MONTANA
STATE UNIVERSITY

Network measurement using Akamai's infrastructure

Mike P. Wittie

Overview

- Akamai has lots of servers close to users and lots of users close to servers
- Let's put their hands together (Of course we're not the first)

- Clever ways of using Akamai's infrastructure
  - Ping through CDN Proxies (**pcp**) [ICCCN'15]
  - Passive detection of cellular middleboxes [PAM'16]
  - Justifying mobile IPv6 content [Mobicom'16]

- Best practices for Web content delivery
  - Third-party Trailing Ratio (TPTR) [PAM'17]
  - Multiple connections of HTTP/2 [submission]

Overview

- Akamai has lots of servers close to users and lots of users close to servers
- Let's put their hands together (Of course we're not the first)

- Clever ways of using Akamai's infrastructure
  - Ping through CDN Proxies (**pcp**) [ICCCN'15]
  - Passive detection of cellular middleboxes [PAM'16]
  - Justifying mobile IPv6 content [Mobicom'16]

  } Network measurement

- Best practices for Web content delivery
  - Third-party Trailing Ratio (TPTR) [PAM'17]
  - Multiple connections of HTTP/2 [submission]

Overview

- Akamai has lots of servers close to users and lots of users close to servers
- Let's put their hands together (Of course we're not the first)

- Clever ways of using Akamai's infrastructure
  - Ping through CDN Proxies (**pcp**) [ICCCN'15]
  - Passive detection of cellular middleboxes [PAM'16]
  - Justifying mobile IPv6 content [Mobicom'16]

  } Network measurement

- Best practices for Web content delivery
  - Third-party Trailing Ratio (TPTR) [PAM'17]
  - Multiple connections of HTTP/2 [submission]

  } Web performance

Overview

- Akamai has lots of servers close to users and lots of users close to servers
- Let's put their hands together (Of course we're not the first)

- Clever ways of using Akamai's infrastructure
  - Ping through CDN Proxies (**pcp**) [ICCCN'15]
  - Passive detection of cellular middleboxes [PAM'16]
  - Justifying mobile IPv6 content [Mobicom'16]

Network measurement

- Best practices for Web content delivery
  - Third-party Trailing Ratio (TPTR) [PAM'17]
  - Multiple connections of HTTP/2 [submission]

Web performance

Credits

Utkarsh Goel →

Methods

- Real-User Monitoring (RUM)
  - Injects Javascript to small fraction of requests
  - Uses Navigation Timing API
    - DNS resolutions
    - TCP connection establishment time
    - Webpage load time (PLT)
- Server TCP logs
  - Latency to client
  - IP addresses (IPv4/IPv6)
  - Cellular ISP name from EdgeScape
- Dynatrace Synthetic Monitoring (formerly Gomez)
  - Desktop and mobile browsers around the world

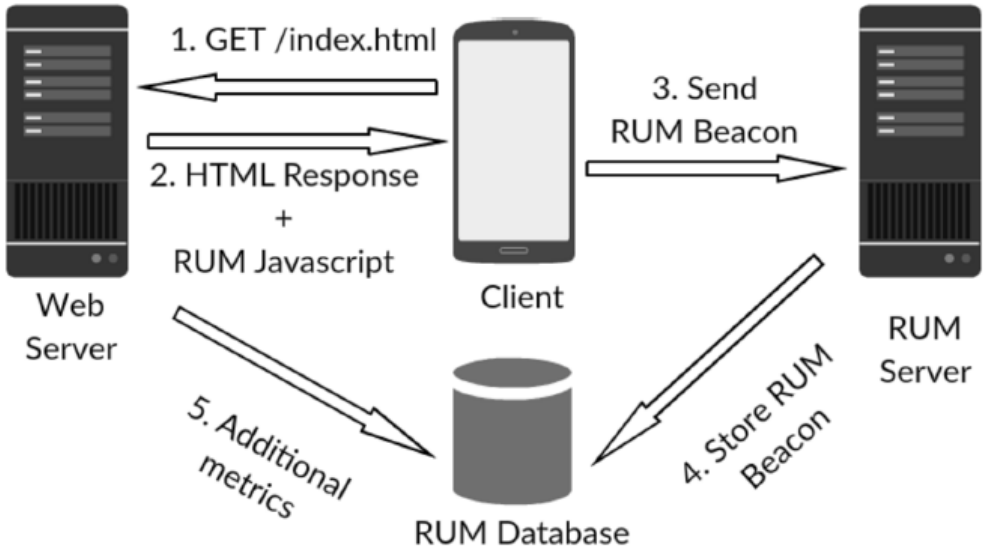

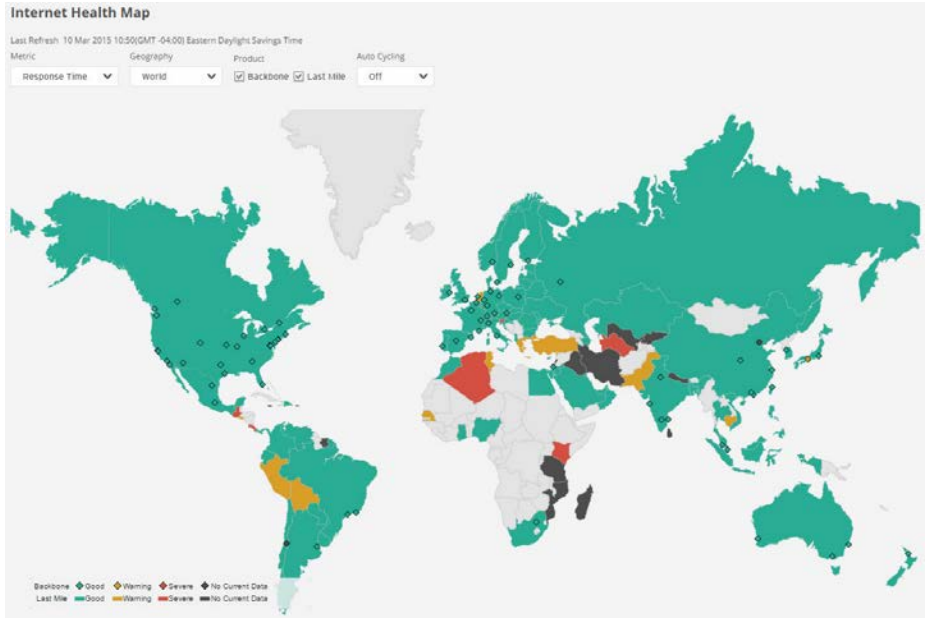

Latency prediction

- How can applications reduce user-perceived latency?

- Server selection
  - Find a server with the lowest latency to a given user
- Clustering
  - Find a group of users with low mutual latency

- Need a reliable, fast, and inexpensive method for latency prediction

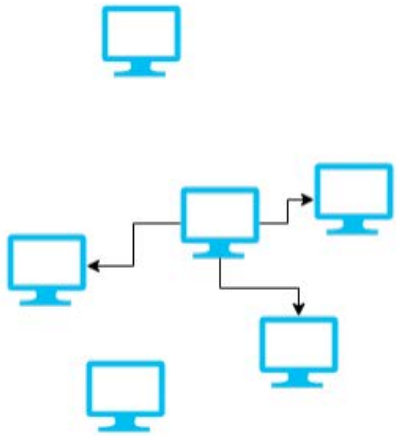

Samuel Micka, Utkarch Goel, Hanlu Ye, Mike P. Wittie, Brendan Mumey. "pcp: Internet Latency Estimation Using CDN Replicas" in International Conference on Computer Communications and Networks (ICCCN), August 2015.

