A Server-to-Server View of the Internet


Bala


Georgios


Arthur


Matthew


KC

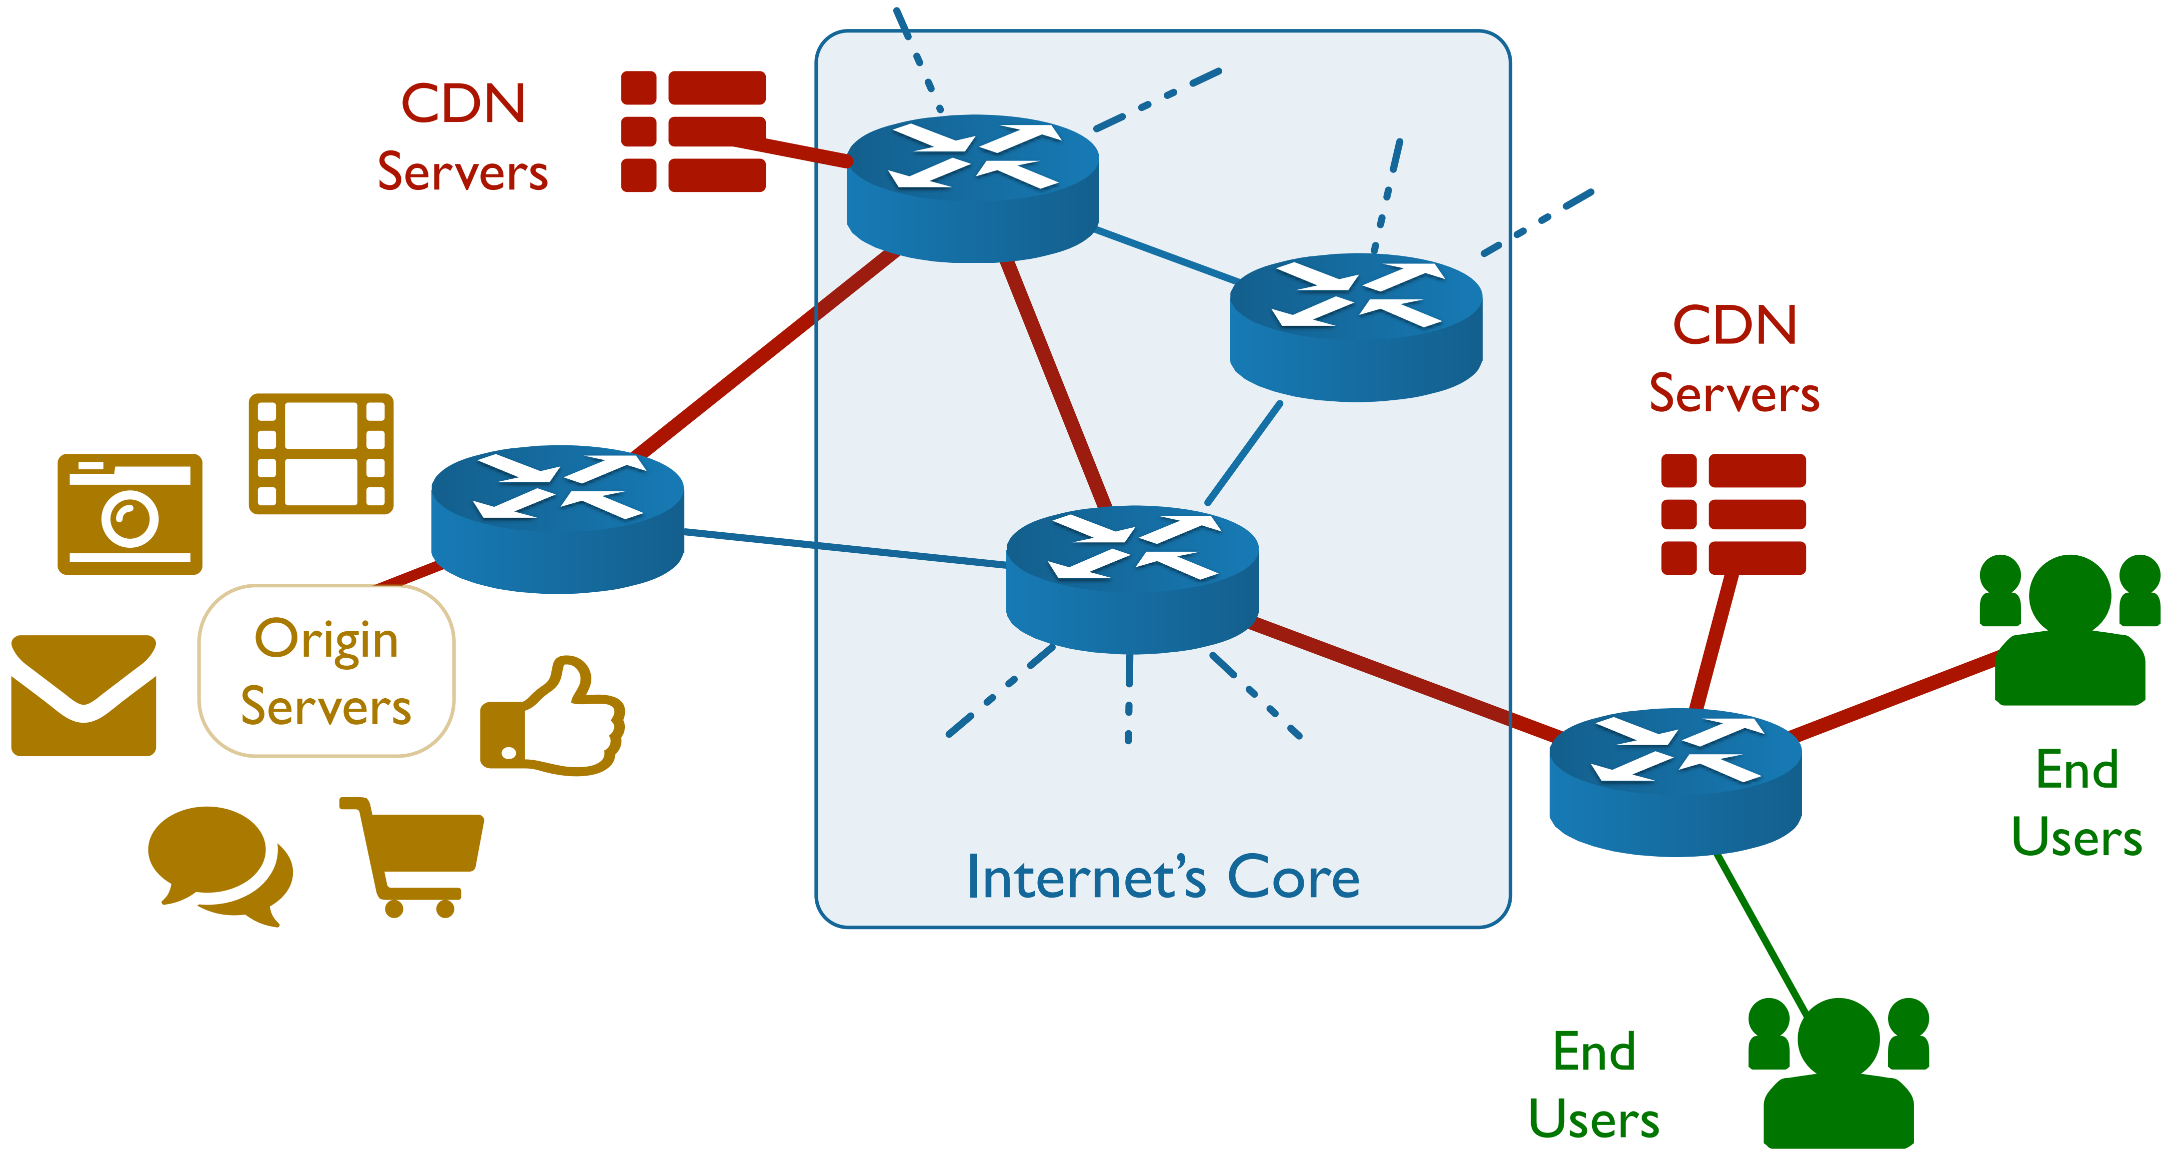

Content moved closer to end users to reduce latency.

Connections from end users are terminated at CDN servers close to the end users.

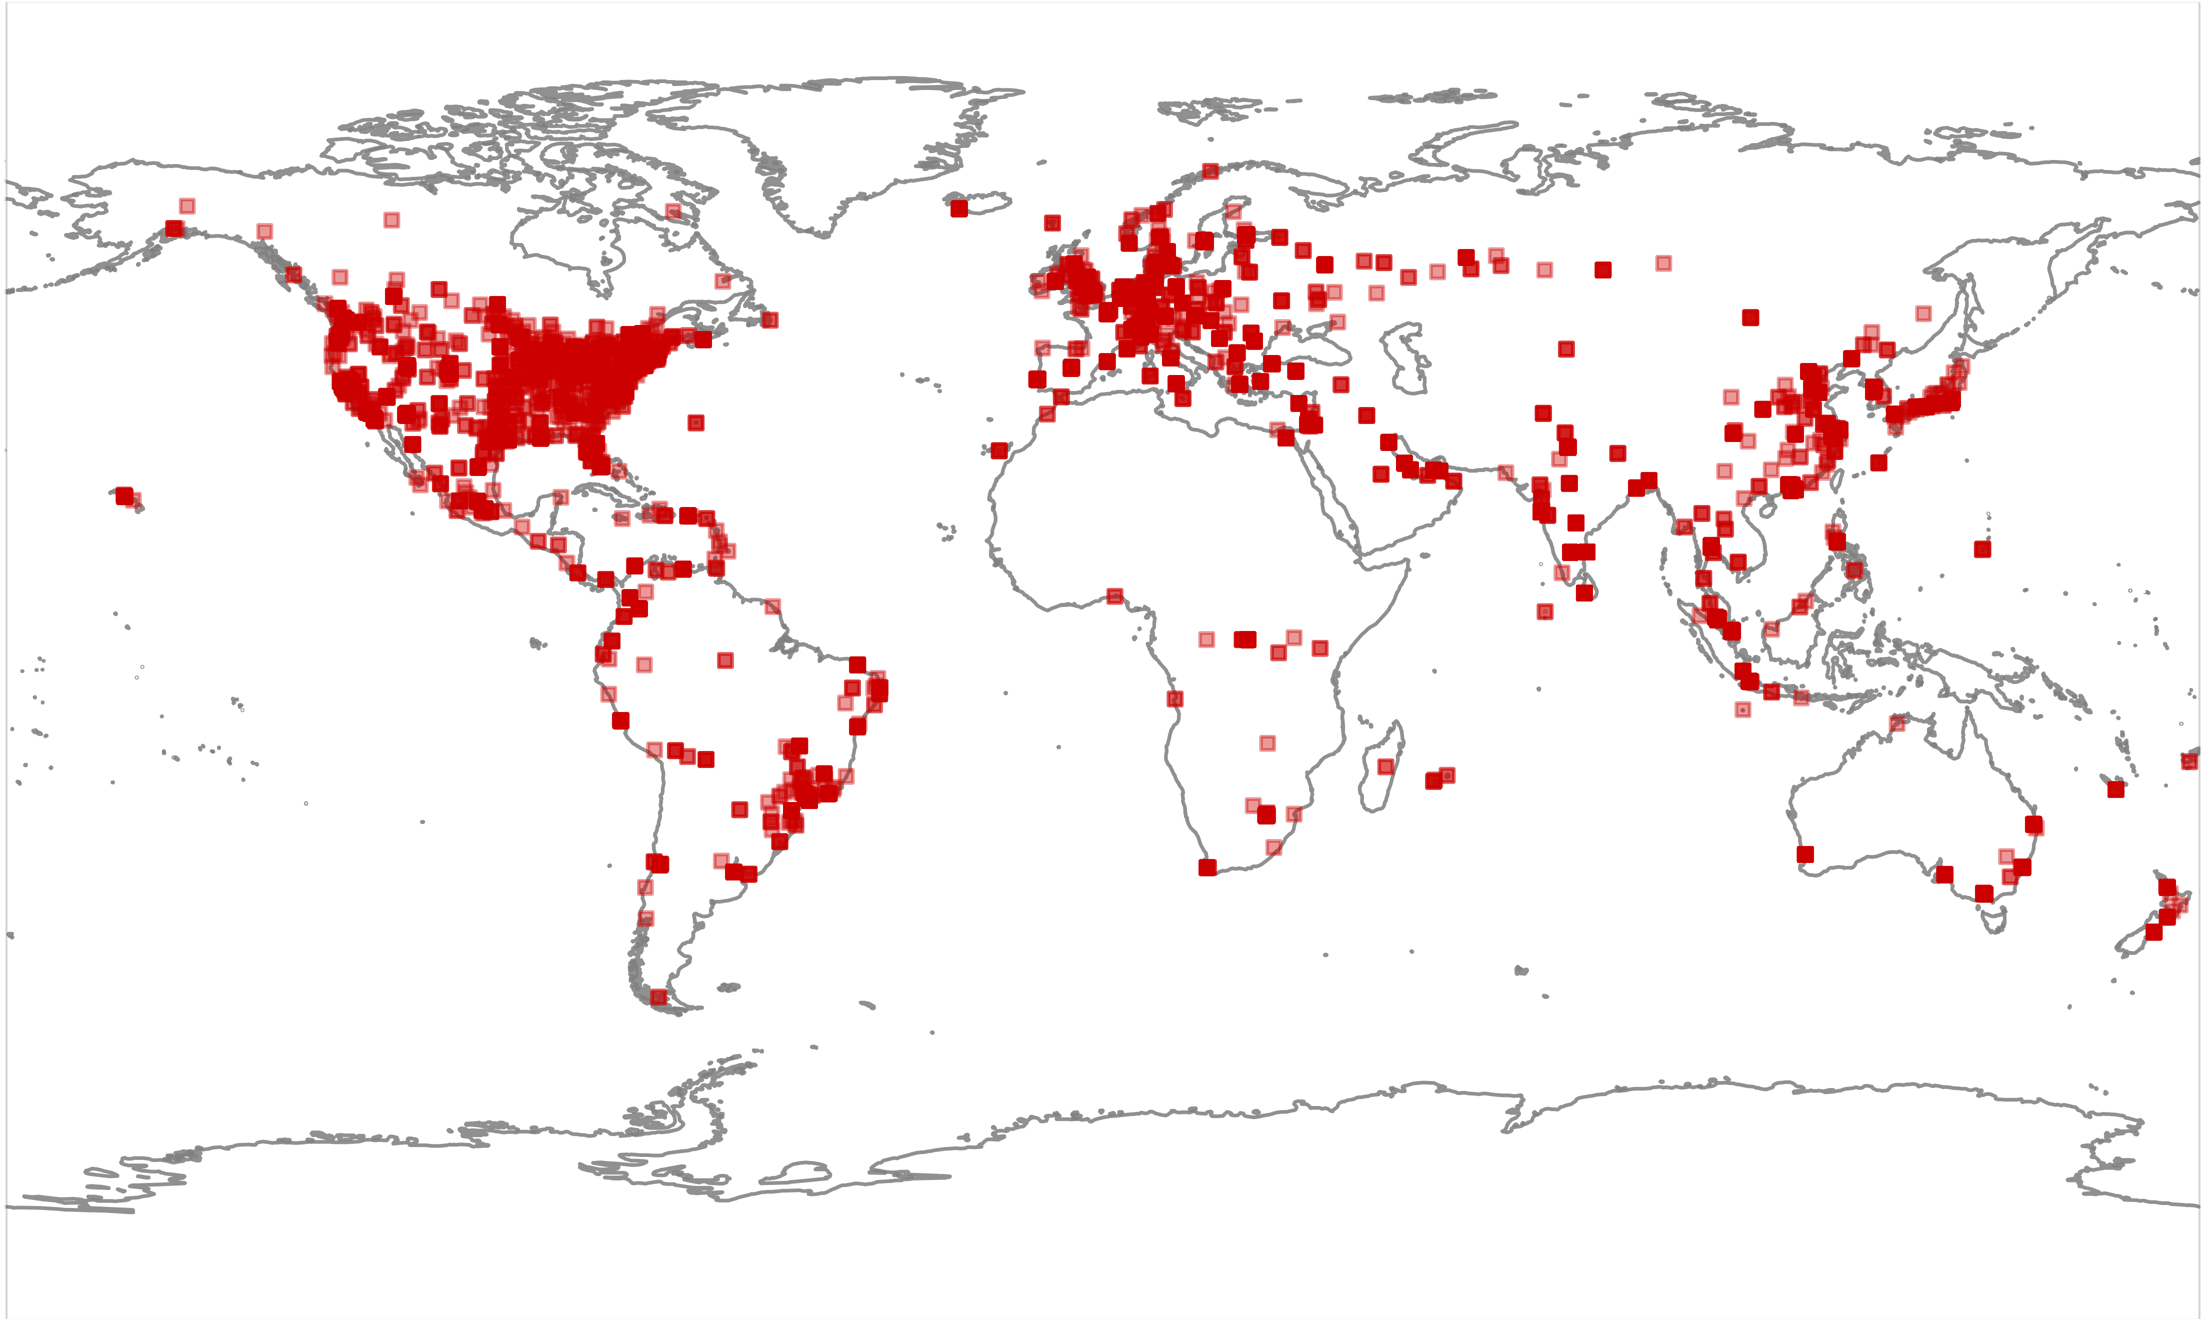

Viewing the Internet's core from the distributed measurement platform of a CDN.

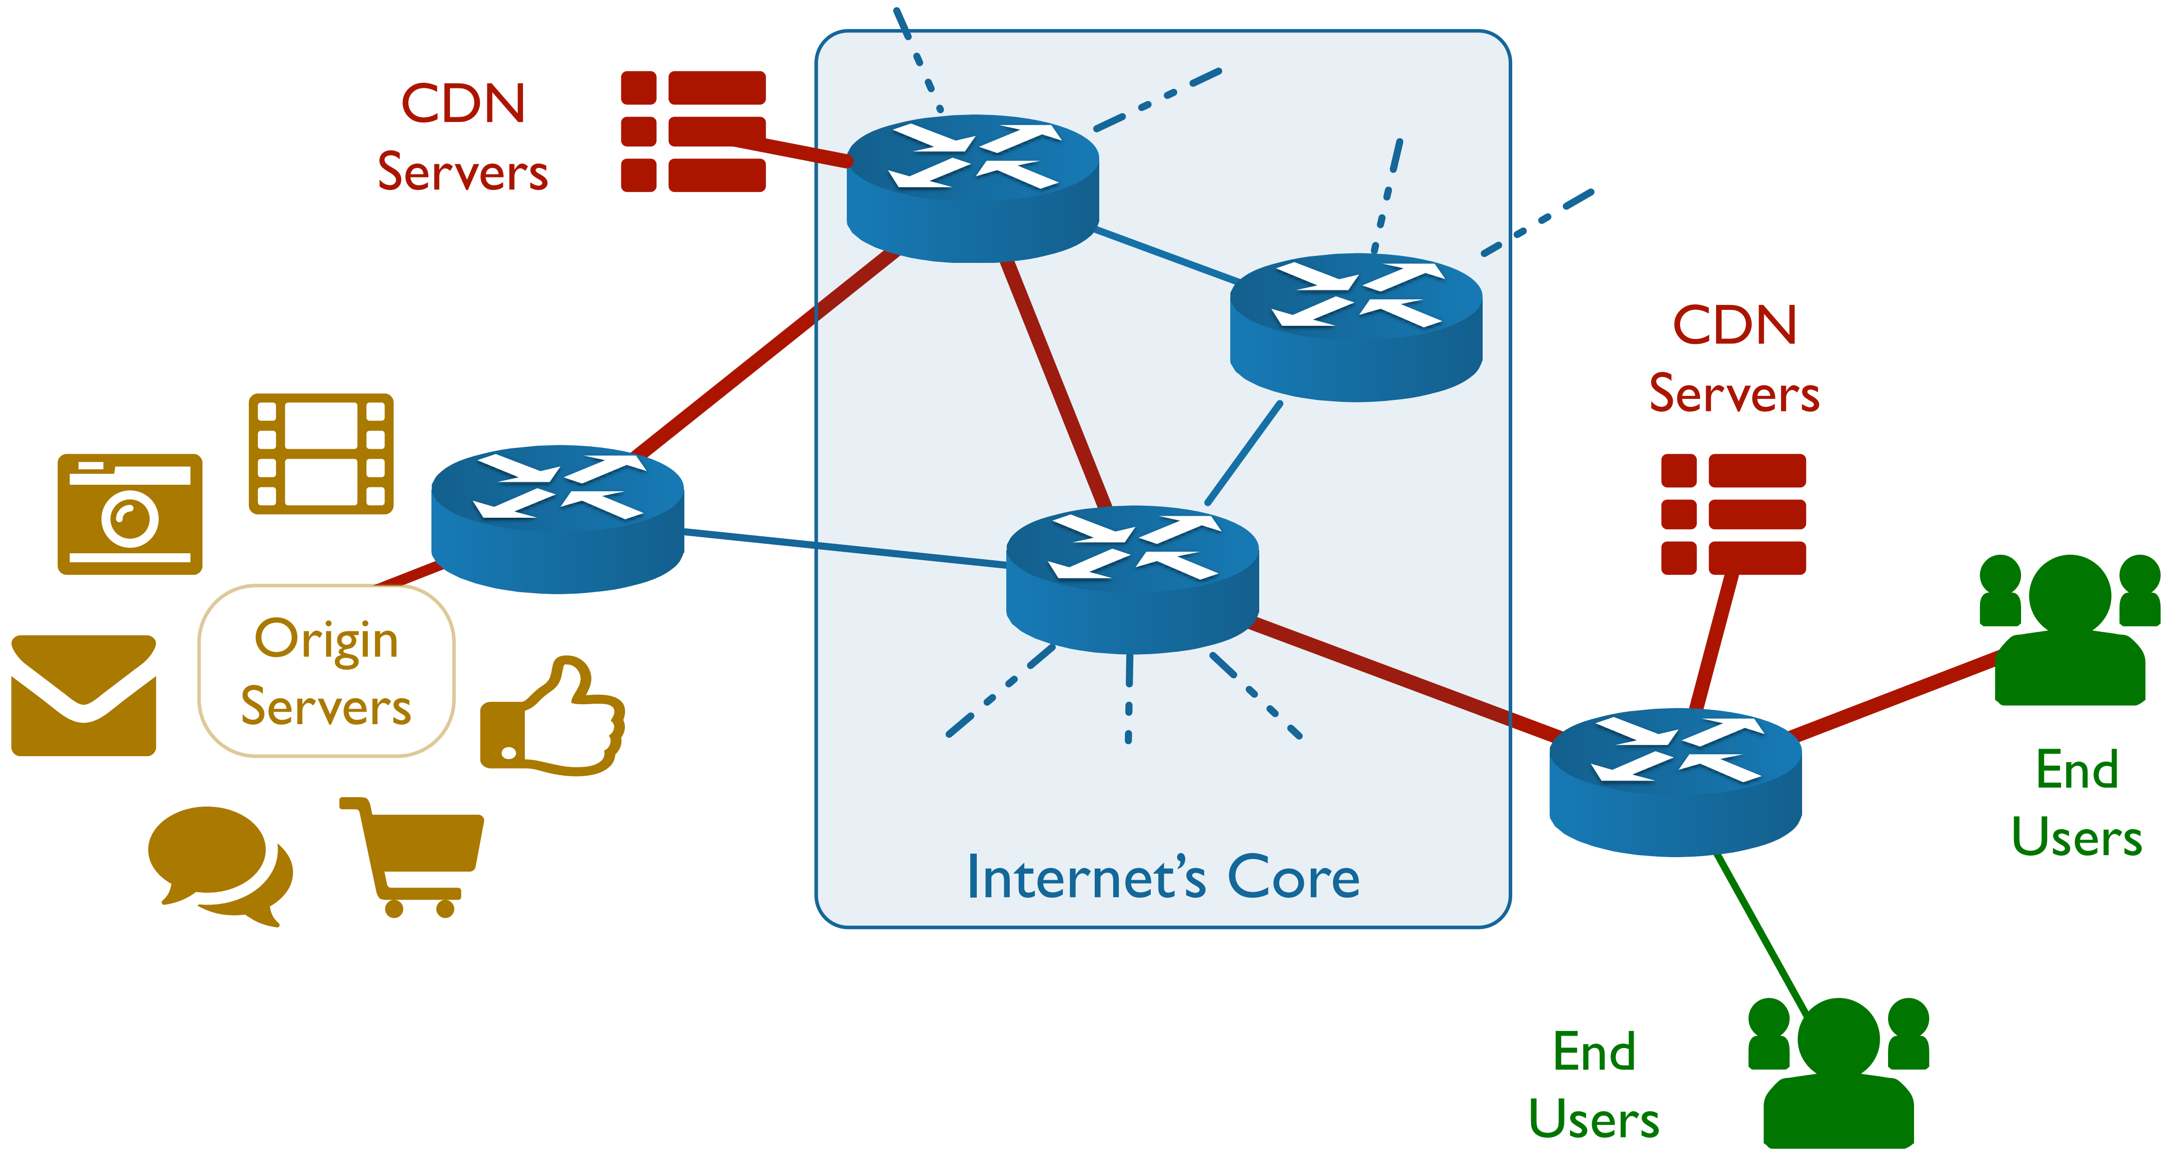

Back-Office Web traffic accounts for a significant fraction of core Internet traffic — *Pujol et al., IMC, Nov. 2014*.

End-user experience is at the mercy of the unreliable Internet and its ***middle-mile*** bottlenecks — *T. Leighton, CACM, Vol. 52. No. 2, Feb. 2009*.

