A Multi-Perspective Analysis of Carrier-Grade NAT Deployment

ACM SIGCOMM Internet Measurement Conference 2016.

Philipp Richter, Florian Wohlfart, Narseo Vallina-Rodriguez,
Mark Allman, Randy Bush, Anja Feldmann, Christian Kreibich,
Nicholas Weaver, and Vern Paxson.

@IETF 99 Prague, Czech Republic

Paper (PDF): https://tinyurl.com/cgnatietf

IPv4 Address Space Exhaustion

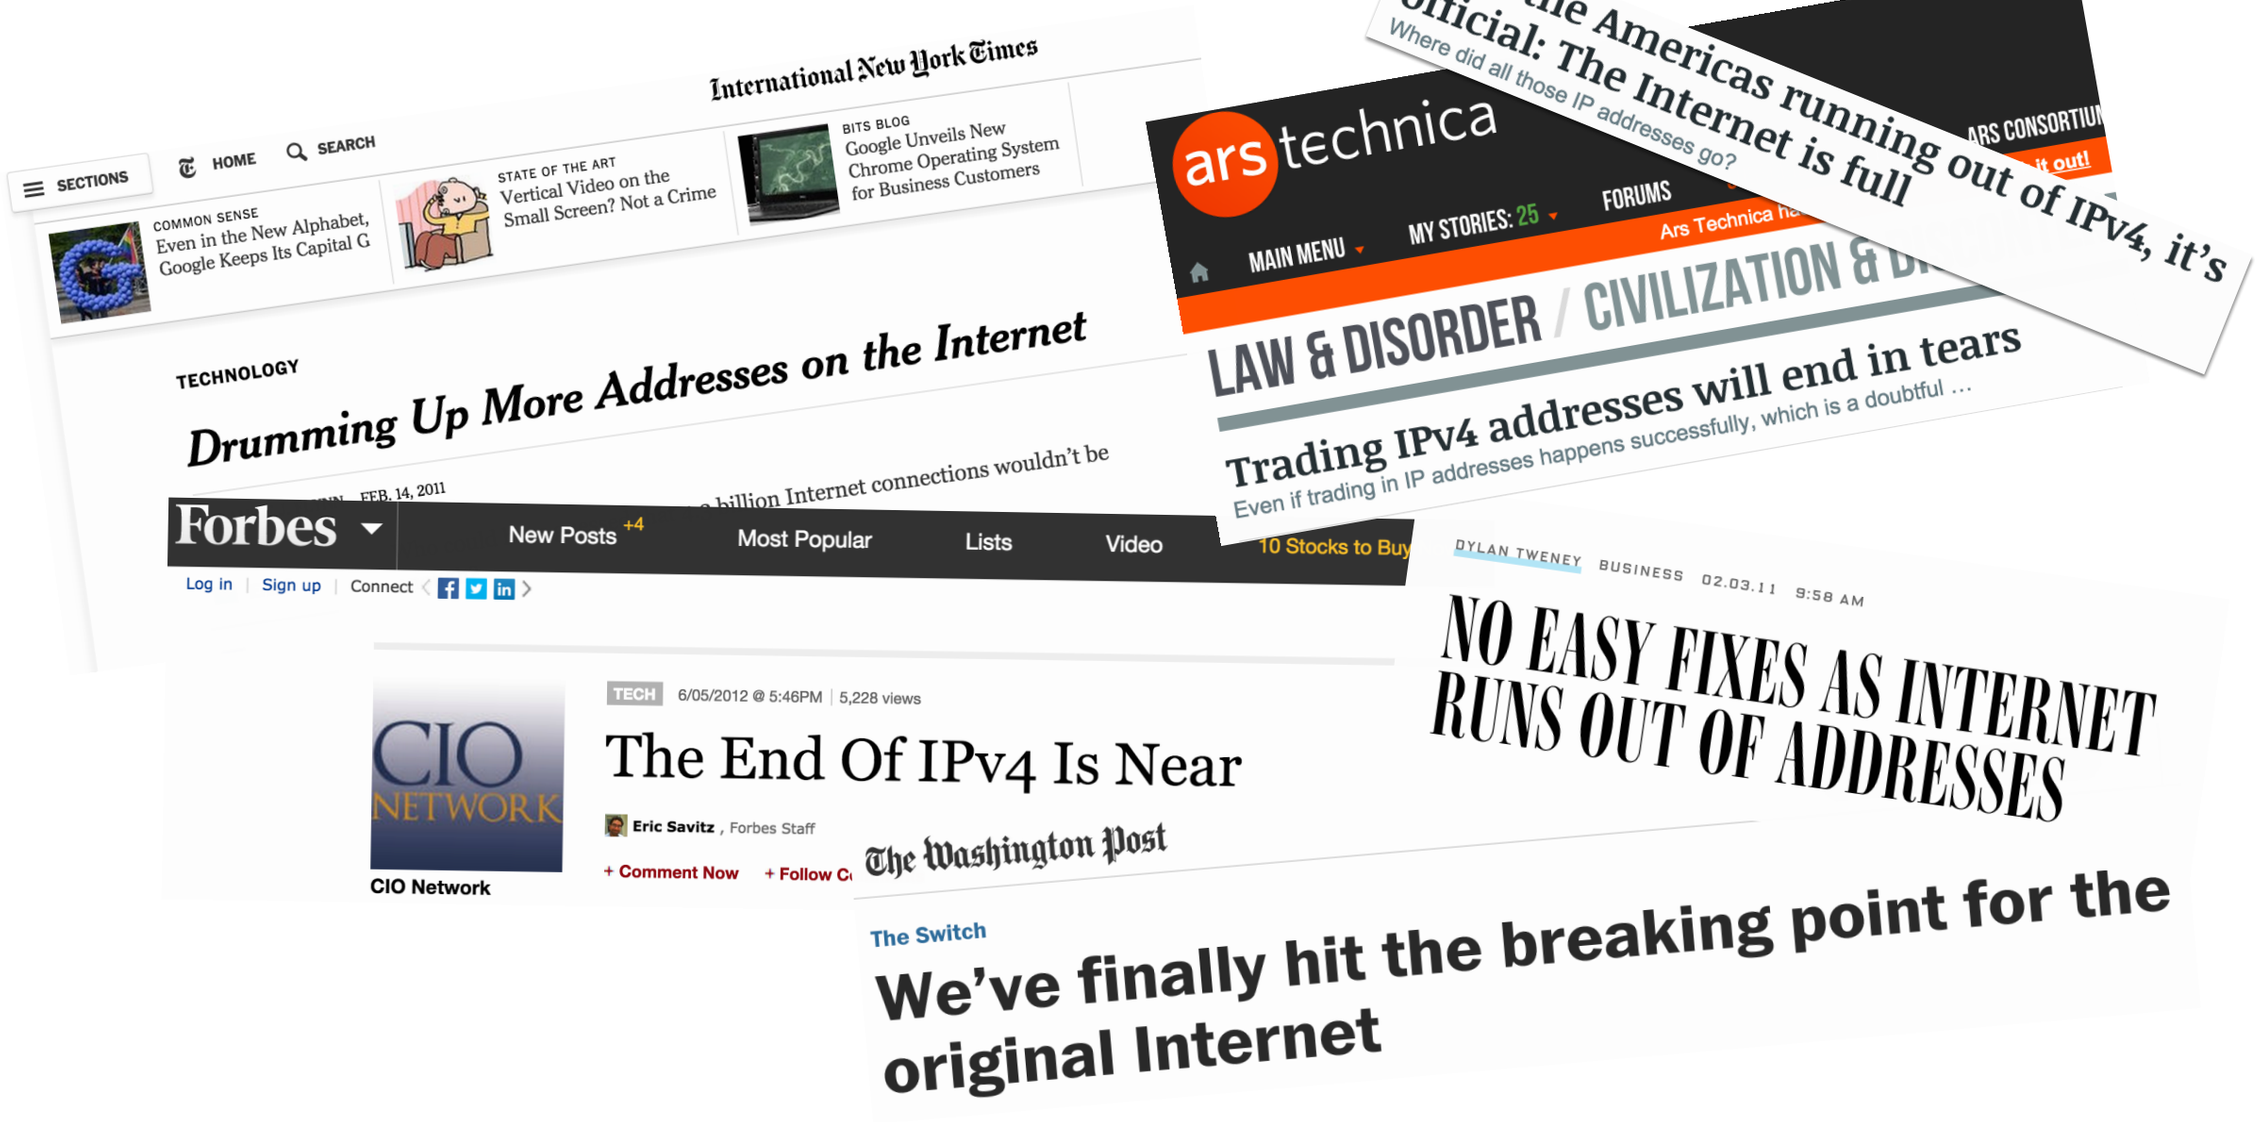

IPv4 Address Space Exhaustion

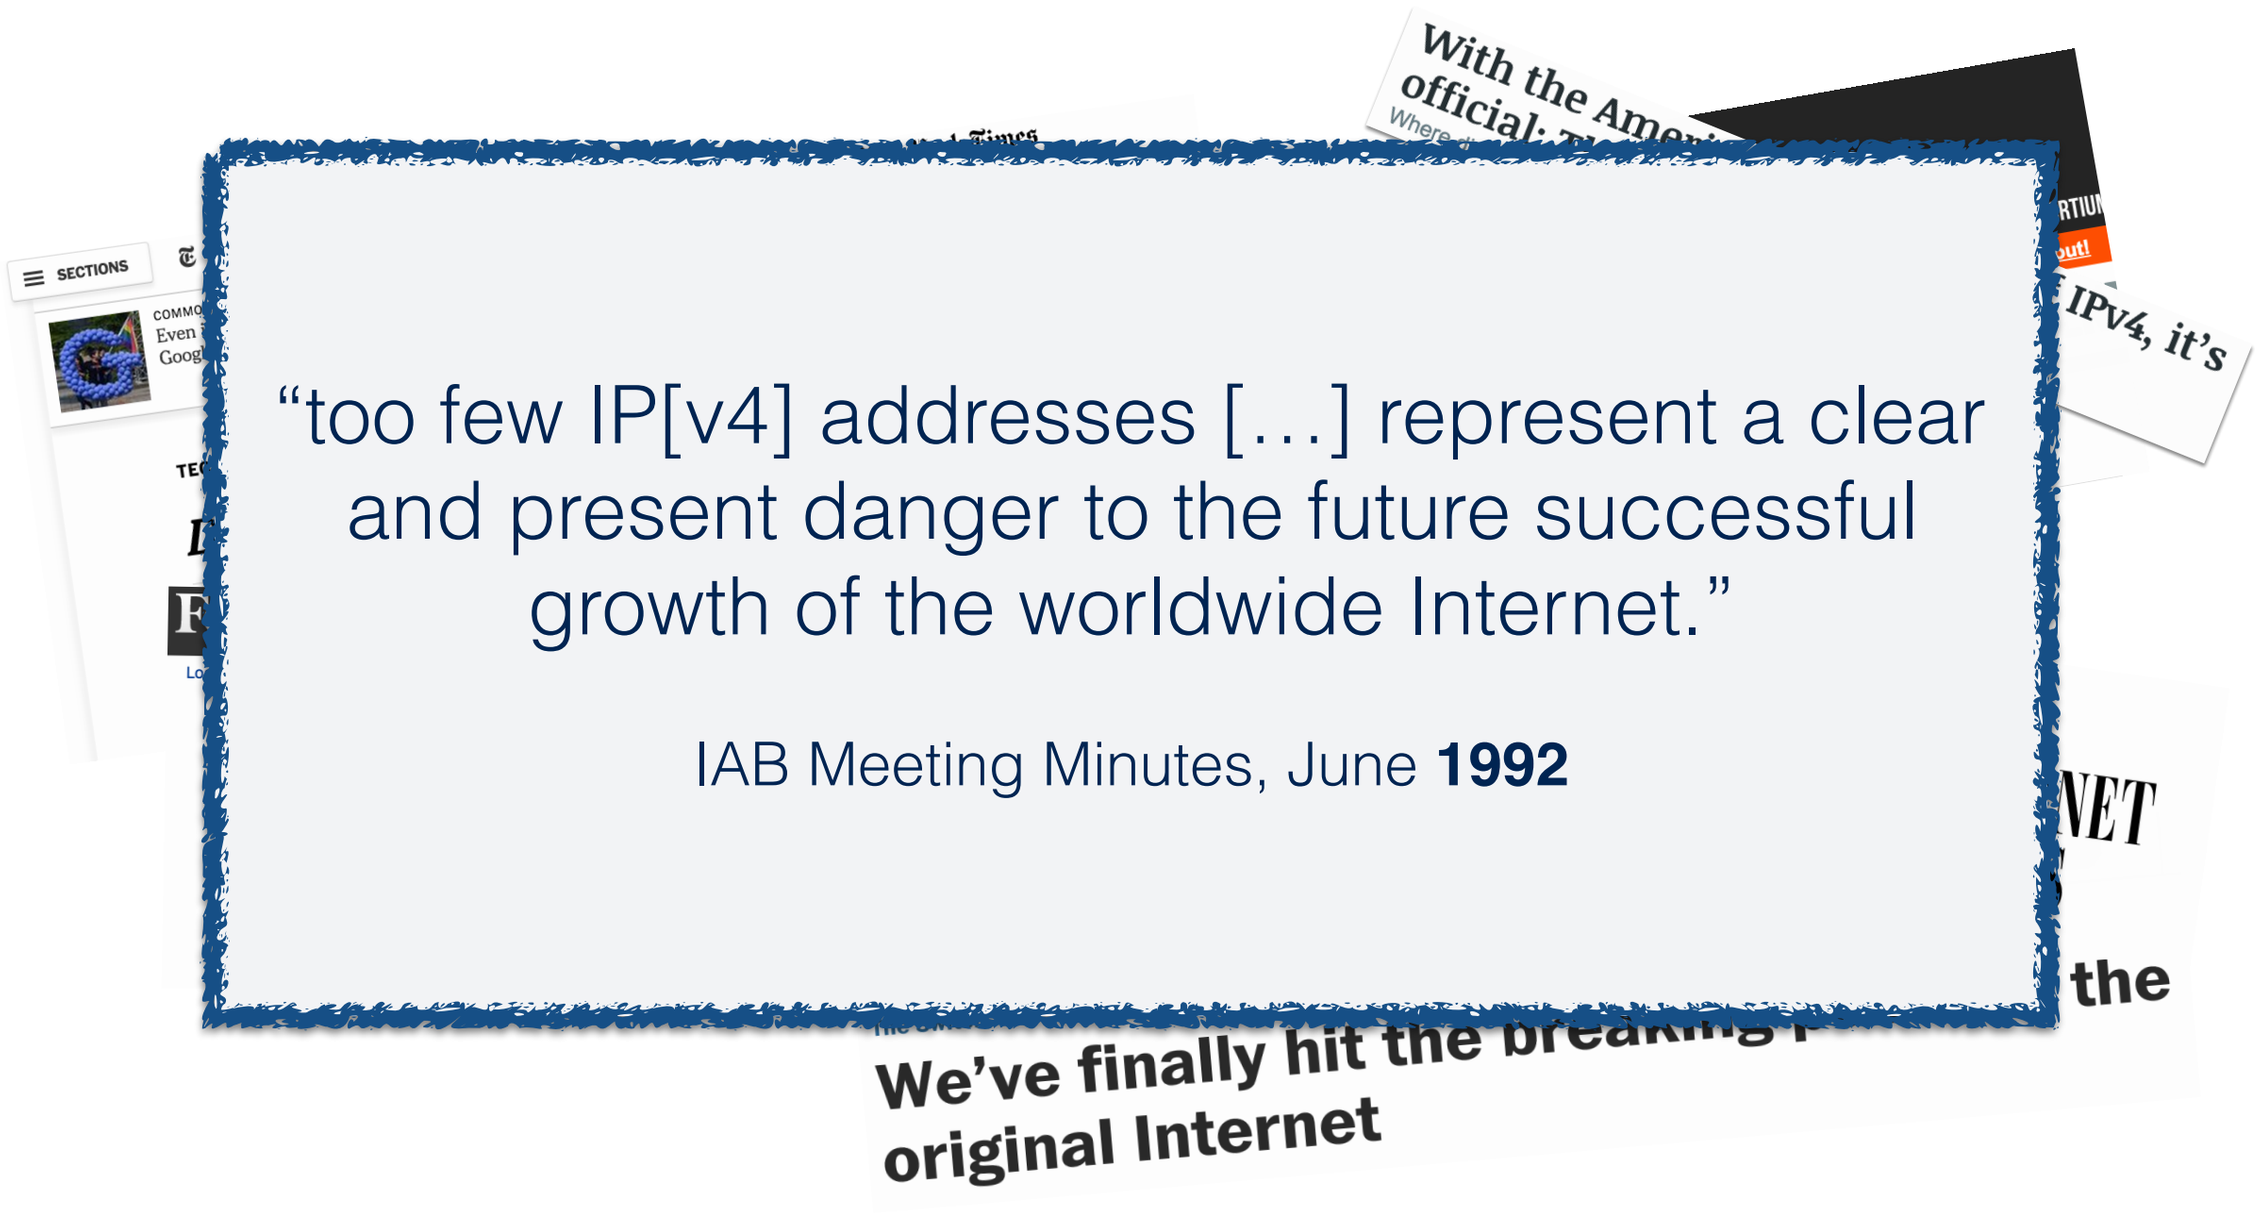

"too few IP[v4] addresses […] represent a clear and present danger to the future successful growth of the worldwide Internet."

IAB Meeting Minutes, June **1992**

Fast Forward to 2017

2017

4 out of 5 RIRs exhausted.

About ~1% of the IPv4 space left unallocated.

Fast Forward to 2017

2017

4 out of 5 RIRs exhausted.

About ~1% of the IPv4 space left unallocated.

