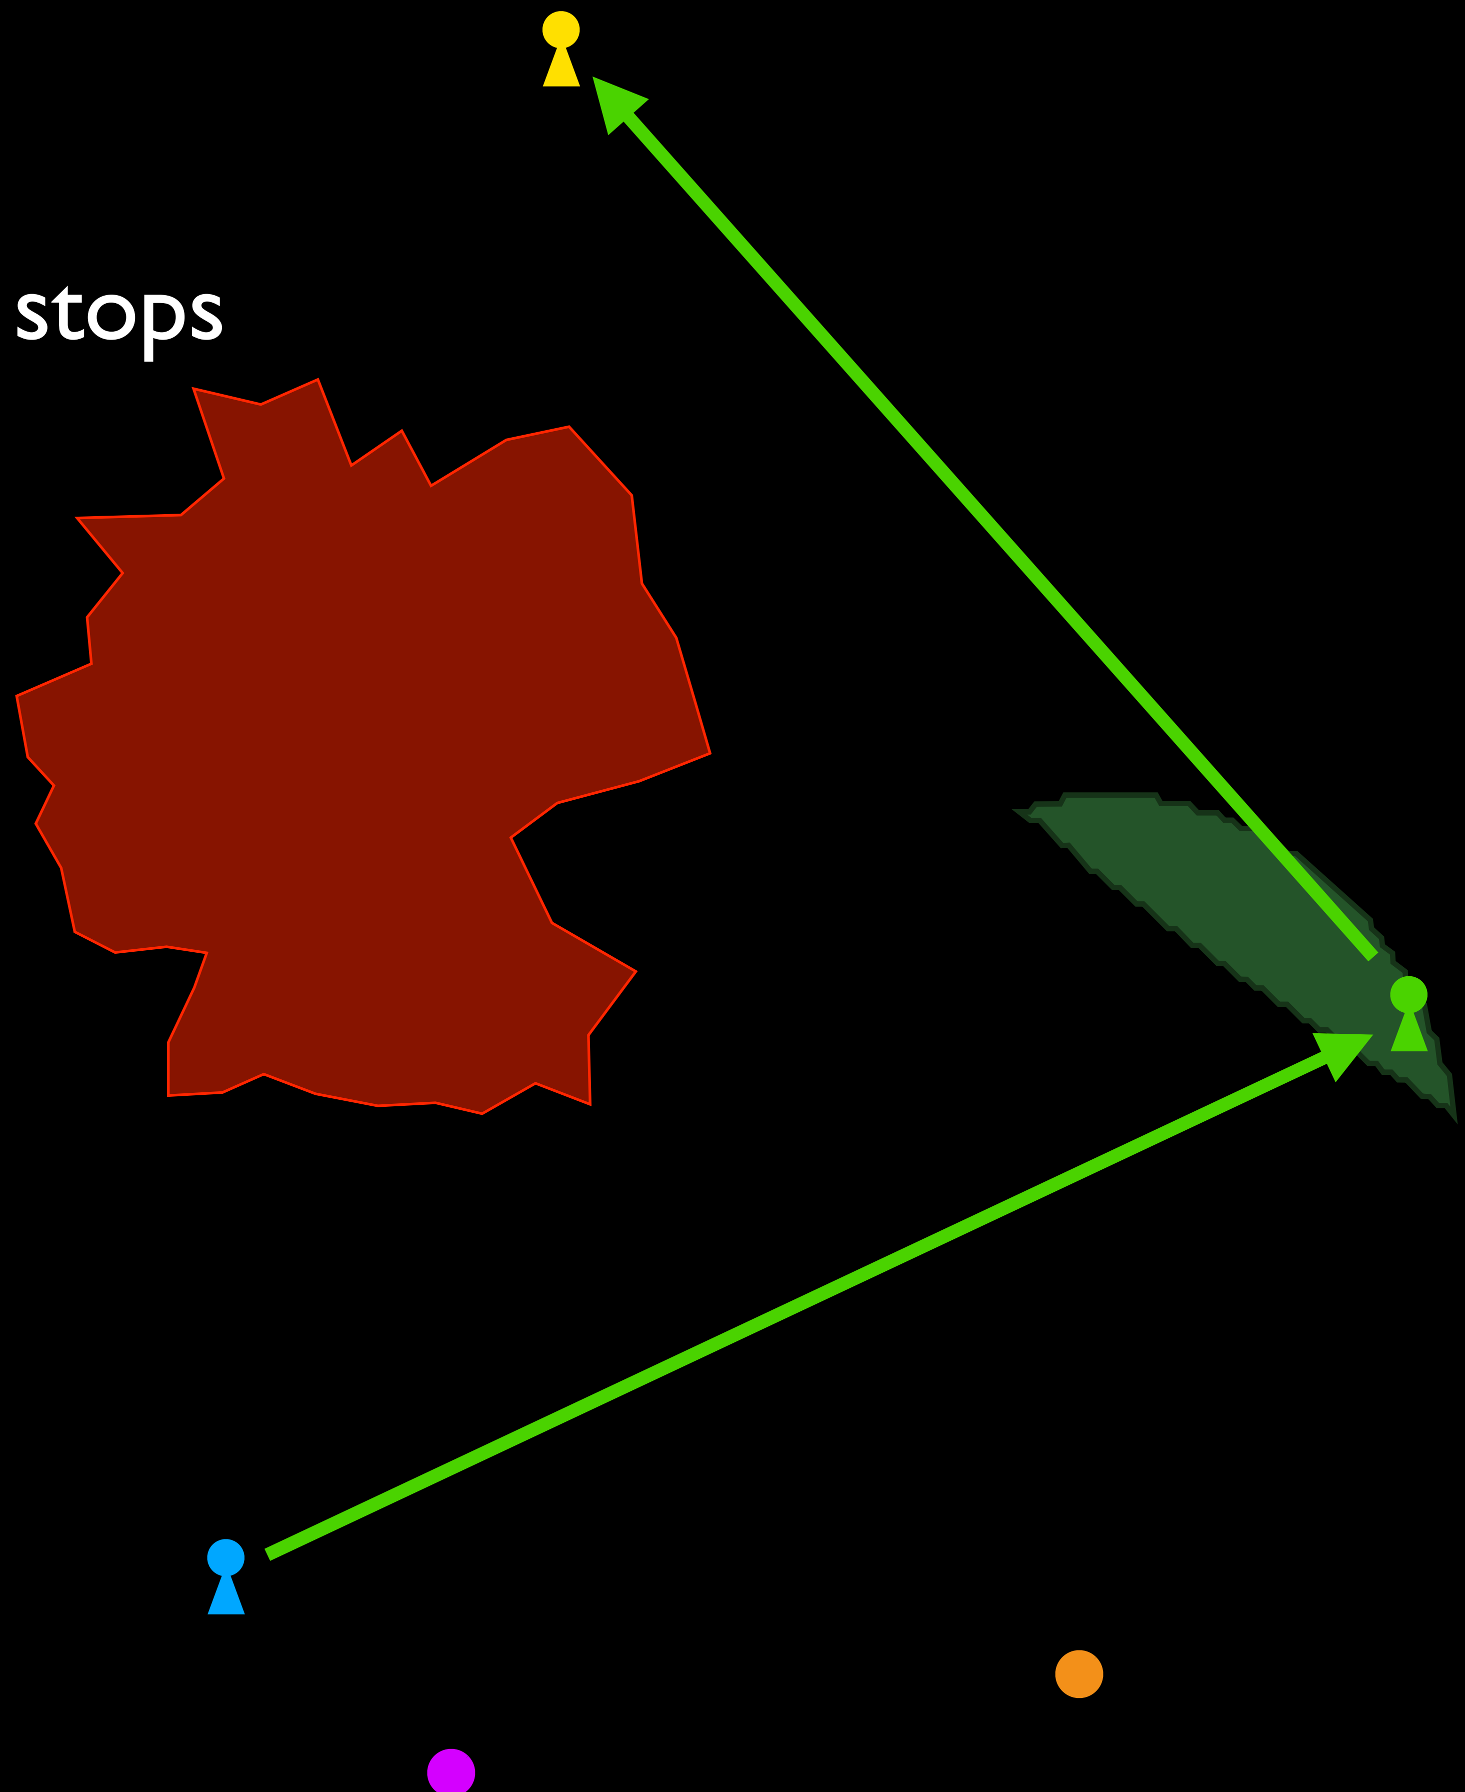

Alibi routing finds *potential* alibis

Proofs of avoidance allow verification

Implementation and Evaluation

Implementation on PlanetLab

425 nodes

Simulation (for scale)

20k nodes

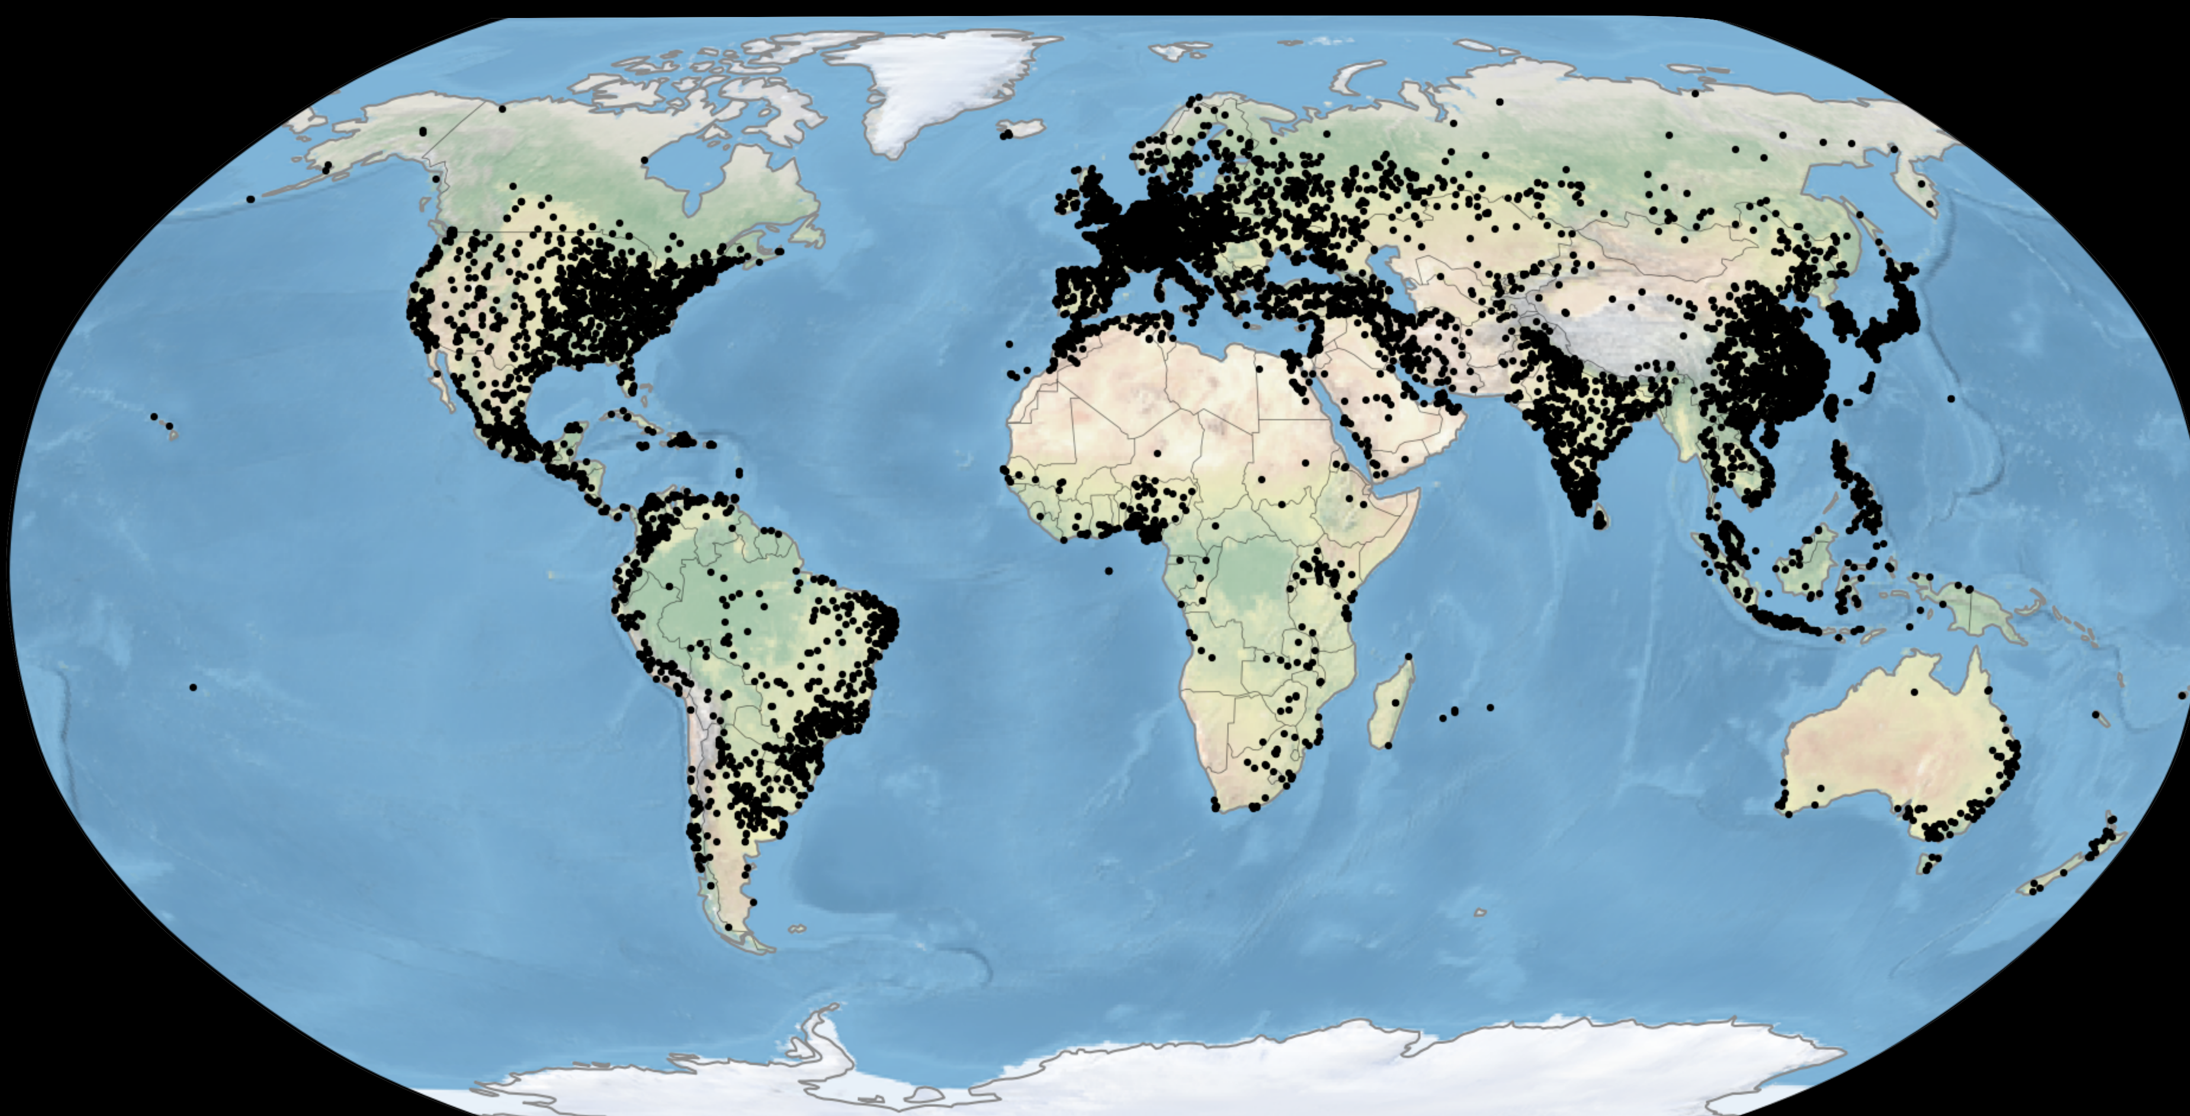

Implementation and Evaluation

Implementation on PlanetLab

425 nodes

Simulation (for scale)

20k nodes

China Iran PR Korea Syria Saudi Arabia

Known censors (Reporters without Borders)

India Japan USA

Most Internet users

Implementation and Evaluation

Implementation on PlanetLab

425 nodes

Simulation (for scale)

20k nodes

China Iran PR Korea Syria Saudi Arabia

Known censors (Reporters without Borders)

