MAPPING INTERDOMAIN CONGESTION

Amogh Dhamdhere, Matthew Luckie
Bradley Huffaker, Young Hyun, Kc Claffy (CAIDA)
Steve Bauer, David Clark (MIT)

amogh@caida.org

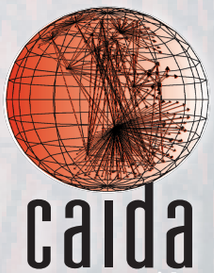

INTERDOMAIN CONGESTION

- We are developing methods to measure the location and extent of interdomain congestion

- **Our goals (1) atlas of interdomain links and their congestion state, (2) improve transparency, empirical grounding of debate**

- This is early work: we just started a 3 year NSF-funded project on topology+congestion measurement

MEASURING LINK CONGESTION

- Internet traffic generally shows diurnal patterns
- When links become congested, queues fill up, network delay and loss rate increases
- Delay increase is related to queue size of congested router

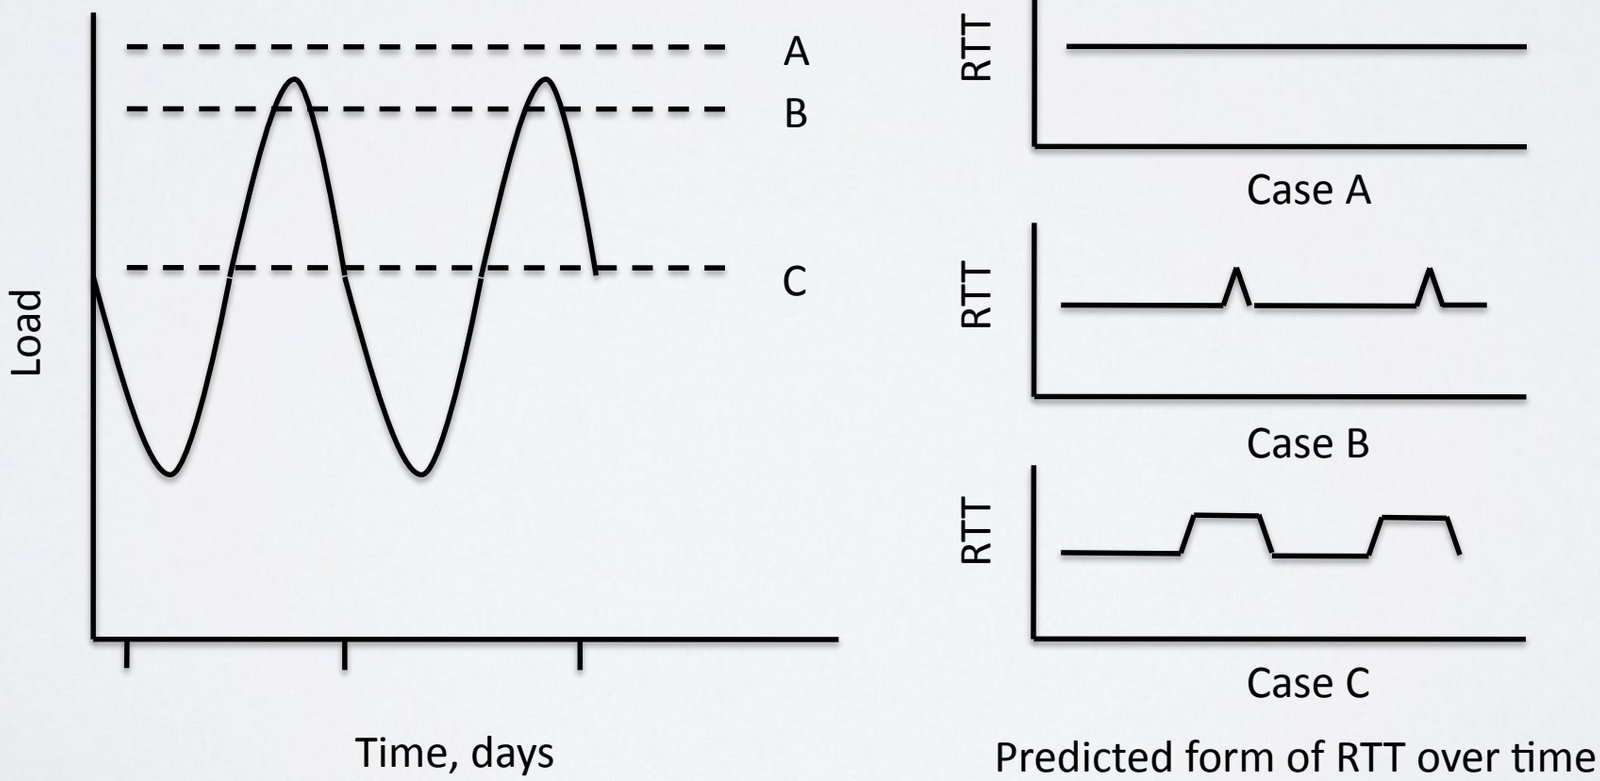

Predicted form of RTT over time

METHOD: TIME SERIES PING

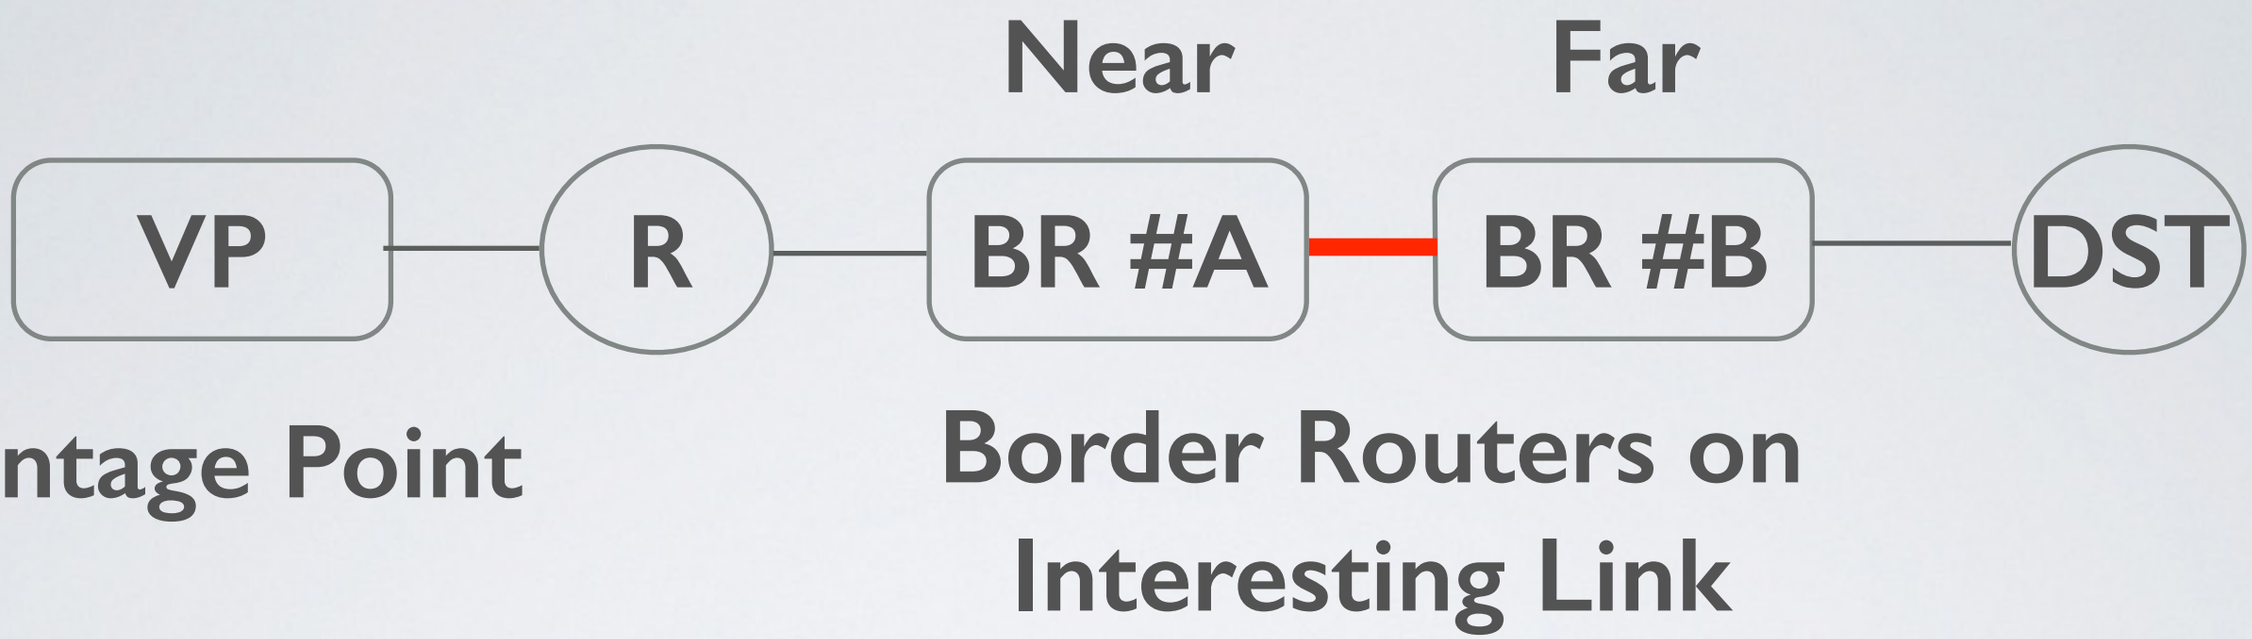

METHOD: TIME SERIES PING

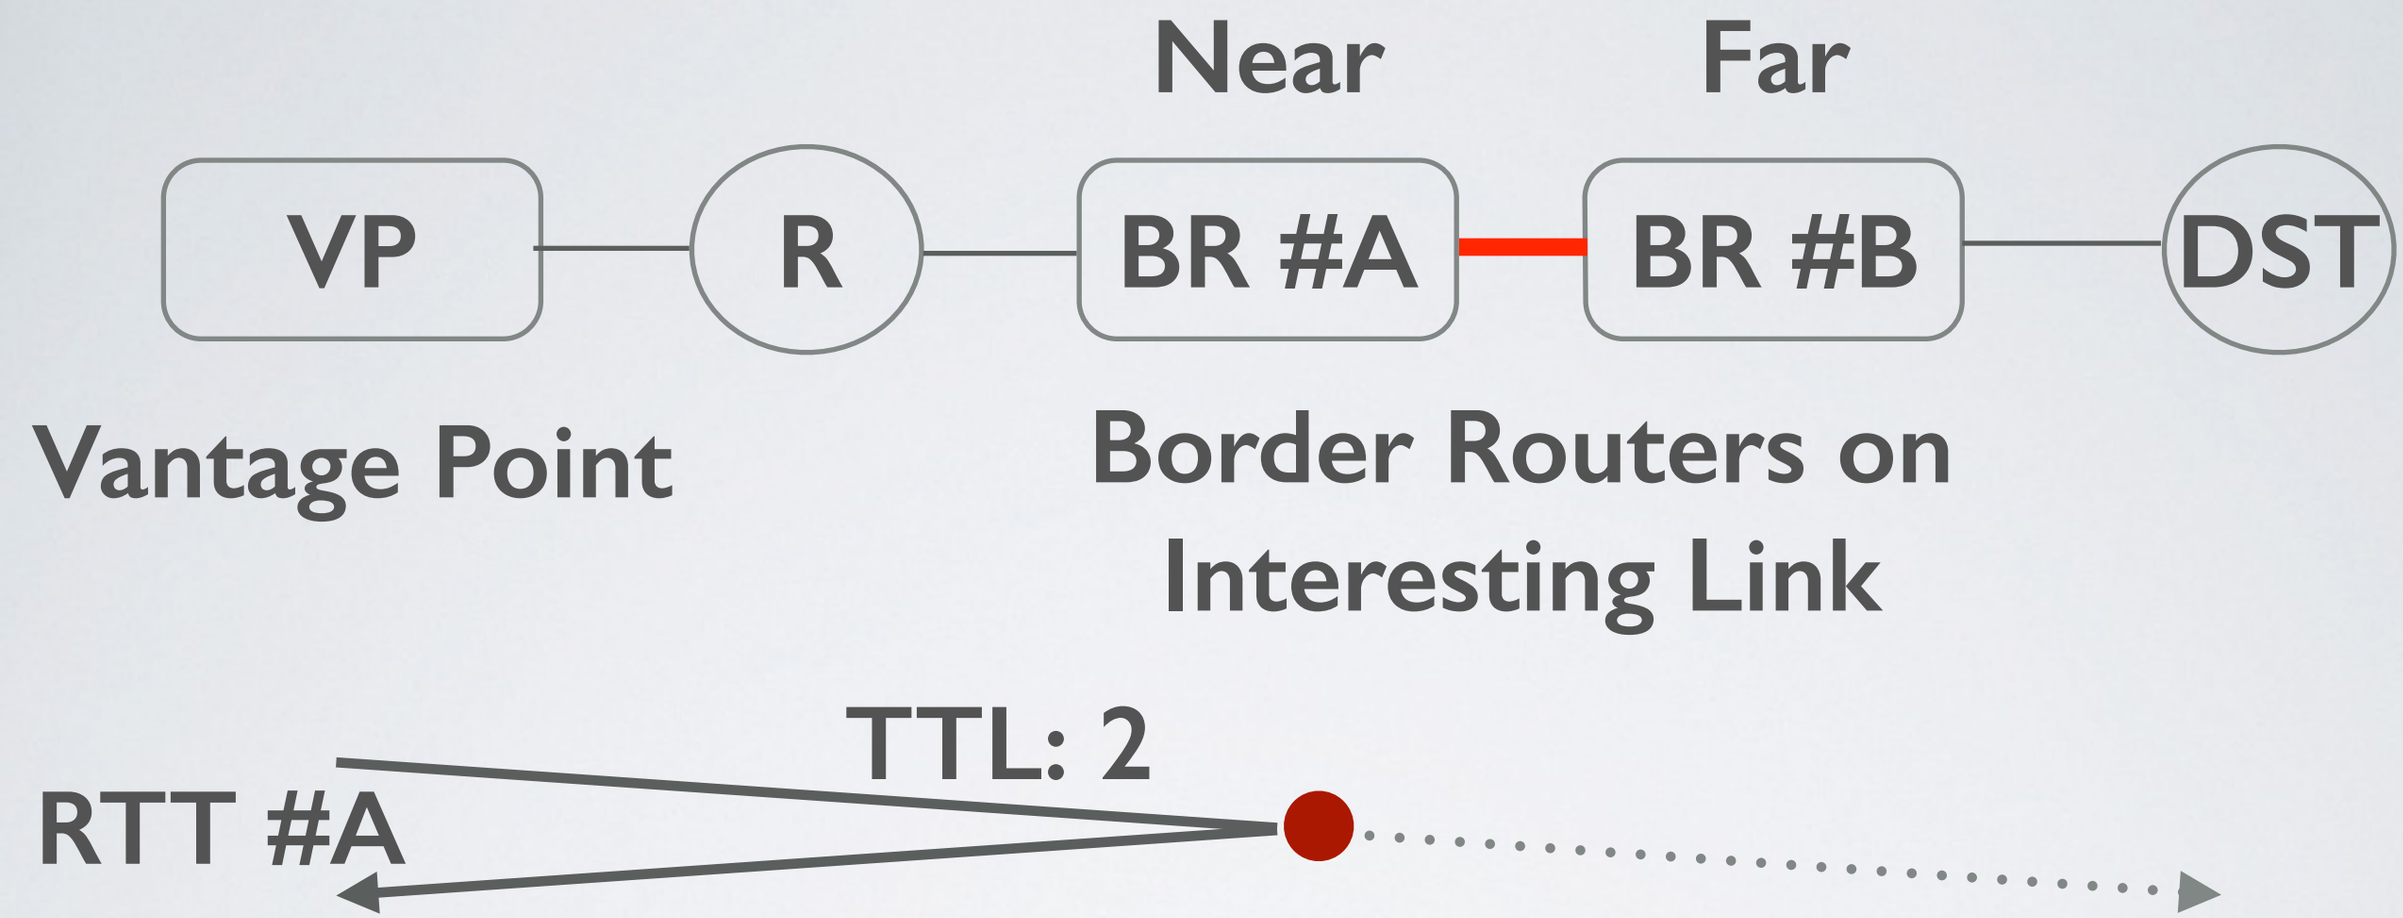

METHOD: TIME SERIES PING

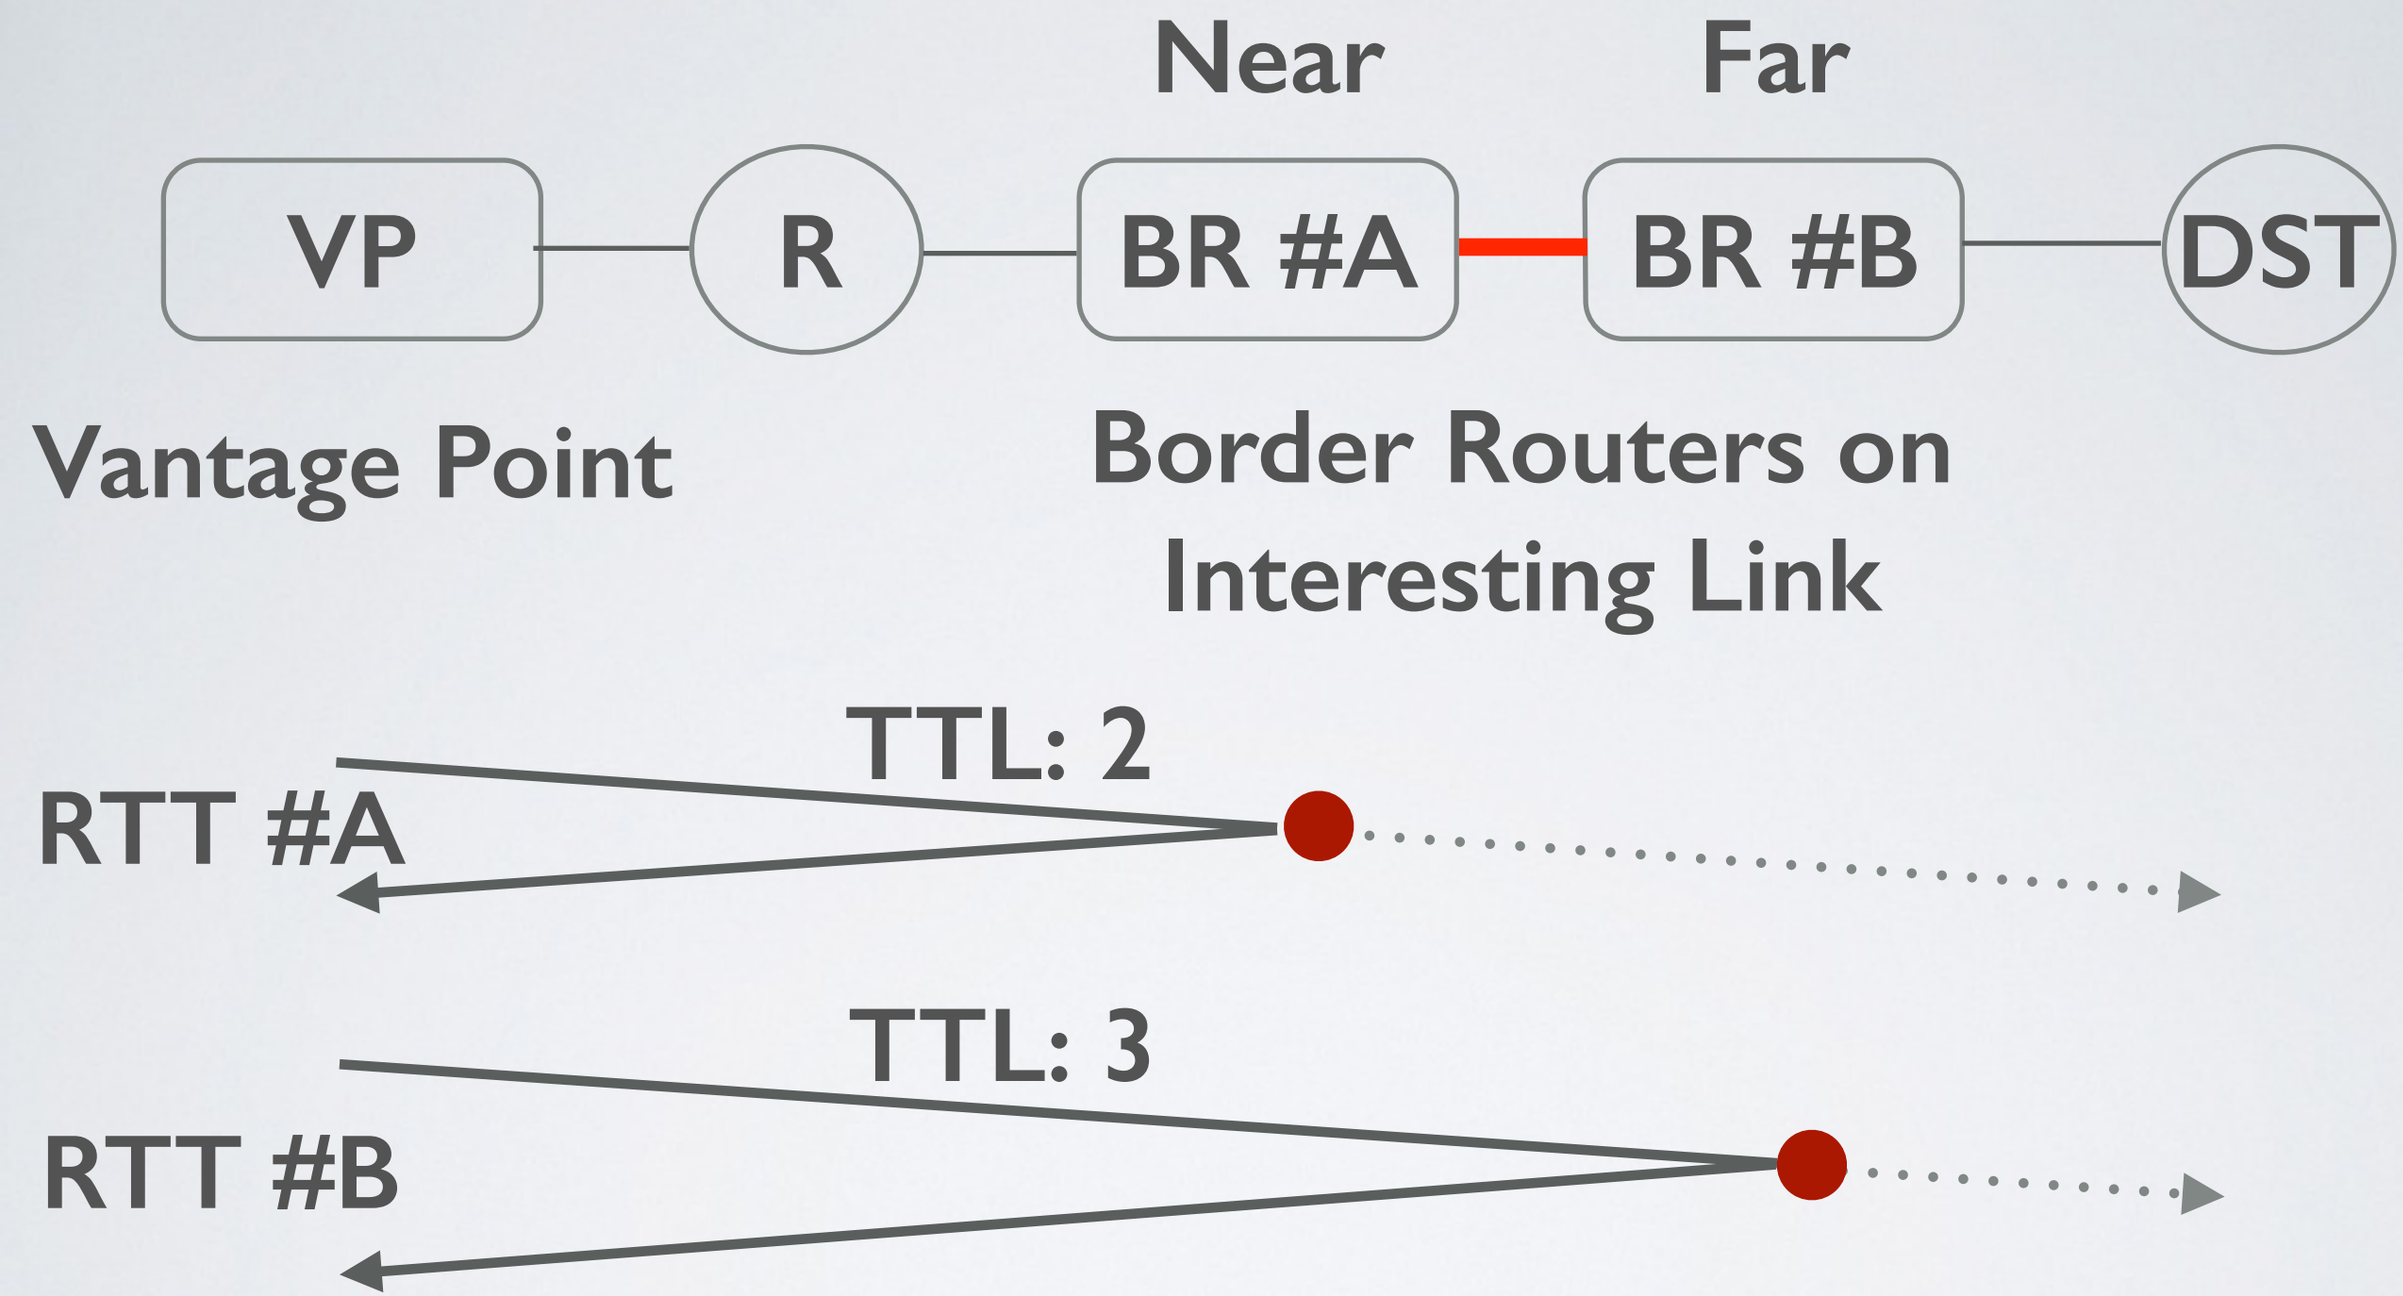

METHOD: TIME SERIES PING

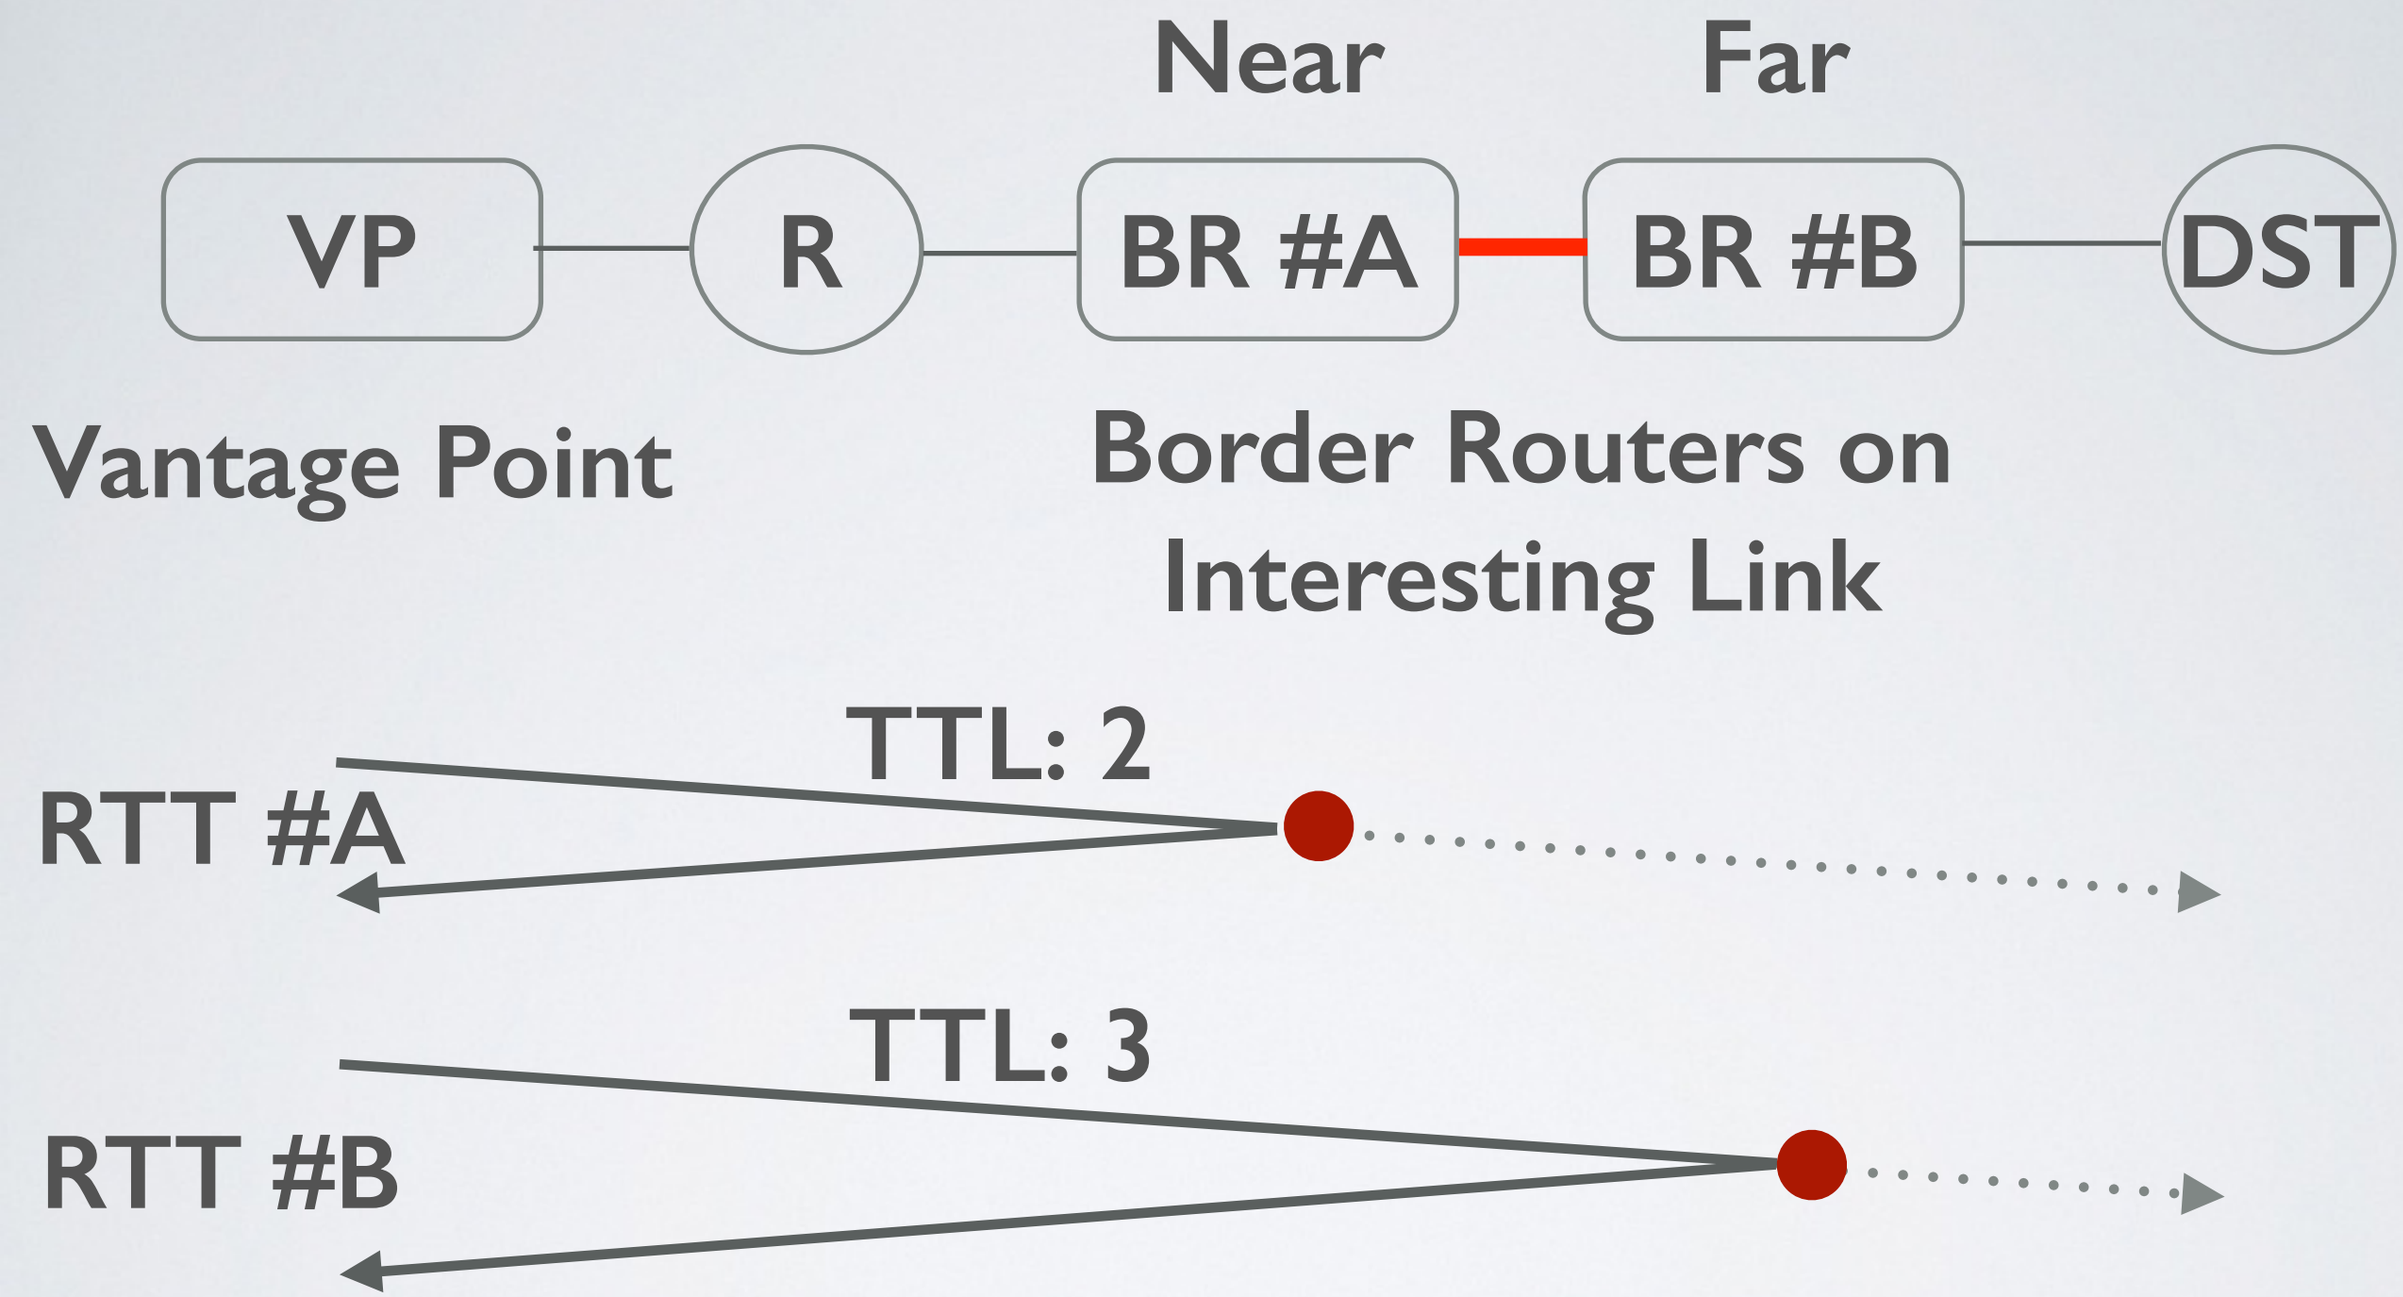

(repeat to obtain a time series)

