Characterizing IPv6 control and data plane stability

Ioana Livadariu, Ahmed Elmokashfi
(Simula Research Laboratory)

Amogh Dhamdhere
(CAIDA/UCSD)

[simula . research laboratory]

caida

IPv4 addresses are running out

- IANA - February 2011
- Last /8 - APNIC (Asia Pacific), RIPE (Europe), LACNIC(Latin America), ARIN (North America)

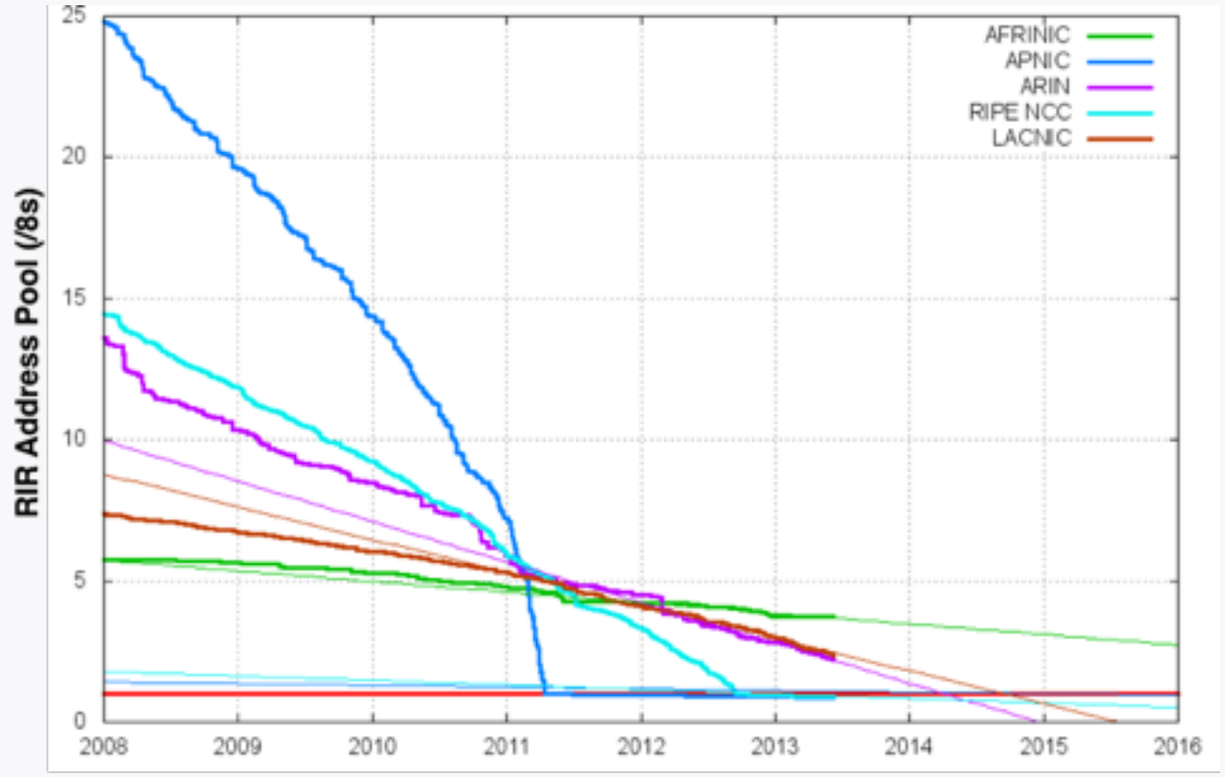

RIR IPv4 Address Run-Down Model [Source: http://www.potaroo.net/tools/ipv4/]

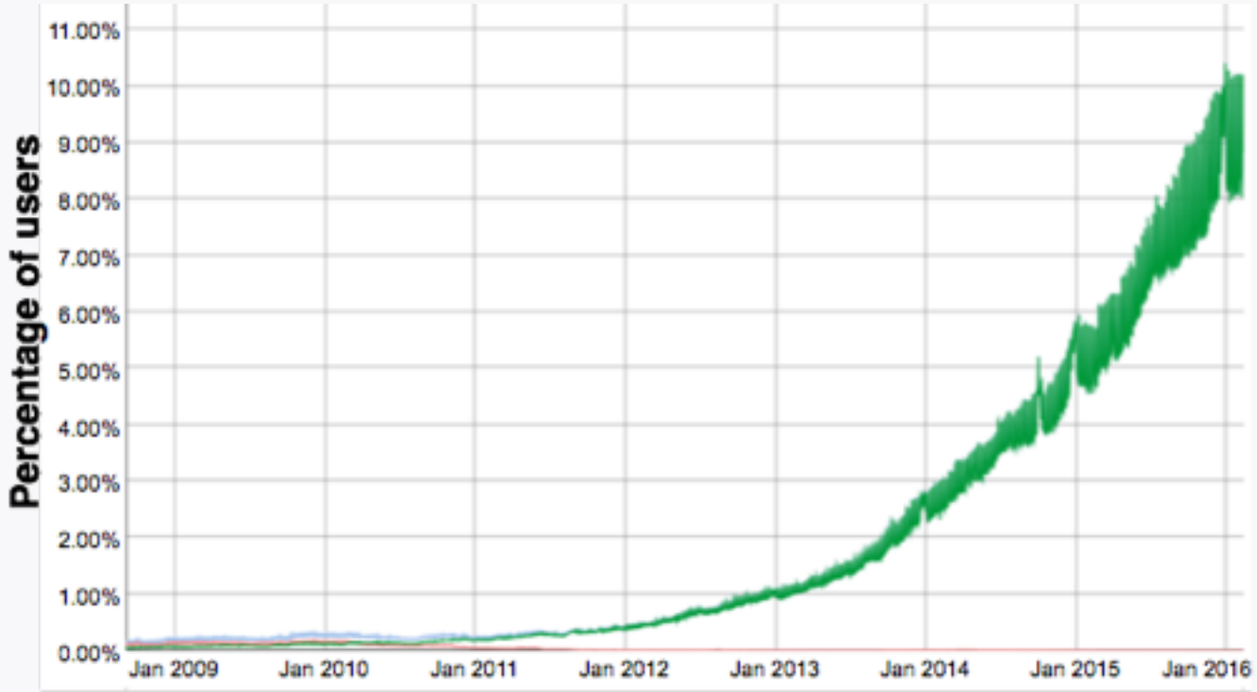

Google IPv6 Statistics [Source: https://www.google.com/intl/en/ipv6/statistics.html]

- Slow adoption of IPv6: only 10% users access Google over IPv6
- Relative performance of IPv6 is a key determinant of a wider adoption of IPv6
- Our focus: IPv6 control and data plane stability

Outline

Goal: Look at performance by analyzing IPv4 and IPv6 **control** and **data plane stability**

1. Control Plane (BGP updates at 5 RouteViews* monitors)
    - How do the routing dynamics differ in IPv4 and IPv6?
    - What are the type of prefixes that contribute to these dynamics?
    - What role does path similarity play in the control plane stability?

2. Data Plane (6 ARK monitors to probe dual-stacked targets)
    - Does the availability of the targets differ over IPv6 than IPv4?
    - Do targets experience higher performance degradation over IPv6 than IPv4?

*RouteViews Project [http://www.routeviews.org]

Approach to study Control plane Stability

Measurement Setup

- BGP updates from five dual-stacked ASes (HE, NTT, Tinet, APAN, IIJ) for quarterly snapshots (January, April, July, October) from 2009 to 2015

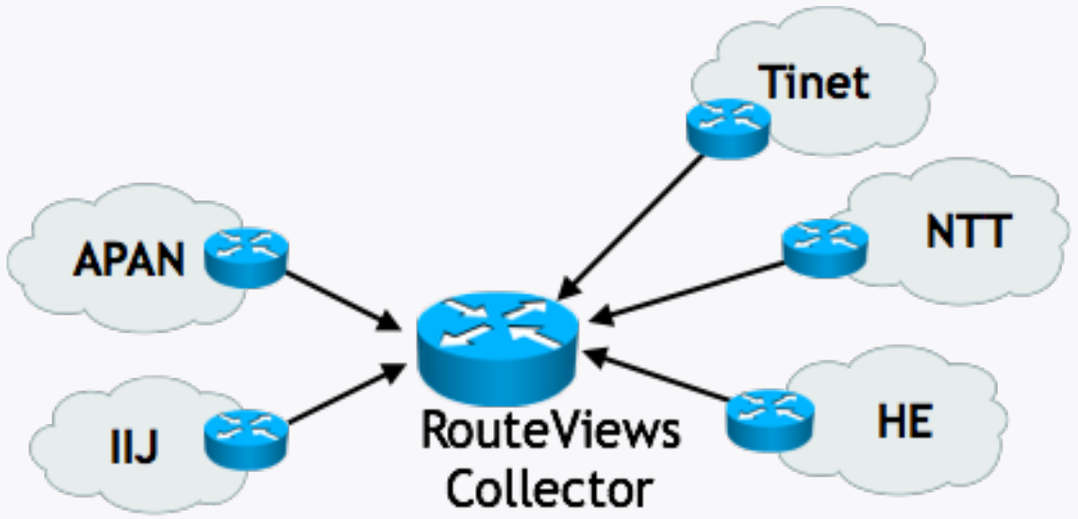

- *Routing event*[*] = consecutive routing updates for the same prefix spaced by 70 seconds or less