India Japan USA

Most Internet users

Alibi Routing success rates

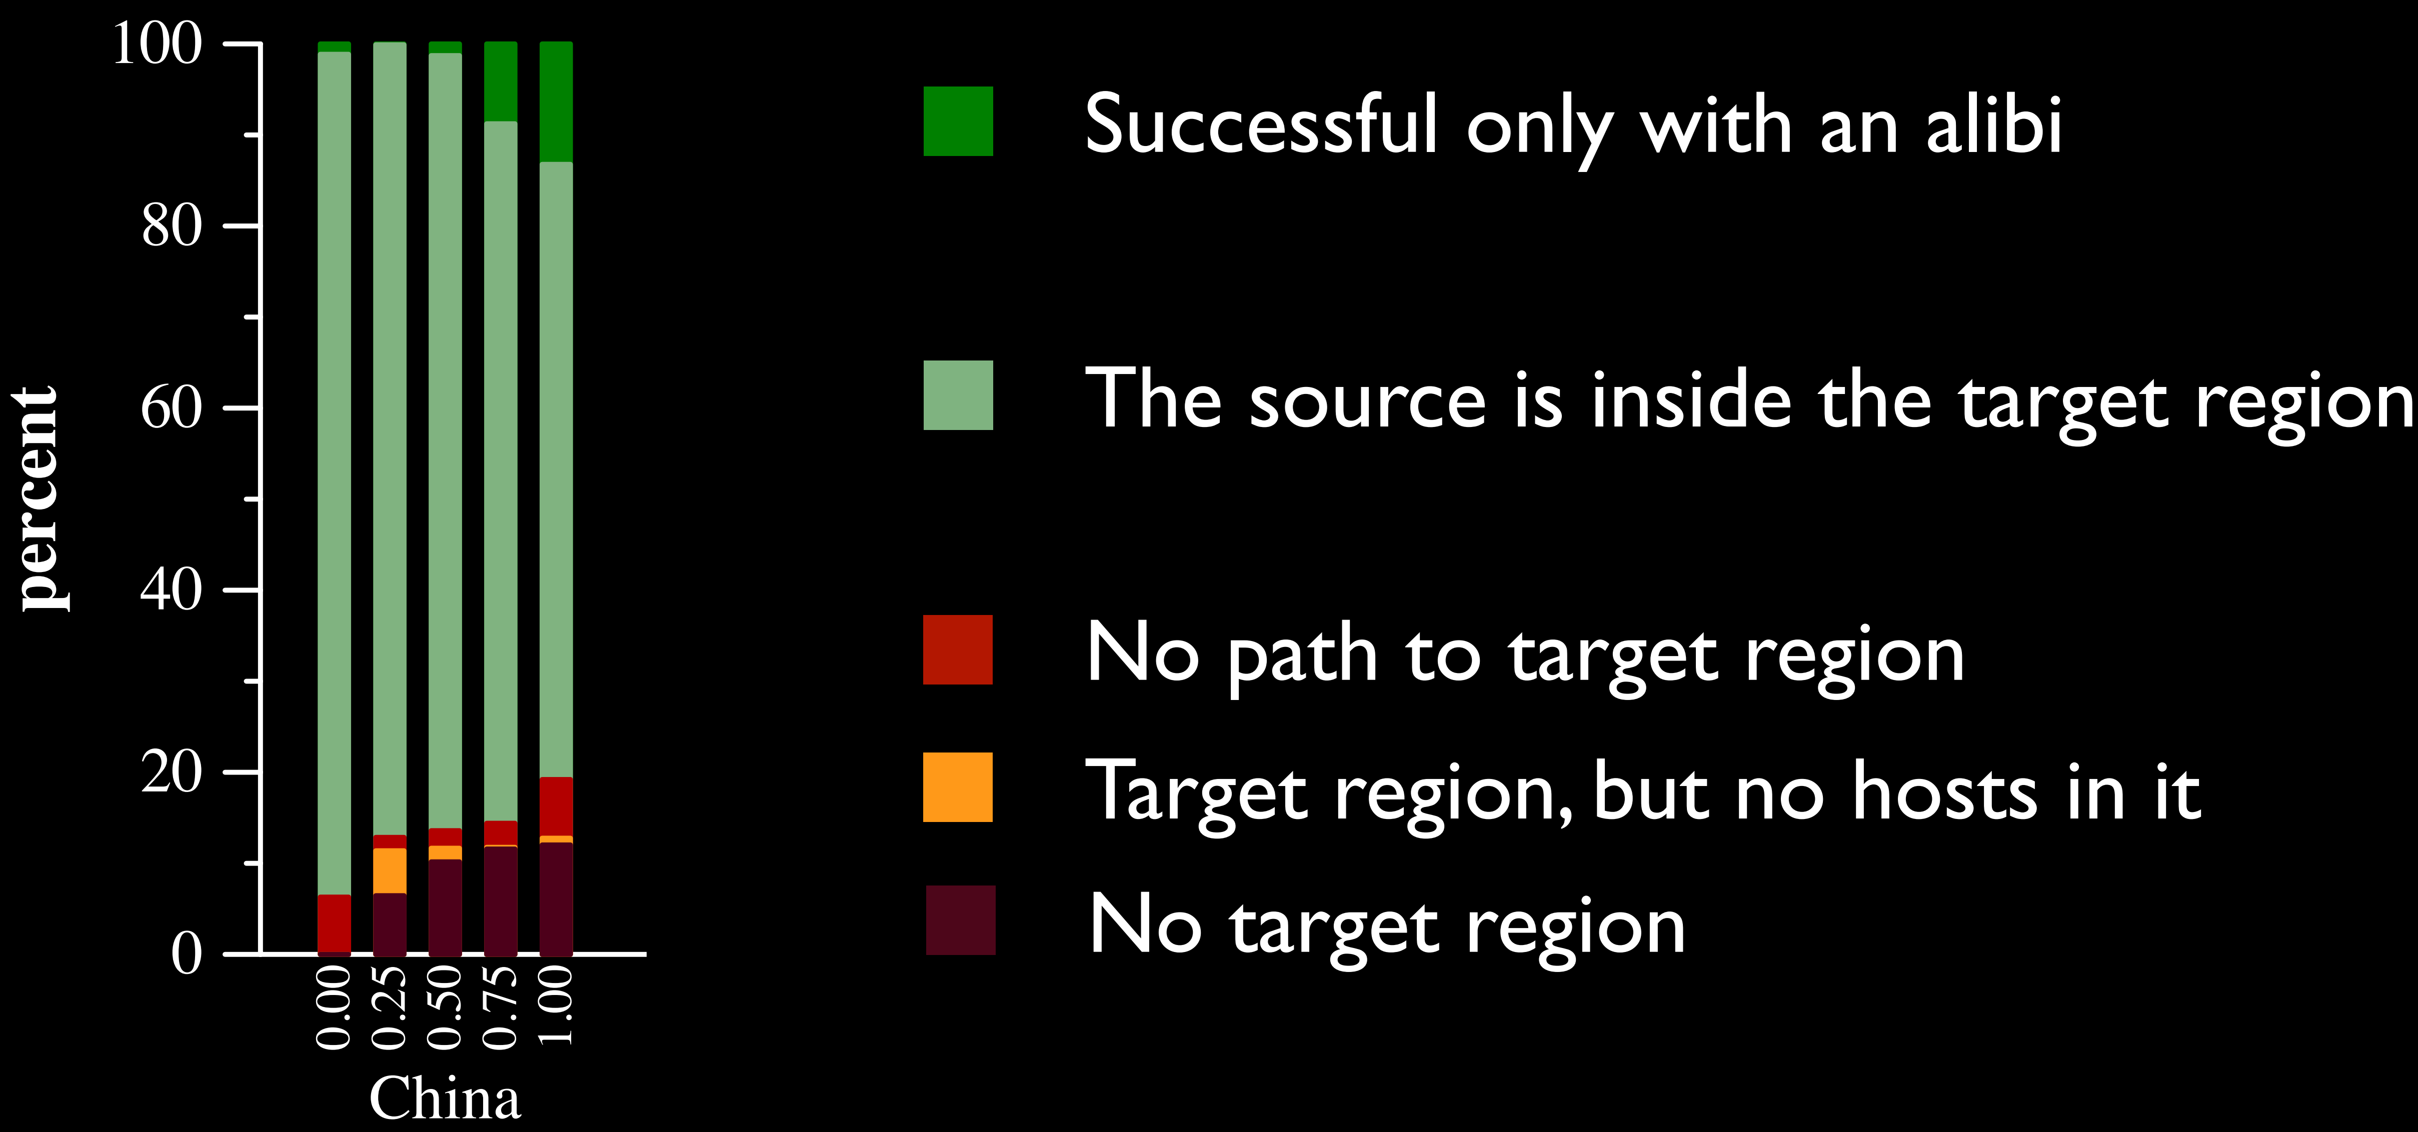

Alibi Routing success rates

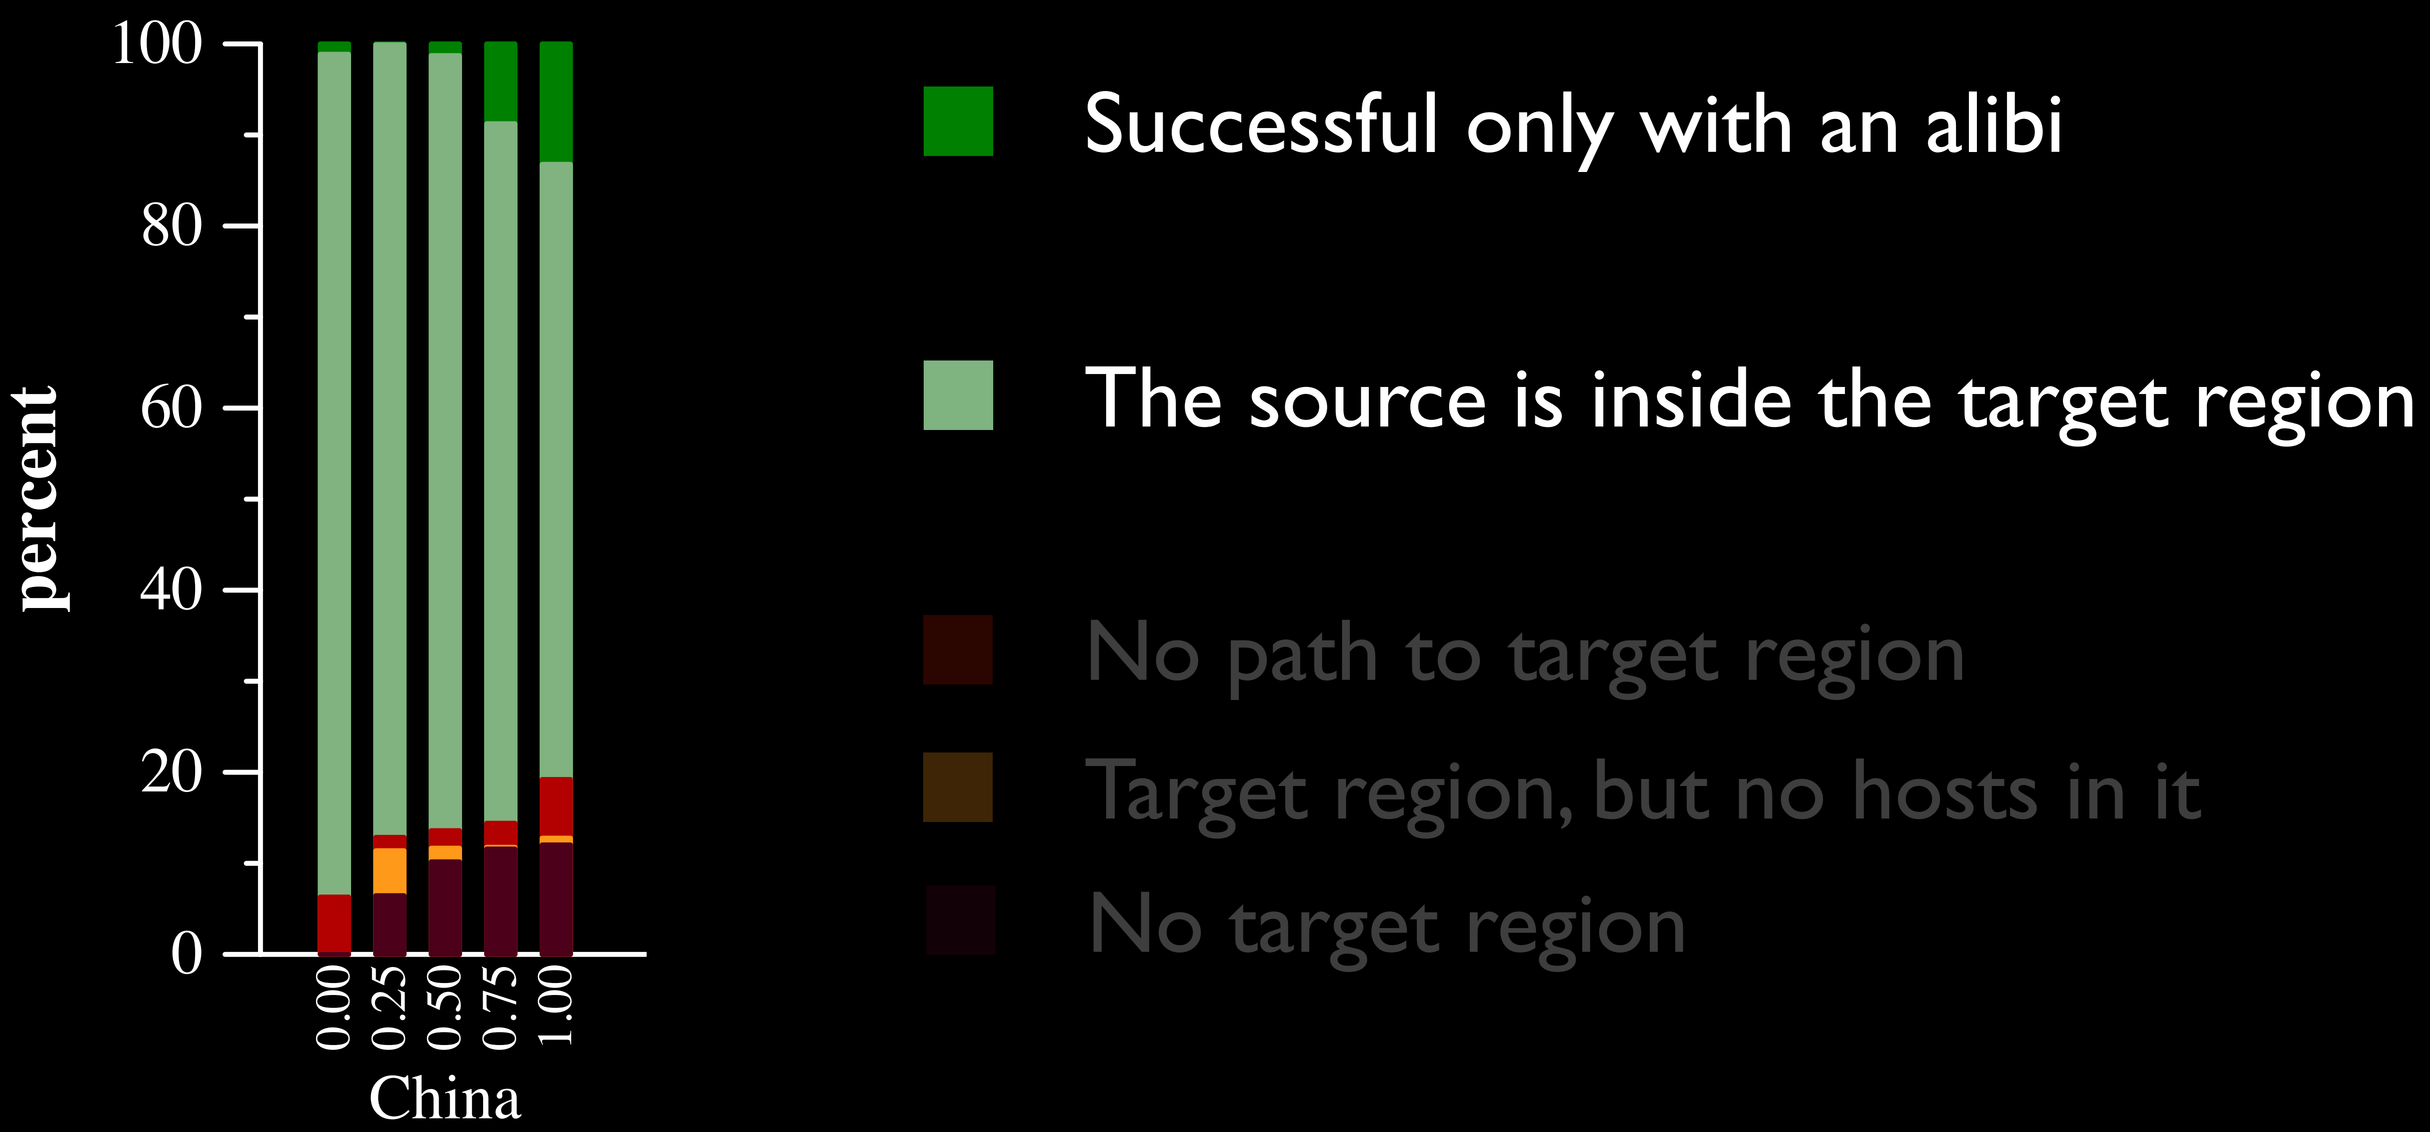

Most src-dst pairs can provably avoid

Alibi Routing success rates

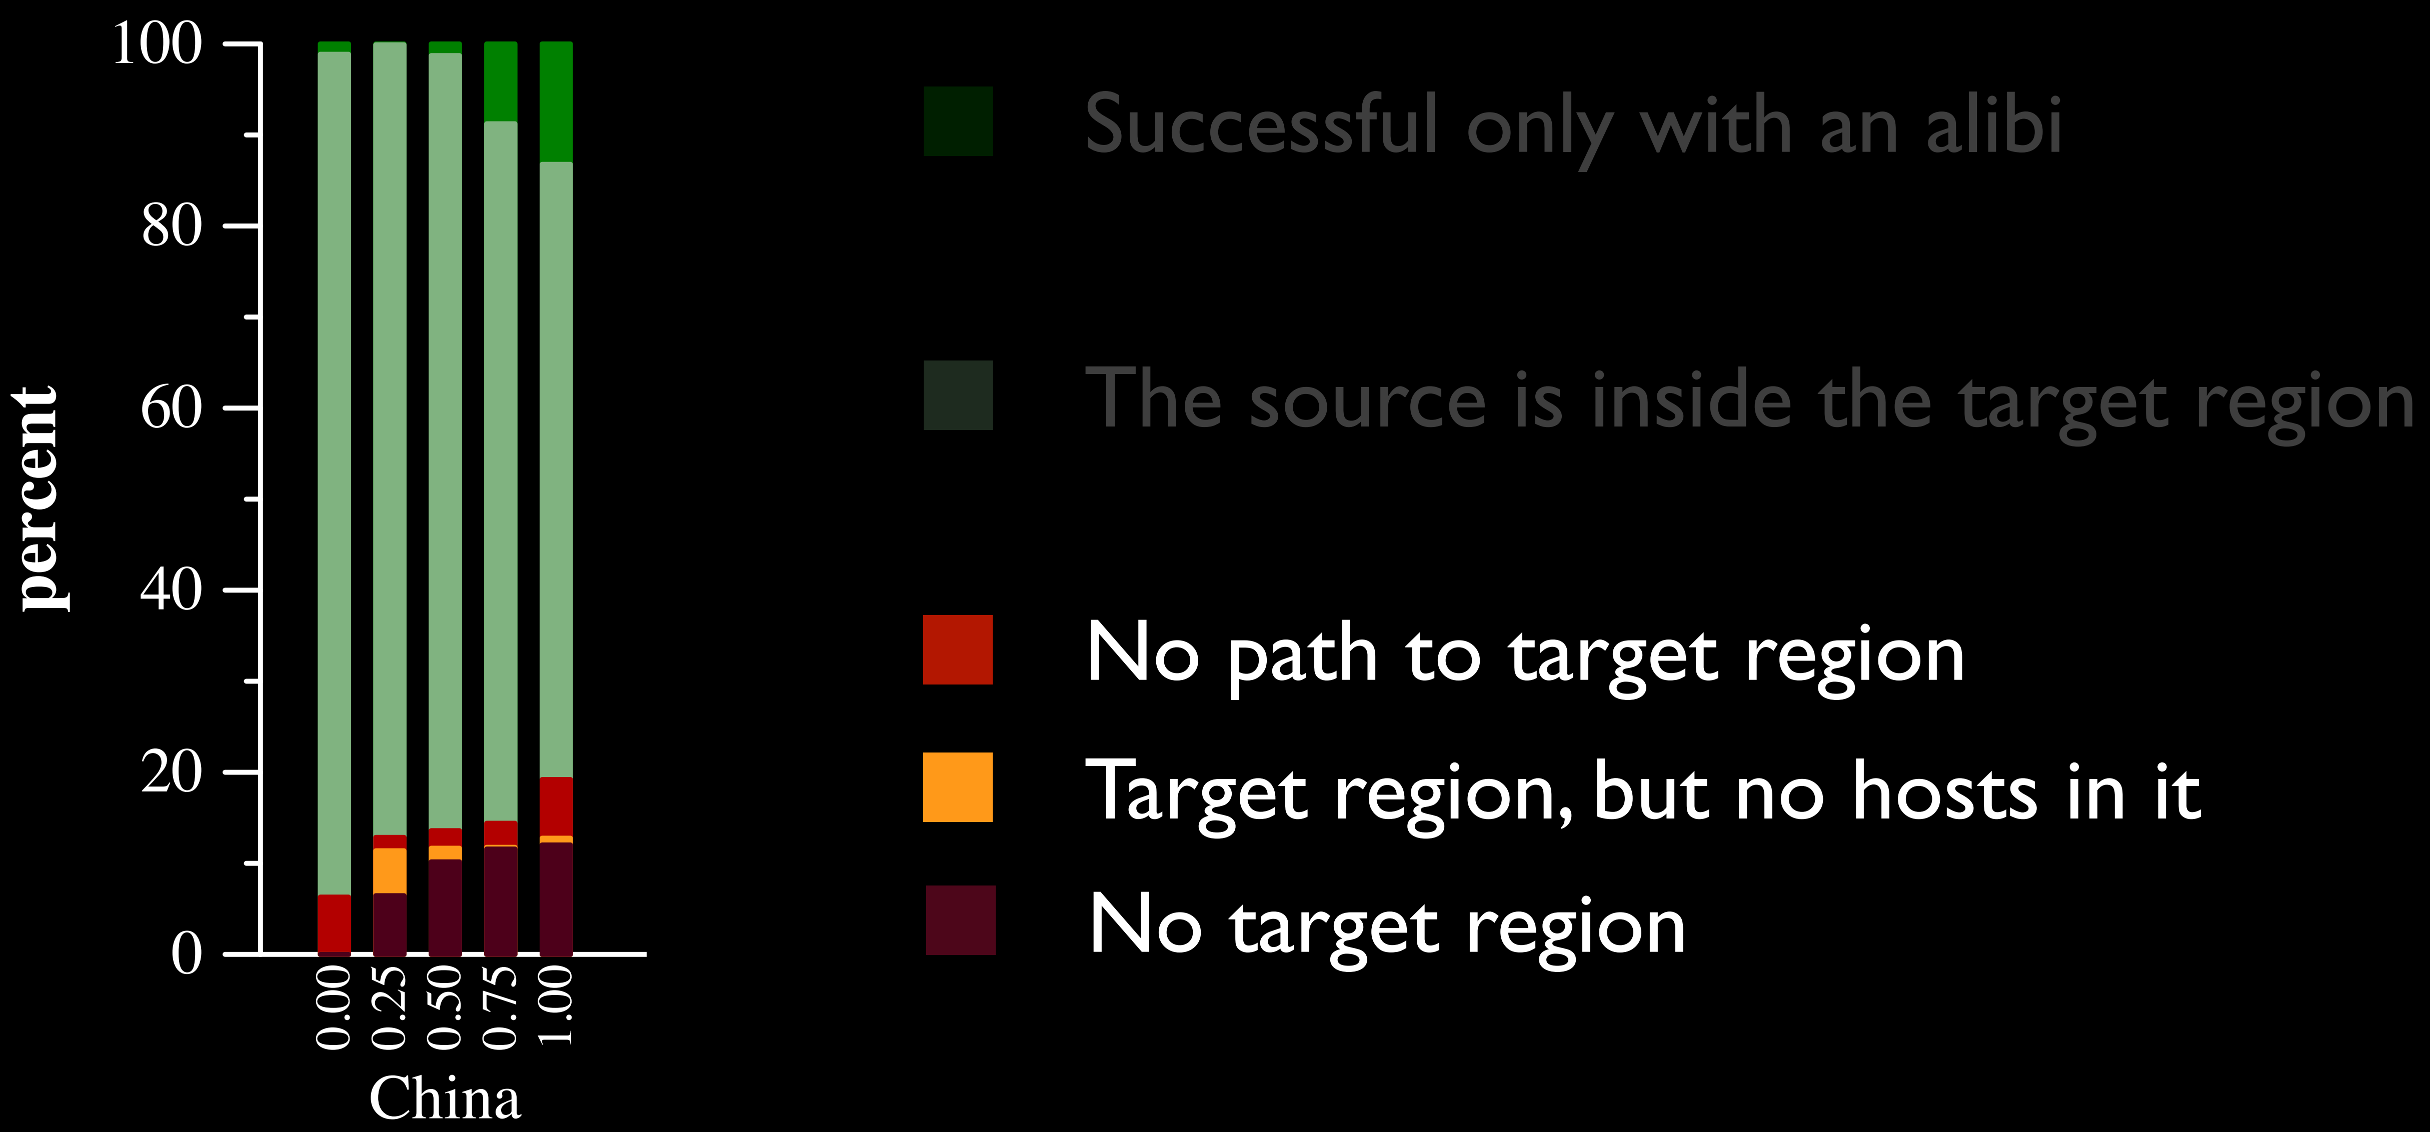

Most src-dst pairs can provably avoid

Failure typically arises when the target region is too small or non-existent

Alibi Routing success rates

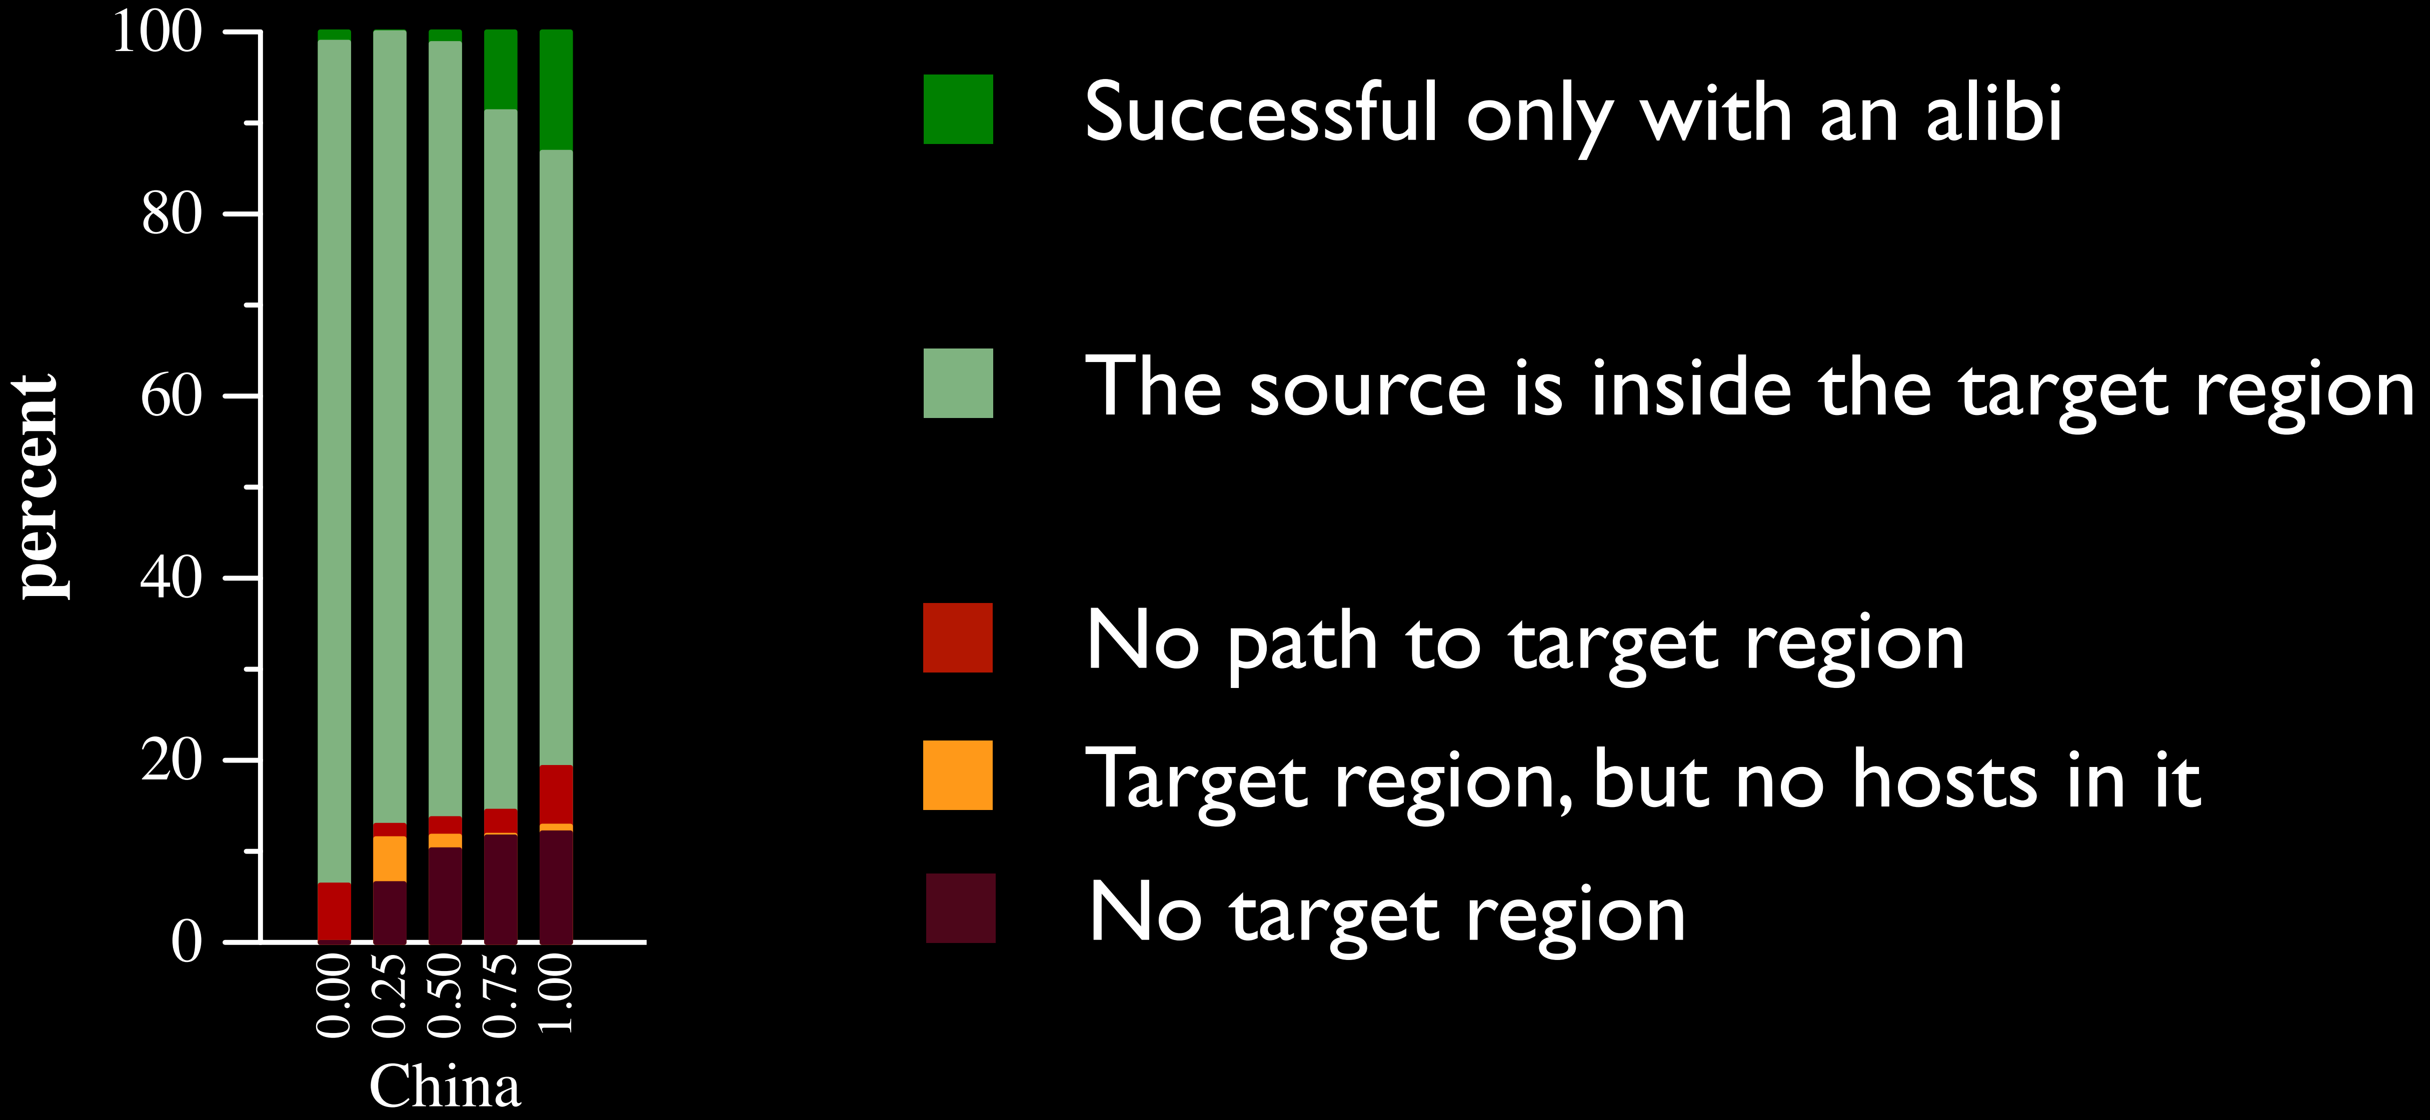

Most src-dst pairs can provably avoid

Failure typically arises when the target region is too small or non-existent

Proximity's effect on target regions

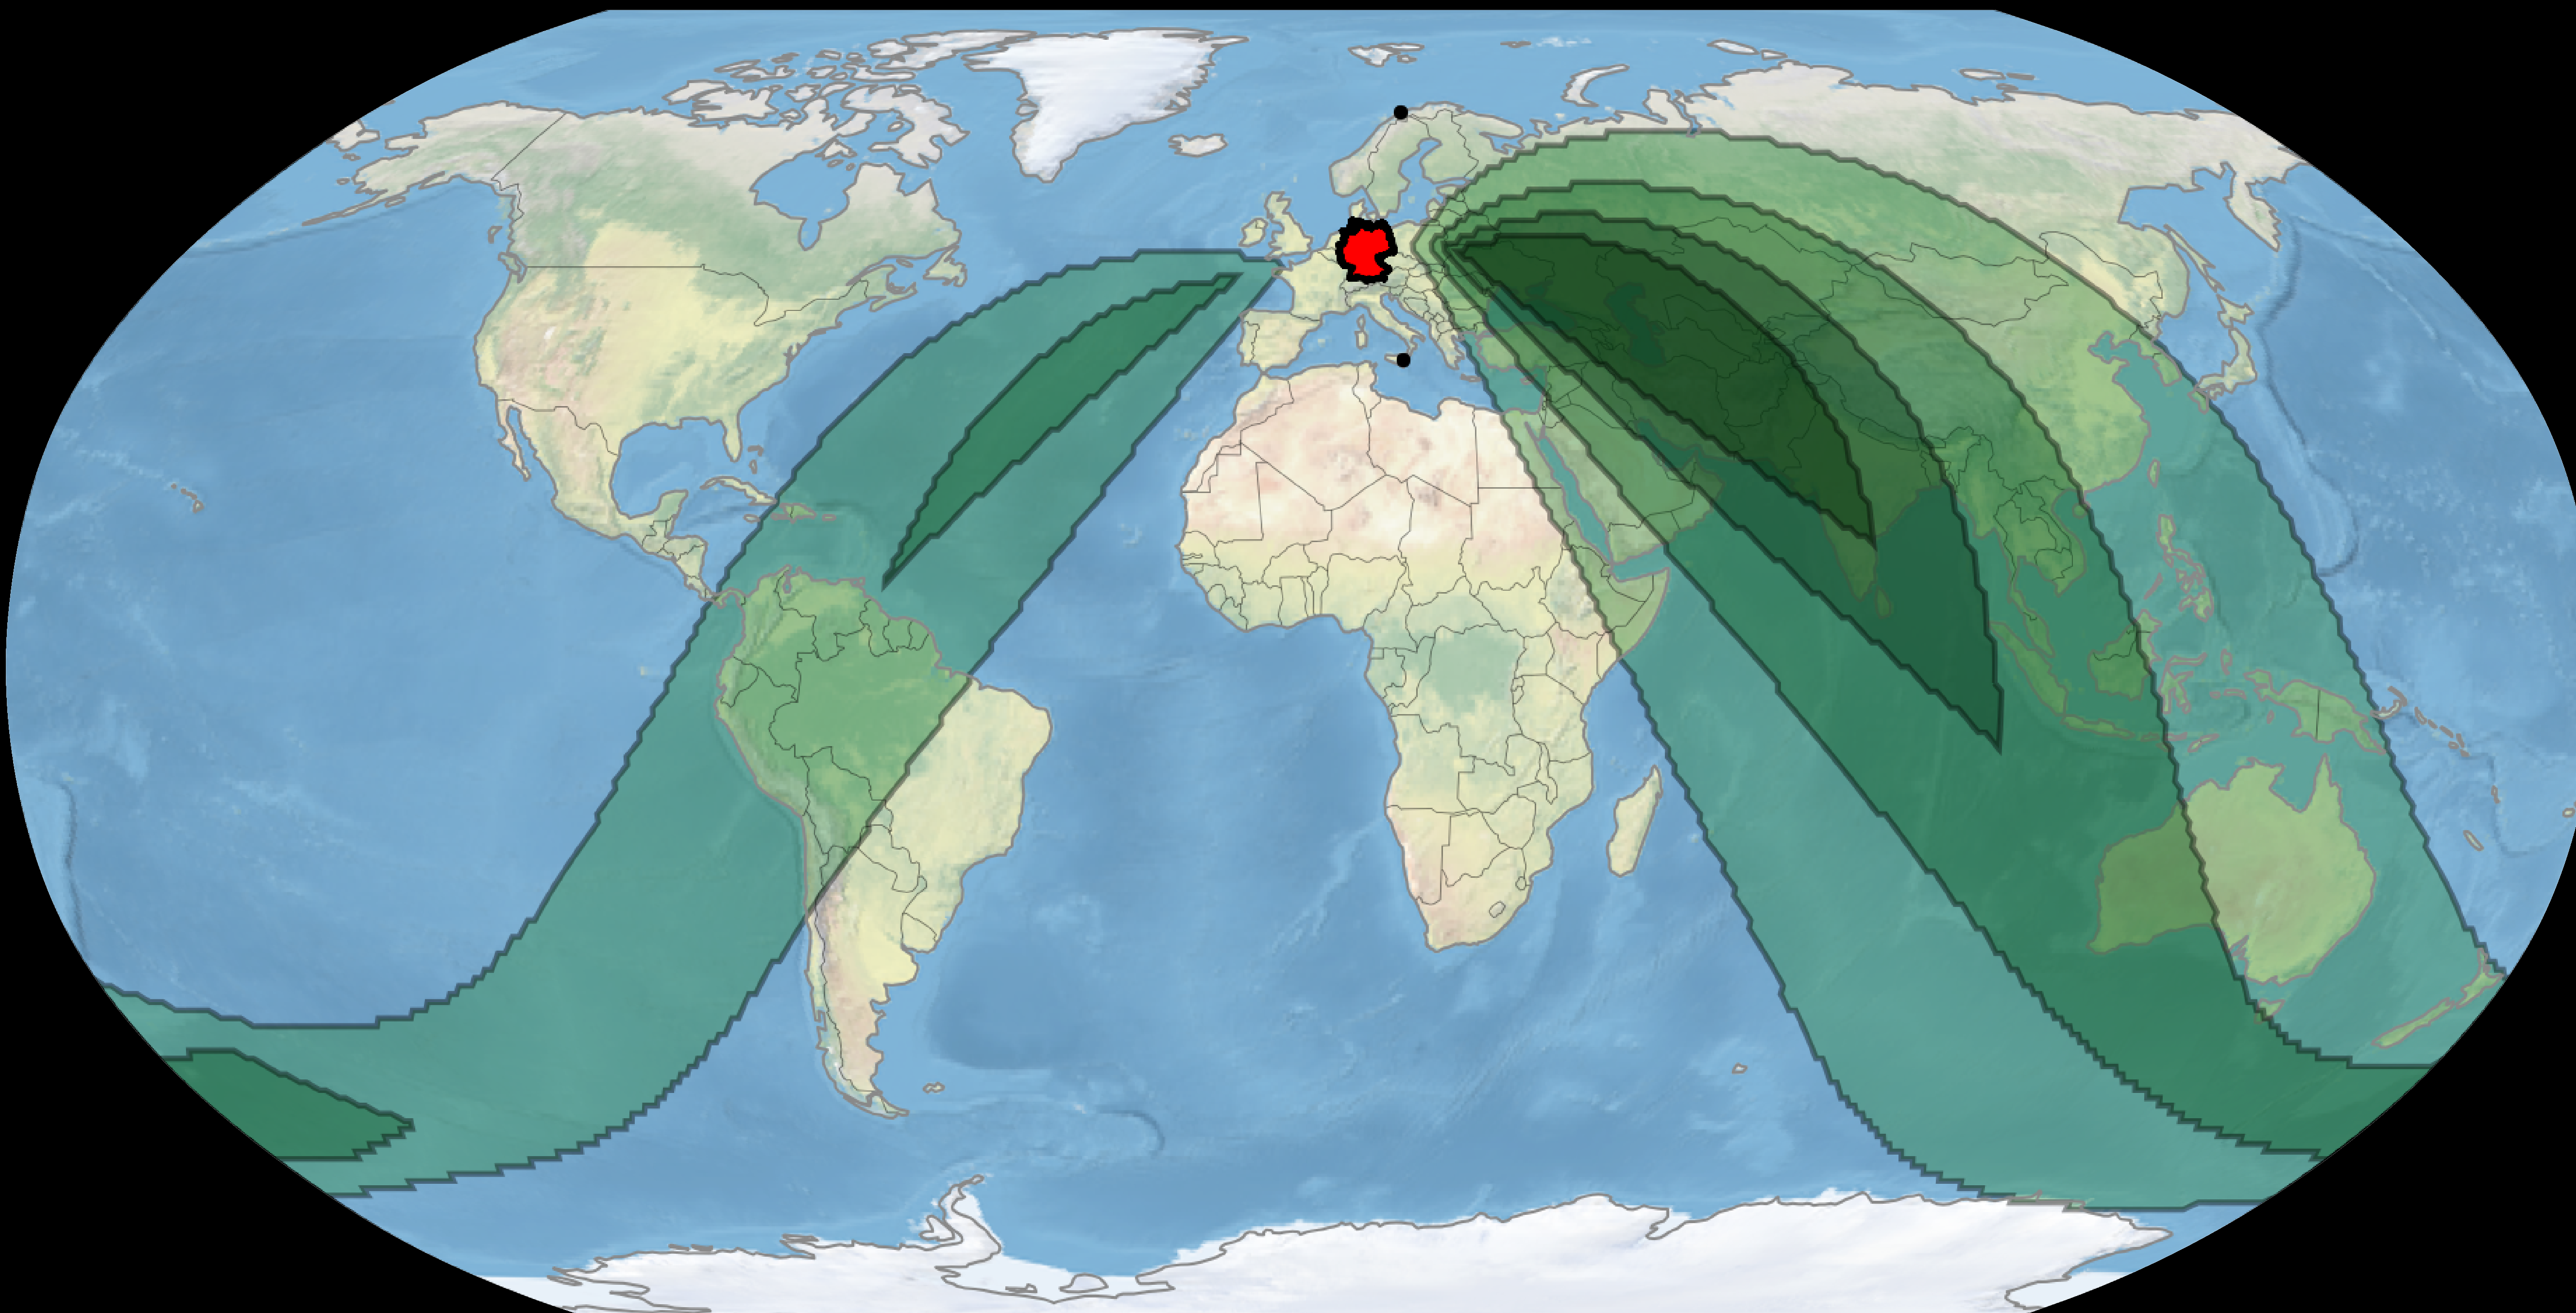

Proximity's effect on target regions

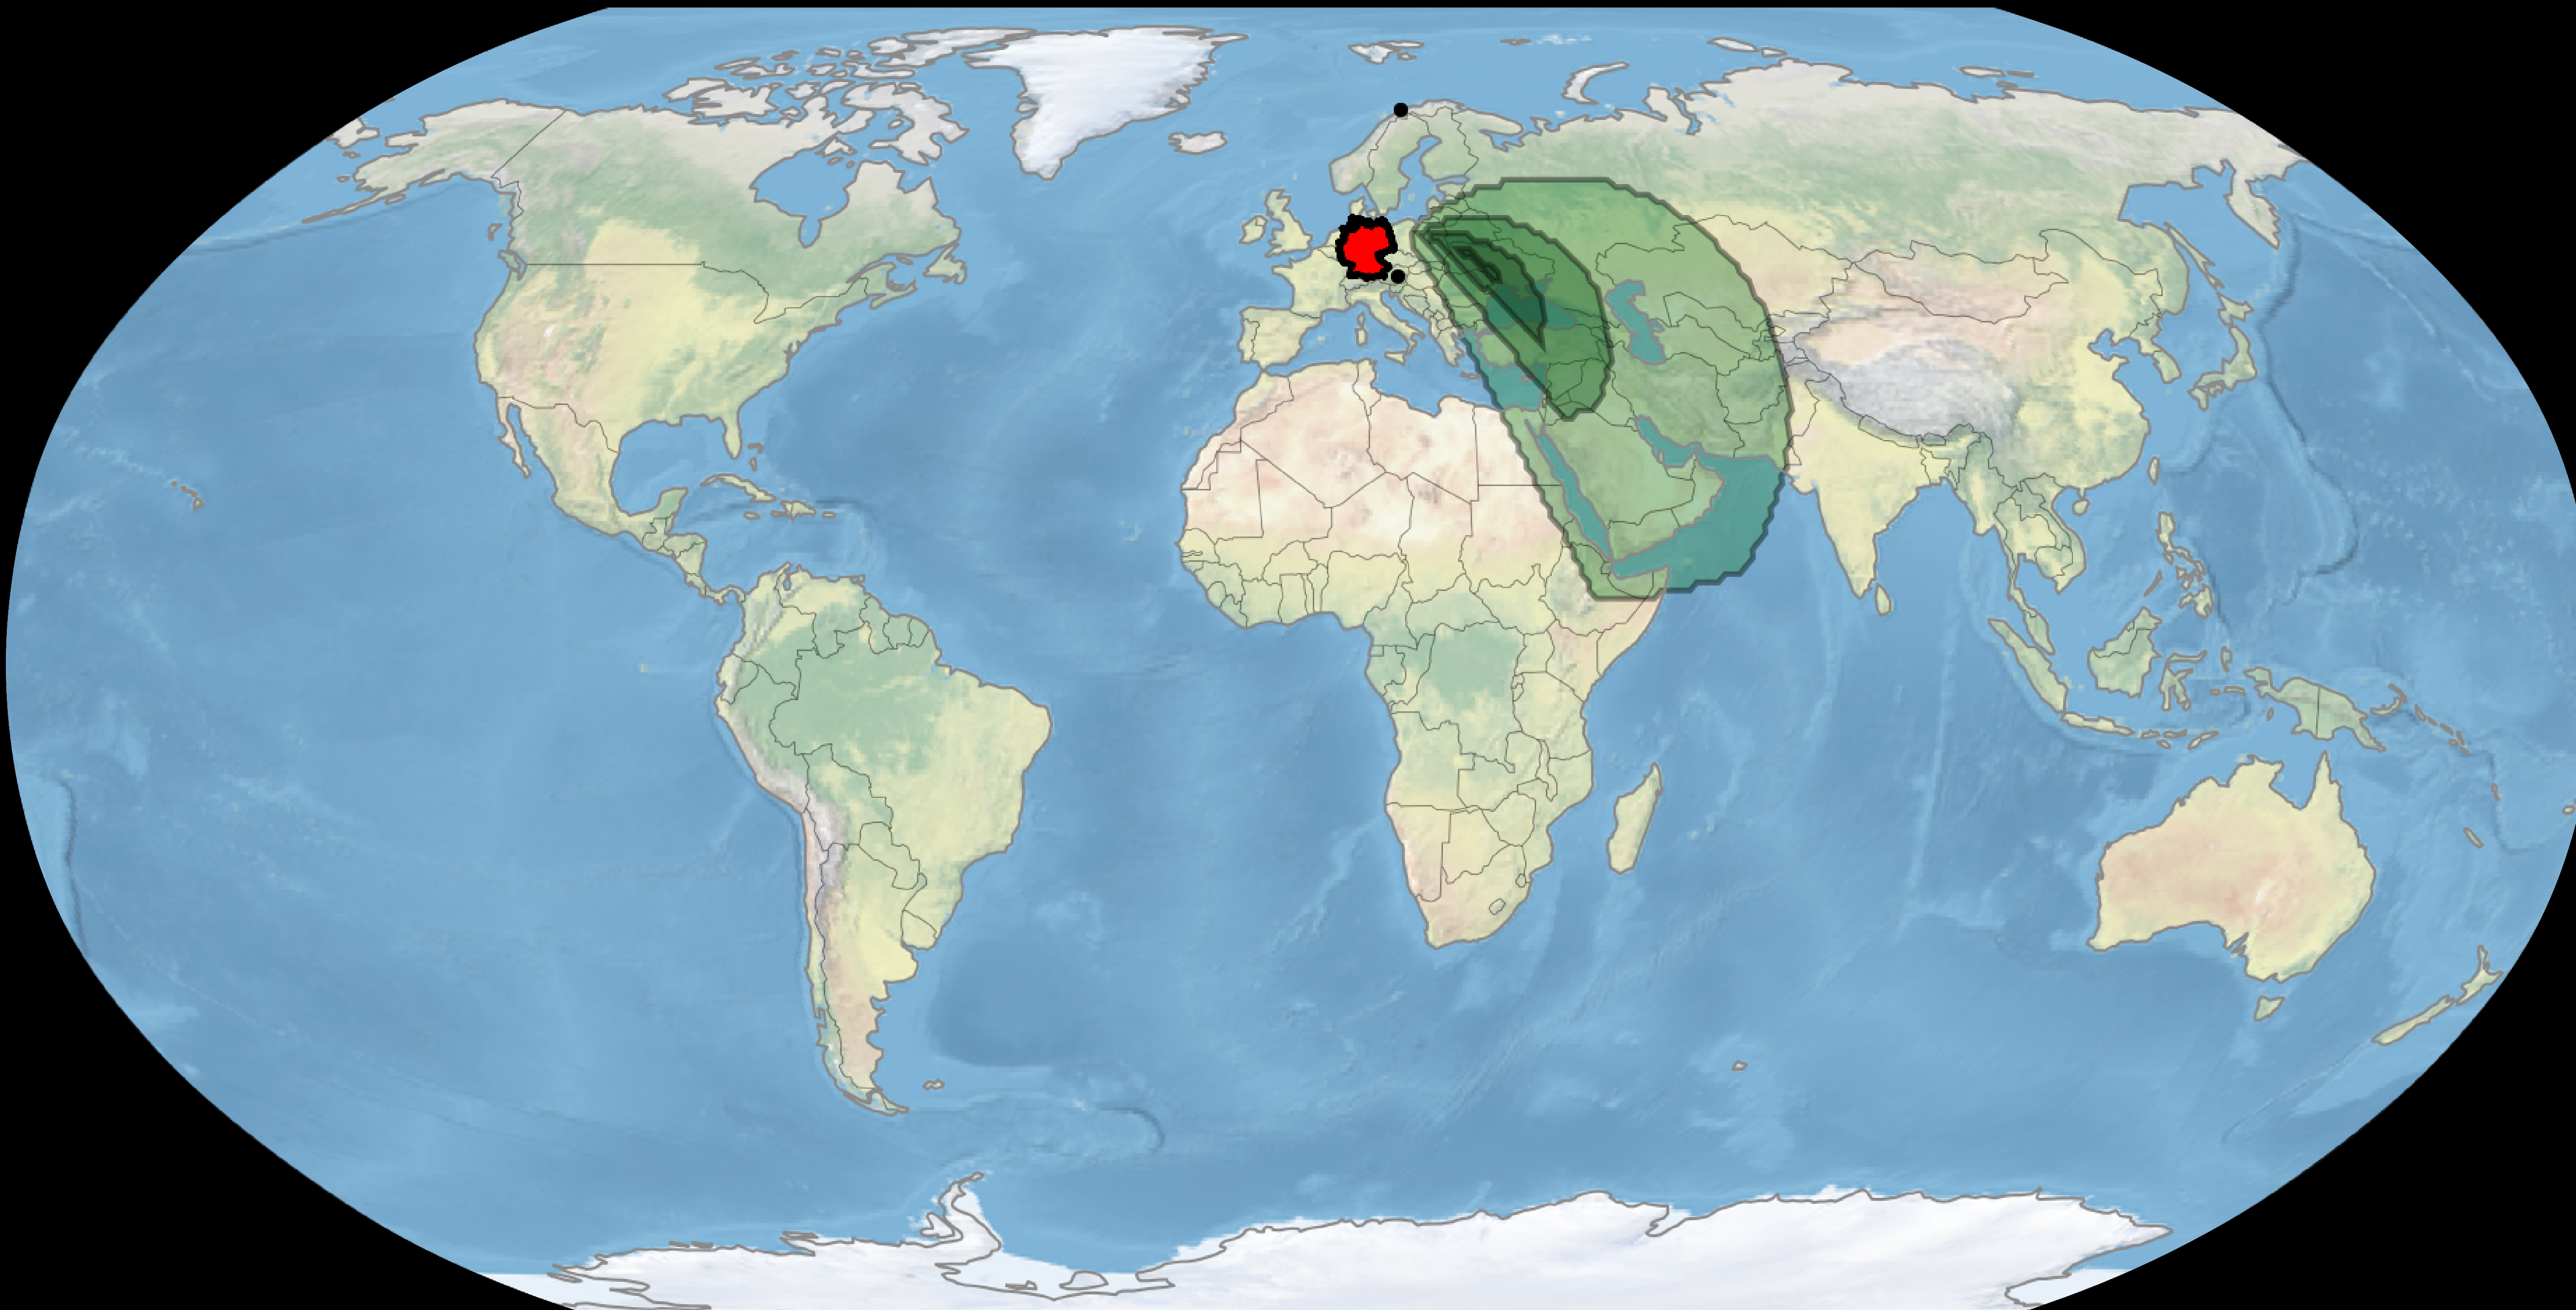

Proximity's effect on target regions

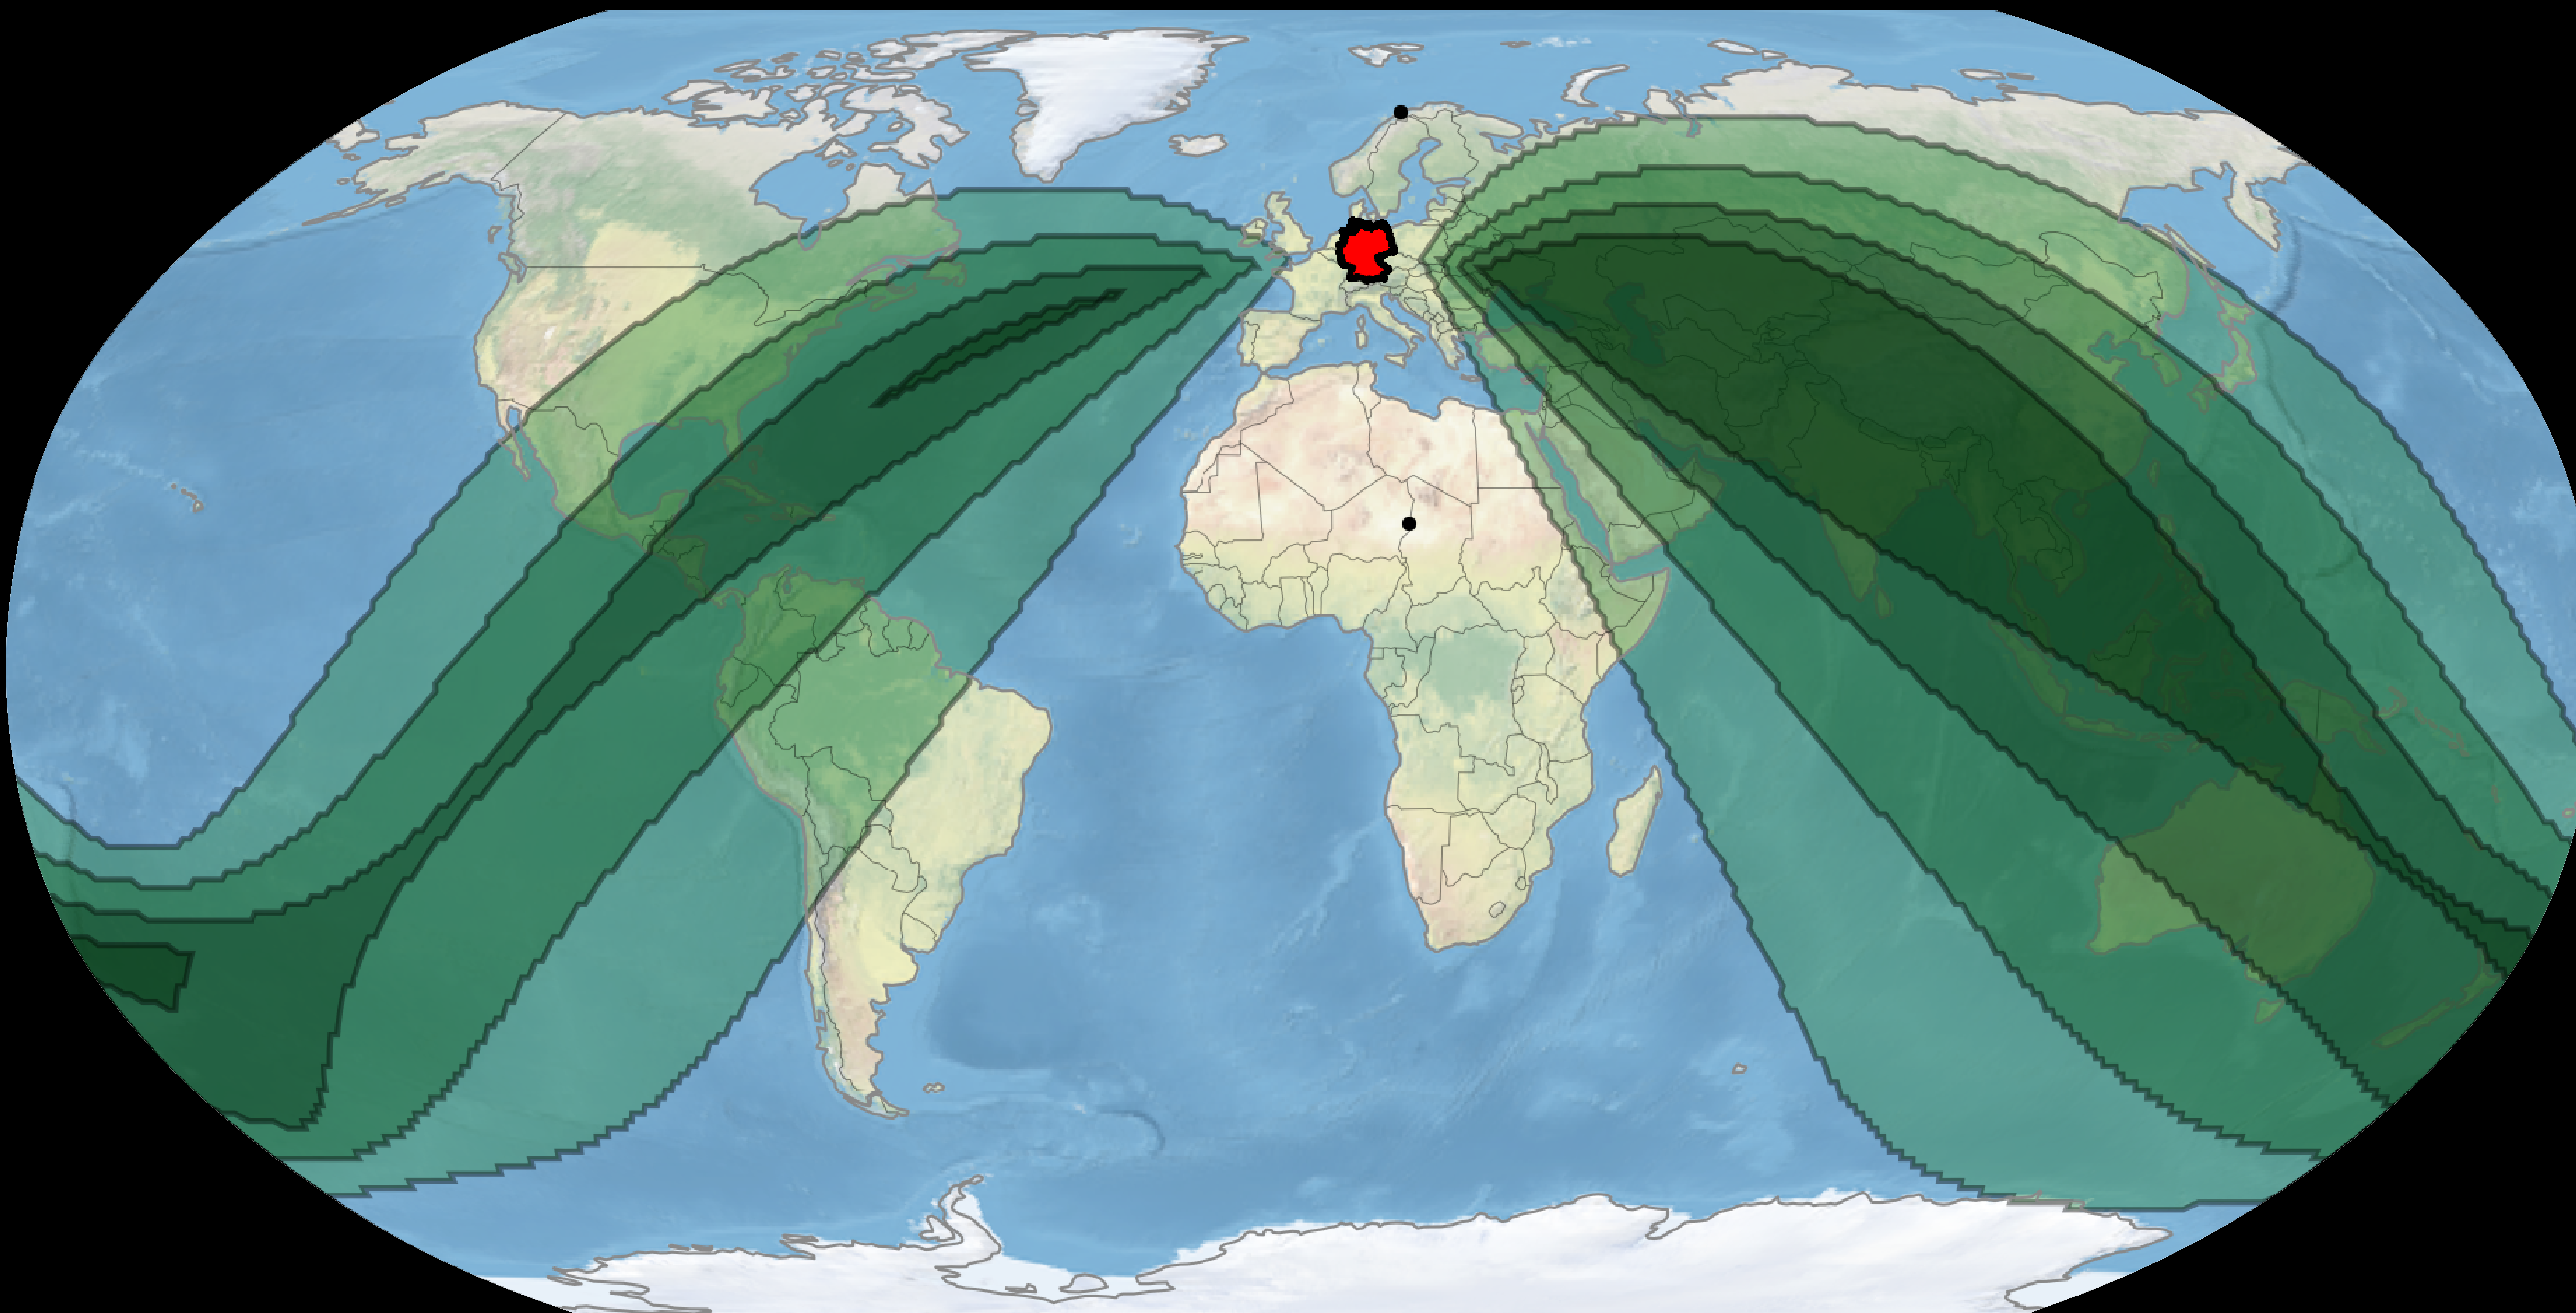

Proximity's effect on target regions

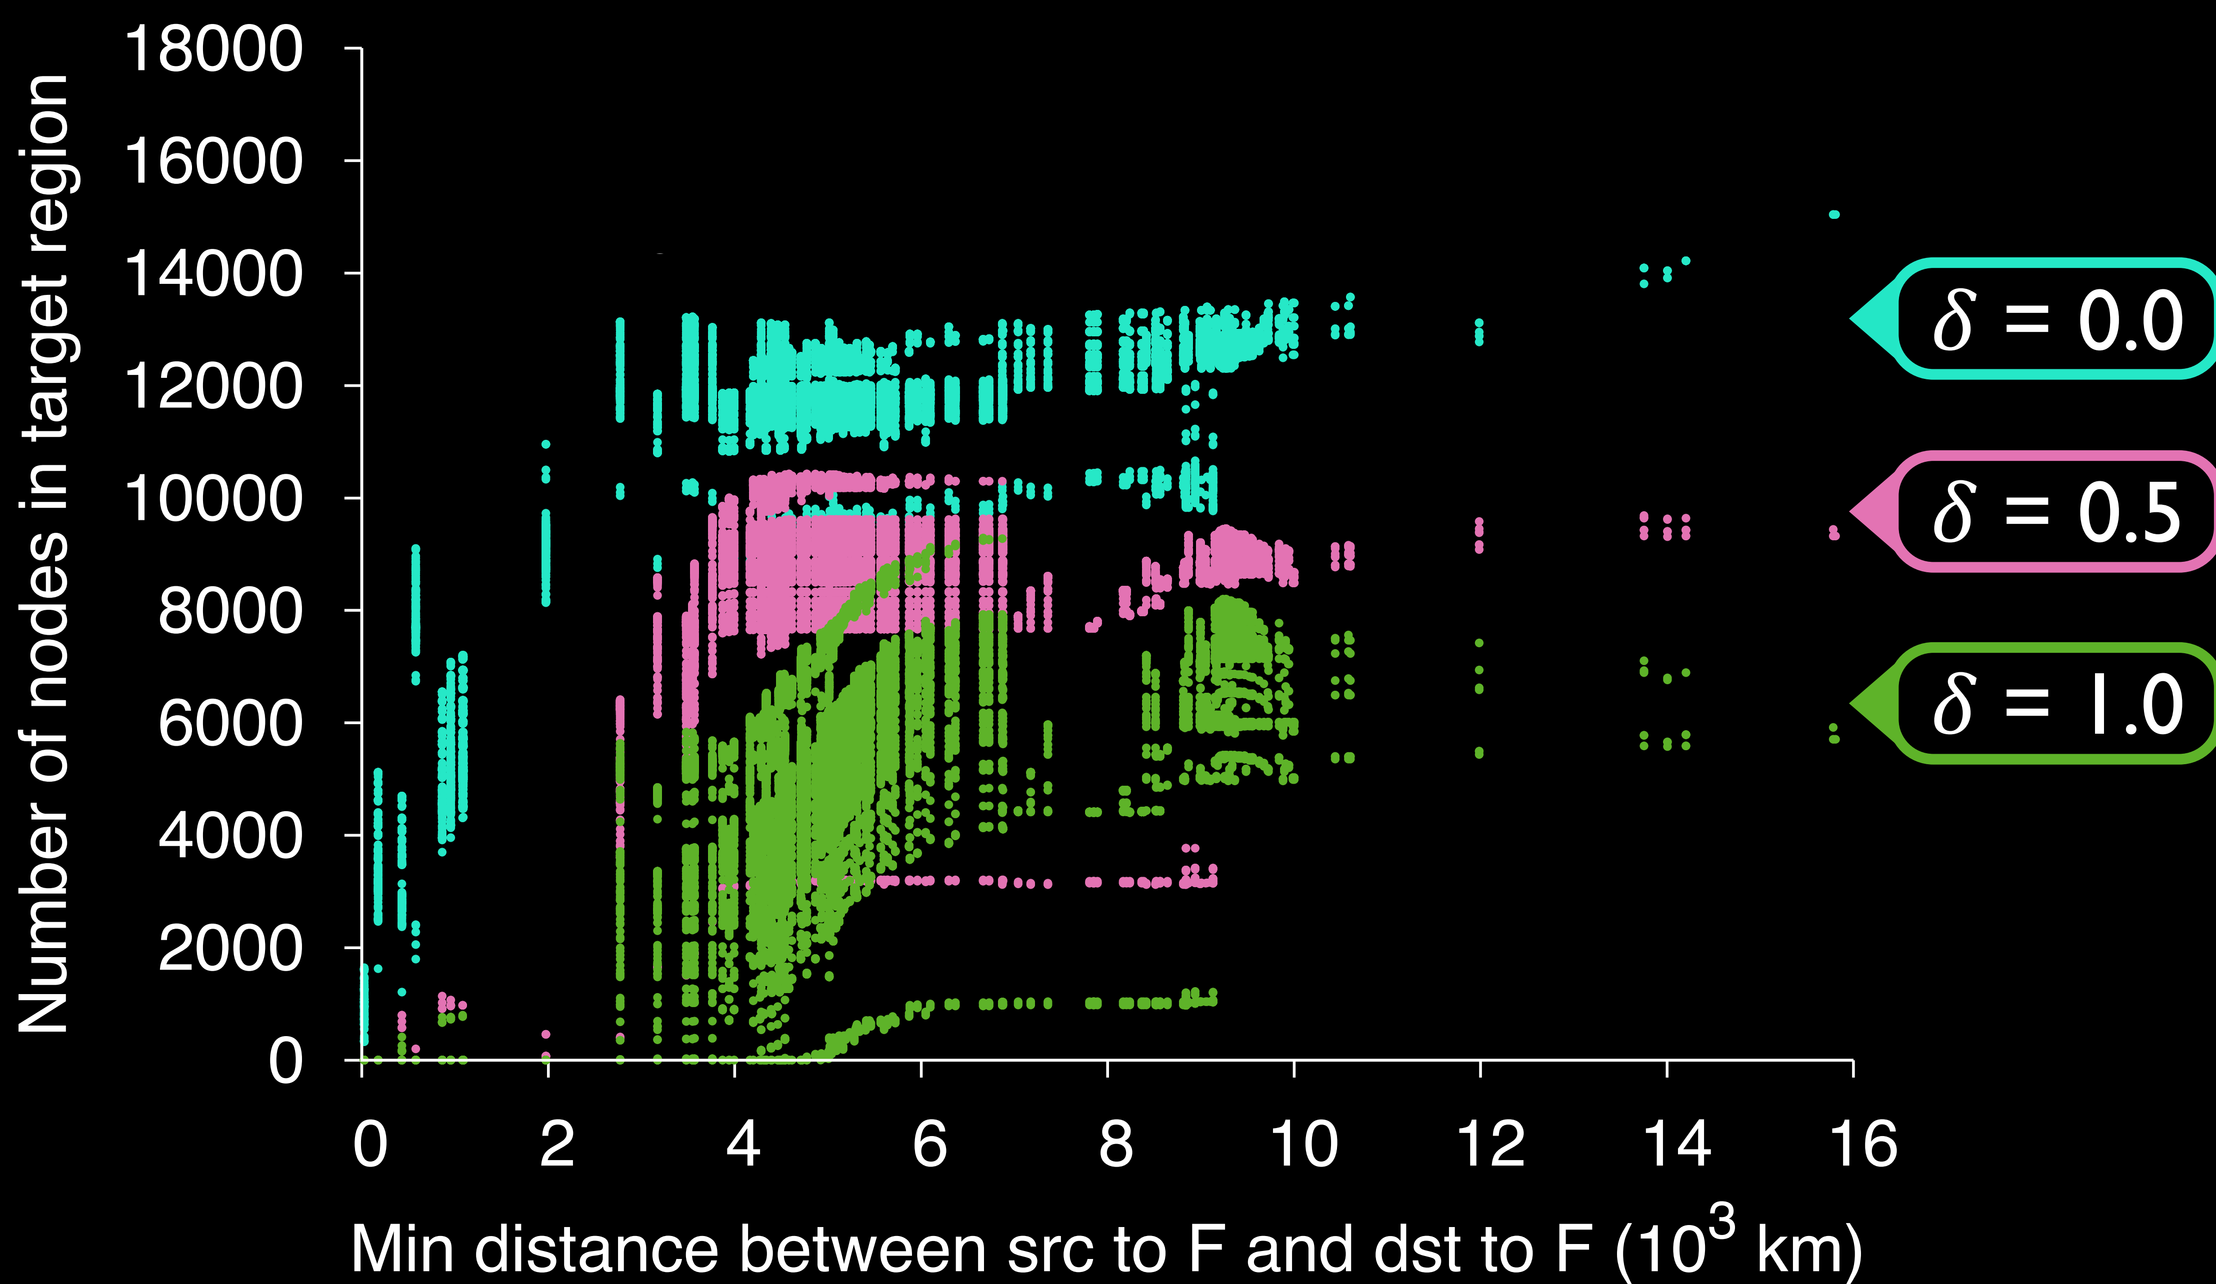

Failure is likely when source or destination are very close to the forbidden region

Other results

- Routes through alibis incur little increase in latency
  - Sometimes even *lower* latencies

- Alibi Routing incurs little communication overhead

- Countries with higher routing centrality are harder, but not impossible, to avoid

Provable avoidance is possible
safely and efficiently

Summary

- Provable avoidance routing
  - Users to specify where they want their packets *not* to go
- "Proof by alibi" makes it possible to provably avoid arbitrary geographic regions without ISP/BGP support
- Alibi Routing finds potential alibis
  - Successfully, so long as src/dst not *too* close
  - At low cost in terms of latency inflation

Code and data available at:

`alibi.cs.umd.edu`

Tor Network

Vulnerable to censorship

[Anon. CCR'12]

Tor Network

Vulnerable to censorship

[Anon. CCR'12]

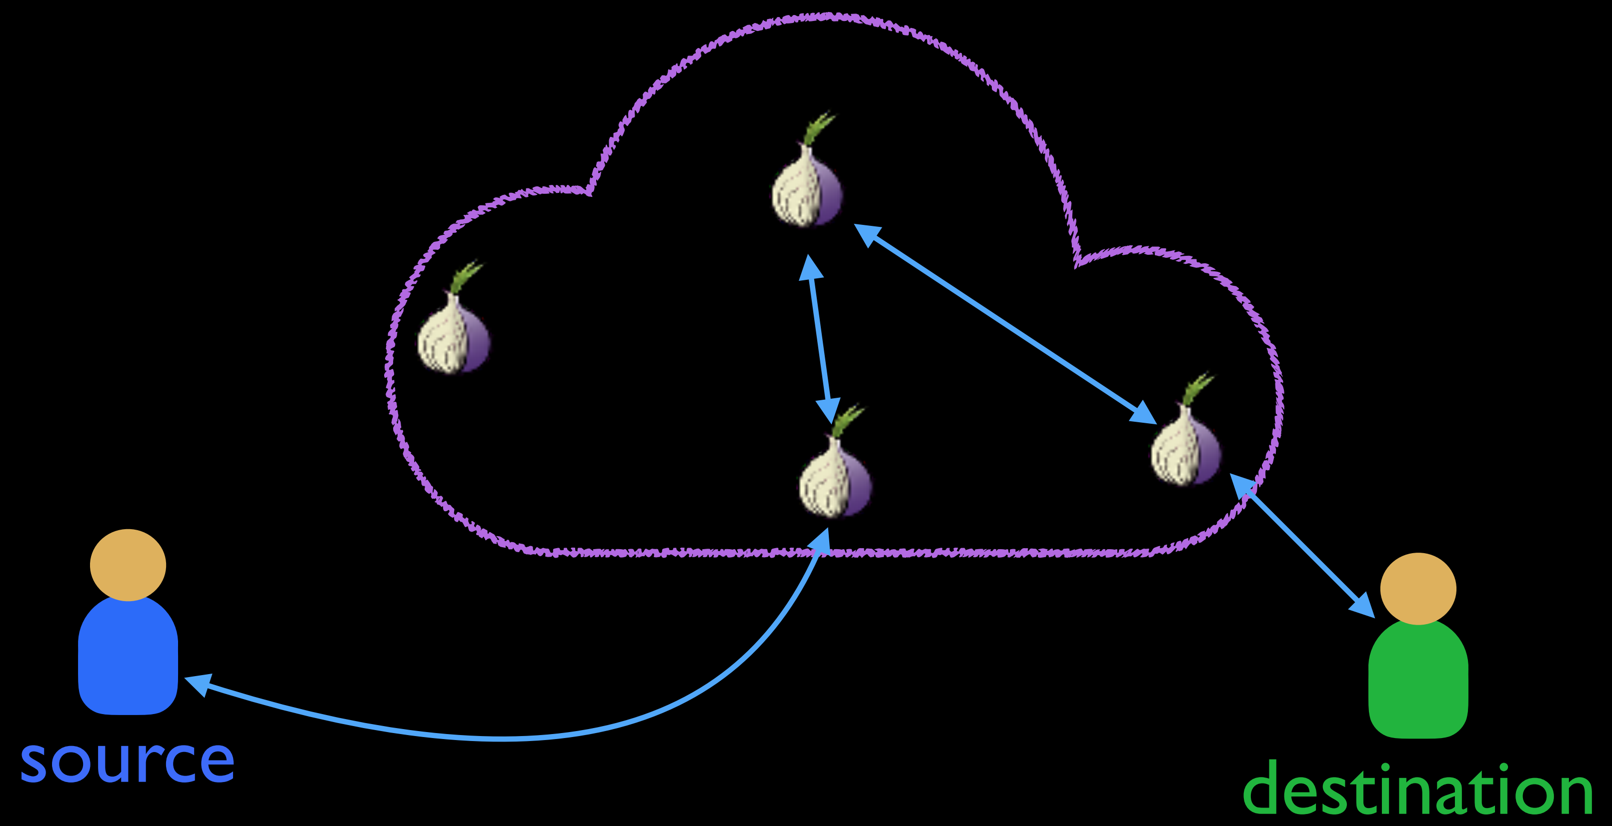

Tor Network

Vulnerable to censorship

[Anon. CCR'12]

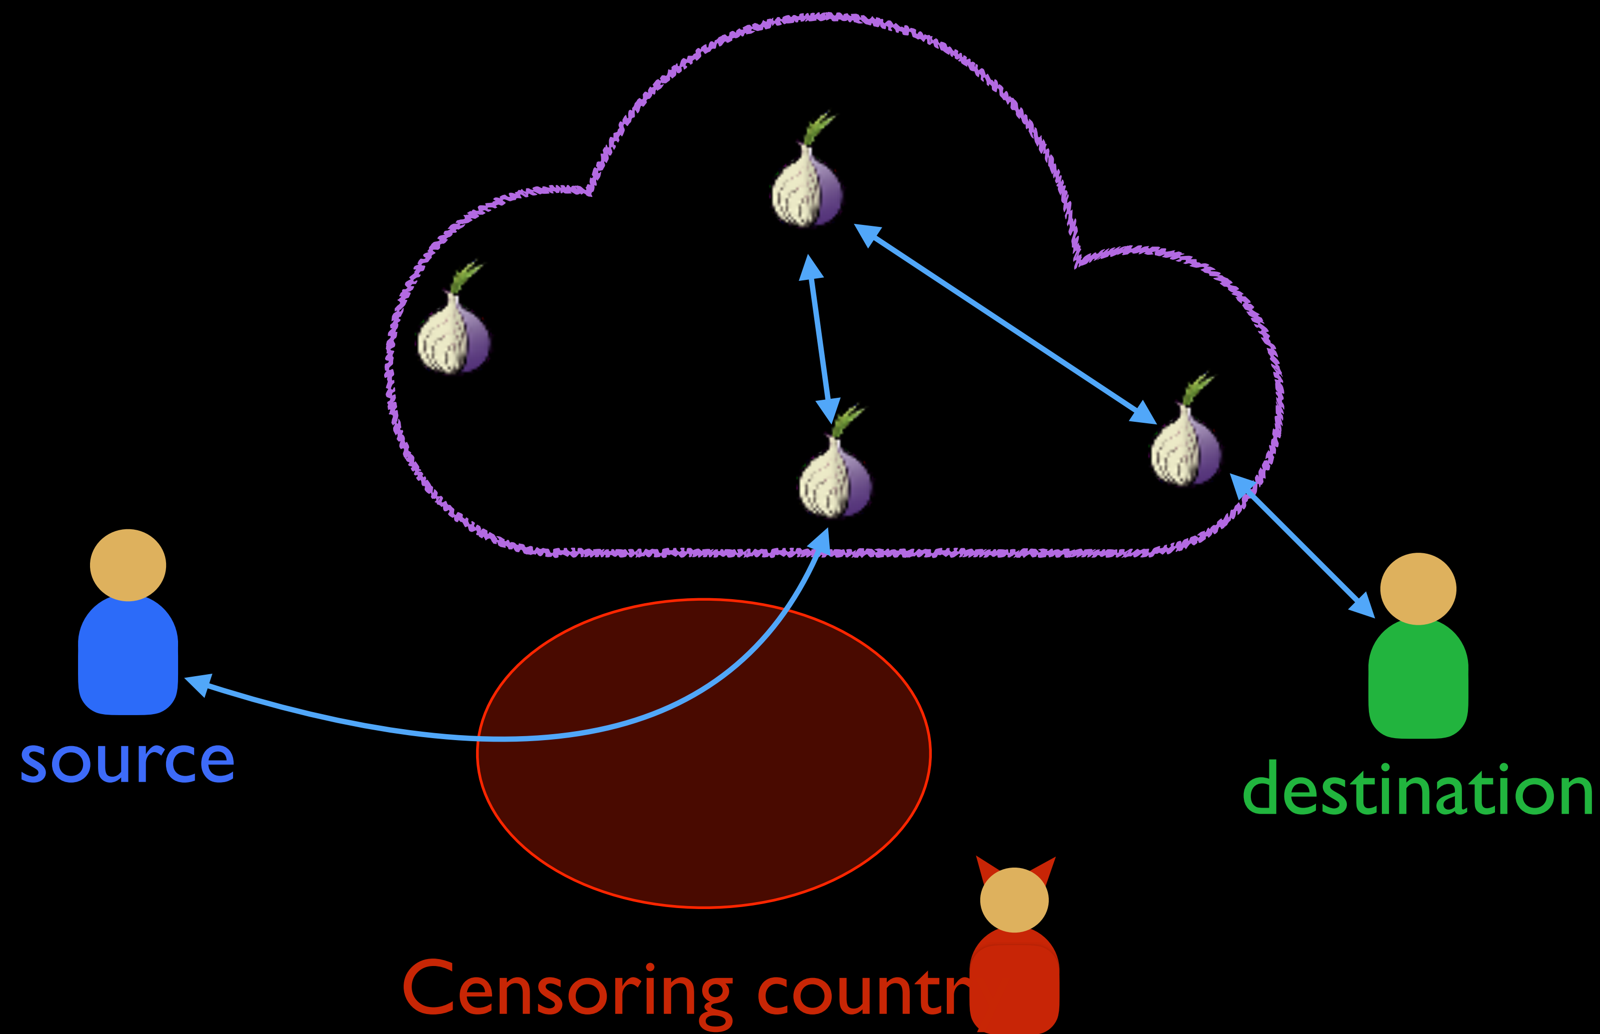

Tor Network

Vulnerable to censorship

[Anon. CCR'12]

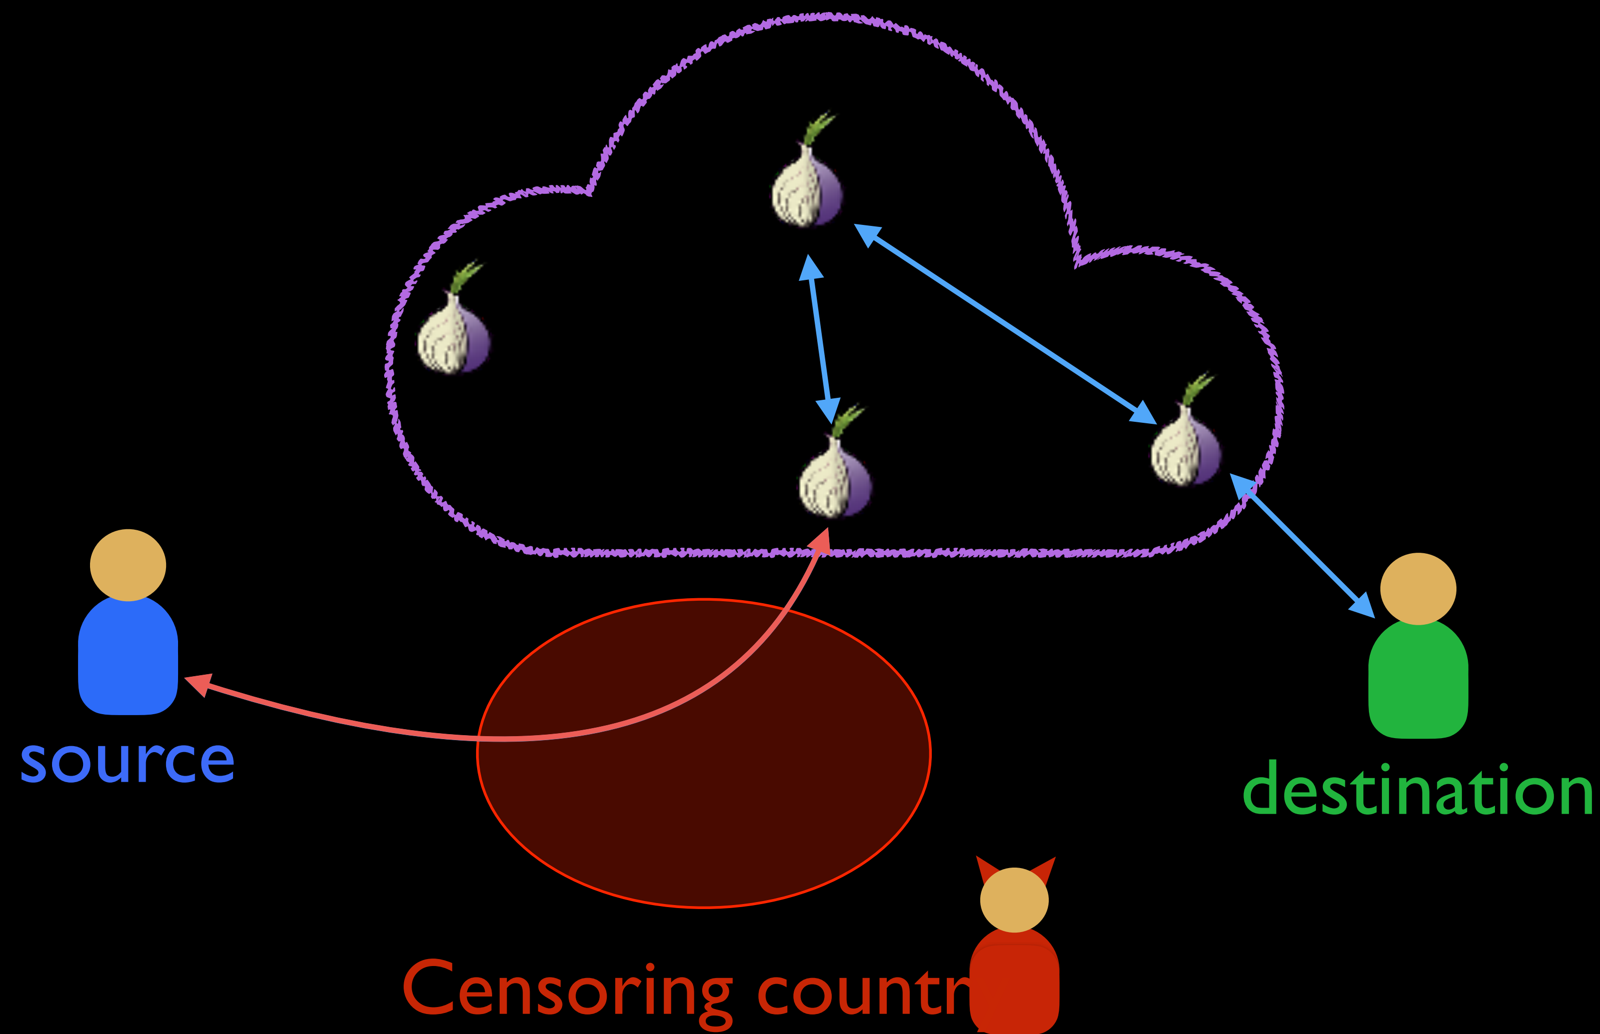

Tor Network

Vulnerable to censorship

[Anon. CCR'12]

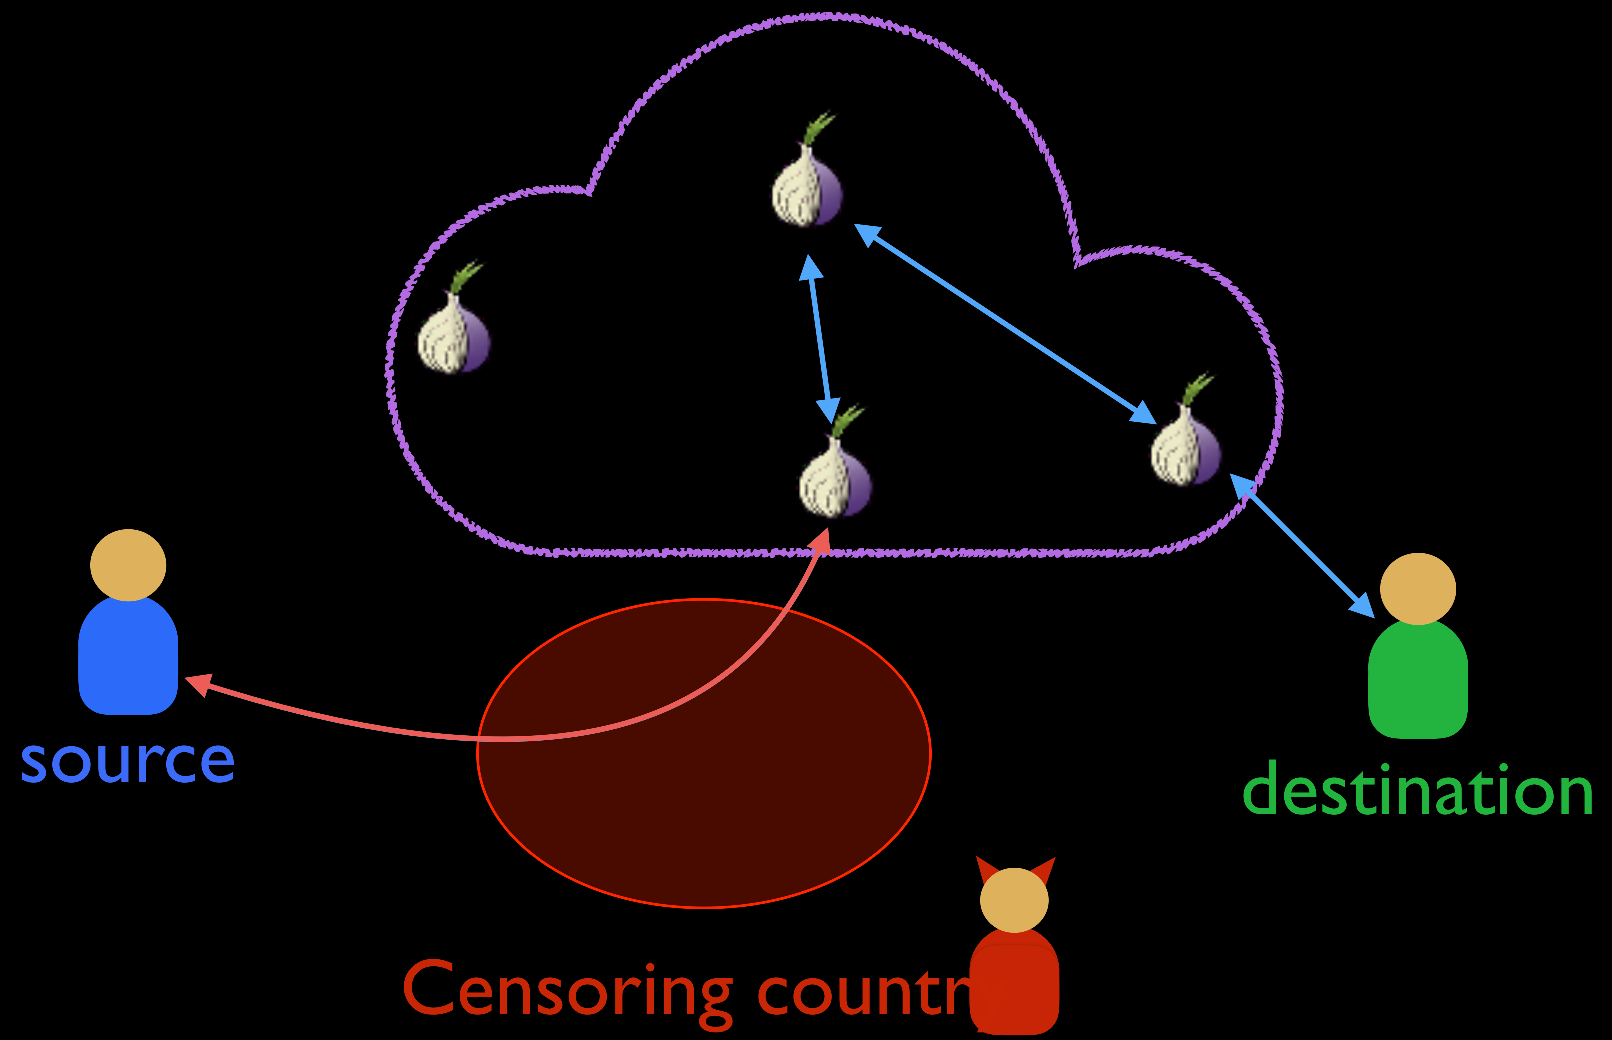

Tor Network

Vulnerable to censorship

[Anon. CCR'12]

