A Multi-Perspective Analysis of Carrier-Grade NAT Deployment

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

in *ACM IMC 2016*.

IPv4 Address Space Exhaustion

4 out of 5 RIRs exhausted.

Less than ~2% of the IPv4 space is still unallocated/"free".

What happens now and what do we know?

Transition to IPv6

→ plenty of measurements and statistics available

Buy IPv4

→ 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
- Range from small rural ISPs in Africa up to Fortune 50 companies

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

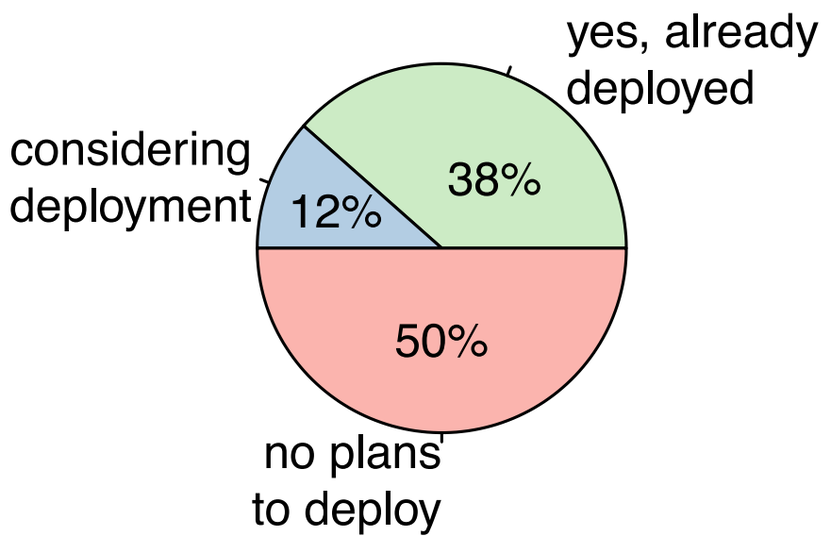

ISP Survey: CGN Specifics

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 and shortage (e.g., RFC1918)

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)

NAT44
(carrier-side)

NAT444
(subscriber-side
and
carrier-side)

