Content distribution in next generation cellular networks

Fabián E. Bustamante

Northwestern U.

An ever-larger and more diverse Internet

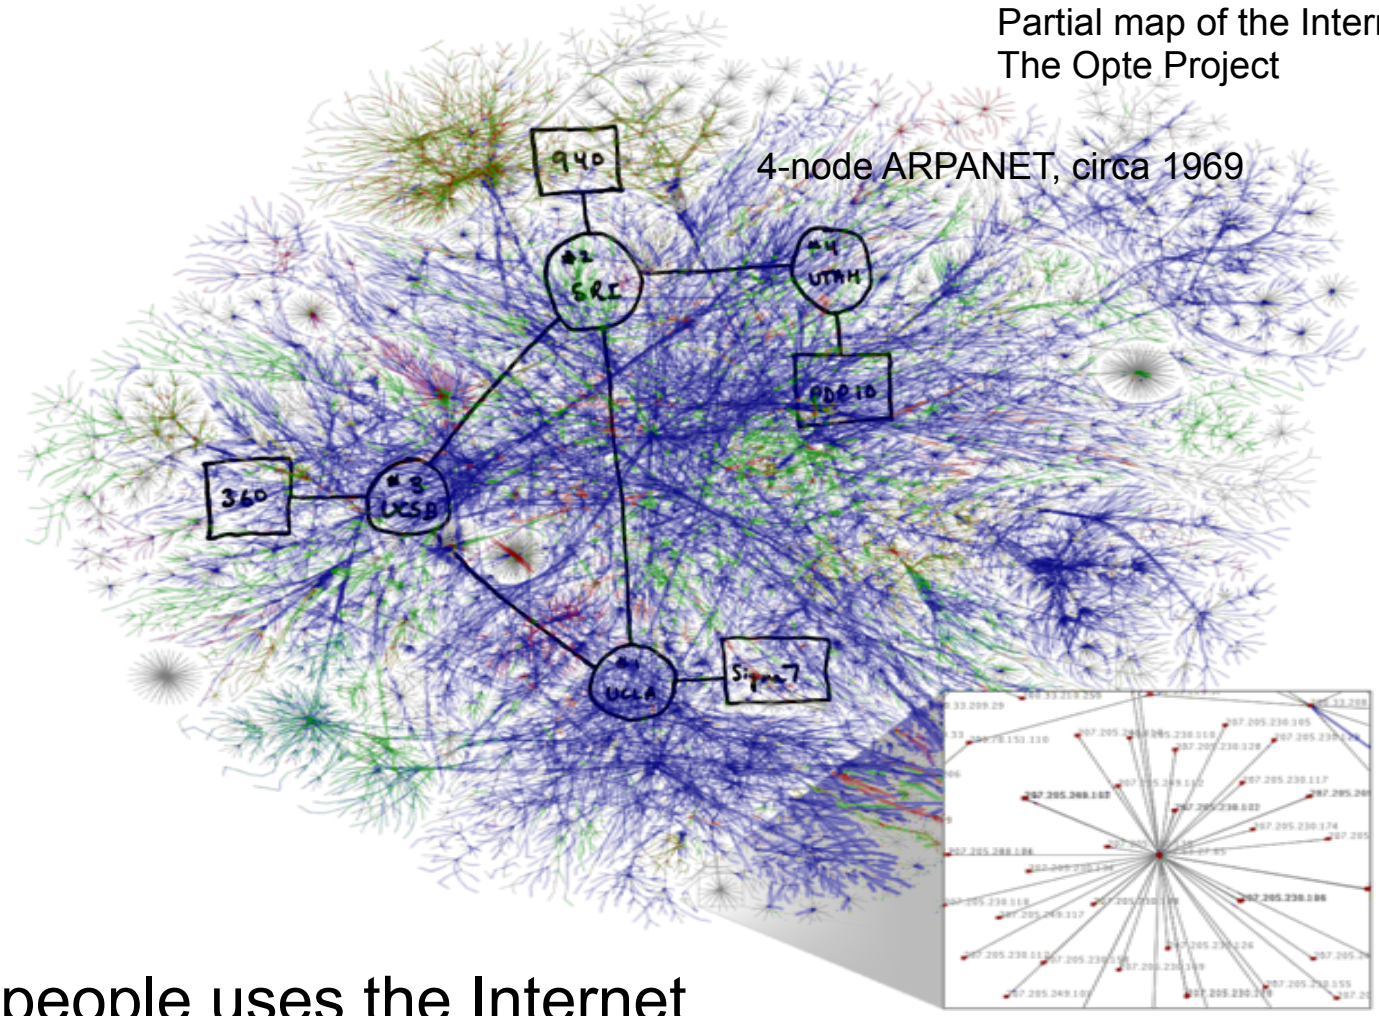

Partial map of the Internet, 2005
The Opte Project

4-node ARPANET, circa 1969

3 billion people uses the Internet
And that's only ~42% of World's population

Internet-scale distributed systems

1.5 billion in Facebook
307 million active users in Twitter
>70 million Skype users online

Challenging our understanding

- *How to build and test large-scale systems*
- *What the underlying network looks like*
- *How systems and their networks interact*
- *How users interact with networked systems*
- *The most appropriate research model*
- *…*

Exploring it all with a focus on the network edge

A sample of projects

- P2P systems and cross-ISP traffic
  - A BitTorrent extension with >1.4 million users

SIGCOMM 08

Ono

- Crowdsourcing event detection
  - Networked service events and user experience

SIGCOMM 10

NEWS

- Content distribution and remote/public DNS
  - A bad pairing and a practical solution

IMC 12

NAMEHELP

- Experiments at the (Internet's) edge
  - Beyond the academic network

NSDI 13

DASU

Content distribution in next generation cellular networks

The value of a few good users at the network edge

The world's on a mobile

- Mobile devices growth on capabilities & numbers
  - Number of devices > world population in 2014
  - By 2019, 11.5 billion devices for ~7.6 billion people

- People spend more time accessing content on mobile device than on PCs

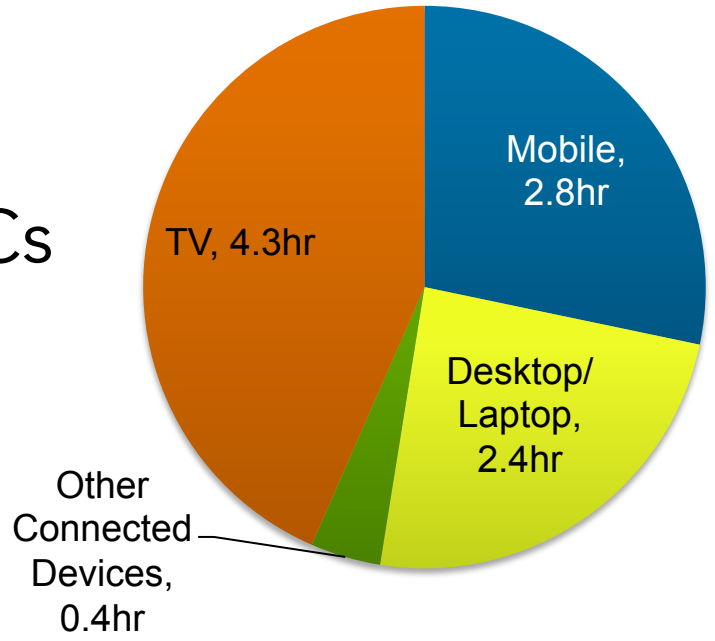

*Nielson Report 2014, +2014 Cisco VNI Report

… getting to content

What do people do on their mobile?

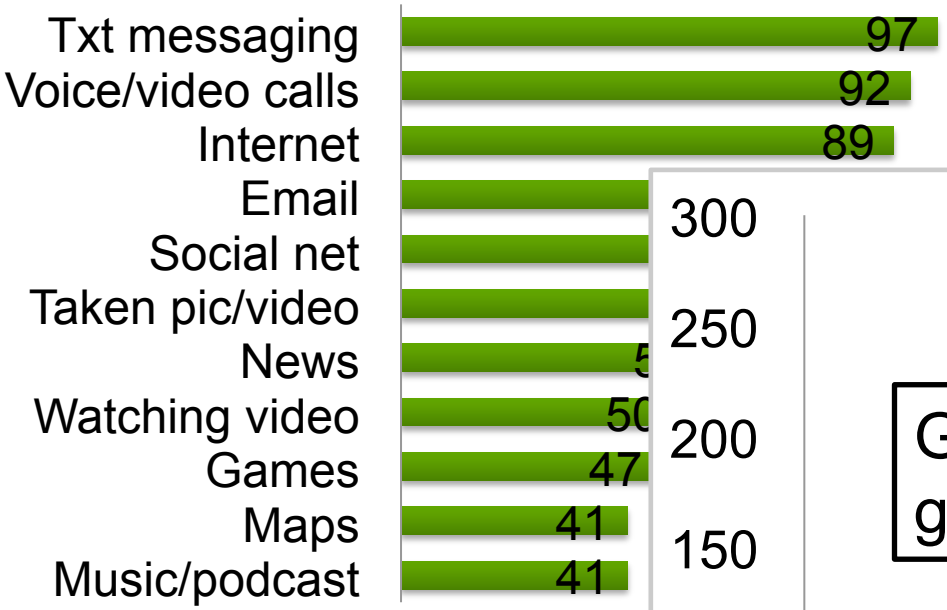

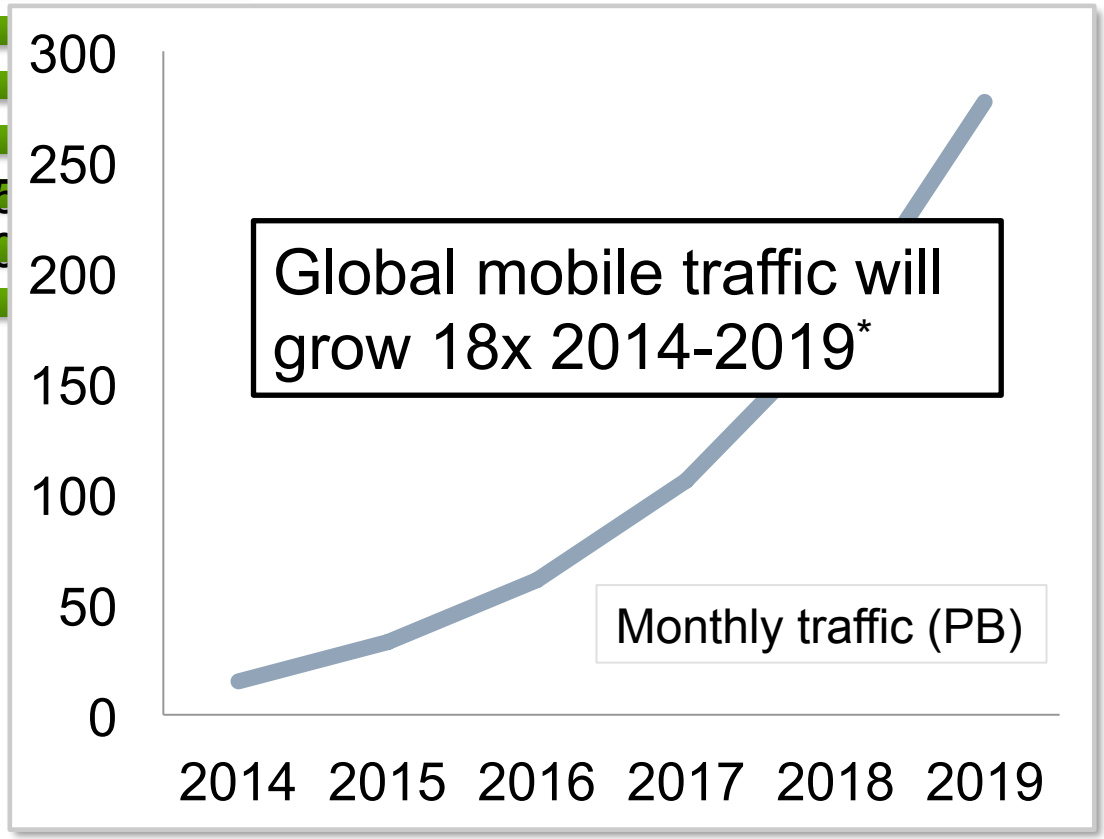

*PewResearch, US Smarthohone Use in 2015

*2015 Cisco VNI Report

Most of it delivered by CDNs

- CDNs replicate content in servers around the world
- Redirect clients to replicas *hopefully* nearby