Transition to IPv6
→ plenty of measurements and statistics available

All Eyes on IPv6

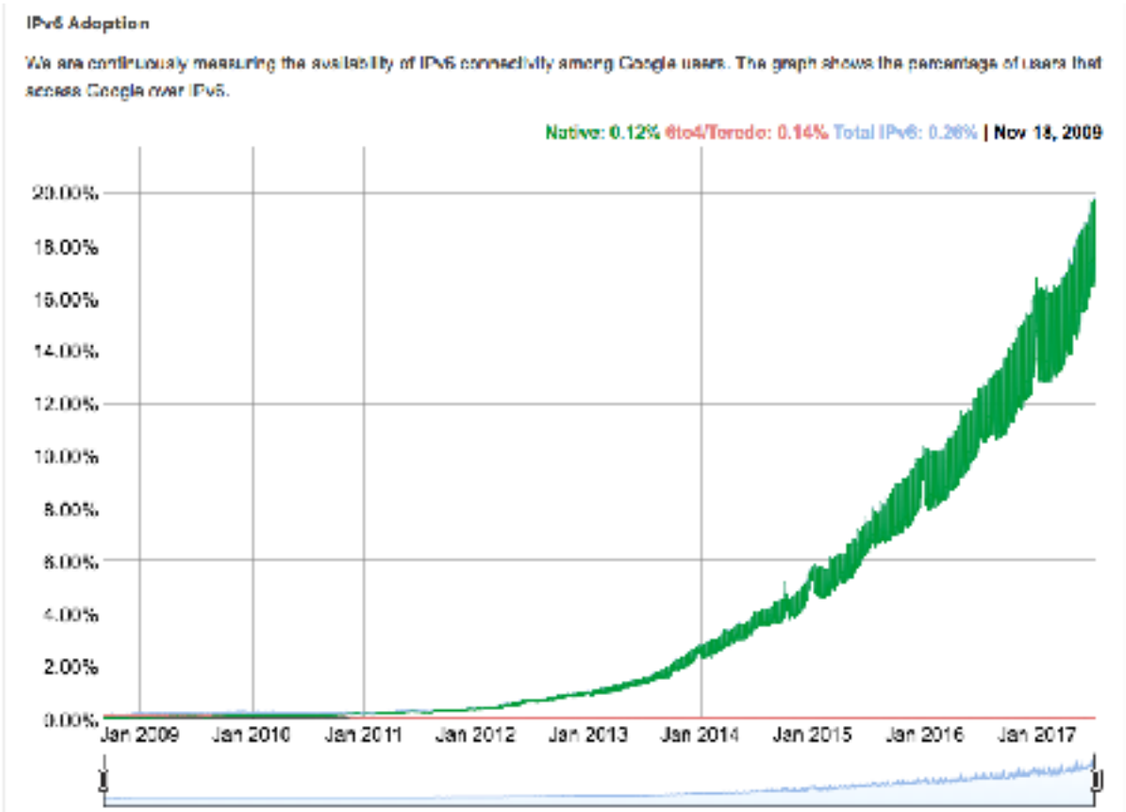

source: Google

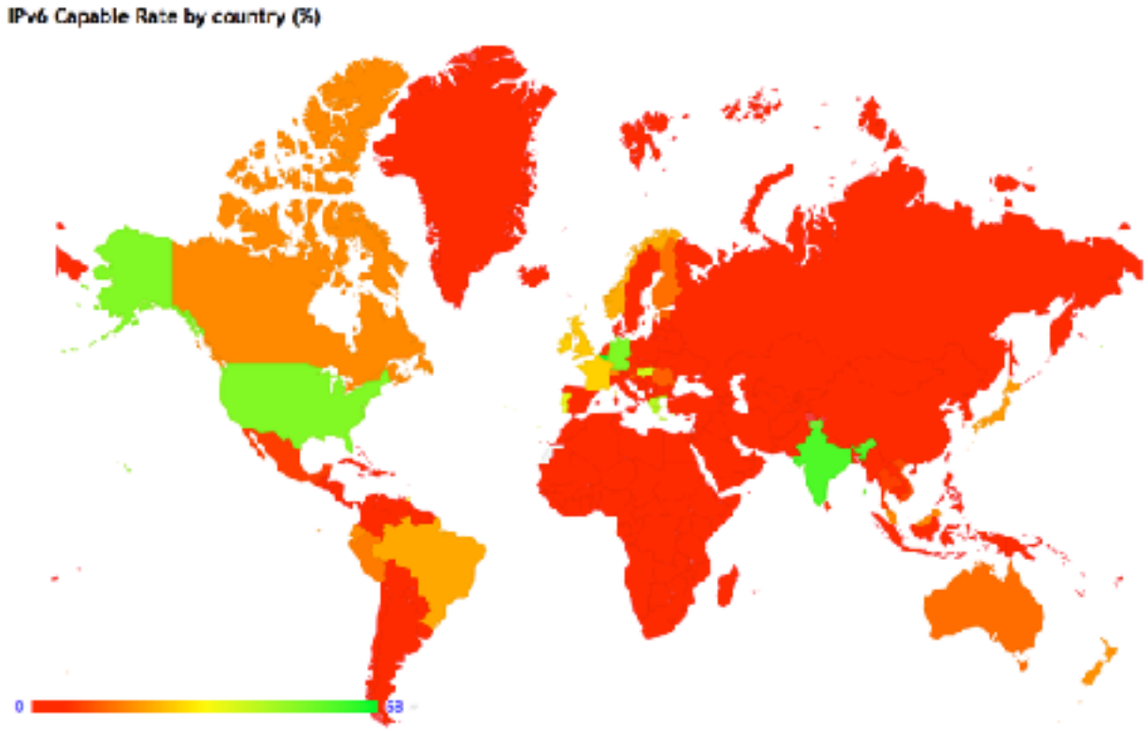

source: labs.apnic.net

Fast Forward to 2017

2017

4 out of 5 RIRs exhausted.

About ~1% of the IPv4 space left unallocated.

Transition to IPv6
→ plenty of measurements and statistics available

Buy IPv4 on Address Markets
→ transfer statistics available from the RIRs

Listed IPv4 Transfers

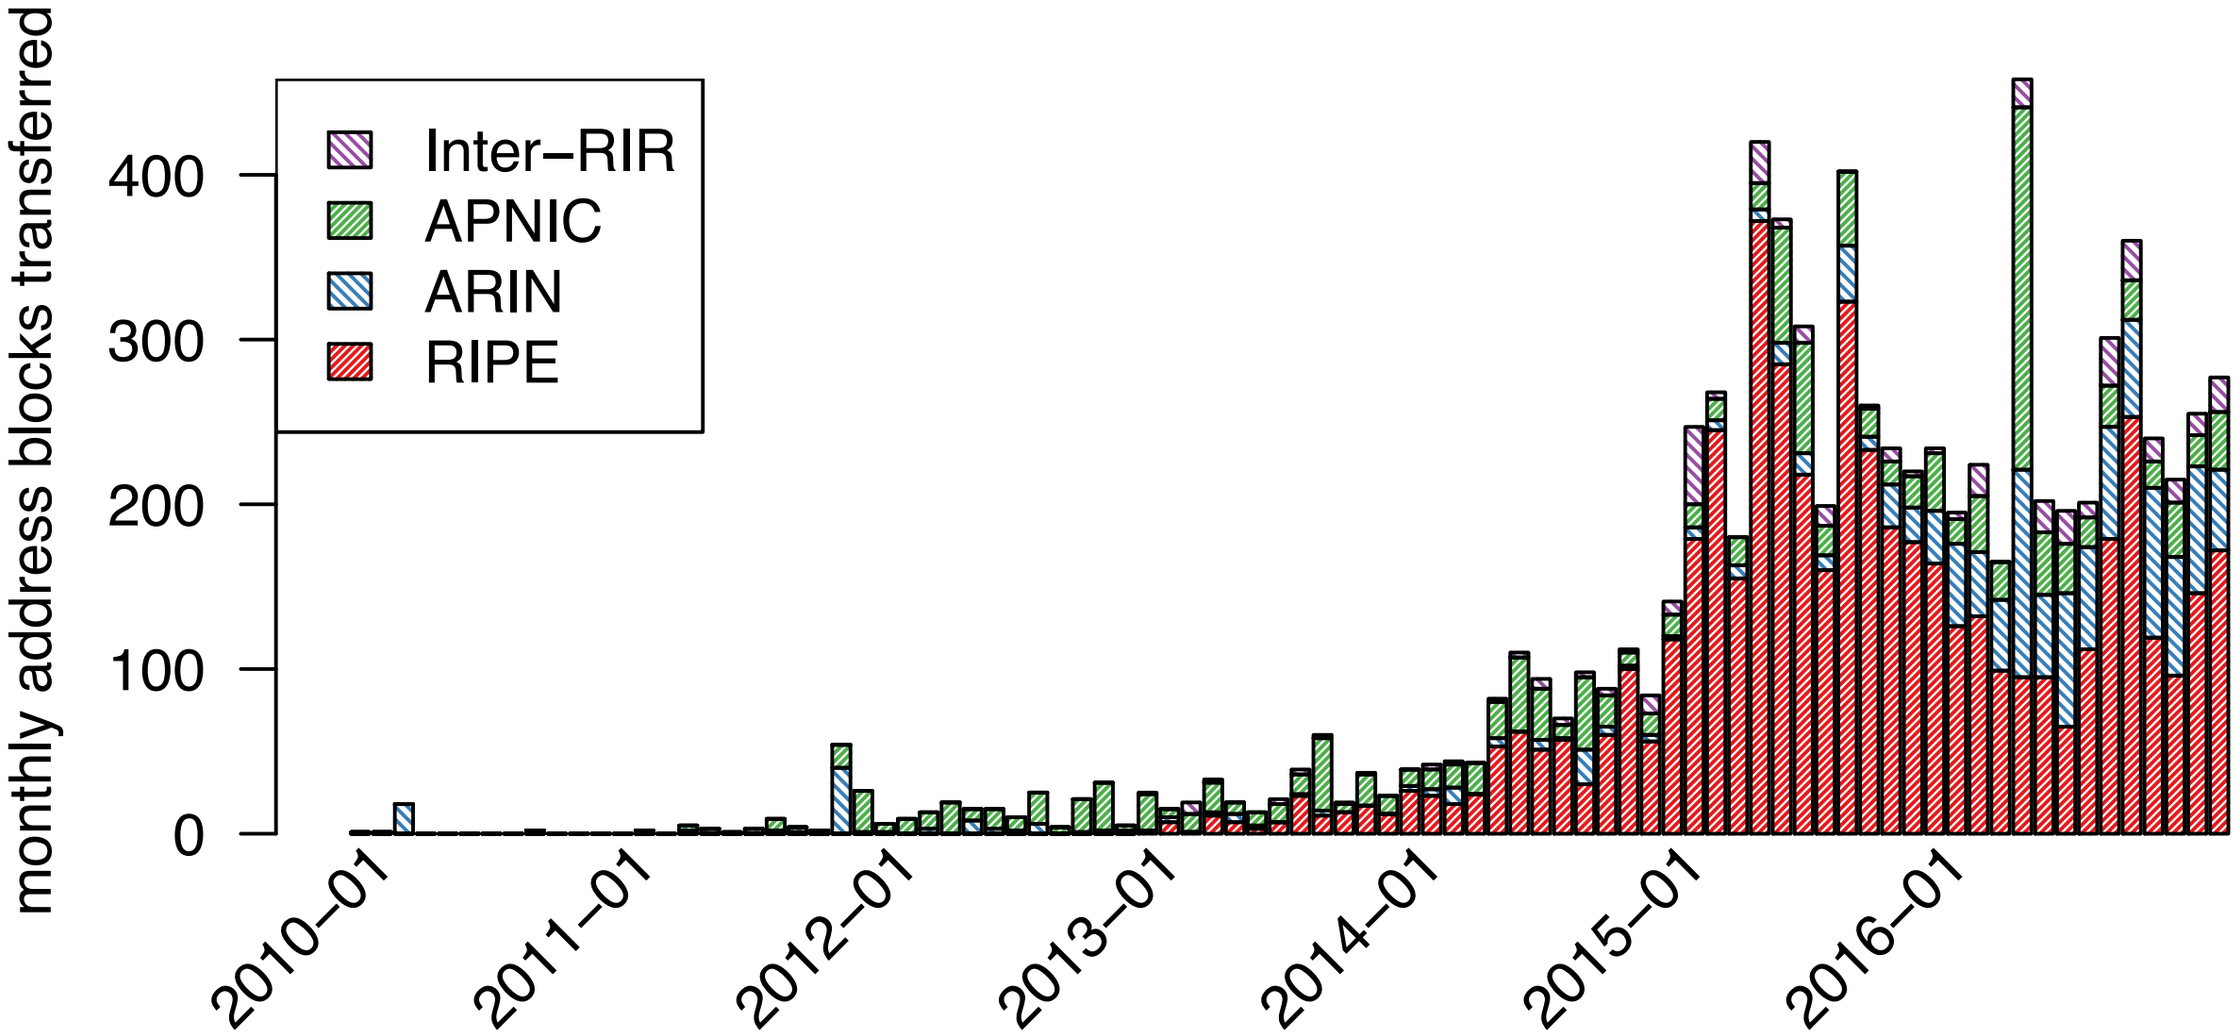

Fast Forward to 2017

2017

4 out of 5 RIRs exhausted.

About ~1% of the IPv4 space left unallocated.

Transition to IPv6

→ plenty of measurements and statistics available

Buy IPv4 on Address Markets

→ transfer statistics available from the RIRs

Use IPv4 Carrier-Grade NAT

→ **no deployment statistics available**

→ **little is known about CGN configurations**

Fast Forward to 2017

2017

4 out of 5 RIRs exhausted.

About ~1% of the IPv4 space left unallocated.

Transition to IPv6
→ plenty of measurements and statistics available

Buy IPv4 on Address Markets
→ transfer statistics available from the RIRs

Use IPv4 Carrier-Grade NAT
→ **no deployment statistics available** → **little is known about CGN configurations**

ISP Survey

We asked ISPs about IPv4 Carrier-Grade NAT
• More than 75 ISPs from all regions of the world replied • Small rural ISPs in Africa up to Fortune 50 companies

ISP Survey

We asked ISPs about IPv4 Carrier-Grade NAT

- More than 75 ISPs from all regions of the world replied
- Small rural ISPs in Africa up to Fortune 50 companies

Did you or do you plan to deploy IPv4 Carrier-Grade NAT?

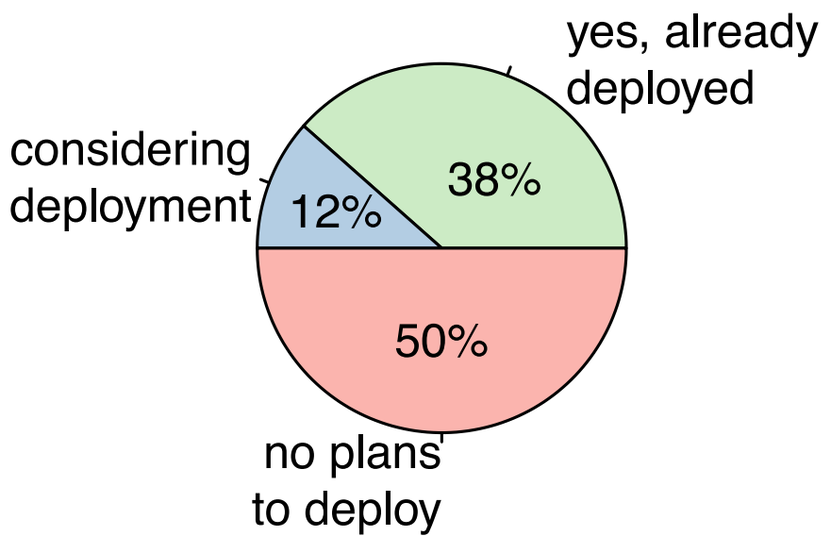

ISP Survey: CGN Specifics

Do you have operational concerns about CGN?
• Subscribers experience problems with application (e.g., gaming) • Traceability of users behind CGN • Issues with CGN IP addresses getting blacklisted

Major challenges/caveats when configuring CGNs?
• Troubleshooting connectivity issues • Resource allocation, quotas and port ranges per subscriber • Internal address space fragmentation/shortage (e.g., RFC1918)

ISP Survey: Comments (Free Text Field)

Do you have operational concerns about CGNs?

- Subscribers experience problems with application (e.g., gaming)
- Traceability of users behind CGN
- Issues with CGN IP addresses getting blacklisting

What challenges/caveats when configuring CGNs?

- Dimensioning CGNs:
  - Allocating IP addresses/ports to subscribers, quotas per subscriber
  - Distributed vs. Centralized CGN Infrastructure
- Troubleshooting connectivity issues
- Hardware limitations (memory/CPU)

“well, NAT s*cks, but there's not much of an alternative”

“CGN is bad enough, but IPv6 is still an afterthought for most and usually quite problematic so it's not worth it yet”

“In Russia, ISPs prefer to just add CGNs when they run out of space and charge a small subset of customers for a public IP address”

Motivation and Objectives

Motivation
• CGNs seems to be widely deployed • ISPs voiced concerns about CGN configuration/operation • No broad and systematic studies available

Objectives
• Develop methods to detect CGN presence “in the wild” • Develop methods to extract properties from detected CGNs • Illuminate the current status of CGN deployment in the Internet

NATs between Subscribers and the Internet

NAT44
(subscriber-side)

