

Drafting Behind Akamai (Travelocity-Based Detouring)

Ao-Jan Su, David R. Choffnes,
Aleksandar Kuzmanovic and
Fabián E. Bustamante





Department of EECS
Northwestern University

ACM SIGCOMM 2006

Drafting Behind Akamai

Drafting



Find Me The Best Priced Trip!

Flight Flight+Hotel Hotel Car

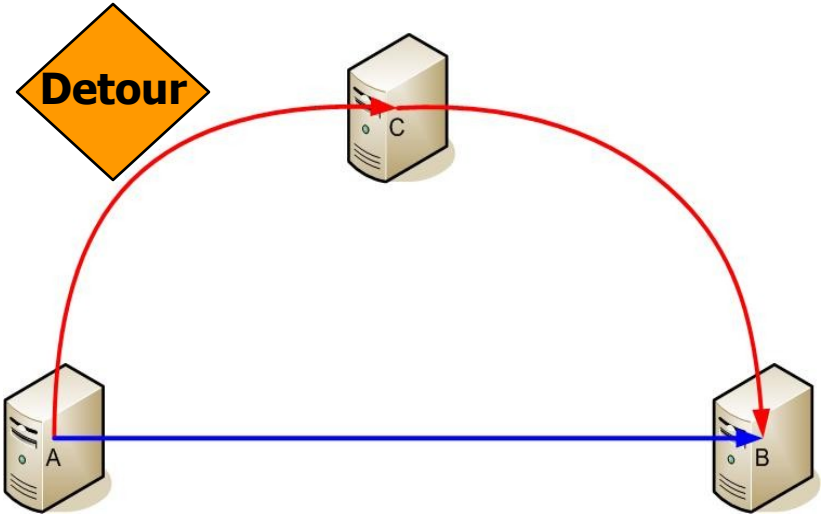
From: To:

Compare surrounding airports ?

Exact dates +/- 1 to 3 days Flexible dates

Depart:

Return:



Motivation

- Growing number of overlay-based systems
 - Can't change IP layer, so change the layers above
 - E.g., end system multicast, anycast, i3.
- Common need for such systems
 - Build a “view” of the underlying network relying on network measurements

Problem

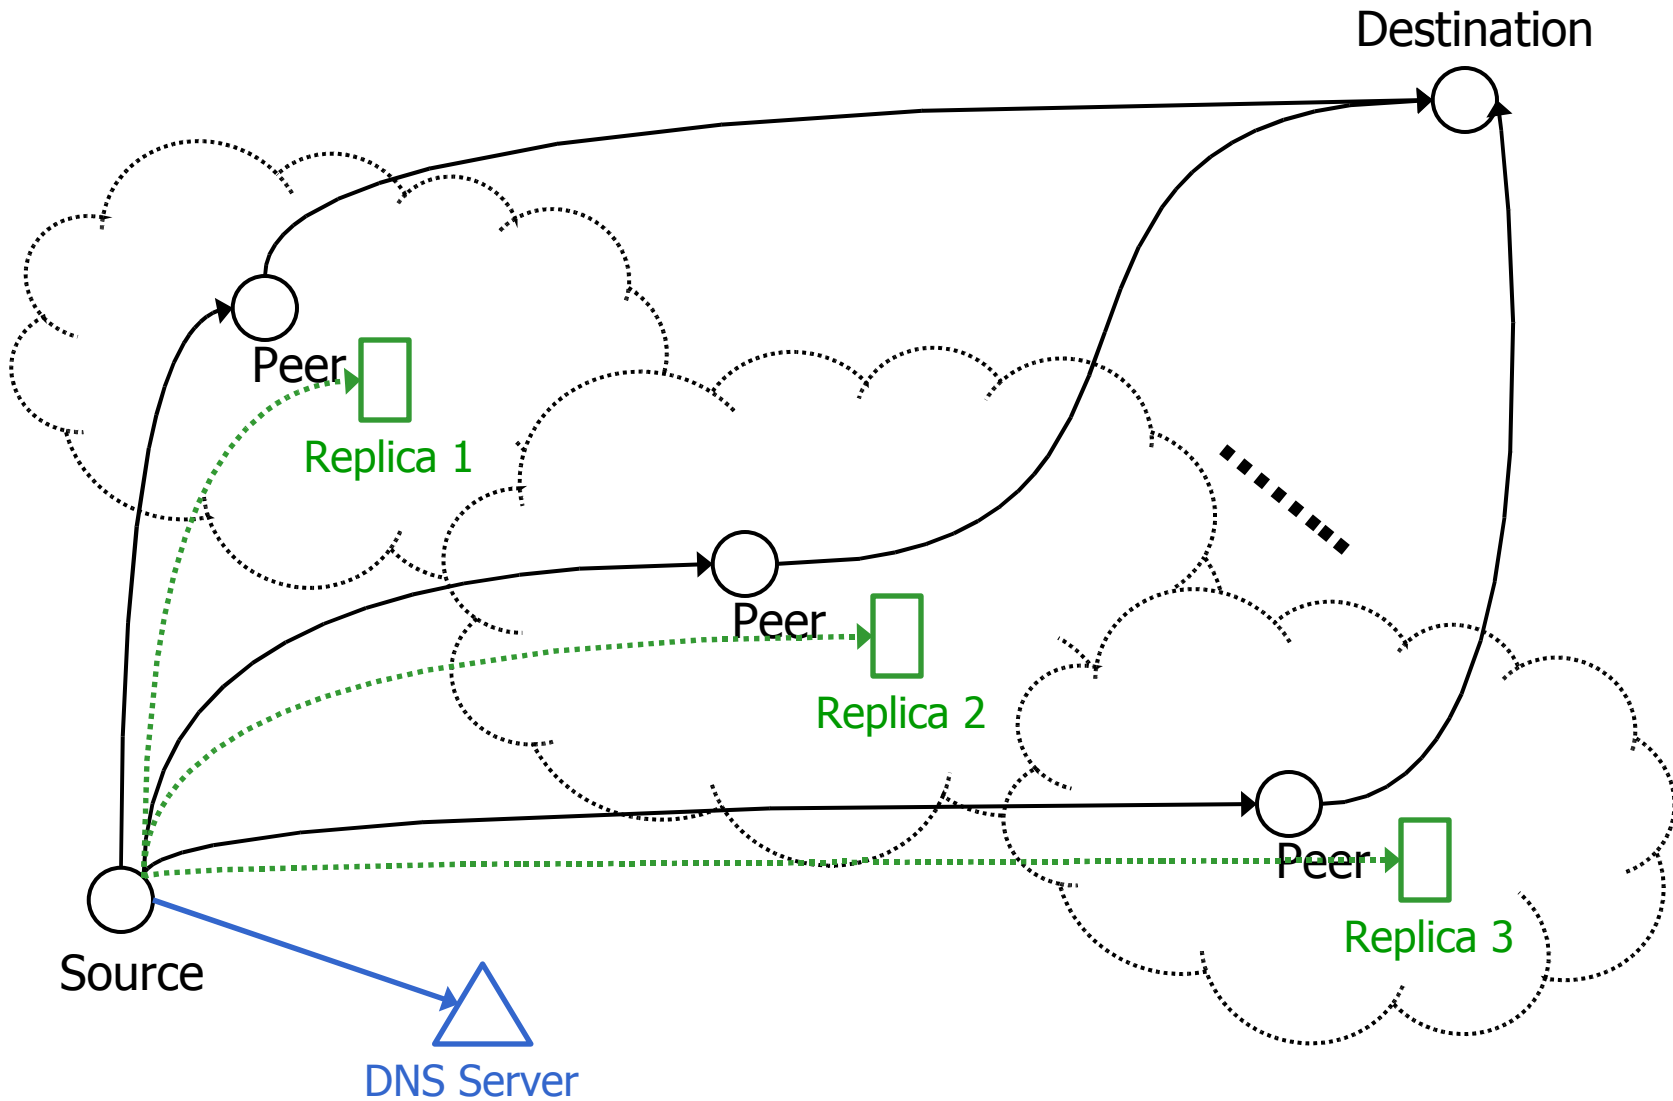
- Independent measurements
 - Redundant
 - Non-scalable
 - Can cause problems
 - E.g., synchronization
- Proposals for common services
 - Knowledge plane
 - A routing underlay for overlays
 - Network weather service