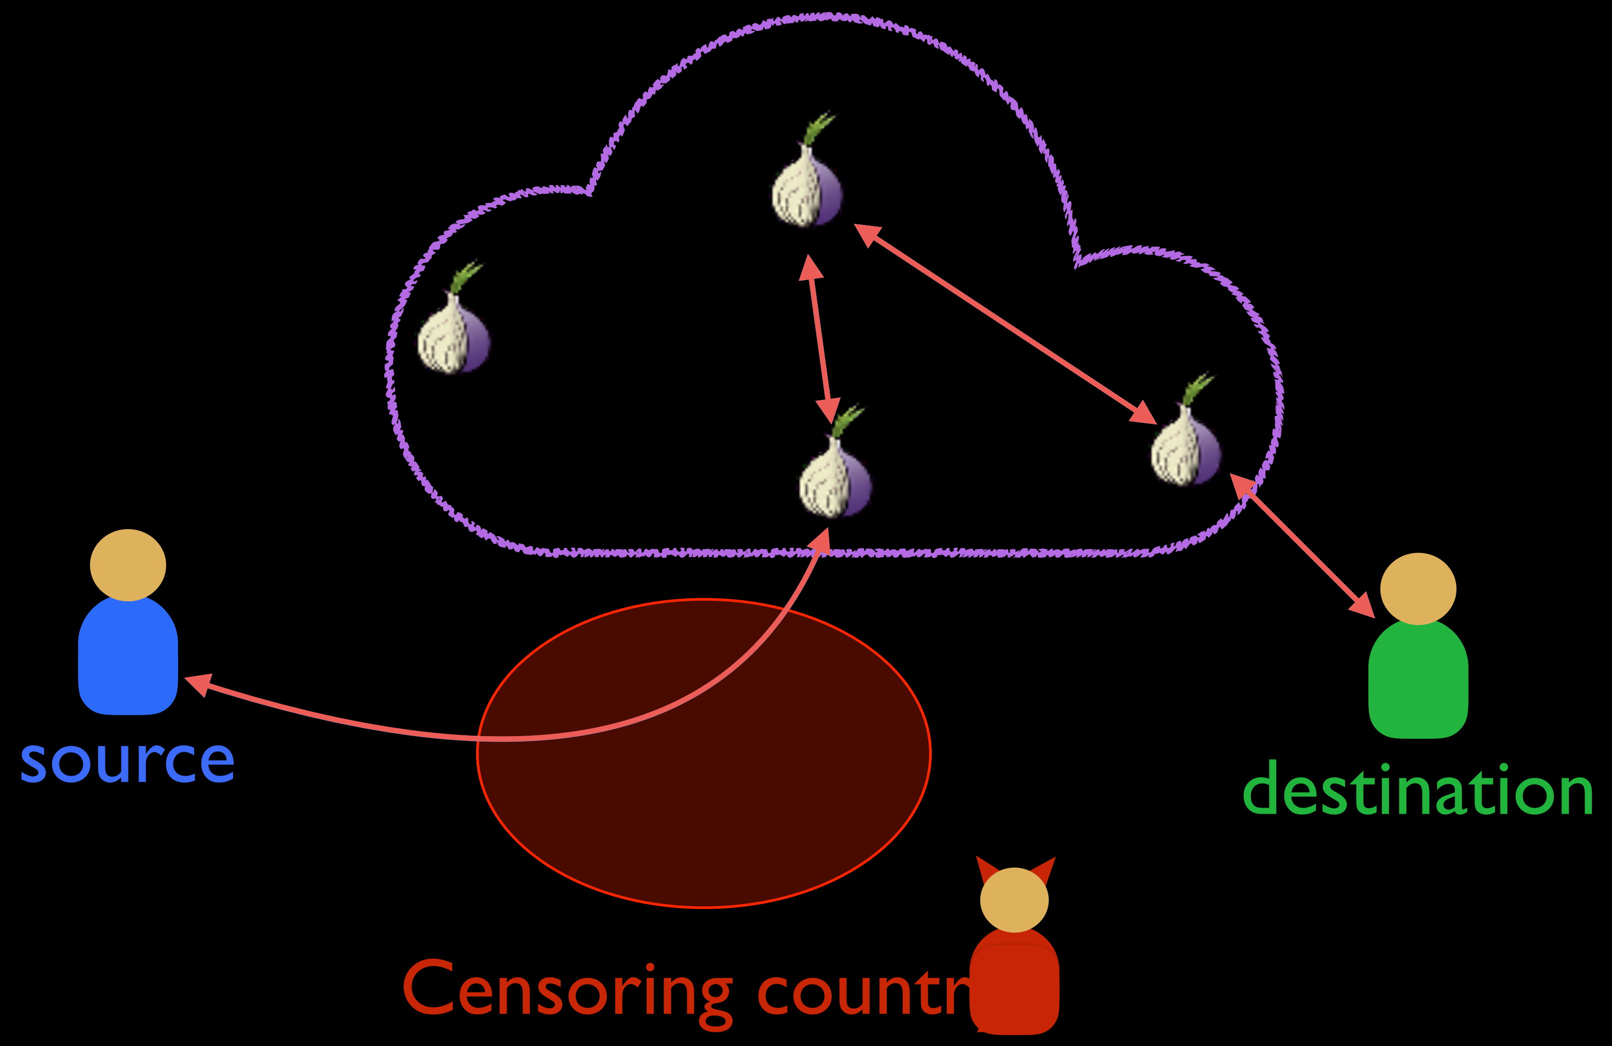

Tor Network

Vulnerable to traffic correlation attacks

[Hopper et al., ACM TISSEC. 2010; Gilad et al., PETS 2012]

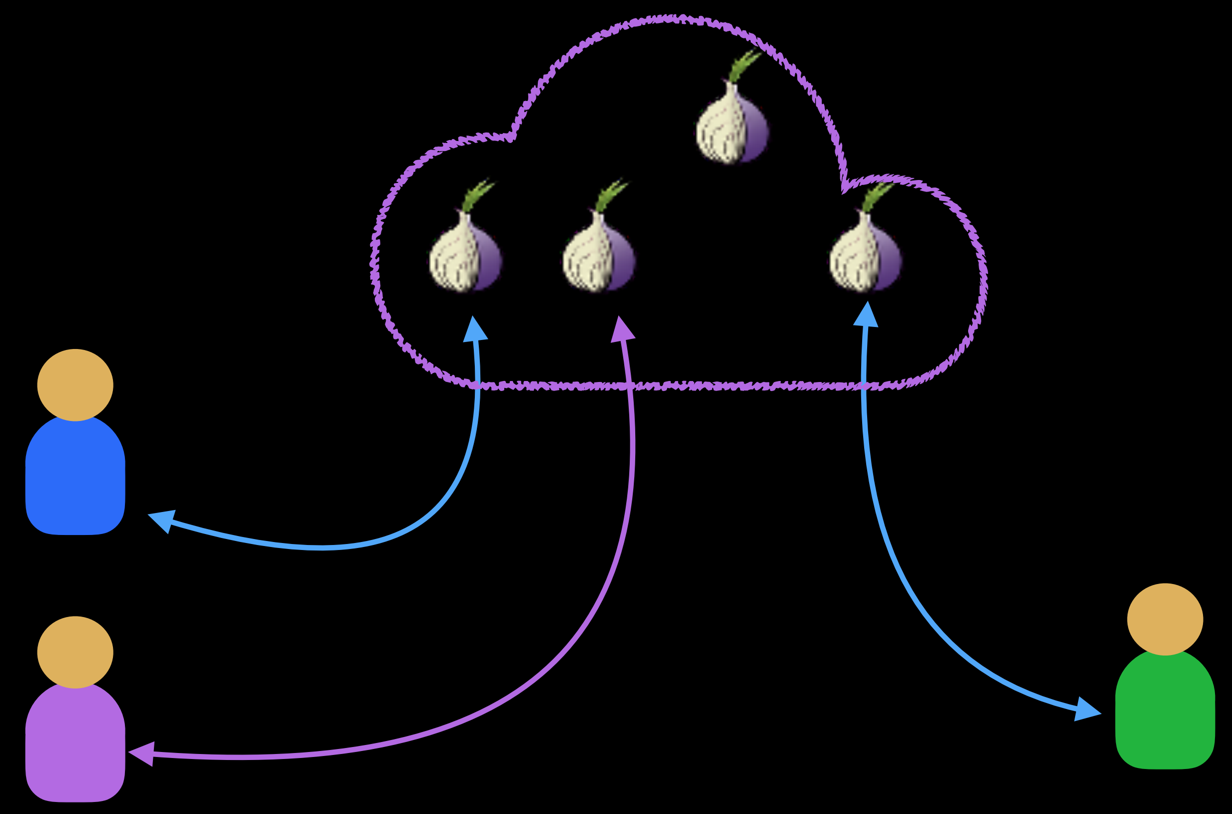

Tor Network

Vulnerable to traffic correlation attacks

[Hopper et al., ACM TISSEC. 2010; Gilad et al., PETS 2012]

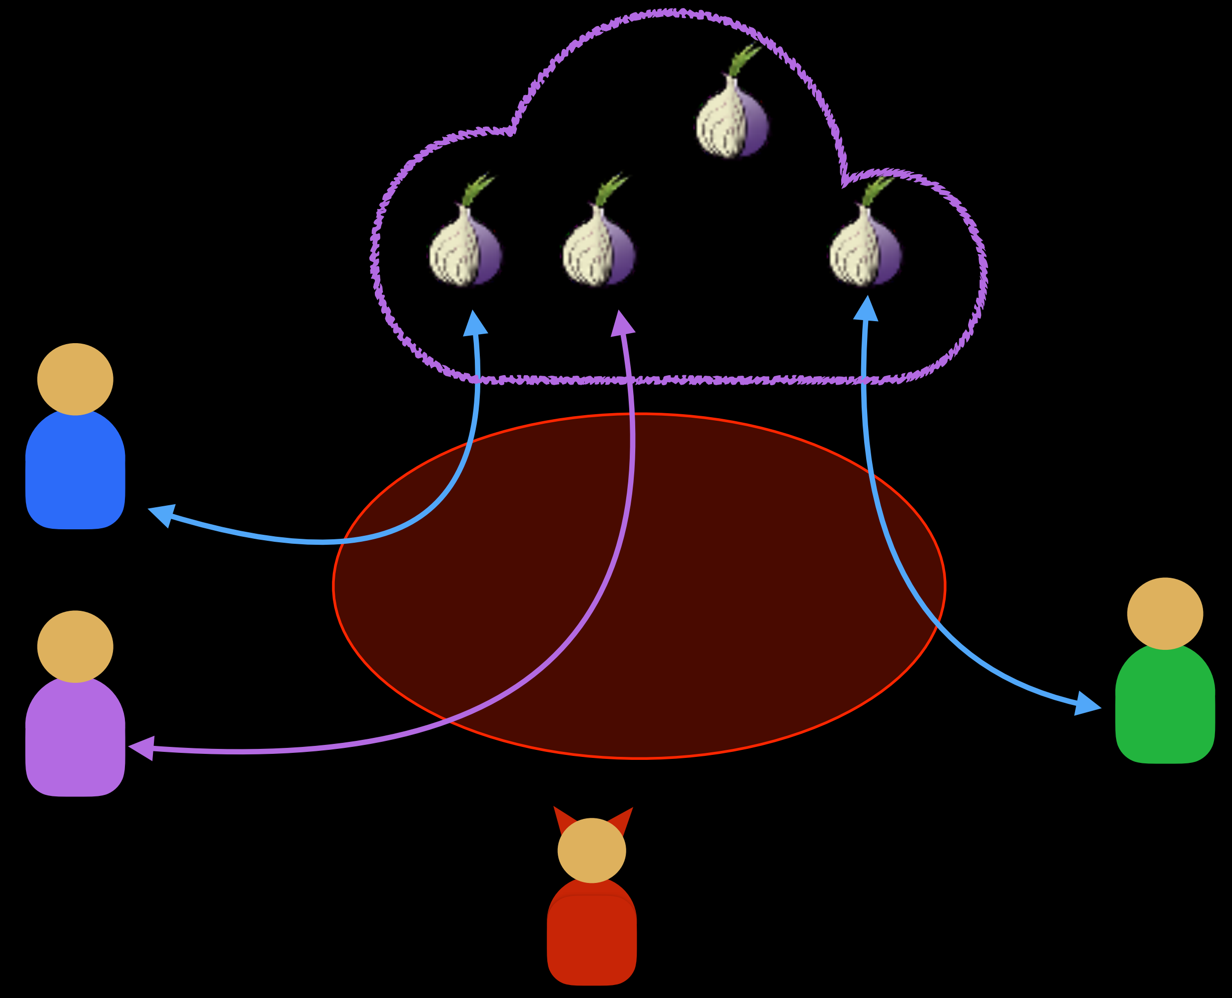

Tor Network

Vulnerable to traffic correlation attacks

[Hopper et al., ACM TISSEC. 2010; Gilad et al., PETS 2012]

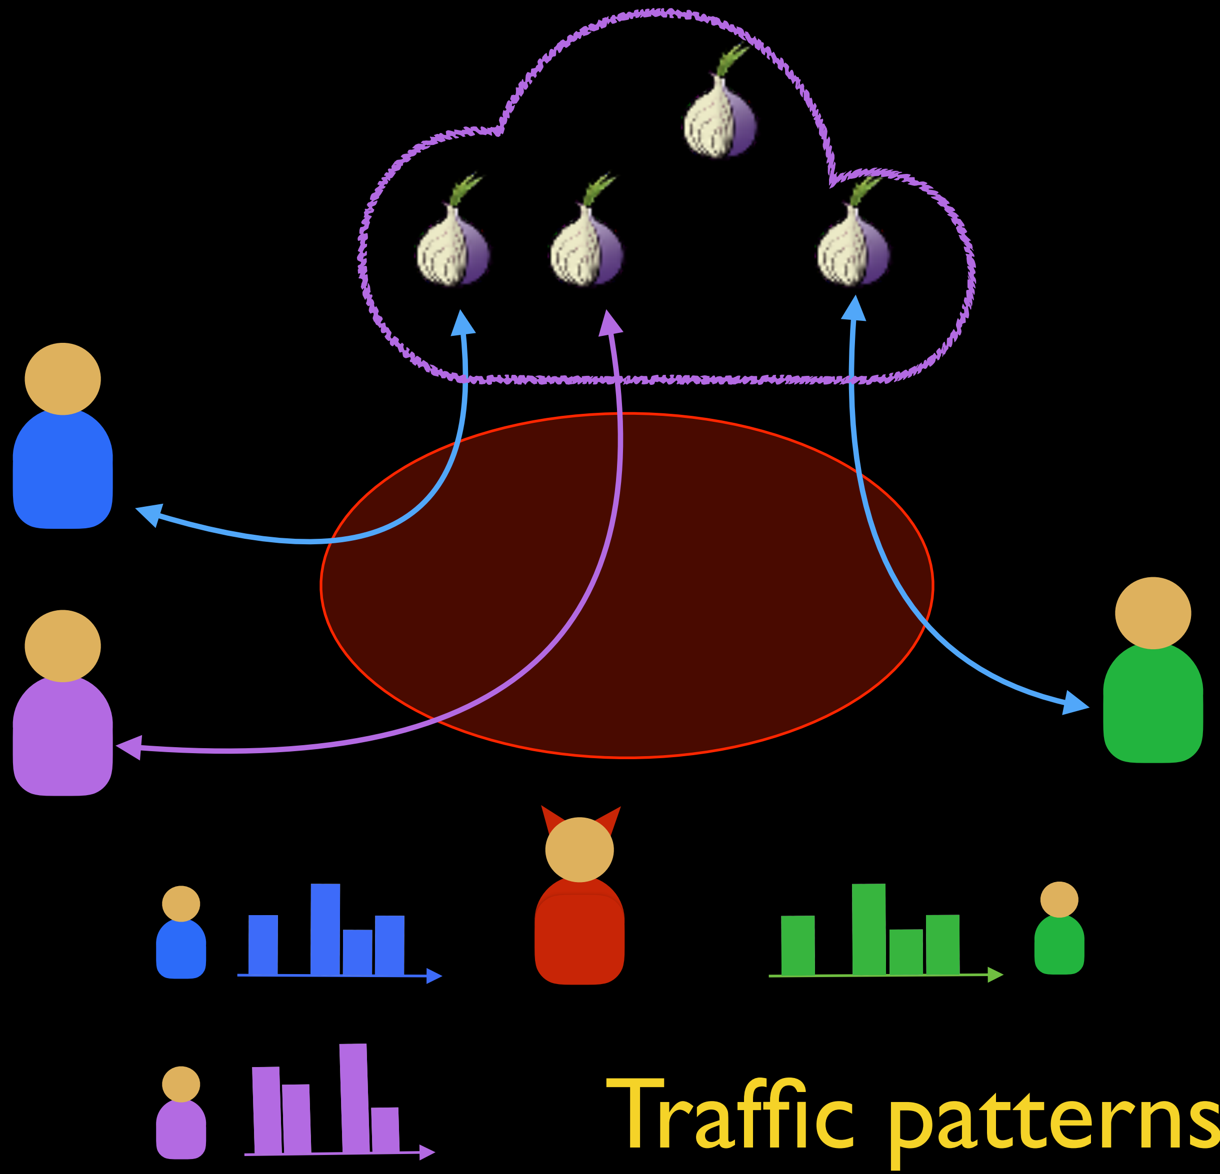

Traffic patterns

Tor Network

Vulnerable to traffic correlation attacks

[Hopper et al., ACM TISSEC. 2010; Gilad et al., PETS 2012]

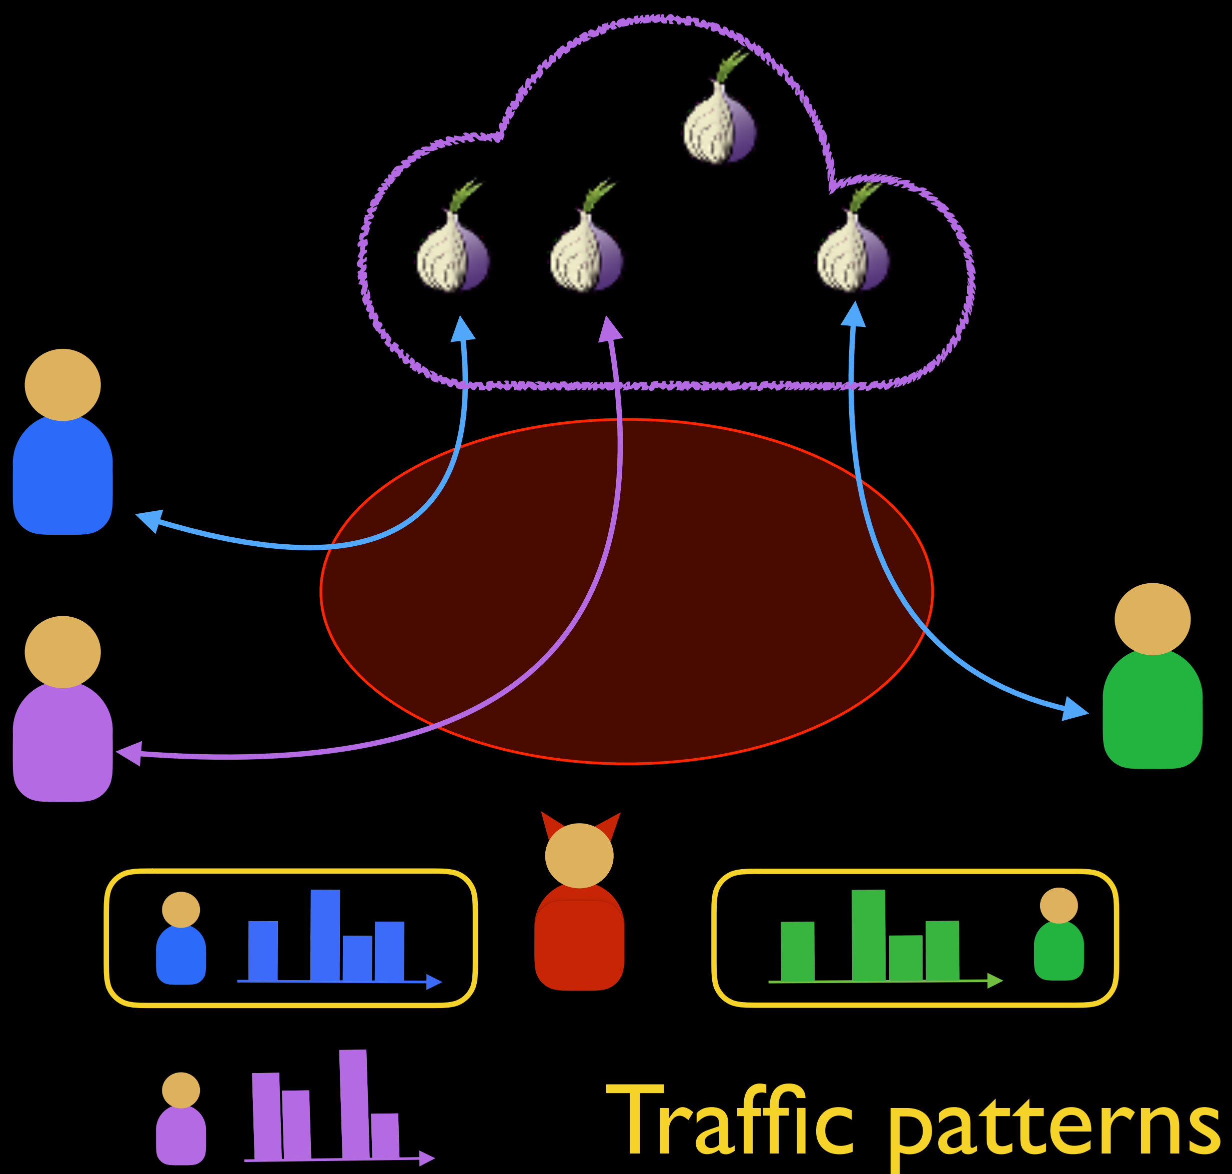

Tor Network

Vulnerable to traffic correlation attacks

[Hopper et al., ACM TISSEC. 2010; Gilad et al., PETS 2012]

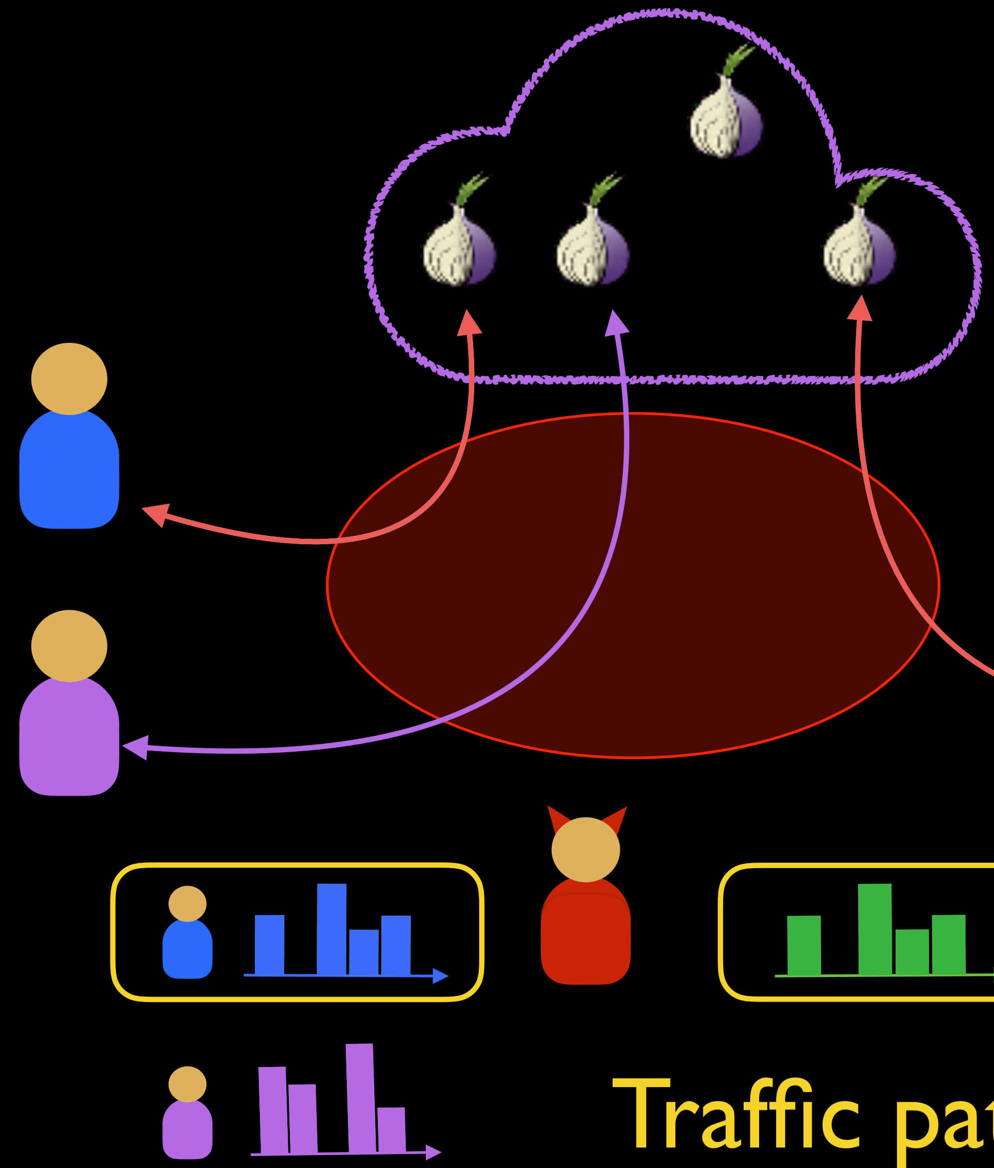

Traffic patterns

Threat model

Nation-state adversary

Threat model
Nation-state adversary

- Adversaries can:
  - launch various attacks when on the path
  - hide from network topology measurement (e.g. *traceroute*)
  - attract routes to their administrative domains

Threat model

Nation-state adversary

- Adversaries can:
  - launch various attacks when on the path
  - hide from network topology measurement (e.g. *traceroute*)
  - attract routes to their administrative domains

Threat model
Nation-state adversary

- Adversaries can:
  - launch various attacks when on the path
  - hide from network topology measurement (e.g. *traceroute*)
  - attract routes to their administrative domains

- Adversaries cannot:

Threat model

Nation-state adversary

- Adversaries can:
  - launch various attacks when on the path
  - hide from network topology measurement (e.g. *traceroute*)
  - attract routes to their administrative domains
- Adversaries cannot:

Fundamental assumption:

We know the geographic boundaries wherein the attackers reside

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once

Never-twice

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once

Never-twice

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once

Never-twice

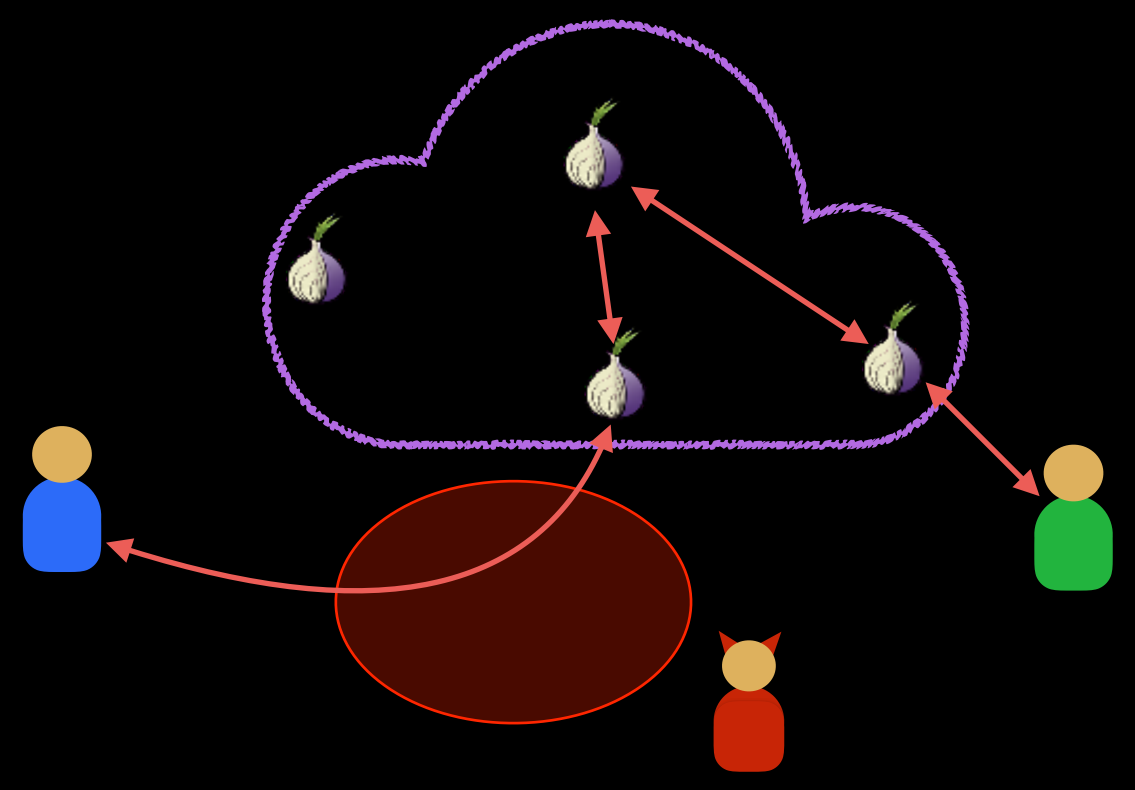

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once

Never-twice

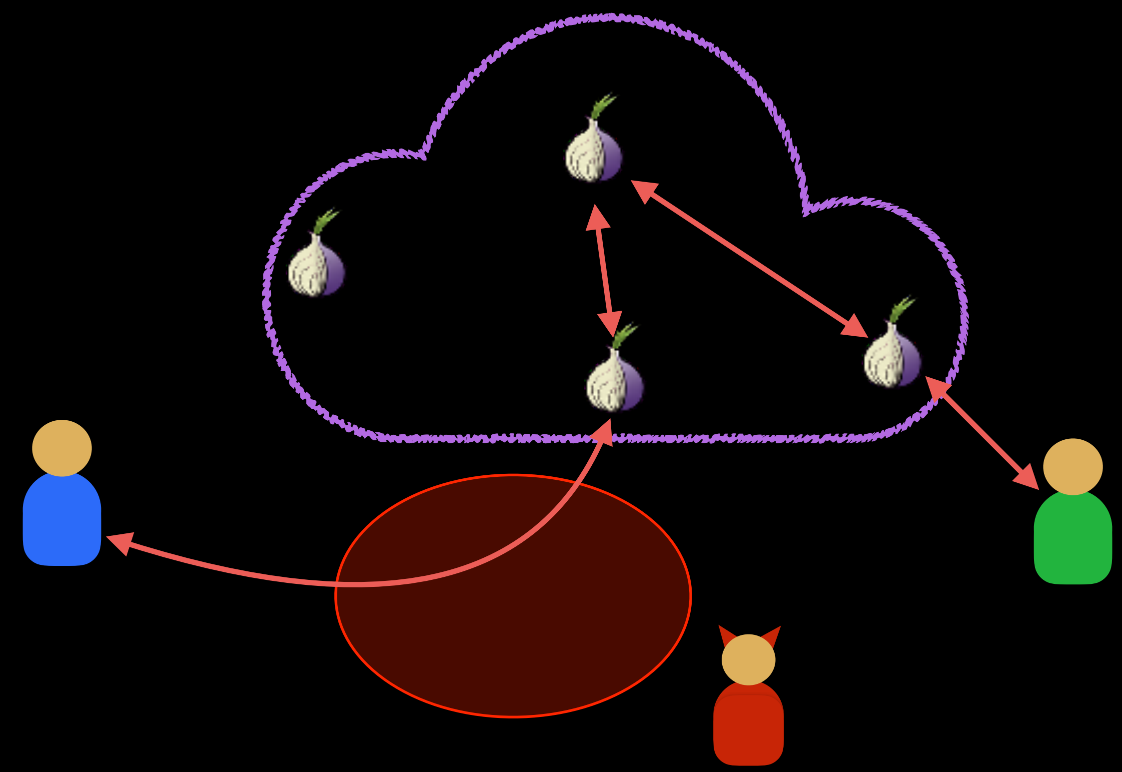

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once

Never-twice

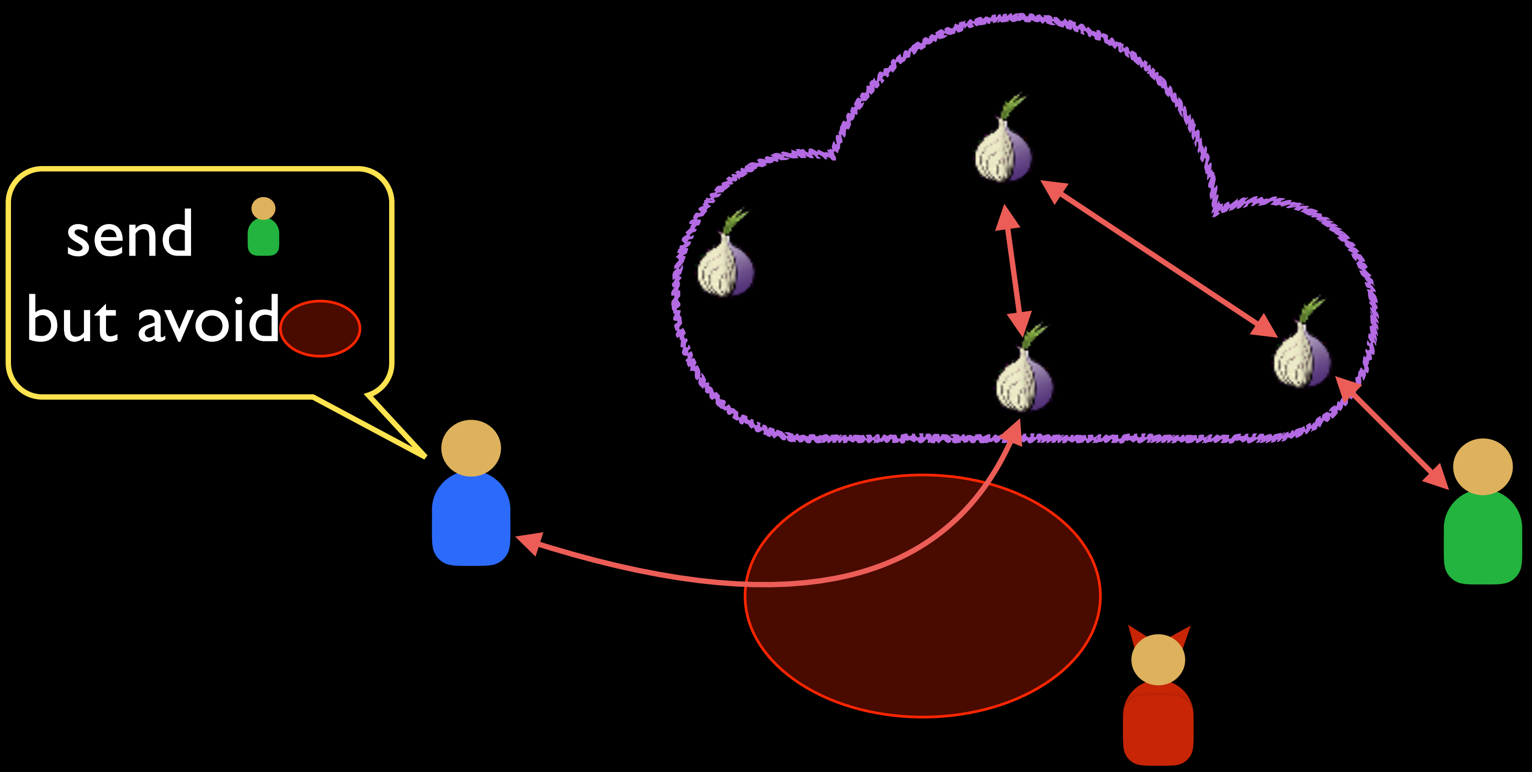

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once

Never-twice

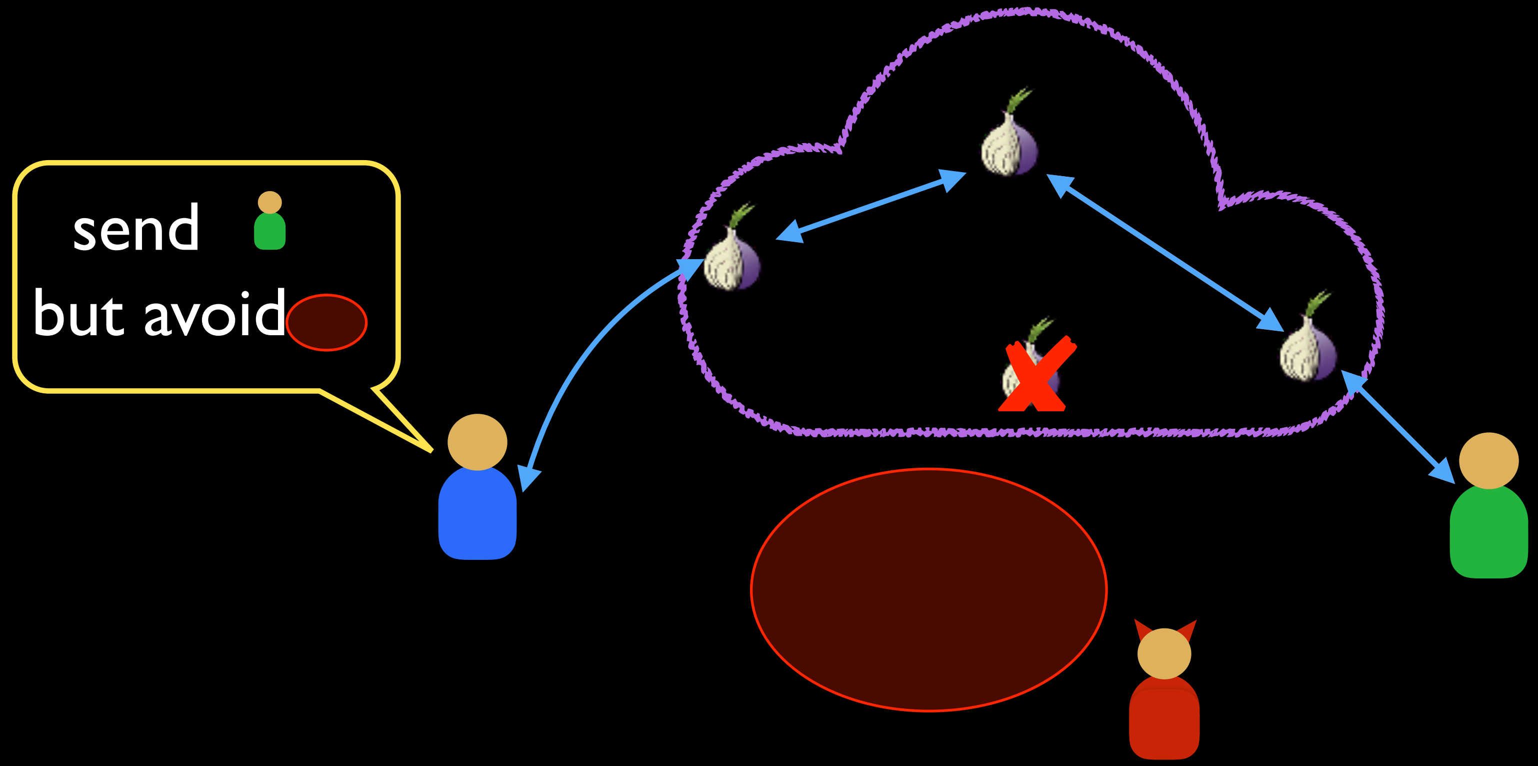

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once

Never-twice

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once

Never-twice

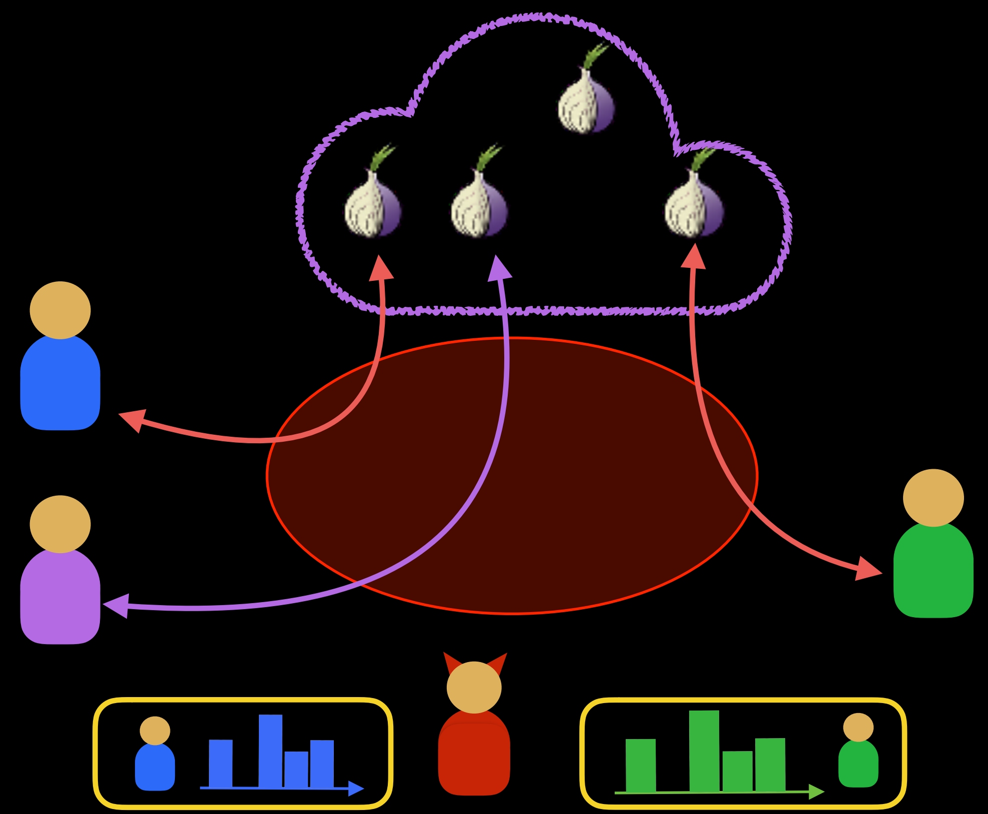

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once

Never-twice

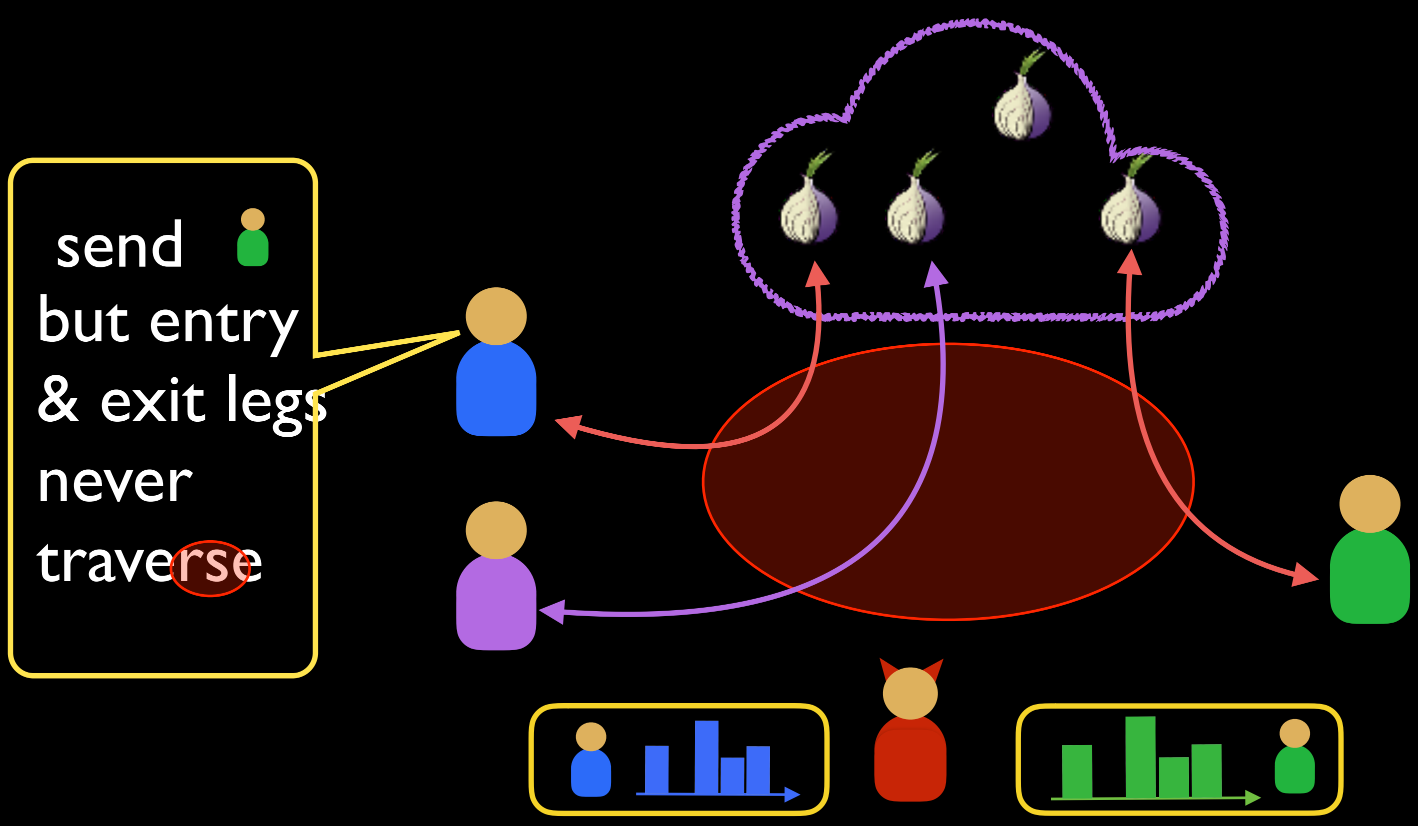

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once

Never-twice

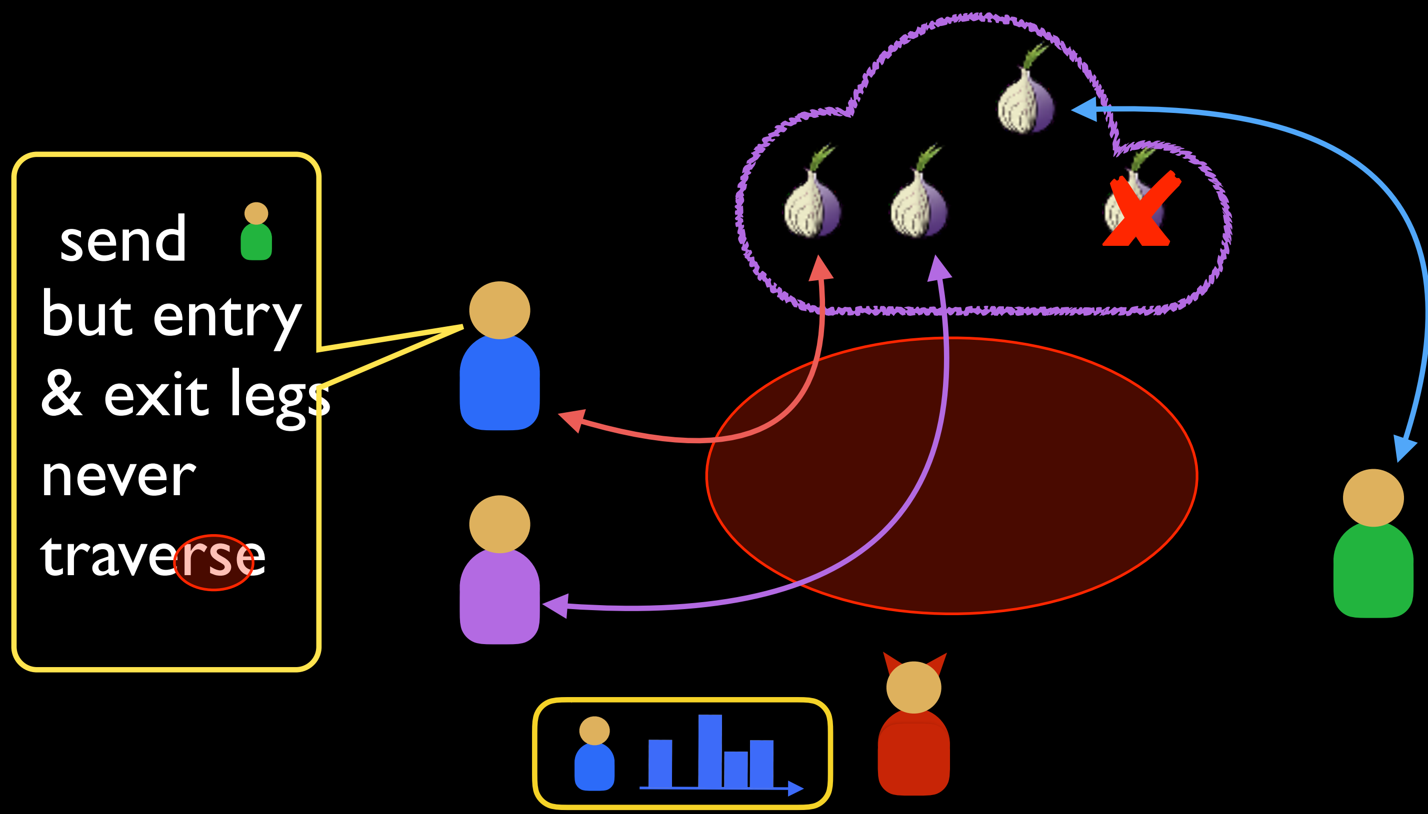

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once

never traverse specified regions

Never-twice

entry & exit legs never traverse

Provide per-packet proof of avoidance

DeTor goals

Deployable

Allow users to avoid adversaries with smart circuits selection

Proof

Provide proofs of avoidance

DeTor goals

Deployable

Allow users to avoid adversaries with smart circuits selection

DeTor goals

Deployable

Allow users to avoid adversaries with smart circuits selection

Without having to know underlying routes

DeTor goals

Deployable

Allow users to avoid adversaries with smart circuits selection

Without having to know underlying routes

Without modifications to Internet routers

DeTor goals

Deployable

Allow users to avoid adversaries with smart circuits selection

- Without having to know underlying routes
- Without modifications to Internet routers
- Without changes to Tor's protocol