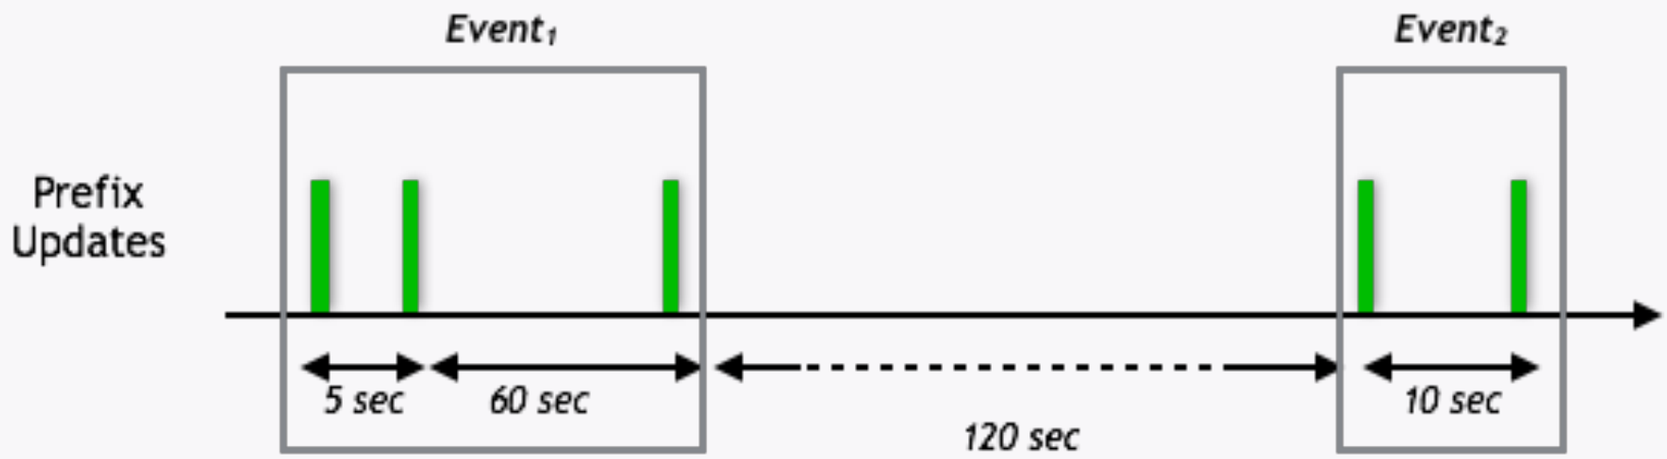

[*] *J. Wu and Z. M. Mao, NSDI 2005*

IPv6 routing system exhibits more routing changes than IPv4

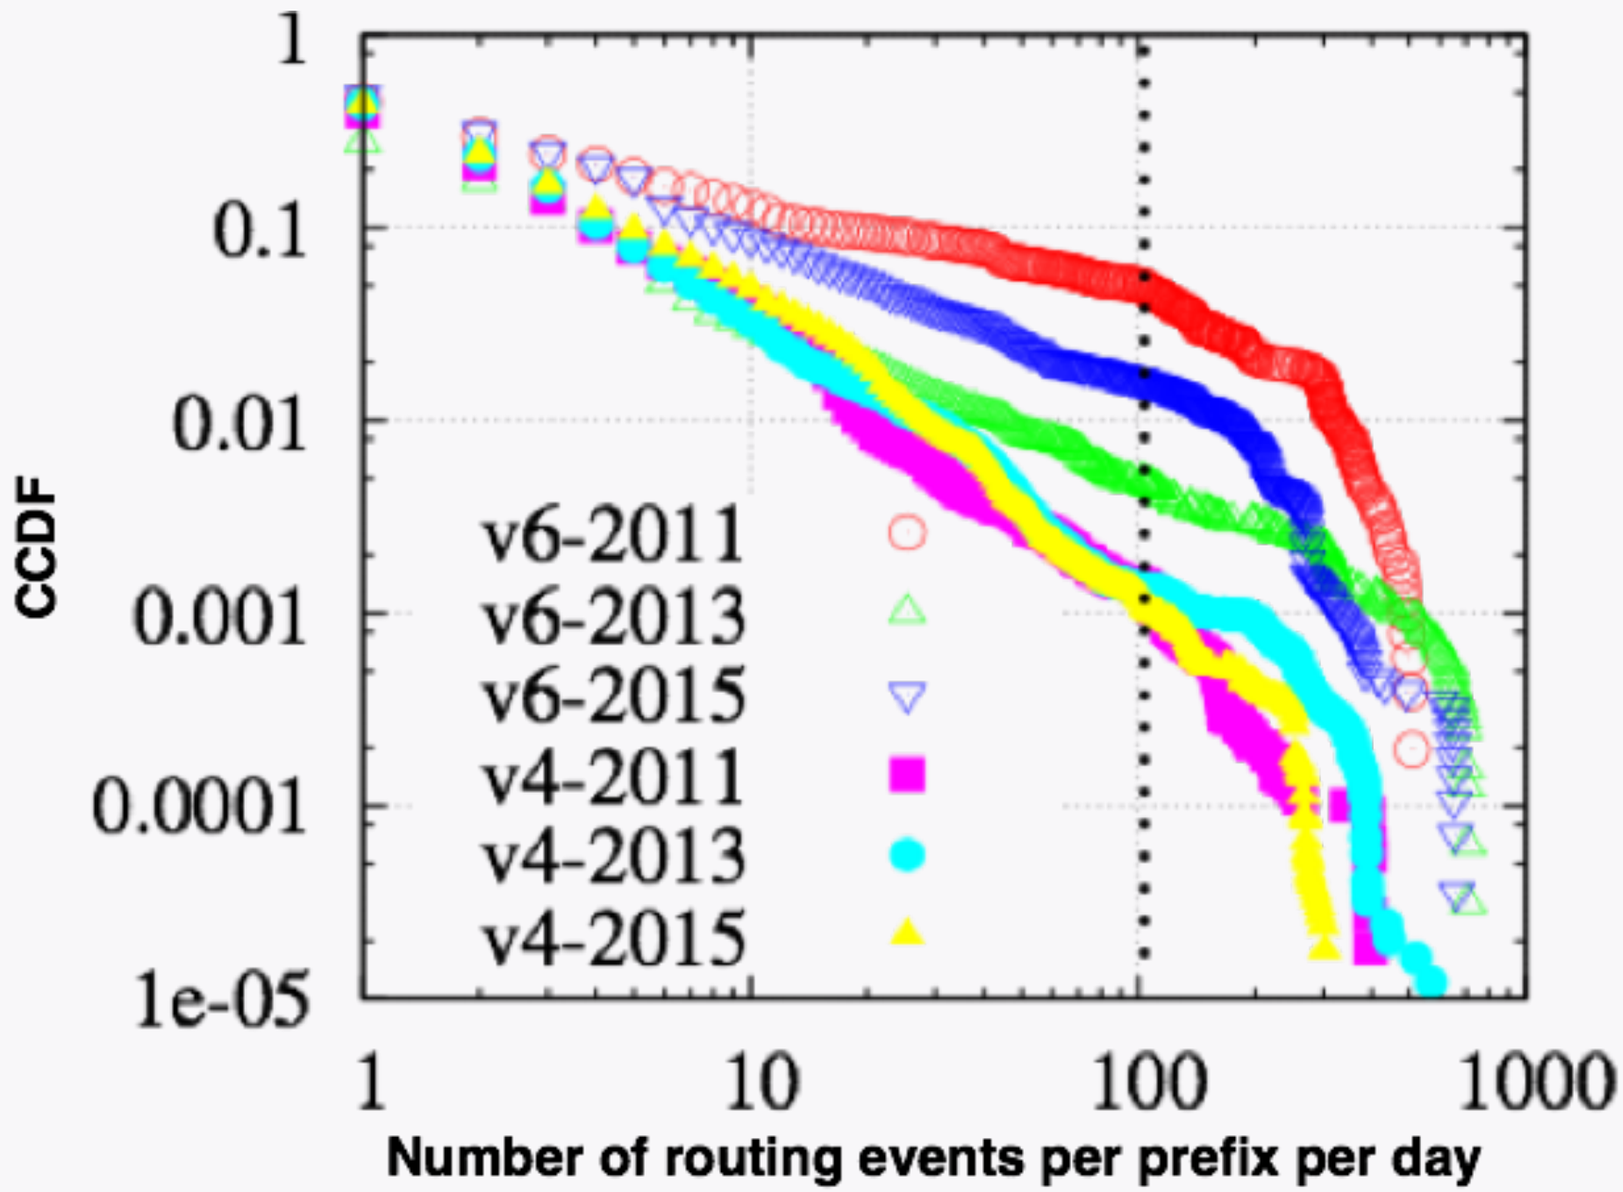

- 0.1% of the IPv4 versus 2% of the IPv6 prefixes experience more than 100 events per day