Our Approach

- Reuse the view of the network gathered by long-running services
 - Significantly reduce the amount of measurements
 - Require no new infrastructure to be deployed
- CDNs (e.g., Akamai) improve web performance by
 - Performing extensive network & server measurements
 - Publishing the results through DNS

Can overlay networks reuse measurements collected by production CDNs?

CDN-Driven One-Hop Source Routing



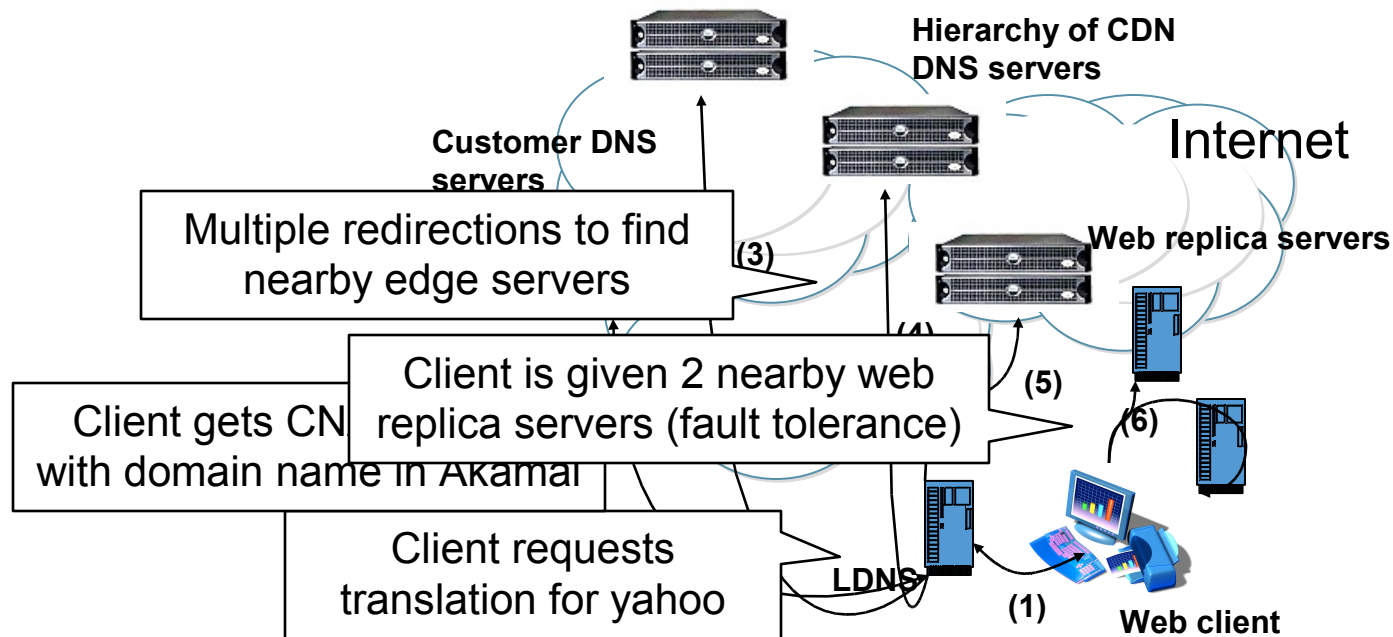
Roadmap

(or *how feasible is all this?*)

- **How does Akamai work?**
- How many web replicas does a client see?
- Impact of different web sites (e.g., Yahoo vs. NY Times)?
- What are the dynamics of DNS redirections?
- What drives redirections - network or server latency?