DeTor goals

Proof

Provide proofs of avoidance

DeTor goals

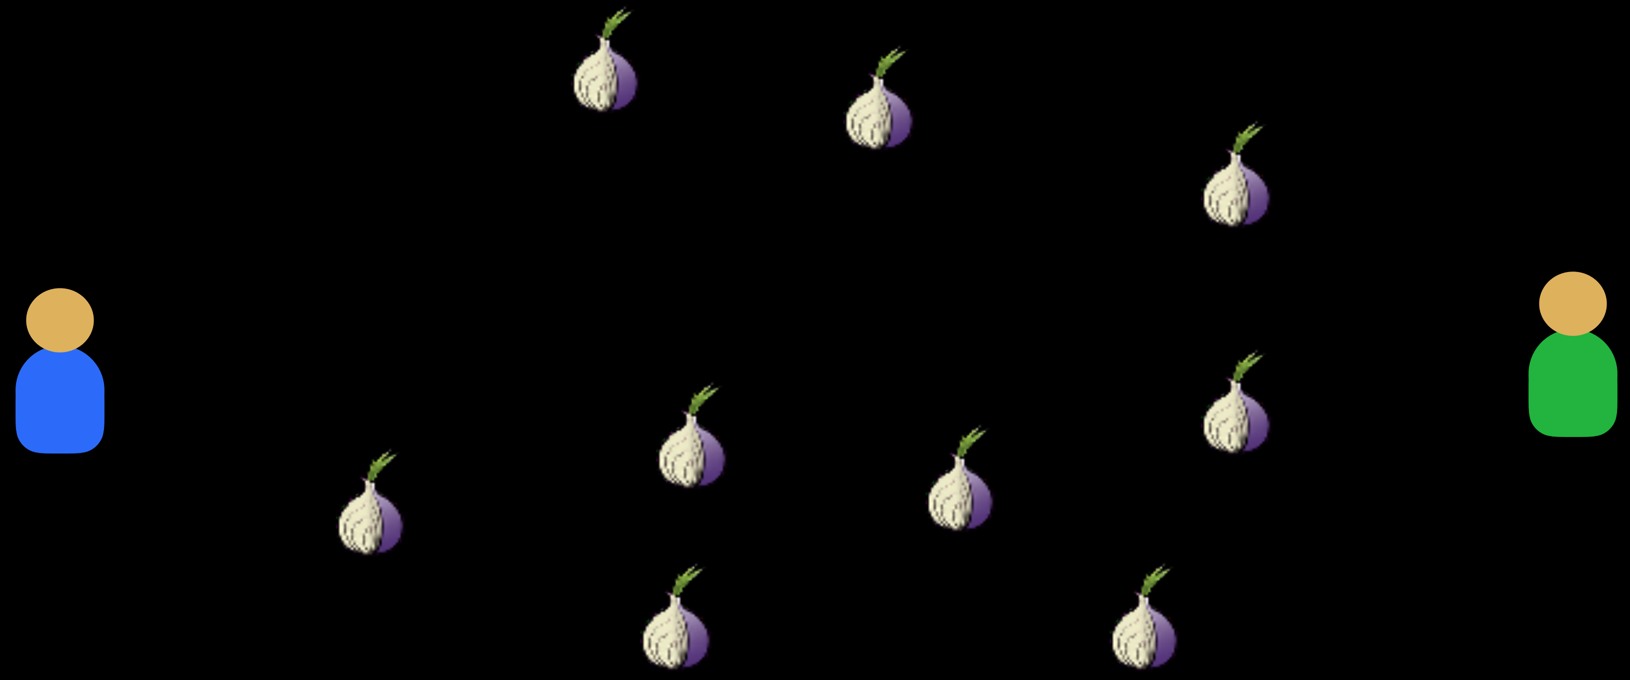

Proof

Provide proofs of avoidance

DeTor goals

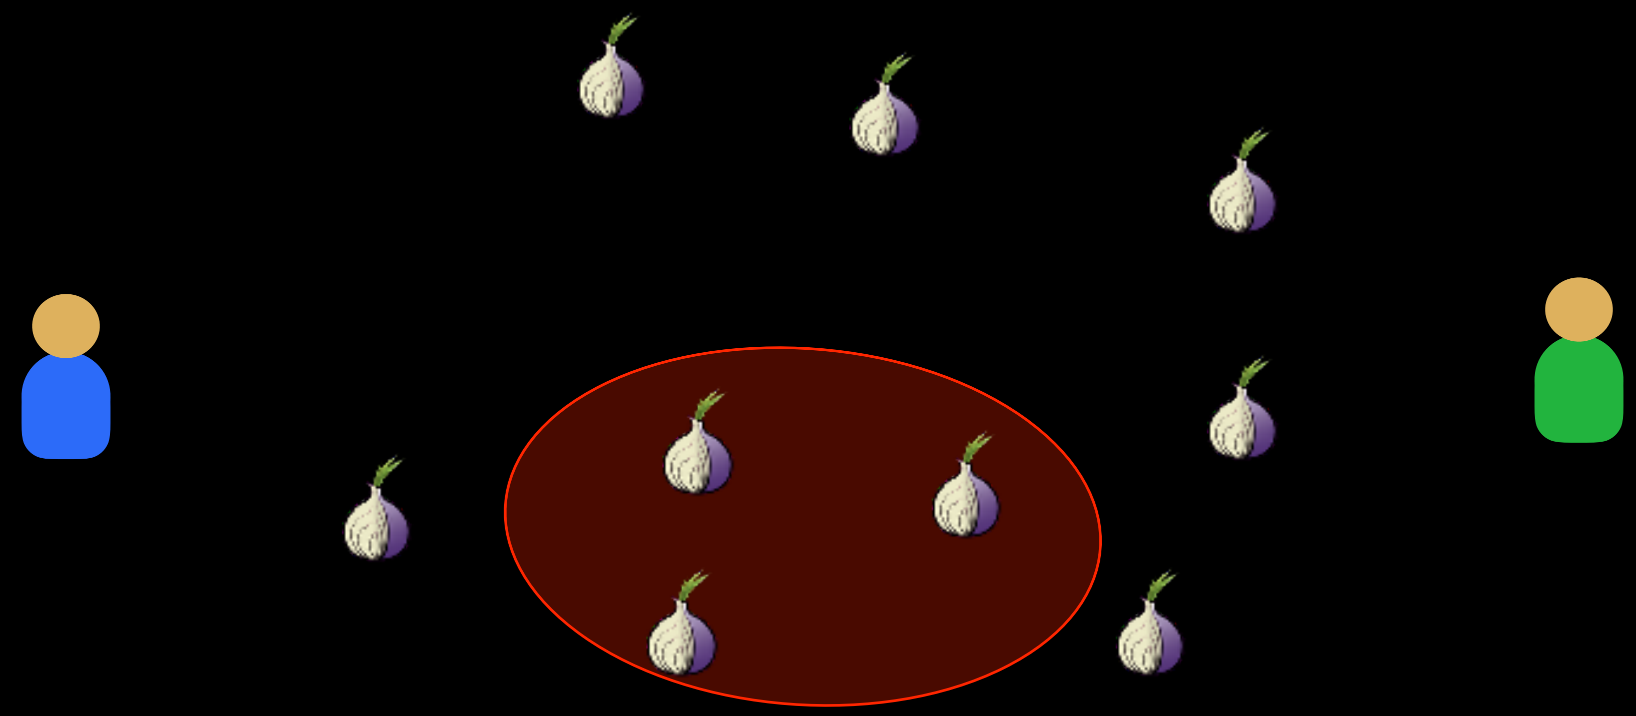

Proof

Provide proofs of avoidance

DeTor goals

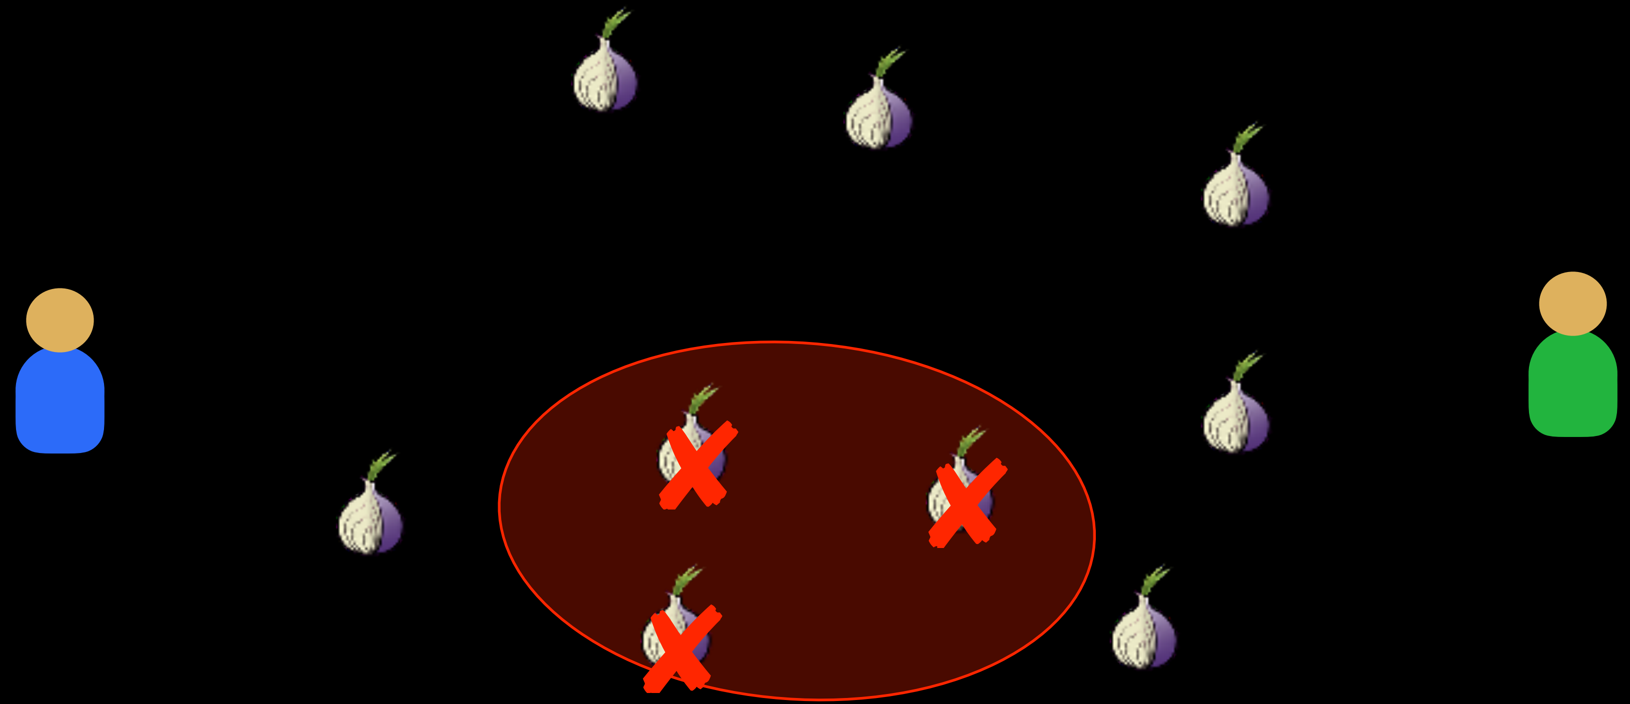

Proof

Provide proofs of avoidance

DeTor goals

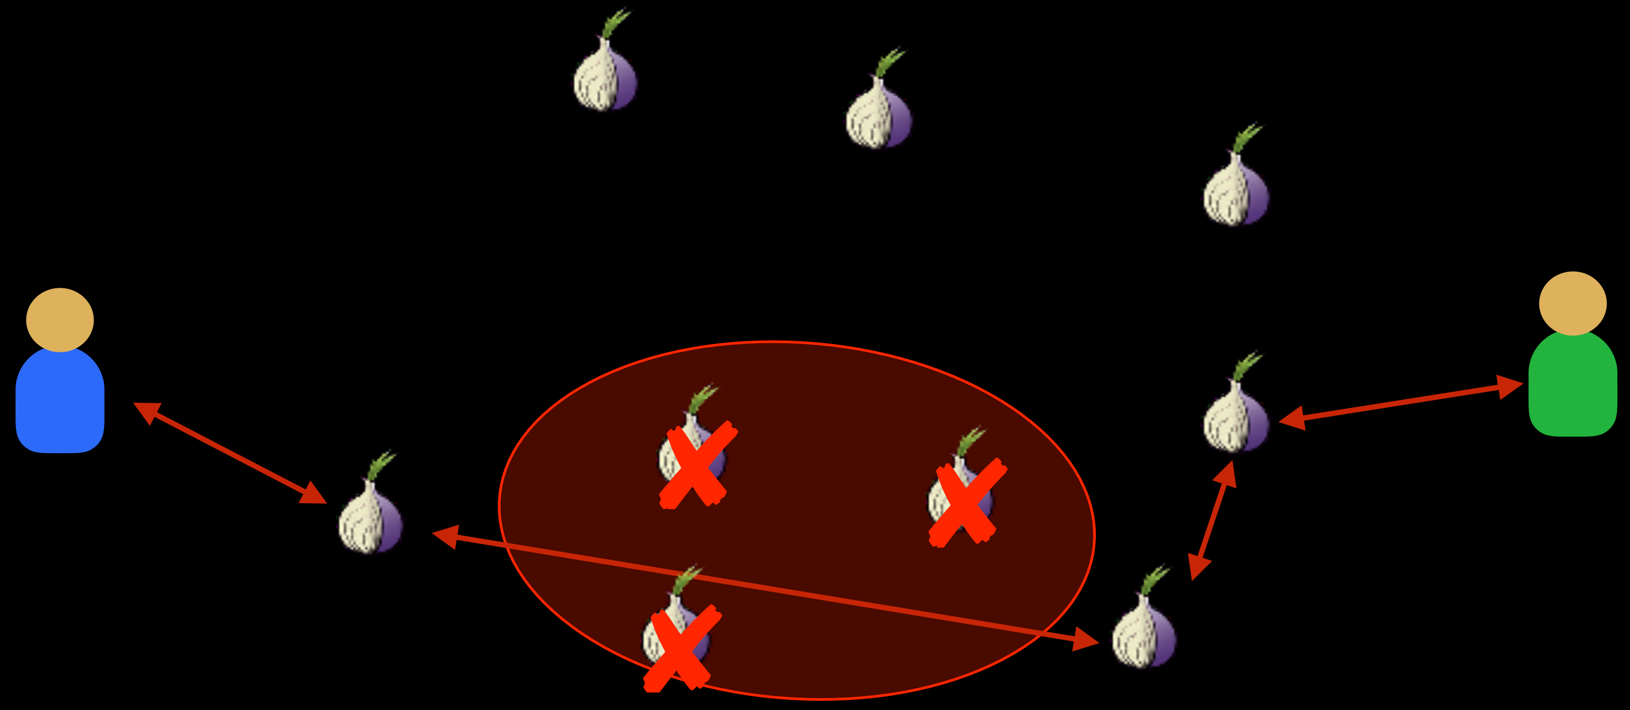

Proof

Provide proofs of avoidance

DeTor goals

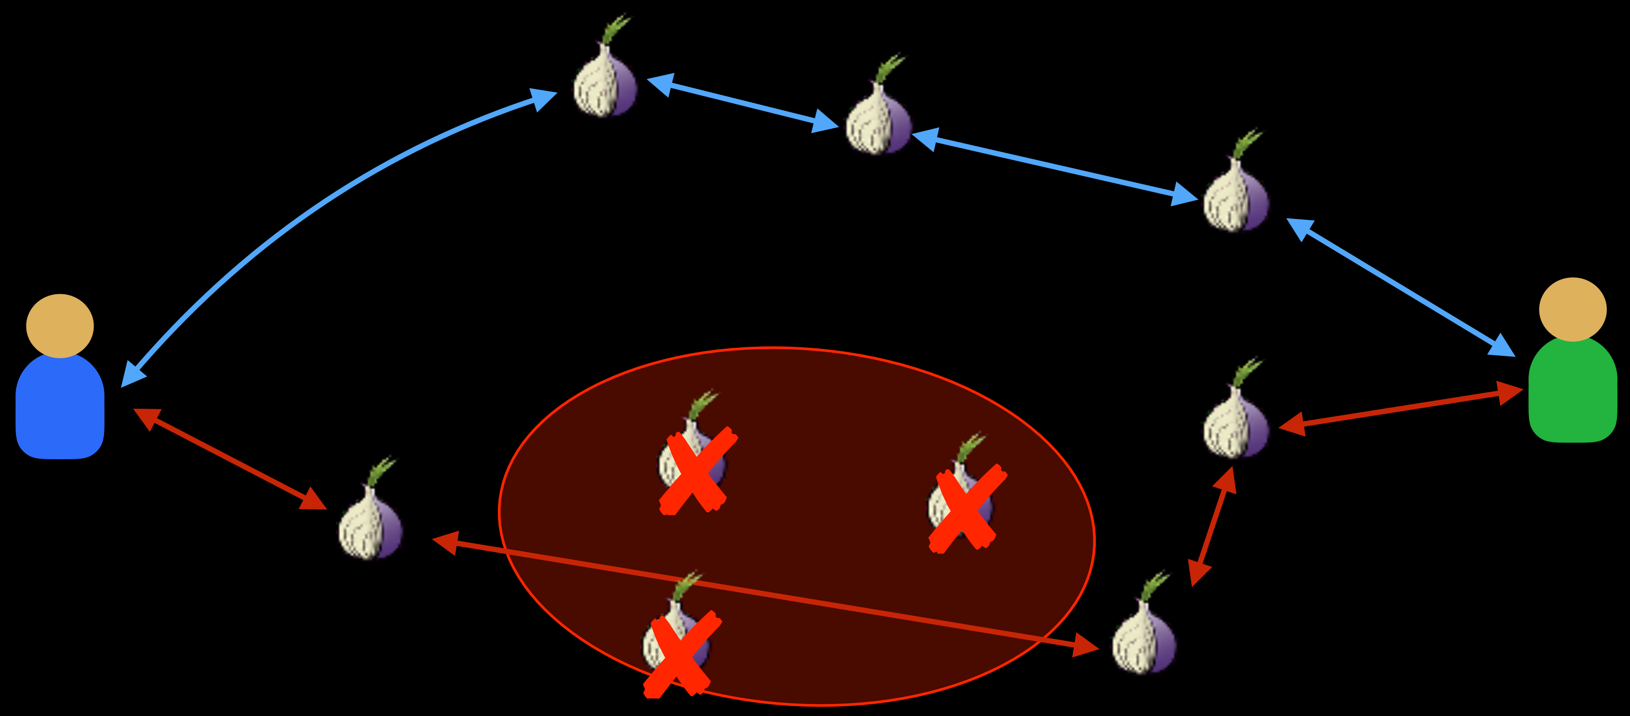

Proof

Provide proofs of avoidance

DeTor goals

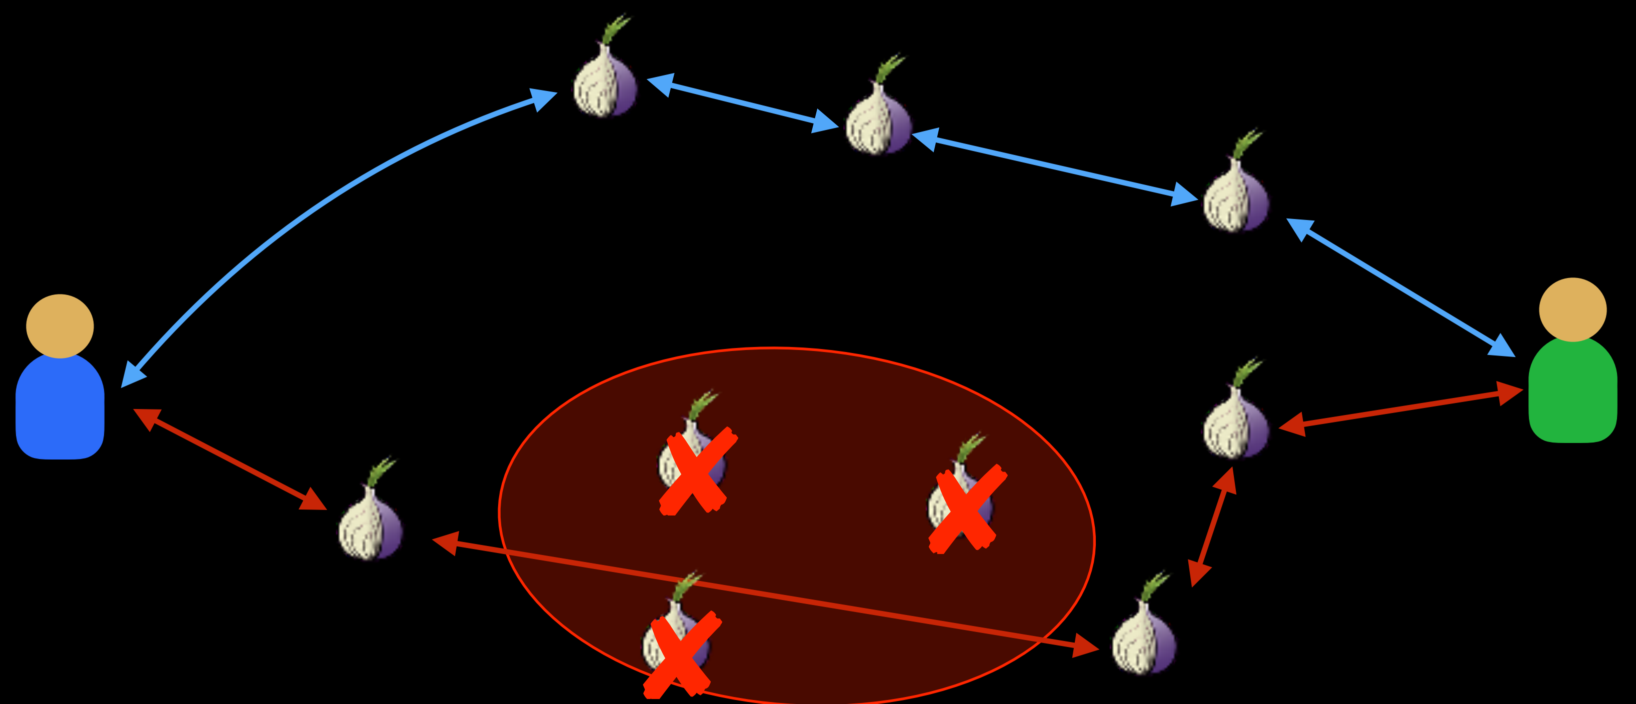

Proof

Provide proofs of avoidance

Measurement of roundtrip time

Provable avoidance

Alibi Routing by Levin et al. in SIGCOMM 2015

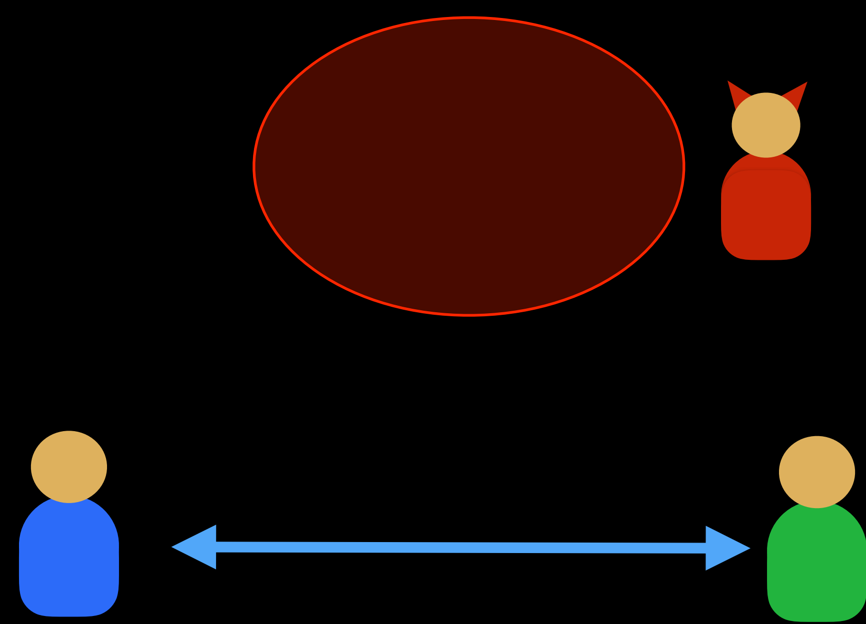

Provable avoidance

Alibi Routing by Levin et al. in SIGCOMM 2015

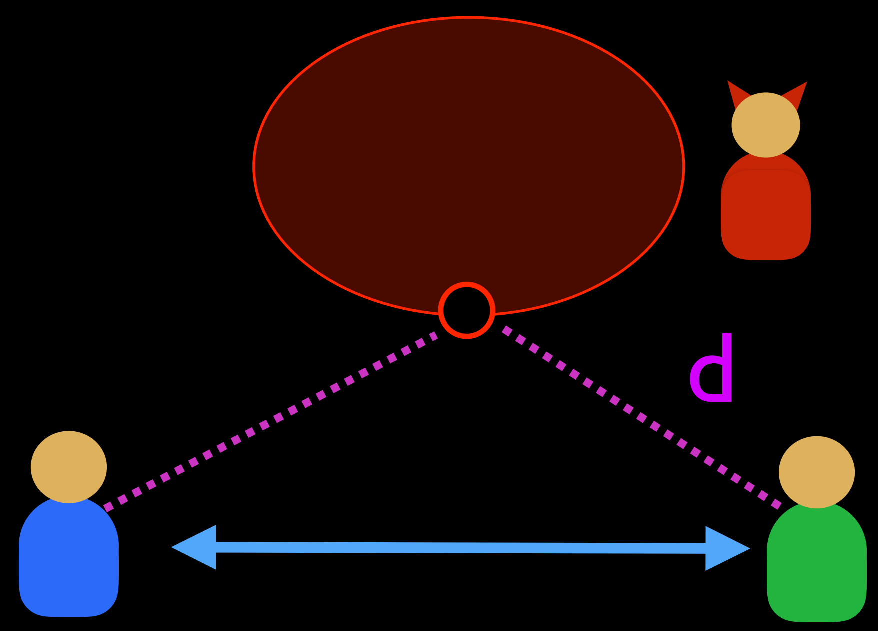

Provable avoidance

Alibi Routing by Levin et al. in SIGCOMM 2015

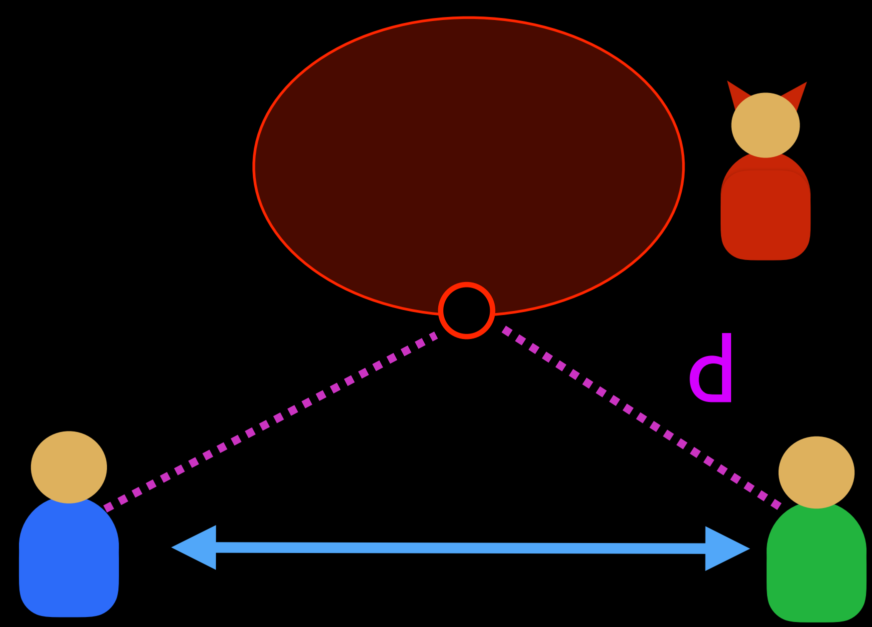

The *shortest* *possible* RTT thru to = 2 d / c

Provable avoidance

Alibi Routing by Levin et al. in SIGCOMM 2015

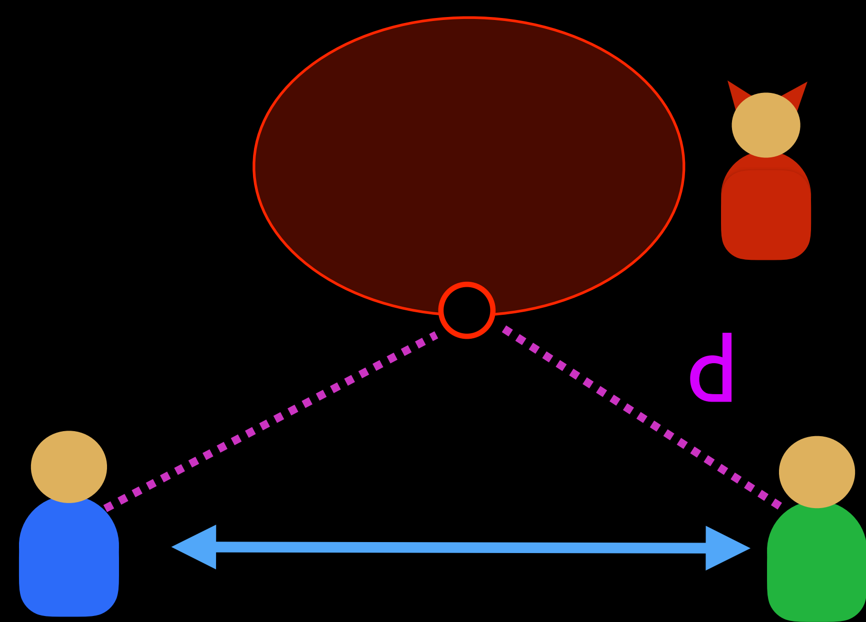

Measured RTT ≪ The shortest *possible* RTT thru to = 2 d / c

Provable avoidance

Alibi Routing by Levin et al. in SIGCOMM 2015

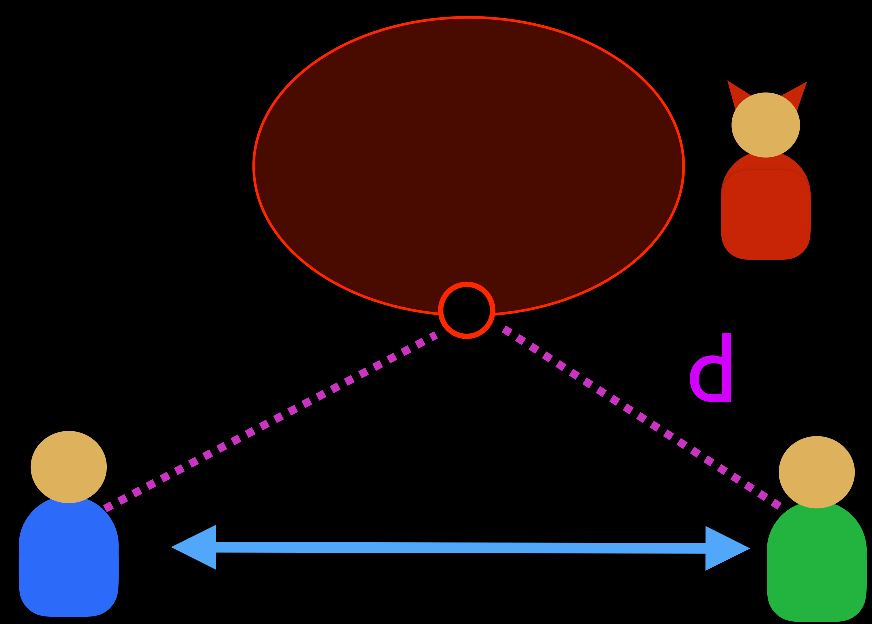

Measured RTT ≪ The shortest *possible* RTT thru to = 2 d / c

⇒ The packet could not have traversed to

Provable avoidance

Alibi Routing by Levin et al. in SIGCOMM 2015

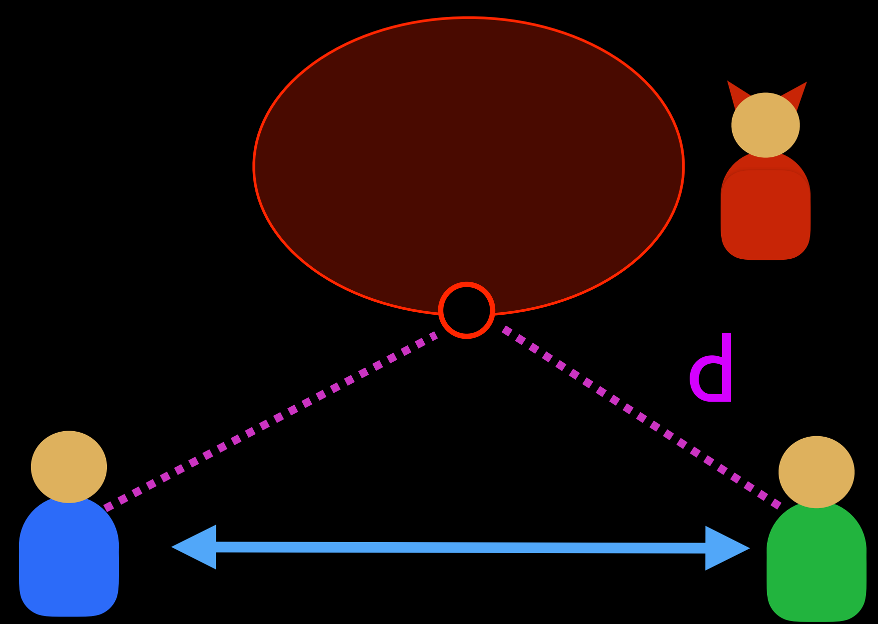

Measured RTT ≪ The shortest *possible* RTT thru to = 2 d / c

⇒ The packet could not have traversed to to

Alibi

Provable avoidance

Alibi Routing by Levin et al. in SIGCOMM 2015

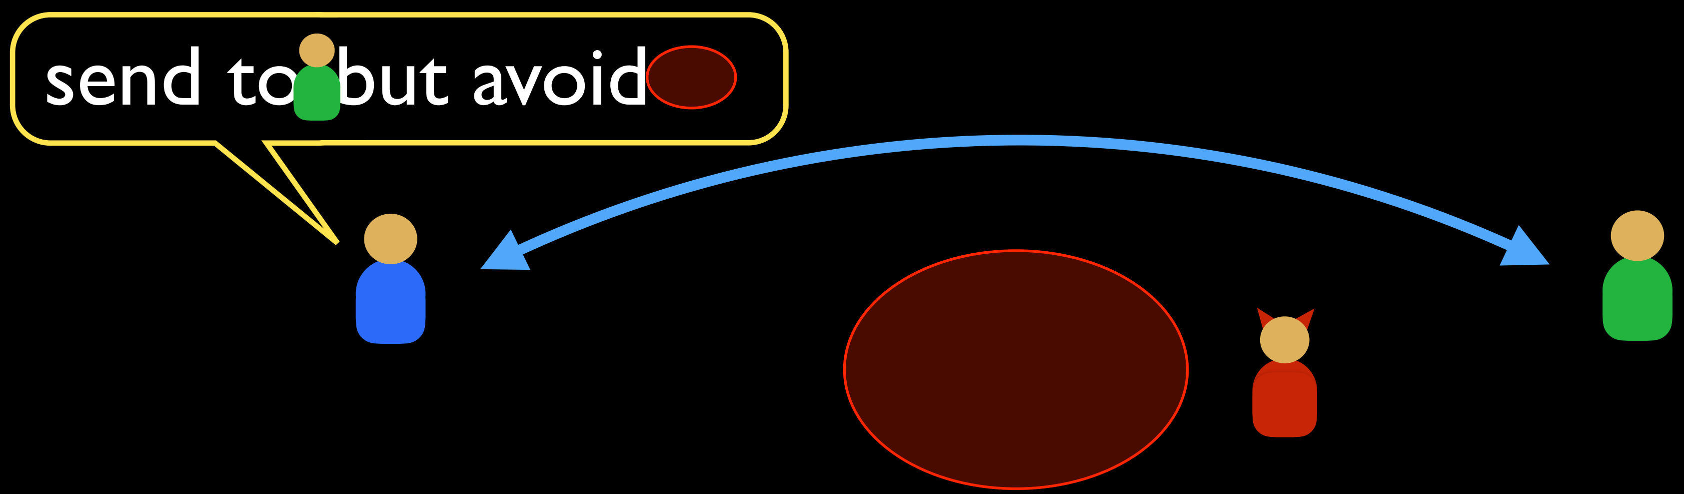

Provable avoidance

Alibi Routing by Levin et al. in SIGCOMM 2015

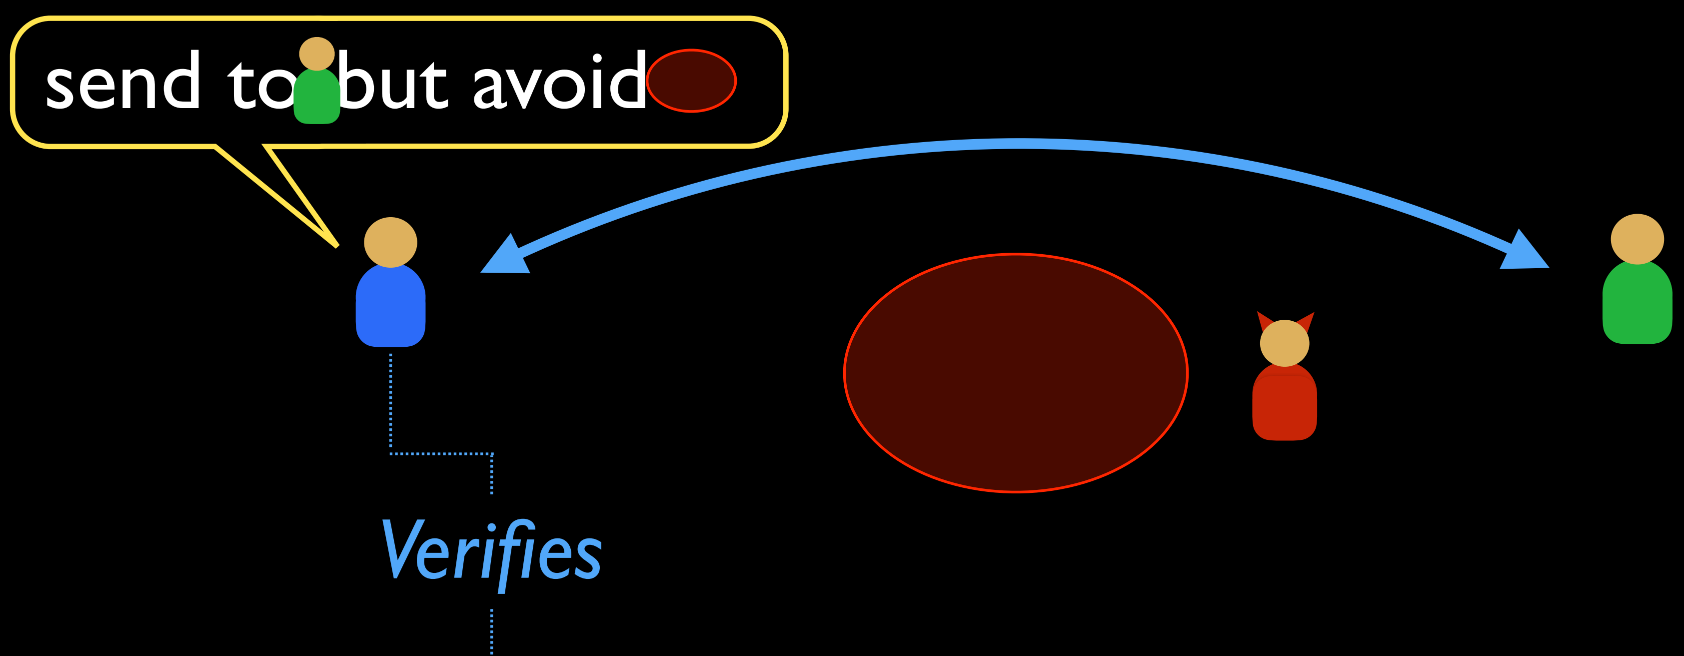

Provable avoidance

Alibi Routing by Levin et al. in SIGCOMM 2015

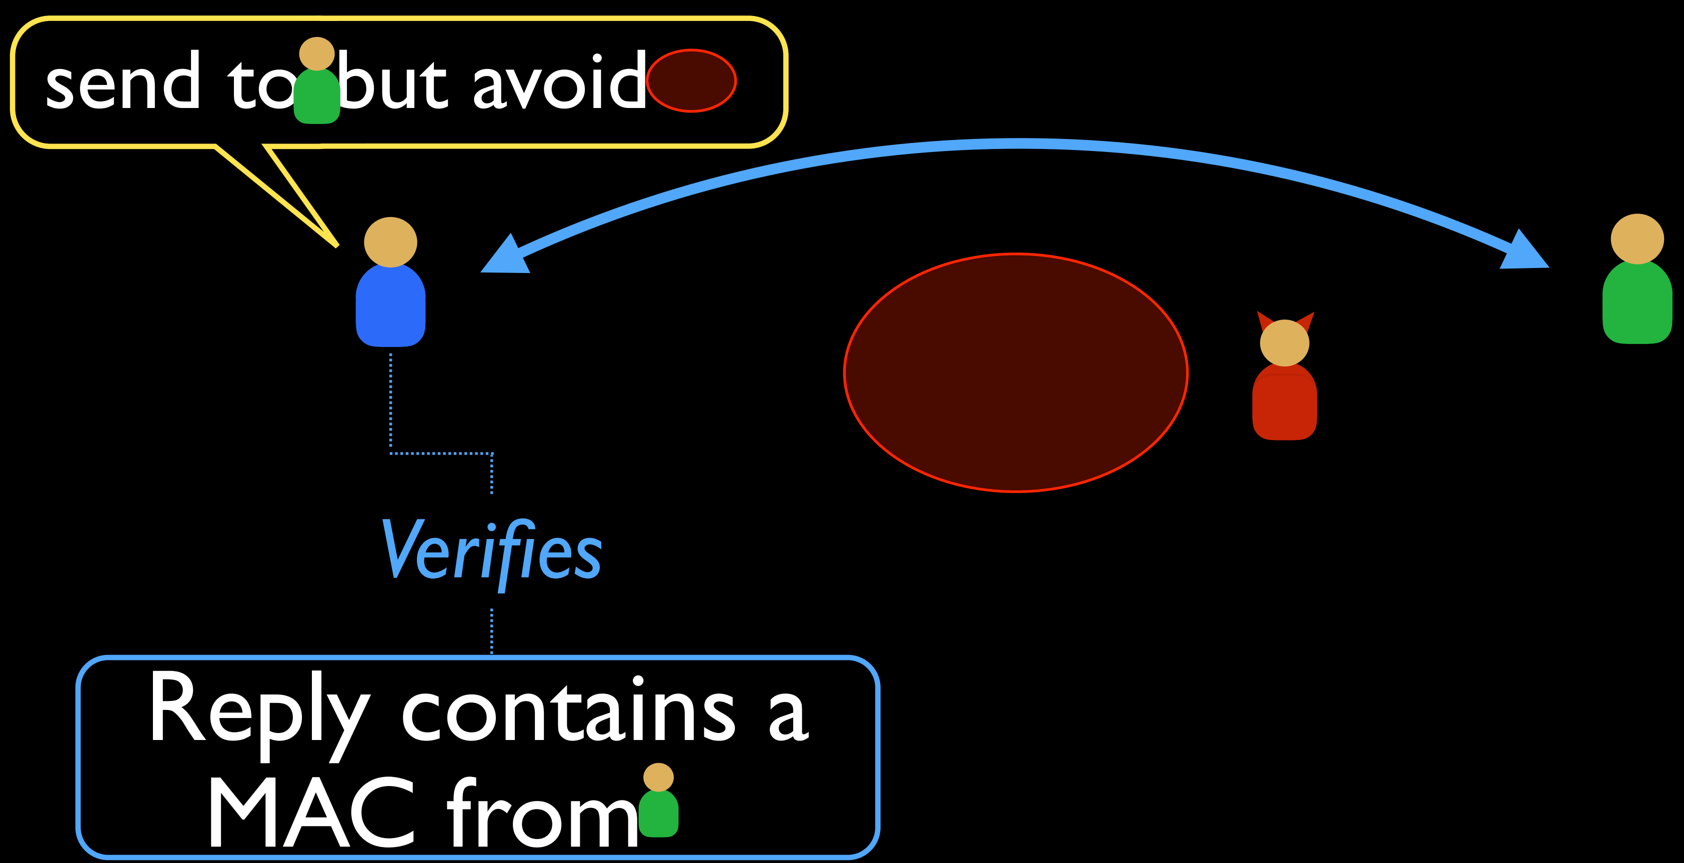

Provable avoidance

Alibi Routing by Levin et al. in SIGCOMM 2015

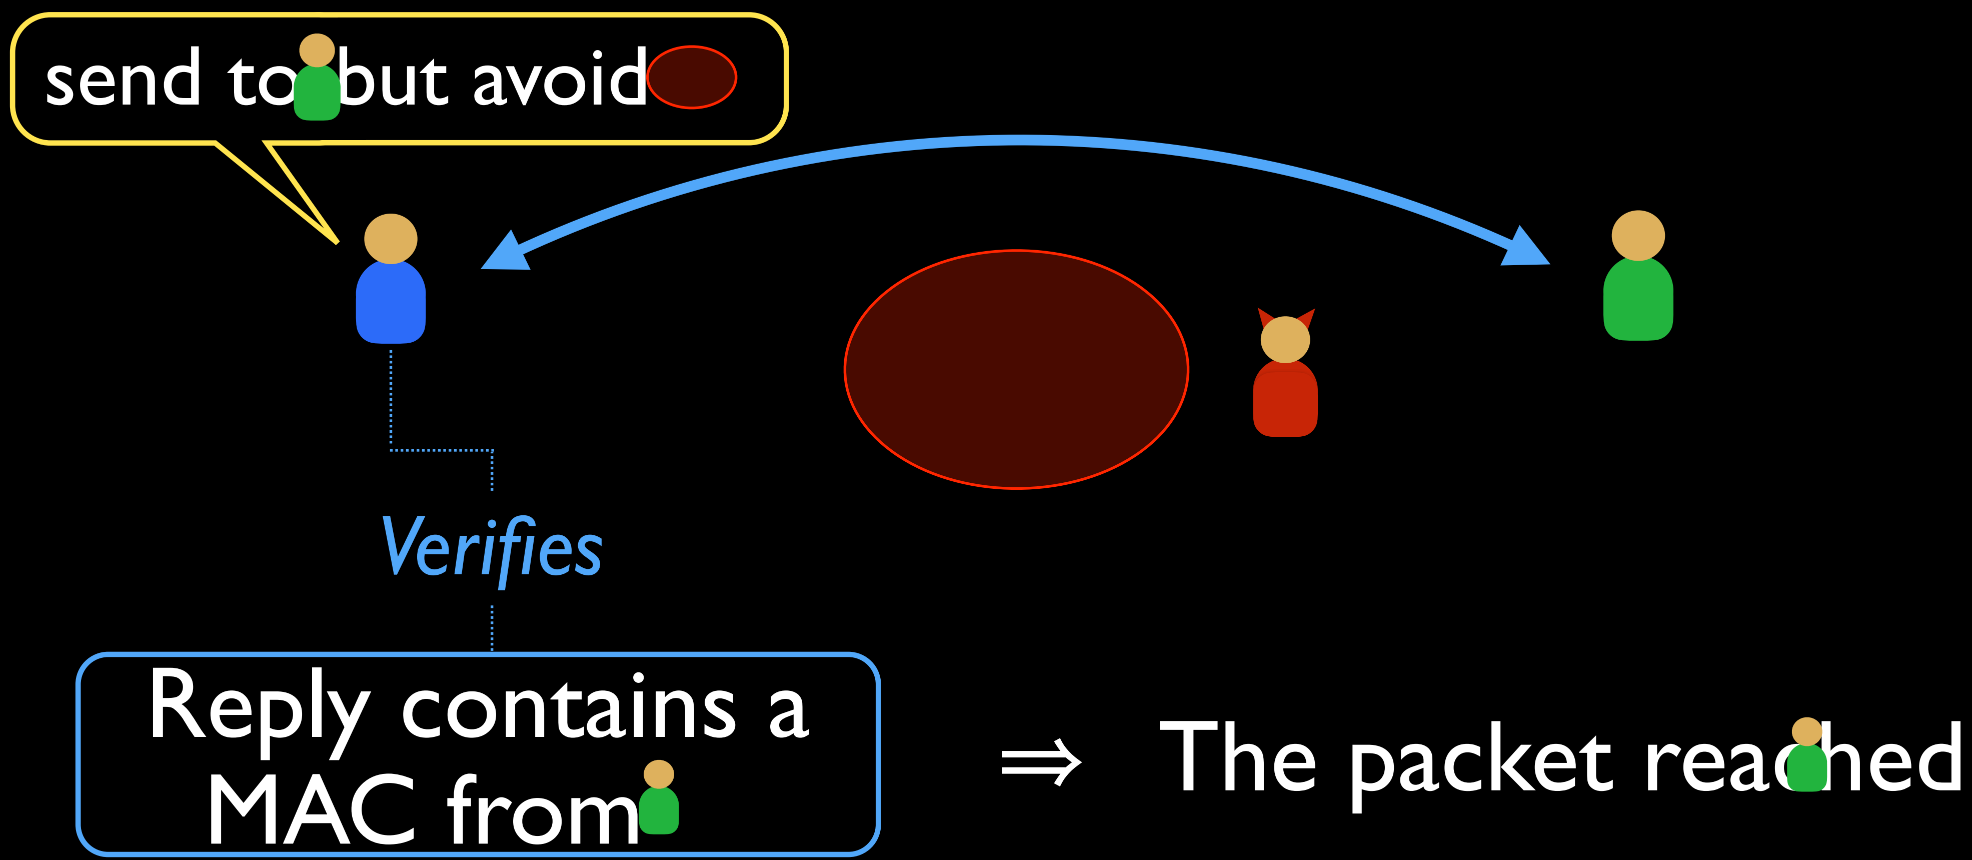

Provable avoidance

Alibi Routing by Levin et al. in SIGCOMM 2015

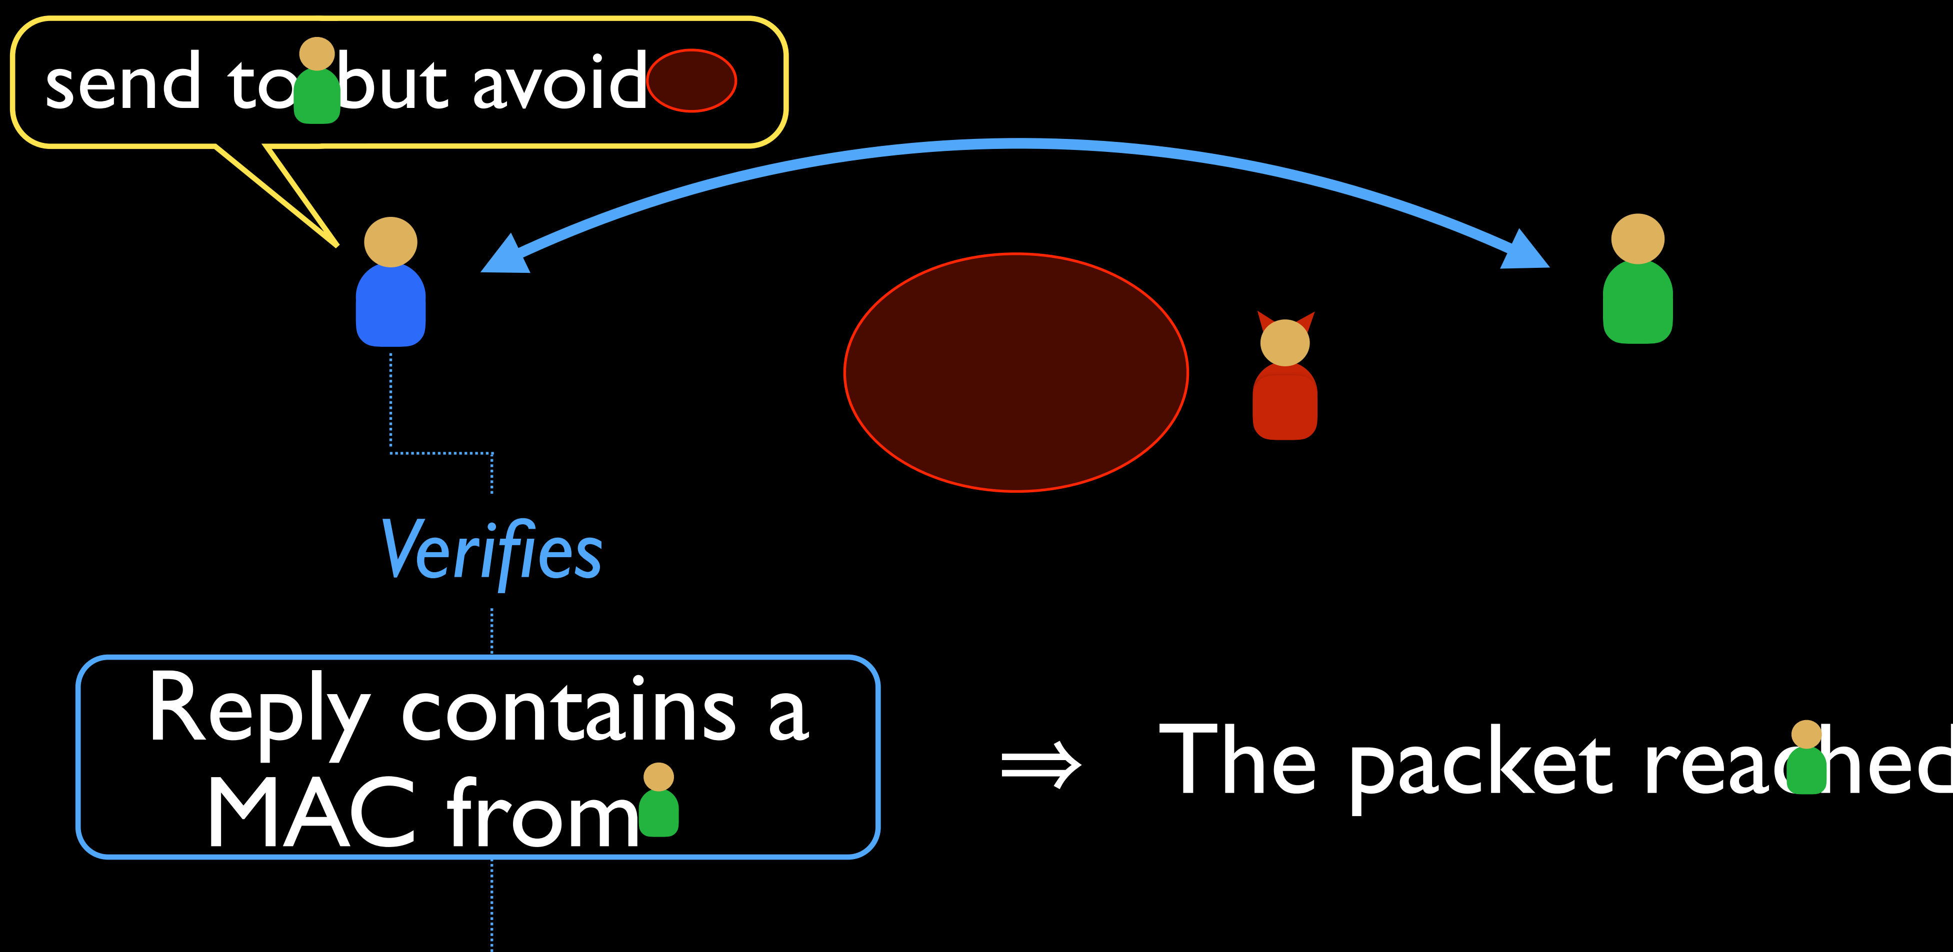

Provable avoidance

Alibi Routing by Levin et al. in SIGCOMM 2015

send to but avoid

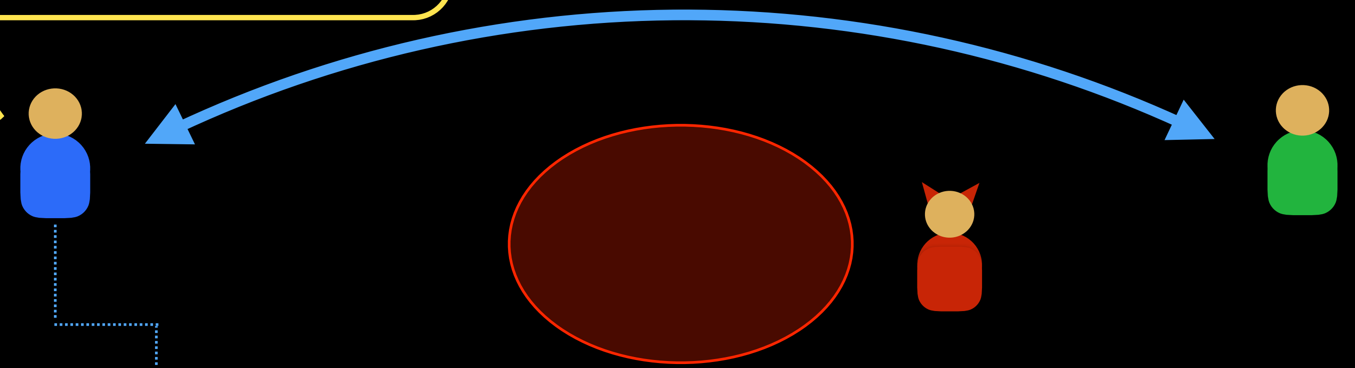

Verifies

Reply contains a MAC from

⇒ The packet reached

RTT less than shortest RTT thru

Provable avoidance

Alibi Routing by Levin et al. in SIGCOMM 2015

send to but avoid

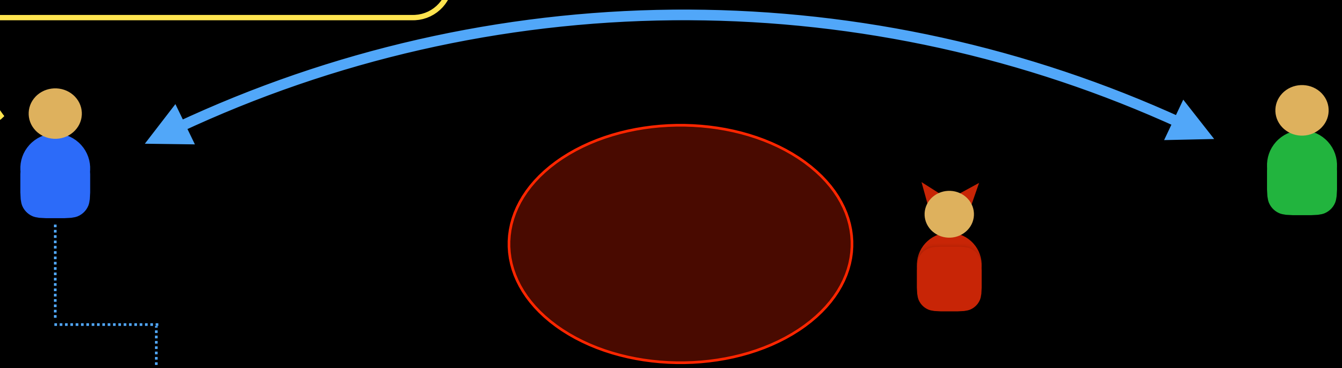

Verifies

Reply contains a MAC from ⇒ The packet reached

RTT less than shortest RTT thru ⇒ The packet could not have traversed &

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once

never traverse specified regions

Never-twice

entry&exit legs never traverse

Provide per-packet proof of avoidance

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once	Never-twice
never traverse specified regions	entry&exit legs never traverse