Overall fraction of active prefixes is higher for IPv6 than IPv4

Active prefix = prefix that experiences a routing change at least once per day

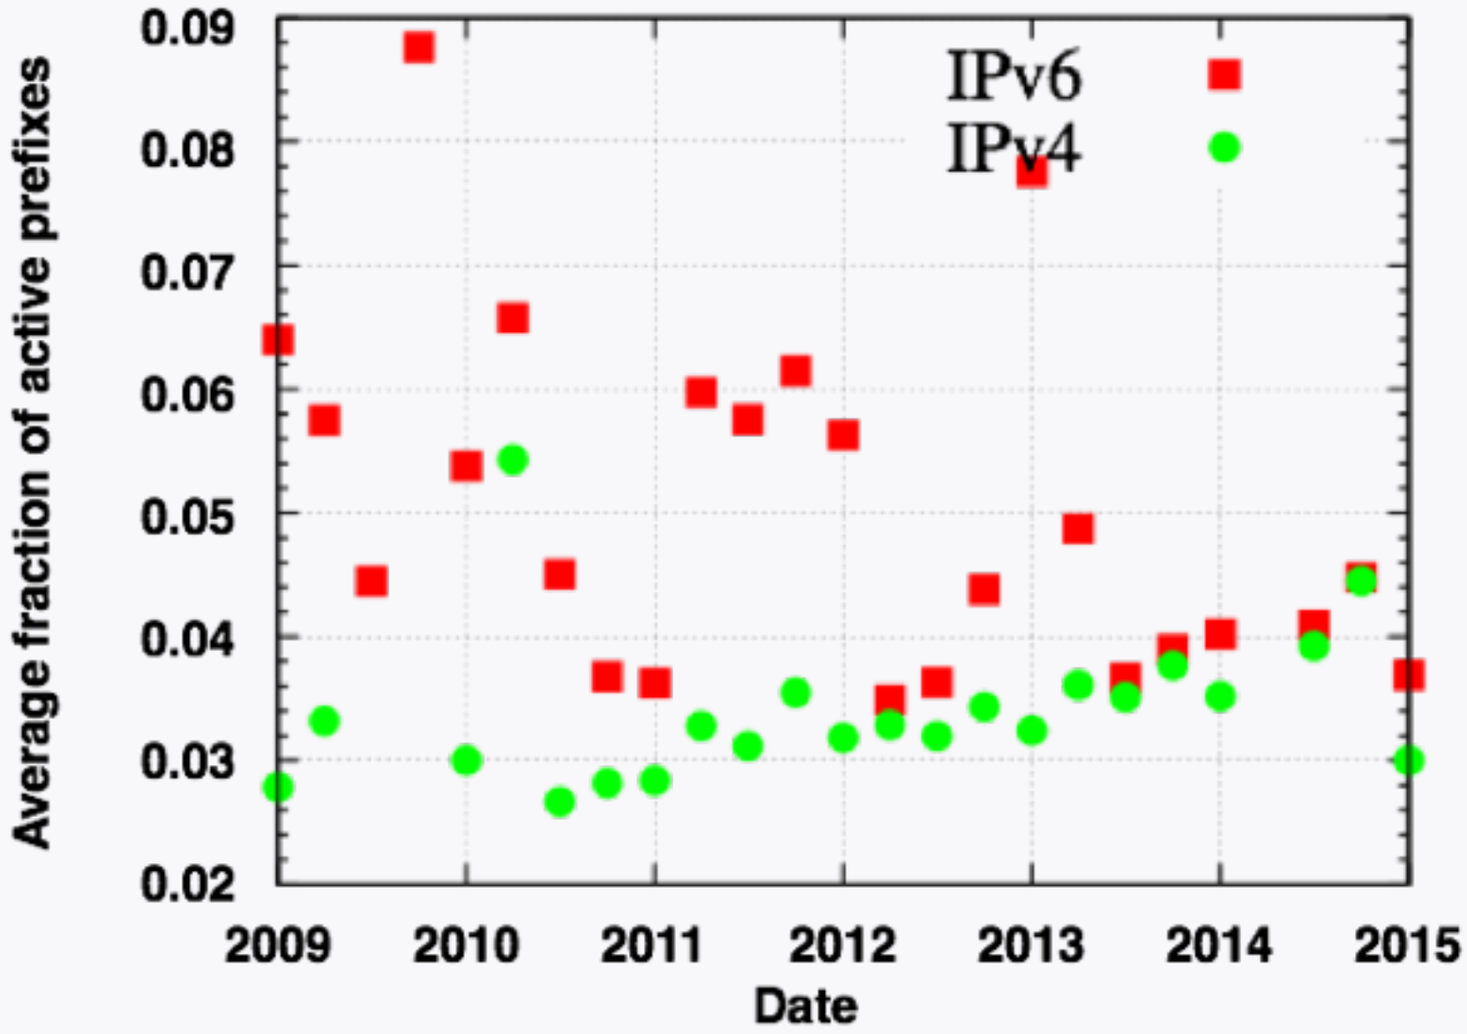

- Fraction of active prefixes is becoming similar in both routing system

Top 1% of active IPv6 prefixes contribute 50% of the overall BGP dynamics

Highly active prefixes = top 1% of the active prefixes in terms of contribution to the BGP dynamics