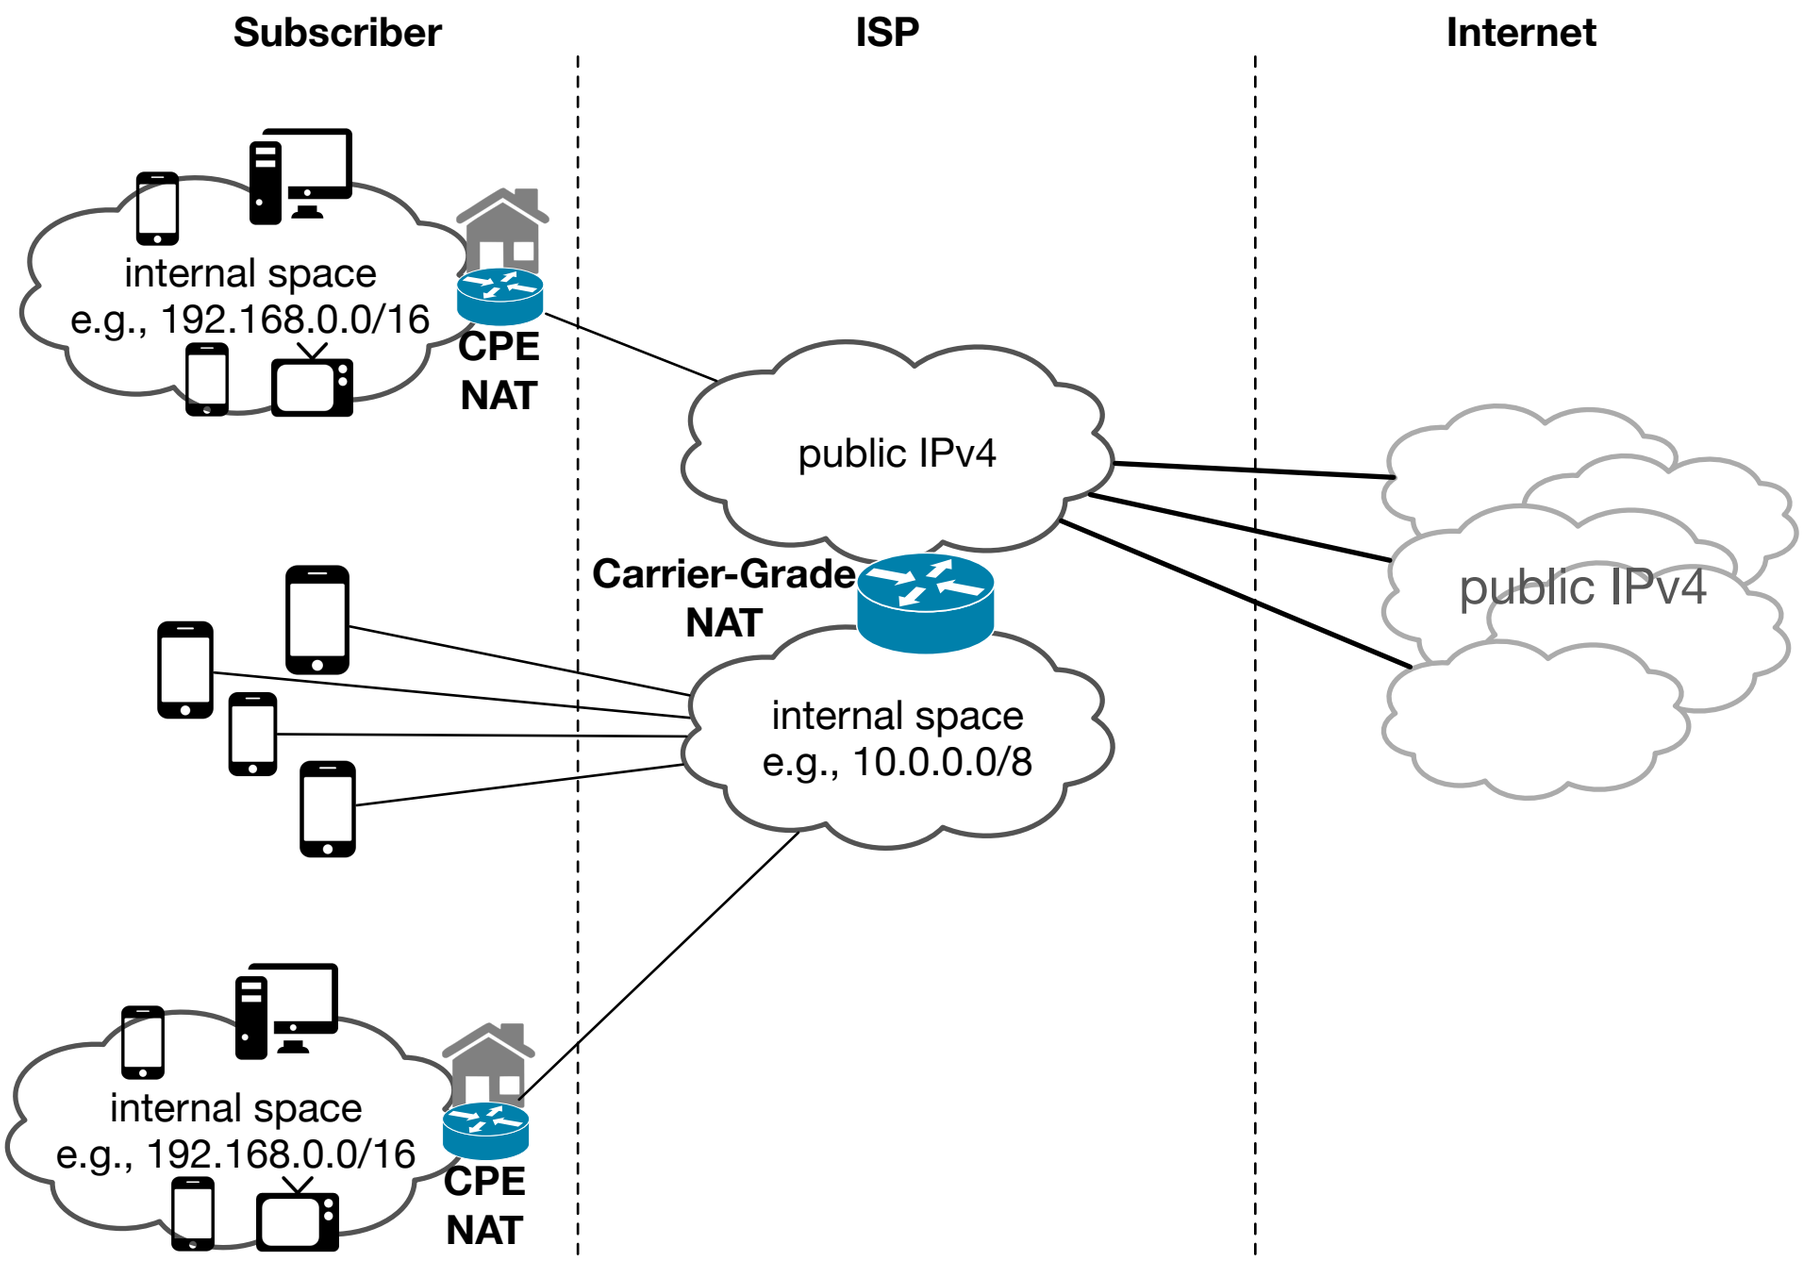

Agenda

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

The BitTorrent DHT

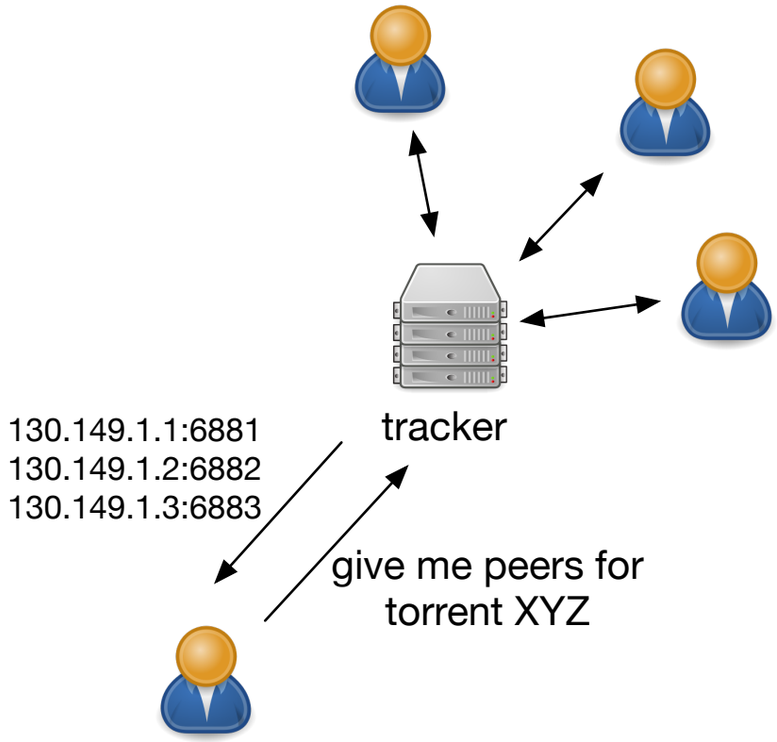

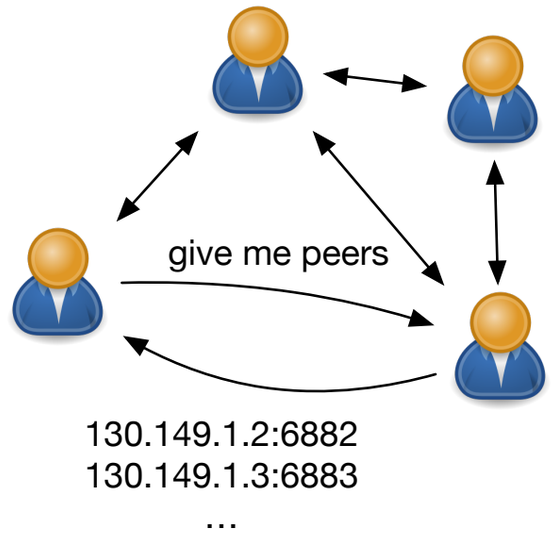

classic BitTorrent
Tracker stores peer contact information (IP:port)