Shortcomings of latency prediction tools

- ICMP ping
  - All to all communication
  - Slow and expensive
  - Often blocked by firewalls

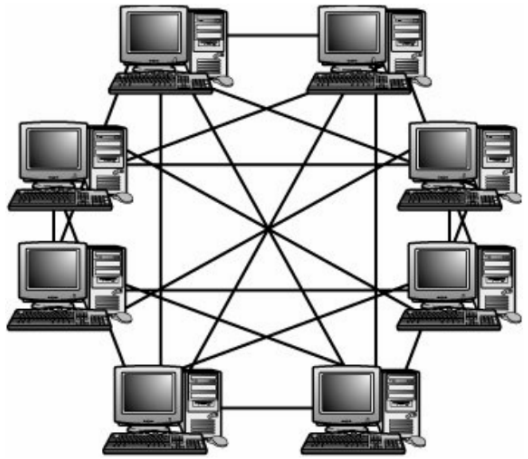

- IP to location databases
  - Locations inaccurate
  - Holes in coverage of IP space
  - Simplistic latency model

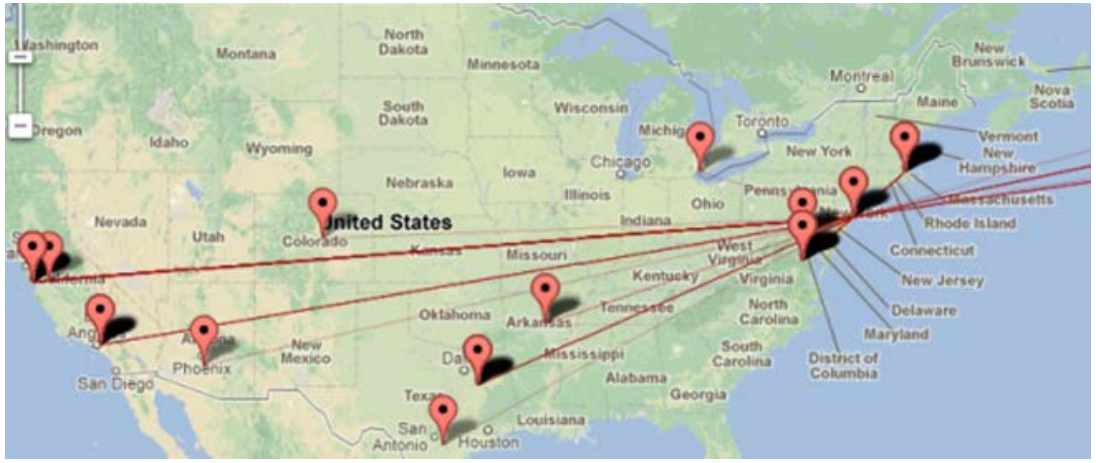

Samuel Micka, Utkarch Goel, Hanlu Ye, Mike P. Wittie, Brendan Mumey. "pcp: Internet Latency Estimation Using CDN Replicas" in International Conference on Computer Communications and Networks (ICCCN), August 2015.

Shortcomings of latency prediction tools

iPlane

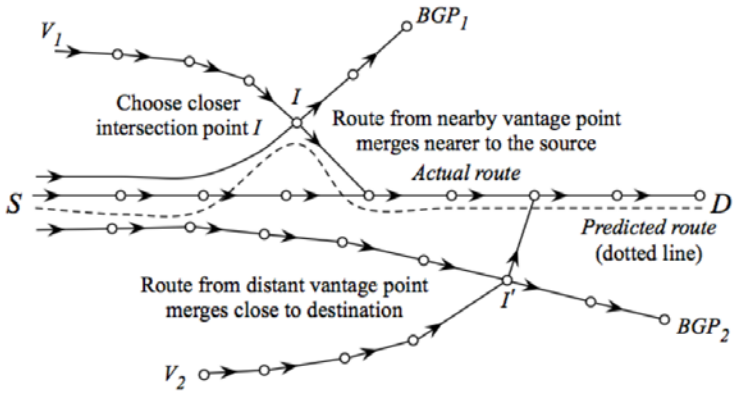

- Predicts latency in a virtual network build from traceroutes
- Measurements out of date
- Holes in the IP space

King

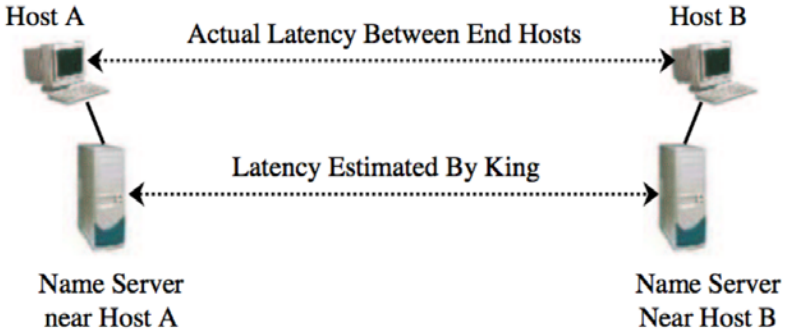

- Predicts P2P latency from latency between name servers
- Requires support for recursive DNS queries

CRP

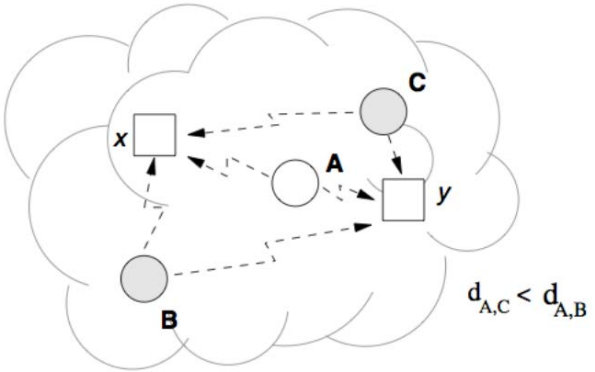

- Ranks node proximity based on similarity of DNS mapping
- Does not predict latency
- Cannot compare nodes without common CDN server mappings

Shortcomings of latency prediction tools

iPlane

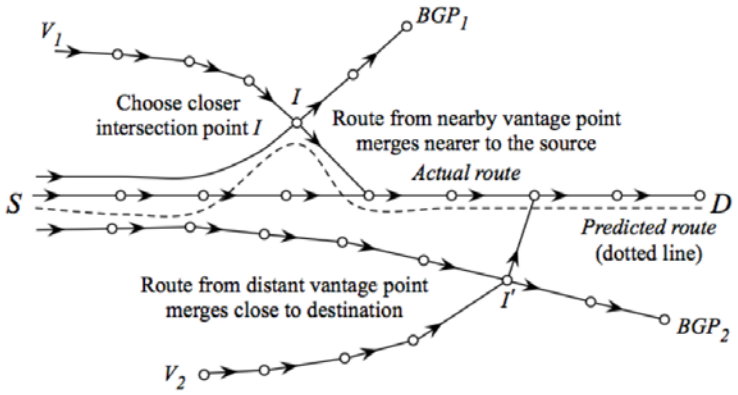

- Predicts latency in a virtual network build from traceroutes
- Measurements out of date
- Holes in the IP space

King

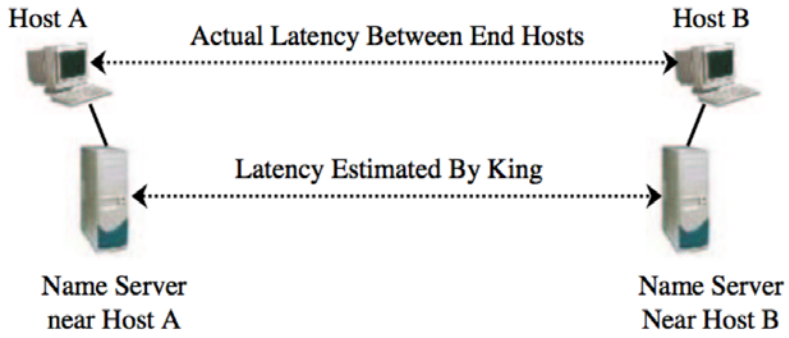

- Predicts P2P latency from latency between name servers
- Requires support for recursive DNS queries

CRP

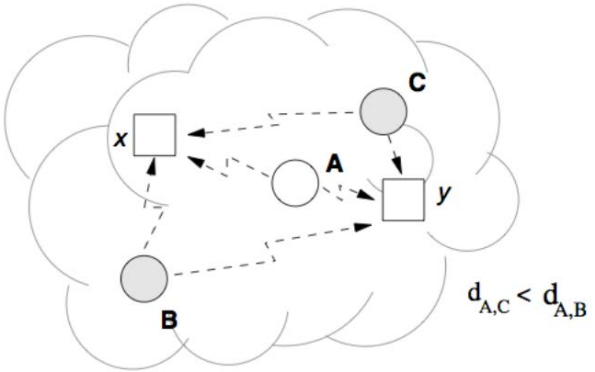

- Ranks node proximity based on similarity of DNS mapping
- Does not predict latency
- Cannot compare nodes without common CDN server mappings