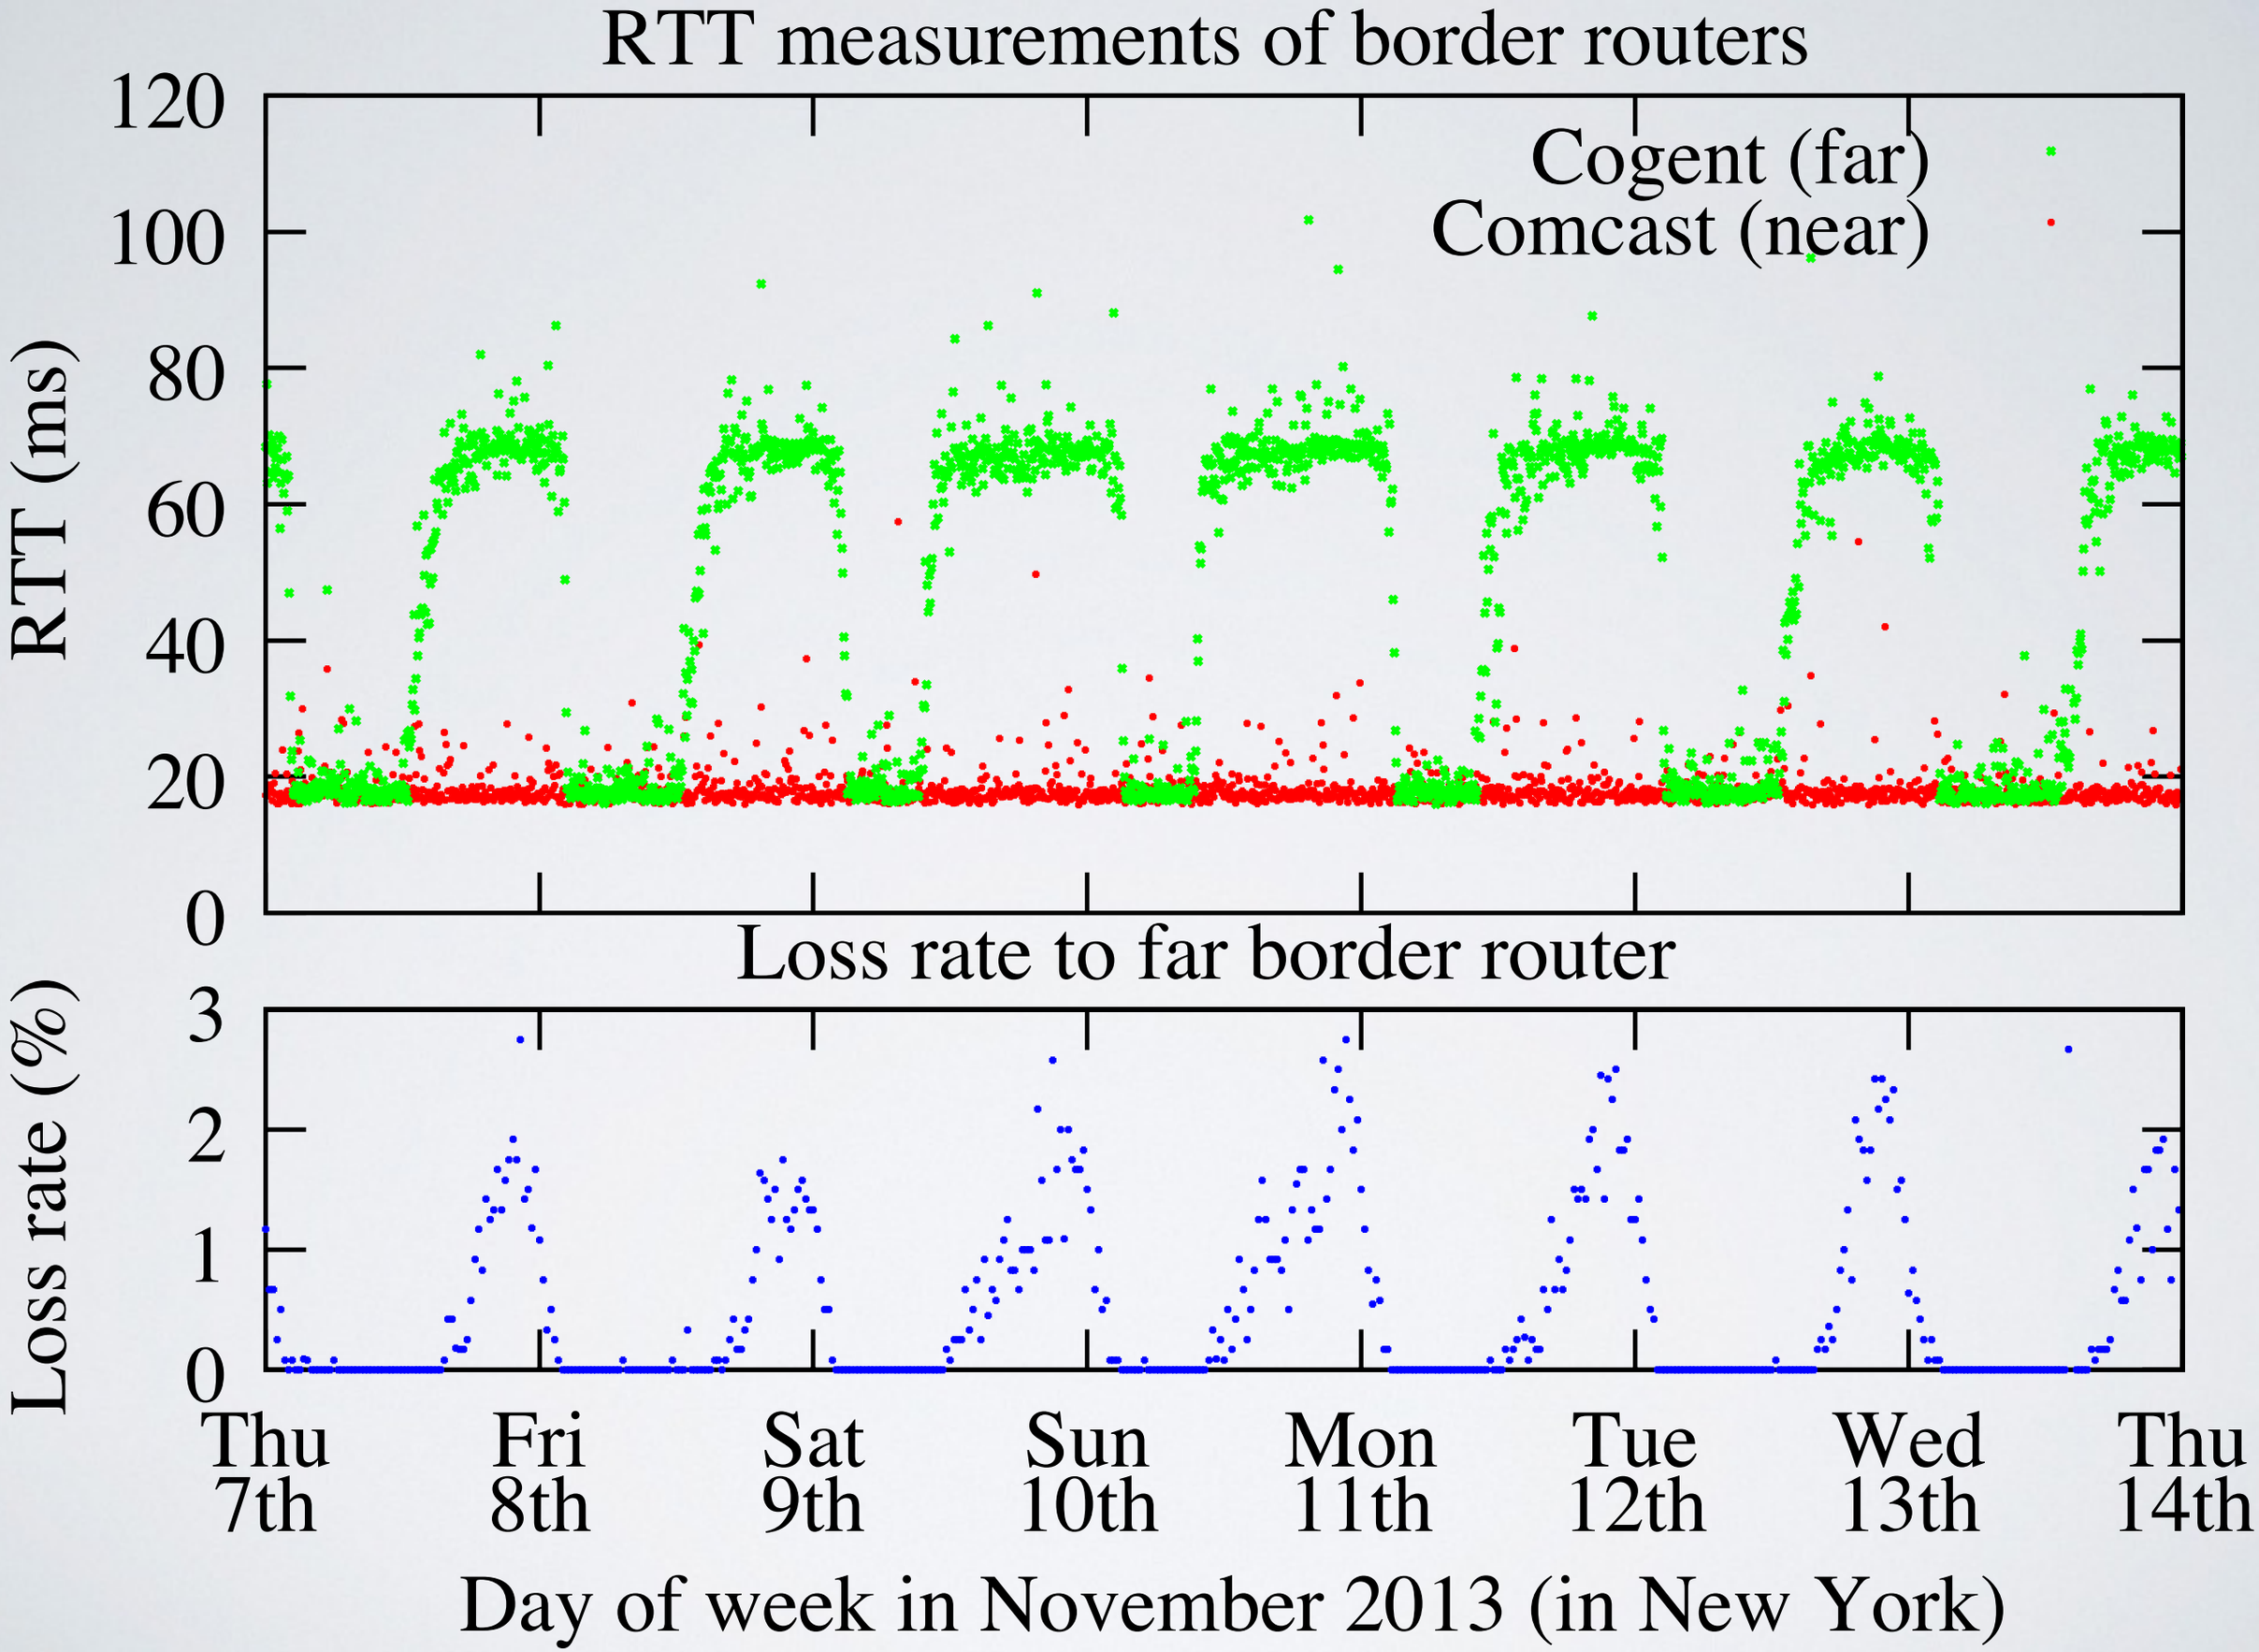
RTT measurements of border routers
RTT (ms)
120
100
80
60
40
20
0
Cogent (far)
Comcast (near)
Loss rate to far border router
Loss rate (%)
3
2
1
0
Thu 7th
Fri 8th
Sat 9th
Sun 10th
Mon 11th
Tue 12th
Wed 13th
Thu 14th
Day of week in November 2013 (in New York)

CHALLENGE: TOPOLOGY

- Mapping the set of interdomain links visible from a VP is a significant challenge:

  - Not trivial to identify which IP link represents the interdomain connection

  - Errors due to third party addresses in traceroutes

  - A single hop seen in neighbor network but not from neighbor's address space

CHALLENGE: TOPOLOGY

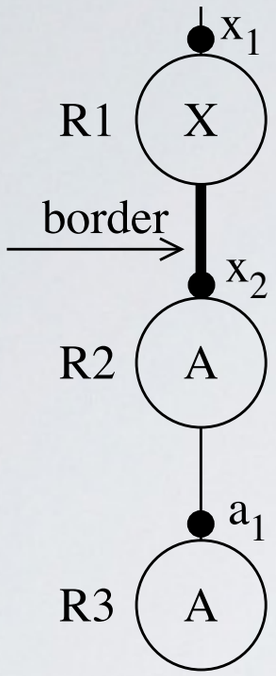

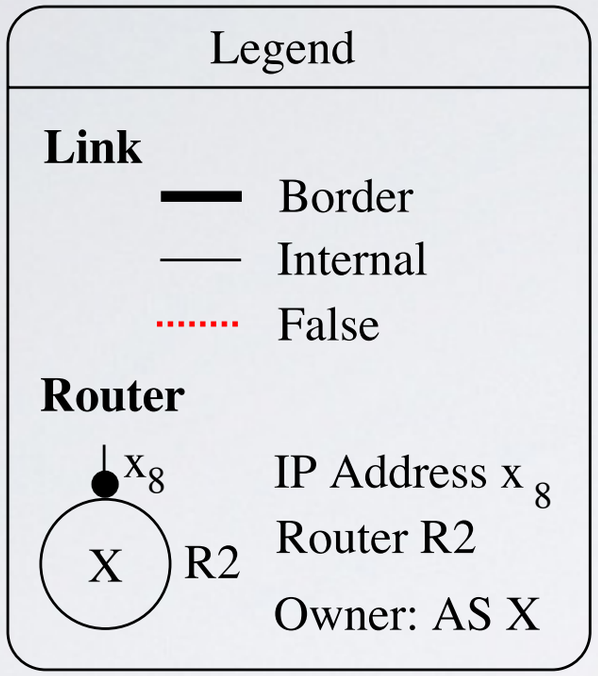

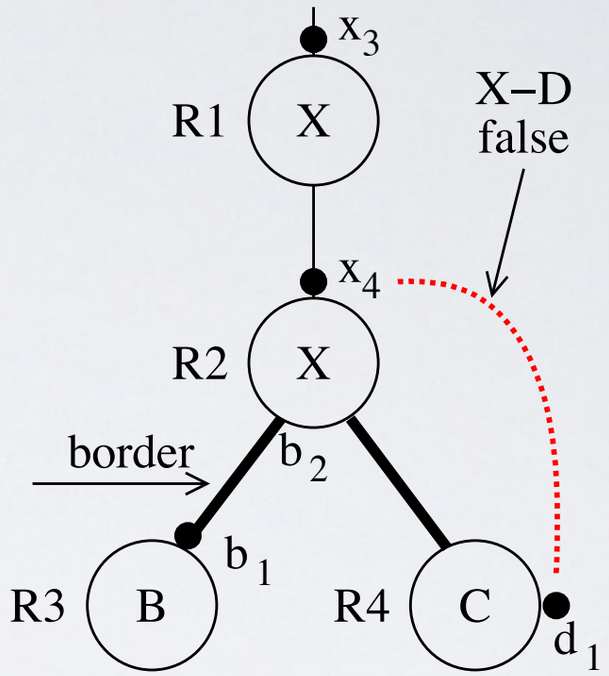

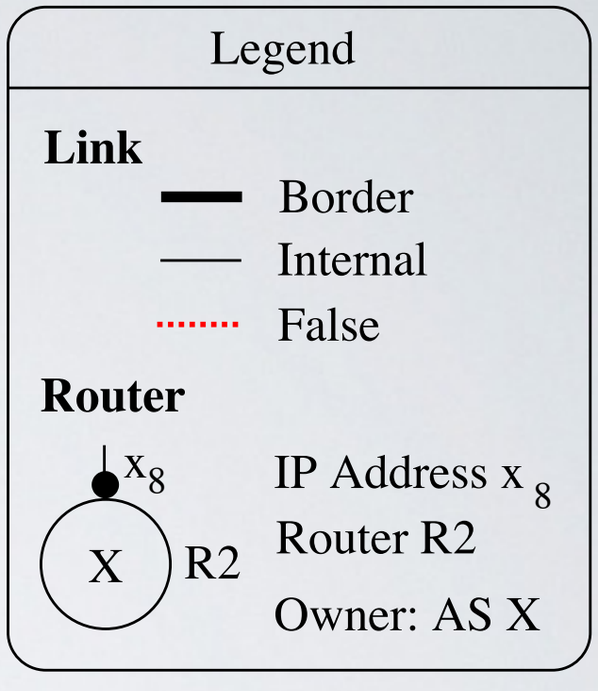

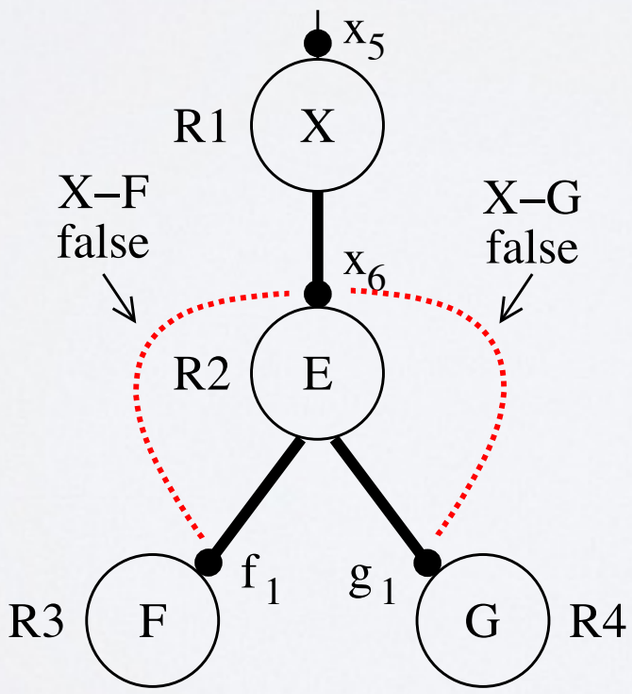

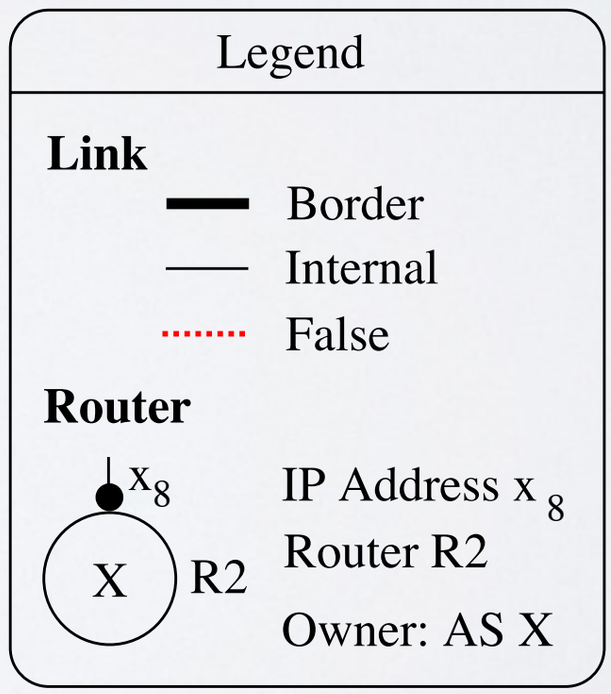

CHALLENGE: TOPOLOGY

- Current status: **Working on a set of heuristics to identify border links of a VP network using active probing from the VP, AS-relationships and alias resolution data**