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

We can use DHT peers as vantage points

Crawling the BitTorrent DHT

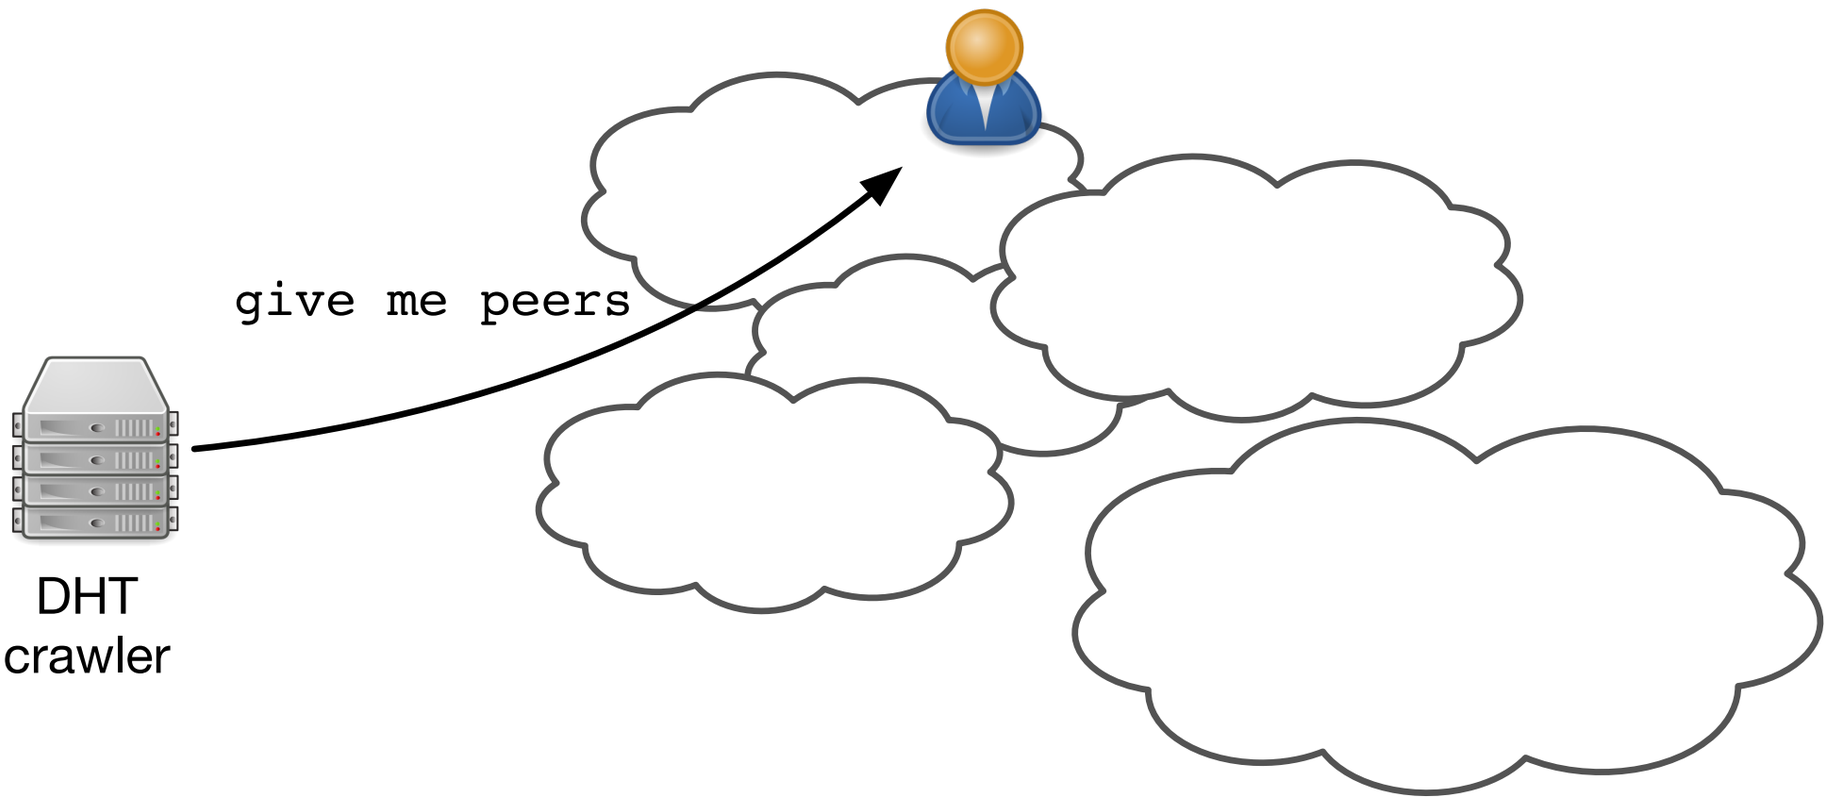

Crawling the BitTorrent DHT

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

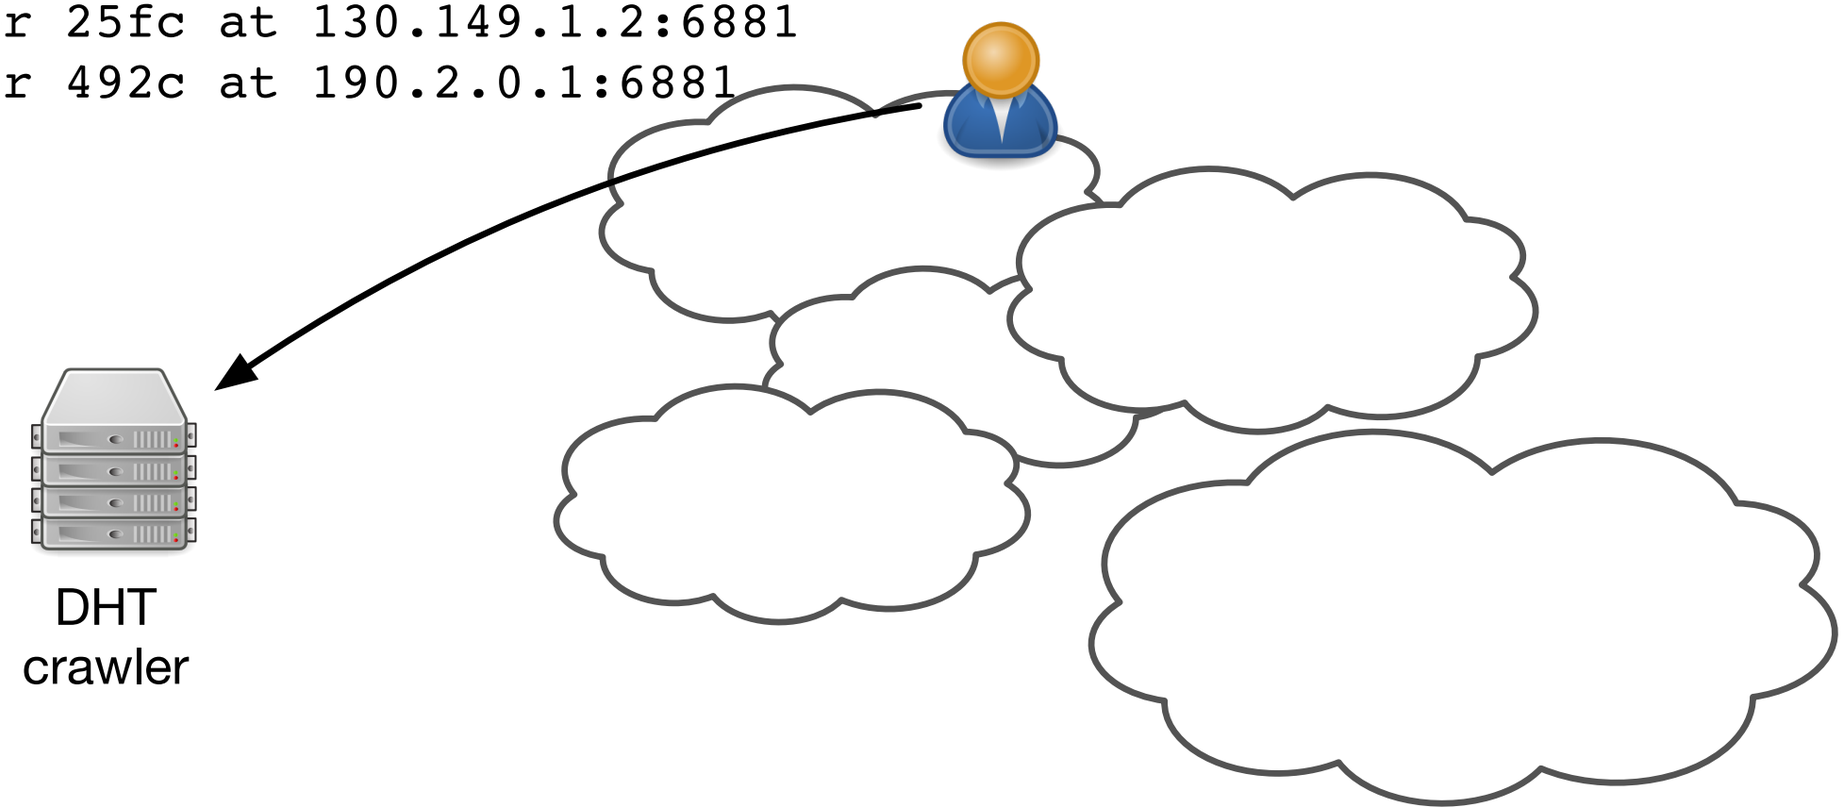