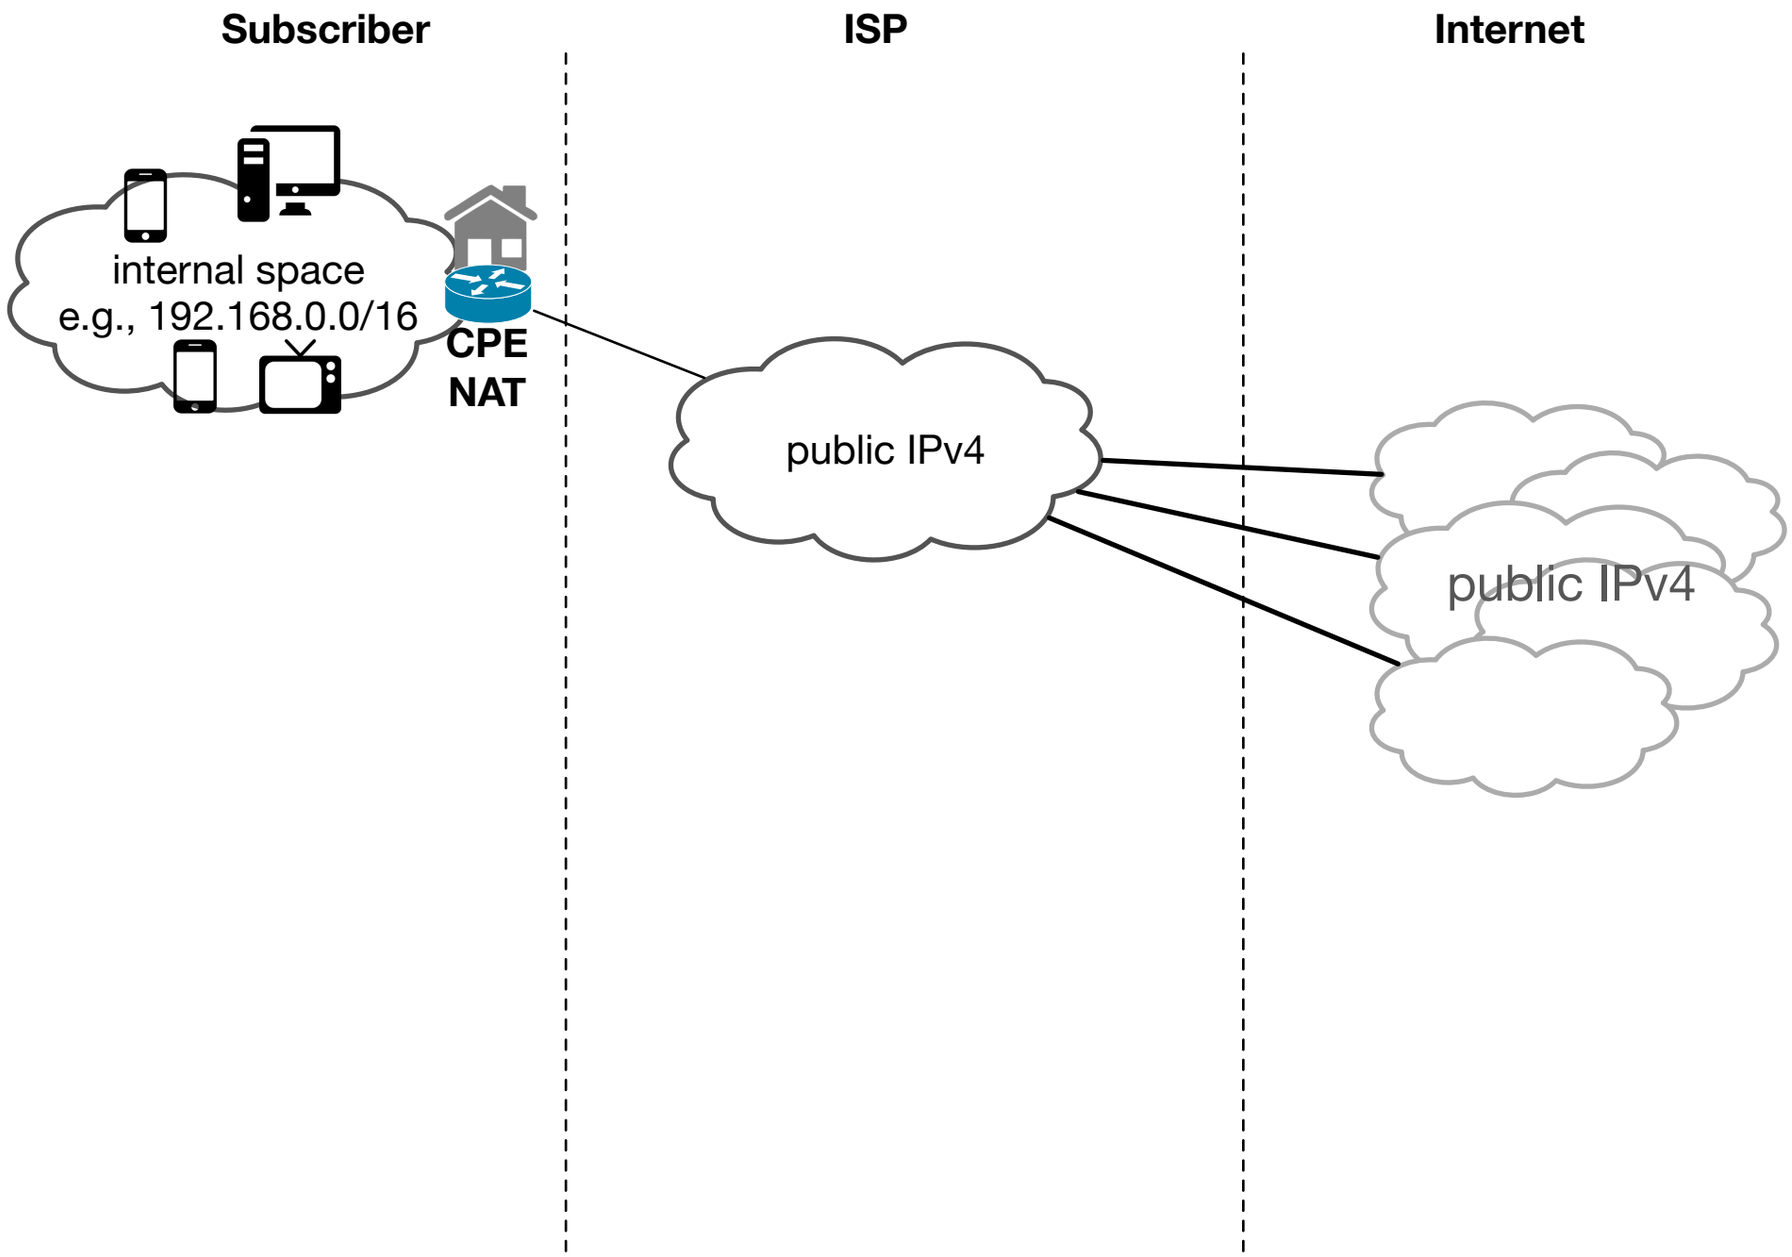

NATs between Subscribers and the Internet

NAT44
(subscriber-side)

NAT44
(carrier-side)

NAT444
(subscriber-side
and
carrier-side)

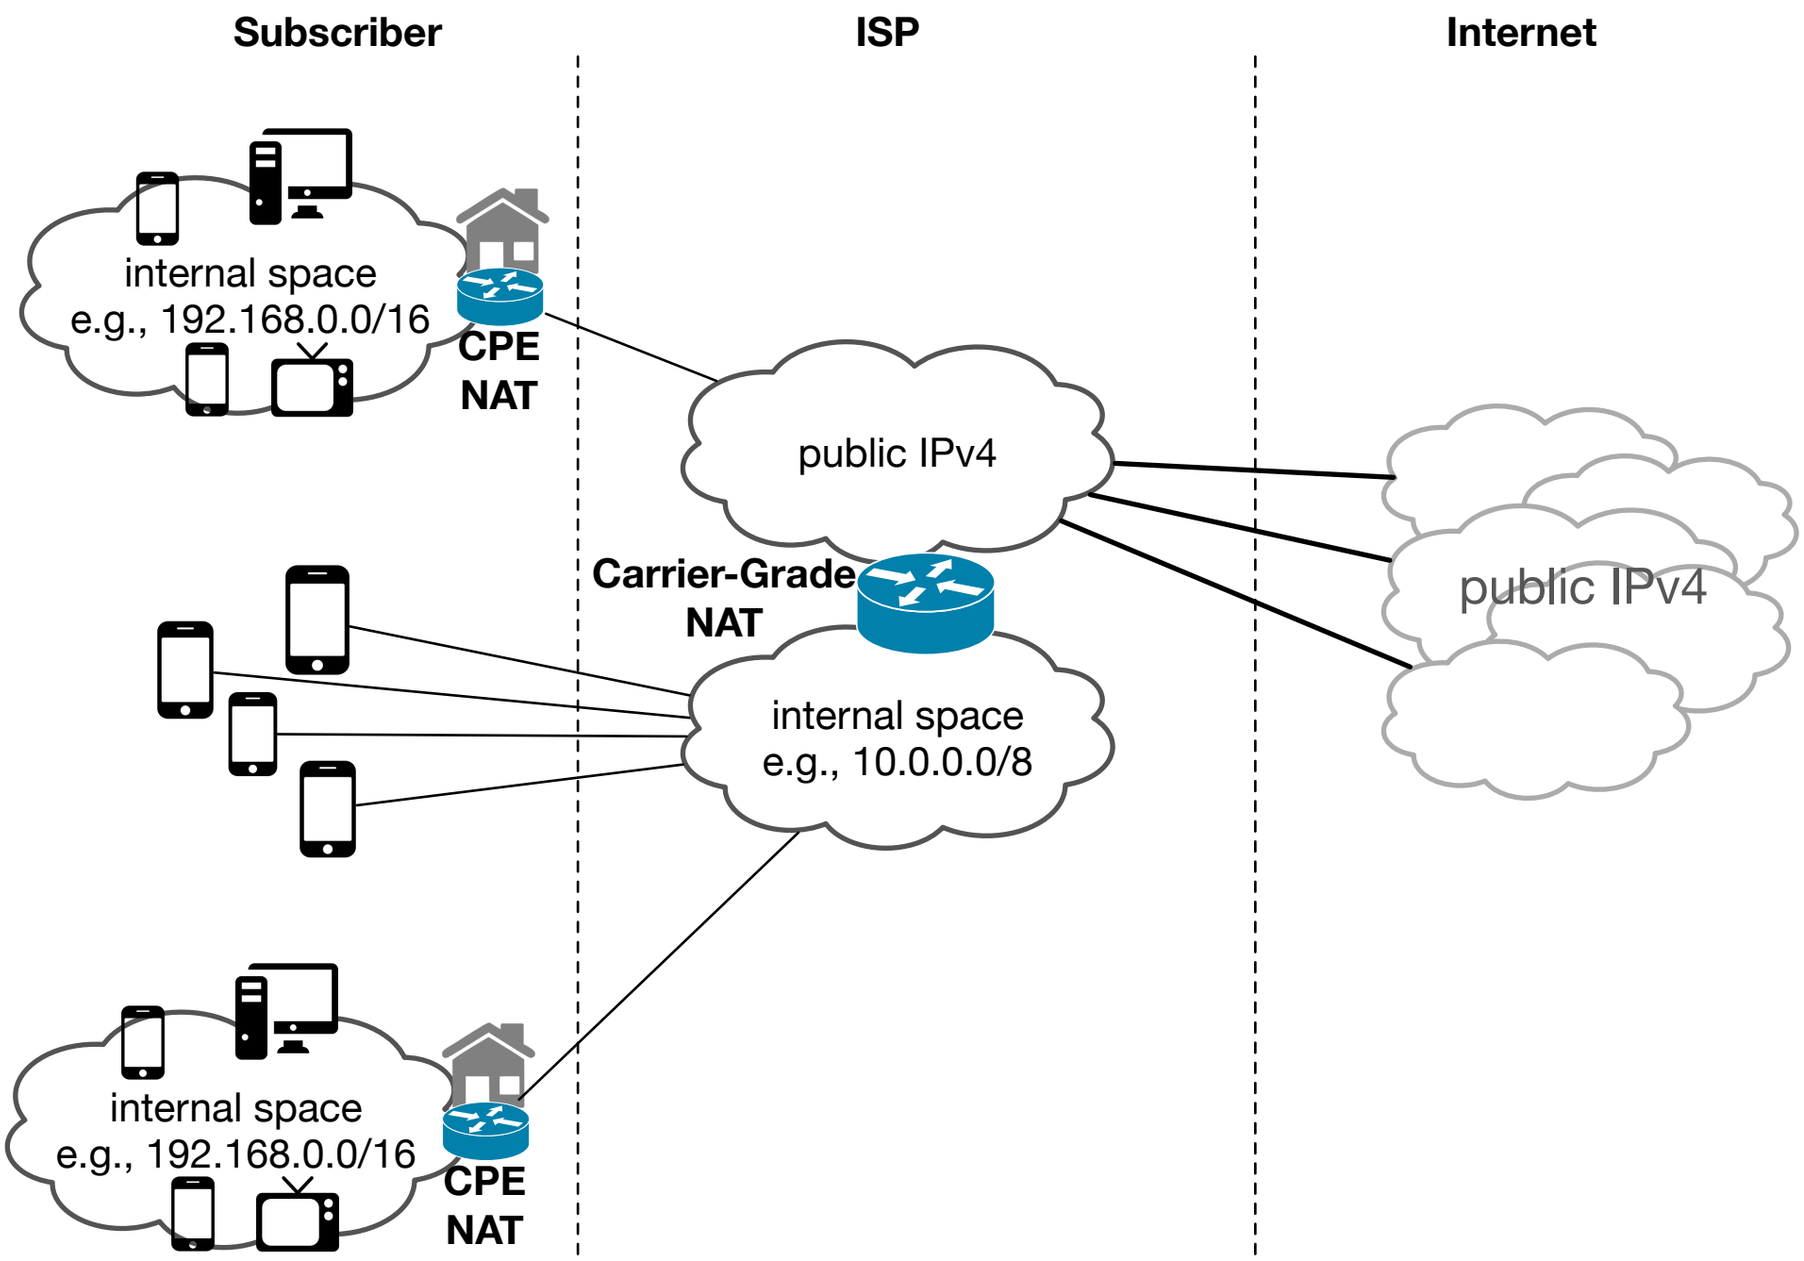

Agenda

- ISP Survey
- Detecting CGN Presence
  - **From the Outside via BitTorrent**
  - From the Inside via Netalyzr
- CGN Deployment Statistics
- CGN Properties
- Conclusion

The BitTorrent DHT

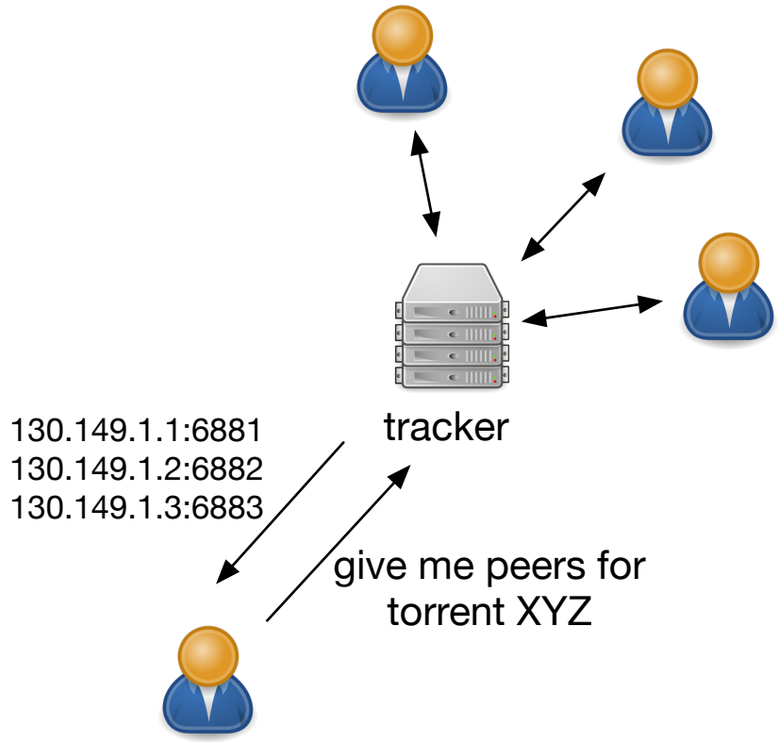

classic BitTorrent

Tracker stores peer contact information (IP:port)

The BitTorrent DHT

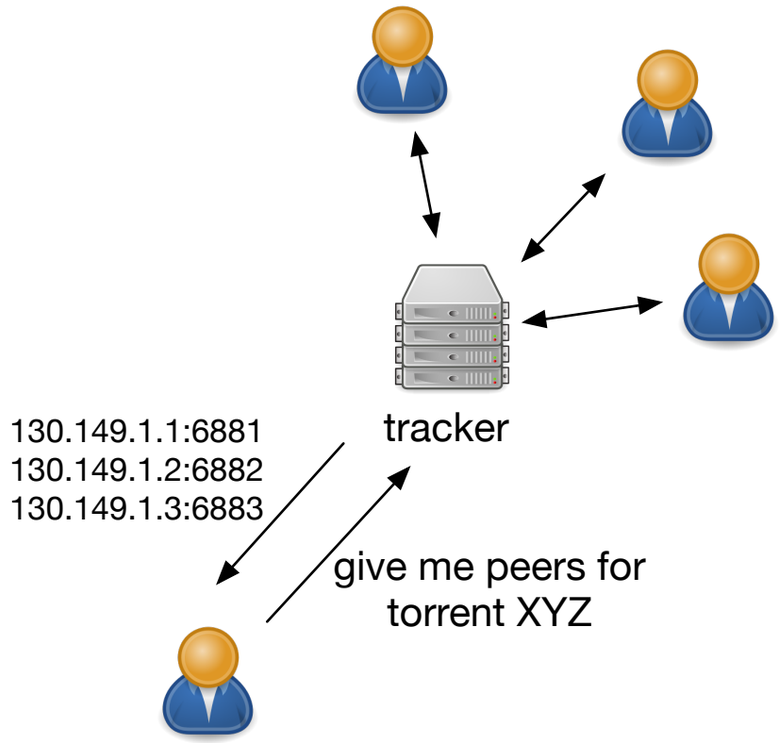

classic BitTorrent

Tracker stores peer contact information (IP:port)

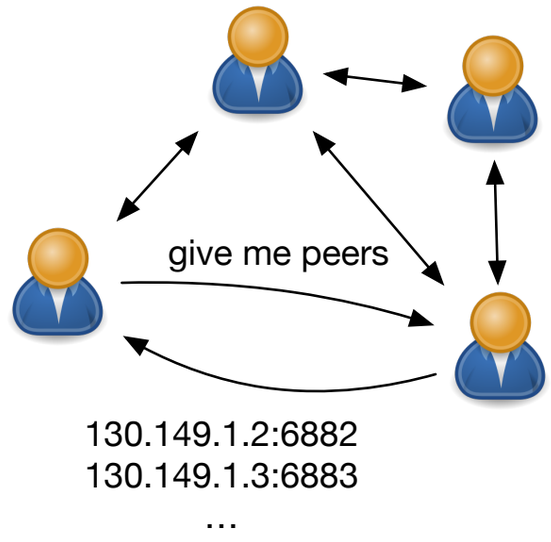

BitTorrent DHT:

Peers store each others' contact information (IP:port, nodeid)

The BitTorrent DHT

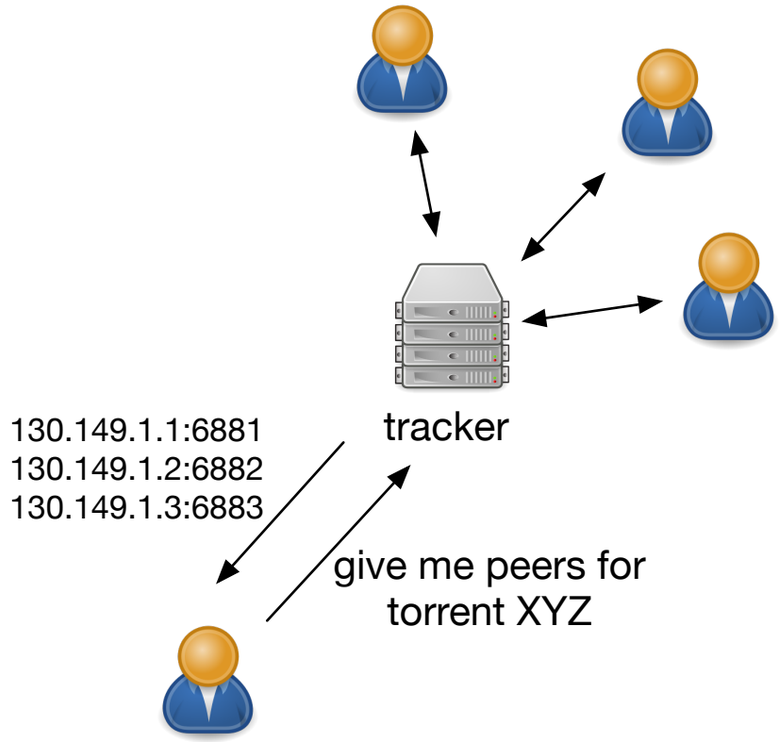

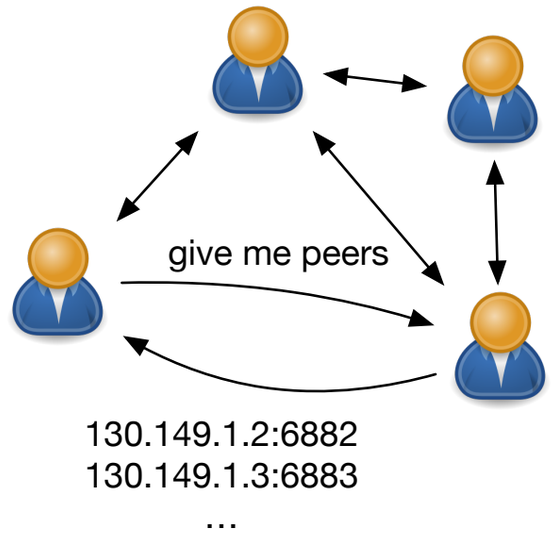

classic BitTorrent
Tracker stores peer contact information (IP:port)

BitTorrent DHT:
Peers store each others' contact information (IP:port, nodeid)