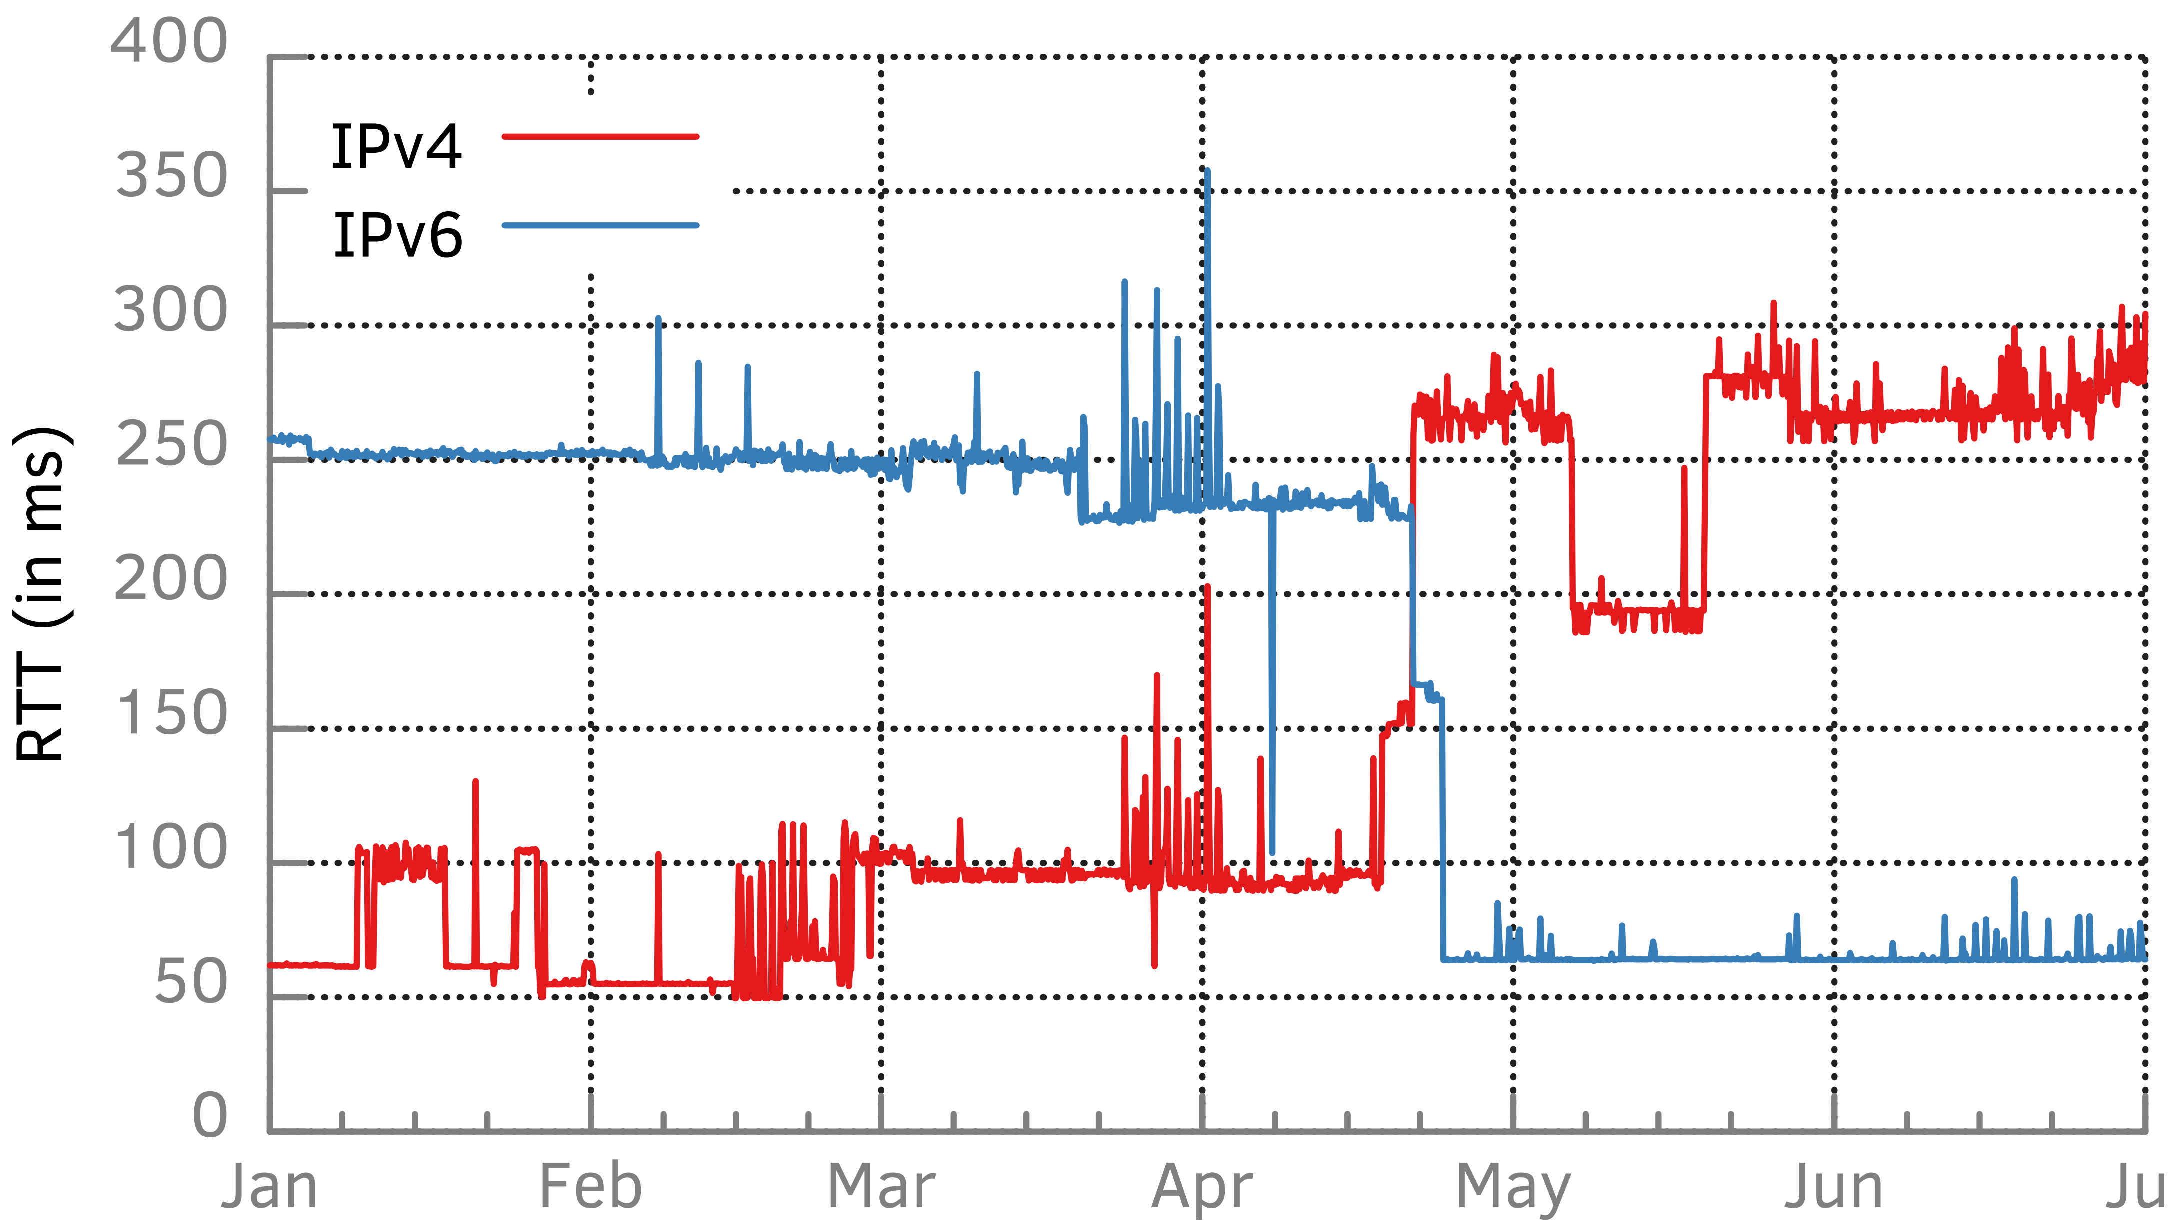

A six-month timeline of RTTs between servers in Honk Kong, HK and Tokyo, JP

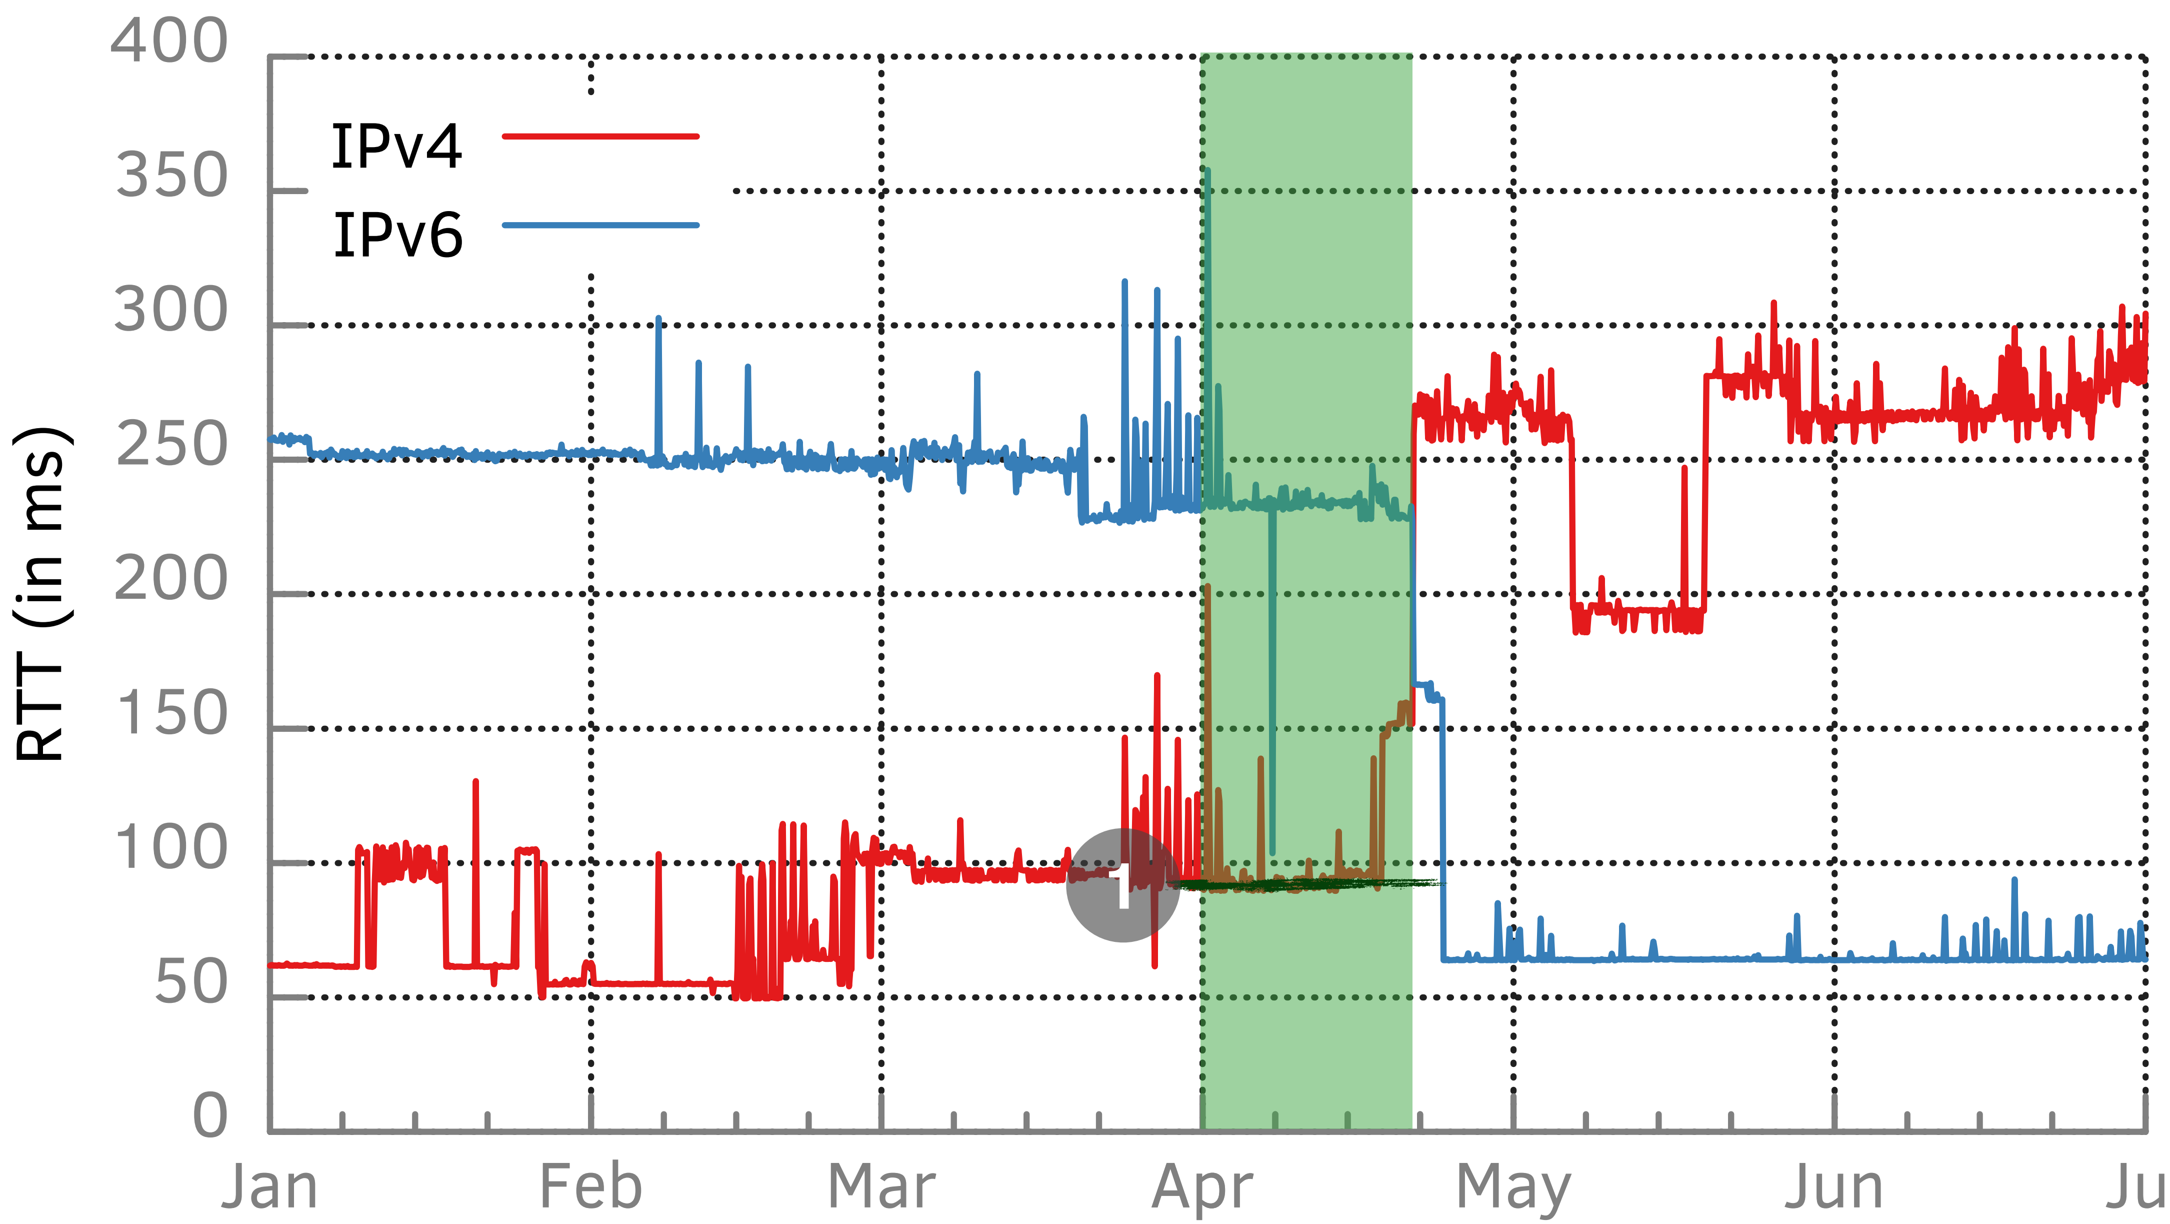

A six-month timeline of RTTs between servers in Honk Kong, HK and Tokyo, JP

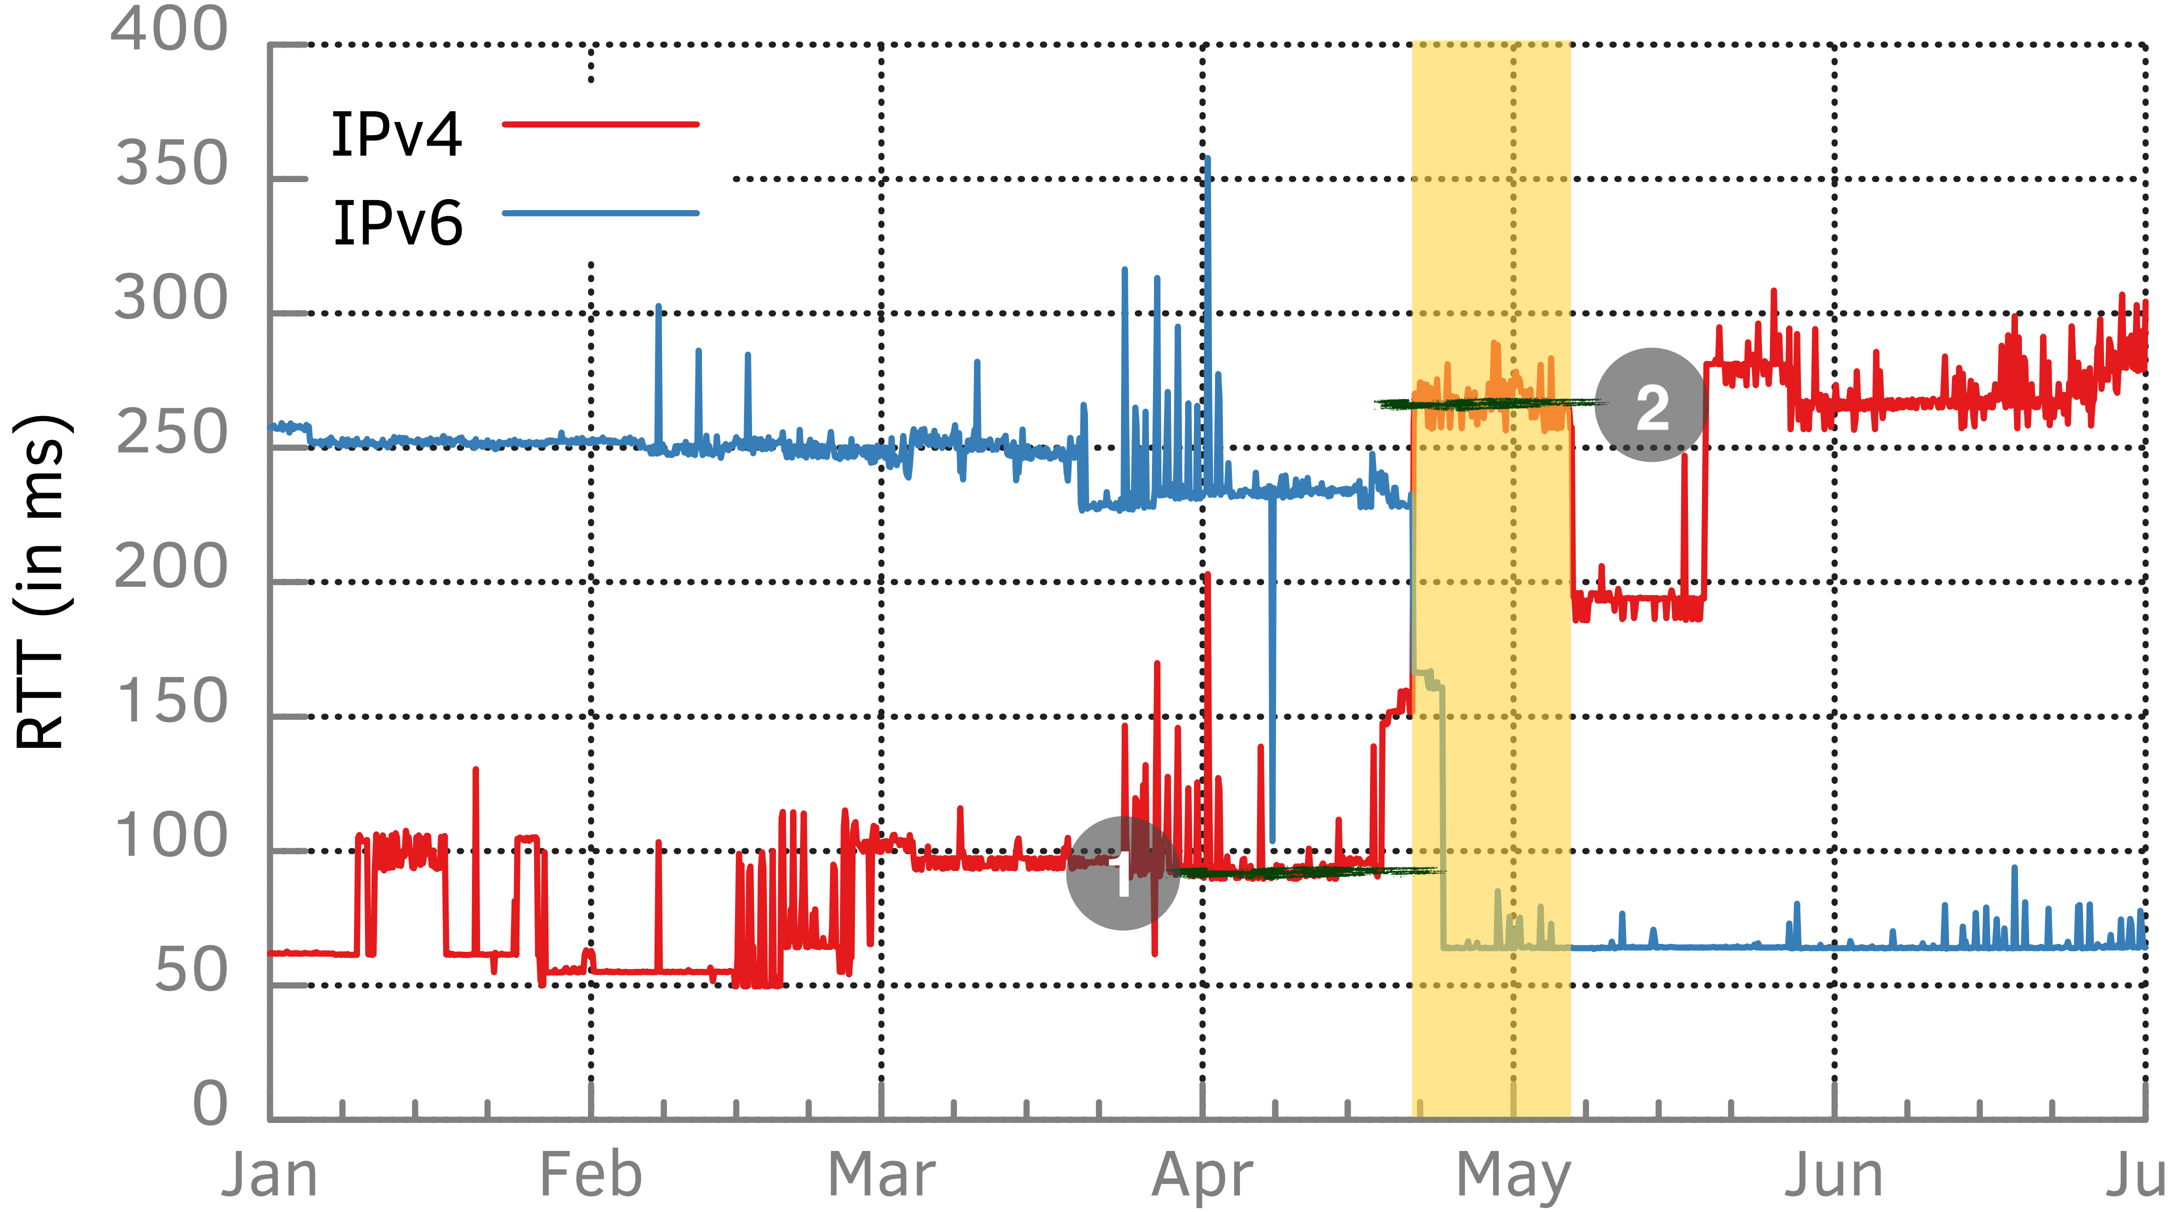

A six-month timeline of RTTs between servers in Honk Kong, HK and Tokyo, JP

Level-shifts in RTTs over both IPv4 and IPv6

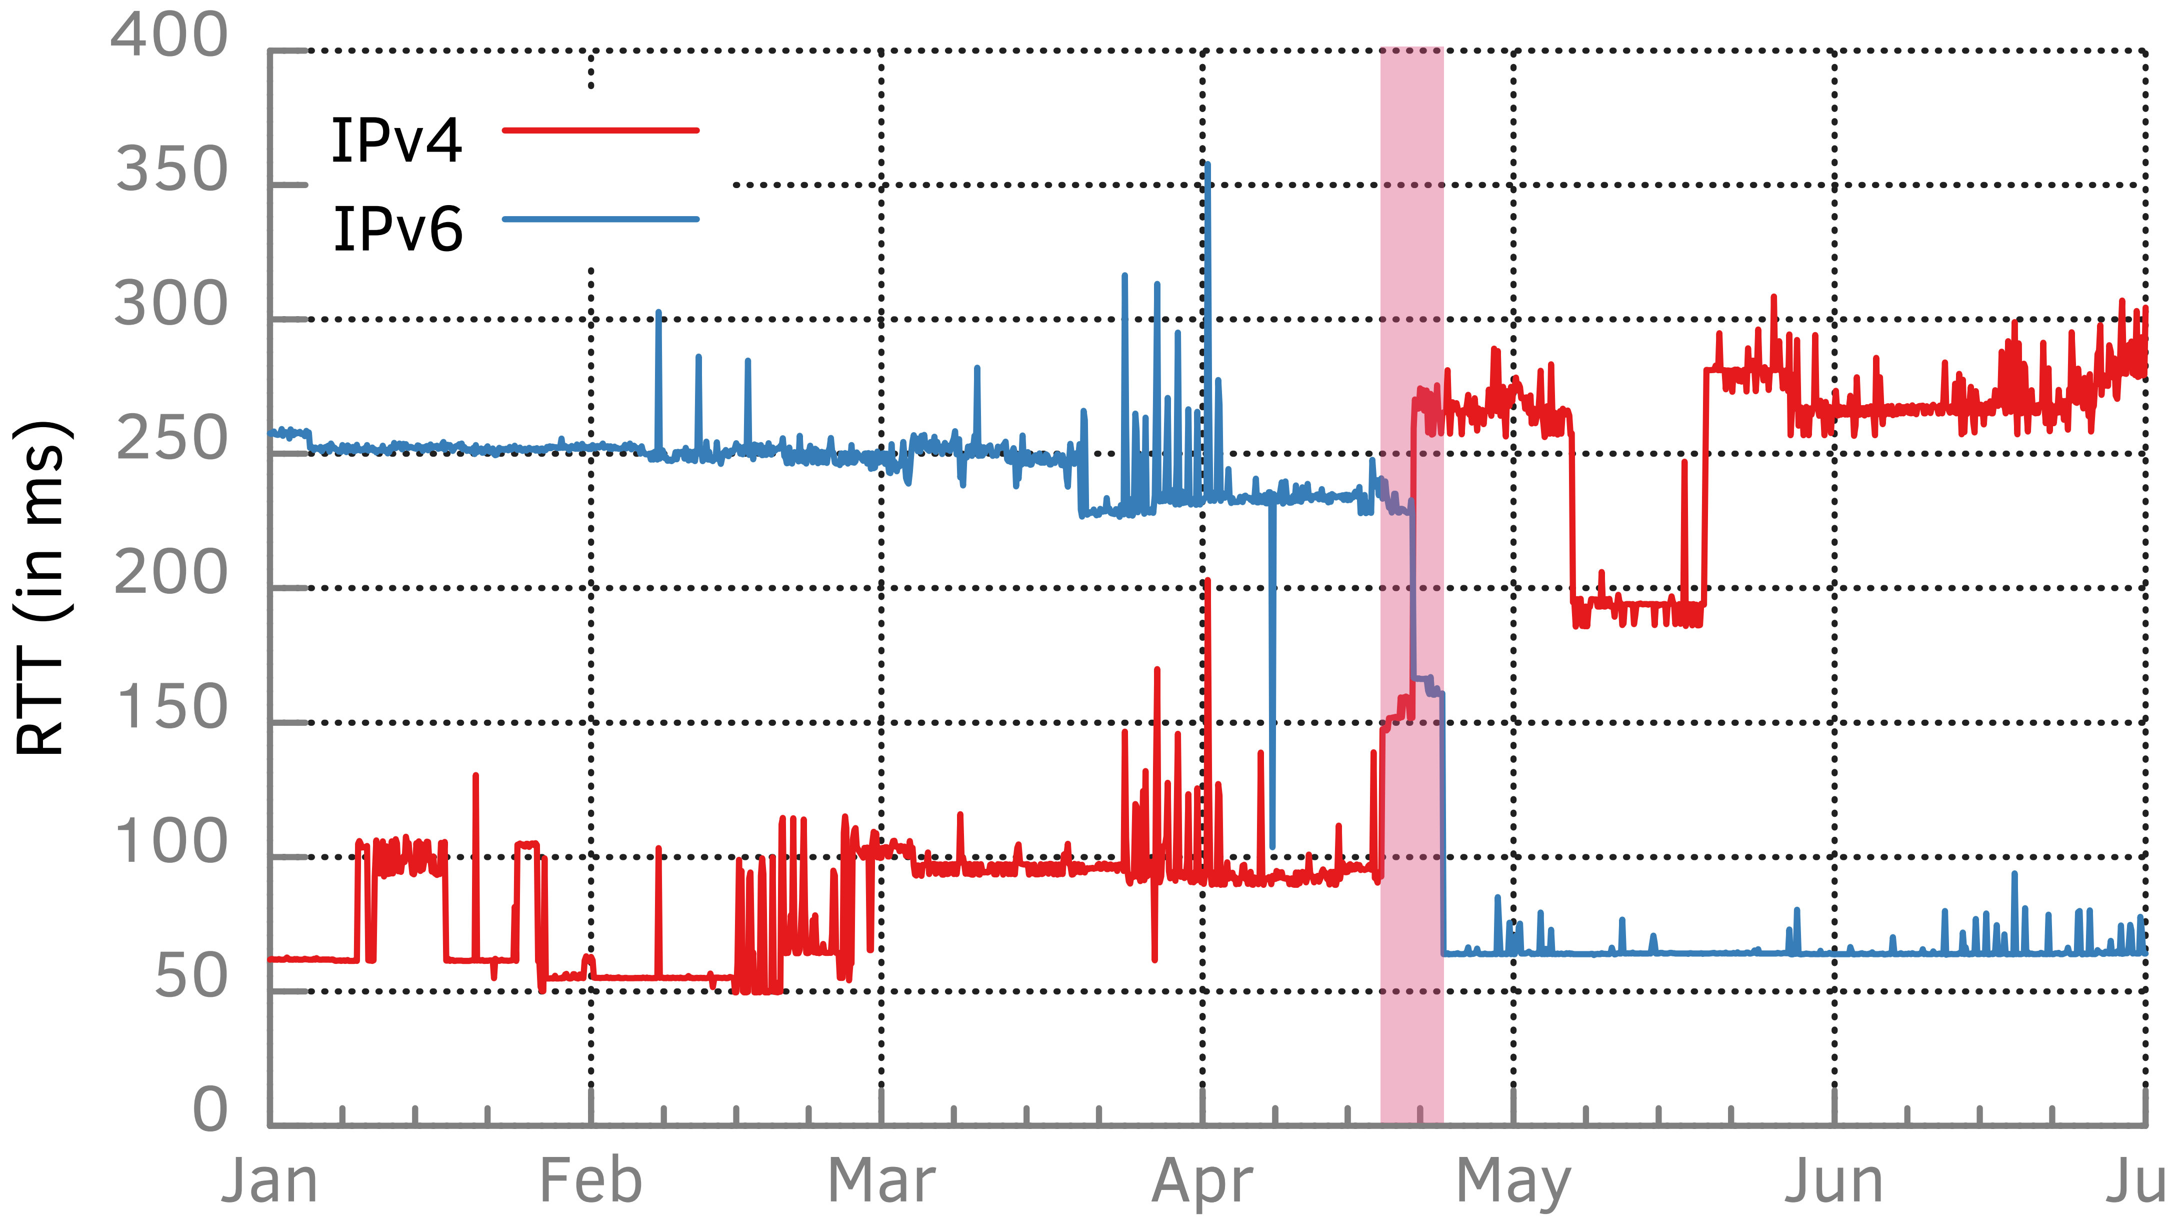

To what extent do changes in the AS path affect round-trip times?

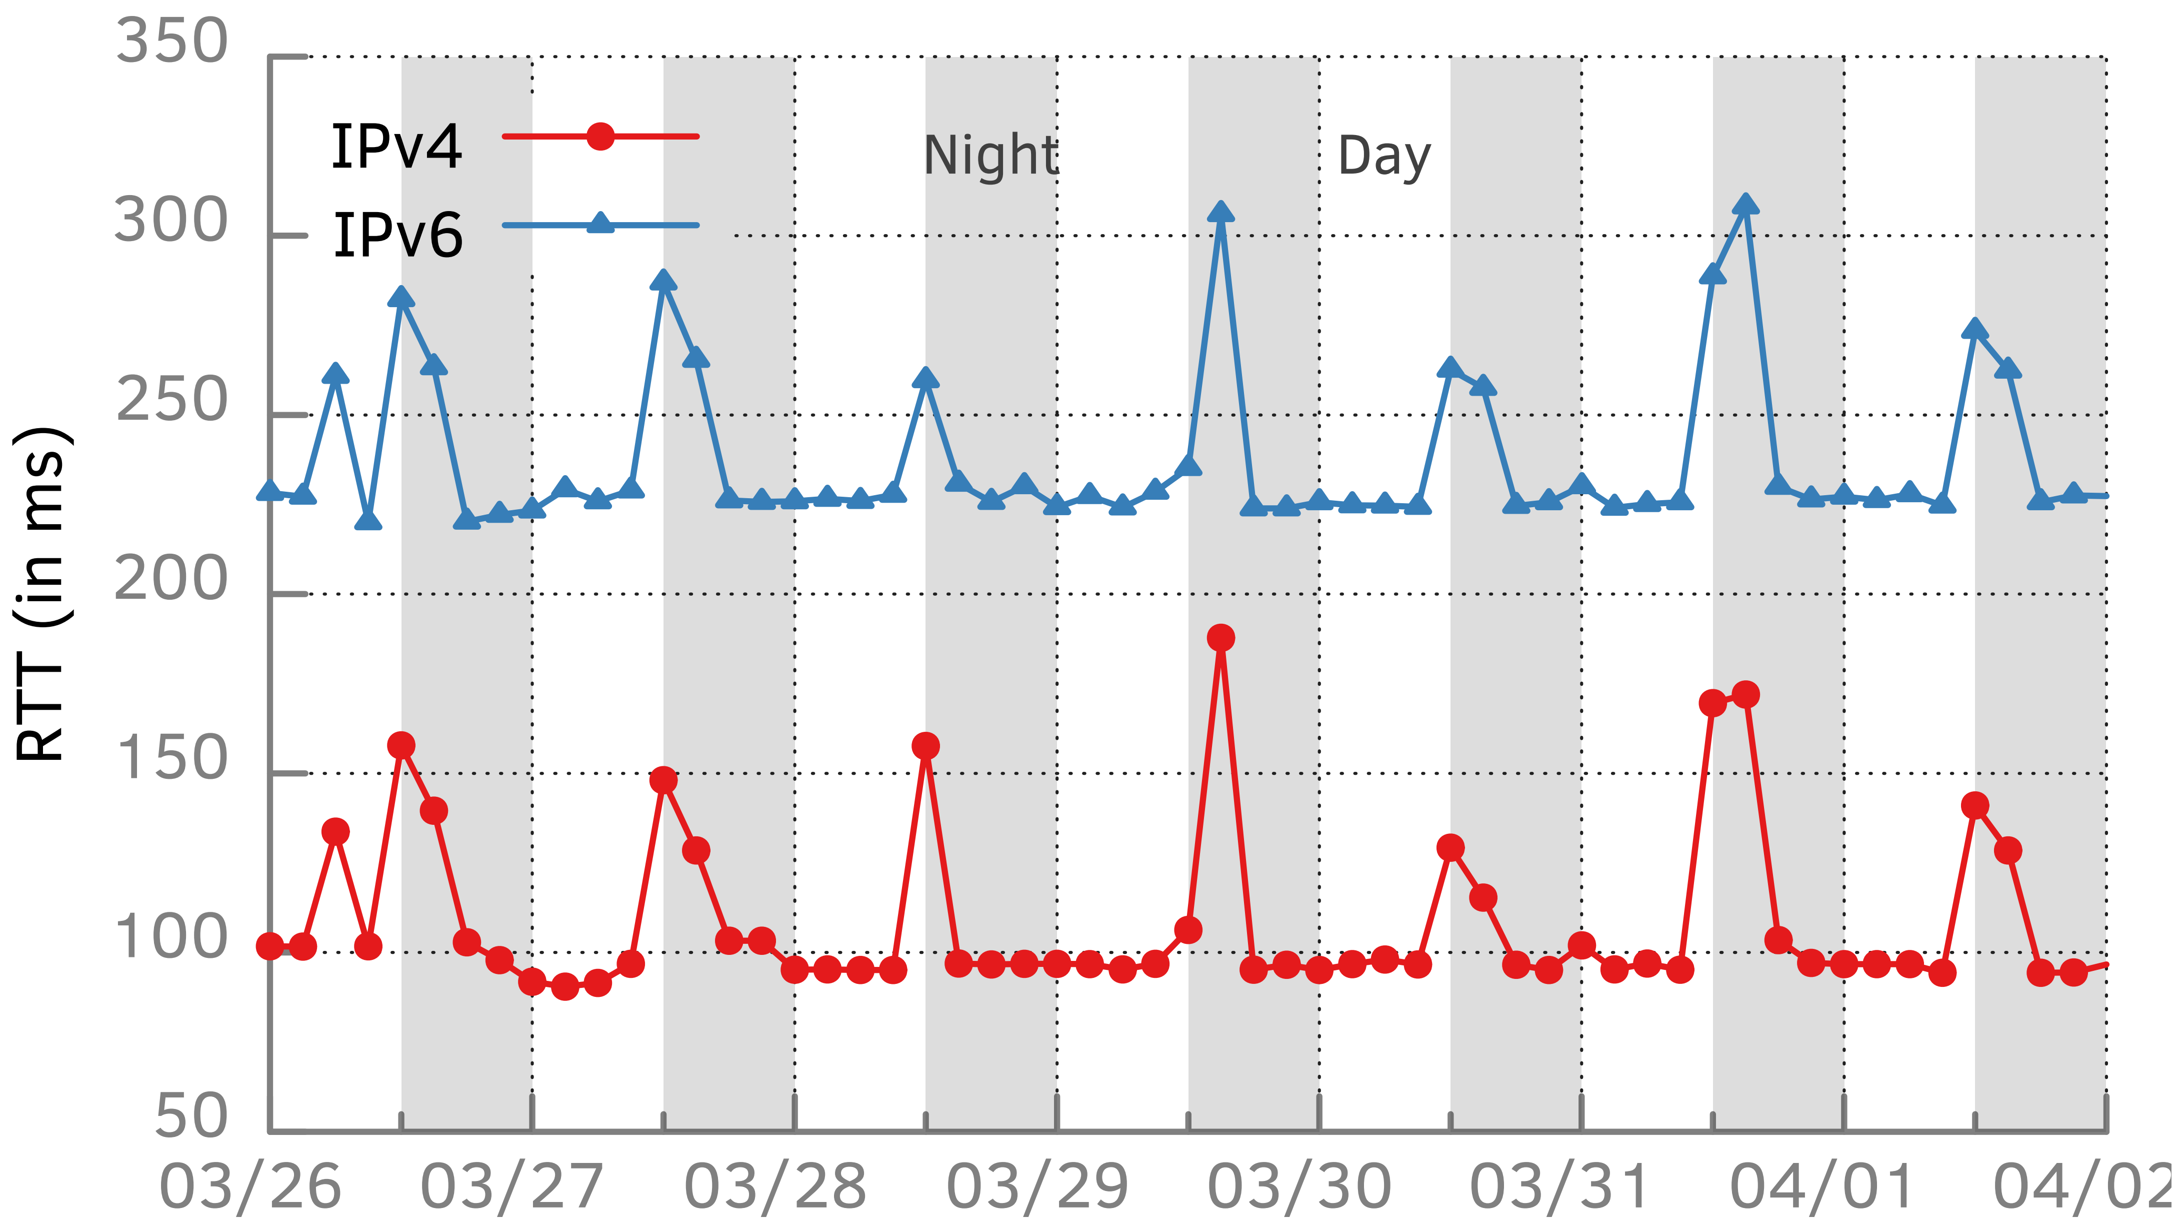

A portion of the timeline of RTTs between servers in Honk Kong, HK and Tokyo, JP.

Daily oscillations in RTT between the servers.

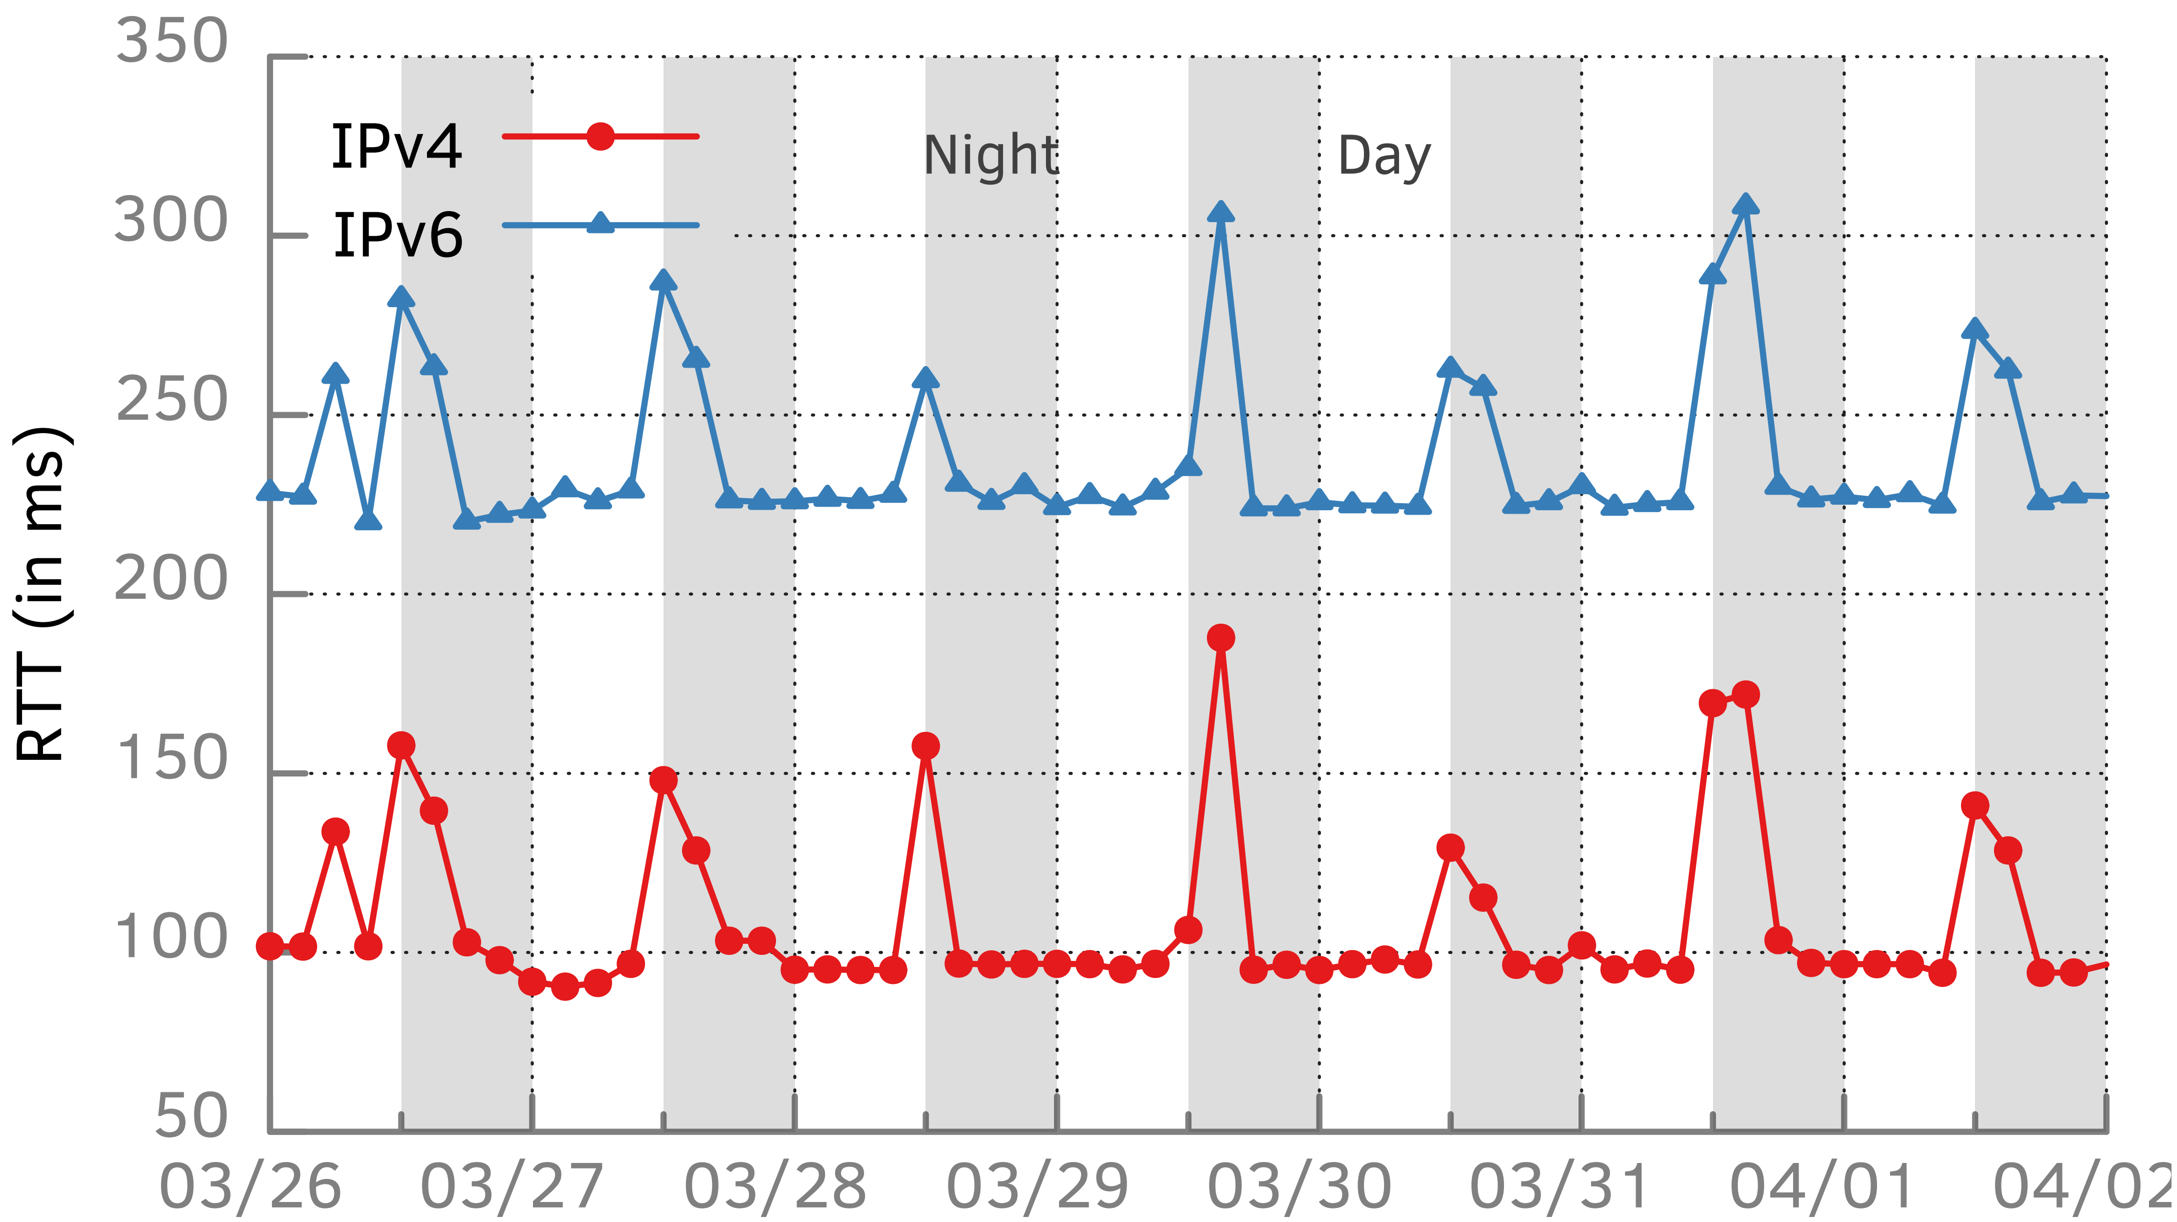

How common are periods of daily oscillation in RTT, and where do they occur?

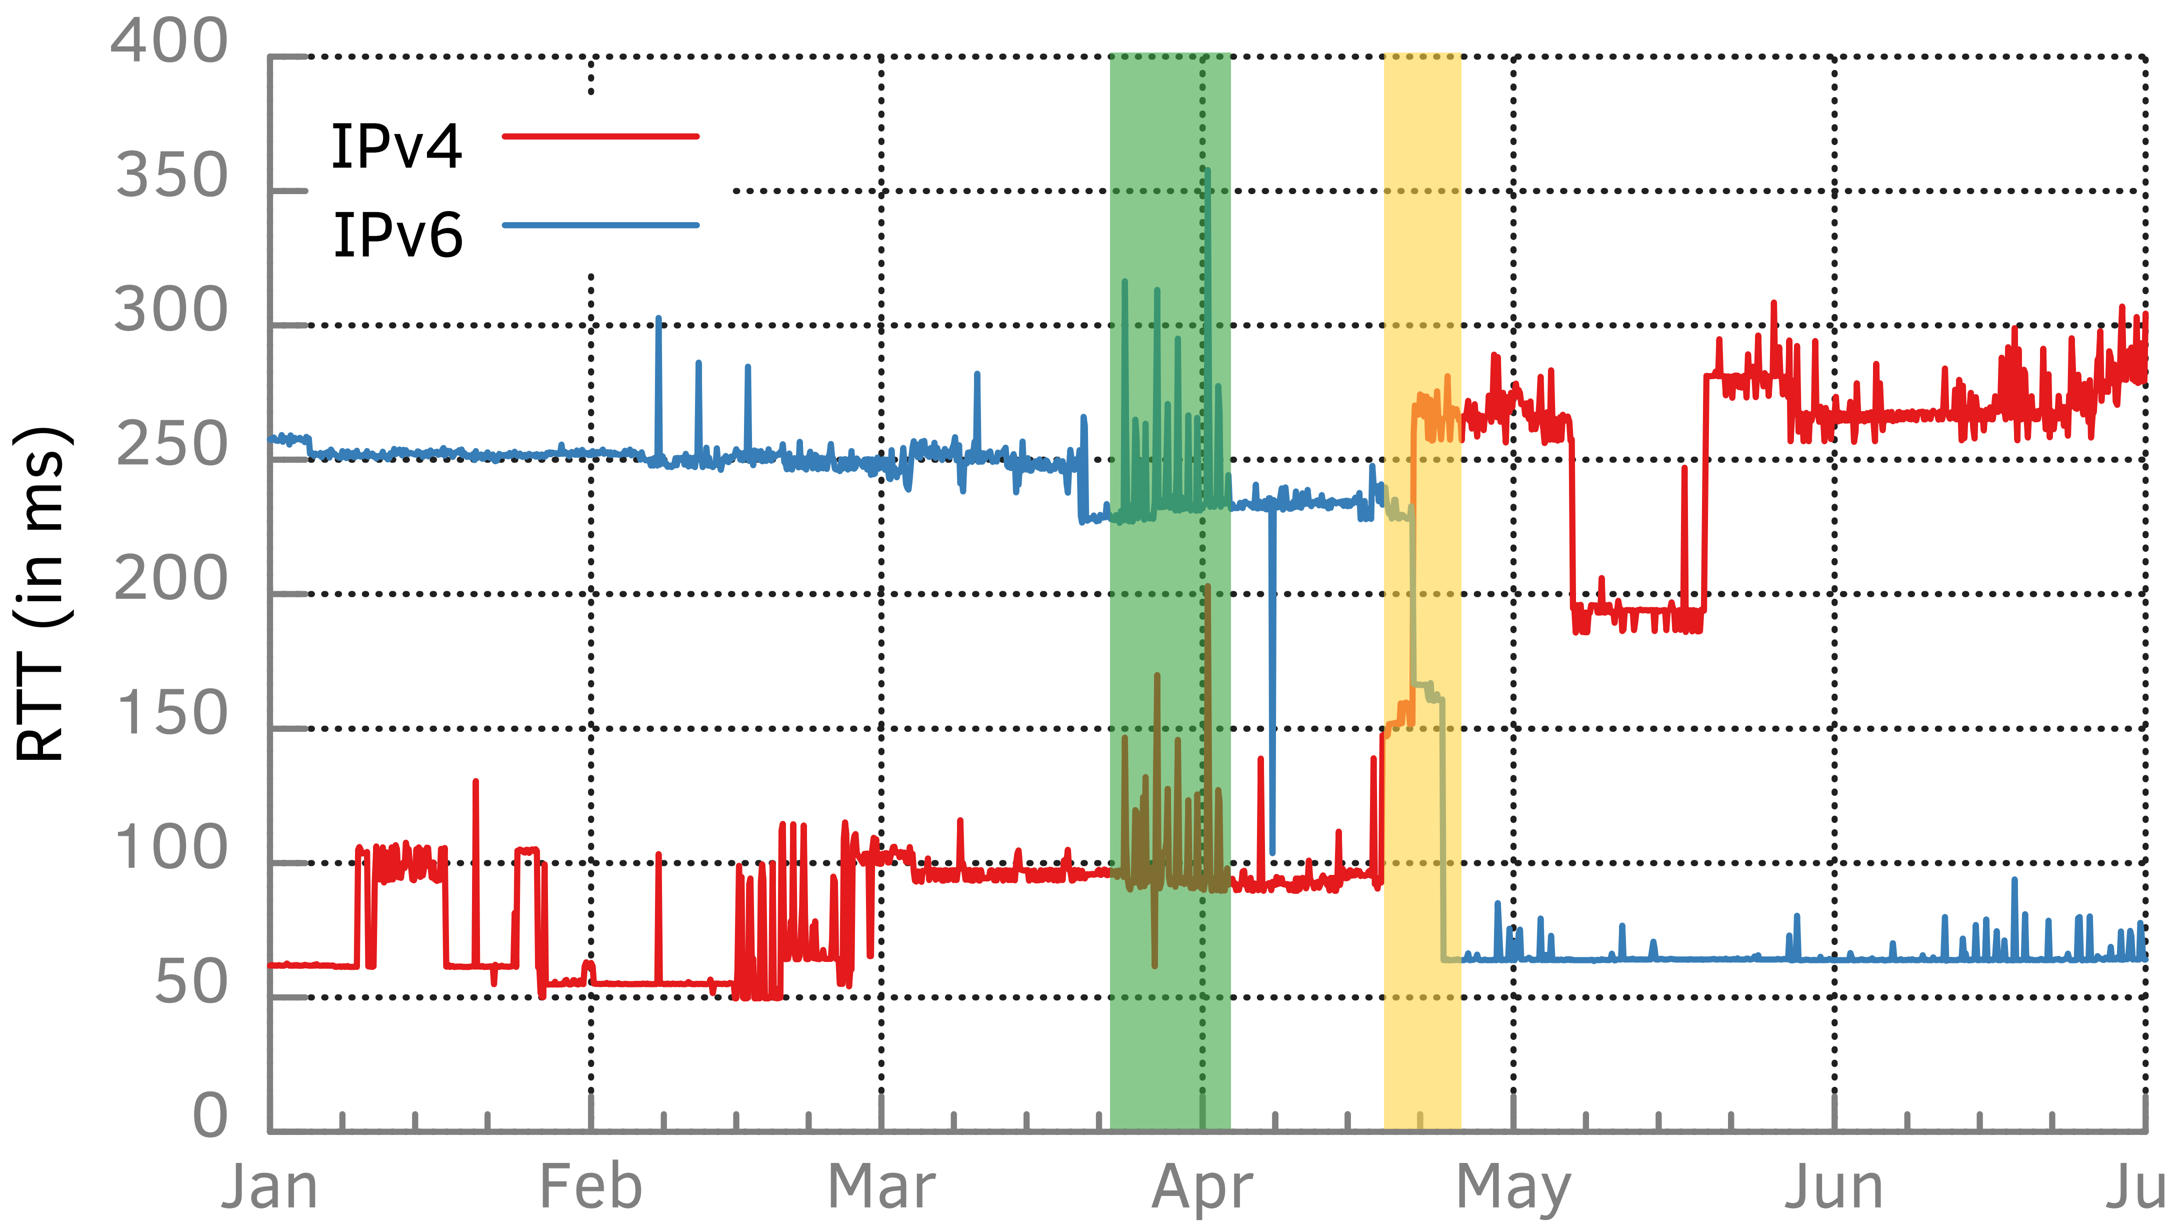

What affects end-to-end RTTs more – routing or congestion?