Crawling the BitTorrent DHT

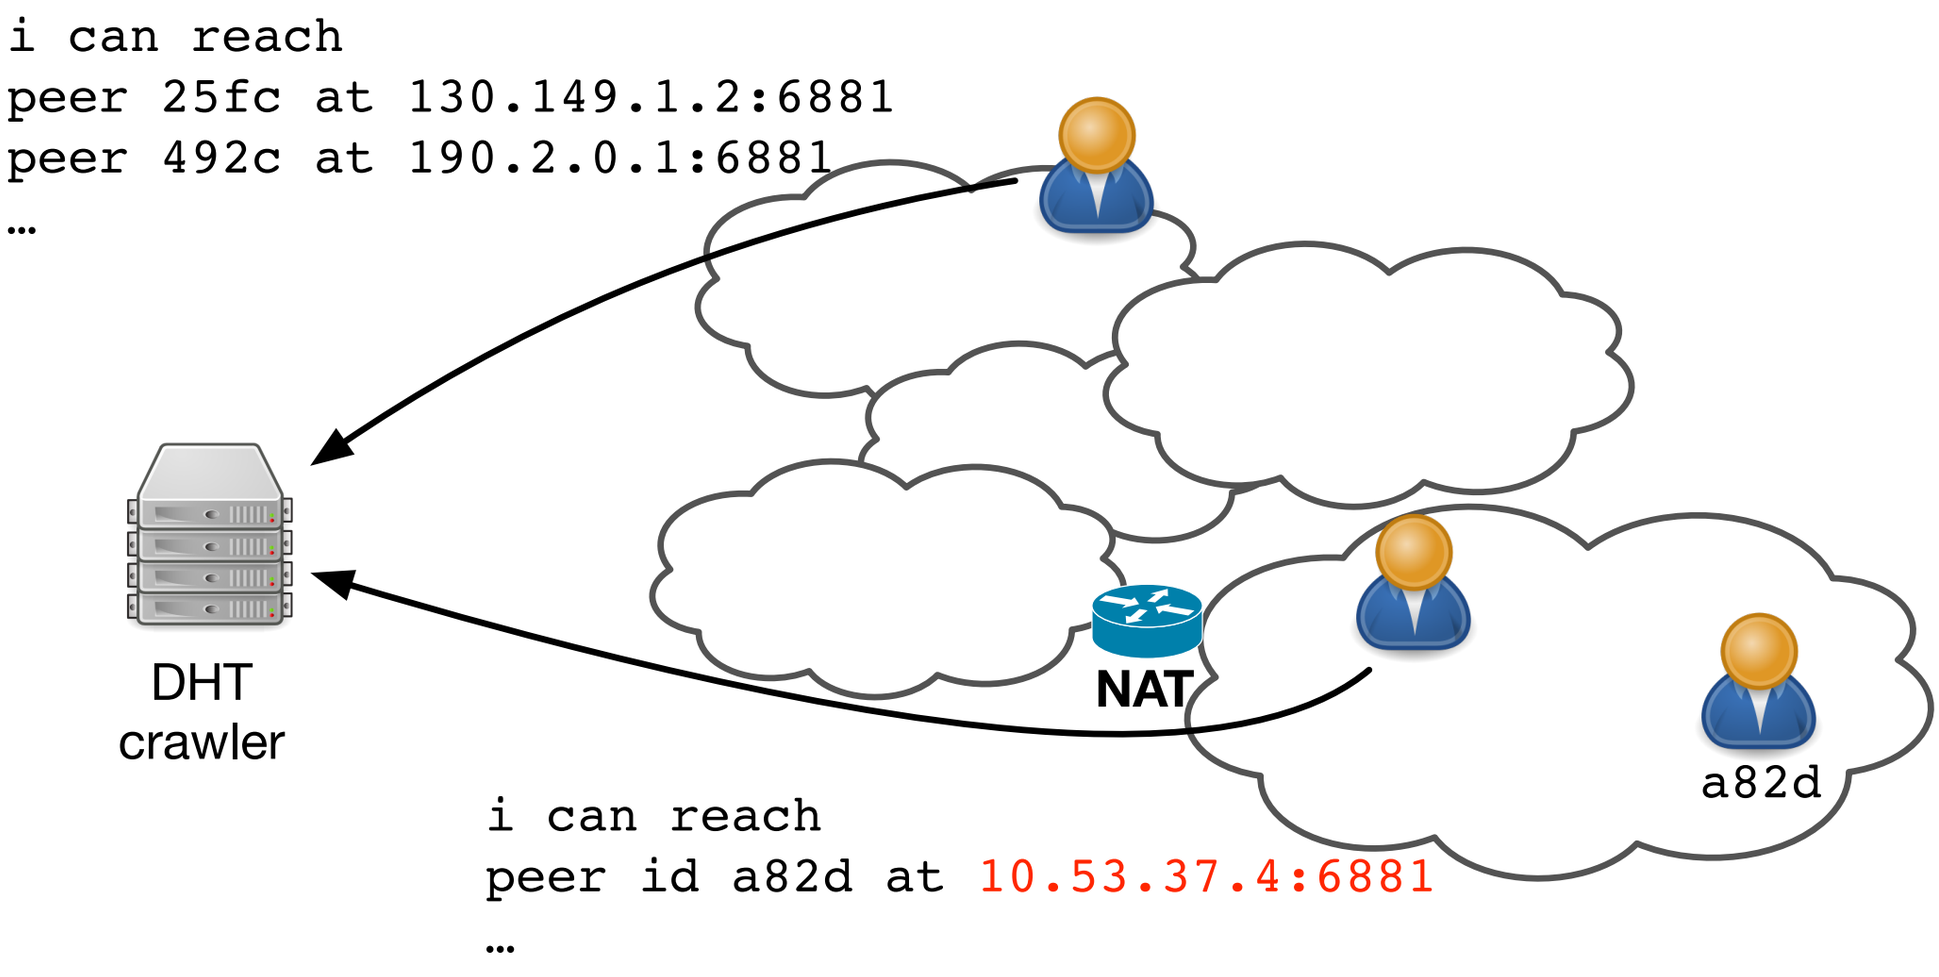

Some peers leak us internal IP addresses of other peers

Crawling the BitTorrent DHT

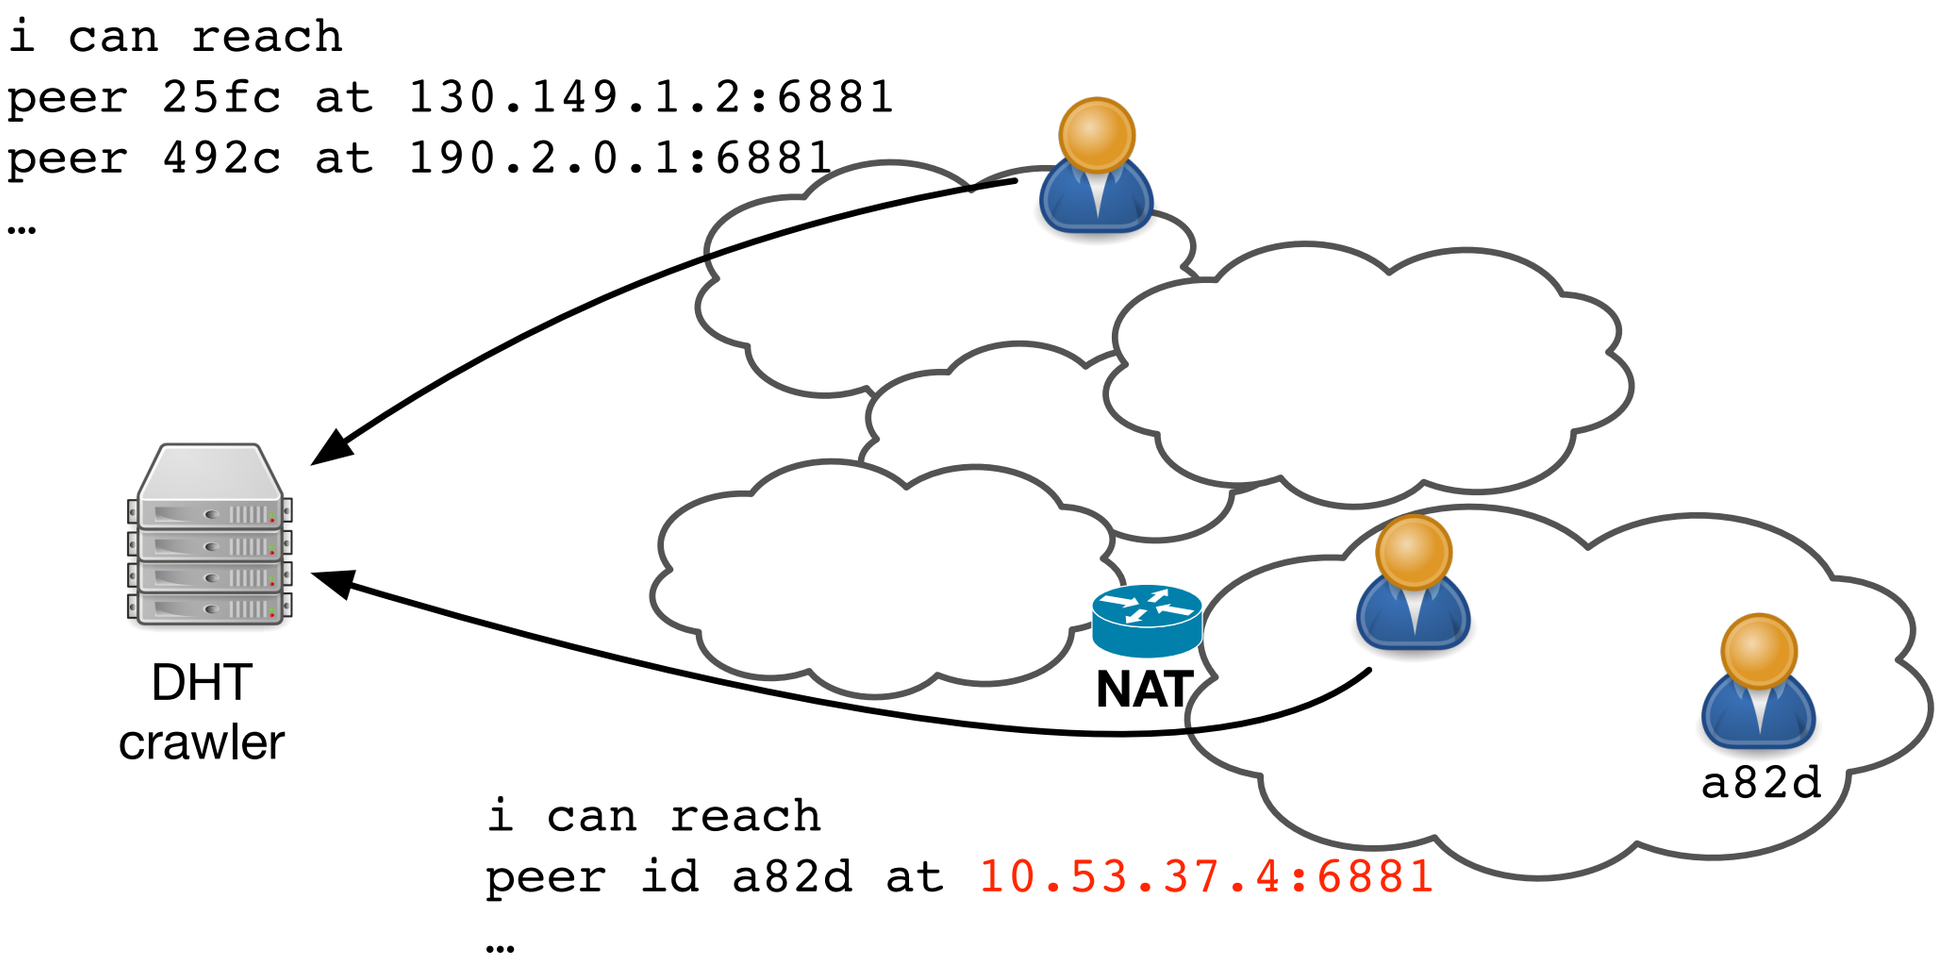

Some peers leak us internal IP addresses of other peers within 1 week: more than 700.000 peers in 5.000 ASes!