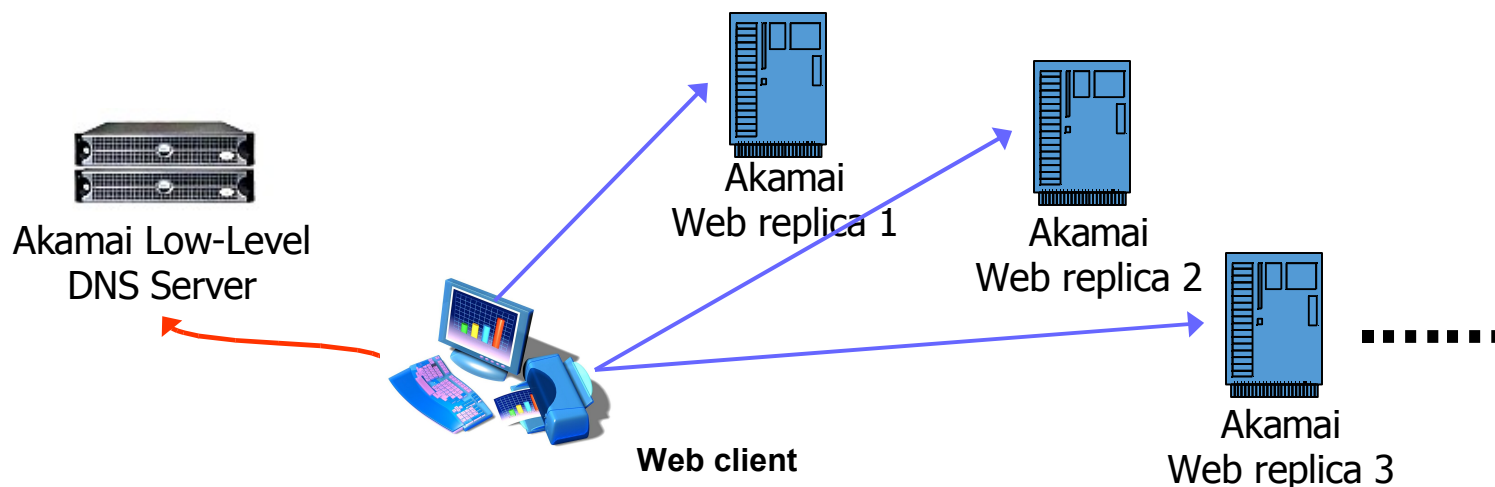
CDNs Basics

- Web client's request redirected to 'close' by server
 - Client gets web site's DNS CNAME entry with domain name in CDN network
 - Hierarchy of CDN's DNS servers direct client to 2 nearby servers



Measuring Akamai

- 2-months long measurement
- 140 PlanetLab nodes (clients)
 - 50 US and Canada, 35 Europe, 18 Asia, 8 South America, the rest randomly scattered
- Every 20 sec, each client queries an appropriate CNAME for
 - Yahoo, CNN, Fox News, NY Times, etc.



Roadmap

(or *how feasible is all this?*)

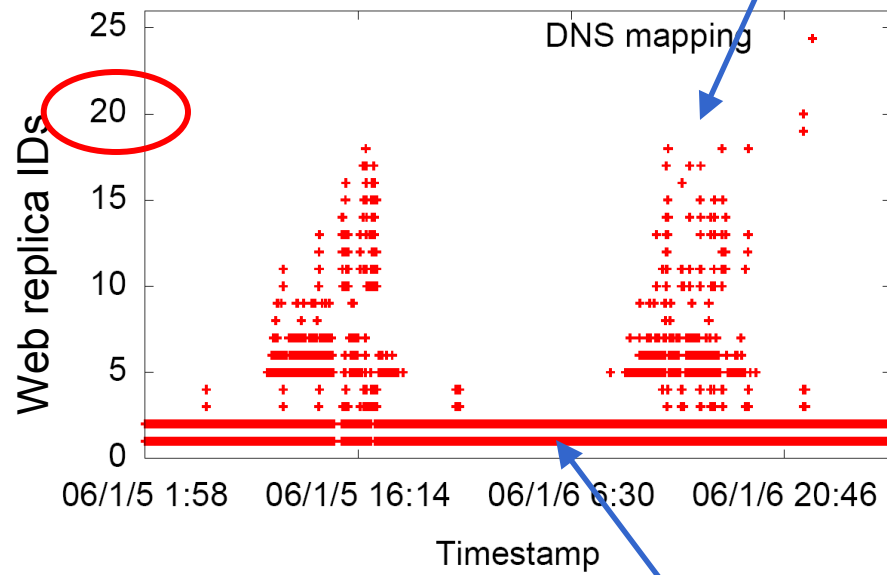
- How does Akamai work?
- **How many web replicas does a client see?**
- Impact of different web sites (e.g., Yahoo vs. NY Times)?
- What are the dynamics of DNS redirections?
- What drives redirections - network or server latency?

Server Diversity

Client 1


Closer to Akamai network

day




Client 2

Further away from Akamai network




Multiple Akamai Web Sites

Different Akamai customers (web sites) get “different” networks



Redirection Dynamics

Do redirection dynamics small enough for network control?




Roadmap


(or *how feasible is all this?*)

- How does Akamai work?
- How many web replicas does a client see?
- Impact of different web sites (e.g., Yahoo vs. NY Times)?
- What are the dynamics of DNS redirections?
- **What drives redirections - network or server latency?**