4G network infrastructure

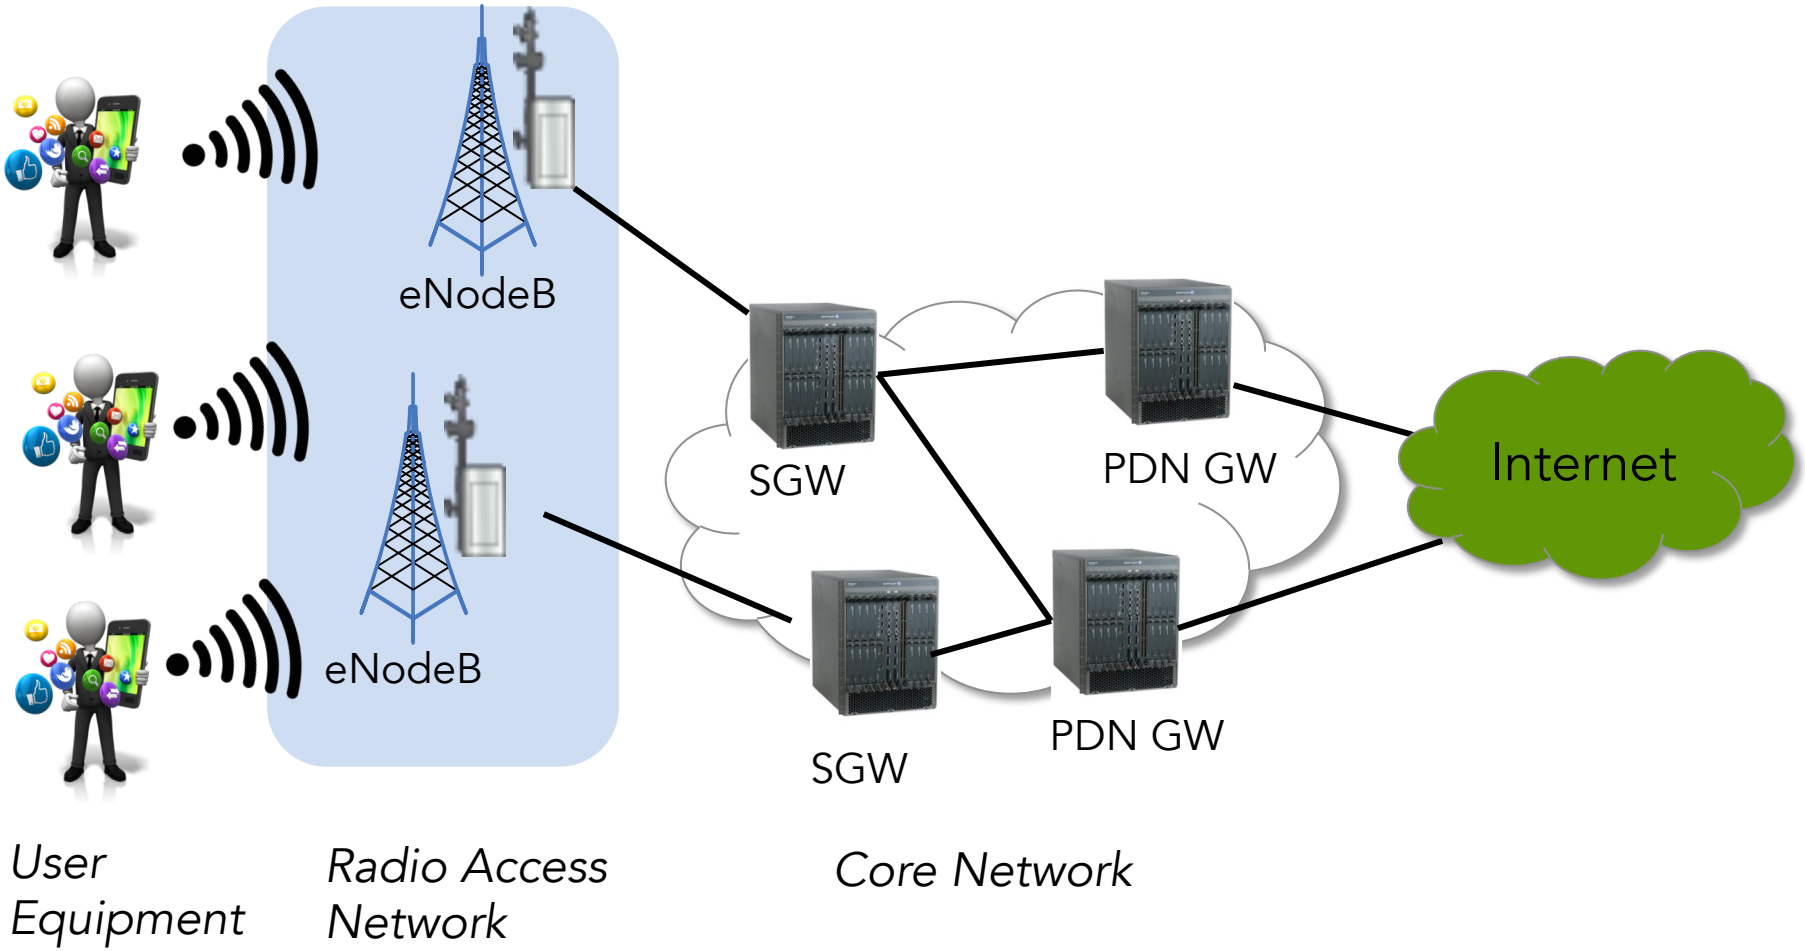

Accessing content

An approach 'inherited' from the wired world

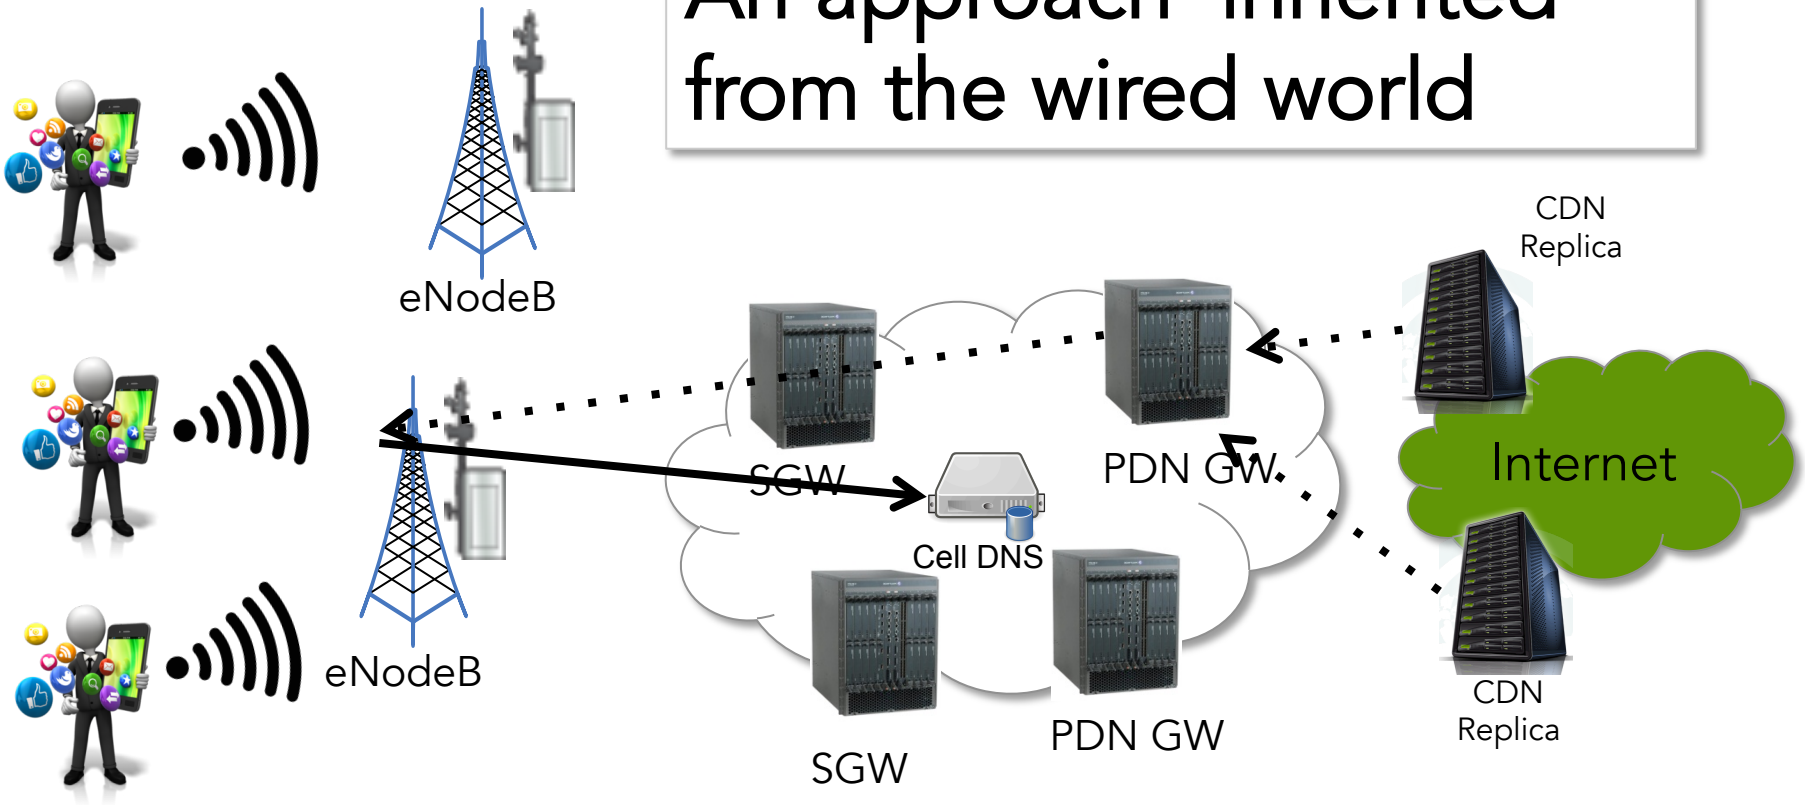

It didn't always work but nobody cared

Before

- 4-5 network egress points
- 100s milliseconds access latencies

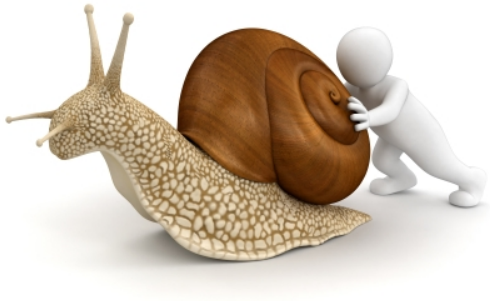

And then came next generation networks

Now

- 10-100s network egress points
- Higher throughput and low latencies ~10-20ms

Latency improvements