Provide per-packet proof of avoidance

DeTor: never-once avoidance

Avoid user specified geographic regions

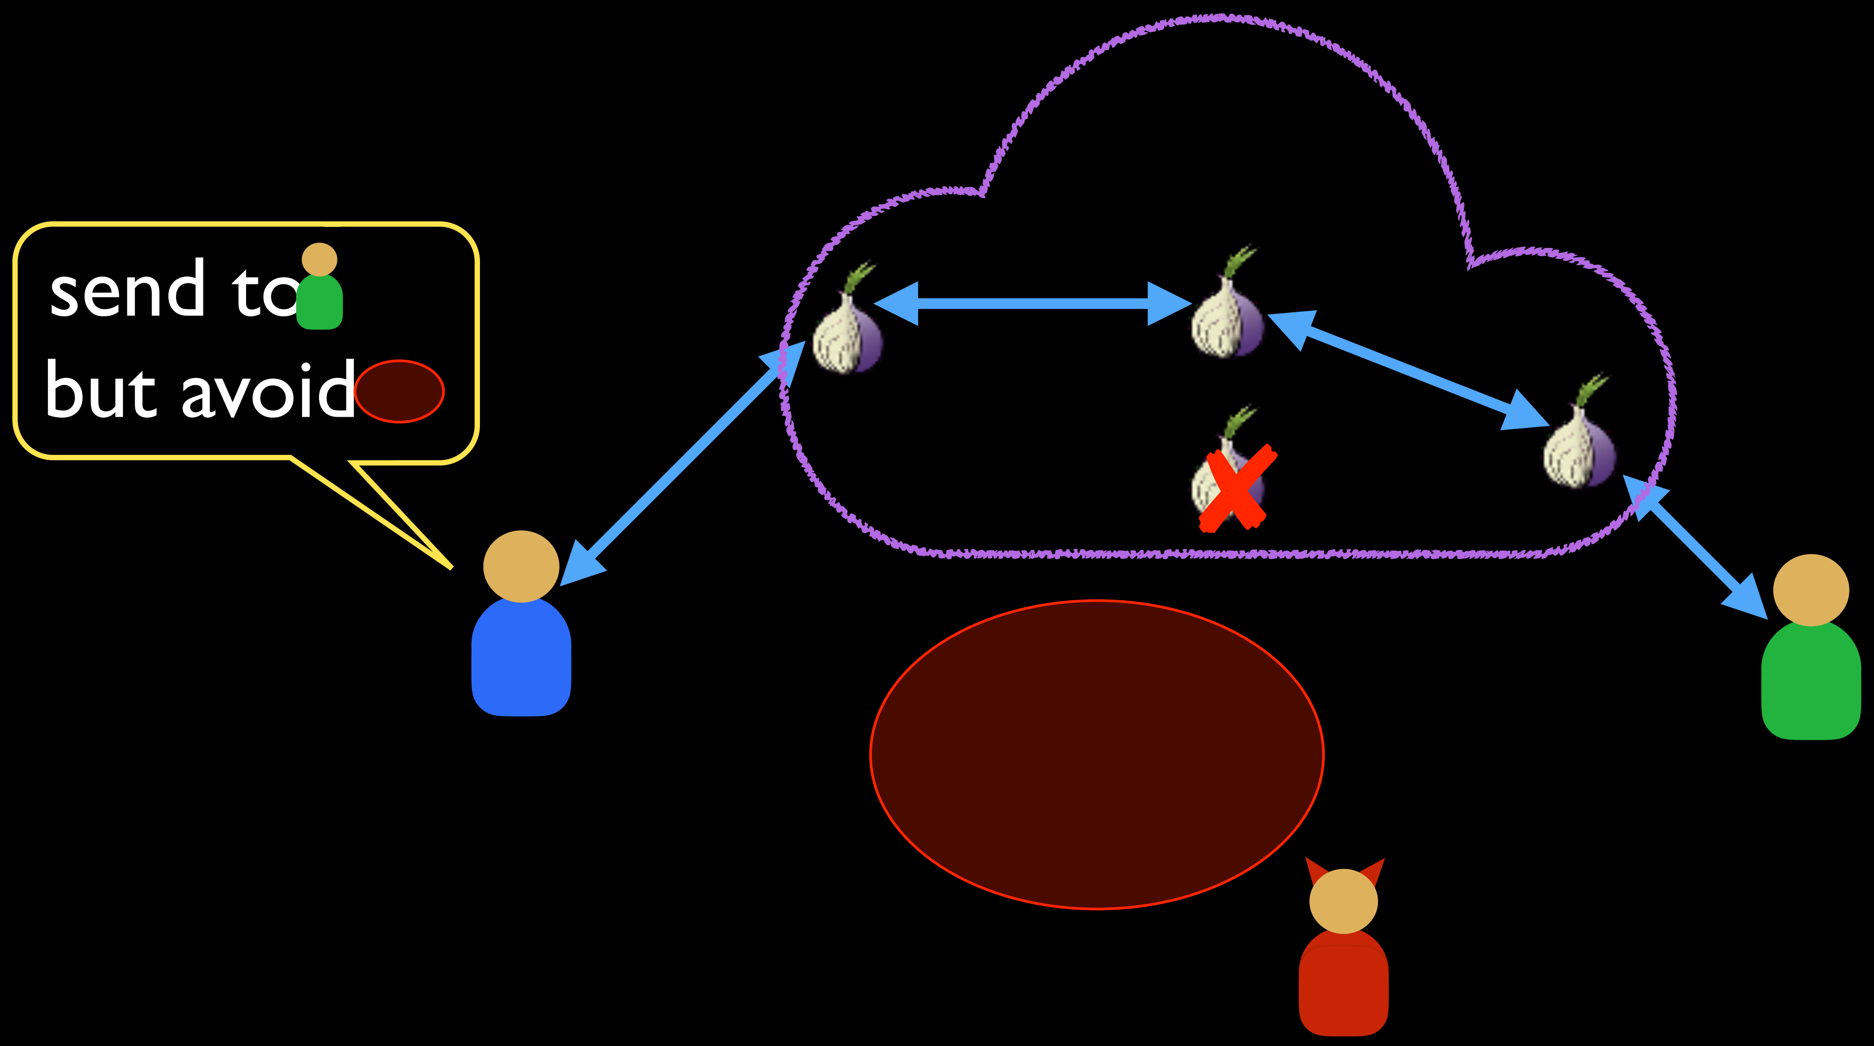

DeTor: never-once avoidance

The *shortest possible* RTT thr and to

shortest distance $= d_l$

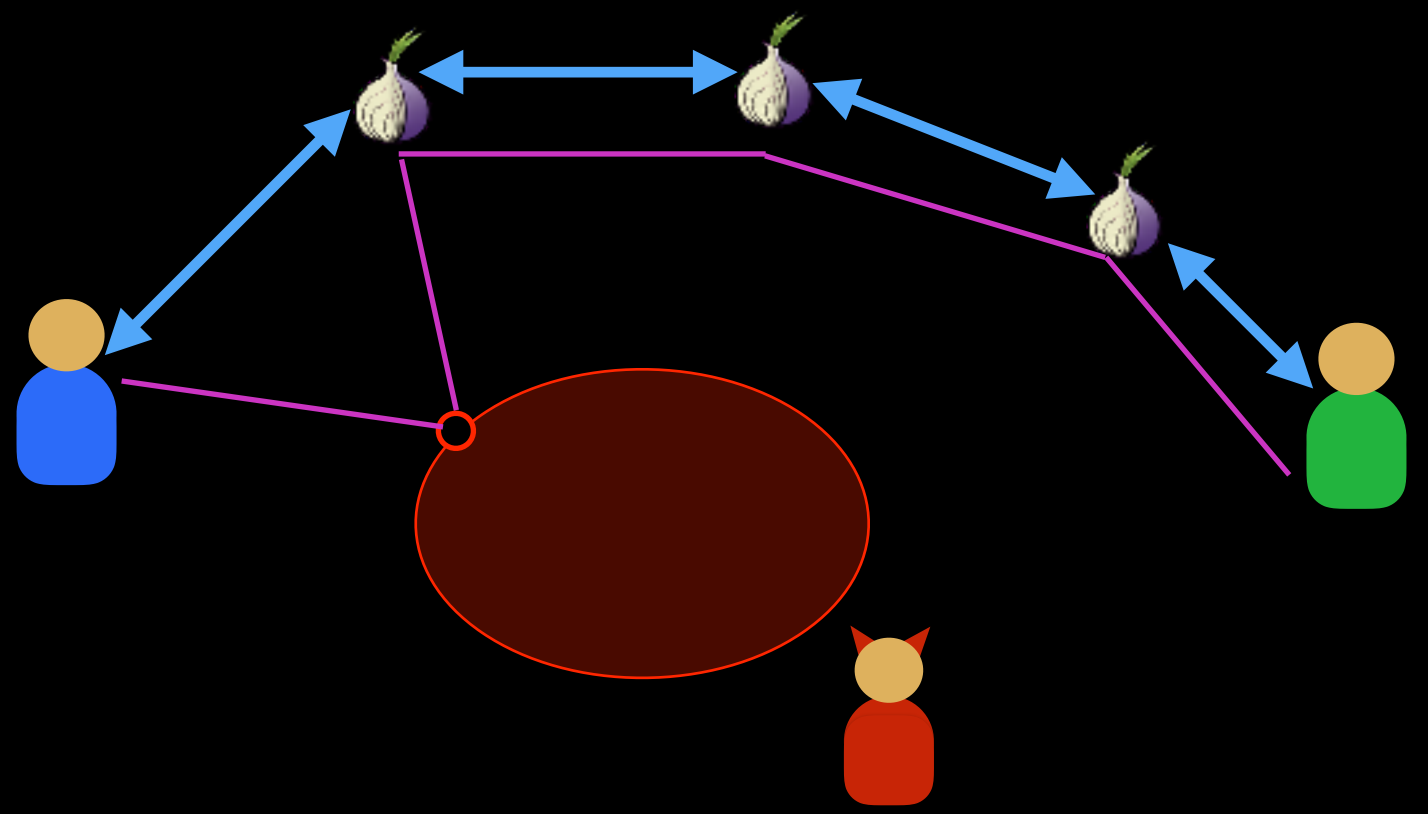

DeTor: never-once avoidance

The shortest *possible* RTT thr and to

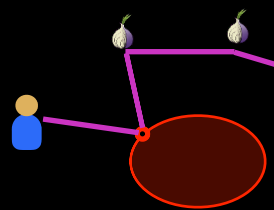

d_1

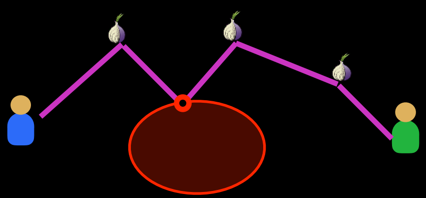

d_2

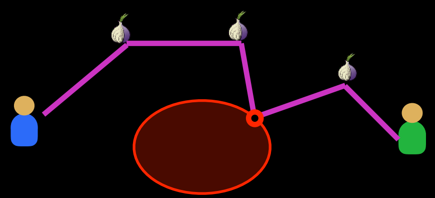

d_3

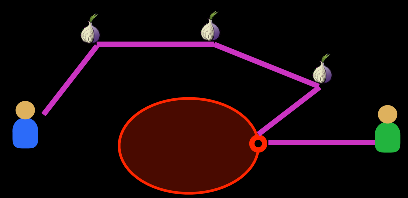

d_4

DeTor: never-once avoidance

The shortest *possible* RTT thru and to

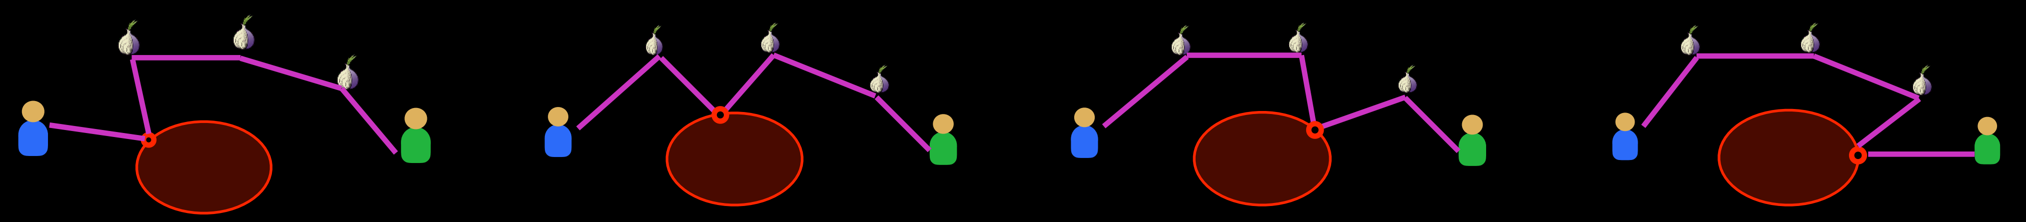

d_1 d_2 d_3 d_4

The shortest *possible* RTT thru and to $= 2 \min\{d_i\}/ c$

DeTor: never-once avoidance

The *shortest possible* RTT thr and to

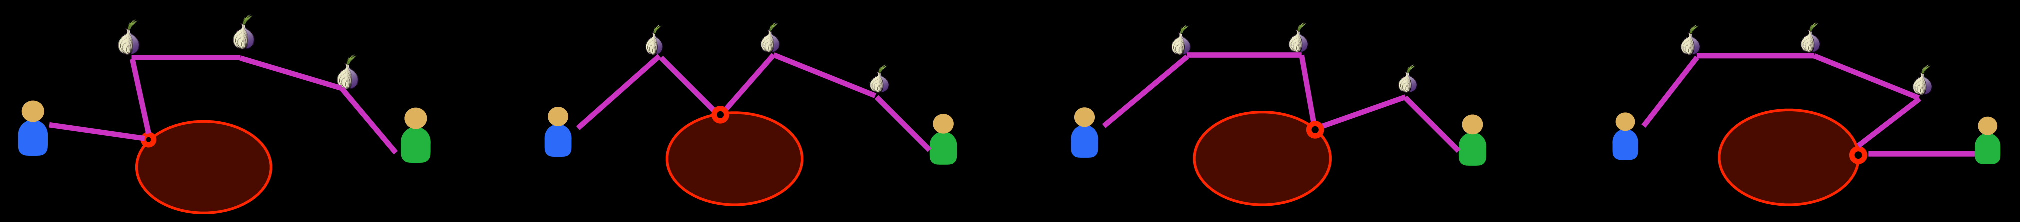

d_1 d_2 d_3 d_4

Measured RTT $\ll$ The shortest *possible* RTT thru and to $= 2 \min\{d_i\} / c$

DeTor: never-once avoidance

The shortest *possible* RTT thr and to

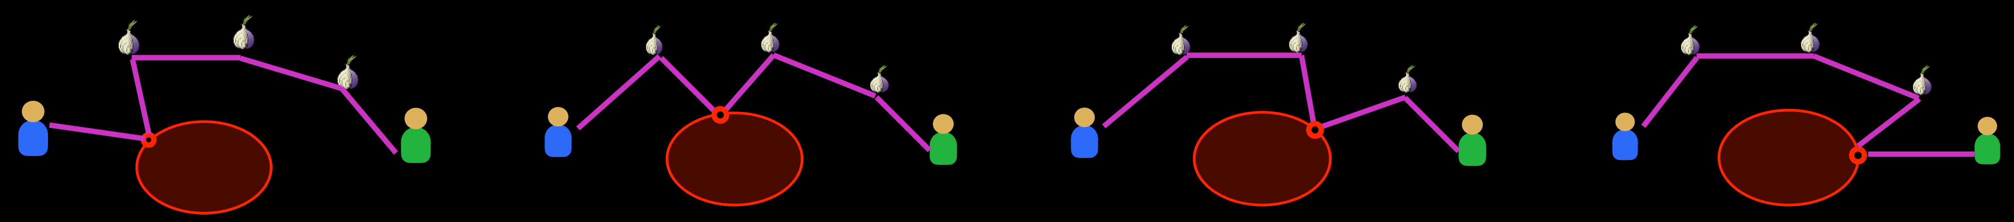

d_1 d_2 d_3 d_4

Measured RTT $\ll$ The shortest *possible* RTT thru and to $= 2 \min\{d_i\}/ c$

$\Rightarrow$ The packet could not have traversed to

DeTor: never-once avoidance

Achieving provable avoidance

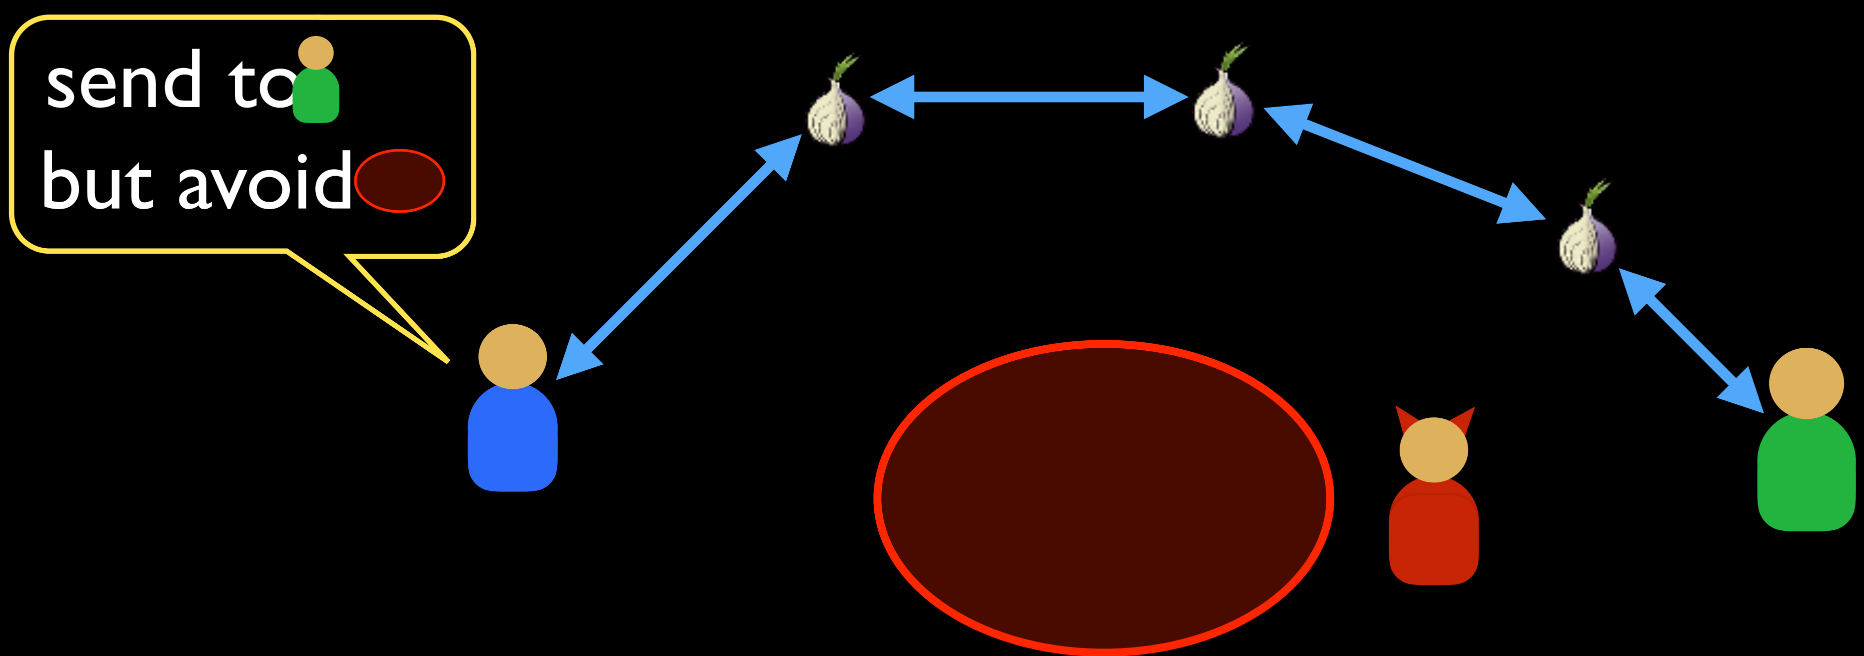

DeTor: never-once avoidance

Achieving provable avoidance

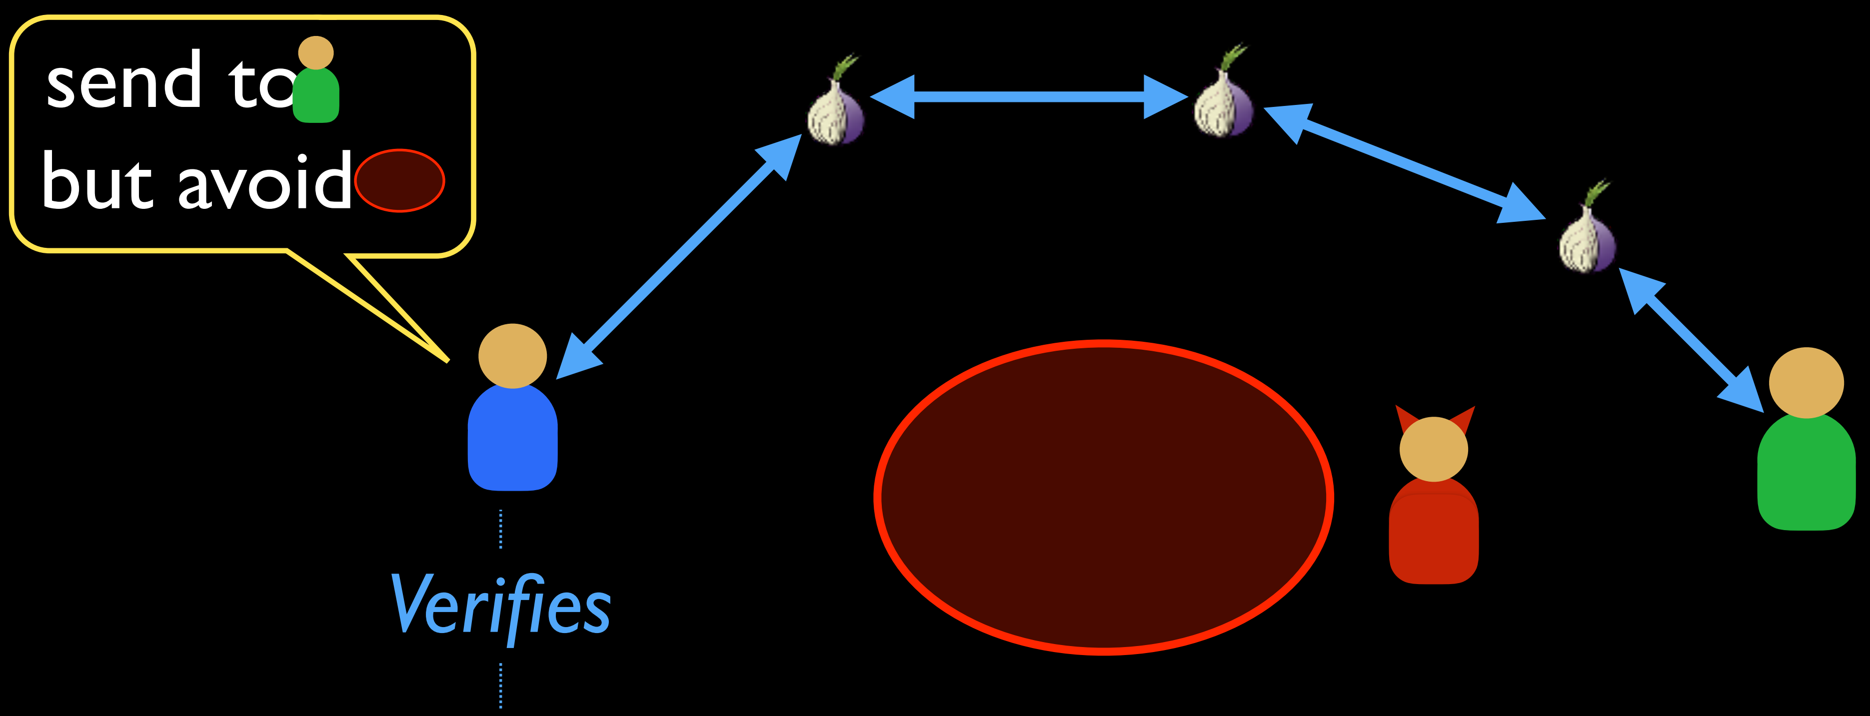

DeTor: never-once avoidance

Achieving provable avoidance

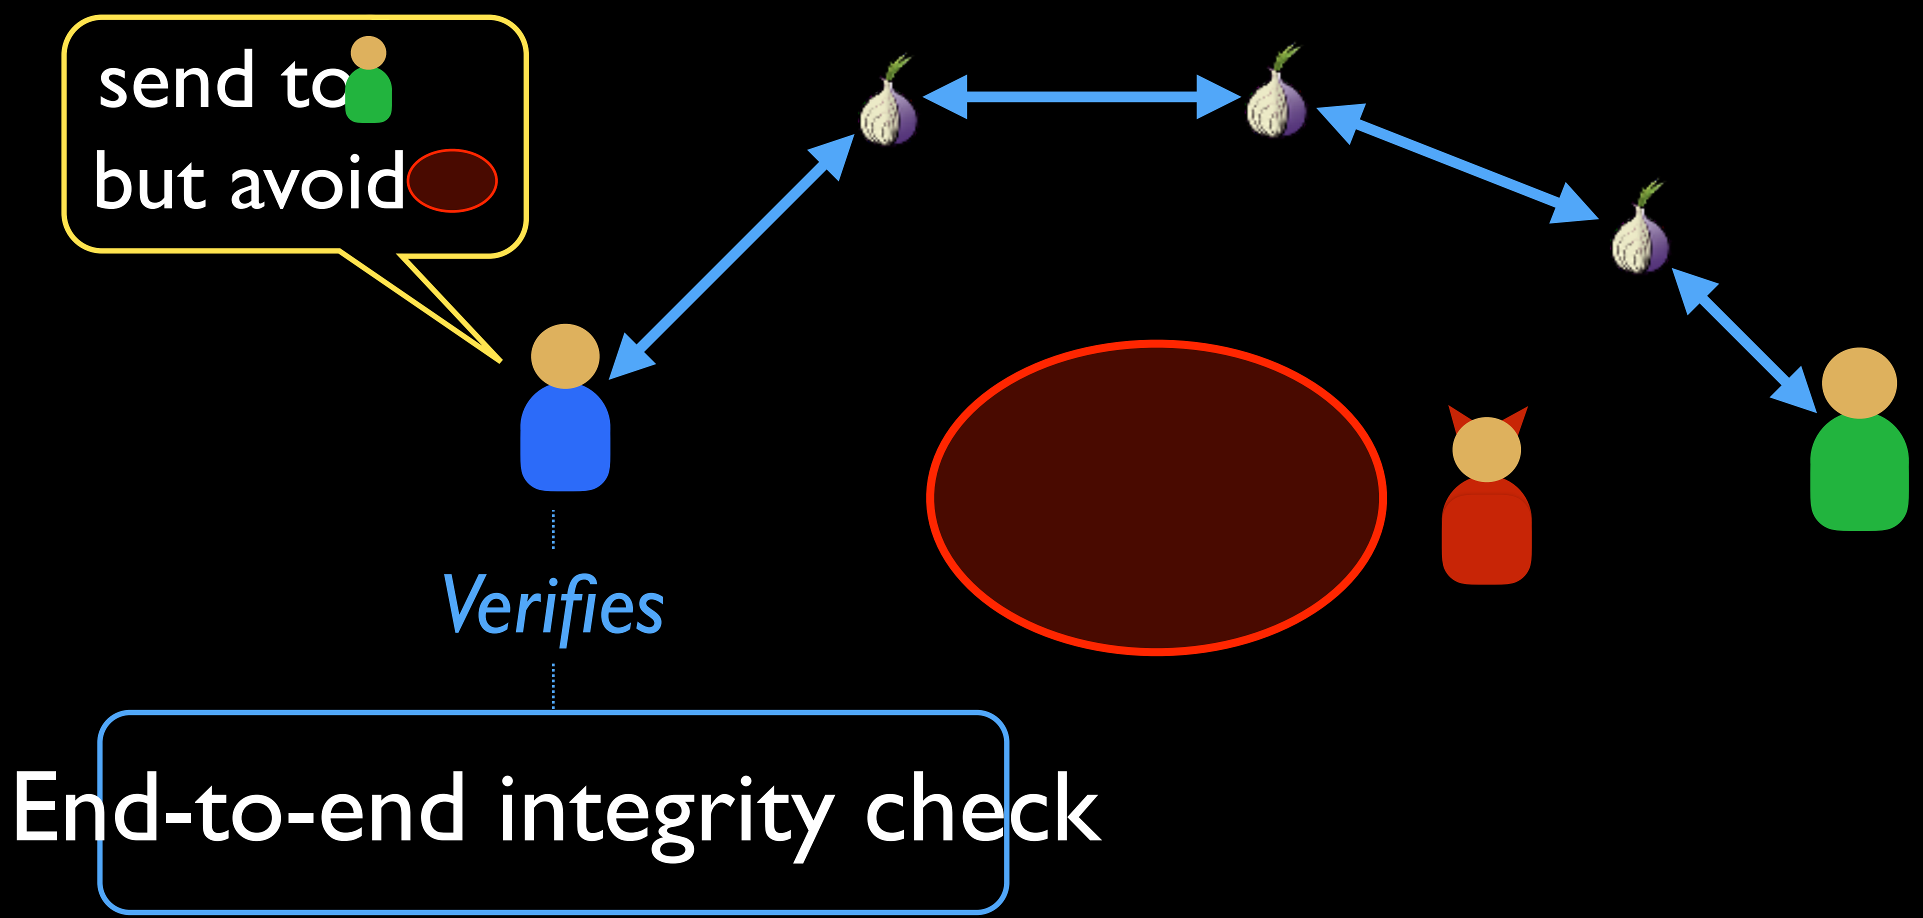

DeTor: never-once avoidance

Achieving provable avoidance

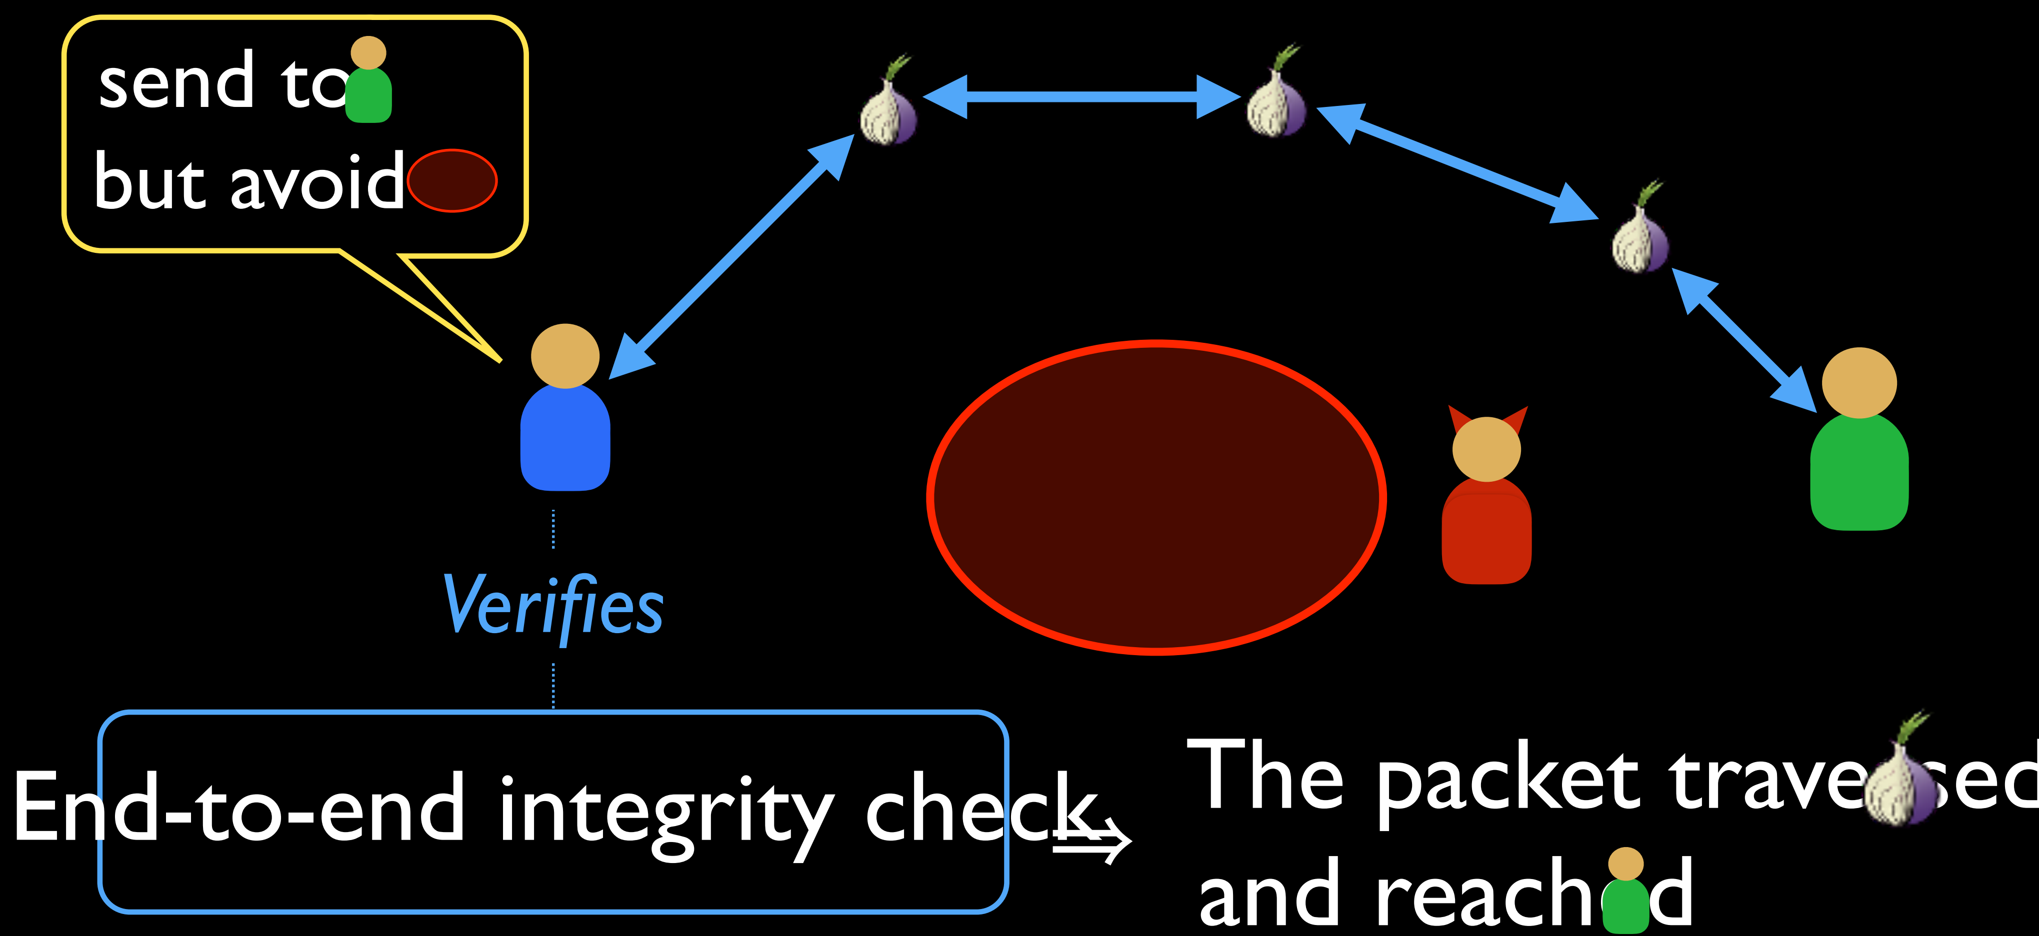

DeTor: never-once avoidance

Achieving provable avoidance

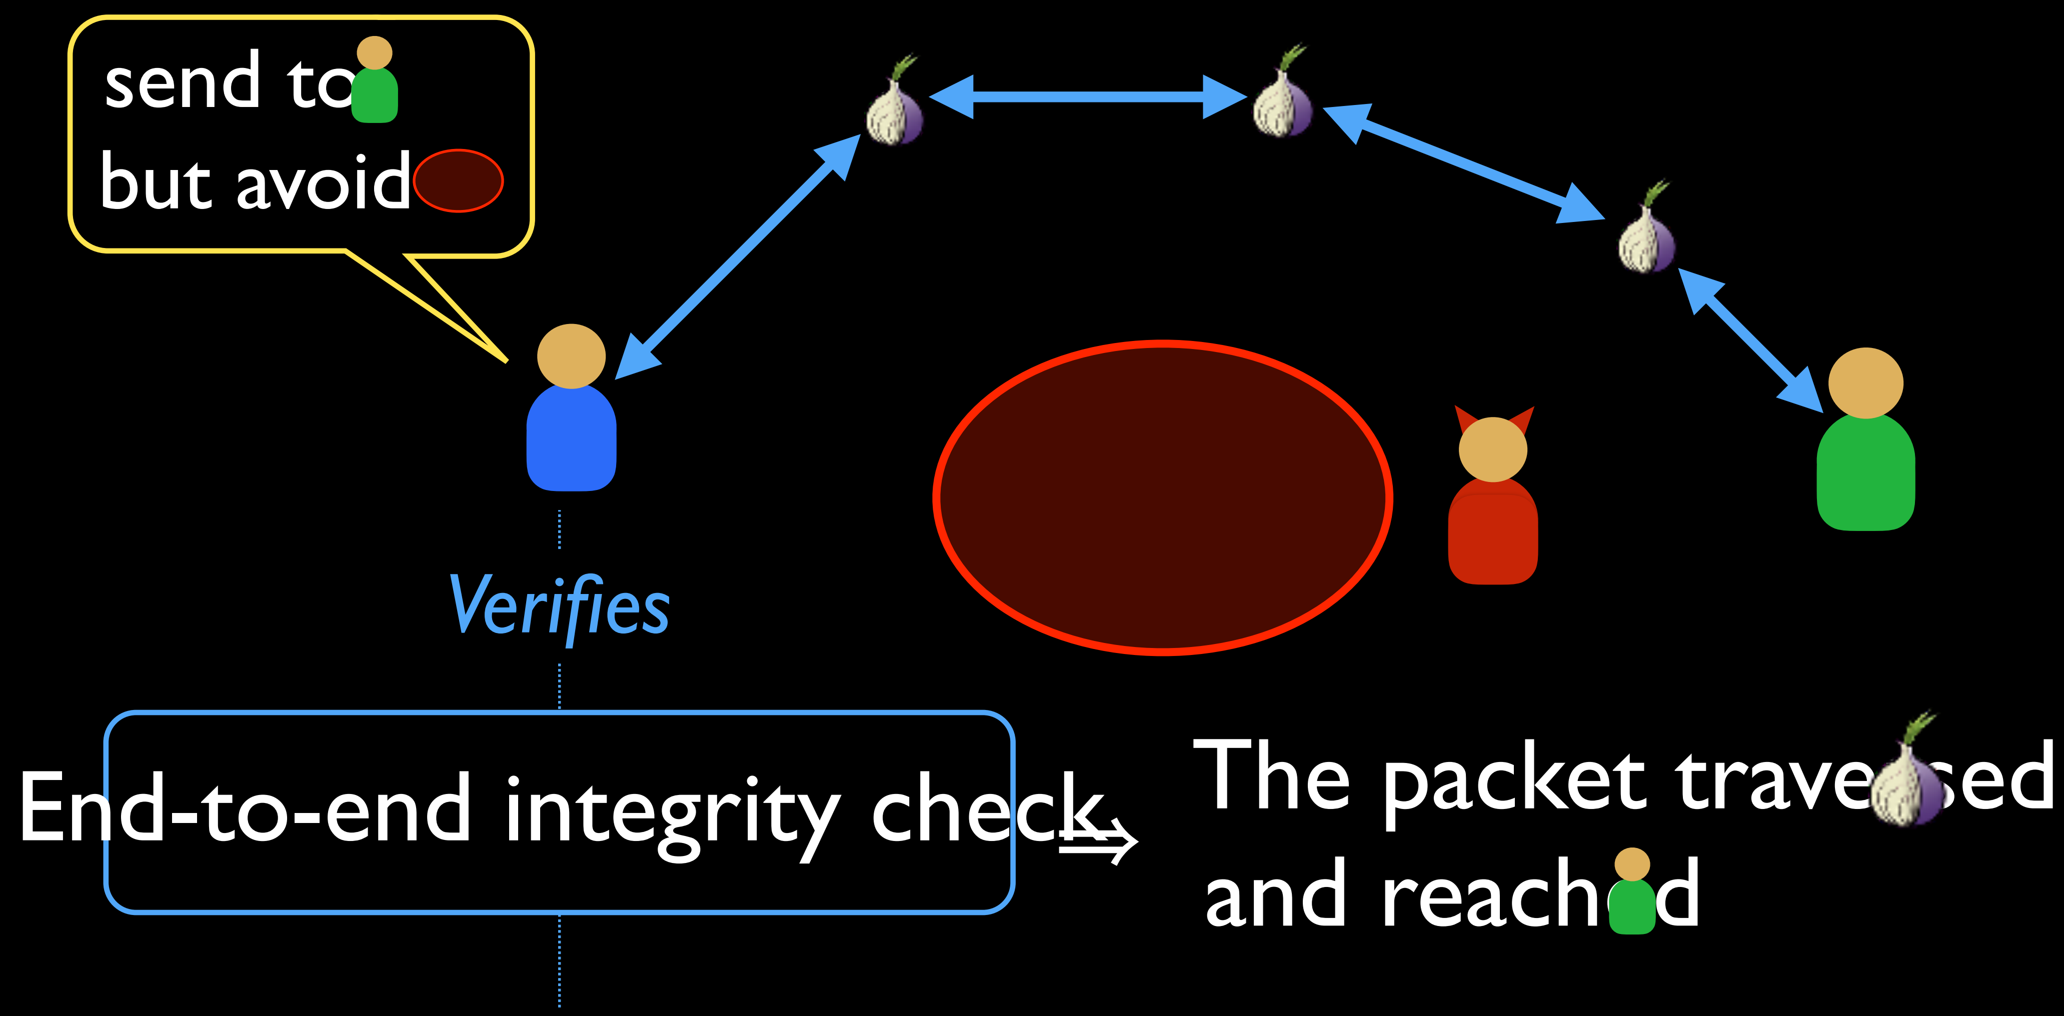

DeTor: never-once avoidance

Achieving provable avoidance

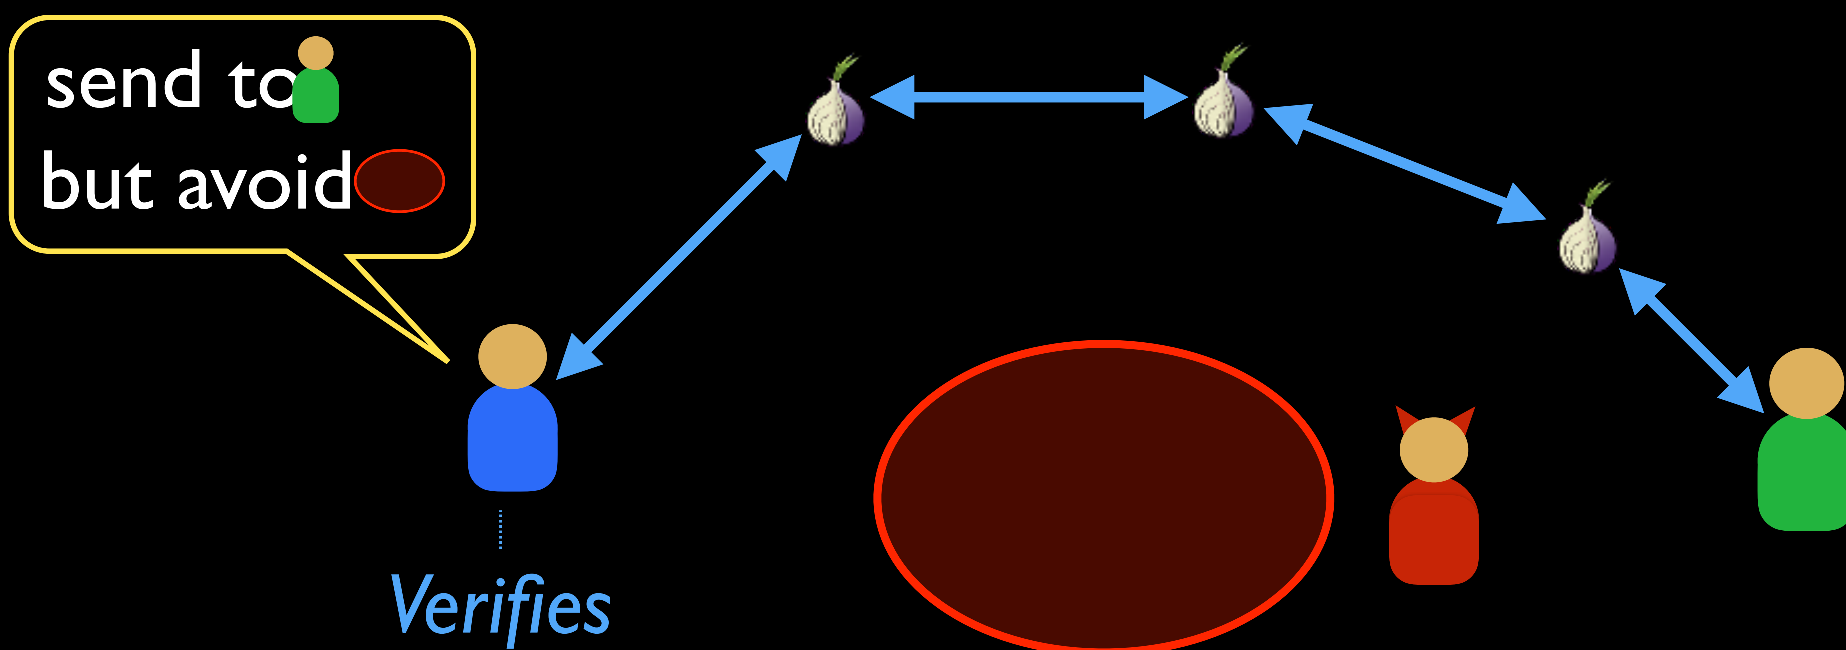

End-to-end integrity check ⇒ The packet traversed and reached

RTT less than smallest RTT thru and

DeTor: never-once avoidance

Achieving provable avoidance

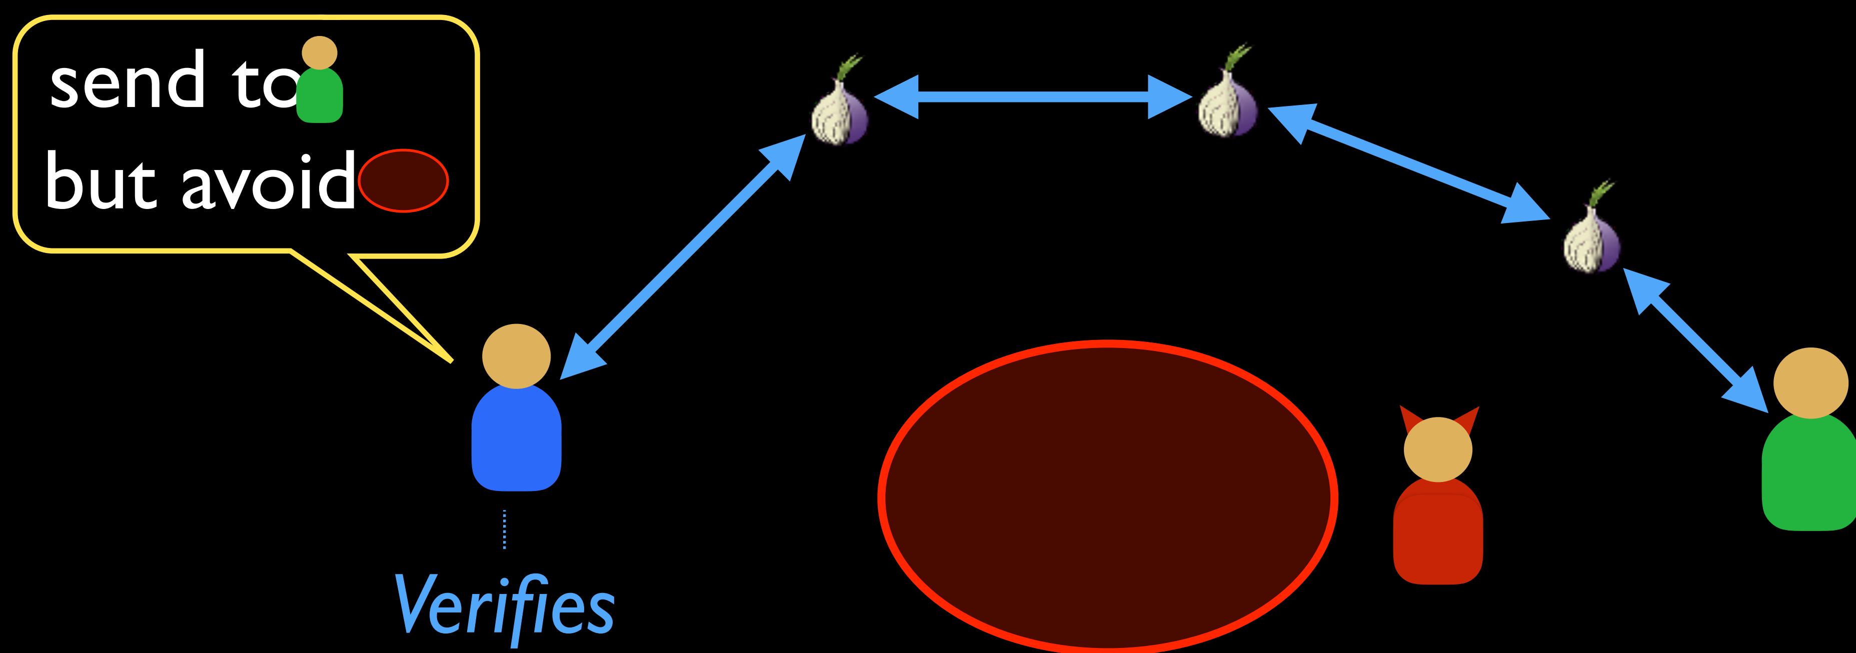

End-to-end integrity check ⇒ The packet traversed and reached

RTT less than smallest RTT thru and ⇒ The packet could not have traversed and

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once	Never-twice
never traverse specified regions	entry&exit legs never traverse