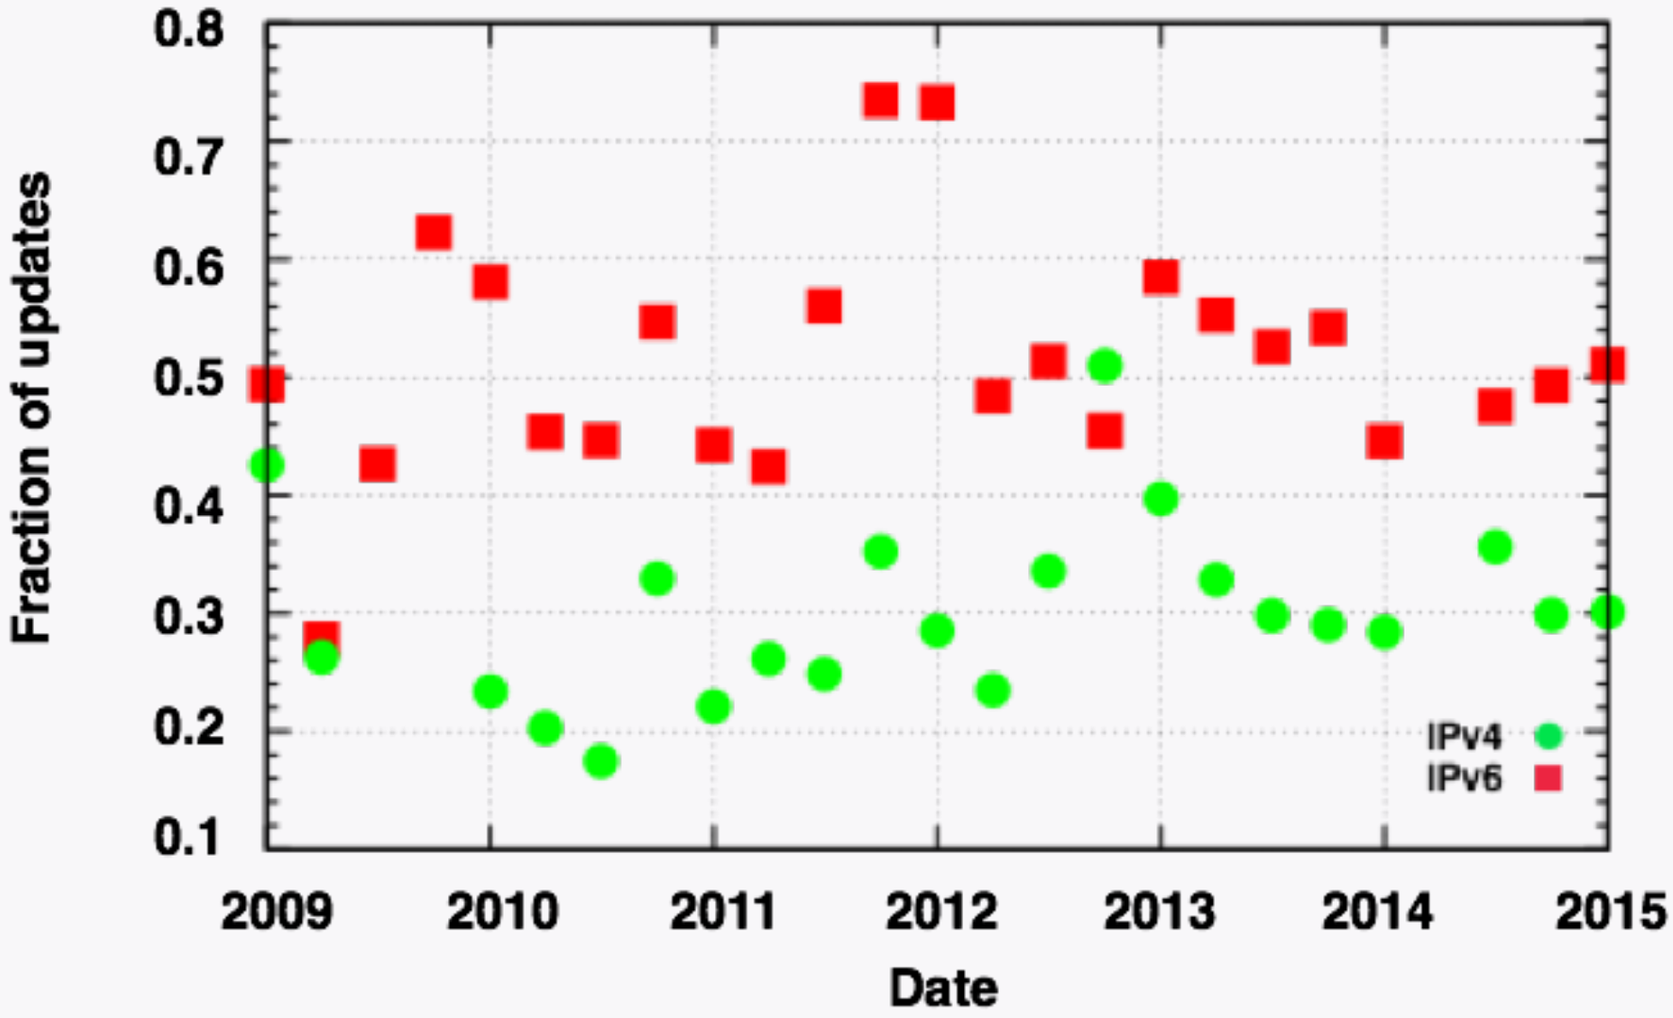

IPv6: 40% - 60%

IPv4: 20% - 30%

- Is the difference in the contribution caused by the relative immaturity of IPv6?

Correlating IPv4 and IPv6 instabilities

Approach:

- Build **instability time windows** by grouping events that affect IPv4 and IPv6 prefixes for the same network

- Determine the overlapping periods between the IPv4 and IPv6 instability time windows for congruent and non-congruent AS paths

- Compute the **correlation fraction** as the fraction of overlapping periods for the same network

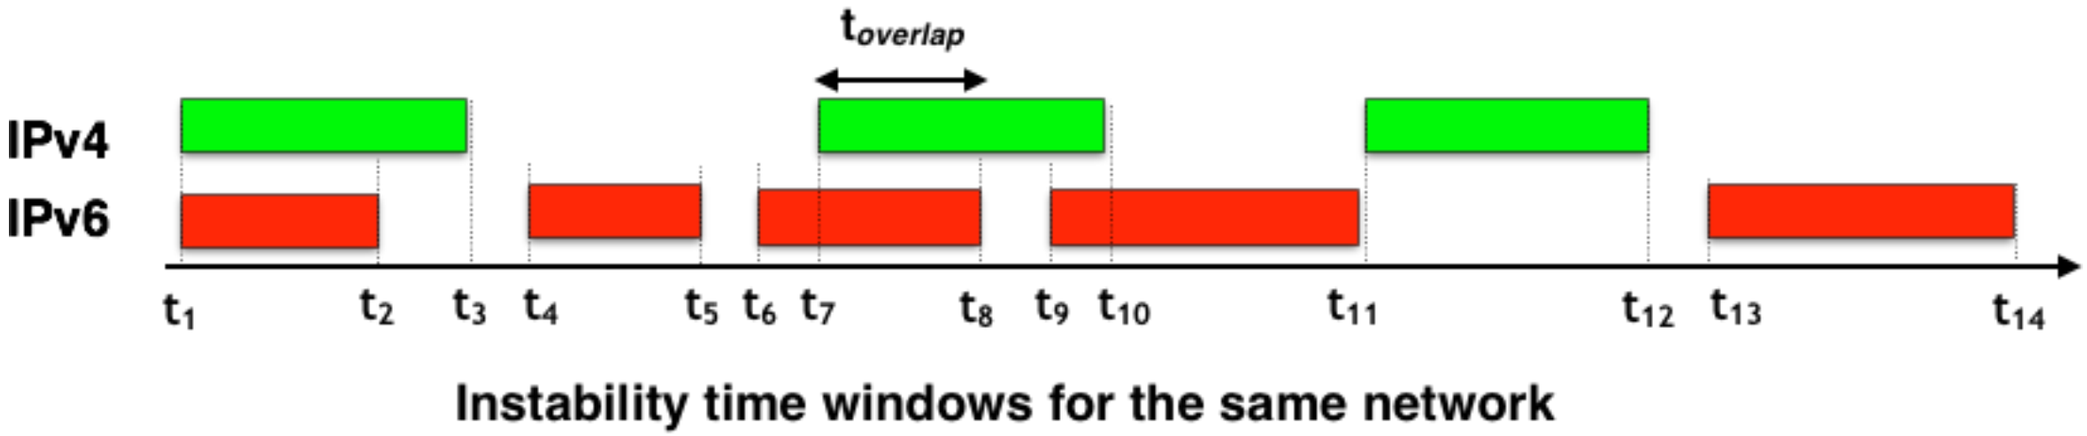

Instability time windows for the same network

IPv4 and IPv6 events show a higher correlation for congruent than non-congruent paths

