Anonymity on Quicksand

Using BGP to compromise Tor

Laurent Vanbever

Princeton/ETH Zürich

HotNets

October, 28 2014

Joint work with

Oscar Li, Jennifer Rexford, Prateek Mittal

WE ARE
ANONYMOUS

WE ARE NOT
ANONYMOUS FOR LONG.

WE ARE NOT
ANONYMOUS FOR LONG.
COURTESY OF BGP.

Internet communications are *not* anonymous

Looking at an Internet communication, one can

- infer who is talking to whom
- infer physical locations
- use that to track behavior and interests

even if the communication is encrypted

Tor aims at preventing adversaries to follow packets between a sender and a receiver

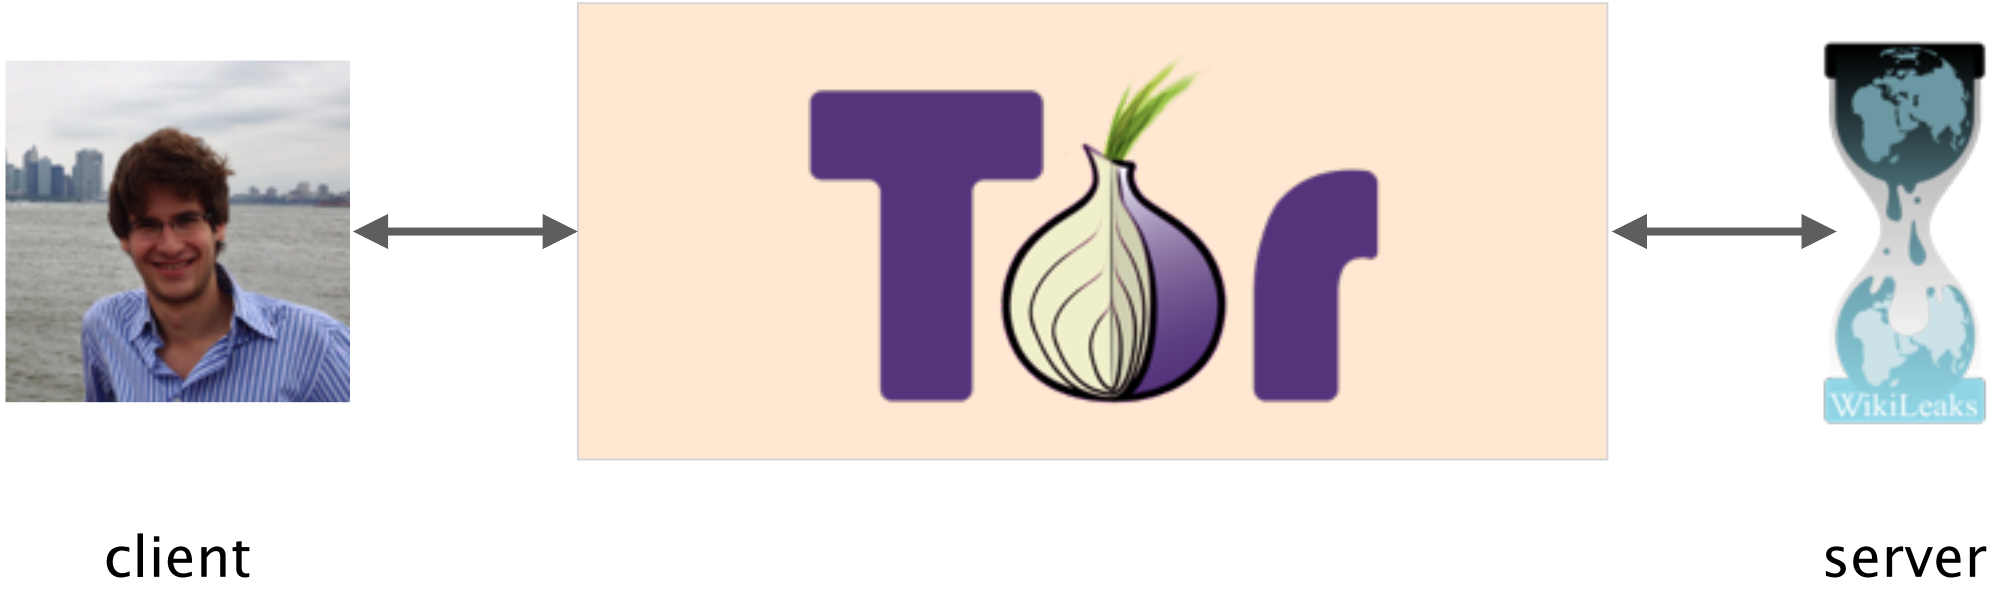

client

server

To do that,
Tor bounces traffic around a network of relays

Tor network

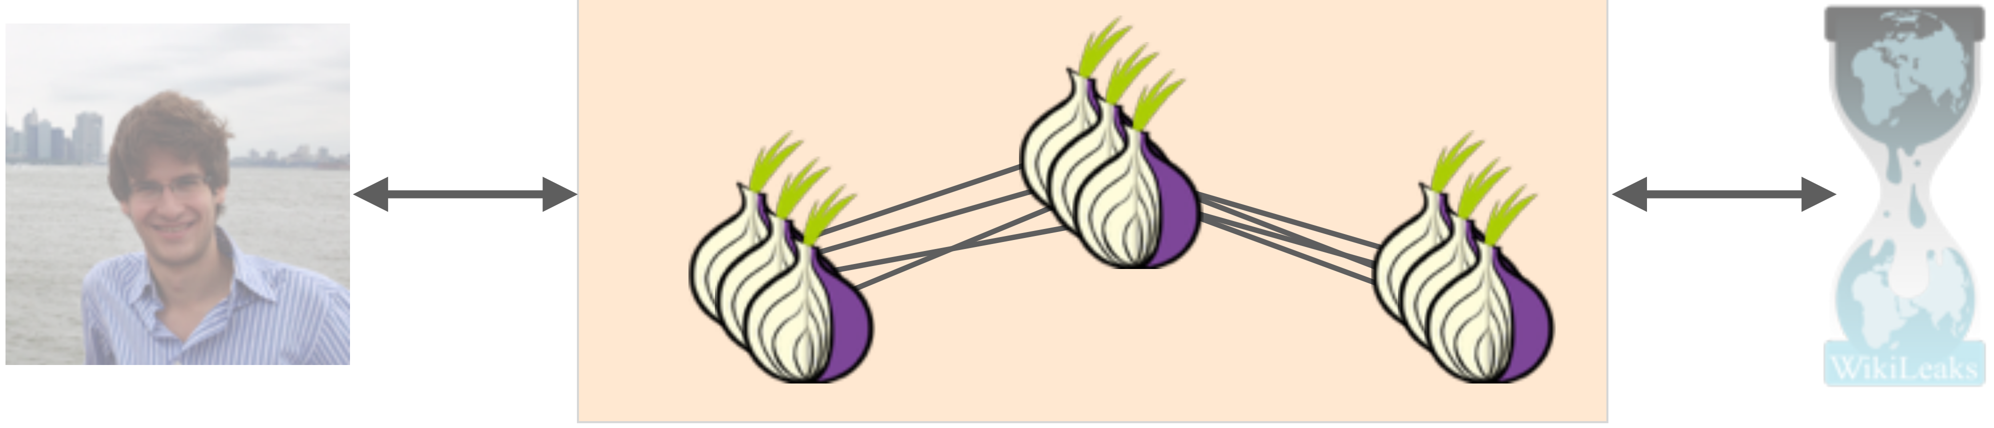

client

entry
point

middle

exit
point

server

Tor clients start by selecting 3 relays, one of each type

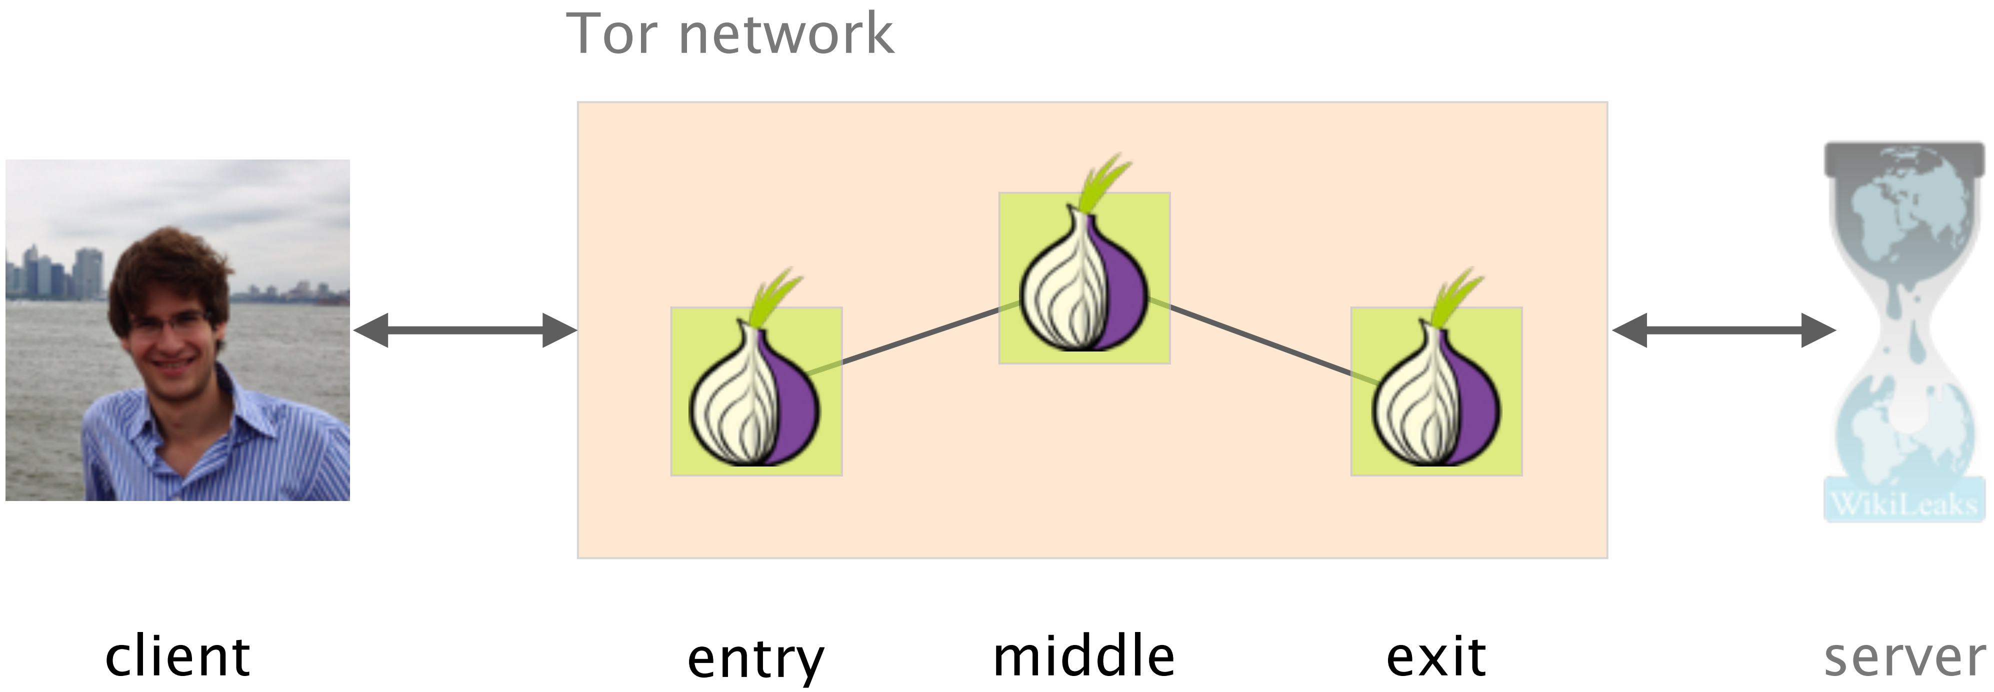

Tor clients then incrementally build encrypted circuits through them

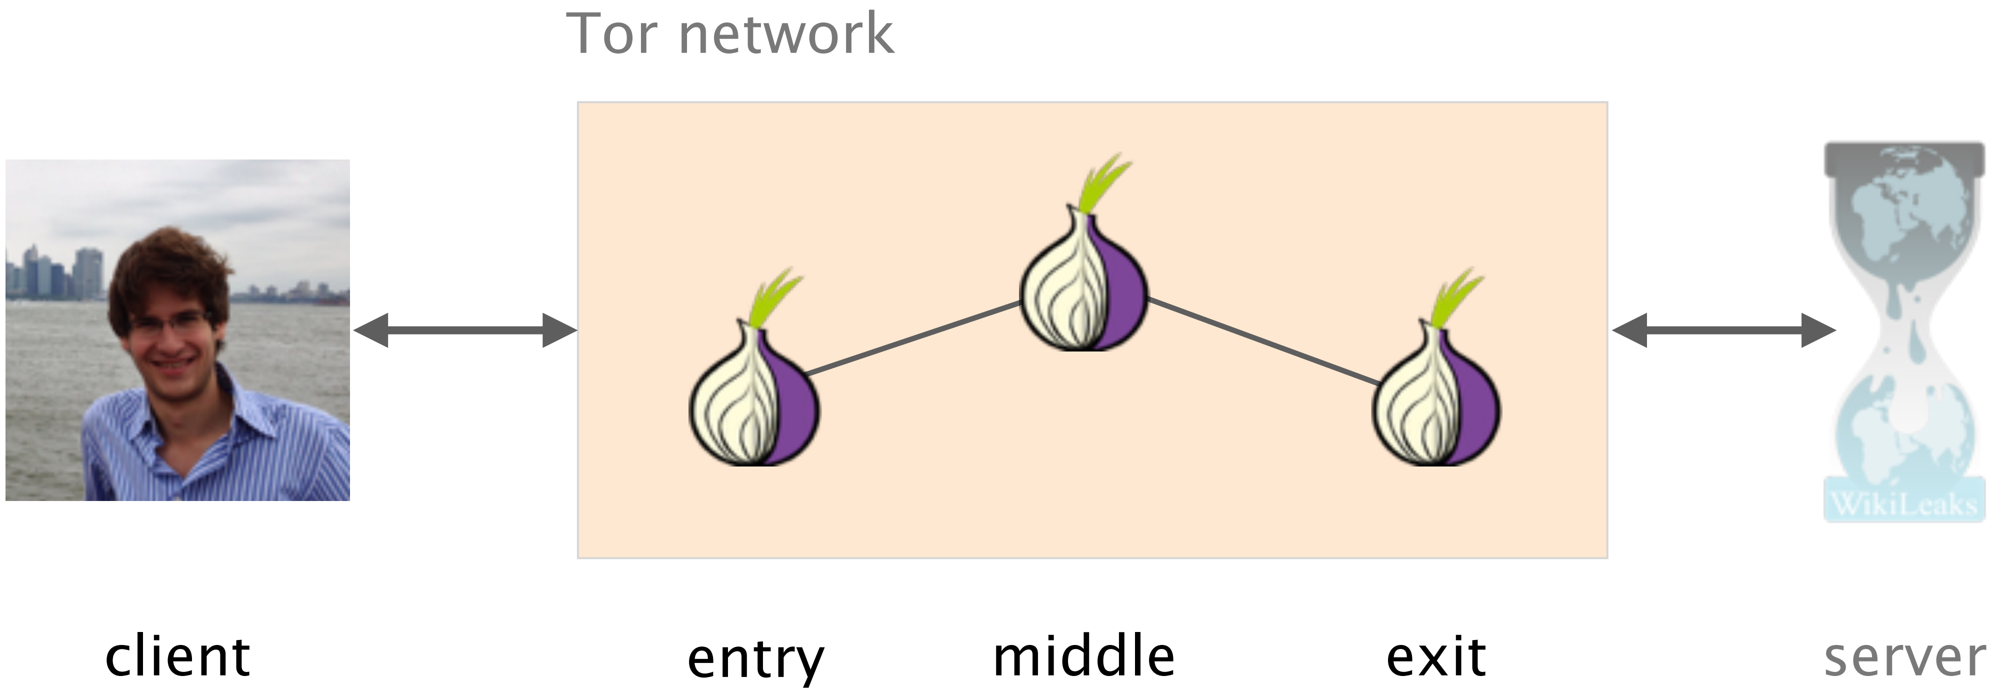

Tor network

client

entry

middle

exit

server

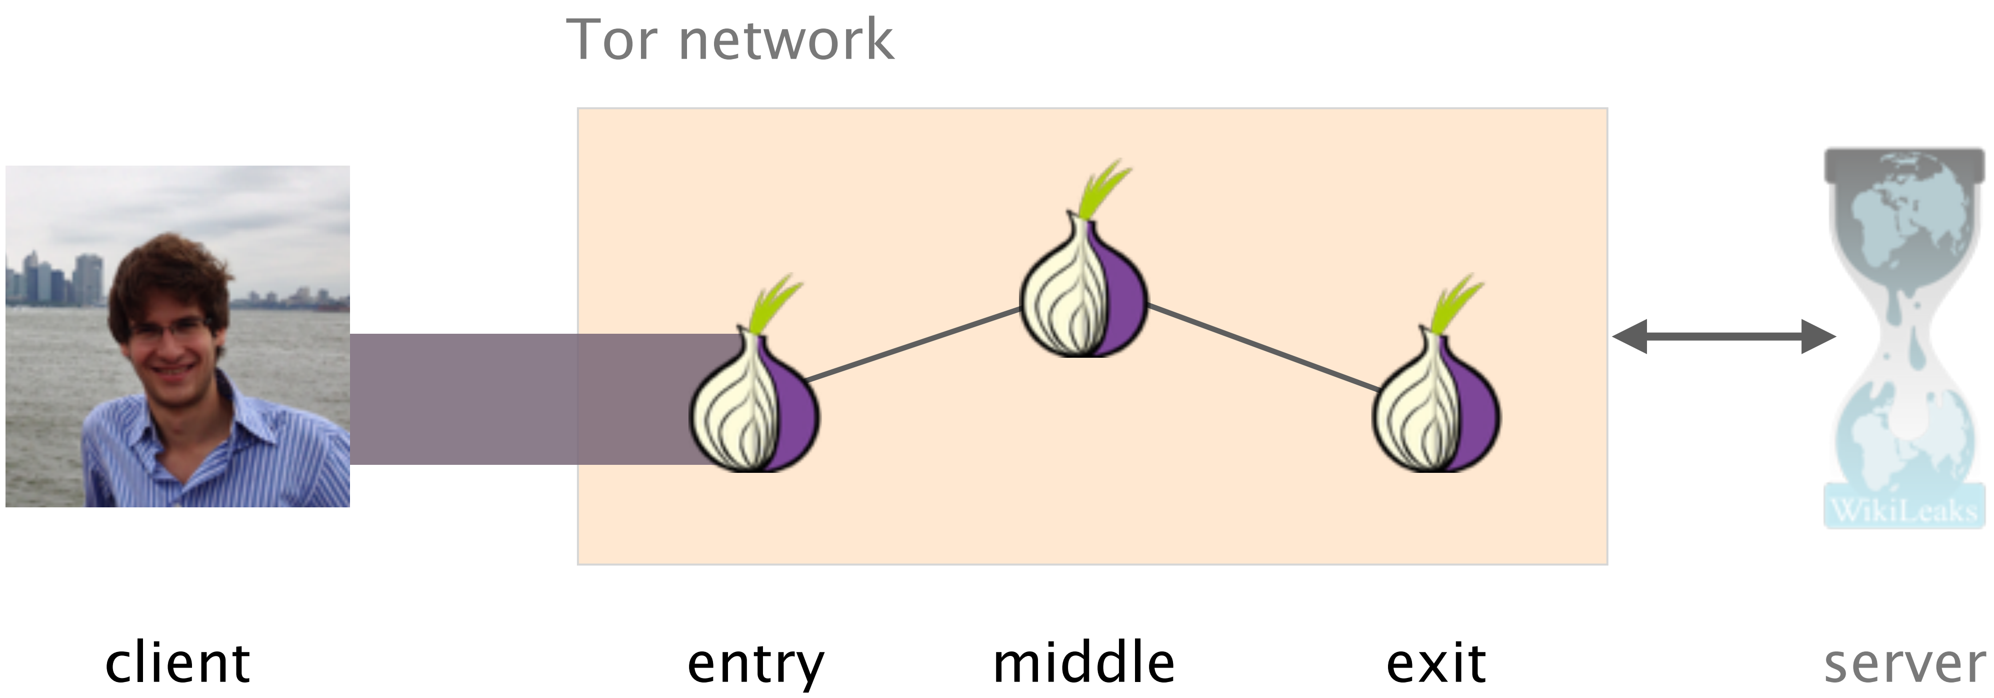

Tor network

client

entry

middle

exit

server

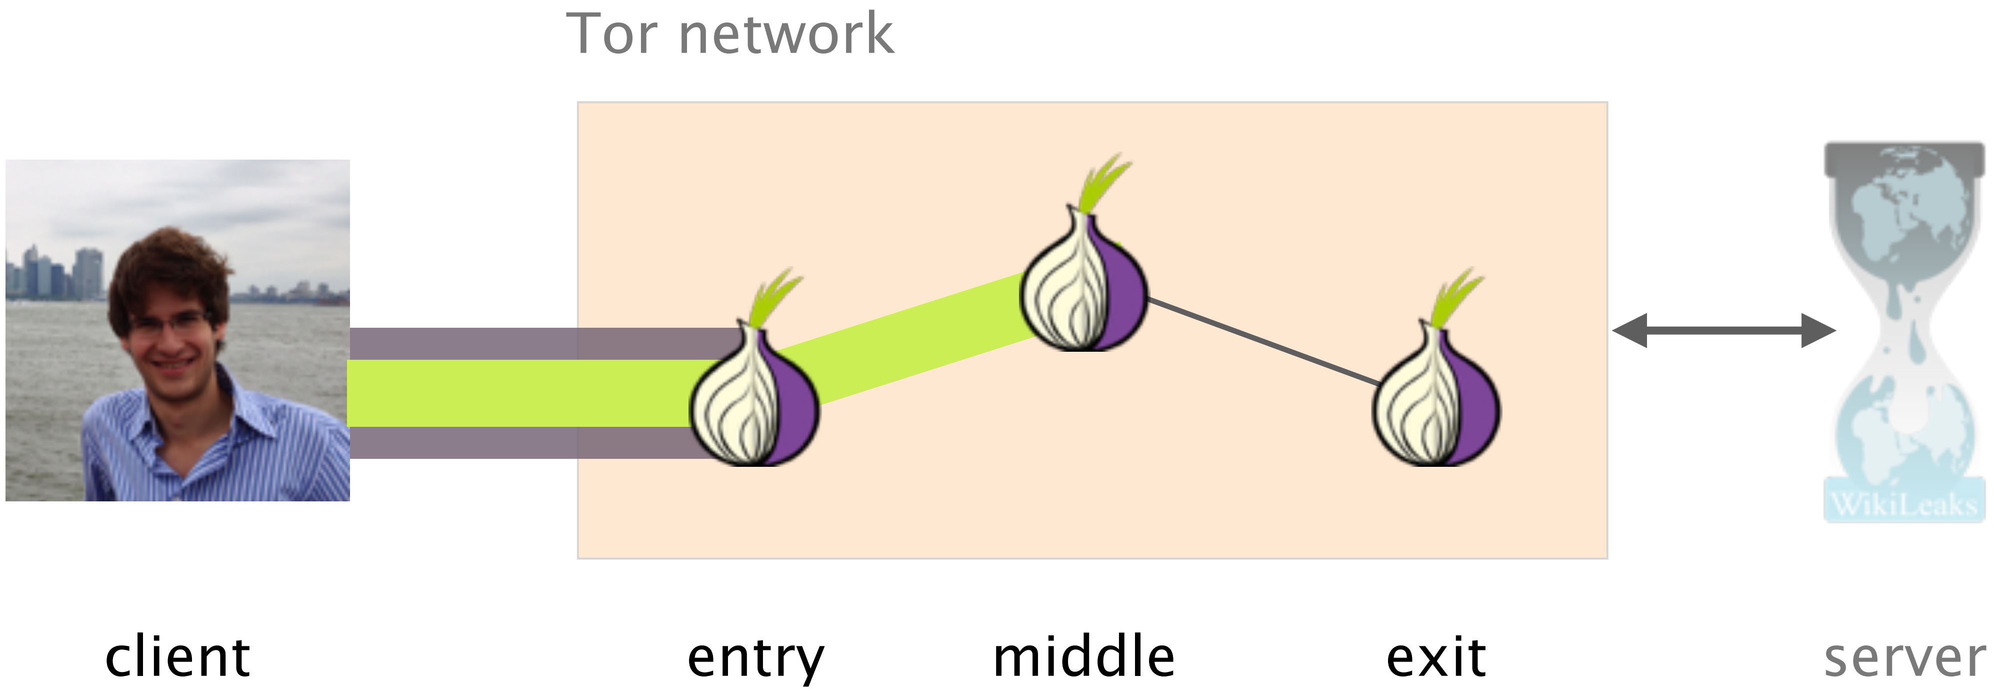

Tor network

client

entry

middle

exit

server

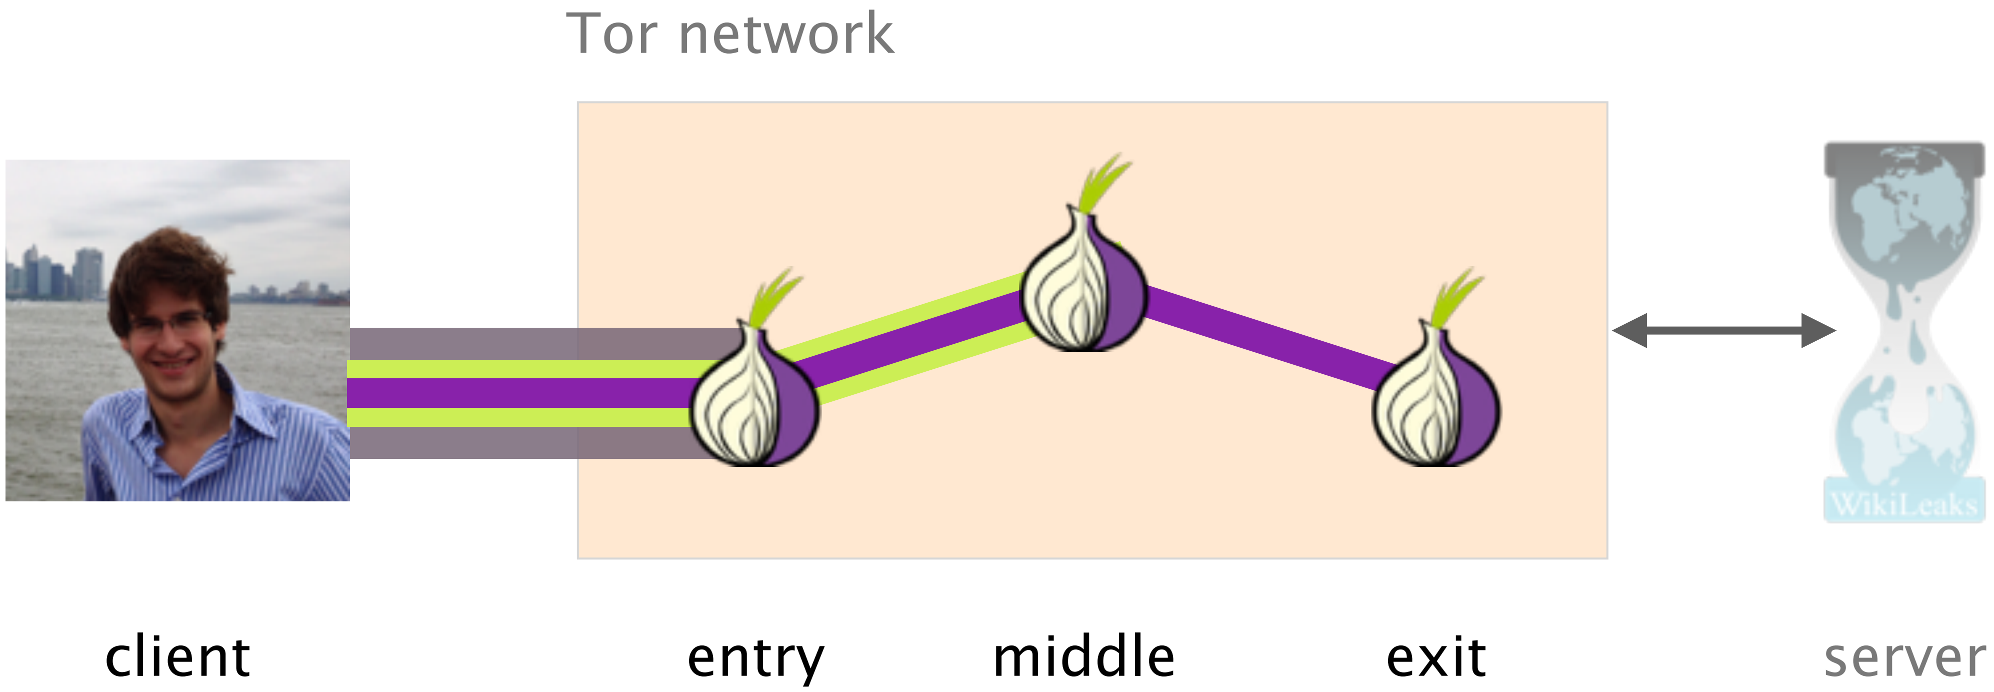

Tor network

client

entry

middle

exit

server

Anonymous communication takes place by forwarding across consecutive tunnels

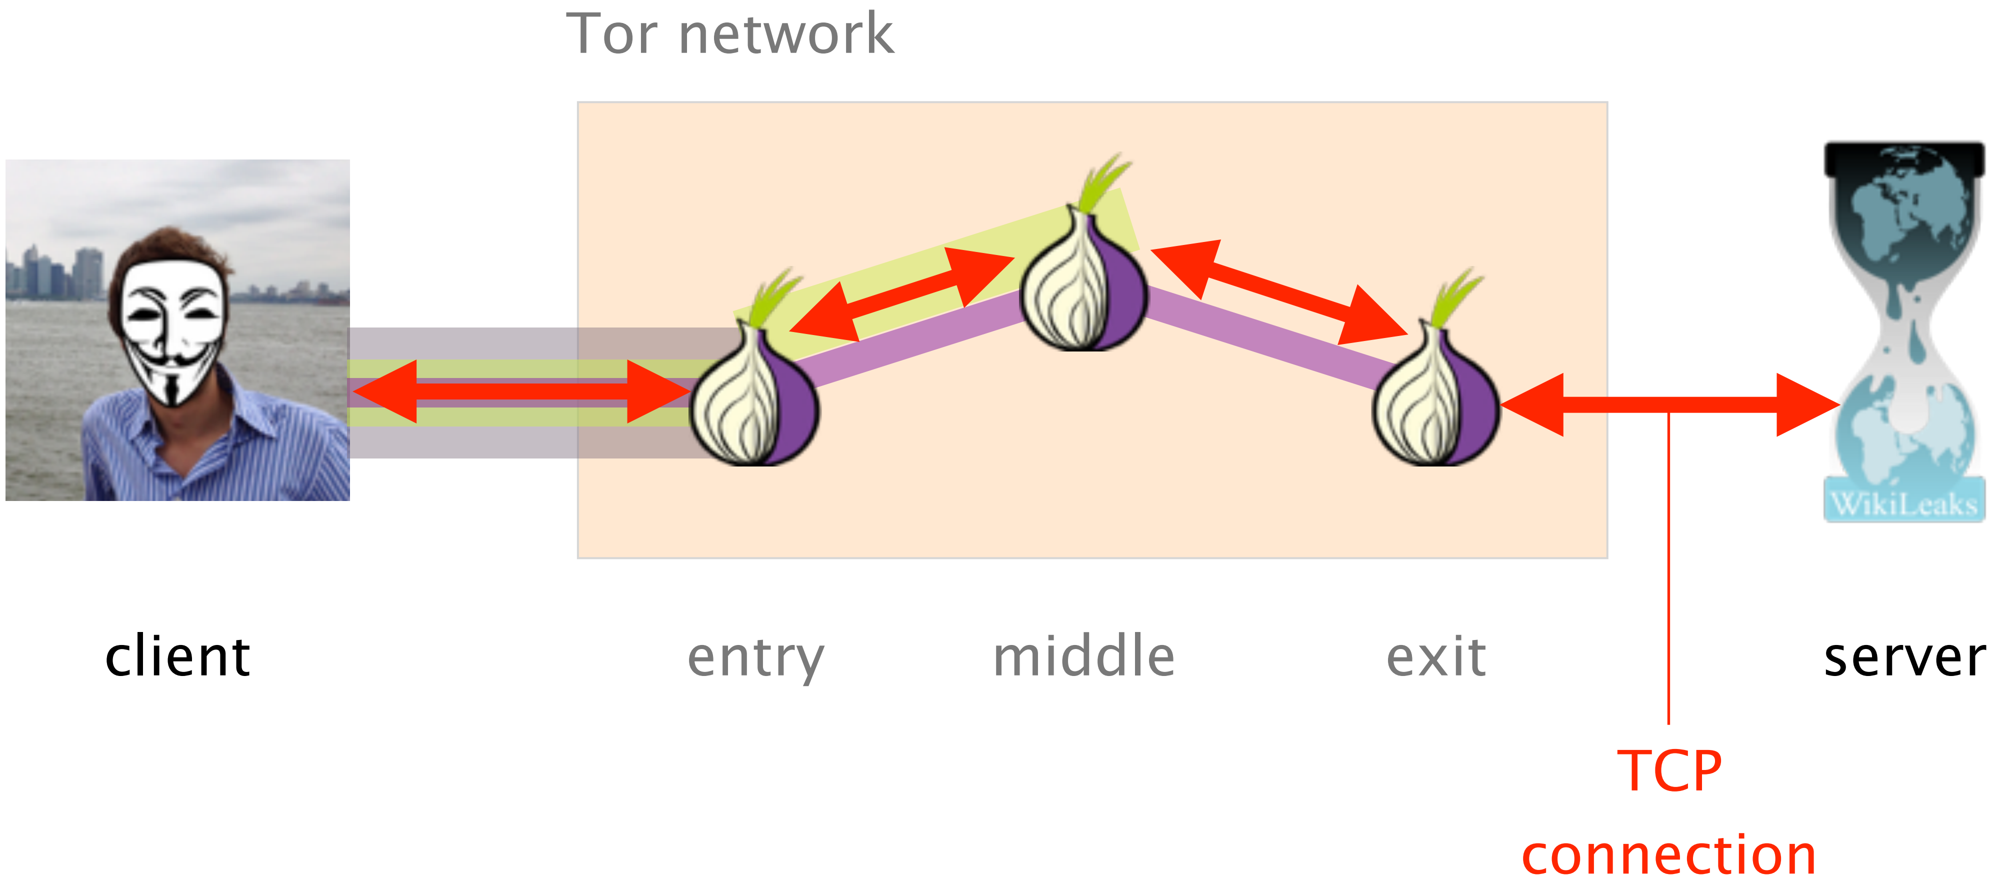

Not a single Tor entity knows the association (client, server)