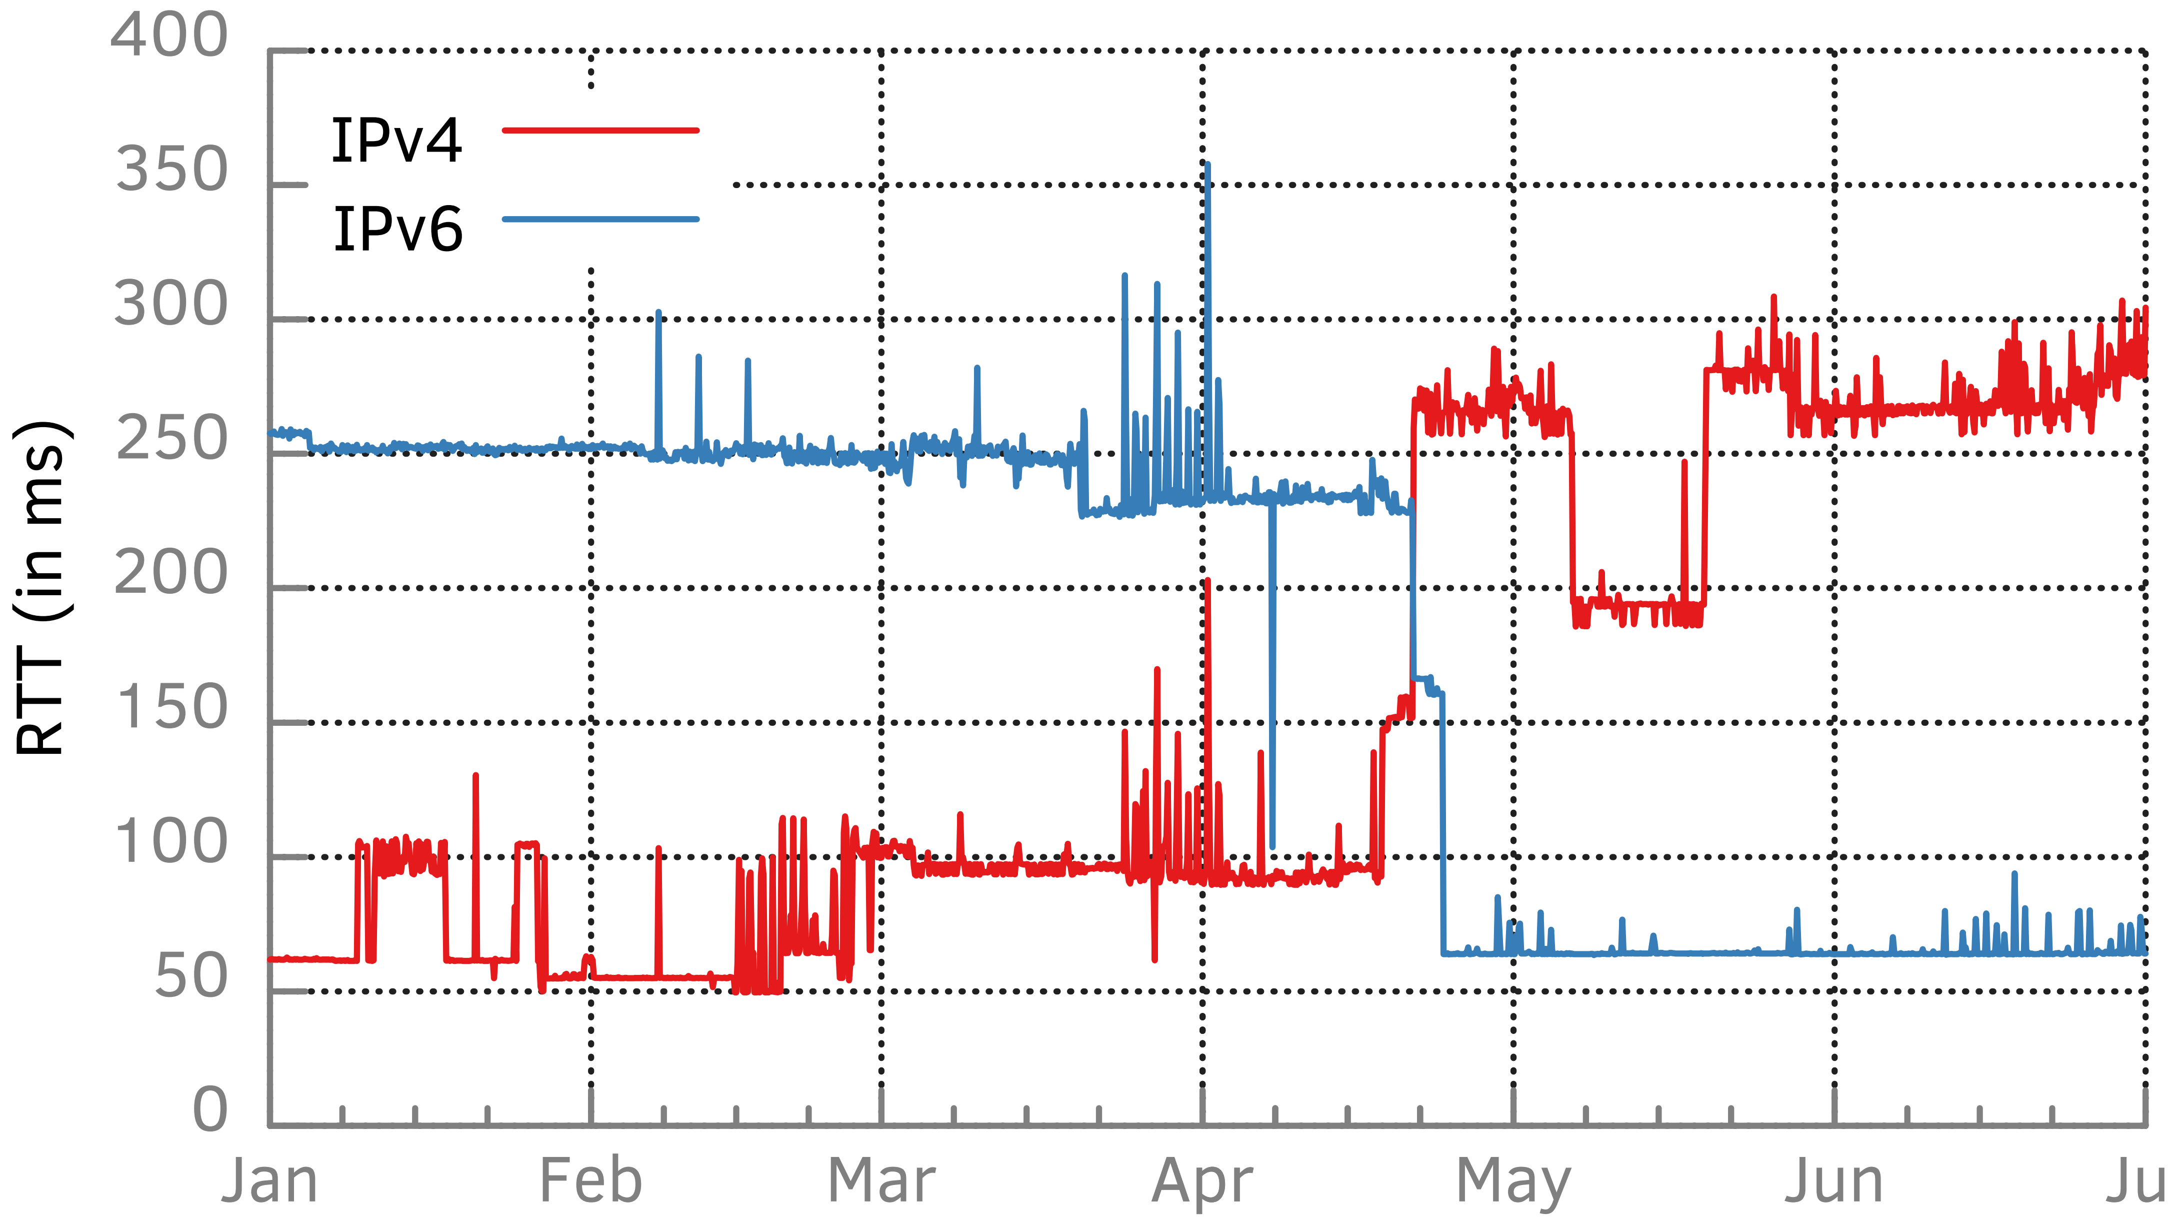

How does IPv4 and IPv6 compare with respect to routing and performance?

1. To what extent do changes in the AS path affect round-trip times?

2. How common are periods of daily oscillation in RTT, and where do they occur?

Effect of *routing changes* on end-to-end RTTs

Data Set: Long Term

- ≈***600*** dual-stacked servers in ***70*** different countries.
  - US, AU, DE, IN, JP, …

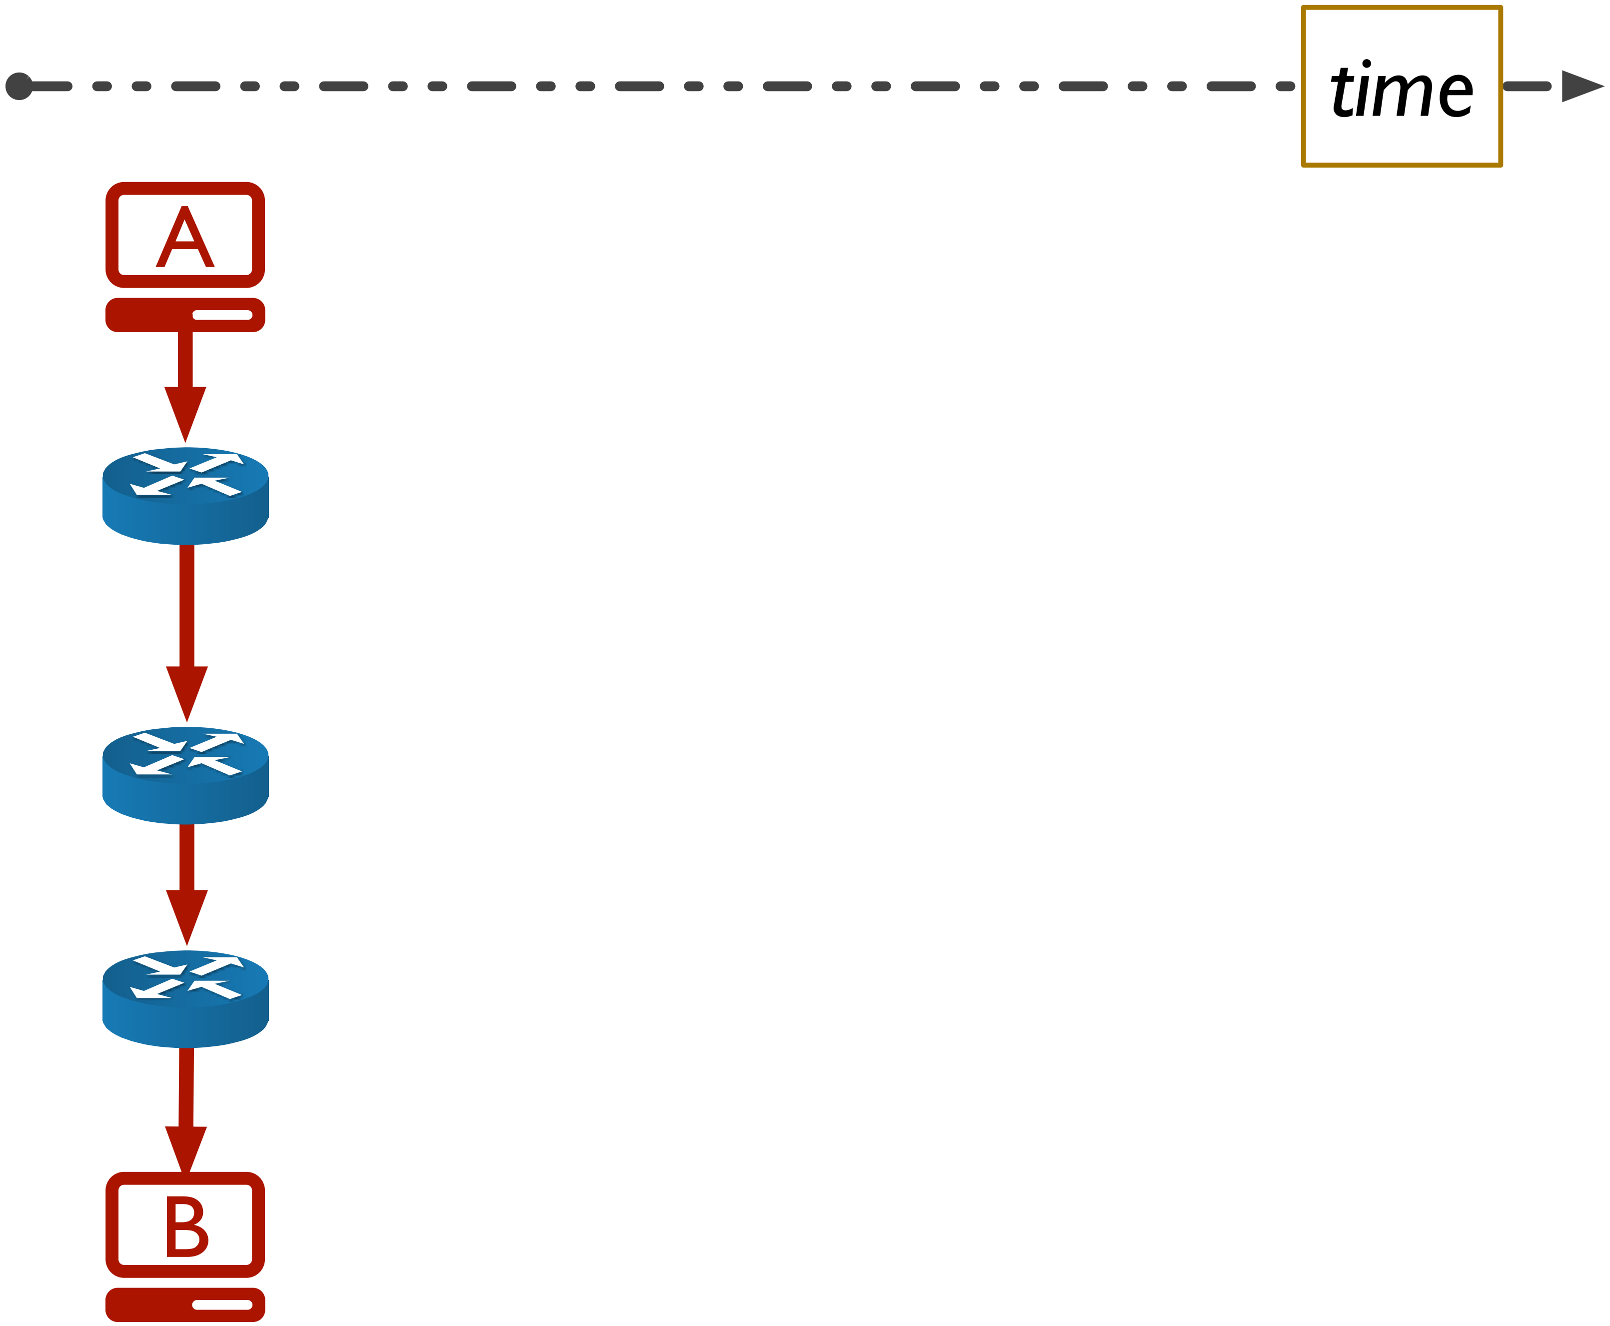

Traceroutes conducted between servers in both directions over both protocols.

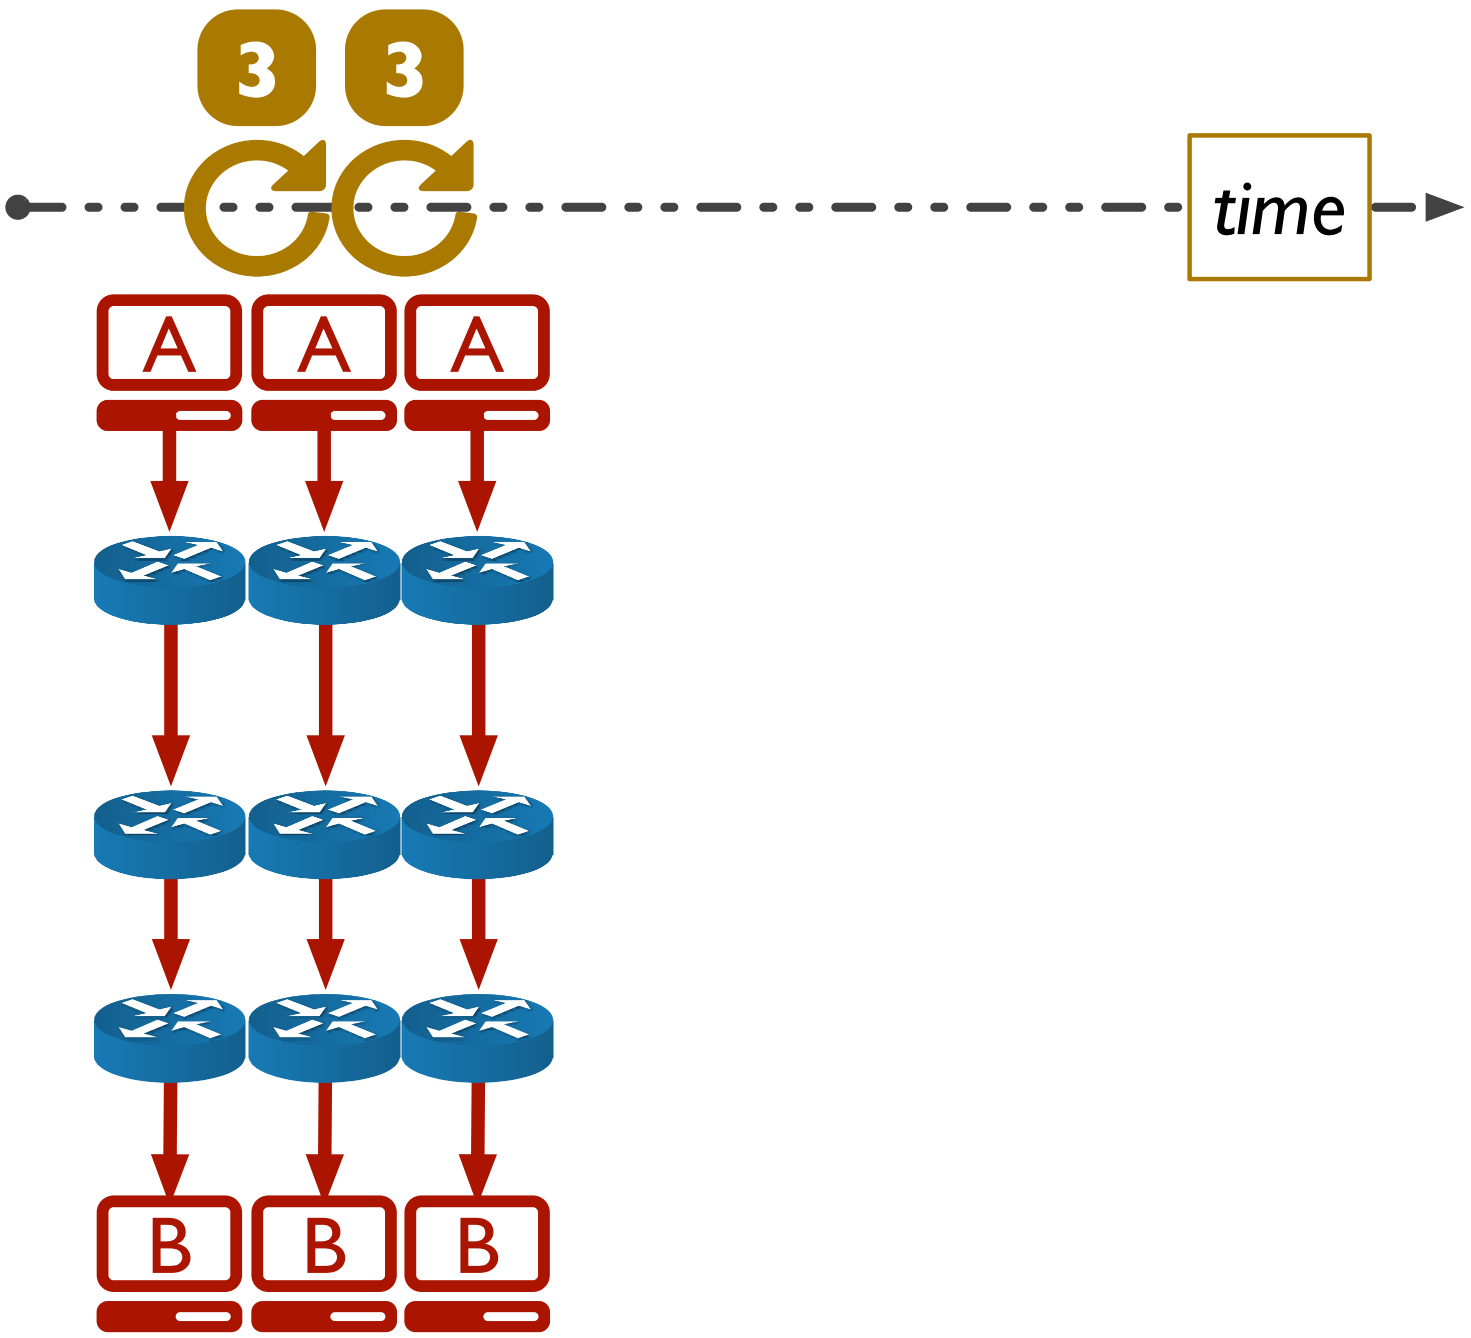

Every ***3*** hours traceroutes done over the *full-mesh*.

All traceroutes in a given ***3*** hour time frame have the same timestamp.

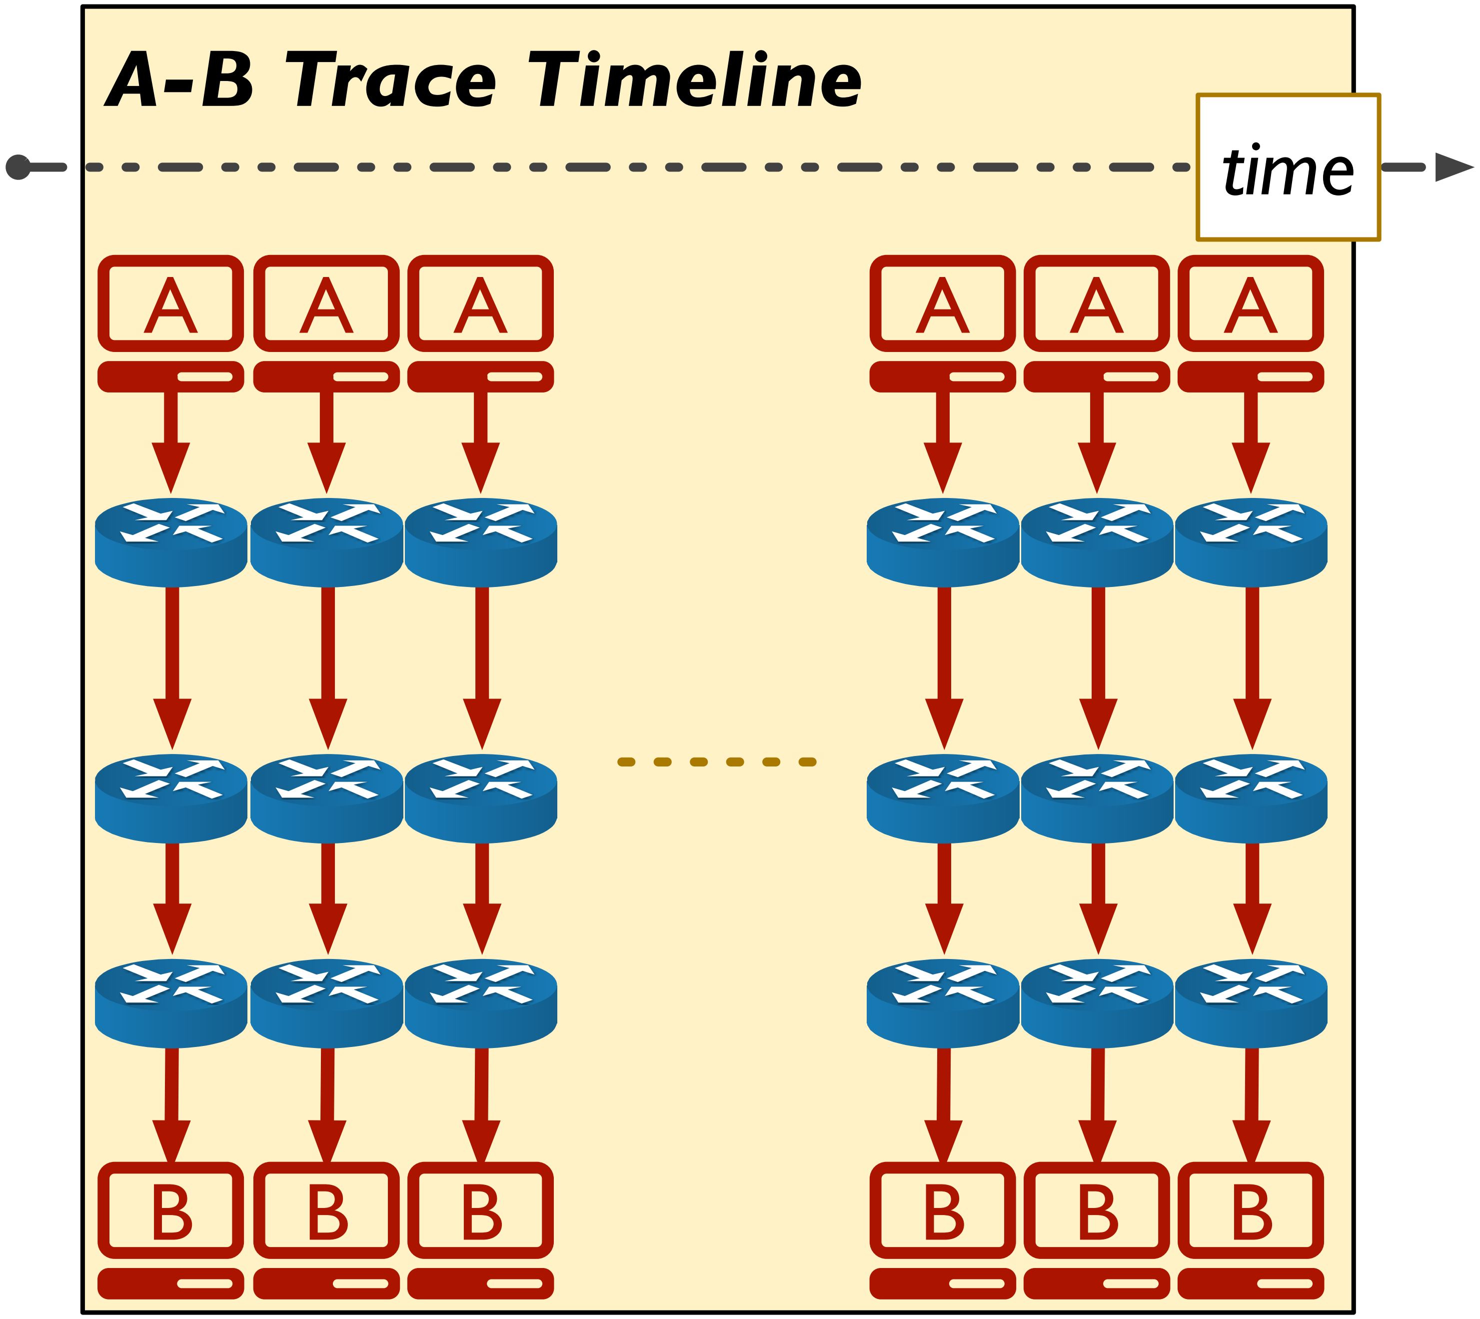

Traceroutes over the *full-mesh* every ***3*** hours for ***16*** months from Jan. 2014 through Apr. 2015.

≈***700M*** IPv4 and ≈***600M*** IPv6 traceroutes

Trace timeline $S_a \rightarrow S_b$ *is different from* $S_b \rightarrow S_a$

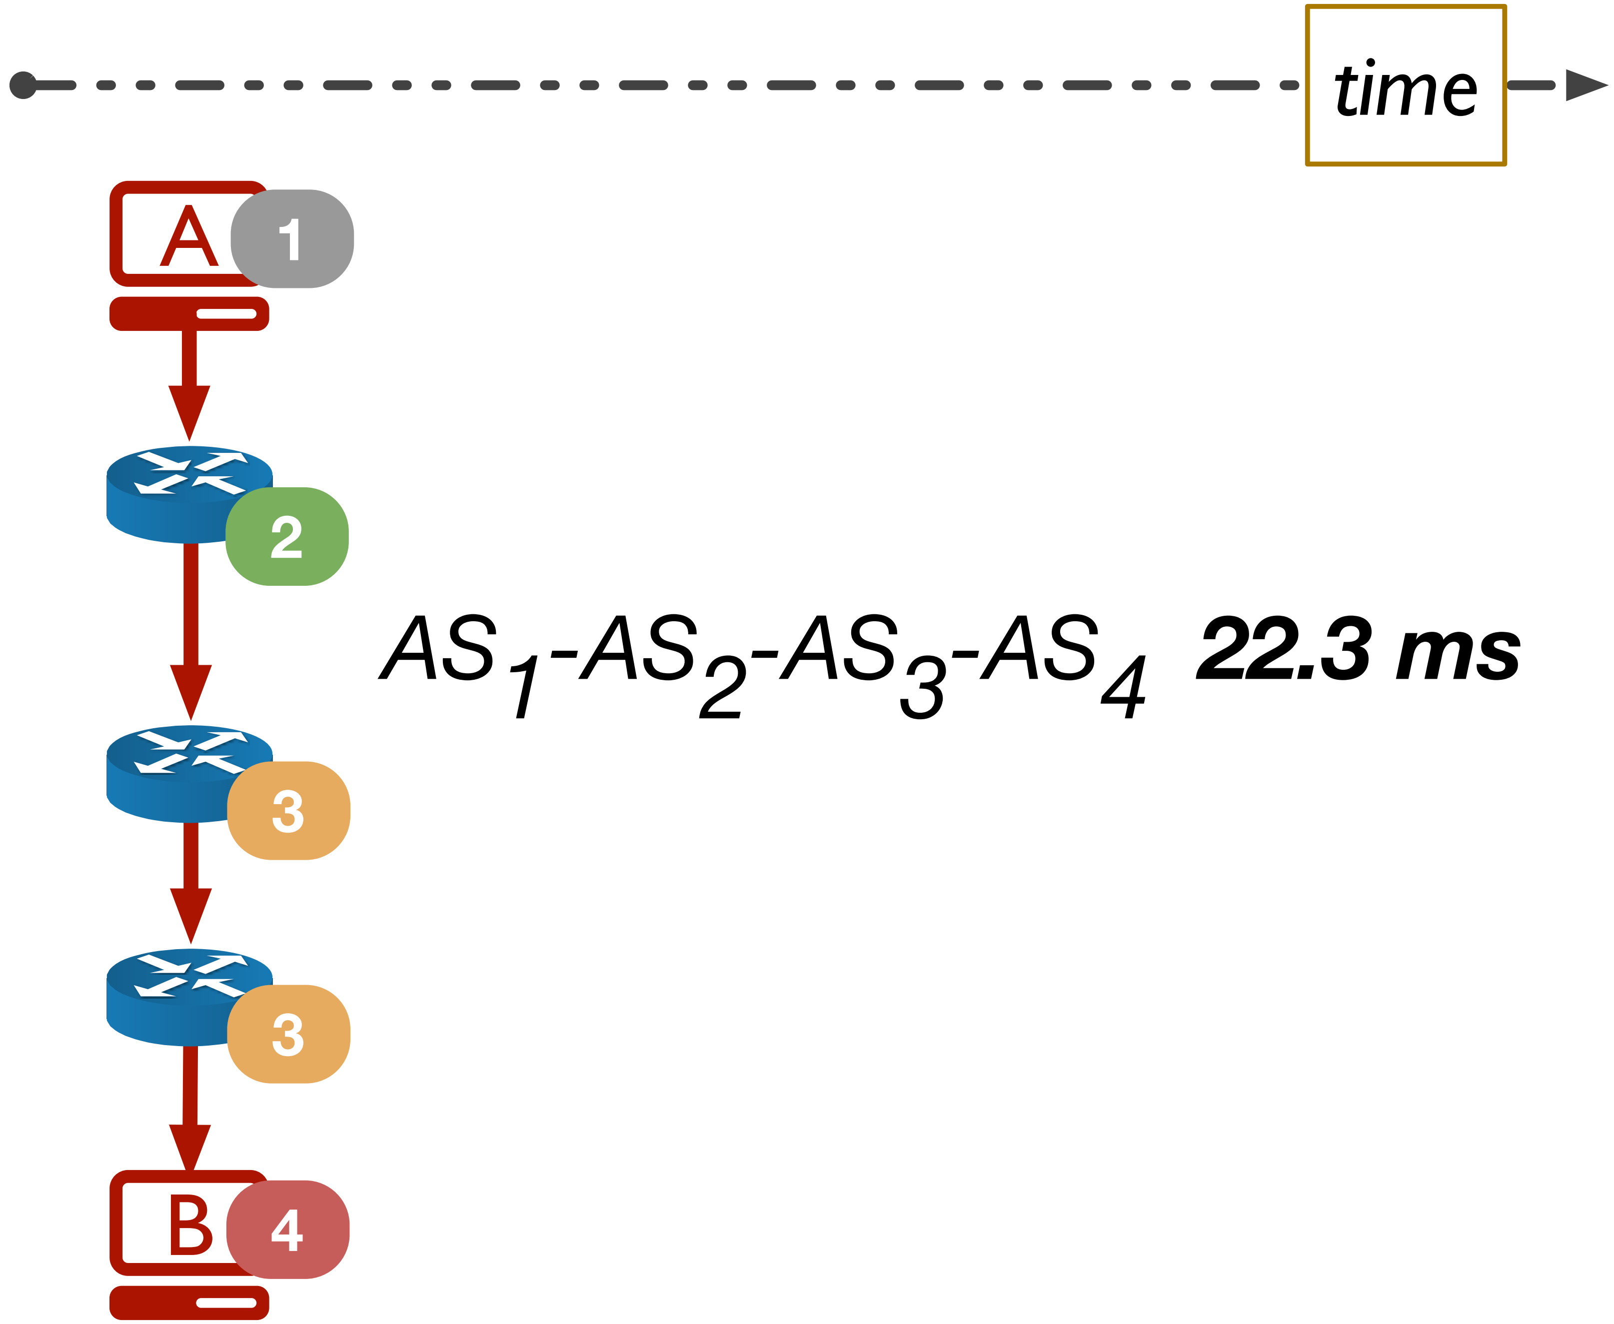

- Extract two pieces of information from each traceroute
  - ***AS path*** inferred from interfaces in the traceroute output
  - ***end-to-end RTT*** between the two servers

time

AS_1	AS_1	AS_1	AS_1	AS_1
AS_2	AS_2	AS_5	AS_2	AS_5
AS_3	AS_3	AS_9	AS_3	AS_9
AS_4	AS_4	AS_4	AS_4	AS_4
22.3	**29.7**	**18.2**	**23.1**	**17.9**

A–B trace timeline

(AS-path, end-to-end RTT) tuples spanning the study period

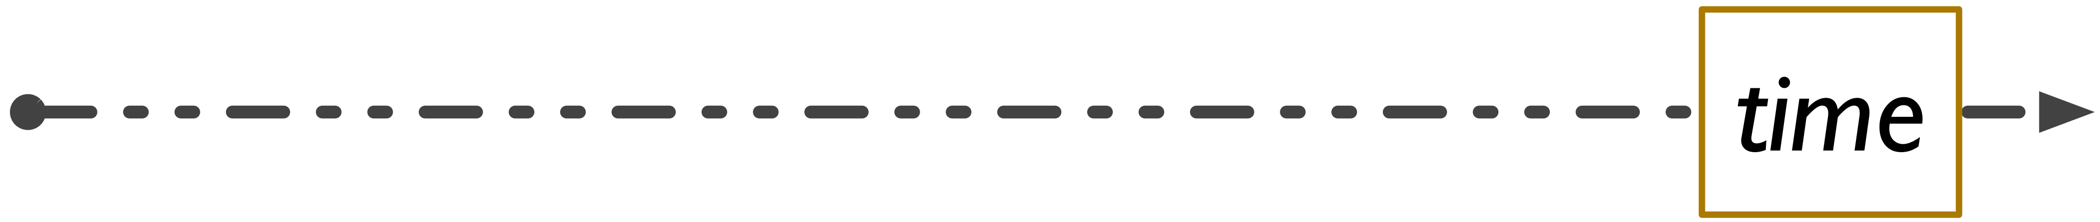

AS_1	AS_1	AS_1	AS_1	AS_1
AS_2	AS_2	AS_5	AS_2	AS_5
AS_3	AS_3	AS_9	AS_3	AS_9
AS_4	AS_4	AS_4	AS_4	AS_4
22.3	***29.7***	***18.2***	***23.1***	***17.9***

