Core Internet Lessons from Today's Network Edge

Fabián E. Bustamante,
Associate Professor, Northwestern U.
Evanston/Chicago, Illinois USA

AquaLab

Today's future

- Several alternative directions to design the future Internet architecture

- *The design of a future Internet needs a better understanding of today's Internet ecosystem*

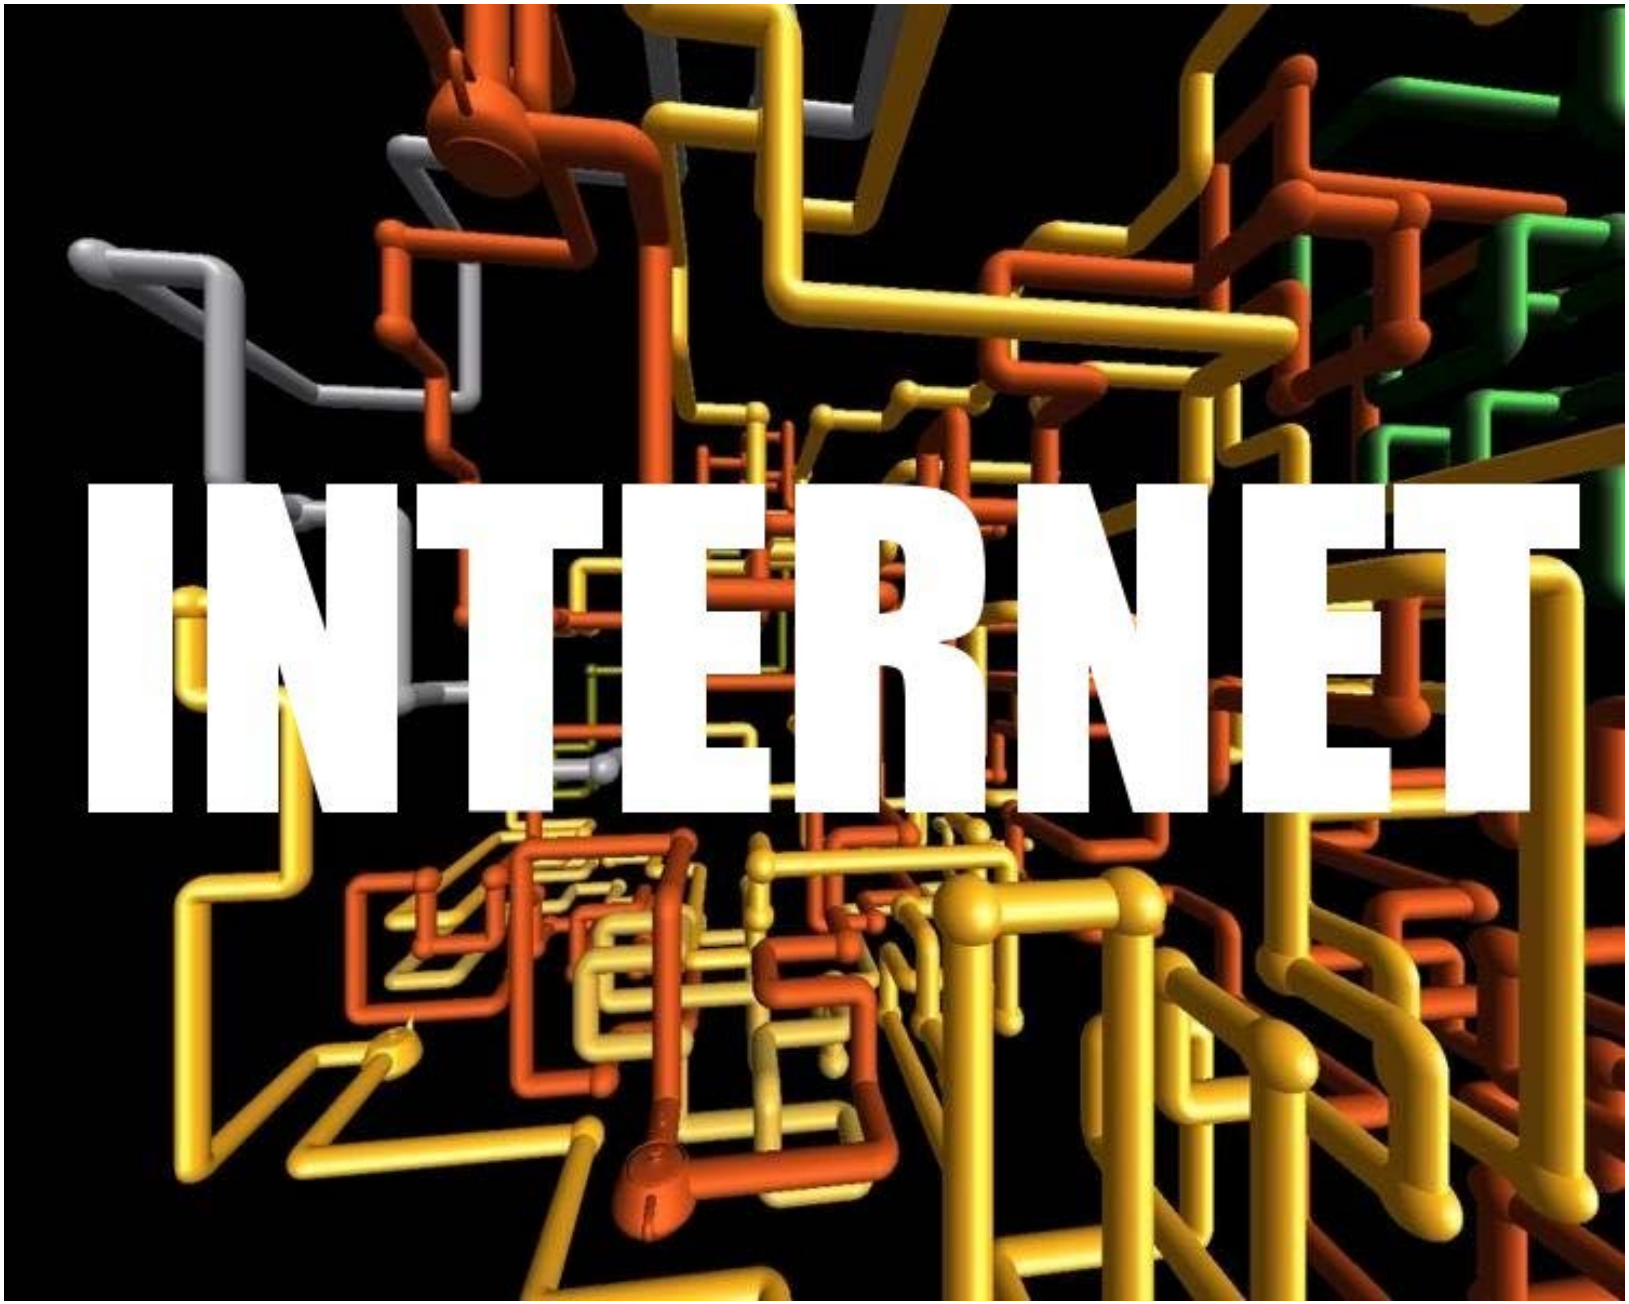
INTERNET

Understanding the Internet ecosystem

- The Internet itself, the global network of networks

  … complex, highly dynamic and opaque by design

- The systems and applications built upon it (e.g. CDN, P2P, VoD, social networks, …)

  … that extend multiple networks, administrative domains, and technologies

- The way in which we use these systems and applications, and thus the underlying network

  … something constantly changing and generally untrackable

Making it concrete with BitTorrent

- A peer-to-peer system that leverages everybody's resources (disk, computation, bandwidth, …) to support content distribution
- Natural scalability, inherent robustness, high flexibility, …
- Over two-thirds of the P2P user population