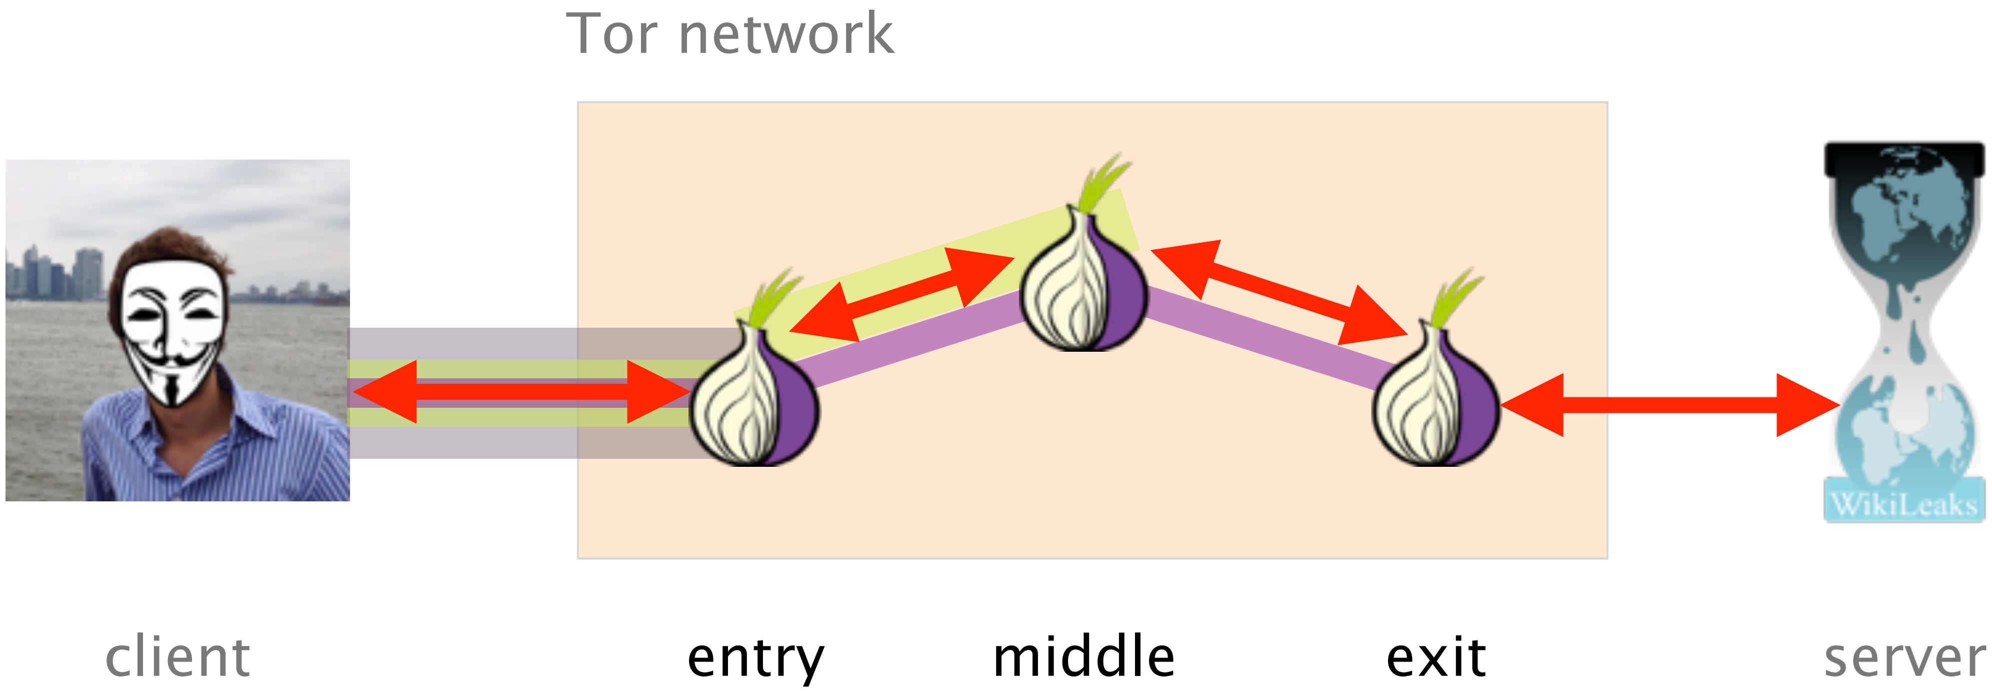

Tor network

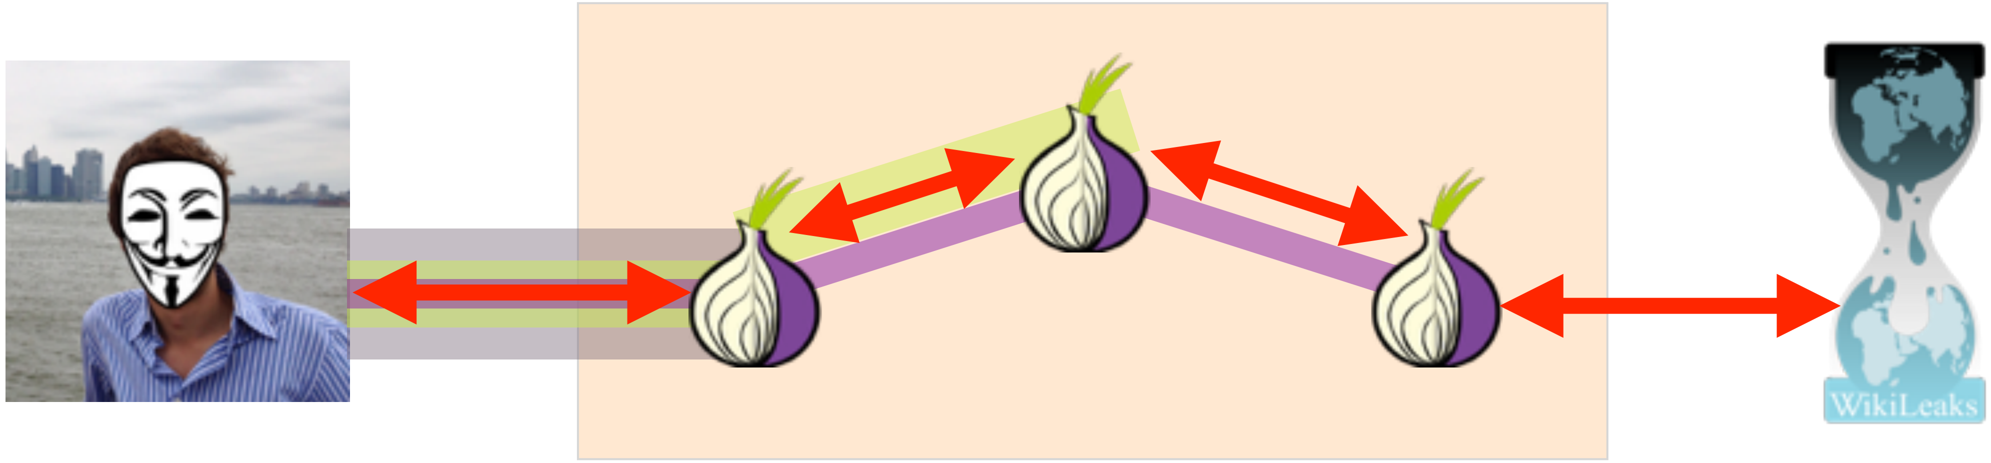

client

entry

middle

exit

server

knows the source,
not the destination

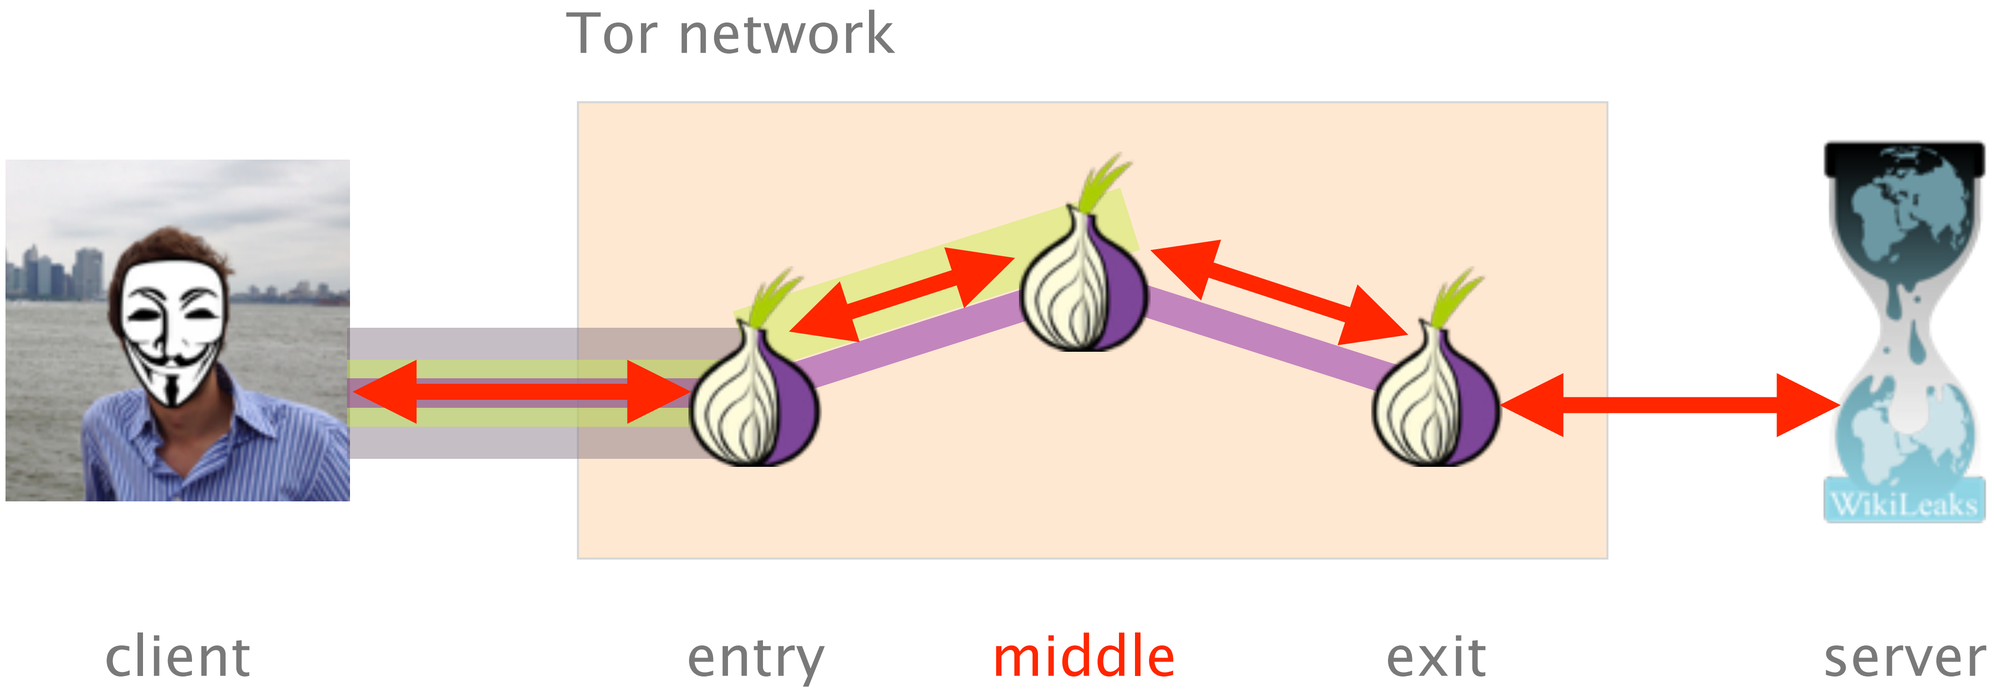

Tor network

client

entry

middle

exit

server

knows neither the source,
nor the destination

Tor network

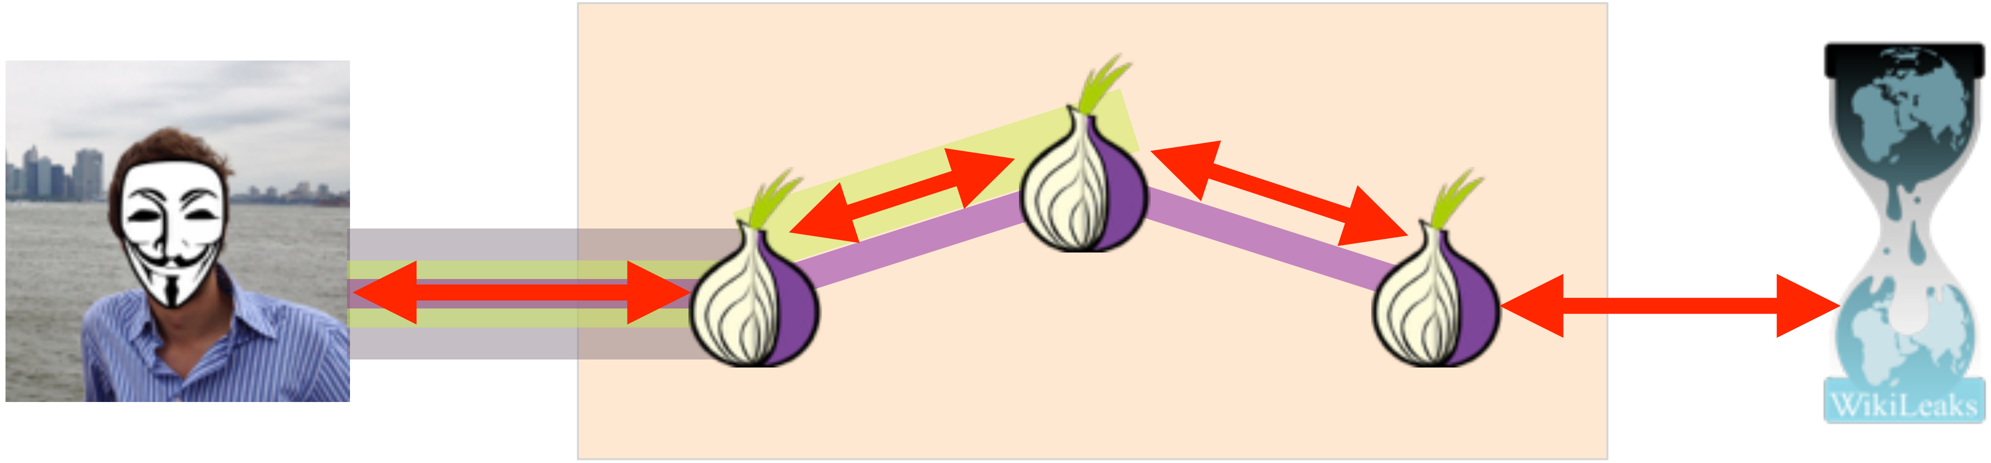

client

entry

middle

exit

server

knows the destination,
nor the source

Traffic entering and leaving Tor is highly correlated

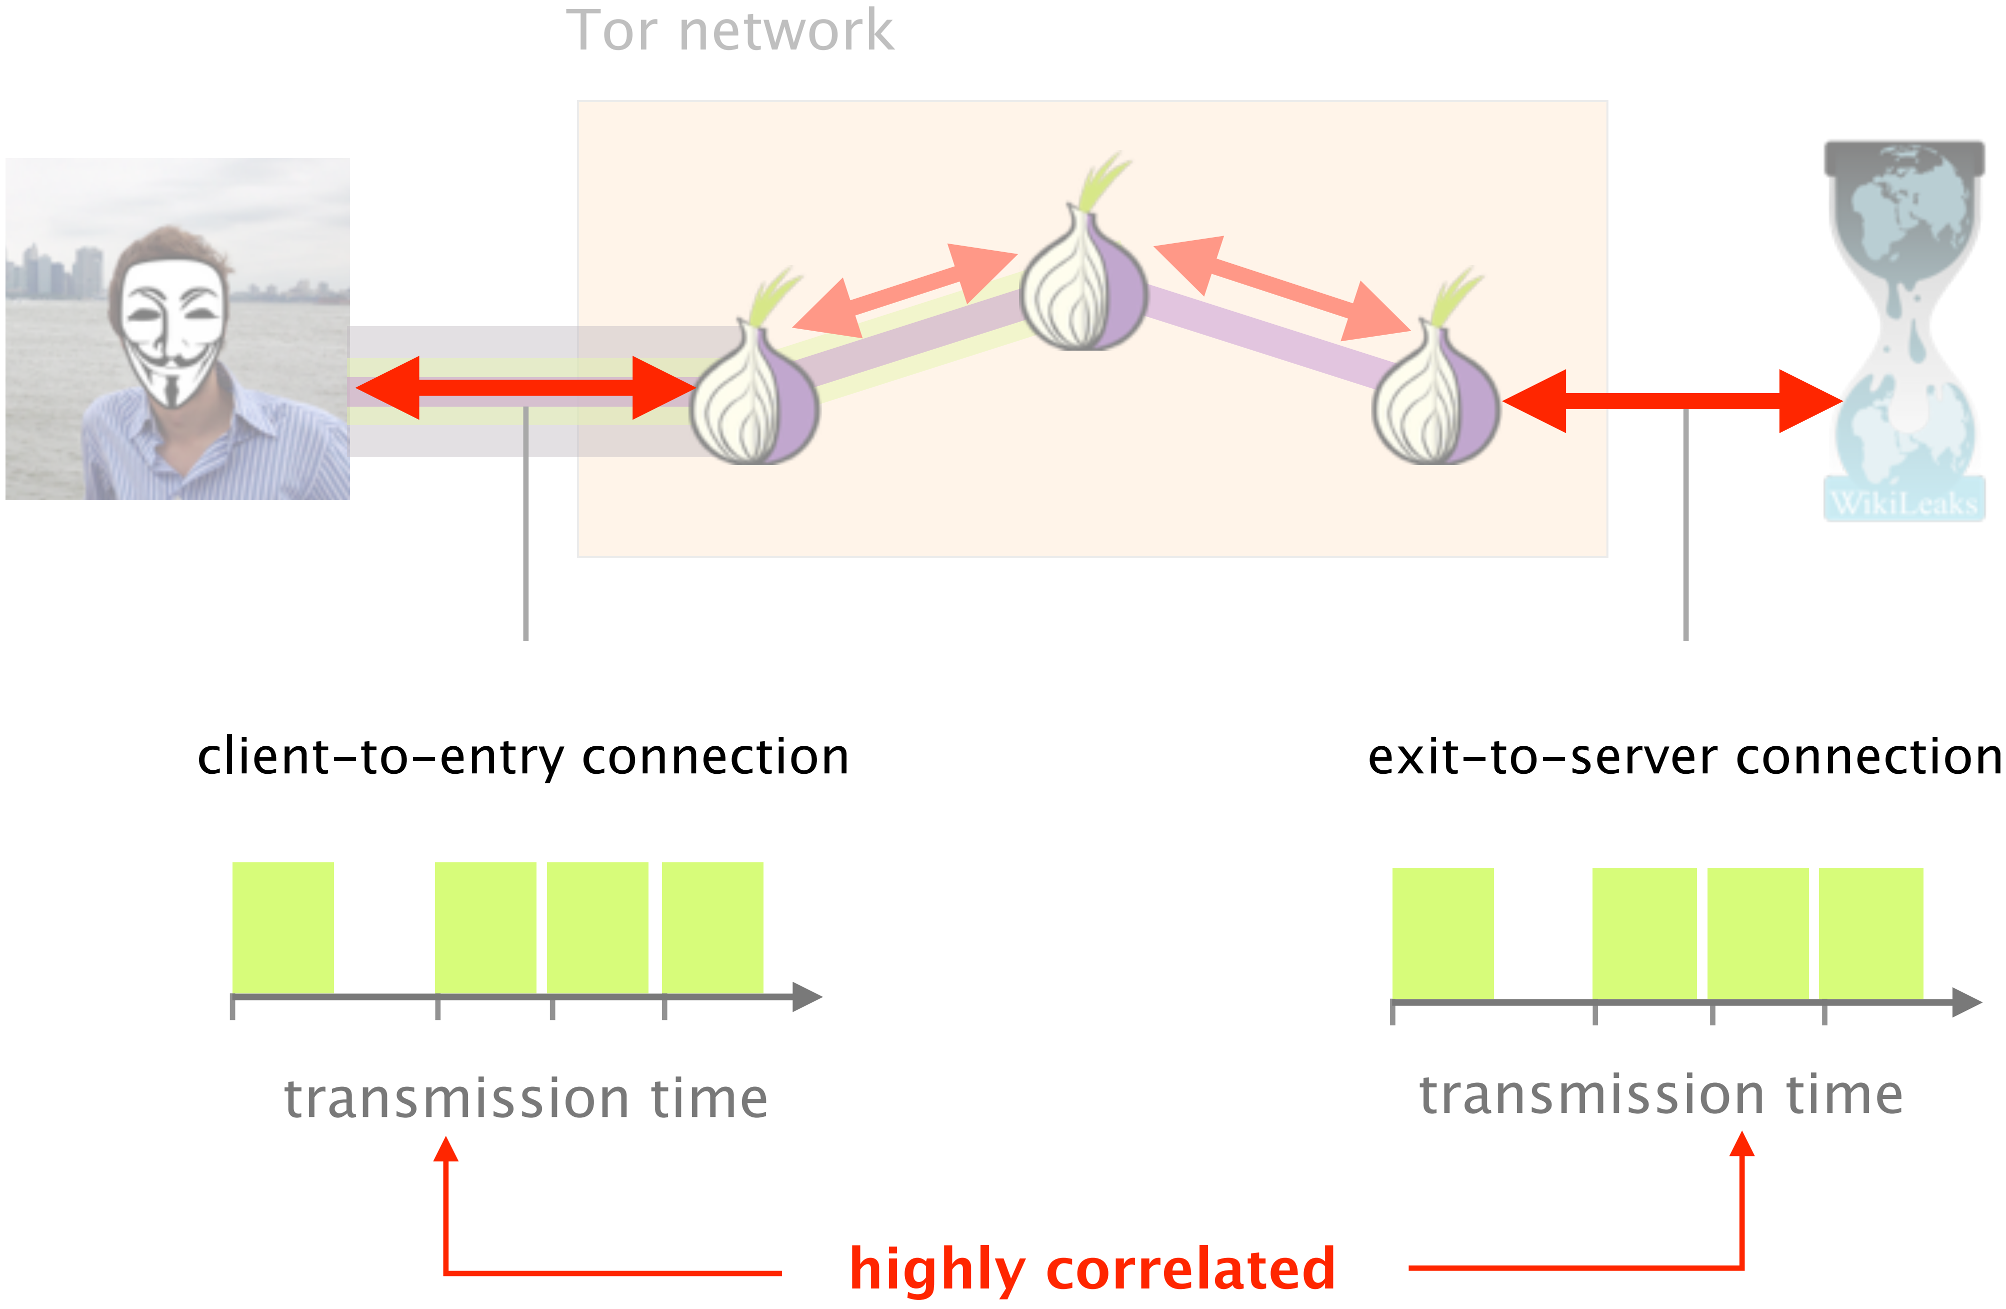

Tor network

client–to–entry connection

transmission time

exit–to–server connection

transmission time

highly correlated

By correlating client-to-entry & exit-to-server flows, one can effectively de-anonymize Tor users

Traffic correlation attacks require to see client-to-entry and exit-to-server traffic

Traffic correlation attacks require to see client-to-entry and exit-to-server traffic

How?

Two ways

Manipulate Tor

malicious relay

Manipulate routing

malicious networks

Two ways

Manipulate Tor

malicious relay

Manipulate routing

malicious networks

This talk

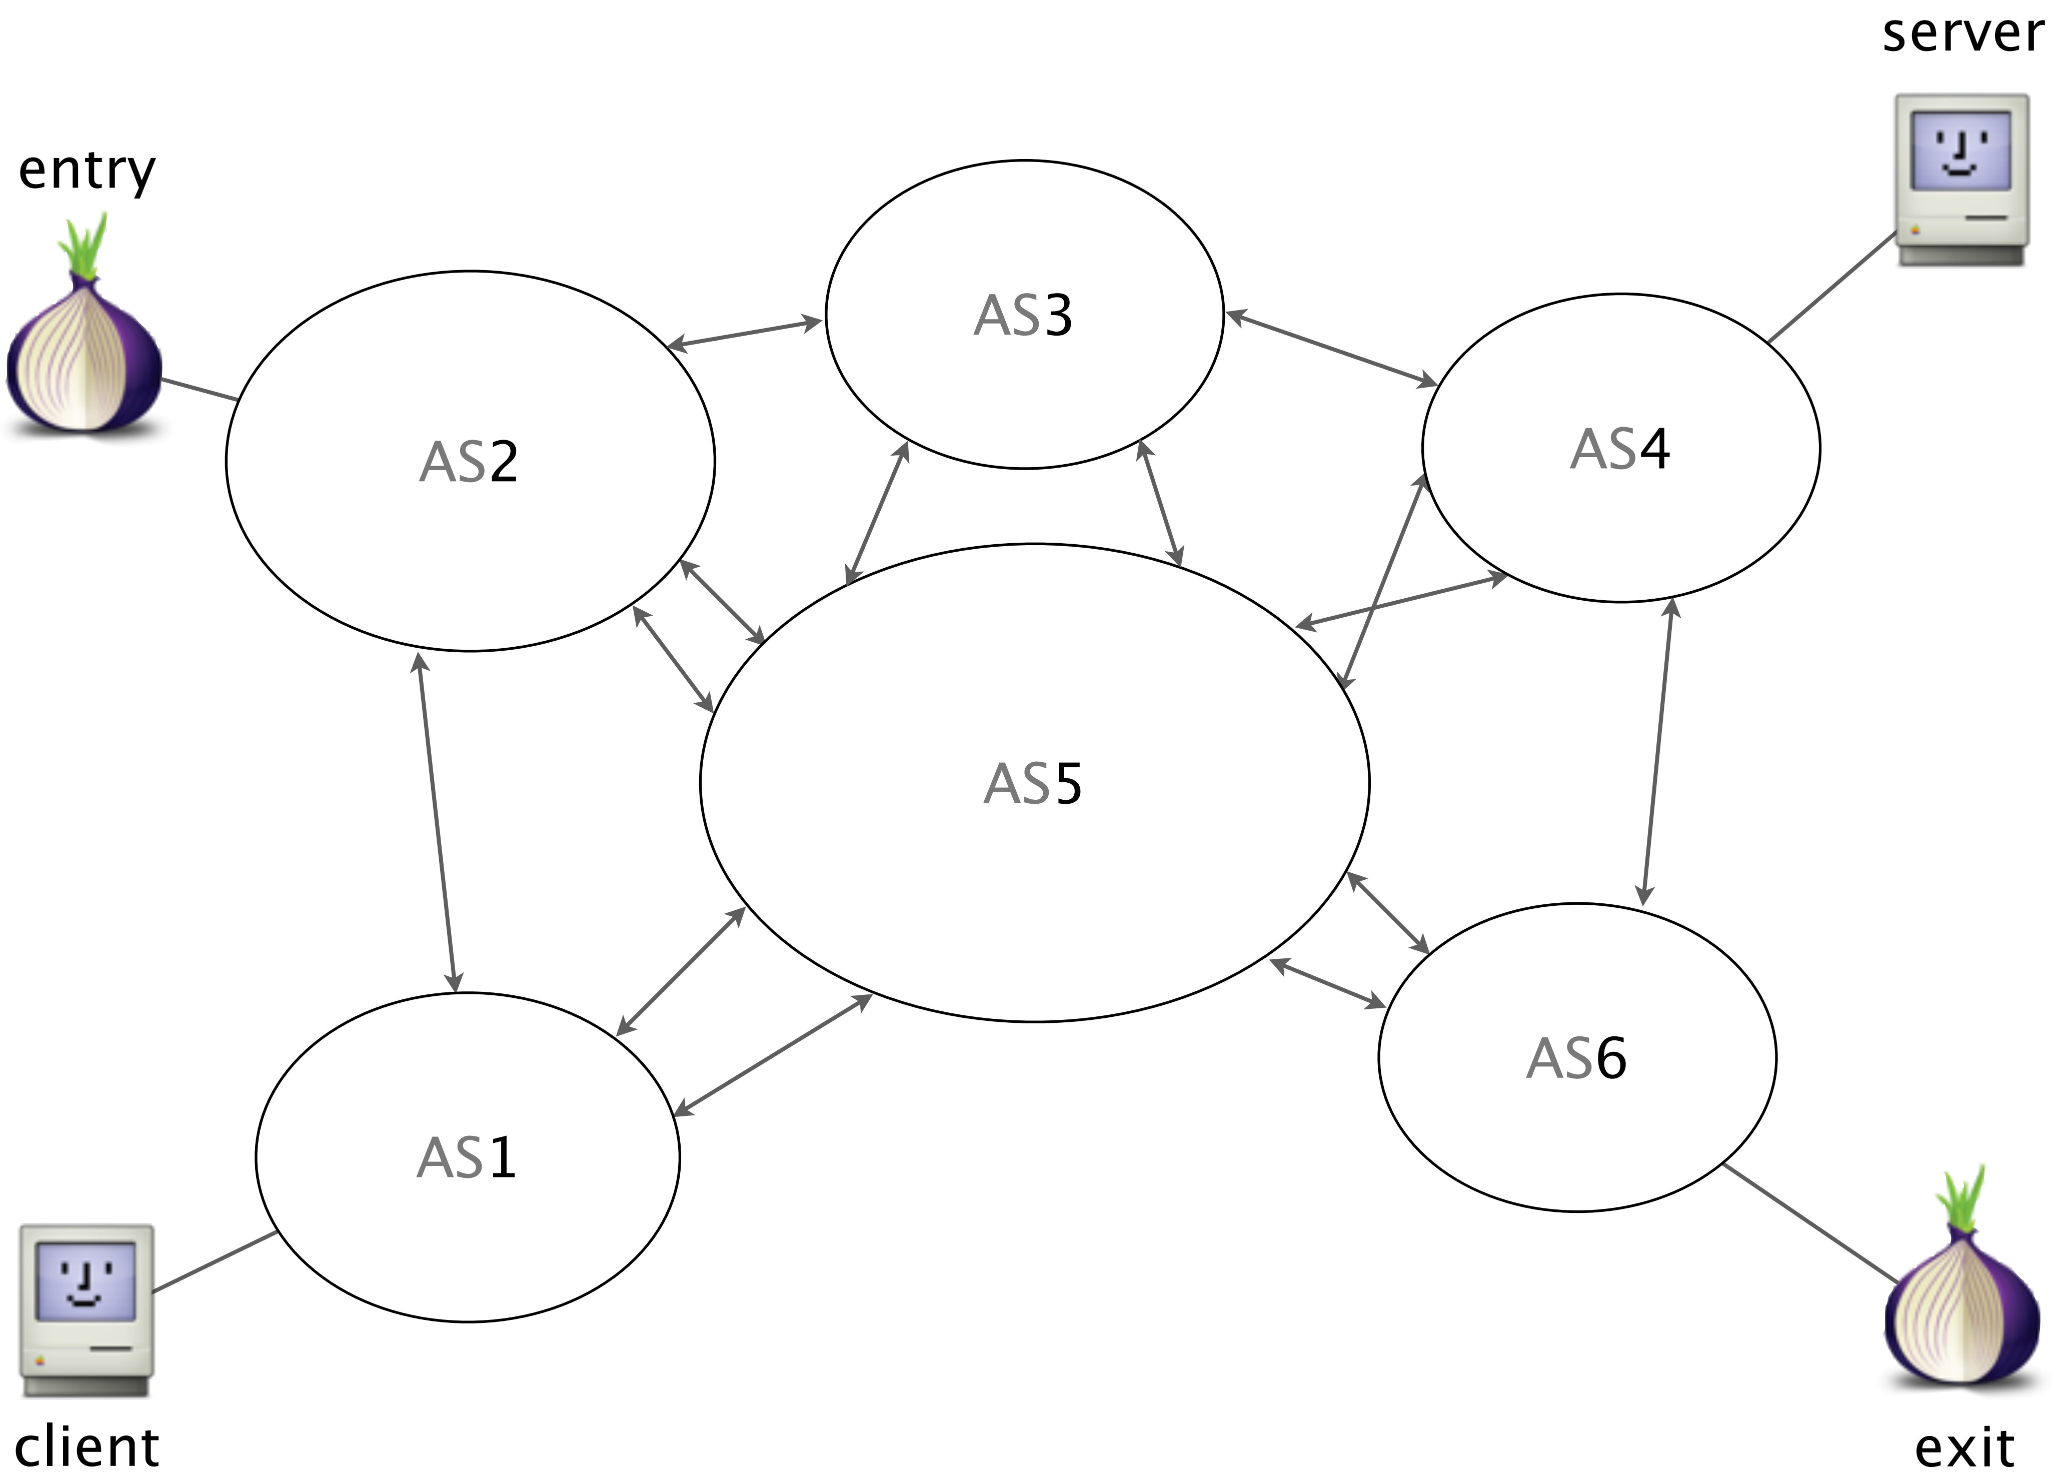
server
entry
AS3
AS2
AS4
AS5
AS6
AS1
client
exit

Tor connections get routed according to BGP

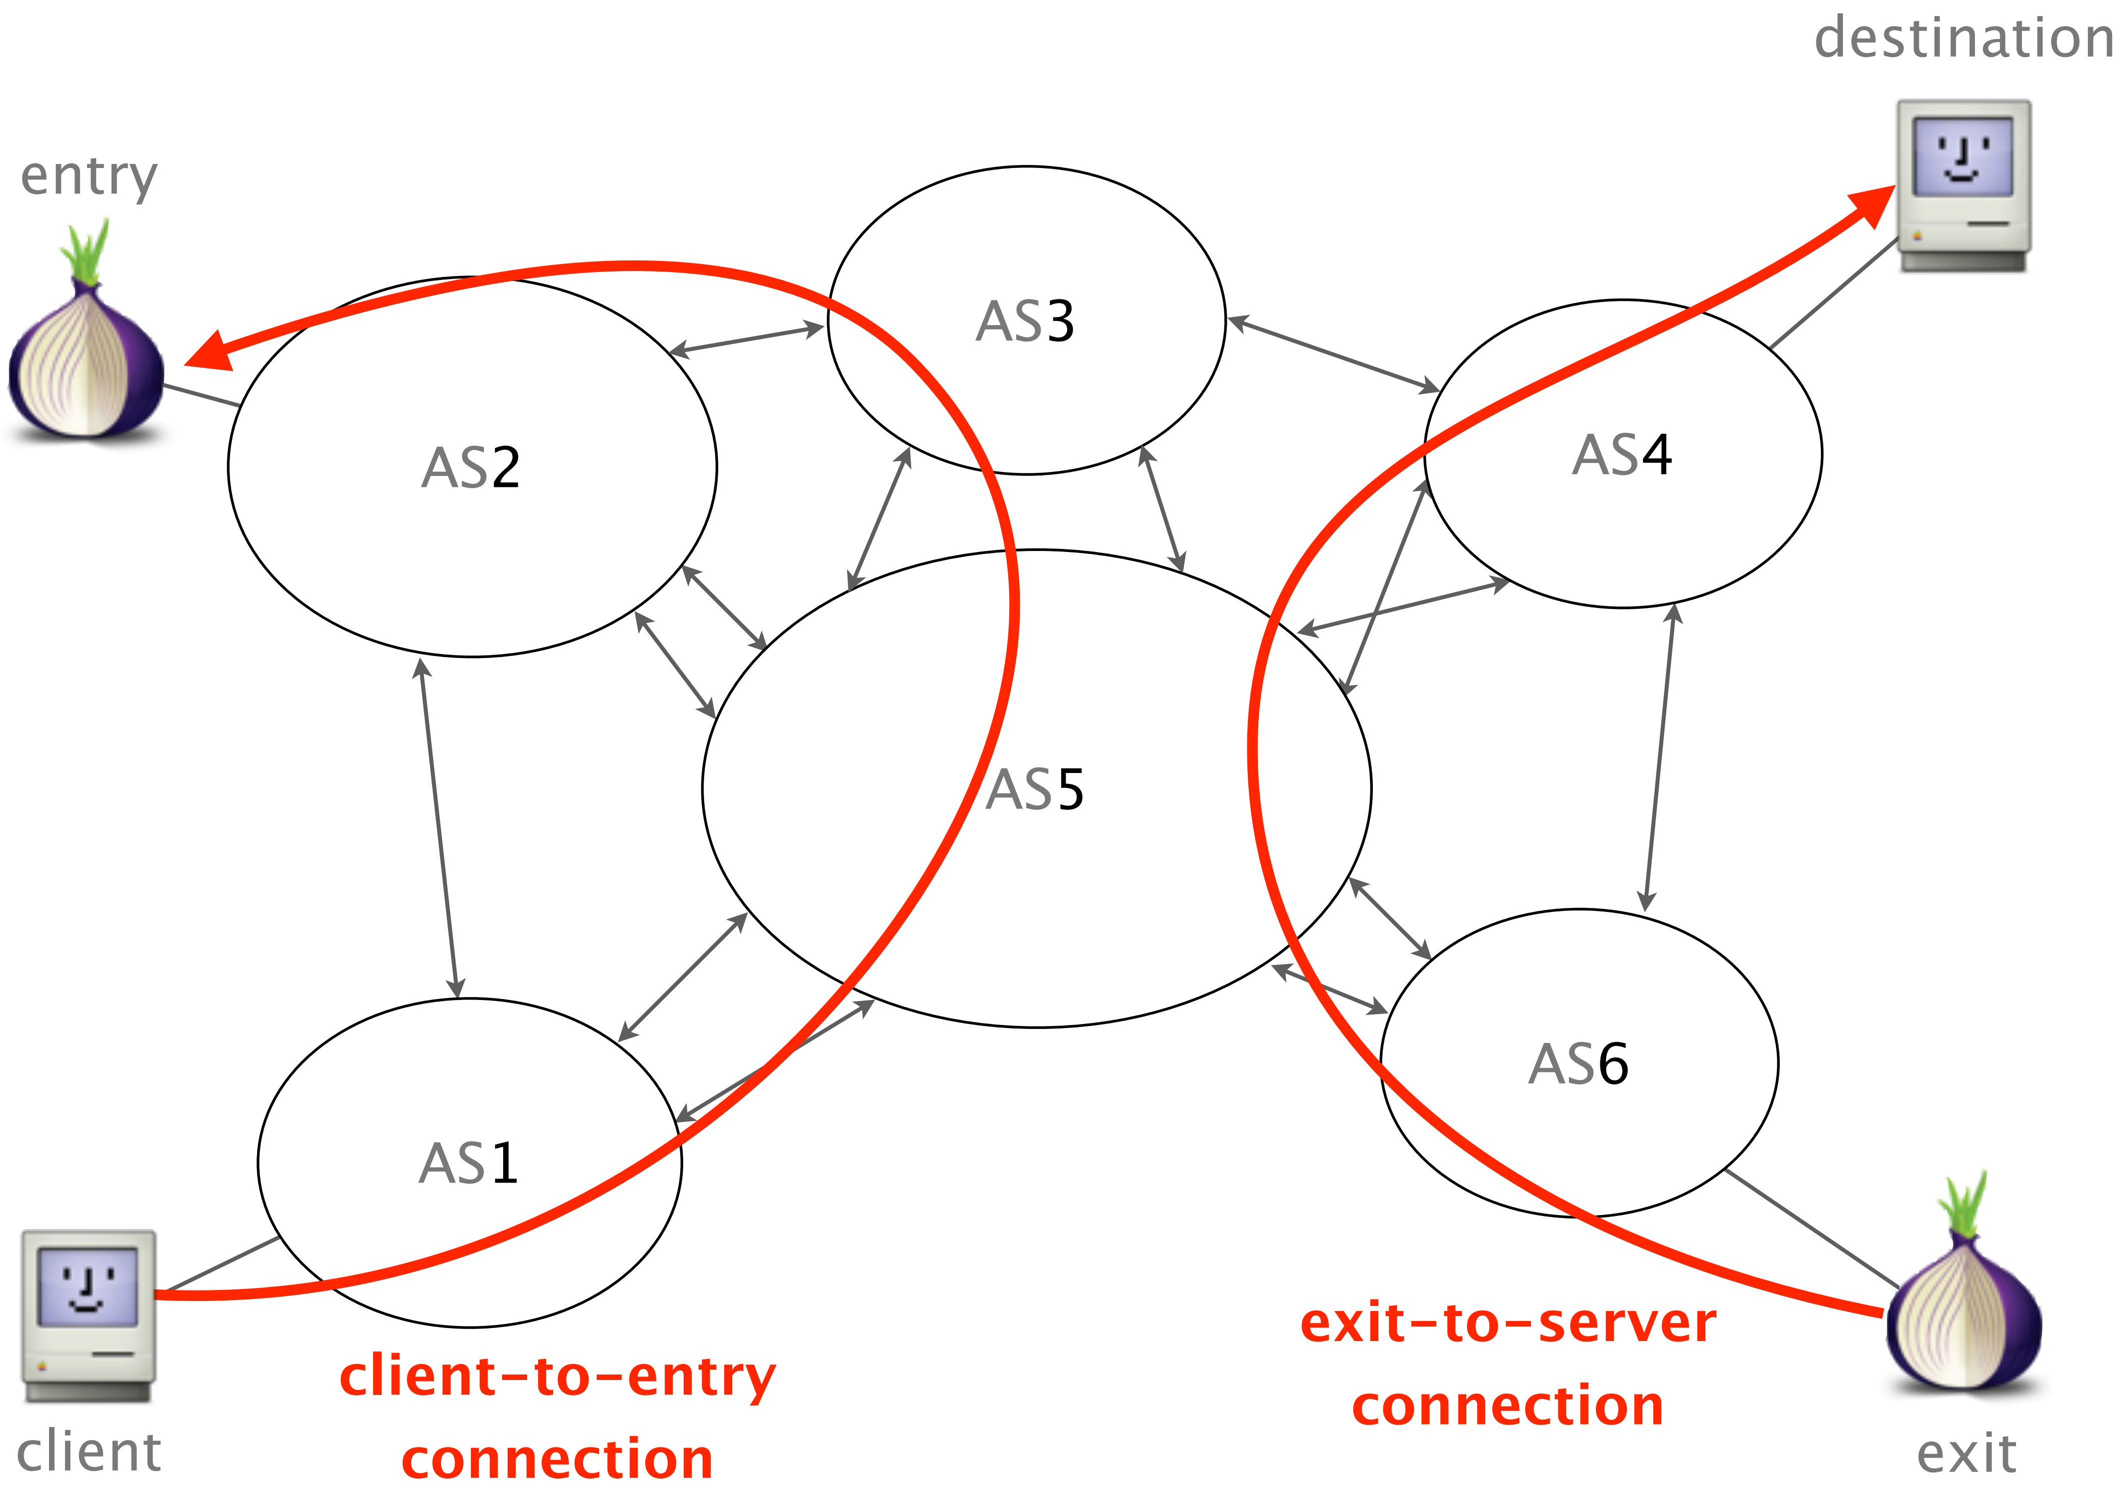

Traffic correlation attacks require to see client-to-entry *and* exit-to-server traffic