We can use DHT peers as vantage points

Crawling the BitTorrent DHT

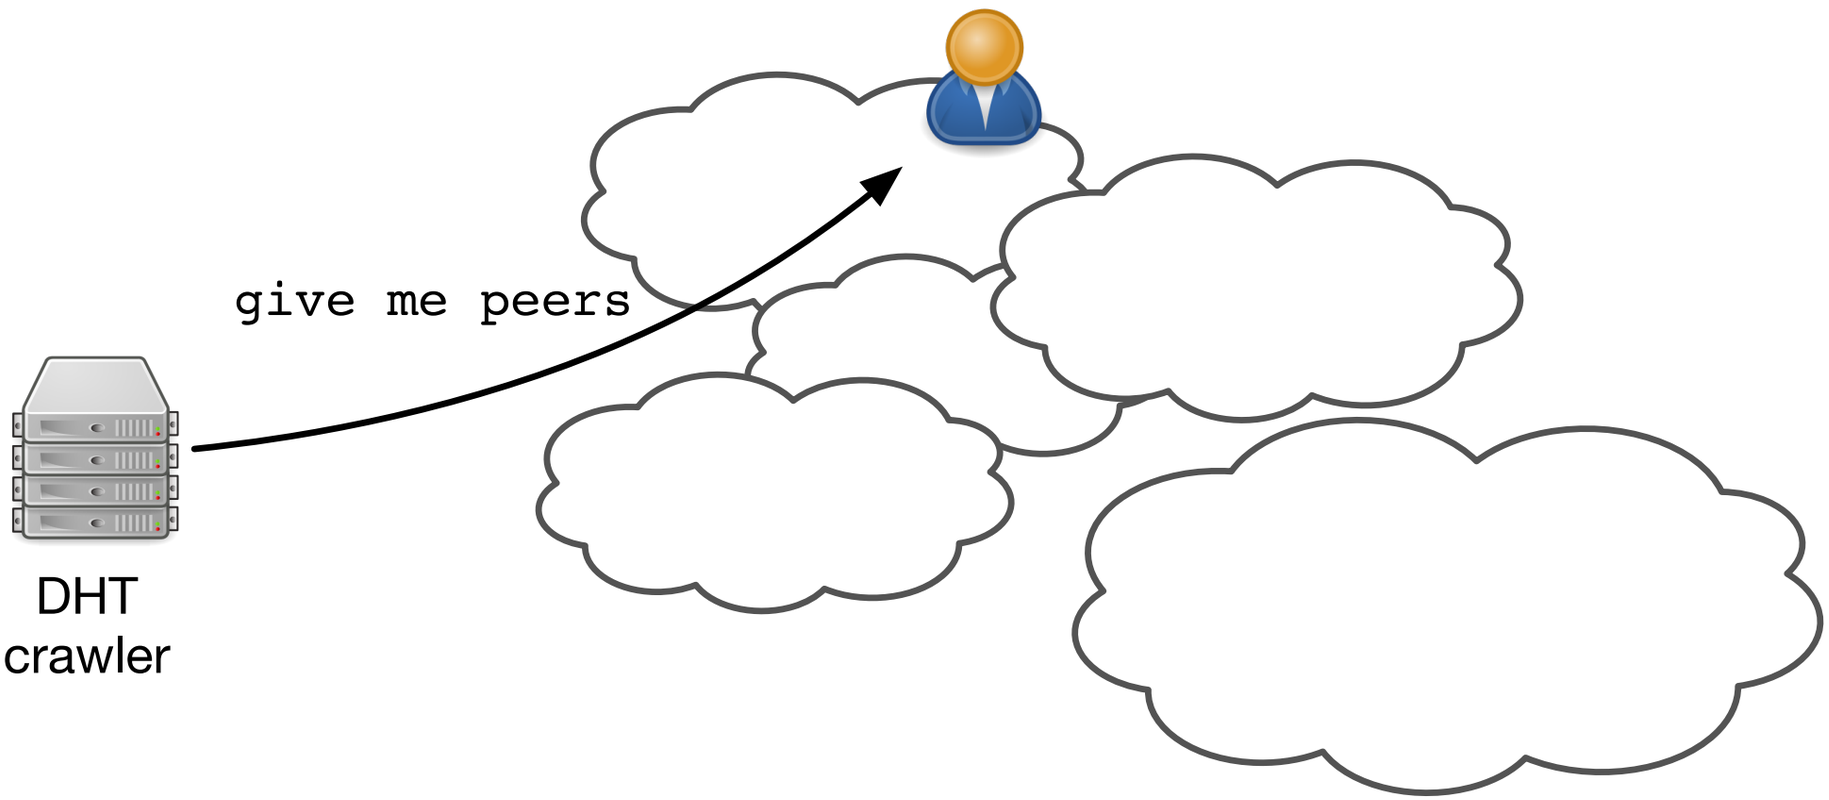

Crawling the BitTorrent DHT

```
i can reach
peer 25fc at 130.149.1.2:6881
peer 492c at 190.2.0.1:6881
…
```

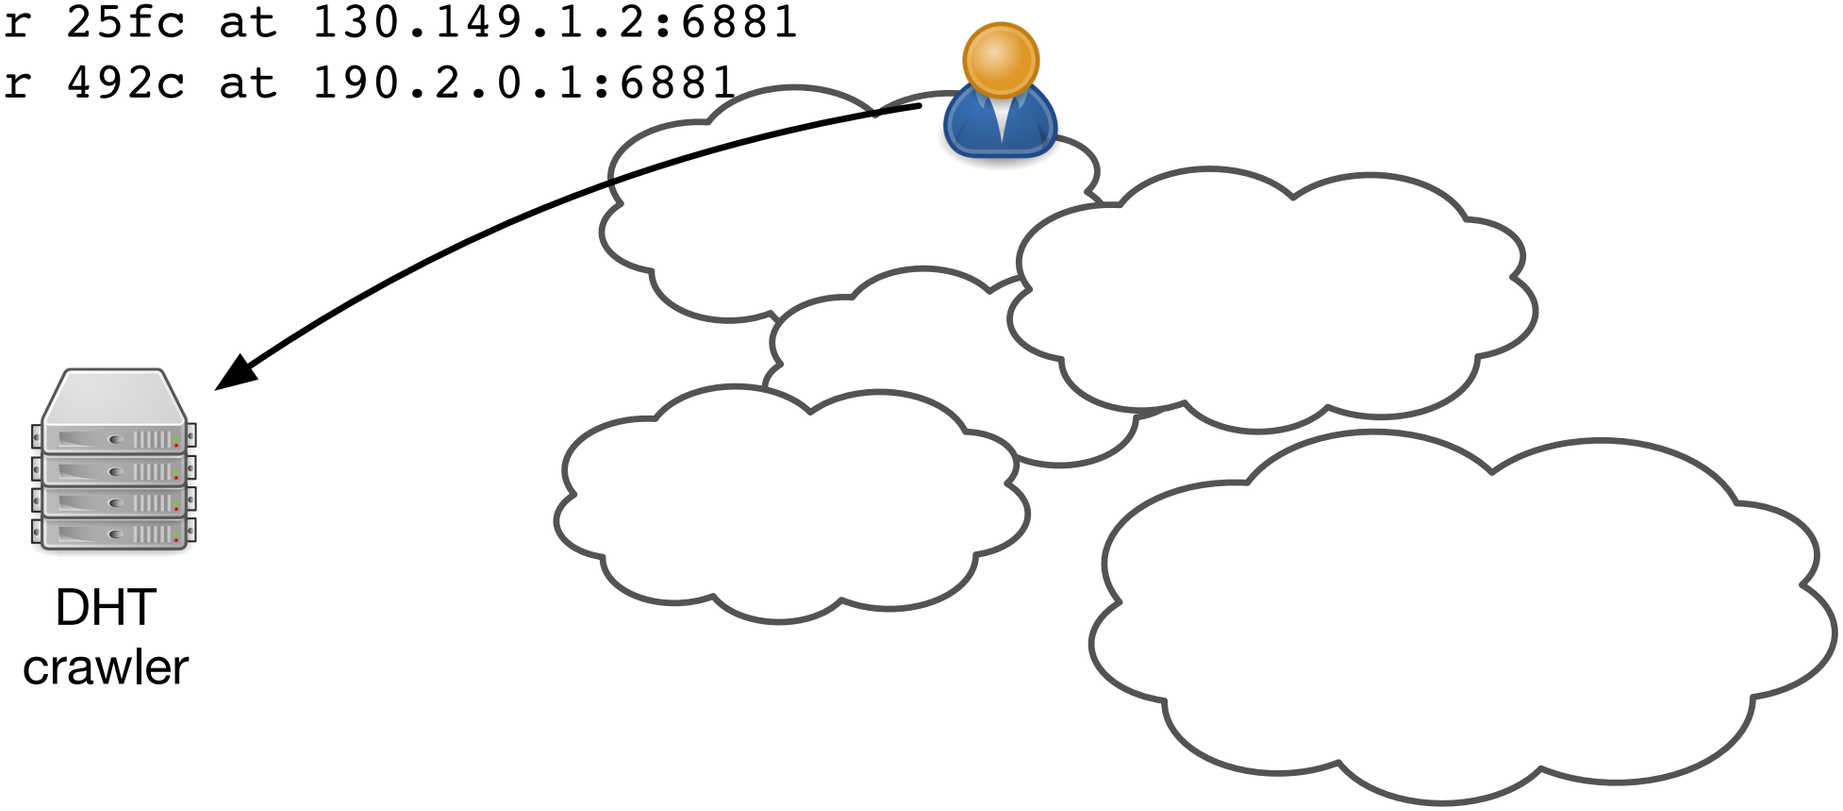

Crawling the BitTorrent DHT

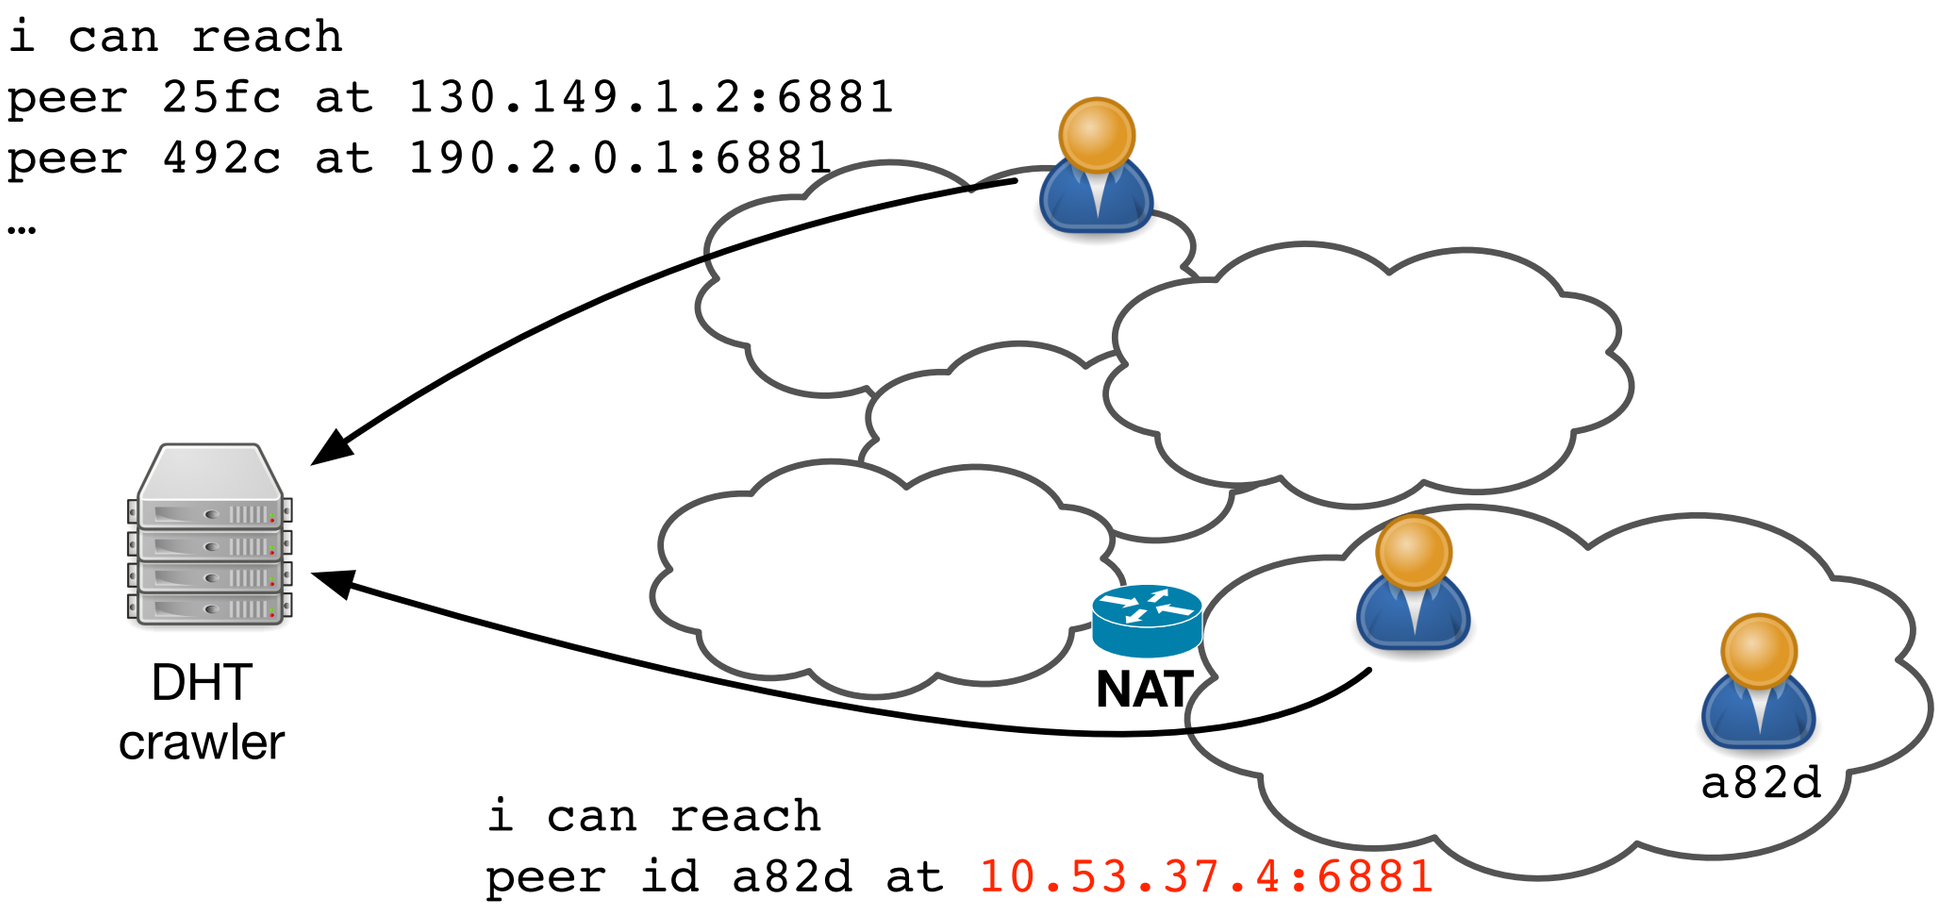

Some peers leak us internal IP addresses of other peers

Crawling the BitTorrent DHT

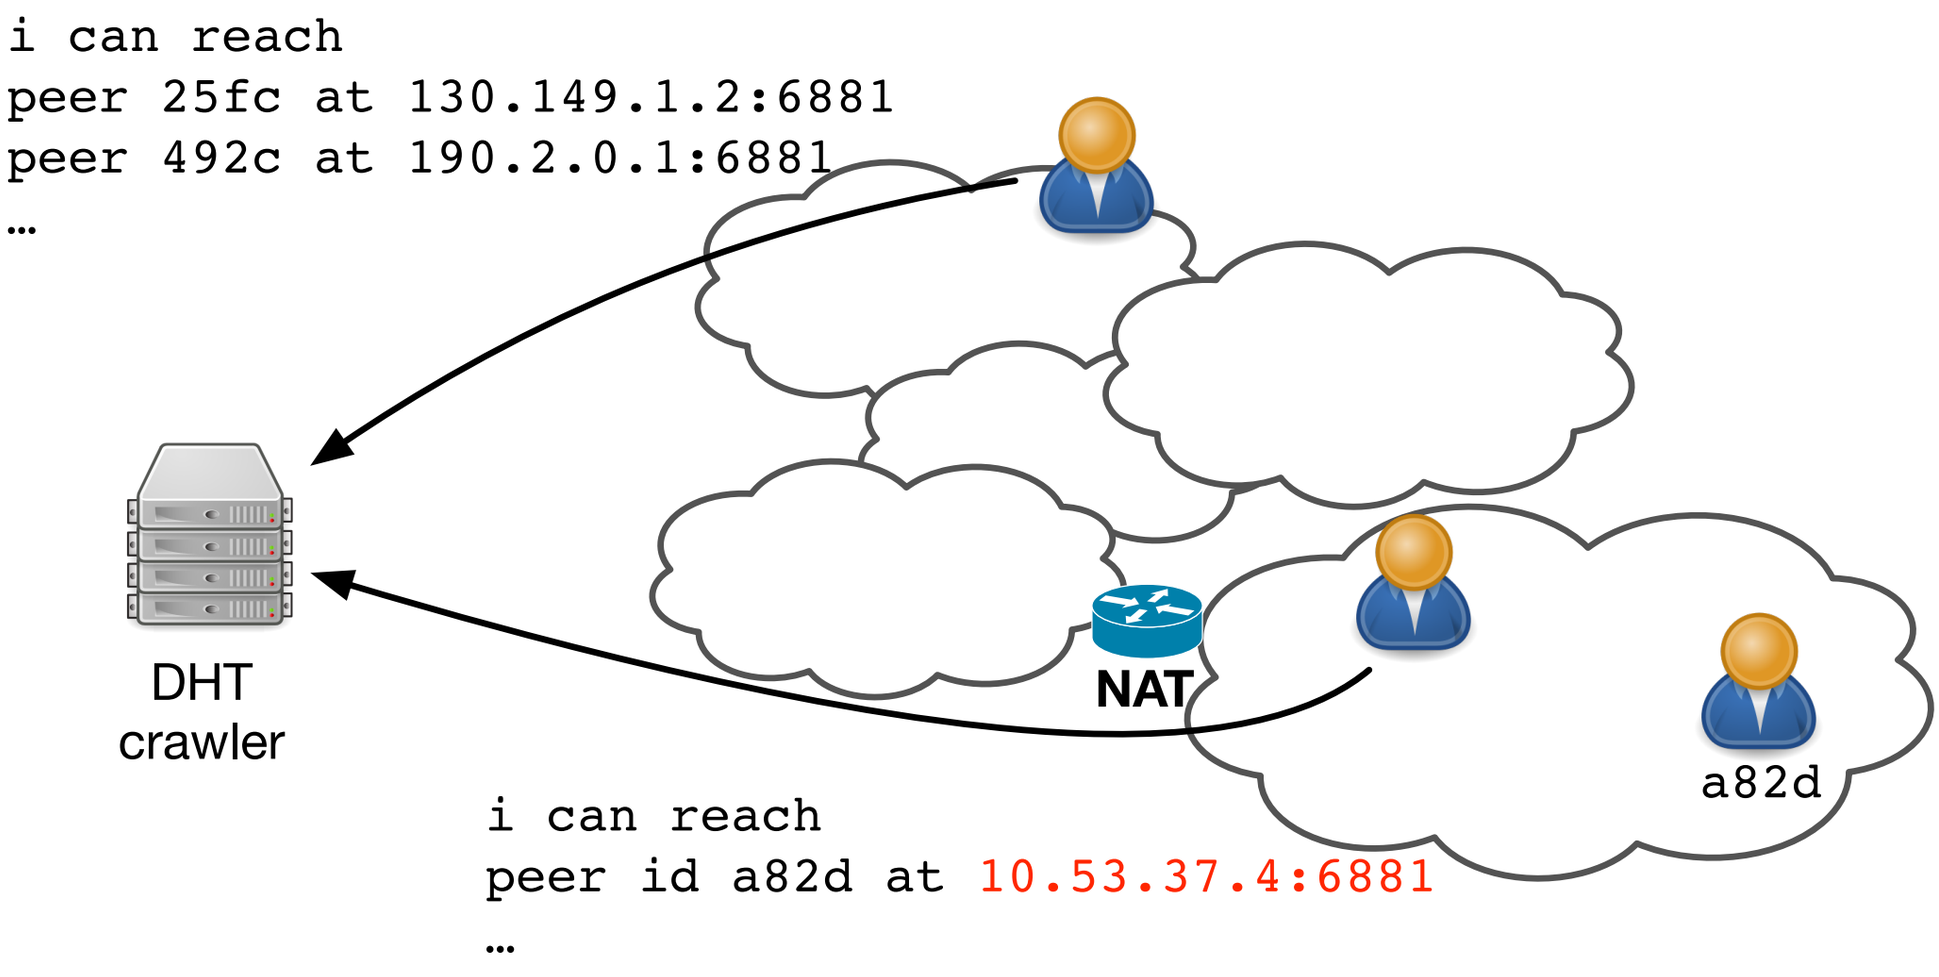

Some peers leak us internal IP addresses of other peers within 1 week: more than 700.000 peers in 5.000 ASes!