BitTorrent

- Distribute content by trading pieces between peers

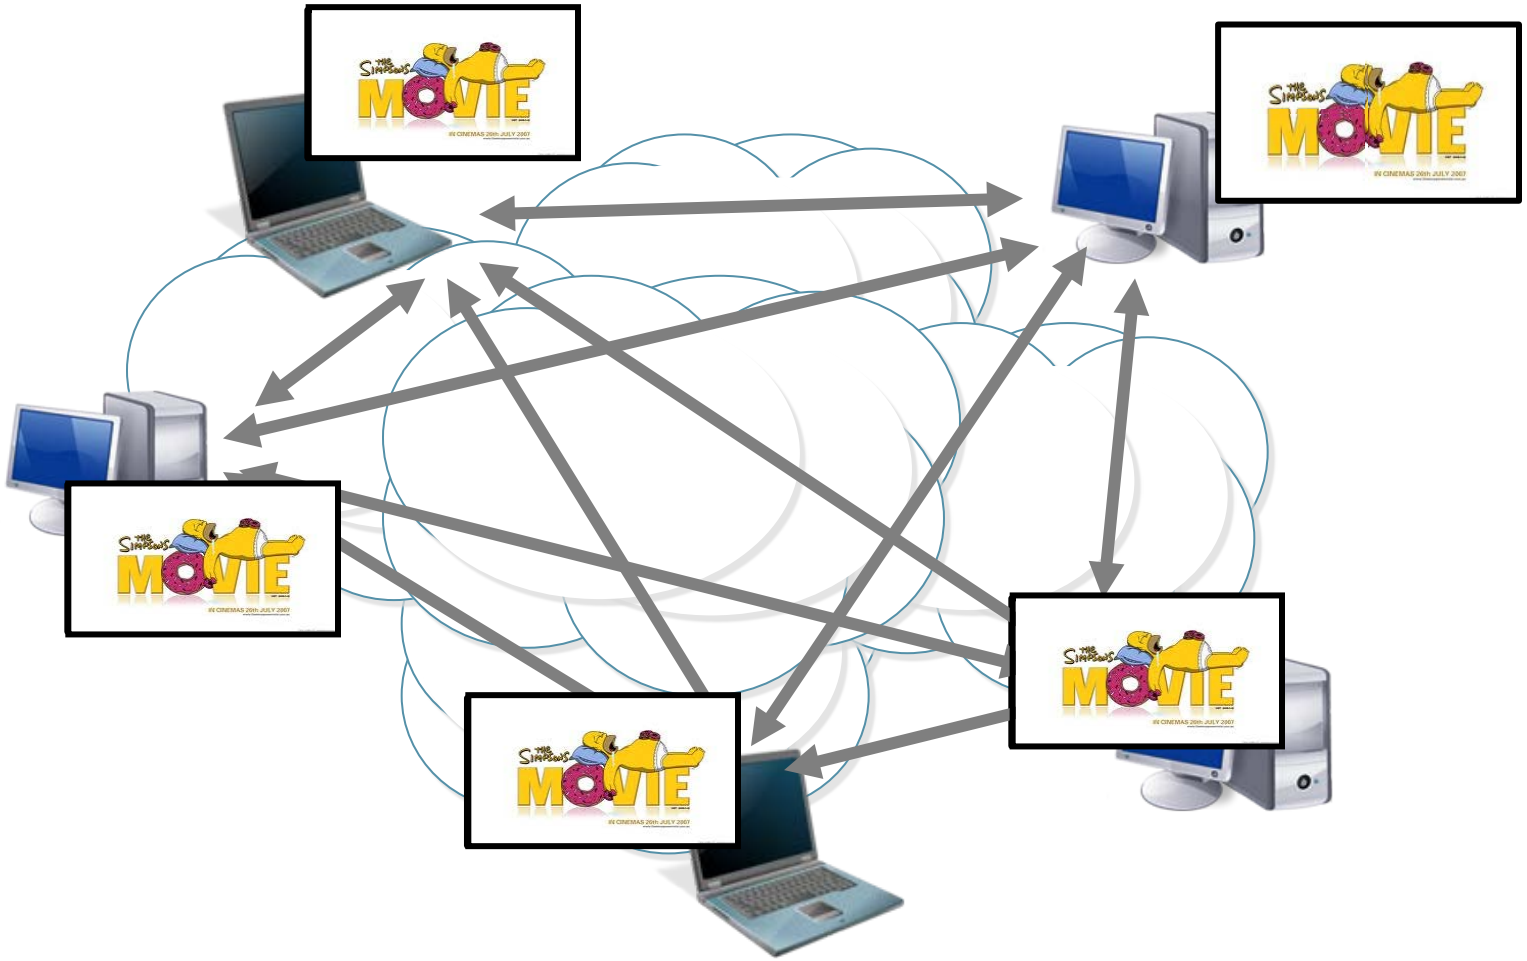

Making it concrete with BitTorrent

- How much of network traffic is from BitTorrent?

What you measure depends on where you measure from

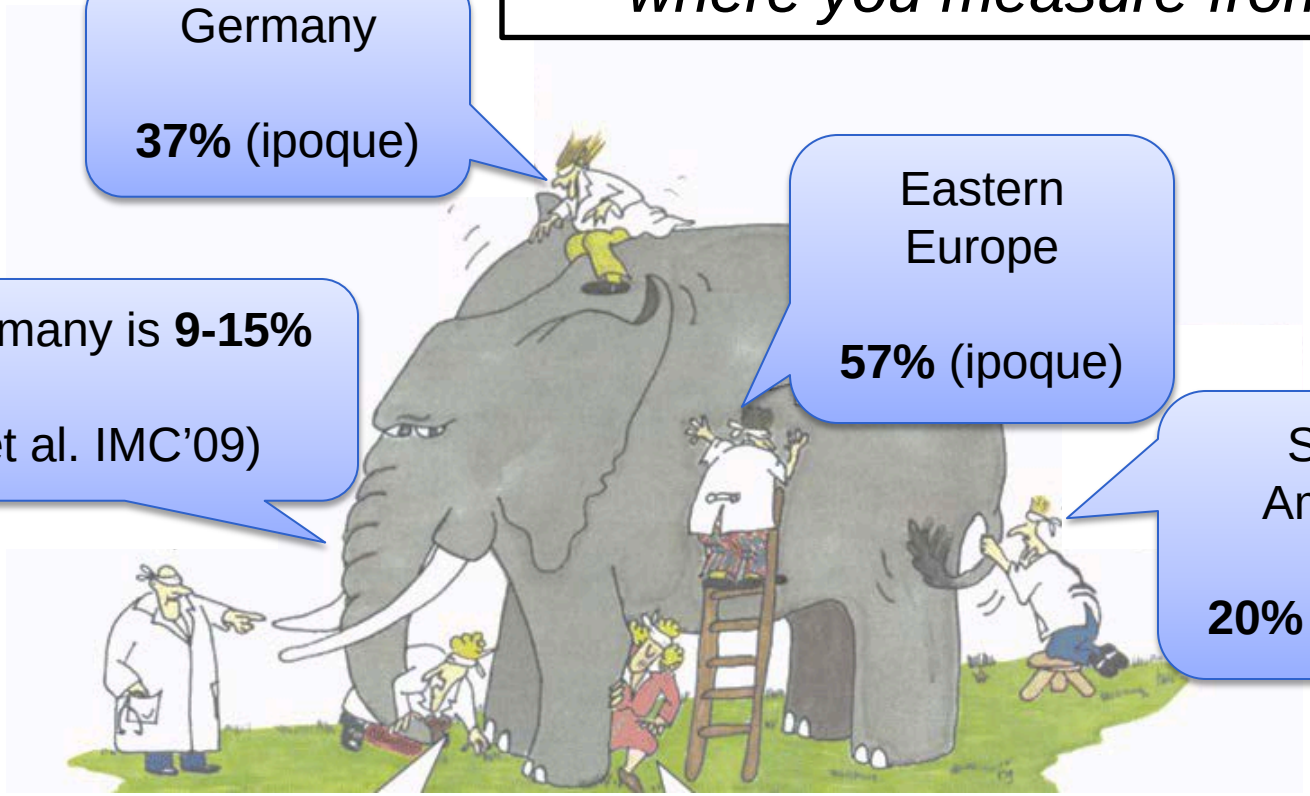

Where the Internet grows

- Internet penetration and growth per region

Region	Penetration	Growth (2000-'11)
North America	78%	152%
Oceania Australia	60%	179%
Europe	58%	353%
Latin America/Caribbean	36%	1,037%
Middle East	32%	1,987%
Asia	24%	707%
Africa	11%	2,527%

Source: http://www.internetworldstats.com

Where we can typically measure

- Percentage of Autonomous Systems (ASes), per type, with a given number of Looking Glass servers