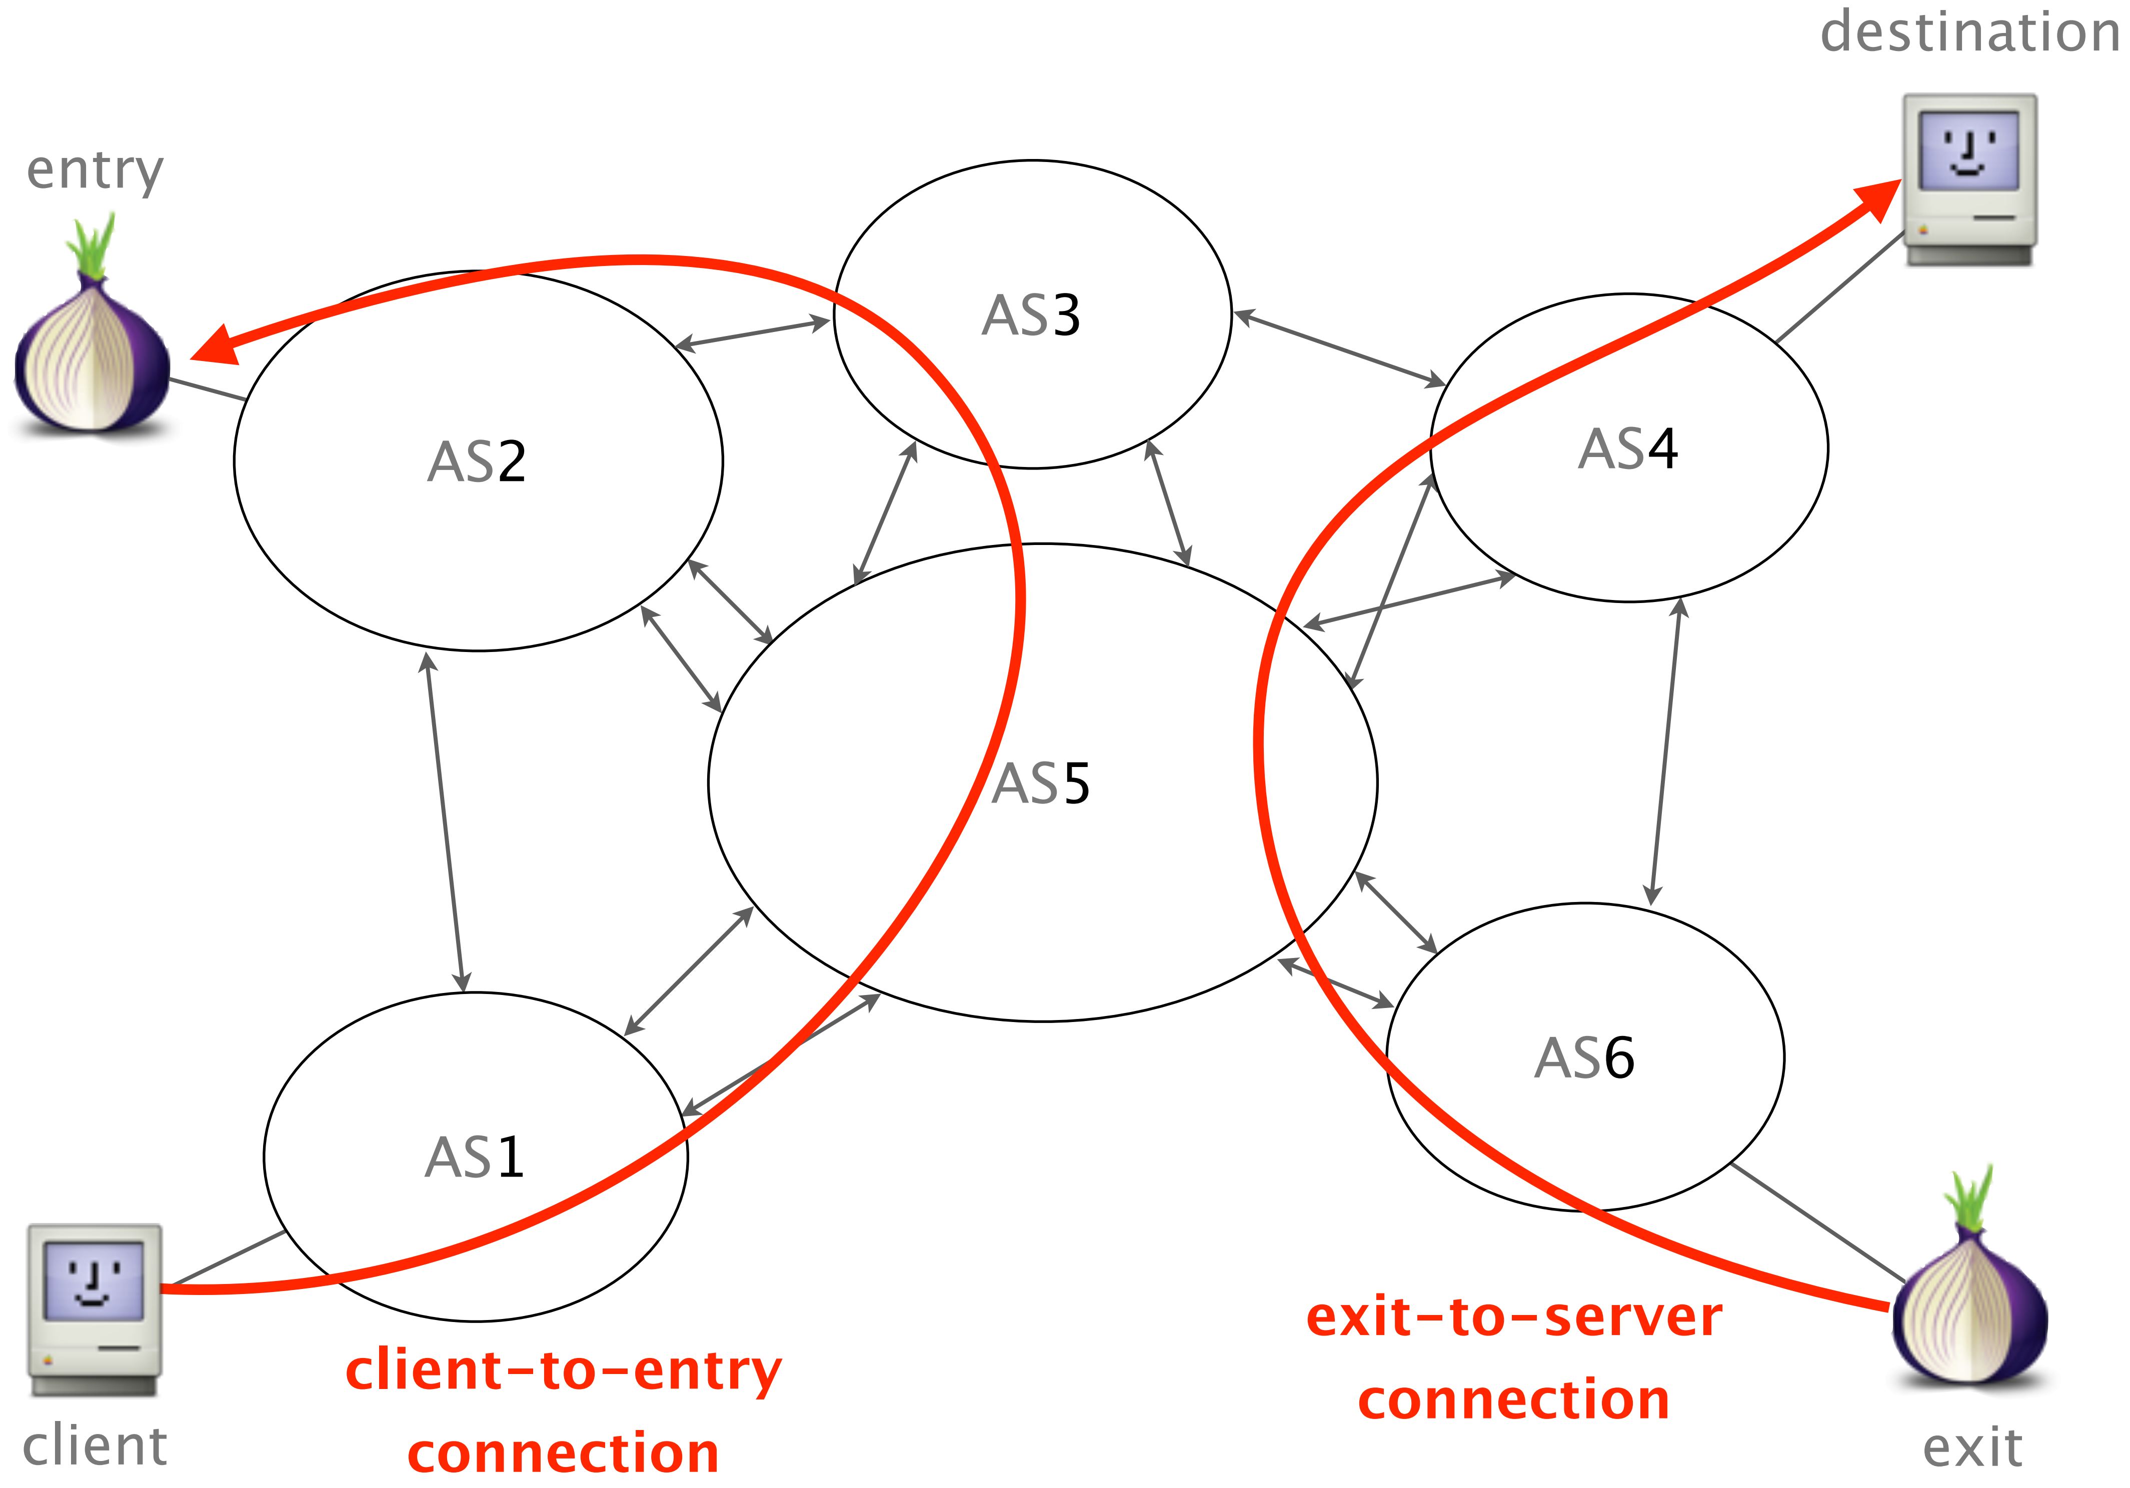

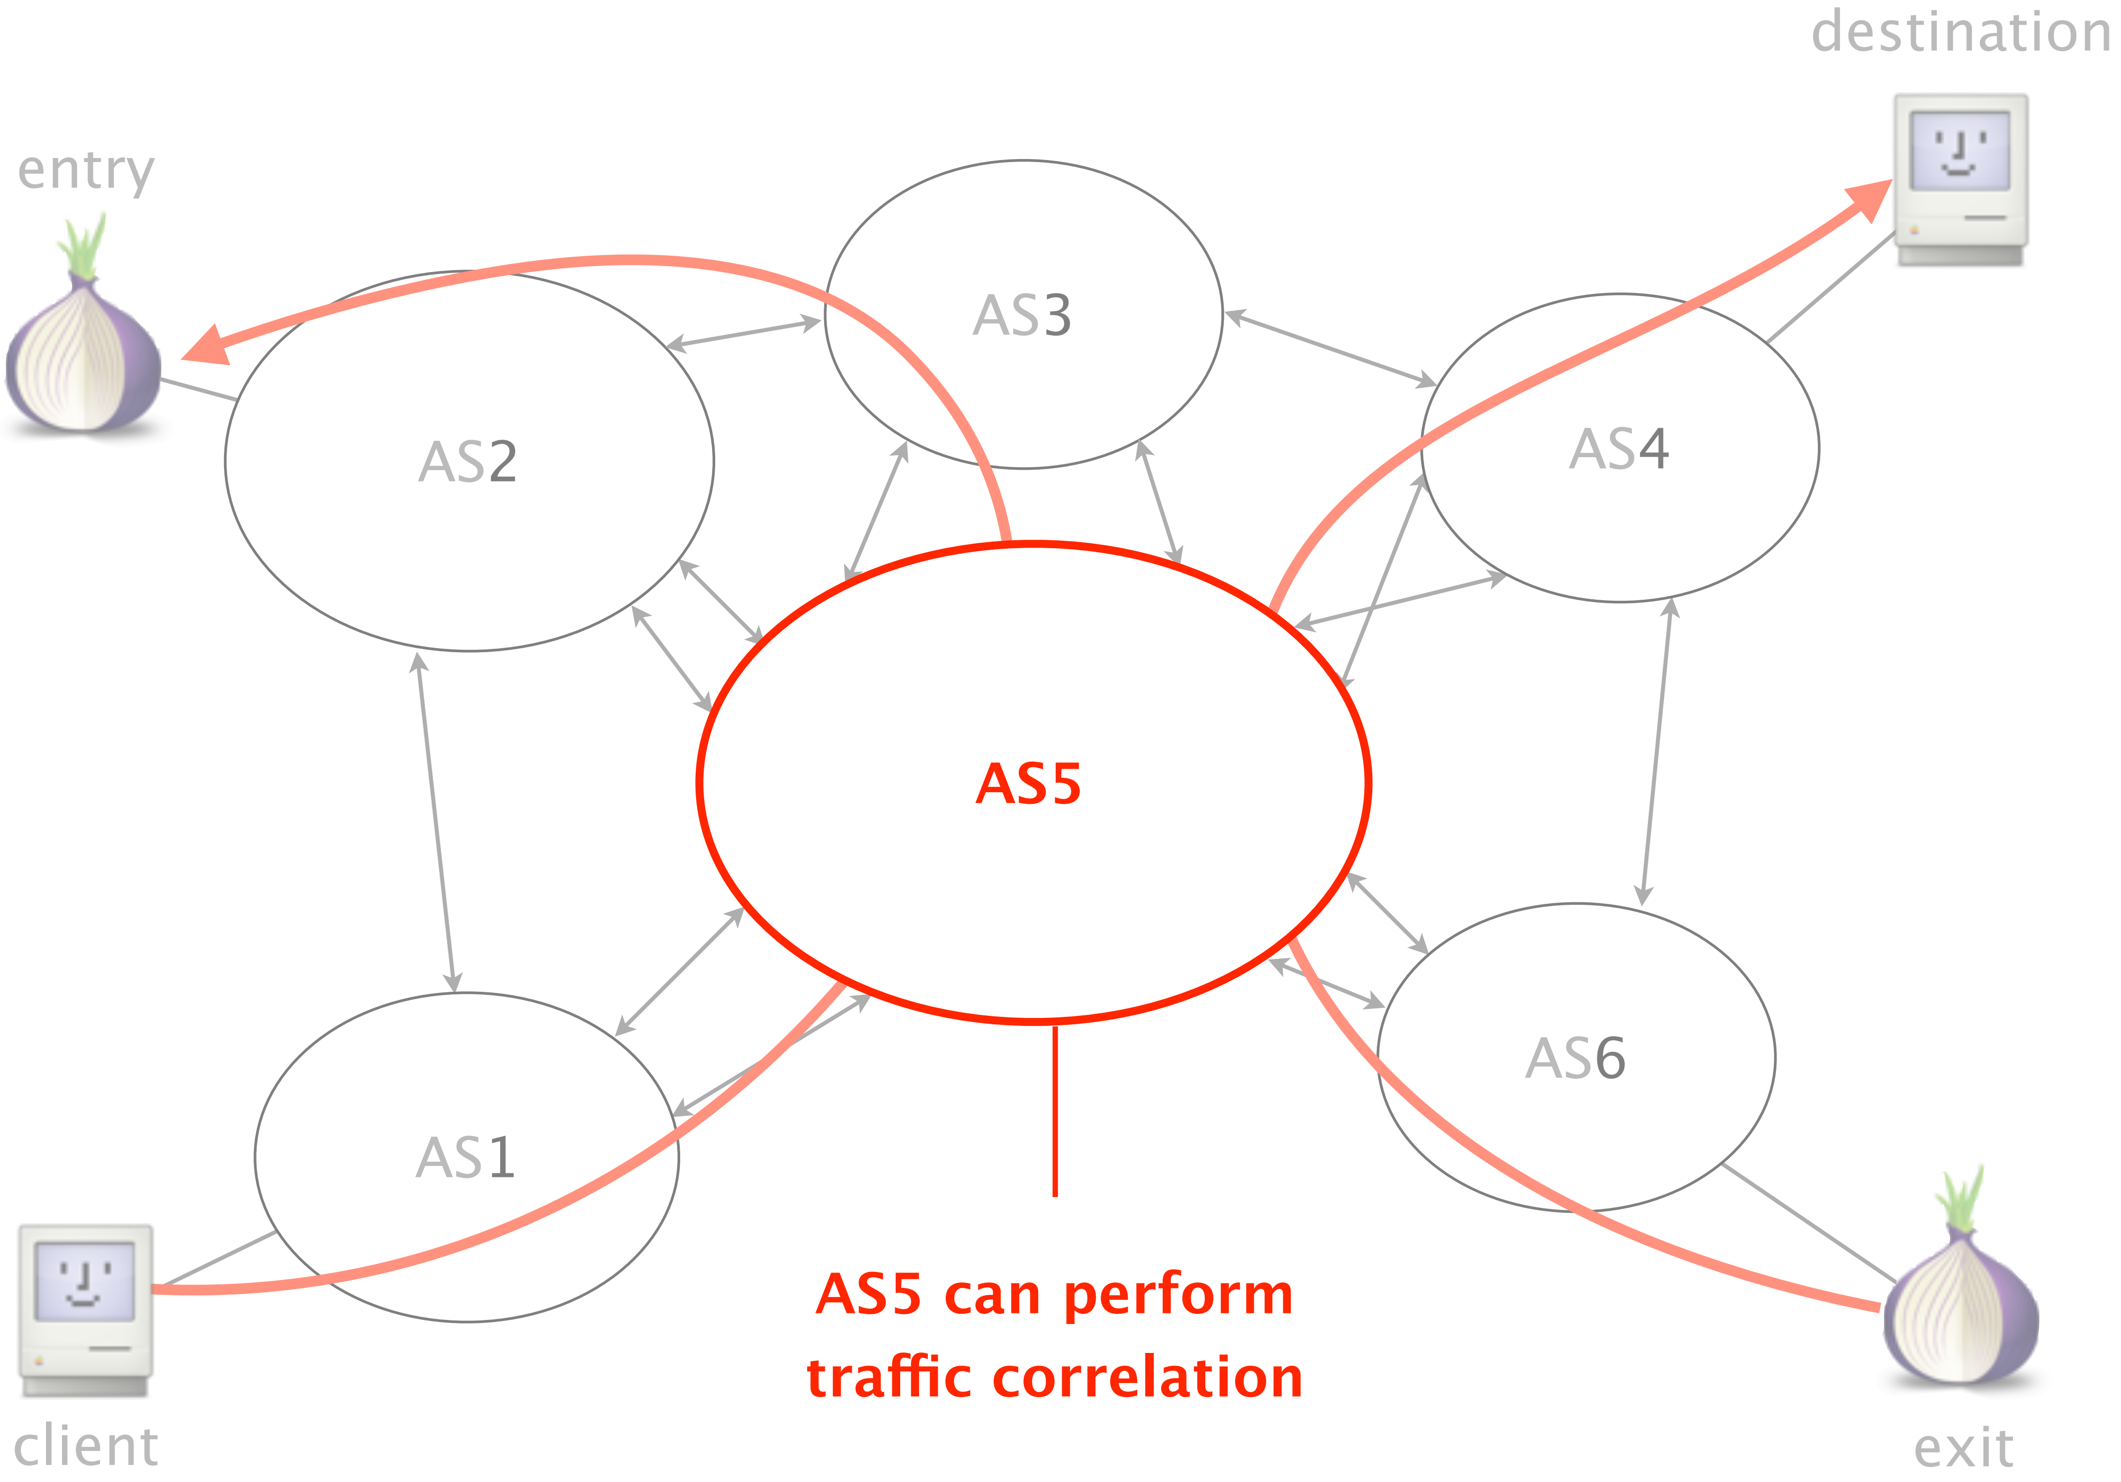

AS5 can perform traffic correlation

Network-level adversaries are a known problem

	Related work	
2004	Location diversity in anonymity networks	Feamster and Dingledine
2007	Sampled traffic analysis by Internet-exchange-level adversaries	Murdoch and Zieliński
2009	AS-awareness in Tor Path Selection	Edman and Syverson
2013	Traffic correlation on Tor by realistic adversaries	Johnson *et al.*

However, these works assume

that the Internet is static

However, these works assume
that the Internet is static

… which is **not** the case

However, these works assume
that the Internet is static

… which is **not** the case

Contribution

What's the impact on Tor?

User anonymity decreases over time due to BGP dynamics

User anonymity decreases over time due to BGP dynamics

3 BGP-induced causes

Natural BGP convergence

policy changes, failures, etc.

Active BGP manipulation

IP prefix hijack, interception (MITM), etc.