CHALLENGE: VALIDATION

- We want to avoid incorrectly inferring a link is congested (or uncongested) given the intense current interest
- For links that show diurnal RTT pattern, how does pattern correlate with traffic data? But peering agreements contain NDA.
- Closest to public data: Level3's blog "Observations of an Internet Middleman"

Anonymous Dallas Link

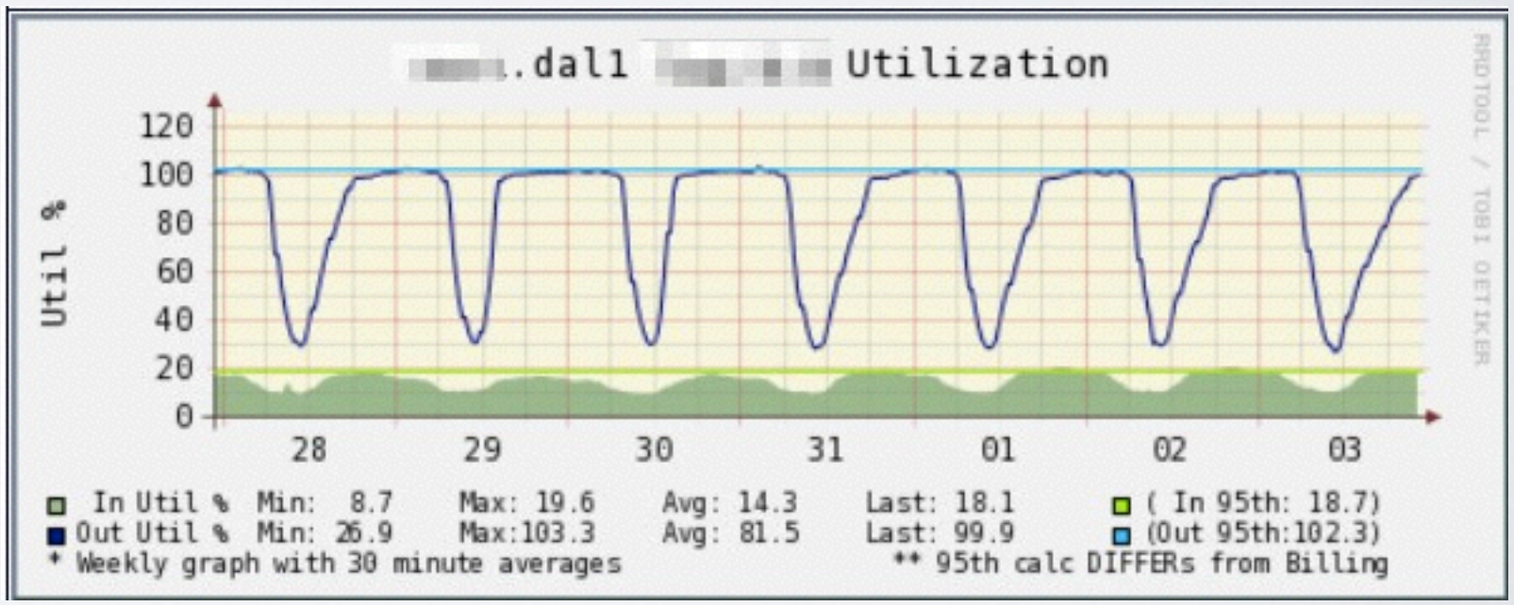

http://blog.level3.com/global-connectivity/observations-internet-middleman/

LEVEL3

(we happen to have a good view of Level3-Dallas)

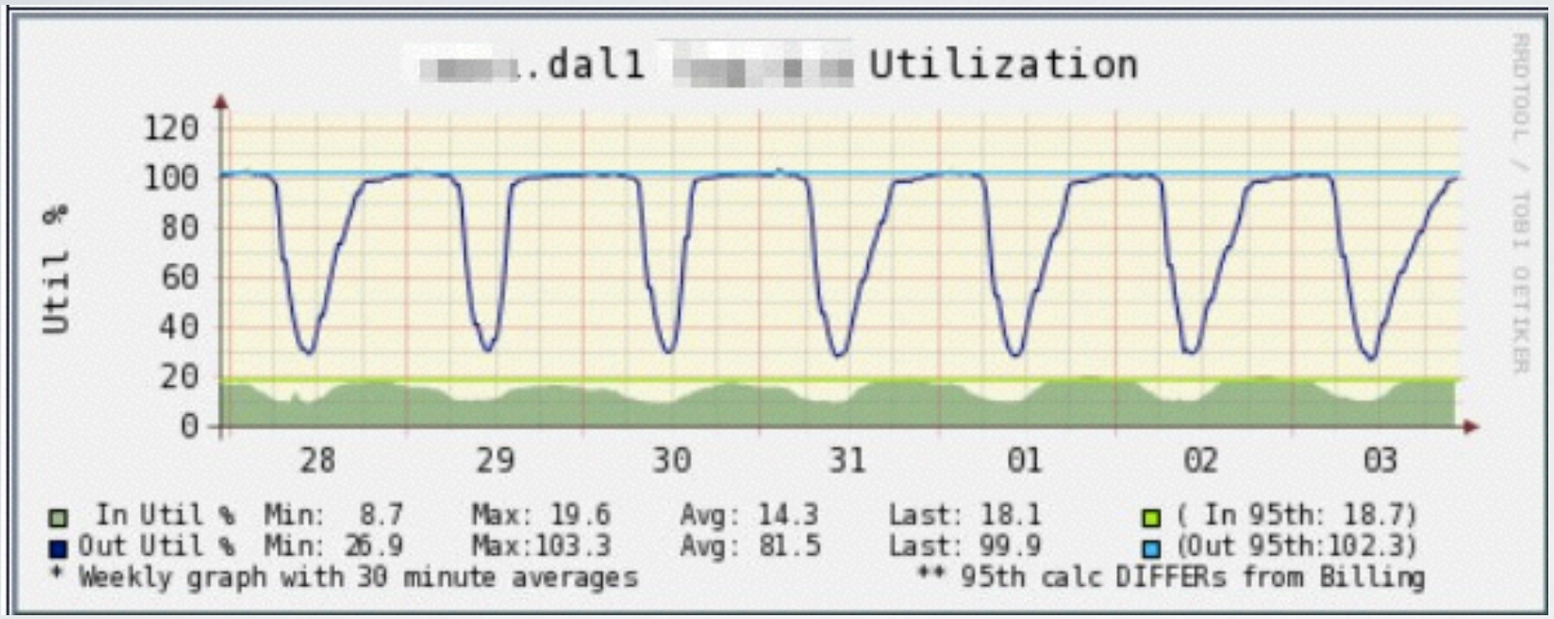

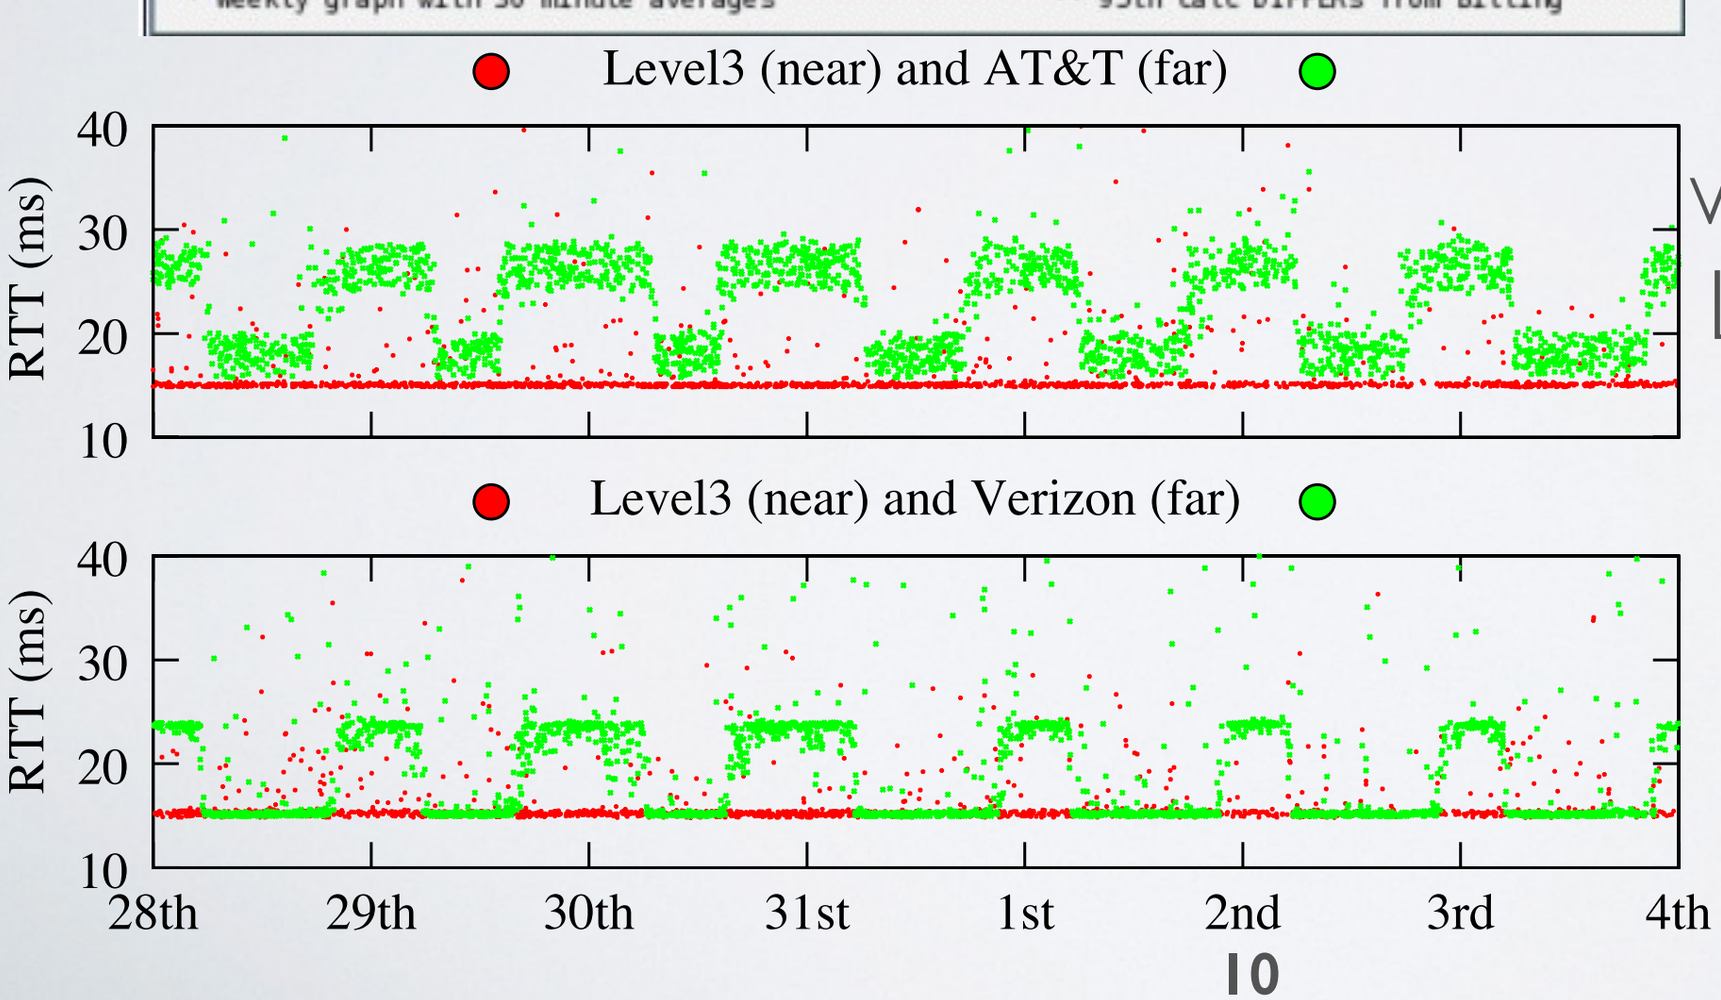

"Ground Truth"