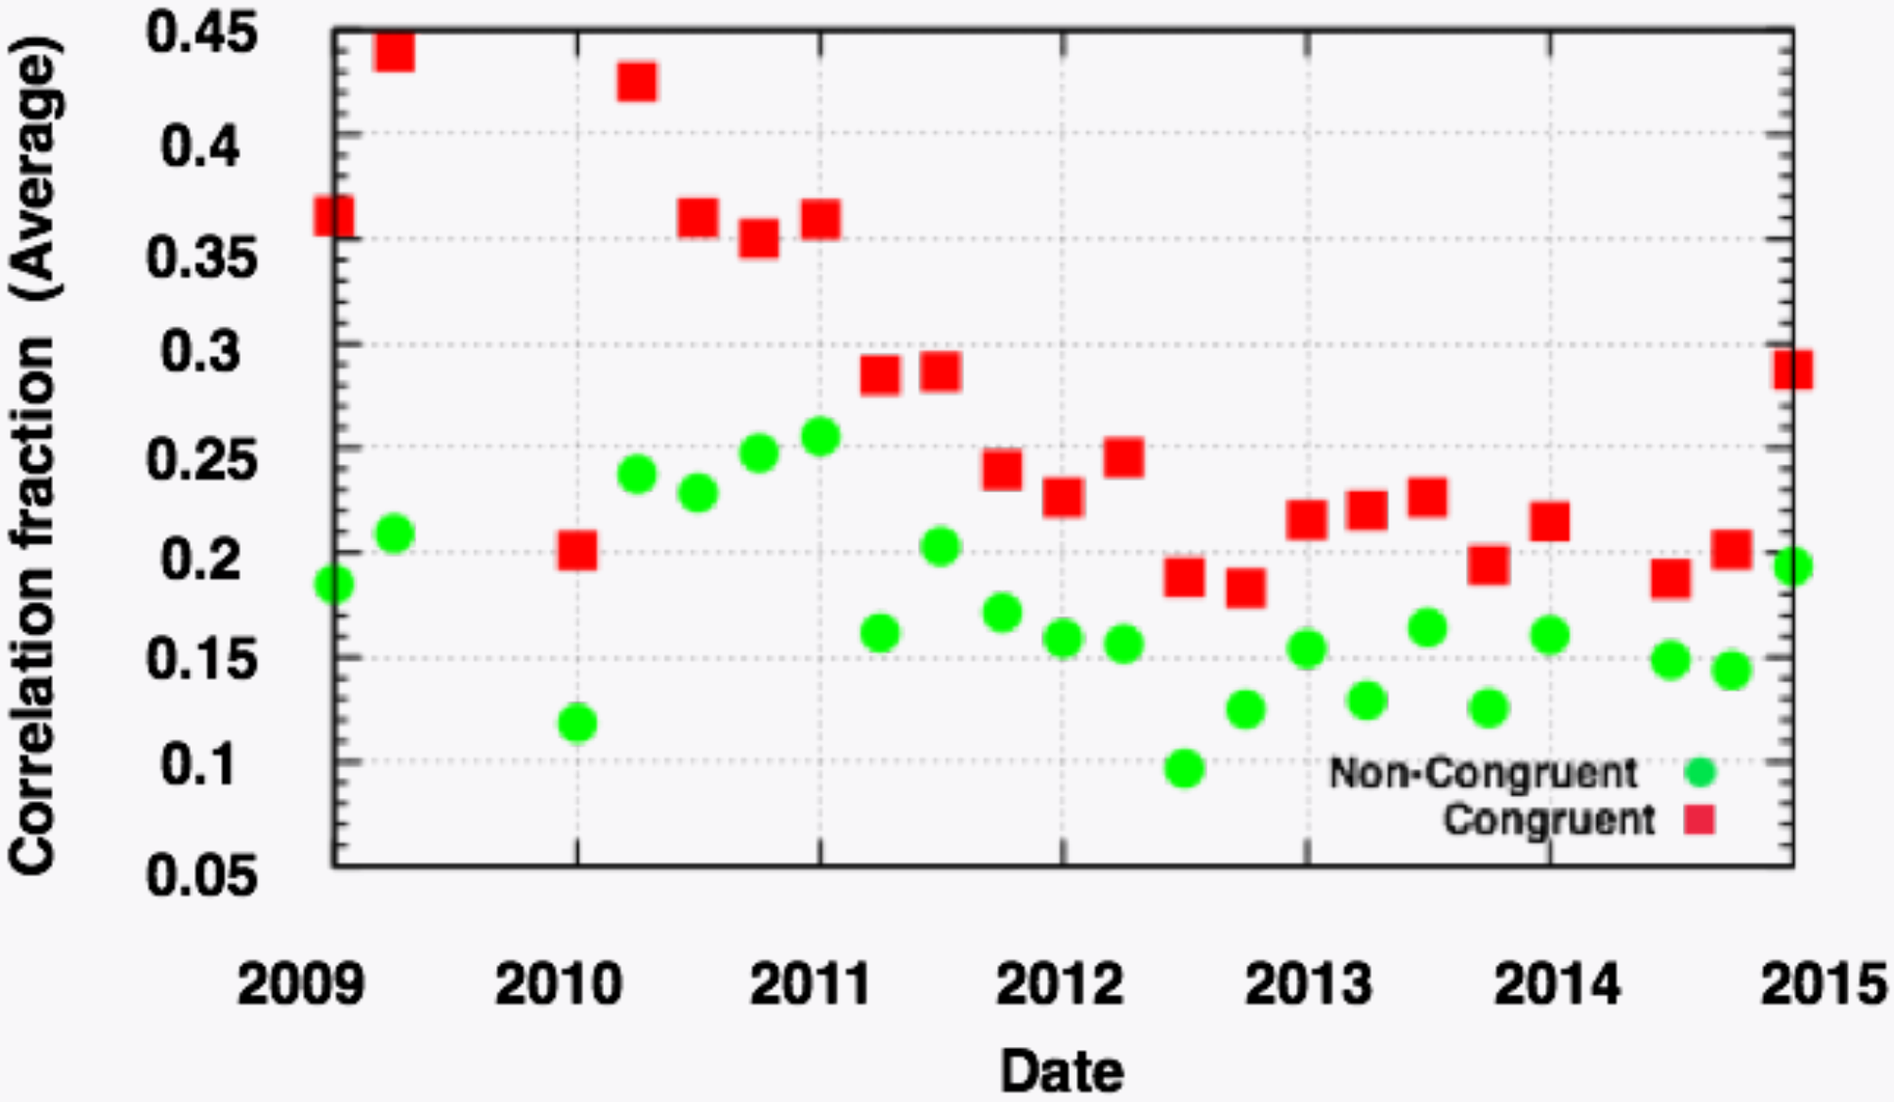

- Overall low correlation fraction (< 0.5) for both congruent and non-congruent paths indicates that IPv4 and IPv6 routing systems do not share the same fate

Takeaways: Control Plane analysis

- IPv6 routing system is less stable than IPv4

- High percentage of the churn in IPv6 is generated by a small set of unstable prefixes

- Low correlation of instability periods in IPv4 and IPv6 for both congruent and non-congruent paths

- Difference in the event composition for IPv4 and IPv6 that hints at the lack of path diversity of IPv6 internetwork (details in the paper)

Measure Data plane Stability

Goal: Study **reachability** and **performance** (relative RTT and RTT instabilities)

Measurement setup:

- Six monitors from the ARK infrastructure to ping (every 5 seconds) and run traceroute (every 2 hours) towards 629 dual-stacked targets* for 1 ½ months (August – September 2014)

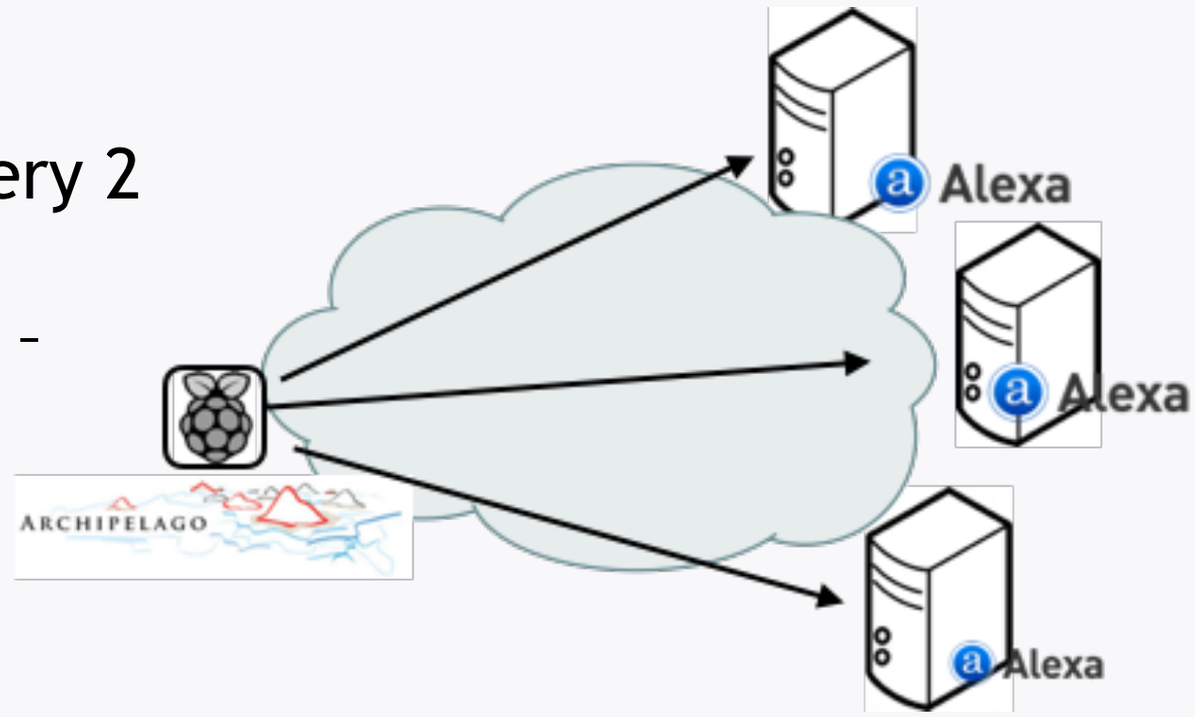

*Alexa [http://www.alexa.com/topsites]

Network reachability higher over IPv4 than IPv6

Reachability period (for a target) = period of time when the probed target was responsive