Popular AS path observed in A–B trace timeline

AS_1-AS_2-AS_3-AS_4 with prevalence ***60%***

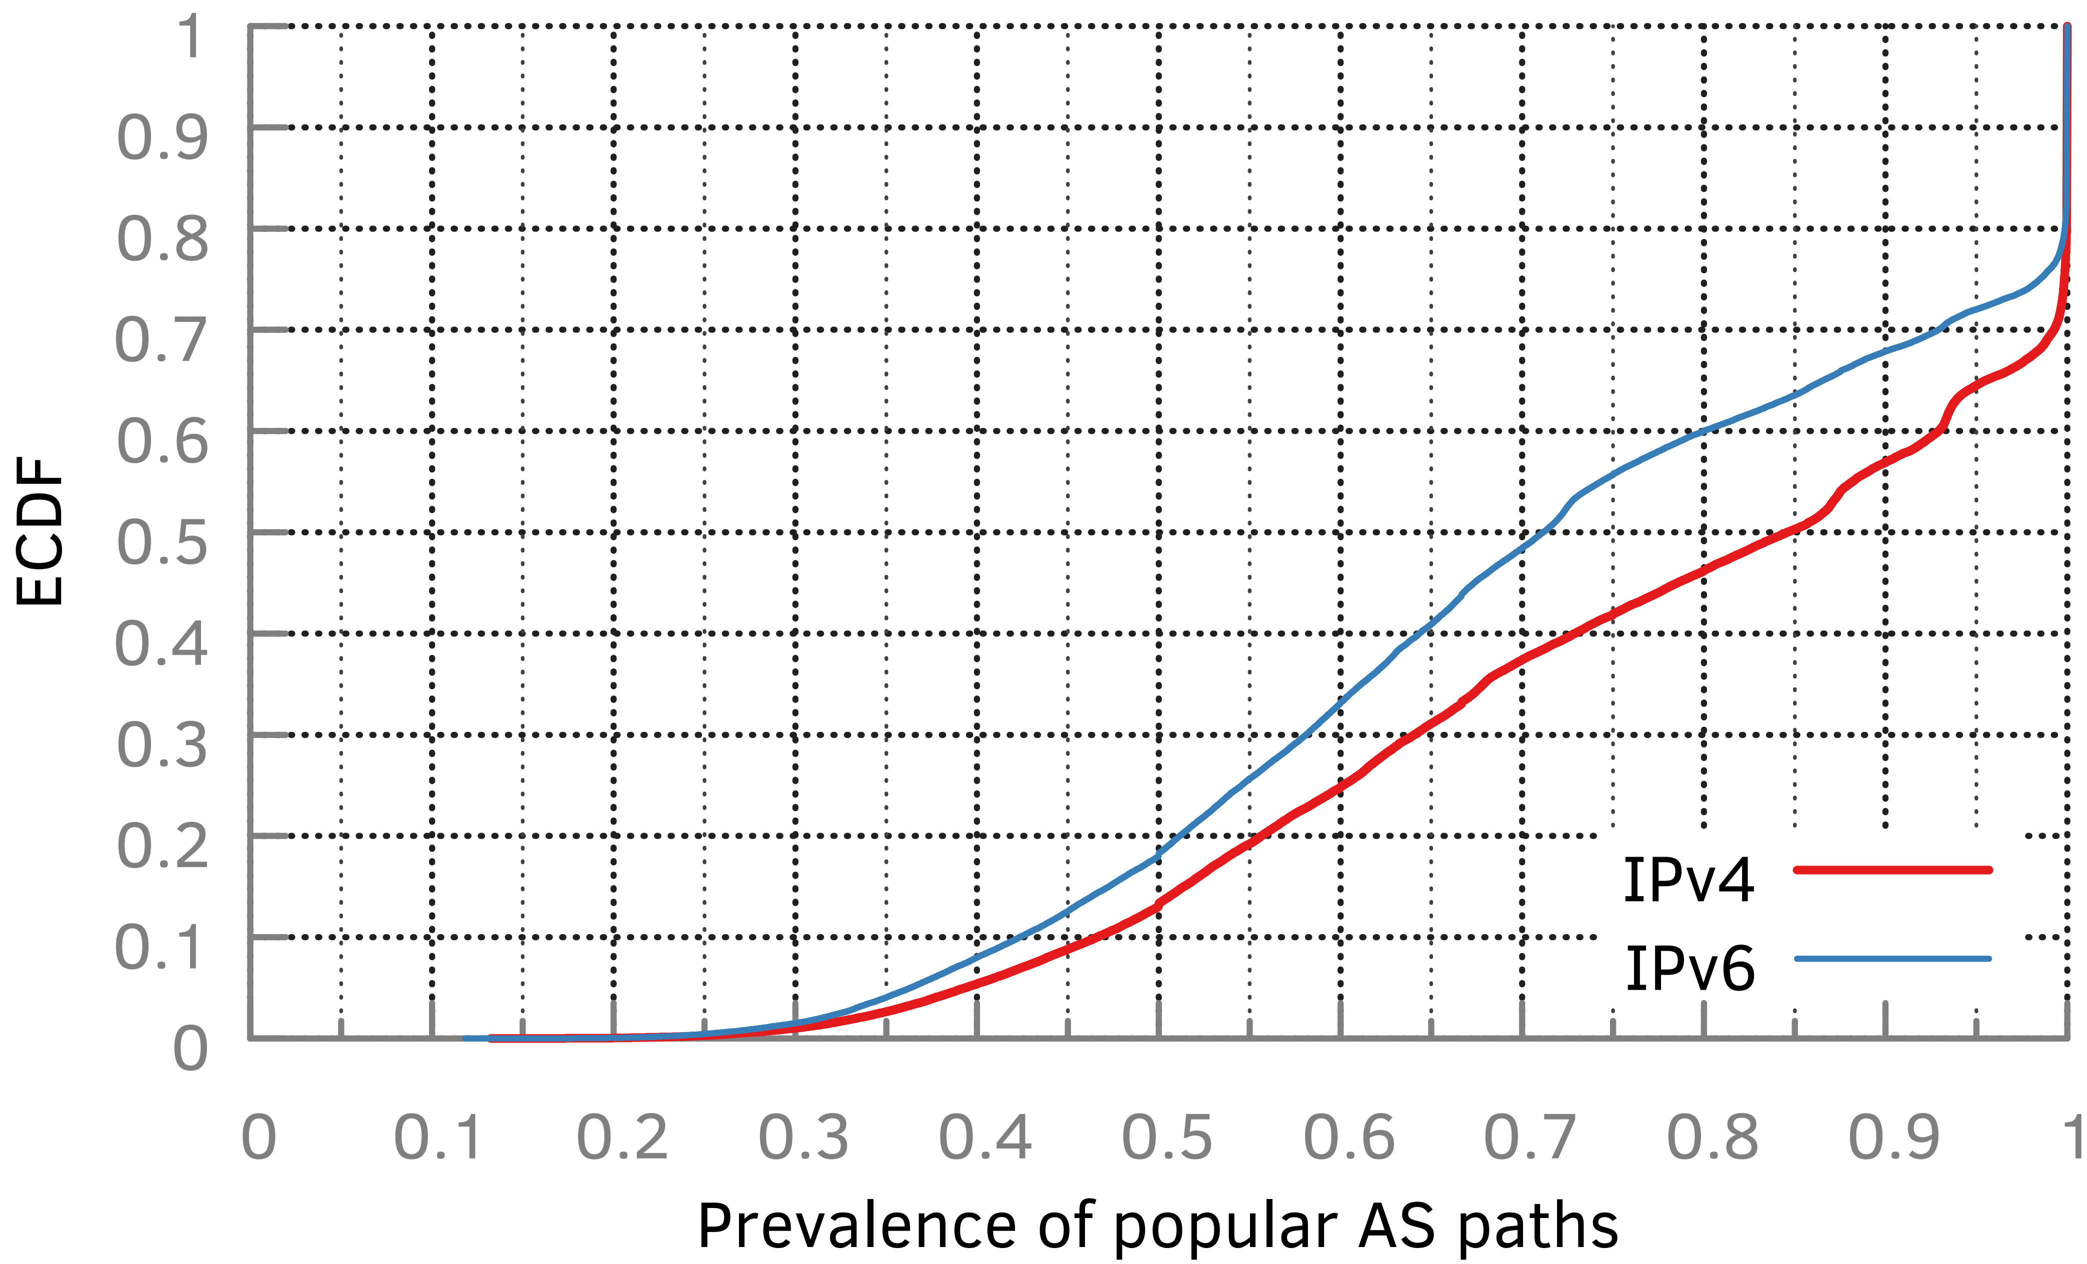

AS path ***prevalence*** — *Vern Paxson, IEEE/ACM Transactions on Networking 1997*

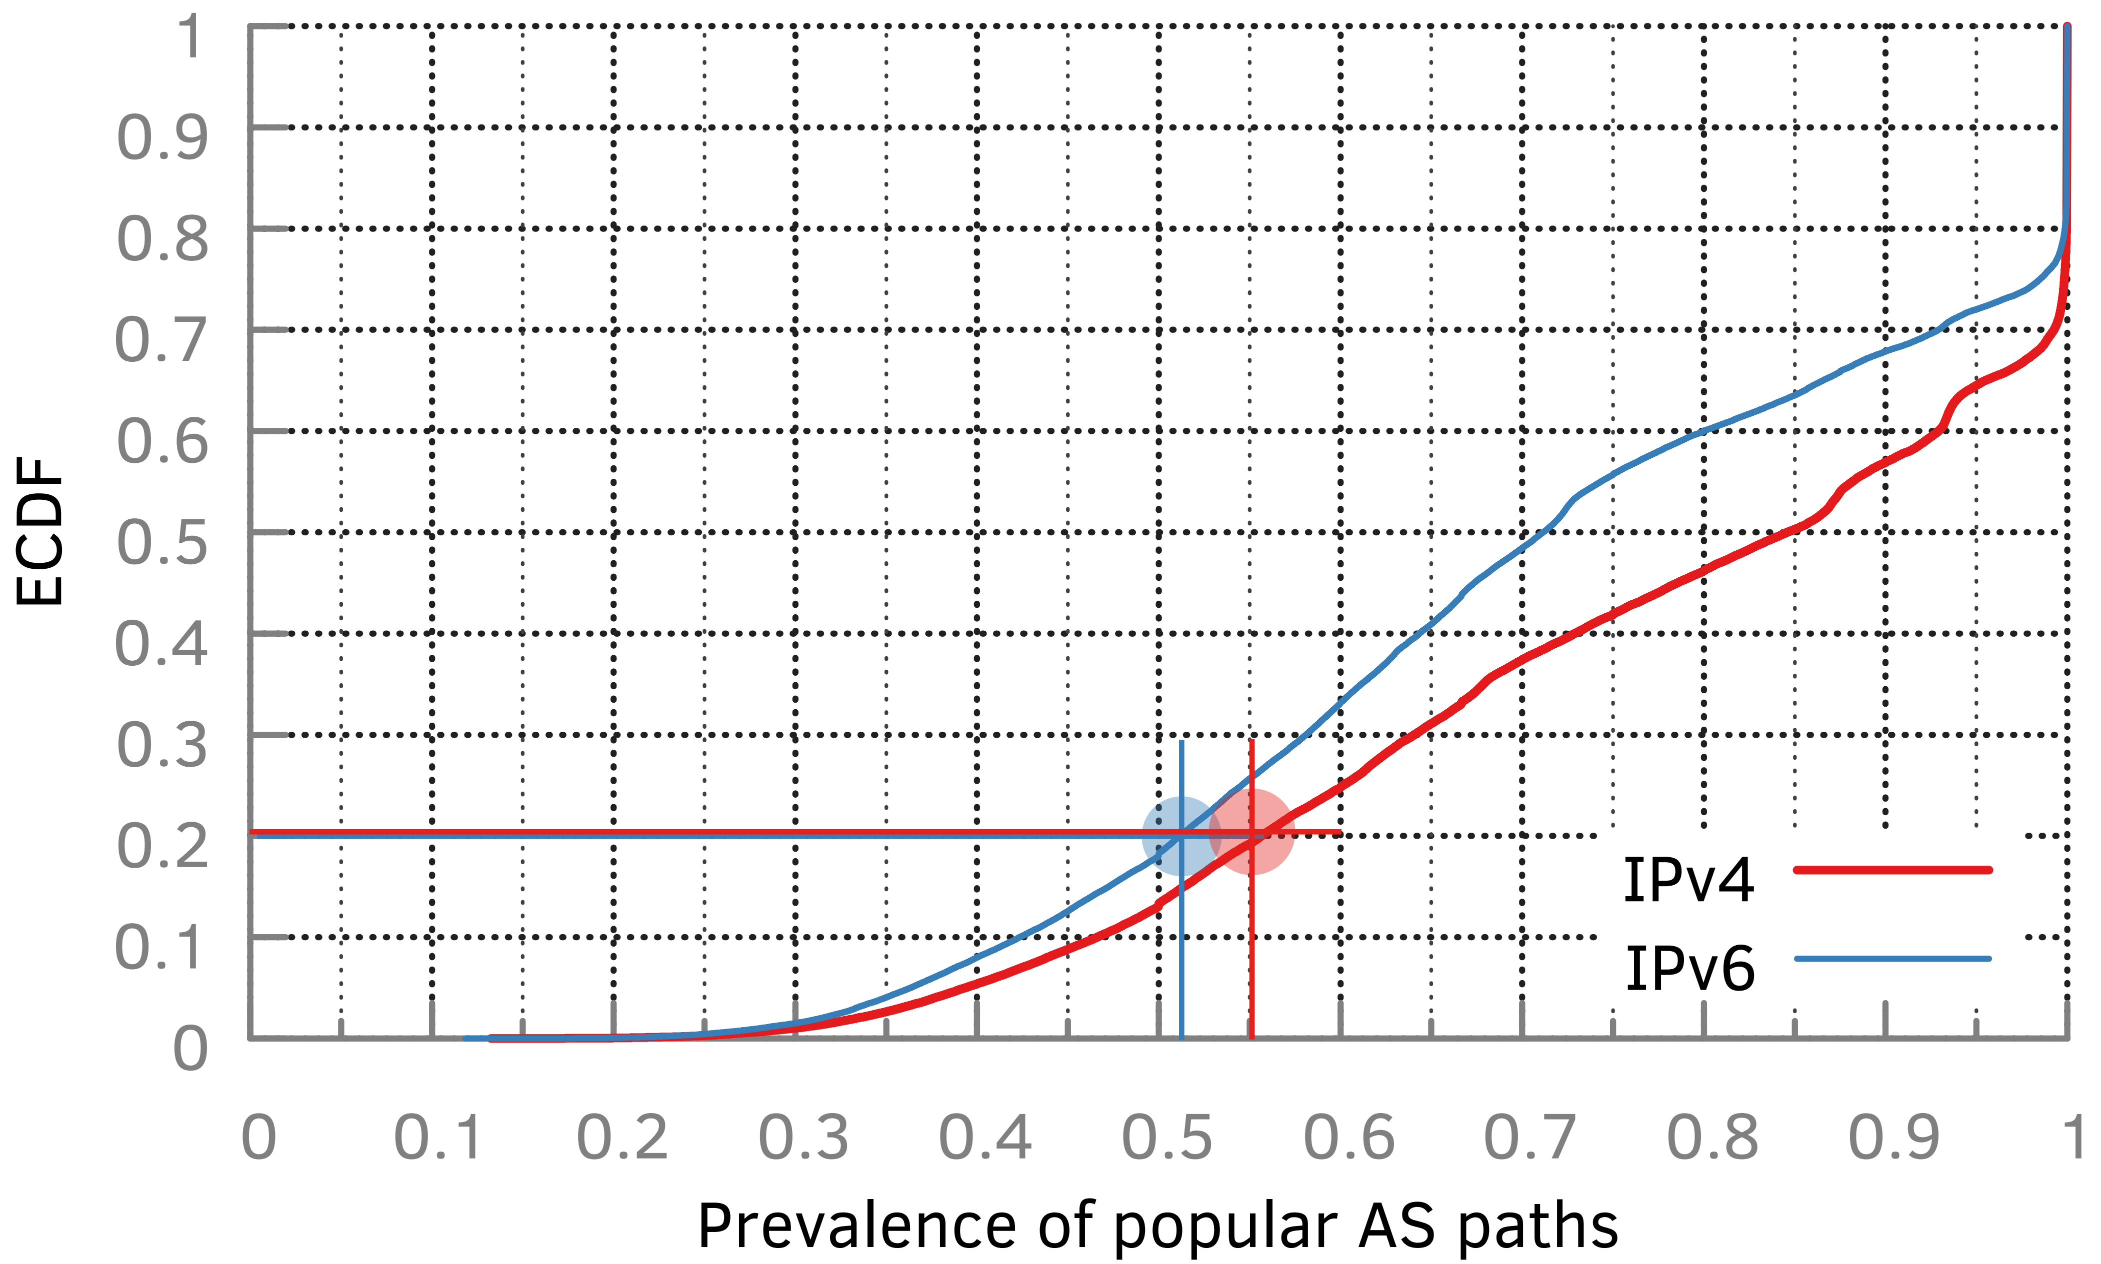

Most paths had one dominant route, with ***80%*** dominant for at least half the period.

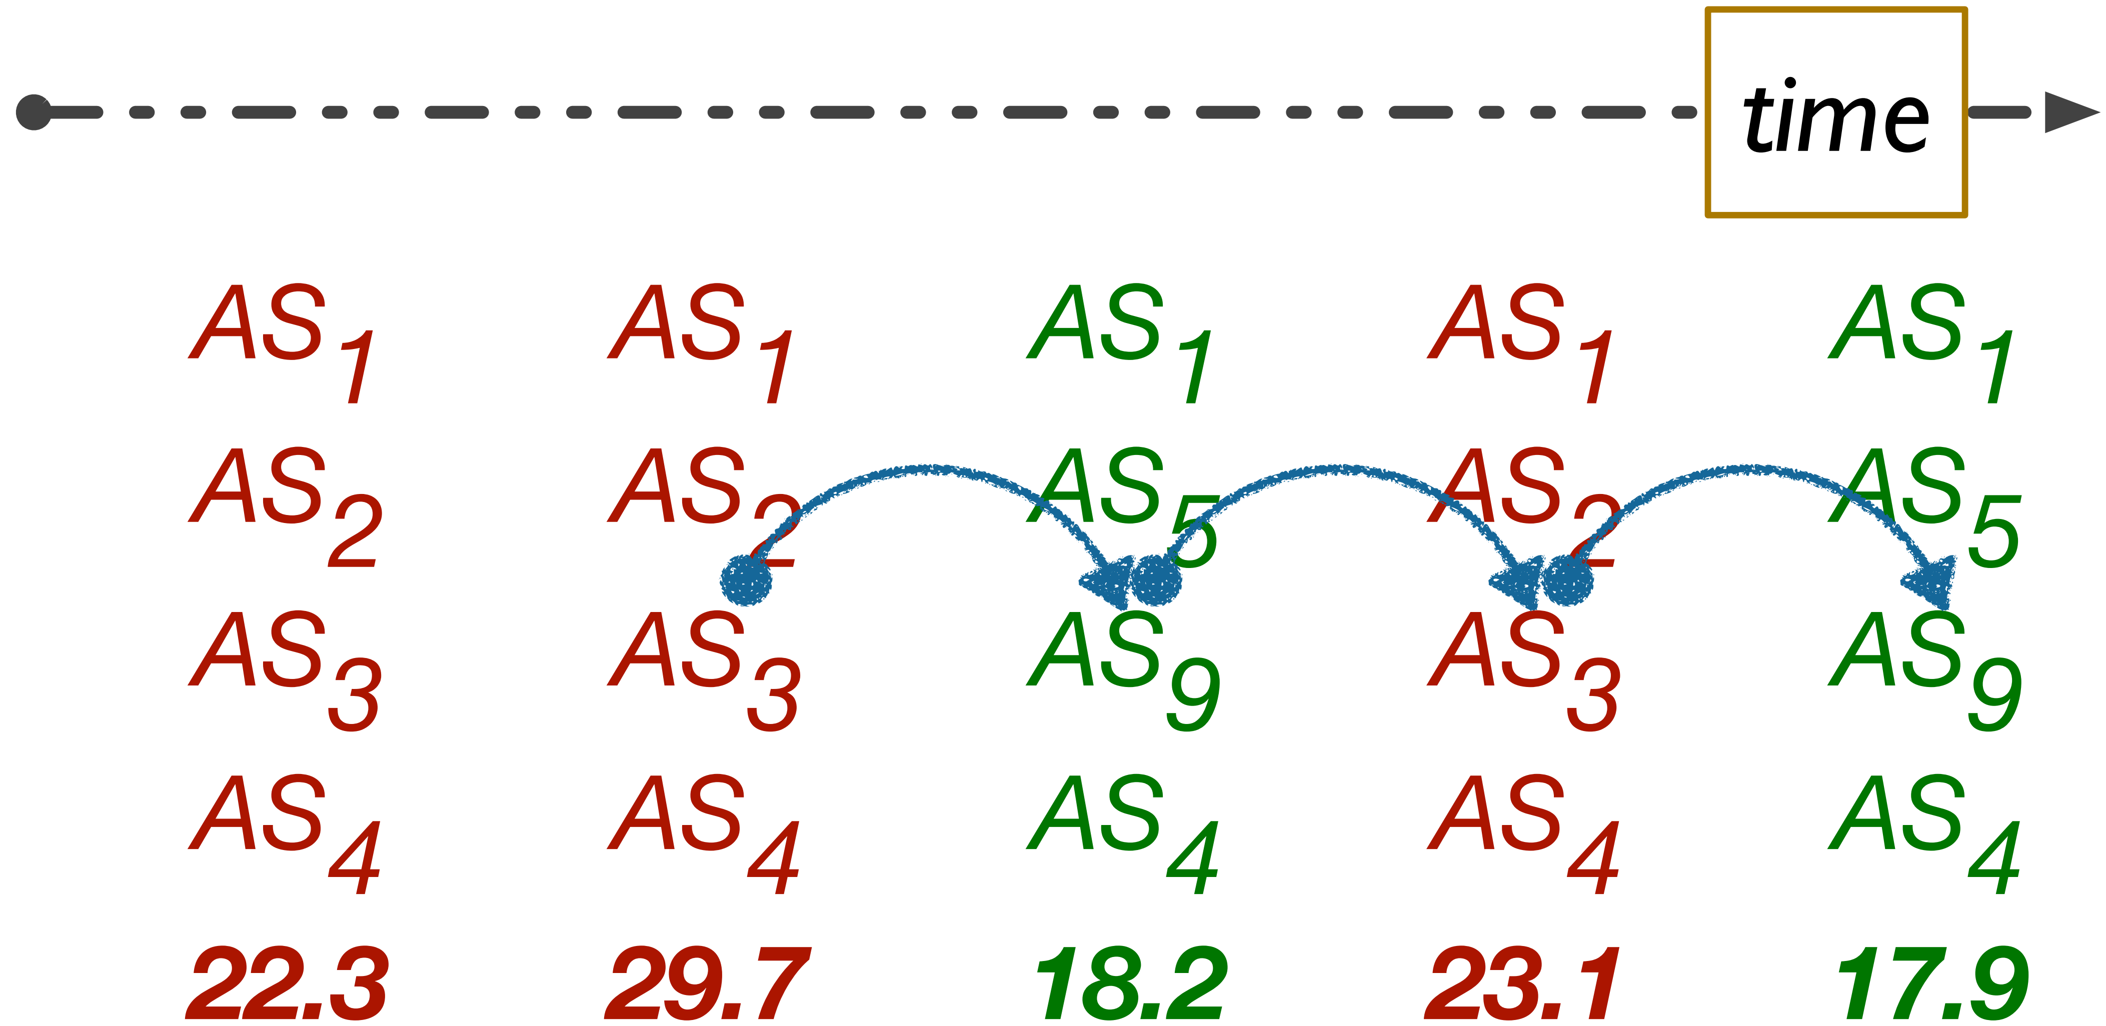

Number of AS-path changes observed in the A–B trace timeline

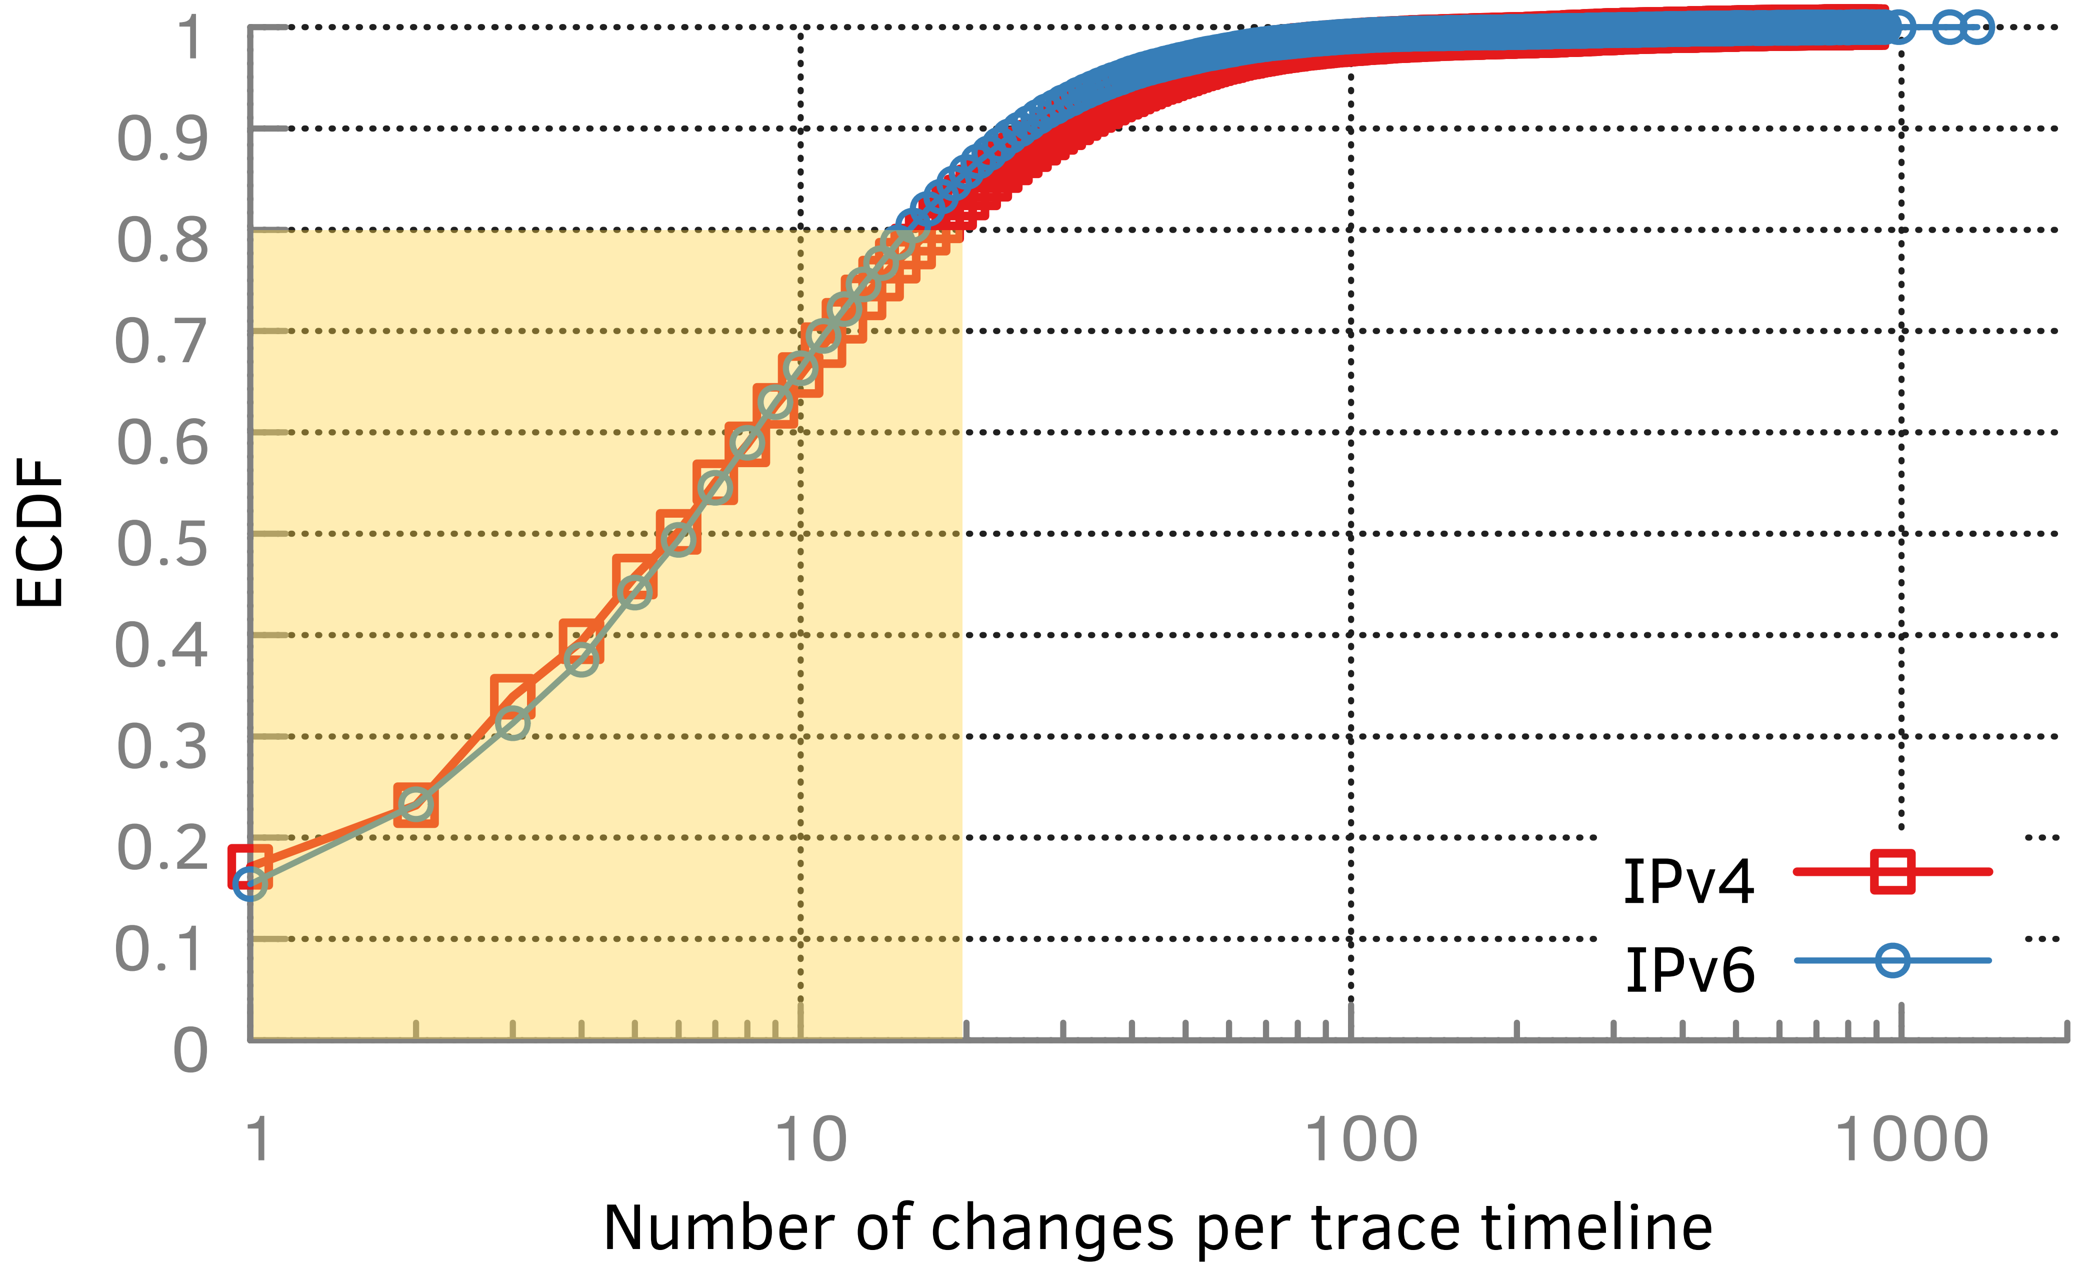

80% of the trace timelines experienced ***20*** or fewer changes over the course of ***16***-months.

How do the AS-path changes affect the ***baseline*** RTT of server-to-server paths?

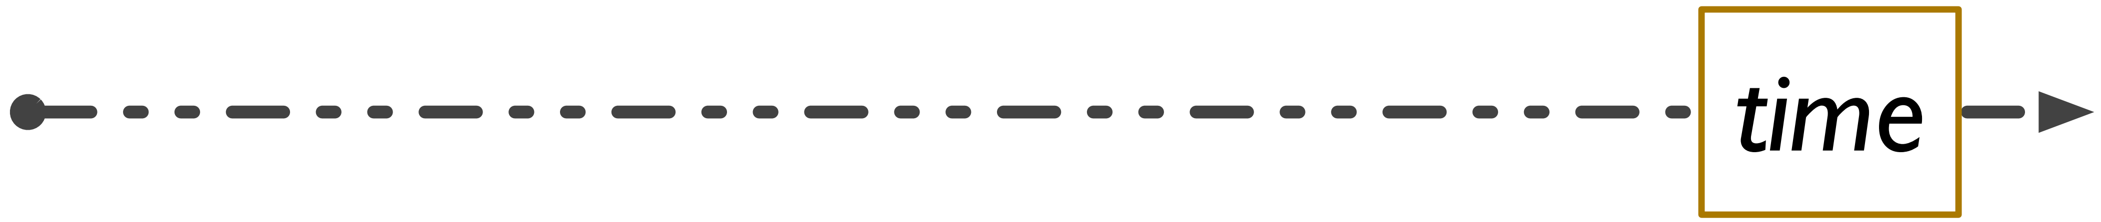

AS_1	AS_1	AS_1	AS_1	AS_1
AS_2	AS_2	AS_5	AS_2	AS_5
AS_3	AS_3	AS_9	AS_3	AS_9
AS_4	AS_4	AS_4	AS_4	AS_4
22.3	**29.7**	**18.2**	**23.1**	**17.9**

Group RTTs by AS paths.

Baseline: ***10th-percentile*** of each AS-path (bucket).

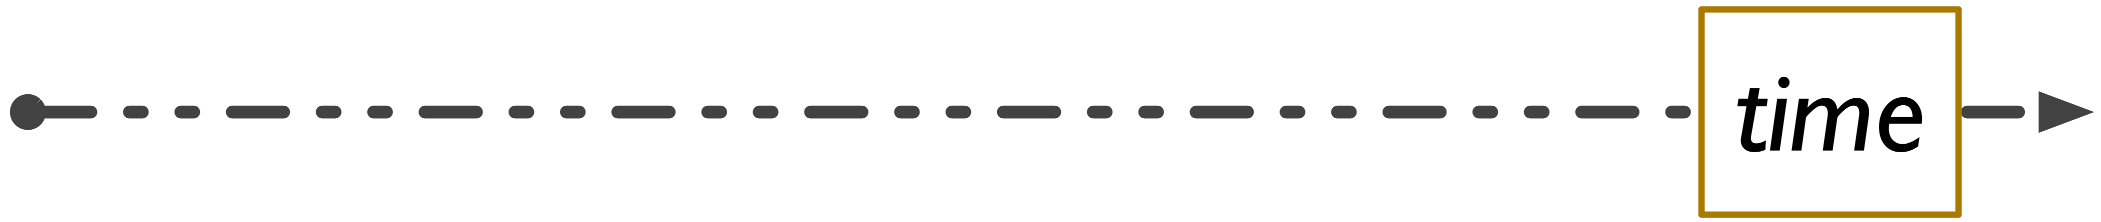

AS_1	AS_1	AS_1	AS_1	AS_1
AS_2	AS_2	AS_5	AS_2	AS_5
AS_3	AS_3	AS_9	AS_3	AS_9
AS_4	AS_4	AS_4	AS_4	AS_4
22.3	**29.7**	**18.2**	**23.1**	**17.9**

Optimal Path: path with lowest baseline.

Optimal: ***AS_1-AS_5-AS_9-AS_4***

Sub-Optimal: ***AS_1-AS_2-AS_3-AS_4***

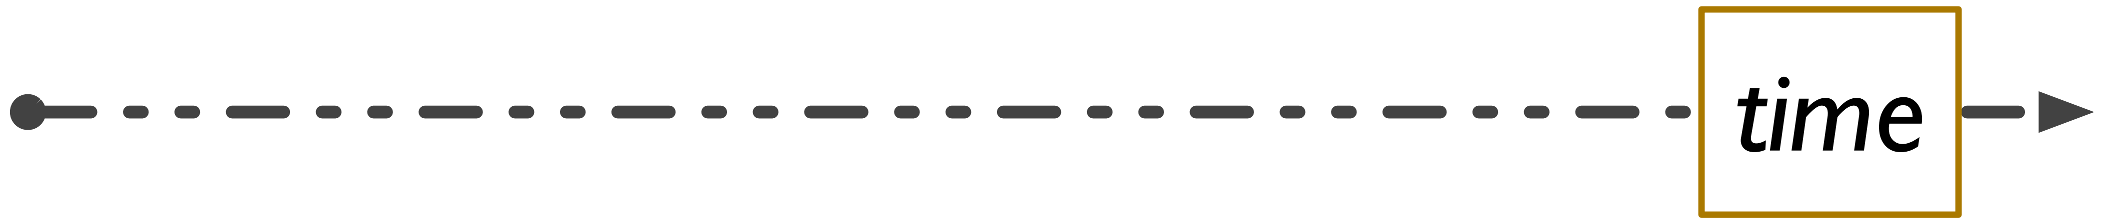

AS_1	AS_1	AS_1	AS_1	AS_1
AS_2	AS_2	AS_5	AS_2	AS_5
AS_3	AS_3	AS_9	AS_3	AS_9
AS_4	AS_4	AS_4	AS_4	AS_4
22.3	***29.7***	***18.2***	***23.1***	***17.9***

Baseline of sub-optimal path with prevalence of ***60%*** is ***~4.5 ms*** increase in end-to-end RTT.