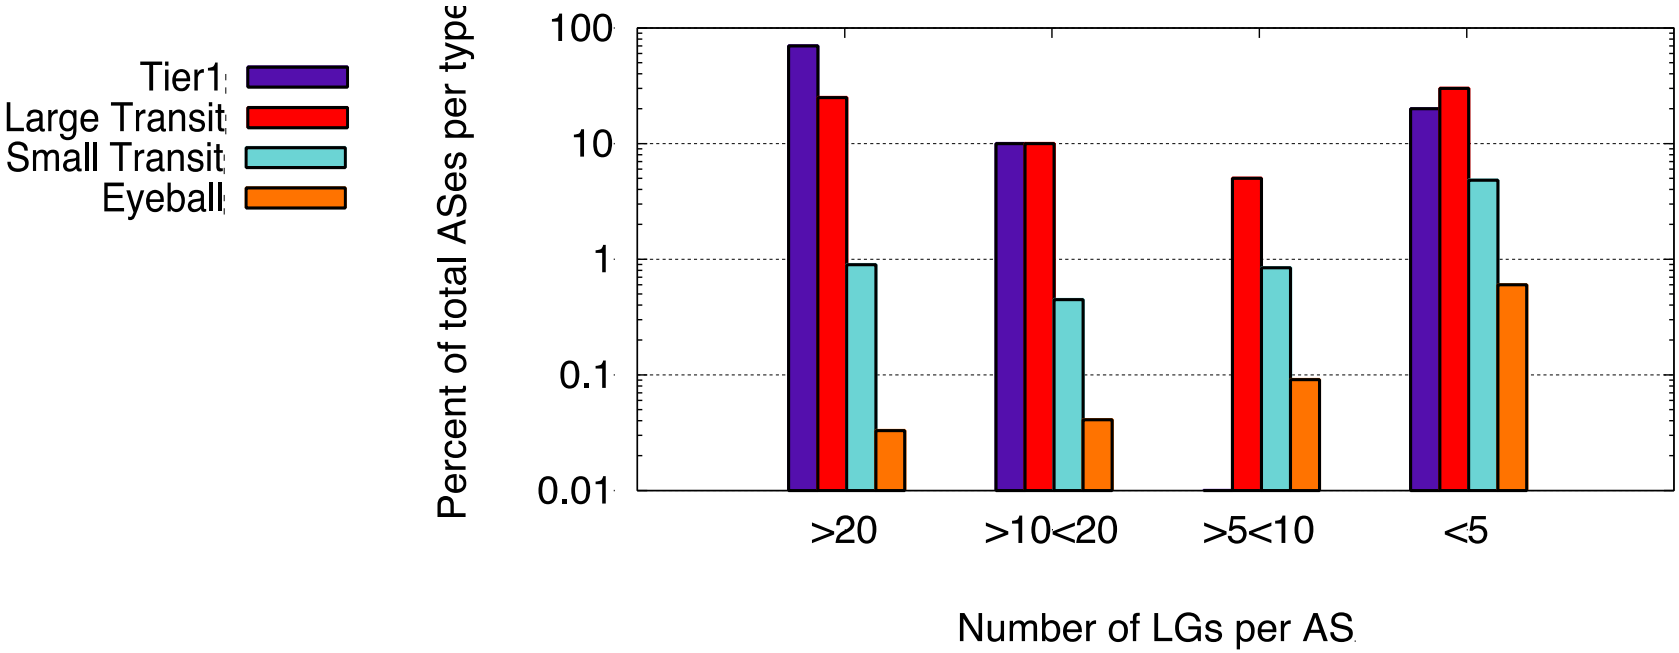

- *Less than 1% of eyeball Ases have any LG coverage*
- *PlanetLab is not better with ~0.5% Eyeballs covered*

And where BitTorrent traffic flows

- Fraction of traffic that reaches tier X

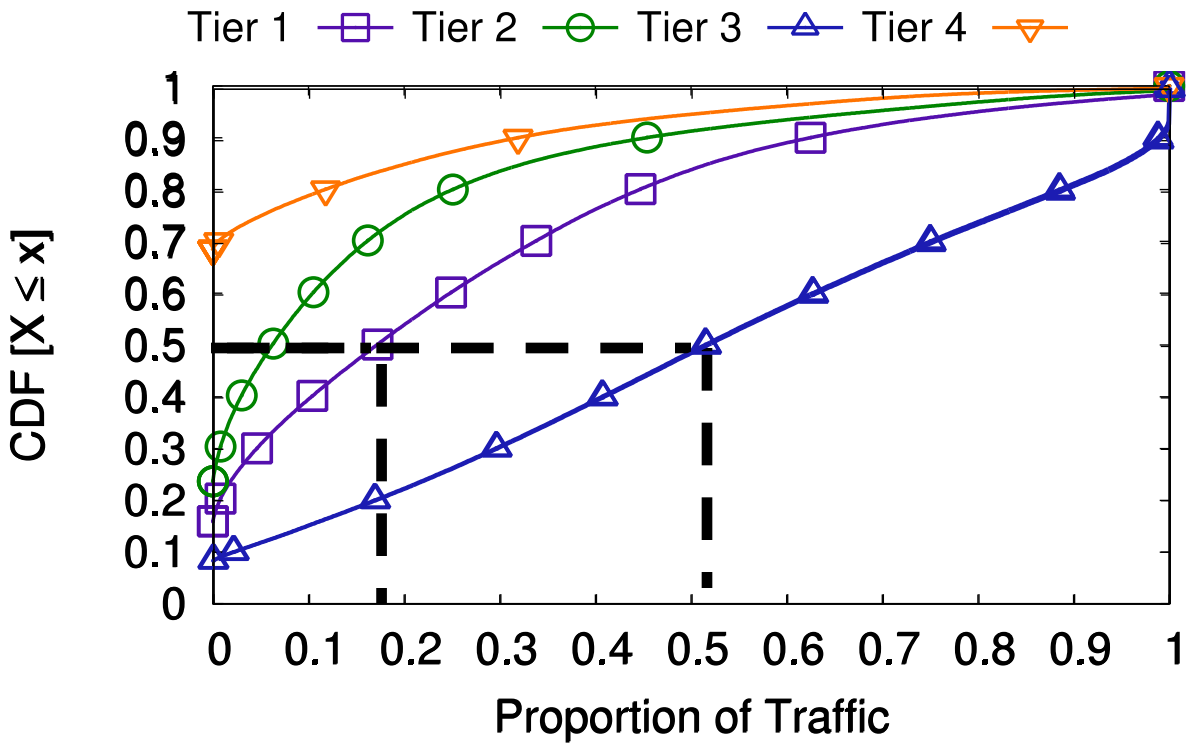

- Most traffic stays at or below Tier 3
- Large fractions never reaches Tiers 1 or 2
  - Typically missed by in-network monitoring studies from the core

How expensive is this traffic?

- Consider the commonly used 95th-percentile billing
  - Aggregate link volume for each 5 minute bin
  - Cost based on 95th-percentile bin's value

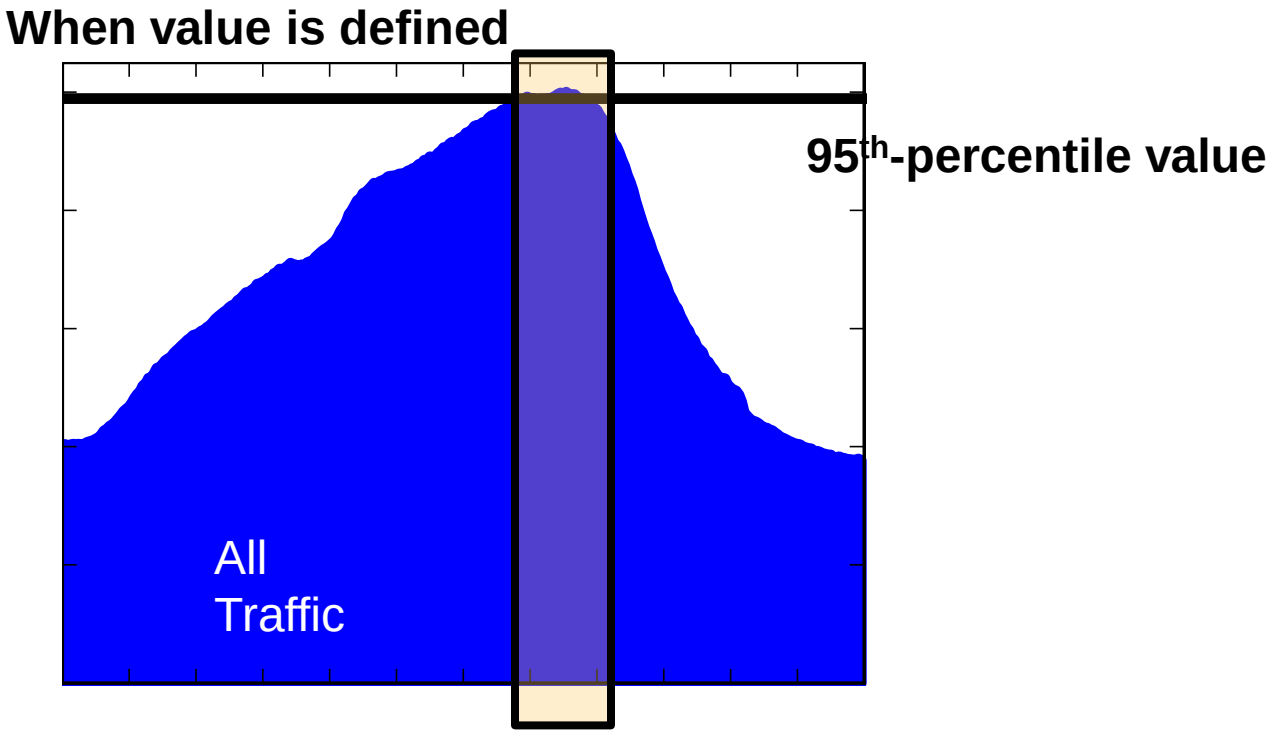

- *Some bytes are much more expensive than others*

How expensive is this traffic?