Still need a reliable, fast, and inexpensive method for latency prediction

Ping through CDN Proxies (pcp)

- Goals
  - Accuracy/reliability
  - Speed
  - Scalability/low cost

- pcp
  - Clients observe RTTs to nearby CDN servers **during routine Web browsing**
  - pcp constructs a virtual topology based on reported RTTs
  - Latency between clients estimated based on shortest path in the virtual topology

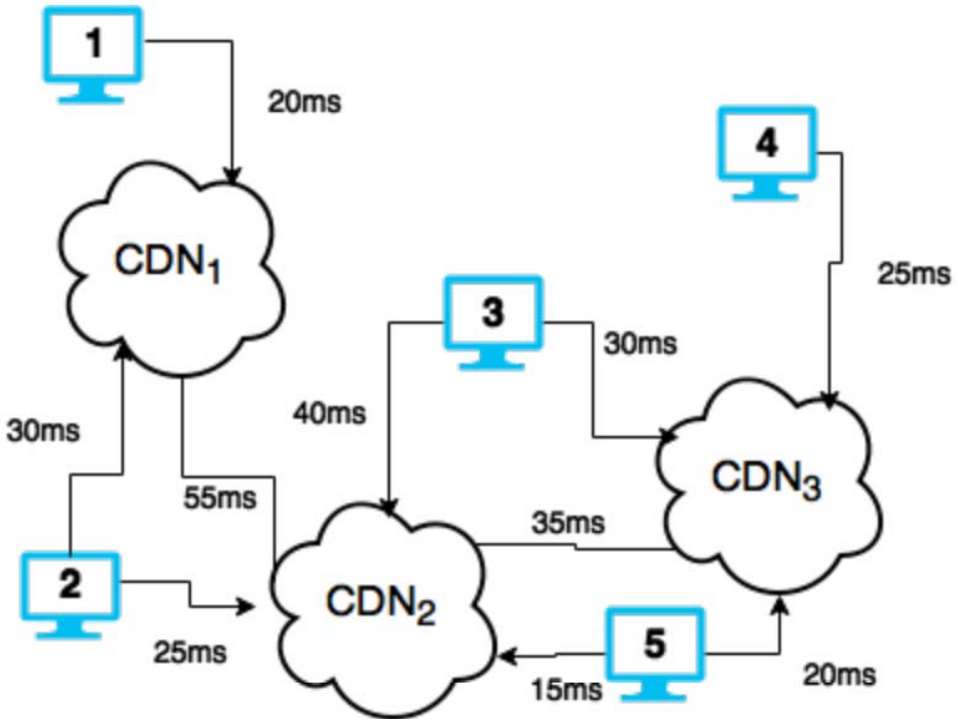

Samuel Micka, Utkarch Goel, Hanlu Ye, Mike P. Wittie, Brendan Mumey. "pcp: Internet Latency Estimation Using CDN Replicas" in International Conference on Computer Communications and Networks (ICCCN), August 2015.

Ping through CDN Proxies (pcp)

- Goals
  - Accuracy/reliability
  - Speed
  - Scalability/low cost

- pcp
  - Clients observe RTTs to nearby CDN servers **during routine Web browsing**
  - pcp constructs a virtual topology based on reported RTTs
  - Latency between clients estimated based on shortest path in the virtual topology

```
L(c1, c4) = L(c1, cdn1)
          + L(cdn1, cdn2) + L(cdn2, cdn3)
          + L(cdn3, c4)
```

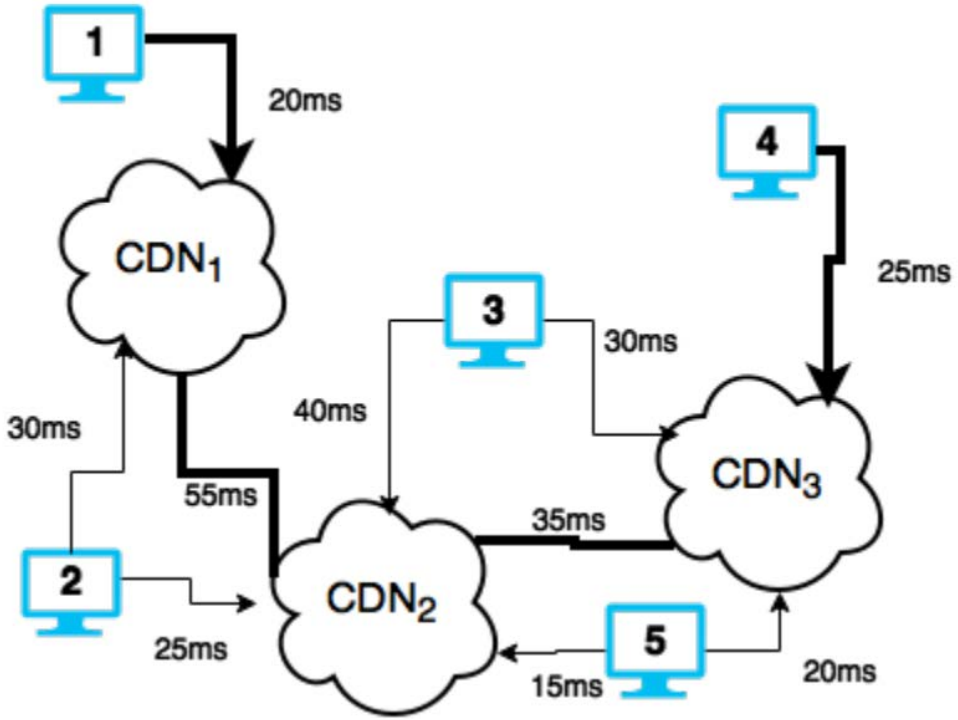

Samuel Micka, Utkarch Goel, Hanlu Ye, Mike P. Wittie, Brendan Mumey. "pcp: Internet Latency Estimation Using CDN Replicas" in International Conference on Computer Communications and Networks (ICCCN), August 2015.

Ping through CDN Proxies (pcp)