Understanding Leakage Relationships

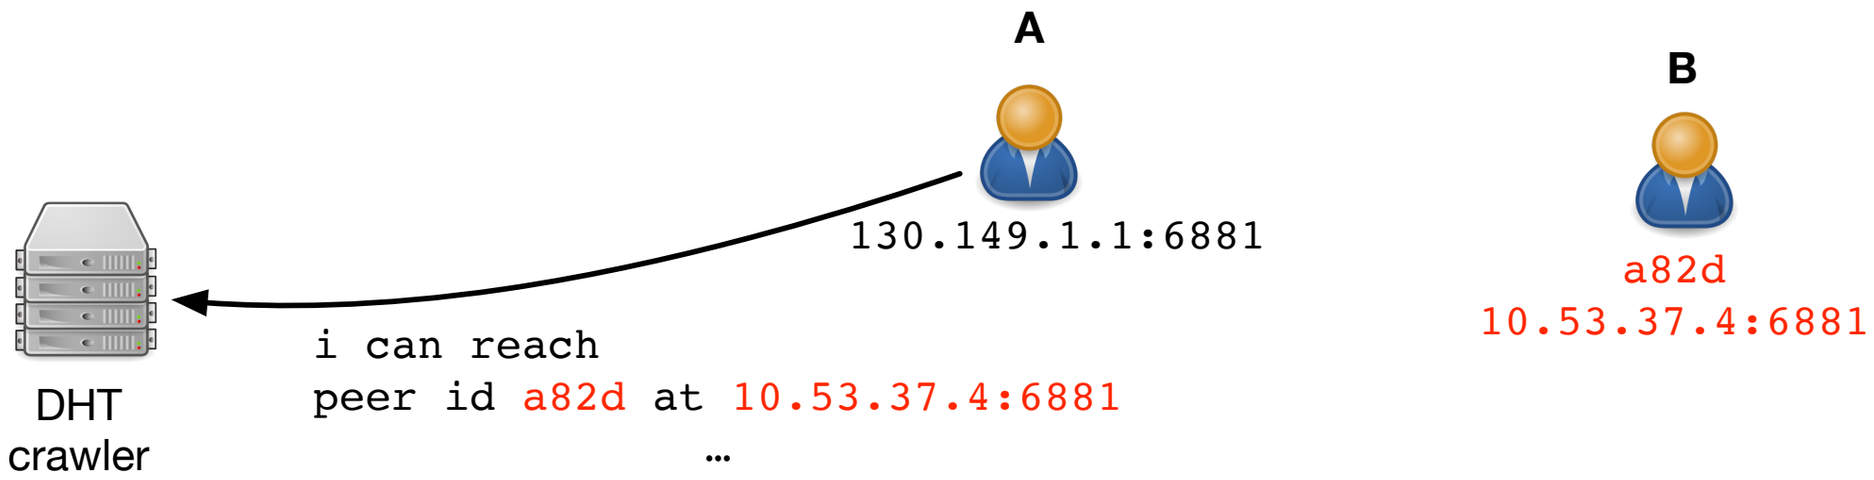

we construct a graph of leaking relationships

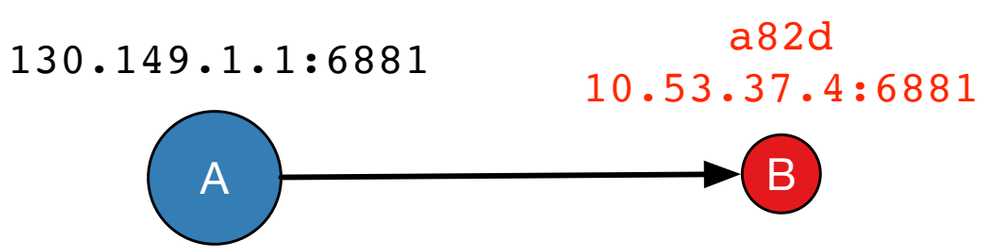

…now we look these graphs on a per-AS basis

BitTorrent Peer Leakage Graph

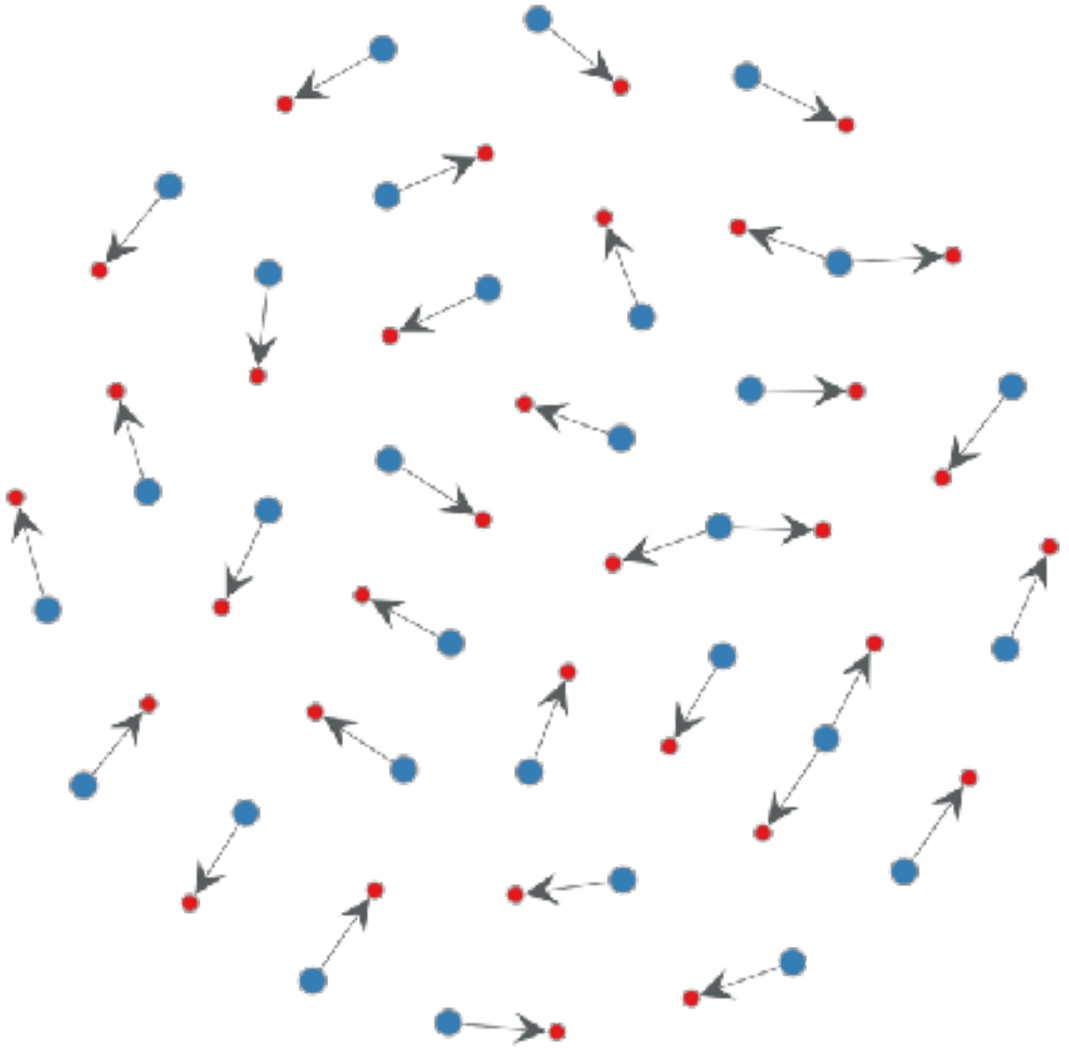

In this AS:
no CGN detected

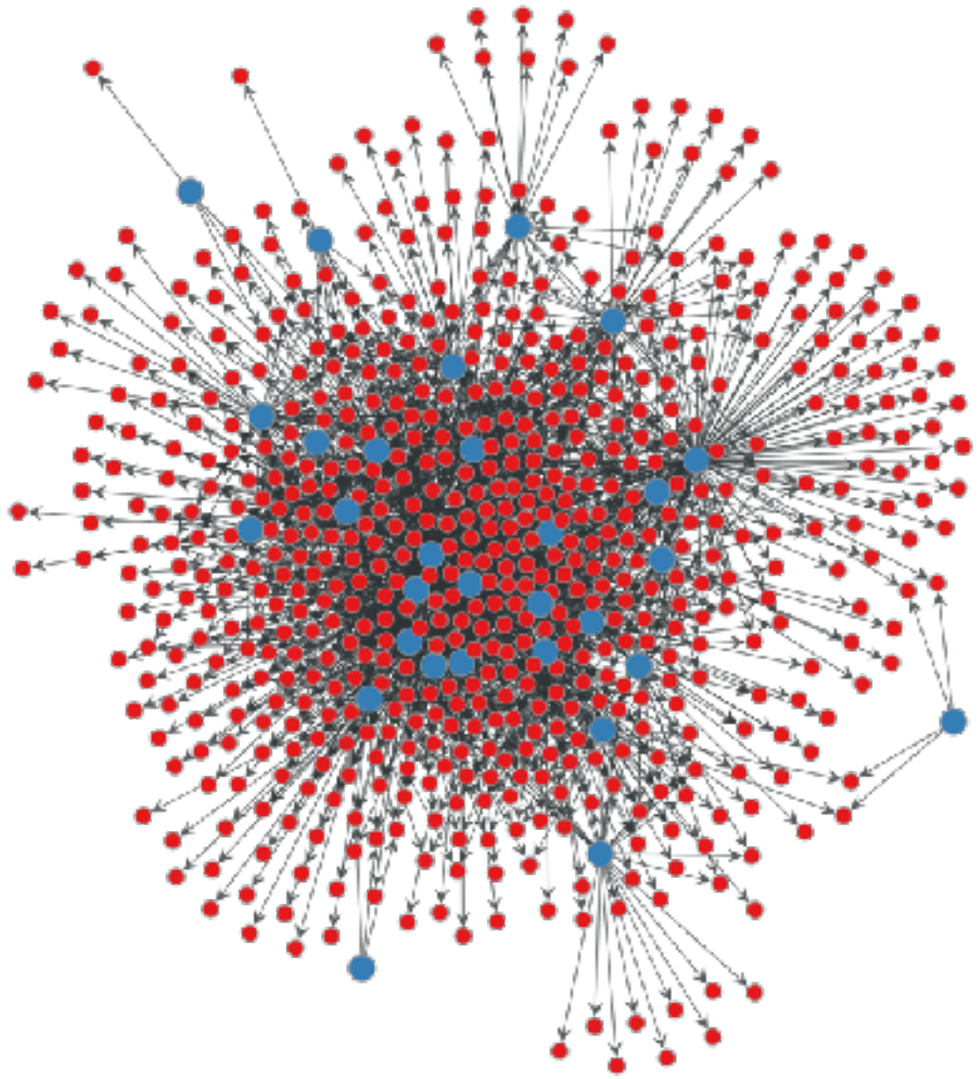

In this AS:
CGN detected

Detecting CGNs with BitTorrent

- We test more than 2700 ASes with this methodology
- We detect CGN (clusters) in 250+ ASes

Benefits	Caveats
• broad coverage	• need BitTorrent activity
• no probing devices needed	• not all CGNs show up
	• cellular networks?

Agenda

- ISP Survey
- Detecting CGN Presence
  - From the Outside via BitTorrent
  - **From the Inside via Netalyzr**
- CGN Deployment Statistics
- Dominant Characteristics of deployed CGNs
- Conclusion

Netalyzr

What is Netalyzr?

- Network Troubleshooting Suite developed by ICSI Berkeley
- Available as Android App, Java Applet, CL tool

Netalyzr in this Study

- More than 550K sessions in 1500+ ASes
- Access to device/router/public IP address
- Runs in cellular and non-cellular networks
- Customized tests

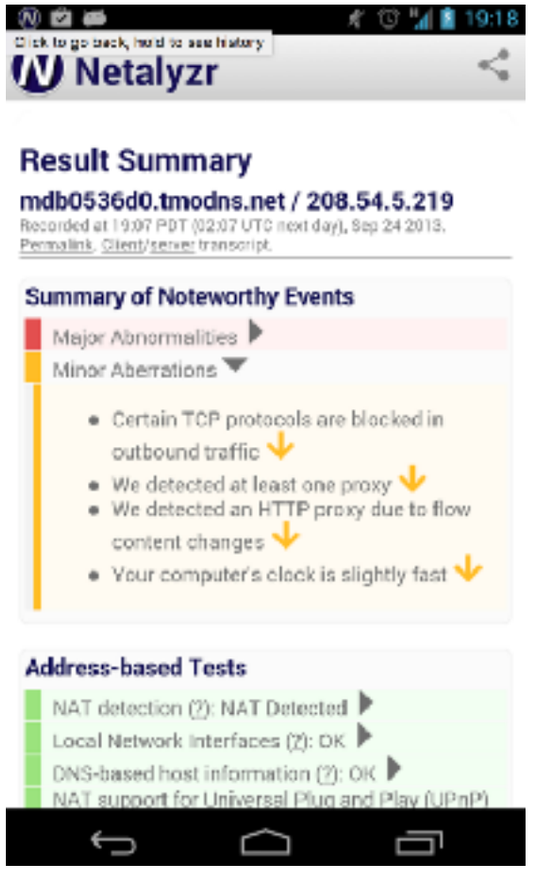

Detecting CGN in Cellular Networks

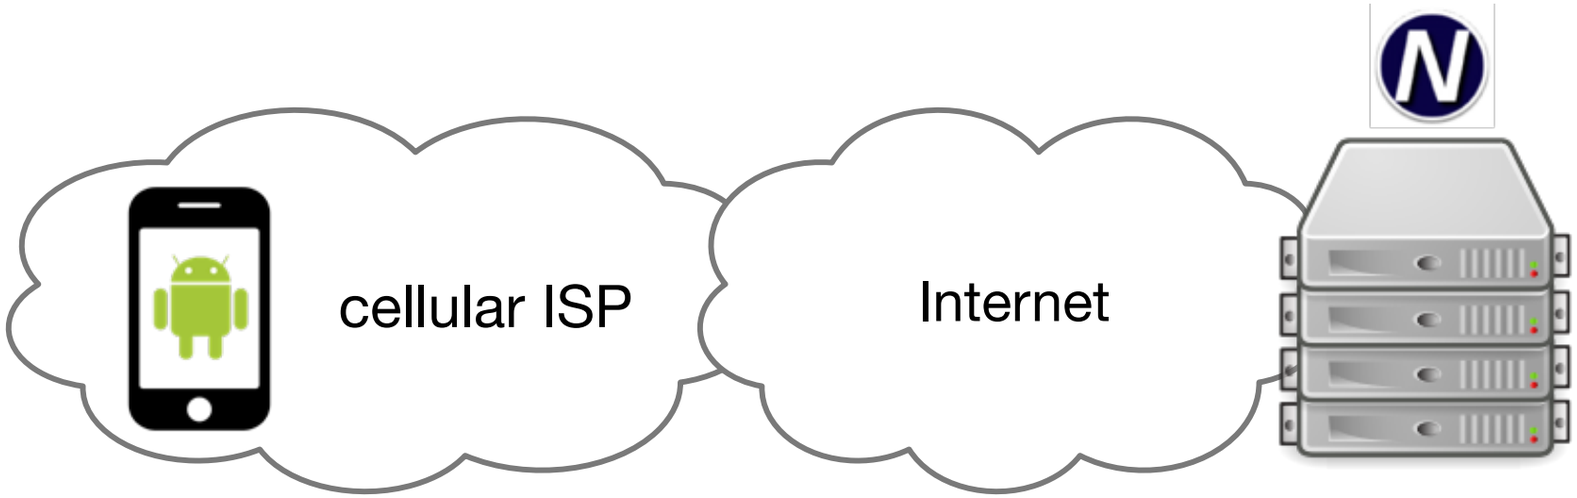

Device IP address assigned directly by the ISP
Device IP ≠ server-side IP → Carrier-Grade NAT

Detecting CGN in Residential Networks

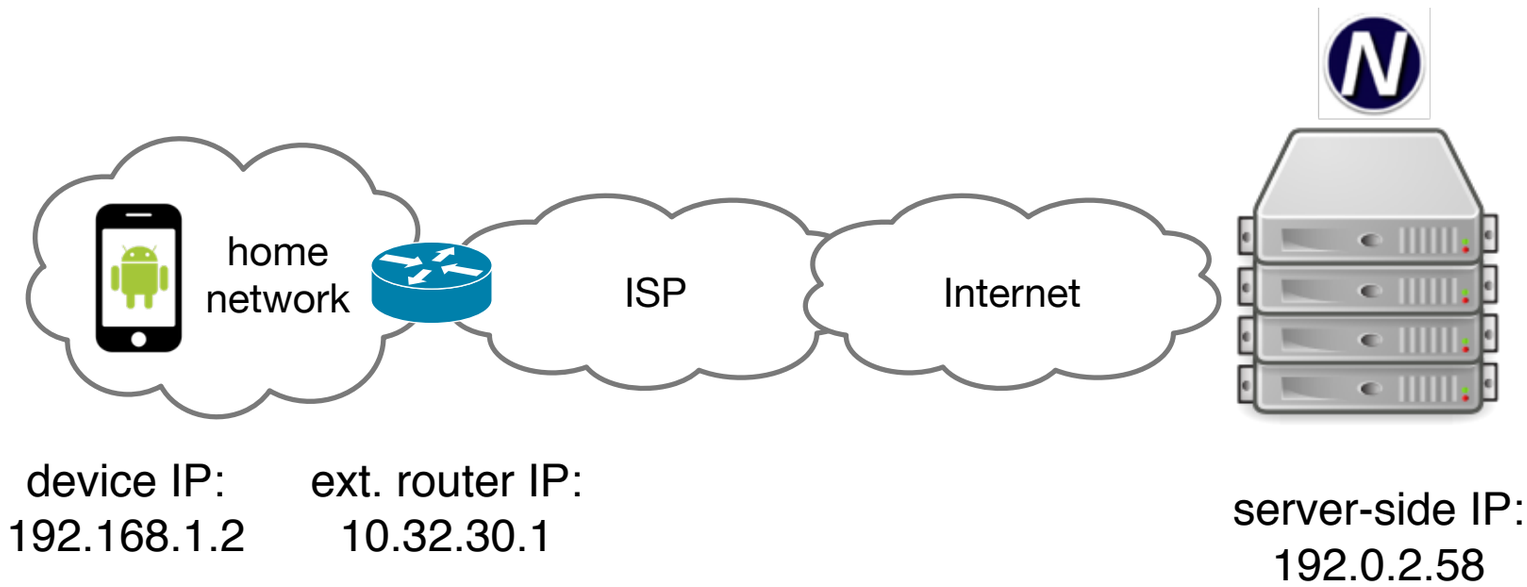

ext. router IP ≠ server-side IP → Carrier-Grade NAT?