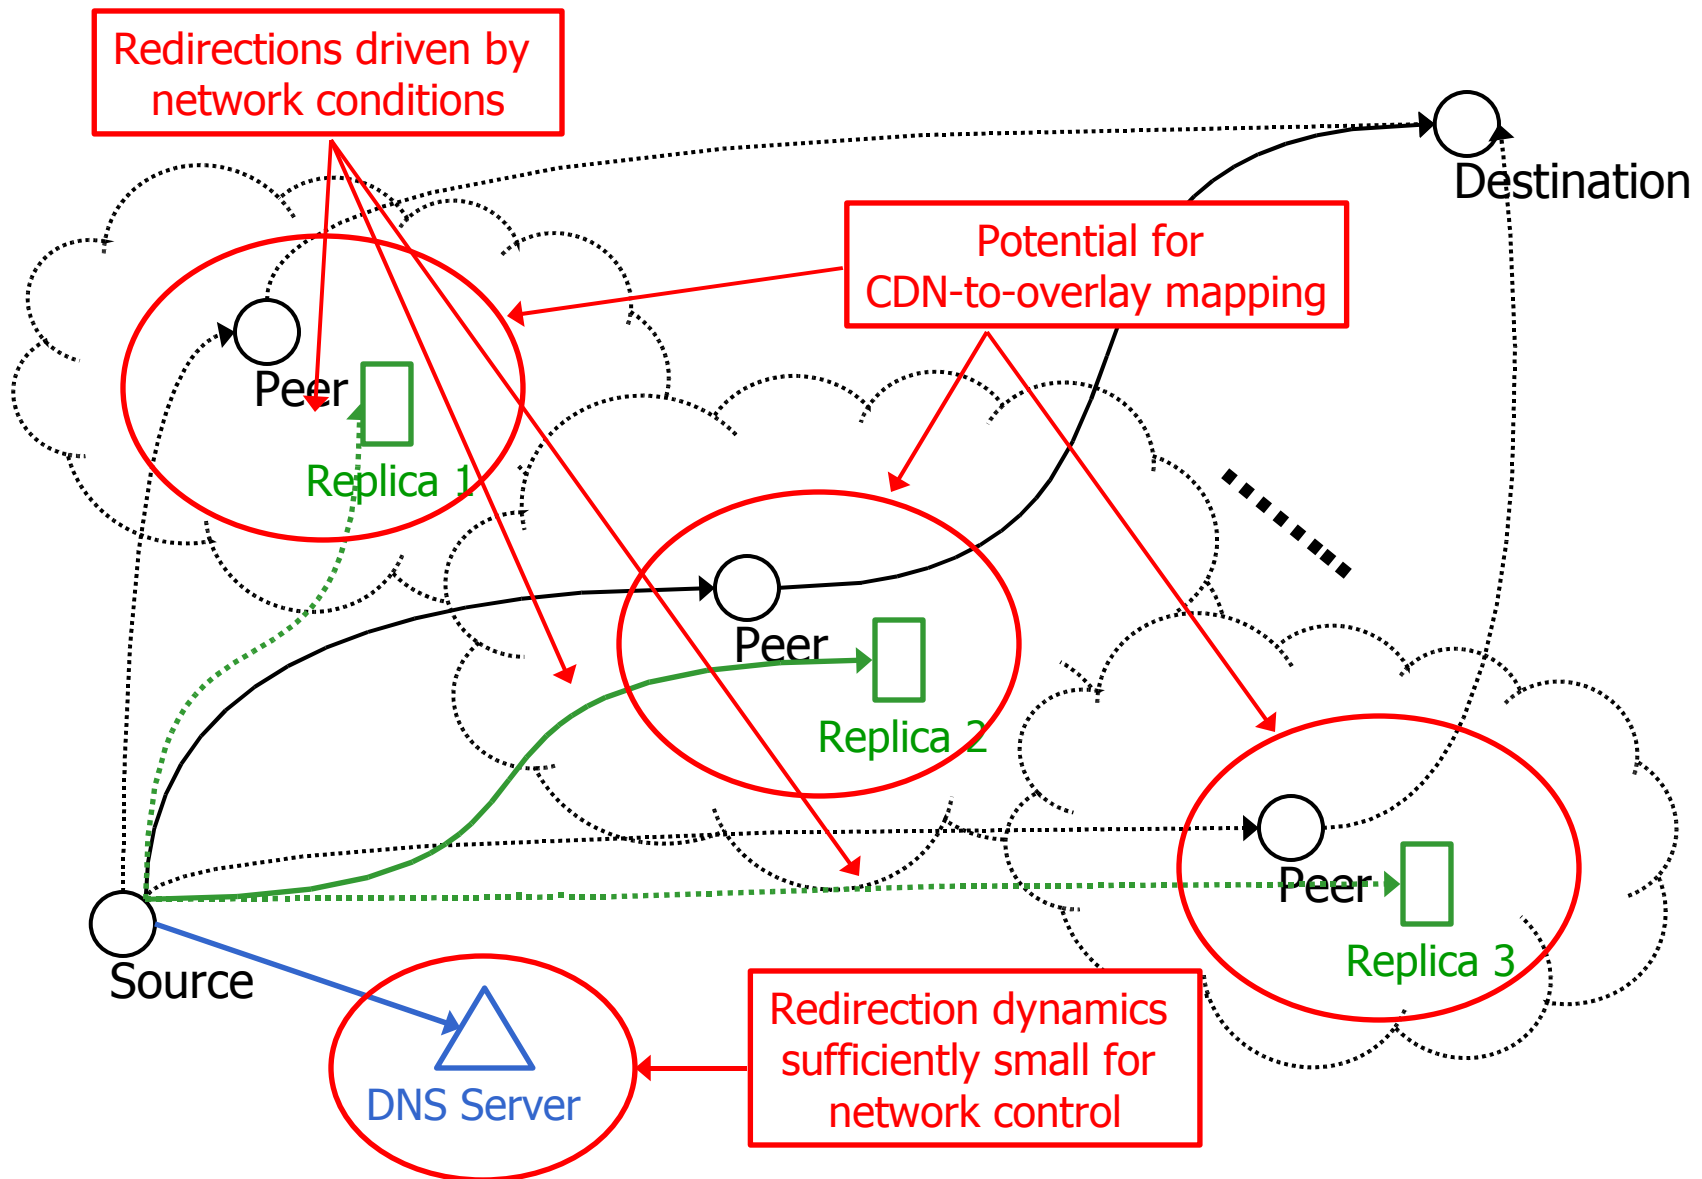
Methodology



Redirections Reveal Network Conditions

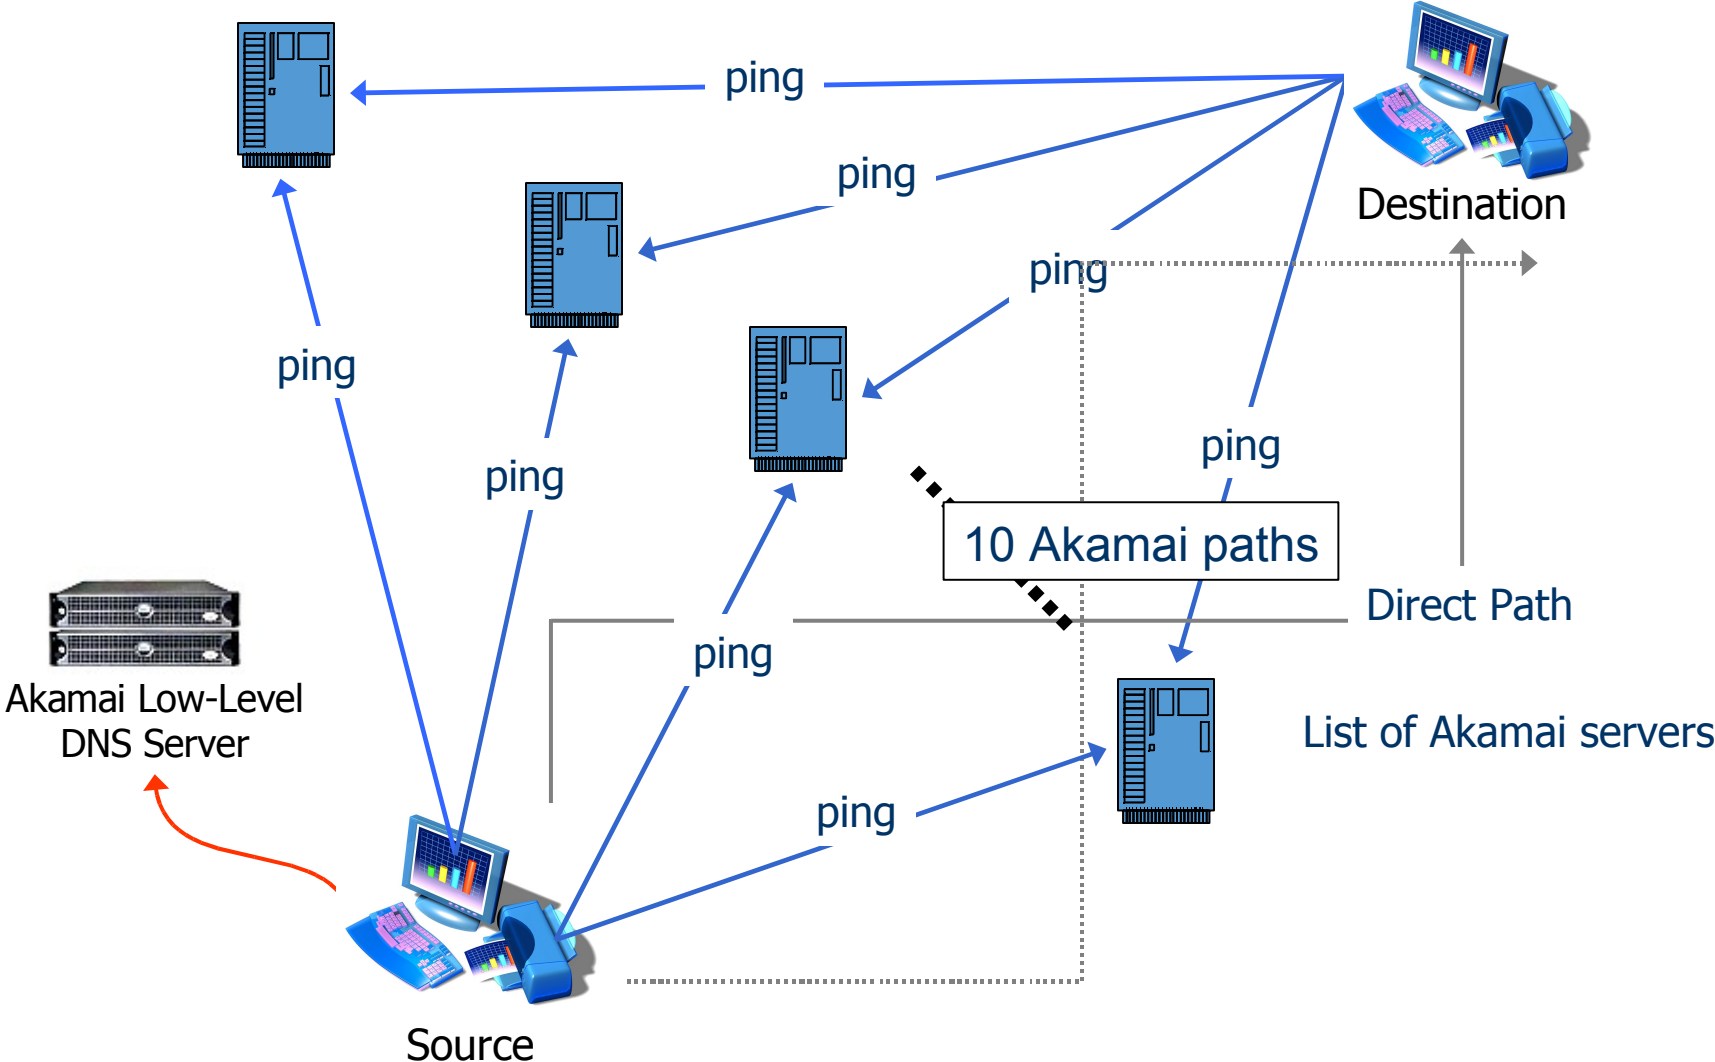


Akamai-Driven One-Hop Source Routing

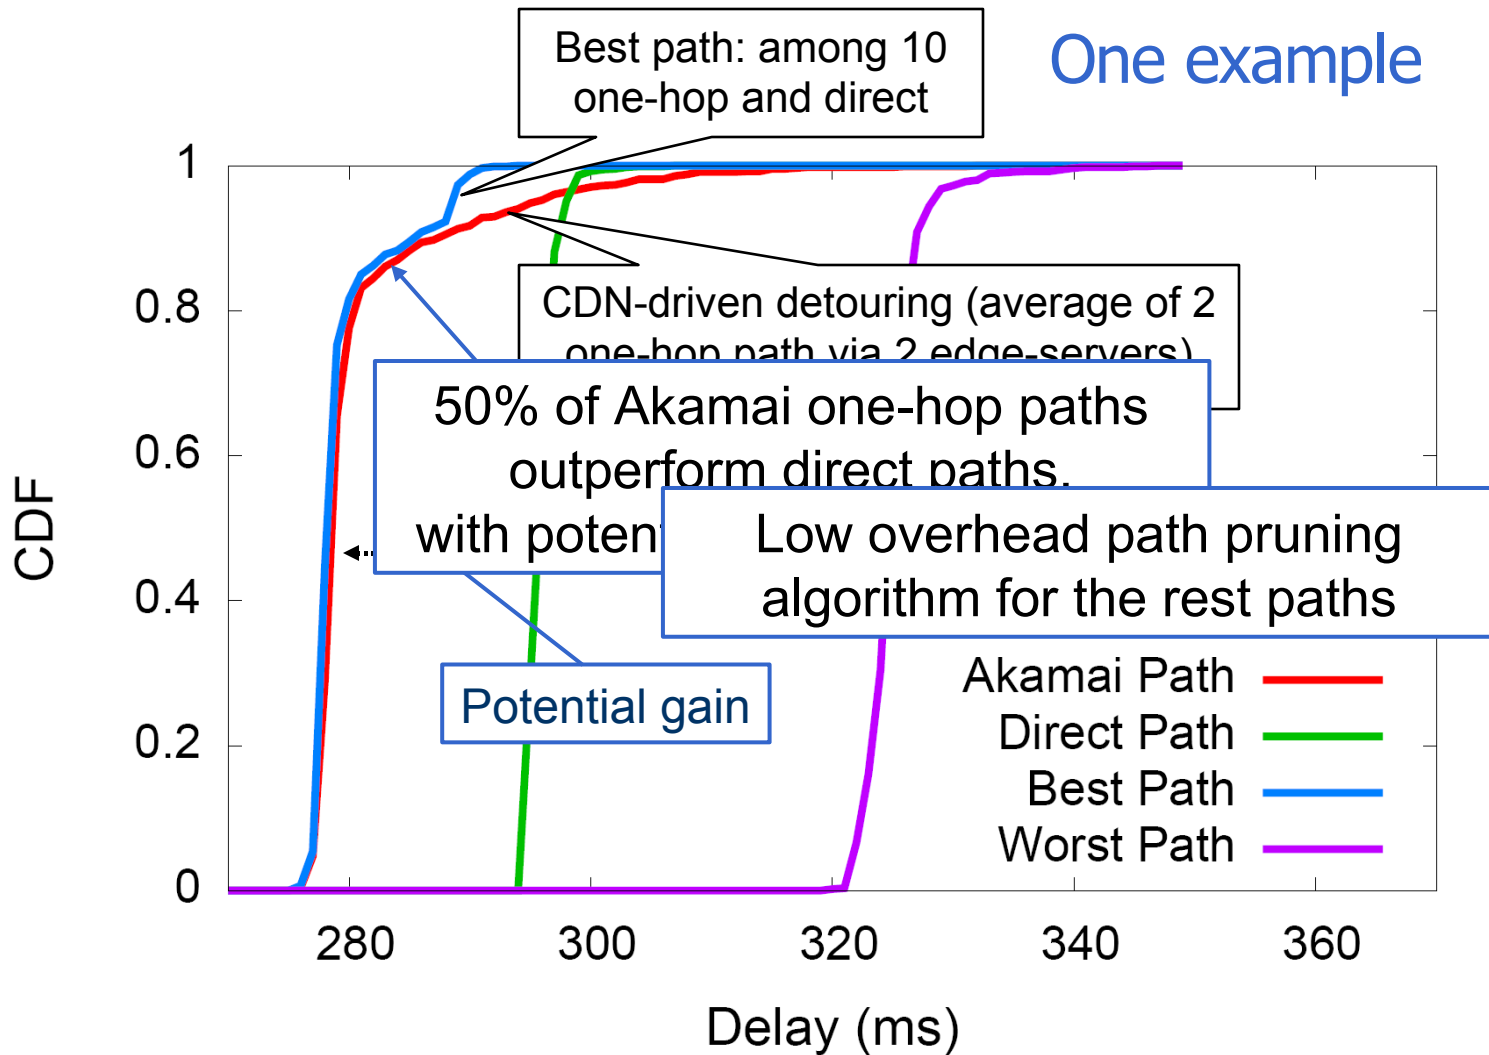


Methodology

Estimate end-to-end latency(rtt) by adding rtt of 2 path segments