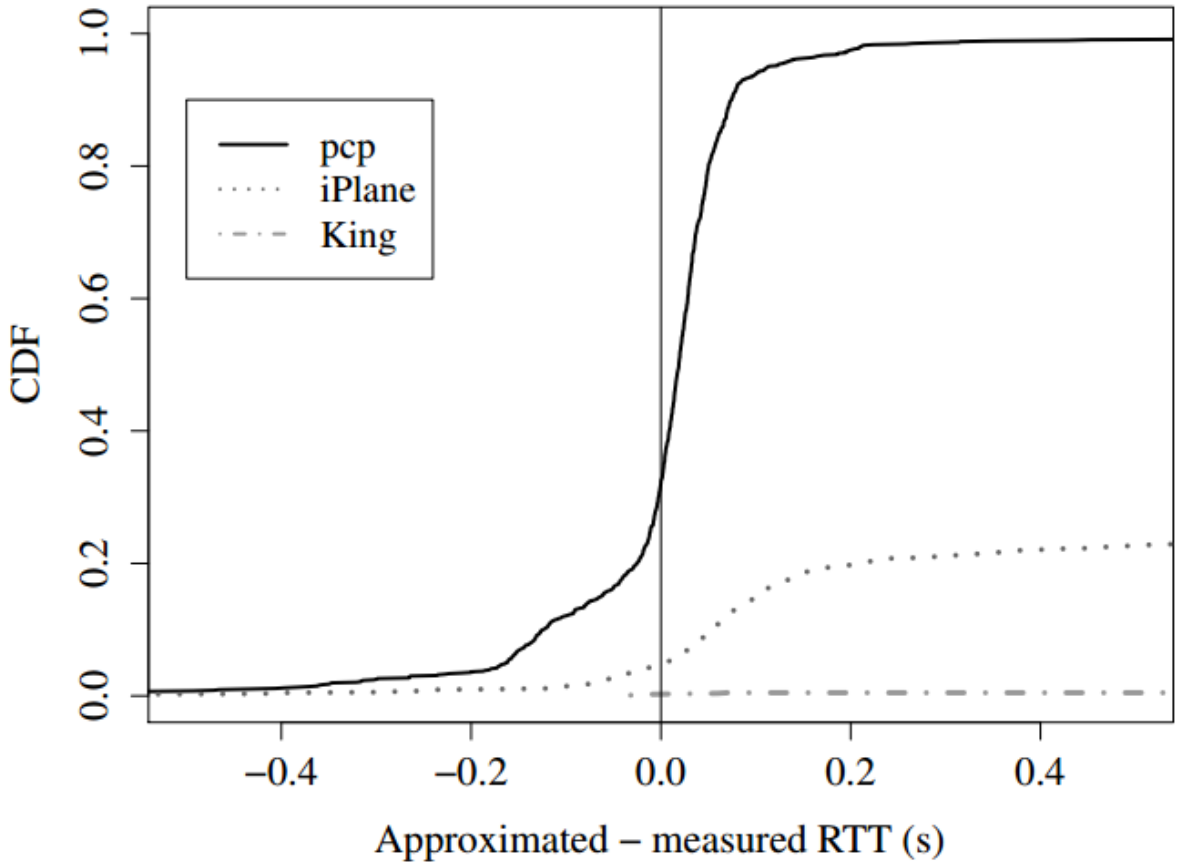

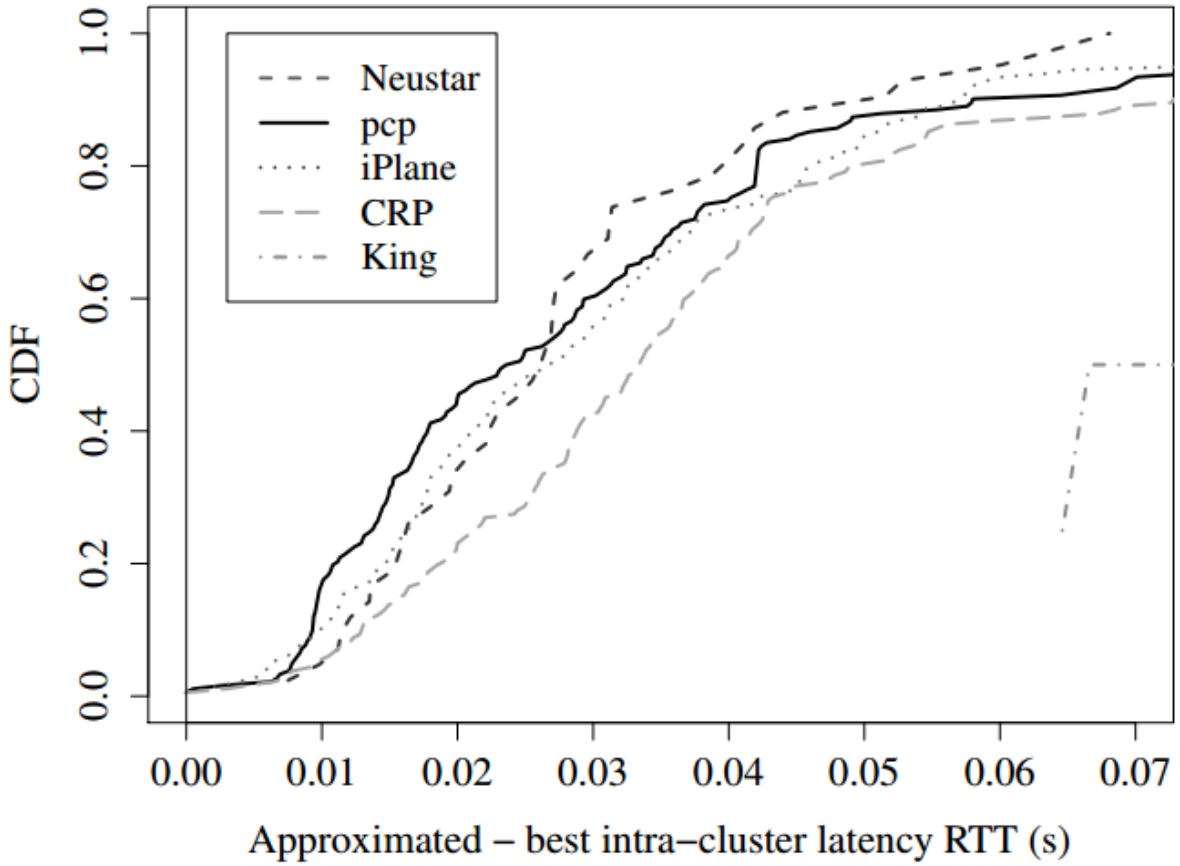

Samuel Micka, Utkarch Goel, Hanlu Ye, Mike P. Wittie, Brendan Mumey. "pcp: Internet Latency Estimation Using CDN Replicas" in International Conference on Computer Communications and Networks (ICCCN), August 2015.

Detecting Middle-boxes

- How can CDNs know if they are communicating with a client or a middlebox?

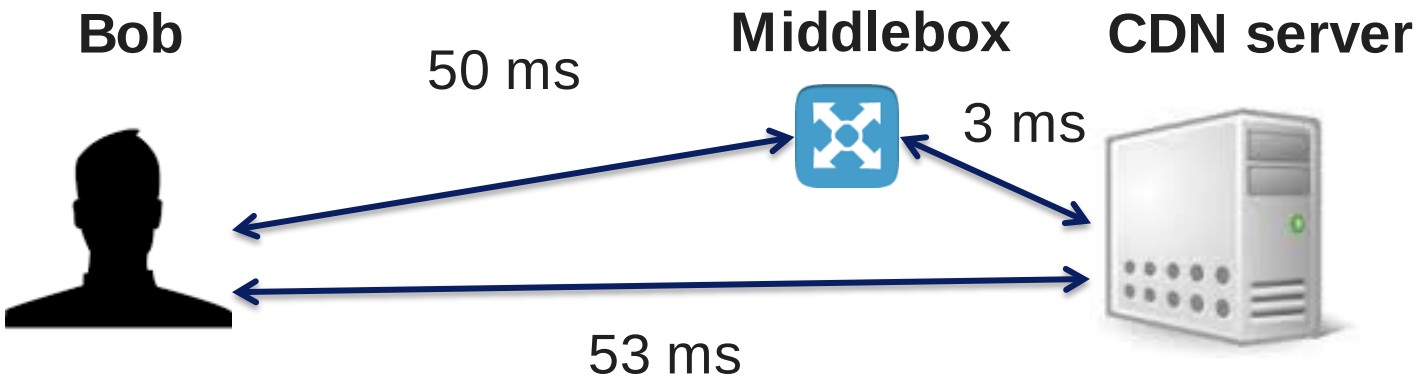

- Compare **latency** seen by servers and clients for both HTTP and HTTPS sessions.
- Compare **packet loss** seen on connections with and without middleboxes, only from the server TCP logs.
- Compare **TCP SYN characteristics** observed for port 80 and 443.

Utkarsh Goel, Moritz Steiner, Mike P. Wittie, Martin Flack, and Stephen Ludin. Detecting Cellular Middleboxes using Passive Measurement Techniques. *in ACM Passive and Active Measurements Conference (PAM) 2016.*

Results

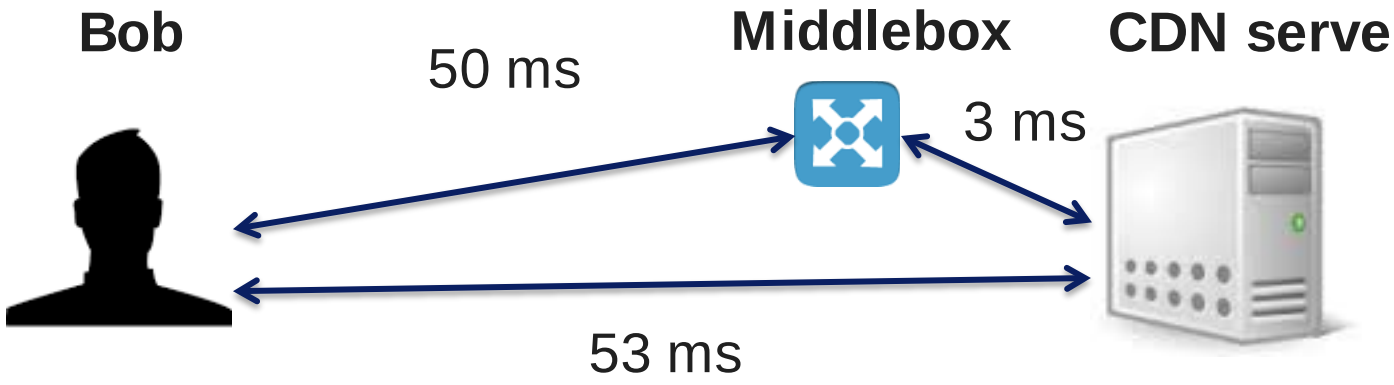

CC	Carrier	Protocol	Hits	Client RTT			Server RTT			Proxy?
				p25	p50	p75	p25	p50	p75	
US	AT&T	HTTP	1.7M	37	47	67	3	4	8	✓
US	AT&T	HTTPS	686K	45	60	89	52	75	114	X
US	Verizon W.	HTTP	1.9M	36	45	69	5	10	21	✓
US	Verizon W.	HTTPS	471K	44	60	87	48	65	87	X
US	T-Mobile	HTTP	2.1M	40	59	85	19	68	157	Limited
US	T-Mobile	HTTPS	459K	45	65	98	59	94	180	–