Understanding Leakage Relationships

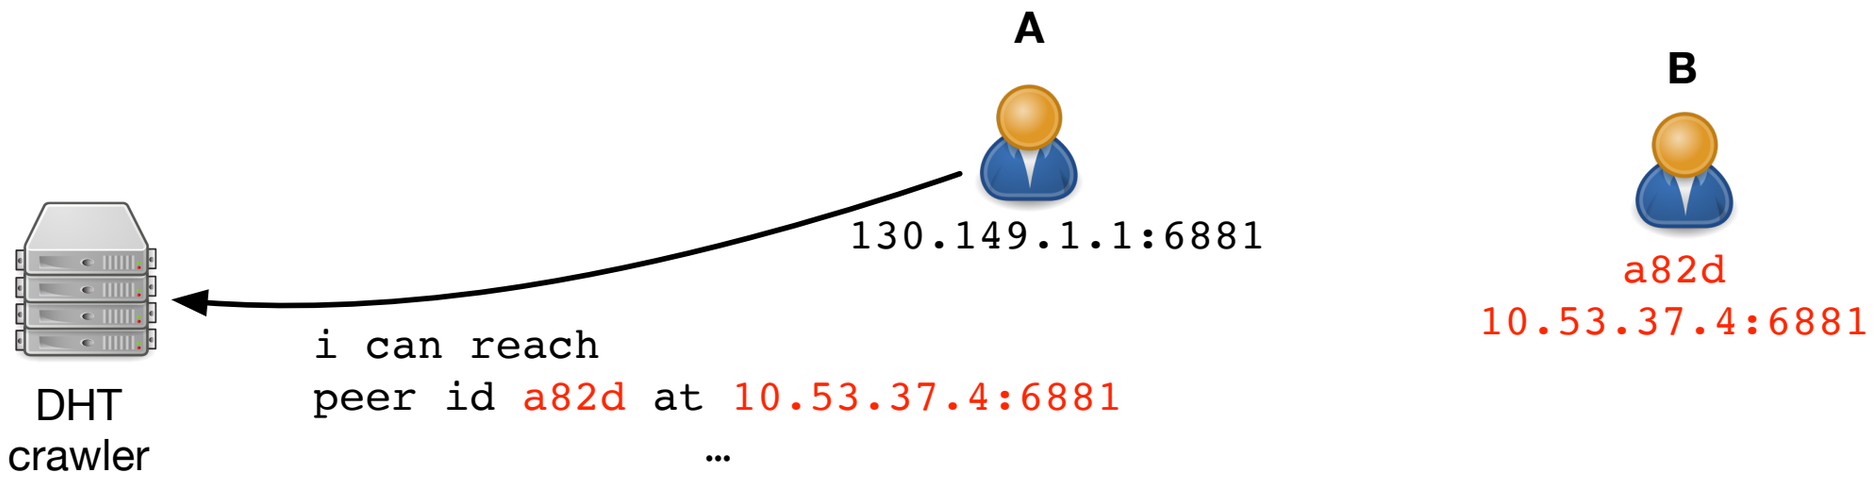

we construct a graph of leaking relationships

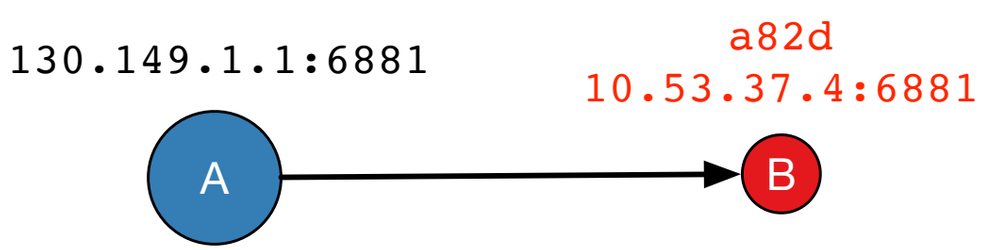

…now we look these graphs on a per-AS basis

BitTorrent Peer Leakage Graph

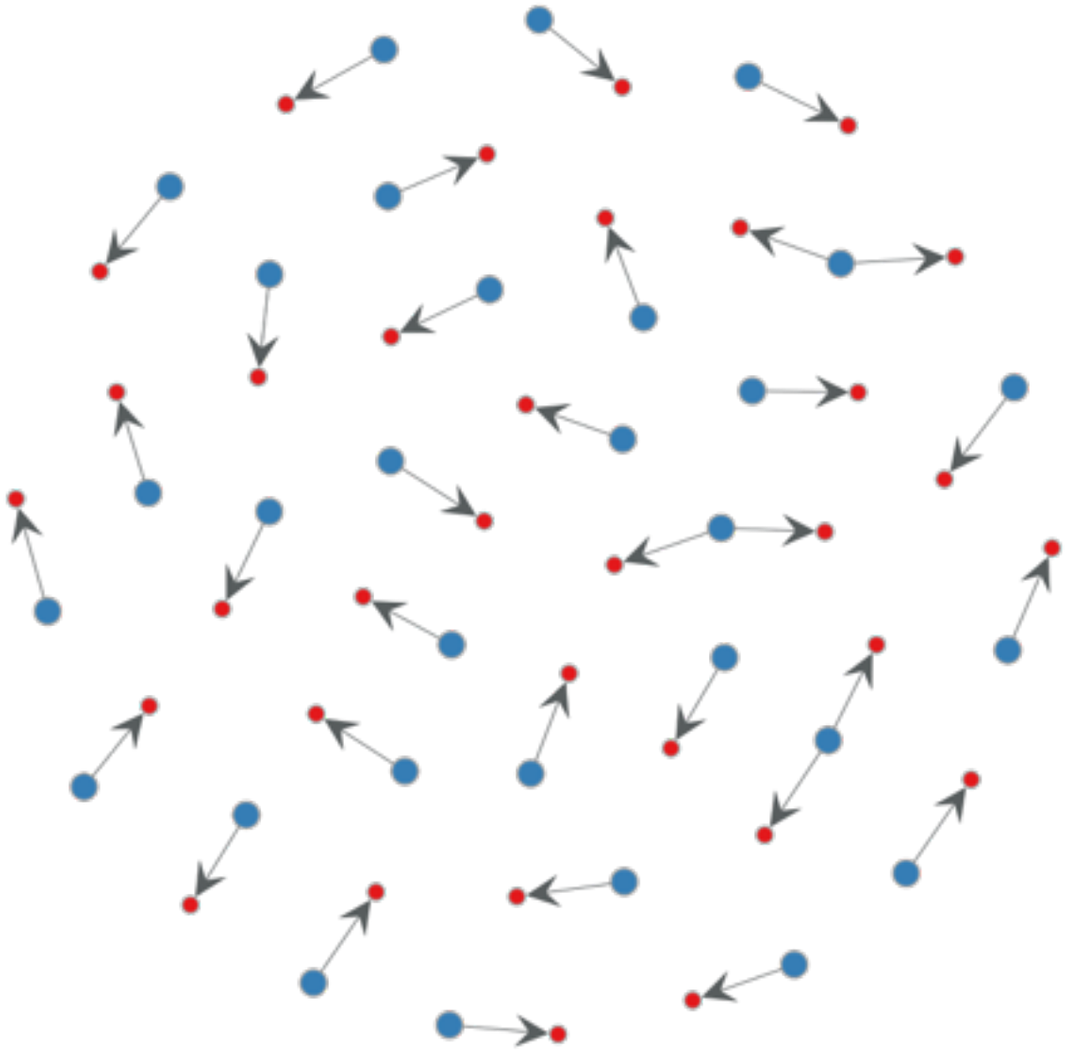

In this AS:
no CGN detected

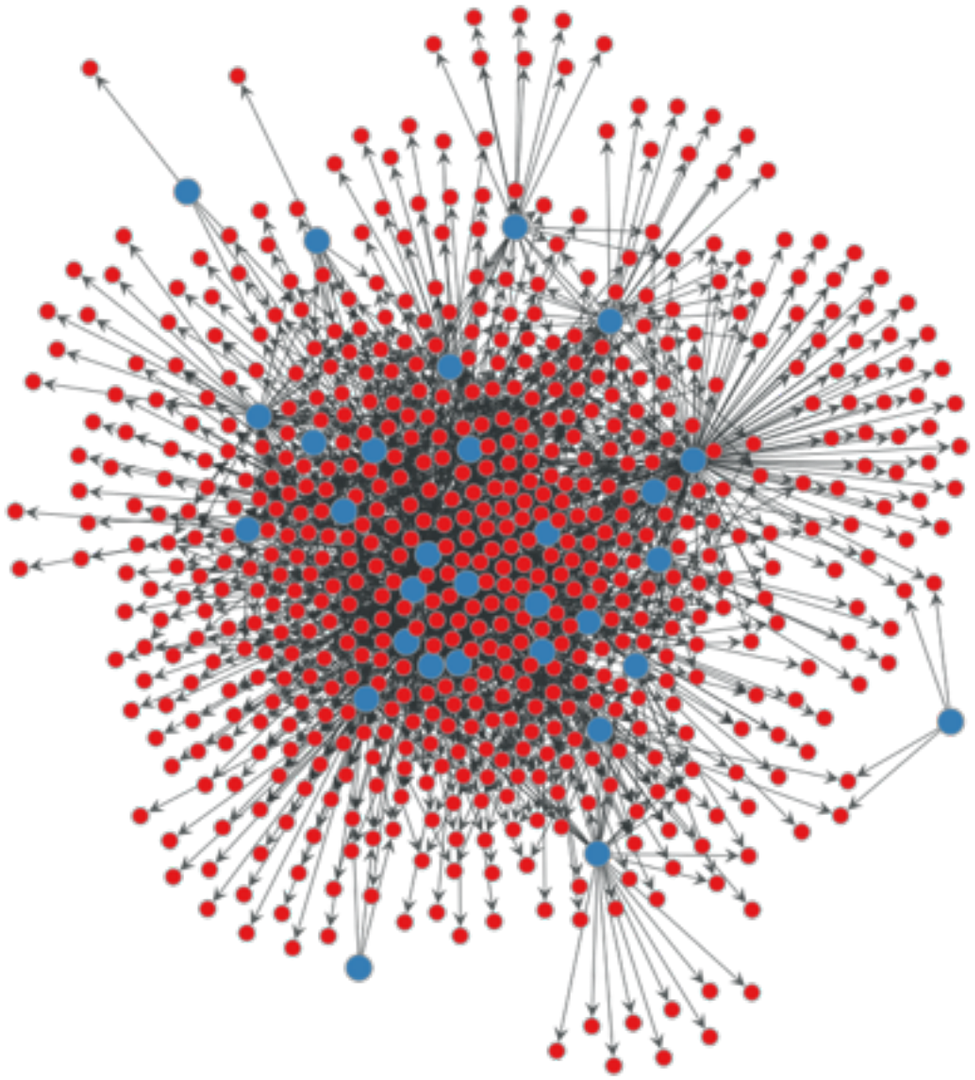

In this AS:
CGN detected

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