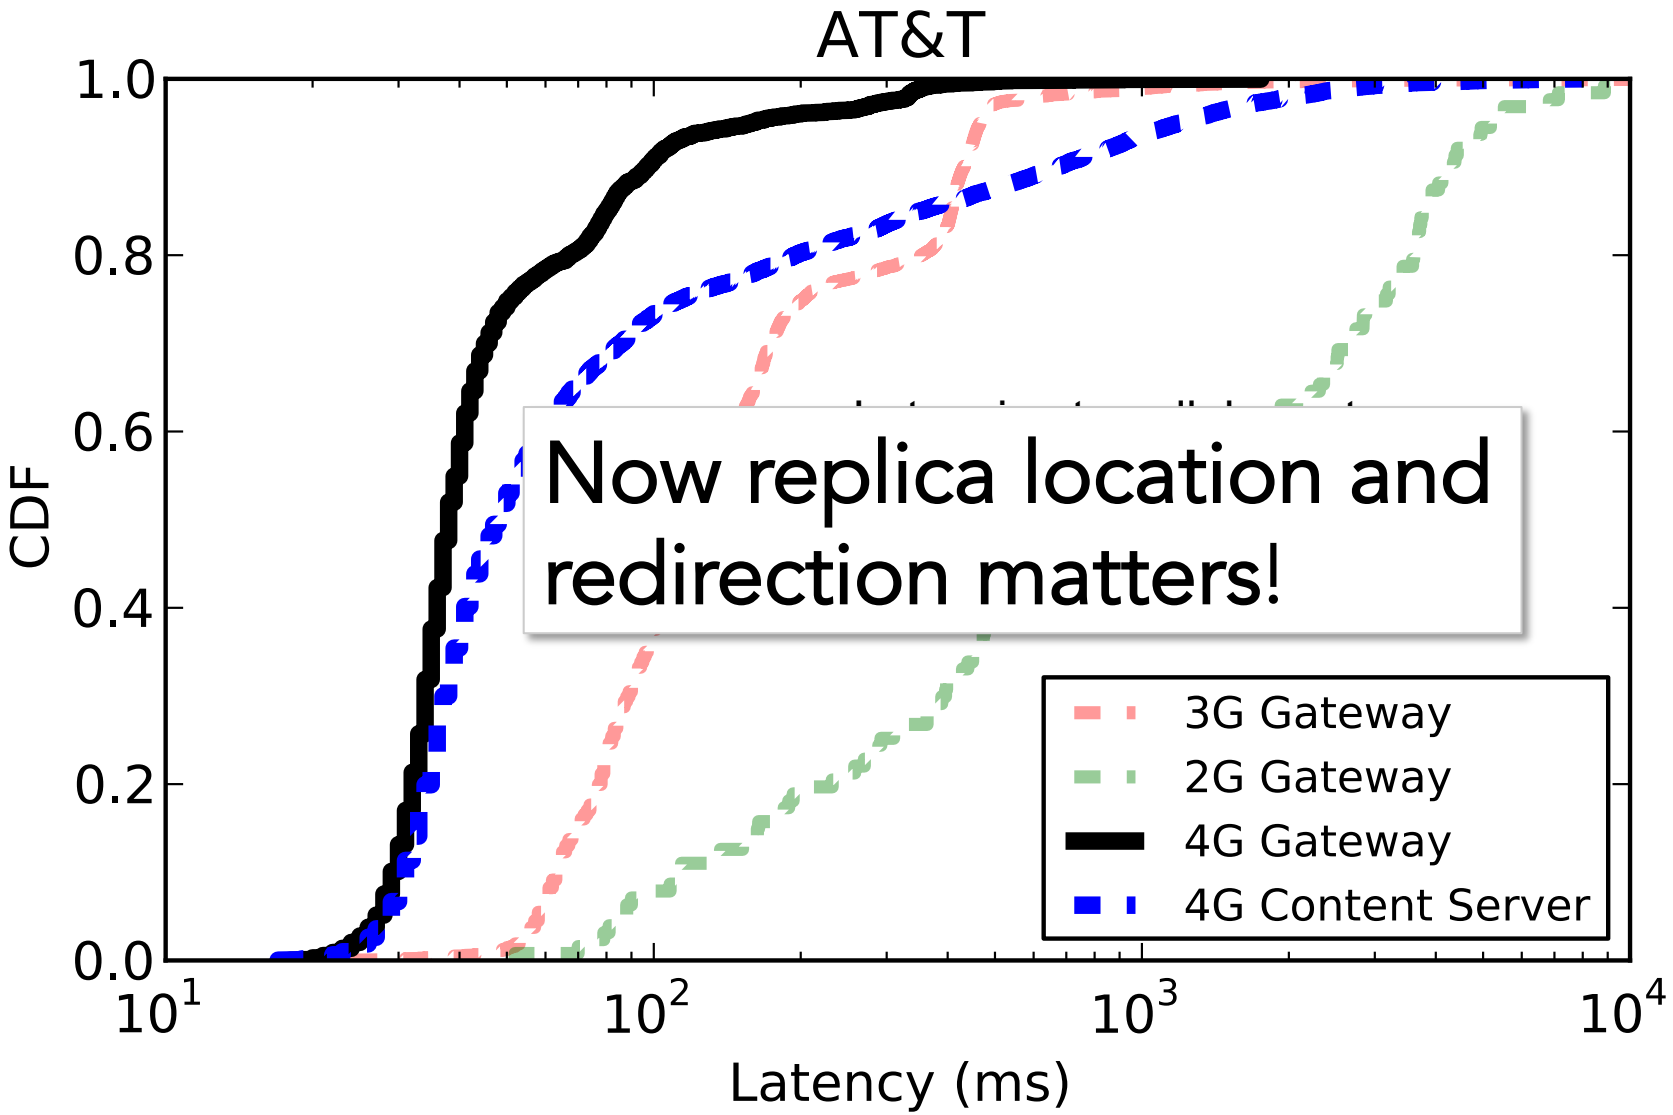

Crowd perspective

- Two apps and a shared experimentation library

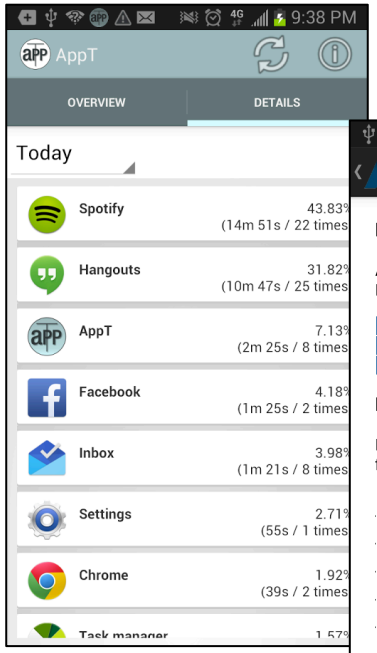

Application Time (AppT)

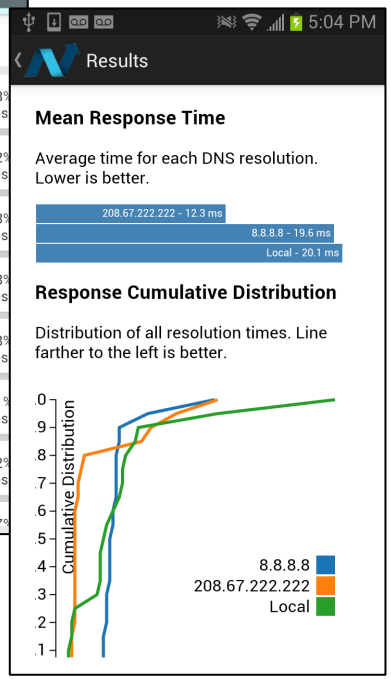

Namehelp Mobile

ALICE

A Lightweight Interface for Controlled Experiments

Current ALICE coverage

- 1216 unique users
- 98 countries
- >300 providers
- 997K experiment runs

Too many replicas – how to pick one

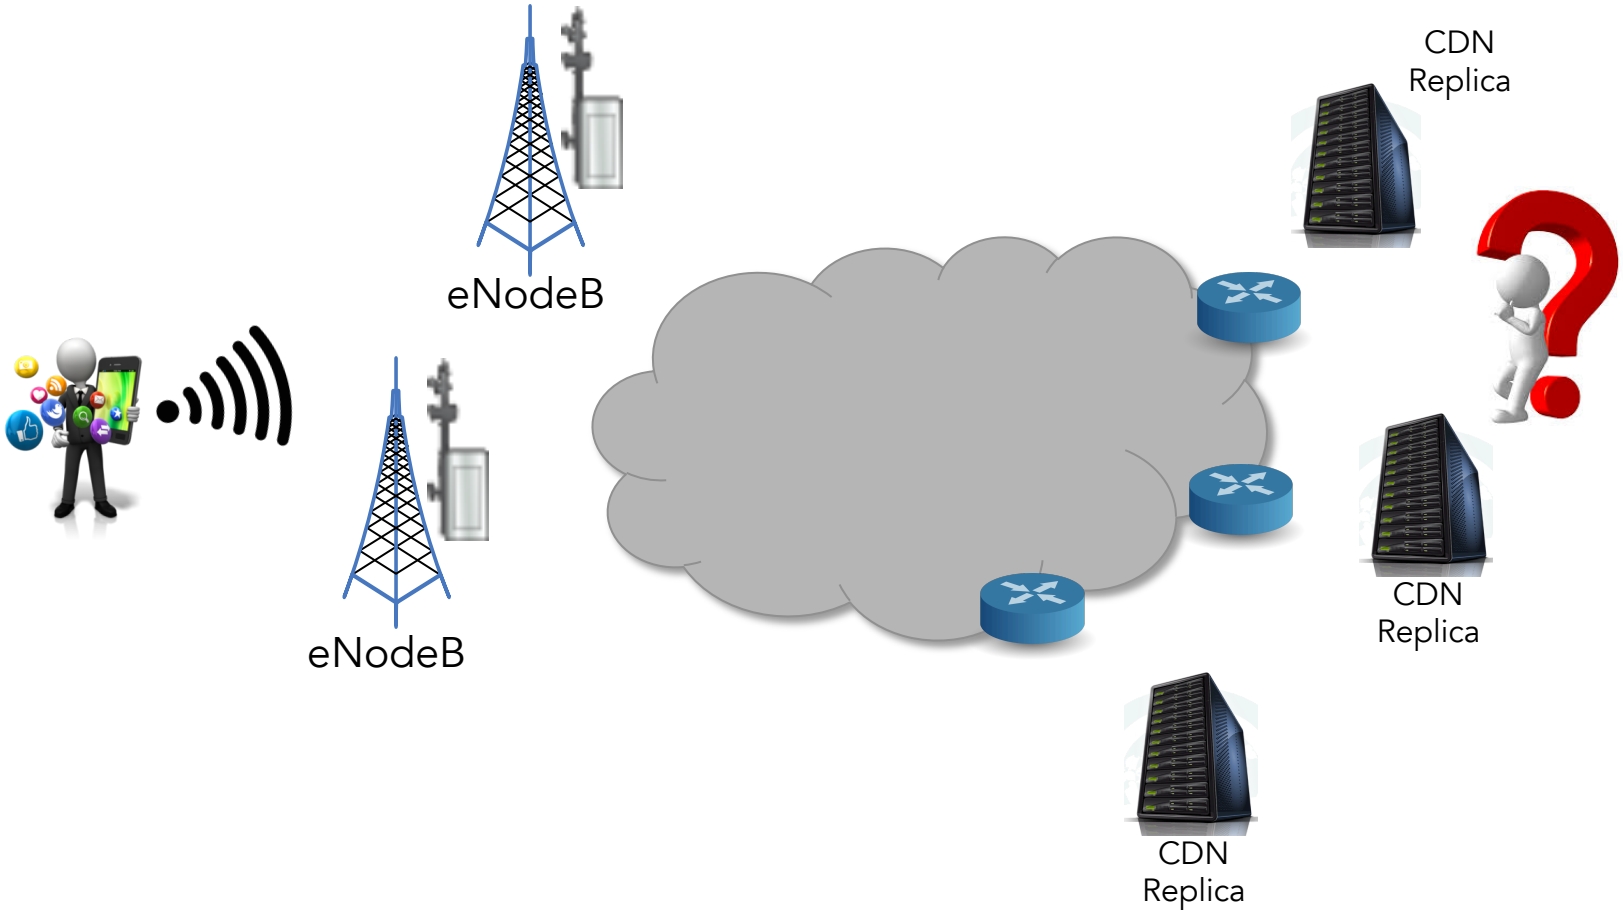

Local DNS as hint for replica selection?