Utkarsh Goel, Moritz Steiner, Mike P. Wittie, Martin Flack, and Stephen Ludin. Detecting Cellular Middleboxes using Passive Measurement Techniques. *in ACM Passive and Active Measurements Conference (PAM) 2016.*

Results

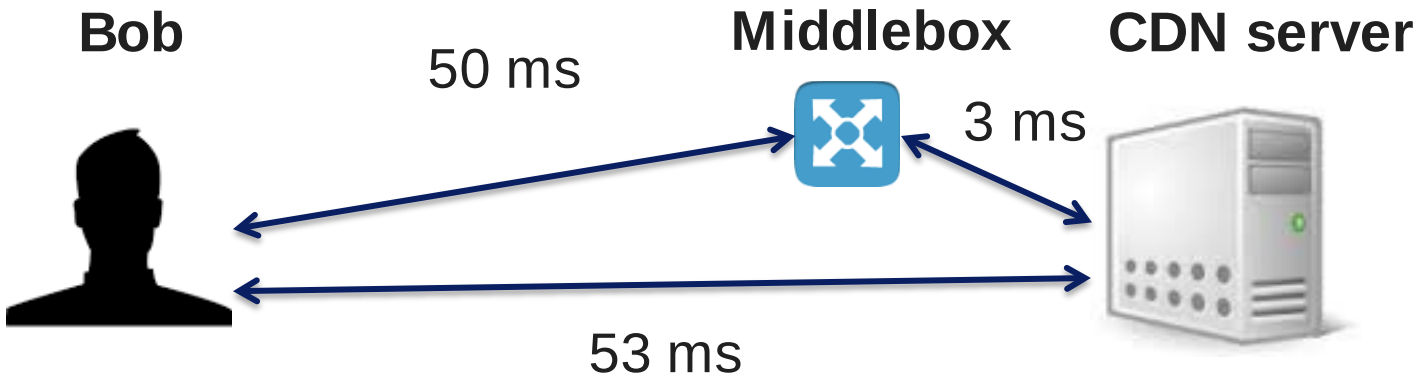

State	Domain Type	Client RTT			Server RTT			Proxy?
		p25	p50	p75	p25	p50	p75	
CA	Clothing website	37	51	75	2	3	3	✓
CA	e-Commerce website	40	56	80	2	2	3	✓
CA	Health Care website	40	56	90	40	80	175	X
CA	Ticketing website	37	49	65	43	93	186	X

Utkarsh Goel, Moritz Steiner, Mike P. Wittie, Martin Flack, and Stephen Ludin. Detecting Cellular Middleboxes using Passive Measurement Techniques. *in ACM Passive and Active Measurements Conference (PAM) 2016.*

Results

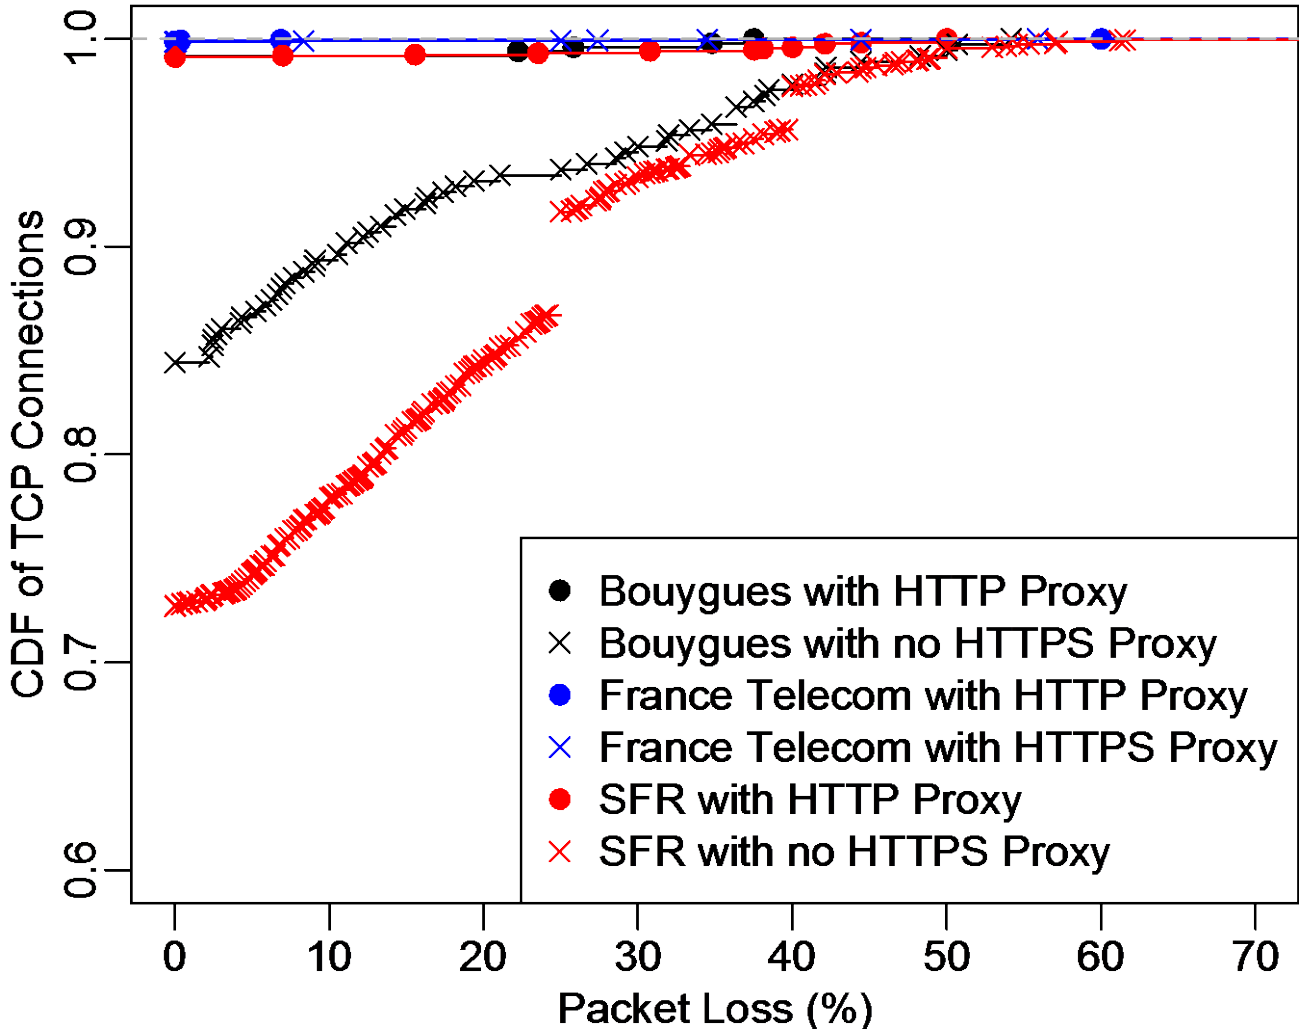

Utkarsh Goel, Moritz Steiner, Mike P. Wittie, Martin Flack, and Stephen Ludin. Detecting Cellular Middleboxes using Passive Measurement Techniques. *in ACM Passive and Active Measurements Conference (PAM) 2016.*