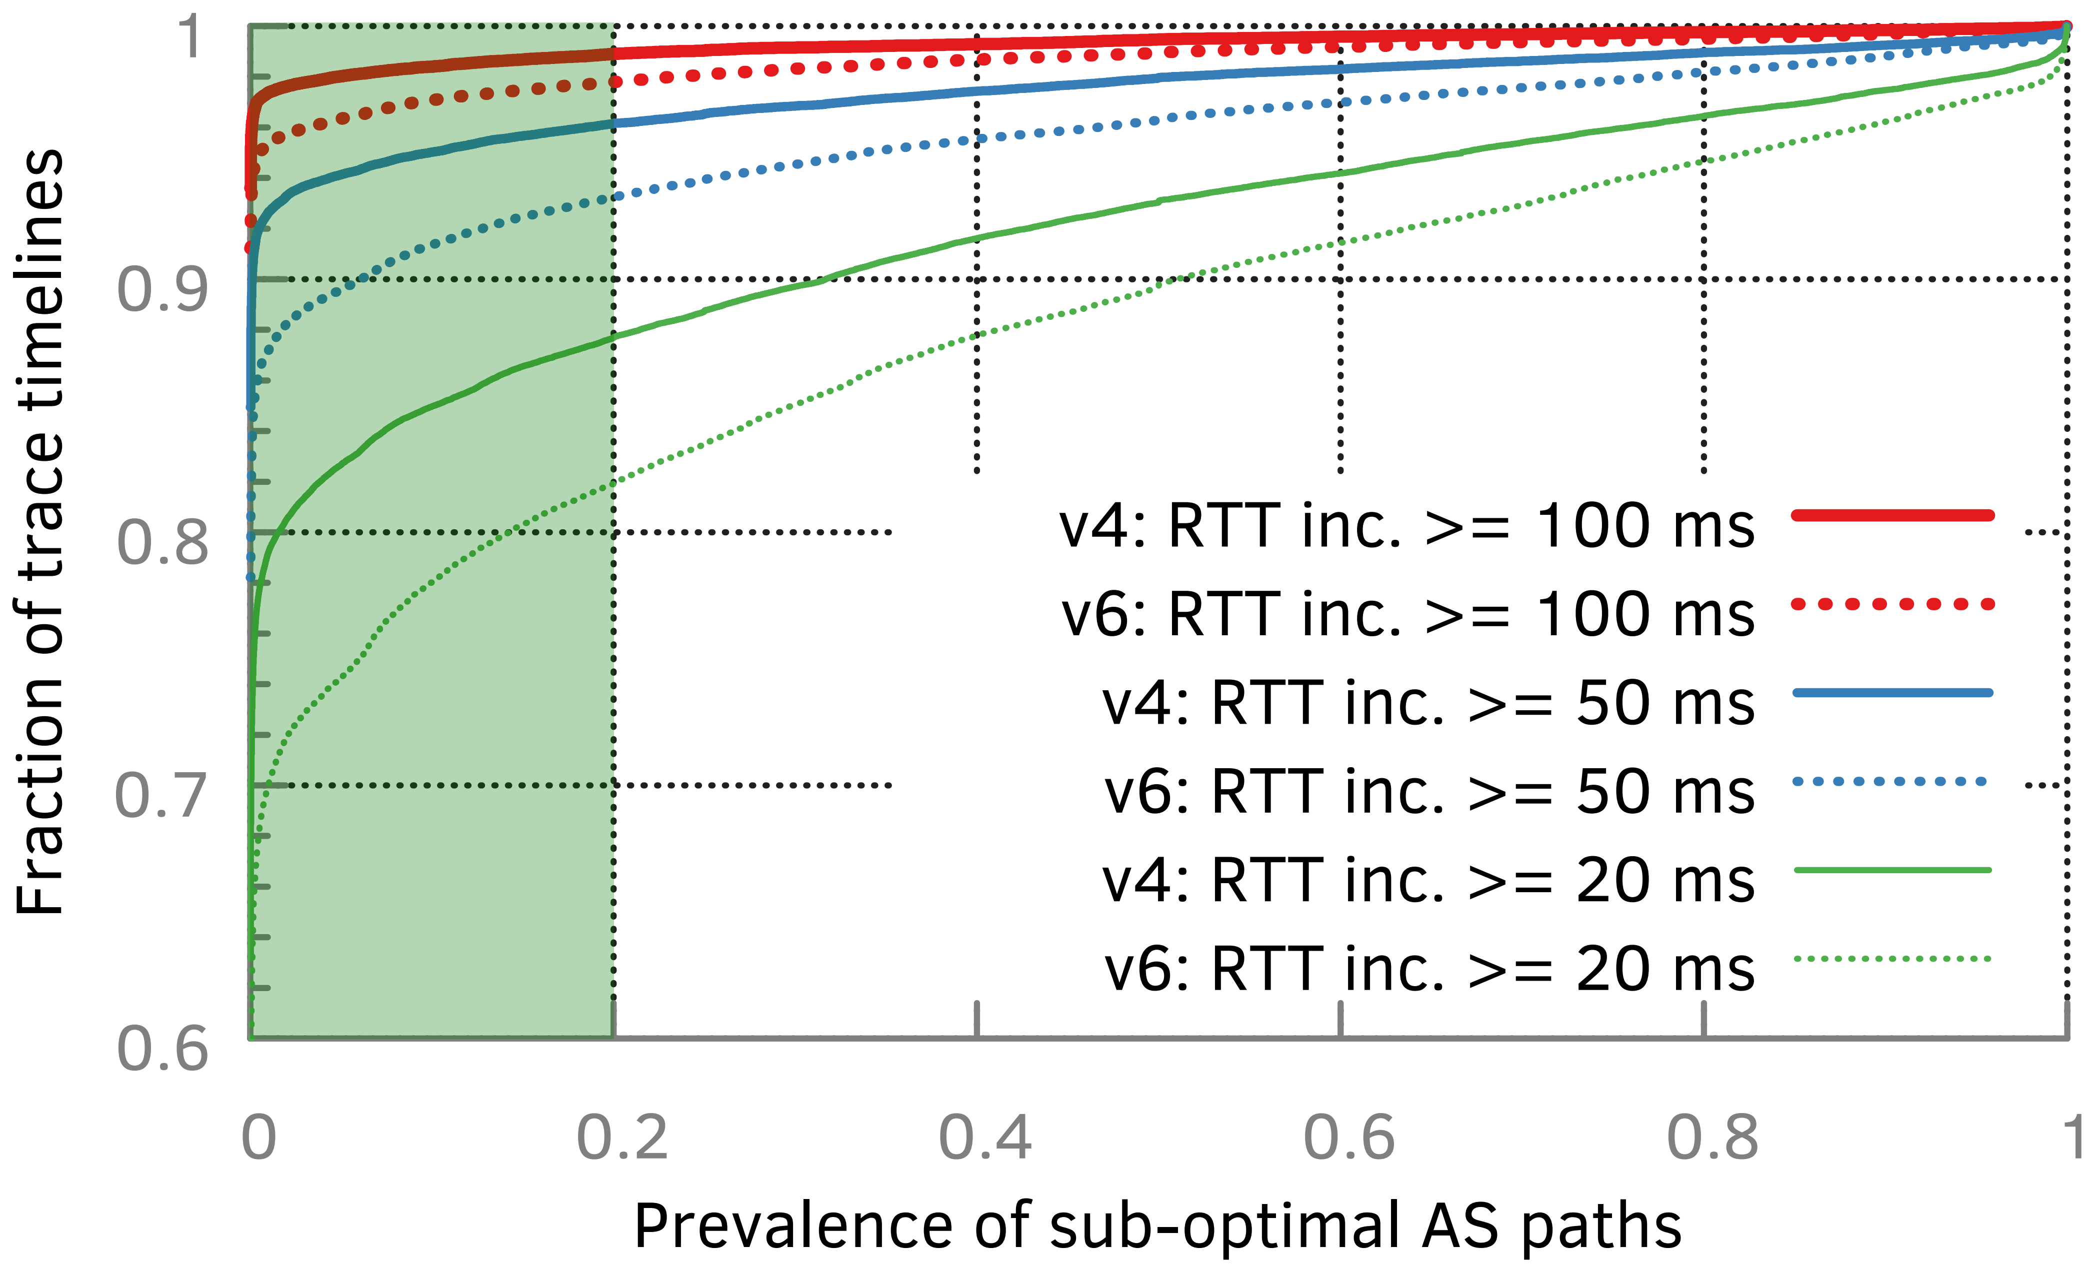

Typically a routing change causes only a small change in RTT.

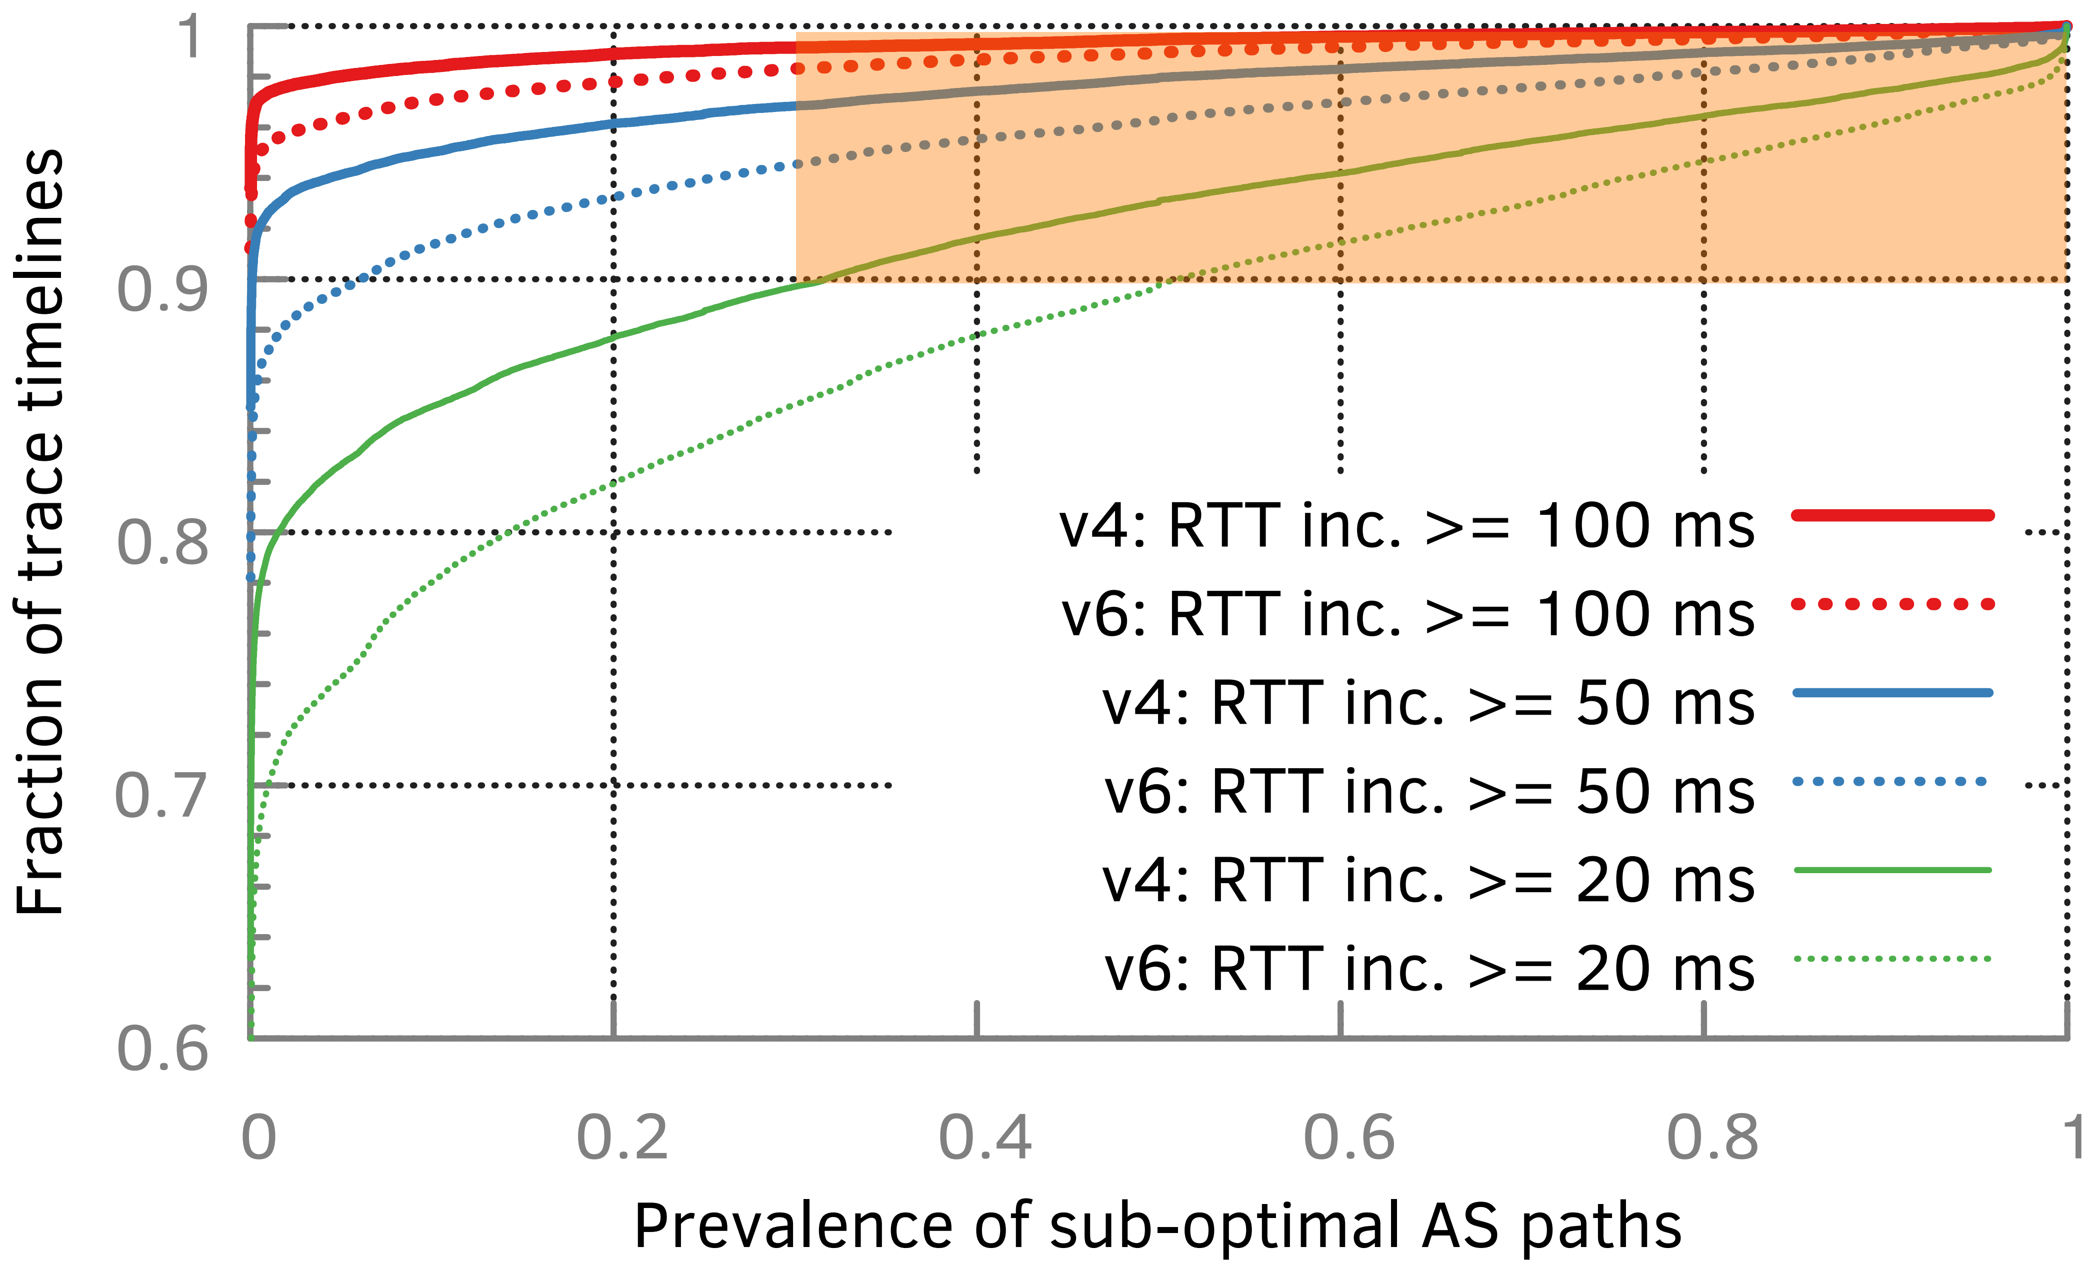

But for a minority of cases, the change can be significant.

10% of trace timelines over IPv4 the (sub-optimal) AS paths that led to at least a ***20* ms** increase in RTTs had a prevalence of at least ***30%***

Effect of *periods of daily oscillation* on end-to-end RTTs

Data Set: Short Term

- ≈***3,500*** server clusters in ***1,000*** locations in ***100*** different countries.

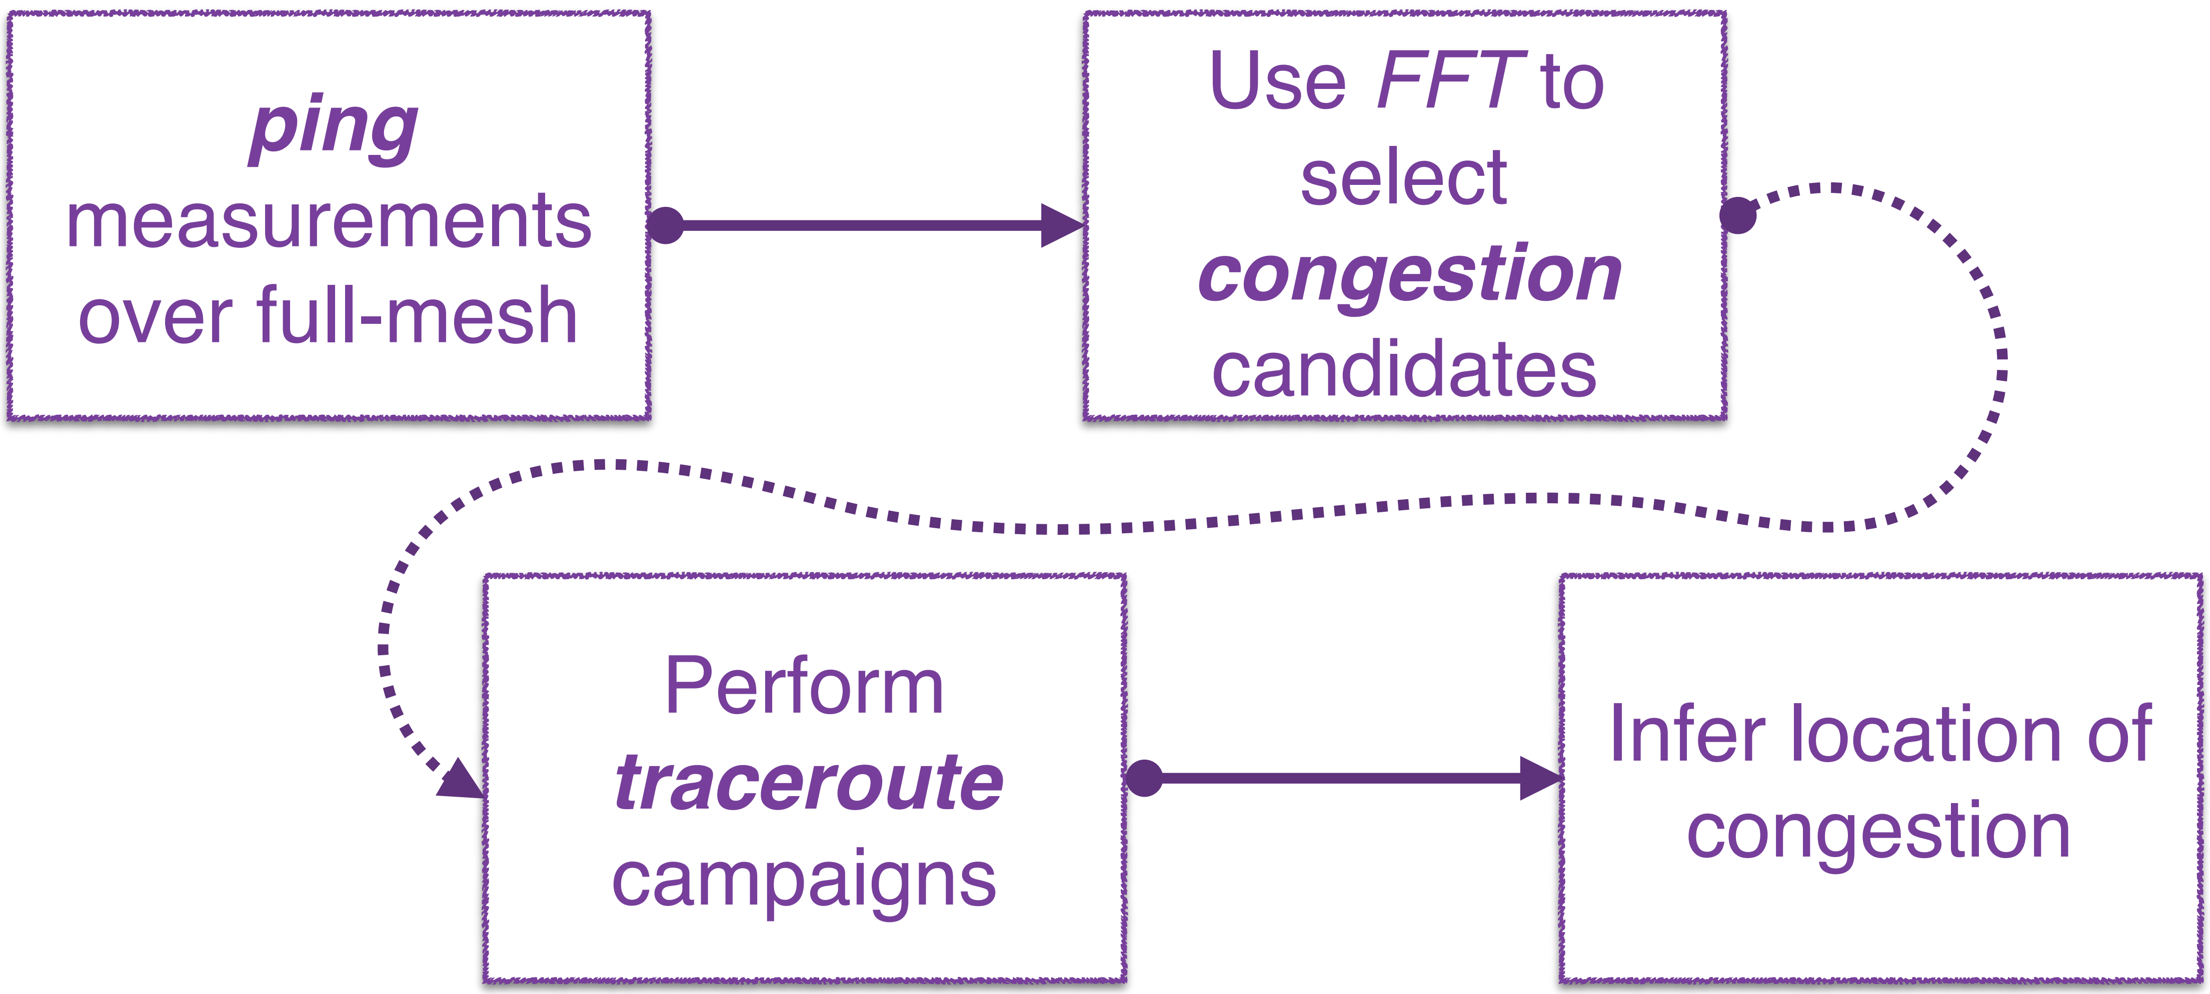

ping measurements every ***15 minutes*** for one week from Feb. 22, 2015 through Feb. 28, 2015.

≈***2.9M*** IPv4 and ≈***1M*** IPv6 server pairs

Based on *Time Sequence Latency Probes* by *Luckie et al., IMC 2014*

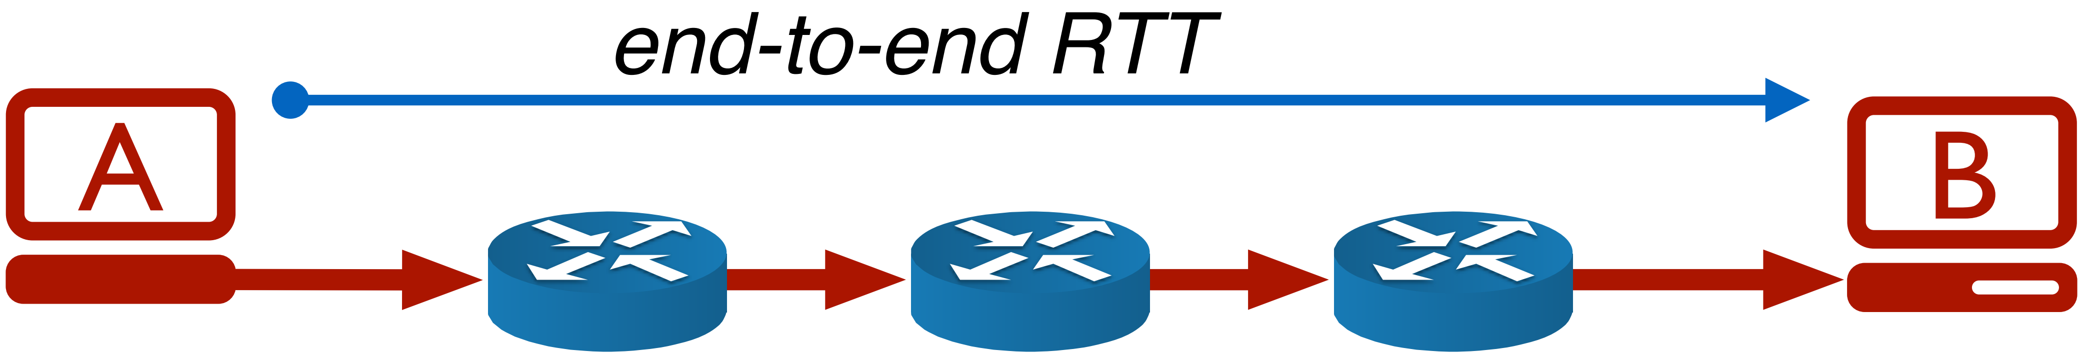
end-to-end RTT
A
B

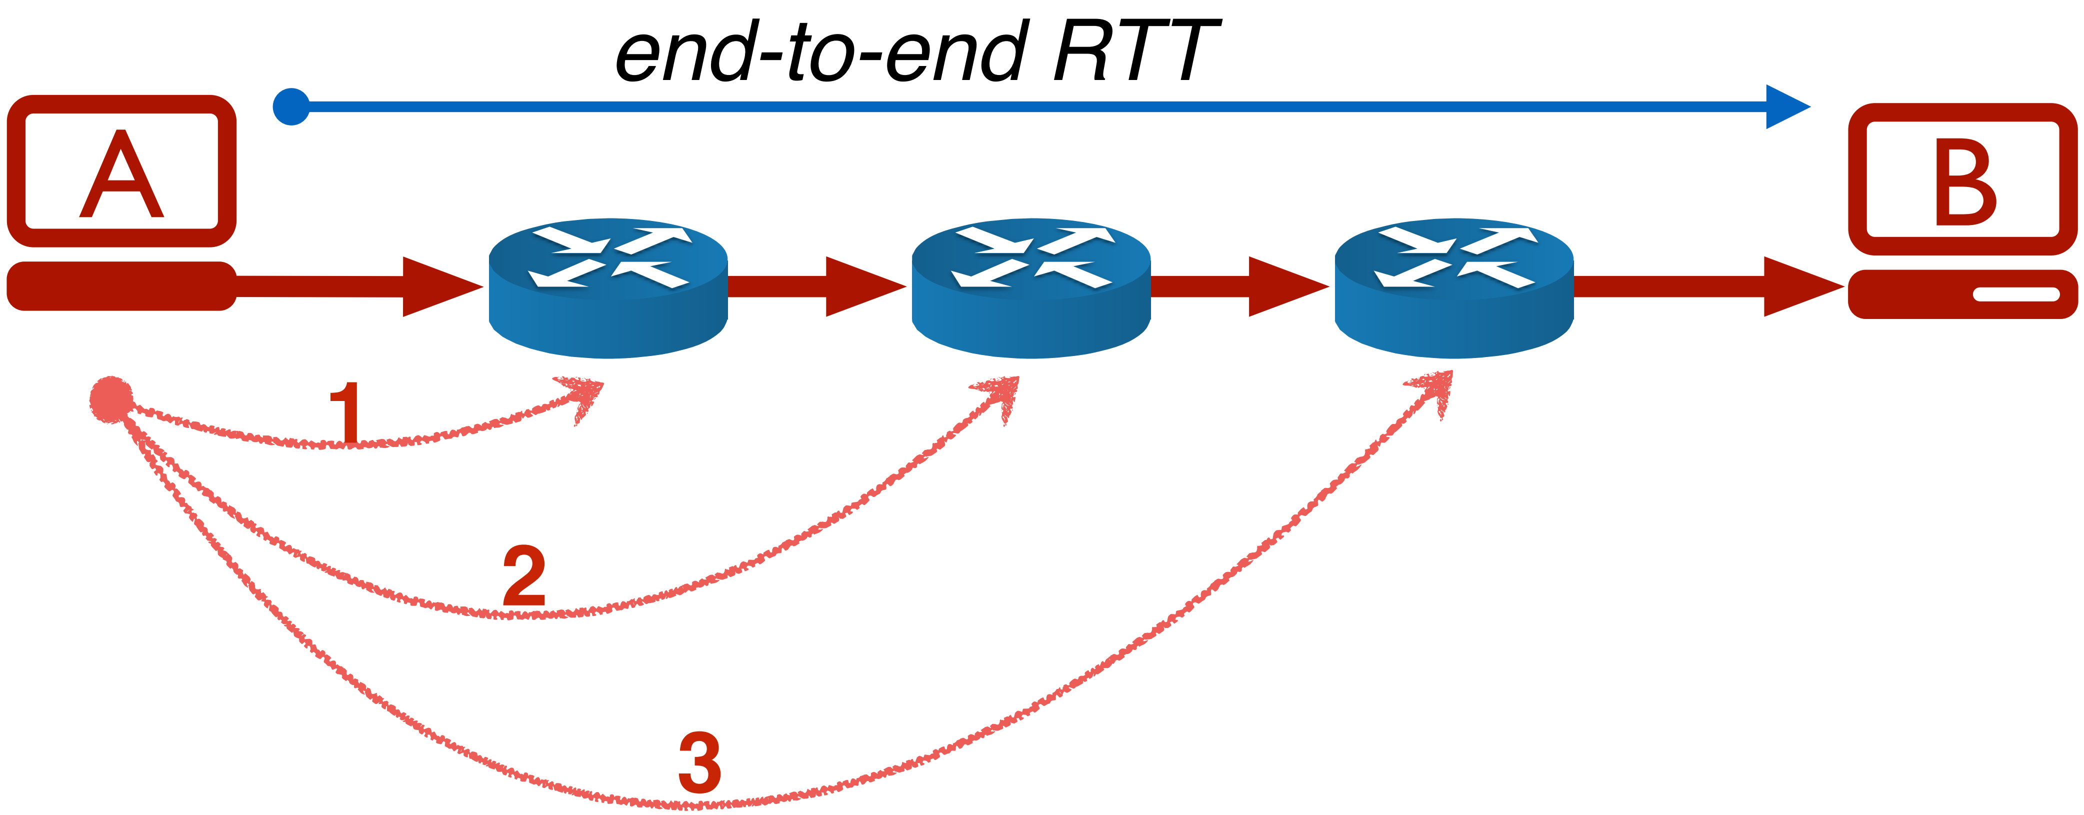

Identify first ***segment*** with high-correlation with end-to-end RTT?

Highlights

3155 *links were congested in our study of IPv4 traceroutes.*

1768 internal & ***1121 interconnection*** links.

Weighting links by the number of server-to-server paths that cross them …

interconnection links are more popular!

Large majority of the interconnection links with congestion were ***private interconnects***.

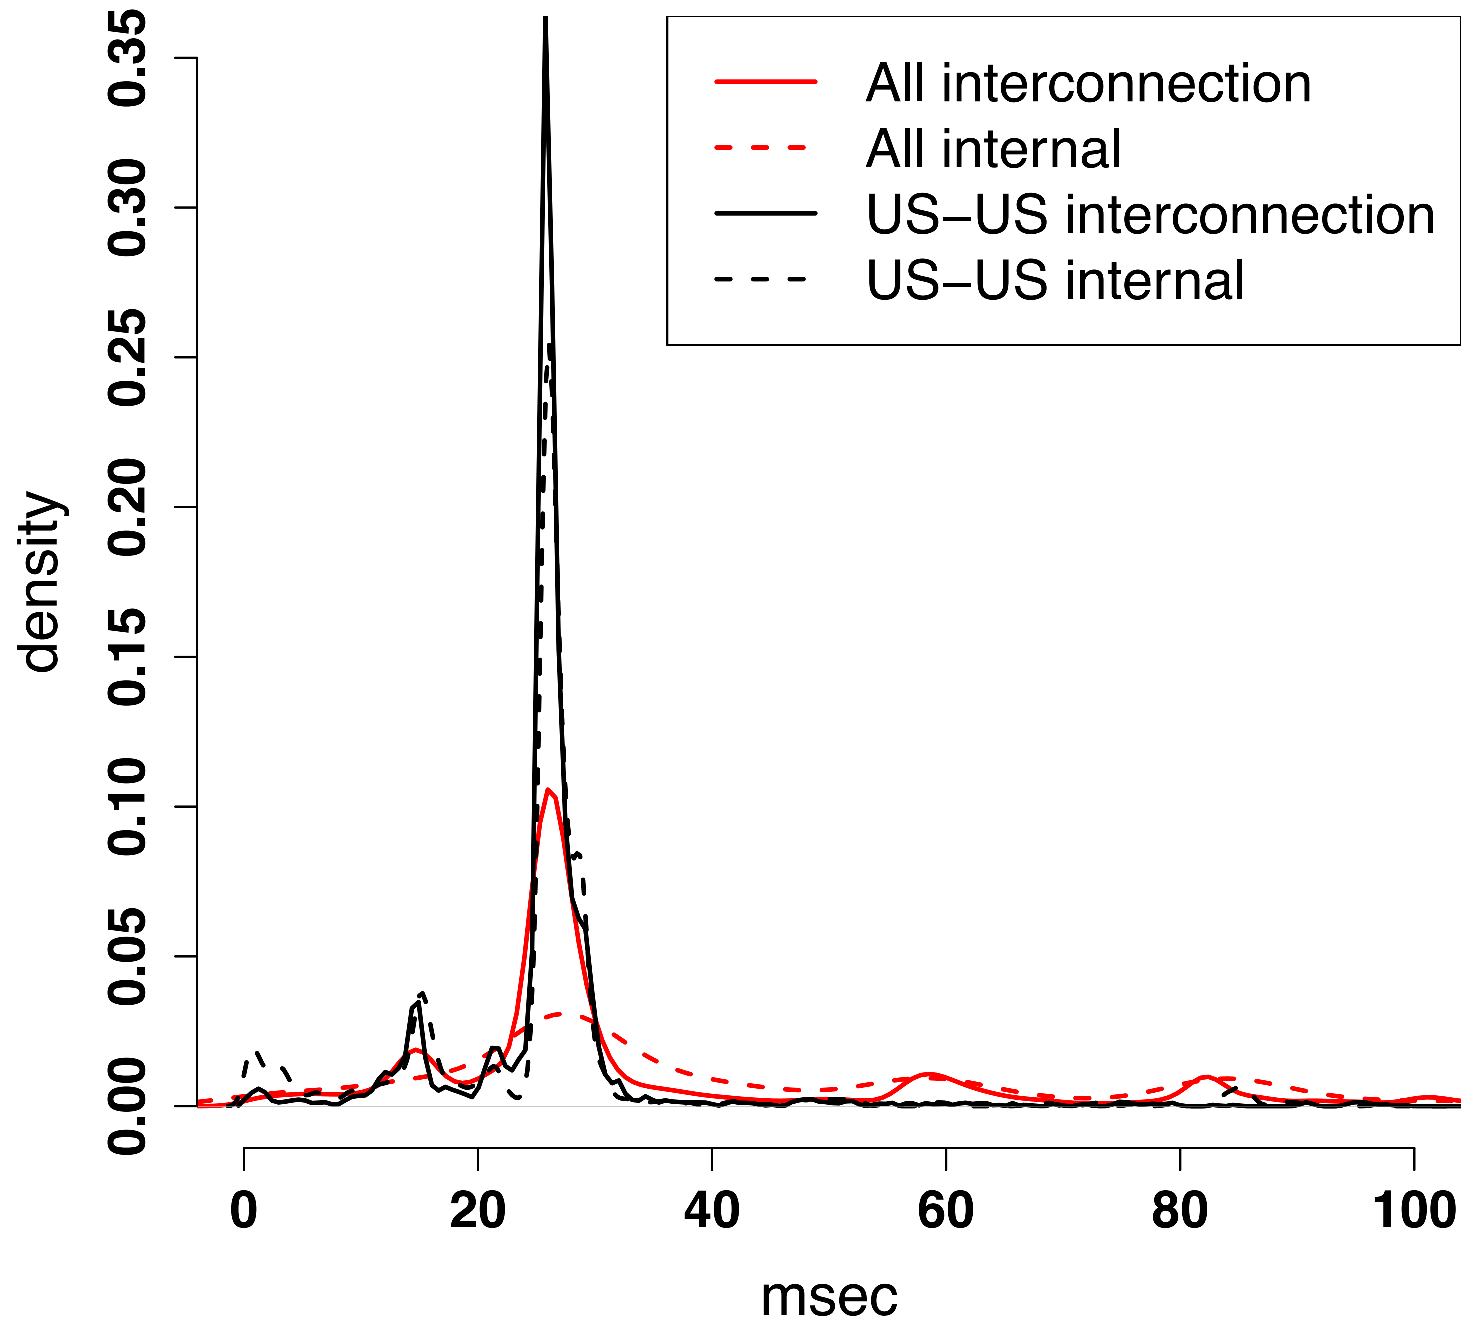

Typical overhead due to congestion is ***20-30 ms***.