Detecting CGN in Residential Networks (2)

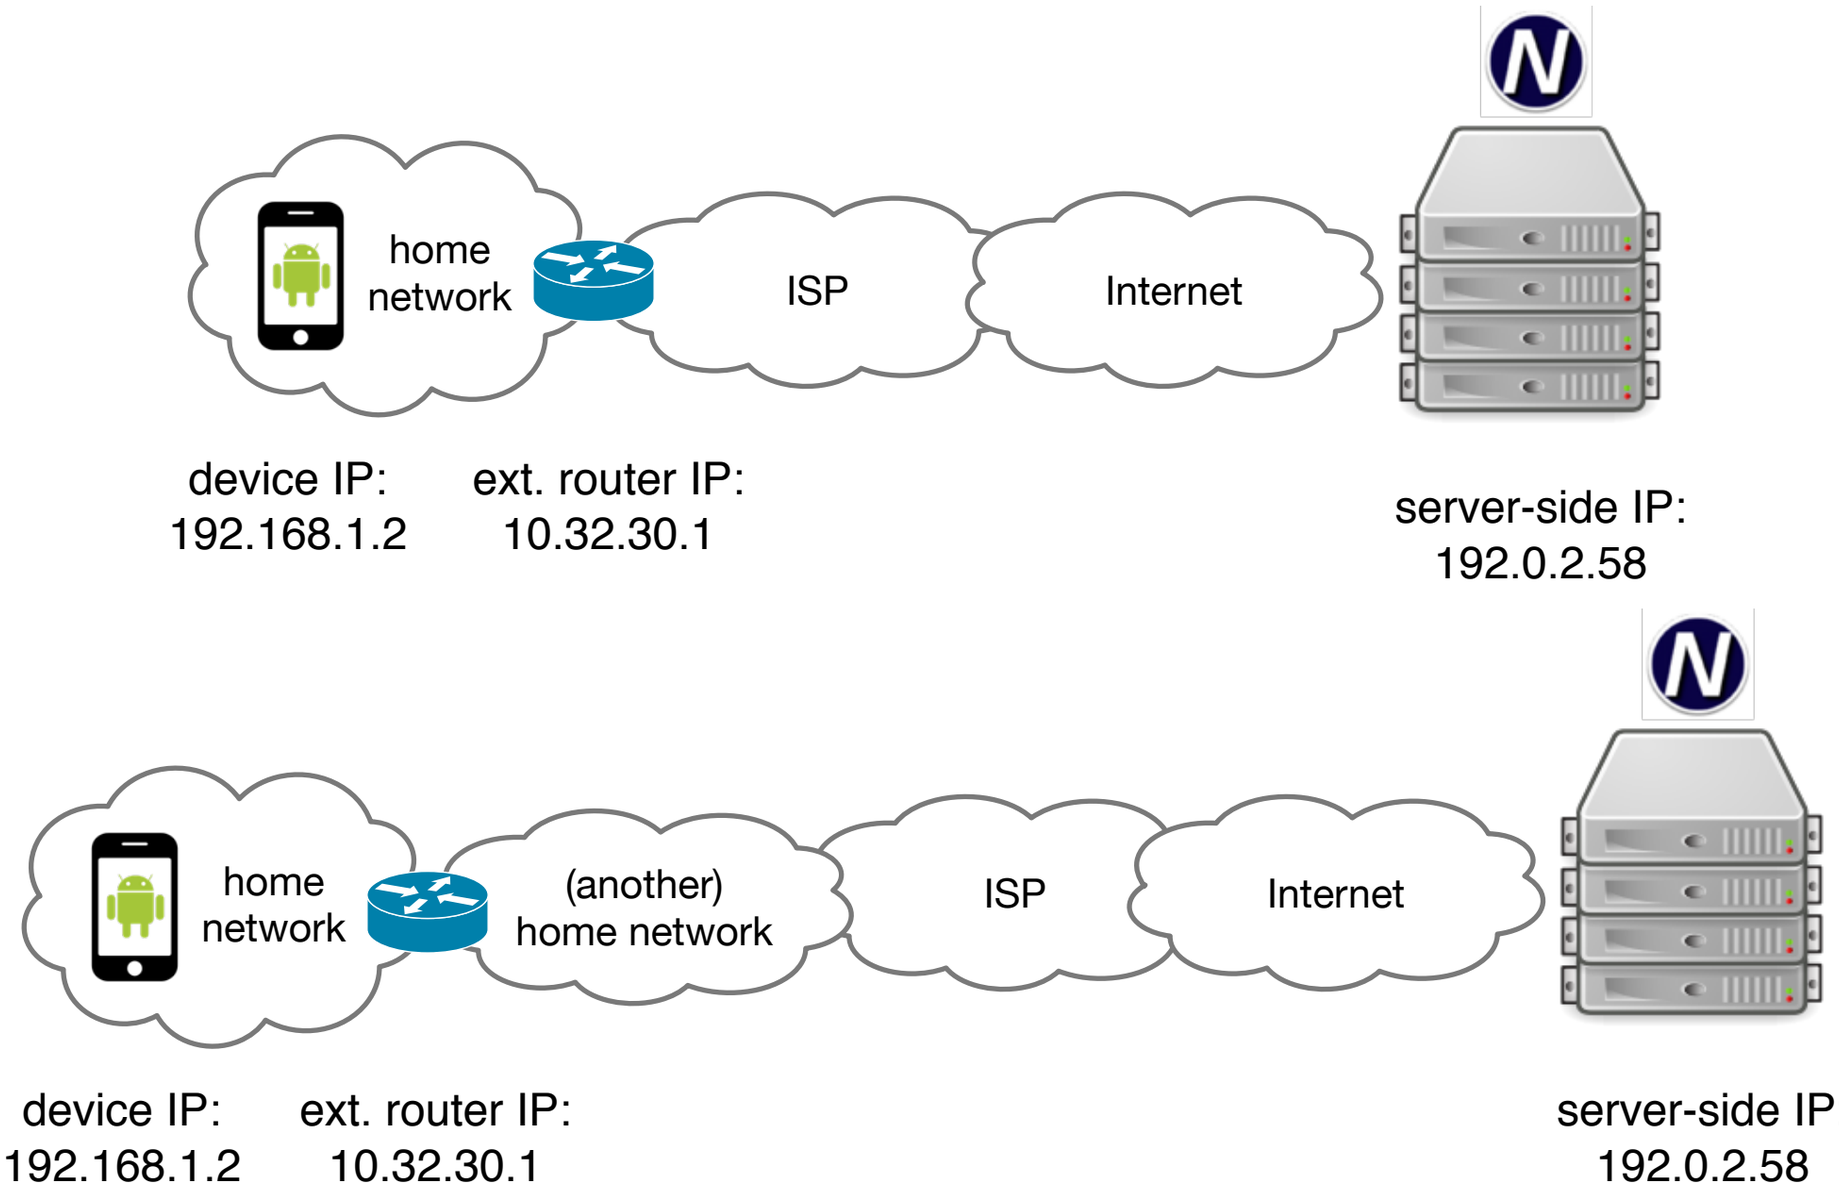

Up to 7% of sessions with chained home NATs

Detecting CGNs with Netalyzr

- We test 1500+ ASes
- We detect CGN in 194 non-cellular and 205 cellular ASes

Benefits
direct IP addressing data
cellular and non-cellular
more customized tests

Caveats
partial visibility, crowdsourced (need users to run Netalyzr)

Agenda

- ISP Survey
- Detecting CGN Presence
  - From the Outside via BitTorrent
  - From the Inside via Netalyzr
- **CGN Deployment Statistics**
- CGN Properties
- Conclusion

How many Networks do we cover?

Eyeball Networks (Non-Cellular)

- Identify Eyeball ASes: Spamhaus PBL / APNIC “aspop”
- Eyeball AS population: 3K ASes
- Tested with BitTorrent/Netalyzr: 1,791 **(62%)**
- No strong geographic bias

Cellular Networks

- Identify Cellular Networks directly via Netalyzr
- Tested: 218 ASes

How many Networks deploy CGN?

Eyeball Networks (Non-Cellular)

- CGN-positive: **17.1%**
  - ➡ particularly in the European and Asia-Pacific Region

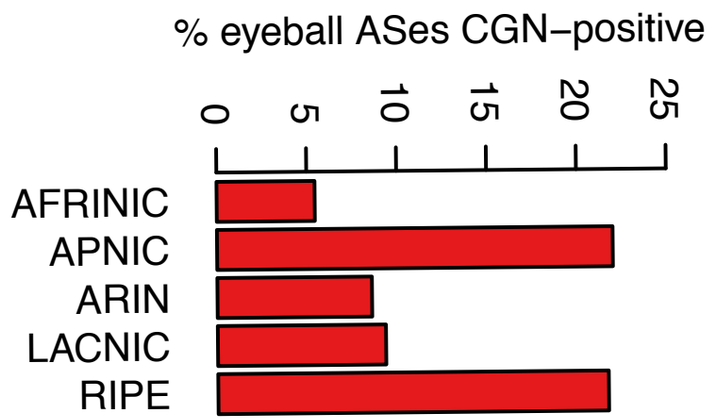

Cellular Networks

- CGN-positive: **94%**
  - ➡ CGN is the norm for cellular

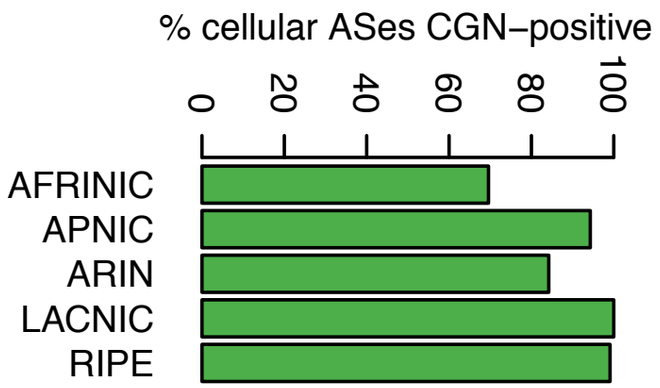

Agenda

- ISP Survey
- Detecting CGN Presence
  - From the Outside via BitTorrent
  - From the Inside via Netalyzr
- CGN Deployment Statistics
- **CGN Properties**
- Conclusion

Per AS: Internal CGN Address Space

Per AS: Internal CGN Address Space

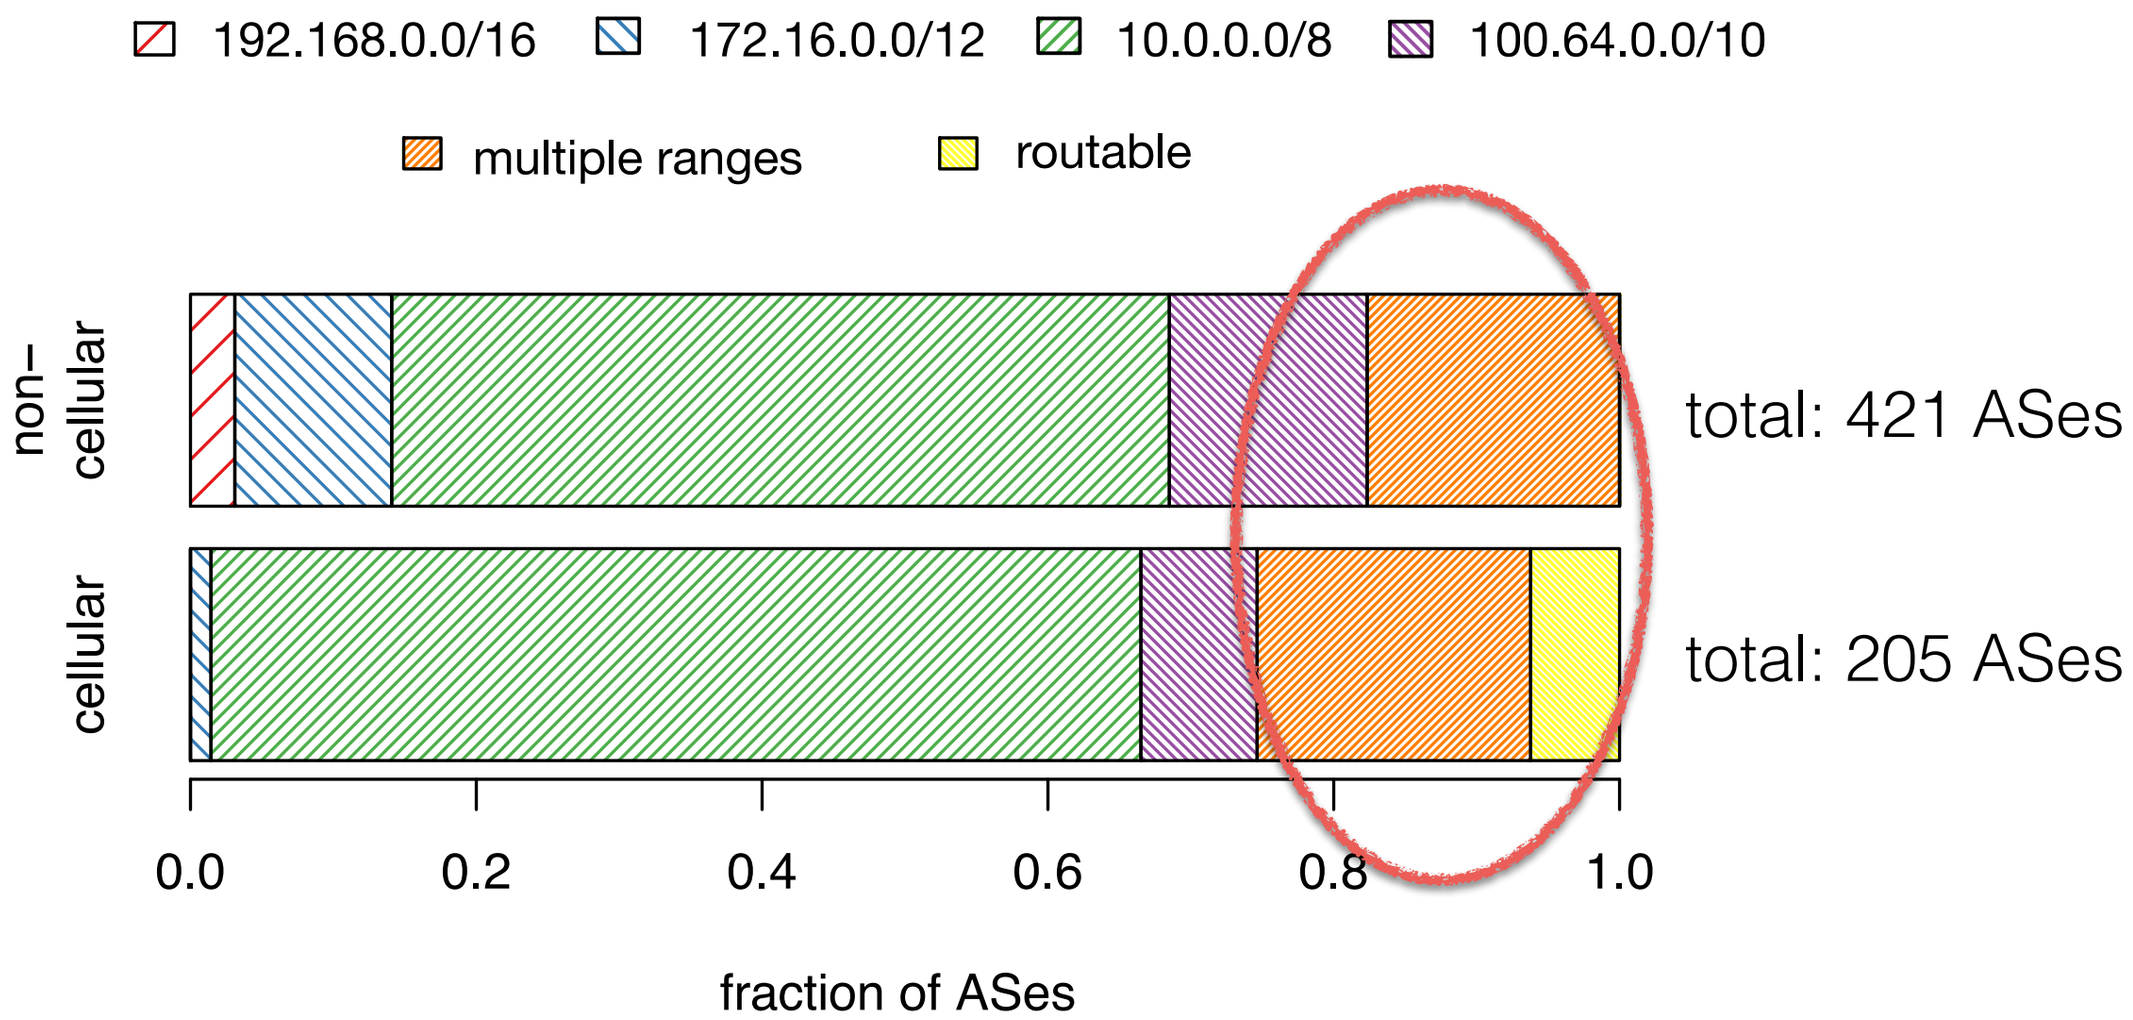

More than 20% of the ASes use multiple internal ranges. Fragmentation/Shortage of Internal Address Space?