Results

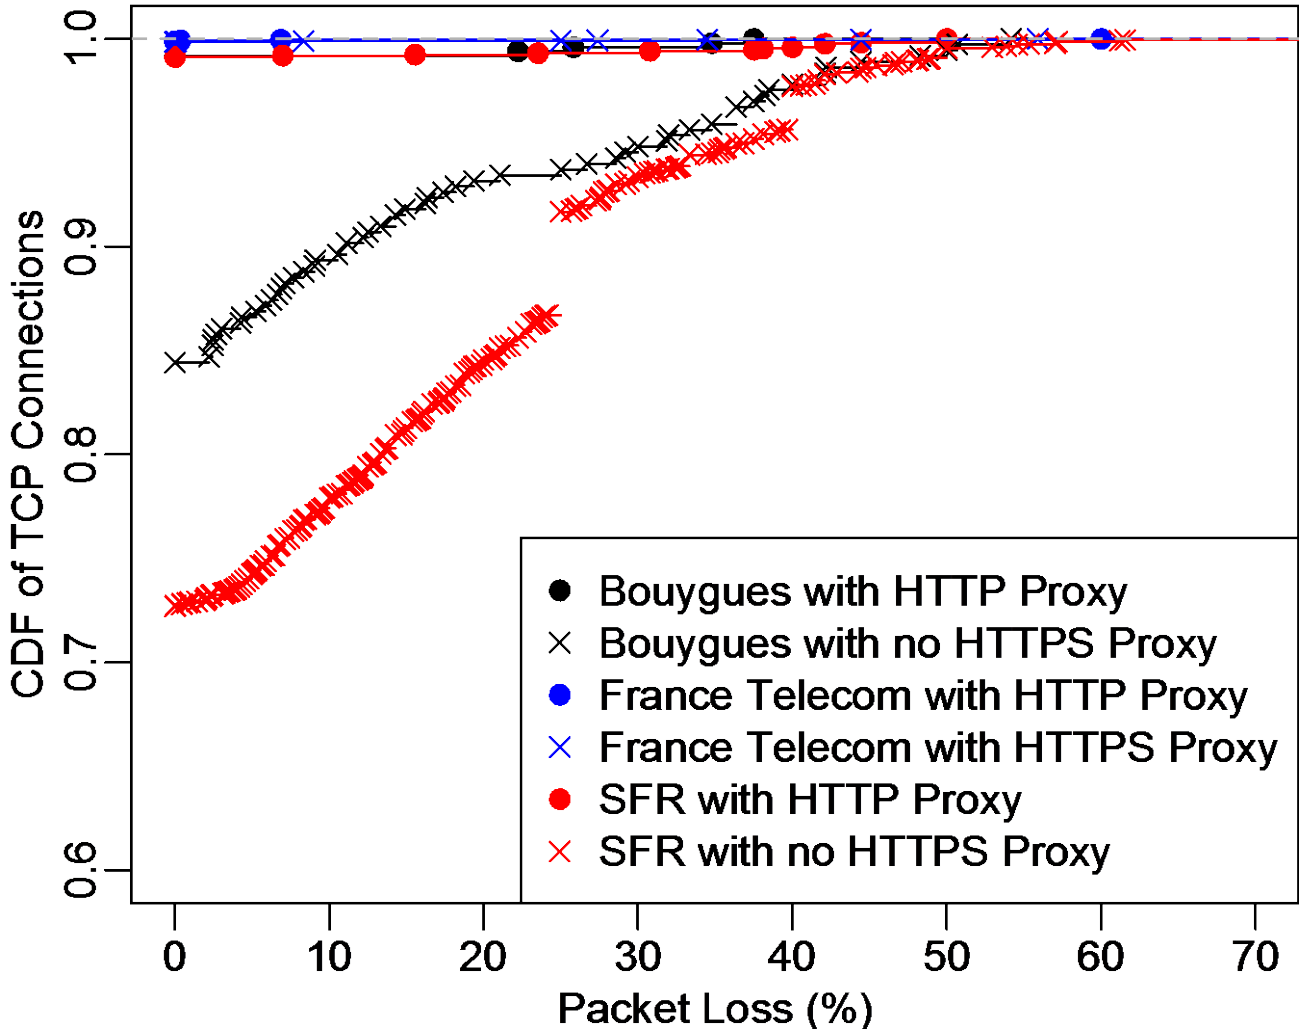

TCP SYN Characteristics of Cellular Proxies differ from mobile devices

- Initial Congestion Window
- Maximum Segment Size
- TCP Timestamp in TCP Options header

Utkarsh Goel, Moritz Steiner, Mike P. Wittie, Martin Flack, and Stephen Ludin. Detecting Cellular Middleboxes using Passive Measurement Techniques. *in ACM Passive and Active Measurements Conference (PAM) 2016.*

Should mobile Web content use IPv6

MONTANA STATE UNIVERSITY

IPv6 paths in cellular networks:

T-Mobile

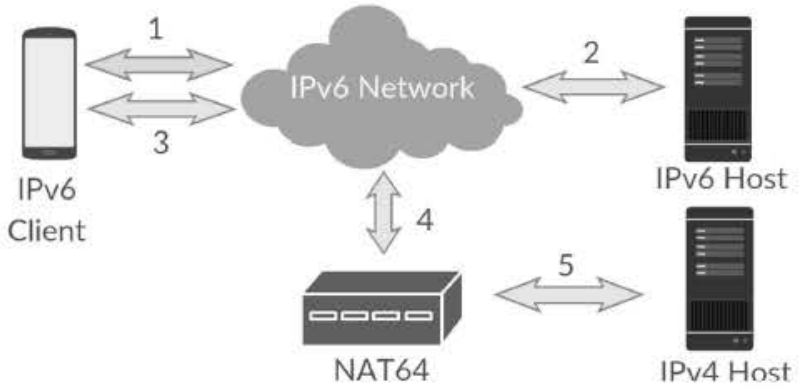

Verizon

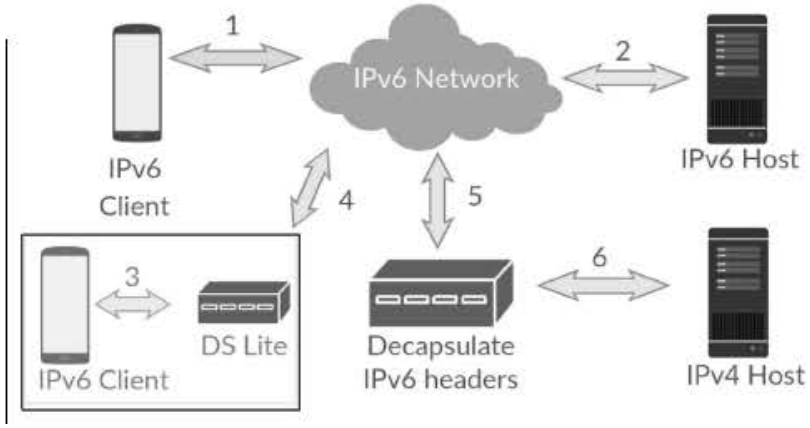

AT&T and Sprint

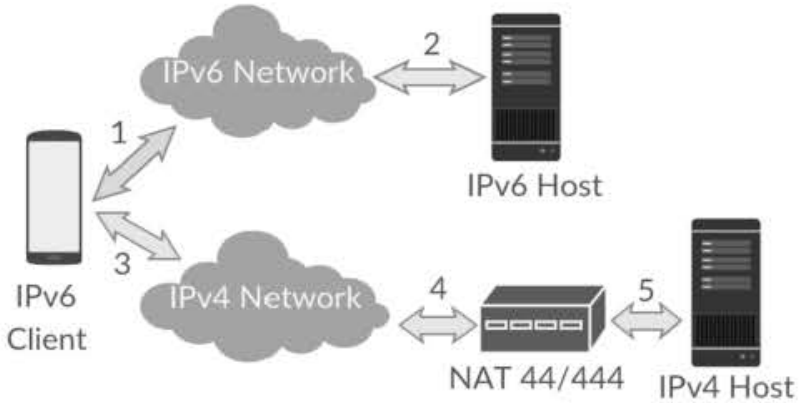

Utkarsh Goel, Moritz Steiner, Mike P. Wittie, Martin Flack, and Stephen Ludin. A Case for Faster Mobile Web in Cellular IPv6 Networks *in ACM Conference on Mobile Computing and Networking (MobiCom) 2016.*