Provide per-packet proof of avoidance

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once

never traverse specified regions

Never-twice

entry&exit legs never traverse

Provide per-packet proof of avoidance

DeTor: never-twice avoidance

Entry and exit legs never traverse the same region

DeTor: never-twice avoidance

Entry and exit legs never traverse the same region

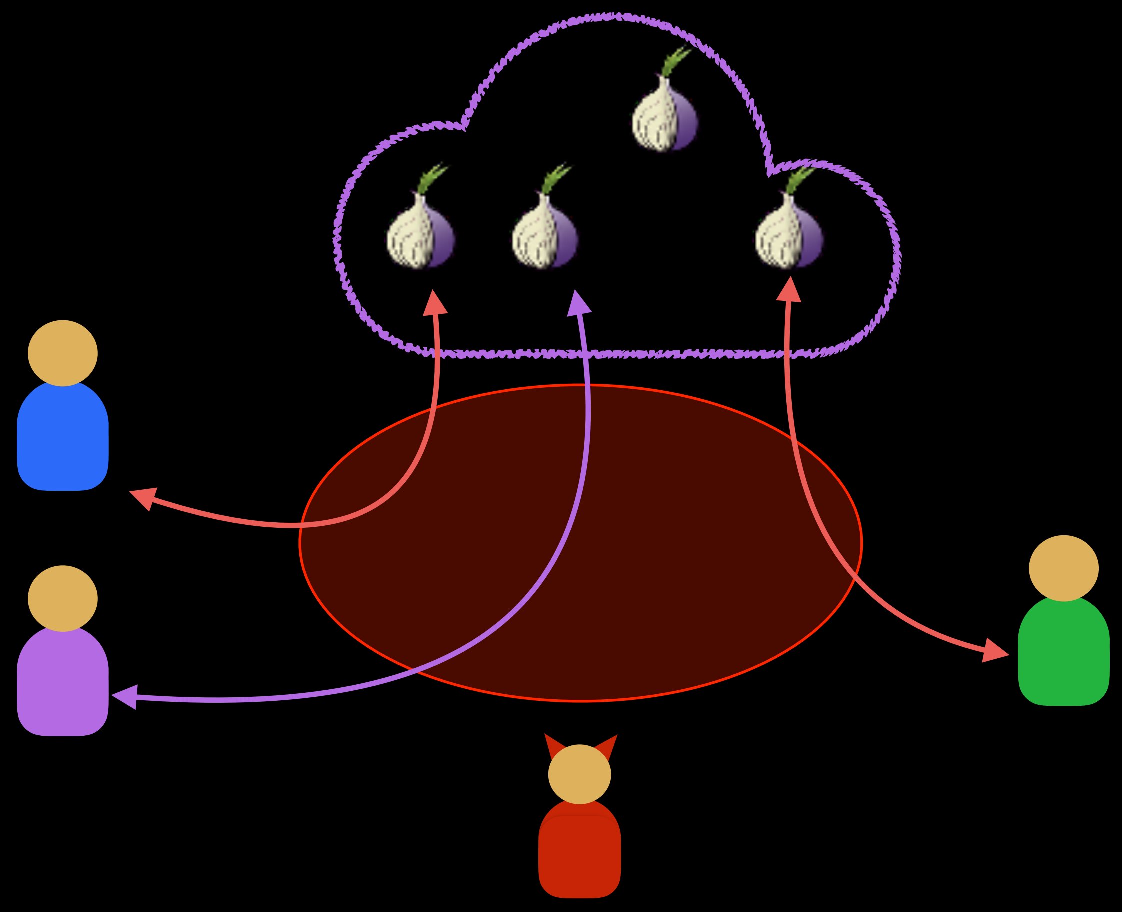

DeTor: never-twice avoidance

Entry and exit legs never traverse the same region

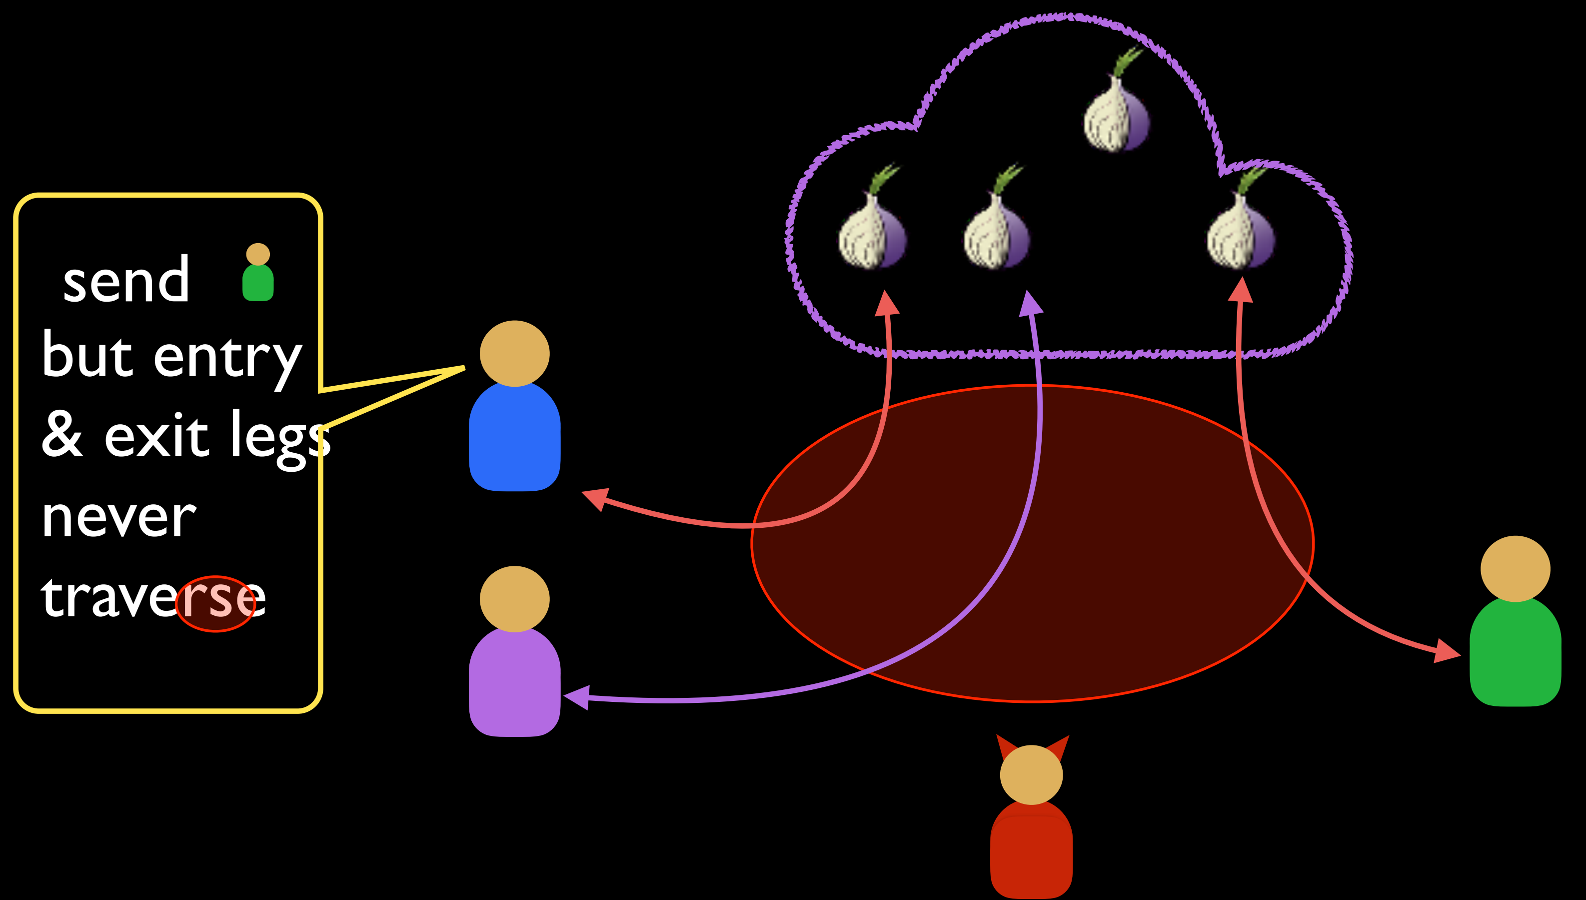

DeTor: never-twice avoidance

Entry and exit legs never traverse the same region

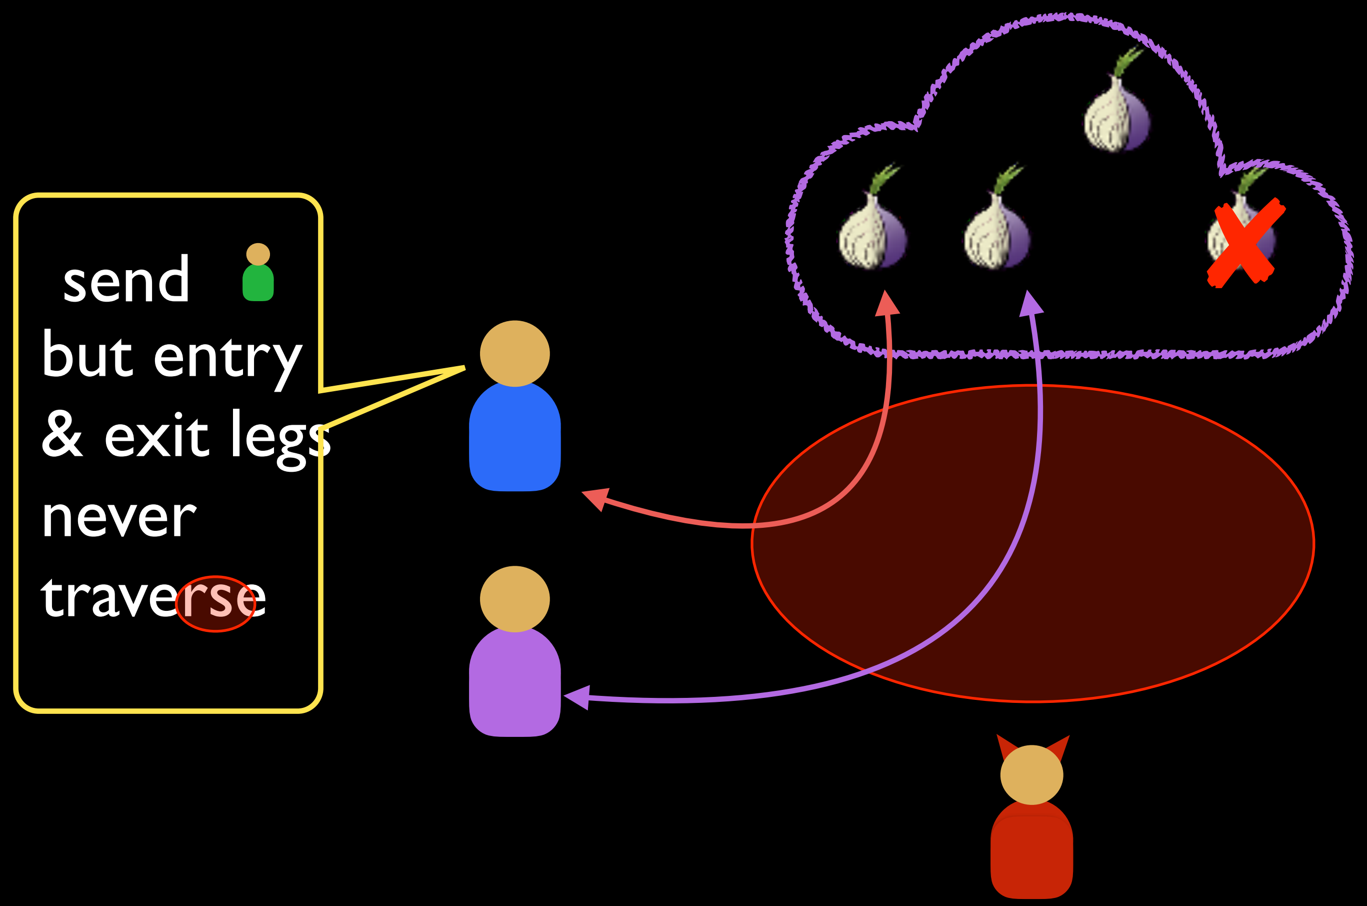

DeTor: never-twice avoidance

Entry and exit legs never traverse the same region

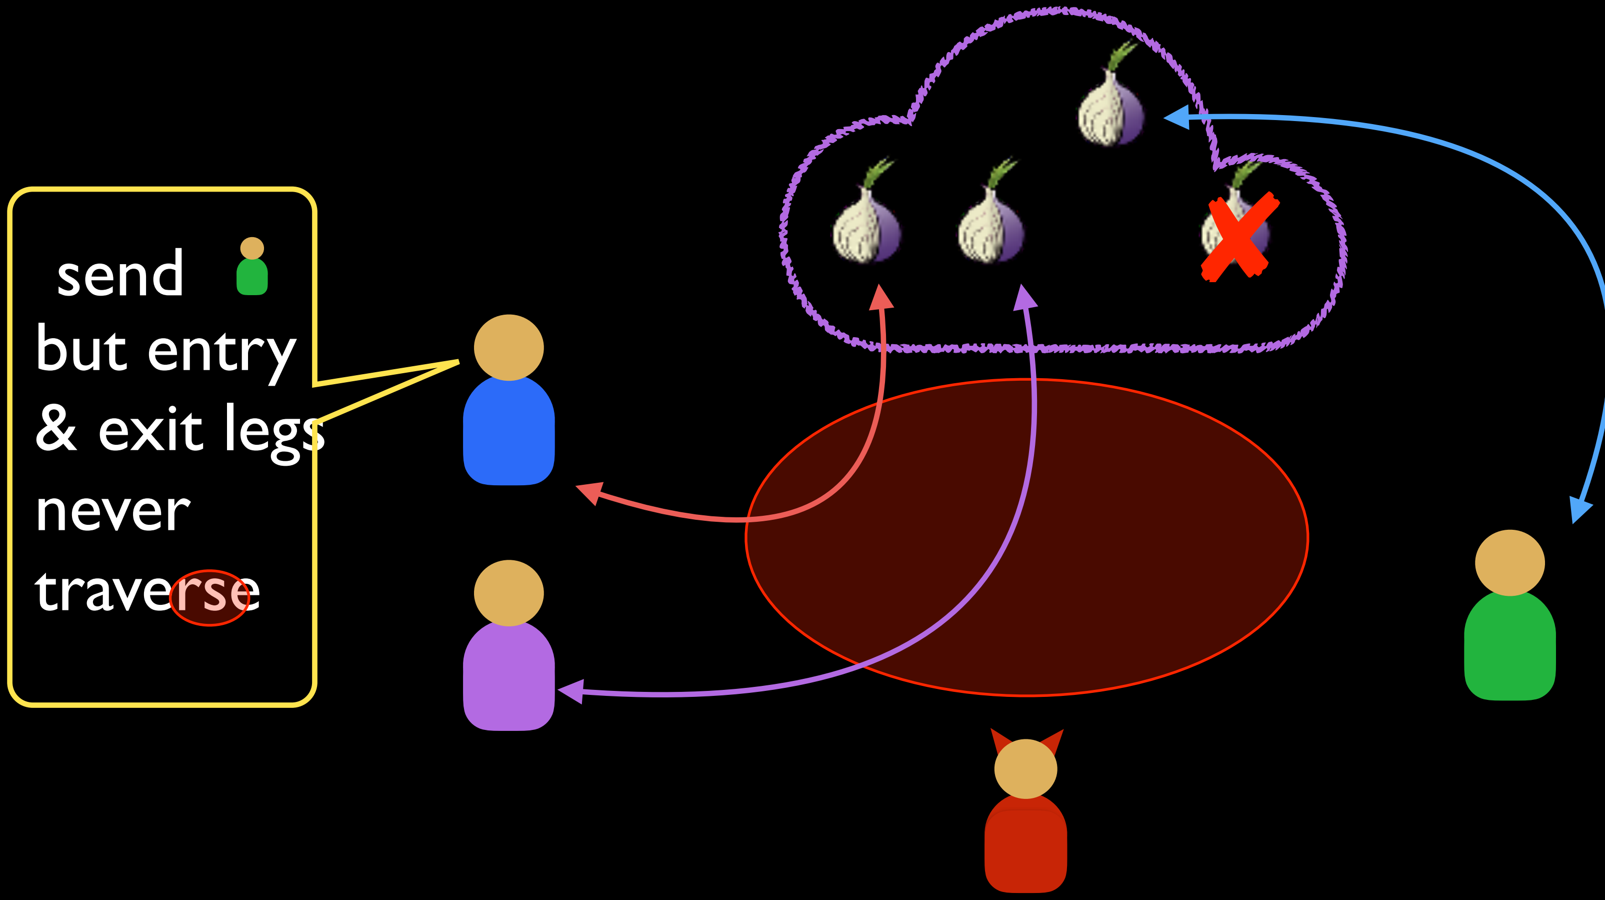

DeTor: never-twice avoidance

Where could packets possibly reach

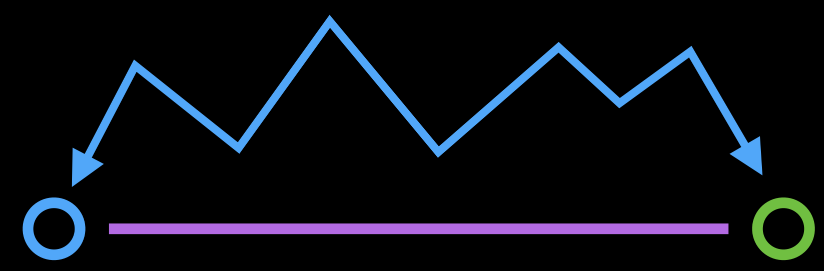

Measured RTT = shortest possible RTT + extra

DeTor: never-twice avoidance

Where could packets possibly reach

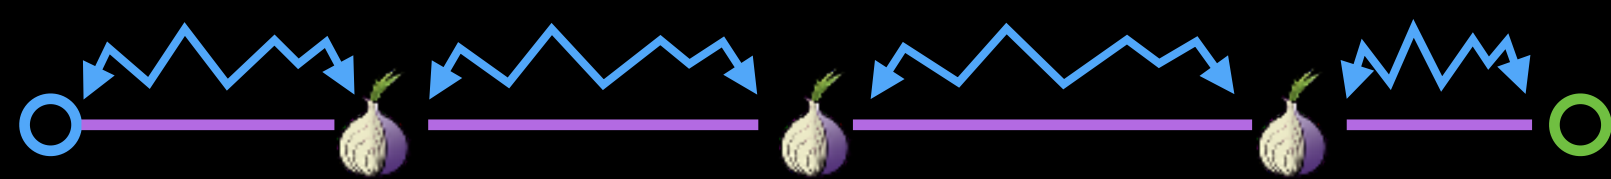

Measured RTT = Σ shortest possible RTT + Σ extra

DeTor: never-twice avoidance

Where could packets possibly reach

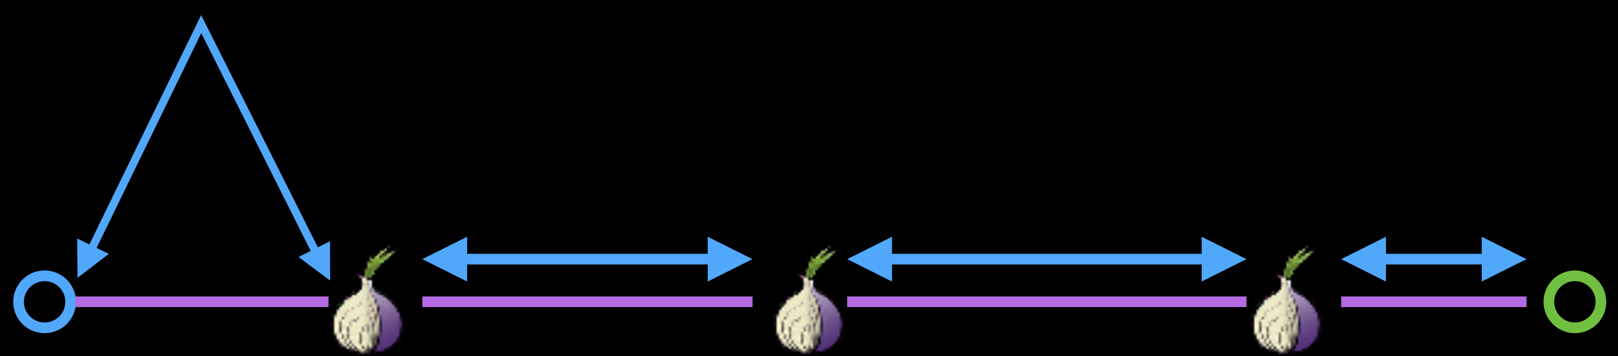

Measured RTT $= \Sigma$ shortest possible RTT $+ \Sigma$ extra

DeTor: never-twice avoidance

Where could packets possibly reach

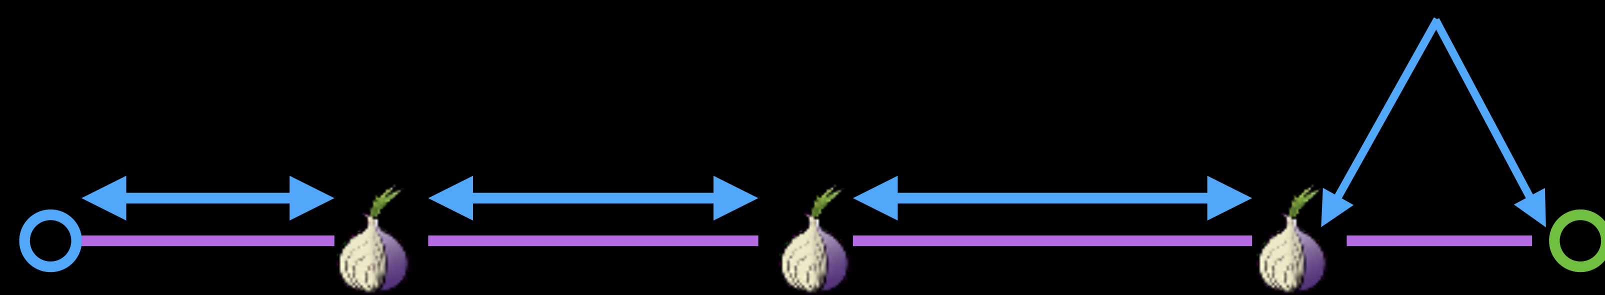

Measured RTT = Σ shortest possible RTT + Σ extra

DeTor: never-twice avoidance

Where could packets possibly reach

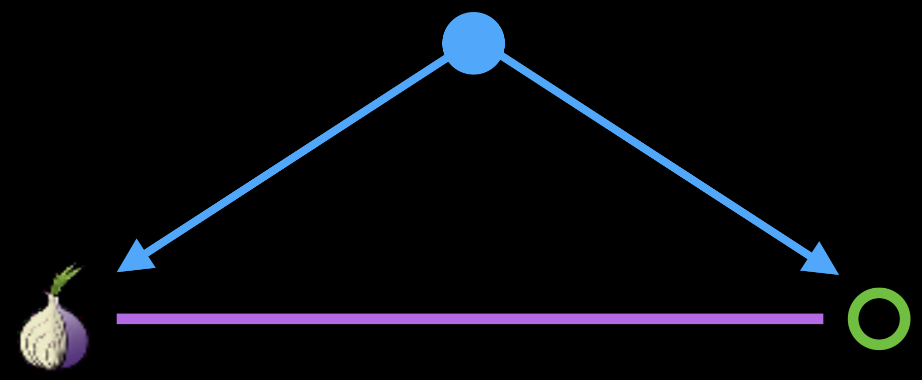

DeTor: never-twice avoidance

Where could packets possibly reach

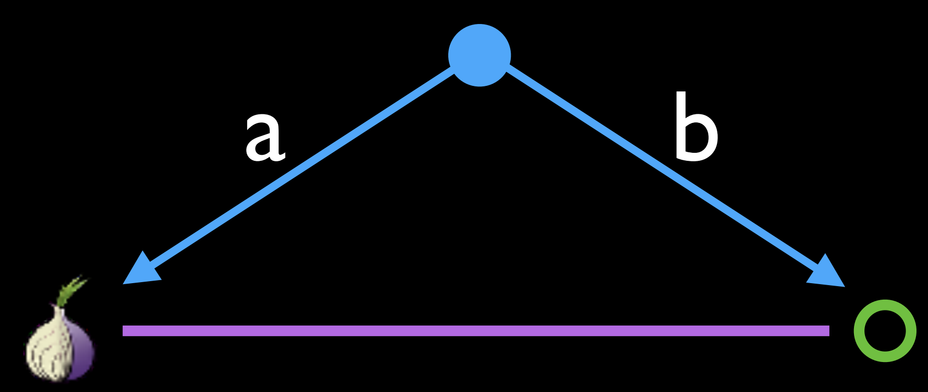

DeTor: never-twice avoidance

Where could packets possibly reach

Upper bound RTT $\geq$ 2 (a+b) / c

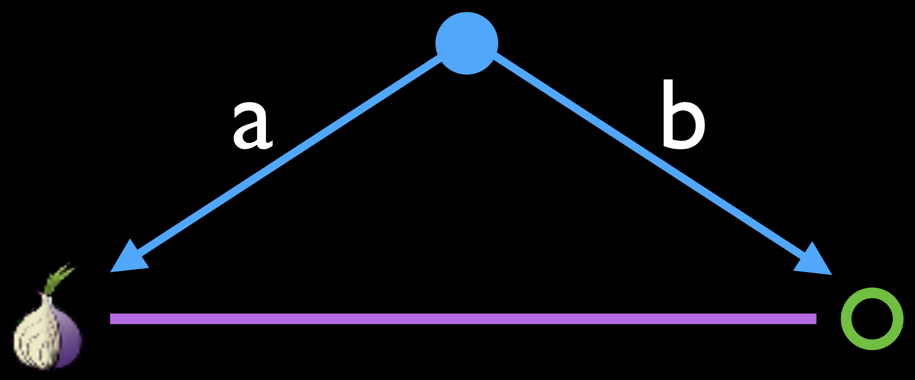

DeTor: never-twice avoidance

Where could packets possibly reach

Upper bound RTT $\geq$ 2 (a+b) / c

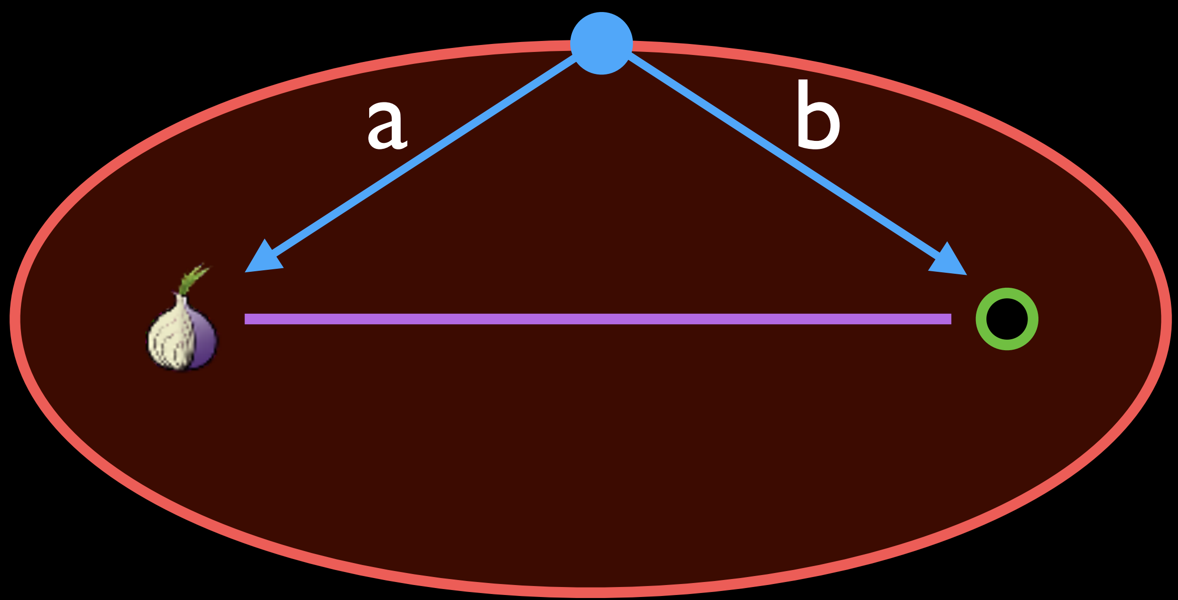

The packet could possibly reach any point in the ellipse

DeTor: never-twice avoidance

Where could packets possibly reach

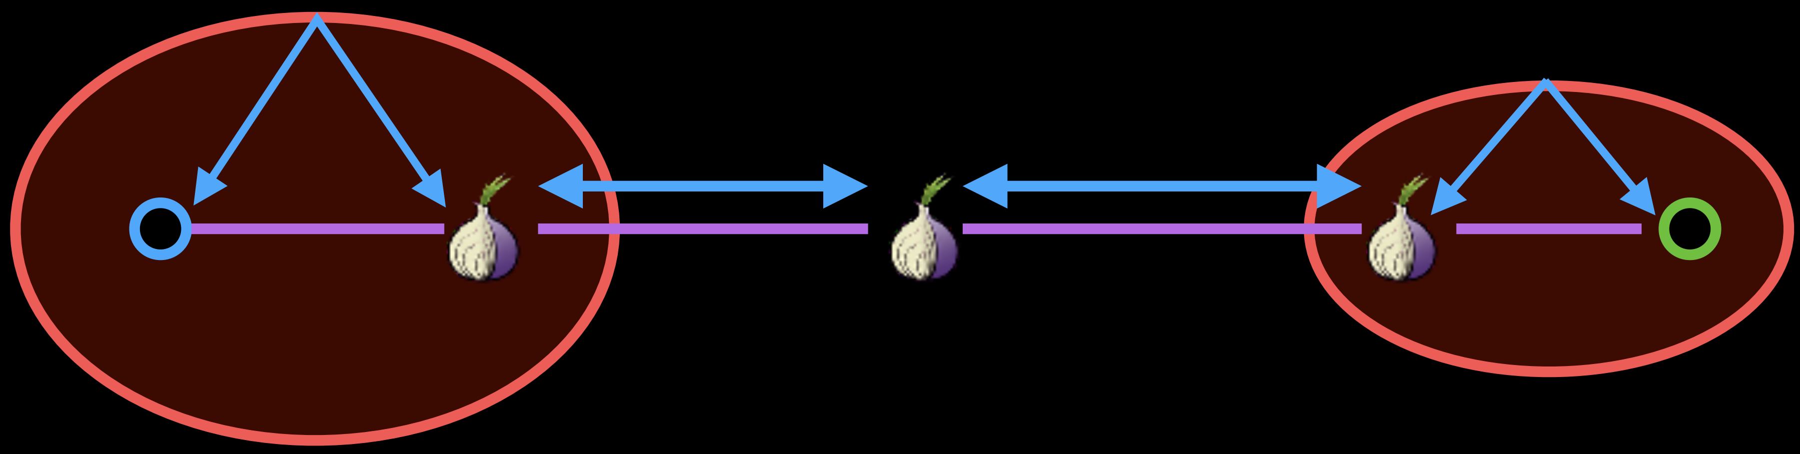

DeTor: never-twice avoidance

Where could packets possibly reach

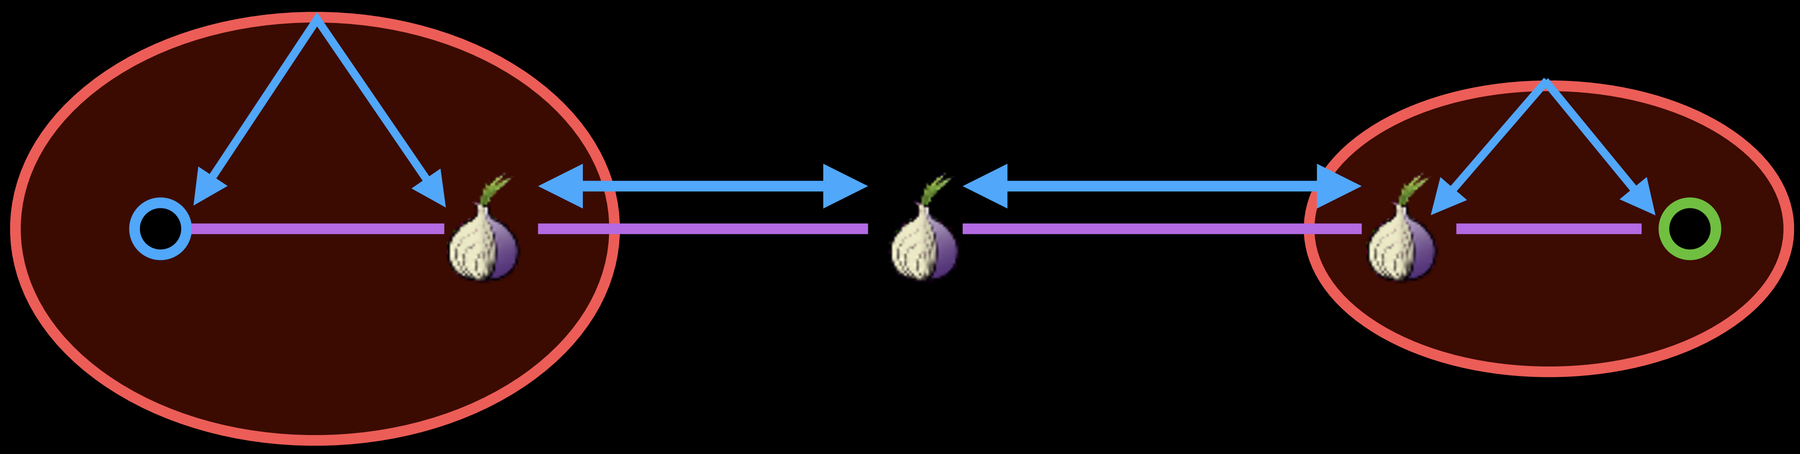

Compute the worst-case scenarios
for *both* entry and exit legs, separately

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

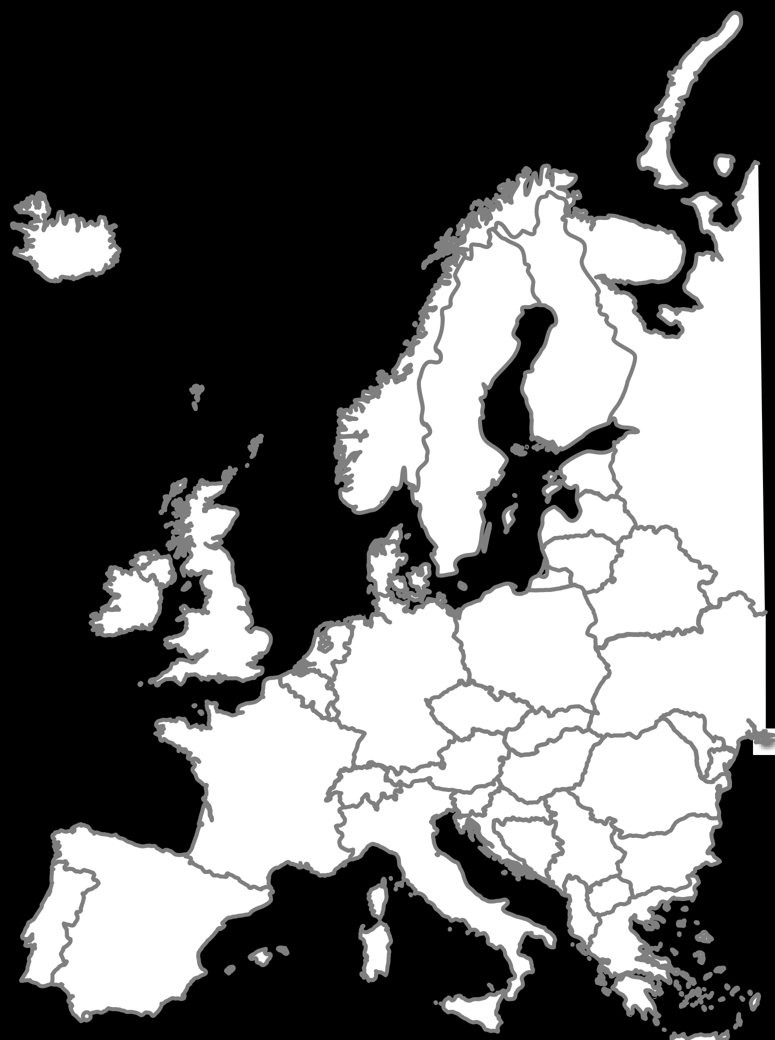

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

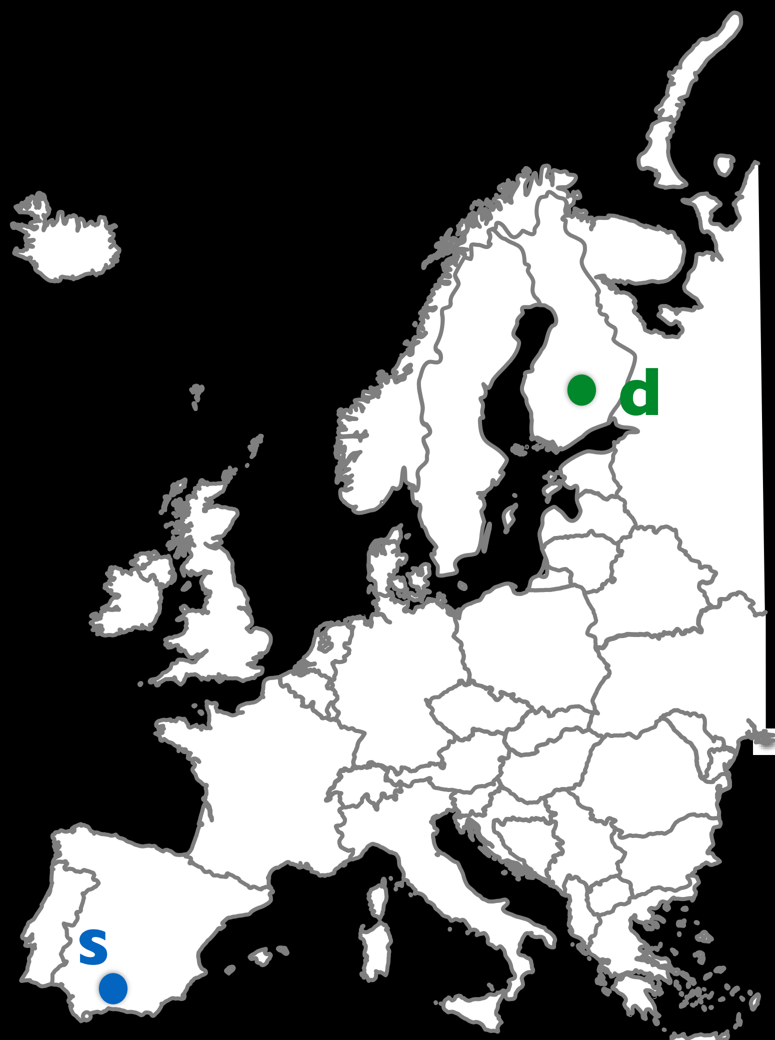

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

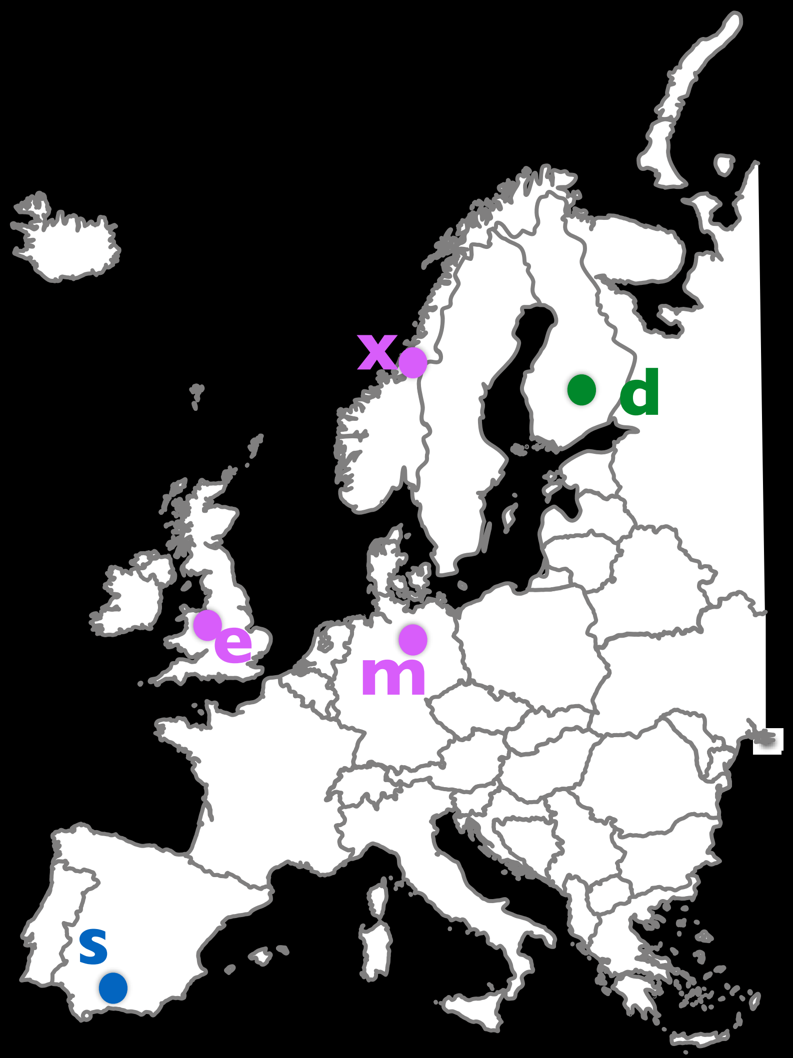

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

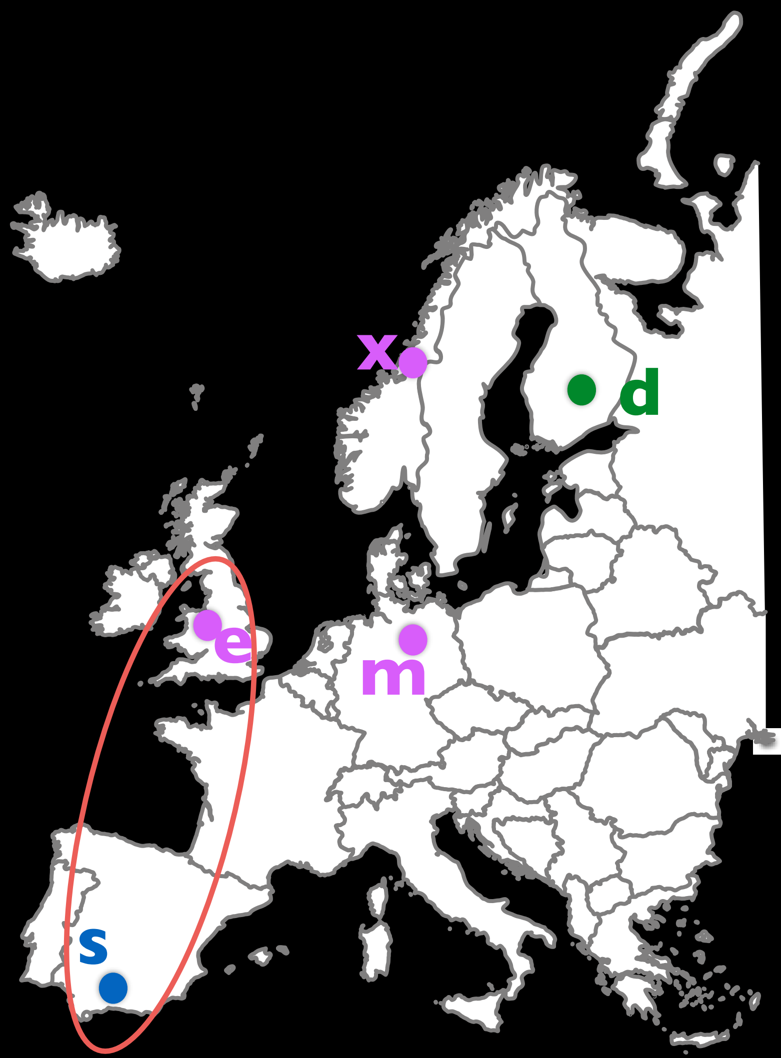

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

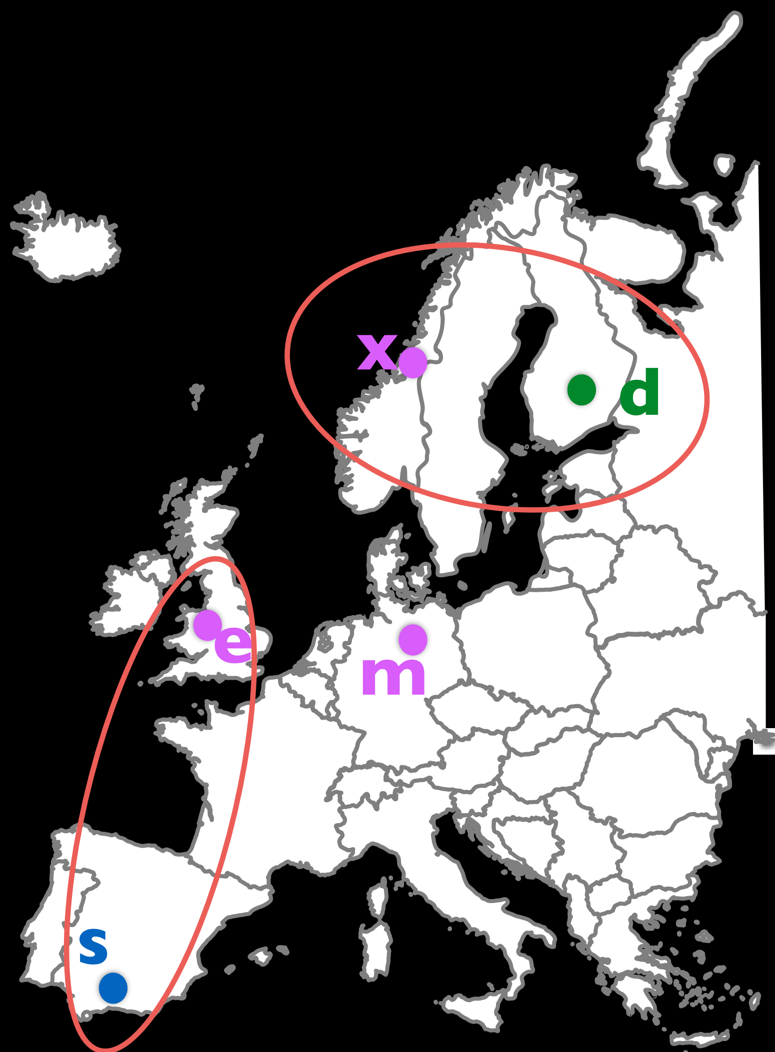

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

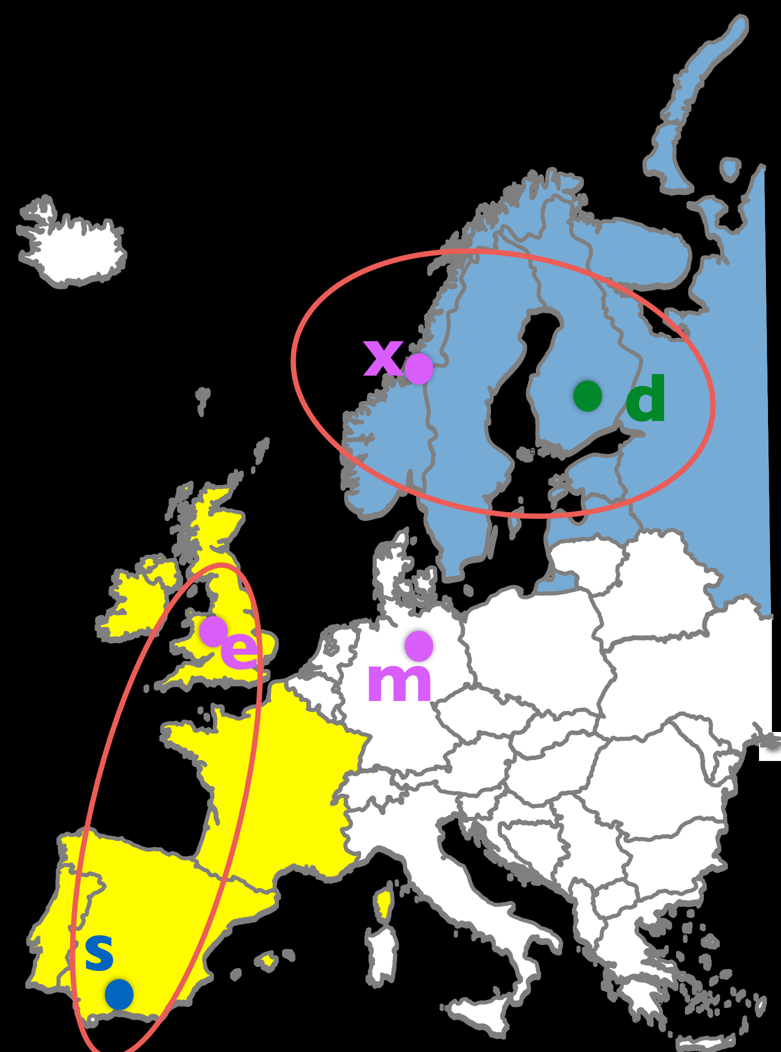

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

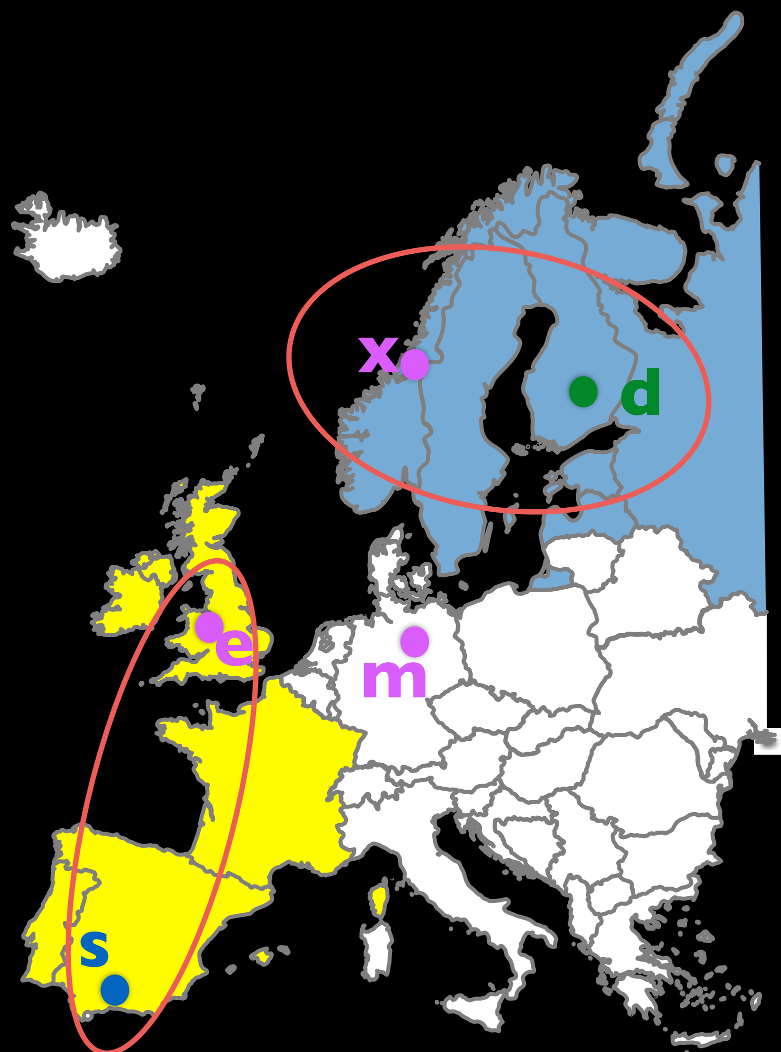

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

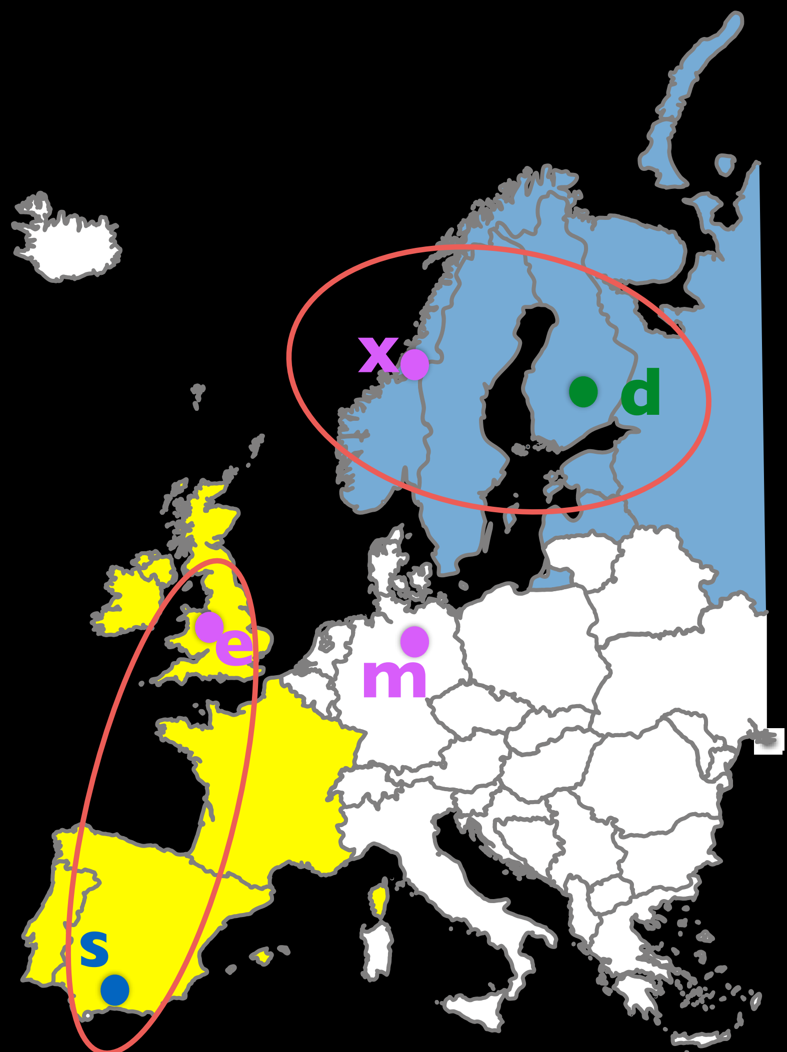

no country
intersects
with both ellipses

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

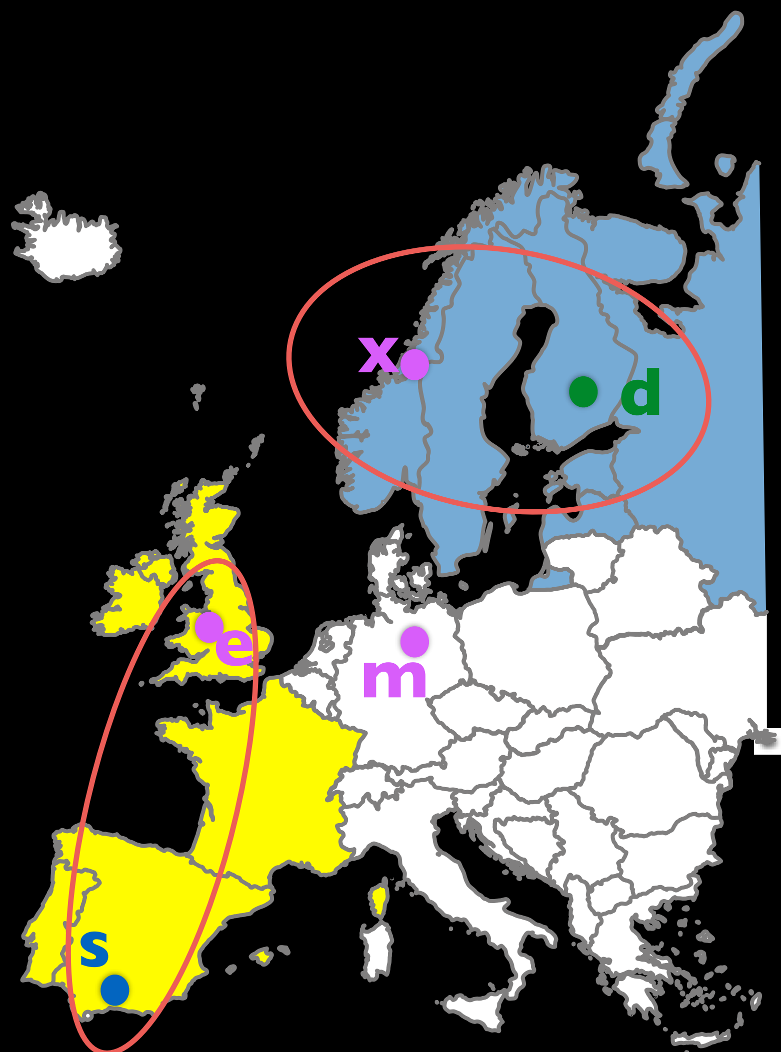

no country
intersects
with both ellipses
⇓

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

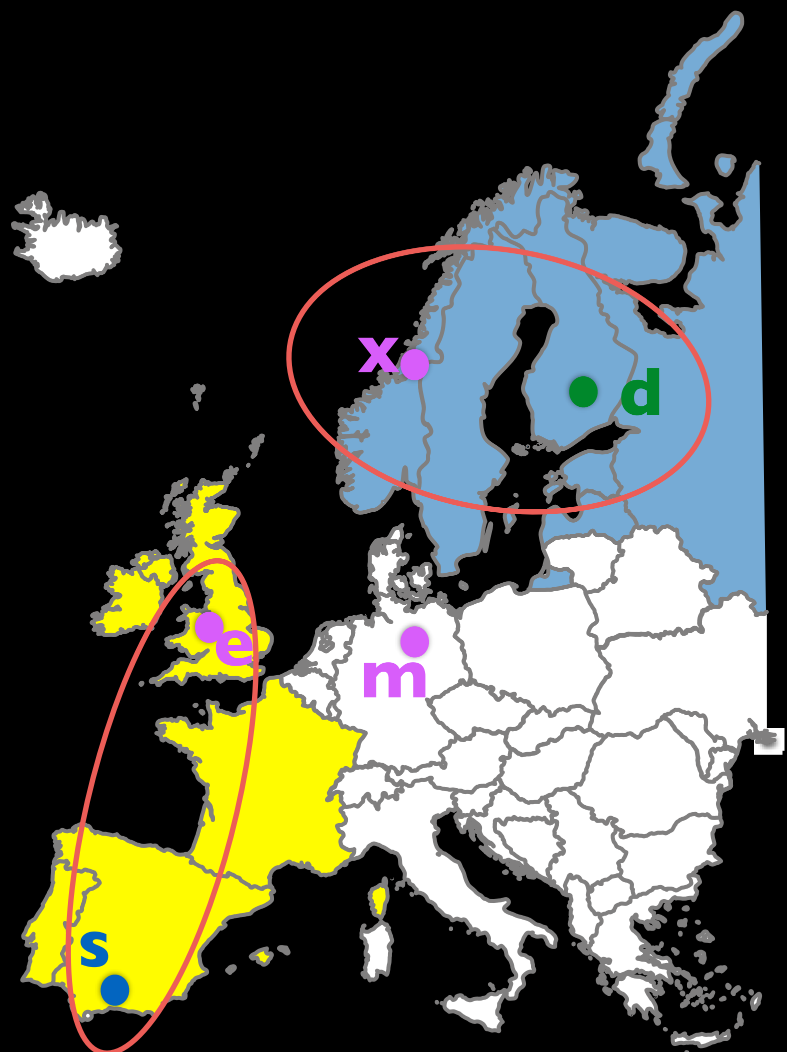

no country
intersects
with both ellipses
⇓
packet over entry/exit legs
could not
have traversed the same
country