CGNs: Routable as Internal Address Space

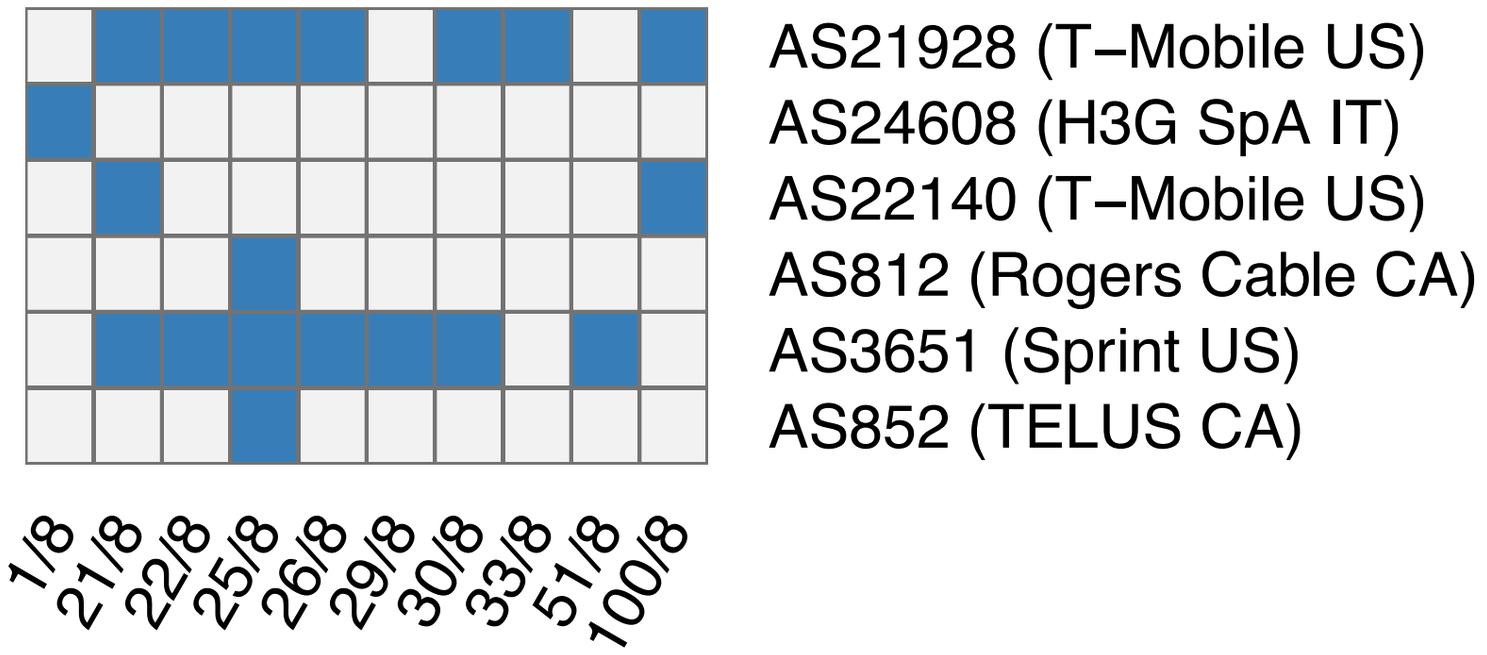

AS21928 (T–Mobile US)
AS24608 (H3G SpA IT)
AS22140 (T–Mobile US)
AS812 (Rogers Cable CA)
AS3651 (Sprint US)
AS852 (TELUS CA)

CGNs: Routable as Internal Address Space

e.g., 25.0.0.0/8: mostly unrouted,
but in internal use by **at least** 4 major networks.
What happens if somebody wants to route it?

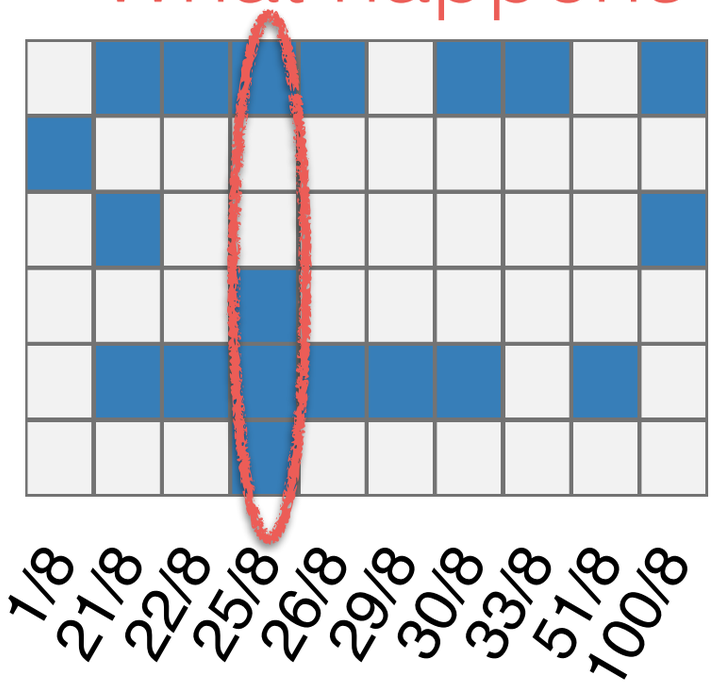

AS21928 (T–Mobile US)
AS24608 (H3G SpA IT)
AS22140 (T–Mobile US)
AS812 (Rogers Cable CA)
AS3651 (Sprint US)
AS852 (TELUS CA)

CGNs: Routable as Internal Address Space

e.g., 25.0.0.0/8: mostly unrouted,
but in internal use by **at least** 4 major networks.
What happens if somebody wants to route it?

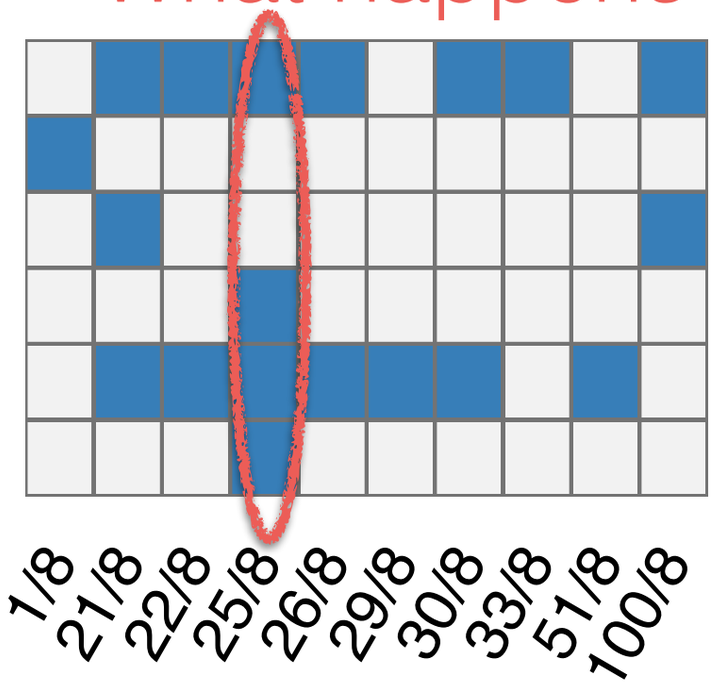

AS21928 (T–Mobile US)
AS24608 (H3G SpA IT)
AS22140 (T–Mobile US)
AS812 (Rogers Cable CA)
AS3651 (Sprint US)
AS852 (TELUS CA)

Consideration for buyers of address space!
Users in major ISPs will likely experience connectivity issues to these address blocks.

CGNs: Extracting More Properties

10 subsequent TCP connections

→ how do CGNs allocate ports and IPs

→ estimate port-chunk per subscriber

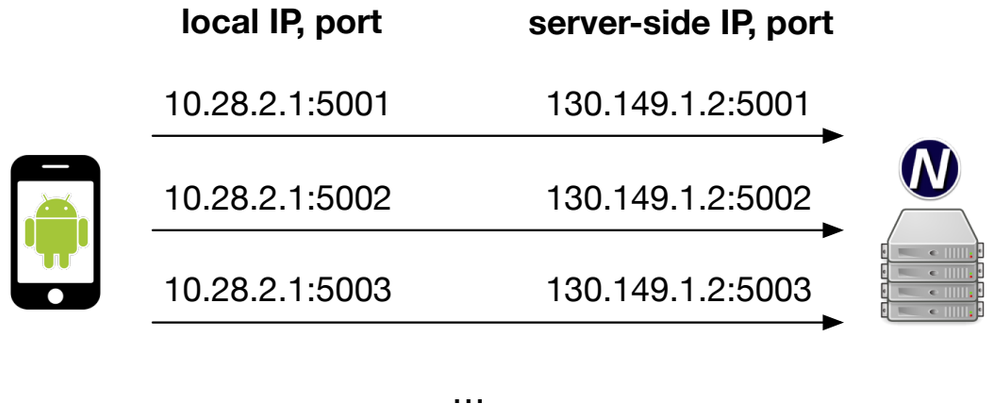

NAT test using TTL-limited probe packets

→ pinpoint the CGN location

→ extract CGN timeout values

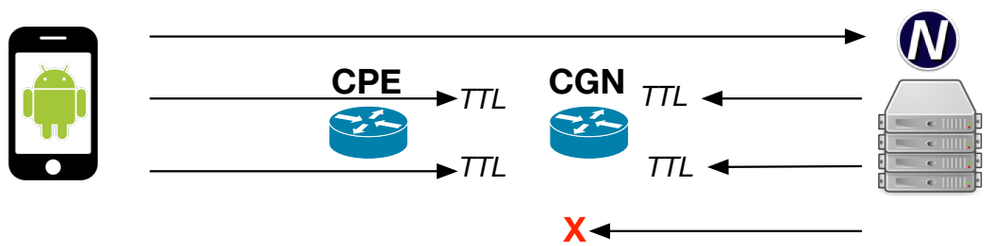

STUN test

→ reason about CGN mapping types

→ compare CGN and CPE mappings

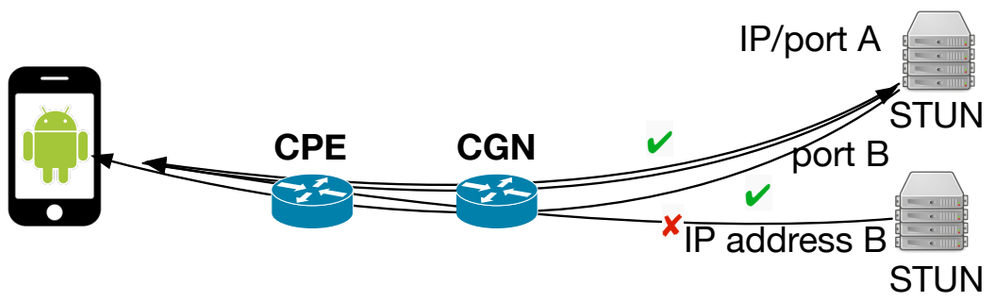

IP Address and Port Allocation

Arbitrary Pooling Behavior

-> Public-facing IP address changes for subsequent connections

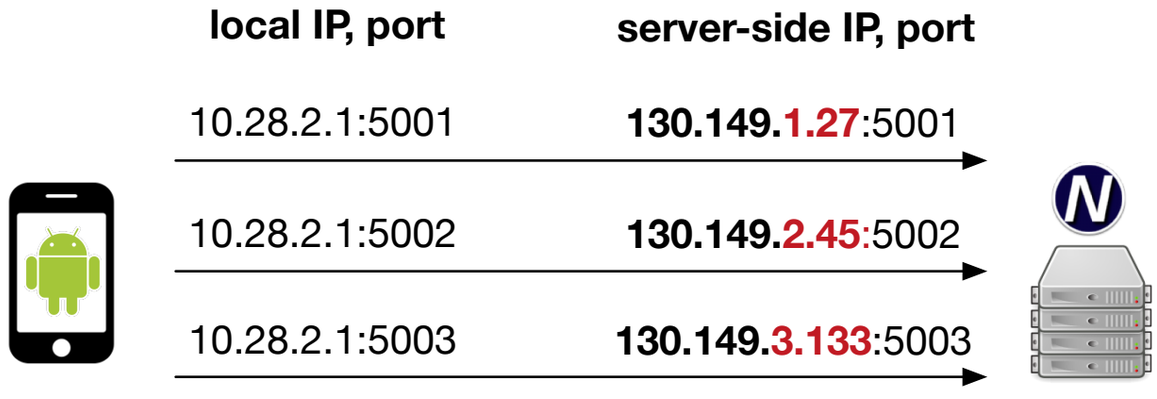

~**20%** of ASes arbitrary pooling

IP Address and Port Allocation

Arbitrary Pooling Behavior

-> Public-facing IP address changes for subsequent connections

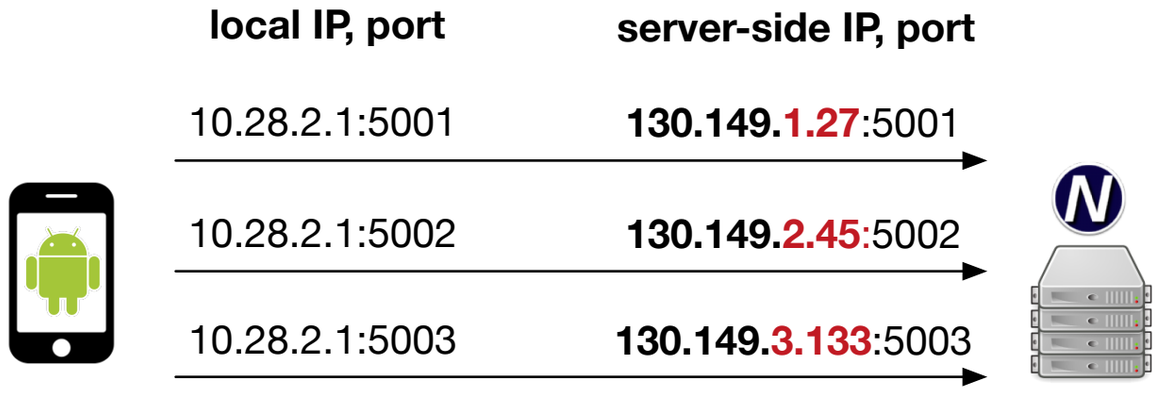

~**20%** of ASes arbitrary pooling

Port Allocation Behavior

-> No dominant strategy; often even inconsistent within the same AS

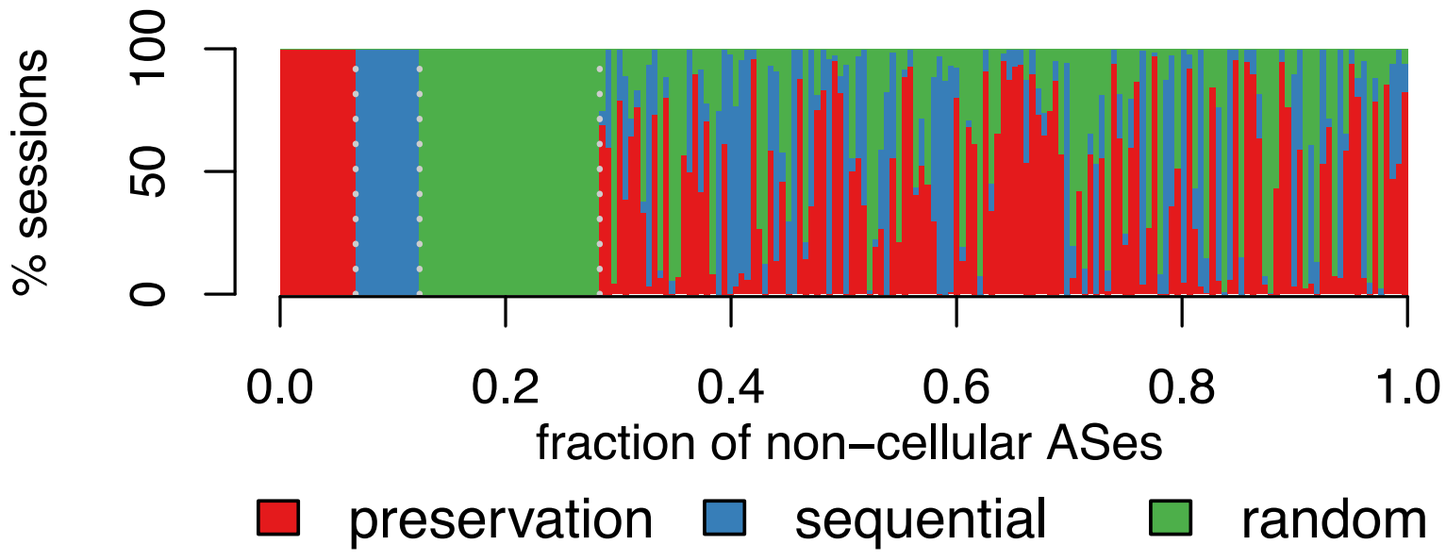

~**70%** of ASes mixed strategies

IP Address and Port Allocation

Huge diversity of address/port allocation strategies.

Most ASes show non-uniform behavior.

Think of Applications, Host Reputation Systems, Attribution.