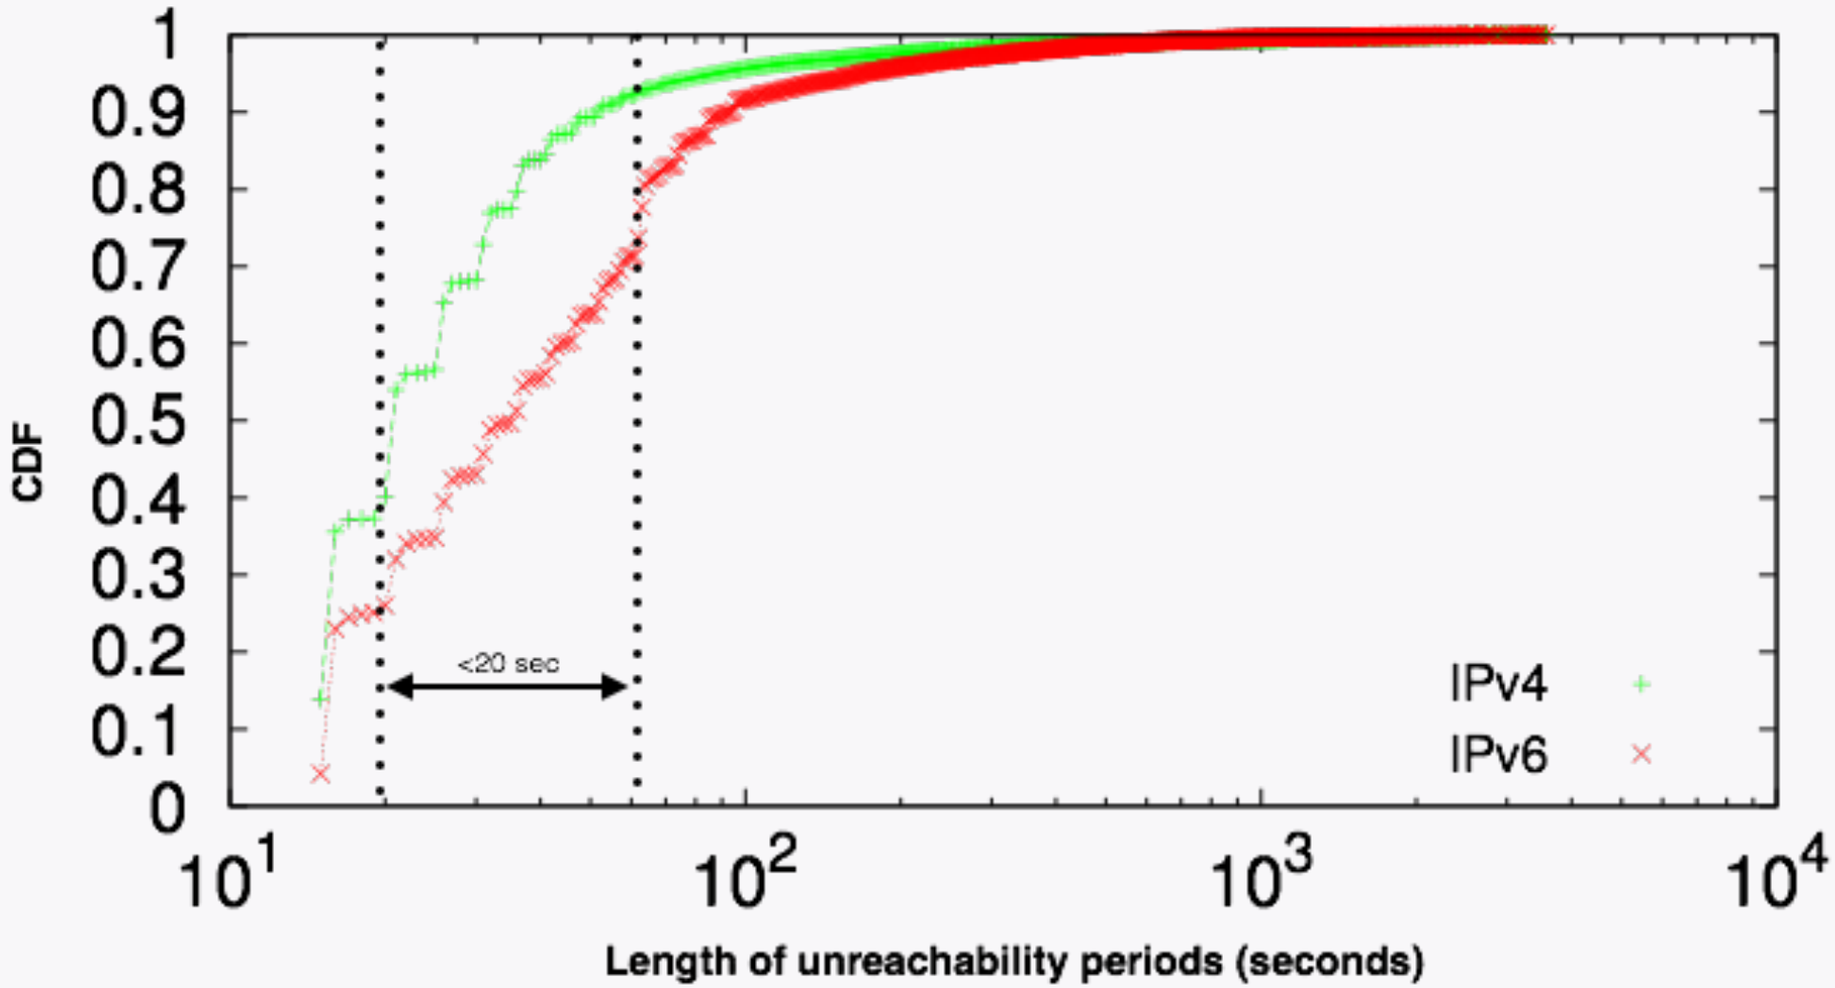

- 91.94% of the targets were reachable over IPv4 and IPv6 in 99% of the probing period
- Longer unreachability intervals over IPv6 than IPv4