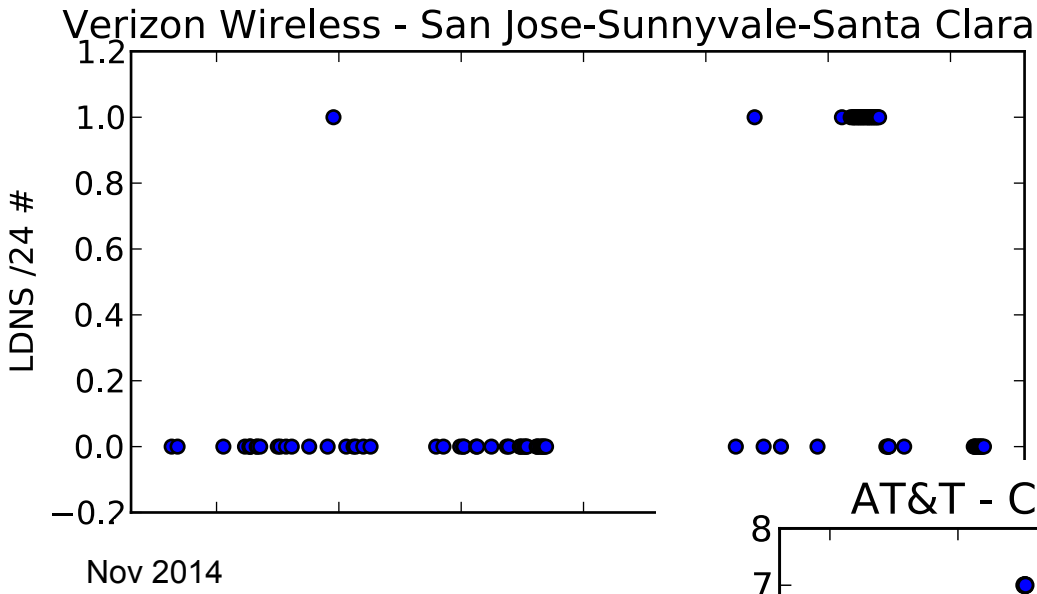

Visible LDNS /24 prefixes for a representative user

Fairly stable

Not quite!