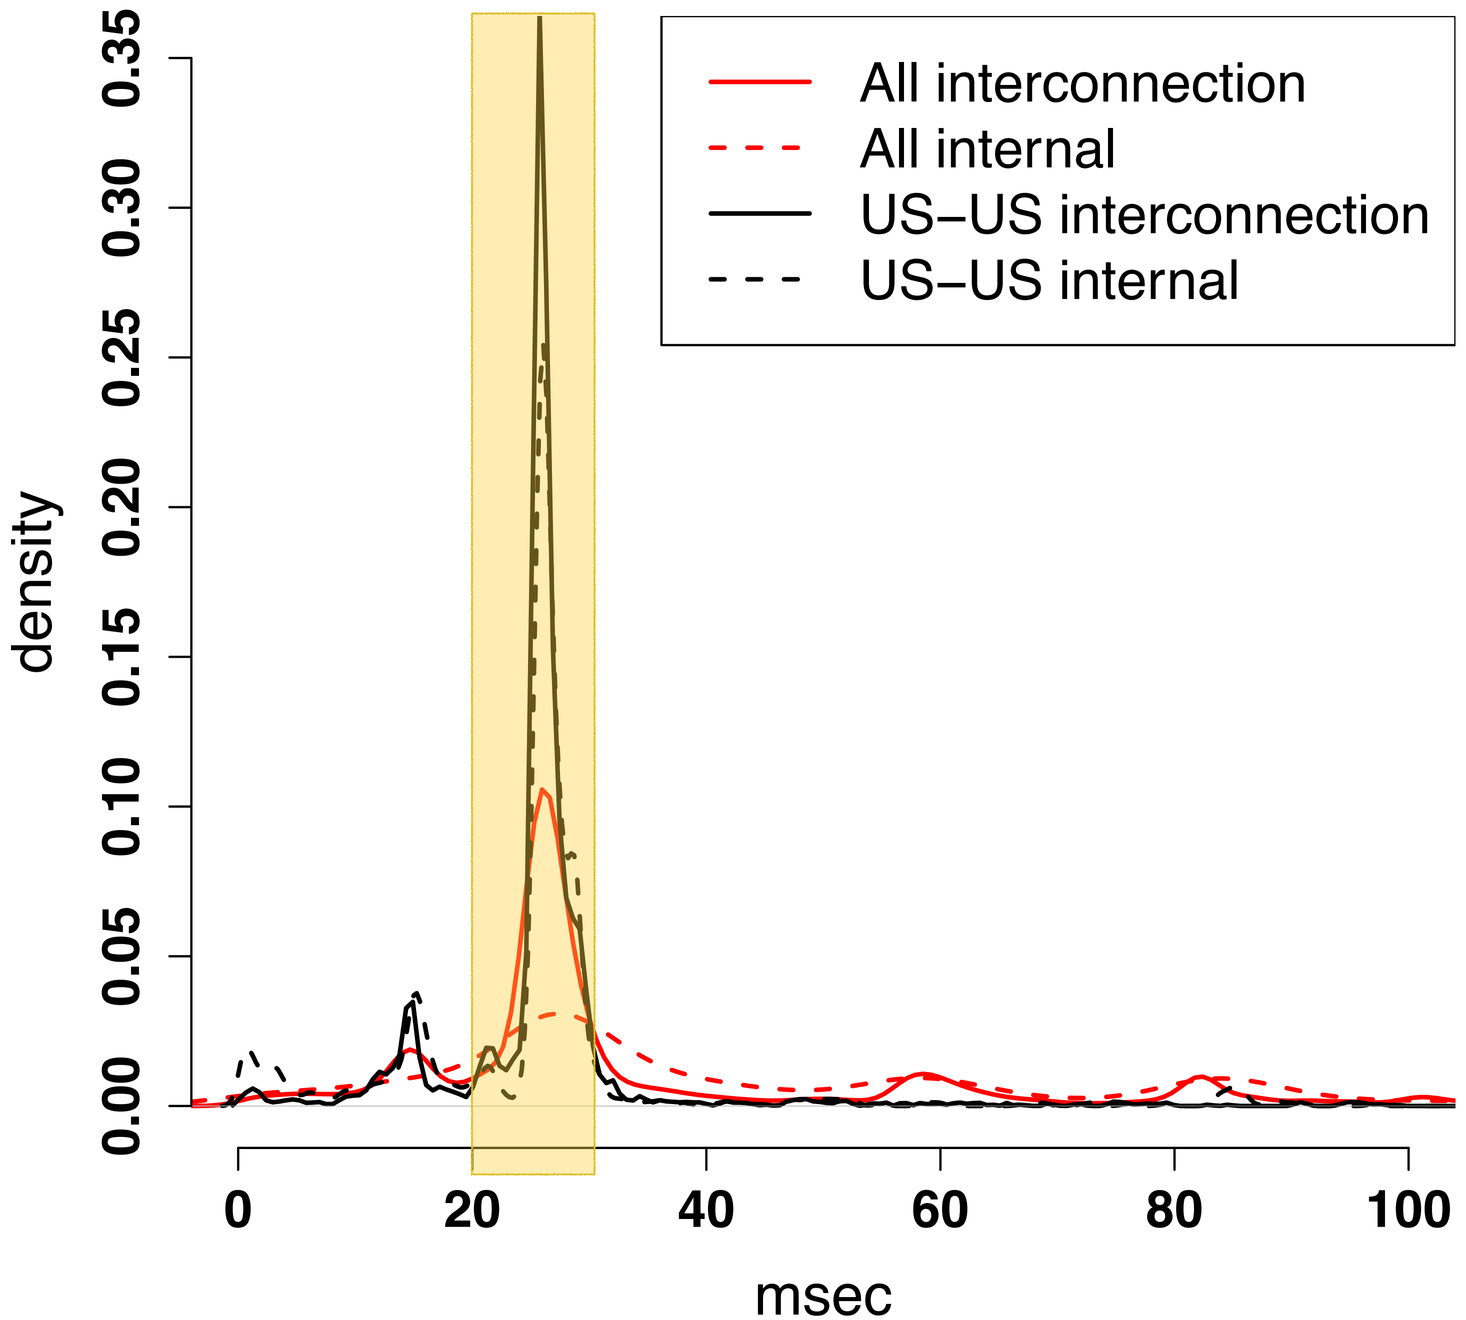

Values between ***20-30 ms*** —

US: accounts for ***90%*** of density.

Europe & Asia: accounts for ***30%*** of density.

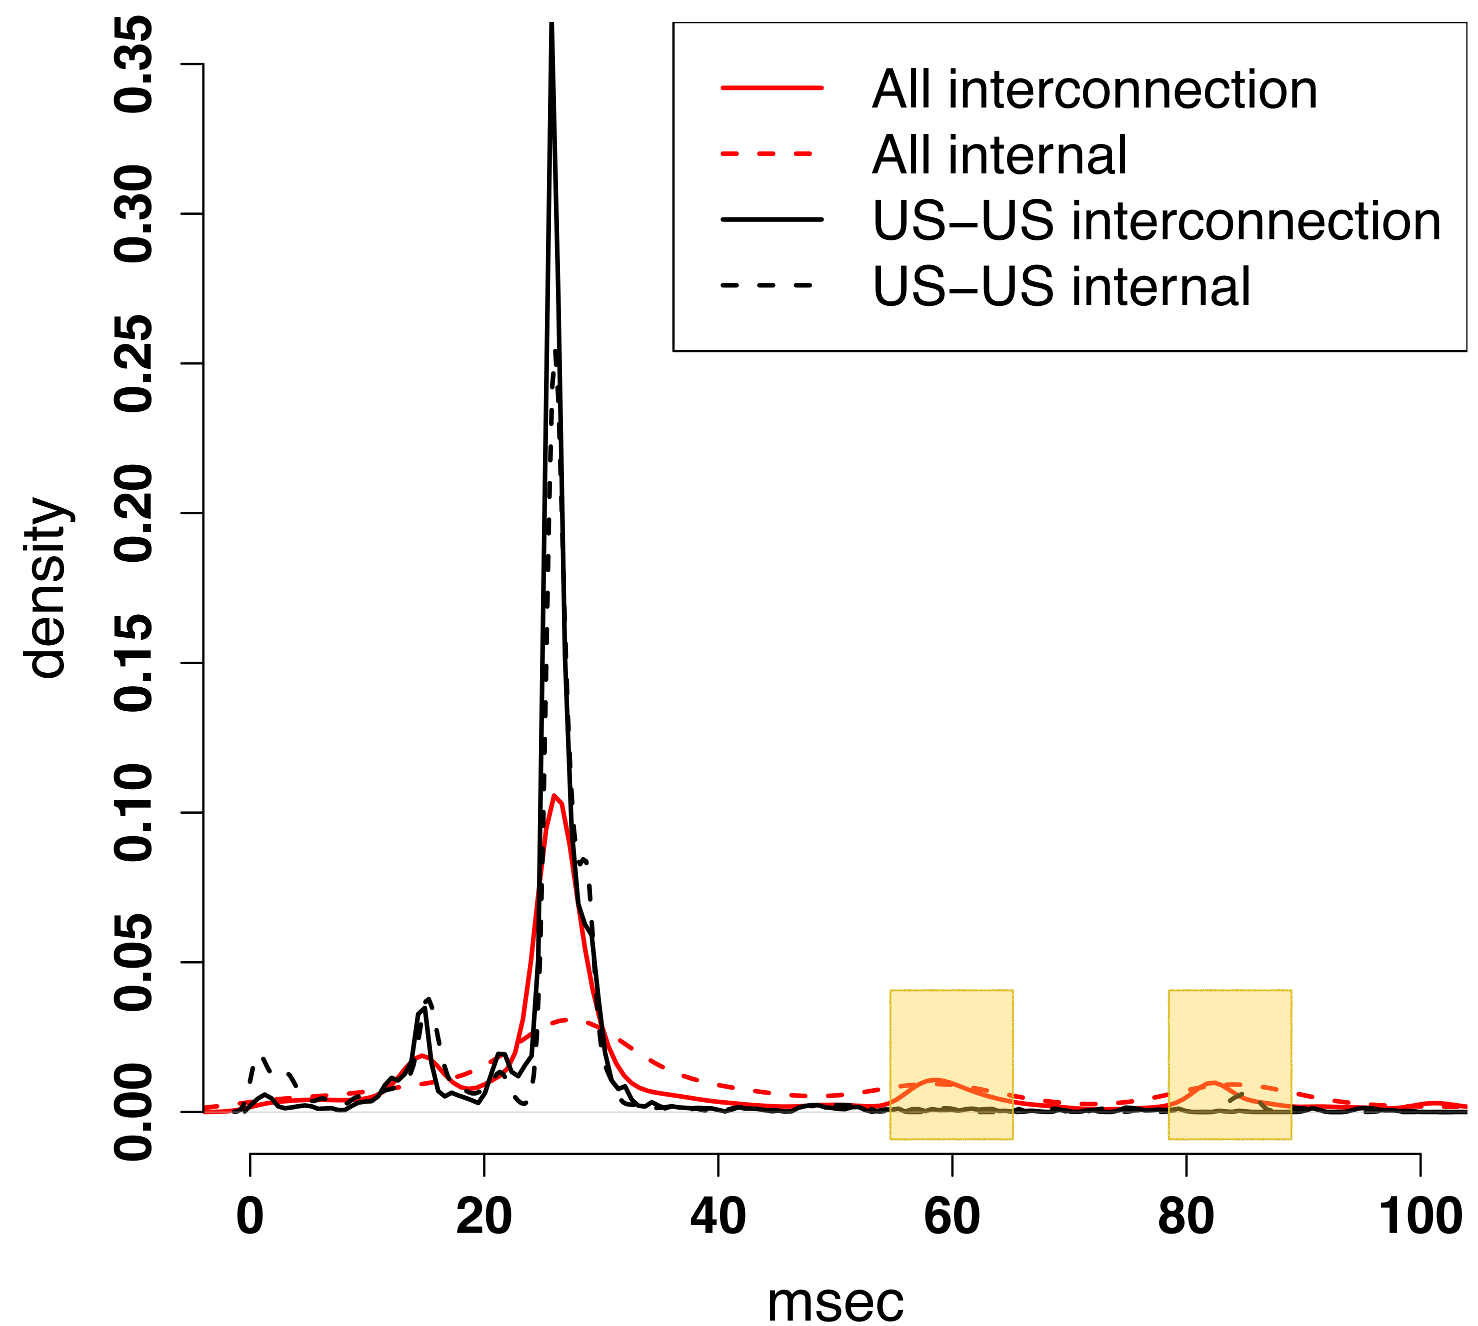

Transcontinental links in ***Europe*** & ***Asia***.

Routing changes typically do not affect end-to-end RTTs.
Congestion is not the norm.

What about ***non-typical*** cases?

Routing

For ***10%*** of server pairs the (sub-optimal) AS paths that led to ***20 ms*** increase in RTTs pertained for at least ***30%*** of the study period for IPv4 & ***50%*** for IPv6.

Congestion

Only ***2%*** of the server pairs over IPv4, and just ***0.6%*** over IPv6, experience a strong diurnal pattern with an increase in RTT of least ***10 ms***.

Routing

10% of trace timelines the (sub-optimal) AS paths that led to at least ***20 ms*** increase in RTTs pertained for at least ***30%*** of the study period for IPv4 & ***50%*** for IPv6.

Congestion

Only

pairs over IPv4, and just

0.6%

experience a strong diurnal pattern with an increase in RTT of least

10 ms

- Focus on bandwidth
- No packet loss measurements; platform limitations
- Explore IPv4 & IPv6 infrastructure sharing

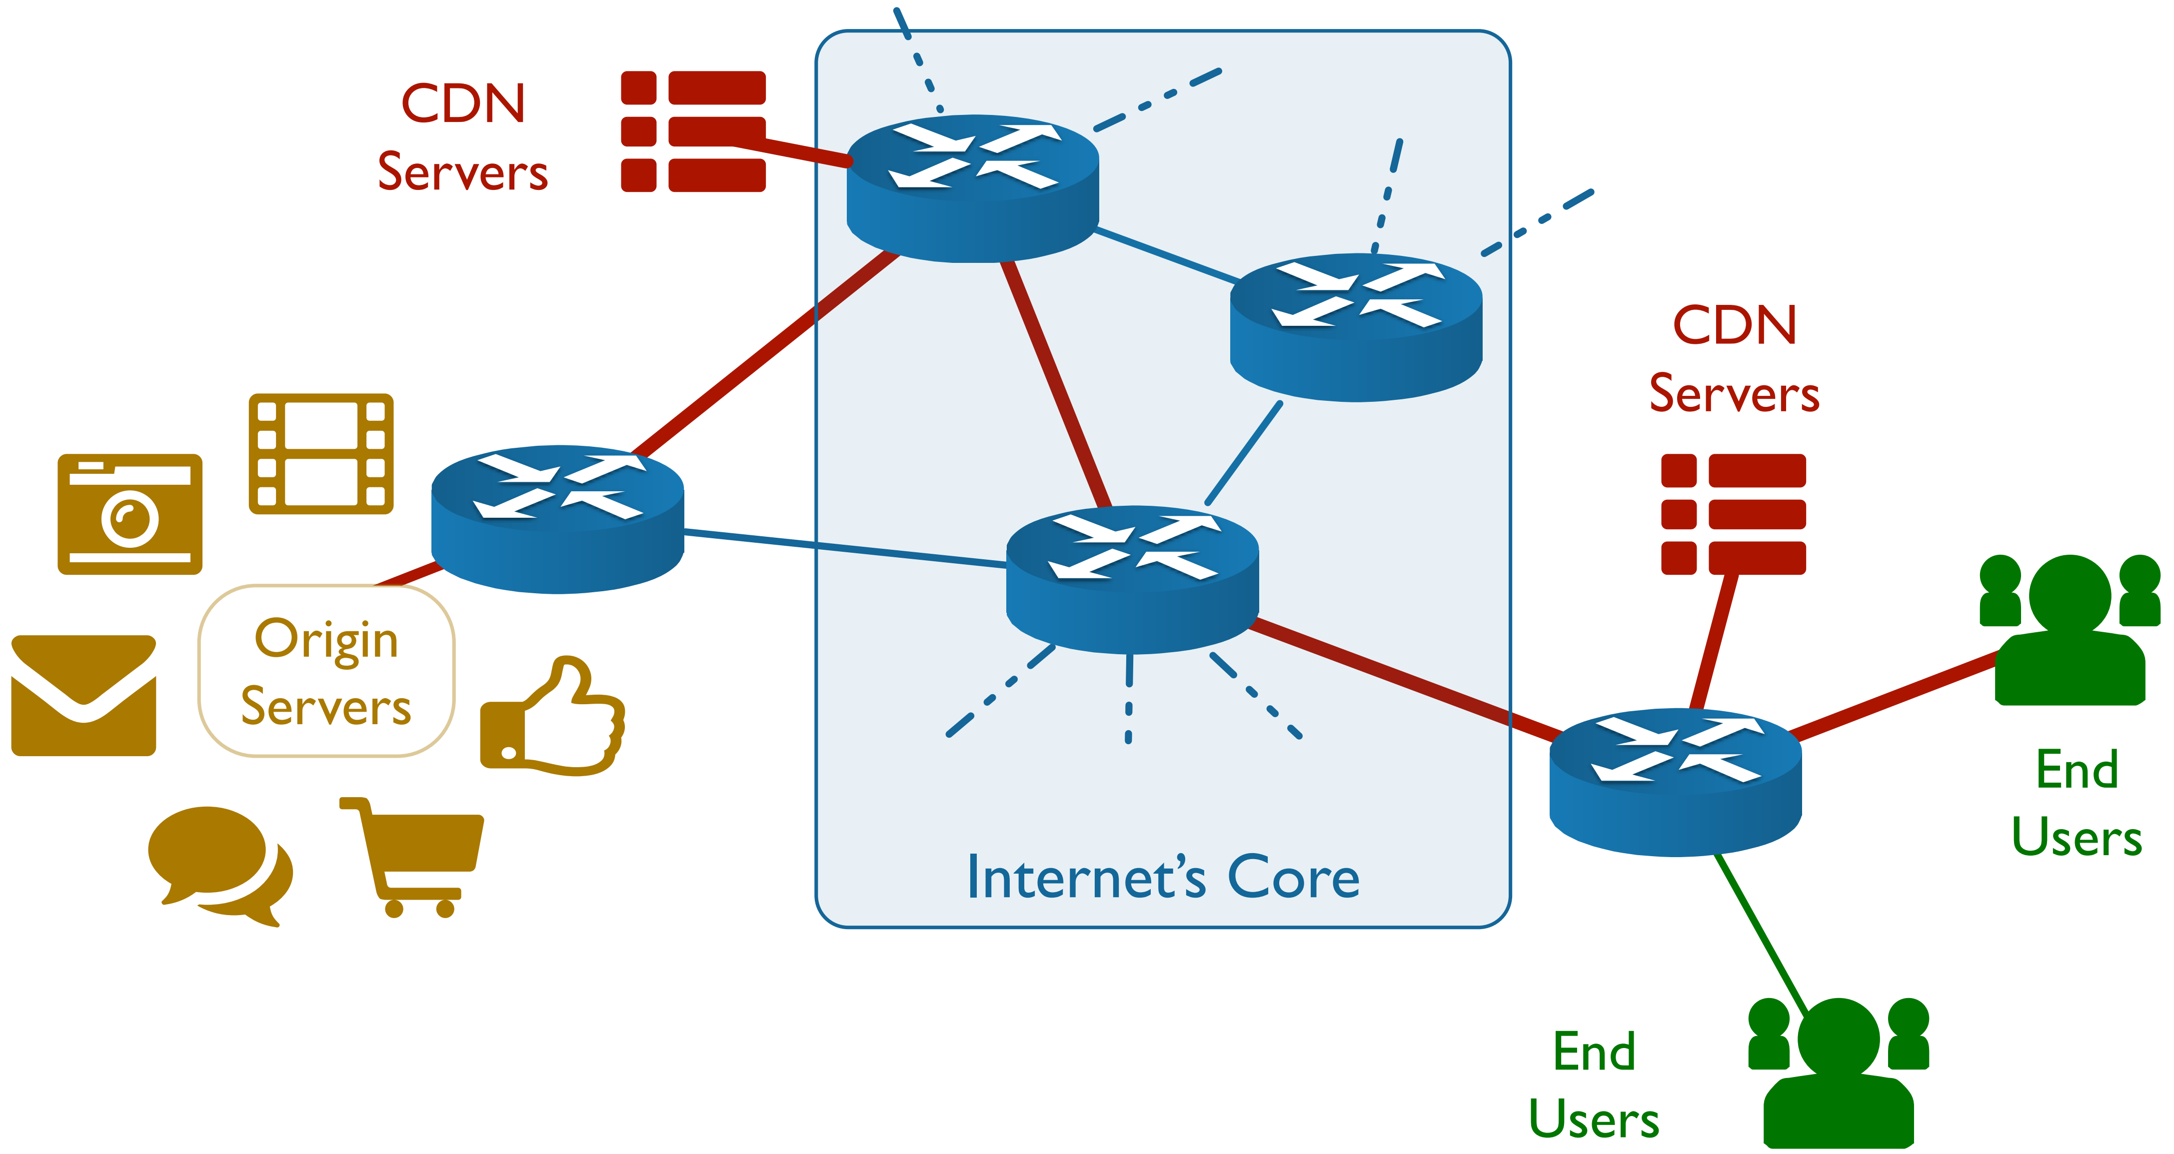

Use measurements over paths between CDN servers to understand the state of the Internet core.

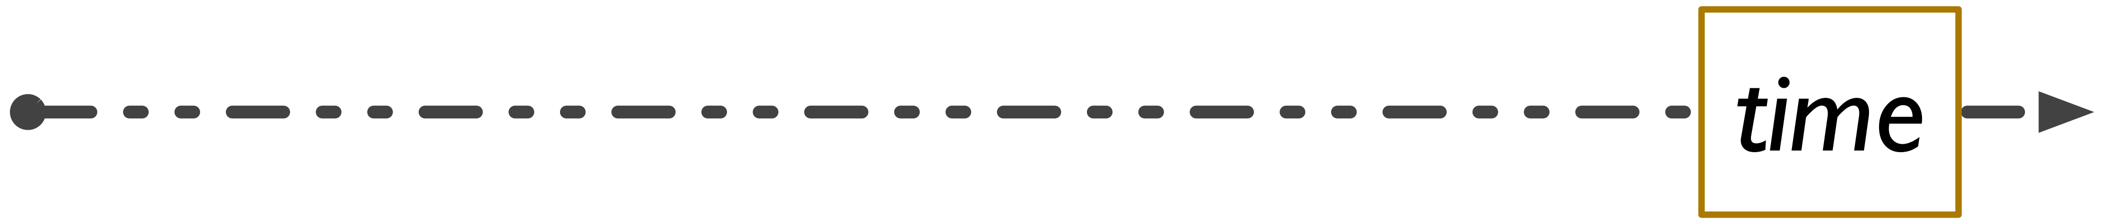

AS_1	AS_1	AS_1	AS_1	AS_1
AS_2	AS_2	AS_5	AS_2	AS_5
AS_3	AS_3	AS_9	AS_3	AS_9
AS_4	AS_4	AS_4	AS_4	AS_4
22.3	**29.7**	**18.2**	**23.1**	**17.9**

Number of unique AS paths observed in the A–B trace timeline

AS_1– AS_2– AS_3– AS_4 and ***AS_1– AS_5– AS_9– AS_4***

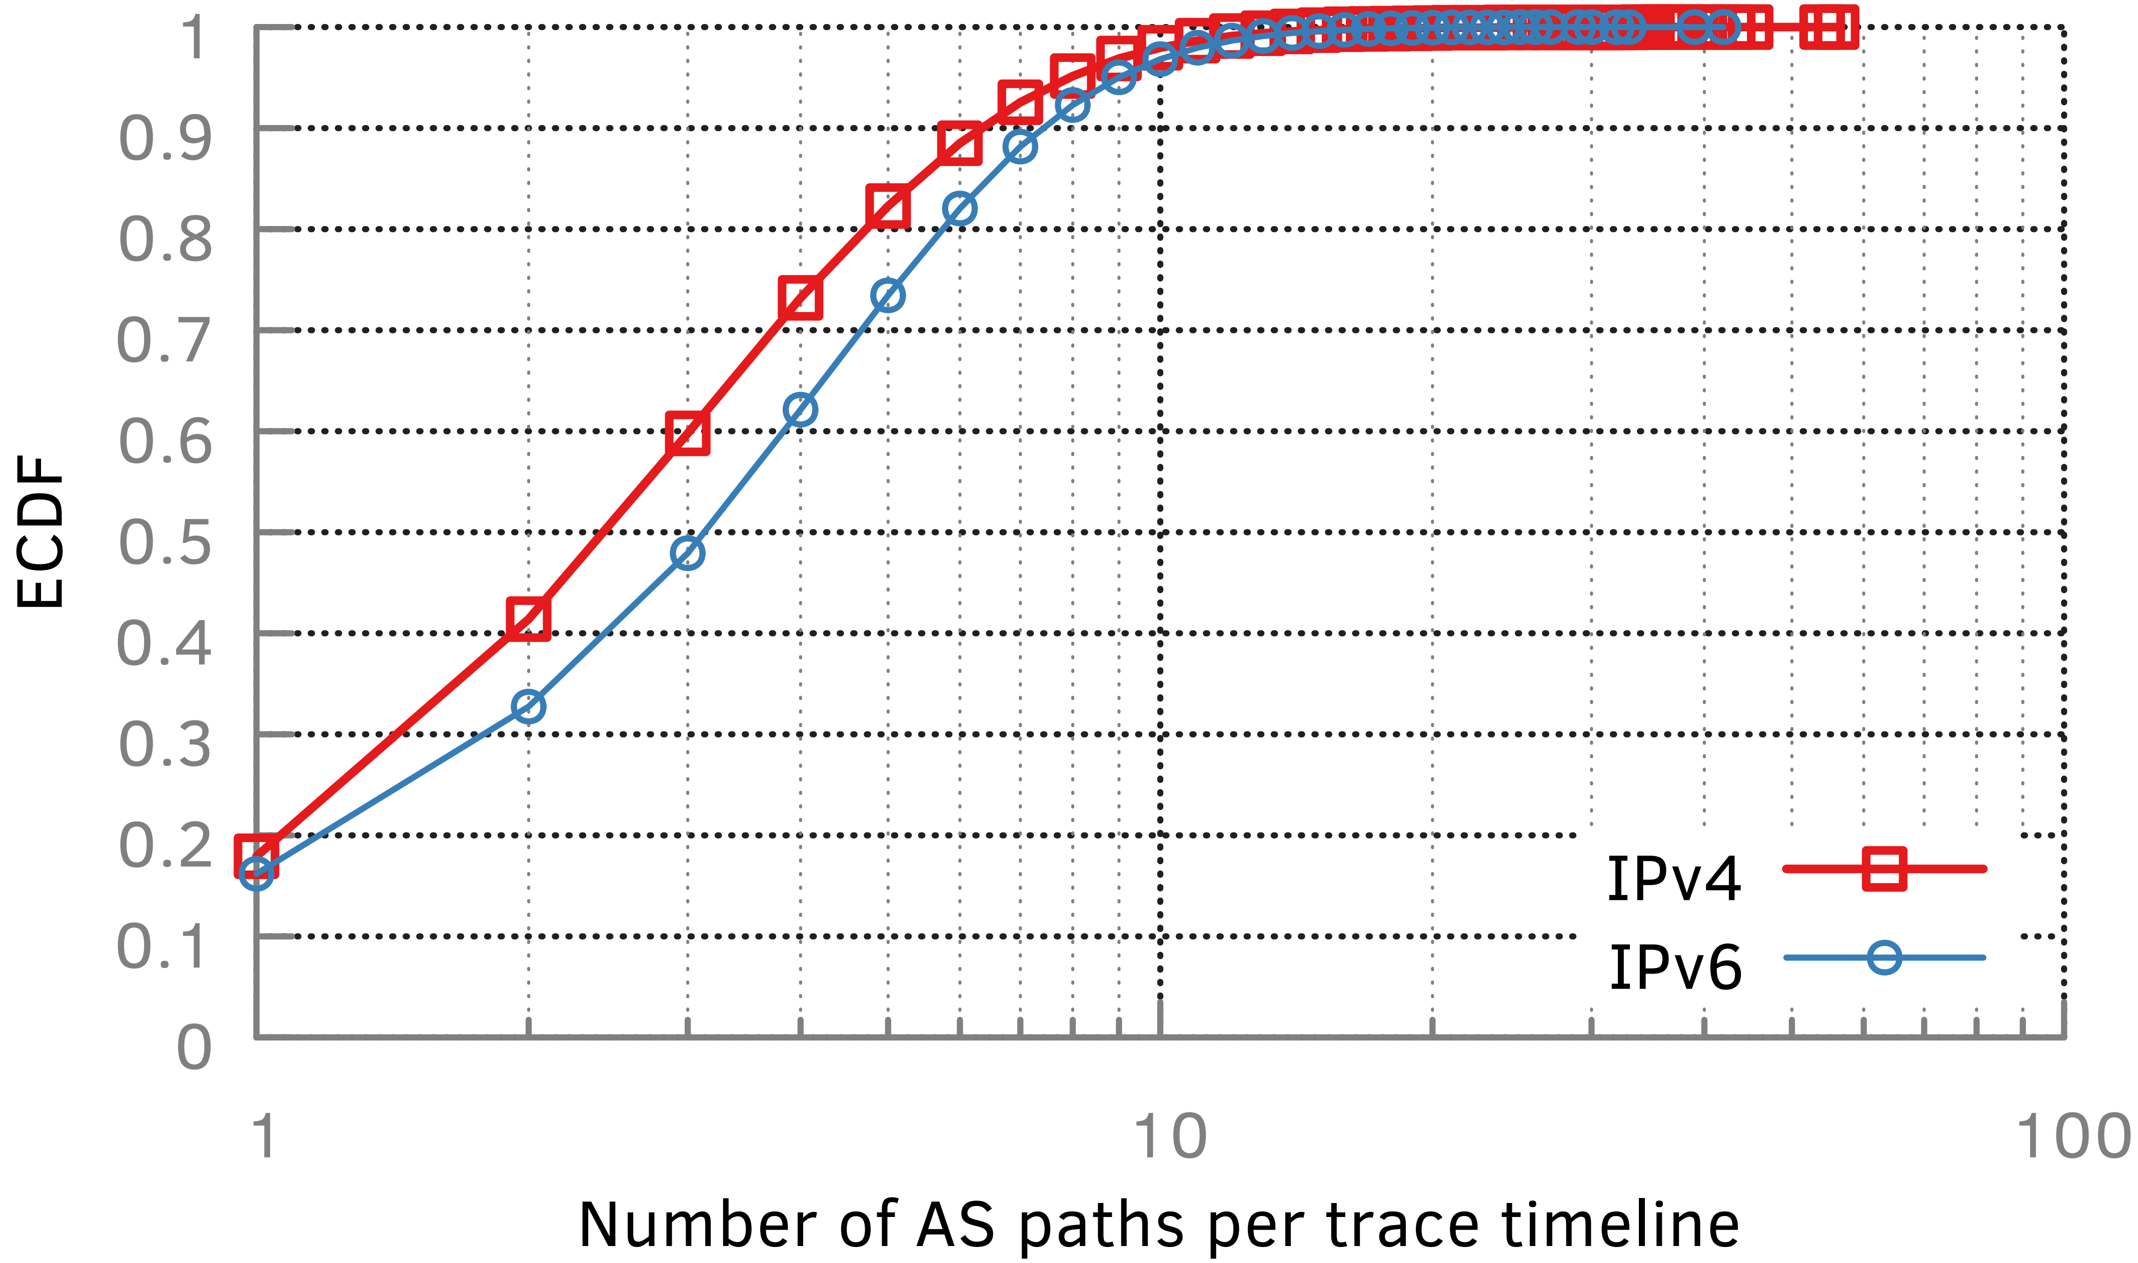

80% of trace timelines have ***5*** or fewer AS paths in IPv4, and ***6*** or fewer in IPv6.

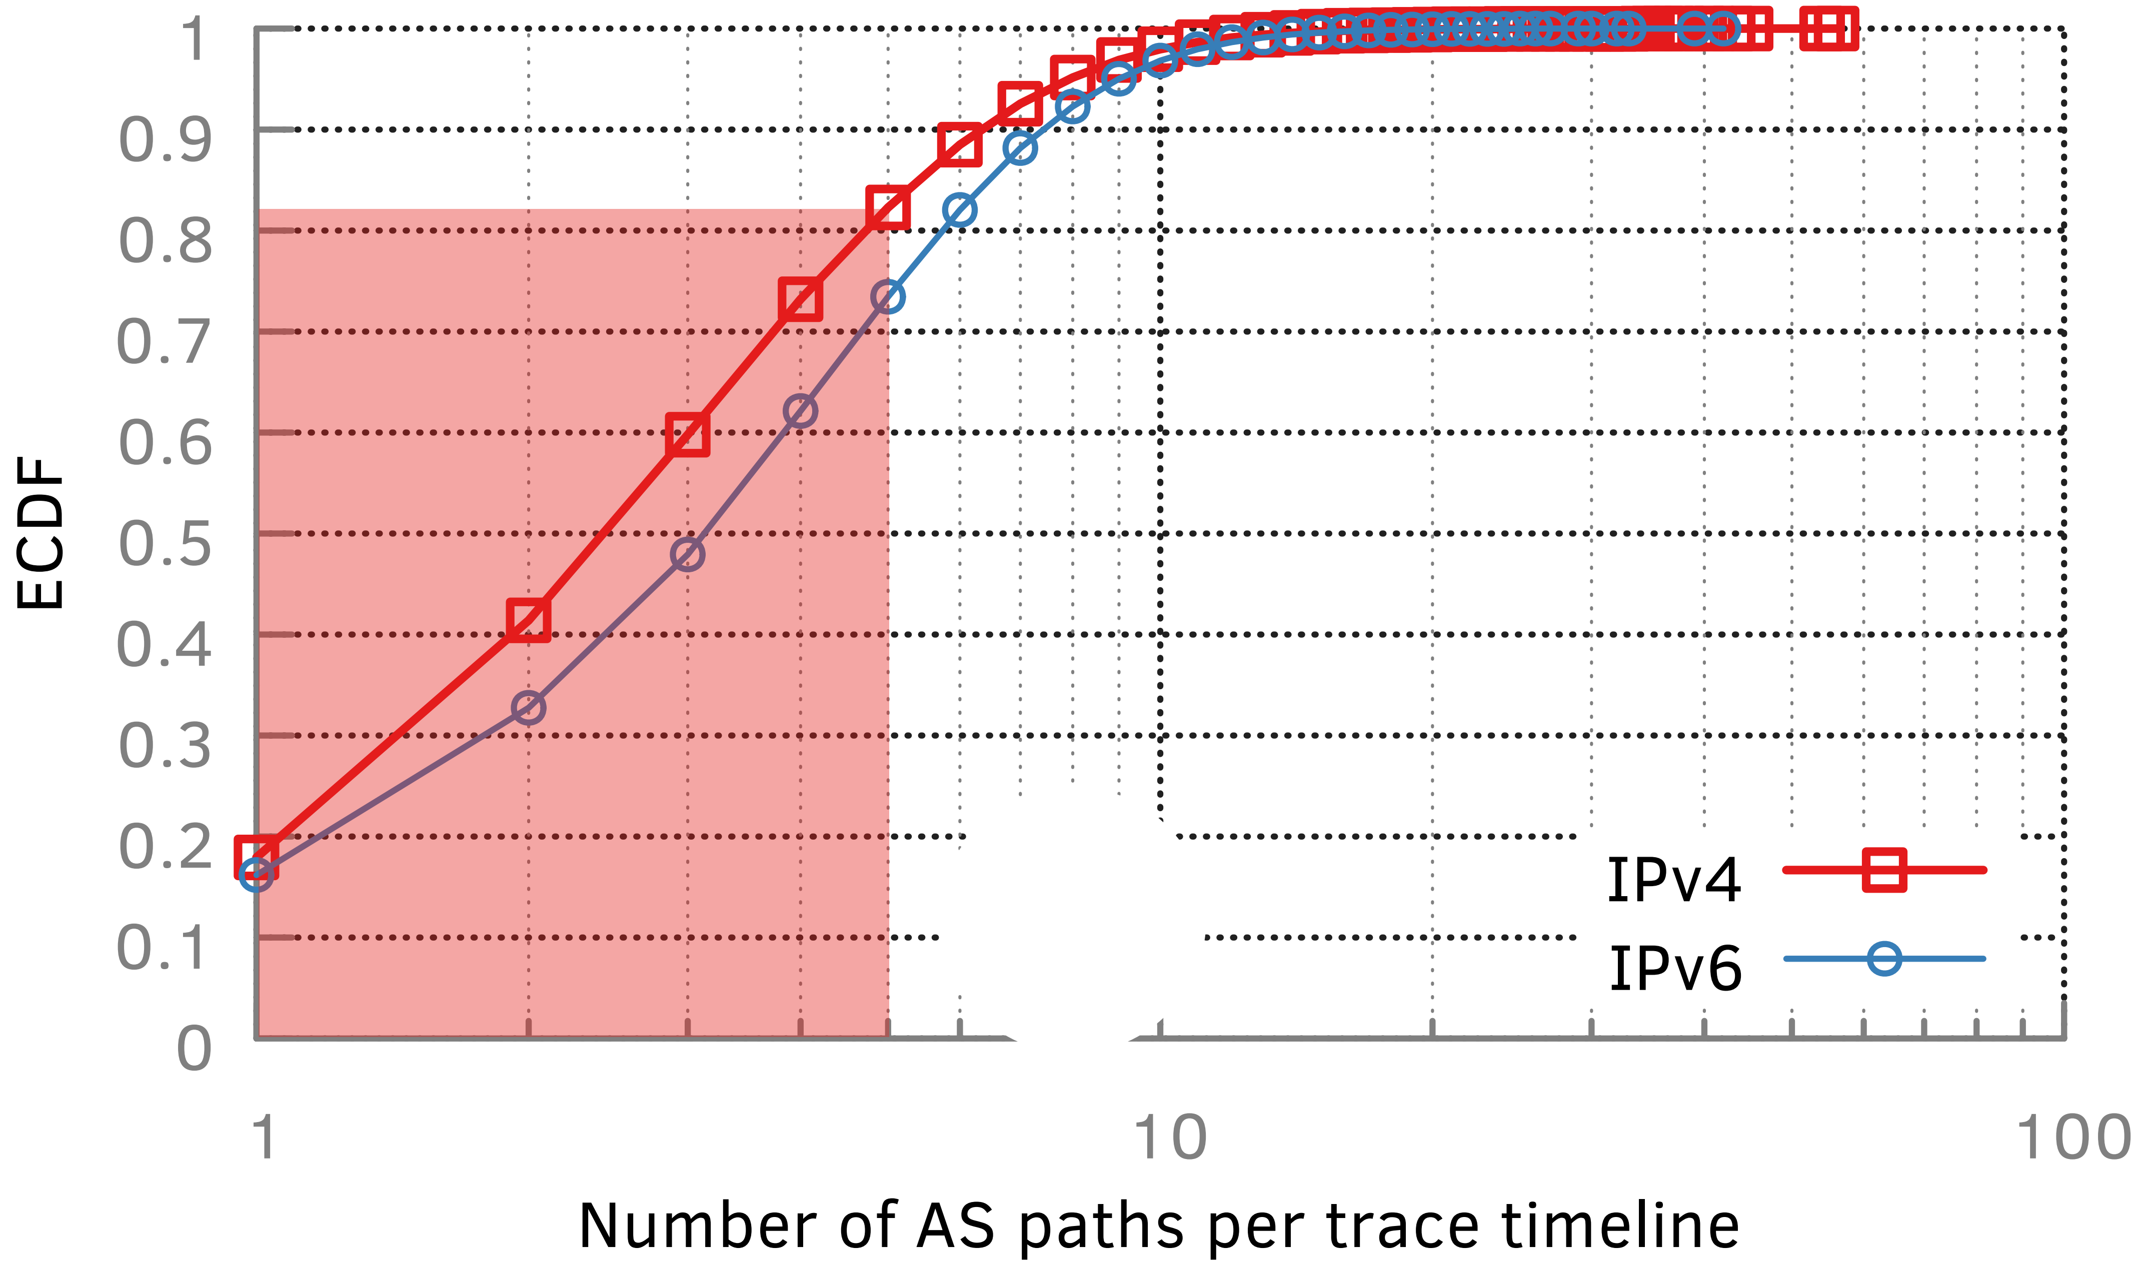

80% of trace timelines have ***5*** or fewer AS paths in IPv4, and ***6*** or fewer in IPv6.