Performance: Relative RTT*

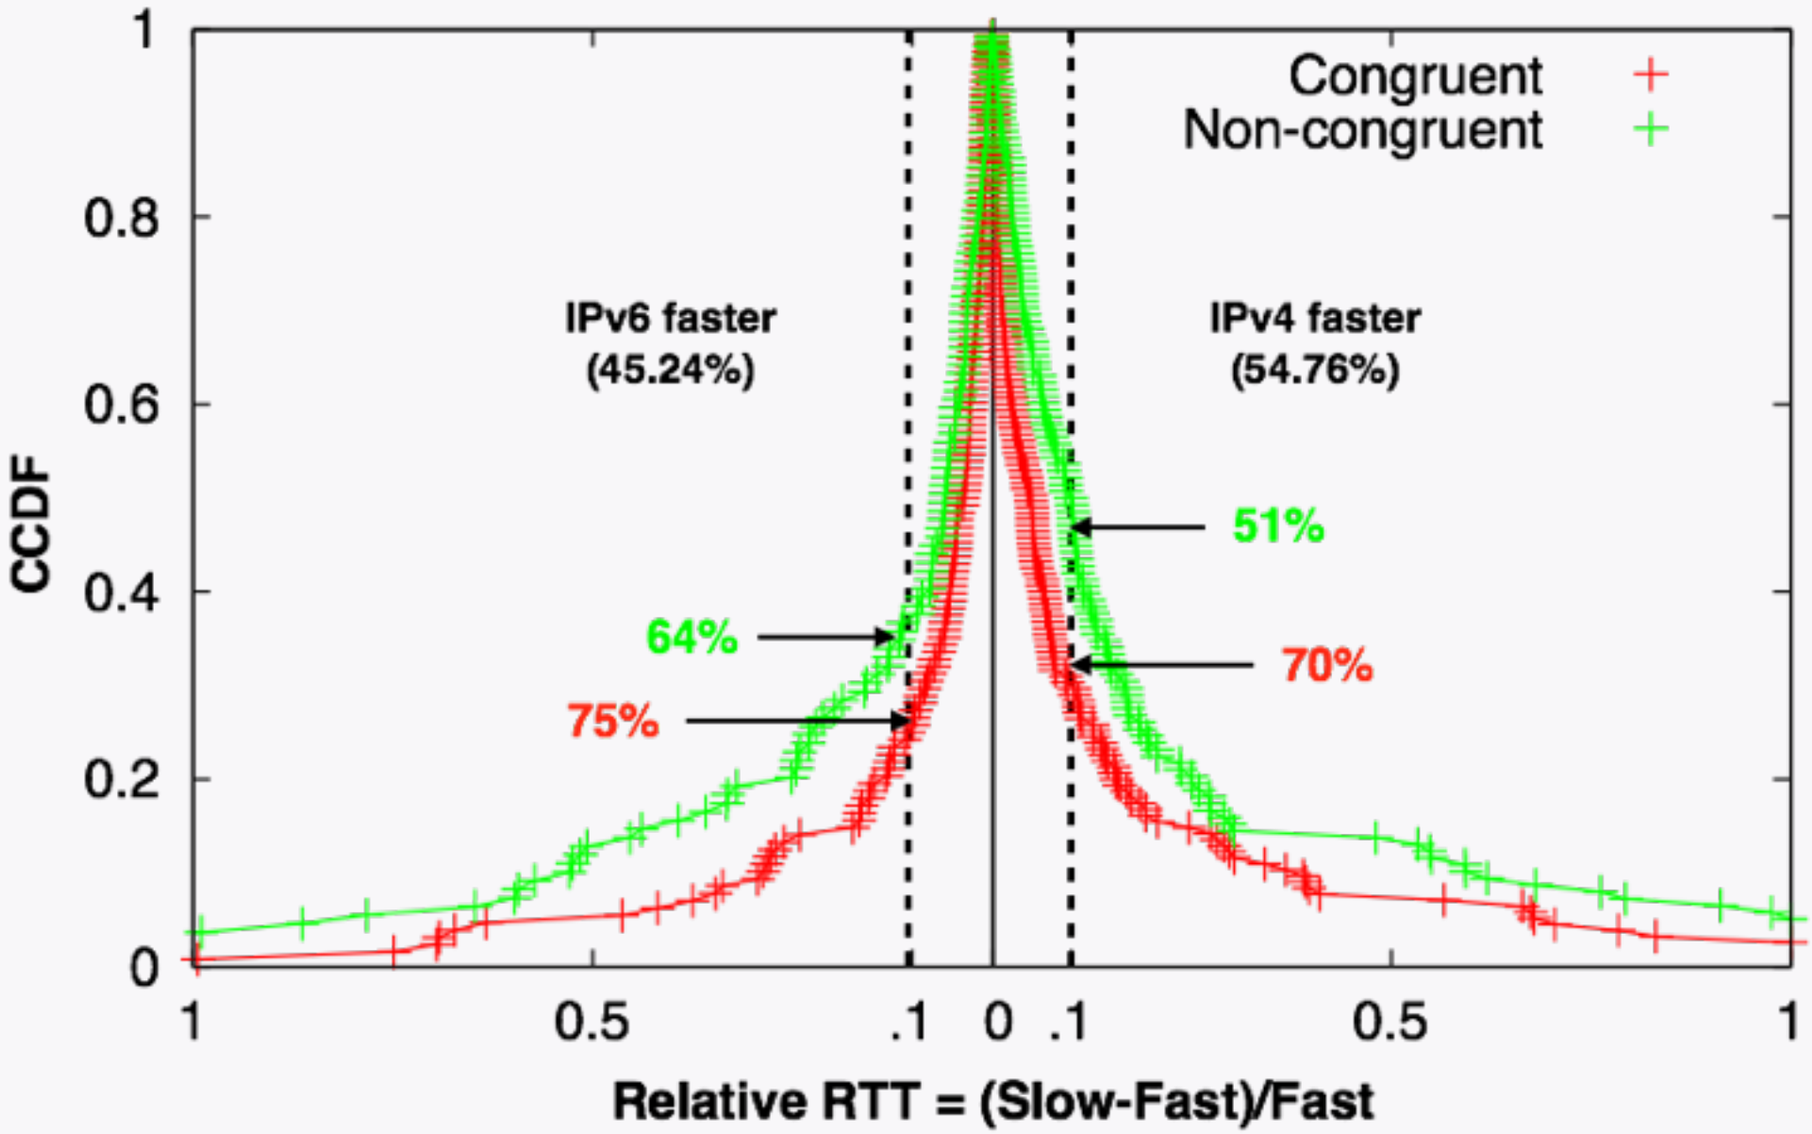

- Previous Study*: IPv6 faster for 22% of the targets
- IPv6 is ***maturing***: comparable number of targets for which IPv6 has similar performance with IPv4, and vice-versa
- Higher percentage of targets with congruent than non-congruent paths experience a similar performance for both IPv4 and IPv6

*Dhamdhere et al., IMC 2012

Performance: RTT Instabilities

- **Detecting** RTT instabilities over IPv4 and IPv6

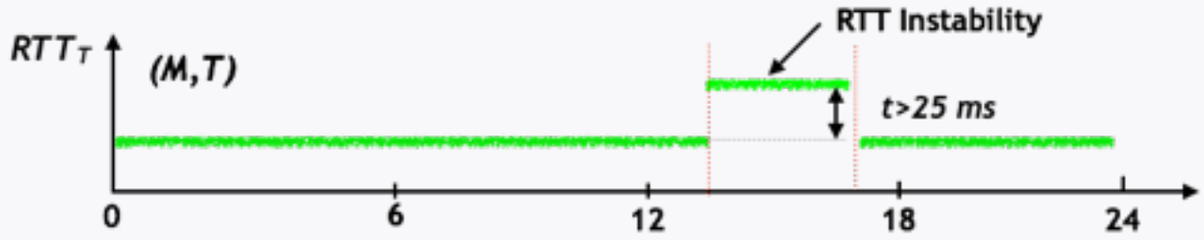

- **Localization**:
  - Identify the first hop on the forward path that shows an increase in the RTT value as the location of the increase
  - Identify the time interval of the day when the RTT instability occurs

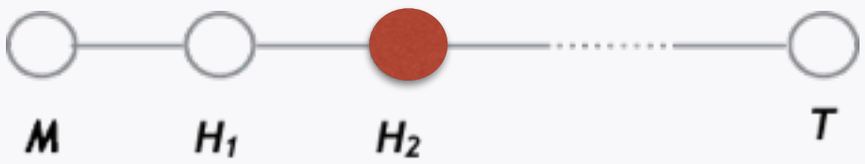

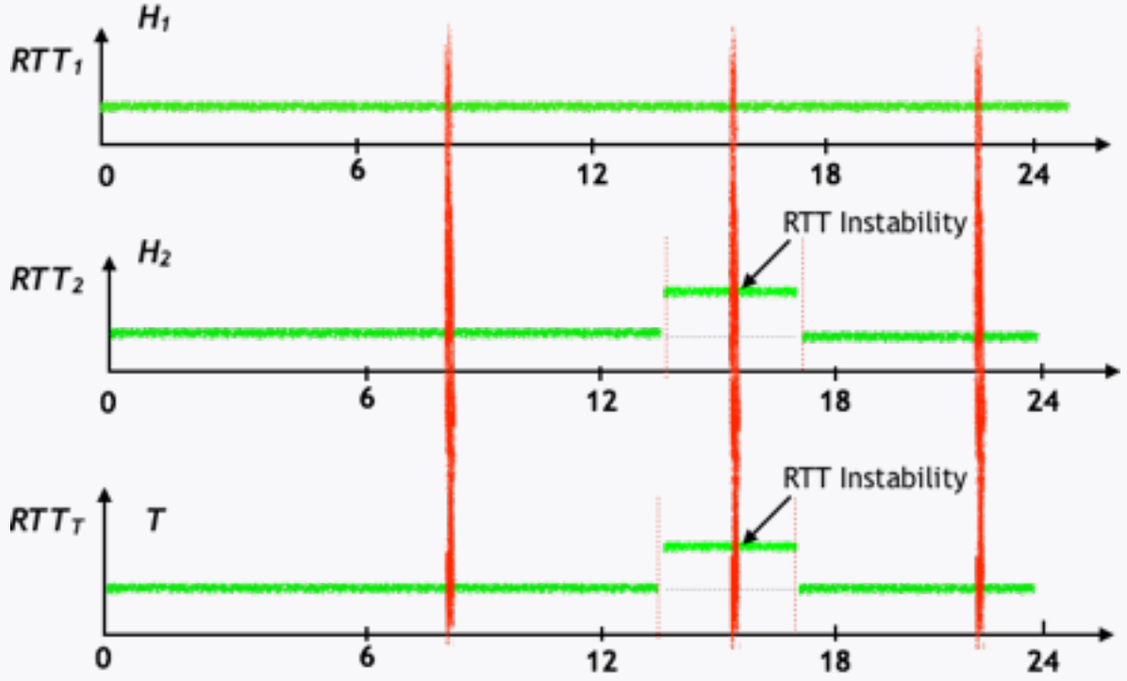

- **Identifying shared infrastructure**: Use DNS data to identify common hop on the IPv4 and IPv6 paths

High number of targets experience RTT instabilities

- 70% of the probed targets experience RTT instabilities; More than half of these experience performance degradation over both IPv4 and IPv6