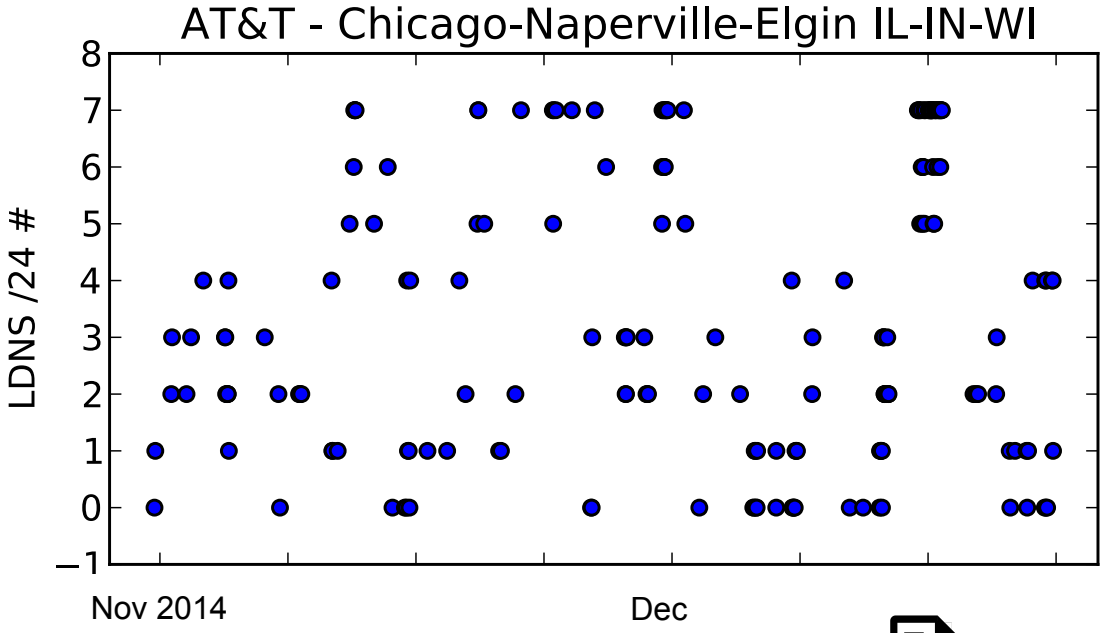

IMC 2014

The cost of inconsistency for CDNs

- CDNs map replica servers to resolver /24 prefix

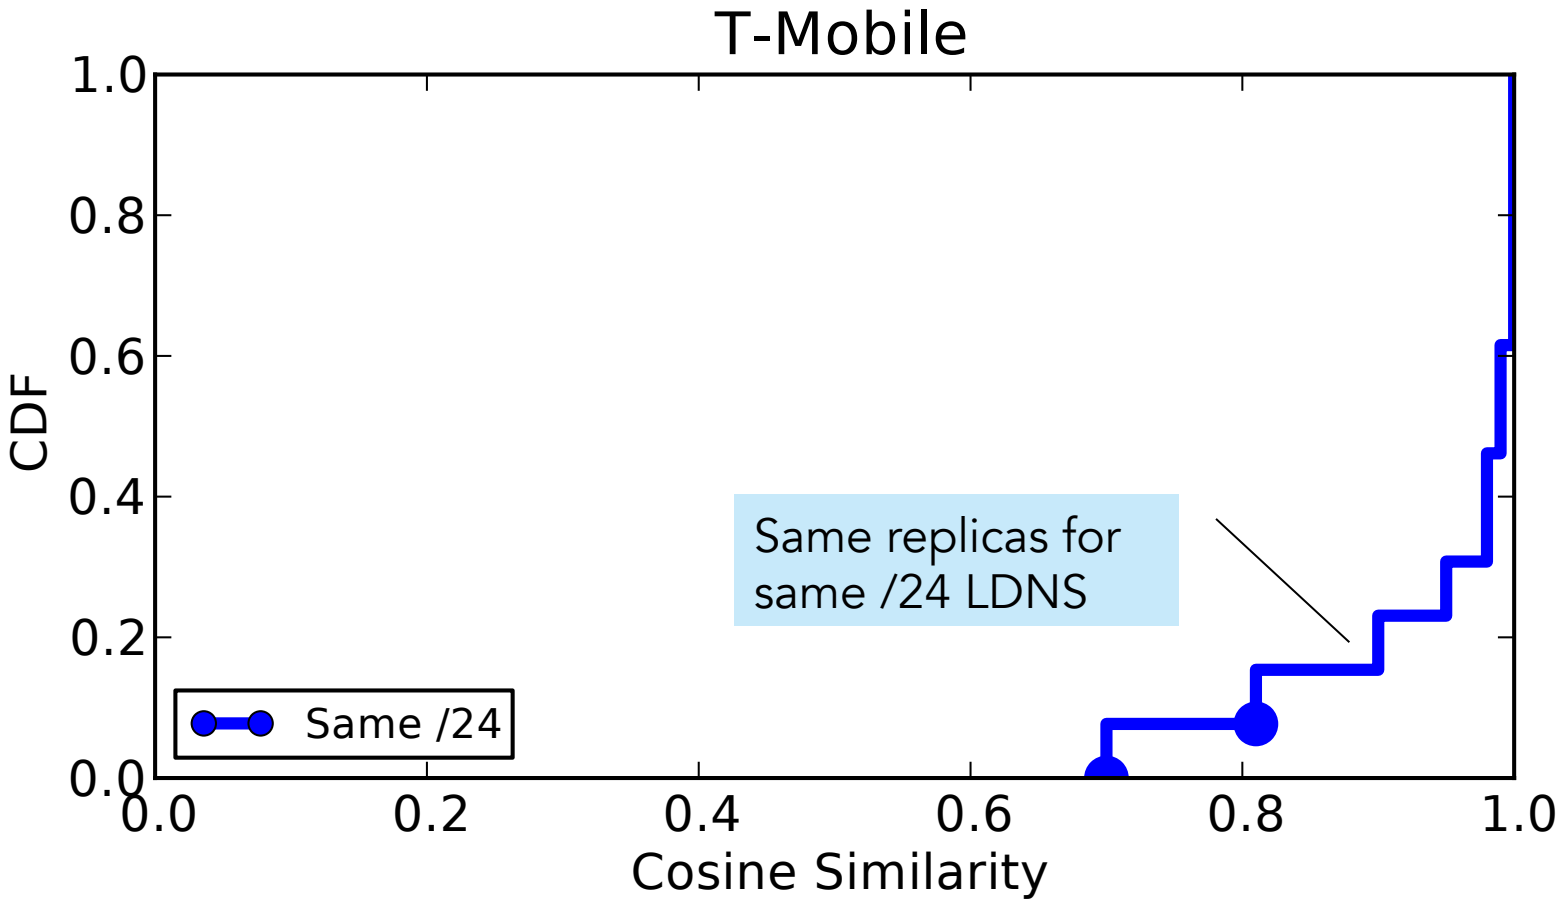

The cost of inconsistency for CDNs

- CDNs map replica servers to resolver /24 prefix

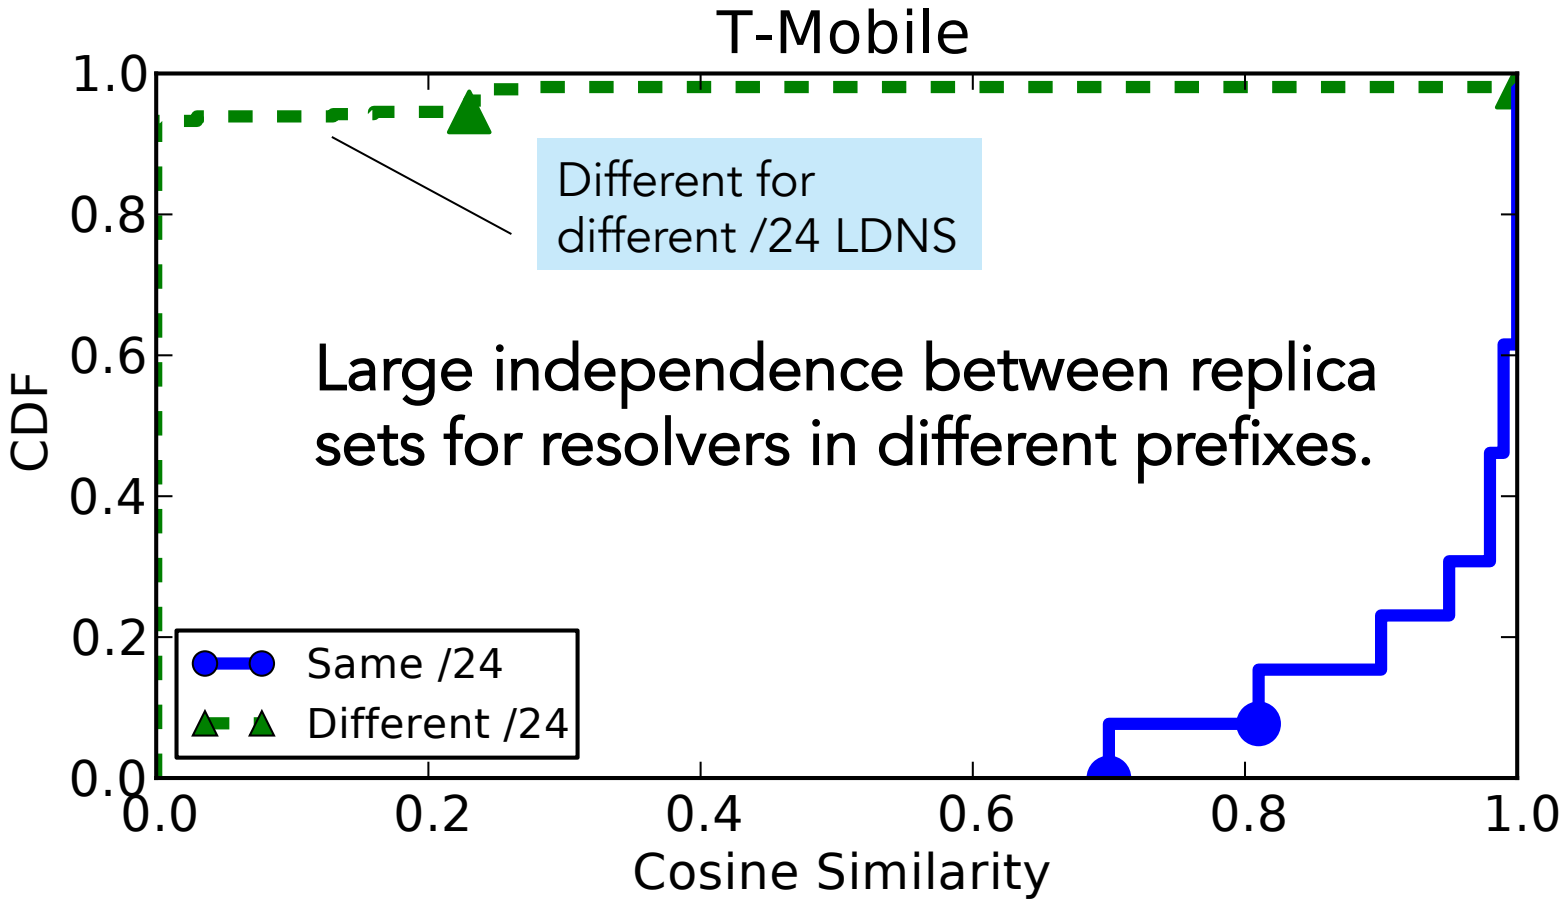

The impact of bad CDN replica selection

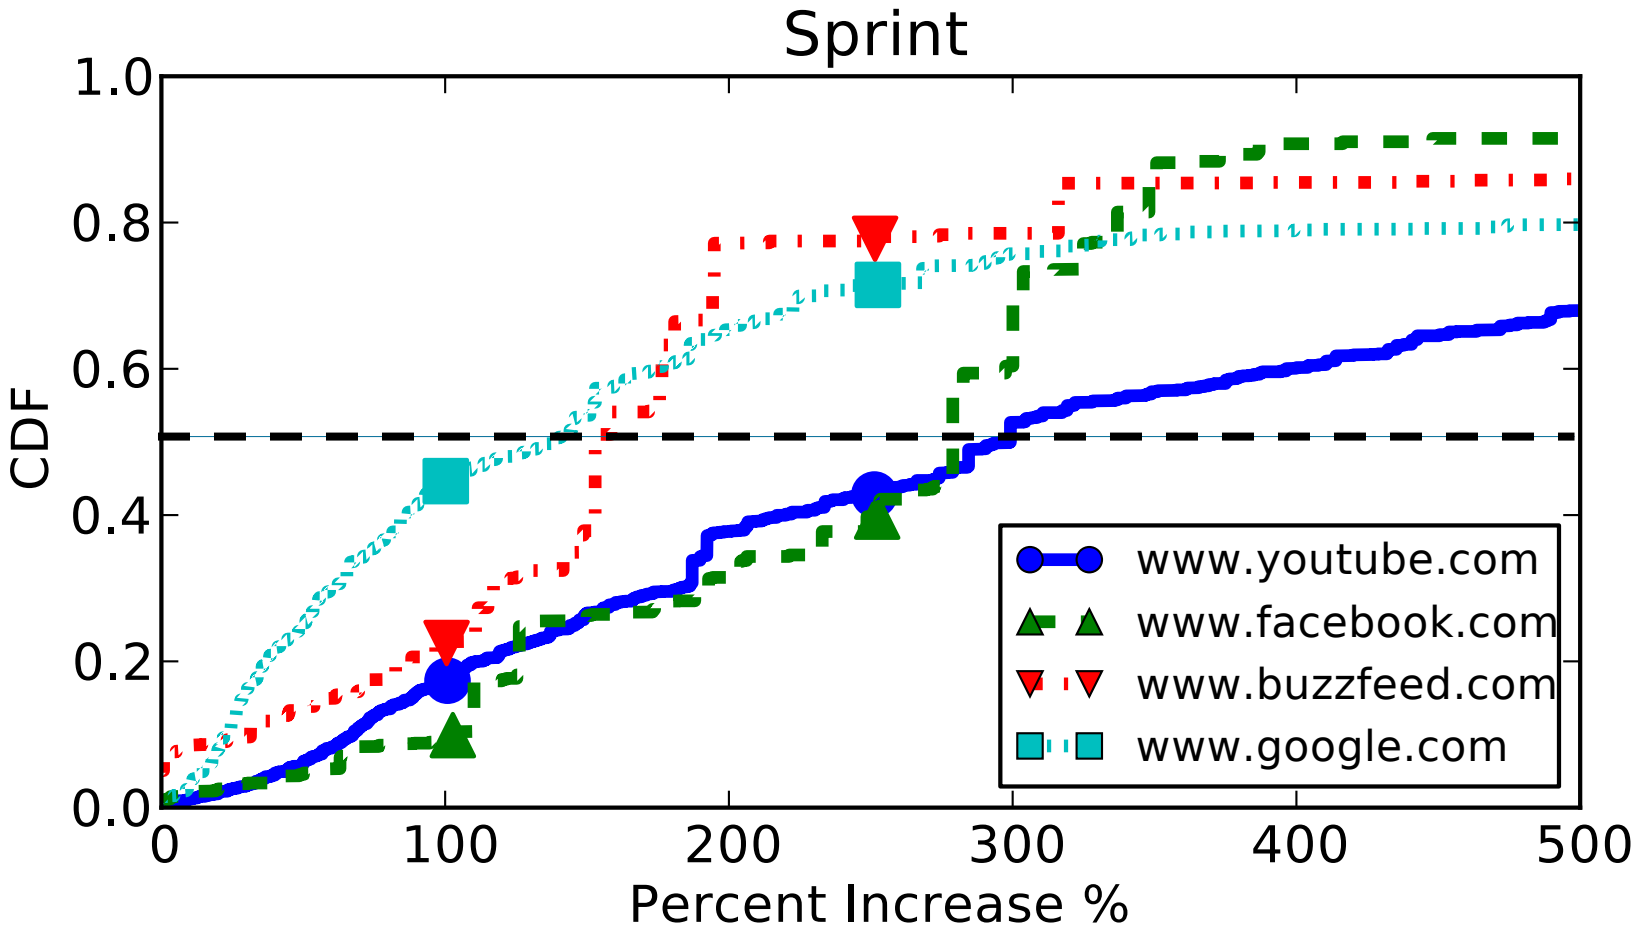

Client IP as hint?

- Enumerated visible client IP addresses for a client

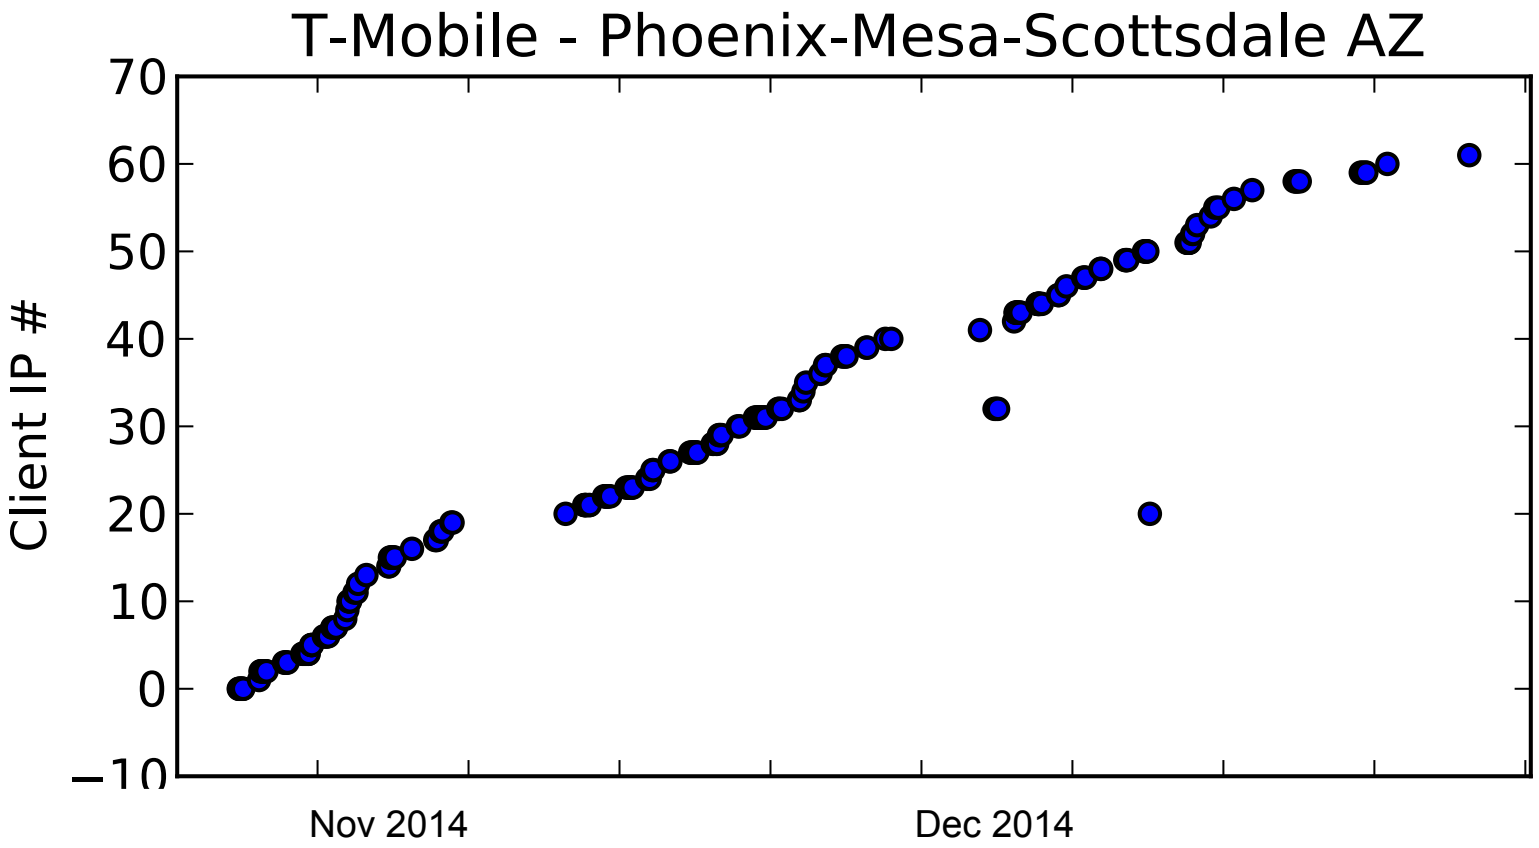

Anycast for routing to replicas?