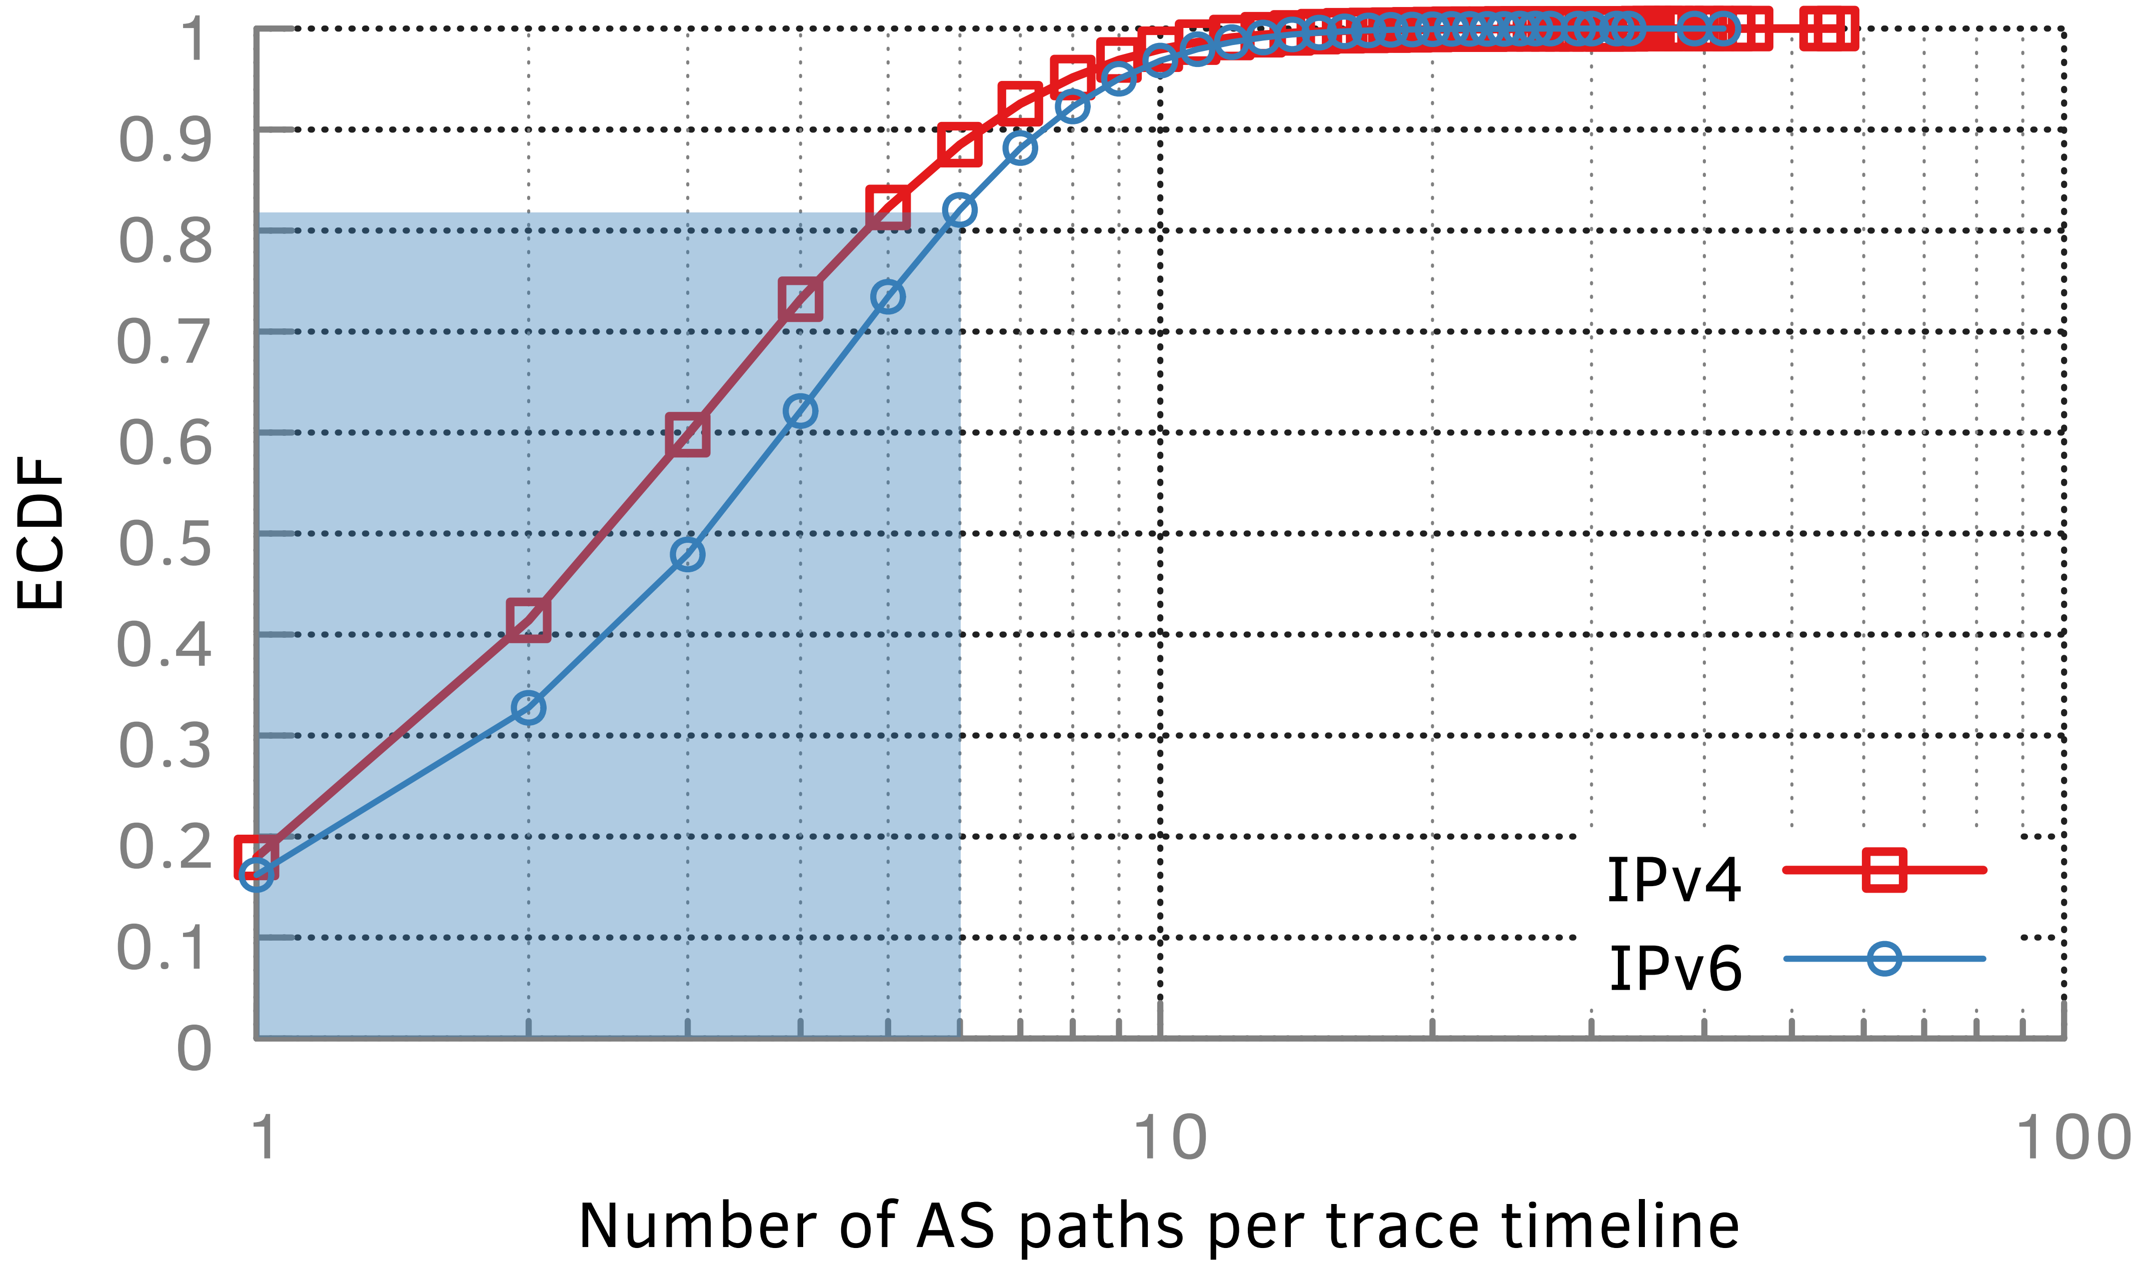

80% of trace timelines have ***5*** or fewer AS paths in IPv4, and ***6*** or fewer in IPv6.

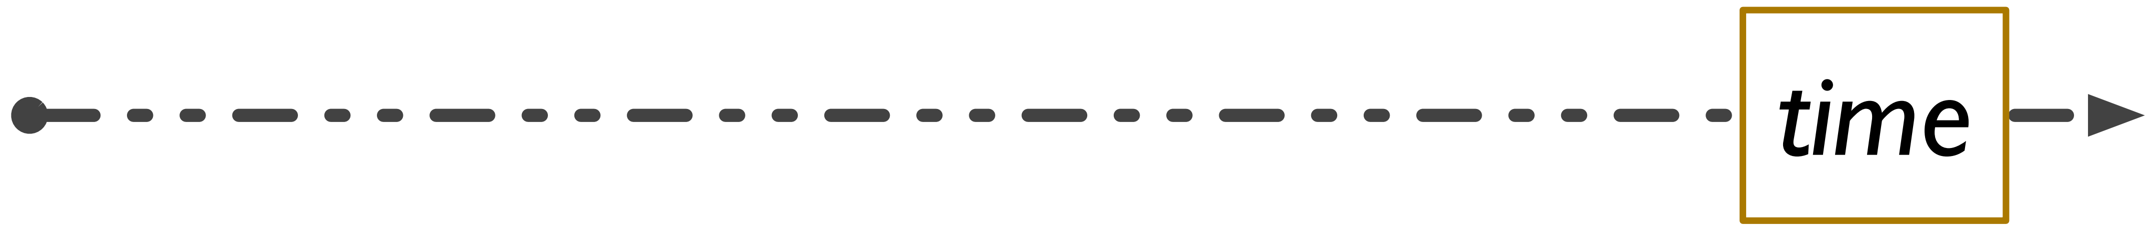

AS_1	AS_1	AS_1
AS_2	AS_2	AS_5
AS_3	AS_3	AS_9
AS_4	AS_4	AS_4
22.3	***29.7***	***18.2***

Combine AS paths observed in the forward direction with

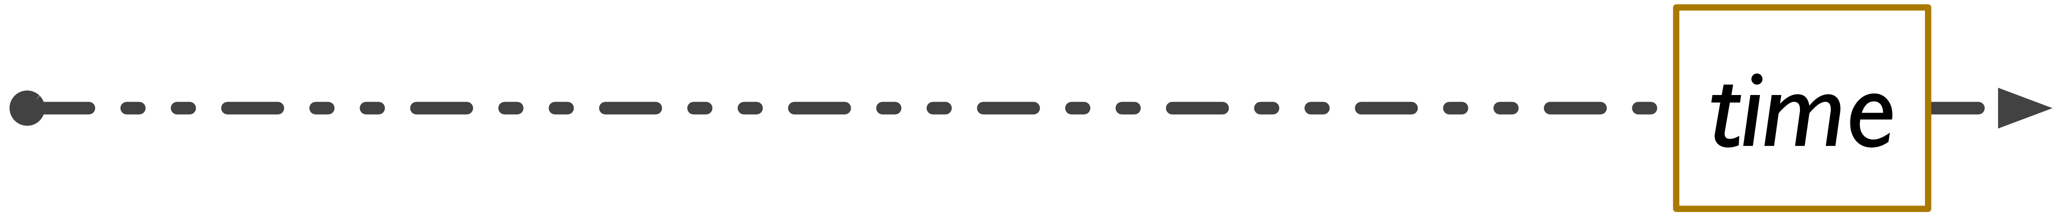

AS_1	**29.1**	AS_1	**24.9**	AS_1	**28.8**
AS_2	AS_1	AS_2	AS_1	AS_5	AS_1
AS_3	AS_8	AS_3	AS_8	AS_9	AS_8
AS_4	AS_3	AS_4	AS_3	AS_4	AS_3
22.3	AS_4	**29.7**	AS_4	**18.2**	AS_4

Combine AS paths observed in the forward direction with those in the reverse direction.

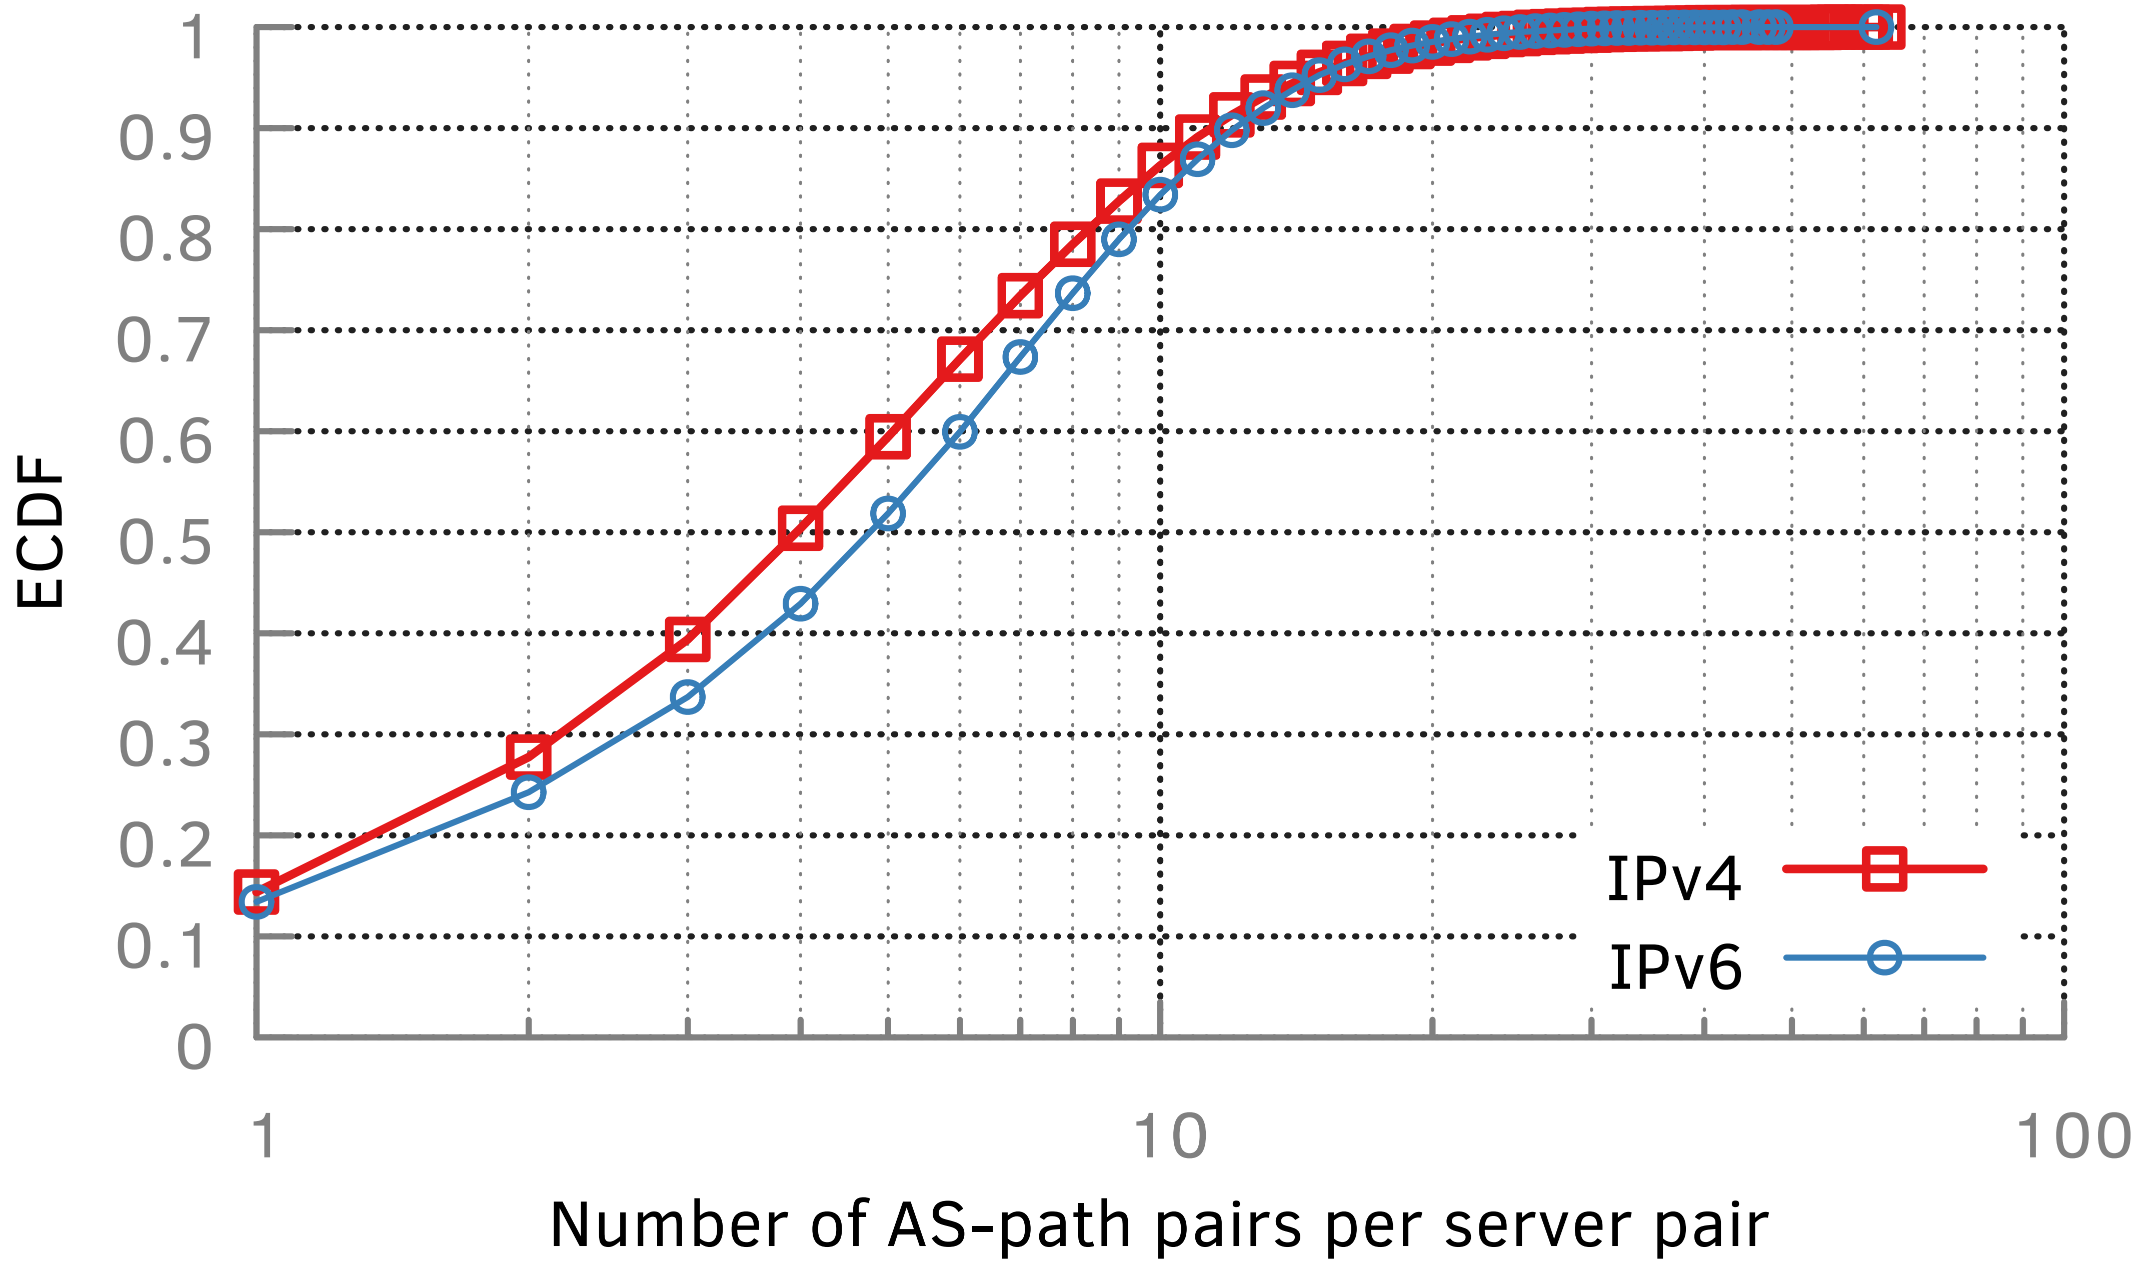

Pairing AS paths in the forward & reverse directions still reveals ***80%*** of server pairs to have ***8*** or fewer path pairs in IPv4, and ***9*** or fewer in IPv6.

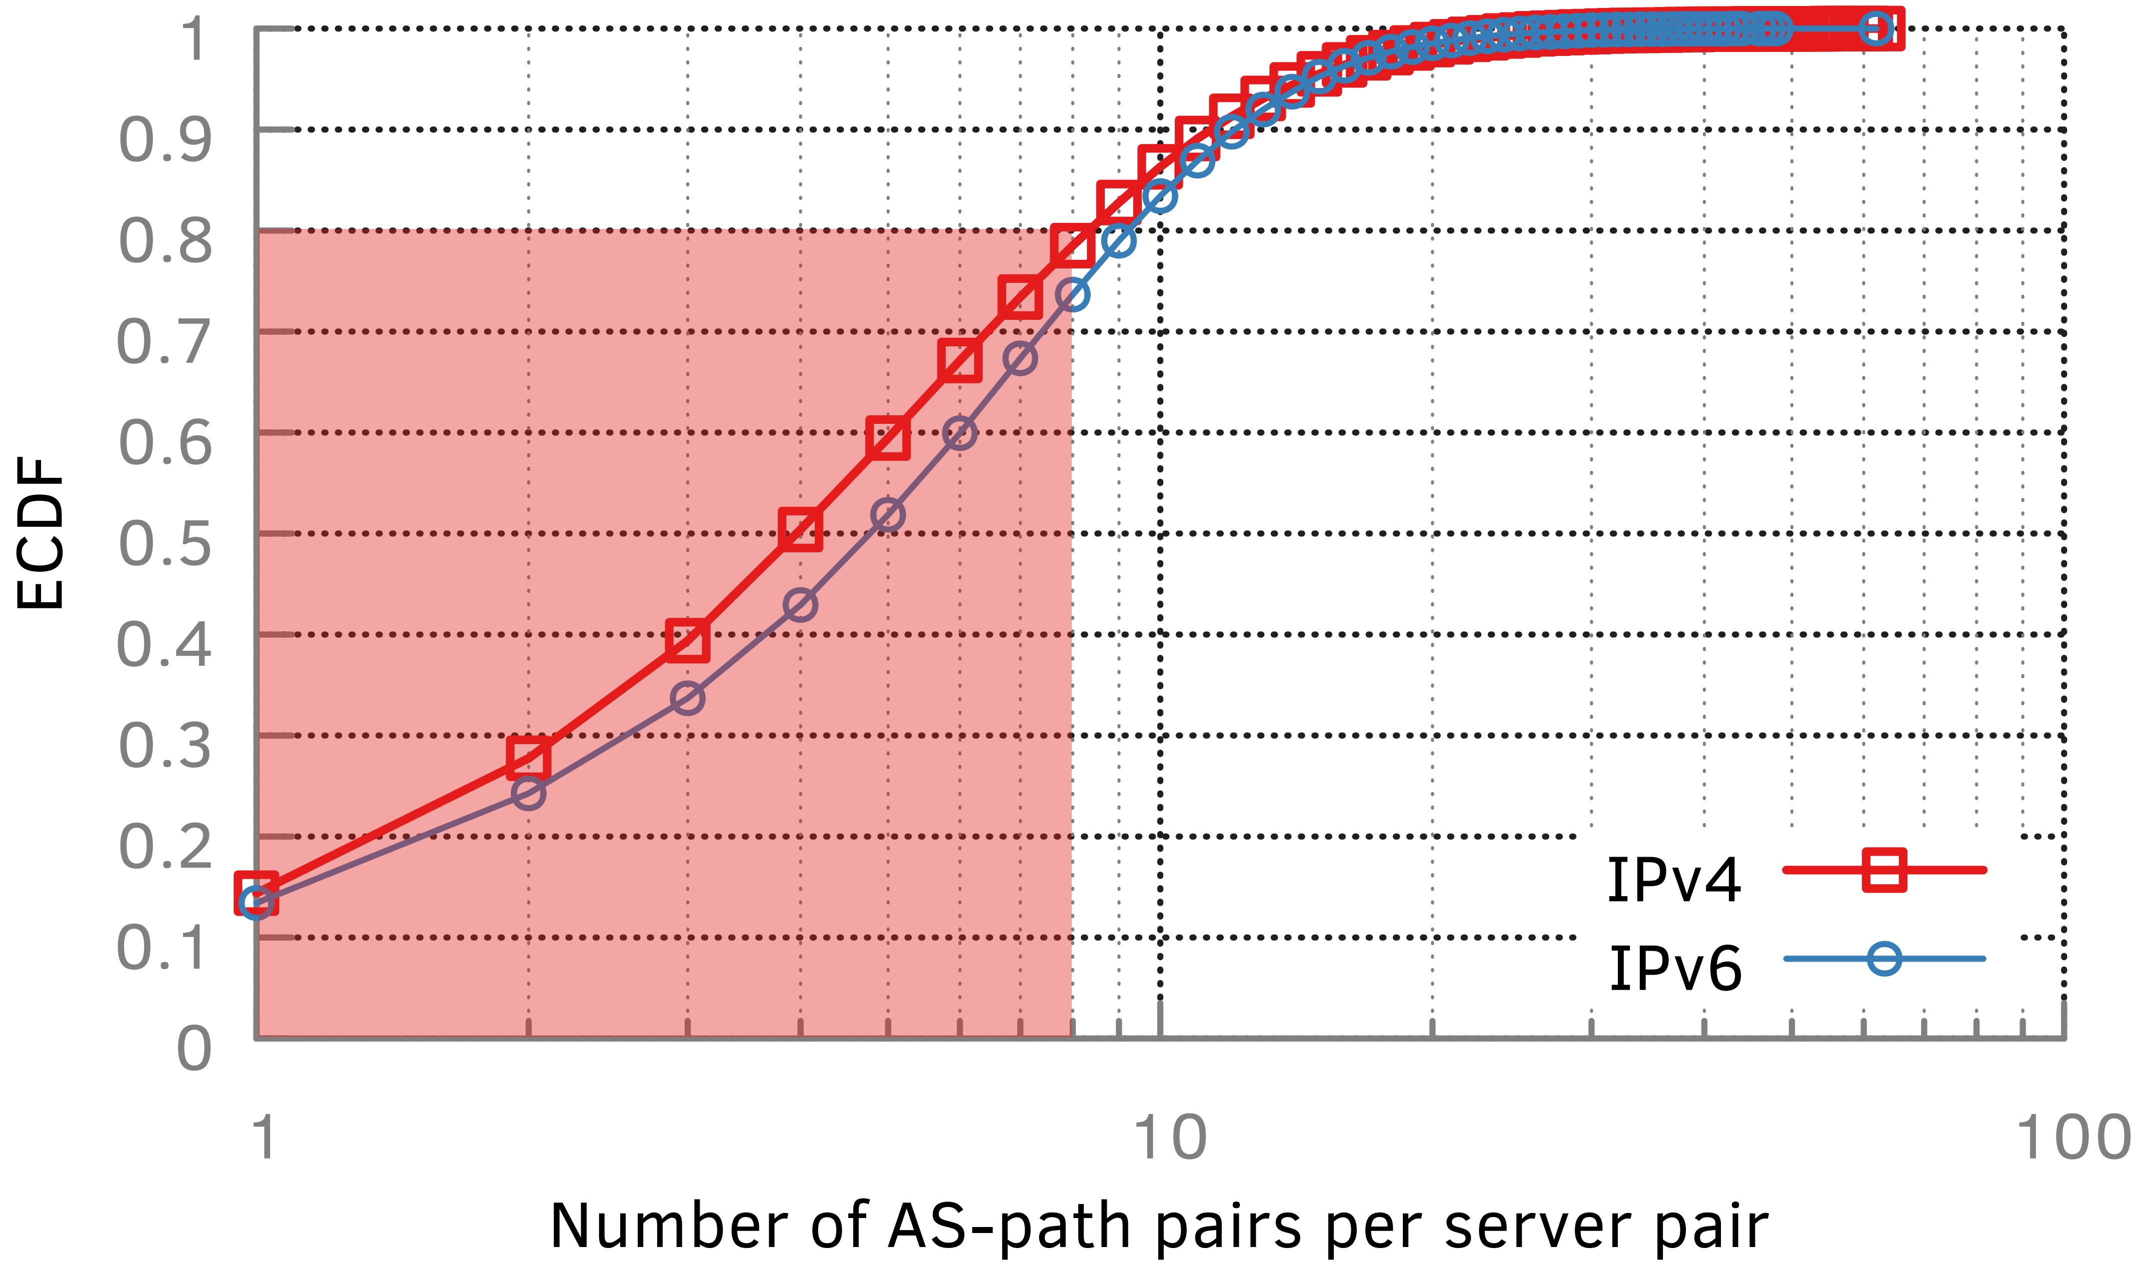

Pairing AS paths in the forward & reverse directions still reveals ***80%*** of server pairs to have ***8*** or fewer path pairs in IPv4, and ***9*** or fewer in IPv6.

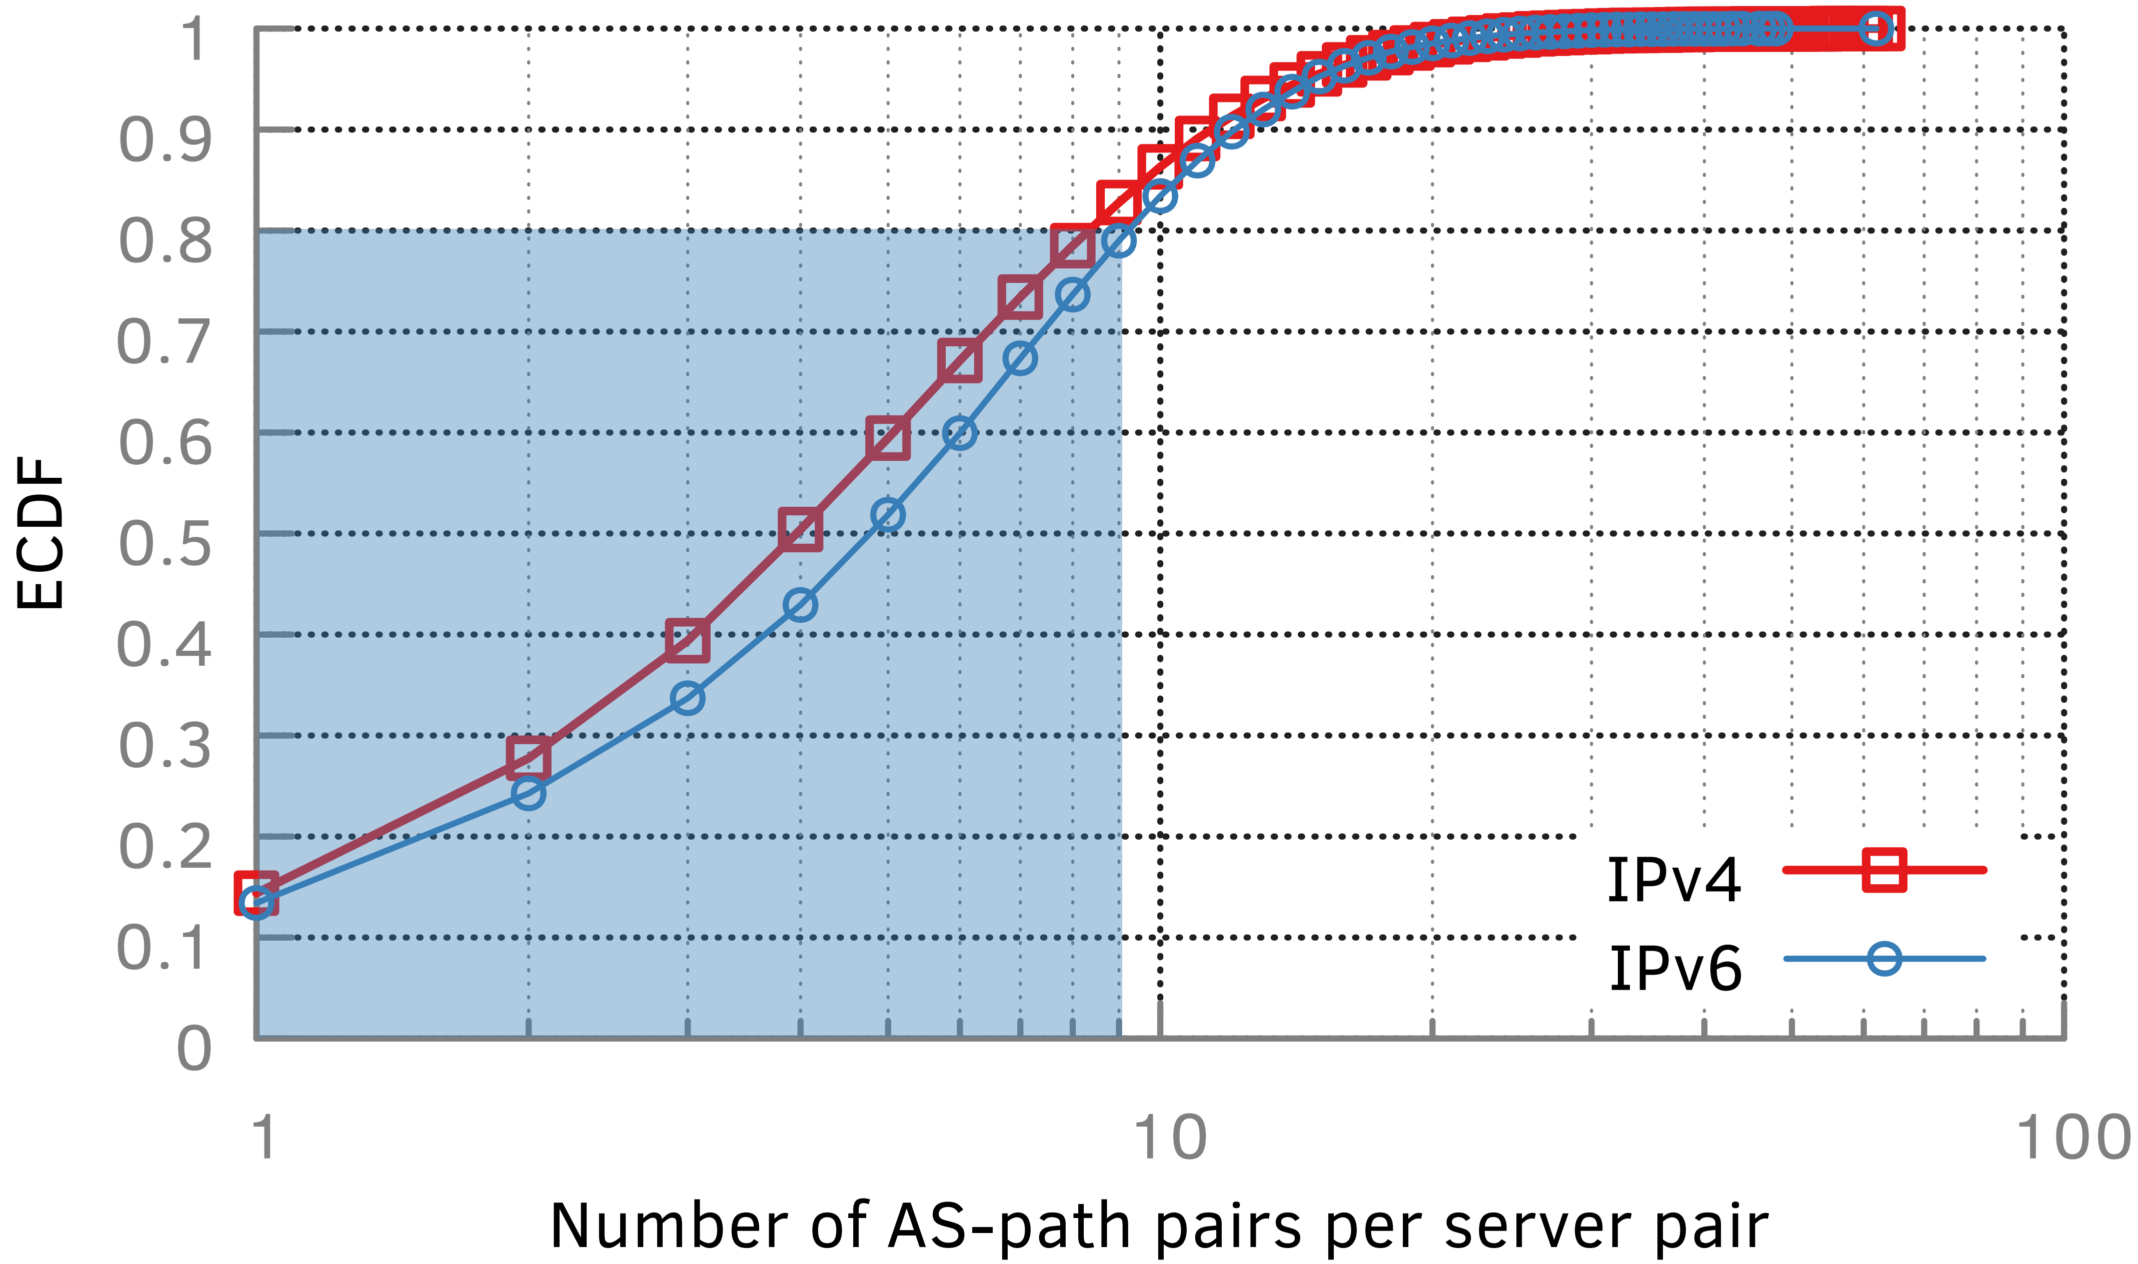

Pairing AS paths in the forward & reverse directions still reveals ***80%*** of server pairs to have ***8*** or fewer path pairs in IPv4, and ***9*** or fewer in IPv6.