We believe both AT&T and Verizon were congested with Level3 in April 2014

Anonymous link is probably AT&T (based on duration of level shift)

CHALLENGES: SYSTEMS AND DATA

- Probing needs to be responsive to change in the network: routing changes, new peering links, connectivity at IXPs

- Need scalable techniques to manage and process data from thousands of interdomain links and automatically detect congestion patterns in (near) real time

- Current status: **Building backend system for adaptive probing, data management, triggering on-demand measurements and data visualization. Investigating FFT and other analysis techniques for automatically detecting congestion patterns**

VP DEPLOYMENTS

- Deployments in various access networks (and other network types, see http://www.caida.org/projects/ark/

- Currently 19 monitors running TSP measurements

- We continue to deploy Ark nodes using Raspberry Pi hardware in homes of our friends (or friends of friends)

- **Future goal: deploy our experiments on other platforms: Bismark, FCC-Samknows**

MEASUREMENT SYSTEM

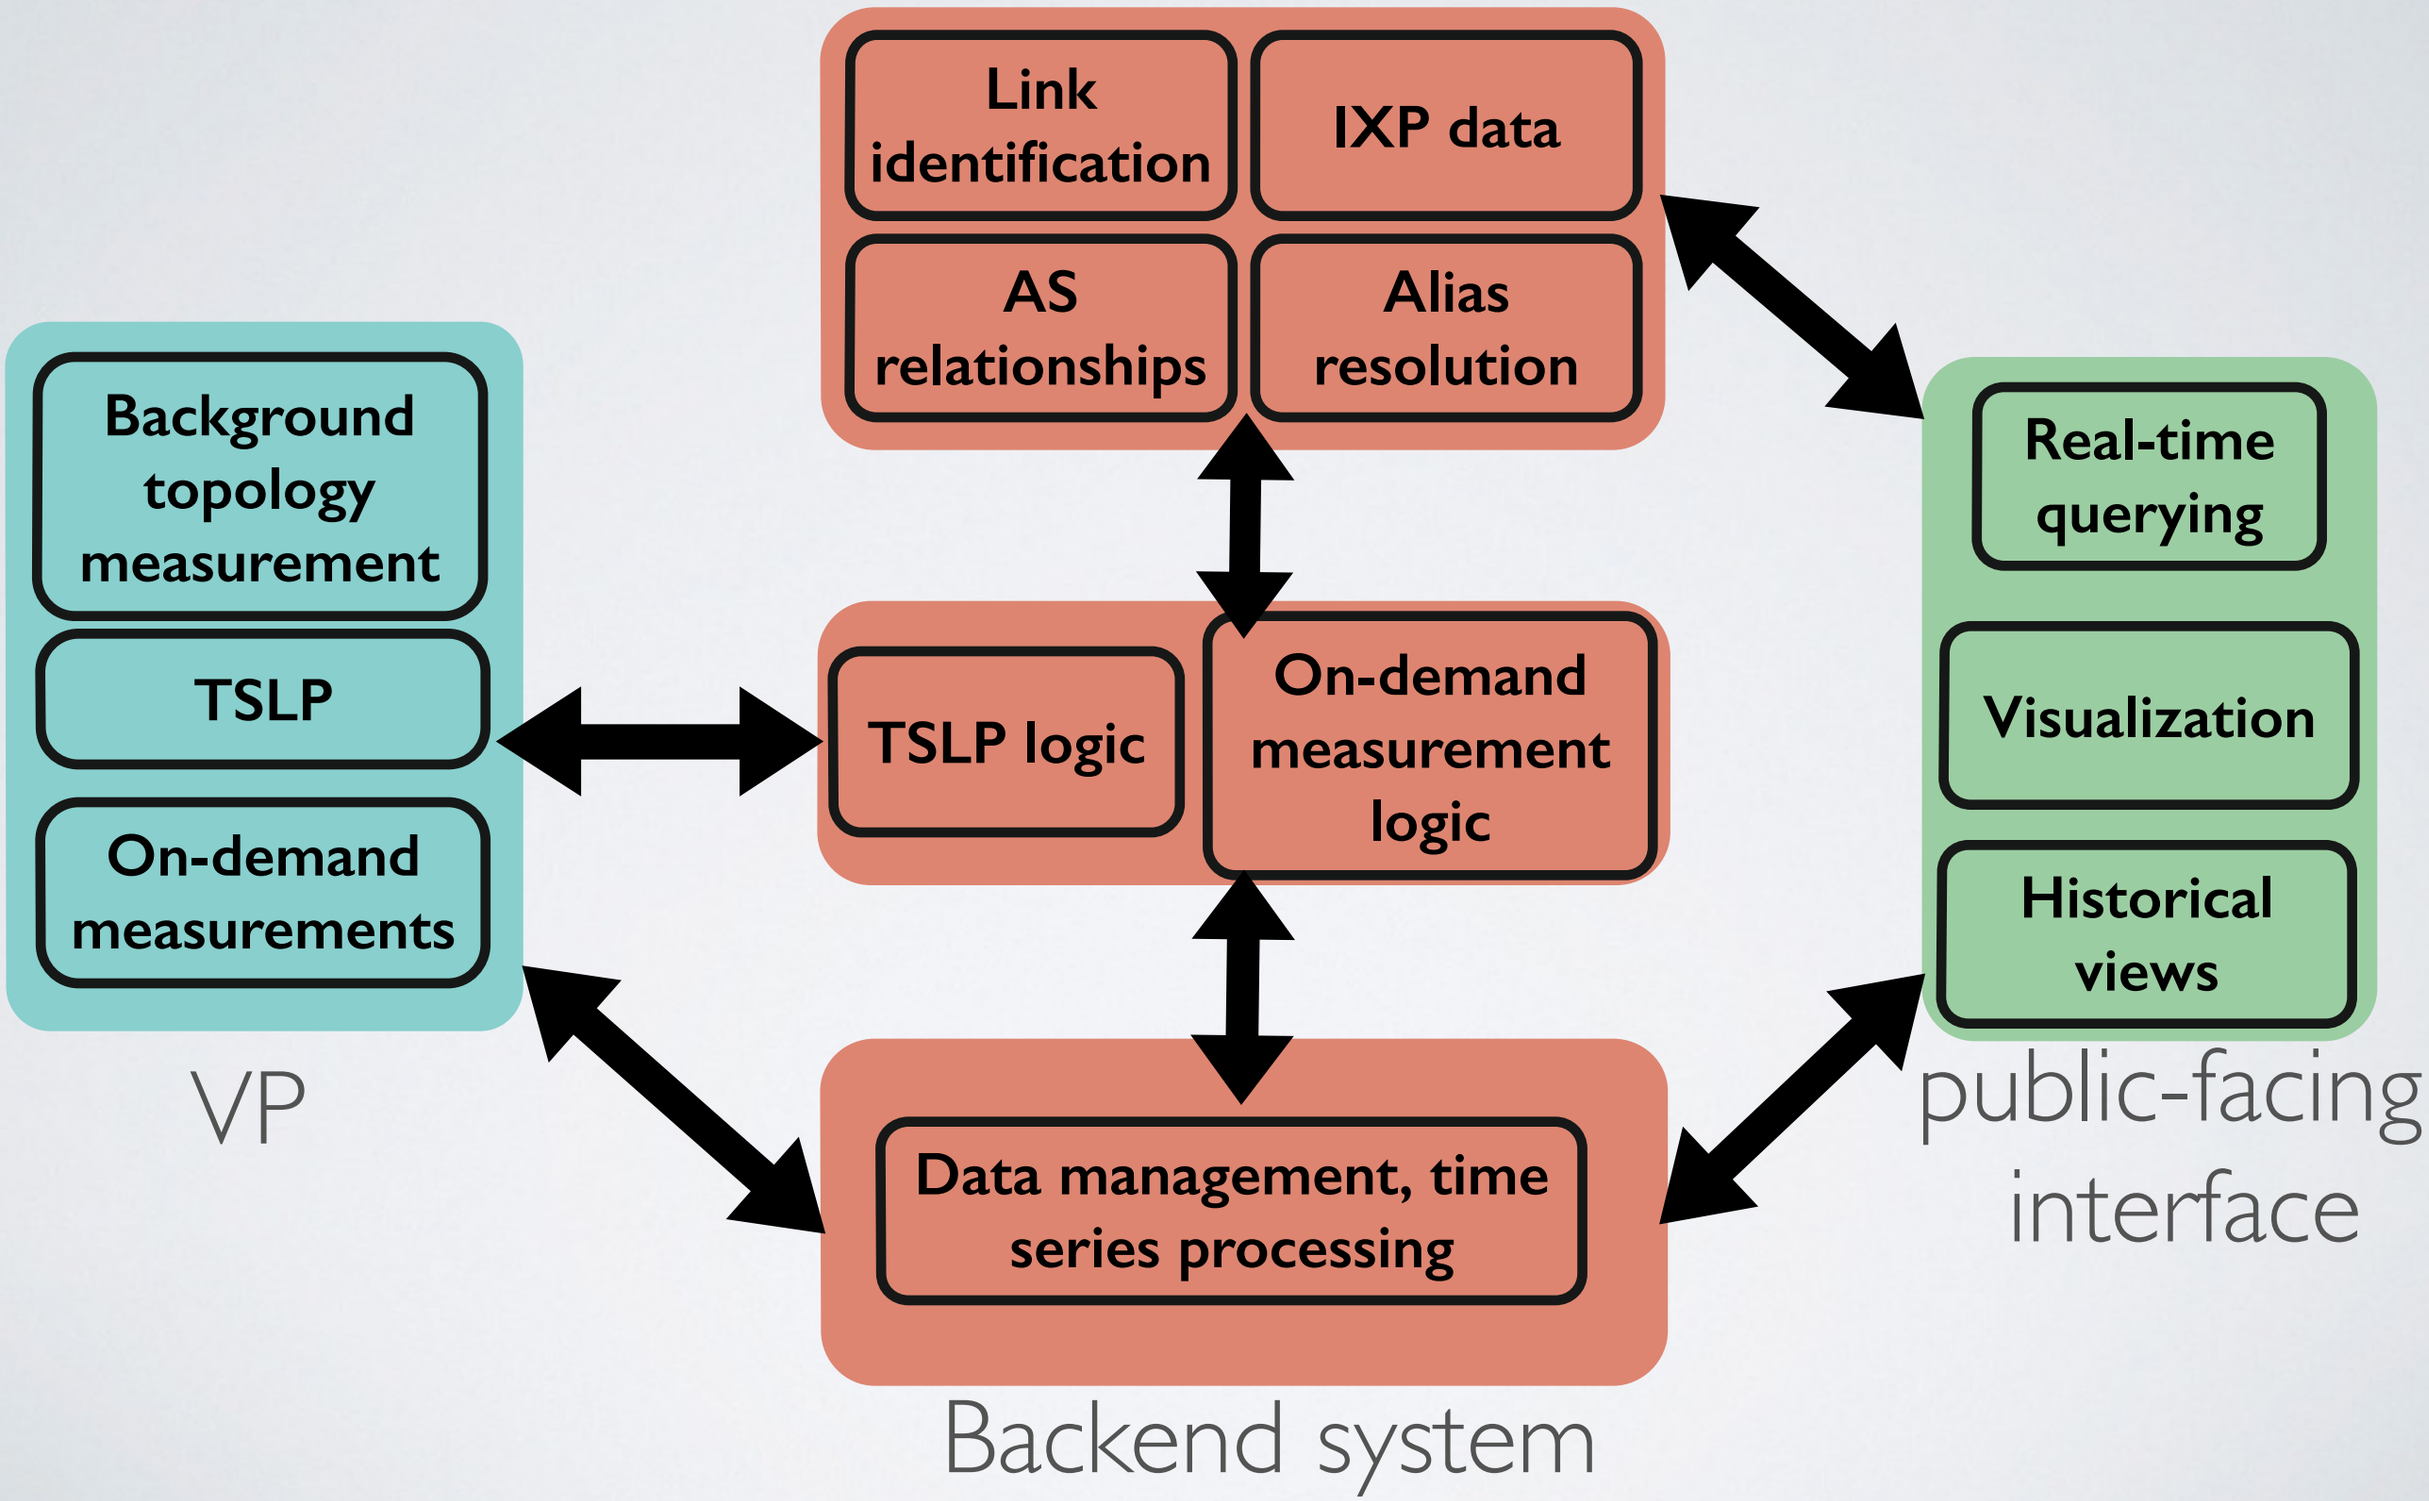

SUMMARY

- **Our goals (1) atlas of interdomain links and their congestion state, (2) improve transparency, empirical grounding of debate**
- Demonstrated a lightweight and easily deployed method to view link congestion patterns
- Currently building the topology+congestion measurement system

Email: amogh@caida.org

CHALLENGE: LINK IDENTIFICATION

- Probe (.1) address in each prefix with ICMP-paris traceroute
- Infer interdomain link when we observe an AS change with ip2as mapping

But which IP-link represents interdomain link?