Asymmetric routing

seeing one direction of the connection is enough

Anonymity on Quicksand

Using BGP to compromise Tor

1 **Attacks**

All your traffic belongs to me

2 **Preliminary results**

Eyes wide open

3 **Countermeasures**

Close the curtains

Anonymity on Quicksand

Using BGP to compromise Tor

1 Attacks

All your traffic belongs to me

Preliminary results

Eyes wide open

Countermeasures

Close the curtains

Attack#1: Natural BGP dynamics increases the number of AS-level adversaries

Initially, only AS5 is seeing traffic

client-to-entry and exit-to-server traffic

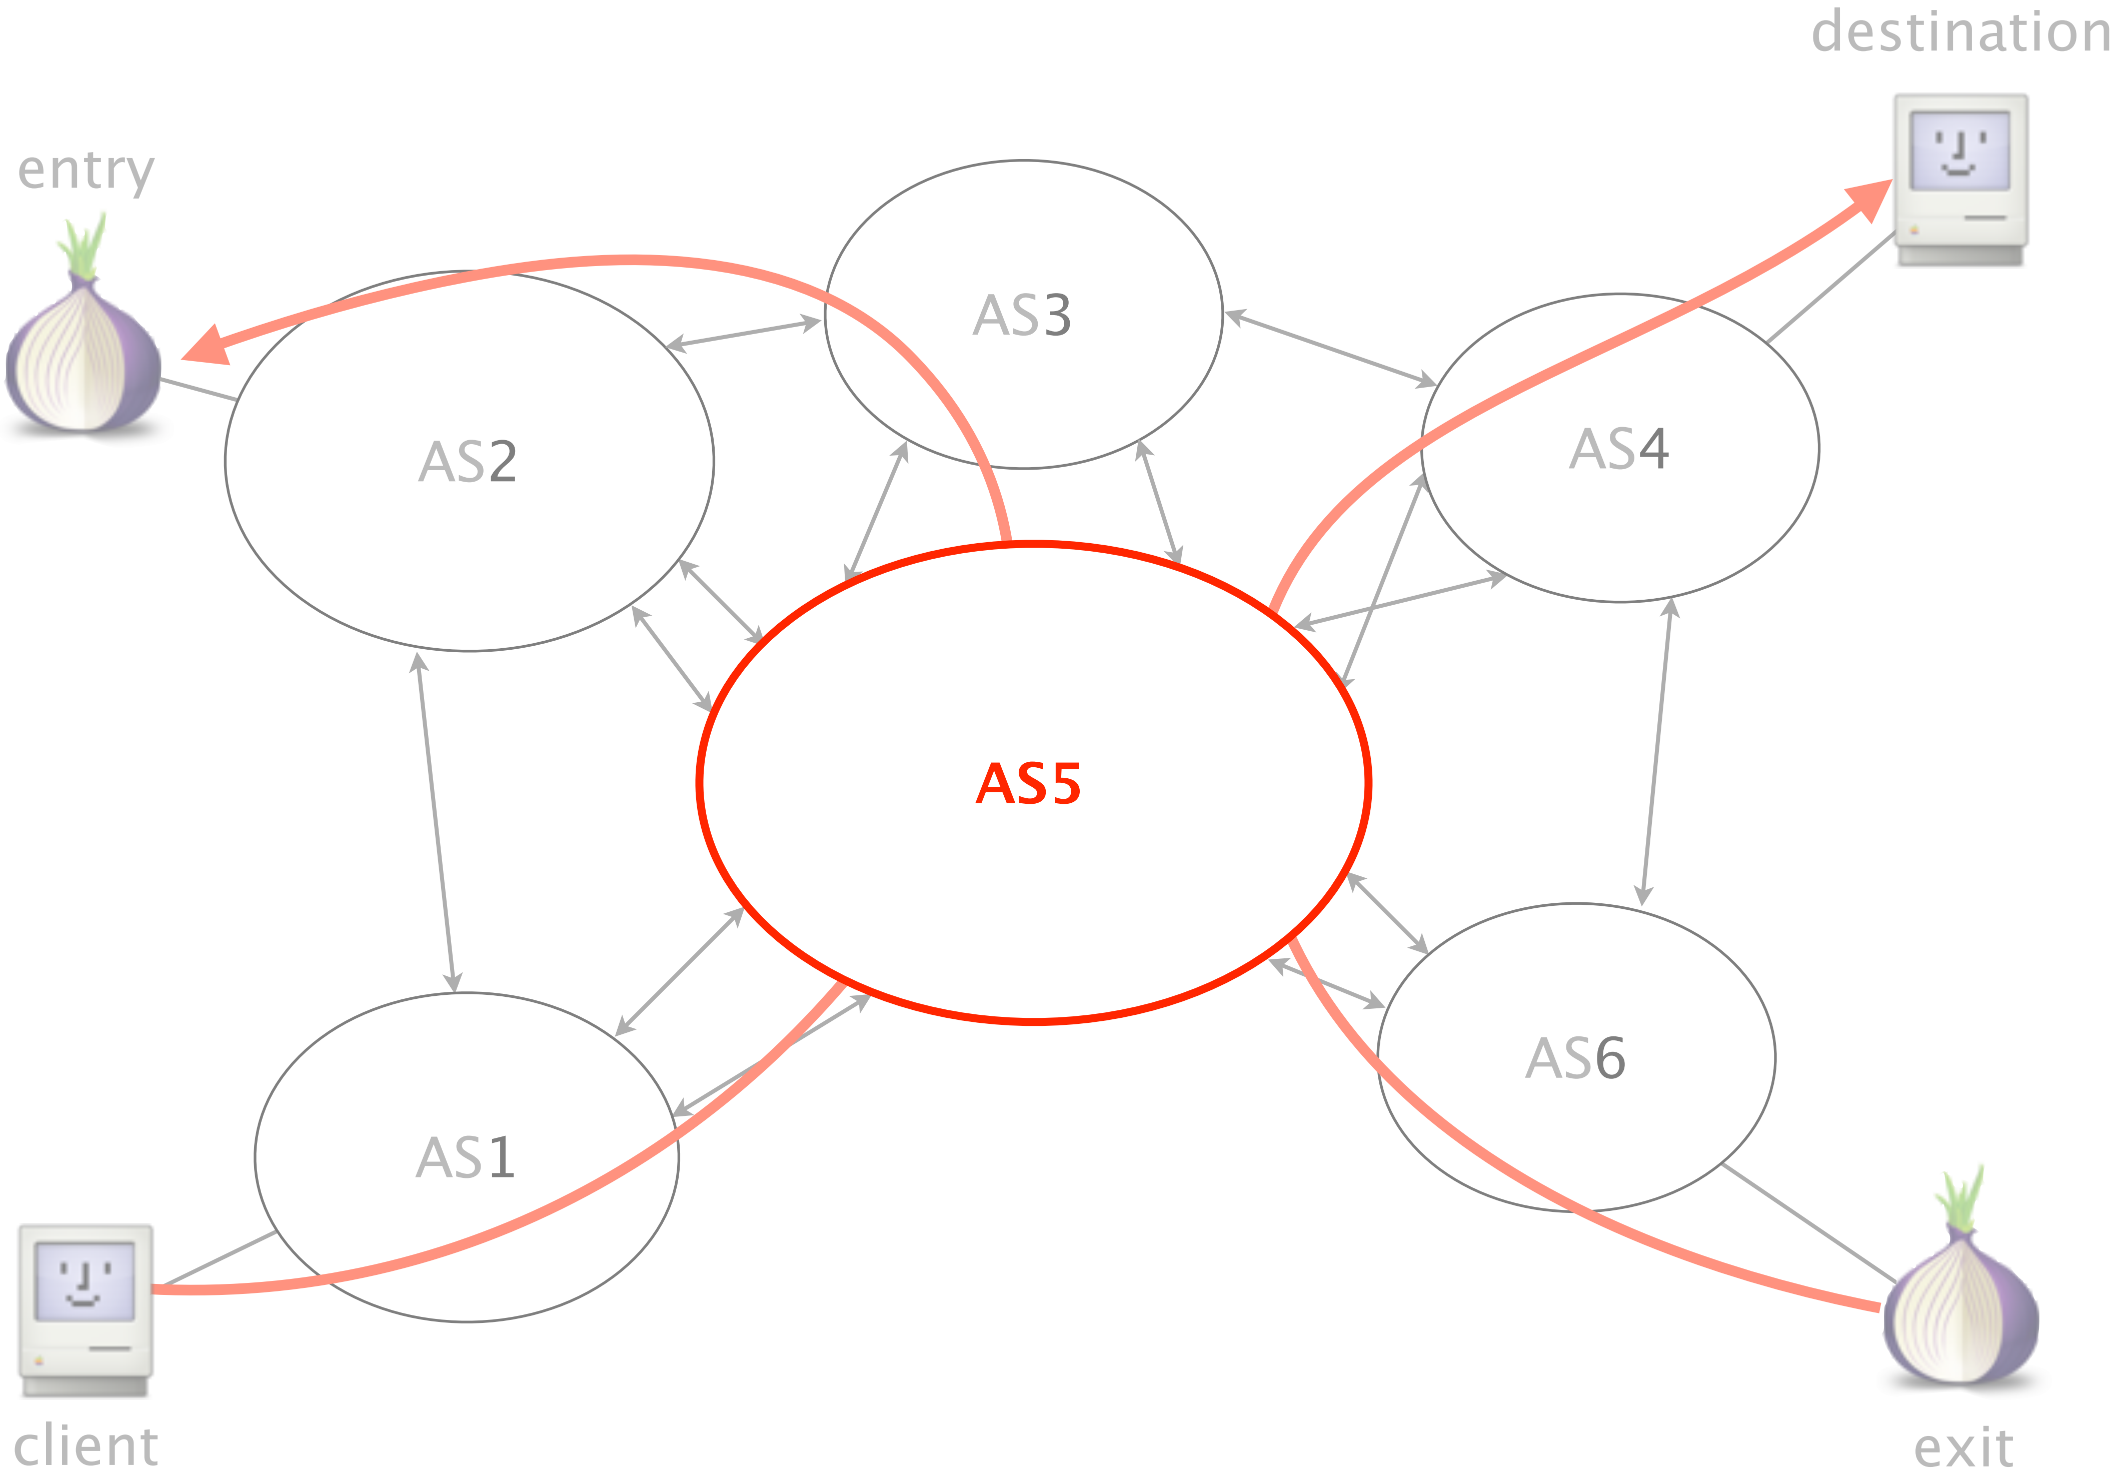

Link between AS4 and AS5 fails

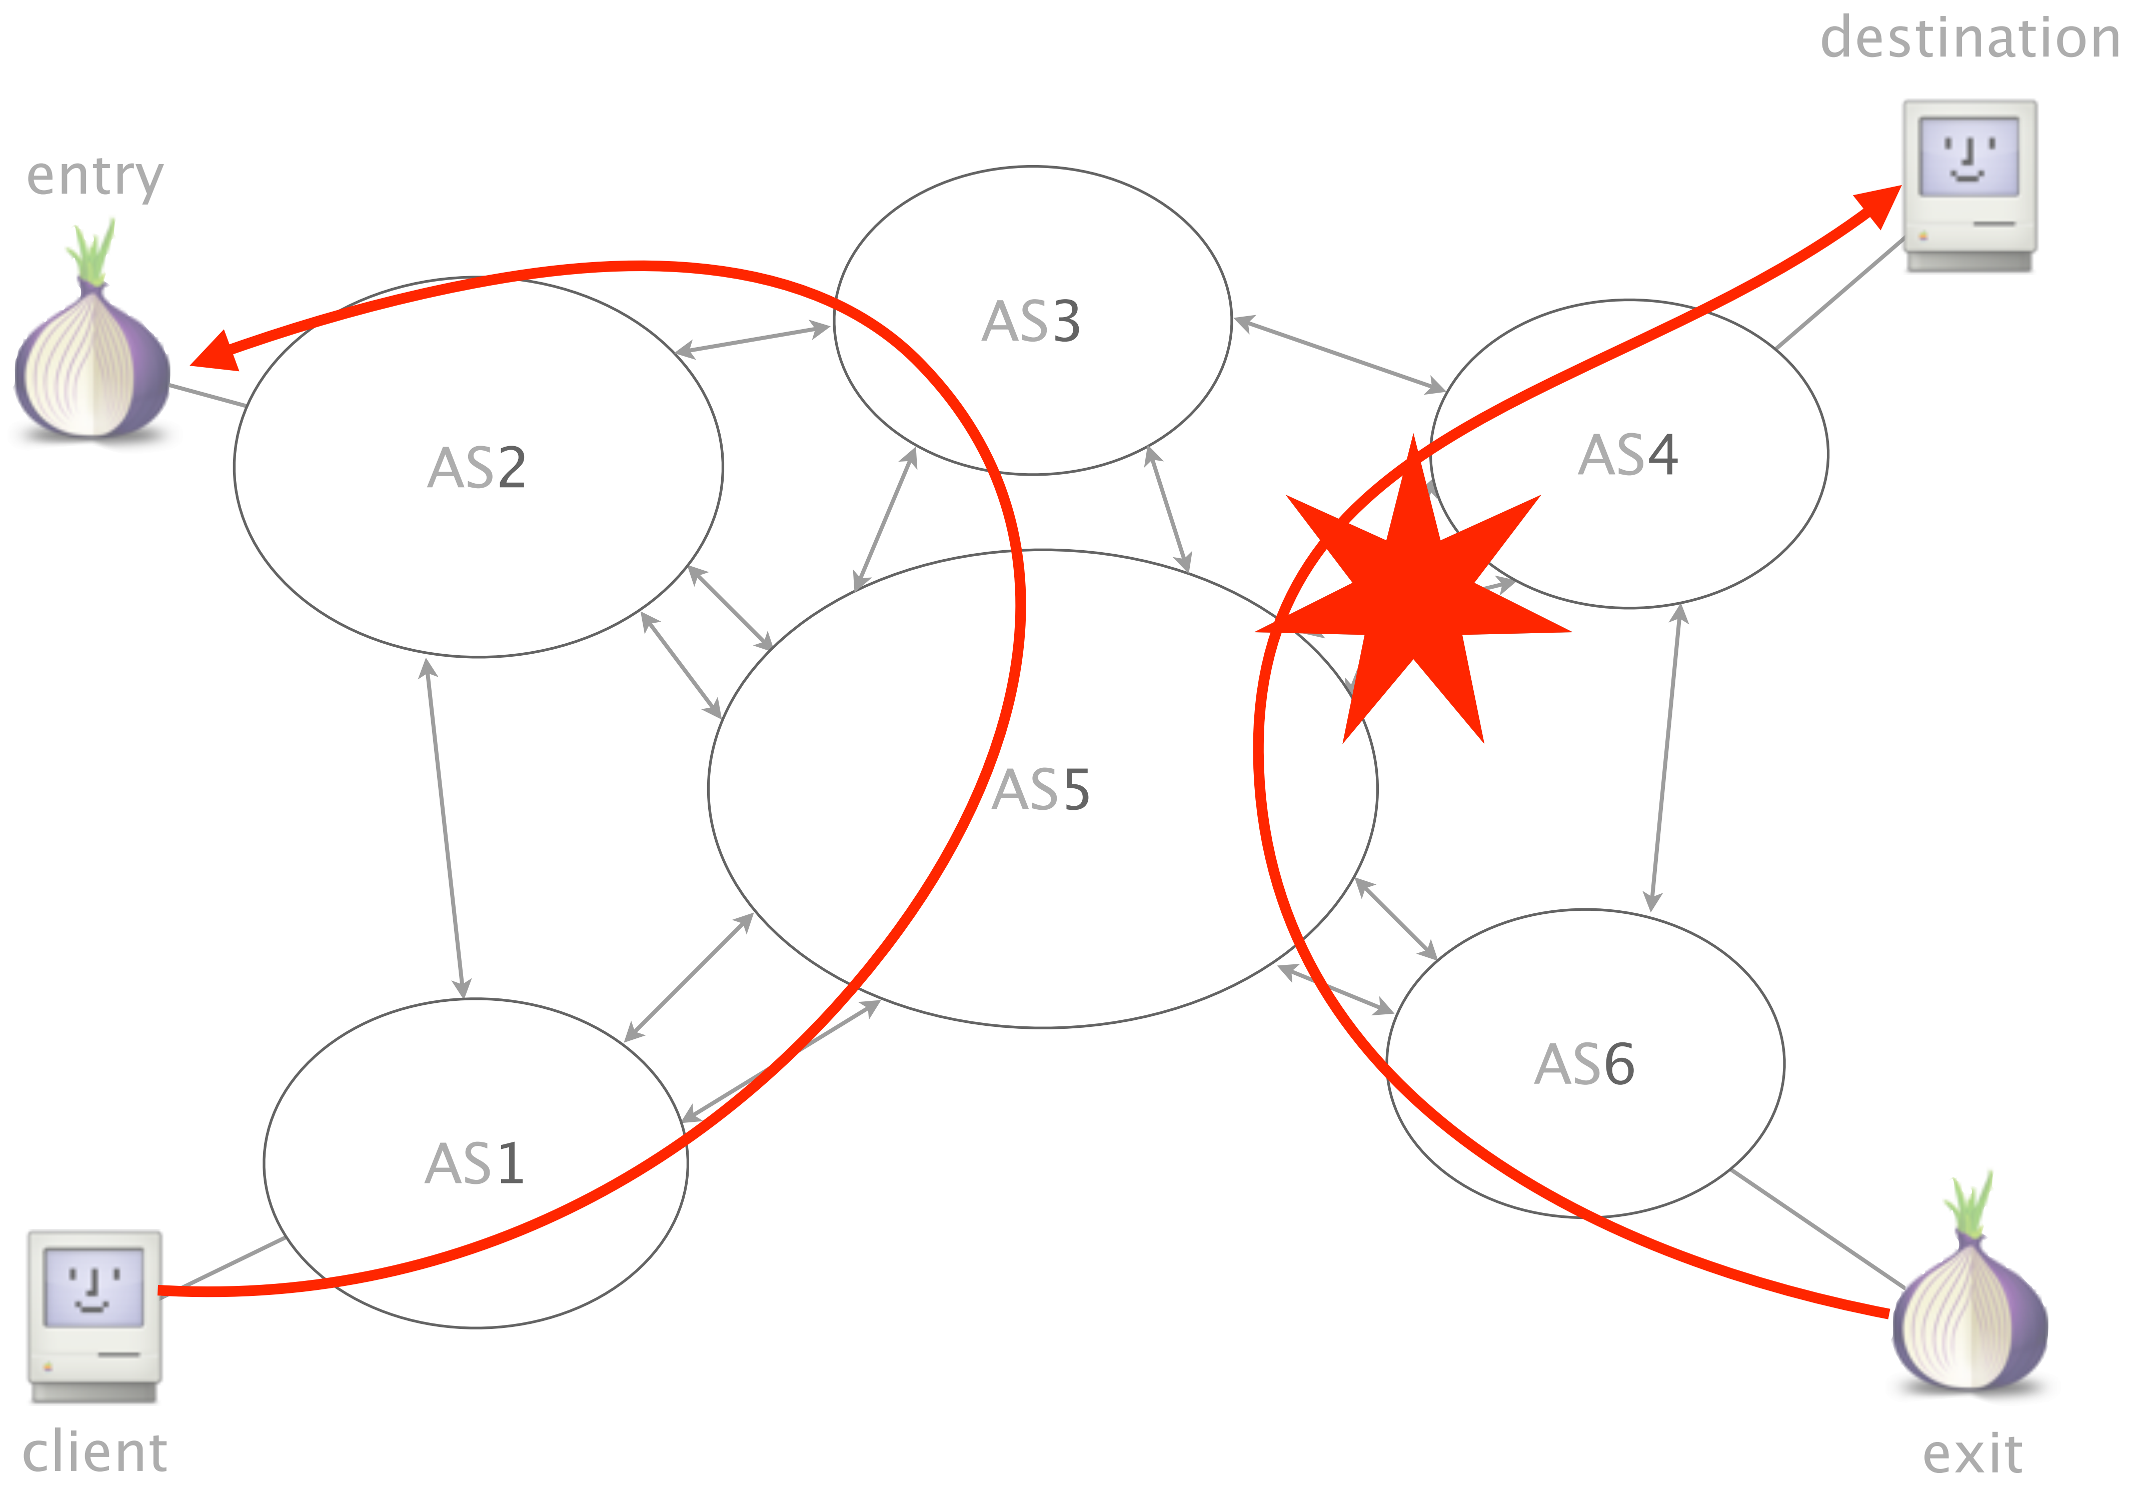

Traffic gets rerouted via AS3

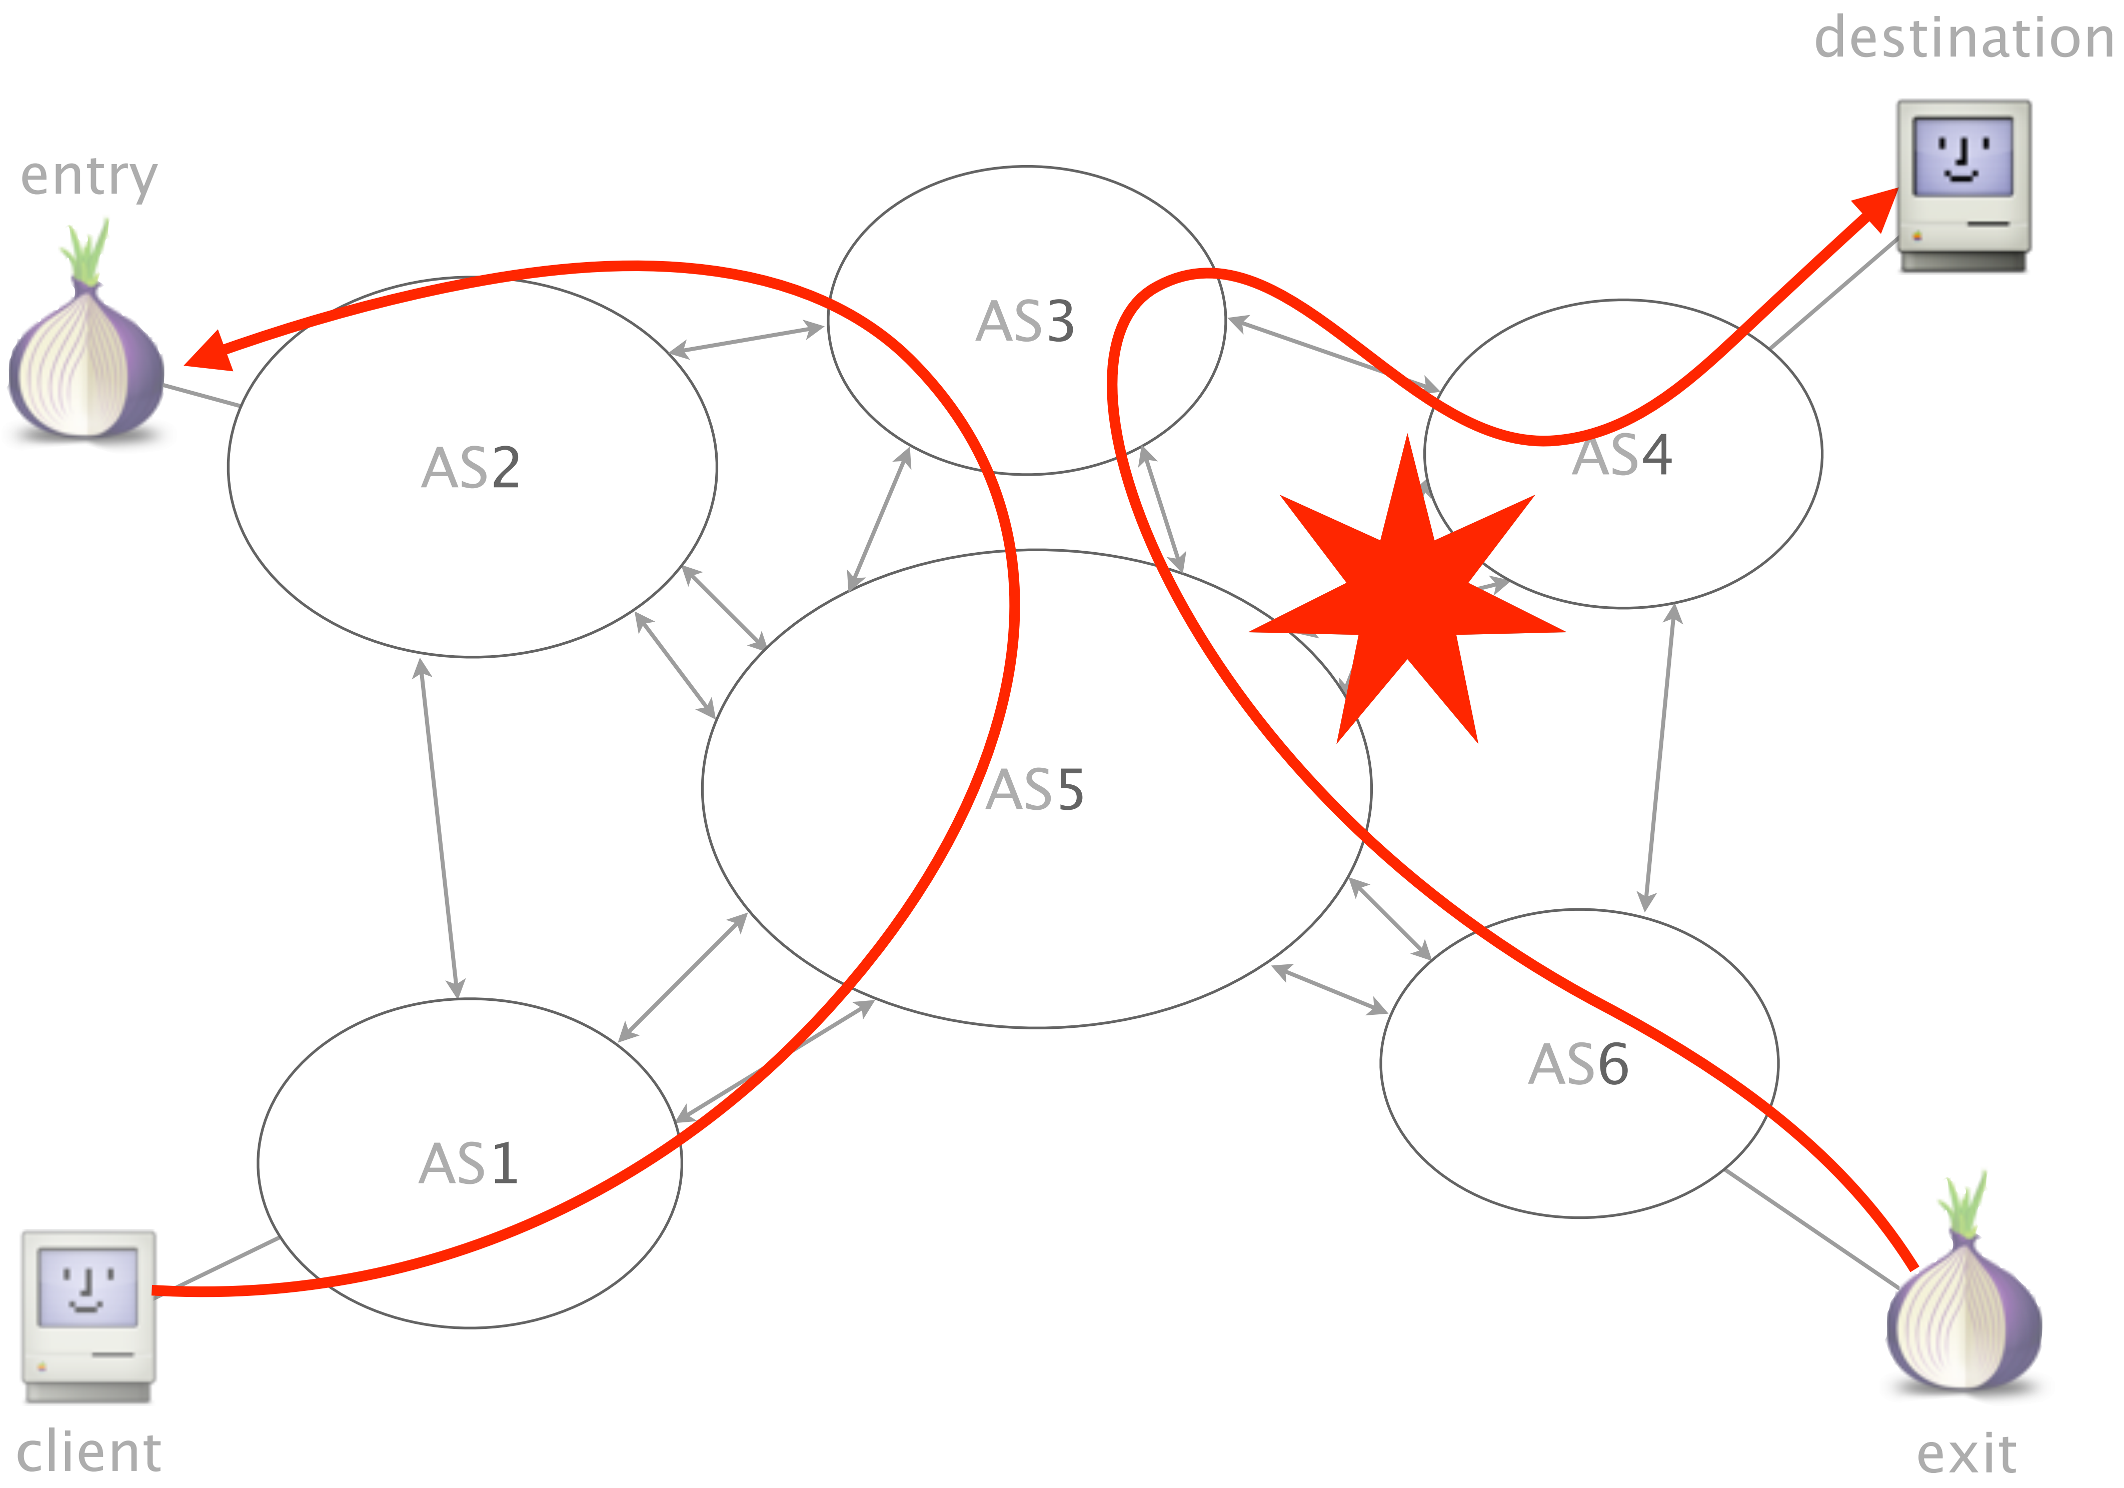

Now, both AS3 and AS5 are seeing client-to-entry and exit-to-server traffic

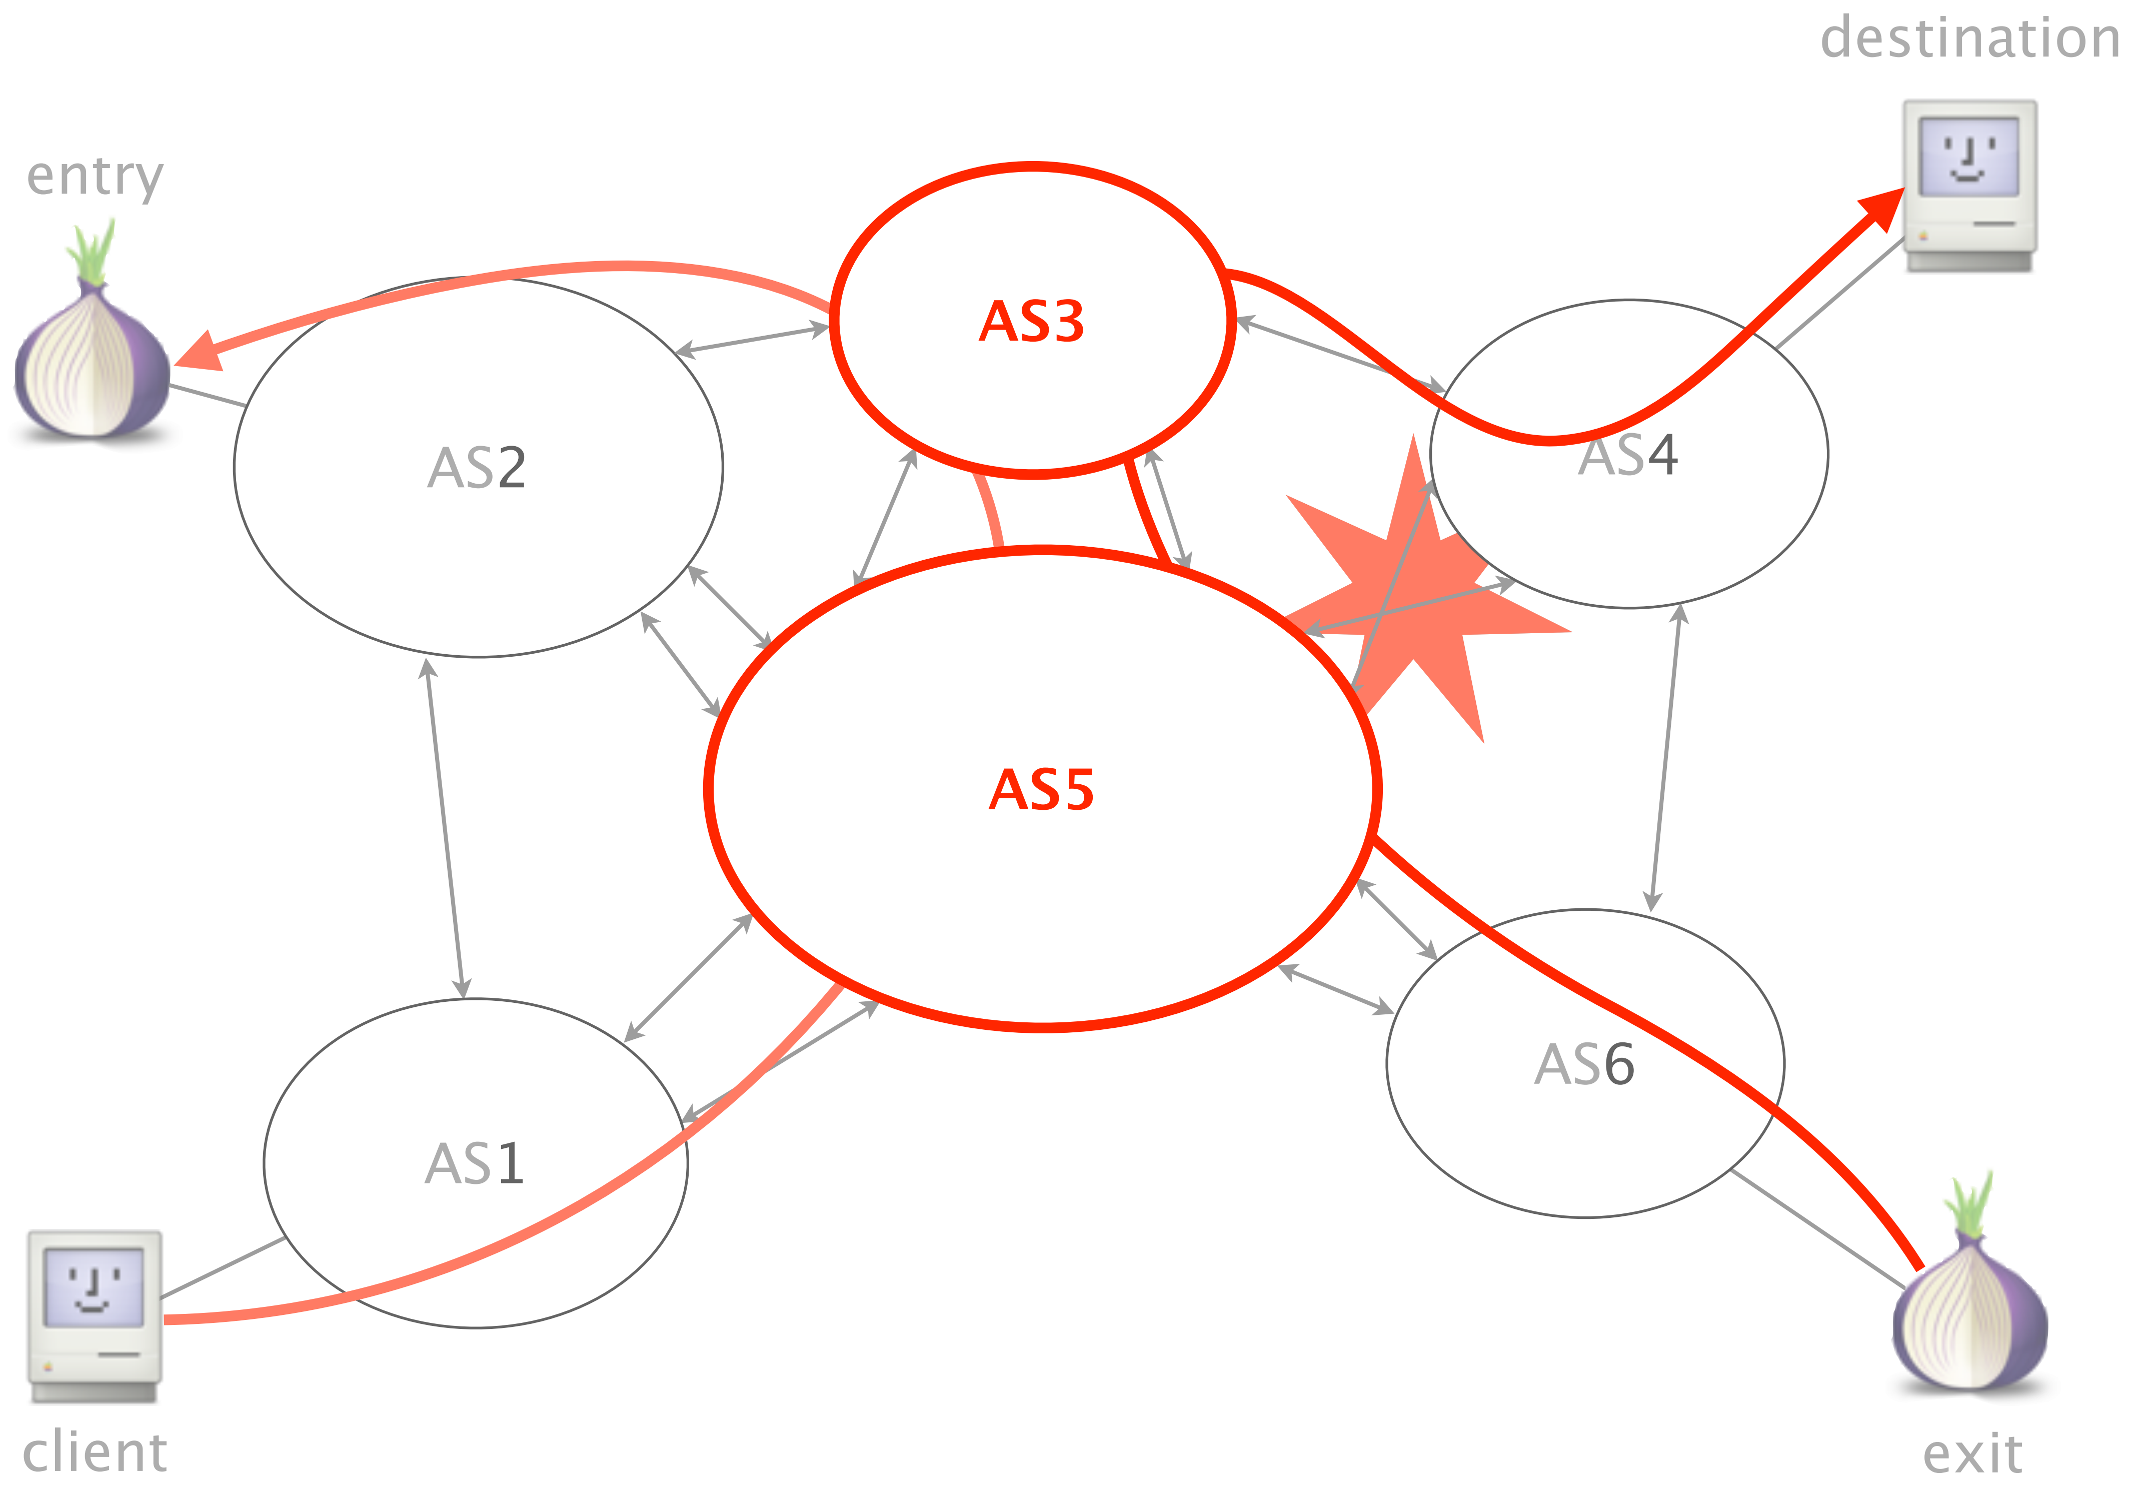

Attack#2: BGP hijacking attacks enable on-demand, fine-grained Tor attacks

Initially, only AS5 is seeing traffic entering and exiting the Tor network

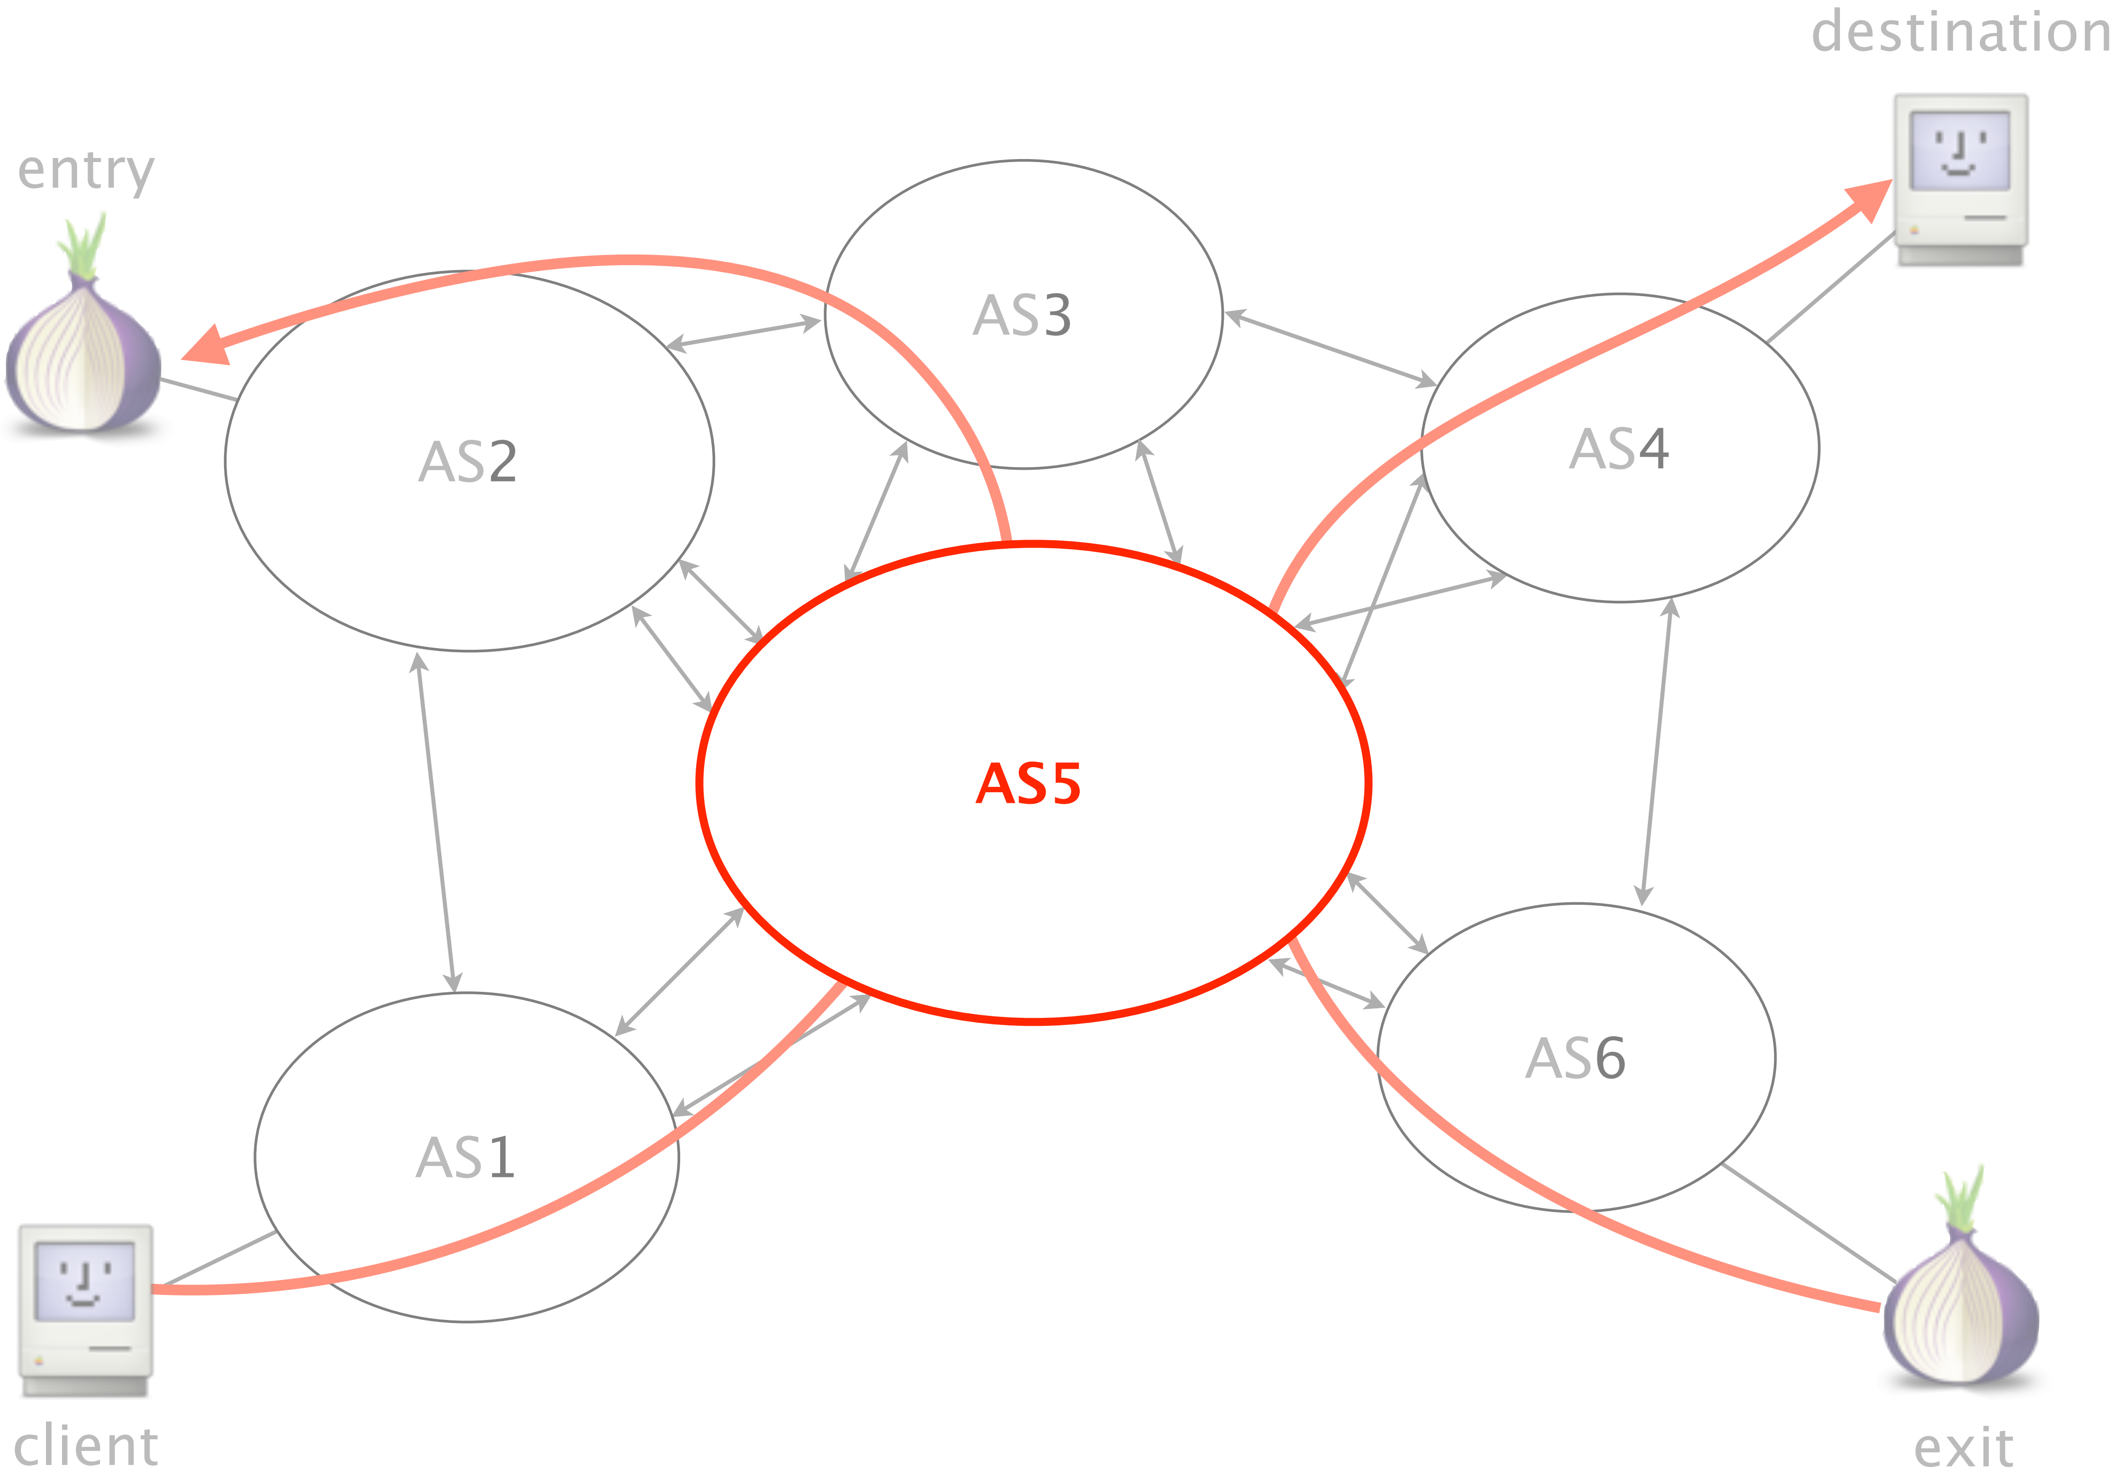

Assume that AS3 is a malicious AS, and wants to observe Tor traffic

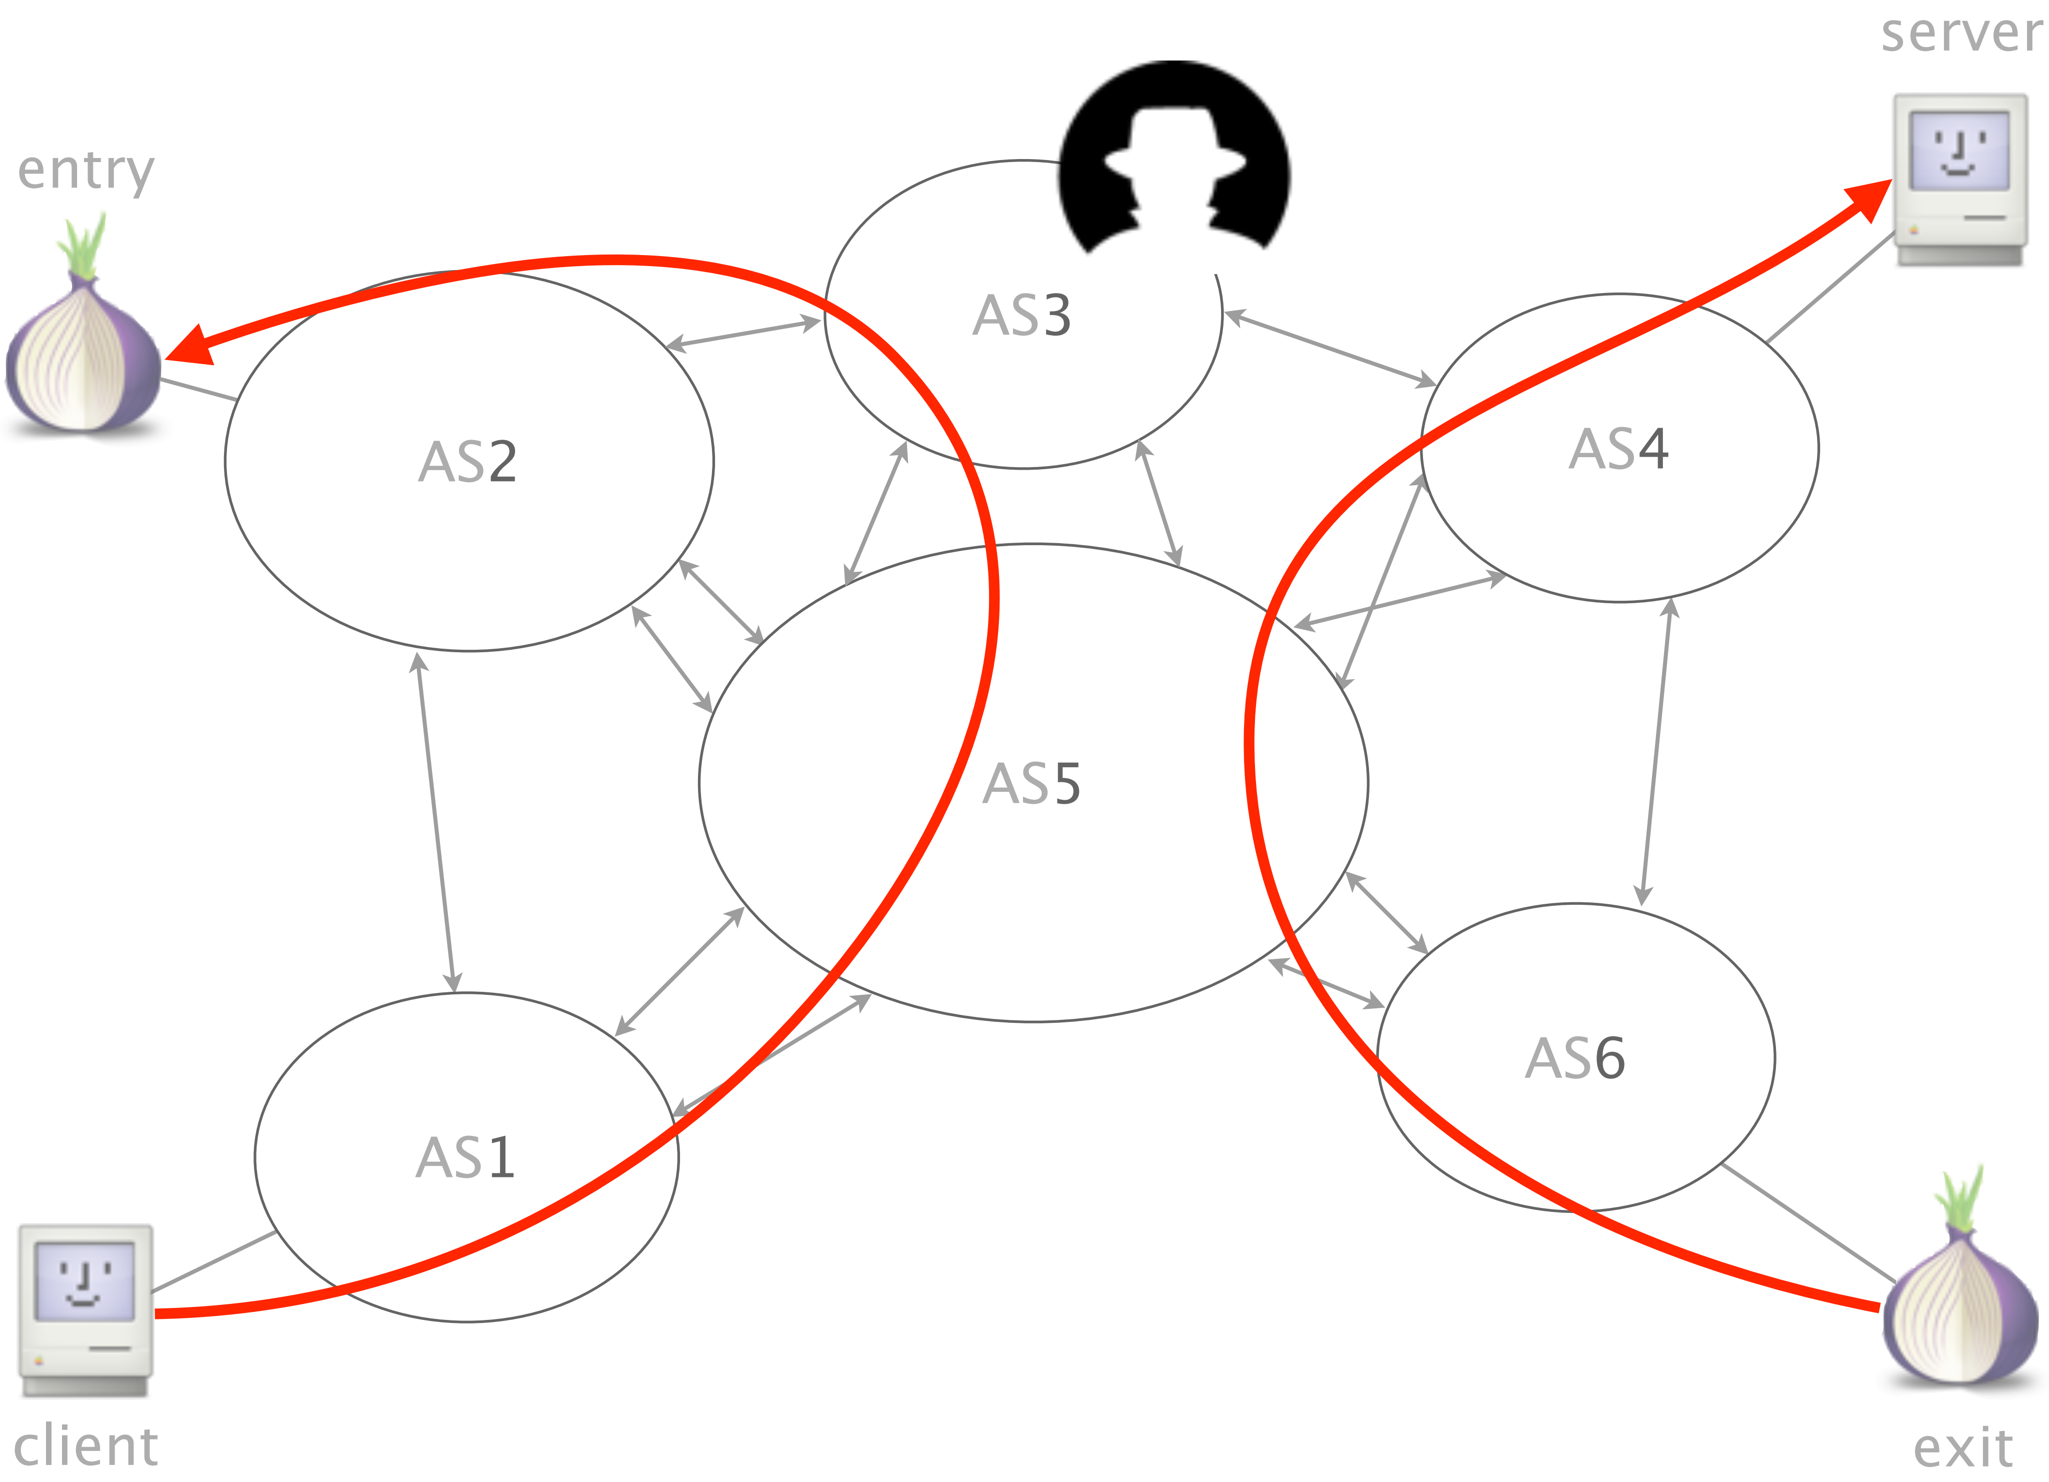

AS3 can put itself on server-to-exit paths by hijacking Tor prefixes

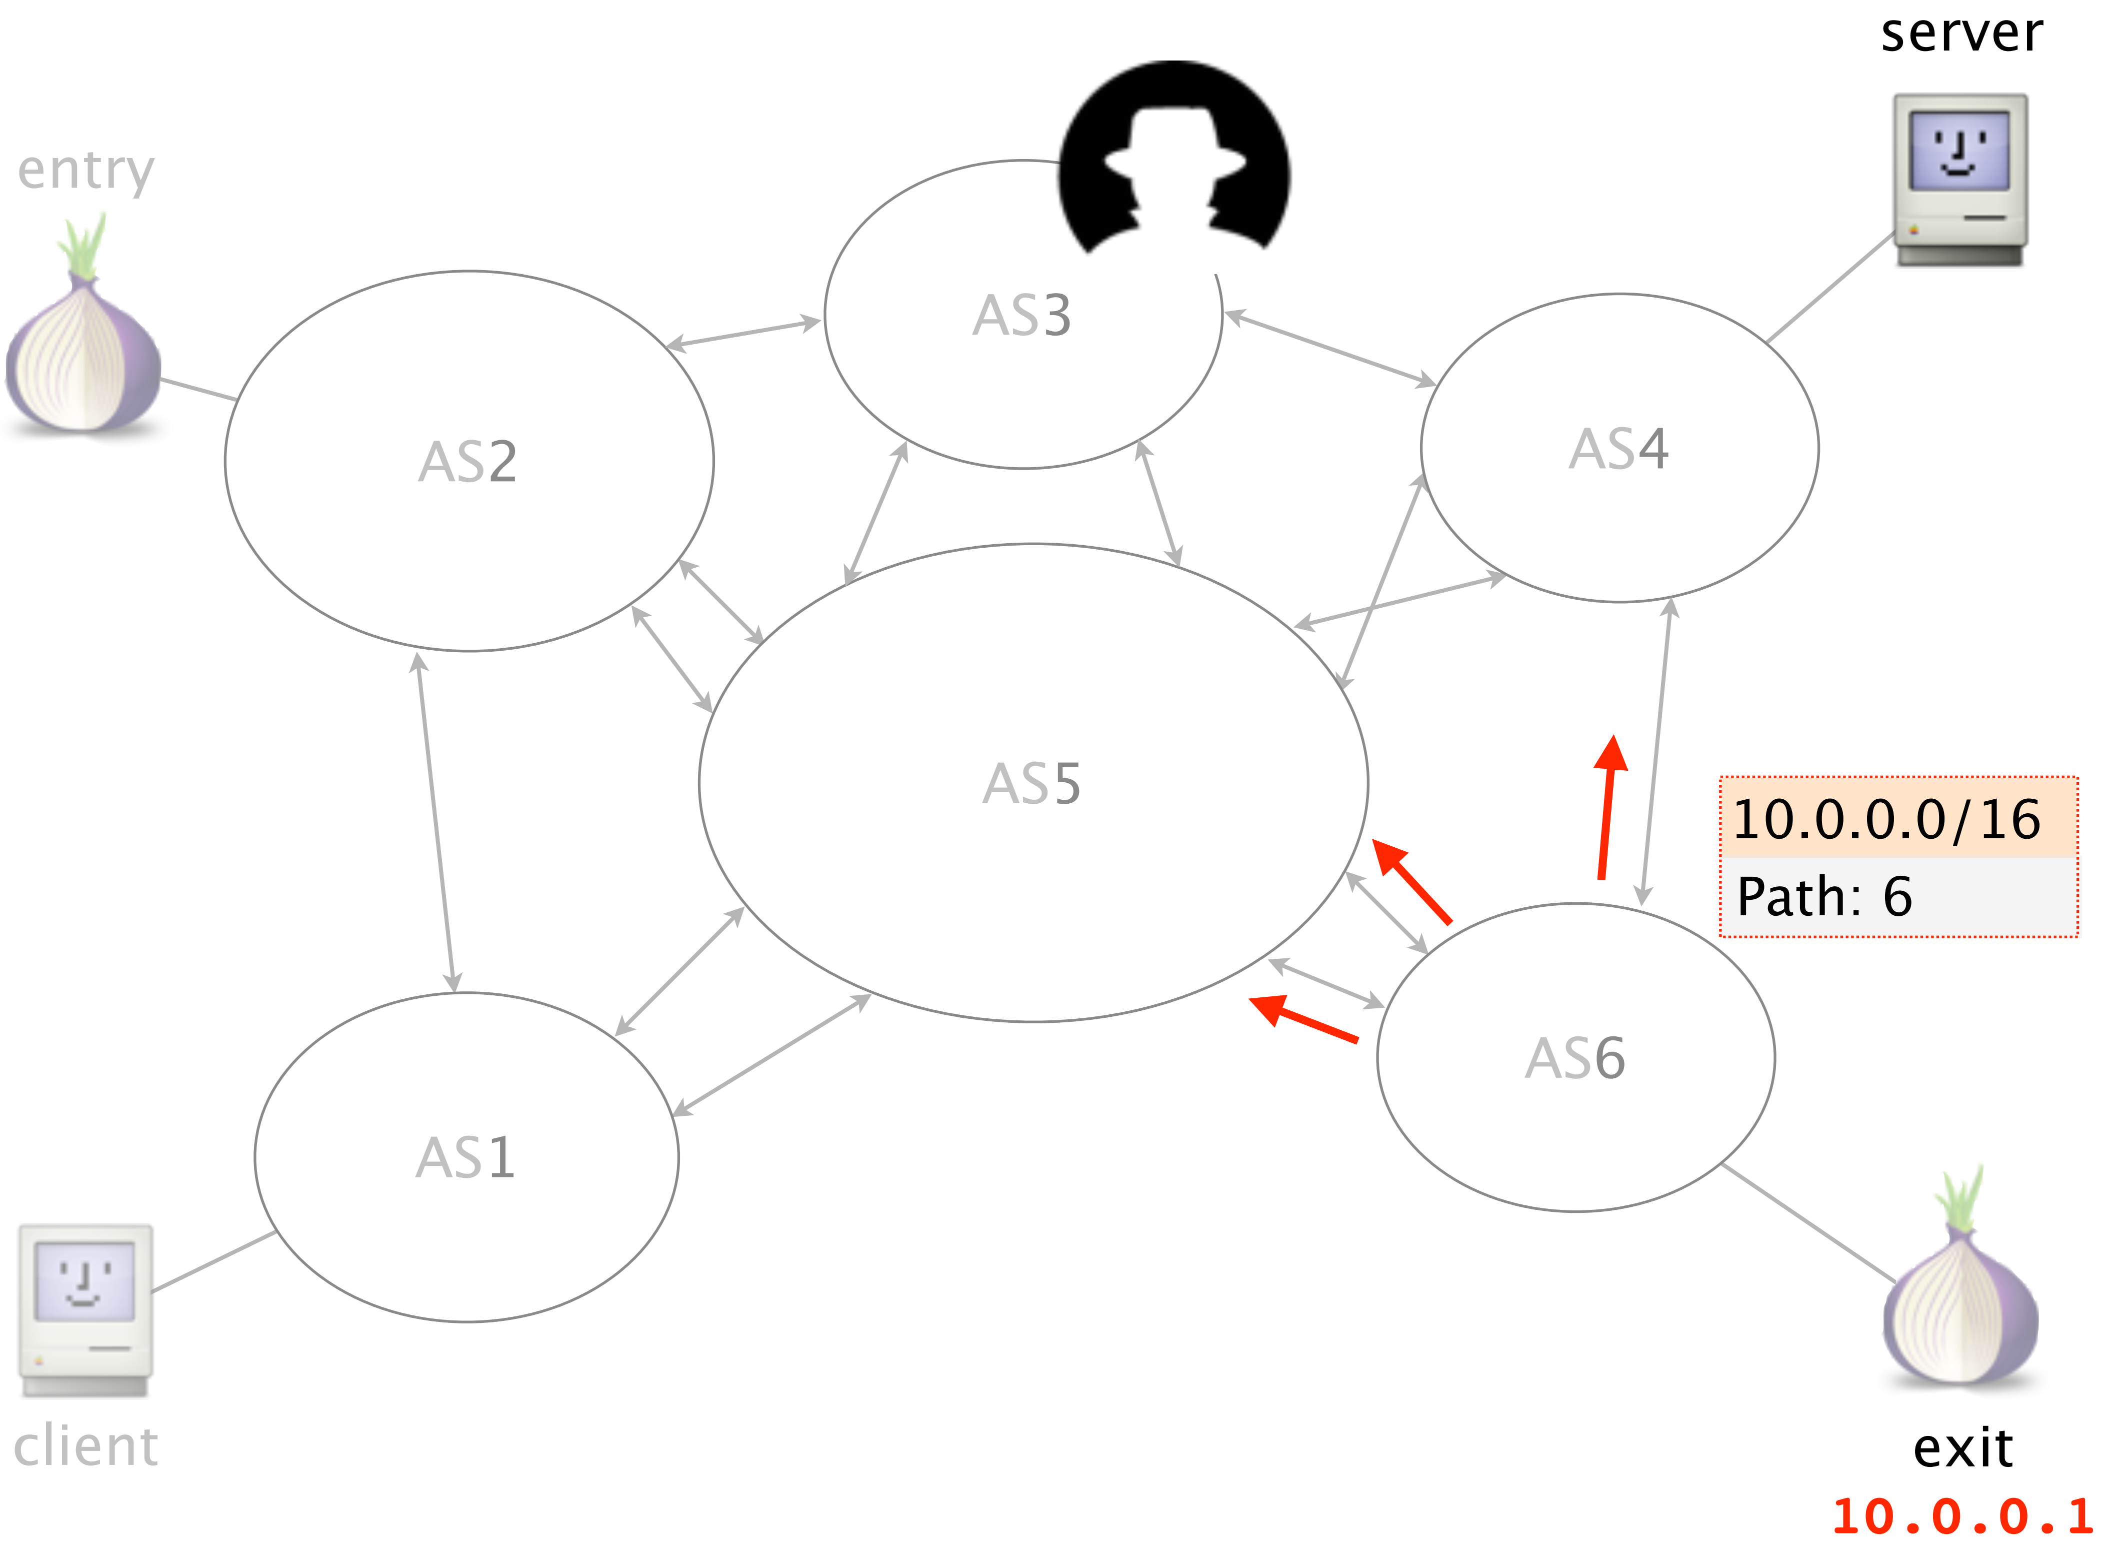