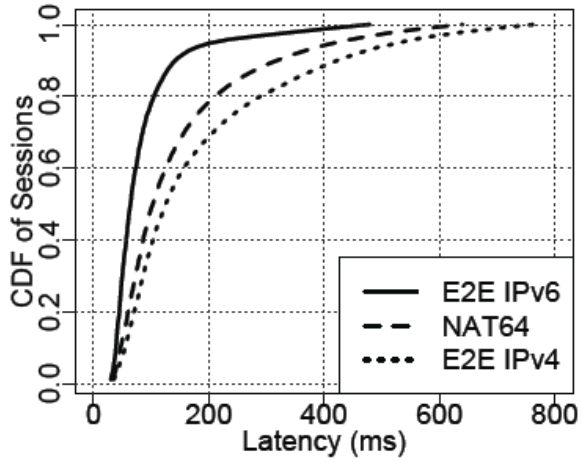

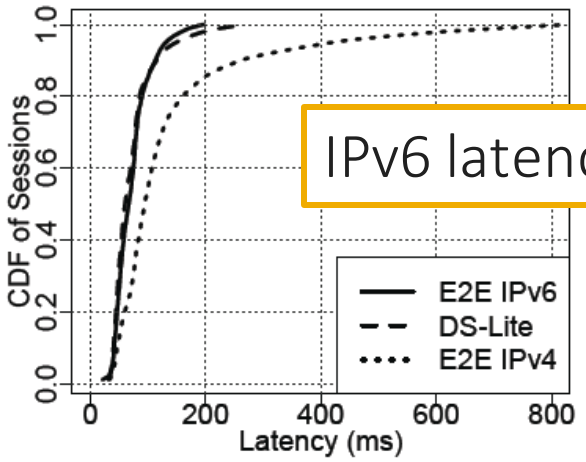

IPv6 latency is faster

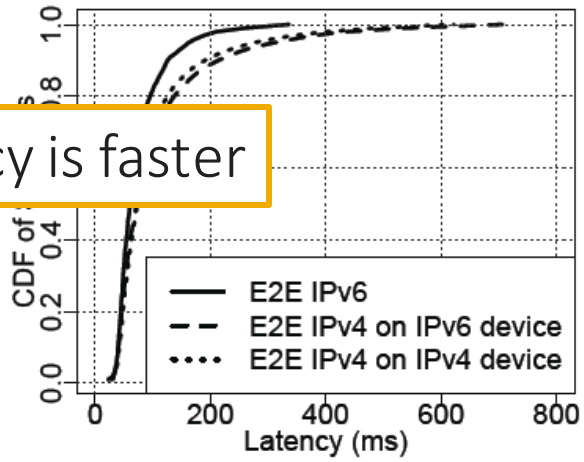

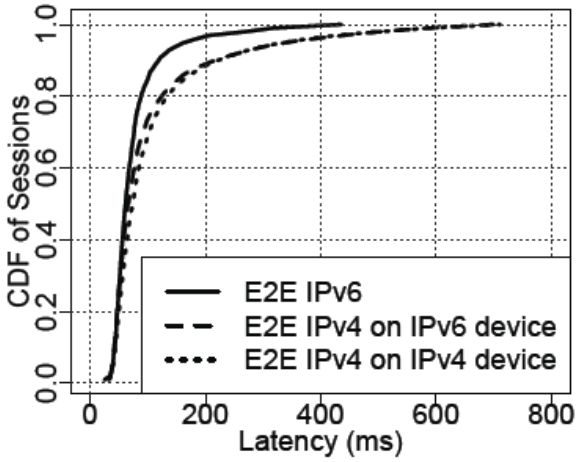

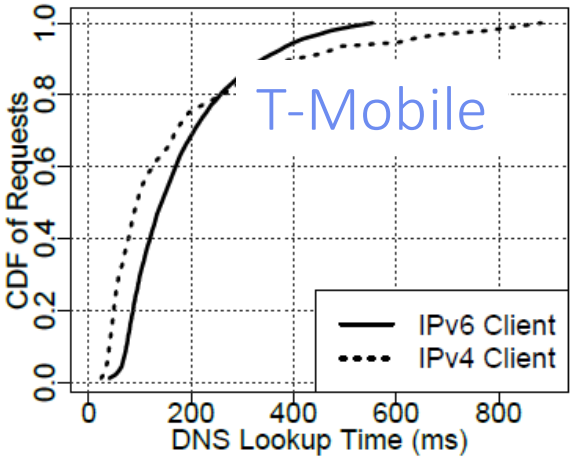

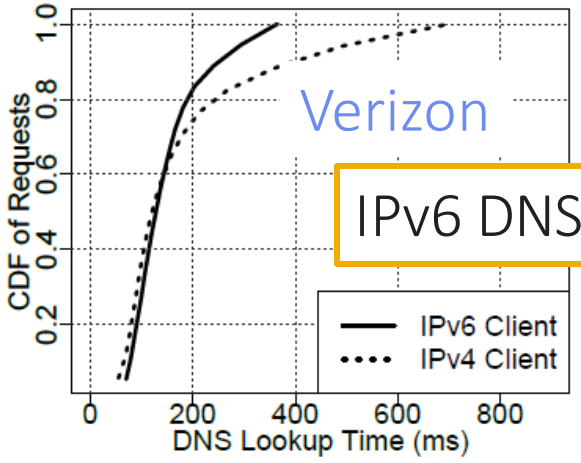

IPv6 DNS is slower

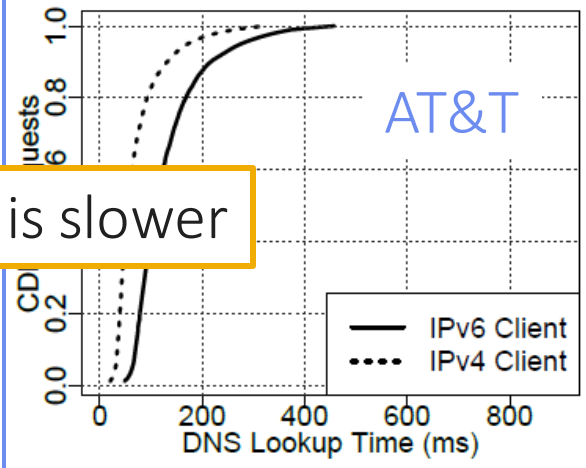

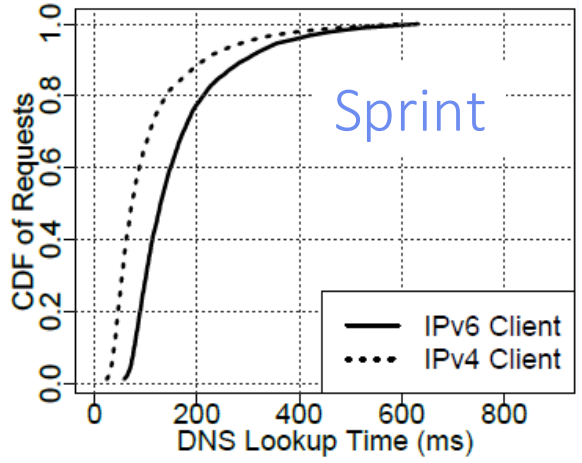

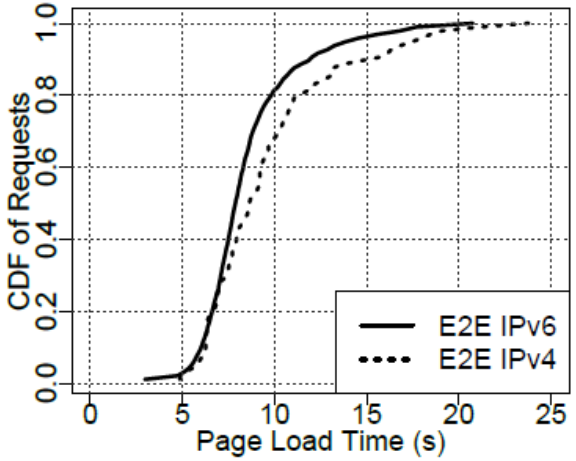

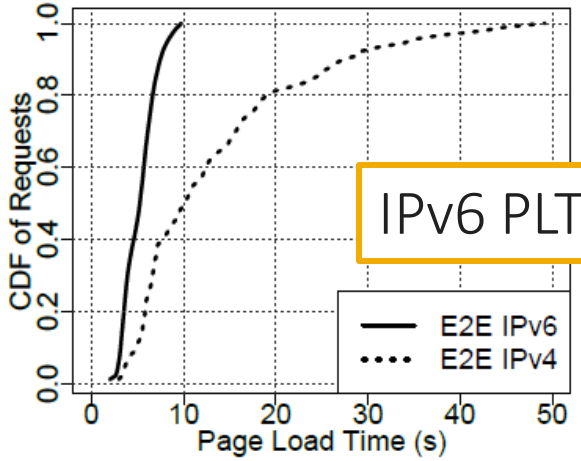