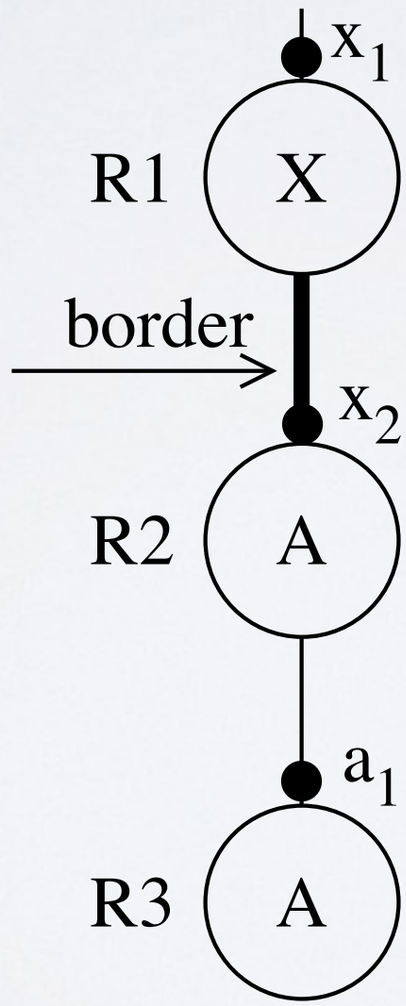

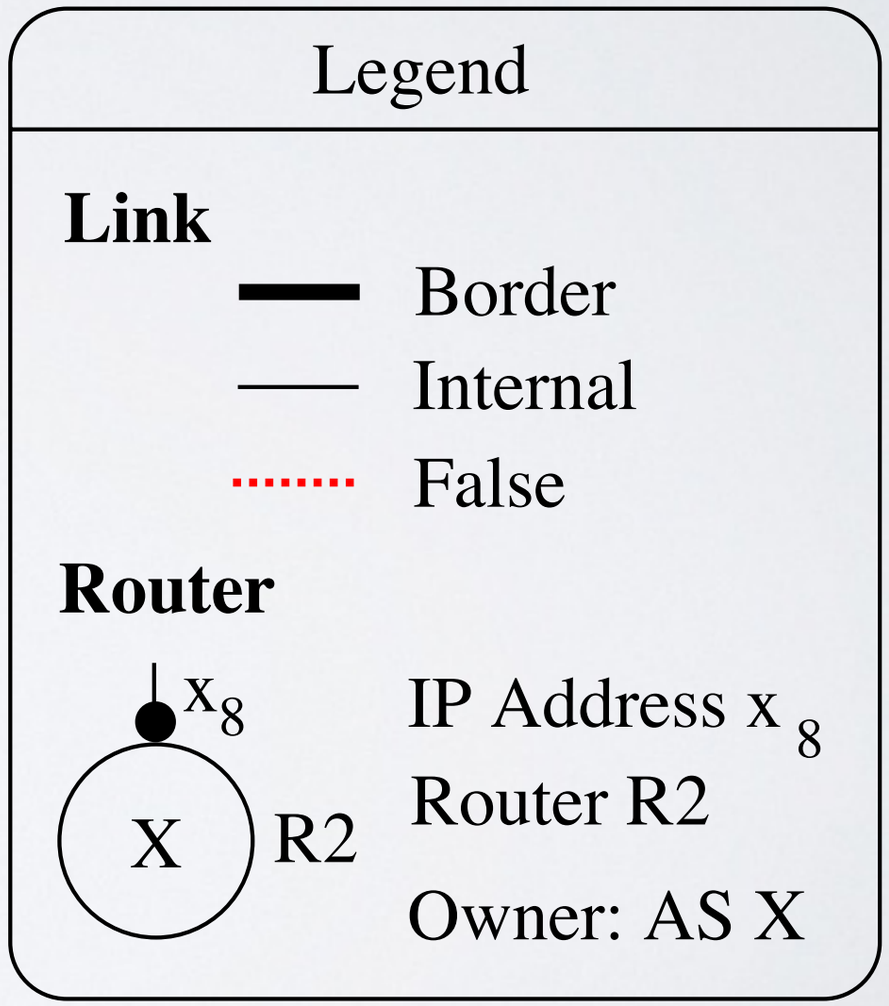

THIRD PARTY ADDRESSES

- Which ip2as mappings represent false mappings?

- Problematic when d_1 is the only address in AS D that responds (i.e. with ICMP echo response)

- Use scamper prefixscan method to prove b_1 in-bound address on pt2pt link

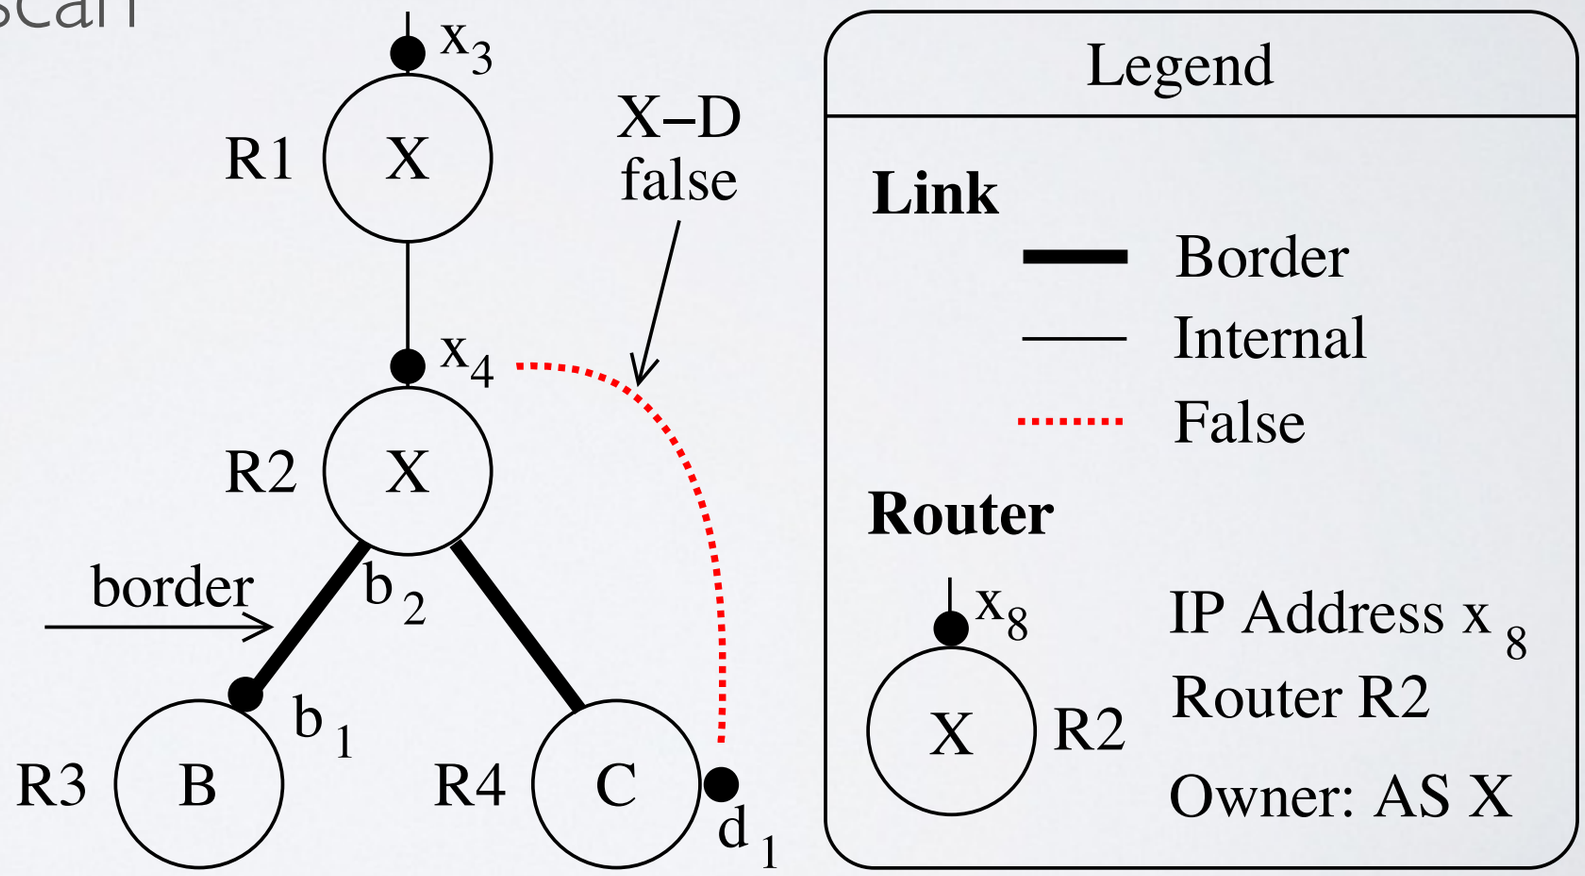

ONE-HOP IN NEIGHBOR

- One hop on router owned by neighbor E, but assigned from X's address space. E's customers directly attached to border router.
- Use CAIDA as-relationship inferences to infer common provider E for F, G.