- Cost of BitTorrent traffic depends on how and when the system is used

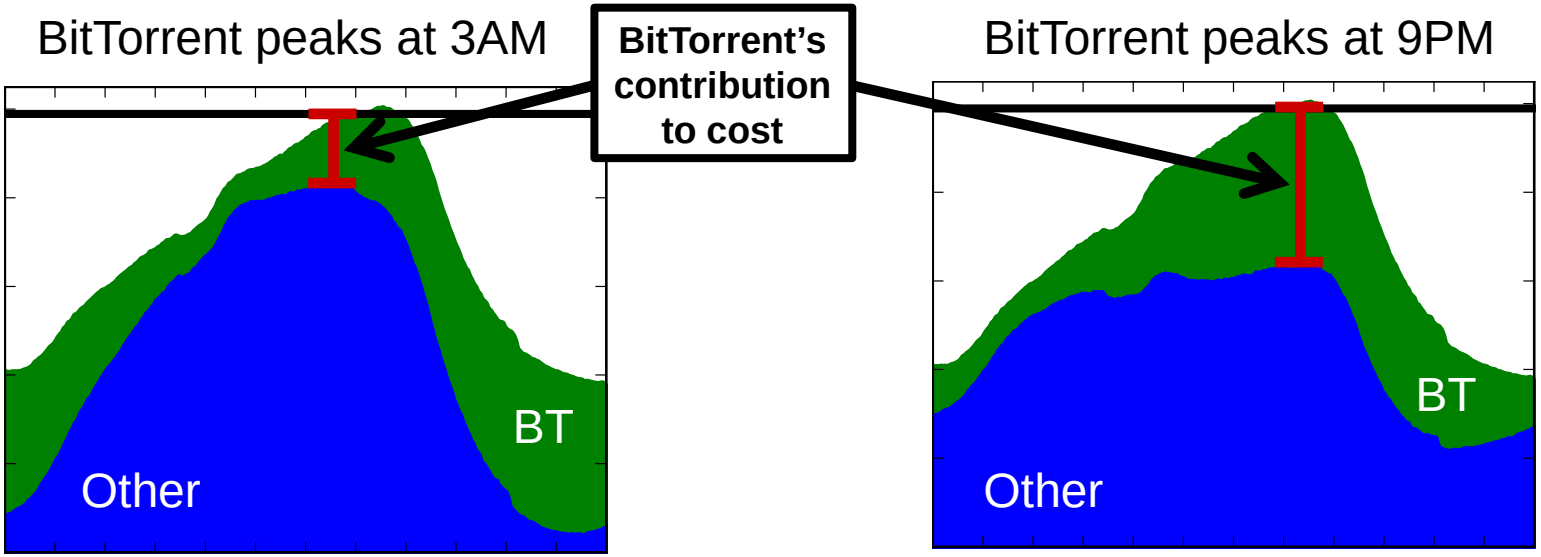

- *BitTorrent at peak hour is more expensive*

Systems' usage patterns change

- BitTorrent trends toward stronger diurnal patterns
- Shift away from overnight use
- Peak usage aligns with evening hours, local time

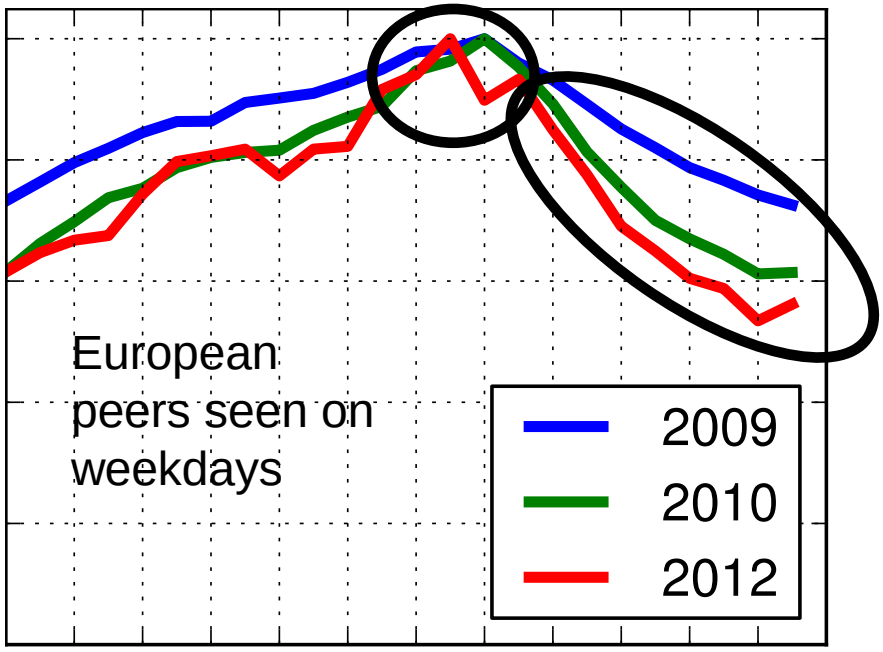

Normalized number of peers seen per hour in Europe, depending on time of day

Relative cost of BitTorrent traffic

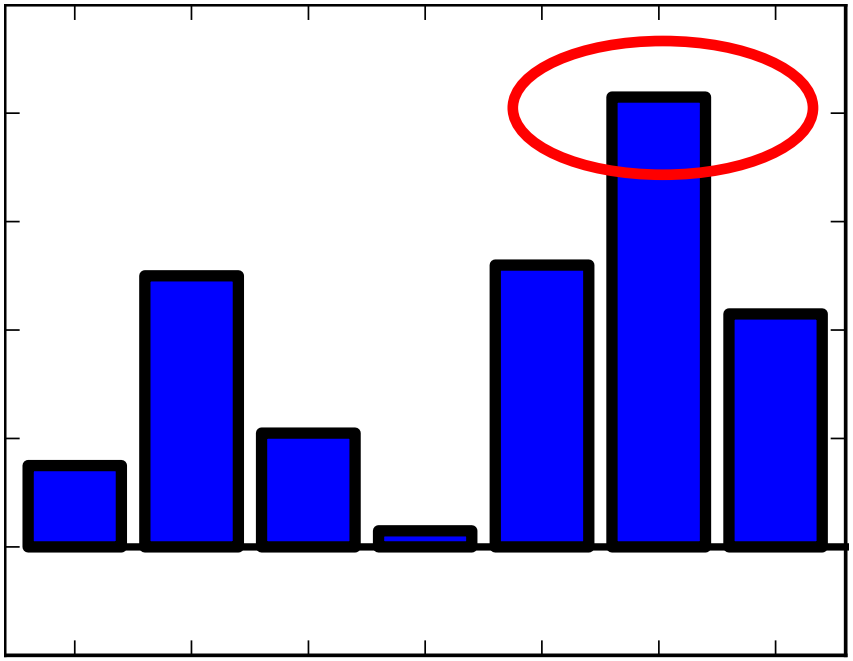

Additional cost of BitTorrent traffic, % above relative cost of 1

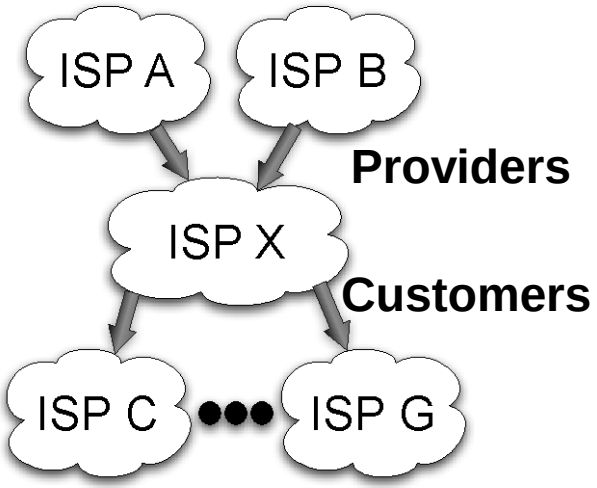

- BitTorrent traffic is generally *more* expensive than other traffic

Common problem – Visibility

(Lack of) visibility

- Into growing networks such as broadband networks and those in developing regions
- Into applications running on end-systems
- Into the behavior of these applications' users

Outline

- Overview
- Challenges to understanding today's Internet
- Our approach – go to the edge for a better view
- Example view – Understanding Public DNS
- Concluding remarks

Our approach – go to the edge

- Capture the edge of the network "recruiting" the help of end users
- Aligning experimenters' goals with user incentives
- By building and releasing software that leverages existing platforms
  - In our case peer-to-peer, particularly BitTorrent

How?