AS3 can put itself on server-to-exit paths by hijacking Tor prefixes

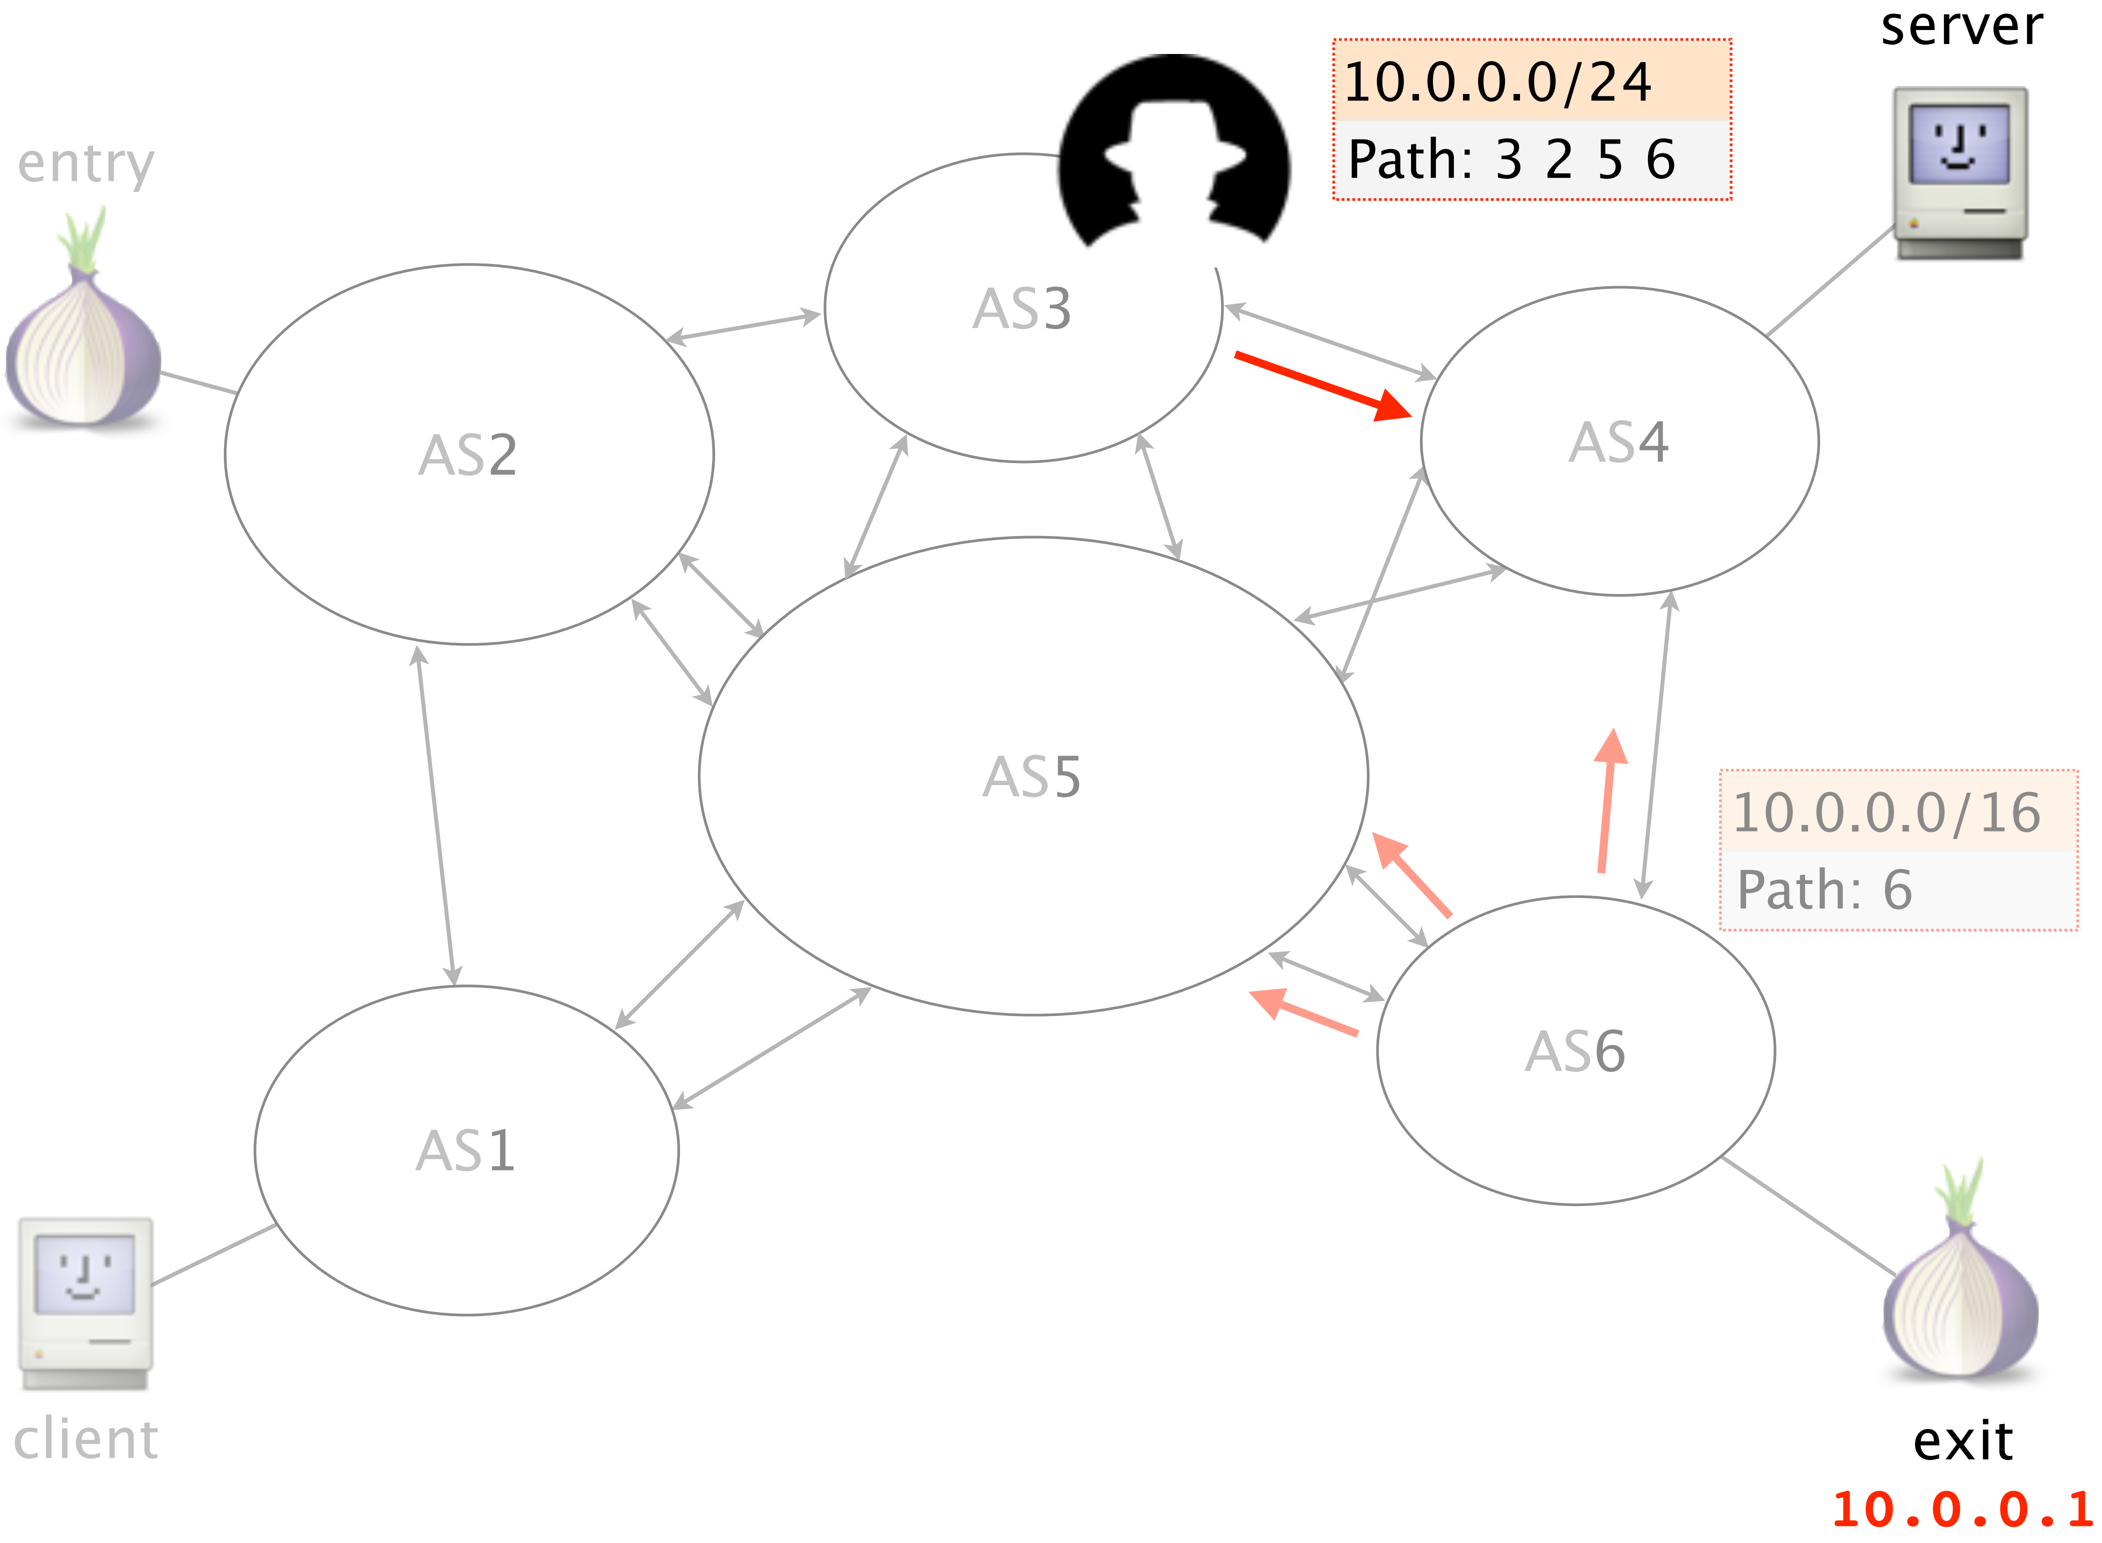

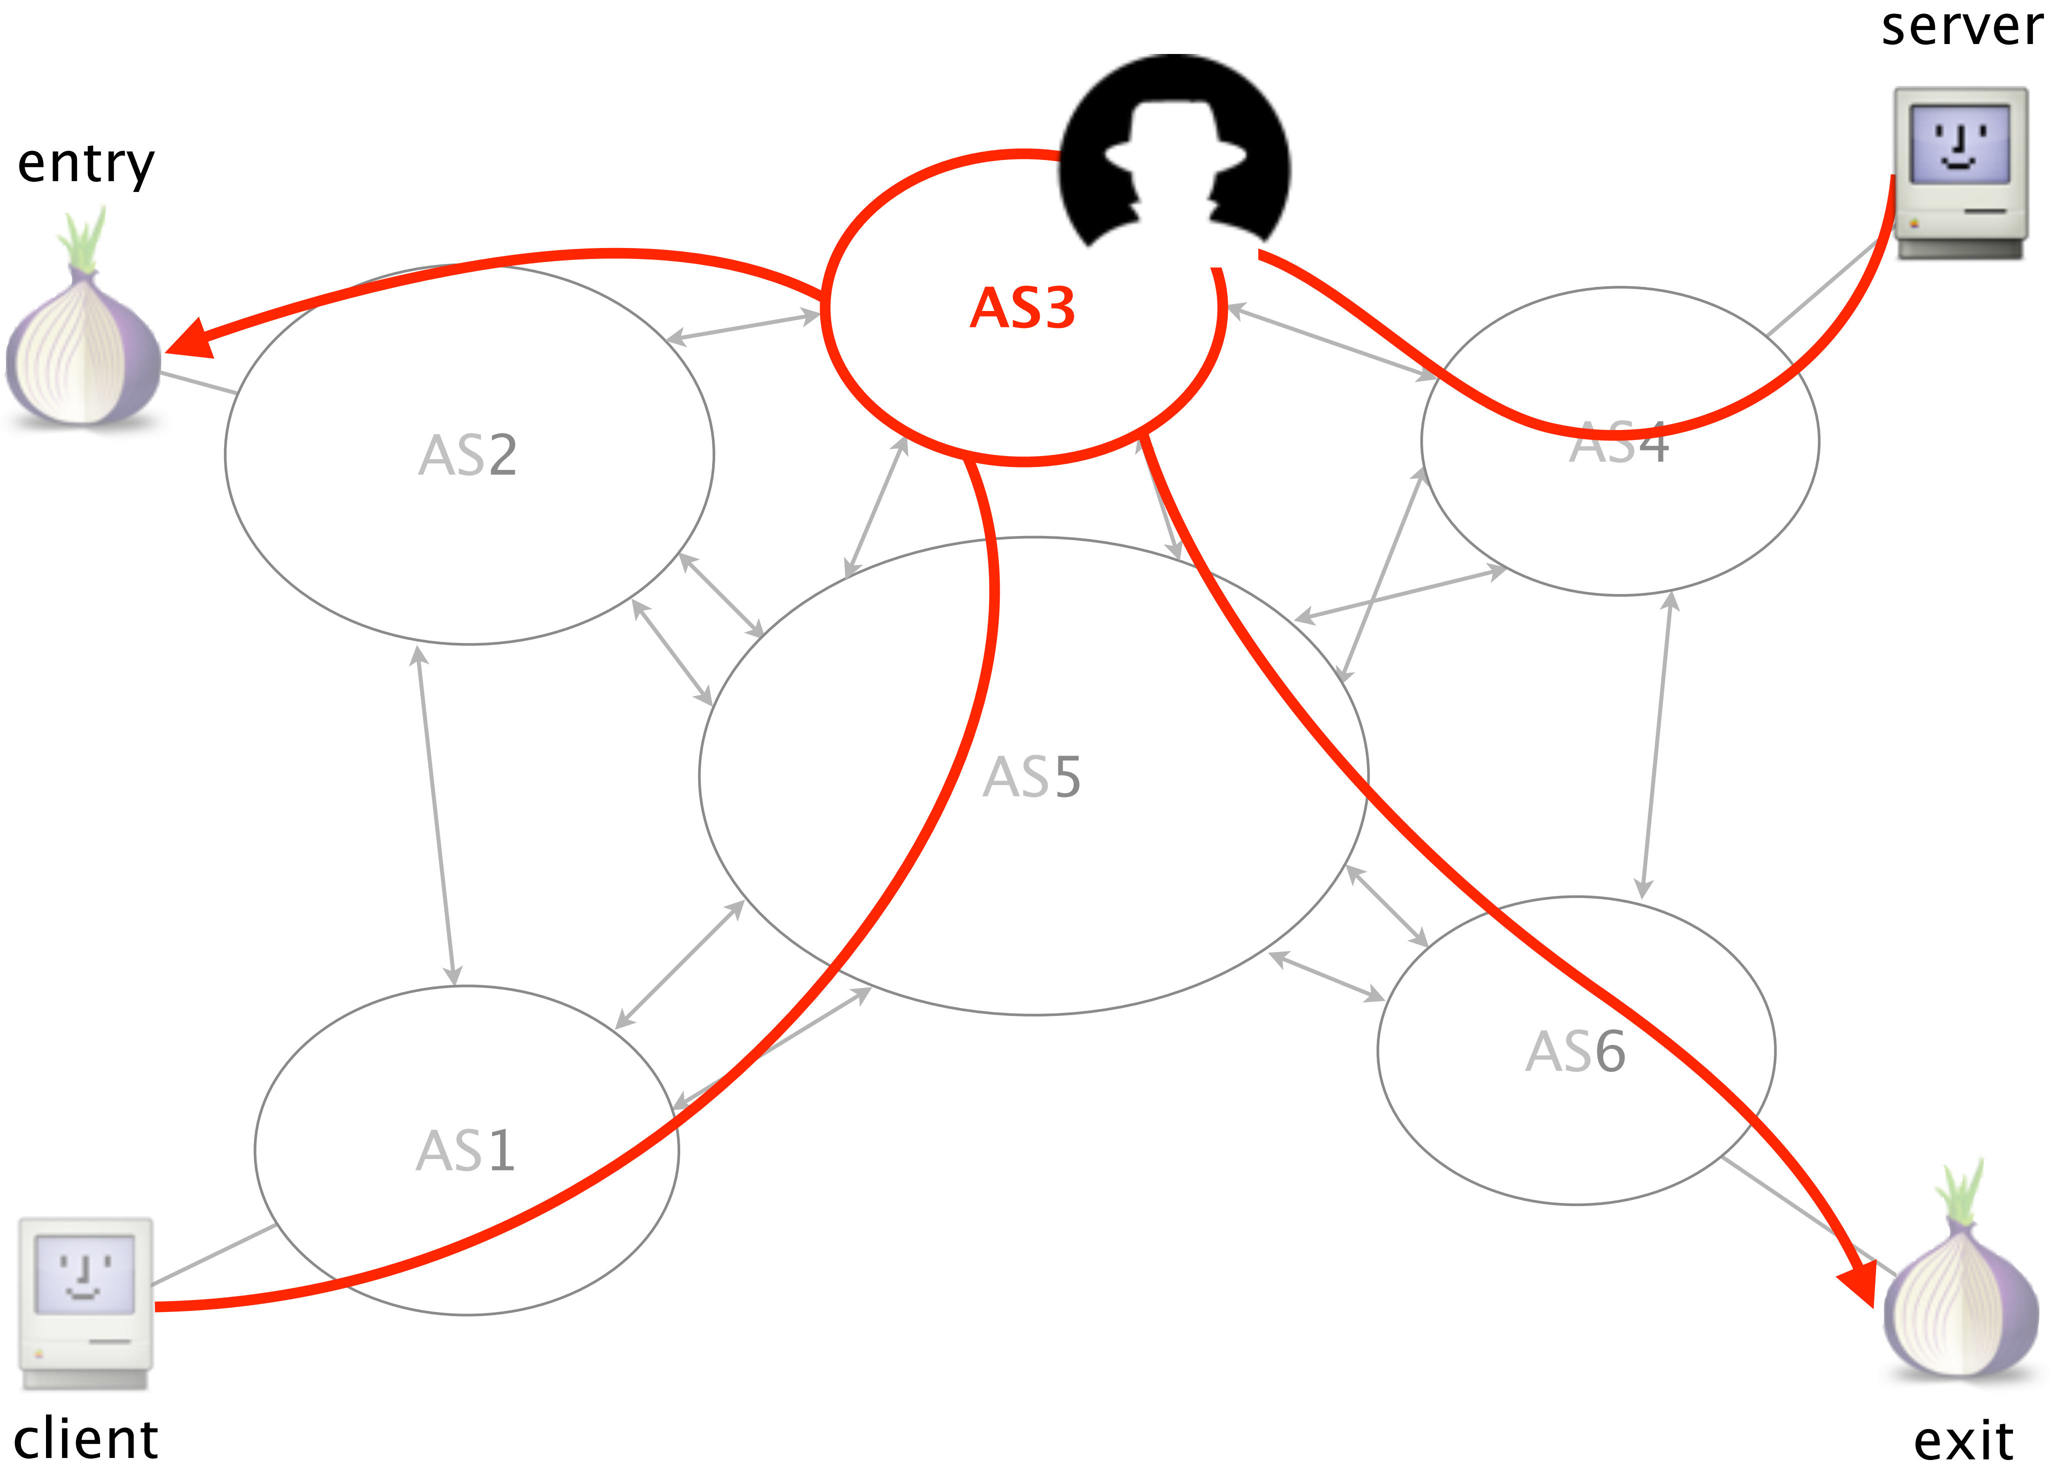
server
entry
AS3
AS2
AS4
AS5
AS1
AS6
client
exit

In November 2010,
China Telecom hijacked 50k prefixes during ~20 min

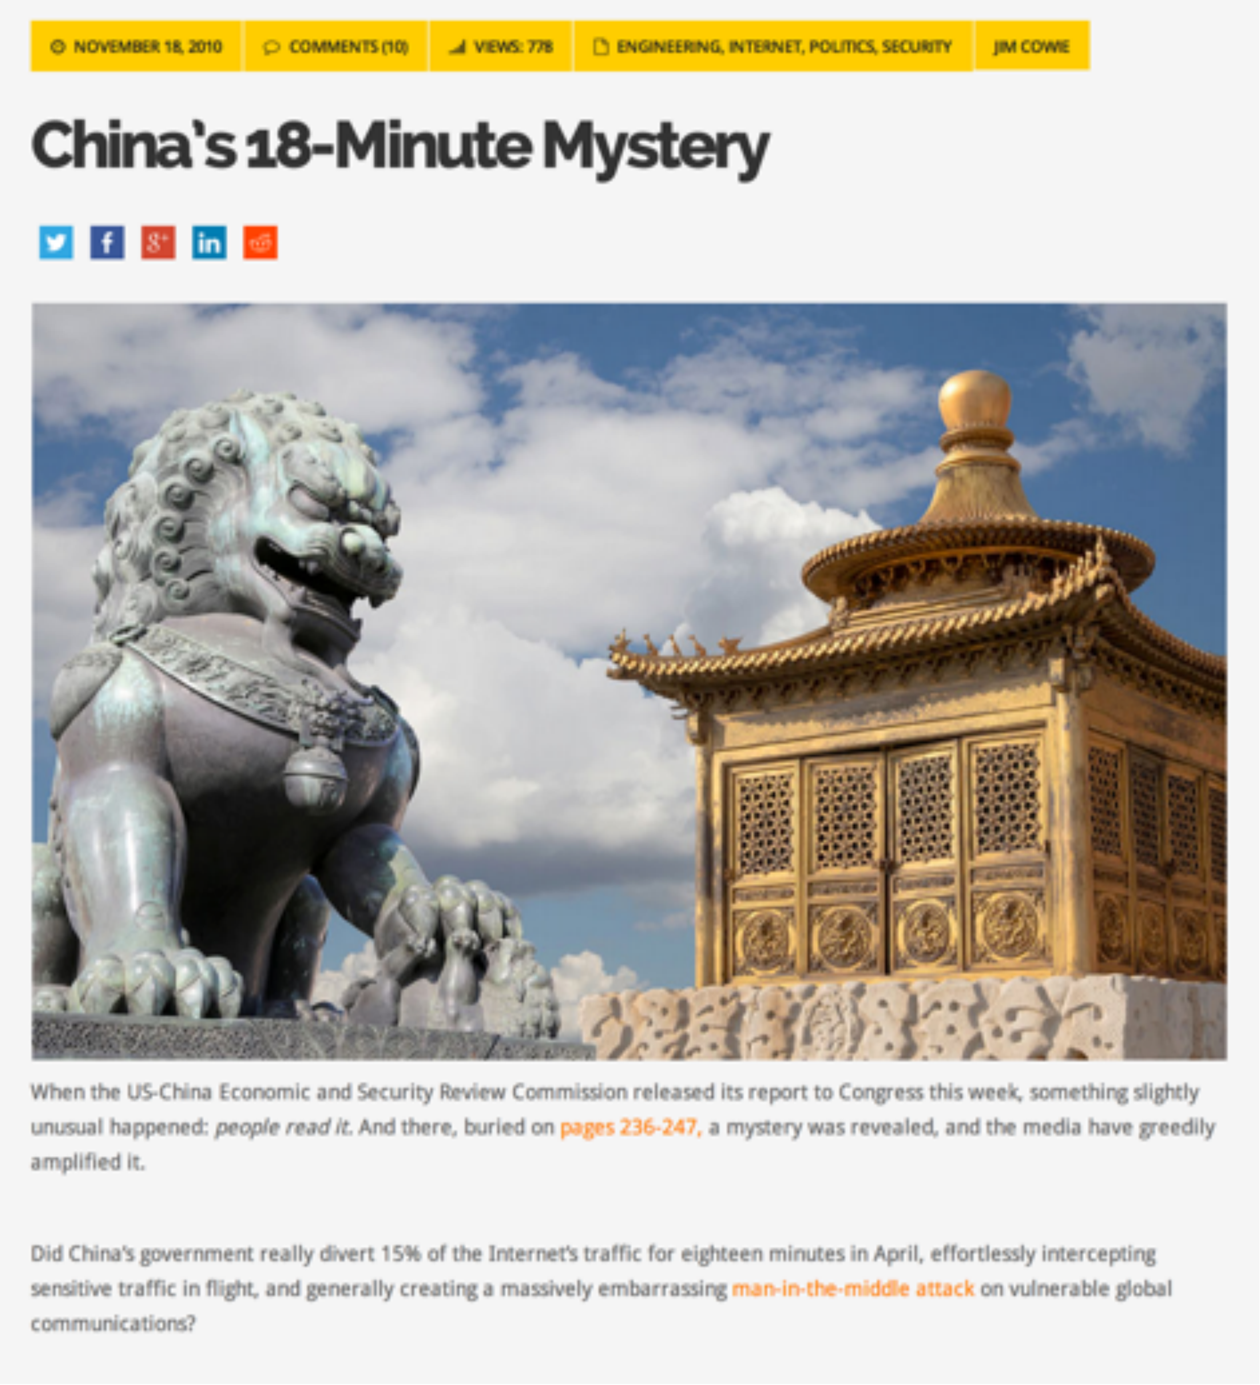

NOVEMBER 18, 2010 | COMMENTS (10) | VIEWS: 778 | ENGINEERING, INTERNET, POLITICS, SECURITY | JIM COWIE

China's 18-Minute Mystery

When the US-China Economic and Security Review Commission released its report to Congress this week, something slightly unusual happened: *people read it*. And there, buried on pages 236-247, a mystery was revealed, and the media have greedily amplified it.

Did China's government really divert 15% of the Internet's traffic for eighteen minutes in April, effortlessly intercepting sensitive traffic in flight, and generally creating a massively embarrassing man-in-the-middle attack on vulnerable global communications?

China Telecom

always sees traffic between
its customer and entry relays

During the attack, it also

saw traffic to/from exit relays
for a non-trivial fraction of traffic

Intentional? No one knows.

Attack#3: Asymmetric routing, too, increases the # of AS-level adversaries

So far, we have considered one side of the Tor traffic: client–to–entry and exit–to–server