Akamai-Driven Source Routing



Conclusions

- Akamai CDN gathers a fairly accurate view of the network
- This view can be reused by overlay networks
 - Significantly reducing the amount of measurements
 - Requiring no new infrastructure to be deployed

Let's draft behind Akamai!



“One more thing.....”

Ono – CDN-based Detouring in BitTorrent

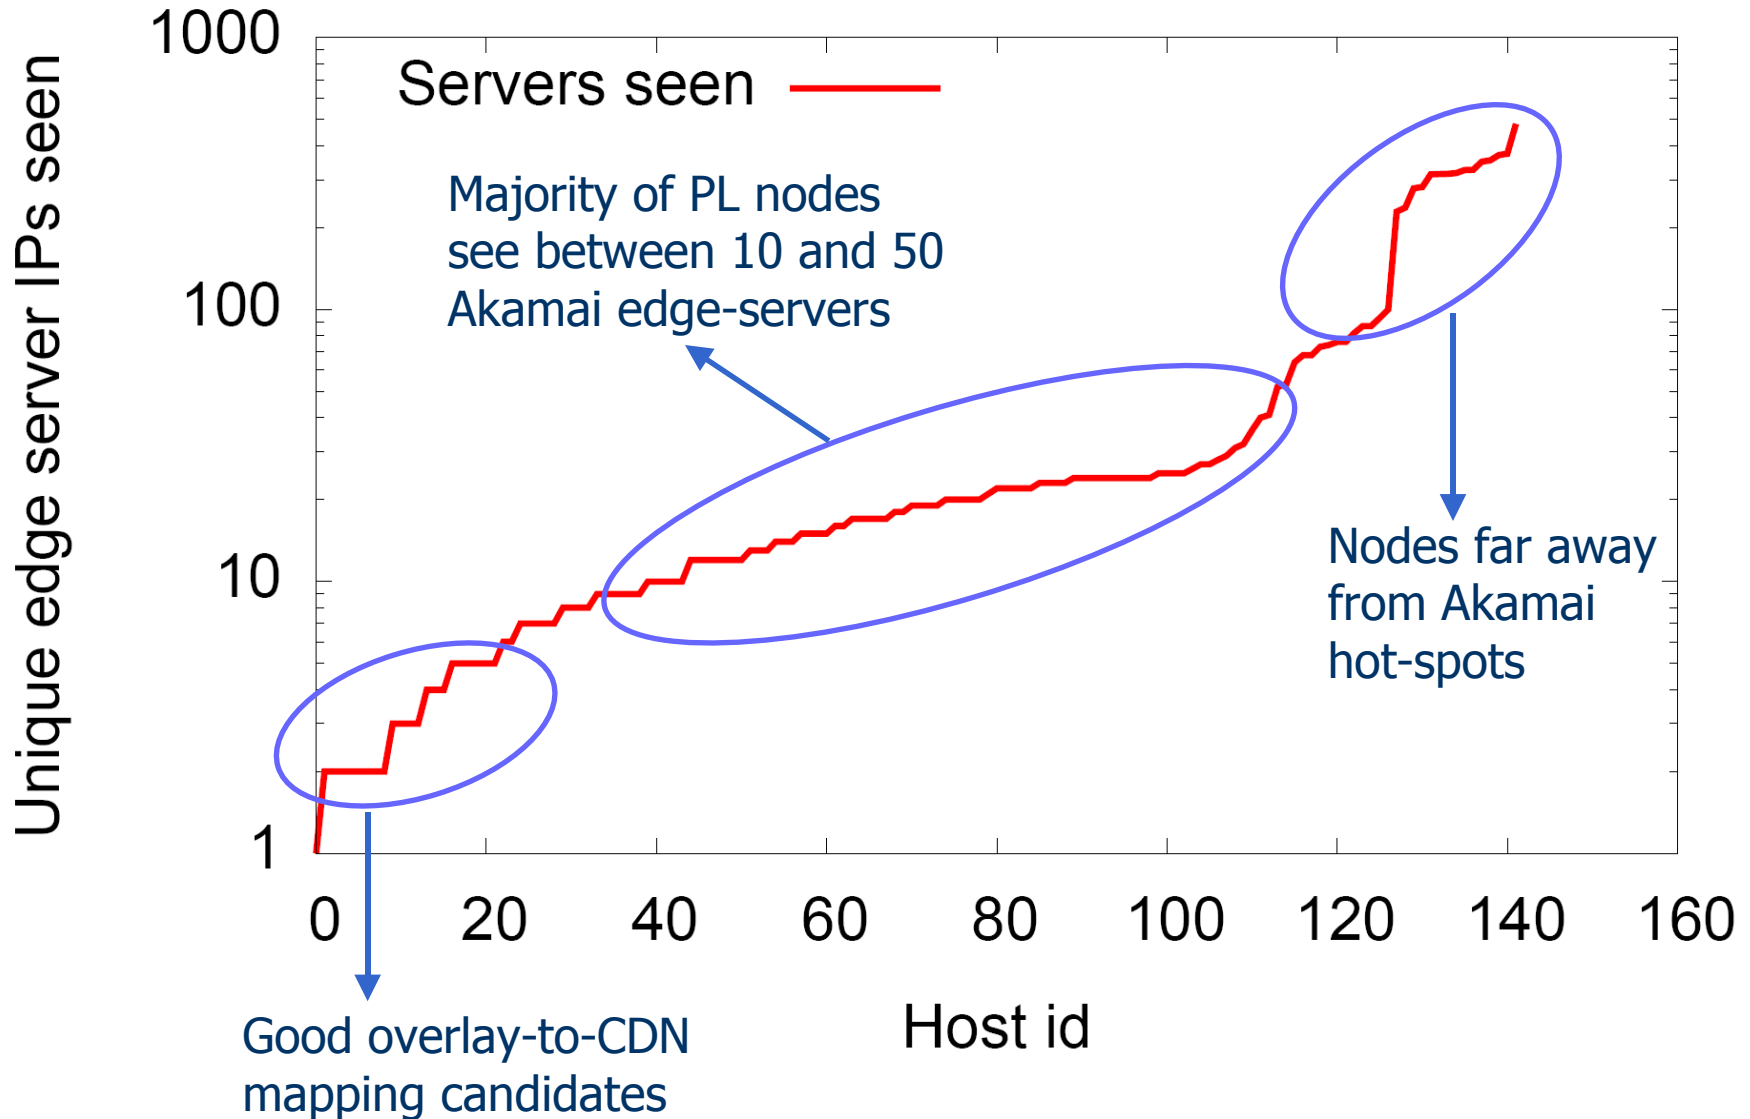
- An Azureus/BitTorrent plugin for you
- Locates quality Internet paths using low-cost DNS queries
- Enables Azureus clients to detour traffic through peers located along lower-latency (& potentially higher throughput) paths.