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

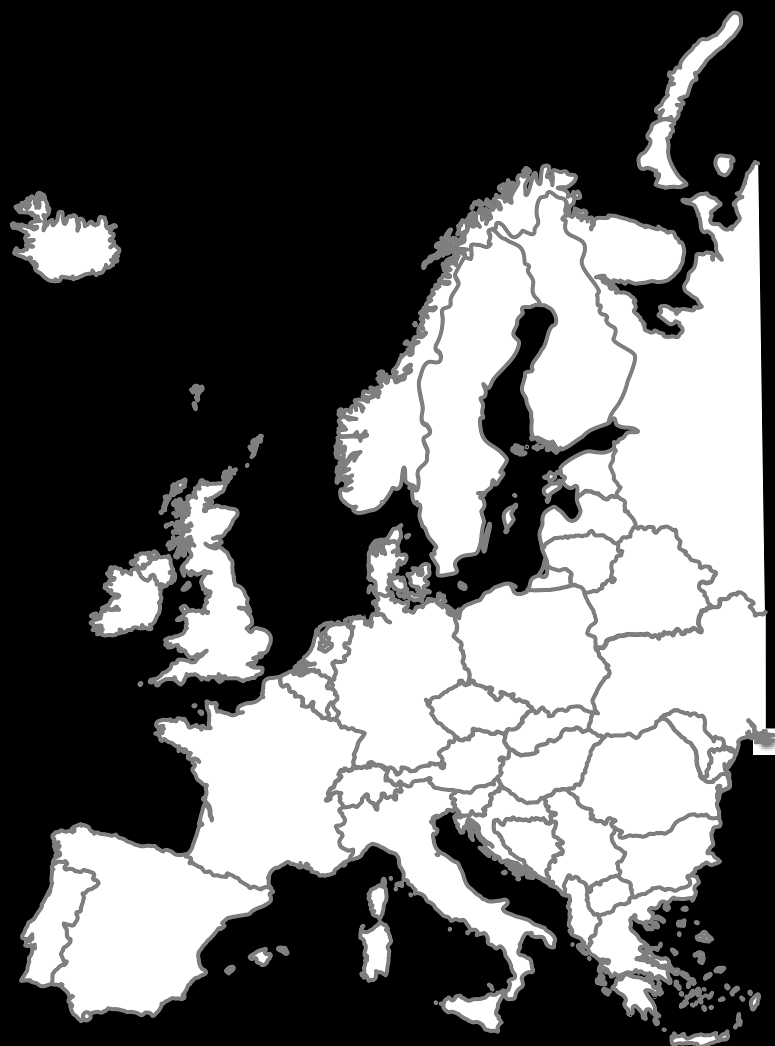

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

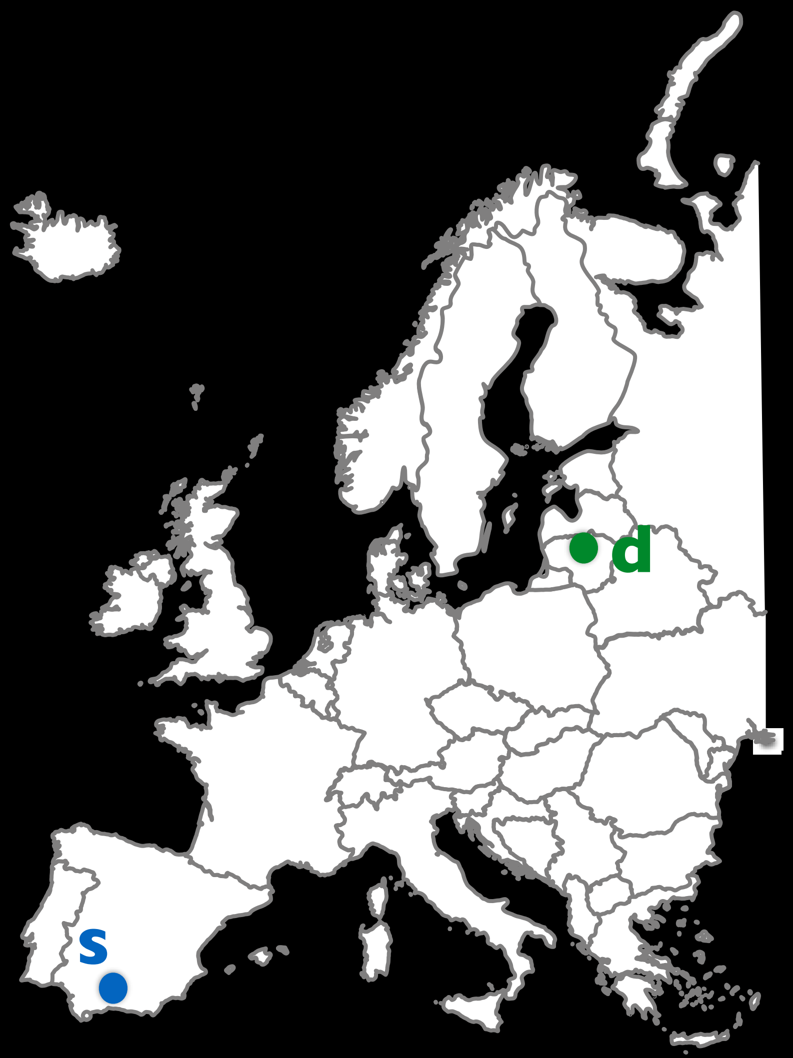

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

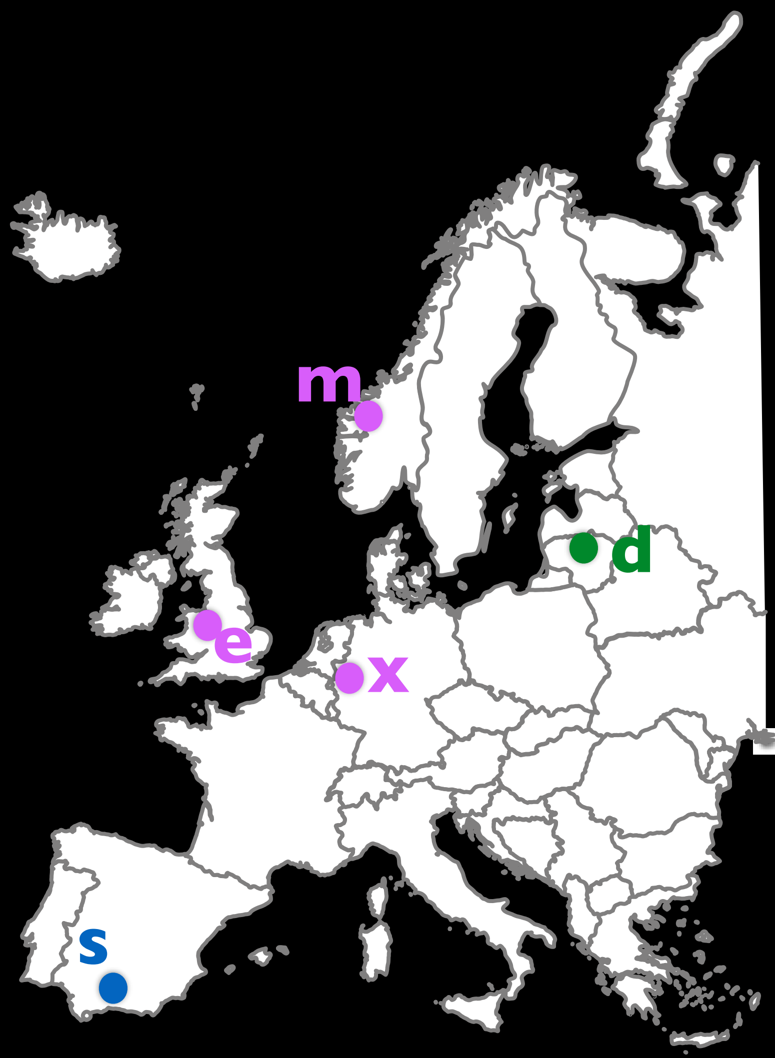

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

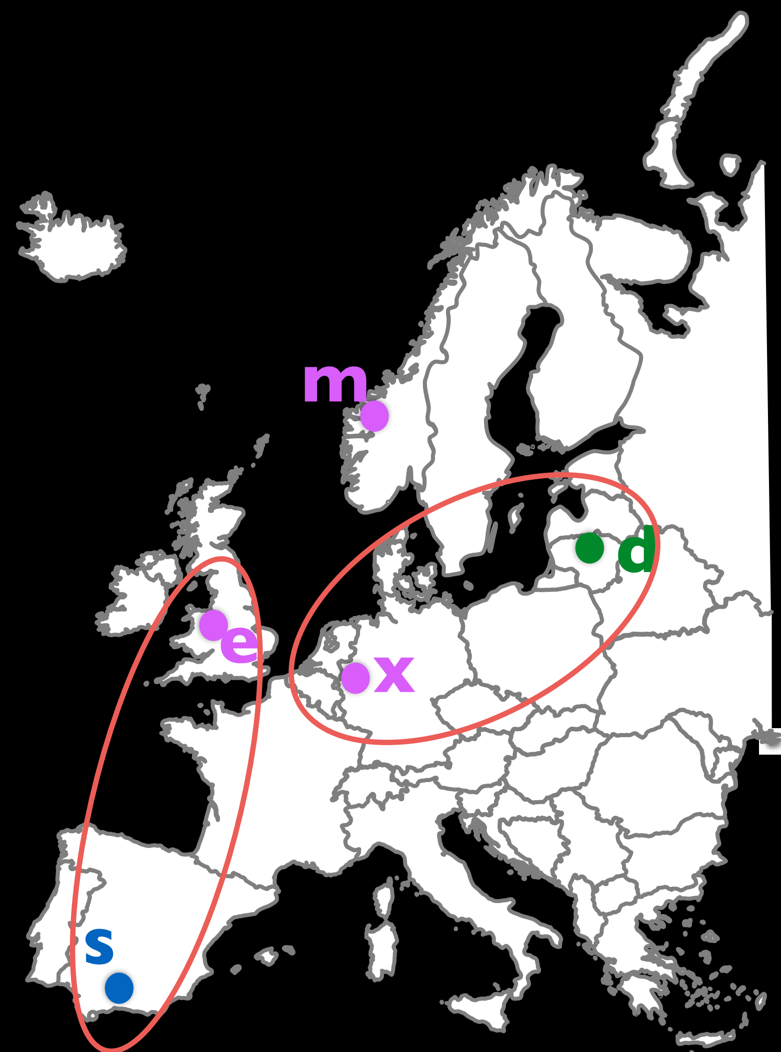

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

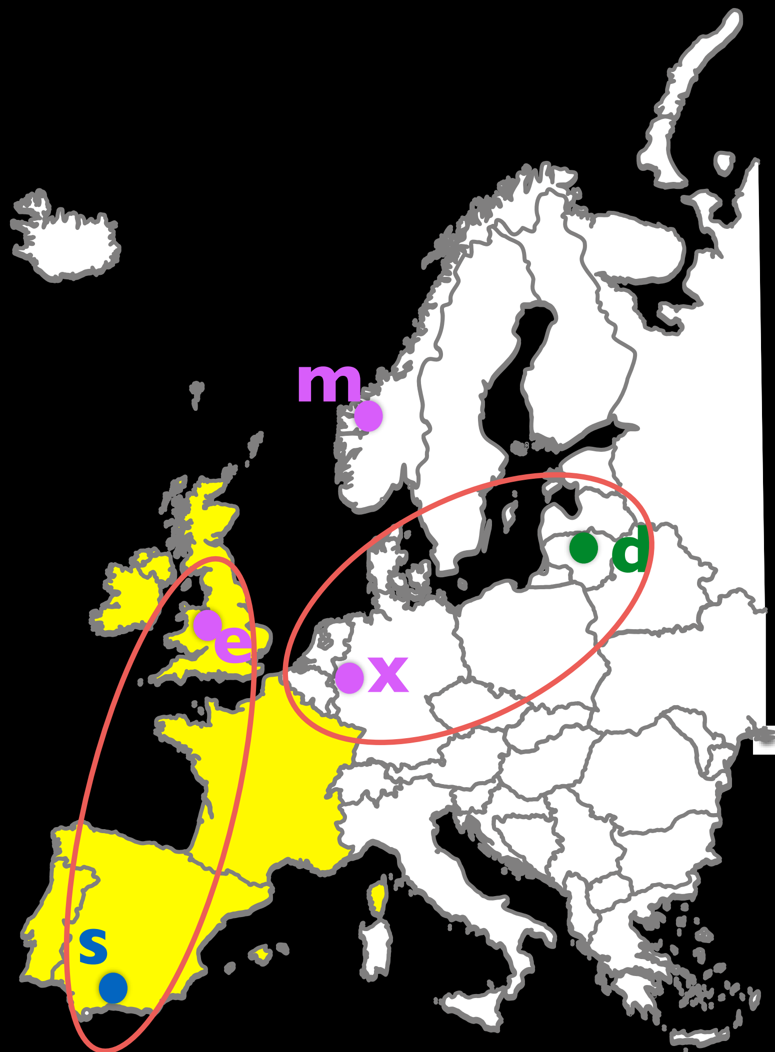

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

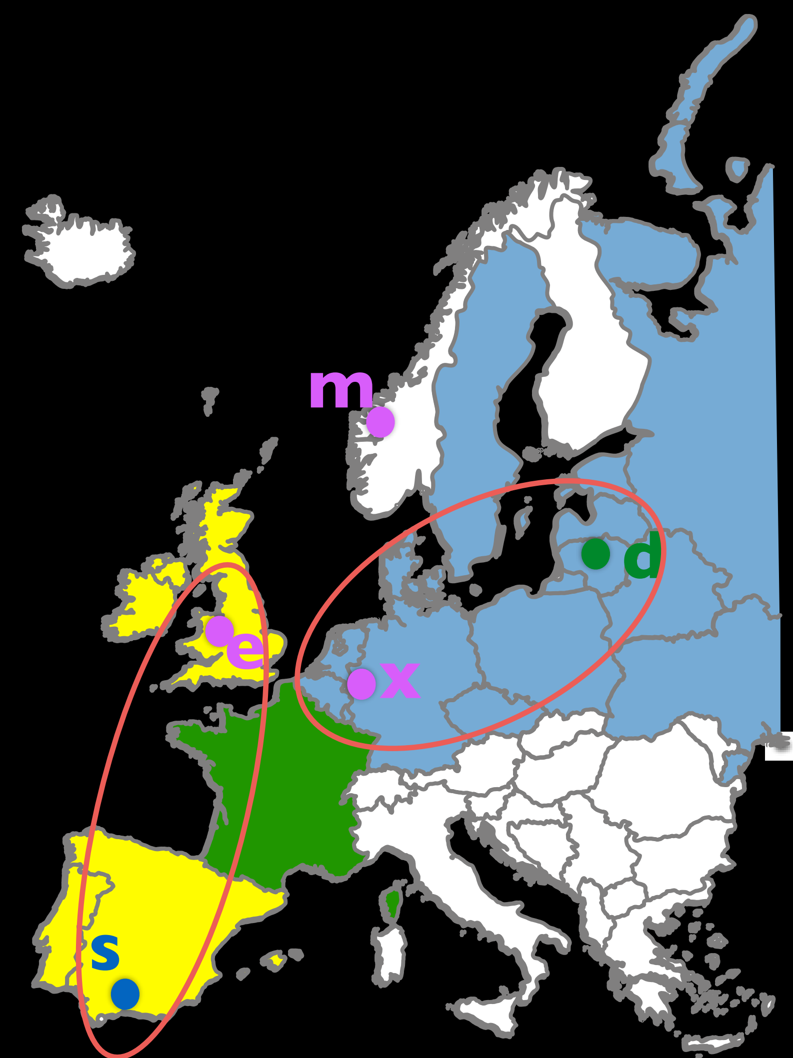

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

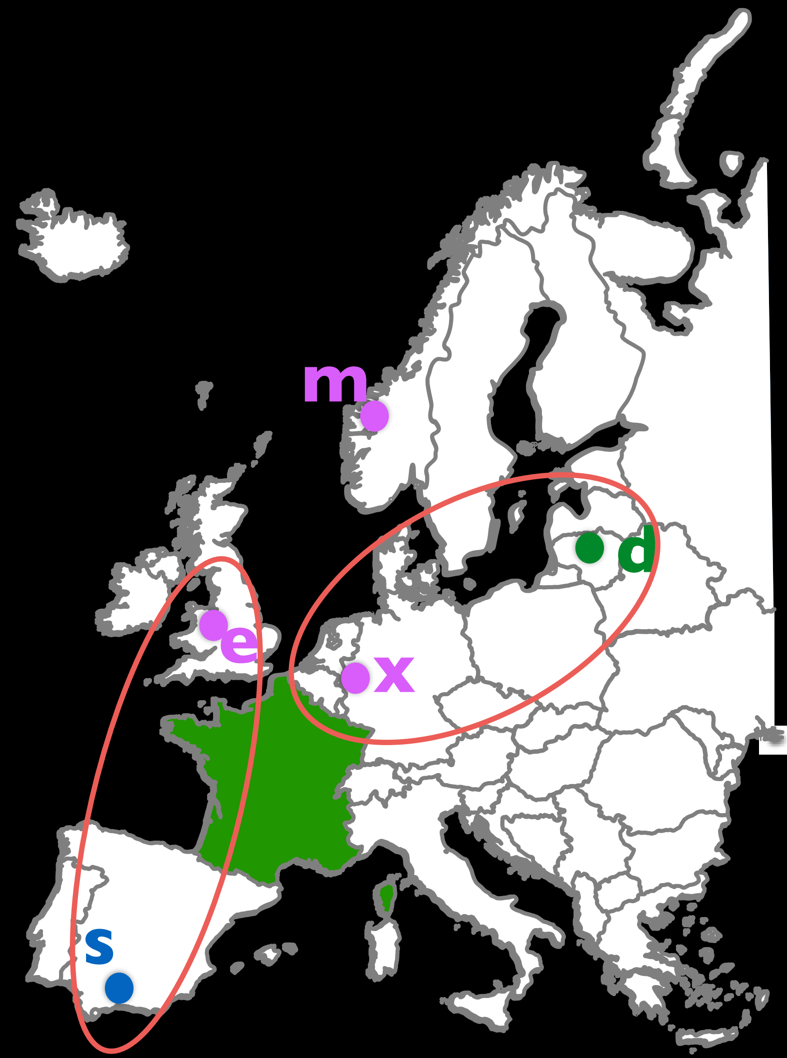

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

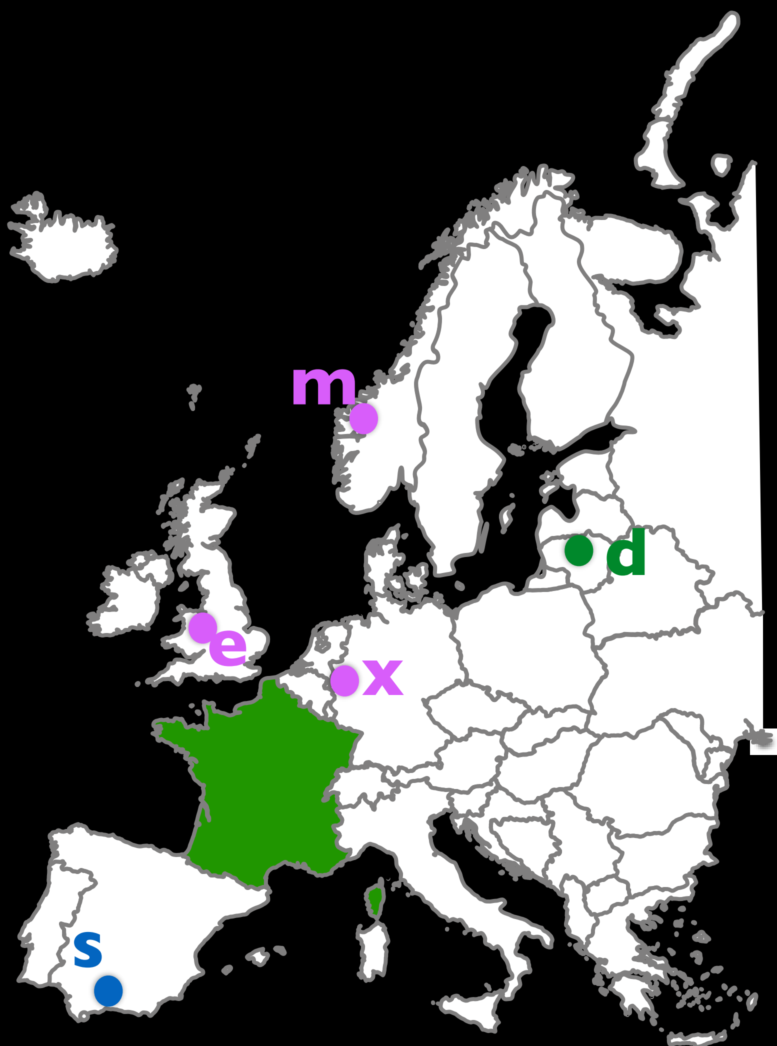

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

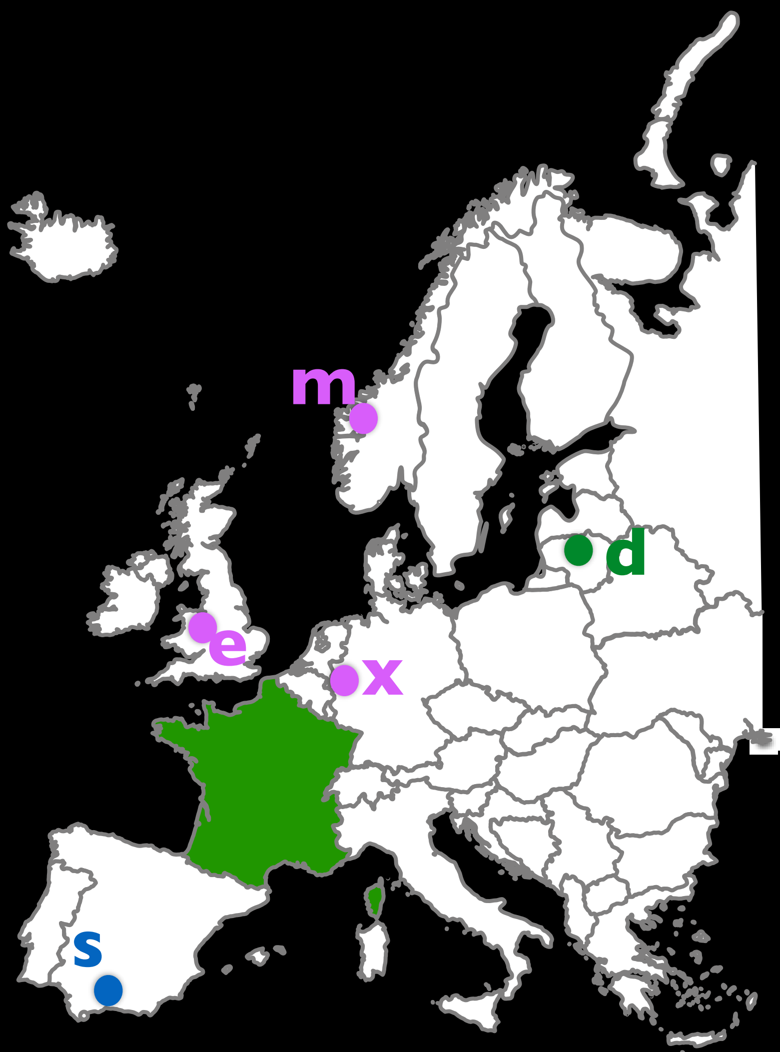

For each country intersects with both ellipses

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

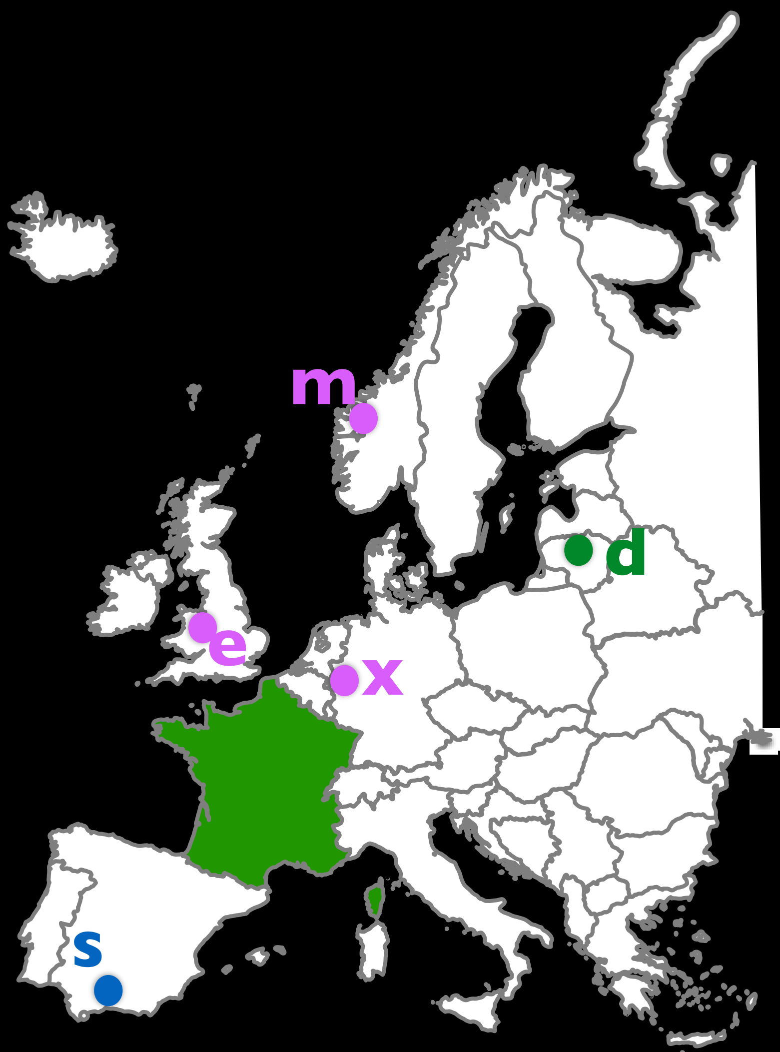

For each country intersects with both ellipses

The *shortest possible* RTT thru Tor and entry & exit legs traverse

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

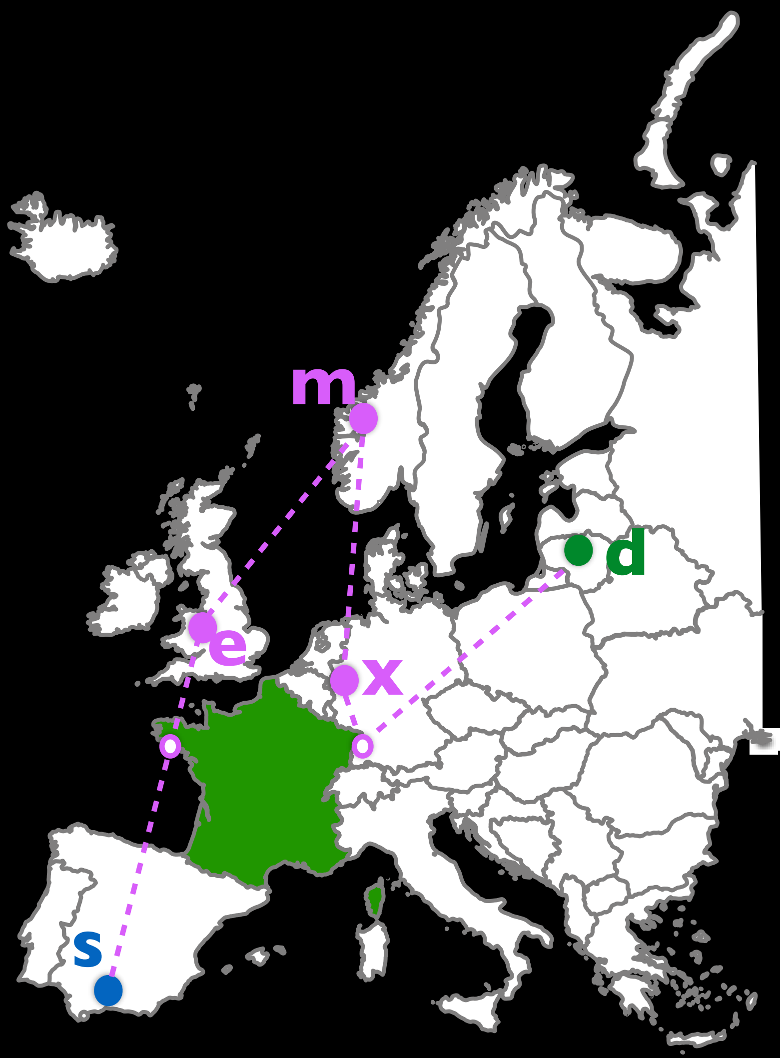

For each country intersects with both ellipses

The shortest *possible* RTT thru Tor and entry & exit legs traverse

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

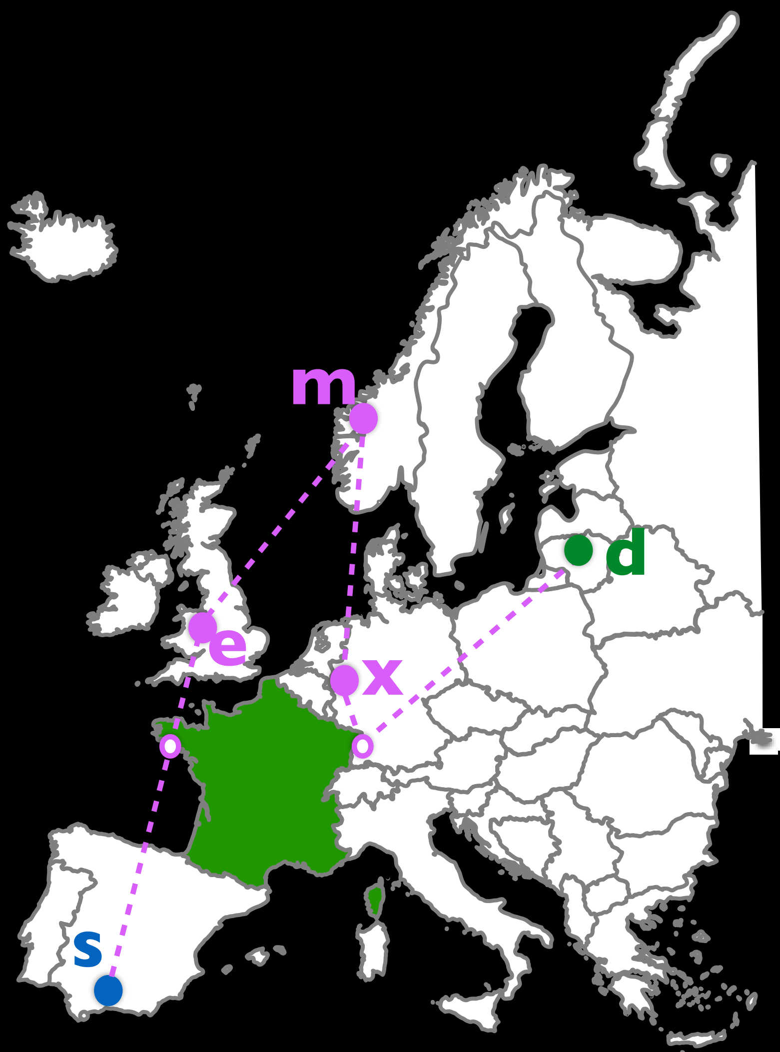

For each country intersects with both ellipses

Measured RTT ≪ The *shortest possible* RTT thru Tor and entry & exit legs traverse

DeTor: never-twice avoidance

Which countries can entry & exit legs reach

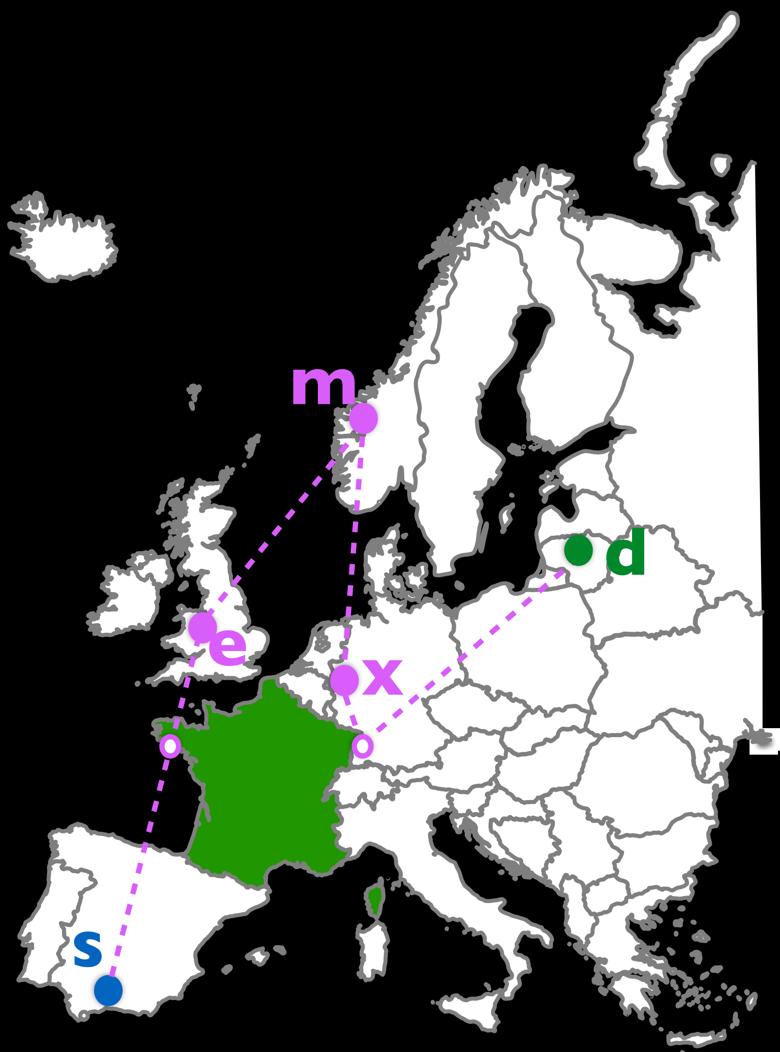

For each country intersects with both ellipses

Measured RTT ≪ The shortest *possible* RTT thru Tor and entry & exit legs traverse

⇒ The packet could not have traversed over entry & exit legs

Evaluation

Through simulation

Evaluation
Through simulation

- 50 random real Tor nodes
  - with GPS locations and pair-wise RTTs using Ting
  - choose sources and destinations among these nodes

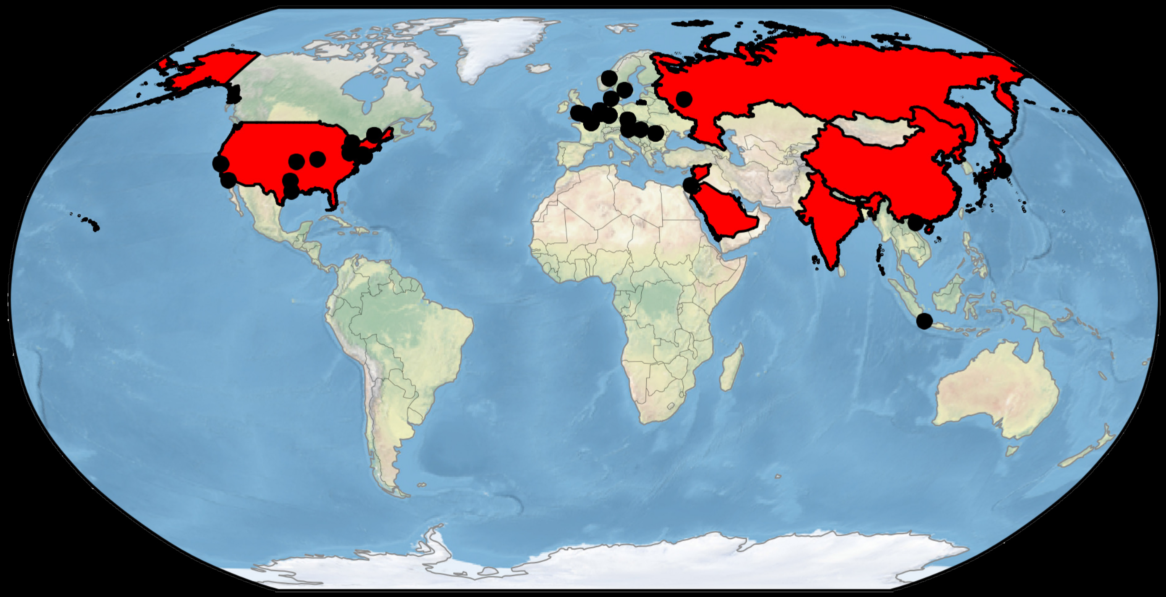

Evaluation
Through simulation

- 50 random real Tor nodes
  - with GPS locations and pair-wise RTTs using Ting
  - choose sources and destinations among these nodes

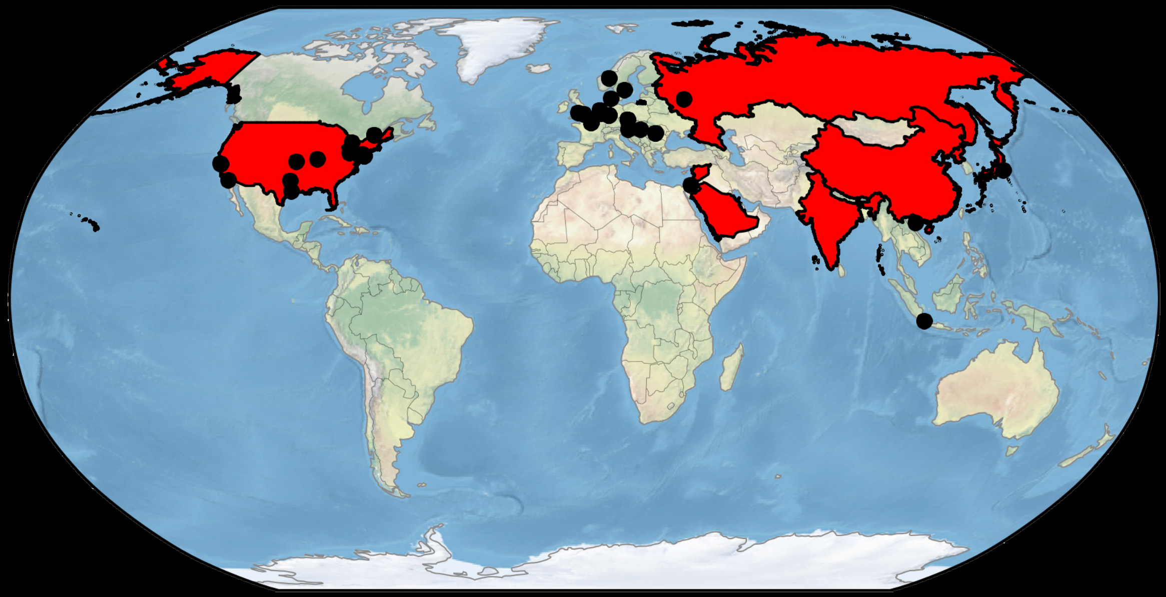

Evaluation

- How successful is DeTor?
- How well do DeTor circuits perform?