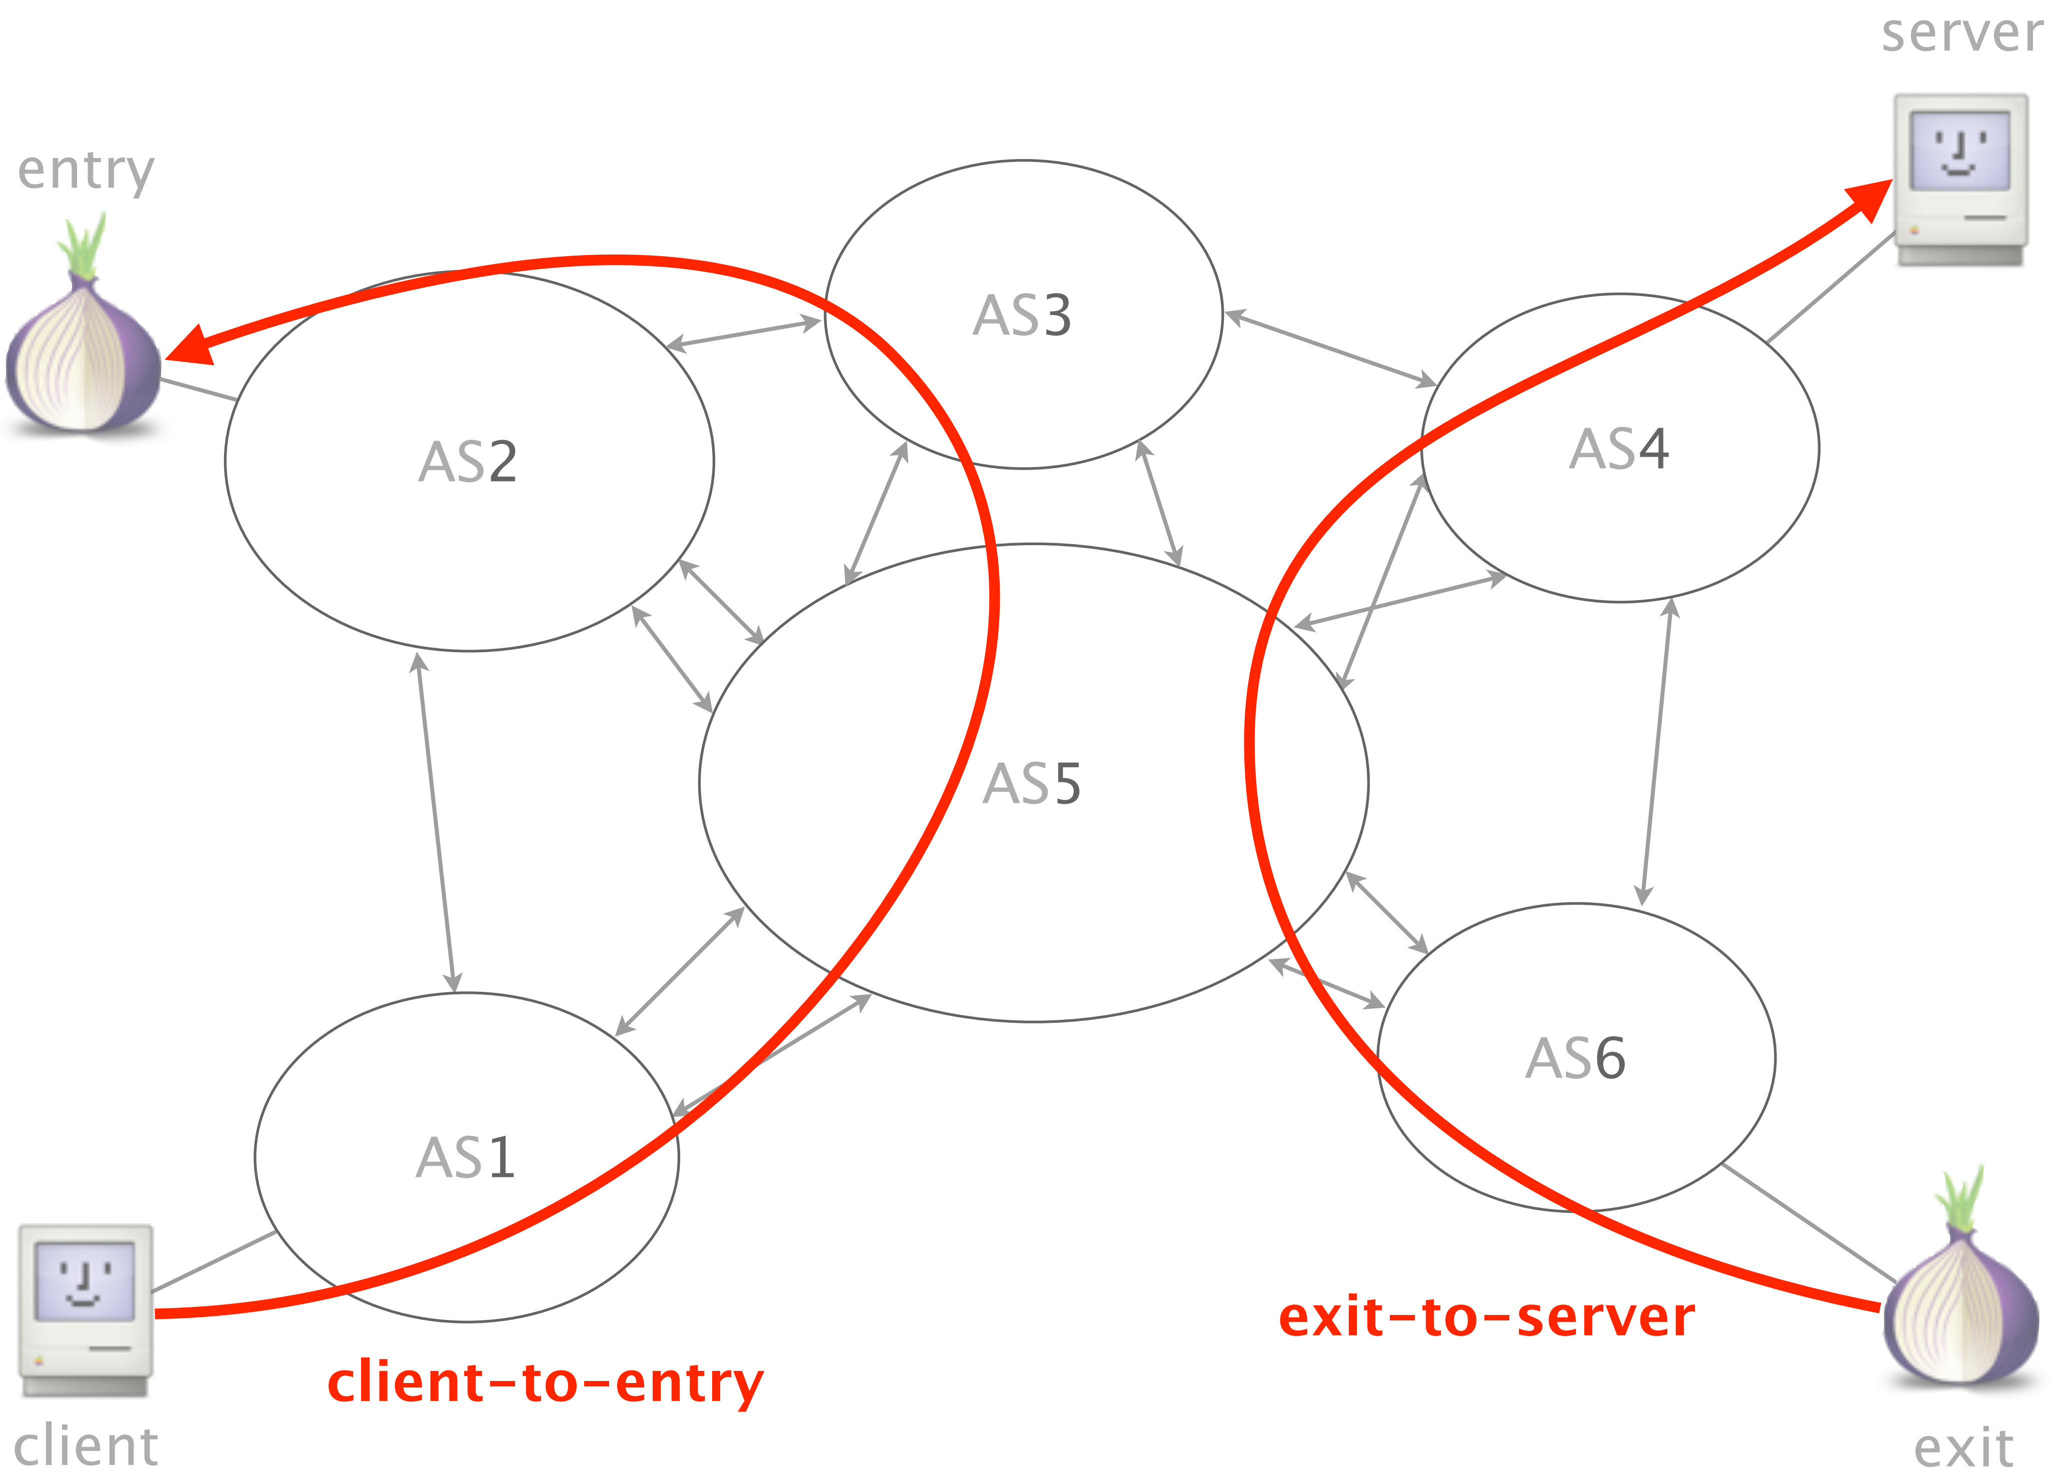

However, because of policies, routing is often *asymmetric*

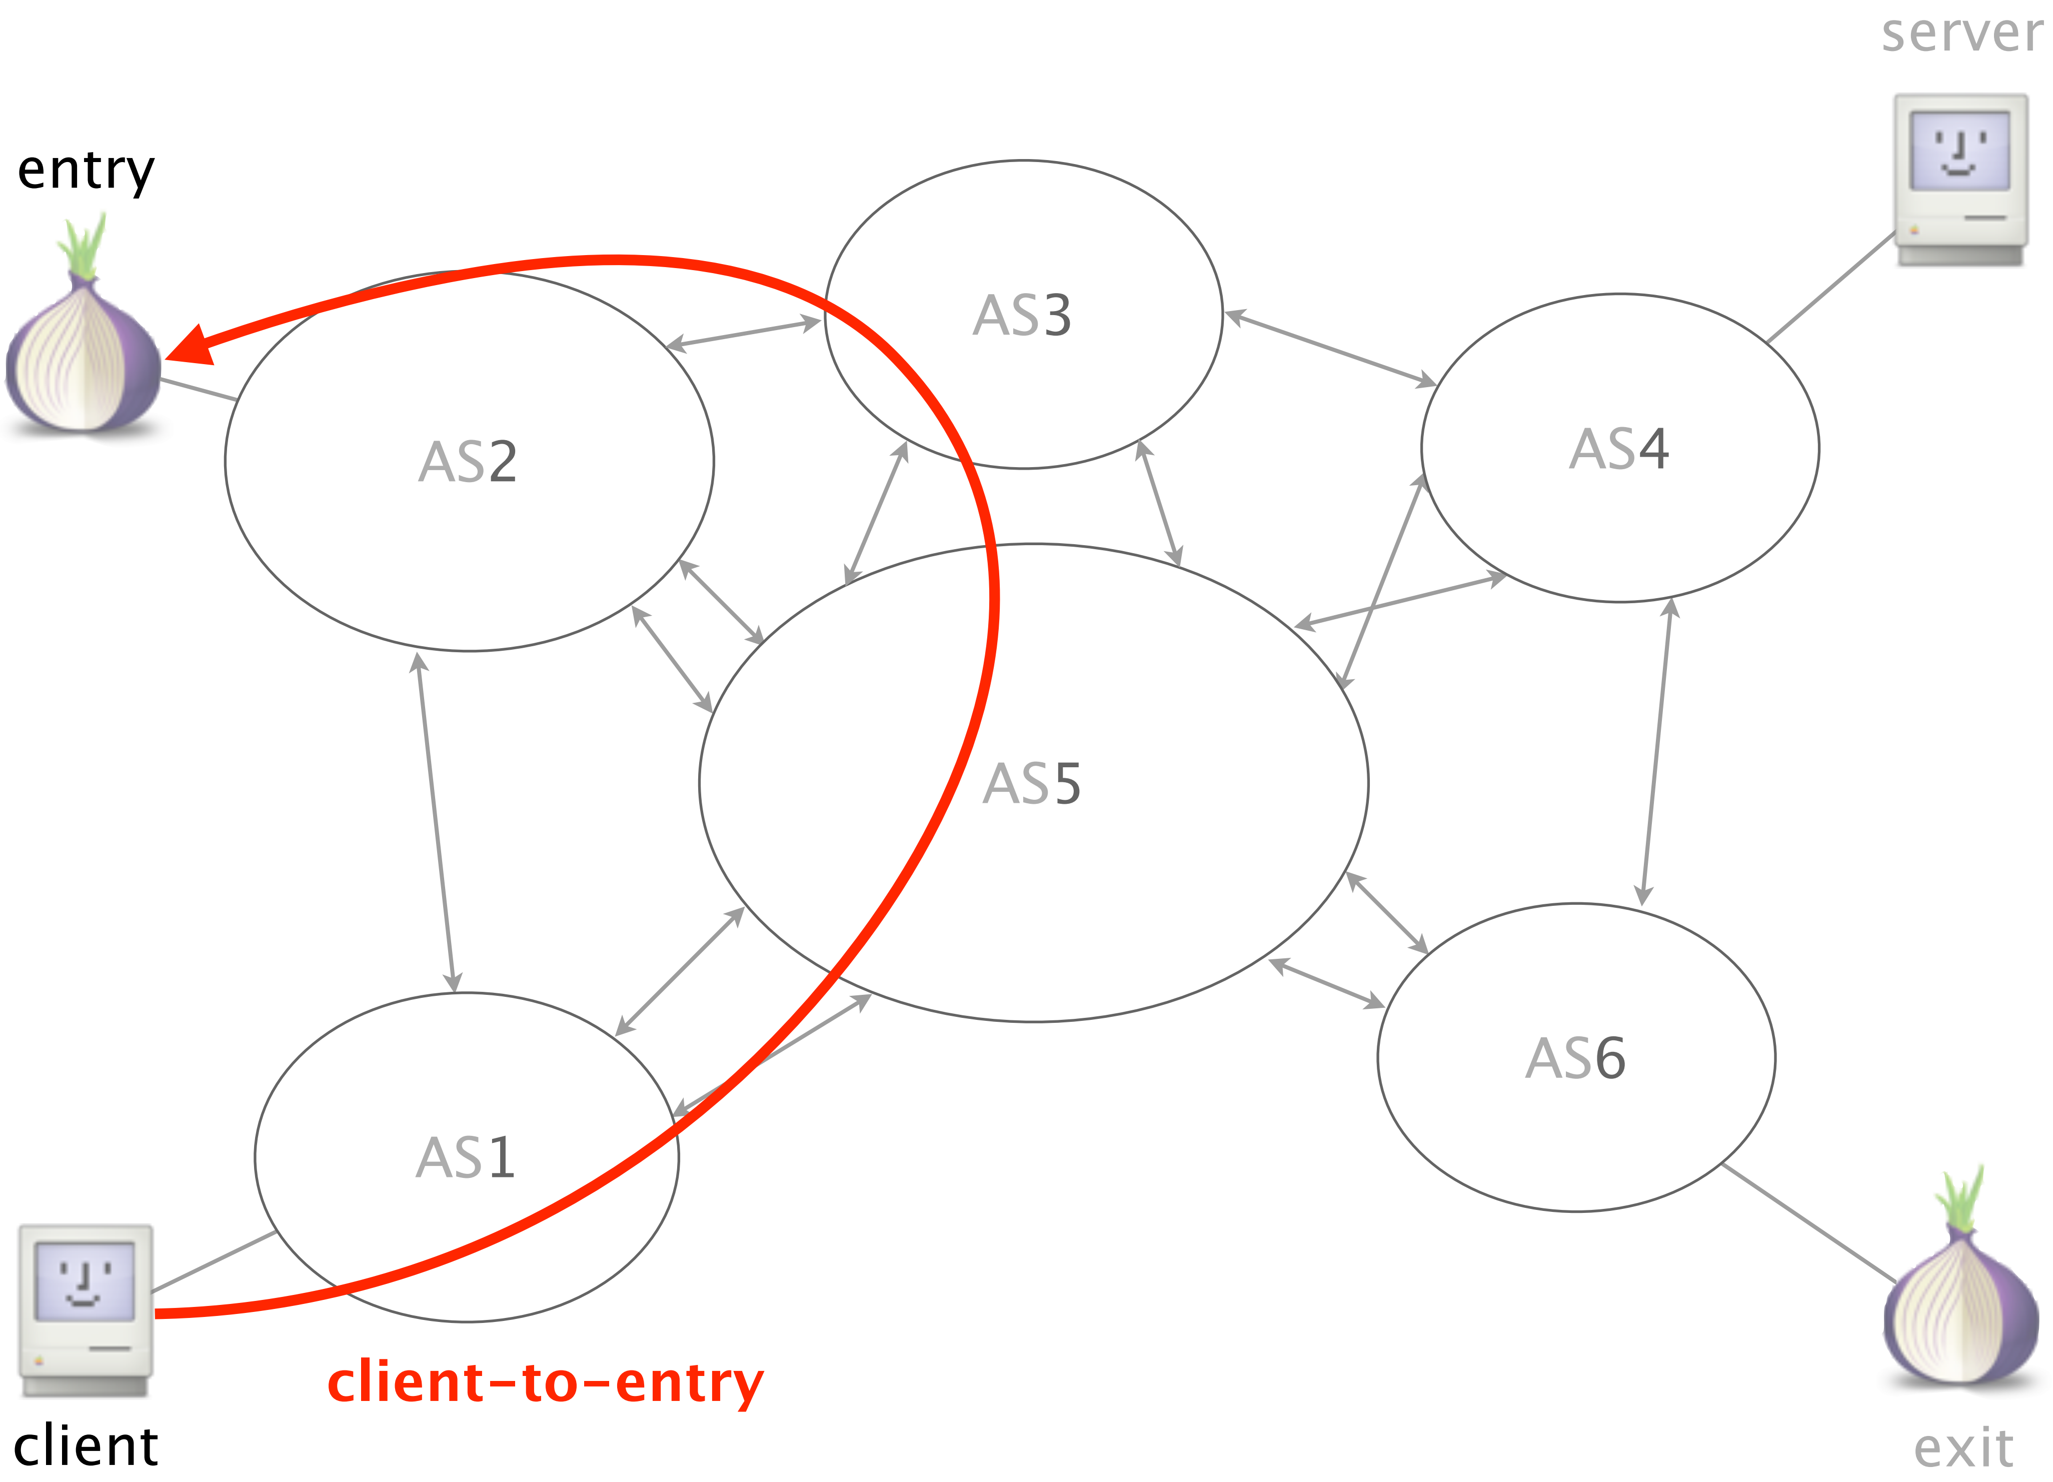

While AS4 does not see client–to–entry traffic, it sees entry–to–client traffic

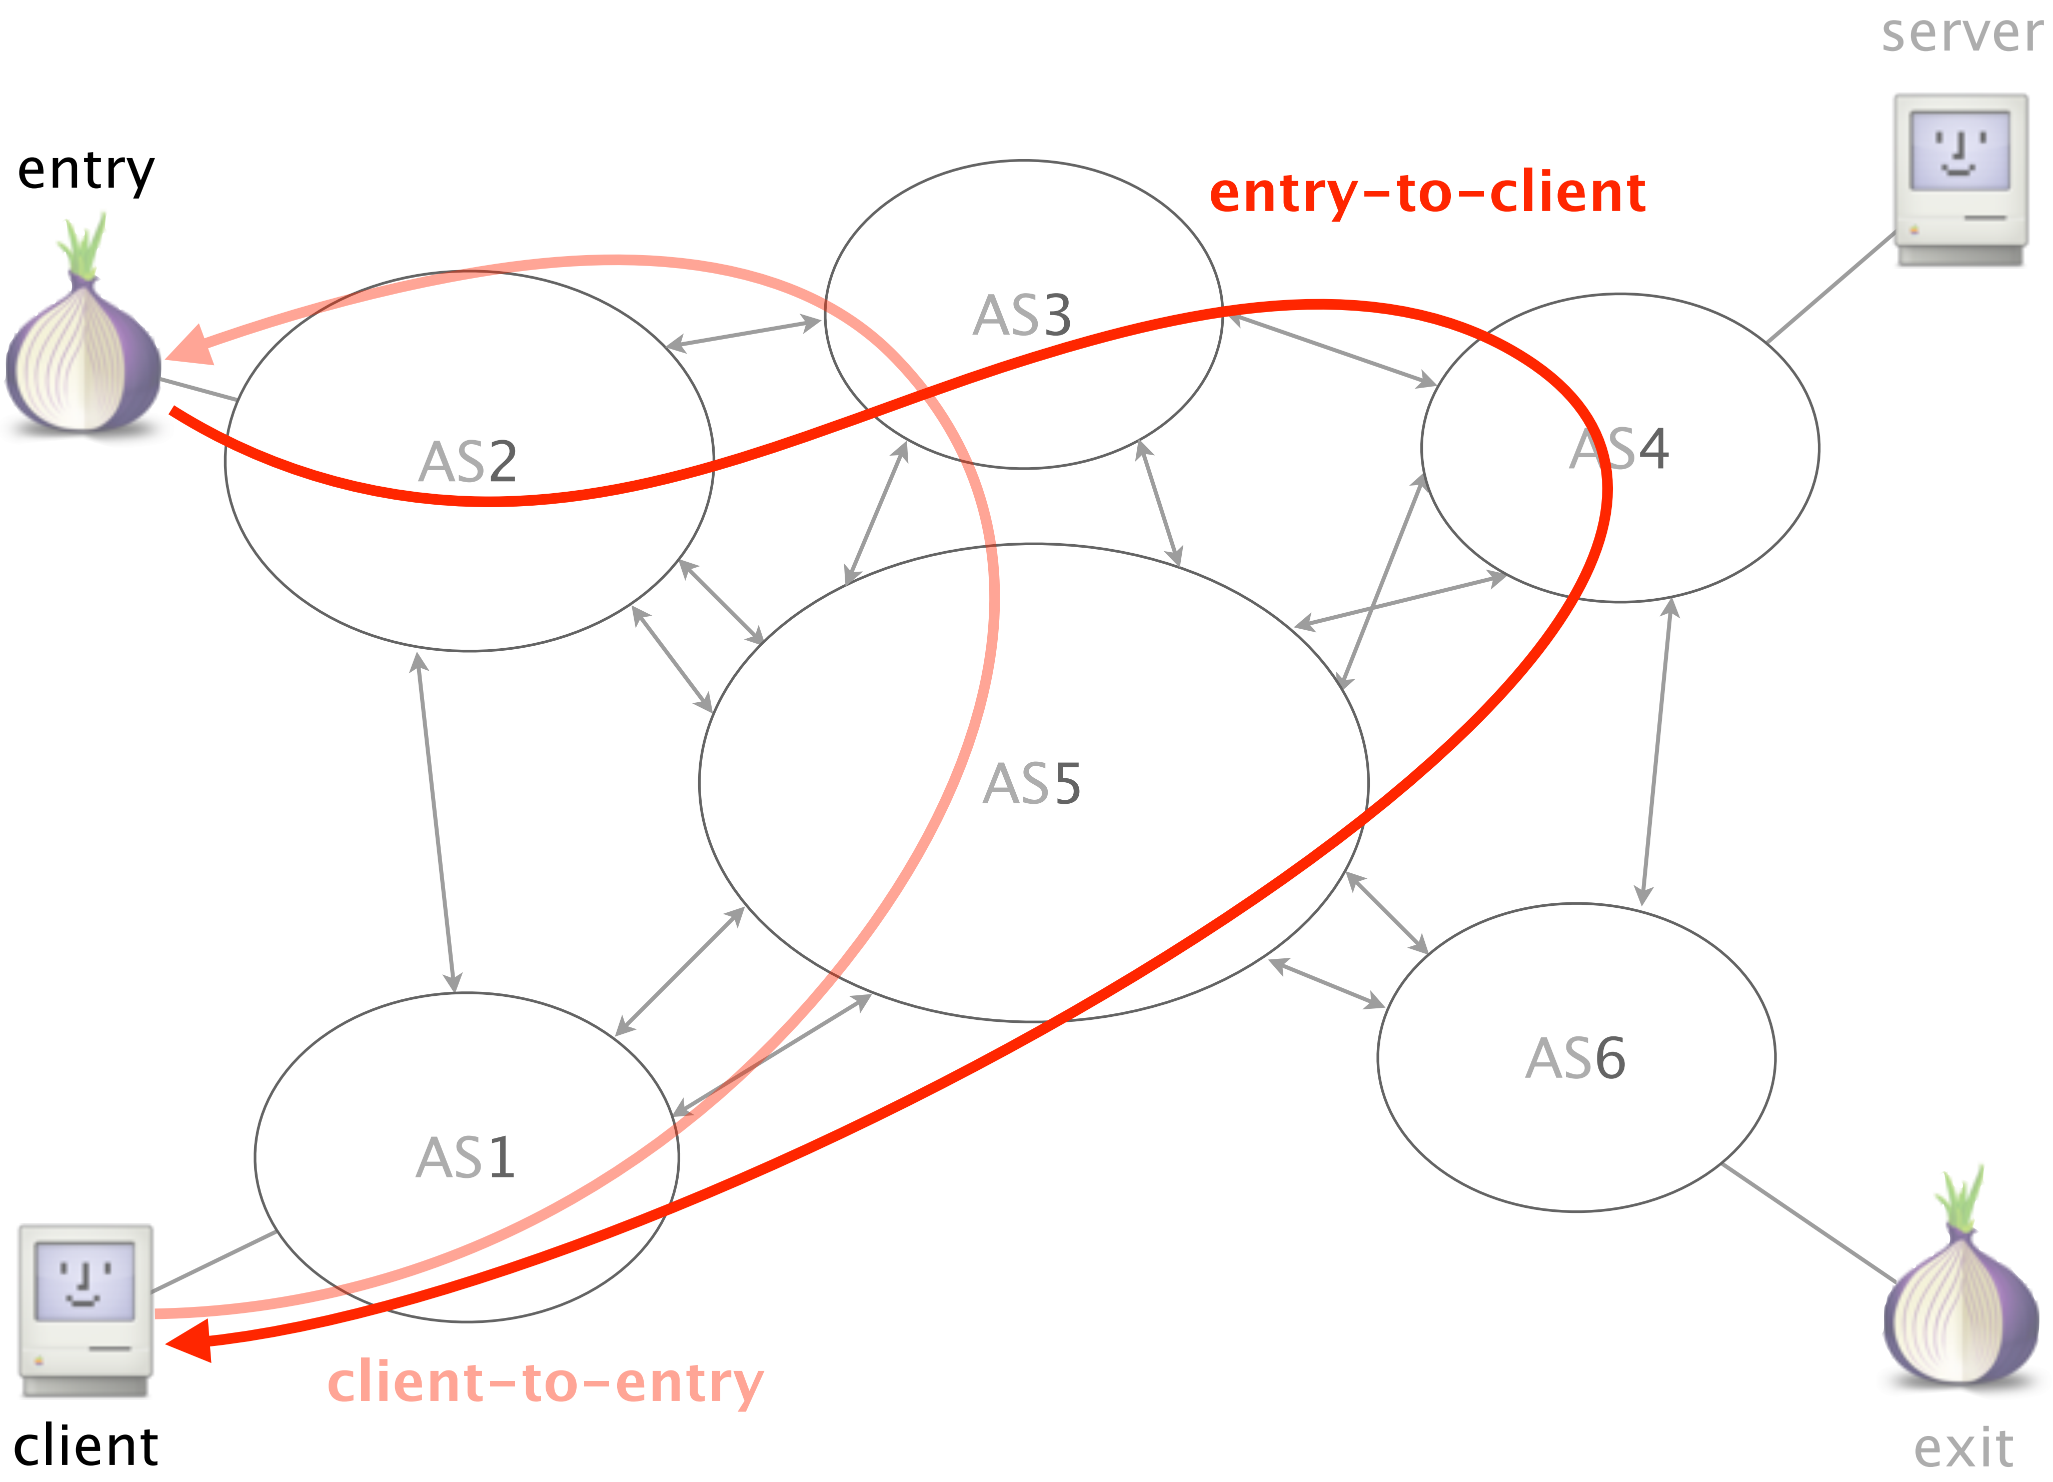

Same applies for server–to–exit traffic

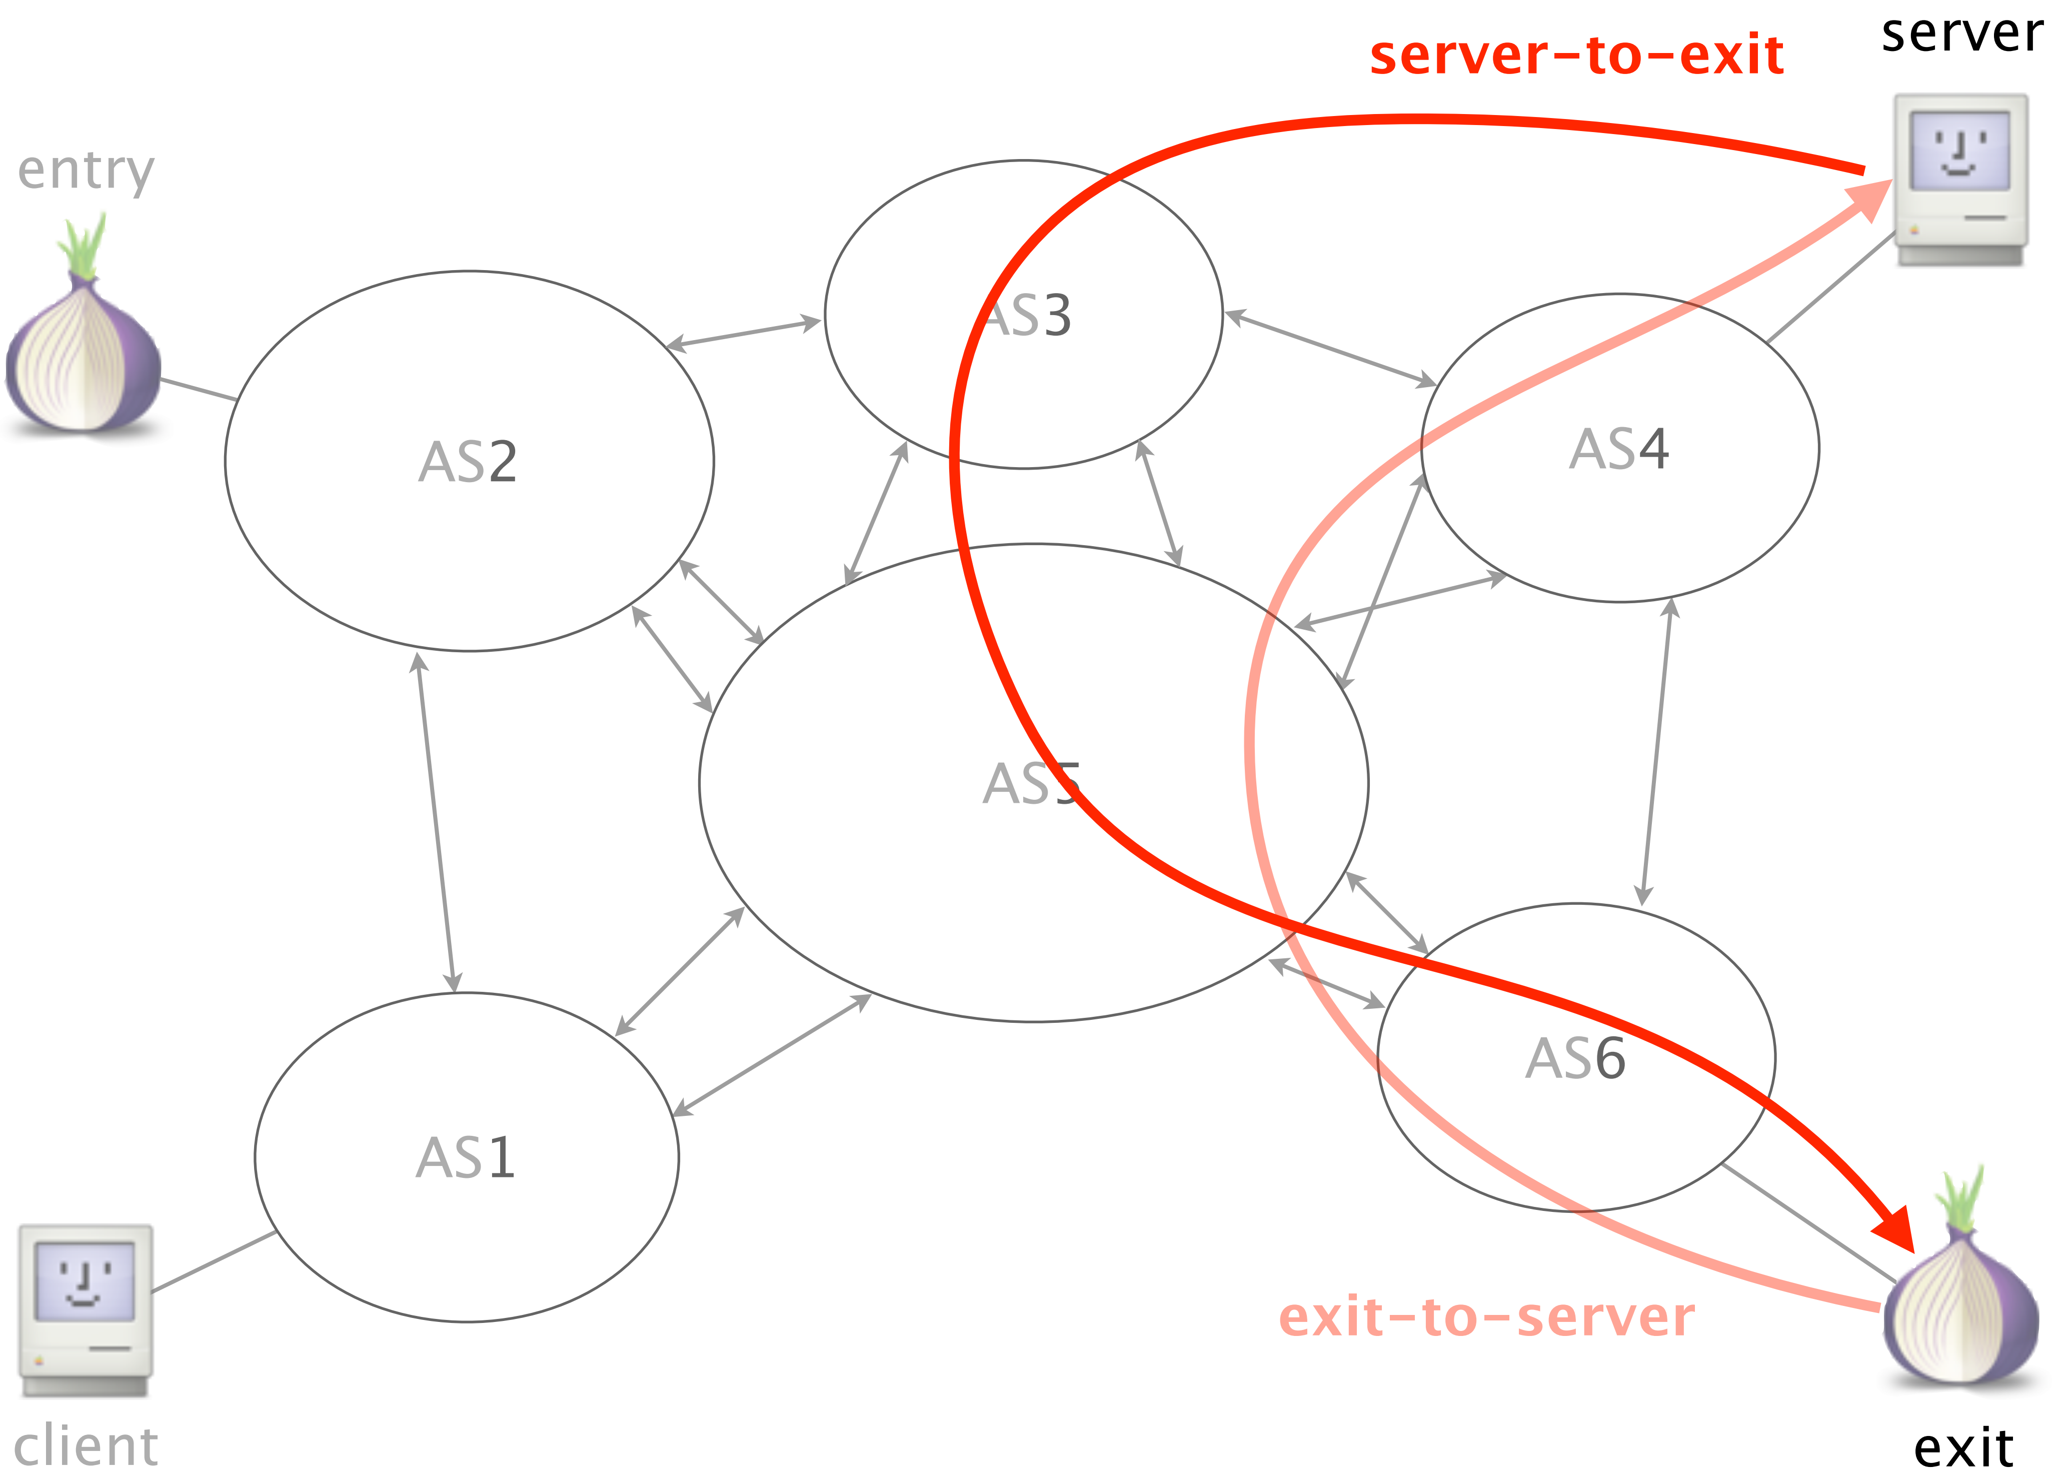

In terms of timing information,
both side of a TCP connection are highly correlated

In terms of timing properties,
both side of a TCP connection are highly correlated

When collecting TCP
timing information,

seeing one direction (*e.g.*, data packets)
is almost equivalent to
seeing two directions (ACKs & data packets)

Considering only one direction, 1 AS is potentially compromising

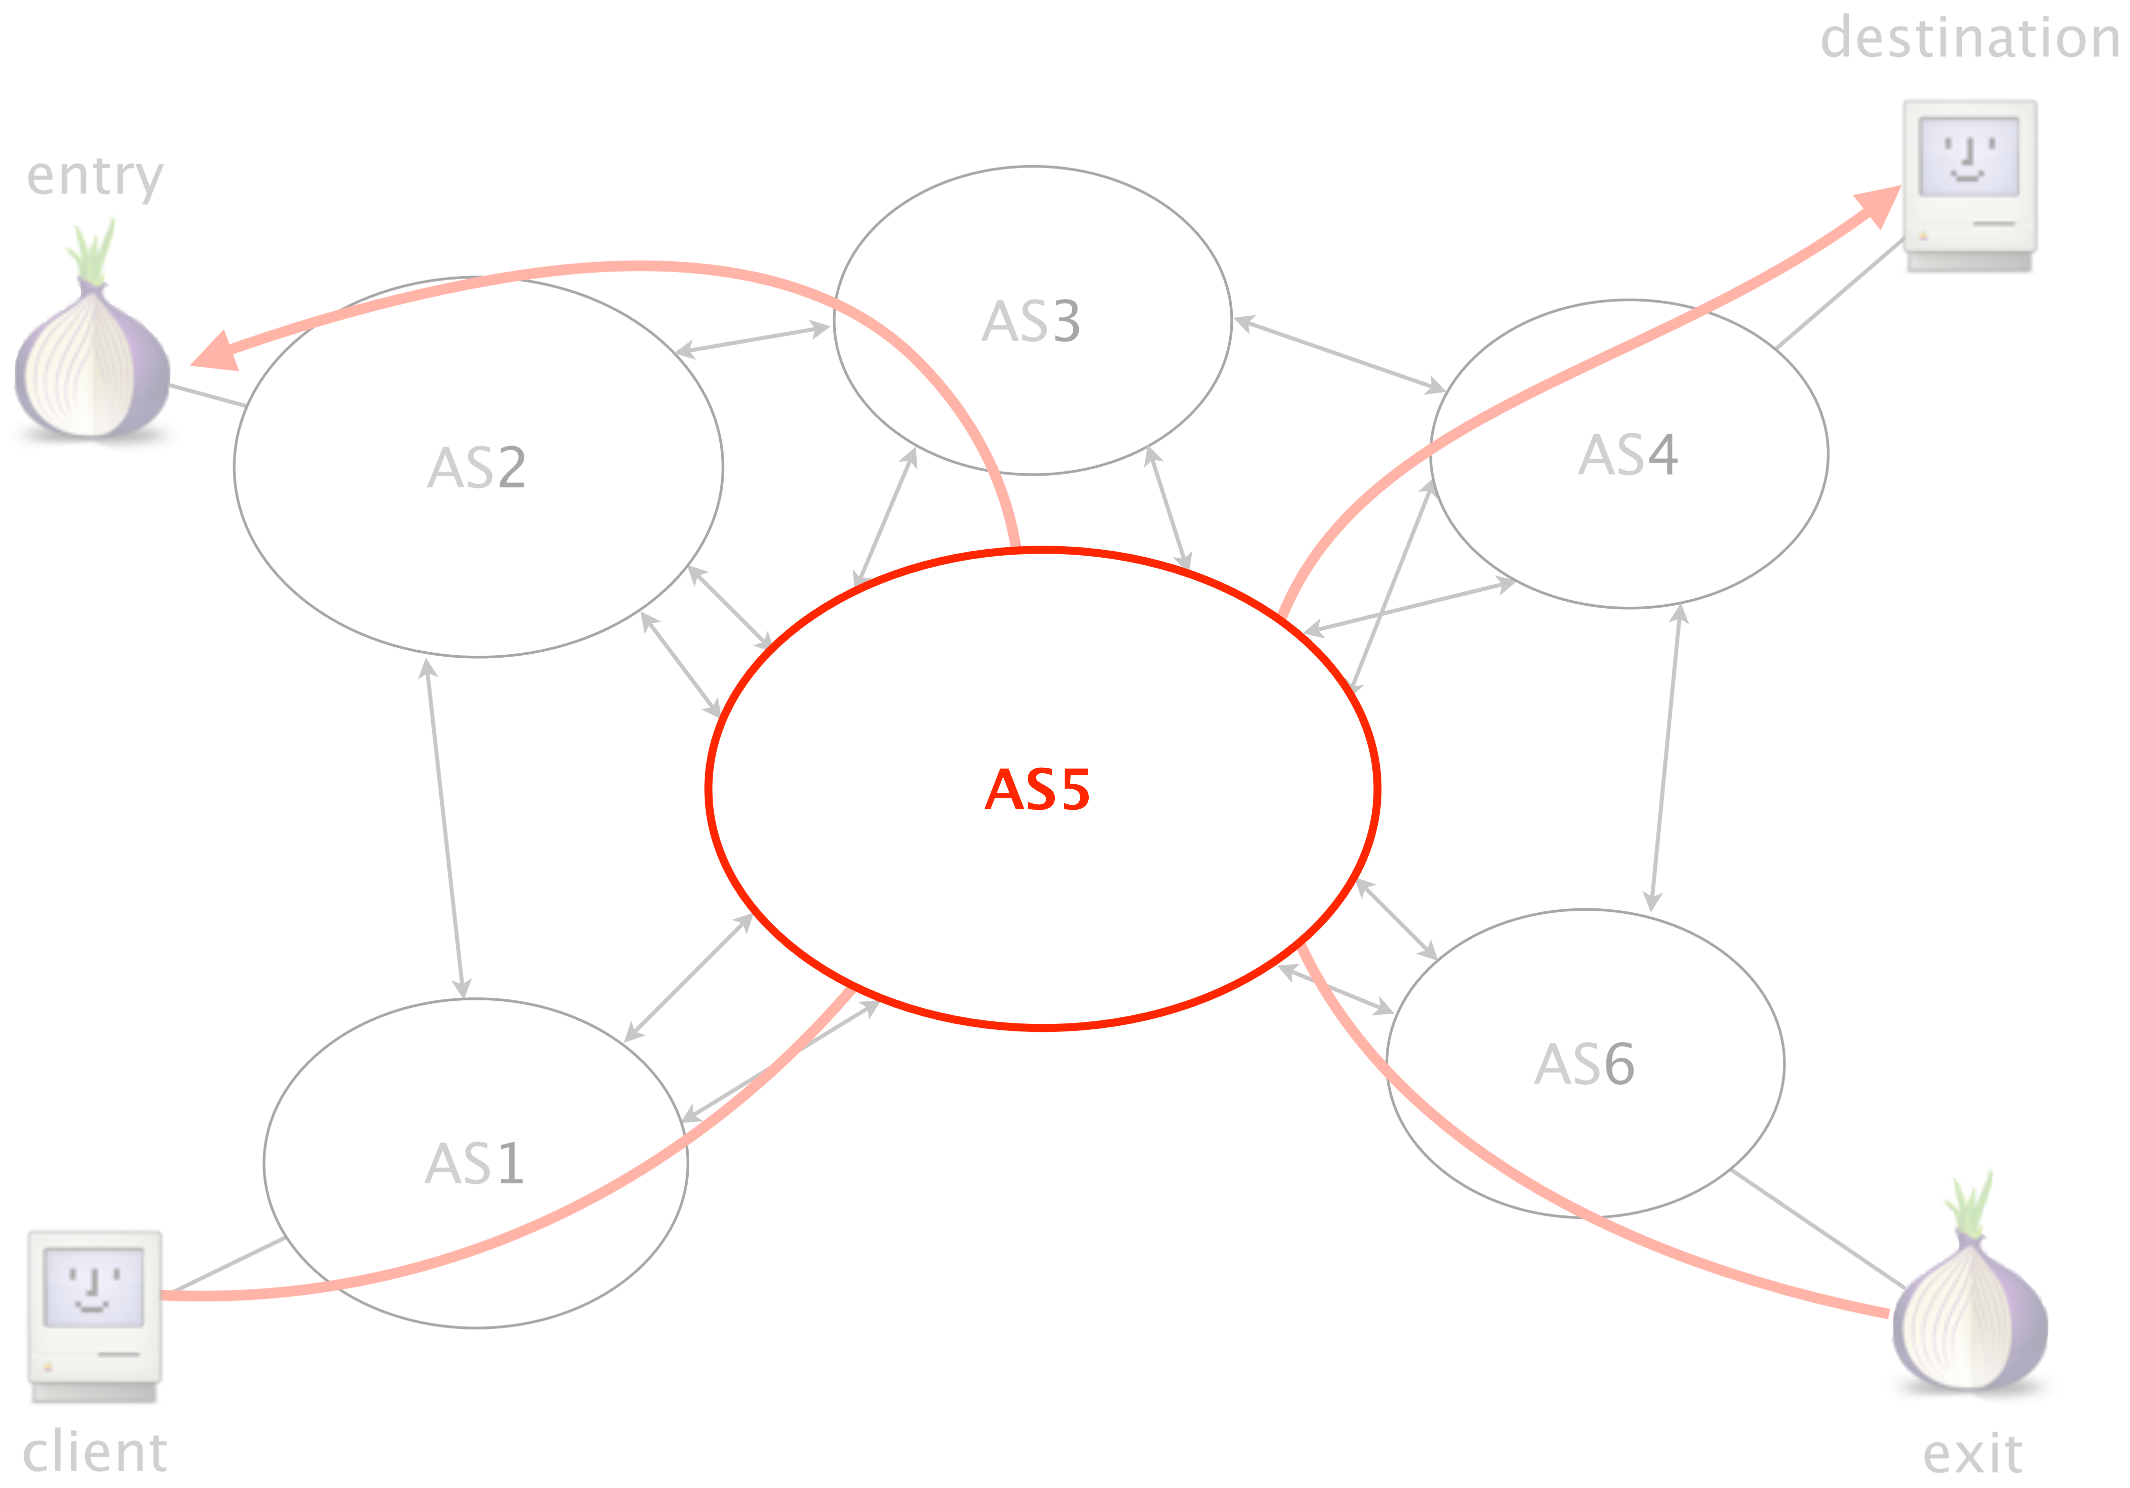

Considering both directions,

3 ASes are potentially compromising

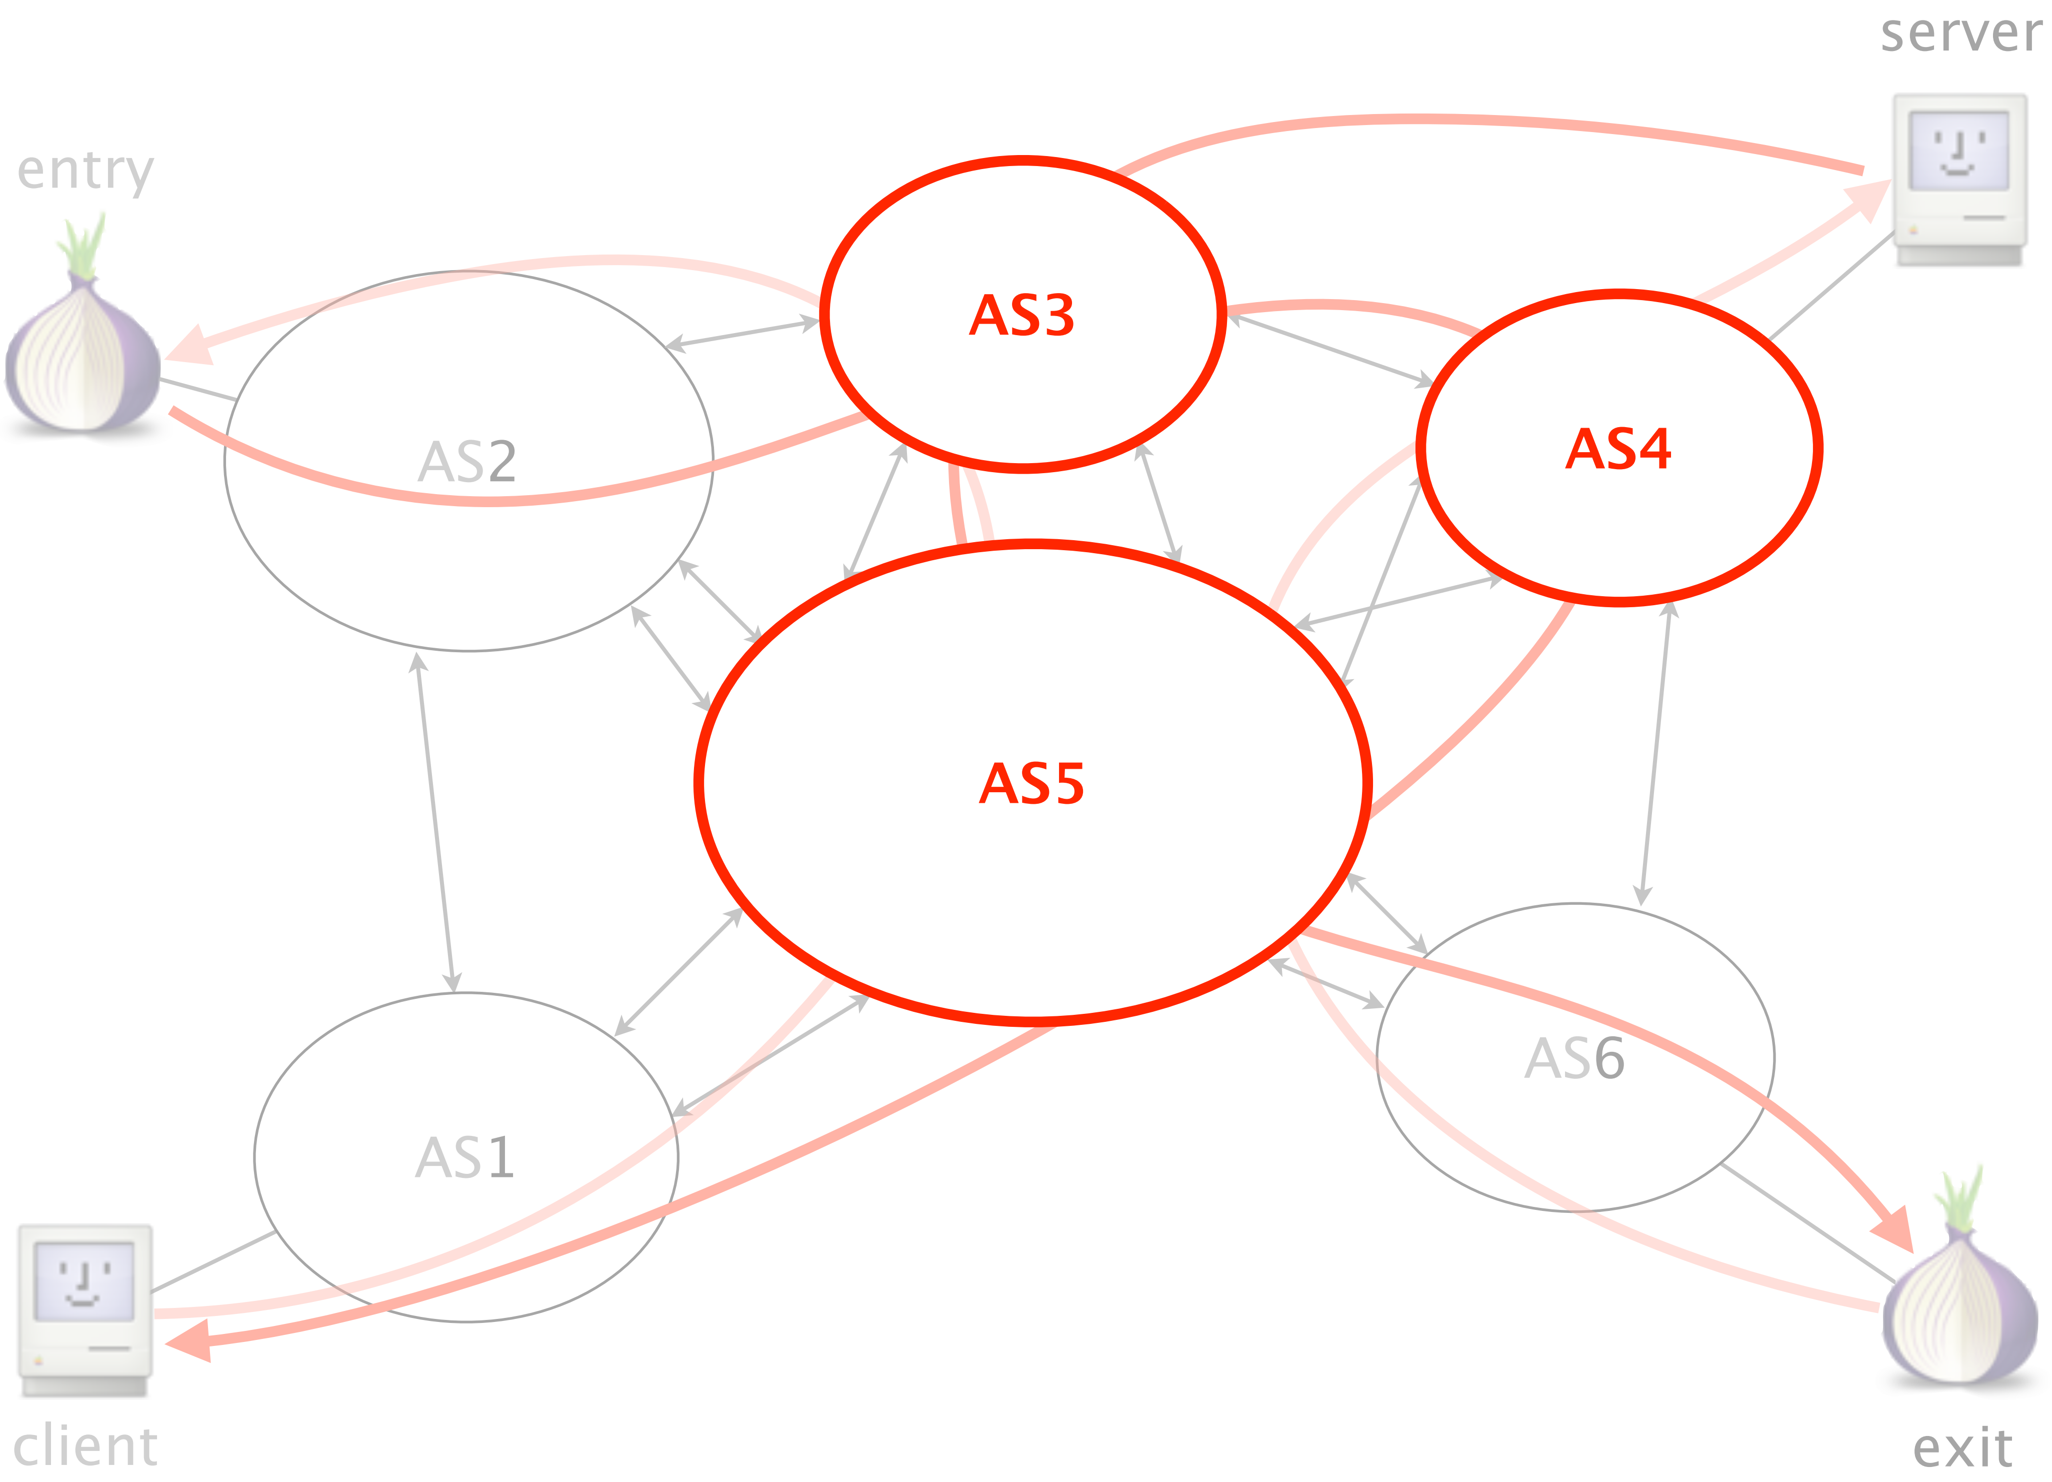

Anonymity on Quicksand

Using BGP to compromise Tor

Attacks

All your traffic belongs to me

2 Preliminary results

Eyes wide open

Countermeasures

Close the curtains

Question#1: How many networks host entry and exit relays?