<http://www.aqualab.cs.northwestern.edu/projects/Ono.html>

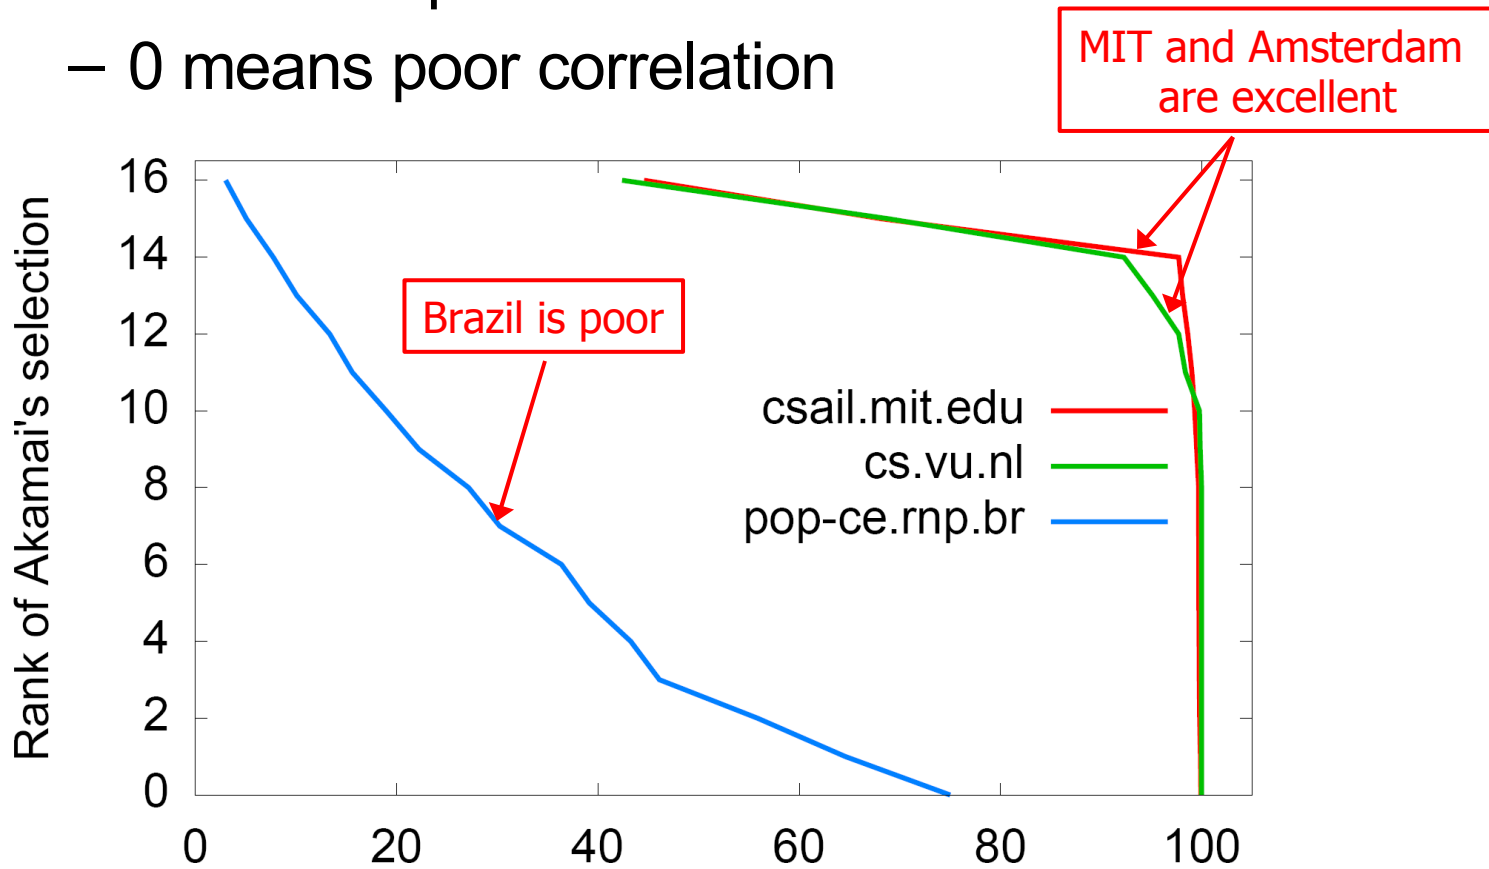
Backup Slides

Server Diversity for Yahoo



Do redirections reveal network conditions?