Never-once success rate

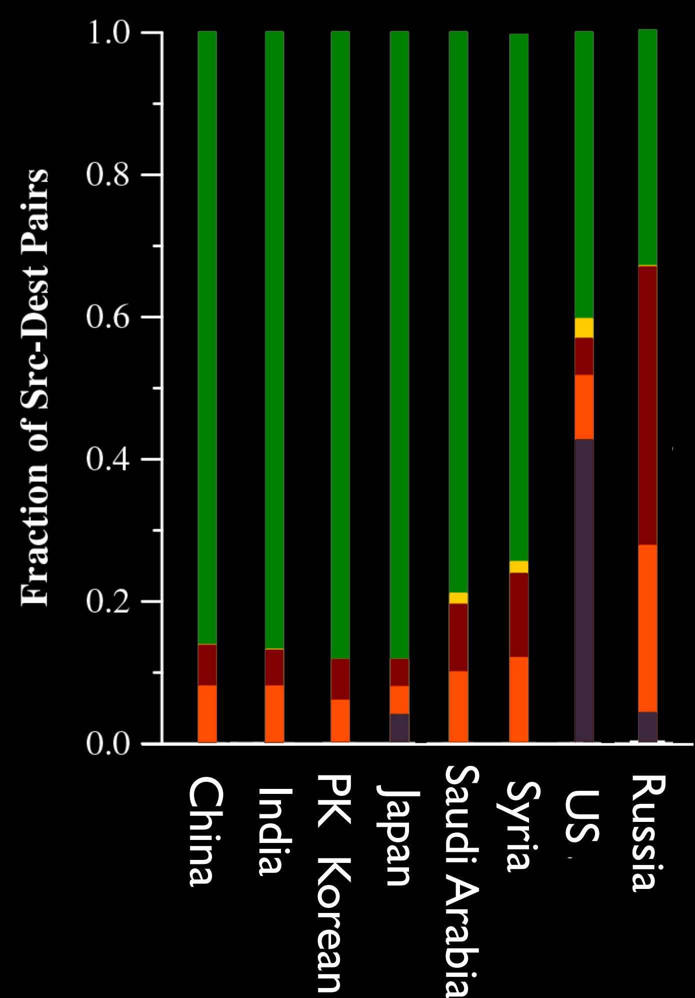

Successful with DeTor

Theoretically avoid, but failed with real RTTs

No circuits could provably avoid

No trusted Tor nodes

Source/Destination in Forbidden region

Never-once success rate

Most src-dst pairs can successful find never-once circuits

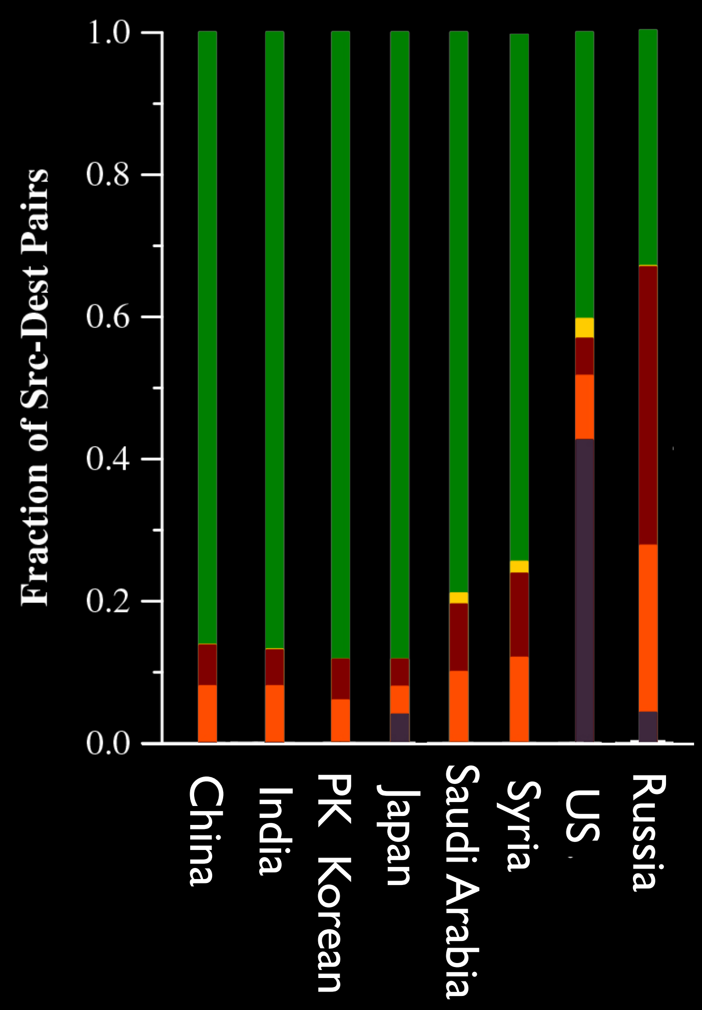

Successful with DeTor

Theoretically avoid, but failed with real RTTs

No circuits could provably avoid

No trusted Tor nodes

Source/Destination in Forbidden region

Never-once success rate

Failure typically arises when users are in or close to the regions to avoid

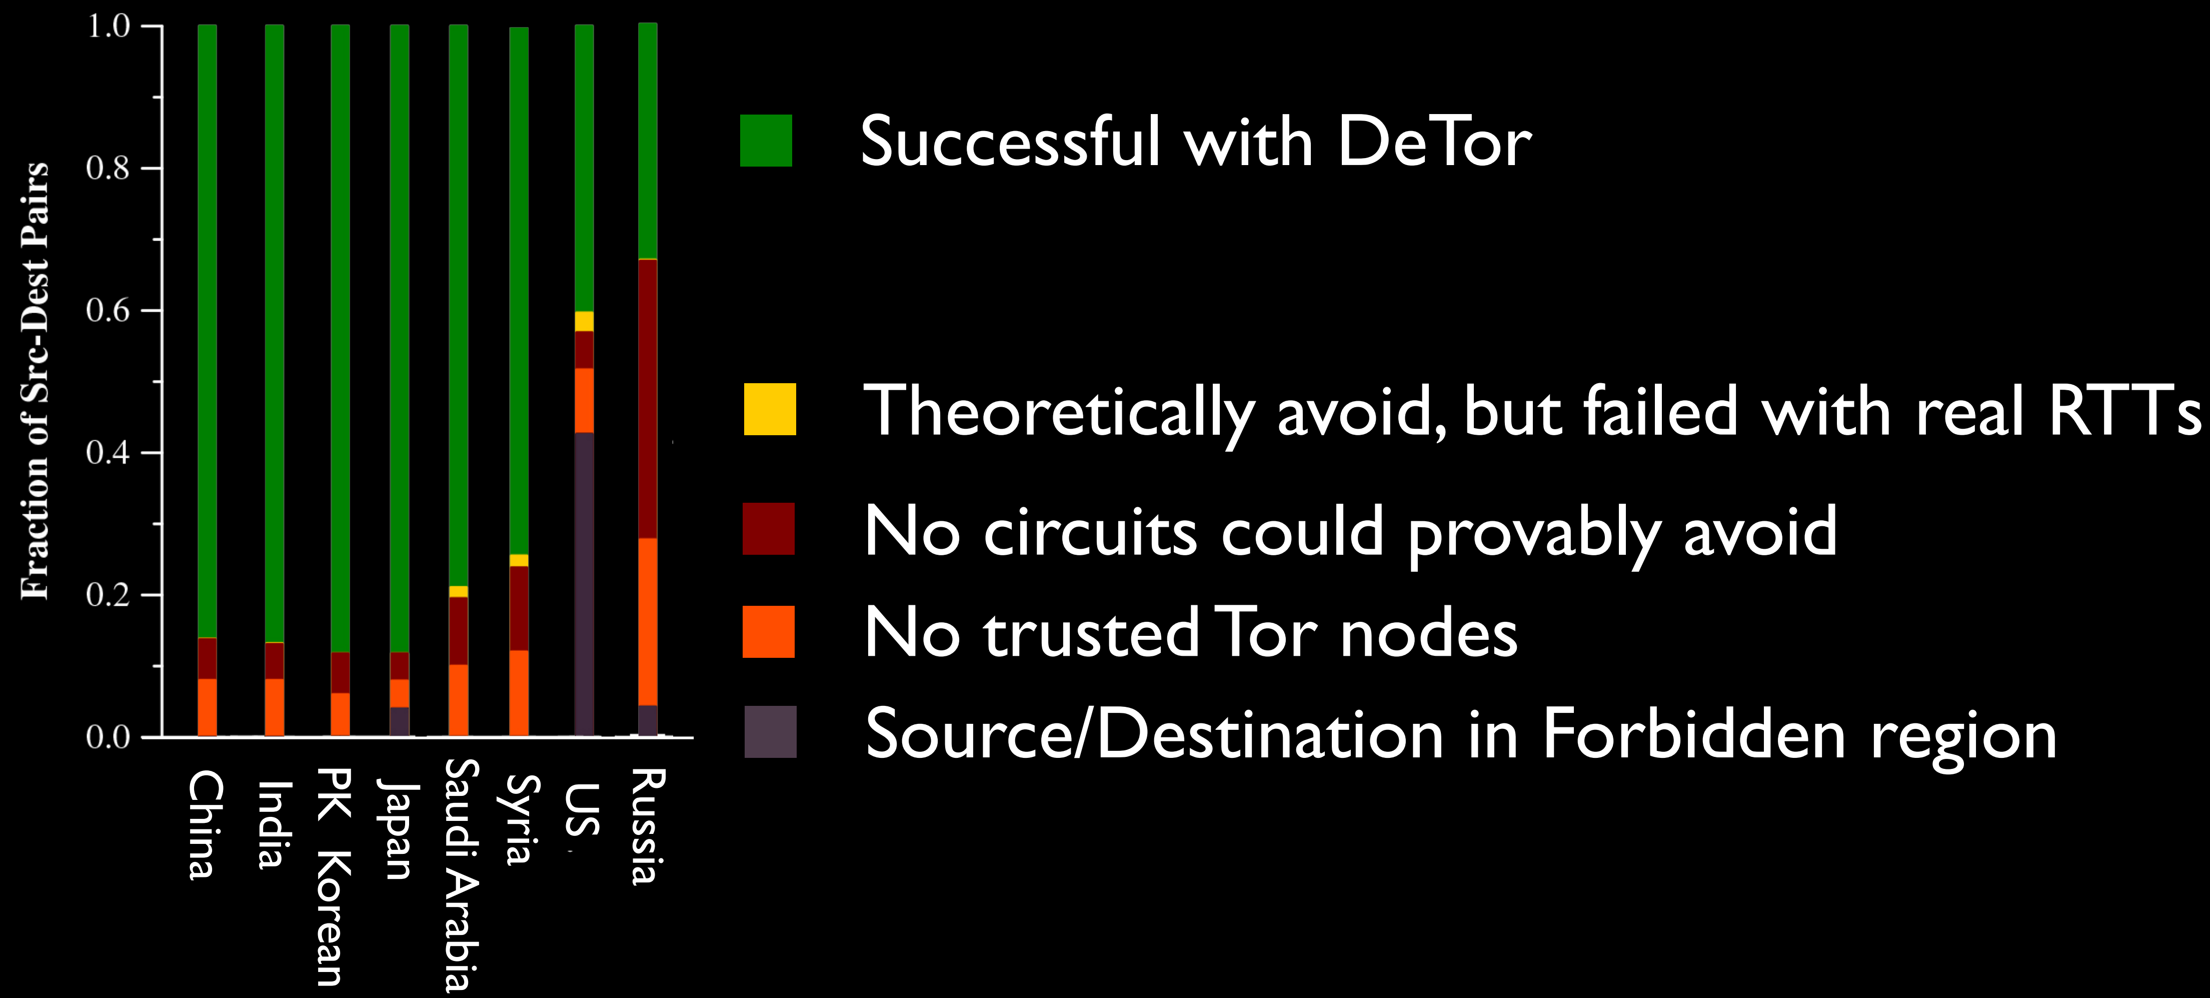

Number of never-once circuits

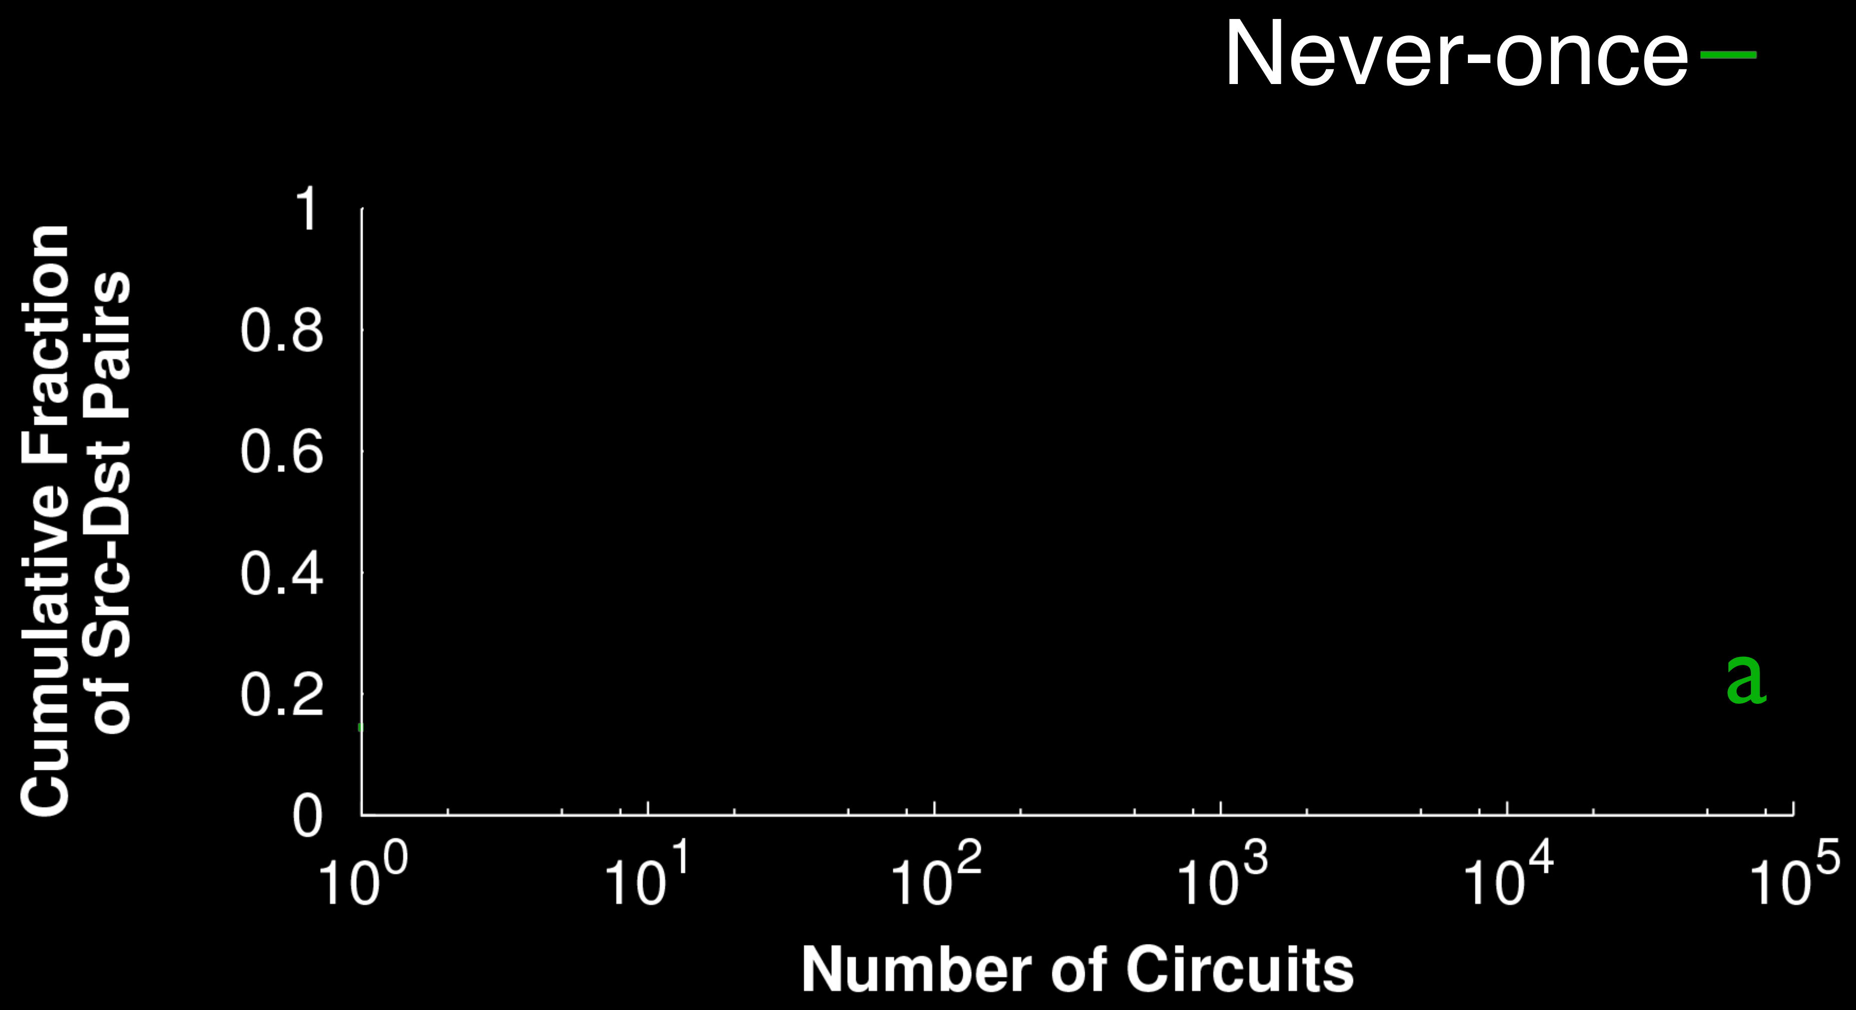

Number of never-once circuits

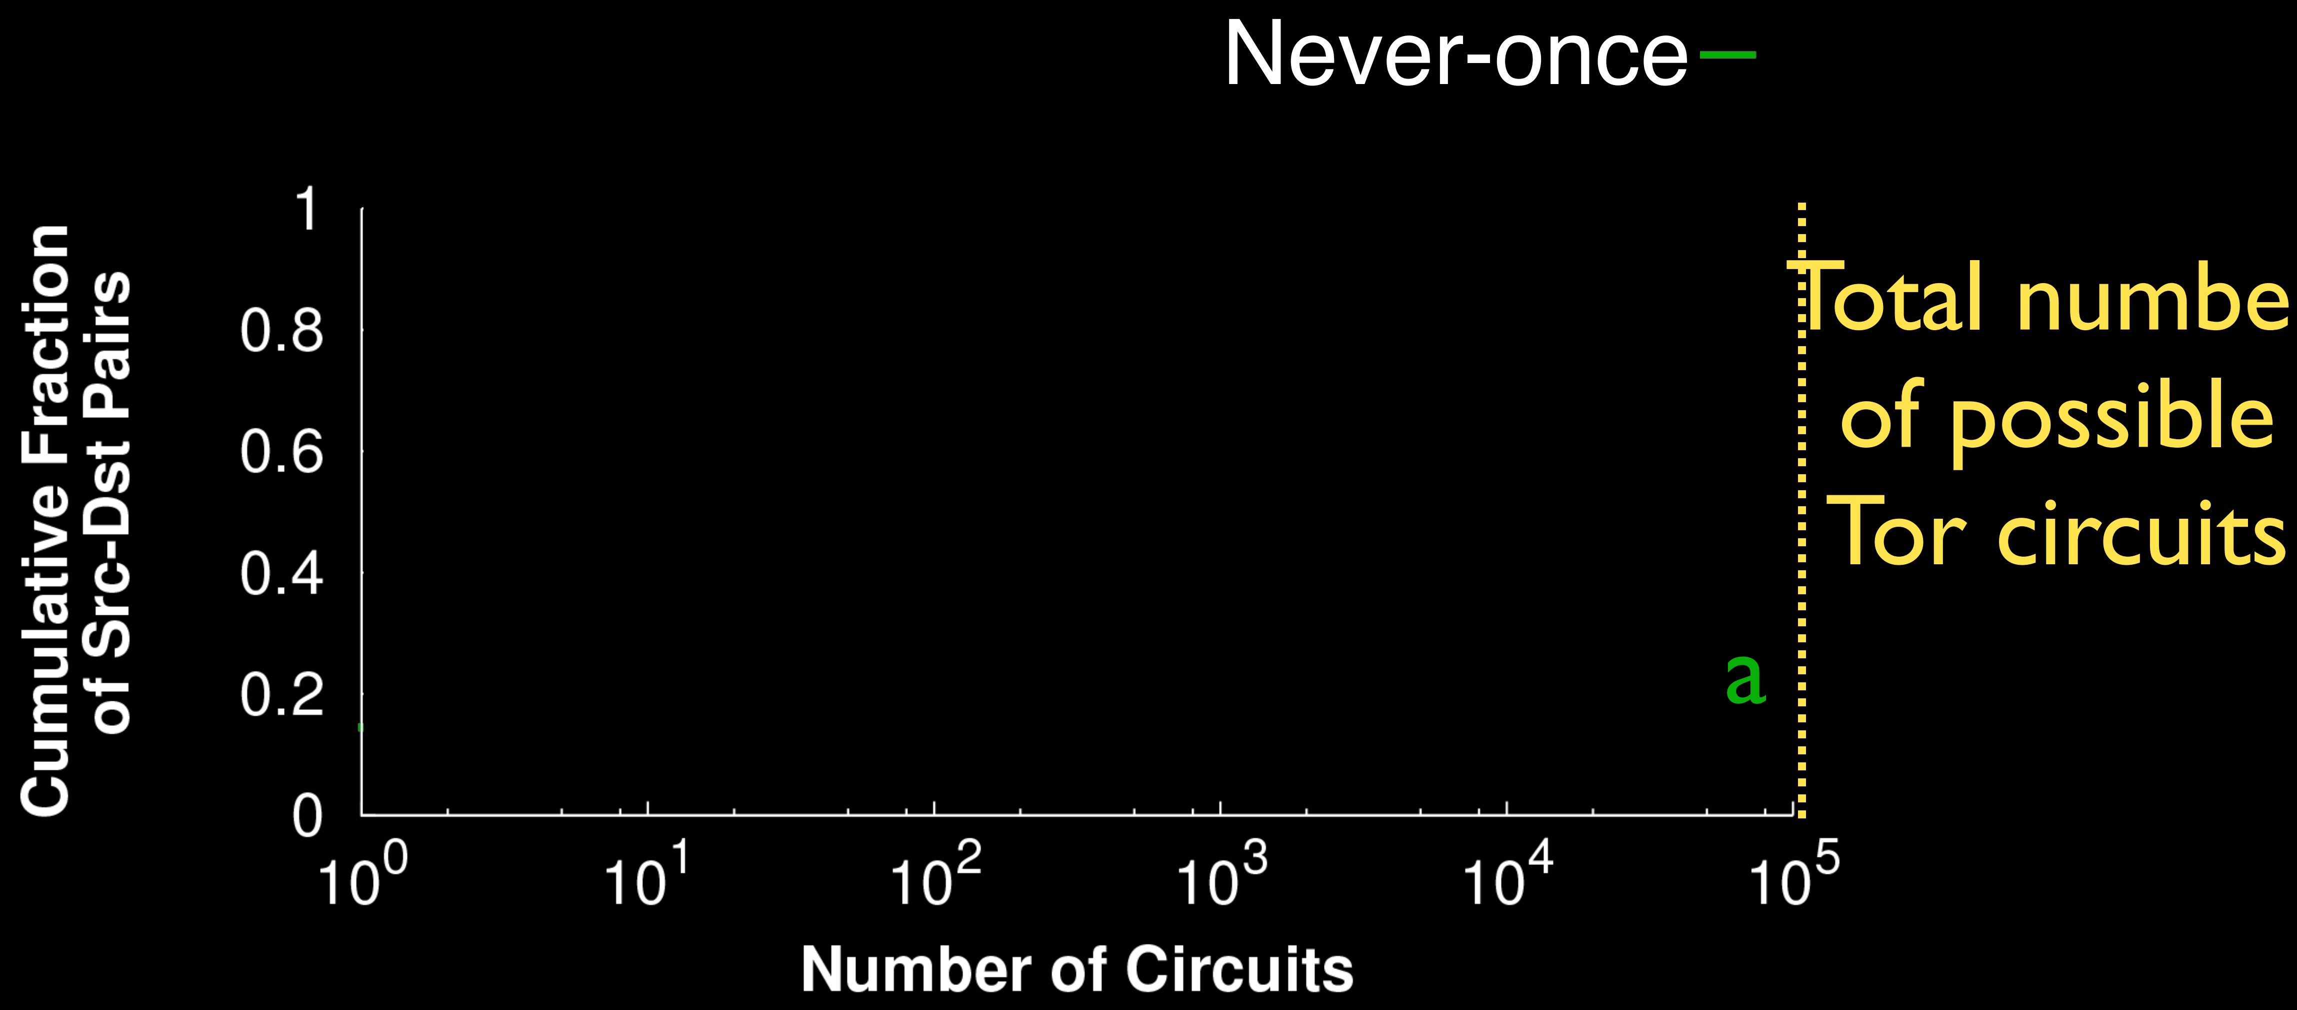

Number of never-once circuits

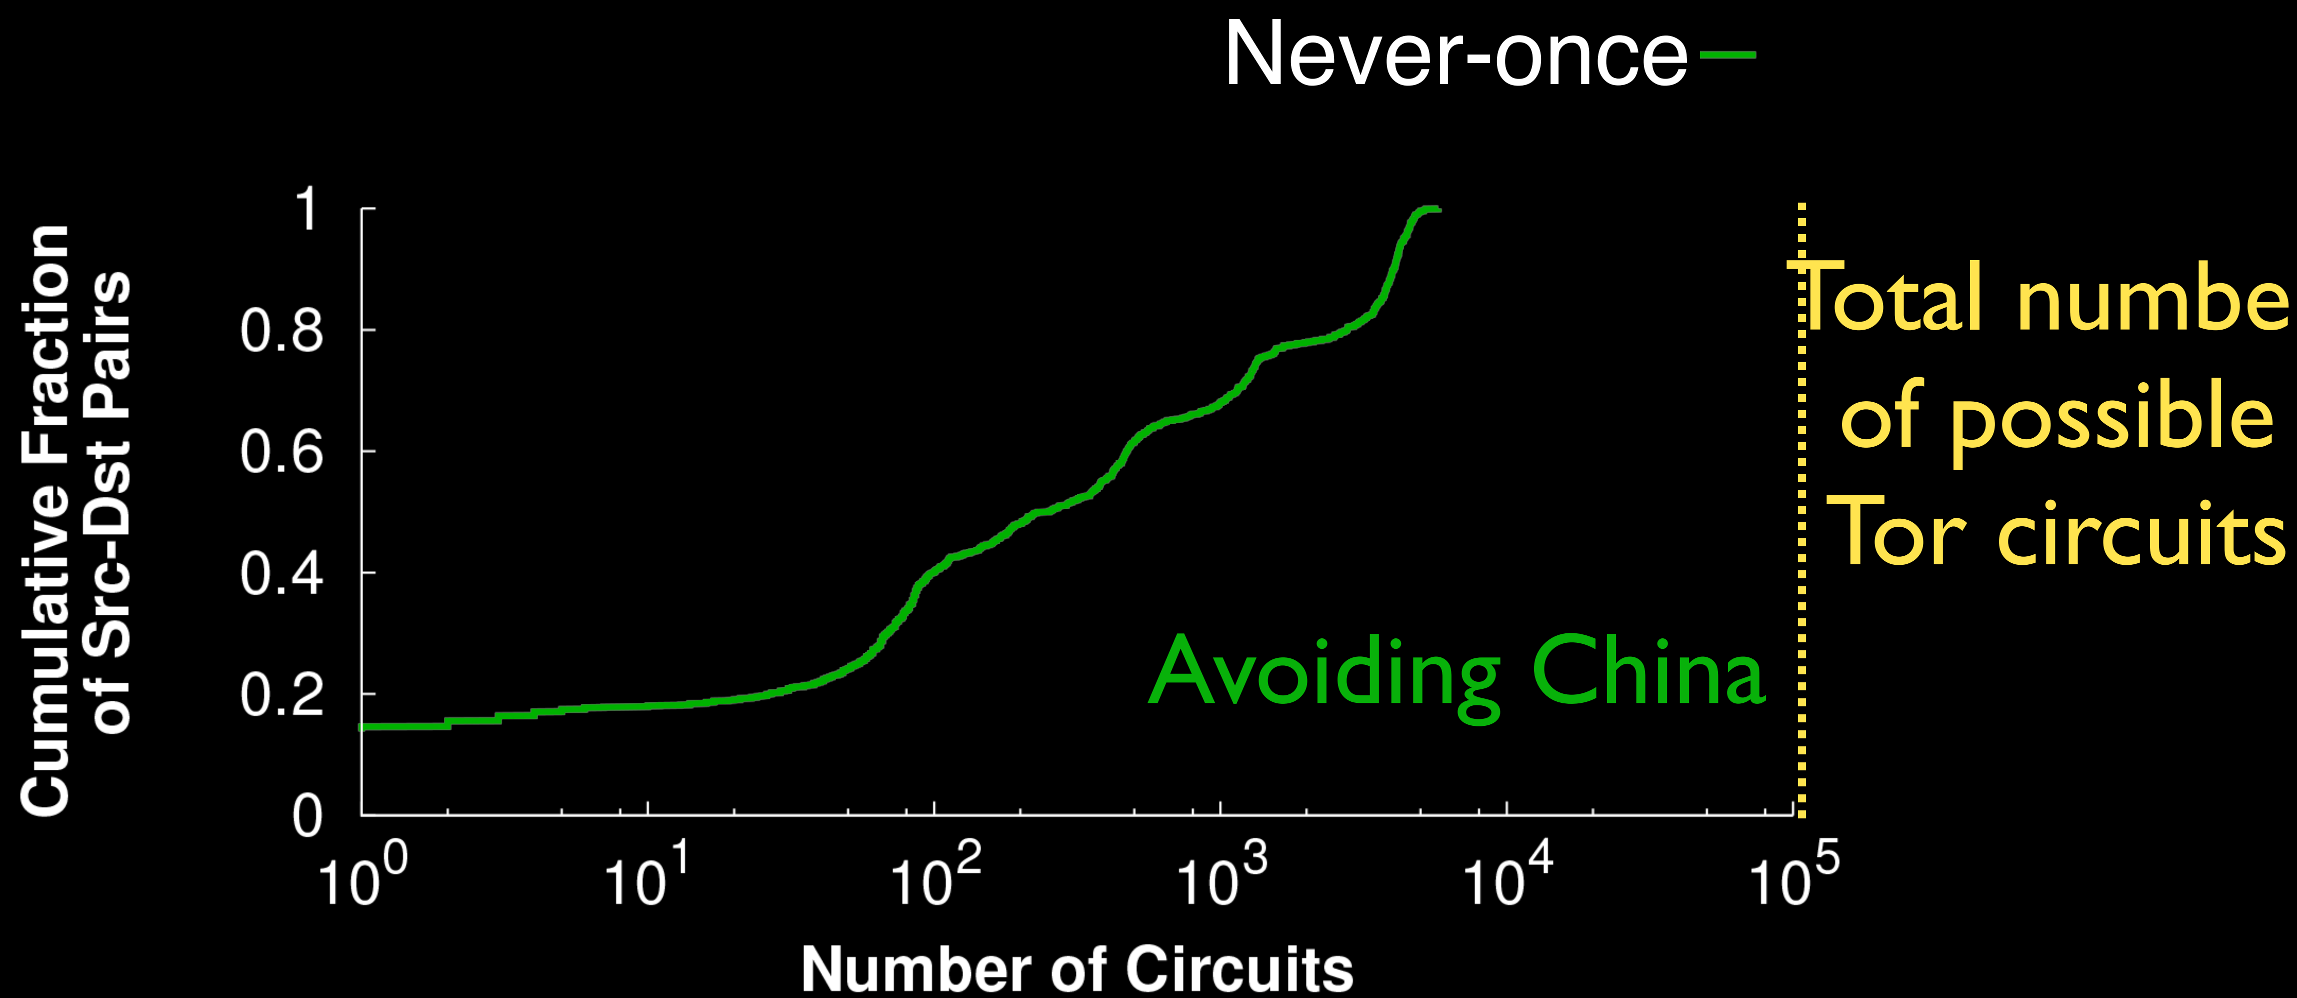

Number of never-once circuits

Half of src-dst pairs have over 500 never-once circuits

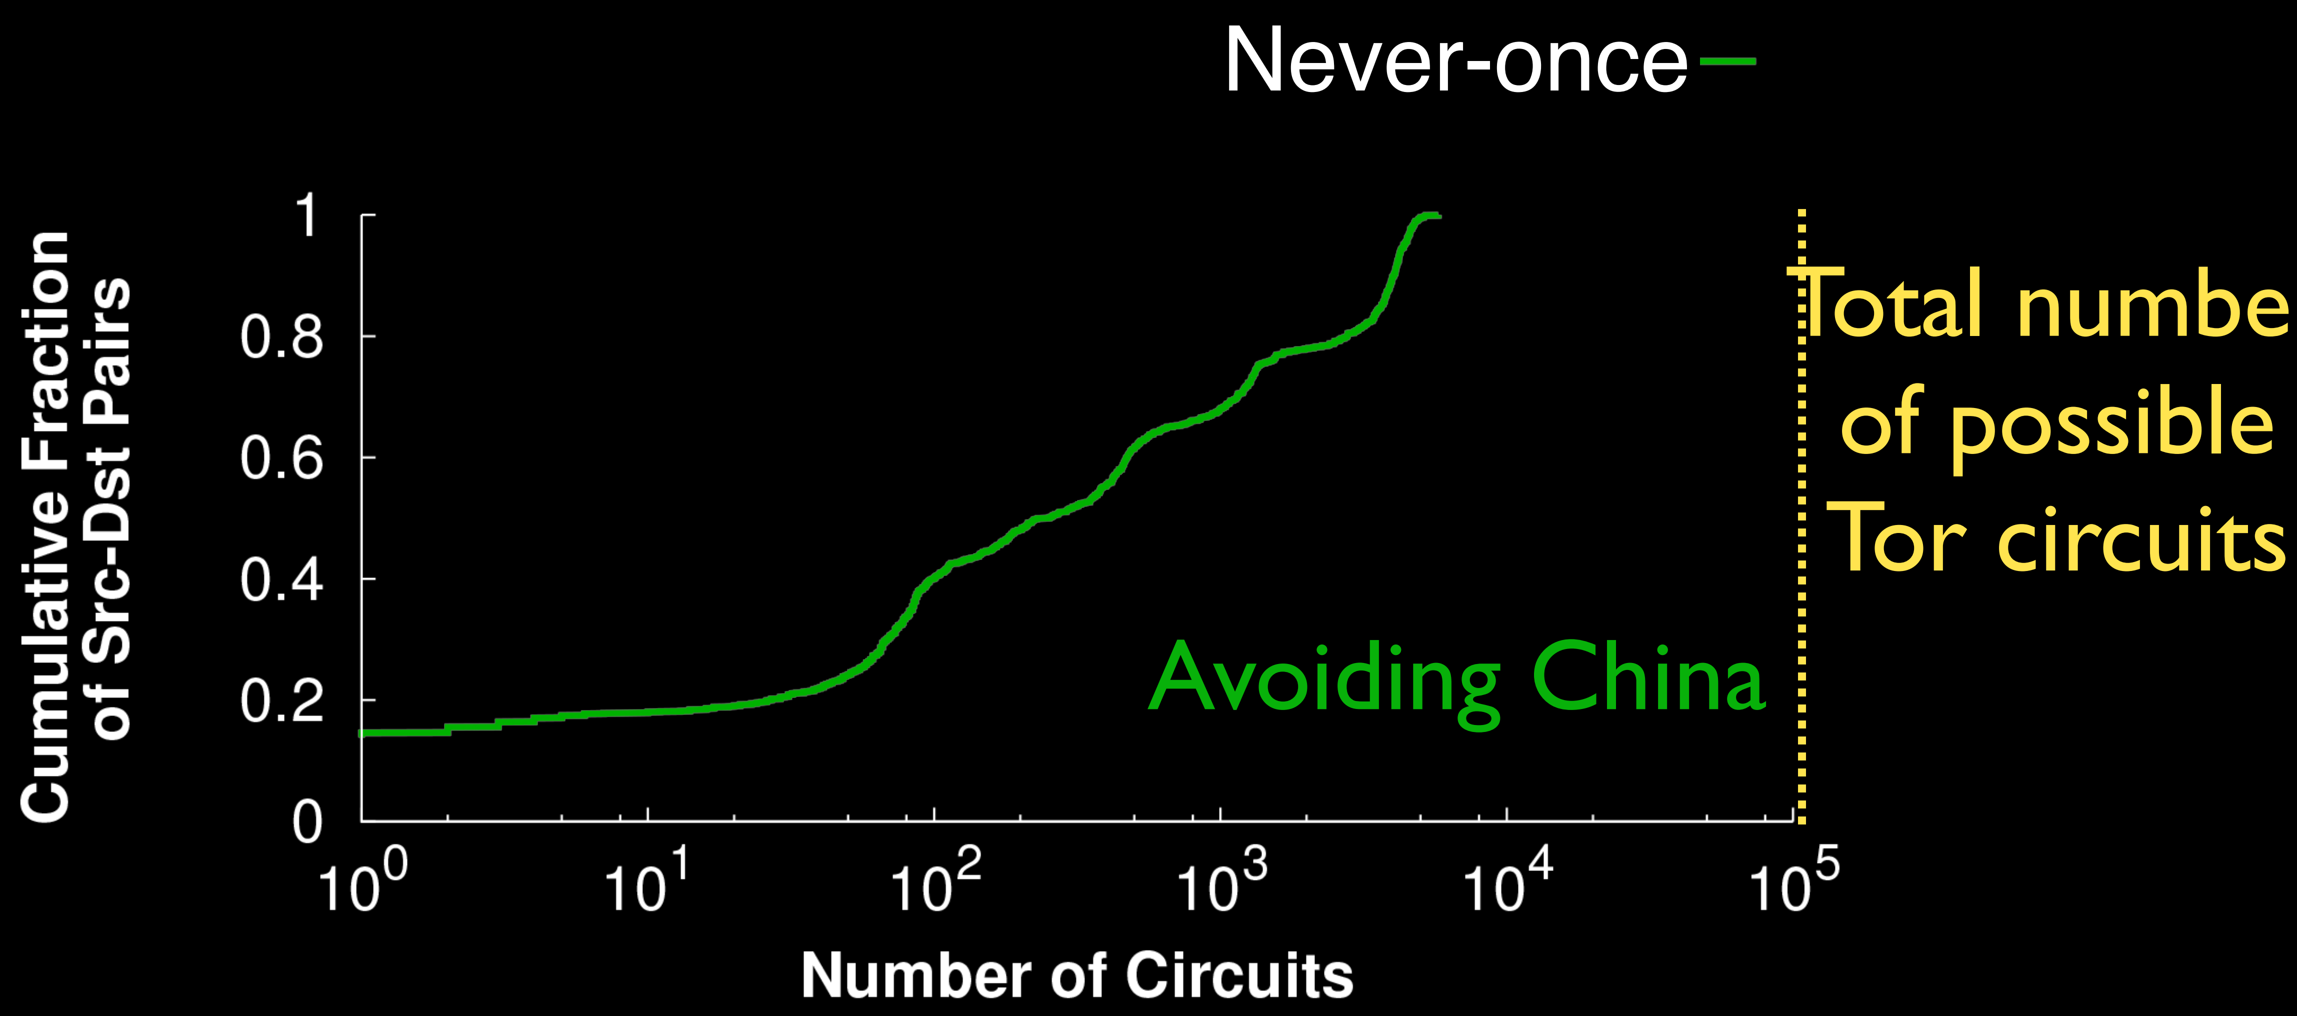

Number of never-once circuits

Tor with no Chinese relays provably avoids China
less than 10% of the time

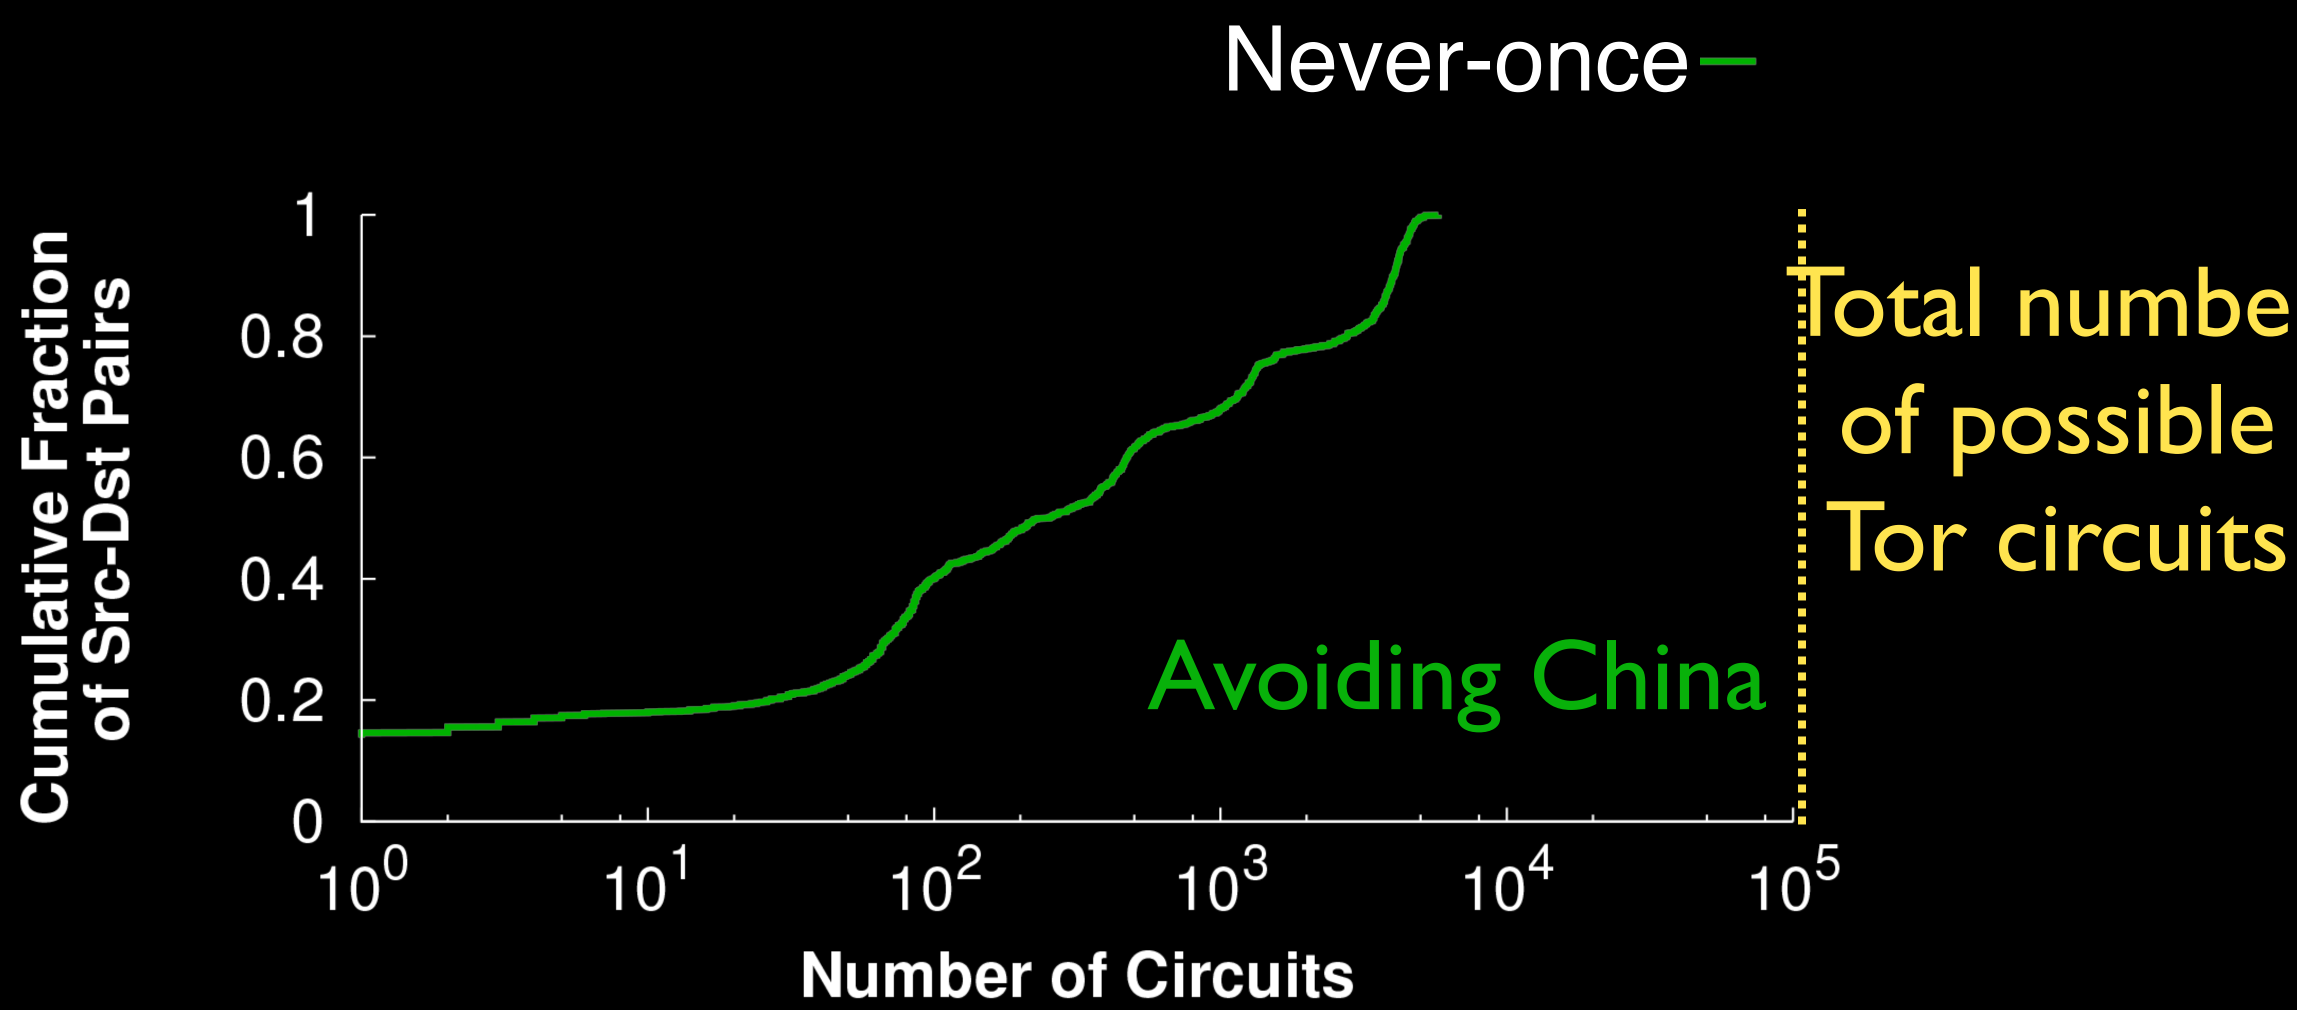

Number of never-twice circuits

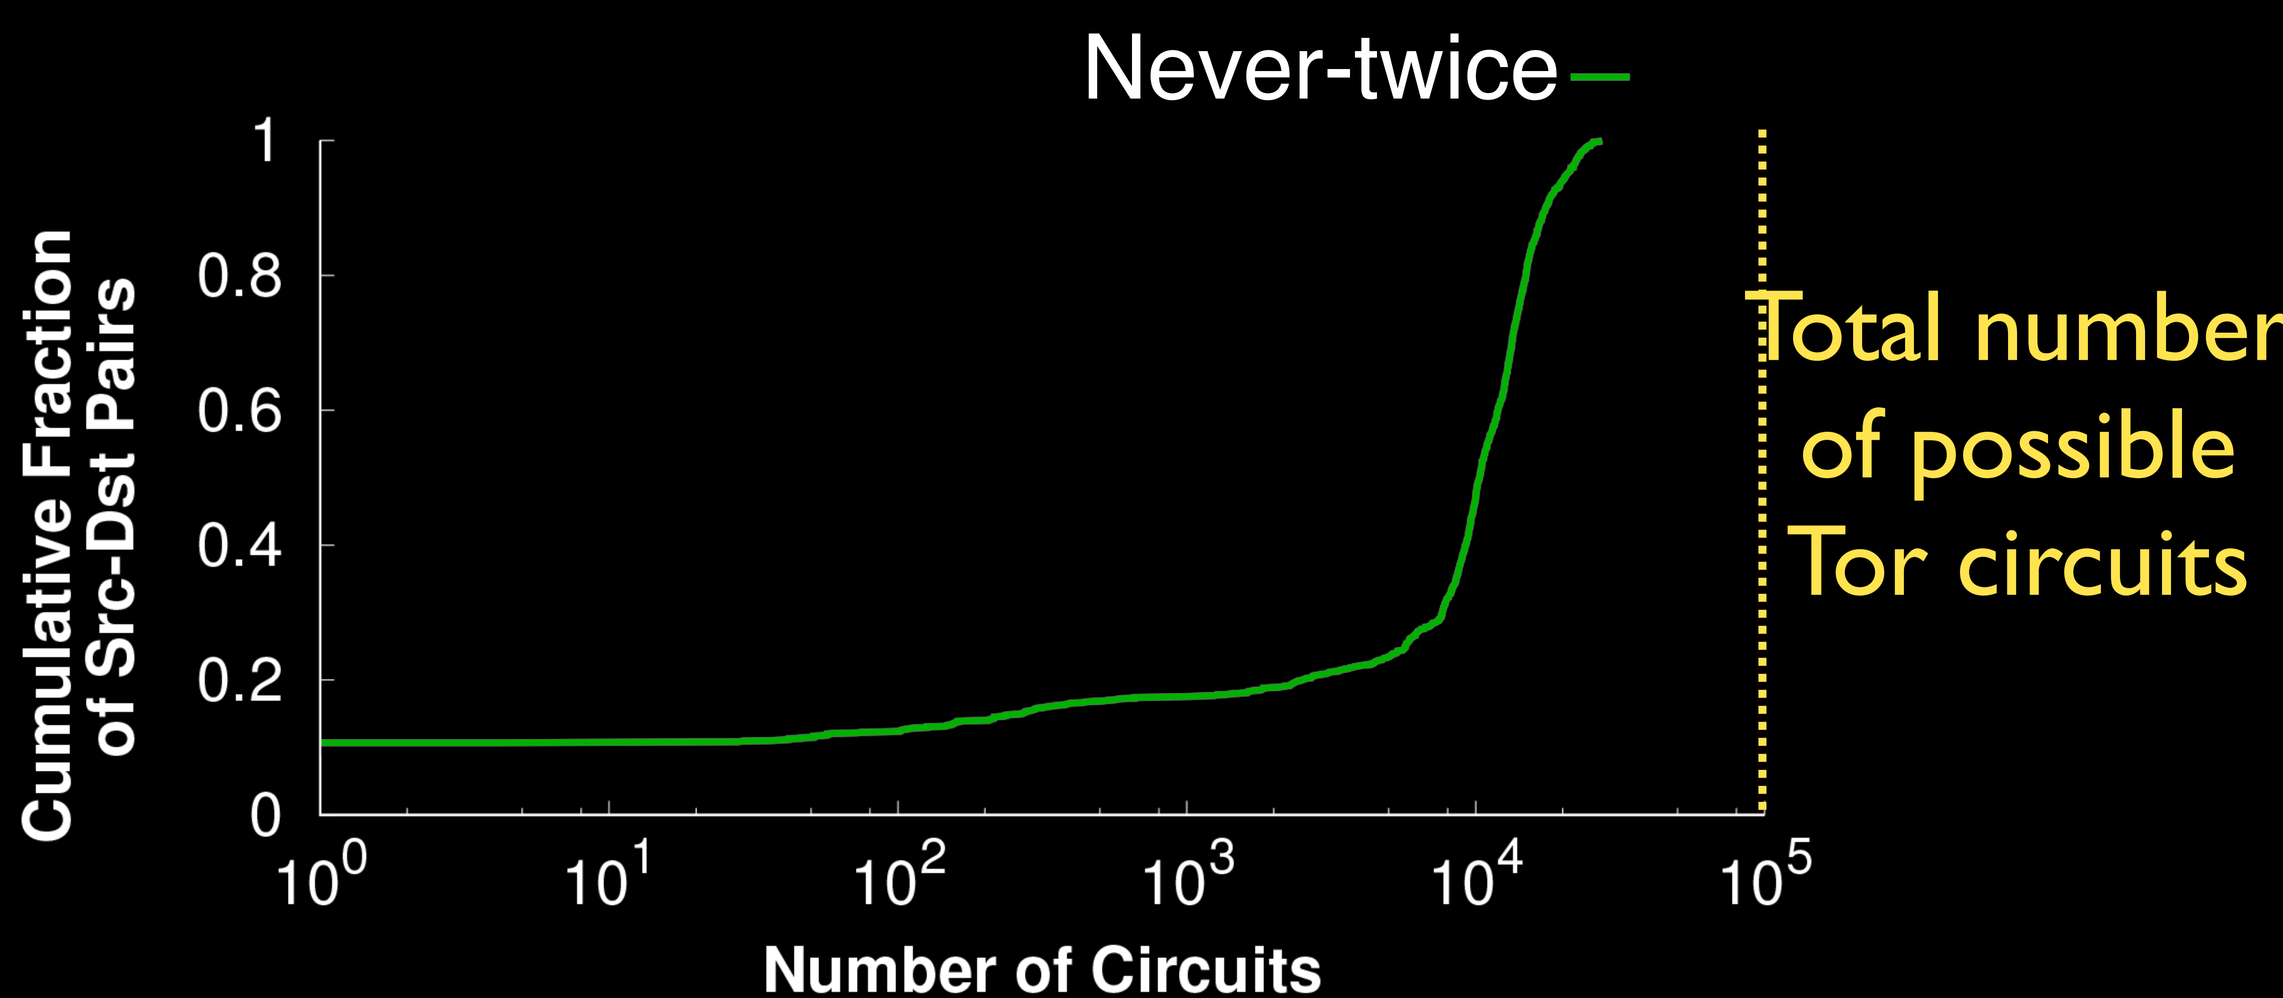

Number of never-twice circuits

Client-side RTTs might be enough to address many attacks

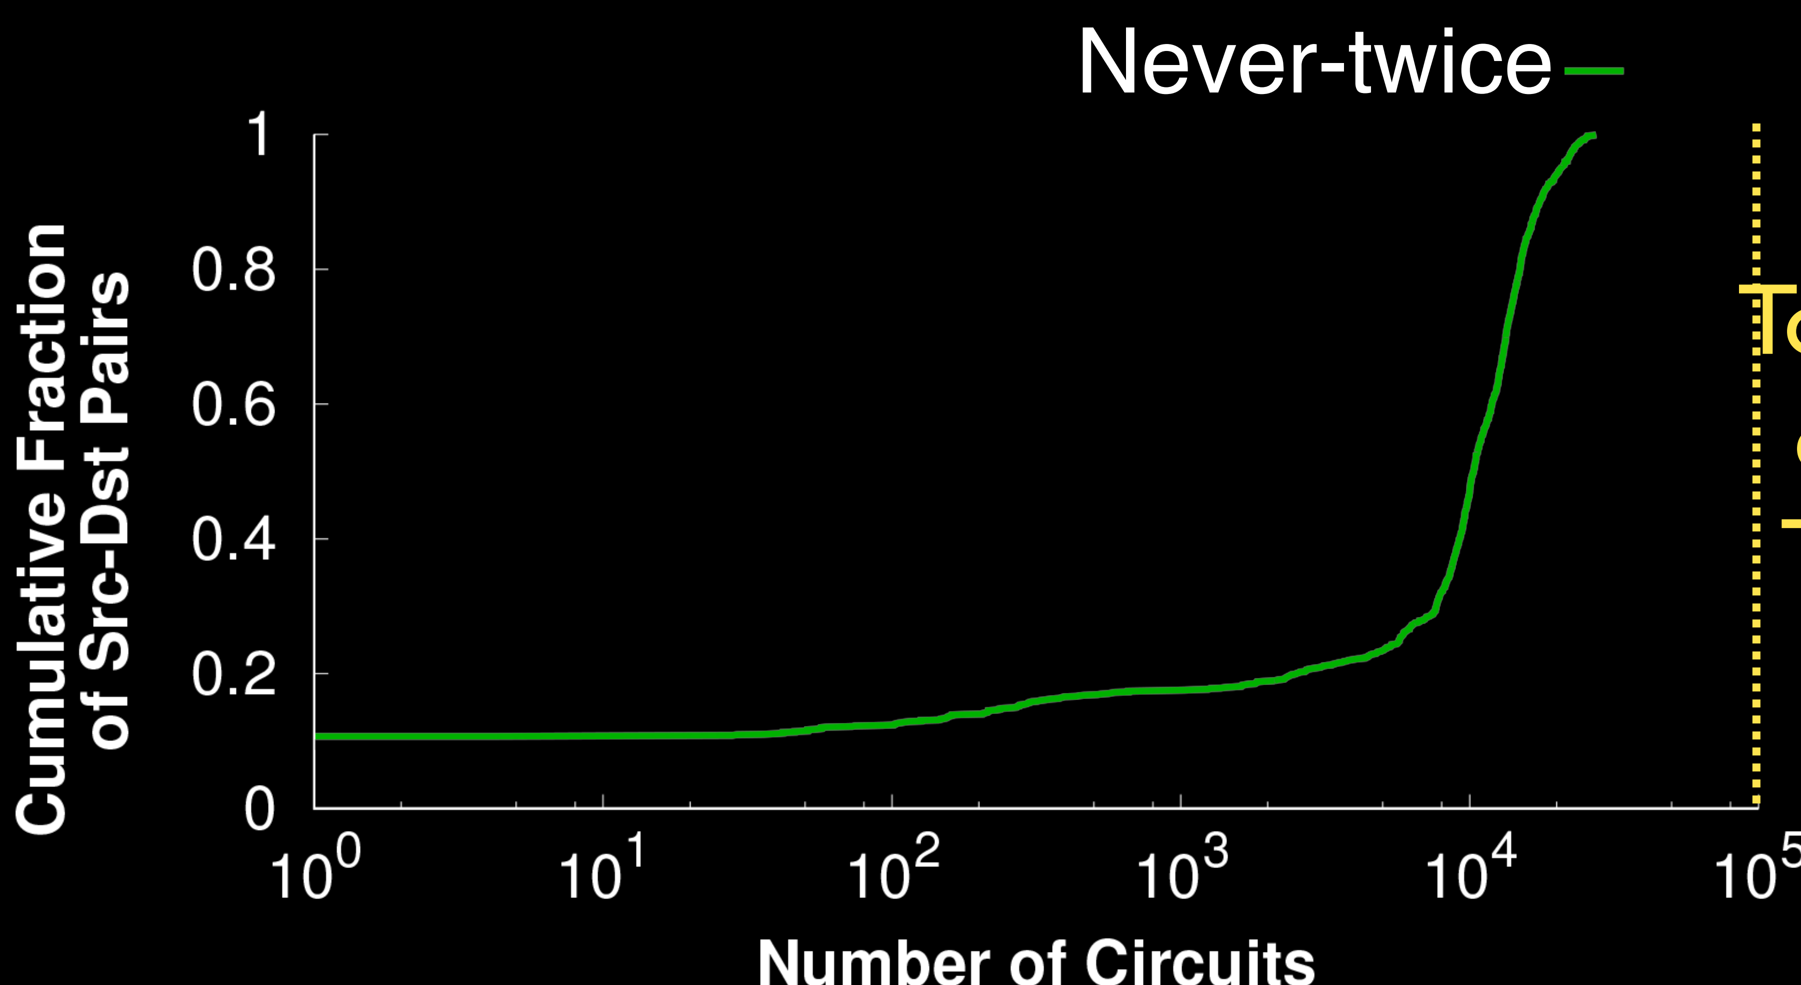

Total number of possible Tor circuits

DeTor circuits tends to have lower RTTs

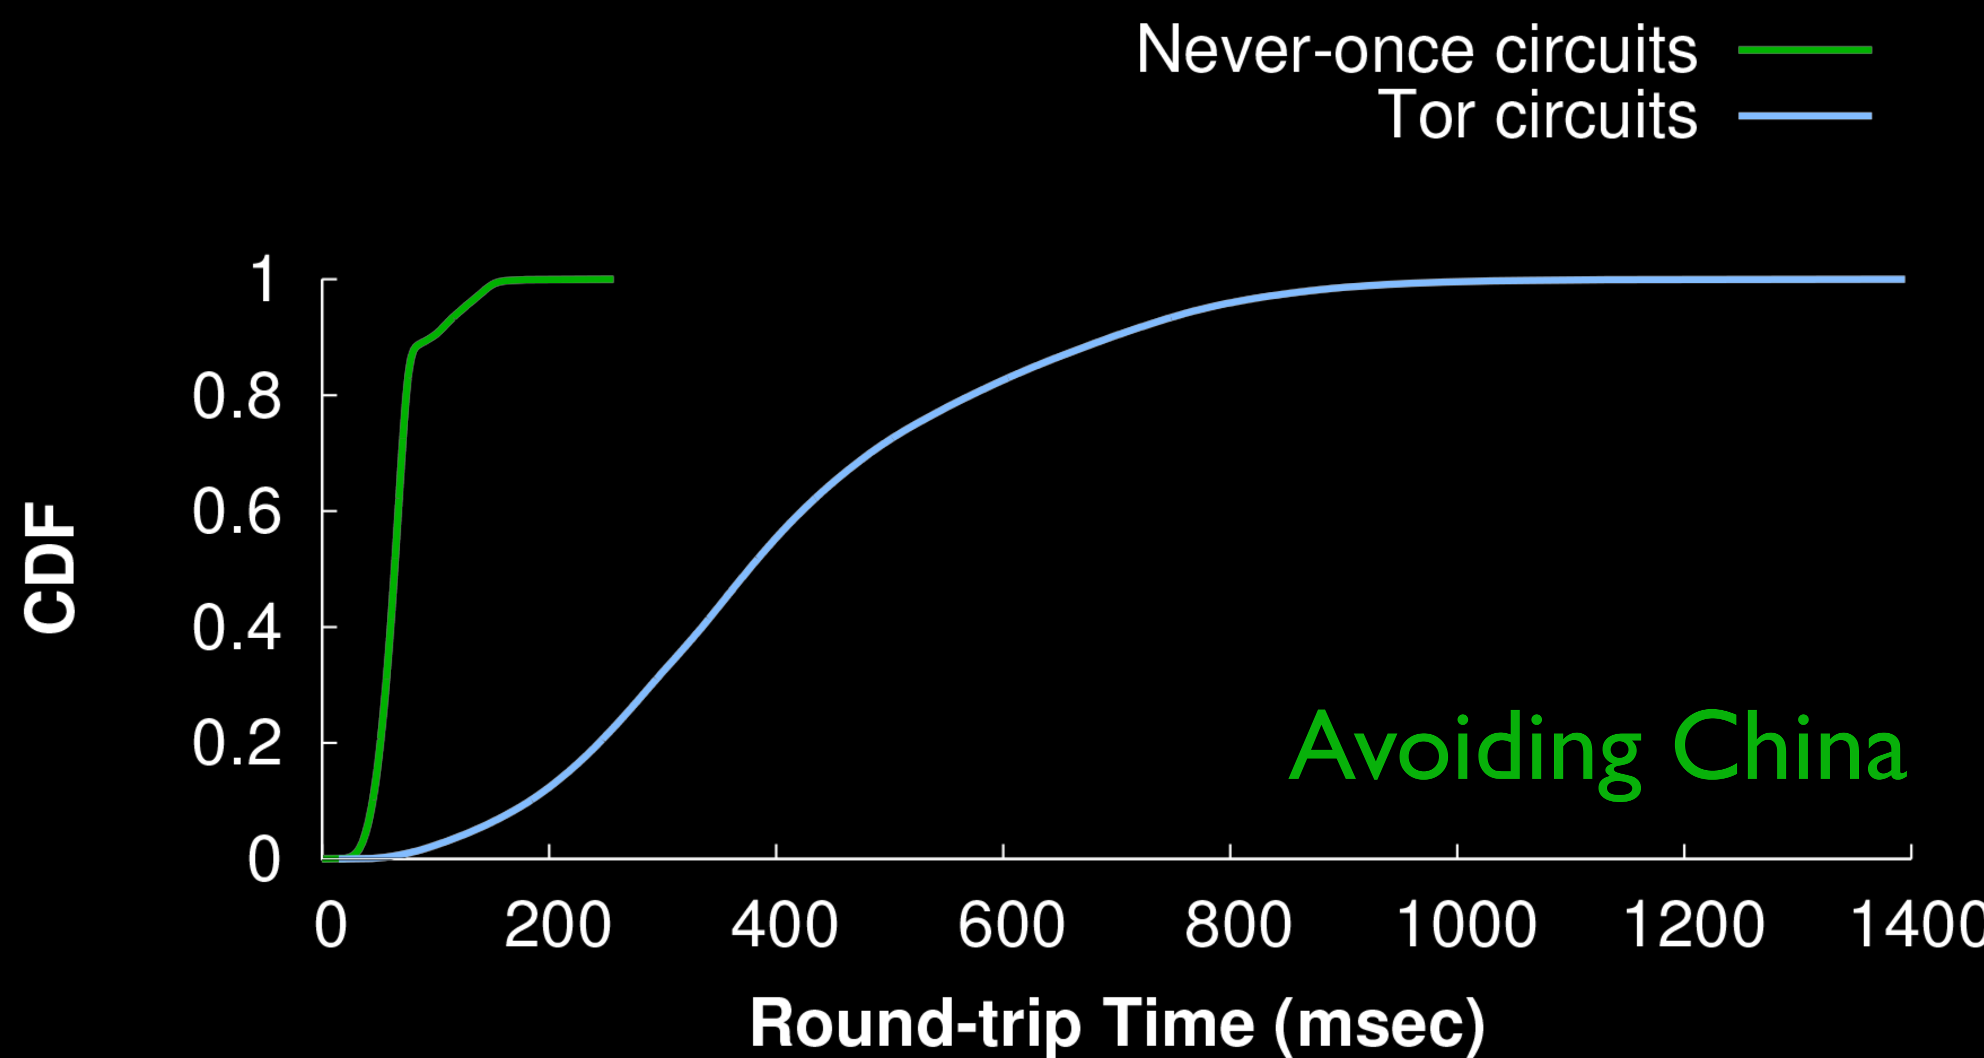

DeTor circuits tends to have lower RTTs

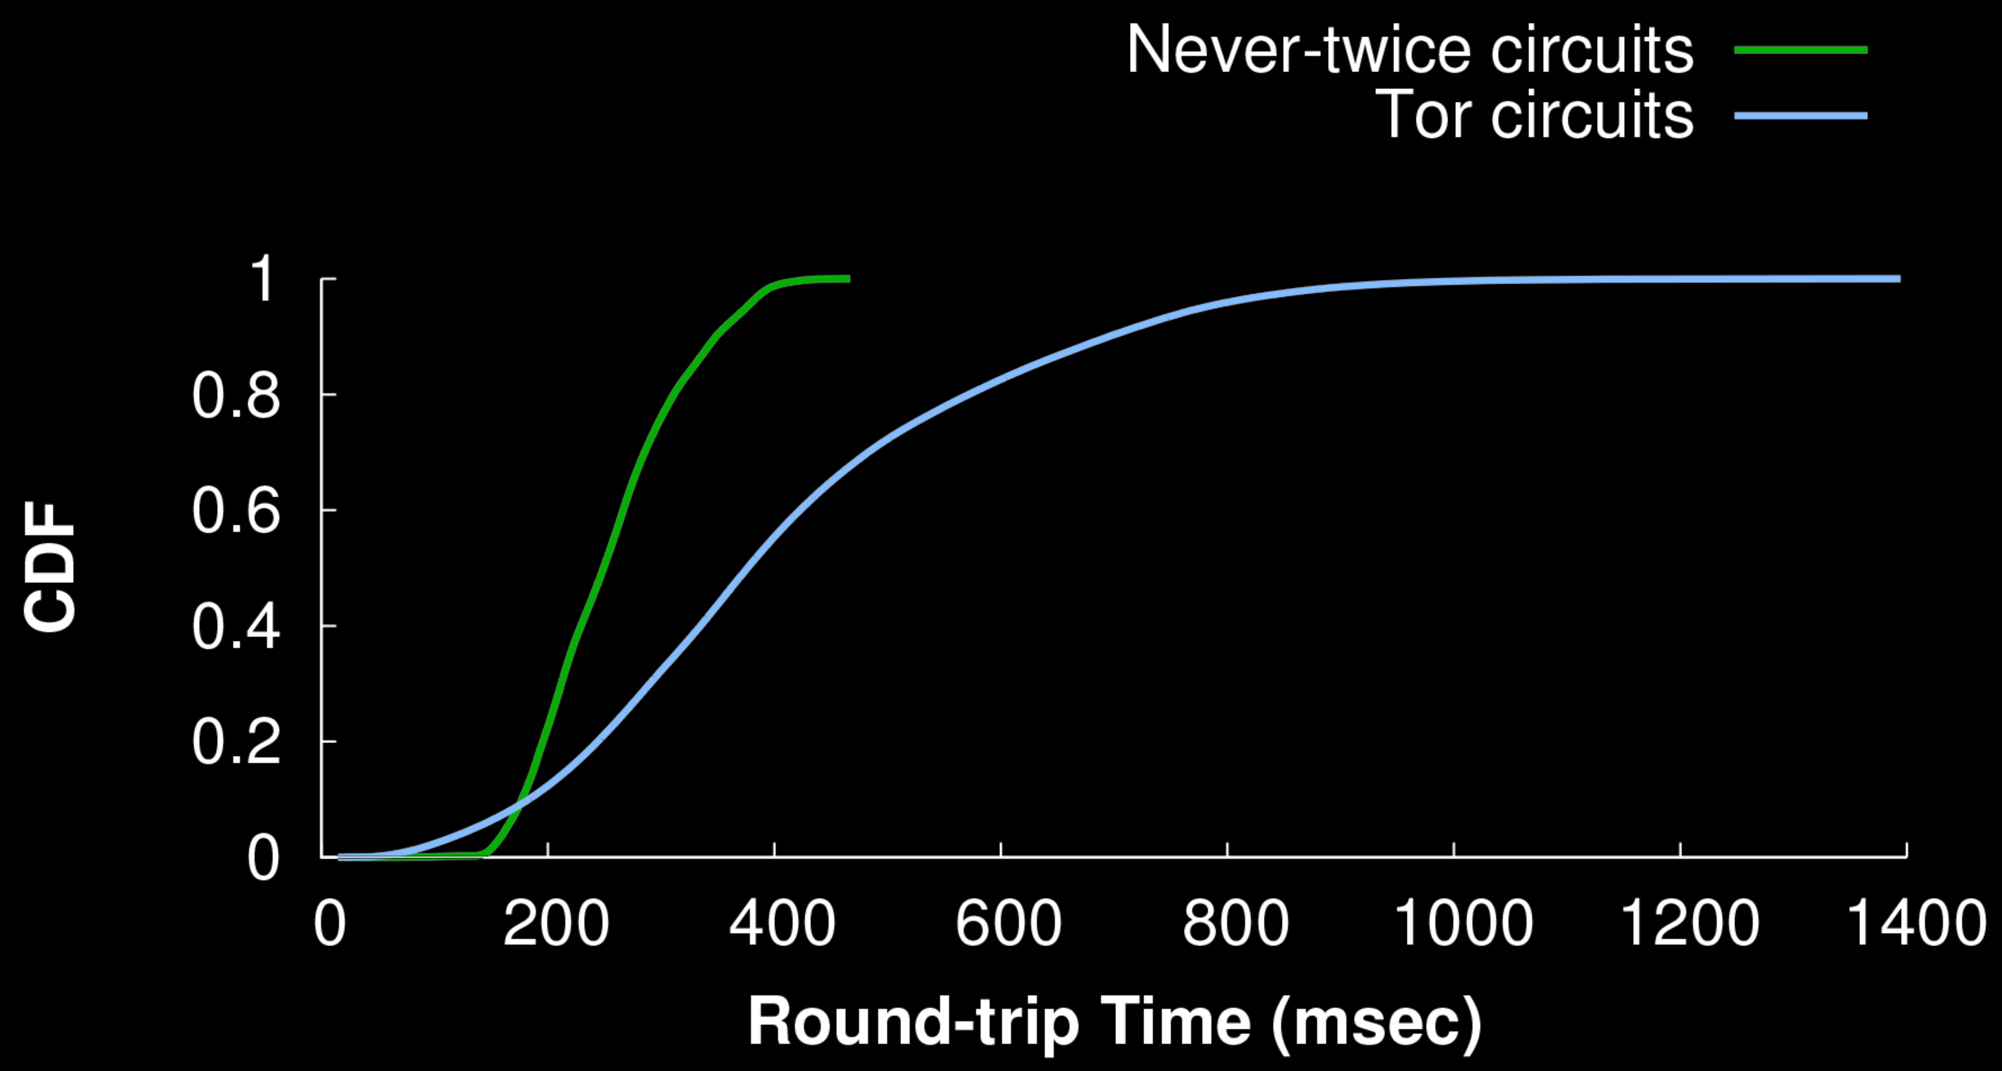

DeTor: never-once avoidance

Achieving provable avoidance

Measured RTT $\ll$ The *shortest possible* RTT thru and $= 2 \min\{d_i\} / c$

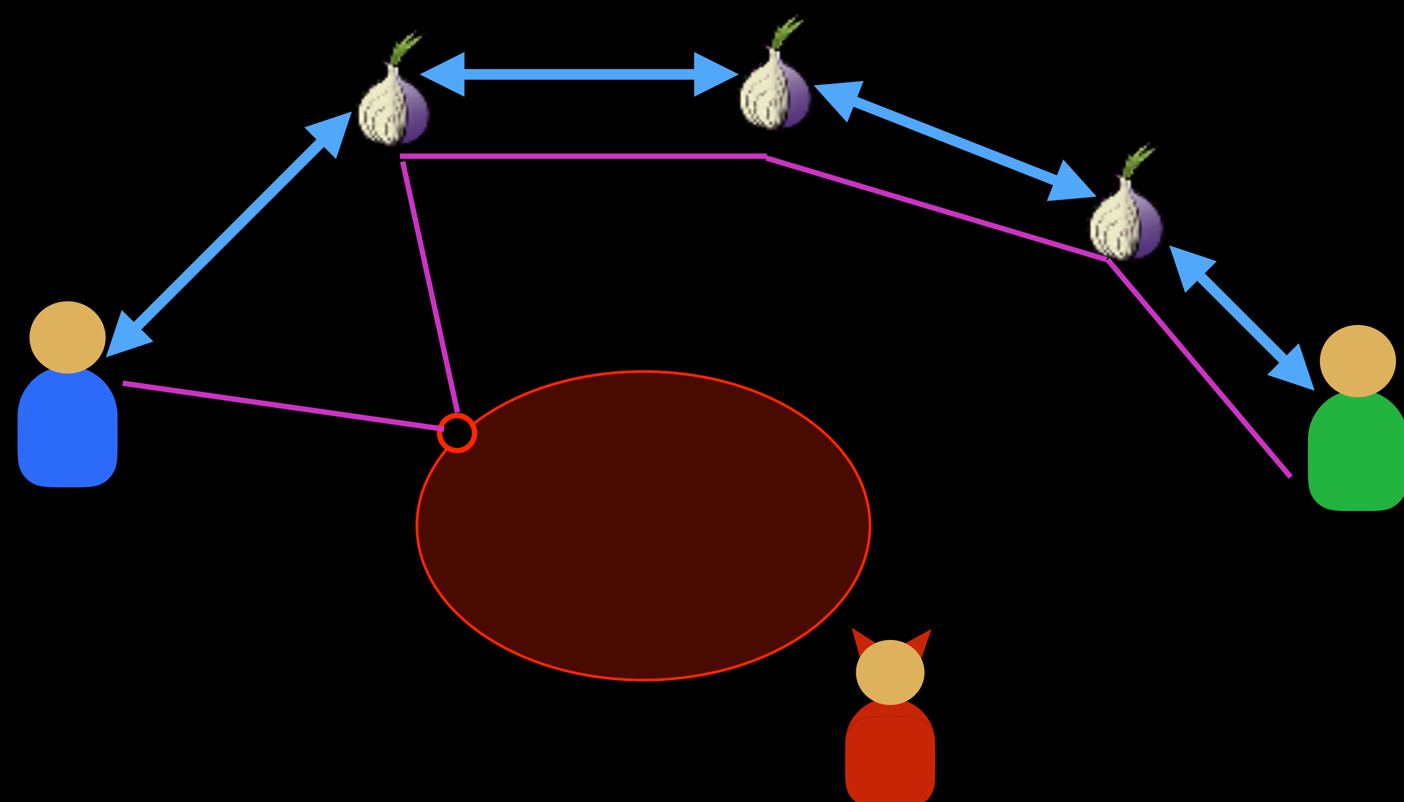

DeTor: never-once avoidance

Achieving provable avoidance

Measured RTT $\ll$ The *shortest possible* RTT thru and $= 2 \min\{d_i\} / c$

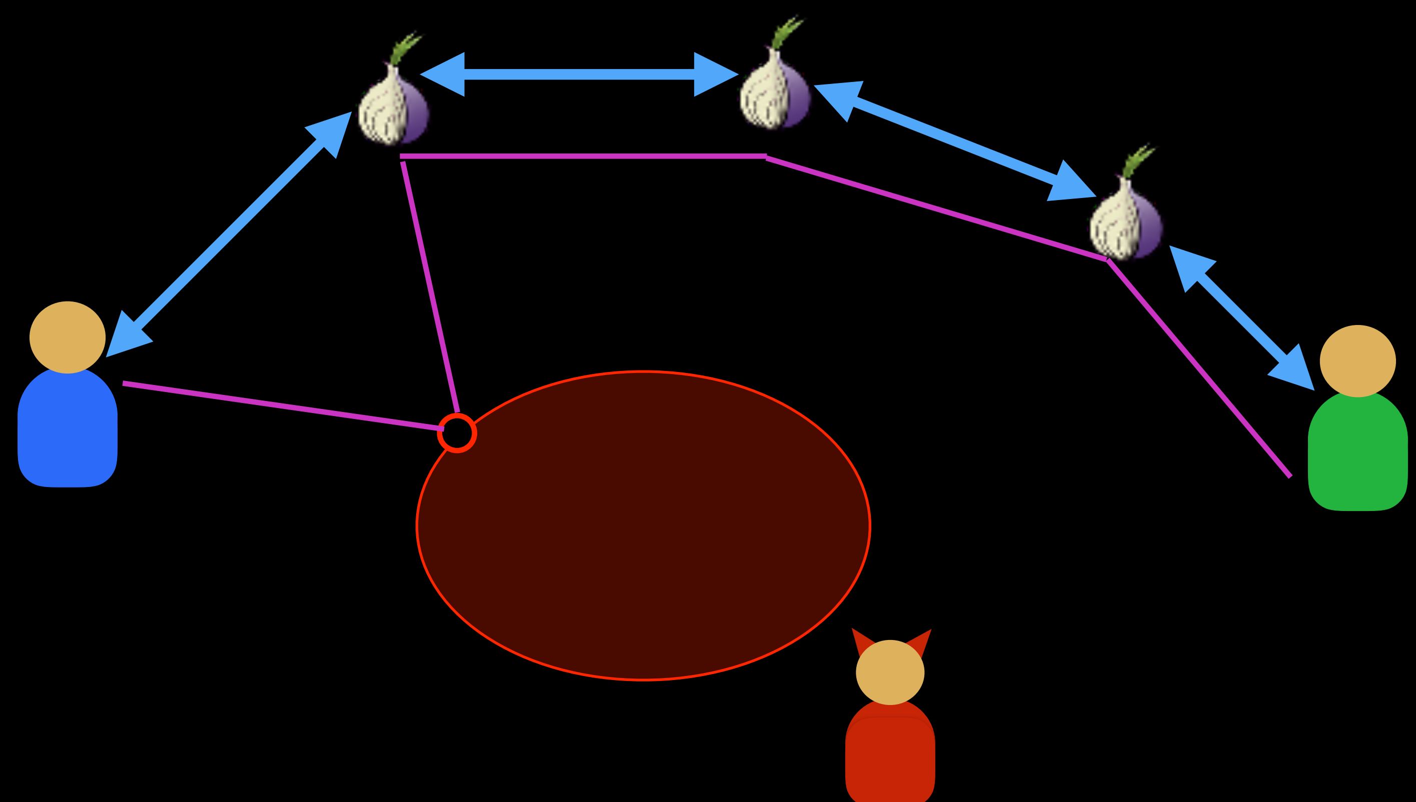

Other results

- DeTor circuits usually have higher bandwidth
- DeTor introduces slight node selection bias
- Most nodes serve on few DeTor circuits
- Possible to predict whether a circuit will

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once	Never-twice
never traverse specified regions	entry & exit legs never traverse

- Proofs of avoidance verify that packets over DeTor circuits have avoided geographic regions

- DeTor circuits
  - are successful for most src-dst pairs
  - have better performance
  - introduce small node selection bias

Code and data available at:
`detor.cs.umd.edu`

DeTor

With smart circuit selection, it is possible to *provably* avoid geographic regions with Tor

Never-once	Never-twice
never traverse specified regions	entry & exit legs never traverse

- Proofs of avoidance verify that packets over DeTor circuits have avoided geographic regions
- DeTor circuits
  - are successful for most src-dst pairs
  - have better performance
  - introduce small node selection bias

Code and data available at:
`detor.cs.umd.edu`