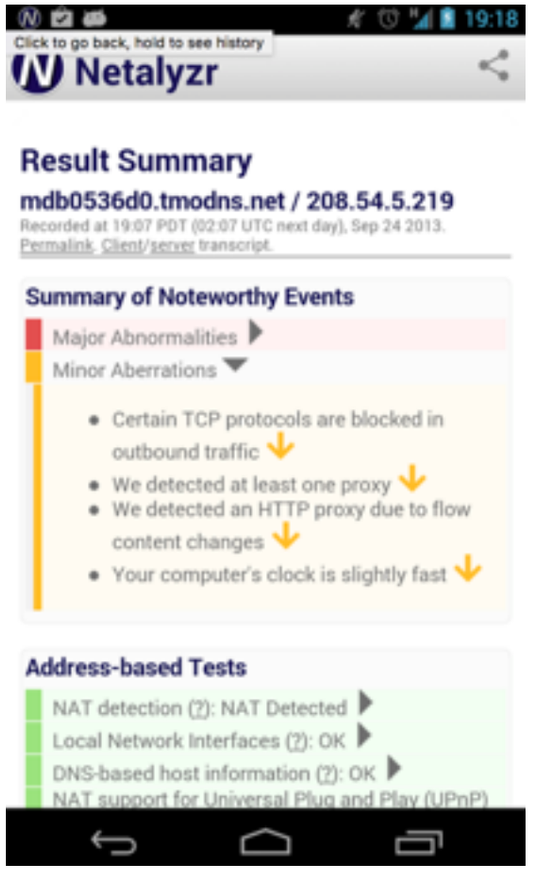

Detecting CGN in Cellular Networks

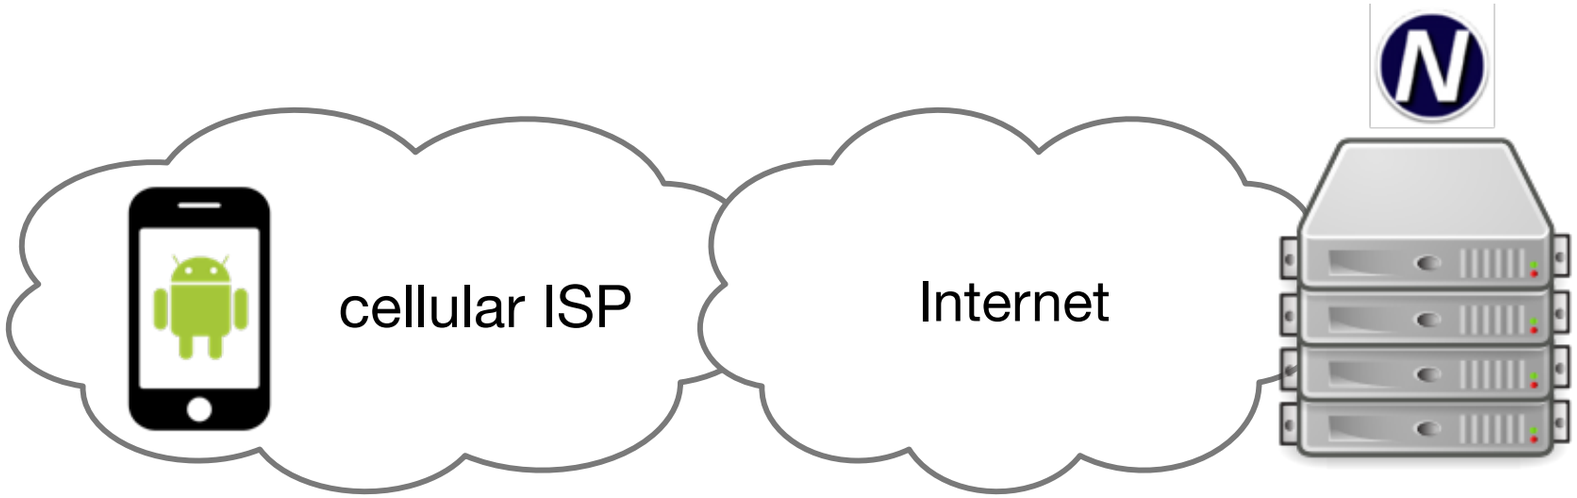

Device IP address assigned directly by the ISP
Device IP ≠ server-side IP → Carrier-Grade NAT

Detecting CGN in Residential Networks

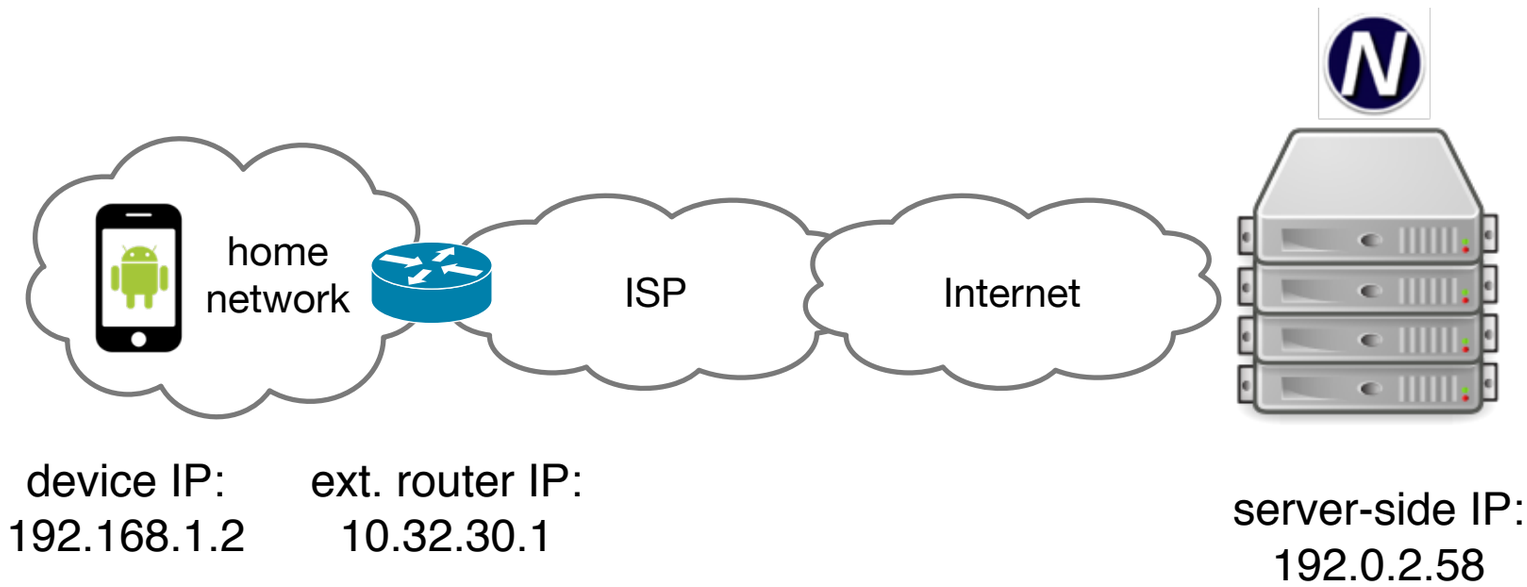

ext. router IP ≠ server-side IP → Carrier-Grade NAT?

Detecting CGN in Residential Networks (2)