IPv6 PLT is faster

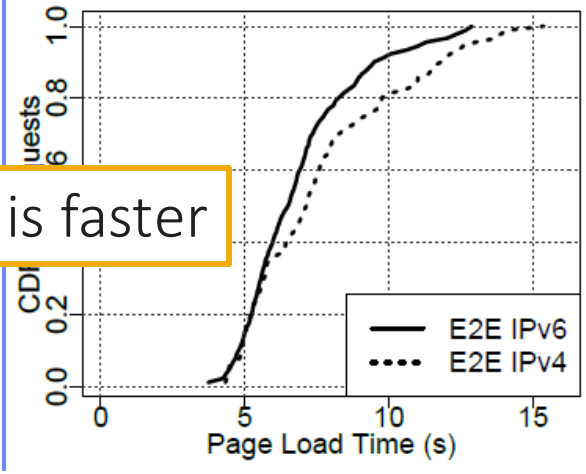

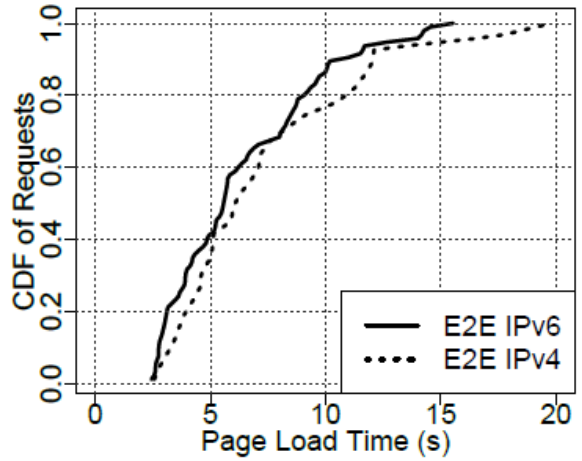

Smarter DNS Infrastructure for IPv6 requests

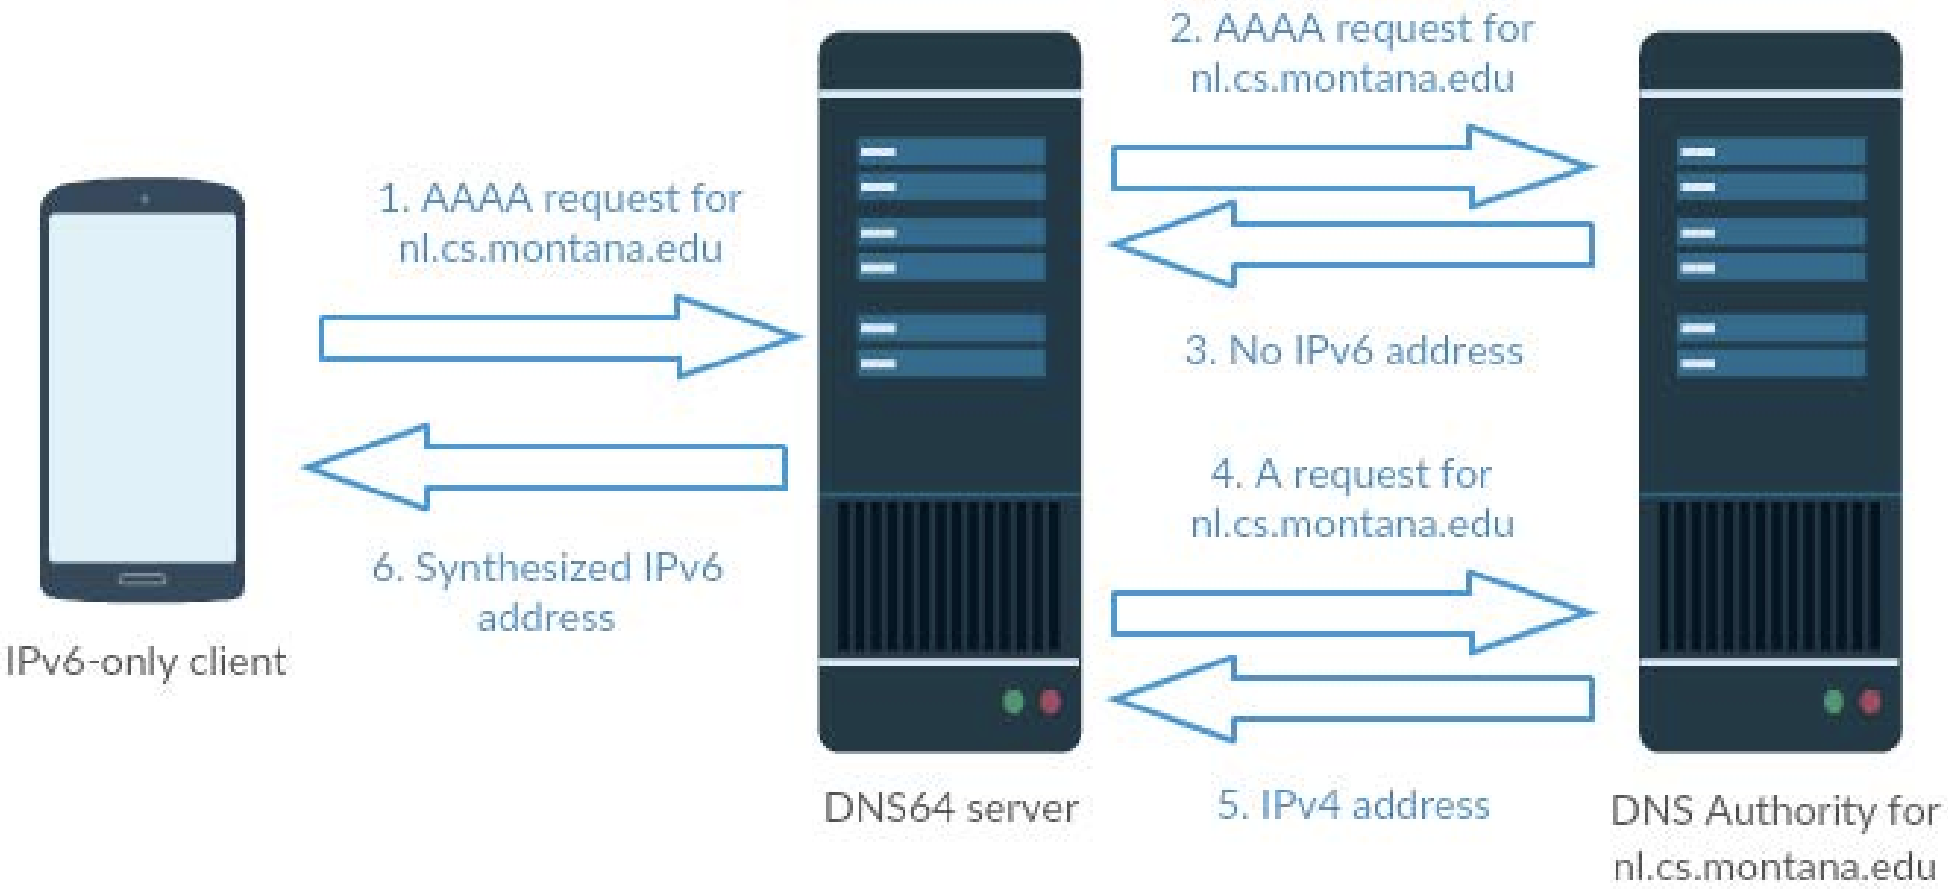

- Eliminate steps 4 and 5
- Send synthetic IPv6 address from the Authority in step 3.

Utkarsh Goel, Moritz Steiner, Mike P. Wittie, Martin Flack, and Stephen Ludin. A Case for Faster Mobile Web in Cellular IPv6 Networks *in ACM Conference on Mobile Computing and Networking (MobiCom) 2016.*