- As extensions to a popular BitTorrent client

- Three most popular ones
  - To improve user download performance and reduce cross-ISP traffic – ***Ono***
  - To collaboratively detect network problems – ***NEWS***
  - To collaboratively characterize broadband services – ***DASU***

- Together

1,500,000 users worldwide

The perspective we have gained

- Peers expand over 218 countries (codes)
- 11,700 networks (ASs)

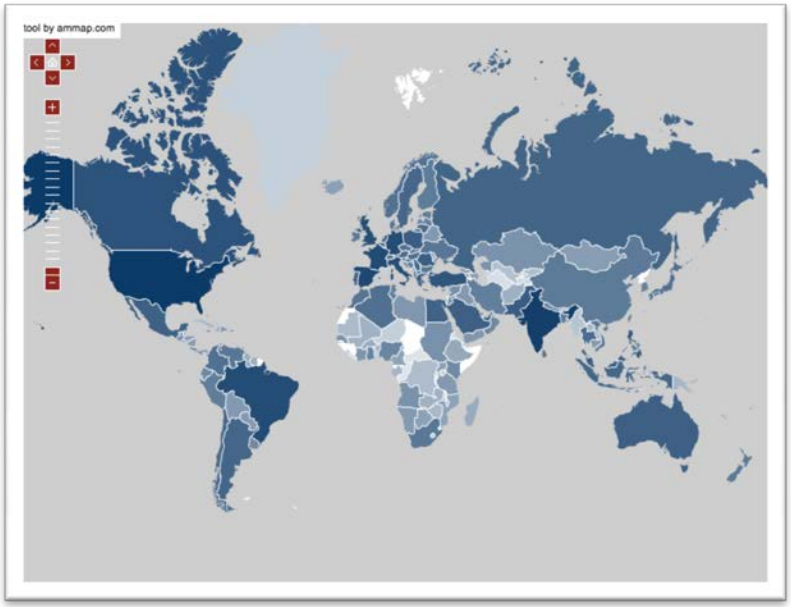

- Each peer contributes to a unique perspective
  - Passive measurements (download/upload rates, latency on the last mile, latency to DNS, …)
  - Some active measurements (traceroutes, pings, DNS resolutions, CDN mappings, web page download, …)

And what we are learning from it

- How people use distributed systems and what it means to us – researchers, designers, network operators and *users*
- What is really happening in the Internet that we may be just missing
- As always, some great ideas do not seem to work that great
- What does the Internet topology actually looks like
- …

Last example – *is Public DNS good for you?*

- Public DNS is fast growing in popularity
  - Our view – 27% annual growth rate

- The appeal of public DNS
  - Low latency ✔
  - High reliability ✔
  - No filtering ✔
  - A fast web ?

(Somewhat) Hidden Interactions

- Content Distribution Networks and DNS interactions
- A 2' CDN primer
  - Push popular content close to web clients at thousands of replica servers worldwide
  - Redirecting web-clients to content replicas nearby

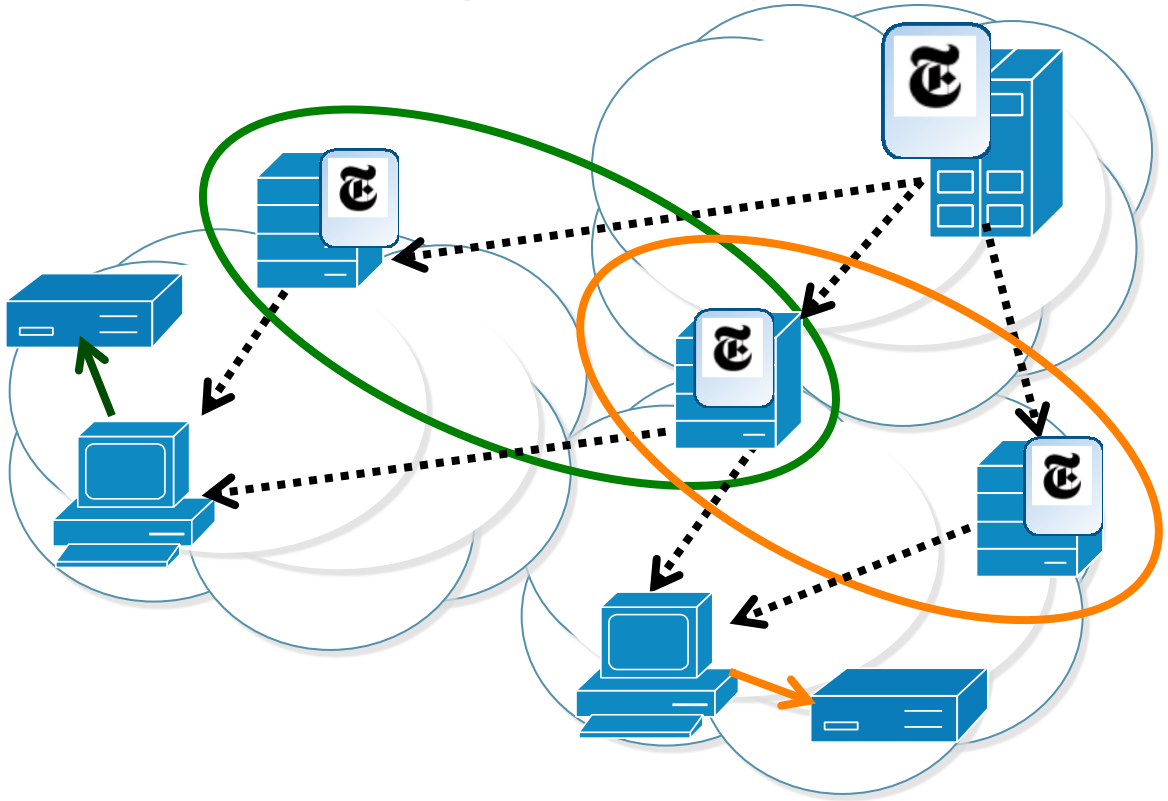

Ubiquity of Content Delivery Networks

- Several motivations for using CDNs
  - Performance, scalability, reliability
- Most popular sites use them

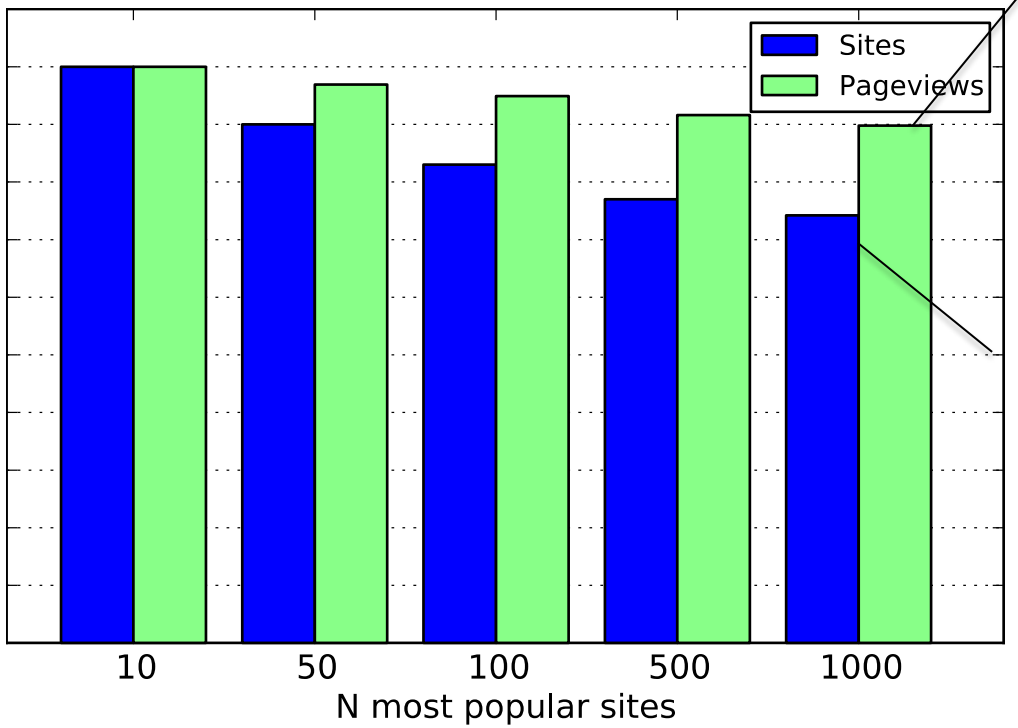

90% of 1000 most visited sites

75% of top 1000 sites

Hidden Interactions and Hidden Costs

- CDNs use DNS to map clients to servers
  - Assume proximity of client to DNS resolver