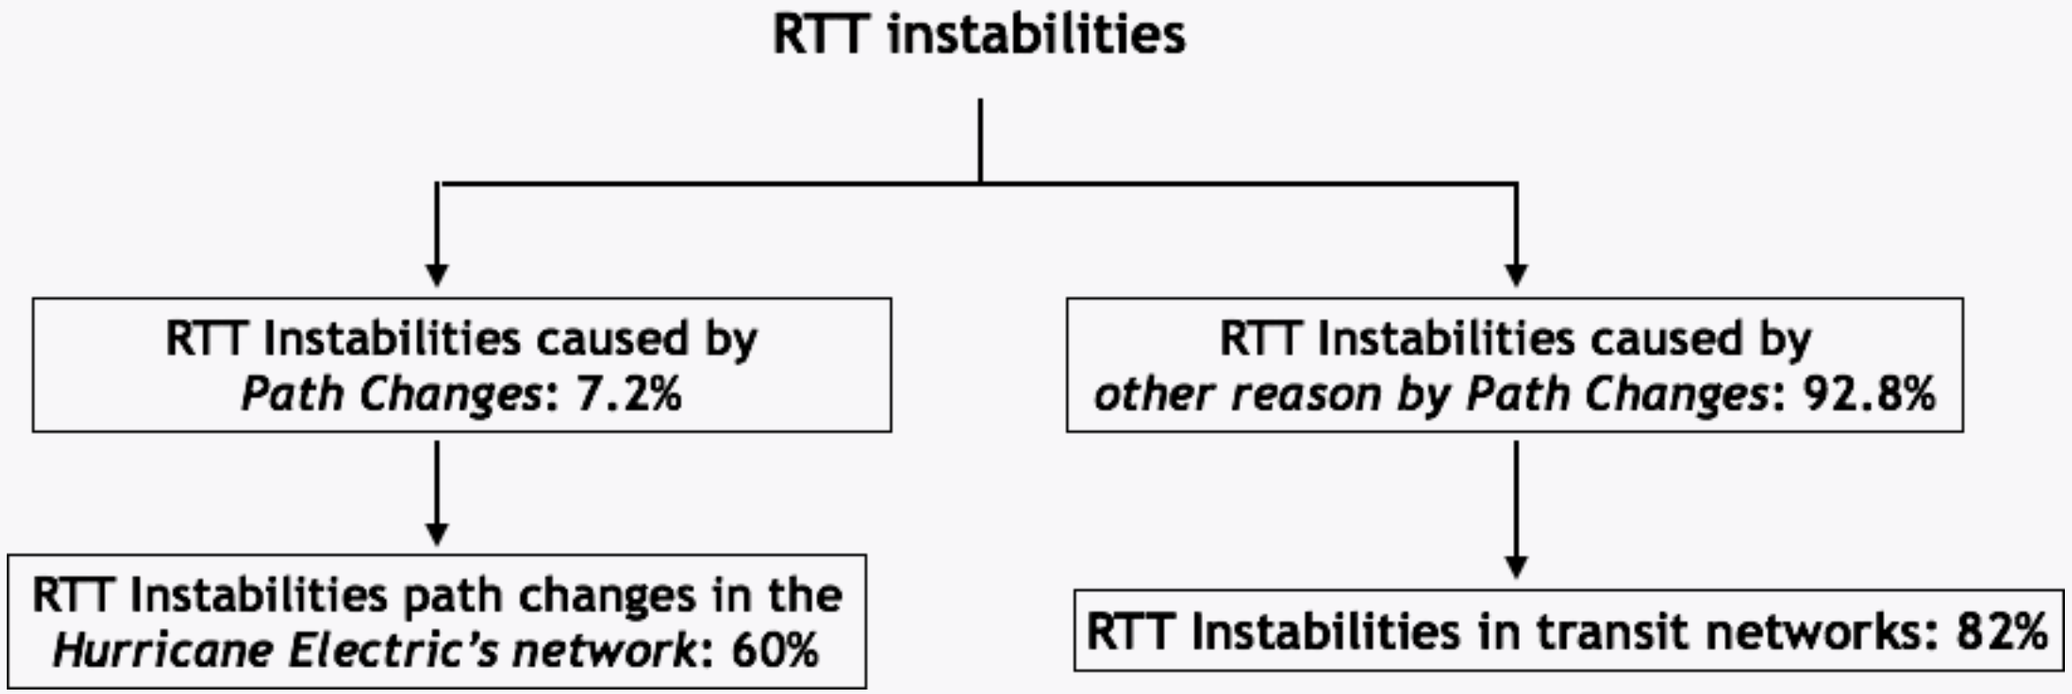

- Changes in HE (Hurricane Electric) have the potential to affect a large number of end-to-end paths

Congruency matters: Shared infrastructure can cause correlated performance degradations

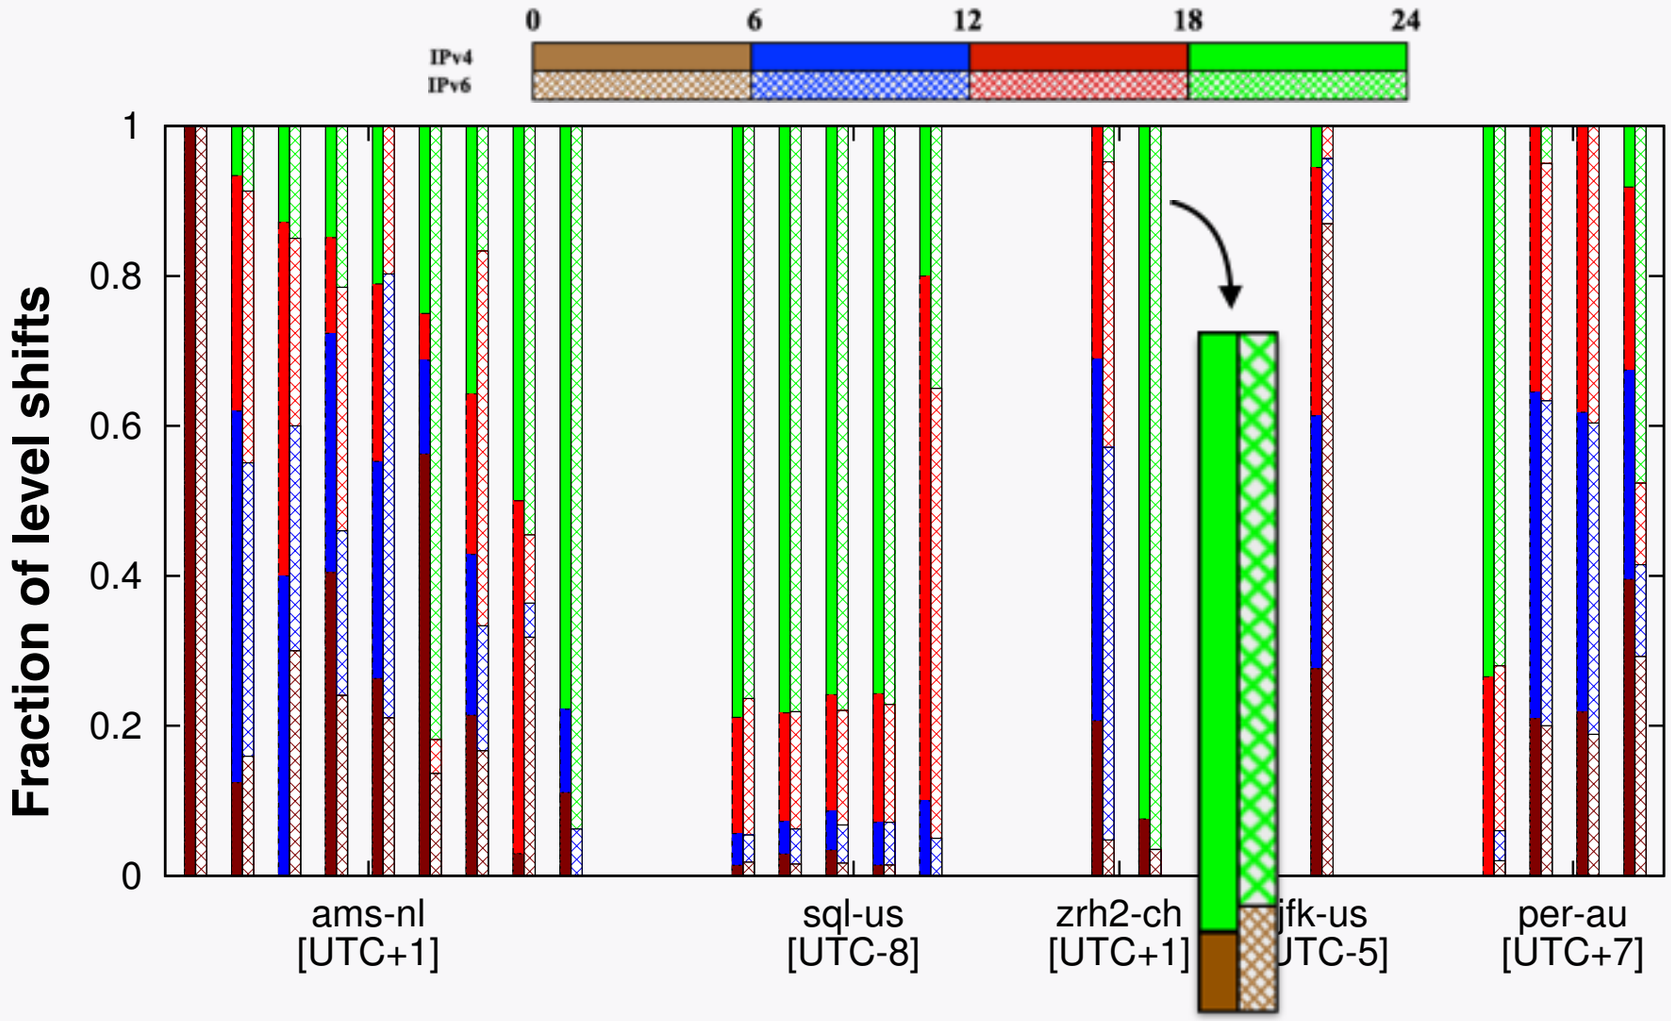

90% of the level shifts over IPv4 and IPv6 occur within the same interval

- RTT increases observed both during peak and non-peak hours* over both IPv4 and IPv6

*FCC, *Measuring Broadband America*

Conclusions

- IPv6 control and data plane stability are comparable to IPv4

- Relative immaturity and topological sparseness of IPv6:
  - Most IPv6 routing dynamics are generated by a small fraction of pathologically unstable prefixes
  - Low correlation of the instability IPv4 and IPv6 events per networks across congruent and non-congruent paths

- RTT performance over IPv6 is becoming markedly similar to IPv4

- Severe RTT degradations are equally likely to affect IPv4 and IPv6 paths

Thank you!