- Anycast variability as seen from /24 prefixes of GoogleDNS resolvers

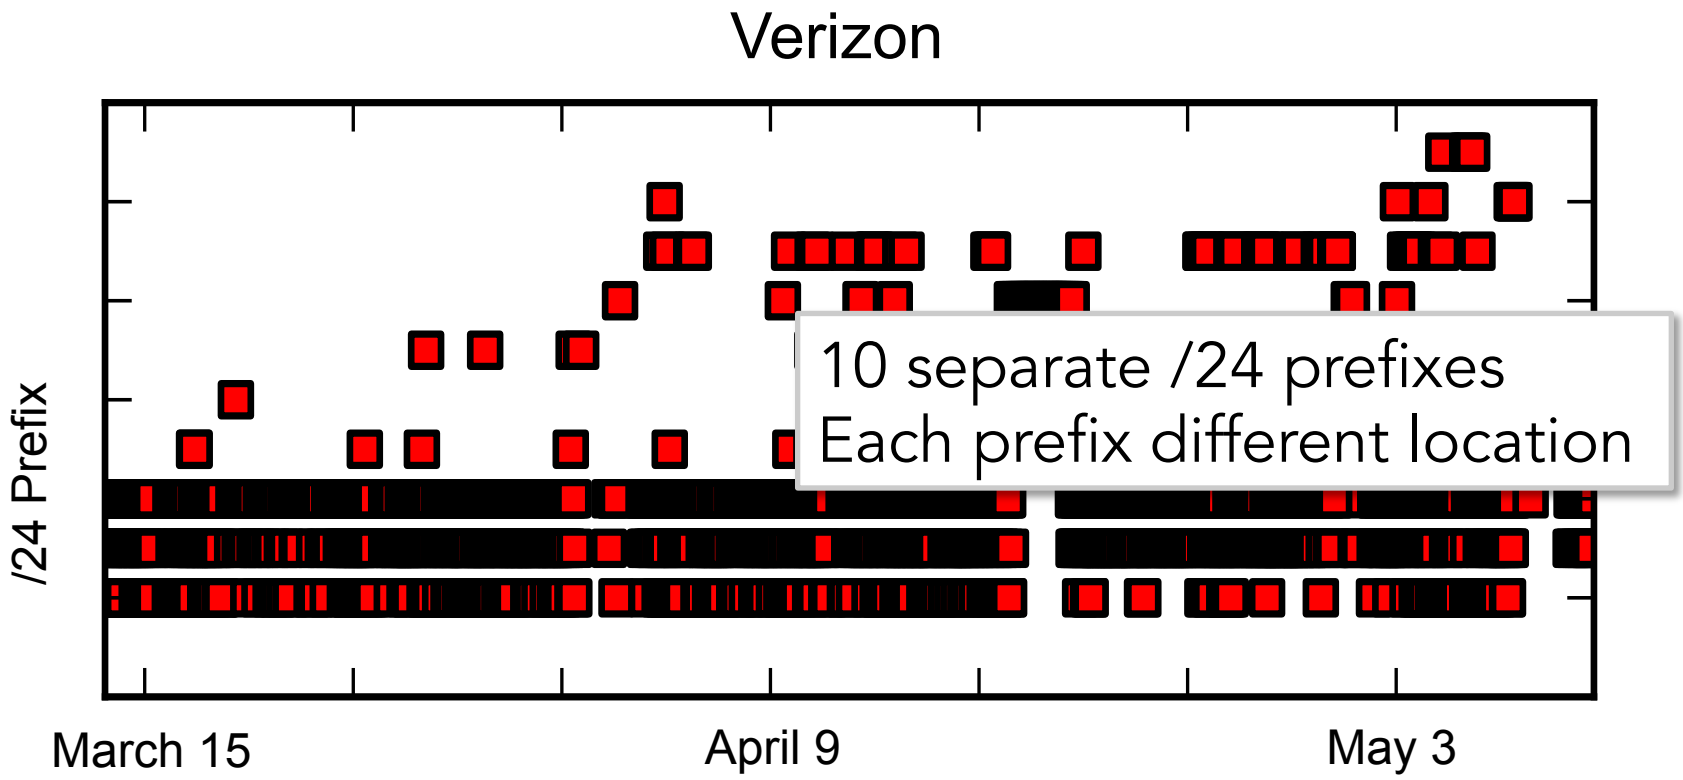

What *does* remain constant?

- In sum …
  - No Local DNS, no client IP, no anycast route …
- What *does* remain constant?!
  - Clients' gateways

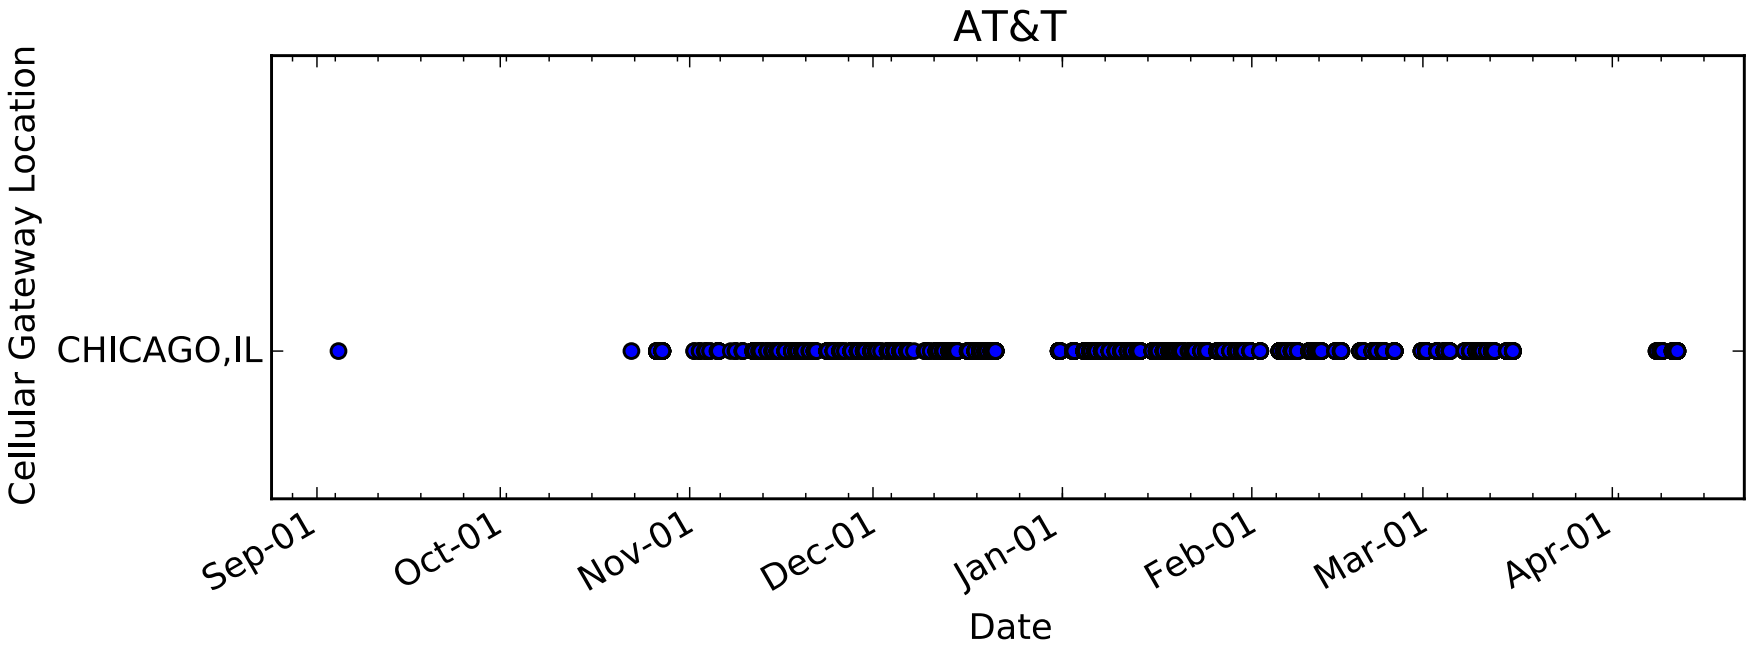

And determines your performance

- … since *all* traffic is routed through gateways
- Large factor of end-to-end performance is server distance to gateway