Arbitrary Pooling Behavior

-> Public-facing IP address changes for subsequent connections

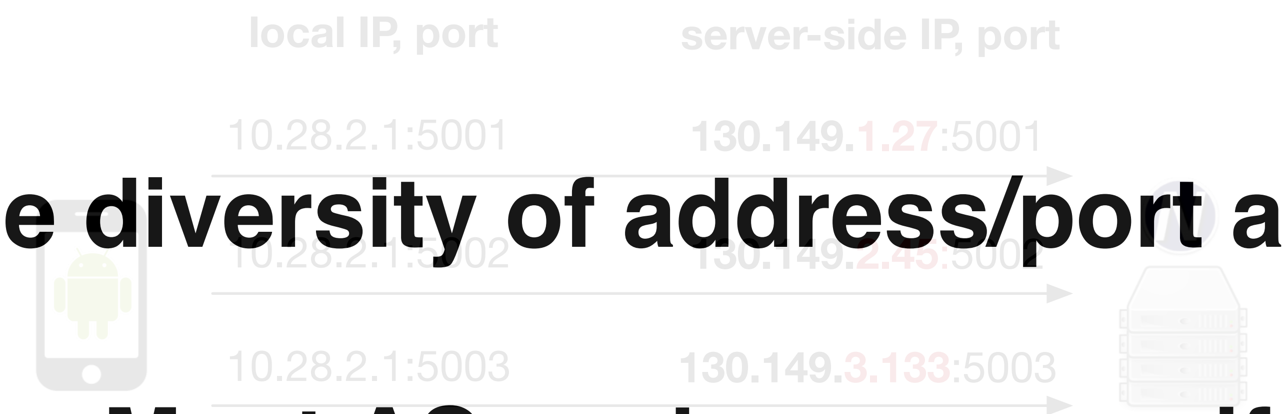

arbitrary pooling

Port Allocation Behavior

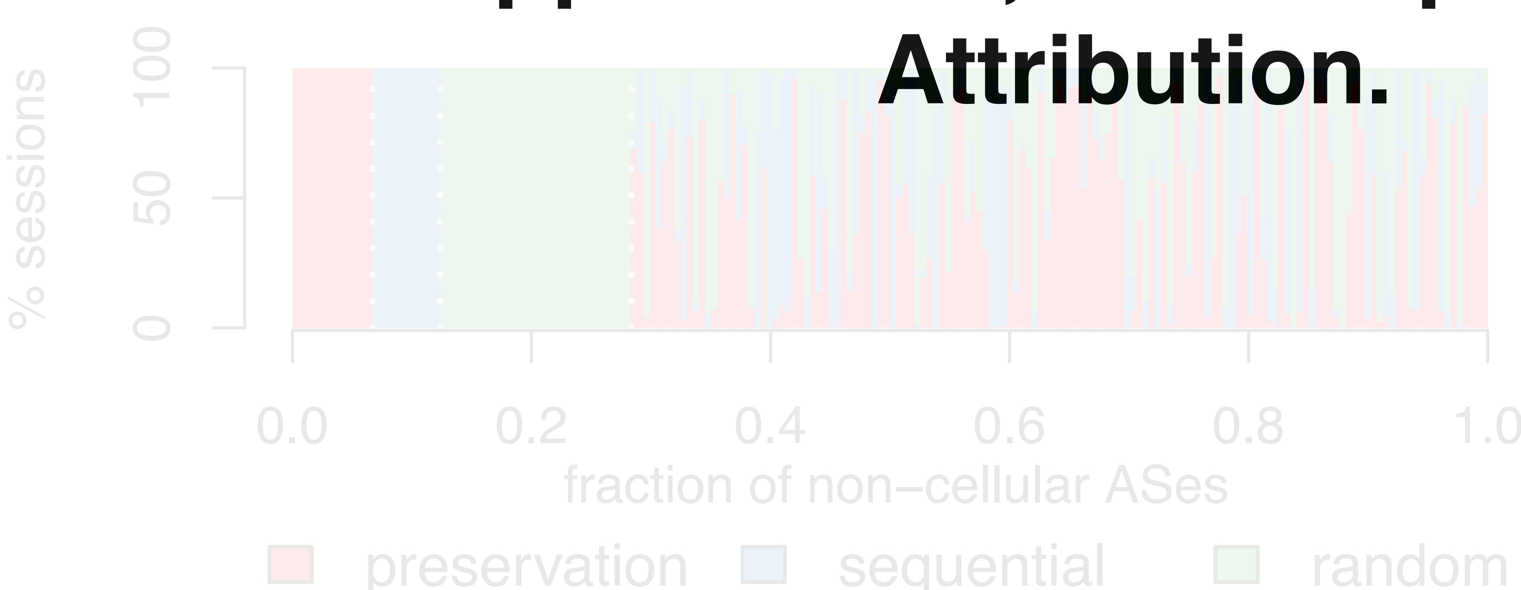

~**70%** of ASes mixed strategies

Chunk-based Port Allocations

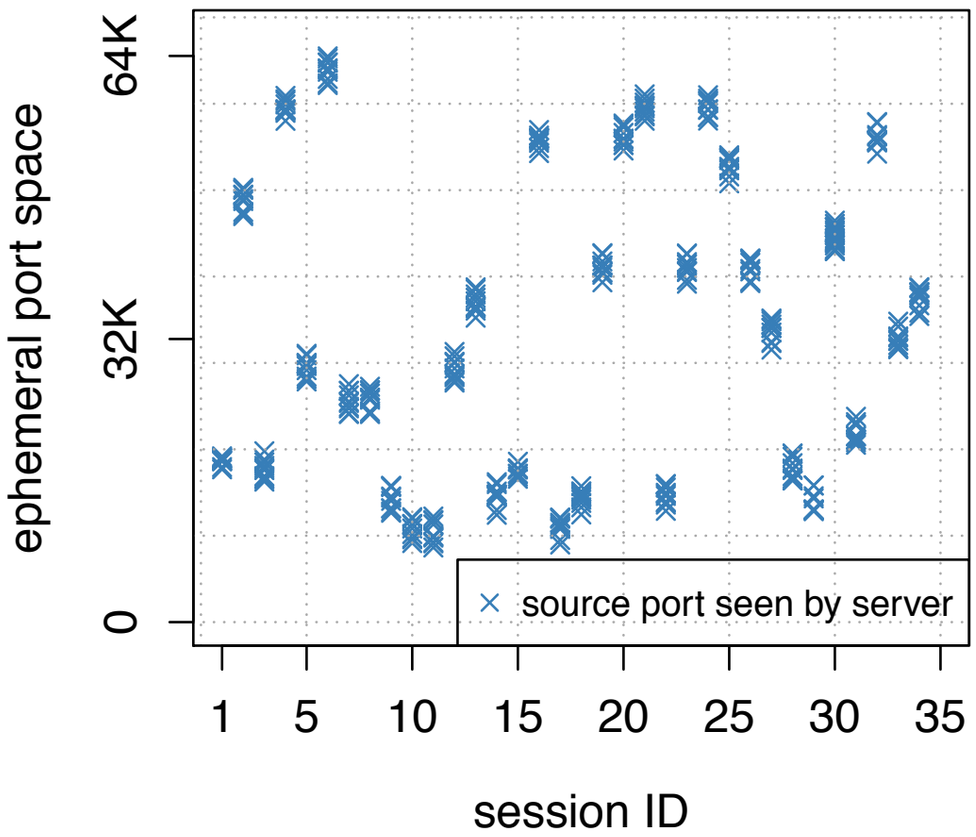

Some ASes: Chunk-based allocation
Down to 512 ports / subscriber -> 128 subscribers per IP

NAT Mapping / Filtering Behavior

CPE NATs

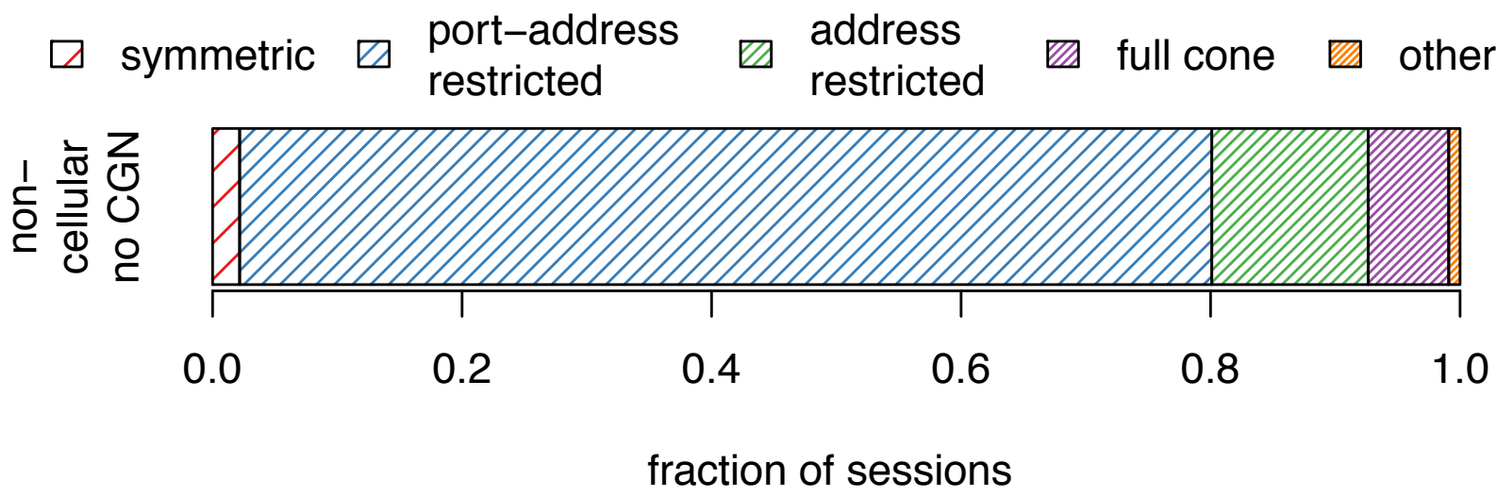

NAT Mapping / Filtering Behavior

CPE NATs

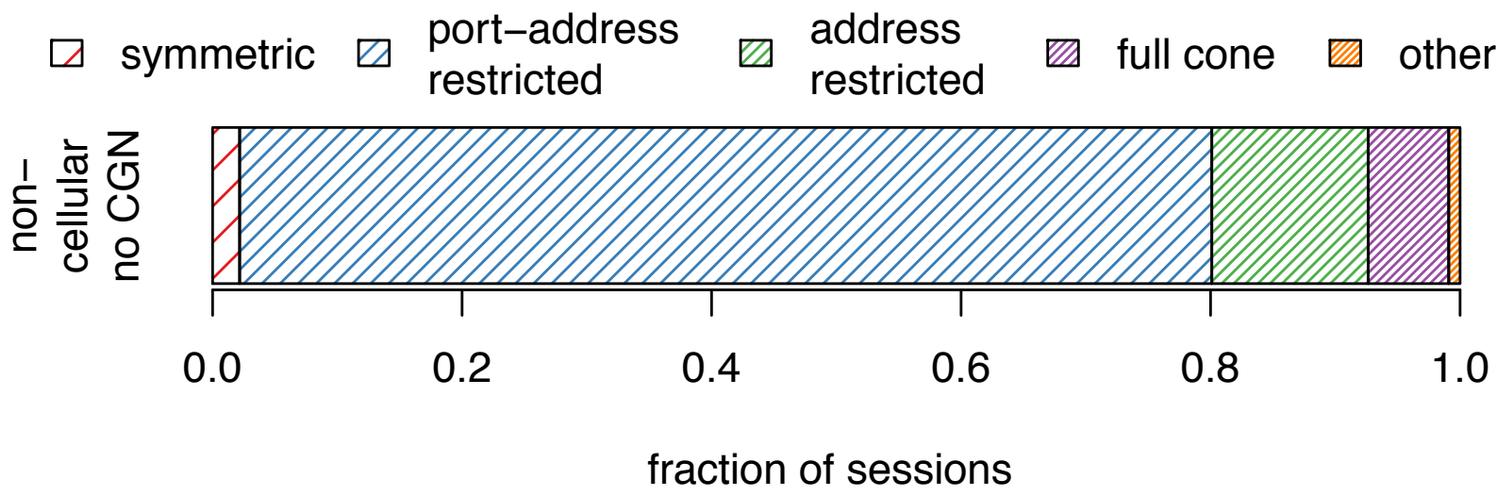

Carrier-Grade NATs

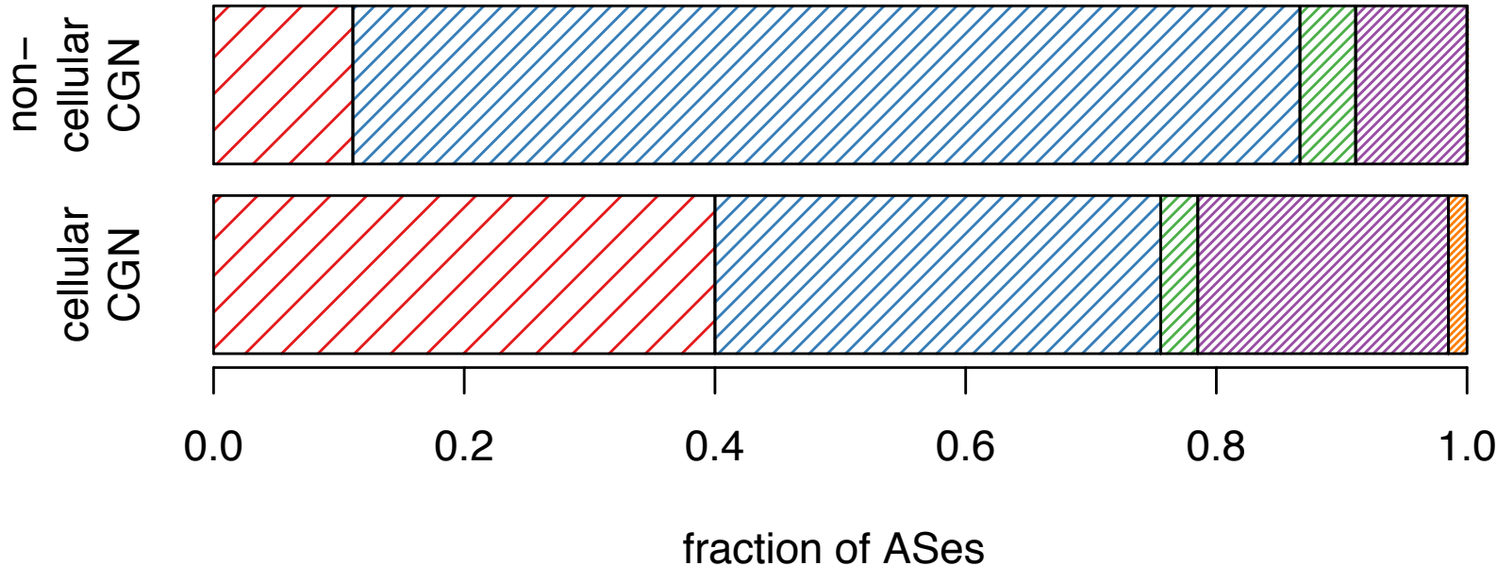

NAT Mapping / Filtering Behavior

CPE NATs

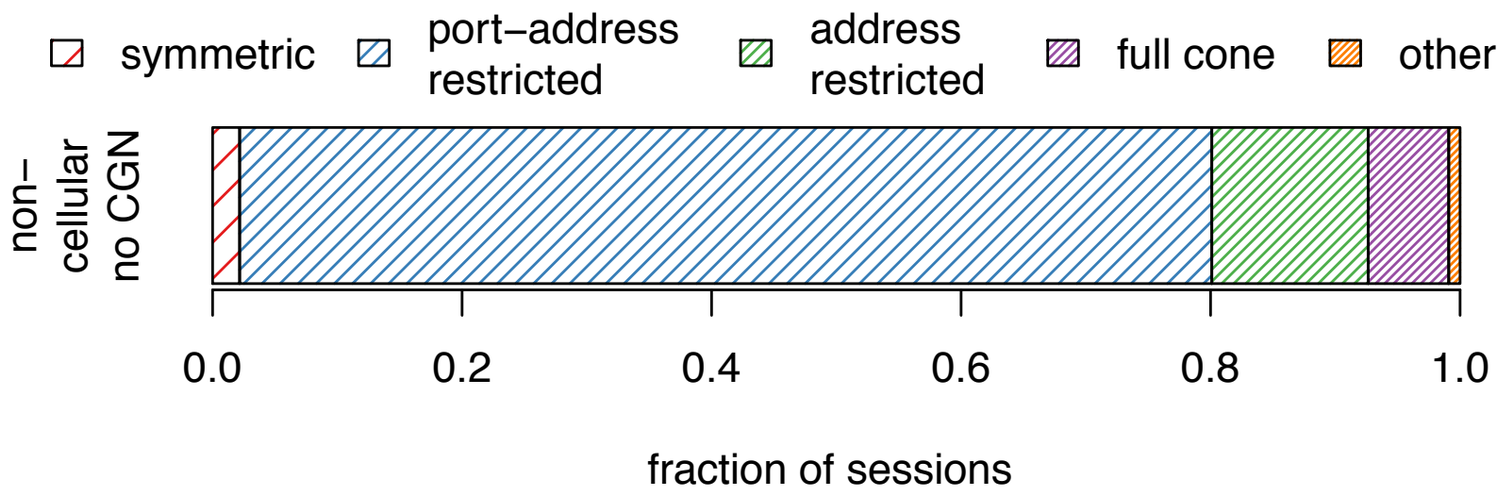

Carrier-Grade NATs

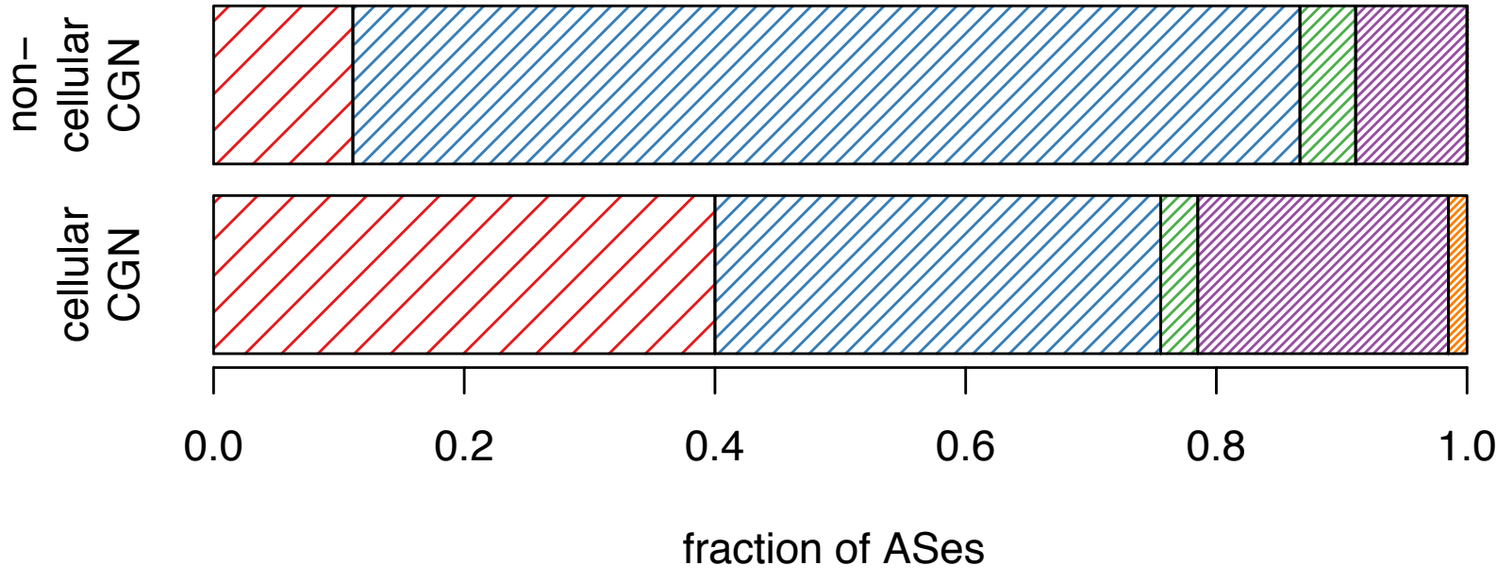

CGN mapping behavior often more restrictive than that of CPE routers

CGN Deployment and Impact

High-Level Overview

- Broadly deployed, CGNs are reality for many users!
- Stunning variety of configurations and setups across ASes and within the same AS
- Degree of resource sharing, IP addresses, ports, varies heavily, down to 512 ports / subscriber
- NAT mappings and timeouts of some CGNs more restrictive compared to CPEs

CGN Deployment and Impact

High-Level Overview

- Broadly deployed, CGNs are reality for many users!
- Stunning variety of configurations and setups across ASes and within the same AS
- Degree of resource sharing, IP addresses, ports, varies heavily, down to 512 ports / subscriber
- NAT mappings and timeouts of some CGNs more restrictive compared to CPEs

CGNs limit the resources available for subscribers
CGN means very different things for different ISPs

CGN Challenges

Measuring End-User Internet Performance

Common metrics

* Speed
* Latency
* Packet Loss

don’t capture limitations imposed by CGNs

New metrics?

* Maximum concurrent connections?
* Types of NAT mappings?

Guidelines / Transparency / Regulation?

- CGNs reduce “how much Internet” subscribers receive
  - Need for guidelines for resource allocation?
  - Need for regulation?

CGN Deployment and Impact

High-Level Overview

- Broadly deployed, CGNs are reality for many users!
- Stunning variety of configurations and setups across ASes and within the same AS
- Degree of resource sharing, IP addresses, ports, varies heavily, down to 512 ports / subscriber
- NAT mappings and timeouts of some CGNs more restrictive compared to CPEs

CGNs limit the resources available for subscribers
CGN means very different things for different ISPs