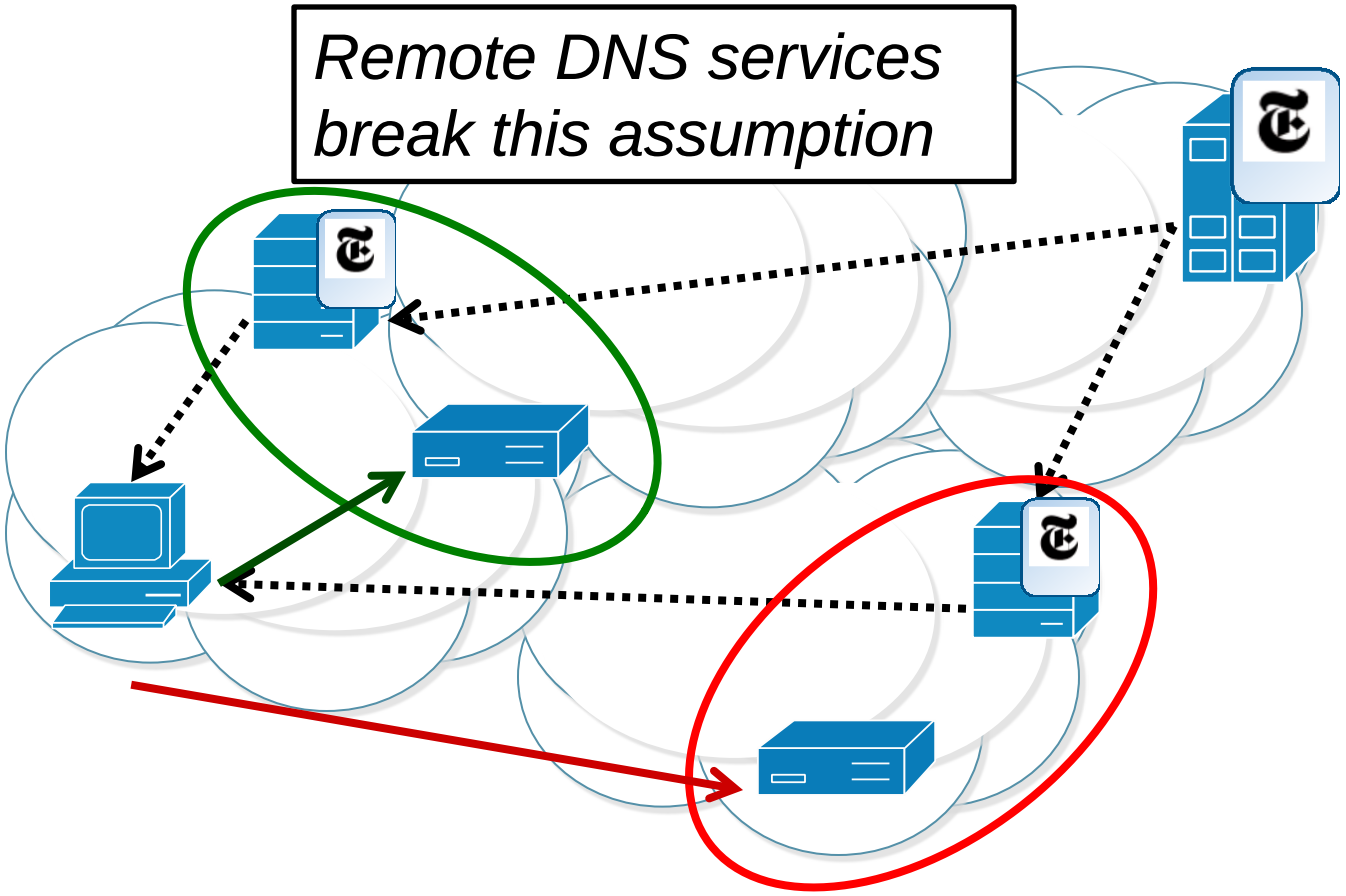

DNS latencies

- Latency difference (%) between public and ISP DNS resolvers

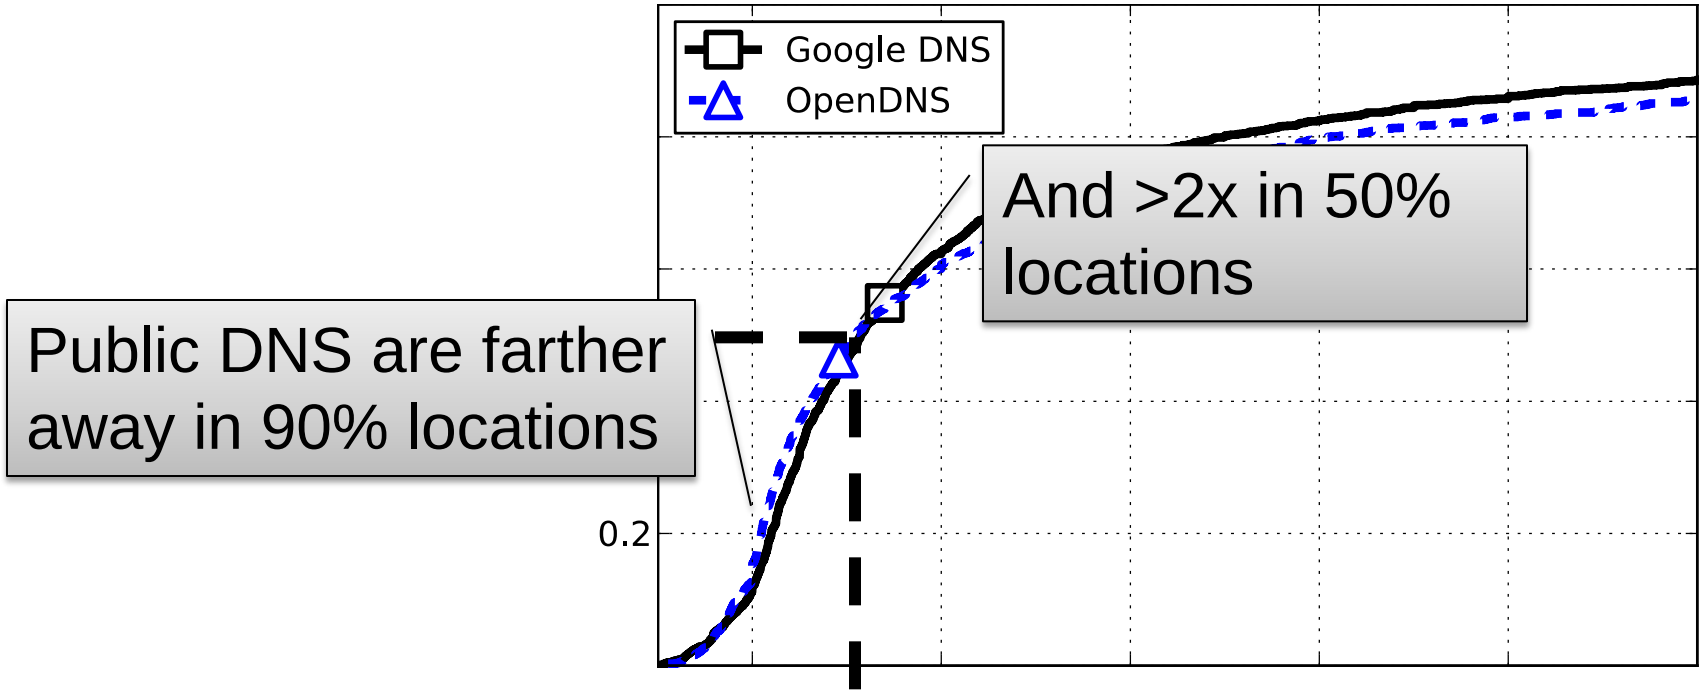

Impact of public DNS on CDN redirections

- Clients using public DNS services see radically different set of replicas

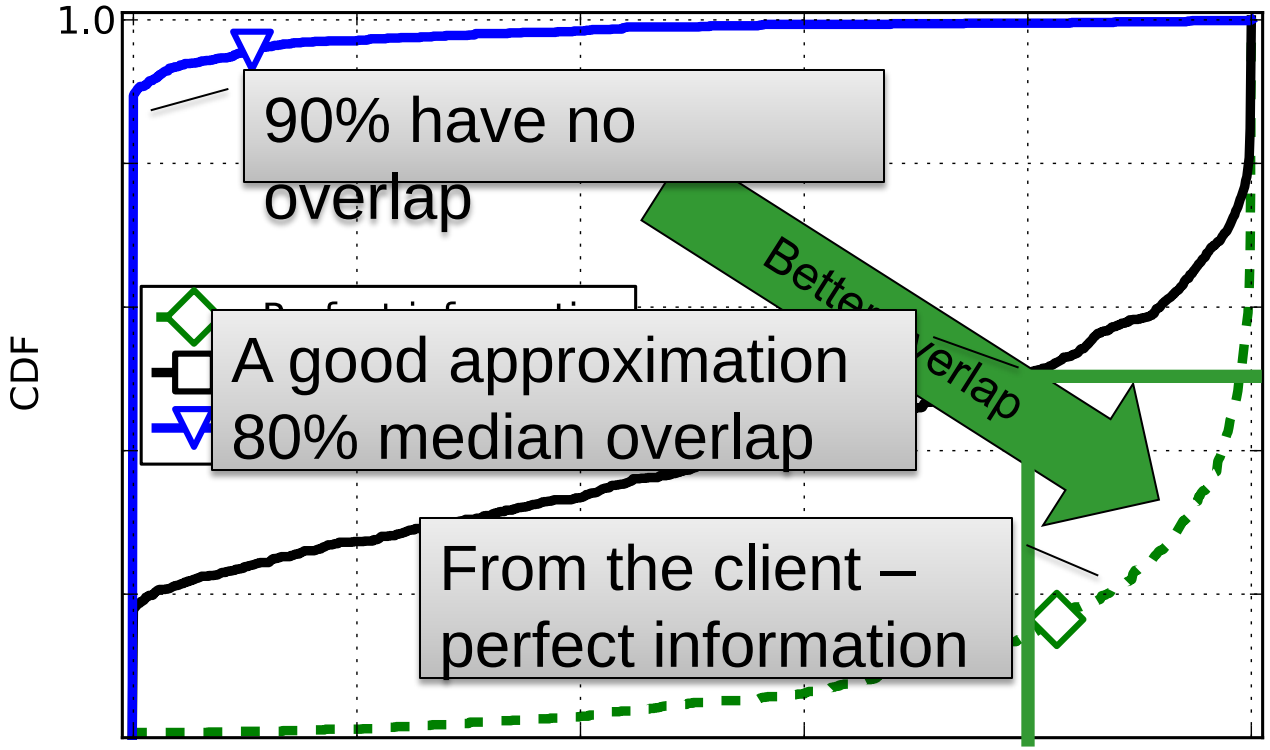

Different impact to different CDNs

- HTTP latency difference (%) using iterative, ISP or public DNS relative to ideal baseline