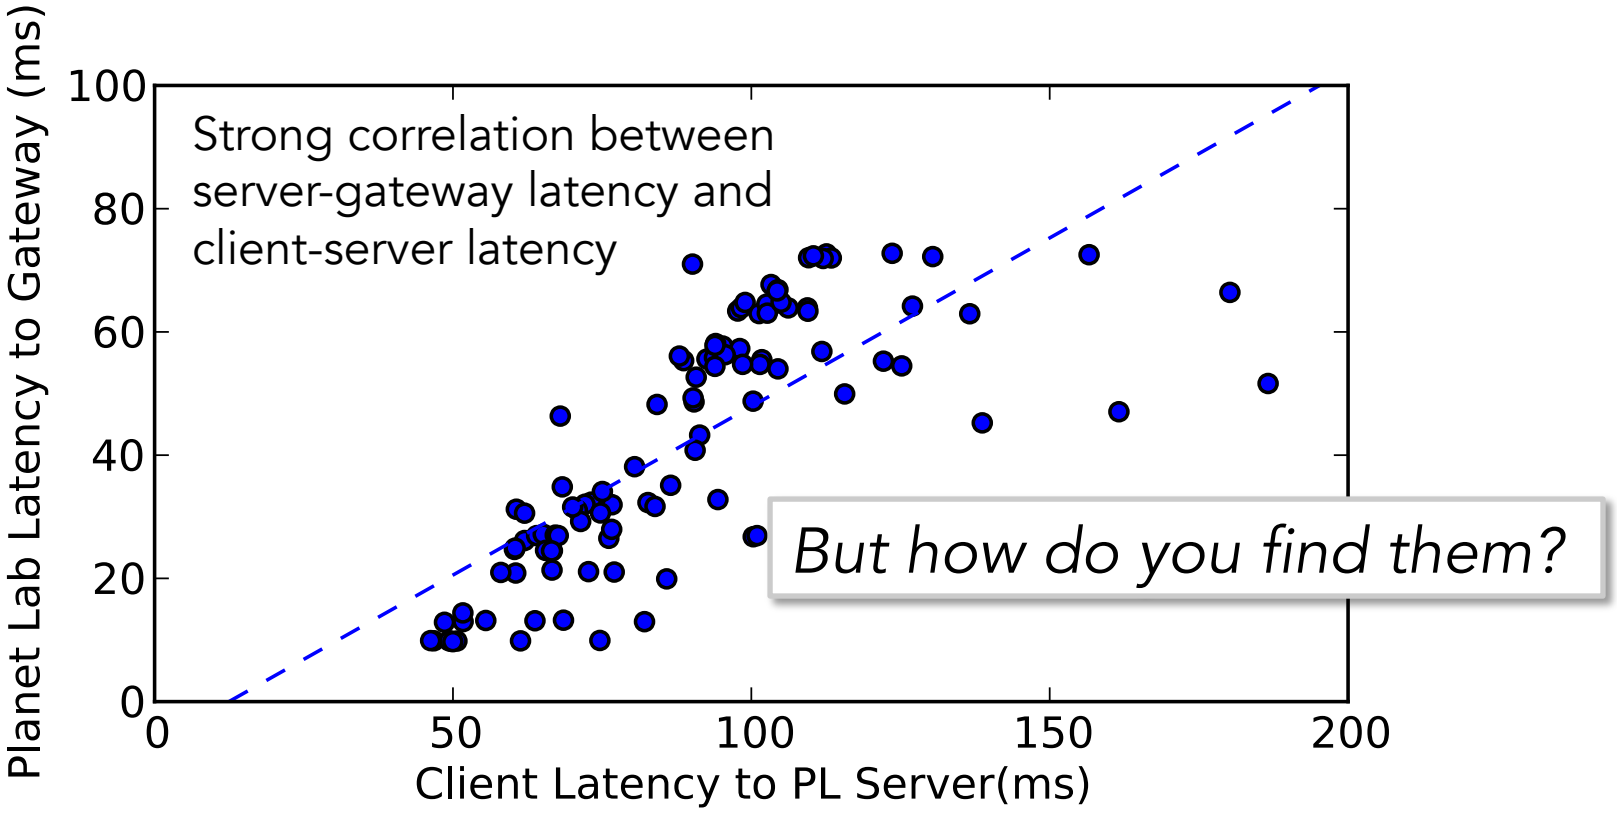

Can we infer clients' gateways?

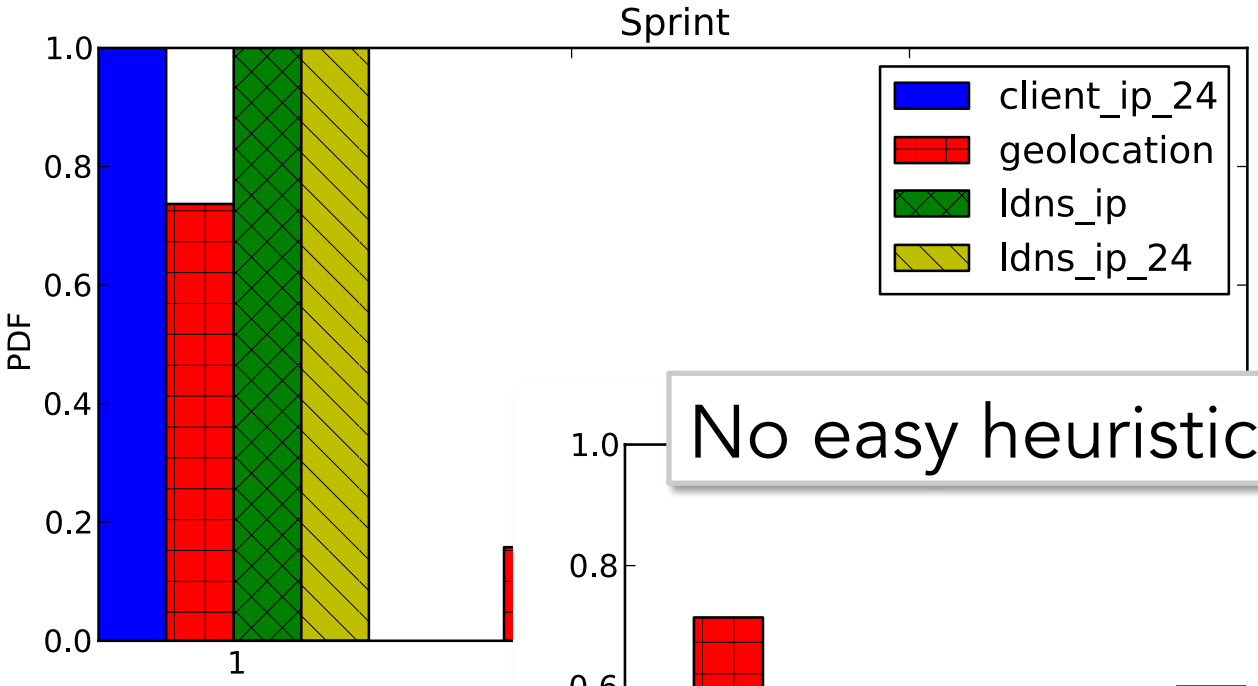

Nicely consistent for client IP /24 and LDNS

No easy heuristics!

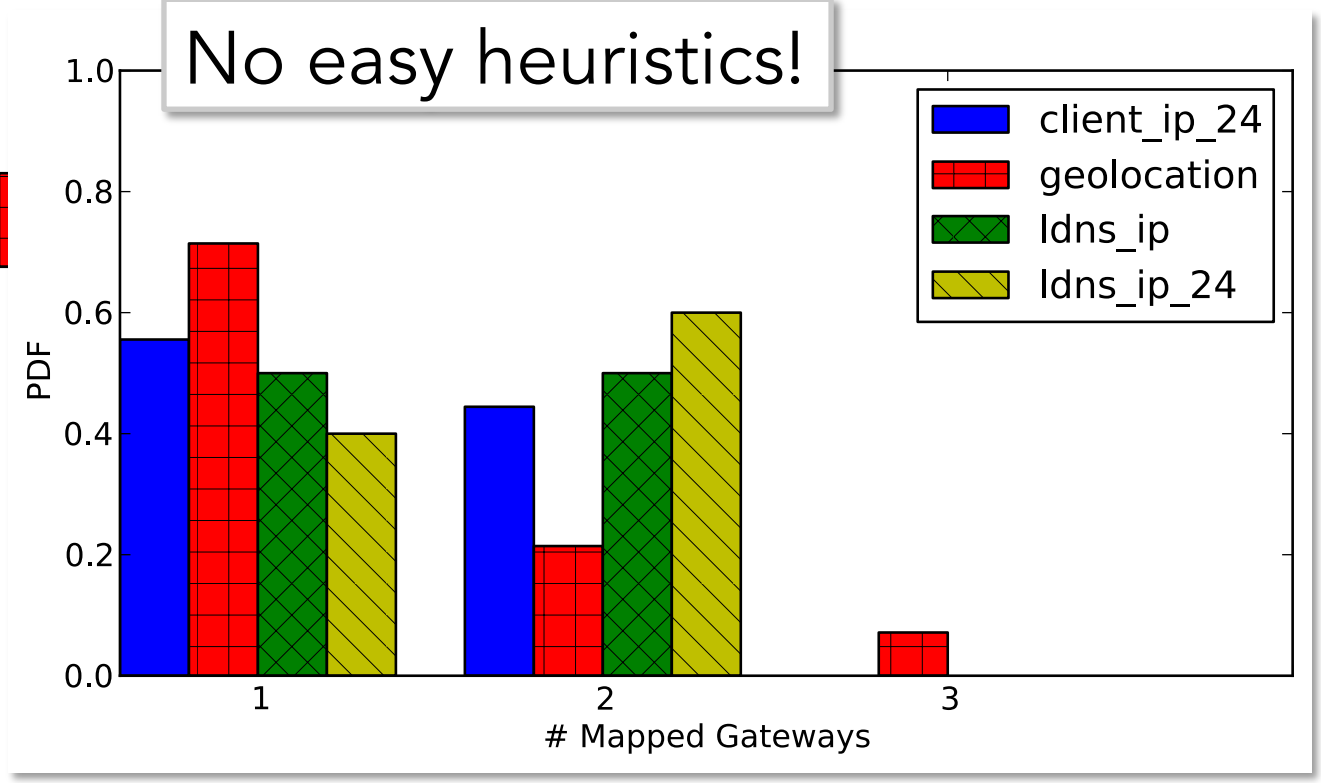

Not quite!

Clients' gateway mapping

- Client-gateway mapping given by client-IP prefix
- But size of the prefix depends on the provider and changes over time

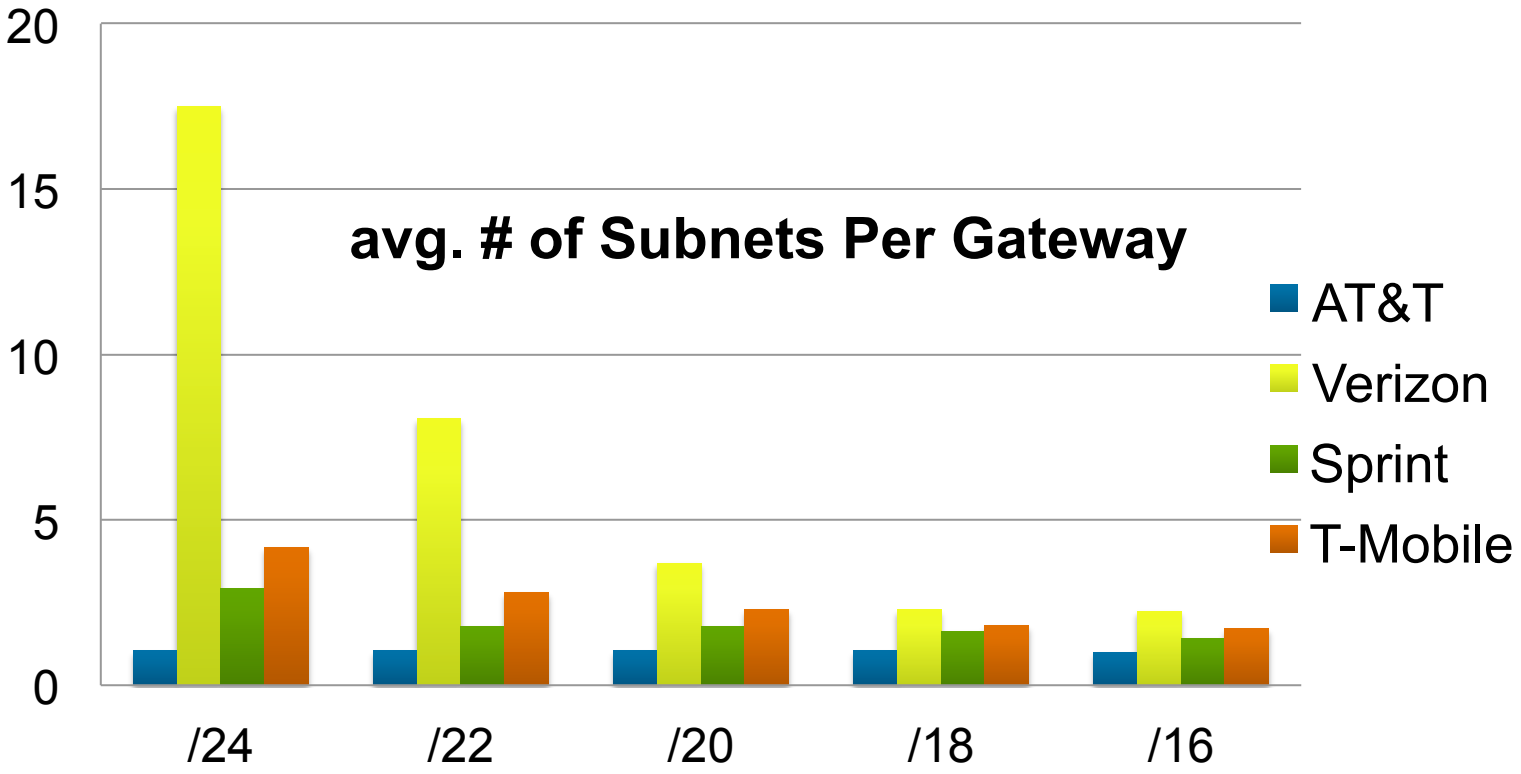

Getting the client involved

- CDN
  - Has coverage but limited visibility
  - Needs client view and contextual information
- Mobile clients
  - High visibility; see full end-to-end path
  - But limited coverage