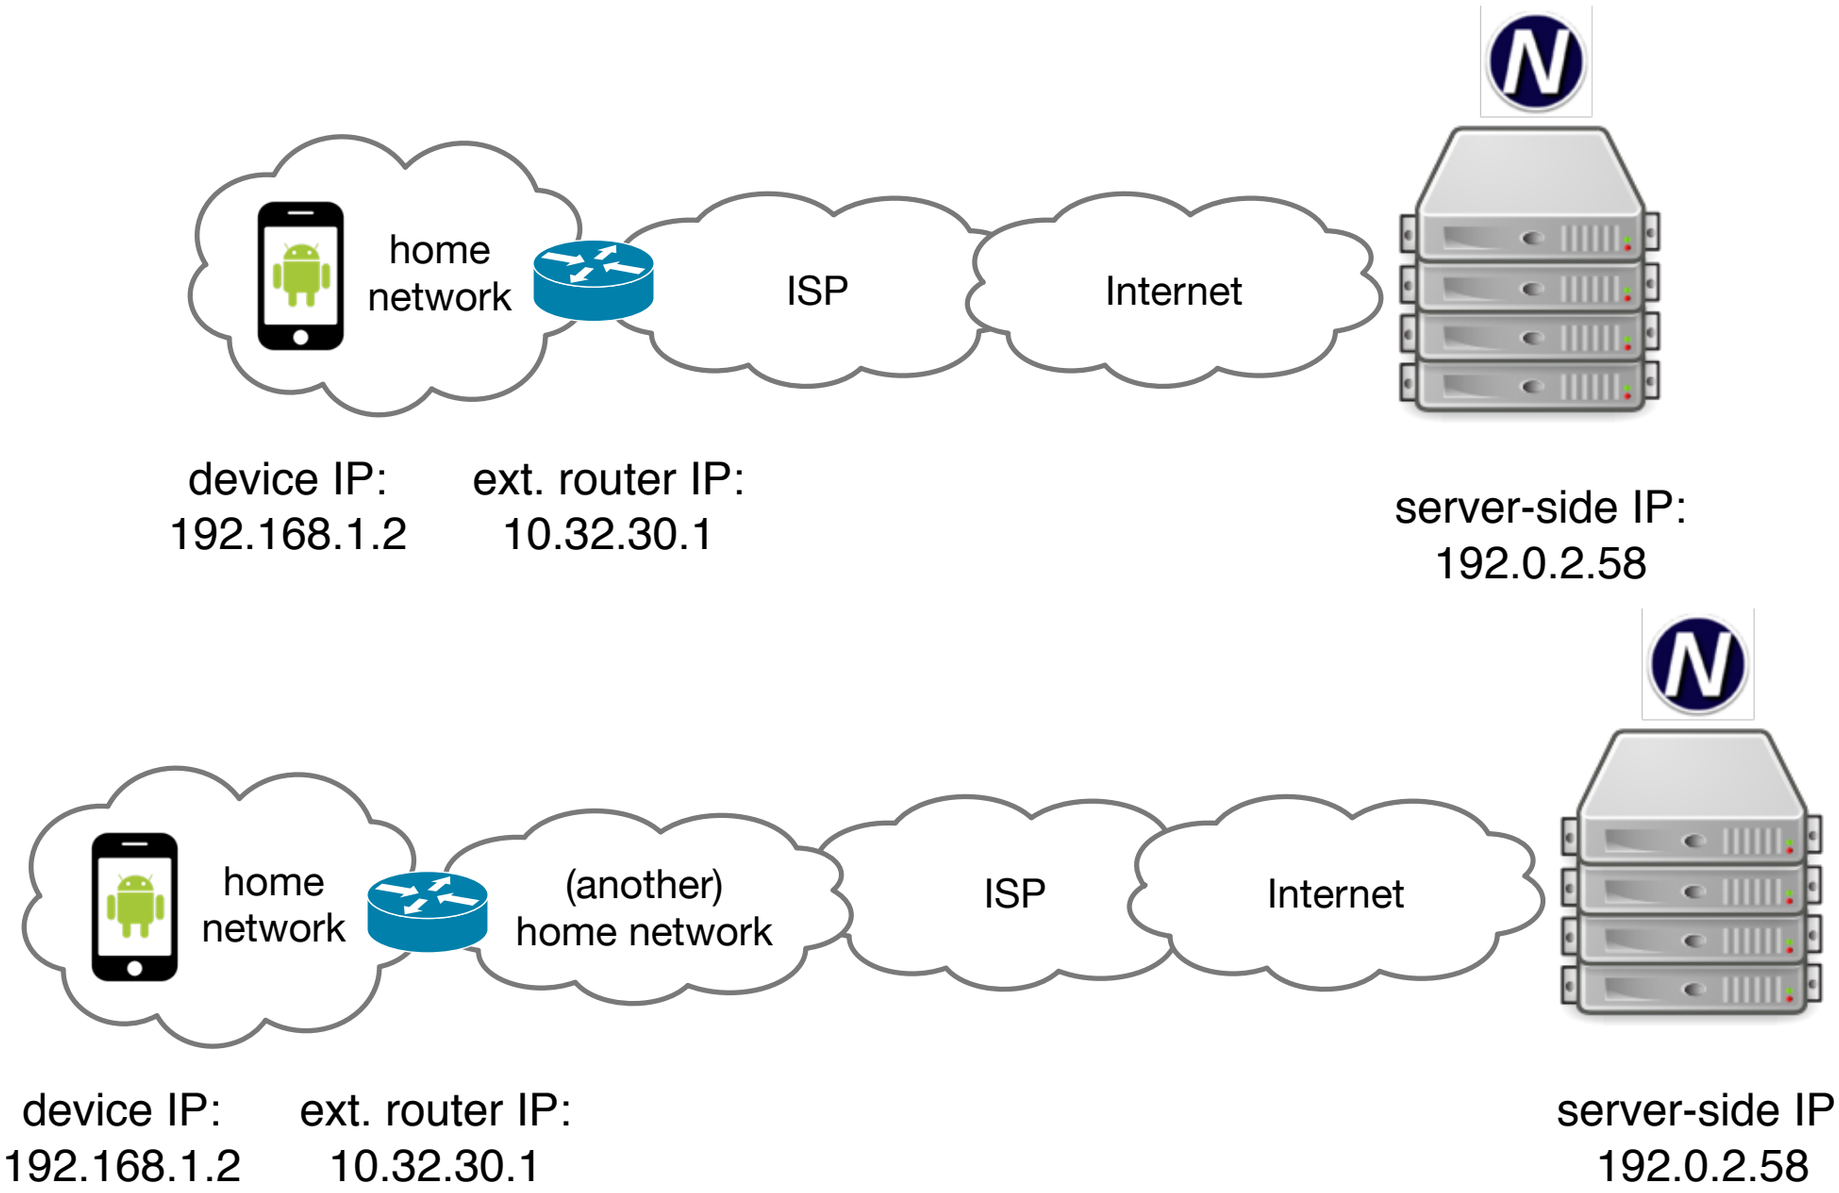

Up to 7% of sessions with chained home NATs

Detecting CGNs with Netalyzr

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

Benefits	Caveats
direct IP addressing data	partial visibility, crowdsourced (need users to run Netalyzr)
cellular and non-cellular	
more customized tests	

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 Labs “aspop”
- Eyeball AS population: 3K ASes
- Tested with BitTorrent/Netalyzr: 1,791 **(62%)**