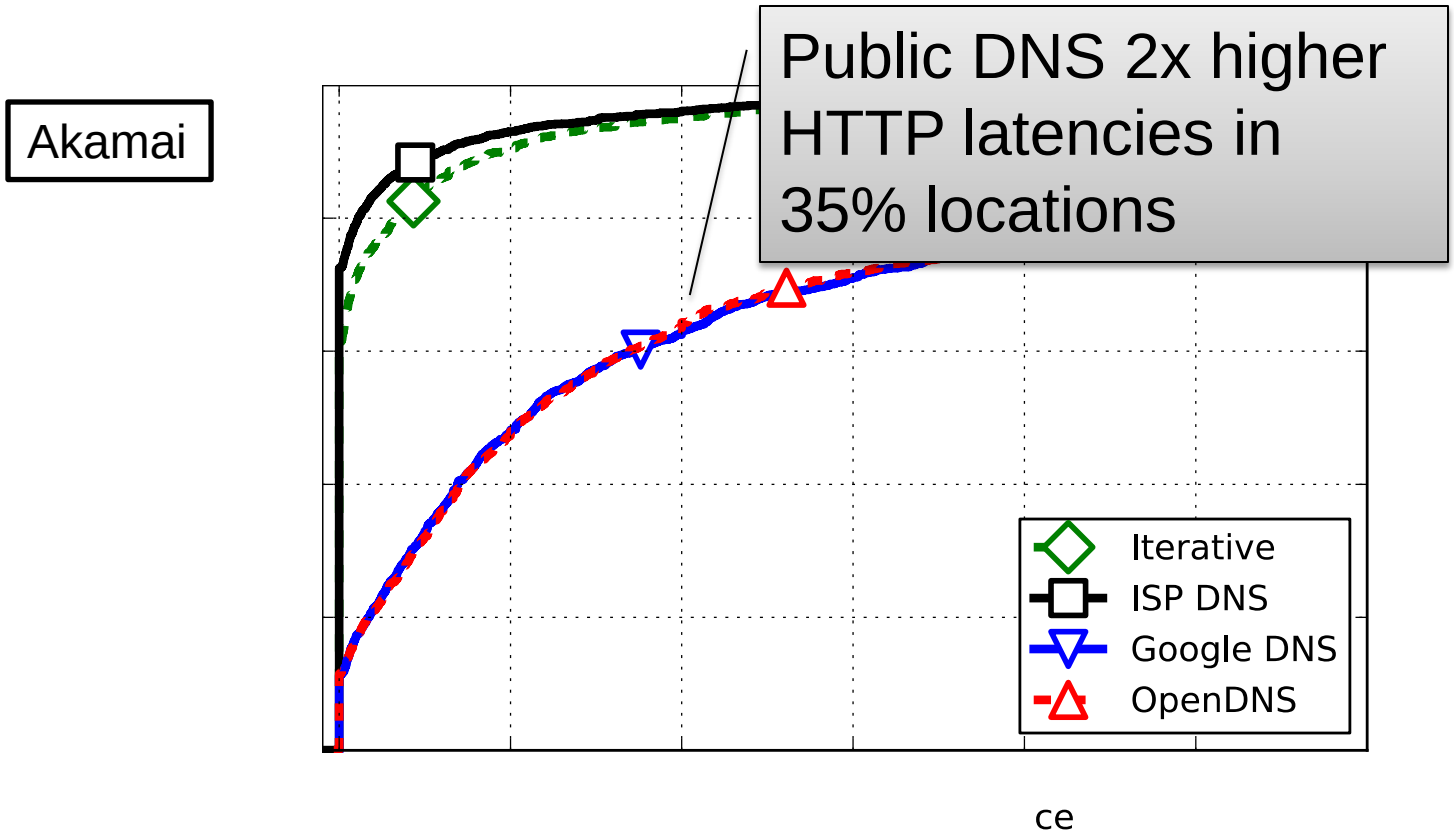

An Industry response – A DNS extension

- Resolver sends client's subnet to CDN's DNS
  - Redirections are based directly on client's location
  - Requires participation of DNS *and* CDN services
- Limited adoption to date

Conservative:
Most from using Google services

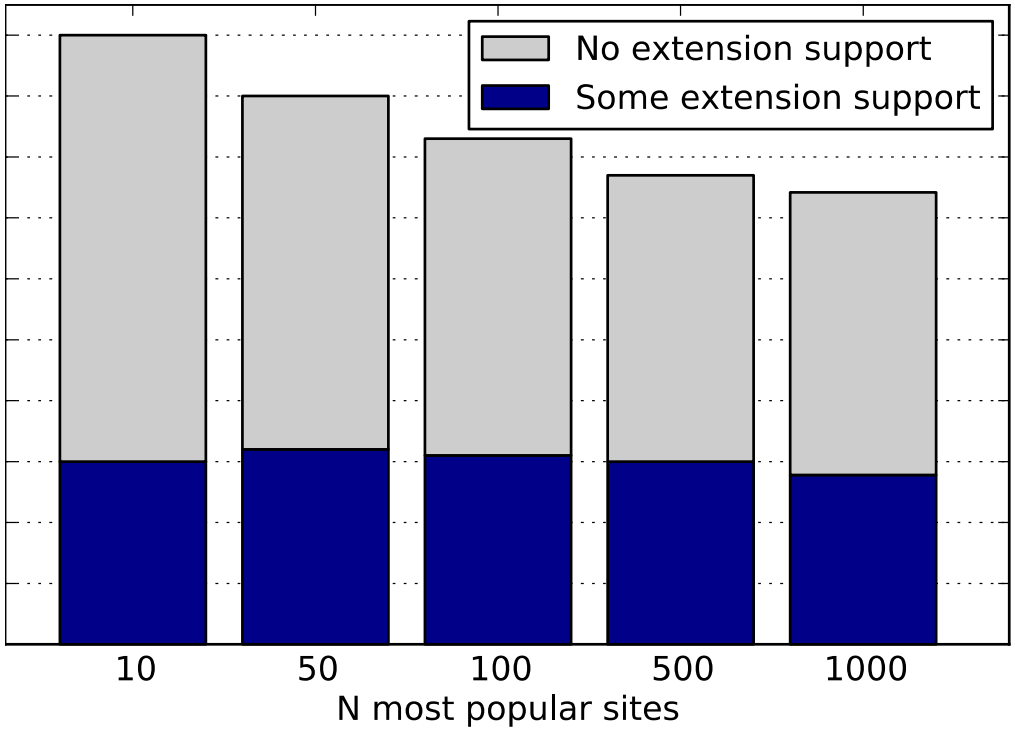

Different benefits to different CDNs

- End-to-end latency % difference, comparing ISP DNS and public DNS with and without the extension