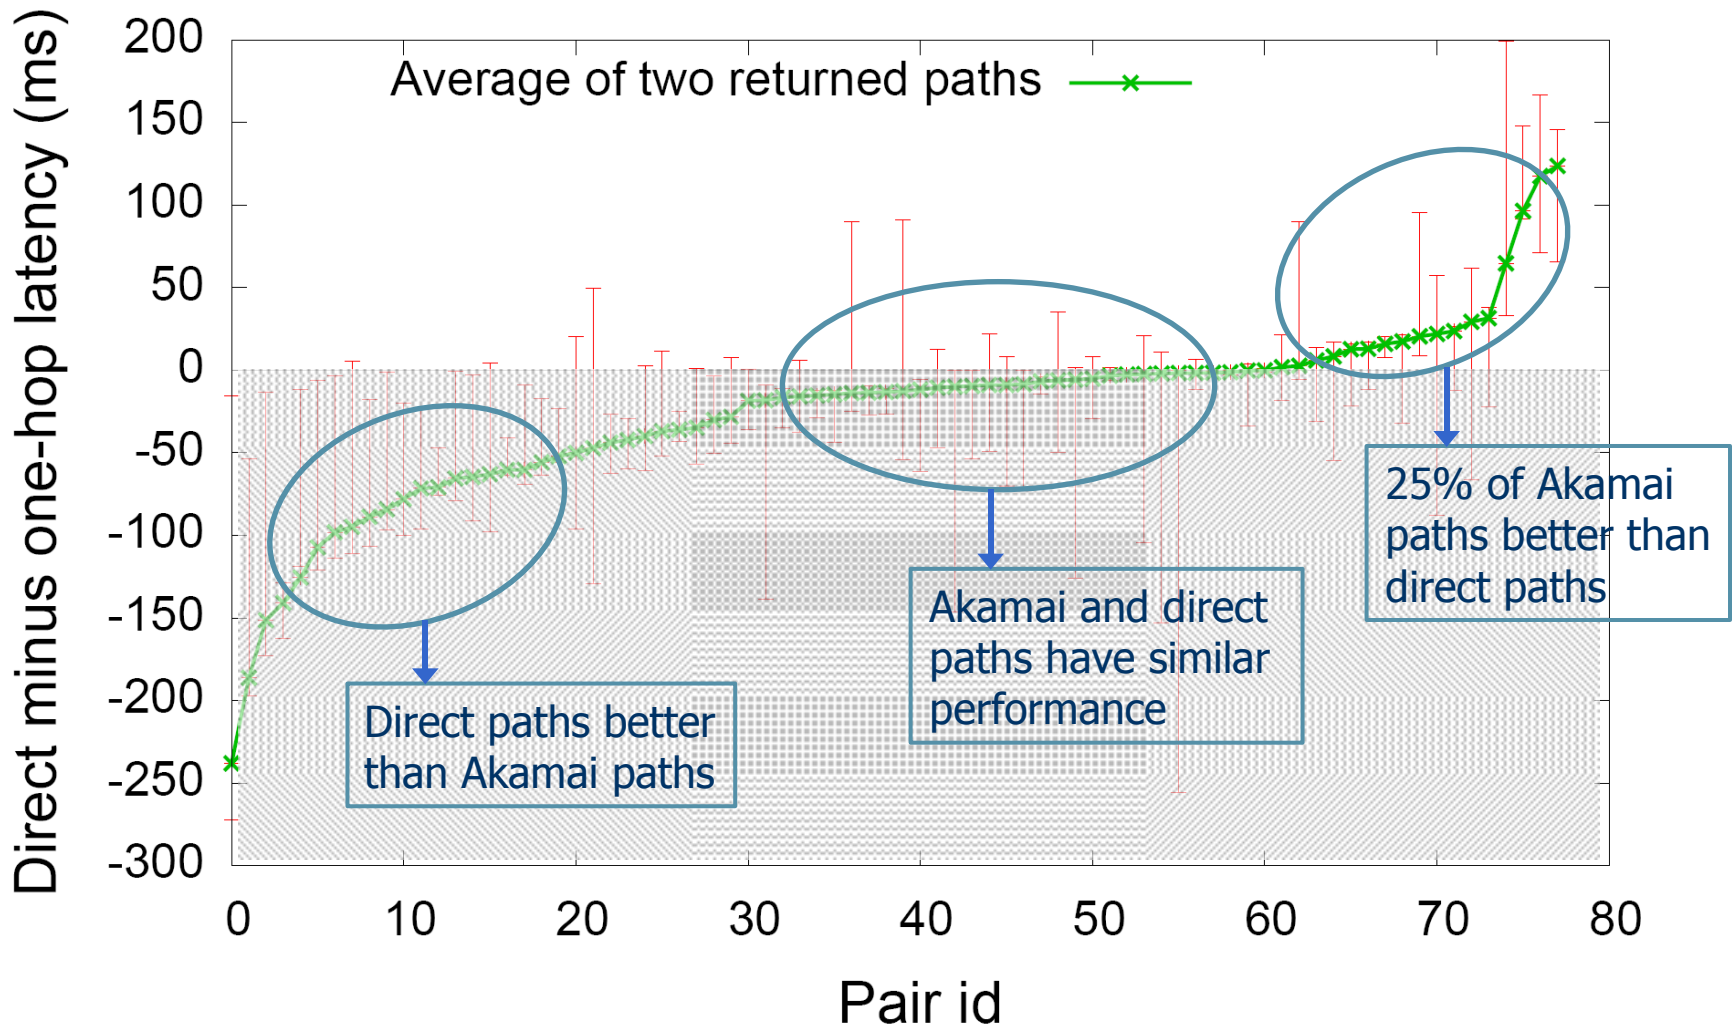
- Rank = $r_1 + r_2 - 1$
 - 16 means perfect correlation
 - 0 means poor correlation



Percentage of time Akamai's selection is better or equal to rank

Akamai-Driven Source Routing

Experiment: US (6), Europe (3), S. America (2), Asia (3)



Path Pruning Result