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

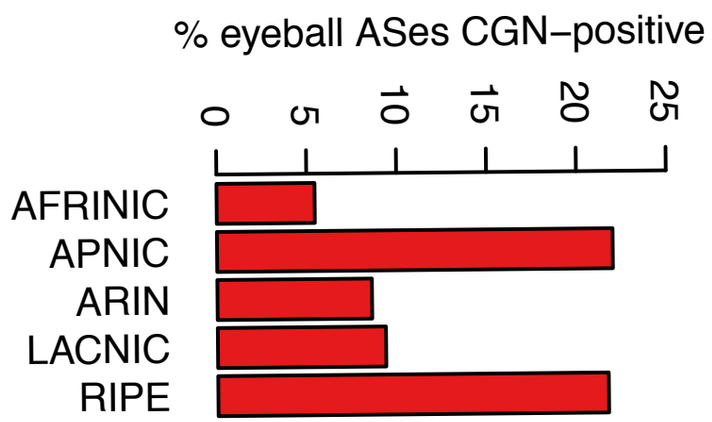

Cellular Networks

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

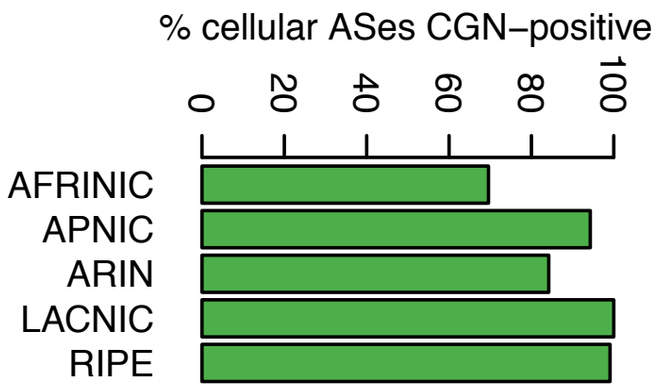

Agenda

- 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

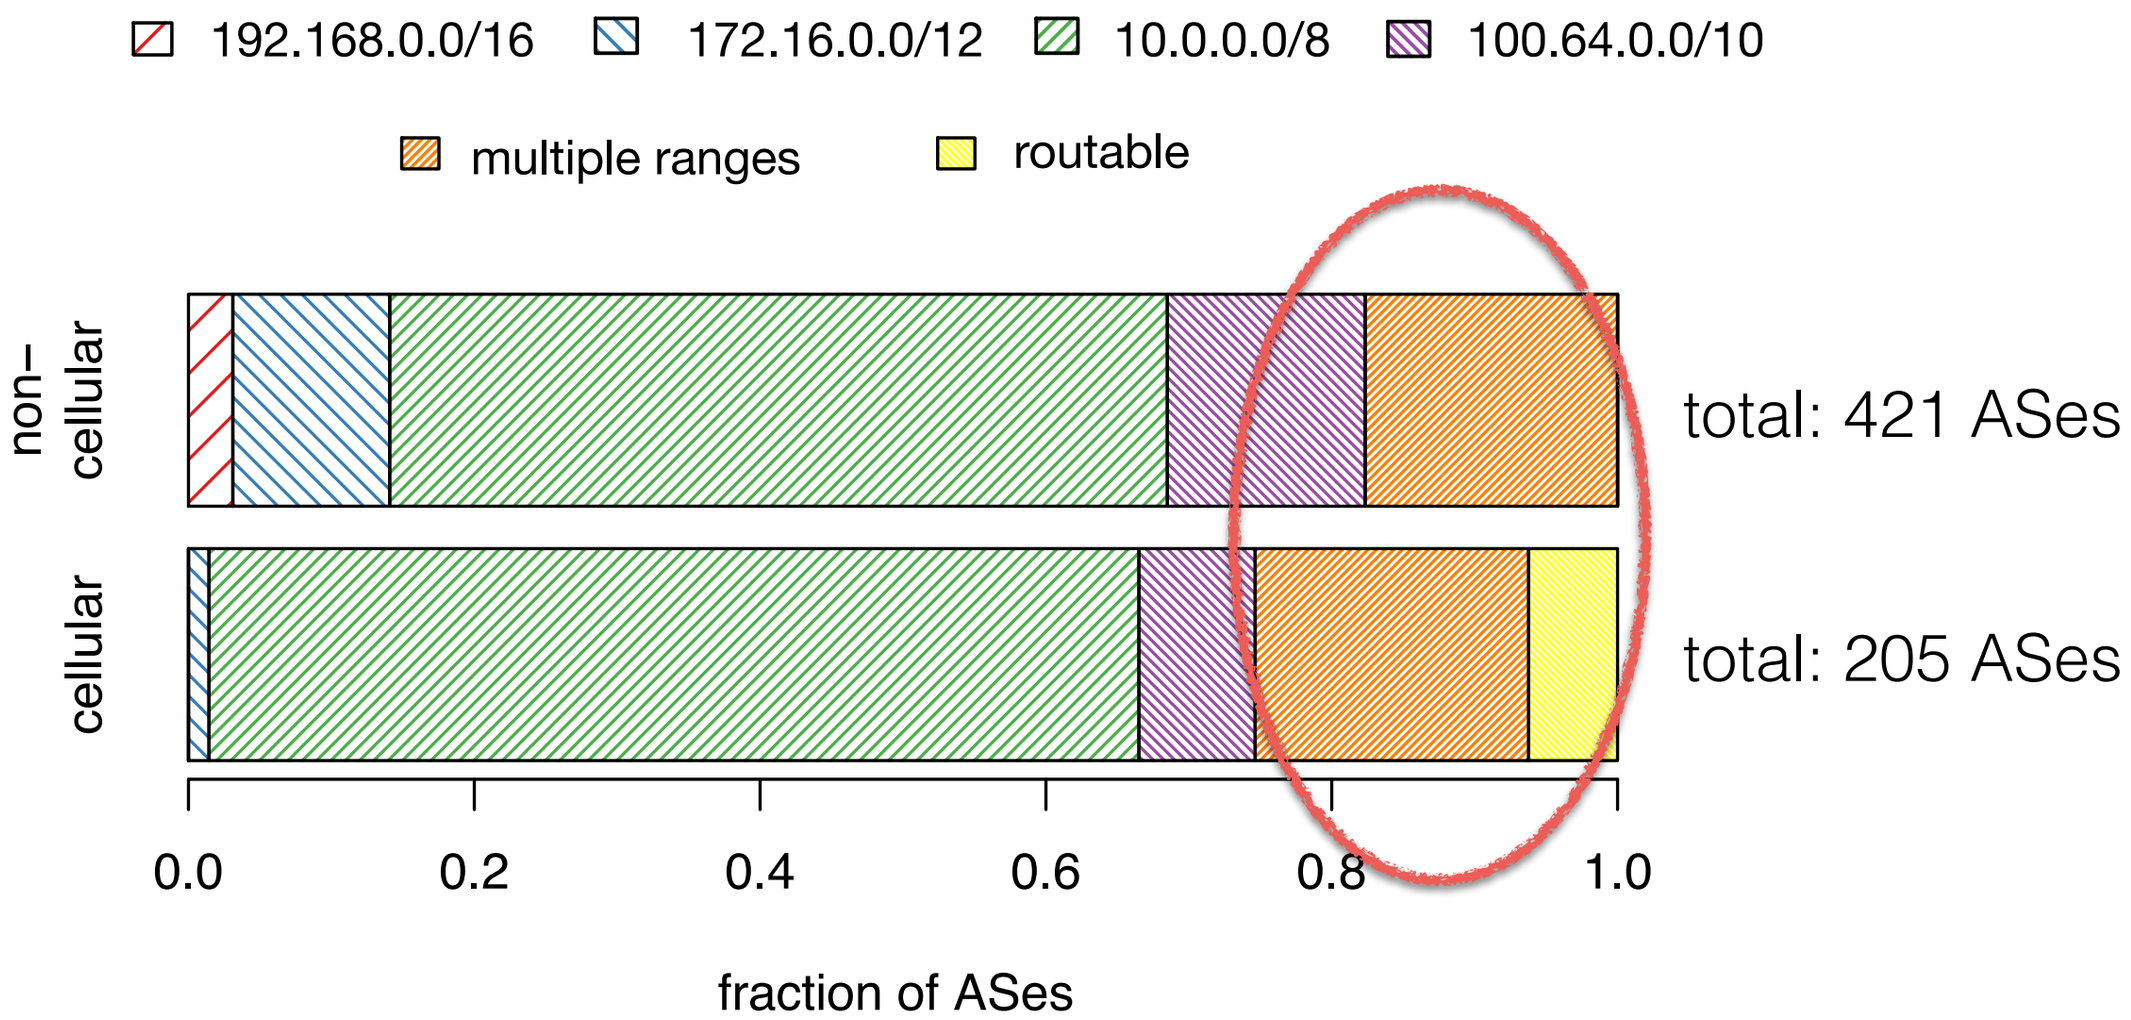

More than 20% of the ASes use multiple internal ranges. Shortage of Internal Address Space?

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?

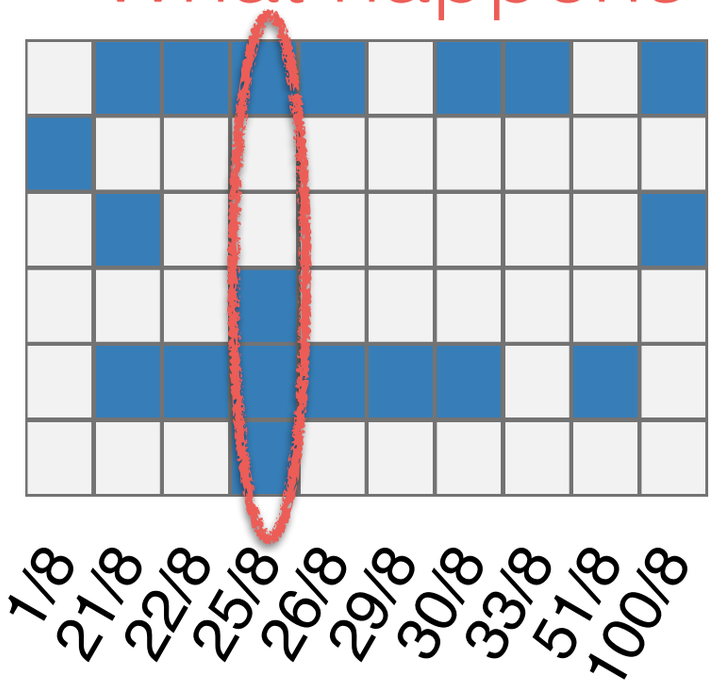

AS21928 (T–Mobile US)
AS24608 (H3G SpA IT)
AS22140 (T–Mobile US)
AS812 (Rogers Cable CA)
AS3651 (Sprint US)
AS852 (TELUS CA)

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

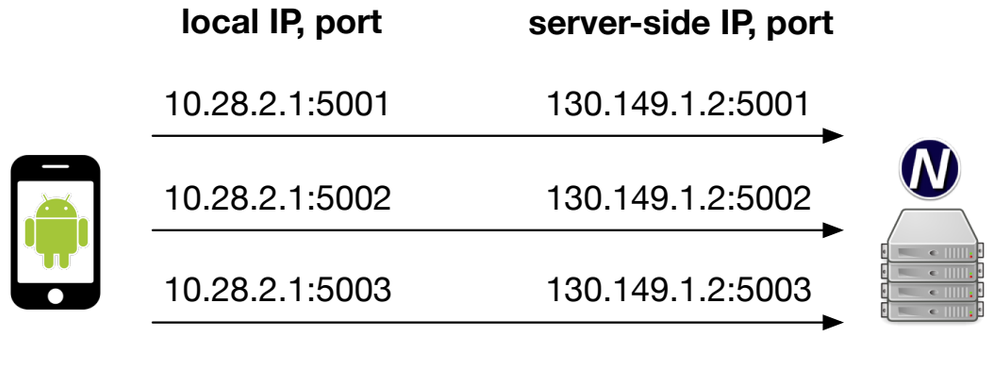

NAT test using TTL-limited probe packets

→ pinpoint the CGN location

→ extract CGN timeout values

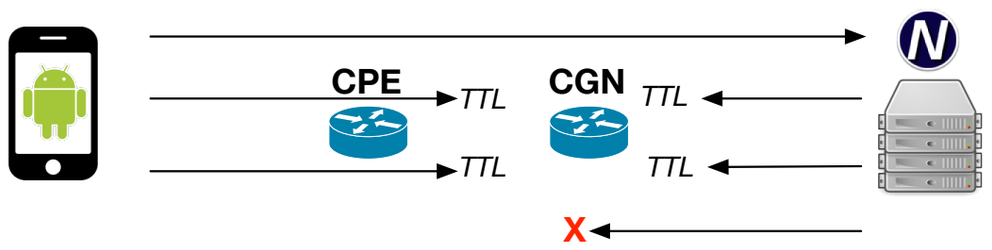

STUN test

→ reason about CGN mapping types

→ compare CGN and CPE mappings

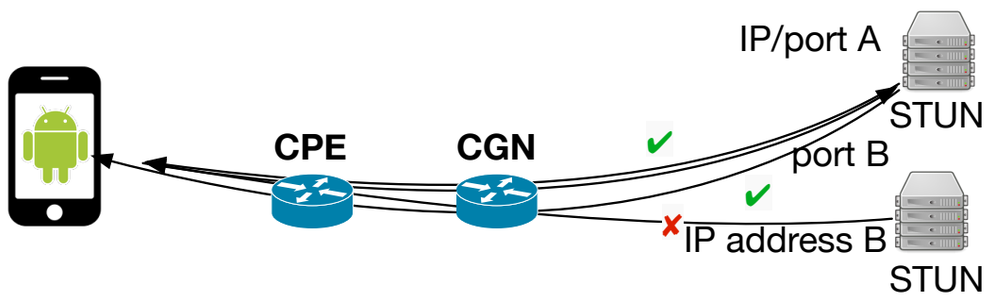

CGN Properties

High-Level Overview

- 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 of some CGNs more restrictive compared to CPEs

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

Summary

- Methods to detect CGN deployment
- Methods to extract properties from CGNs

- More than 500 CGN instances detected and analyzed
- CGN deployment rate
  - >= 17% non-cellular
  - 94% for cellular

CGN Considerations

CGNs are popular

- Consideration when developing applications
- IP address reputation systems, geolocation systems

CGNs are different

- Degree of resource sharing varies heavily across CGNs
- Directly reduce “how much Internet” a subscriber receives

CGNs still poorly understood

- What is an “acceptable” degree of resource sharing?
  - ➡ Measurements needed
  - ➡ Input for best practices for CGN dimensioning, regulations