Akamai

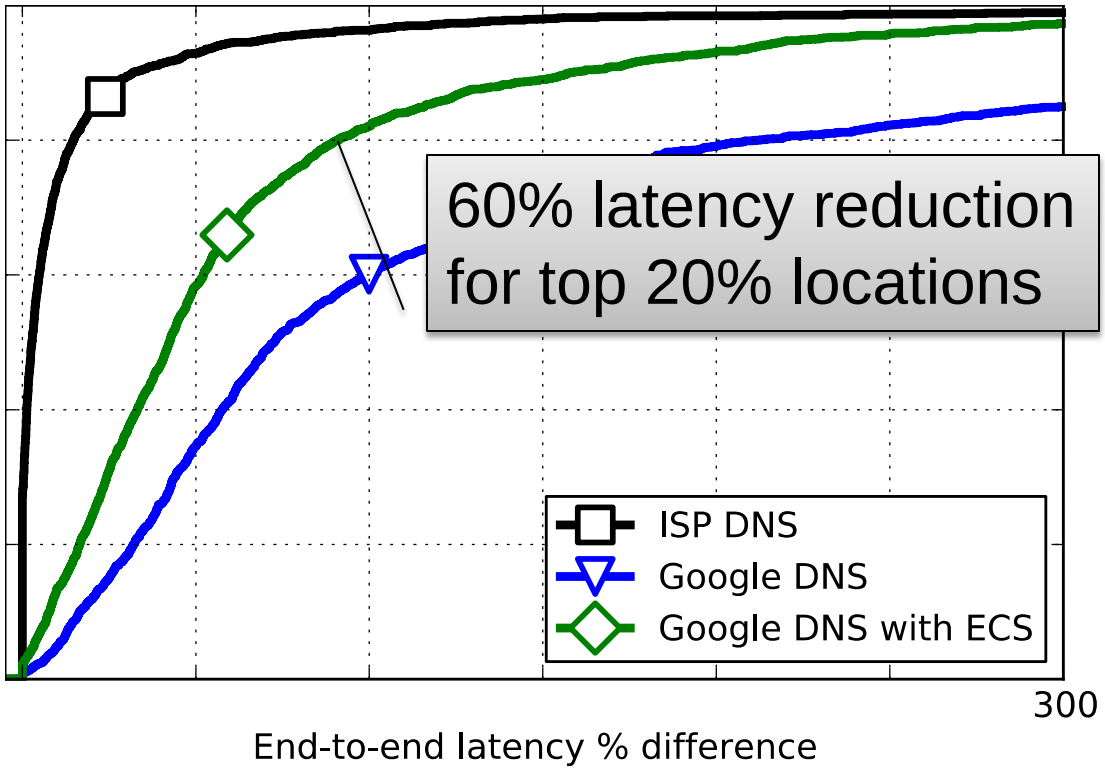

An end host solution – namehelp

- A smart DNS resolver running at the end host
  - Leverage Public or ISP DNS when possible
  - Avoid it when costly, i.e. on the last step

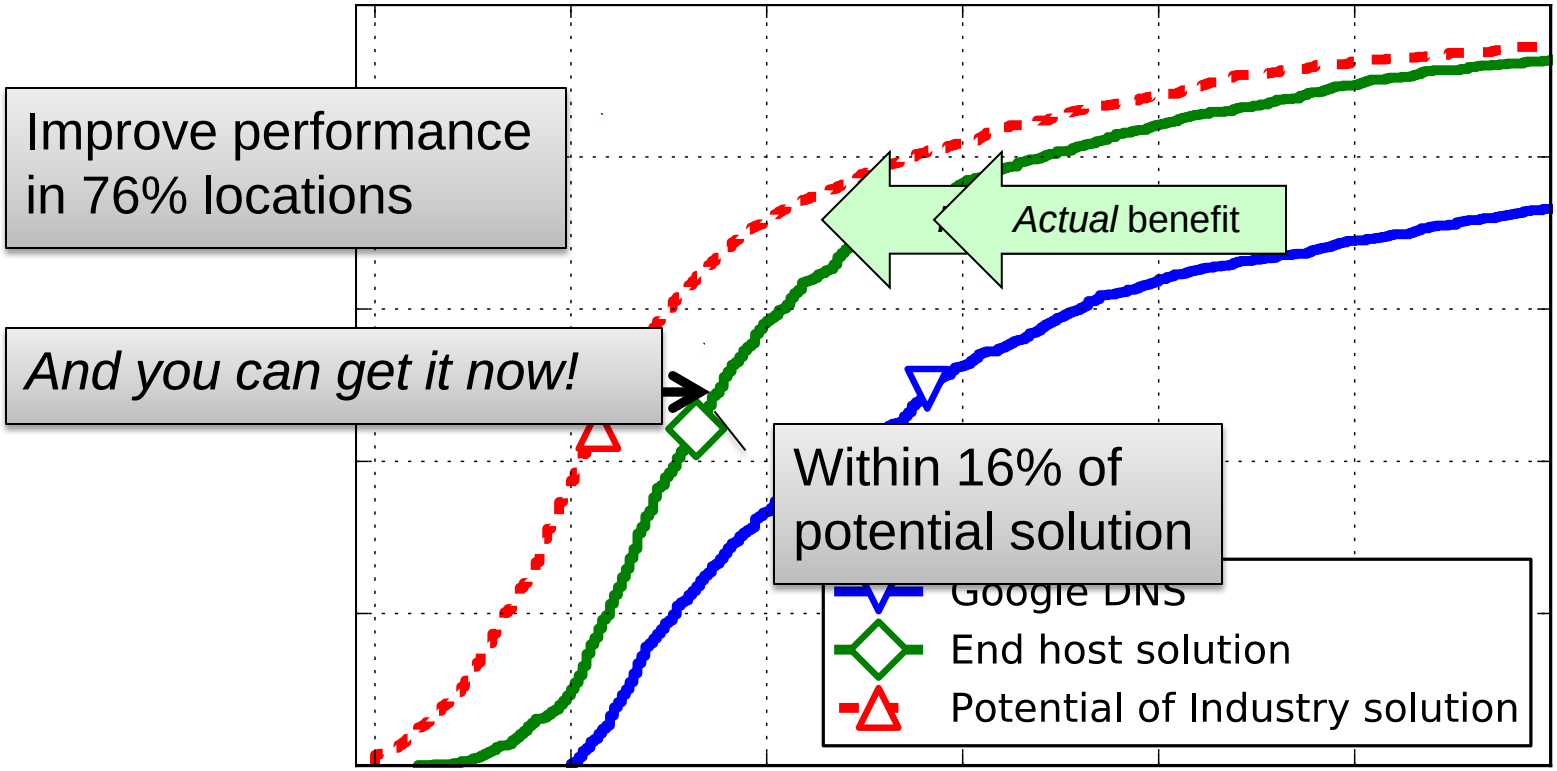

Some closing thoughts

- Discussed an approach to ride the wave of the Internet growth and made suggestion for how to support it
- Can we get ahead of the wave?
- *A new challenge – building experimental platforms that stand at the edge of the network*

Challenges for a Network Edge Platform

- How to get hosts running the edge, e.g. at homes, on-board?
- What sort of experiments should be allowed?
- What programming model should we use?
- How do we coordinate resource usage?
- How do we provide access to researchers?
- How do we prevent accidentals DDoS?
- …

A Network Edge Platform – Dasu

- Broadband characterization as the incentive
  - *Are you getting from your ISP what you are paying for?*
- Many experiments aligned well with this goal
- Declarative language for experiments
- A curator-based model for sharing
- A novel coordination approach for experiments
- …

>75,000 users worldwide

Acknowledgement

- Graduate students
  - David Choffnes (graduated)
  - John Otto
  - Mario Sanchez
  - Zach Bischof
  - John Rula
  - Ted Stein
- Collaborators
  - Yan Chen
  - Zihui Ge (AT&T)
  - Bala Krishnamurthy (AT&T)
  - Alek Kuzmanovic
  - Dan Pei (AT&T)
  - Georgios Siganos (Telefonica)
  - Walter Willinger (AT&T)
- Funding sources
  - National Science Foundation
  - Google
- My family

Core Internet Lessons from Today's Network Edge

Fabián E. Bustamante

Northwestern U.

Evanston/Chicago, Illinois USA