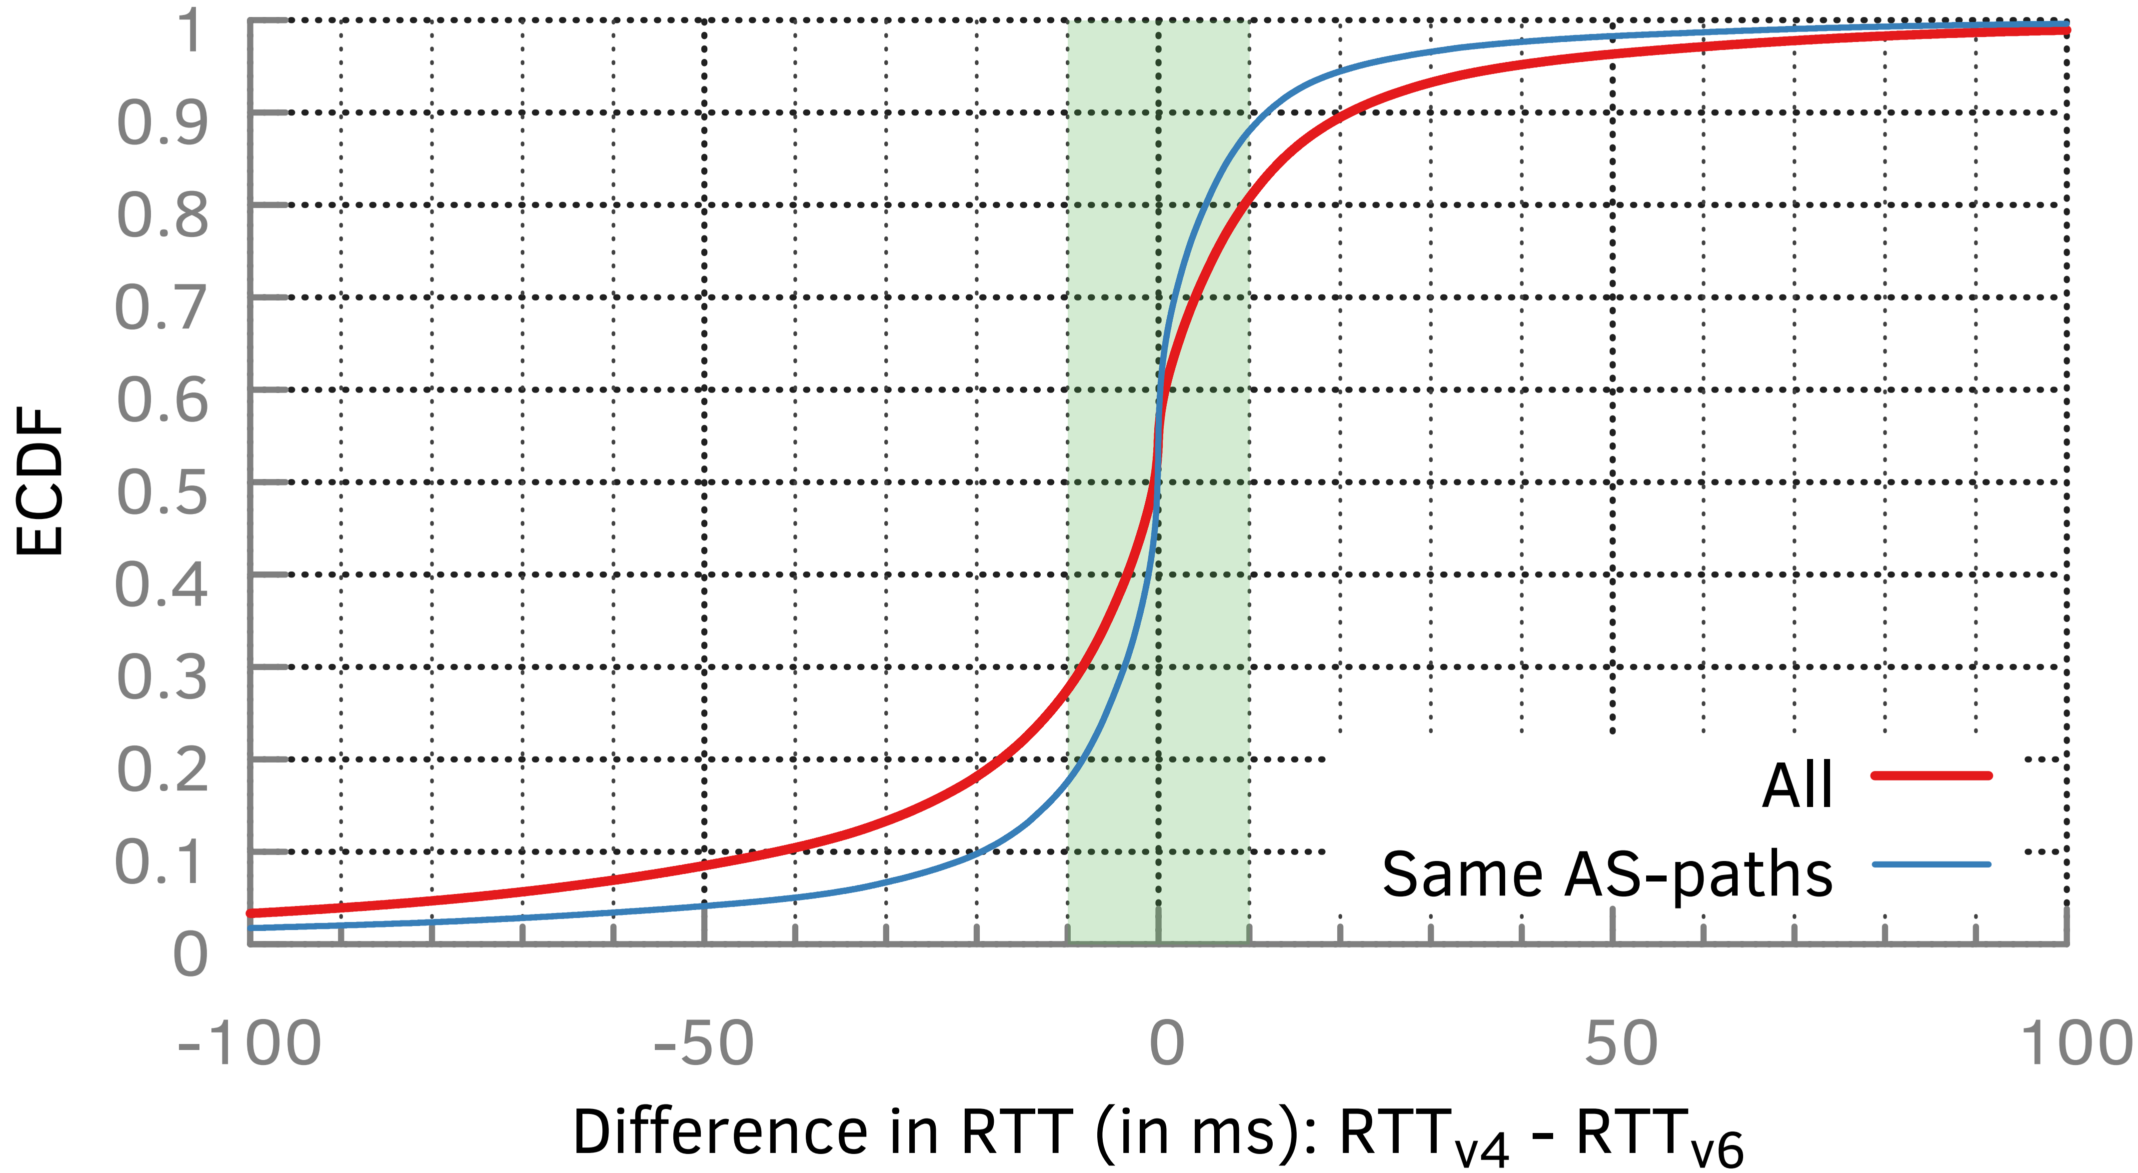

ECDF
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
-100
-50
0
50
100
All
Same AS-paths
Difference in RTT (in ms): RTTv4 - RTTv6

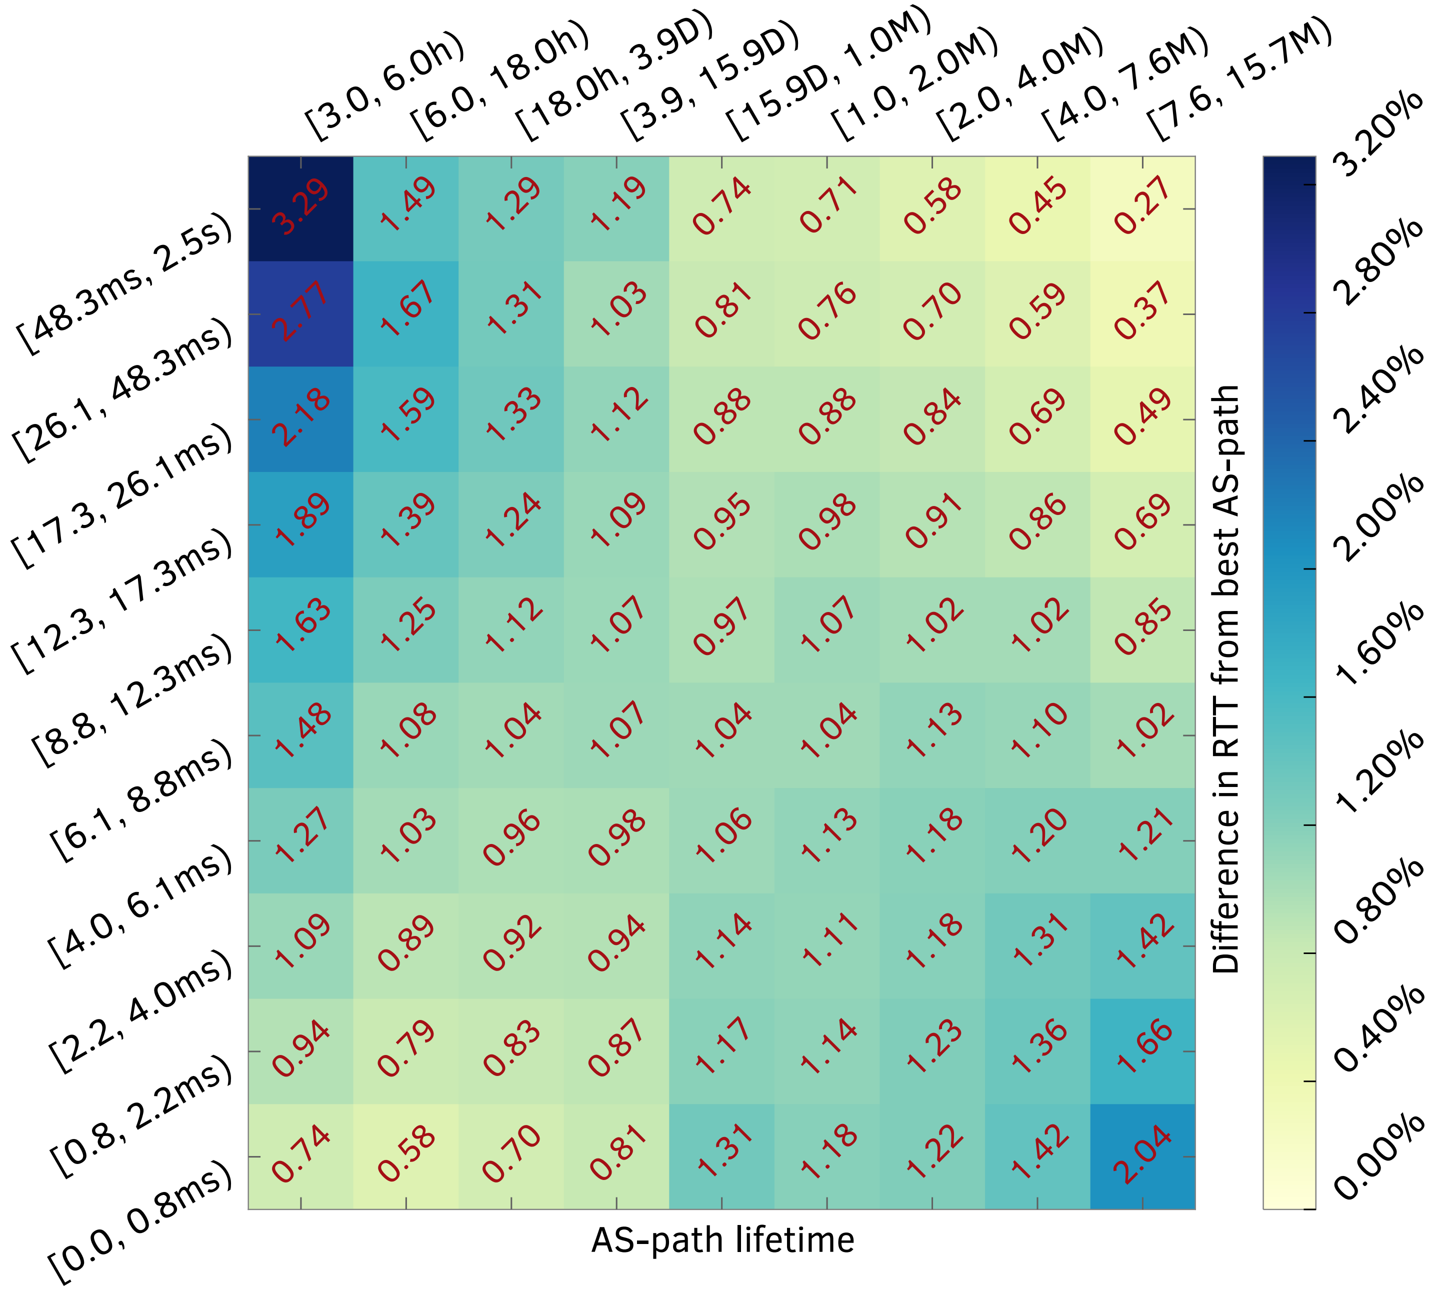

Comparing magnitudes of increase in (baseline) *10th* percentile of RTTs of AS paths (each relative to the best AS path of the corresponding trace timeline) with the lifetime of AS paths …

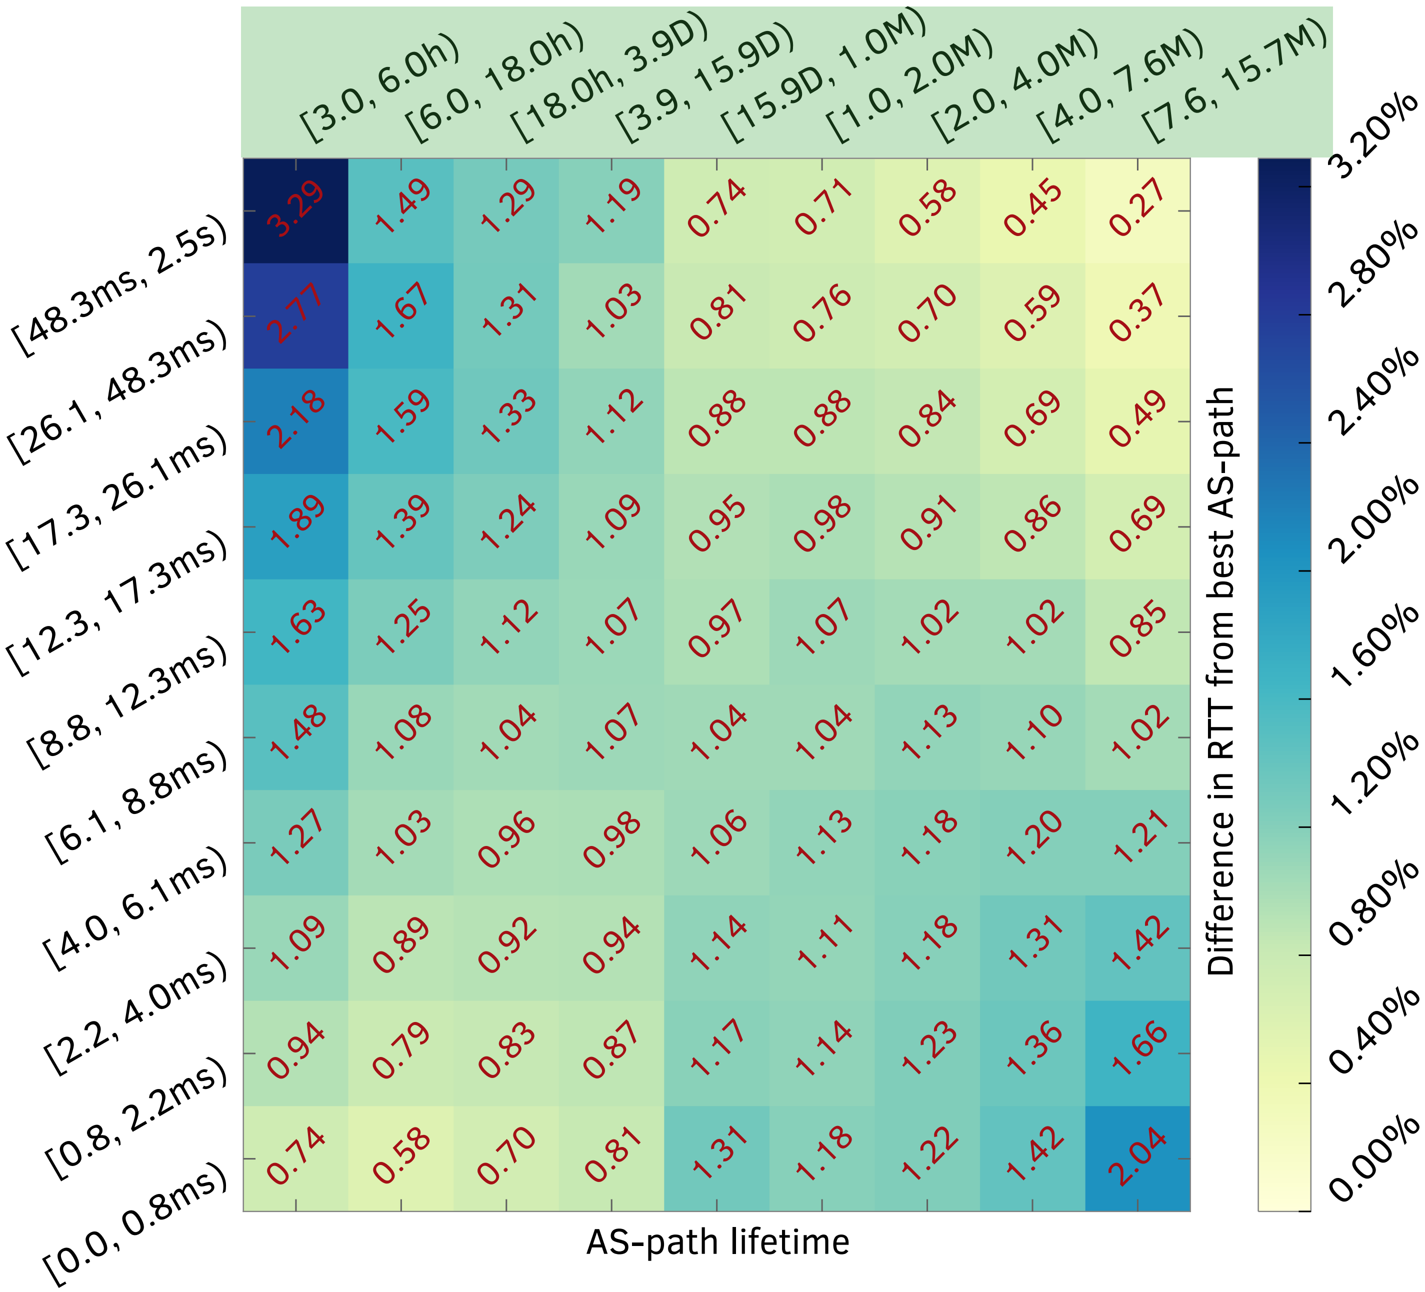

X-axis: deciles of the distribution of AS-path lifetimes.

half-open intervals

[0.0, 3.0h) has no data points

Same value for *0th%* and *10th%* of the AS-path lifetime distribution

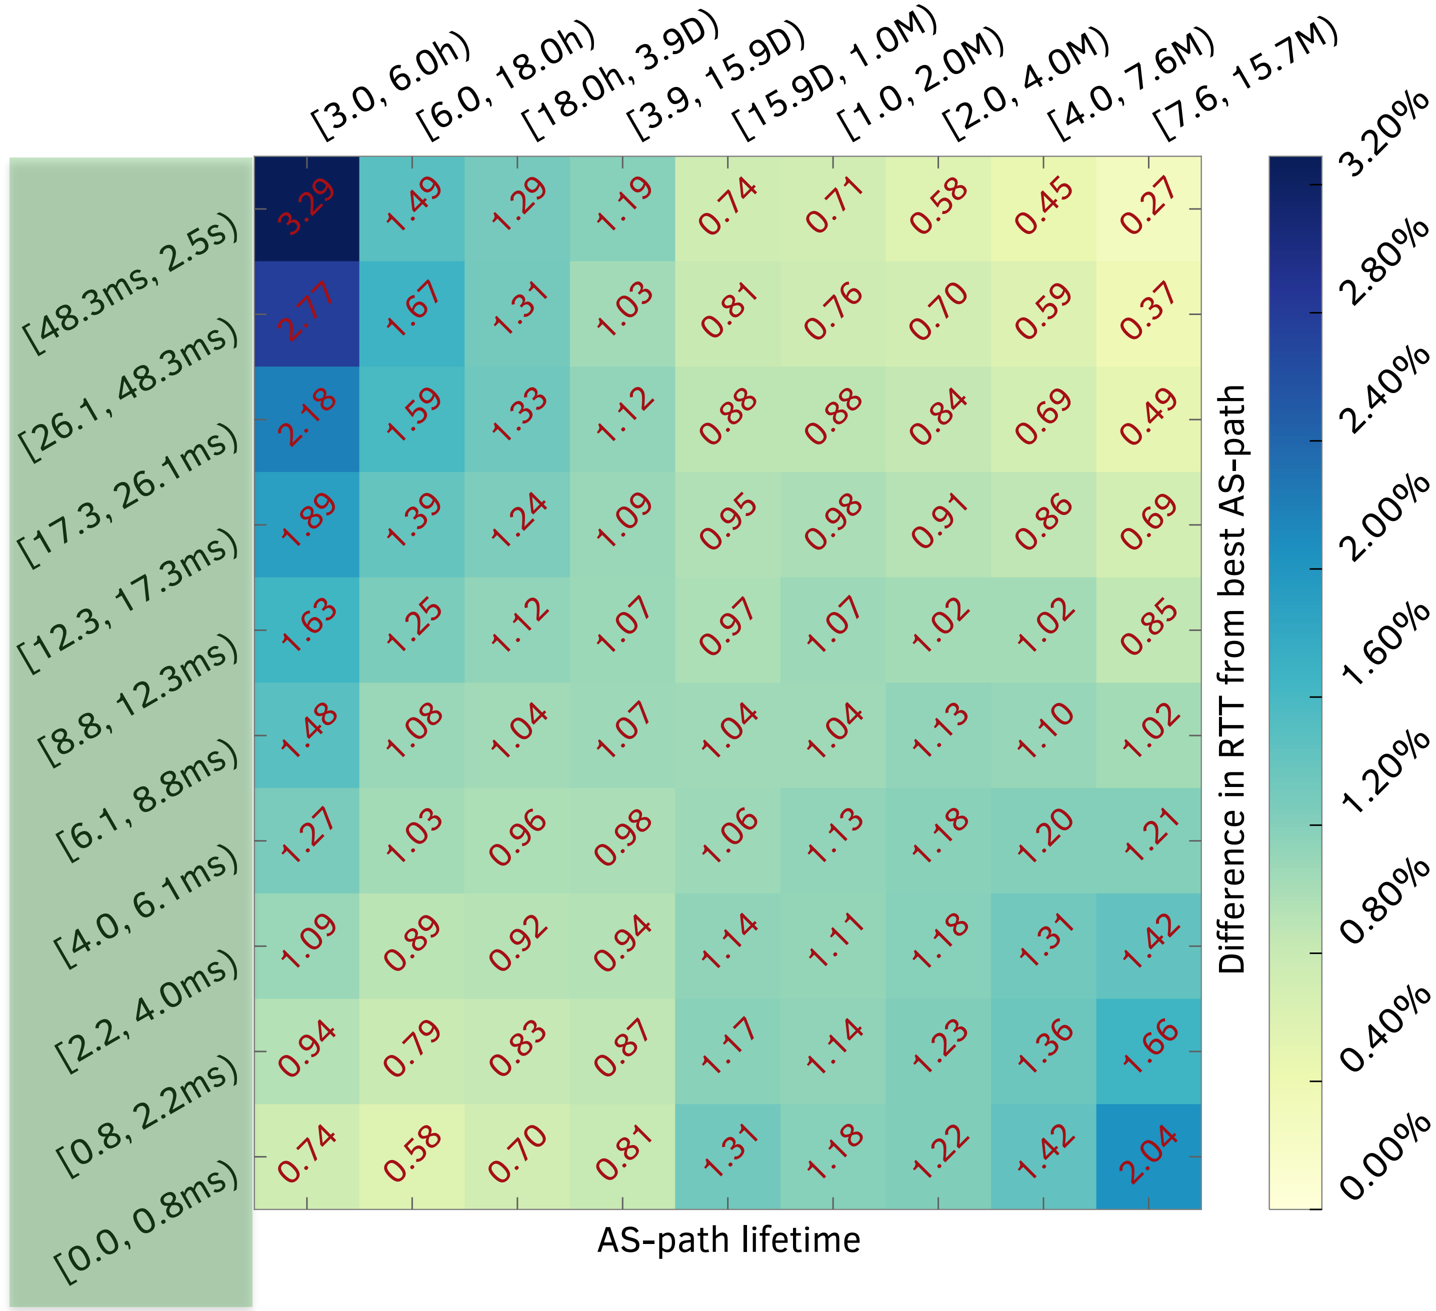

Y-axis: deciles of the distribution of magnitudes of increase in 10th percentile of RTTs of AS paths (each relative to the best AS path of the corresponding trace timeline).

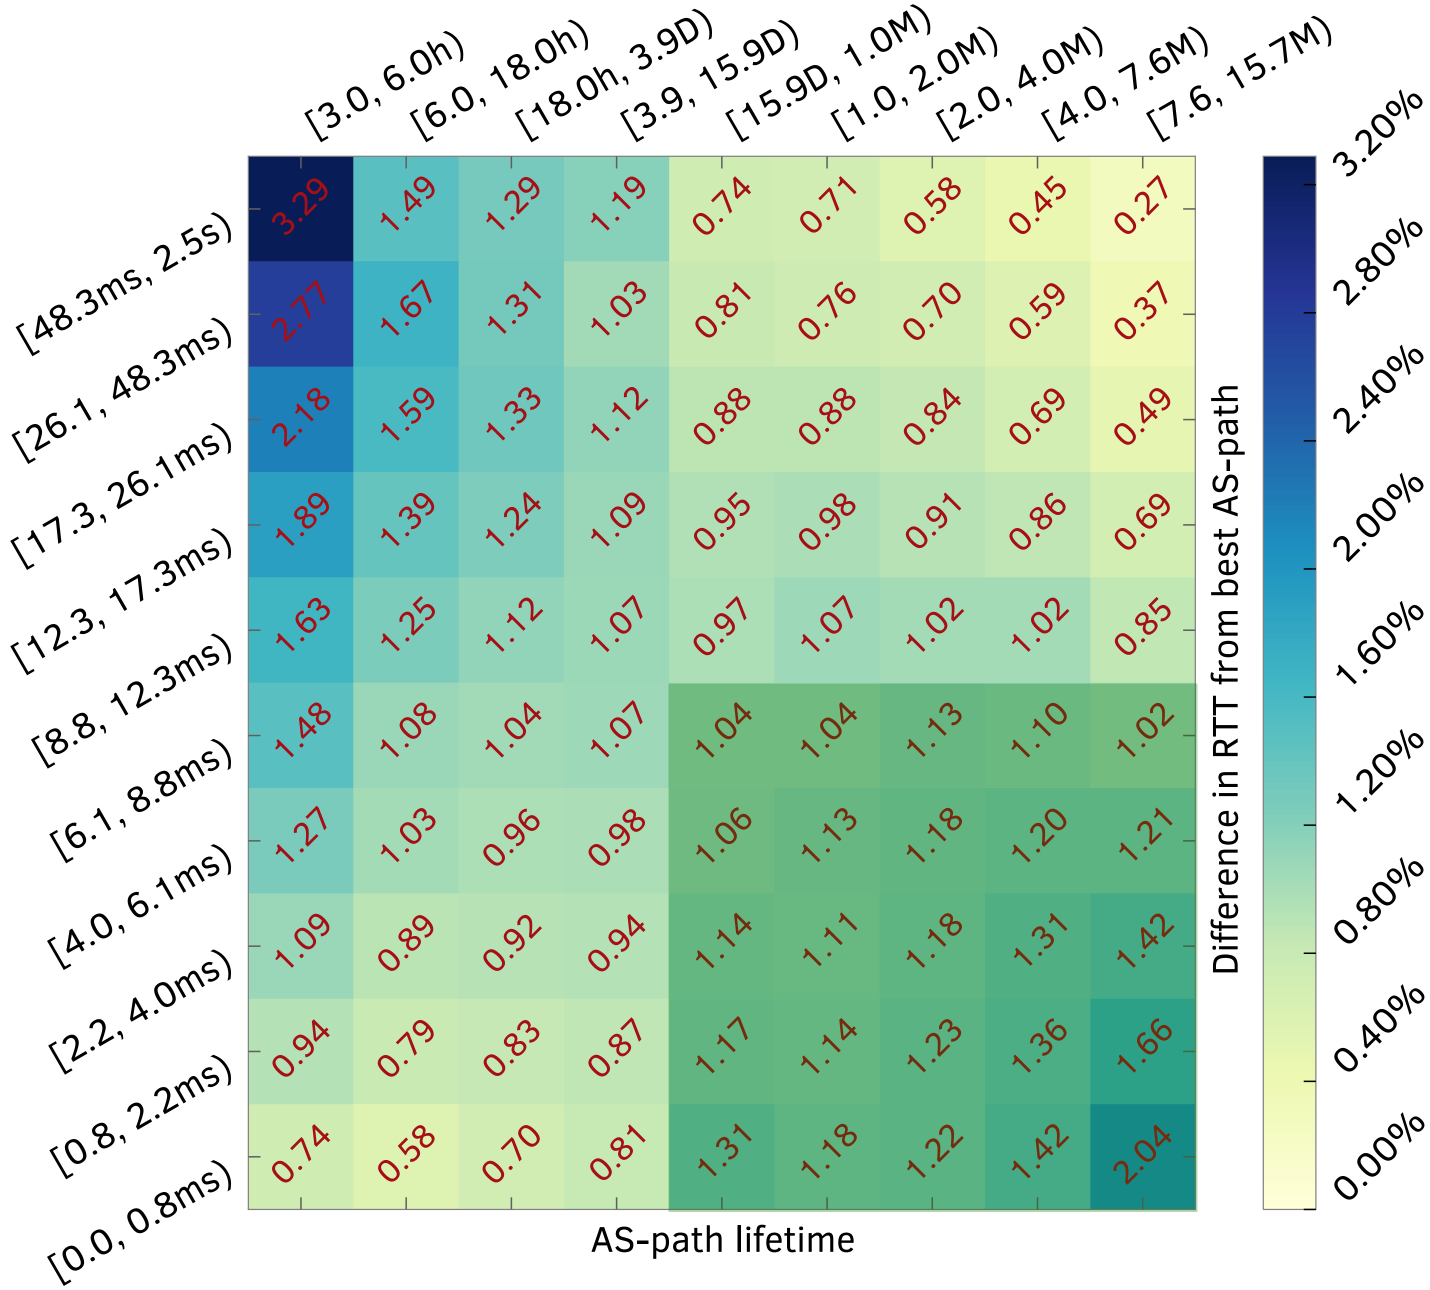

Baseline RTTs of AS paths with longer lifetimes are close in value to that of the best AS path of corresponding trace timelines.

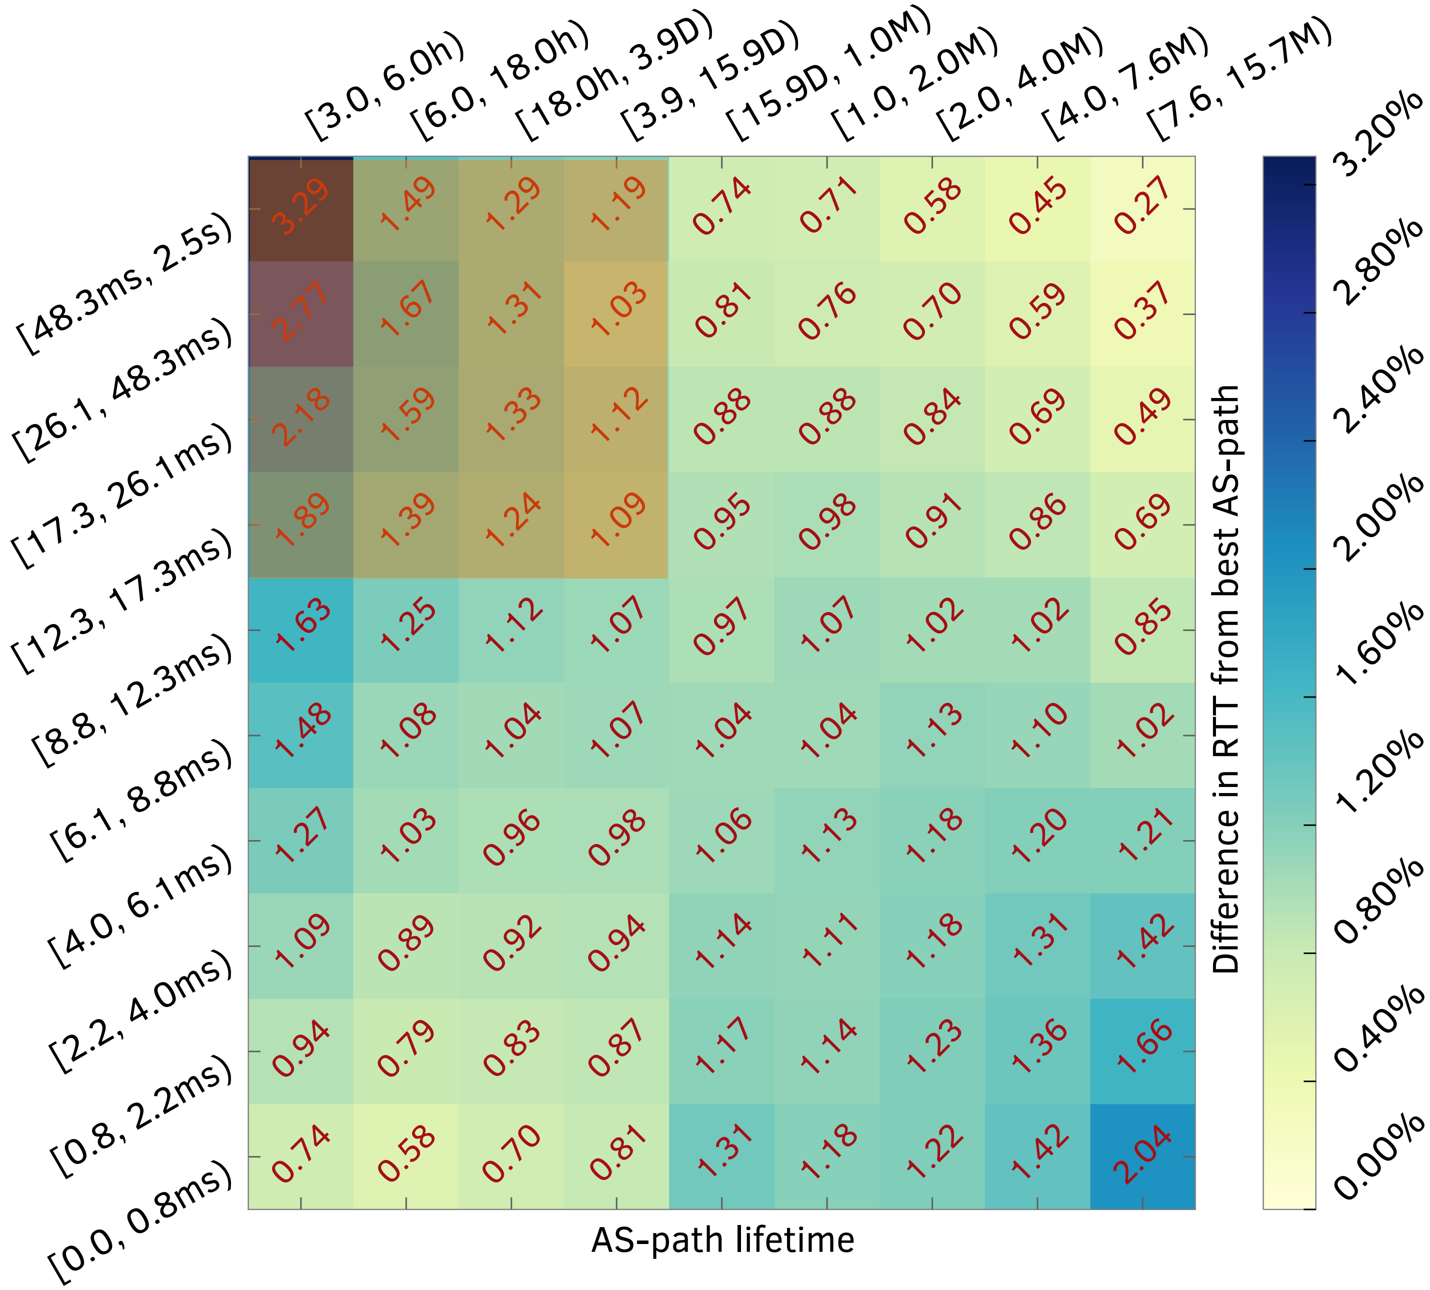

Paths with poor-performance are often those with relatively short lifetimes.

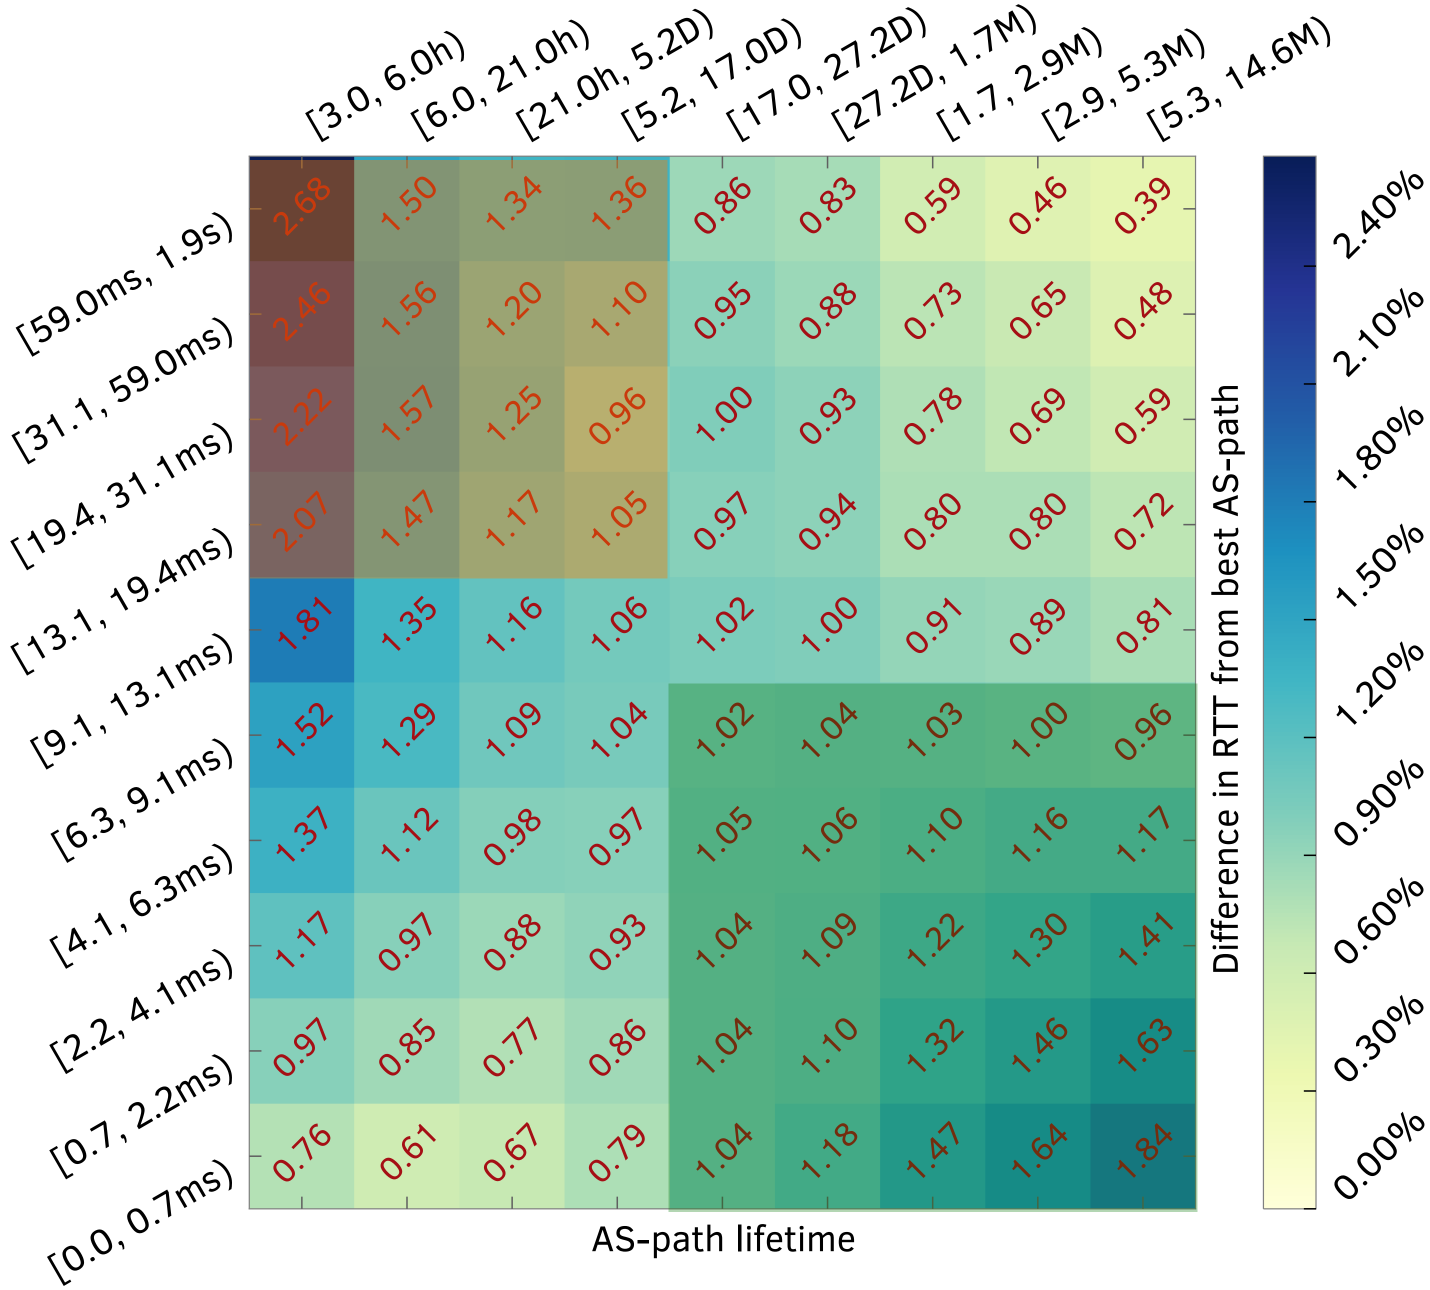

Similar observations from IPv6 traceroutes.