We collected BGP-related information for each Tor entry and exit relay

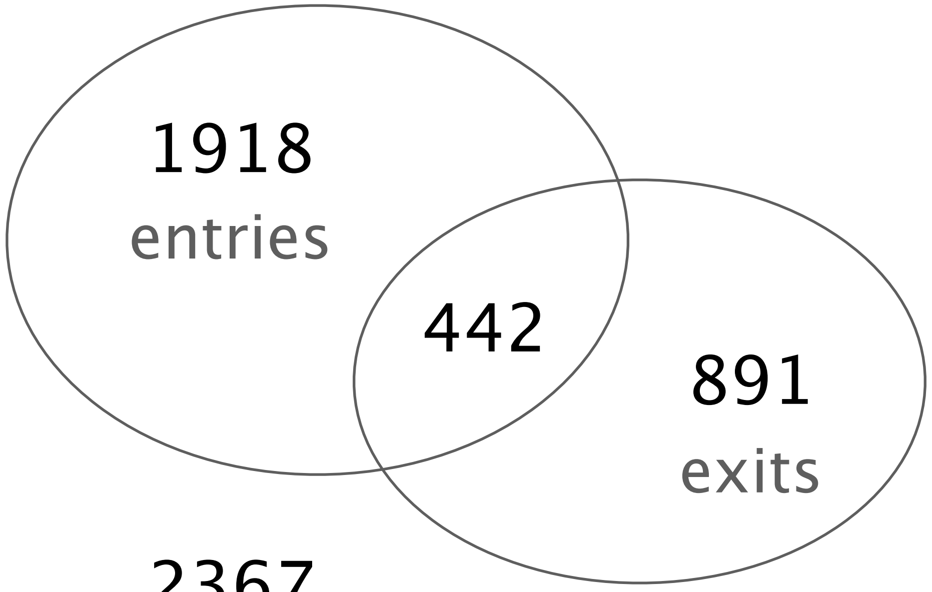

2367

entries + exists

(May’14 data)

BGP-related data

- IP address
- most-specific covering prefix
- advertising AS

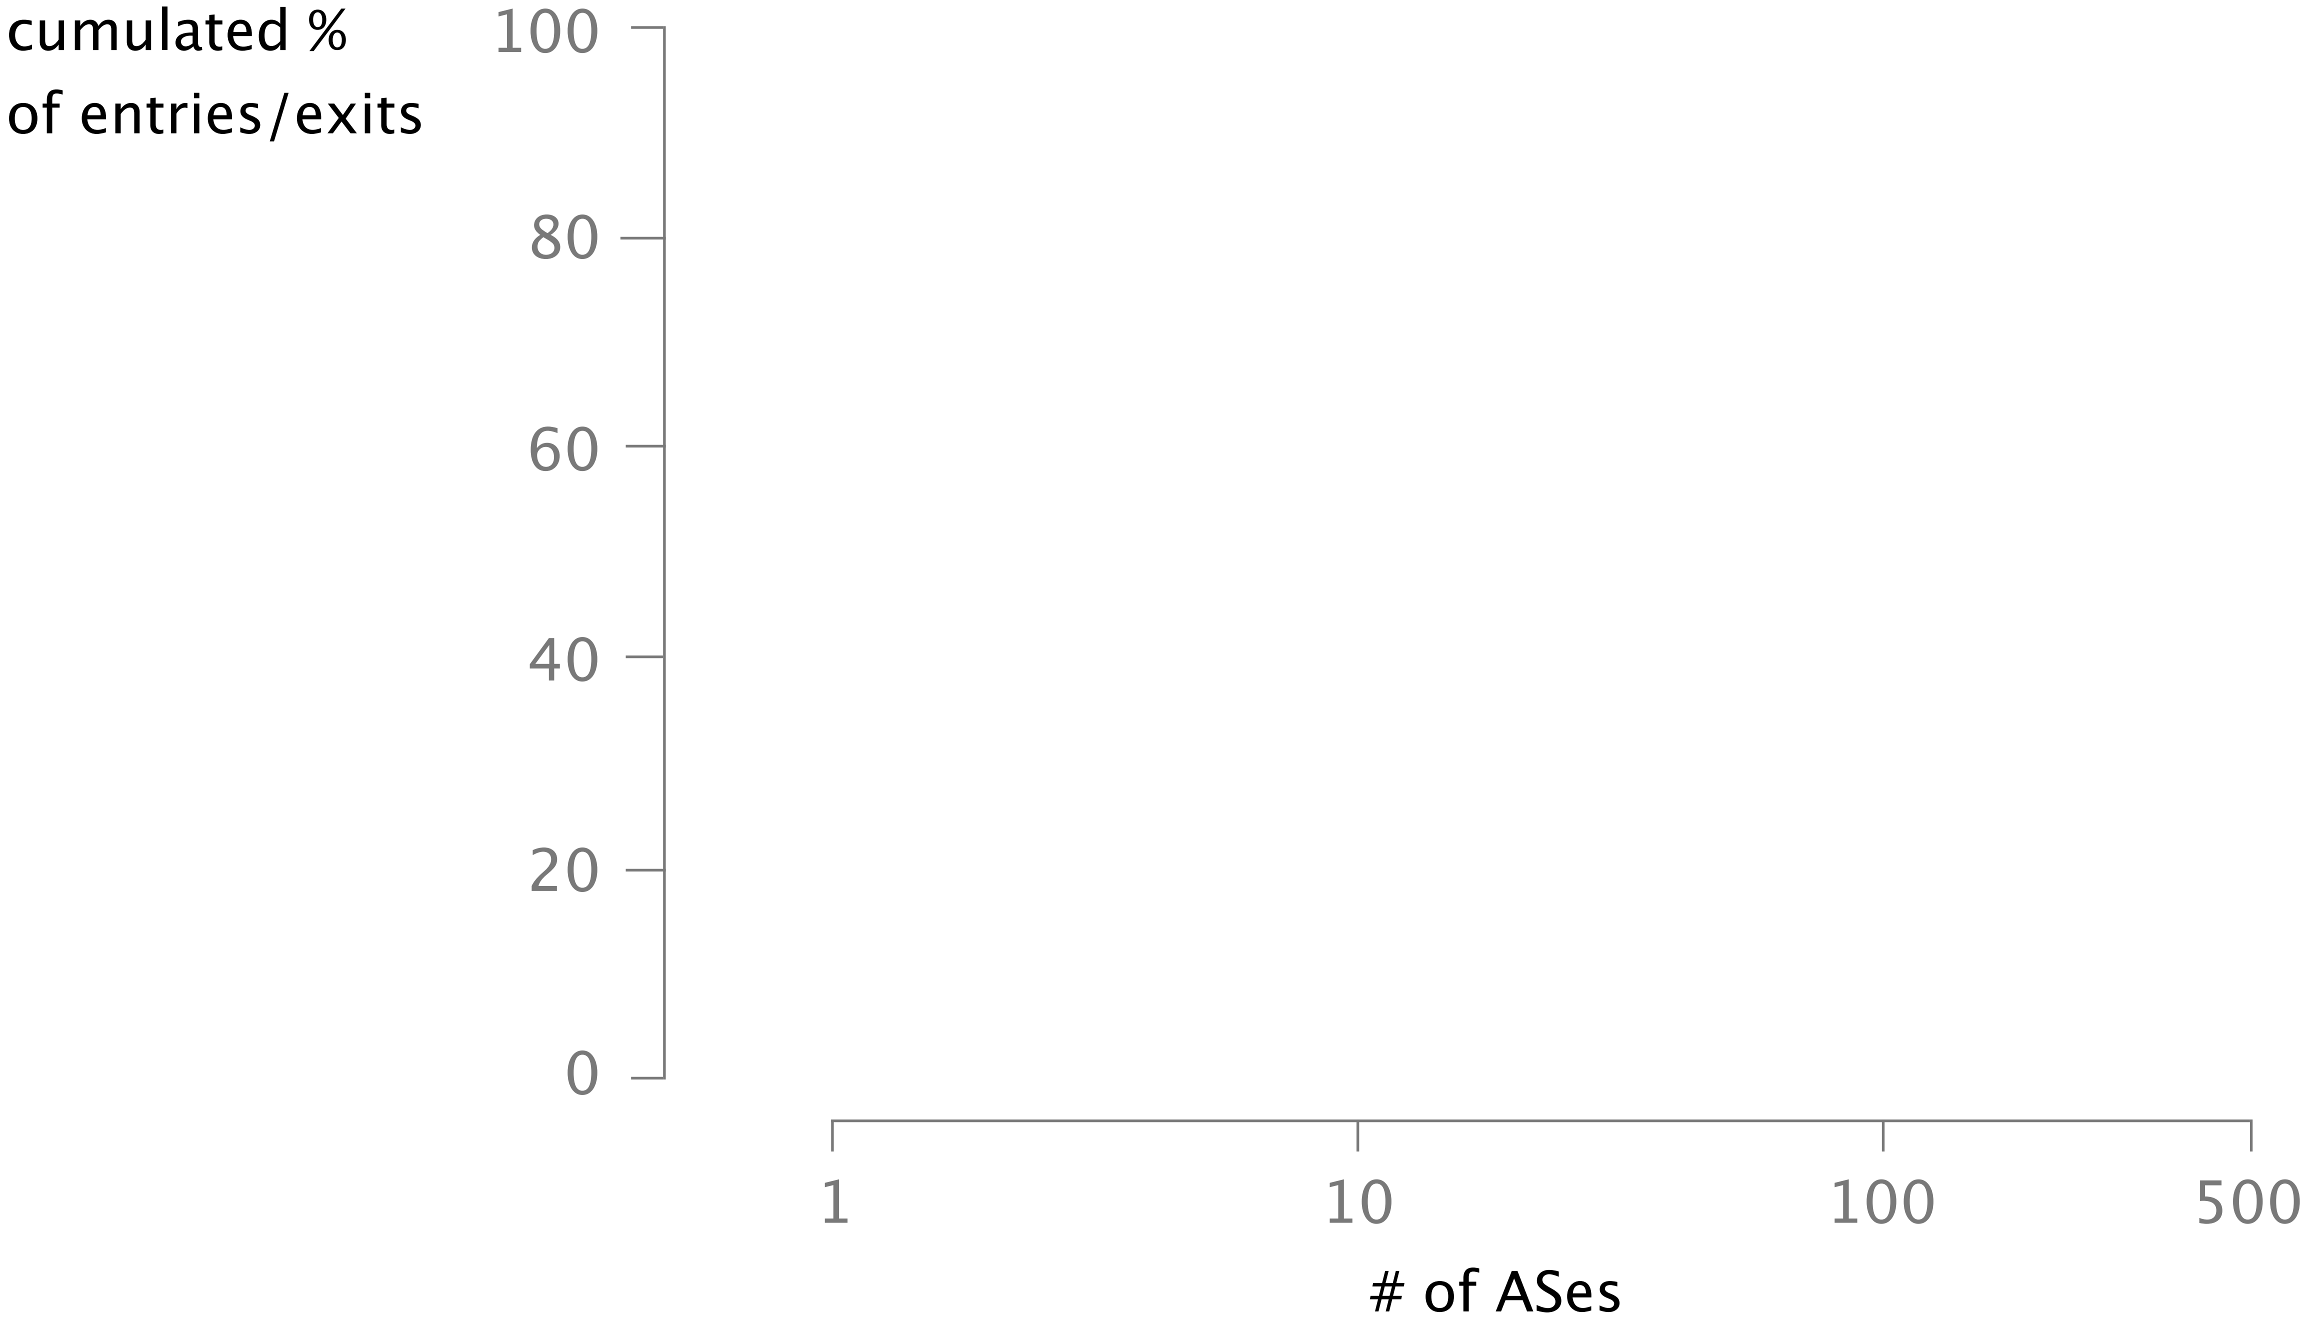
cumulated %
of entries/exits
100
80
60
40
20
0
1
10
100
500
# of ASes

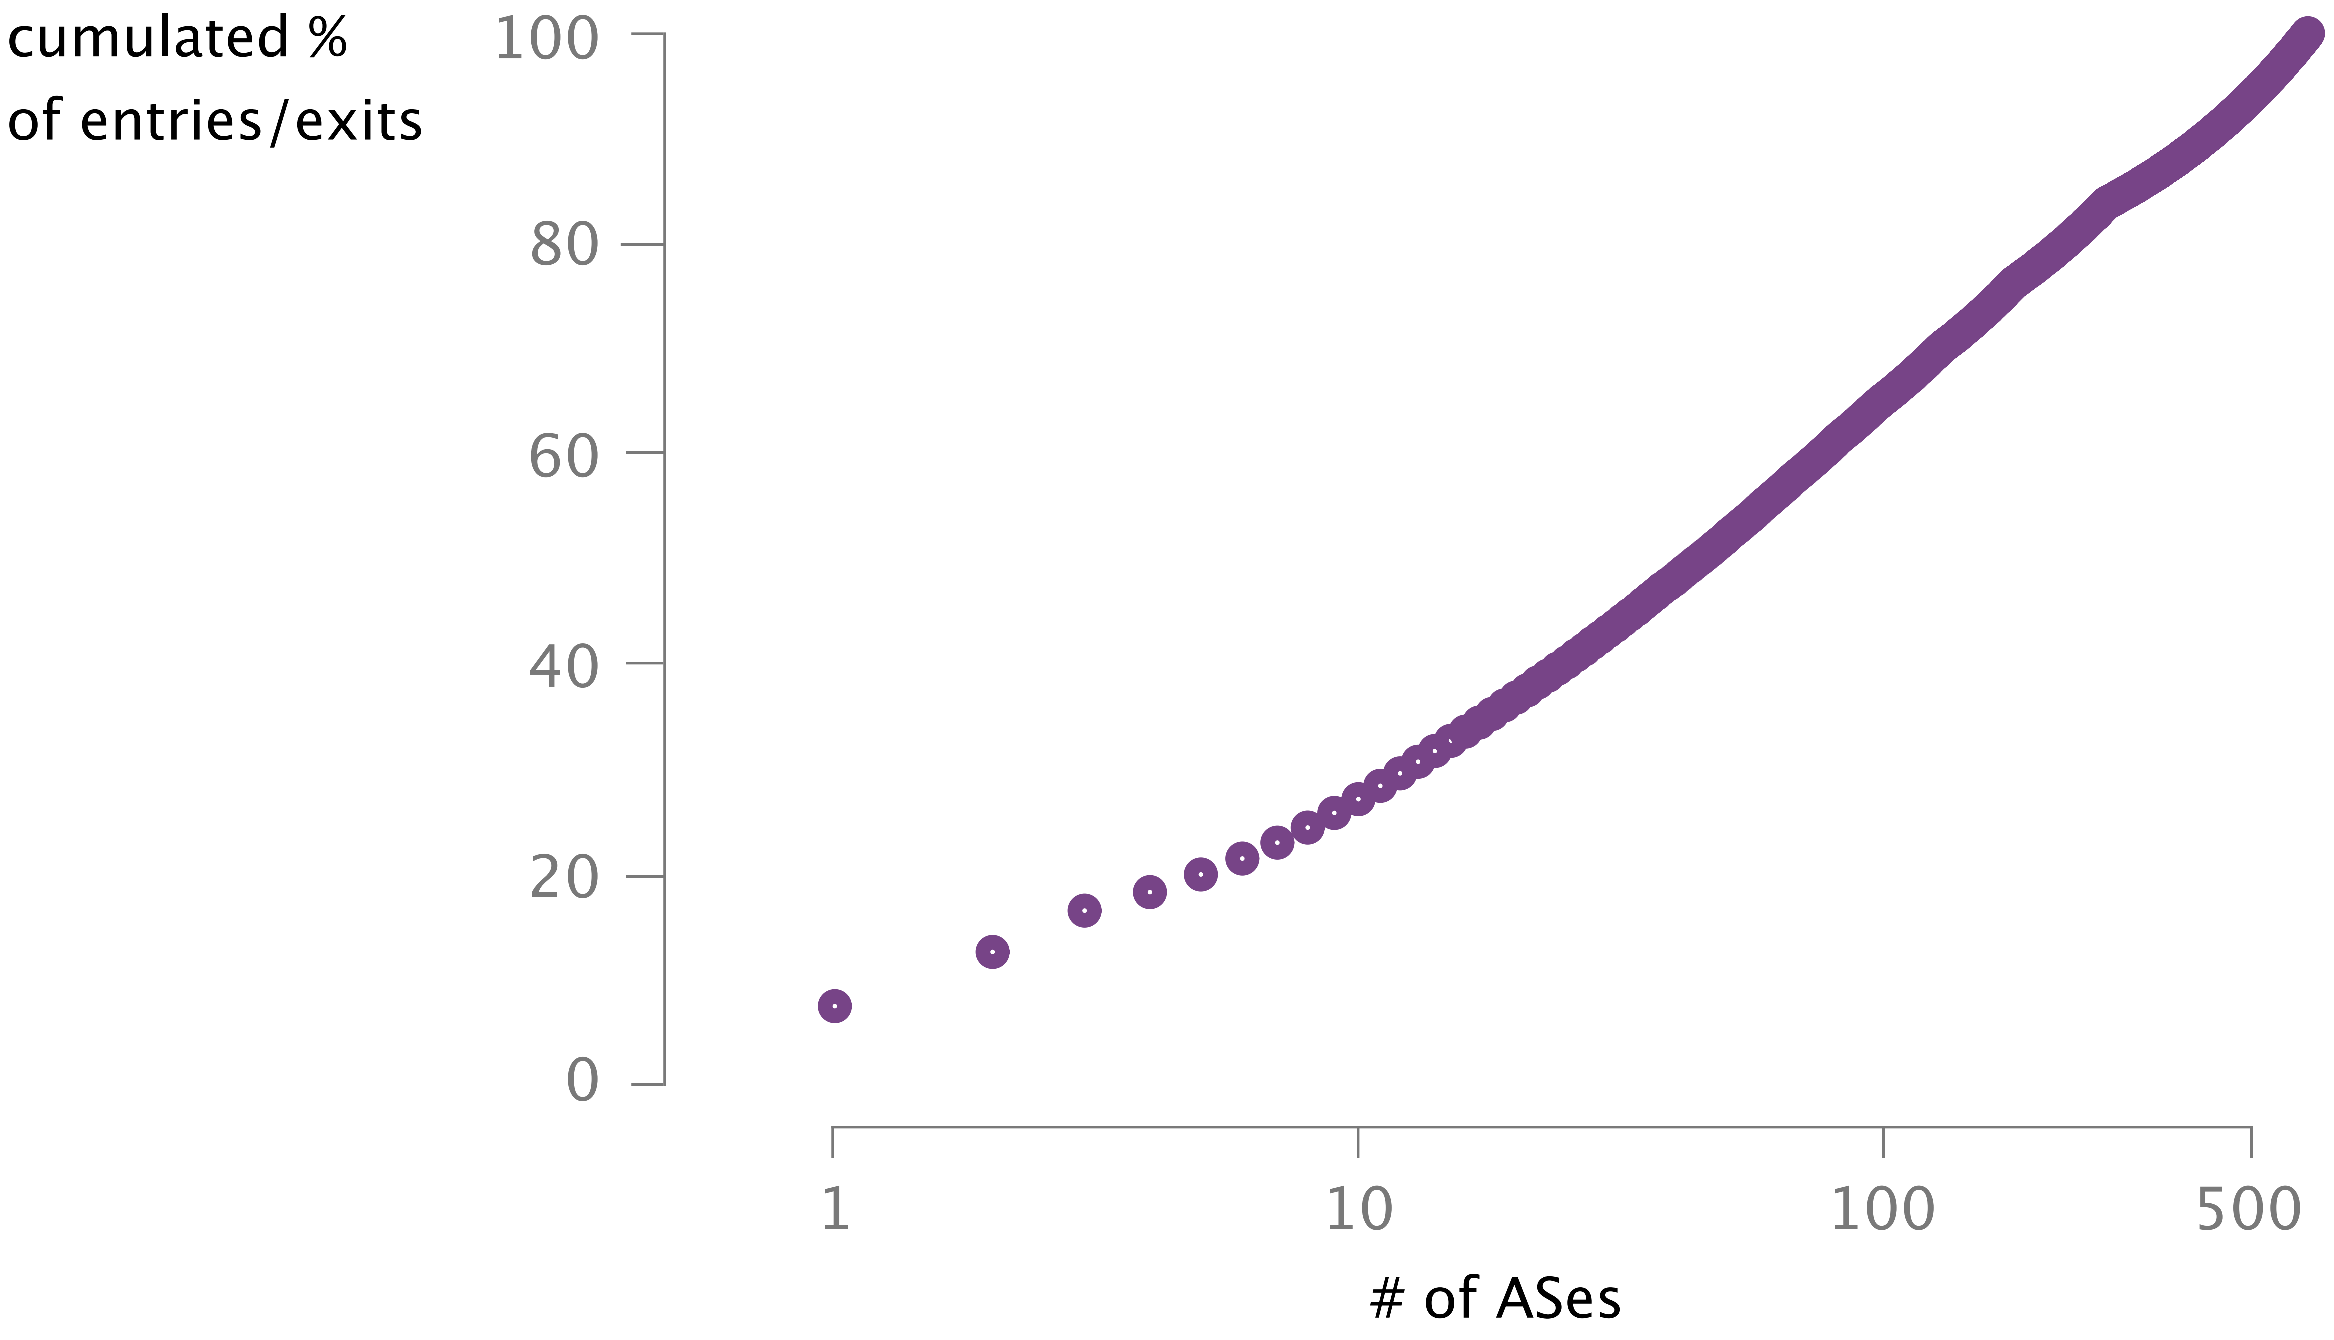
cumulated %
of entries/exits
100
80
60
40
20
0
1
10
100
500
# of ASes