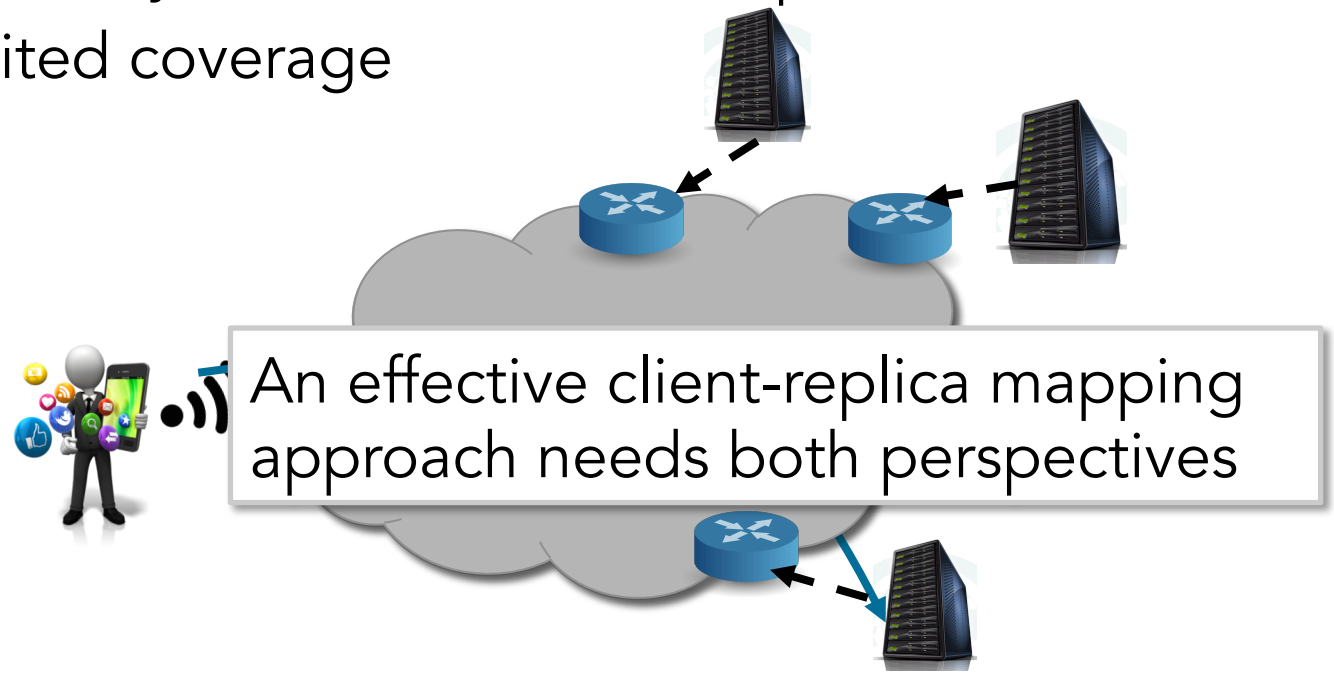

A client/CDN cooperation

- The task: for a network, find the prefix/gateway mapping
  - If pool size is known, one measurement per pool would suffice
  - But changes across providers and over time
- Naïve approach
  - Any/all clients probe and inform the CDN
- More scalable solution
  - Client probing directed by CDNs
  - Leveraging client IP dynamics

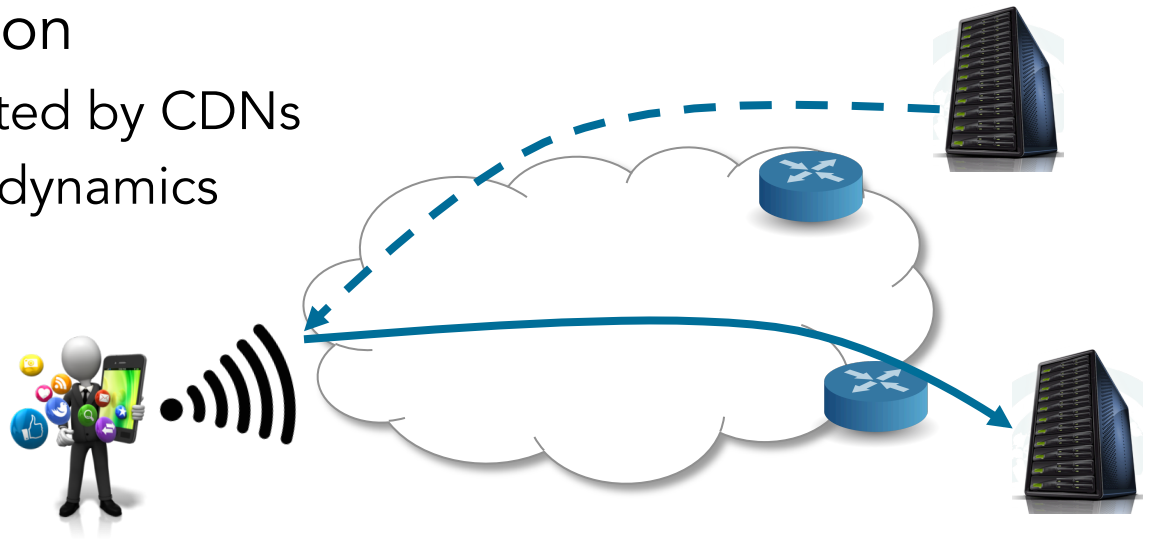

Exploit network dynamics

- Mobile clients assigned multiple public IP addresses
  - NATs (CGNs)
  - Network attachment (possible assignment to different pools)

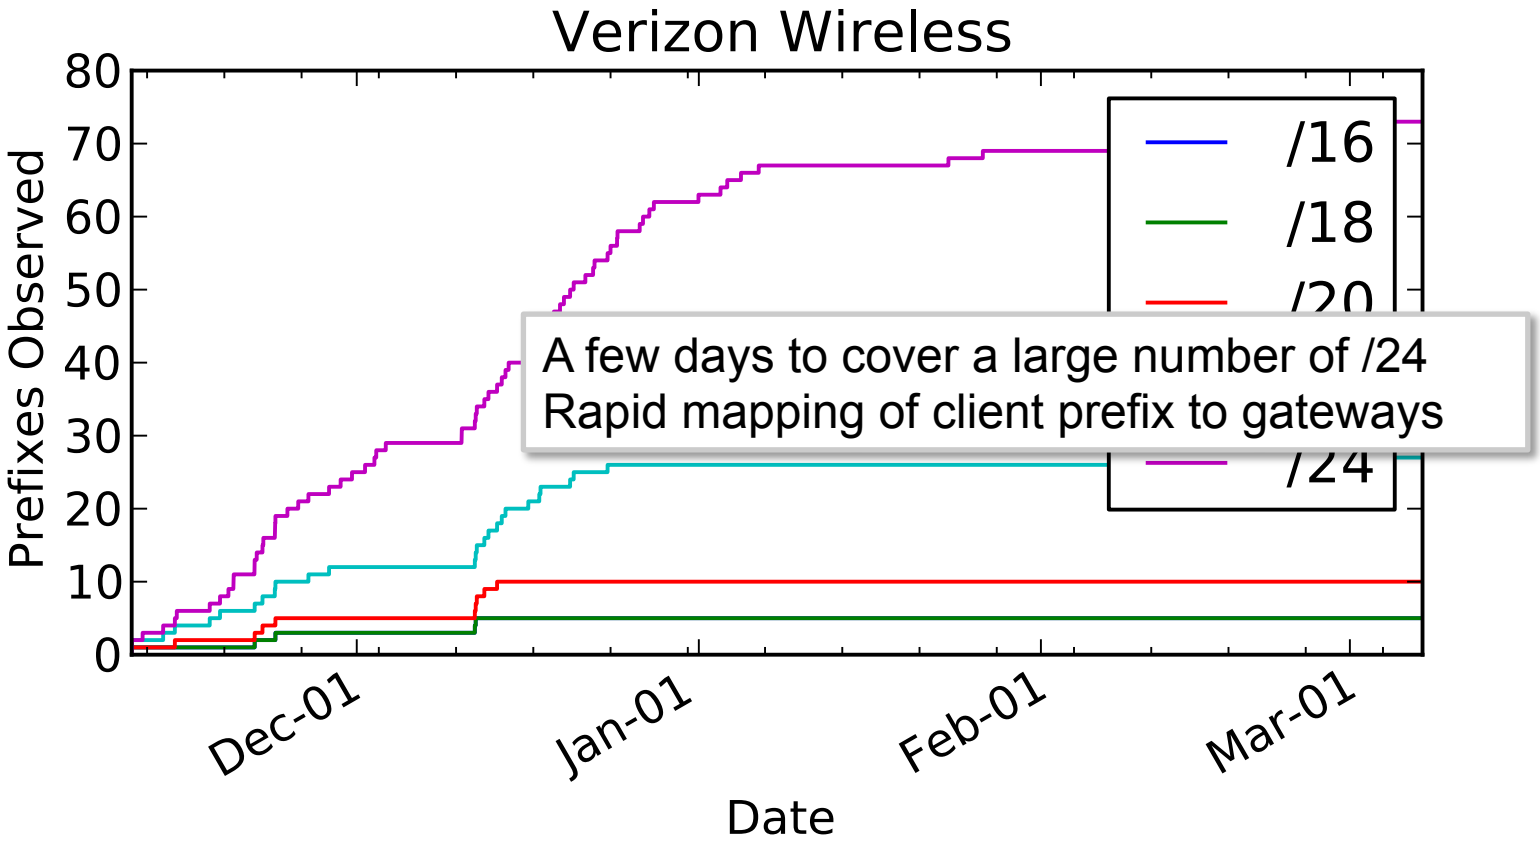

Working on the end-to-end evaluation

- Working with Akamai to evaluate this approach
  - Showing the value of few good users at the edge

- End user involvement not just
  - For measurement
  - Or experimentation
  - Or a crowdsourced solution
  - *But instrumental to an effective approach to client redirection*

Content distribution on next generation cellular networks

Work in collaboration with …
John Rula (*NU*), Mortiz Steiner and Ruomei Gao (*Akamai*), …

AquaLab