Entry & exit relays
are concentrated in few ASes

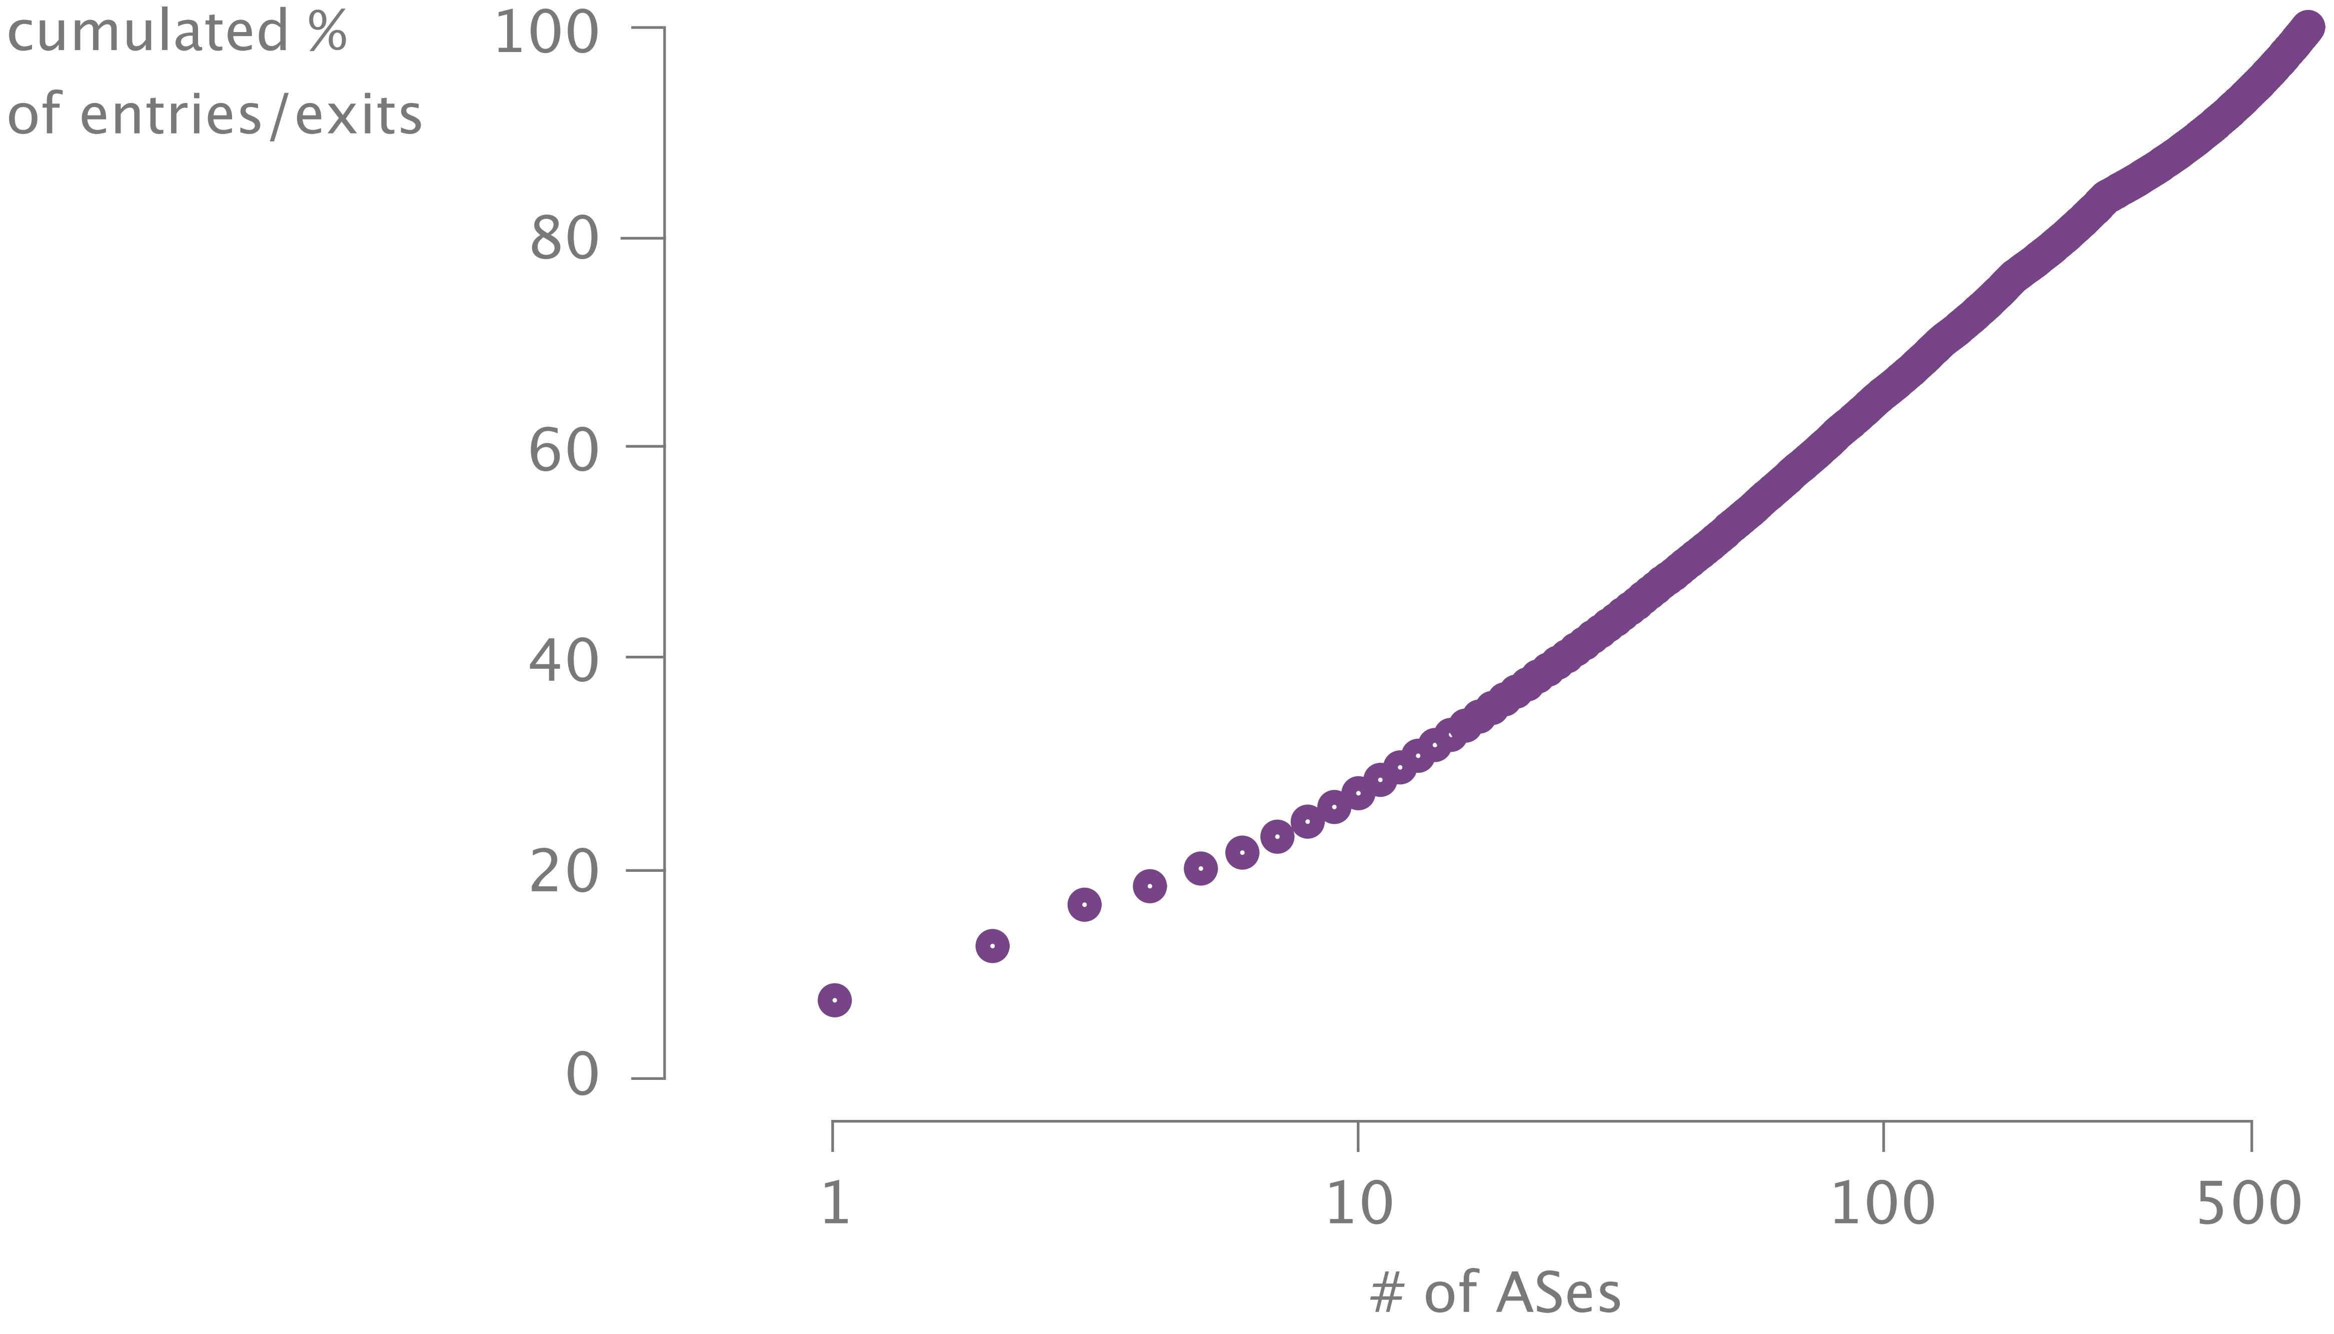

Entry & exit relays are concentrated in few ASes

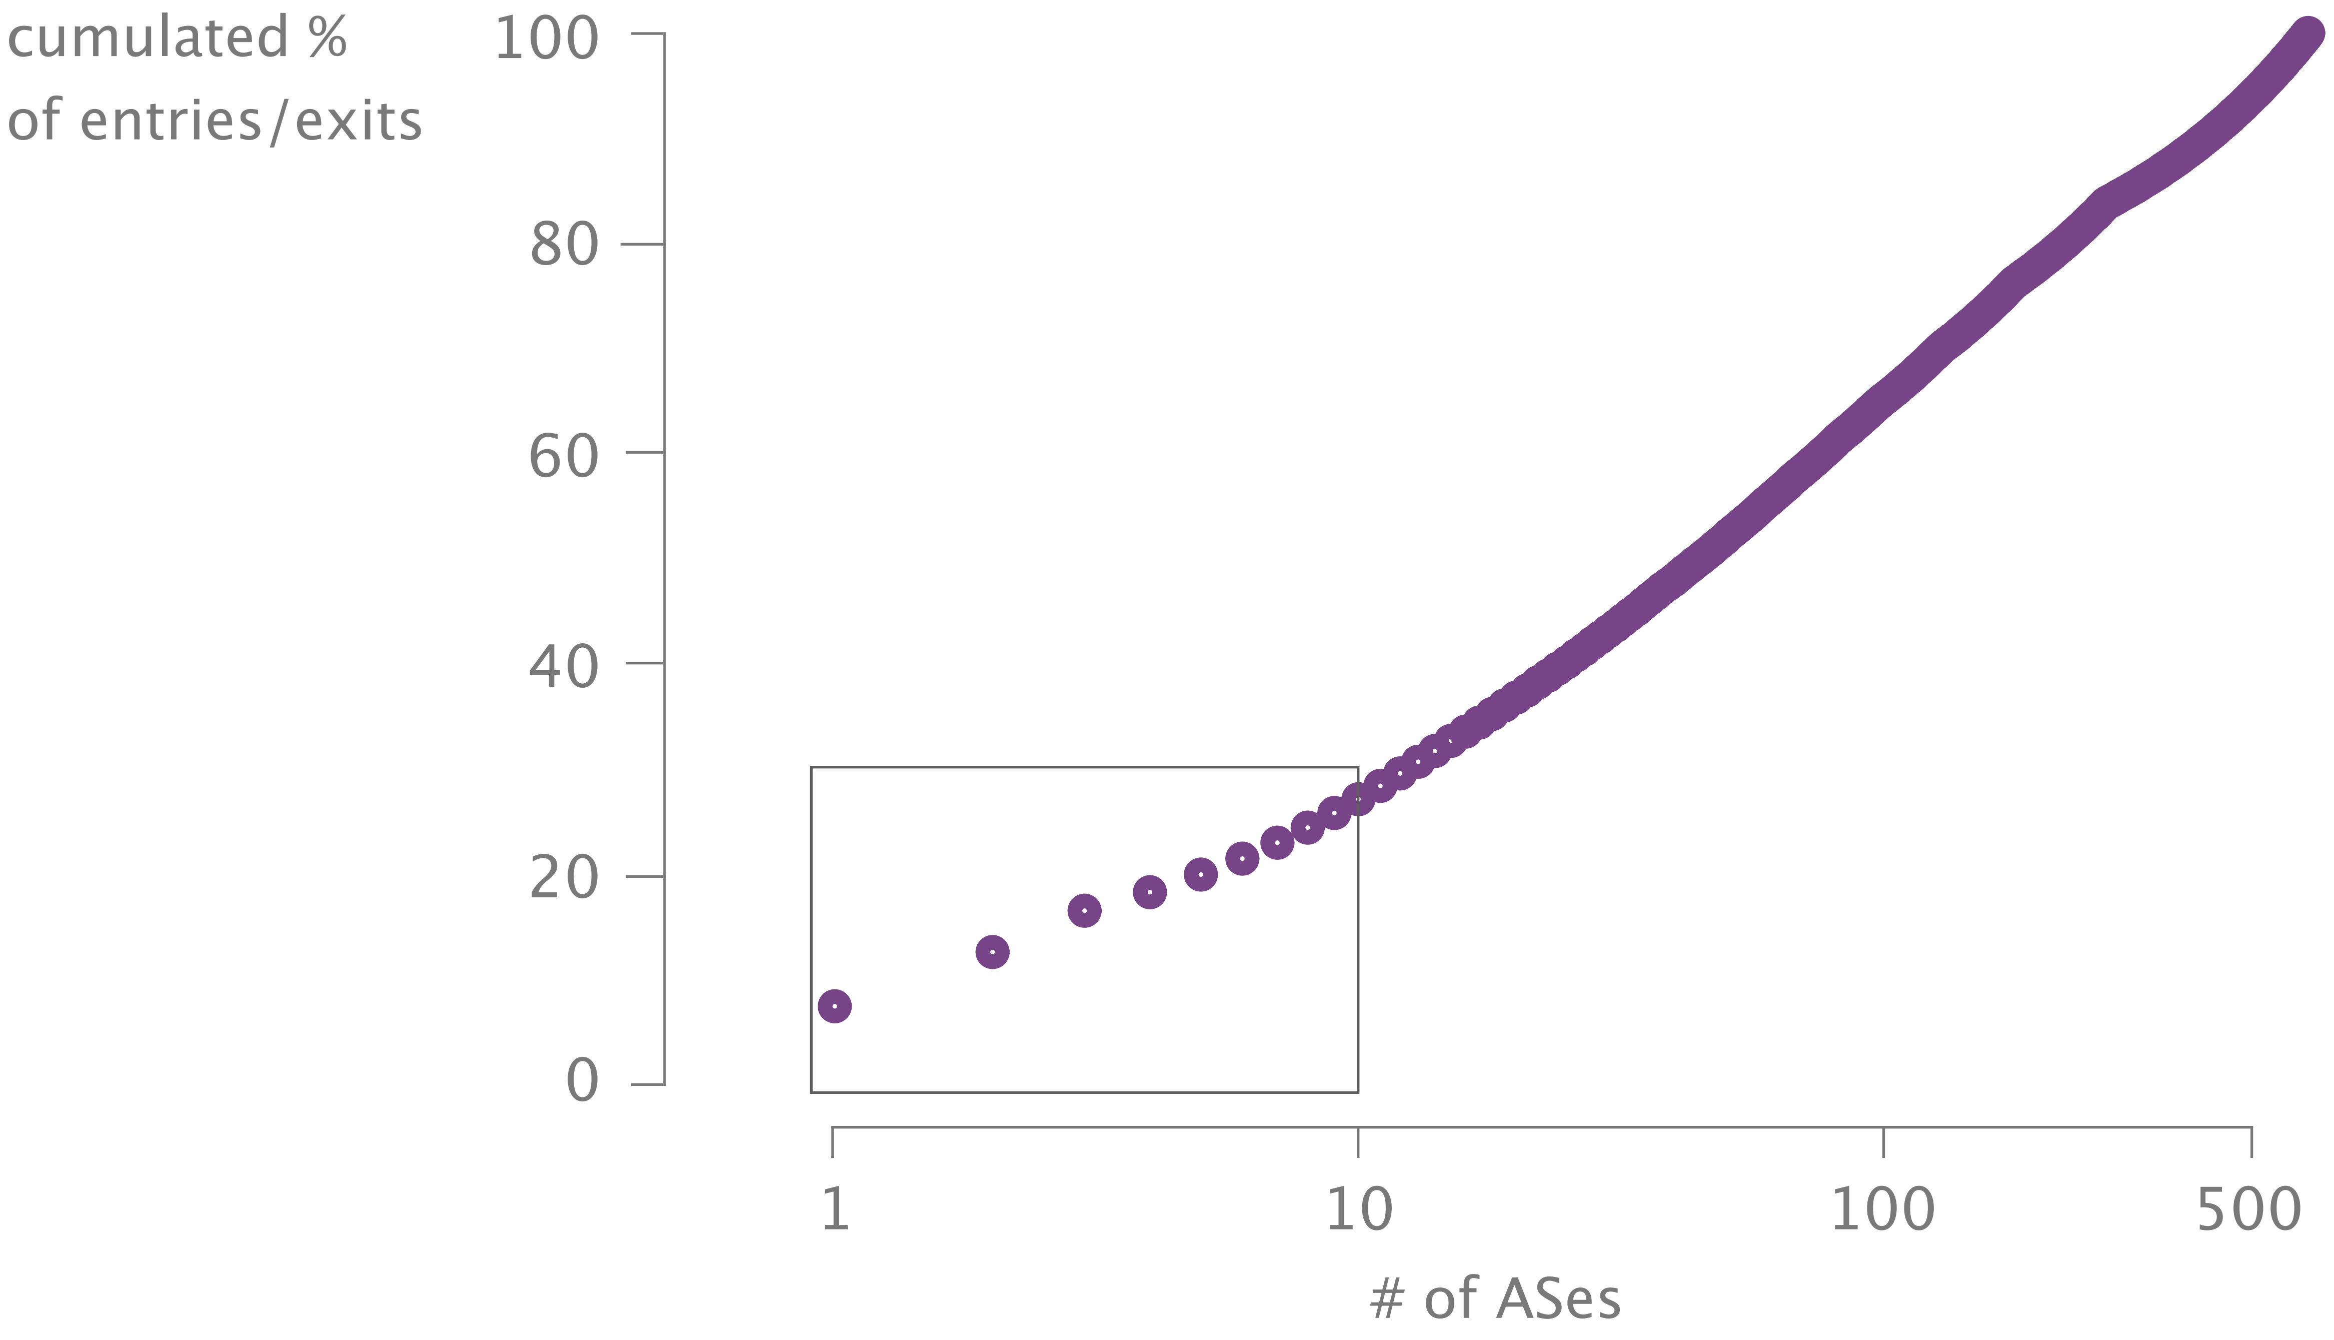

3 ASes host close to 20% of the entry & exit relays

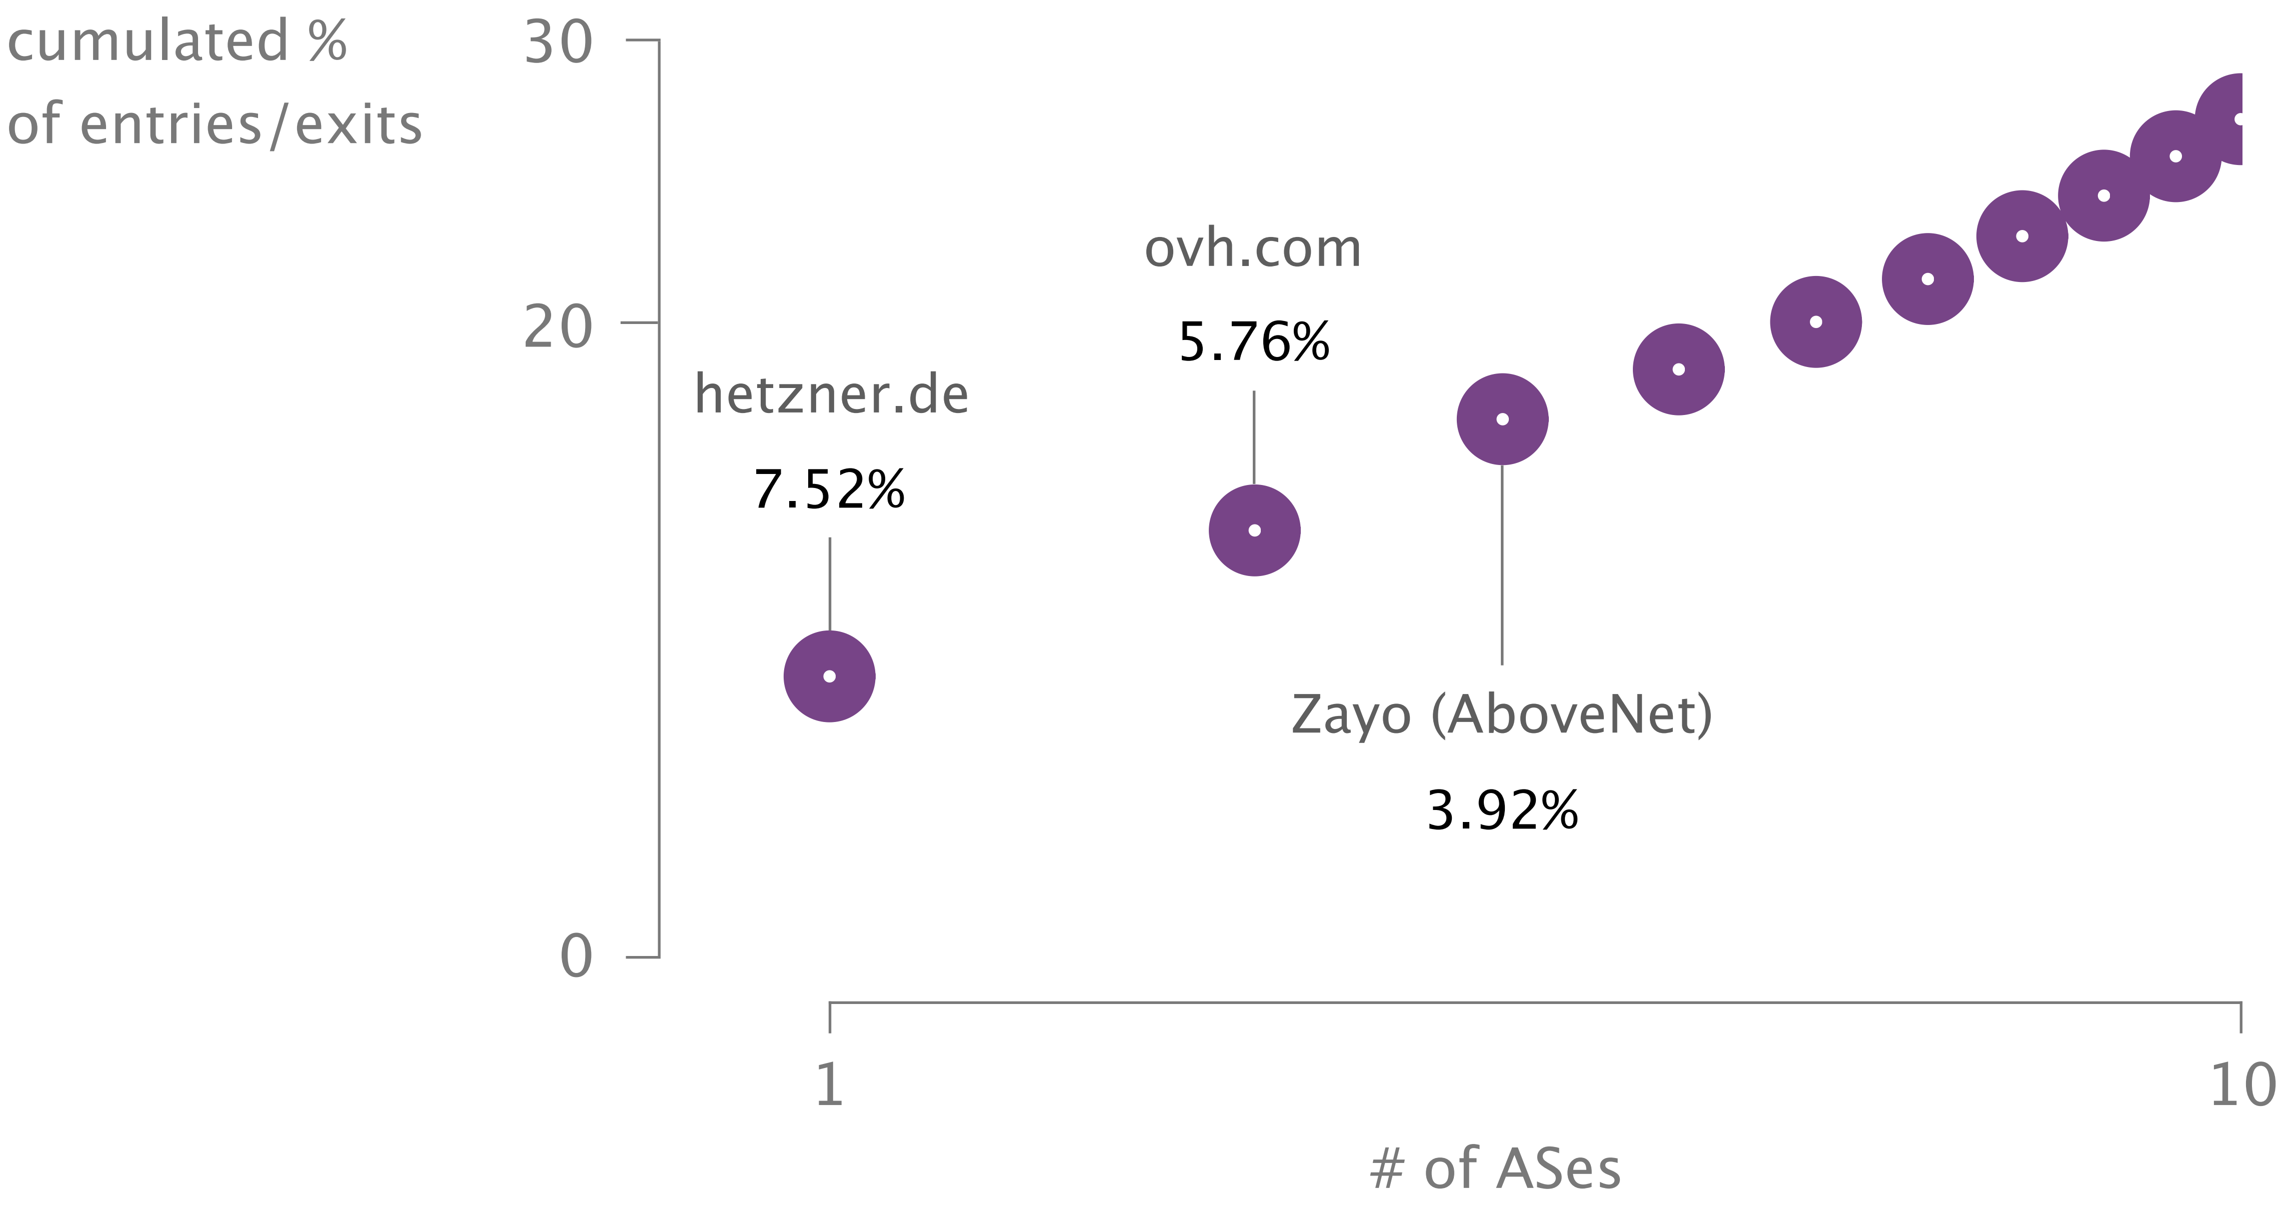

Question#2: How much path changes were Tor prefixes seeing with respect to BGP prefixes?

To measure the effect of BGP dynamics we collected BGP updates over 1 month

# BGP sessions	71
(RIPE RIS collectors)	
# BGP prefixes	1.2k
advertised by	650 ASes
# BGP updates	1.4M
announcements/withdraws	

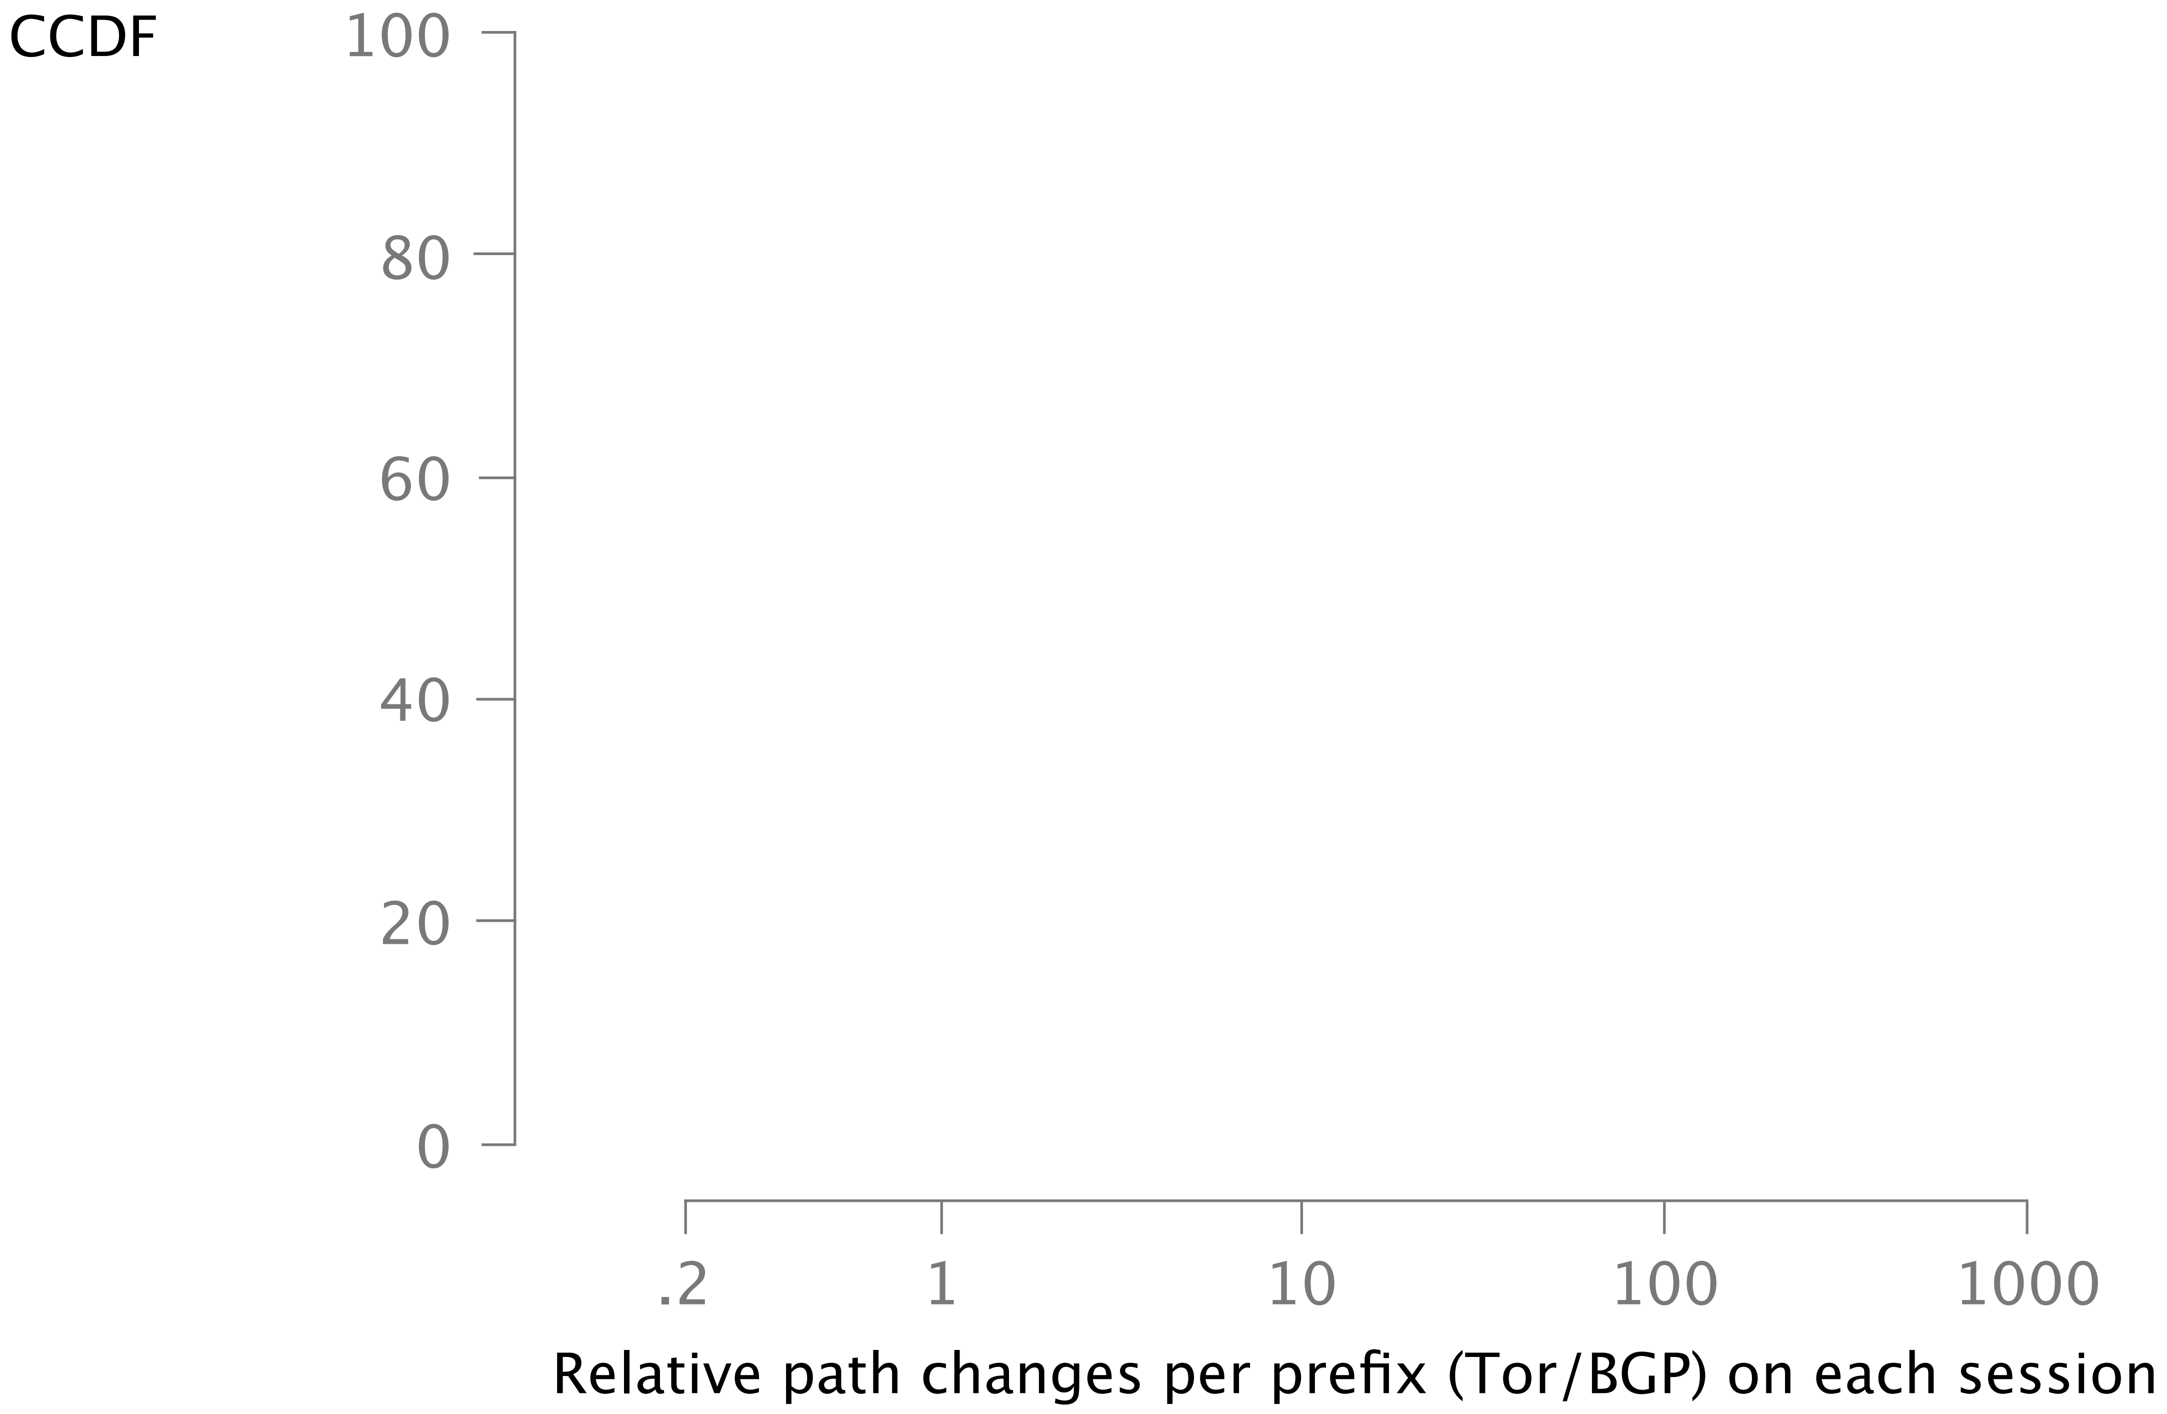
CCDF
100
80
60
40
20
0
.2
1
10
100
1000
Relative path changes per prefix (Tor/BGP) on each session

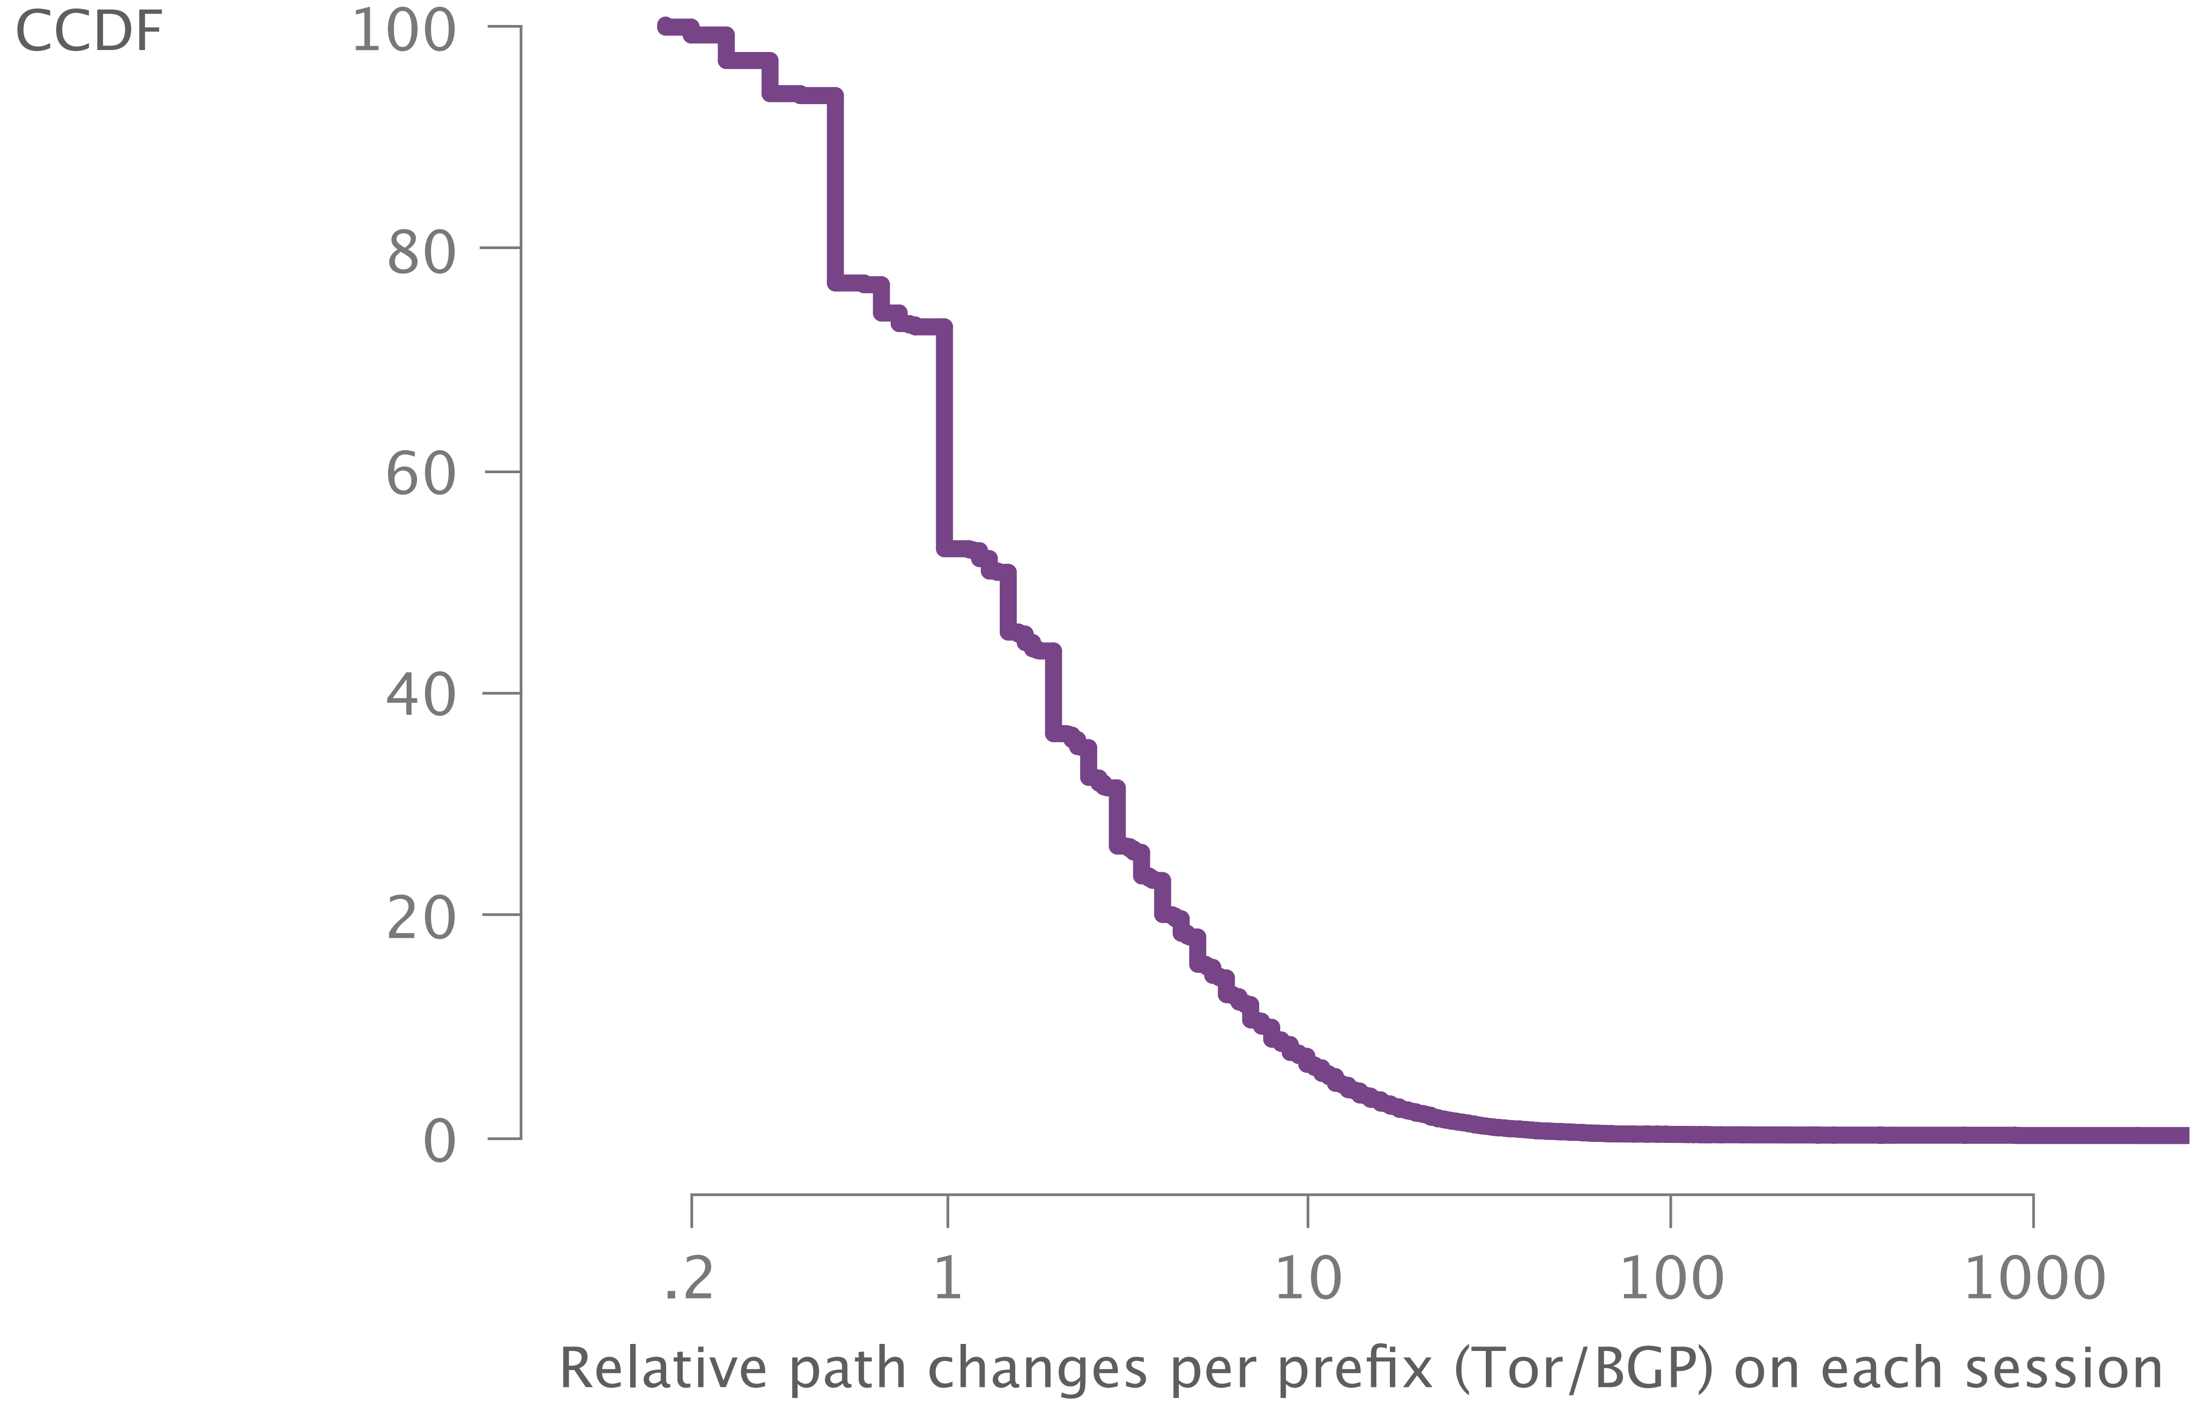
CCDF
100
80
60
40
20
0
.2
1
10
100
1000
Relative path changes per prefix (Tor/BGP) on each session

In 25% of the cases, Tor prefixes saw >3.5 more changes than BGP prefixes on a session

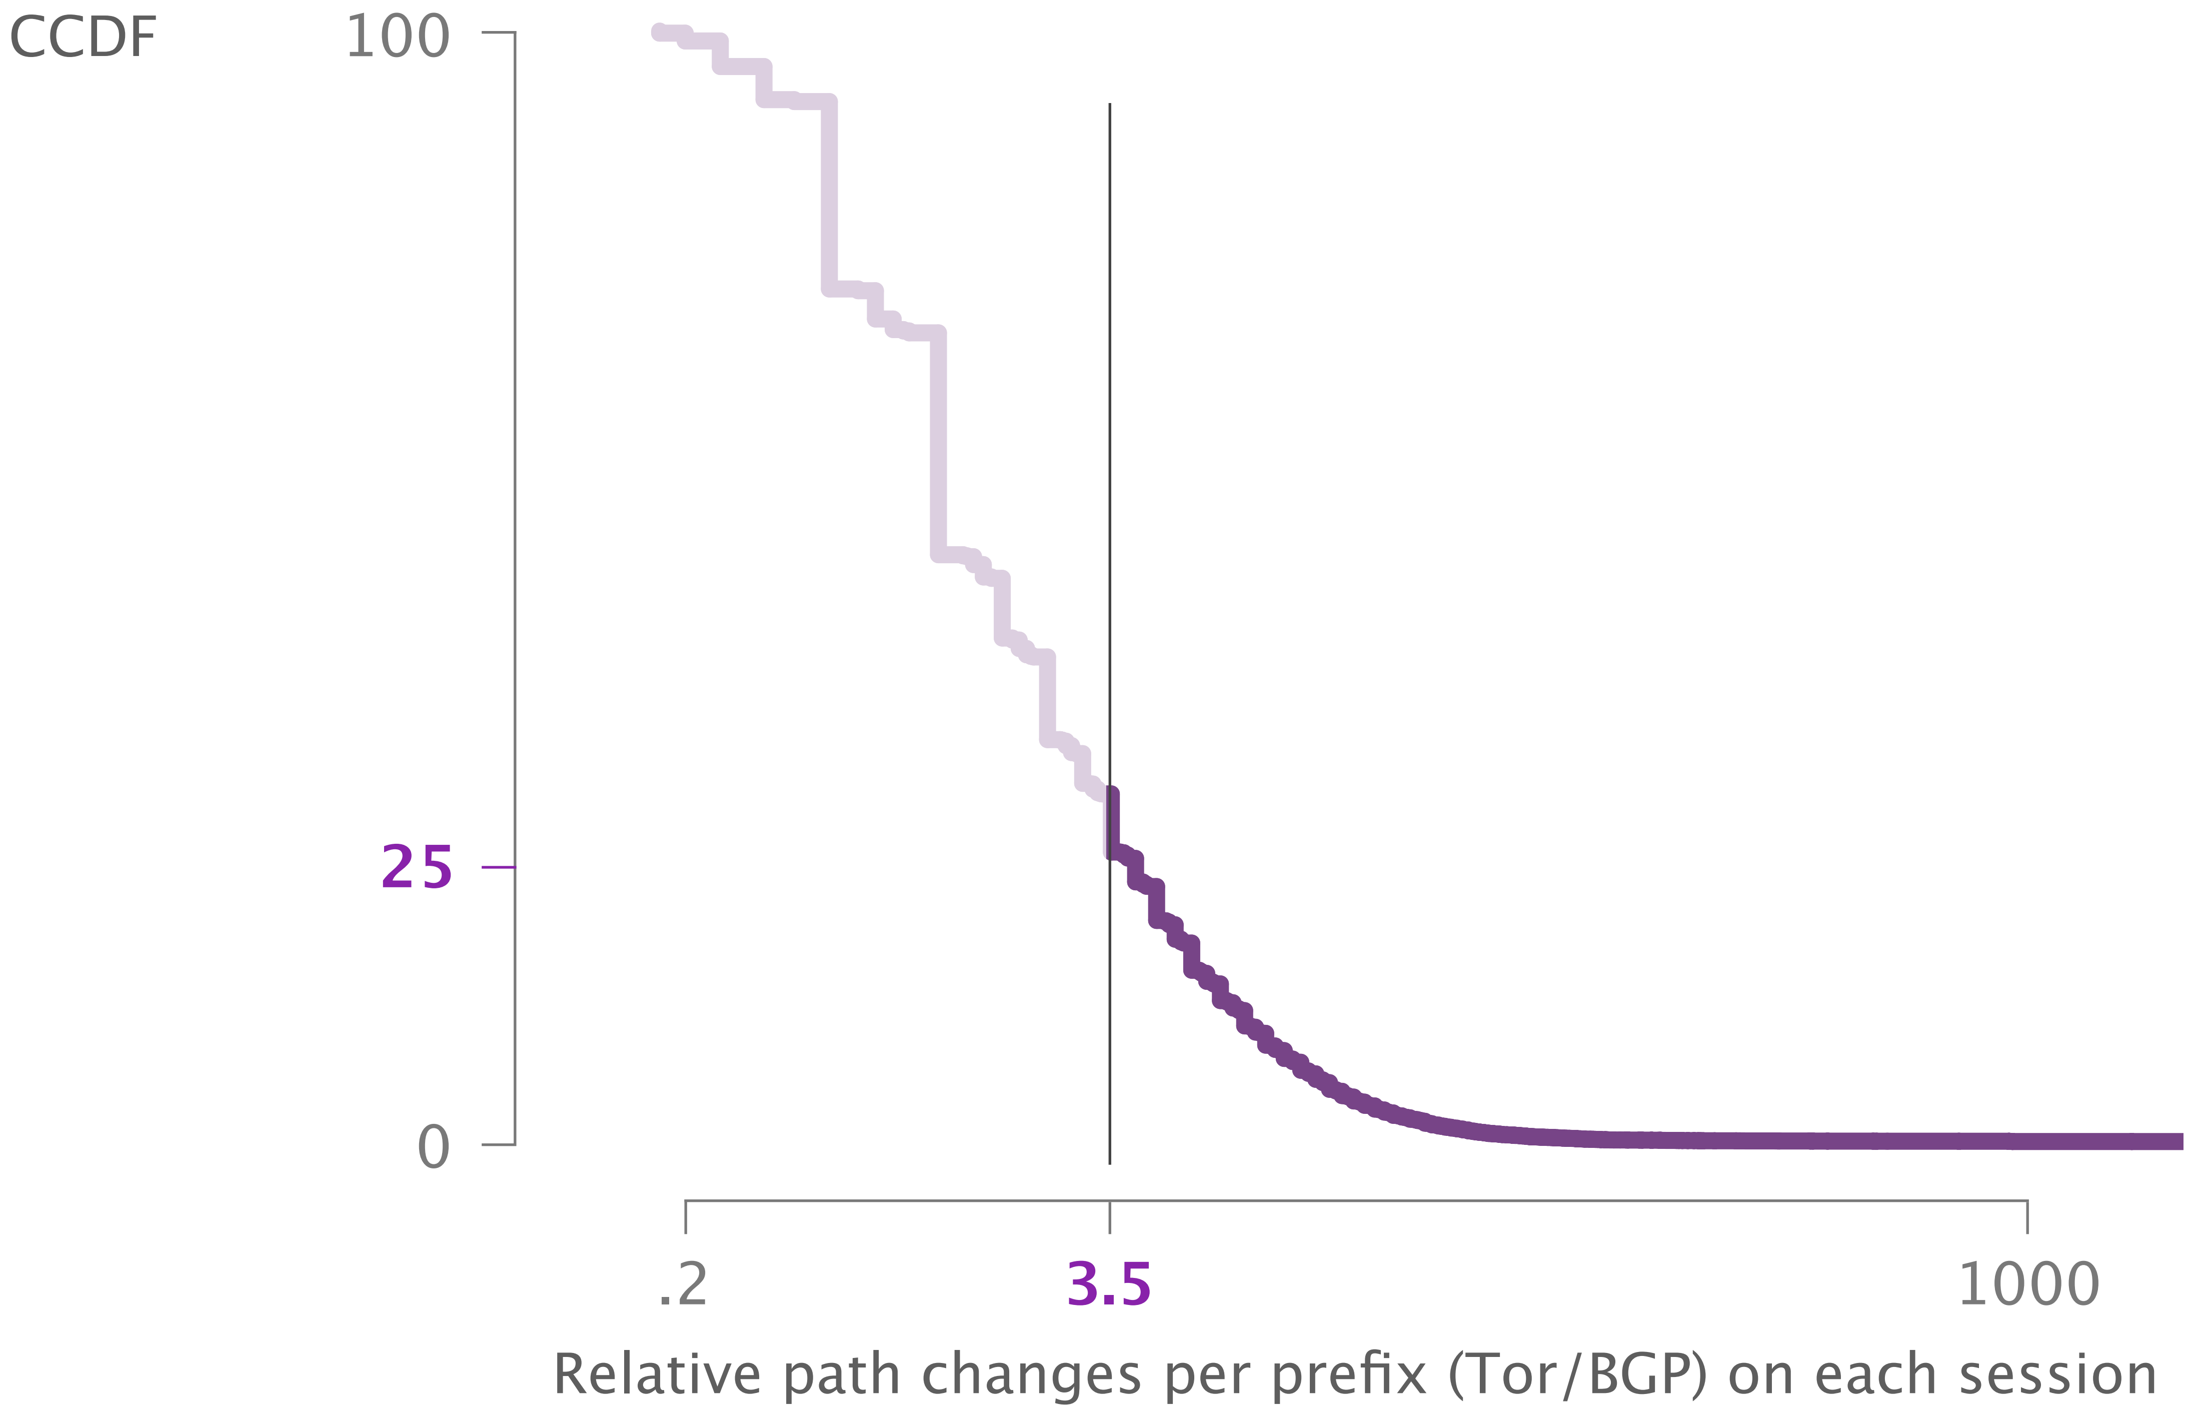

These changes caused a bunch of extra ASes to see Tor traffic

In **60%** of the cases, **>2 extra ASes** receive traffic
over the month because of BGP dynamics

significant as the average # of ASes per path is ~4

Anonymity on Quicksand

Using BGP to compromise Tor

Attacks

All your traffic belongs to me

Preliminary results

Eyes wide open

3 Countermeasures

Close the curtains

To protect itself, Tor should become more aware of the network underlying it

Problems	Countermeasures	Tools
Natural dynamism	prefer stable relays	BGP monitoring
Route manipulation	discard “suspicious” relays prefer close relays	BGP monitoring + BGPsec
Asymmetric analysis	encrypt transport header	IPsec

These countermeasures help, but come with tradeoffs

Problems	Countermeasures	Tradeoffs
Natural dynamism	prefer stable relays	more power to fewer relays
Route manipulation	discard “suspicious” relays prefer close relays	
Asymmetric analysis	encrypt transport header	not widely used (easier to detect)

Anonymity on Quicksand

Using BGP to compromise Tor

Attacks

All your traffic belongs to me

Preliminary results

Eyes wide open

Countermeasures

Close the curtains

BGP is not only a problem for Tor…

THREAT LEVEL

Hacker Redirects Traffic From 19 Internet Providers to Steal Bitcoins

BY ANDY GREENBERG 08.07.14 | 1:00 PM | PERMALINK

… A bitcoin thief redirected a portion of online traffic from no less than 19 Internet service providers, including data from the networks of Amazon and other hosting services like DigitalOcean and OVH, with the goal of stealing cryptocurrency from a group of bitcoin users…

… A bitcoin thief redirected a portion of online traffic from no less than 19 Internet service providers, including data from the networks of Amazon and other hosting services like DigitalOcean and OVH, with the goal of stealing cryptocurrency from a group of bitcoin users…

OVH is the second AS in terms of # Tor relays hosted

… A bitcoin thief redirected a portion of online traffic from no less than 19 Internet service providers, including data from the networks of Amazon and other hosting services like DigitalOcean and **OVH**, with the goal of stealing cryptocurrency from a group of bitcoin users…

Internet routing matters when it comes to user anonymity

BGP dynamics decreases user anonymity over time

natural & induced, exacerbated by asymmetric routing

Initial results illustrate the vulnerabilities

full evaluation is required—and underway

Short-term countermeasures helps, to an extent

need a better understanding on their impacts

Anonymity on Quicksand

Using BGP to compromise Tor

Laurent Vanbever

www.vanbever.eu

HotNets

October, 28 2014