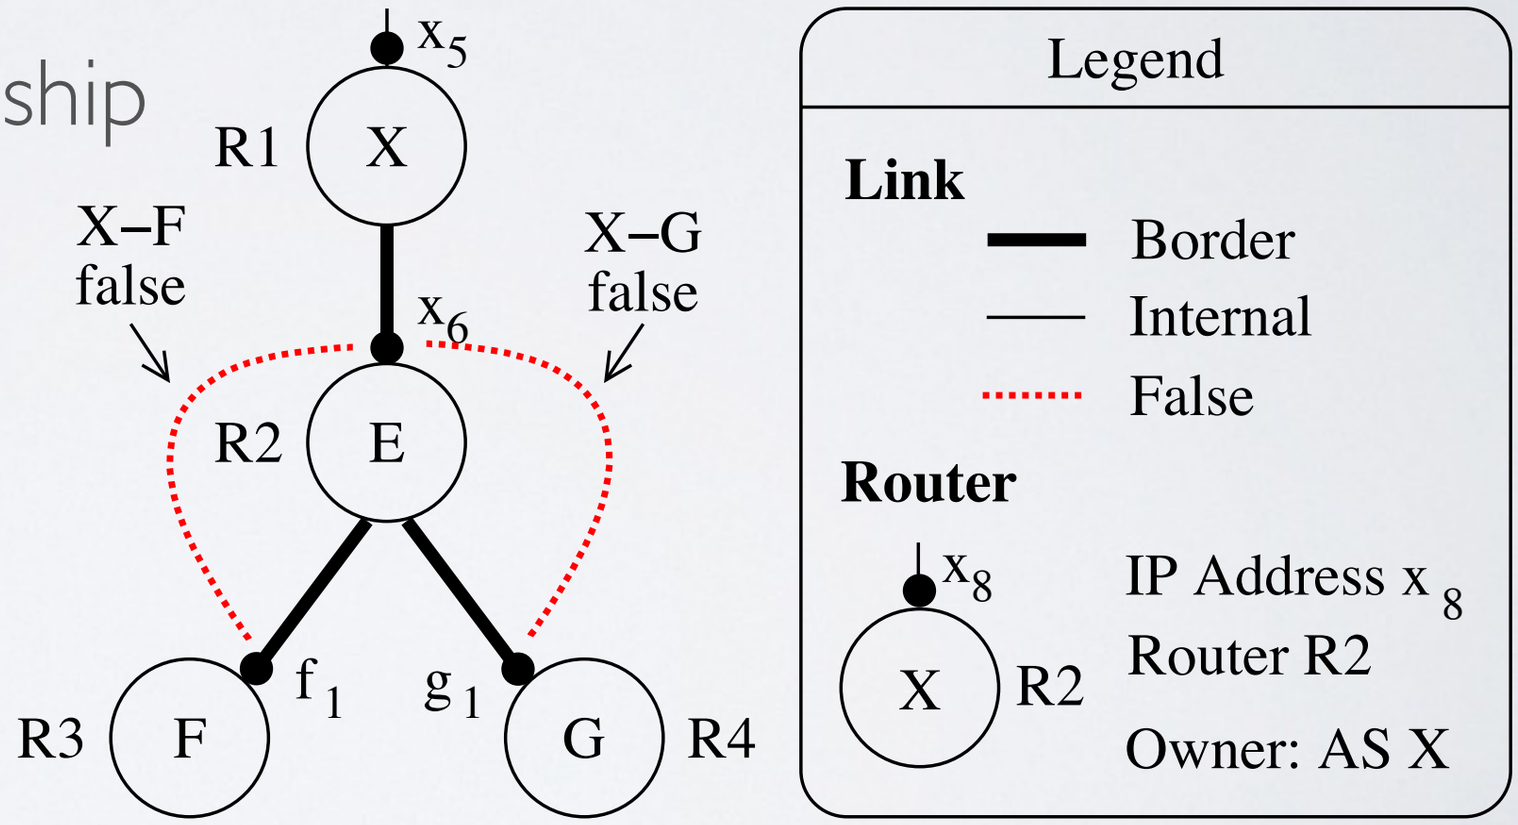

CHALLENGE: REVERSE PATH

- Difficult to know that the response from far router returns over targeted link

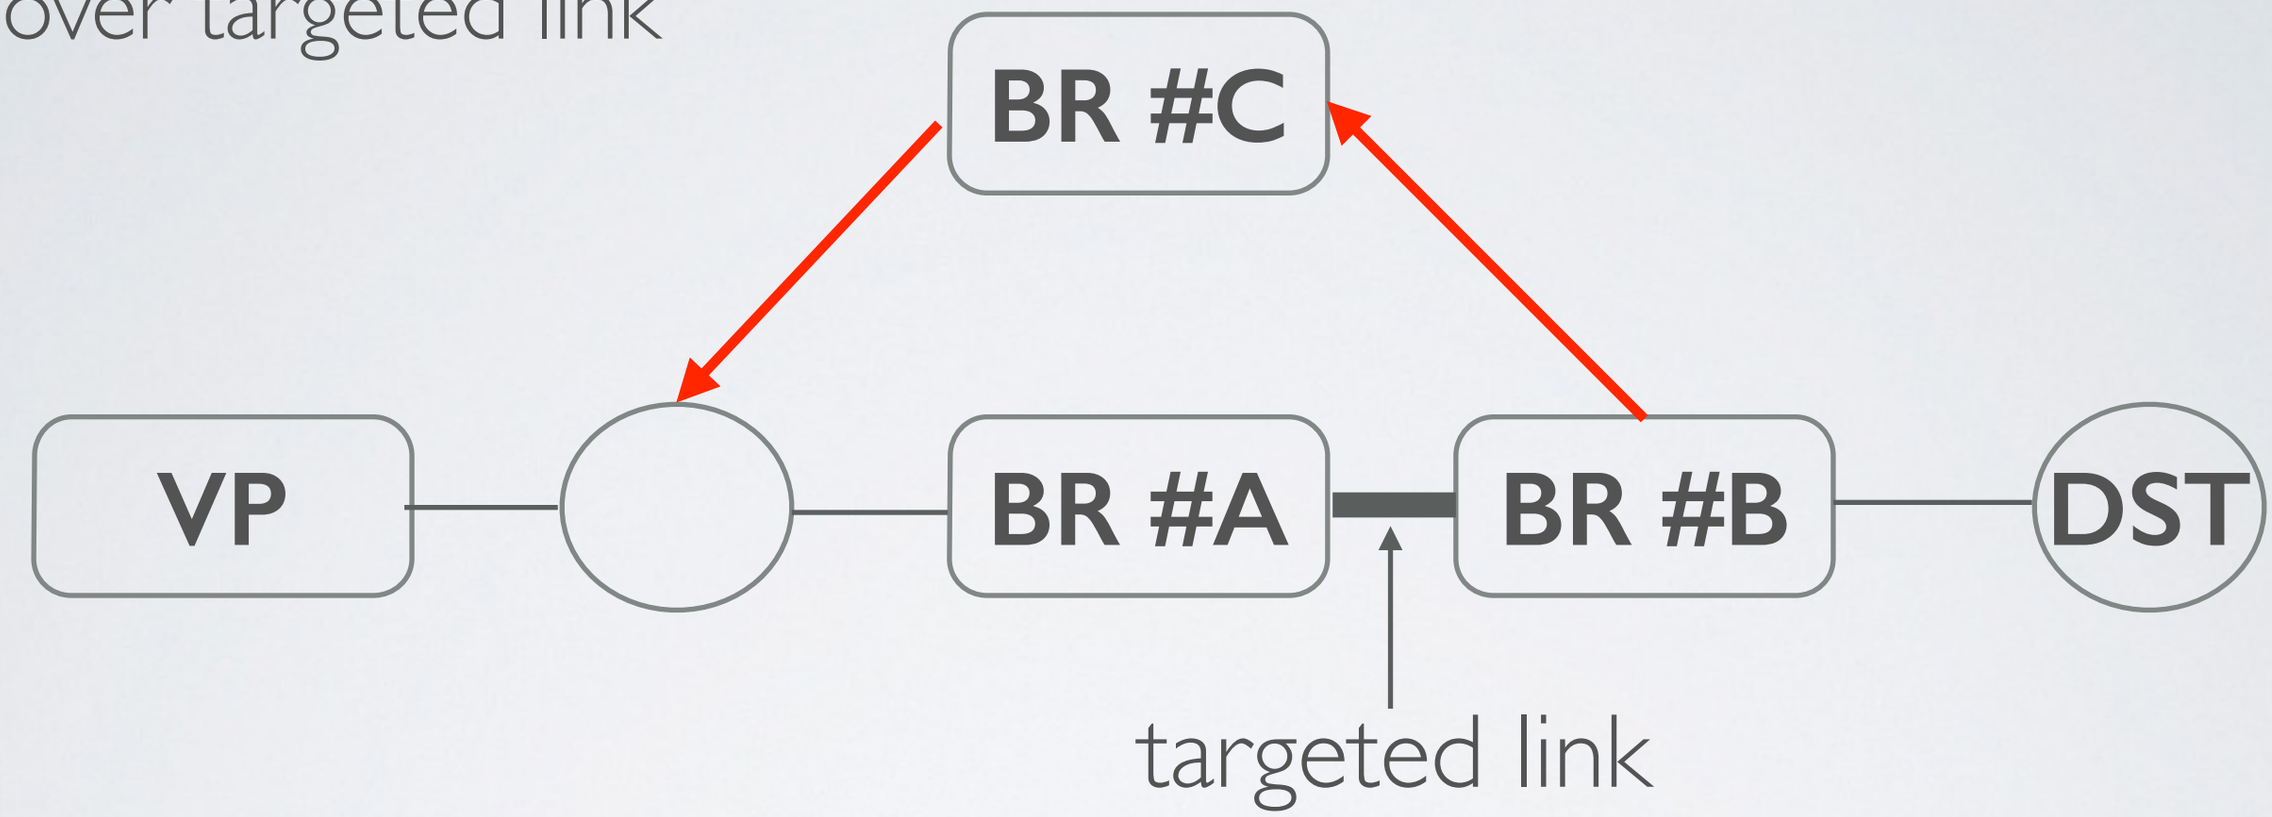

Methods that support inference:
Reverse path traceroute, IP record route,
IP timestamp option, tomography

CHALLENGE: REVERSE PATH

- For a single monitor inside Comcast, can show 30% of return paths traverse the targeted link with record route, or IP timestamp option
  - mostly limited by options support of neighbor routers, or distance of link from testing node
- Can improve with denser deployment of testing nodes

CHALLENGE: PARALLEL LINKS

- Some interdomain connections consist of many parallel links

IP-level links seen: A-B1, A-B2, A-B3, A-Bn

- We are aware of ***link striping*** caused by long lived flows; we hypothesize all parallel links will show same level shift pattern under load.

CHALLENGE: ADAPTING TO CHANGE

- Network configurations and routing change over time
- Need to know the current distance of a link from our VP
- Need to catch routing and peering changes as they happen
- Approach: Background topology discovery process on each VP

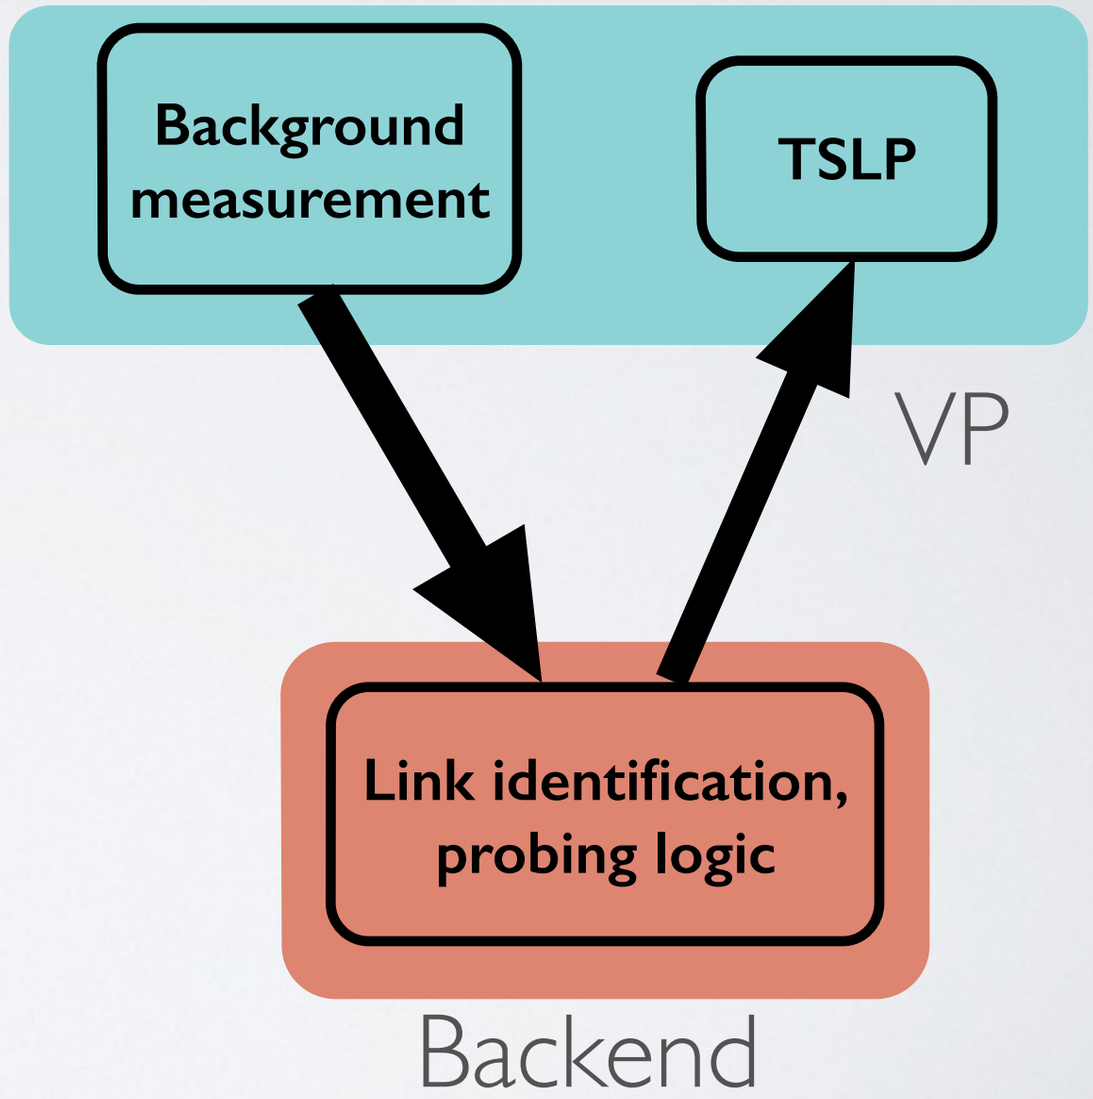

CHALLENGE: IDENTIFYING DIURNAL TRENDS

- We measure thousands of interdomain links
- We need scalable ways of looking through all these time series to find congestion patterns
- Approach: FFT to take advantage of diurnal congestion patterns

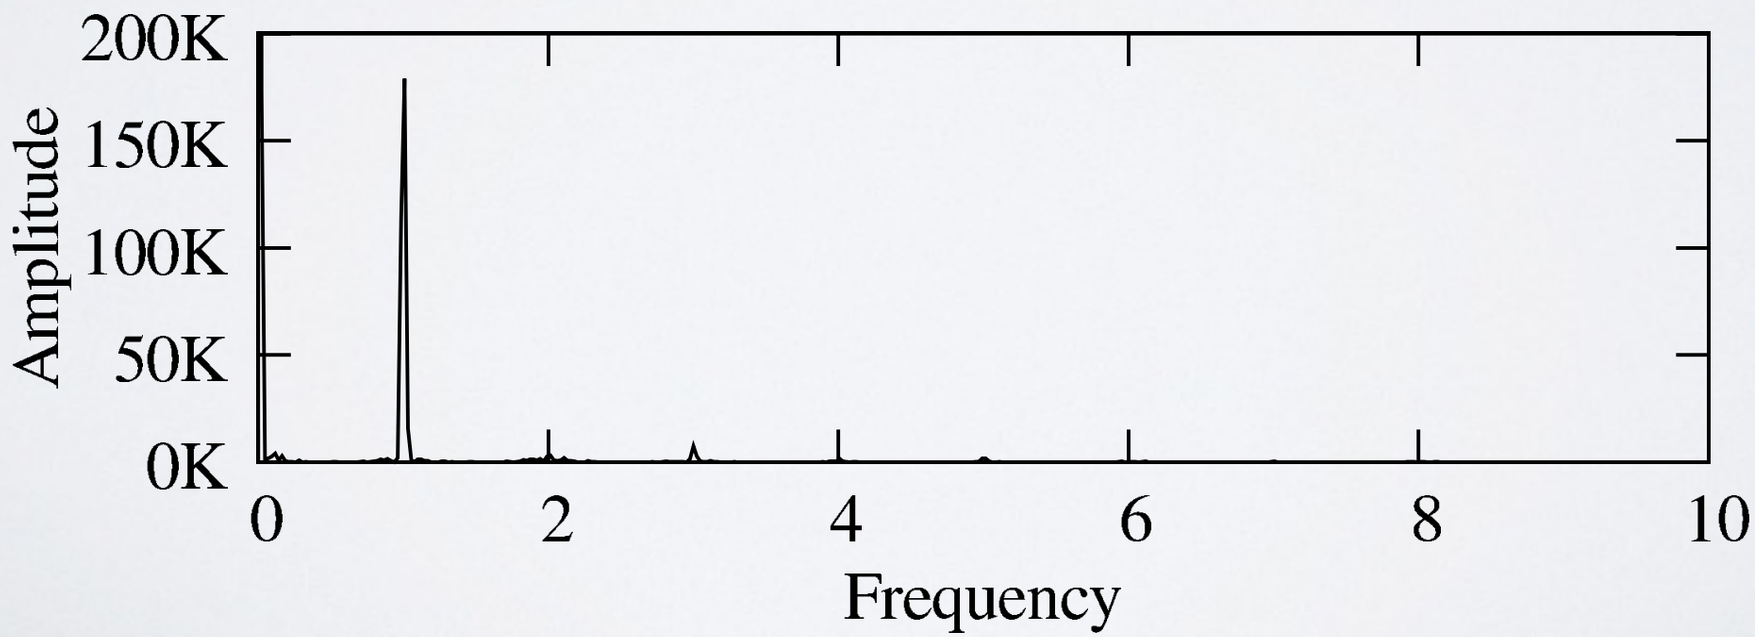

NETFLIX TO COMCAST

(the prior configuration)

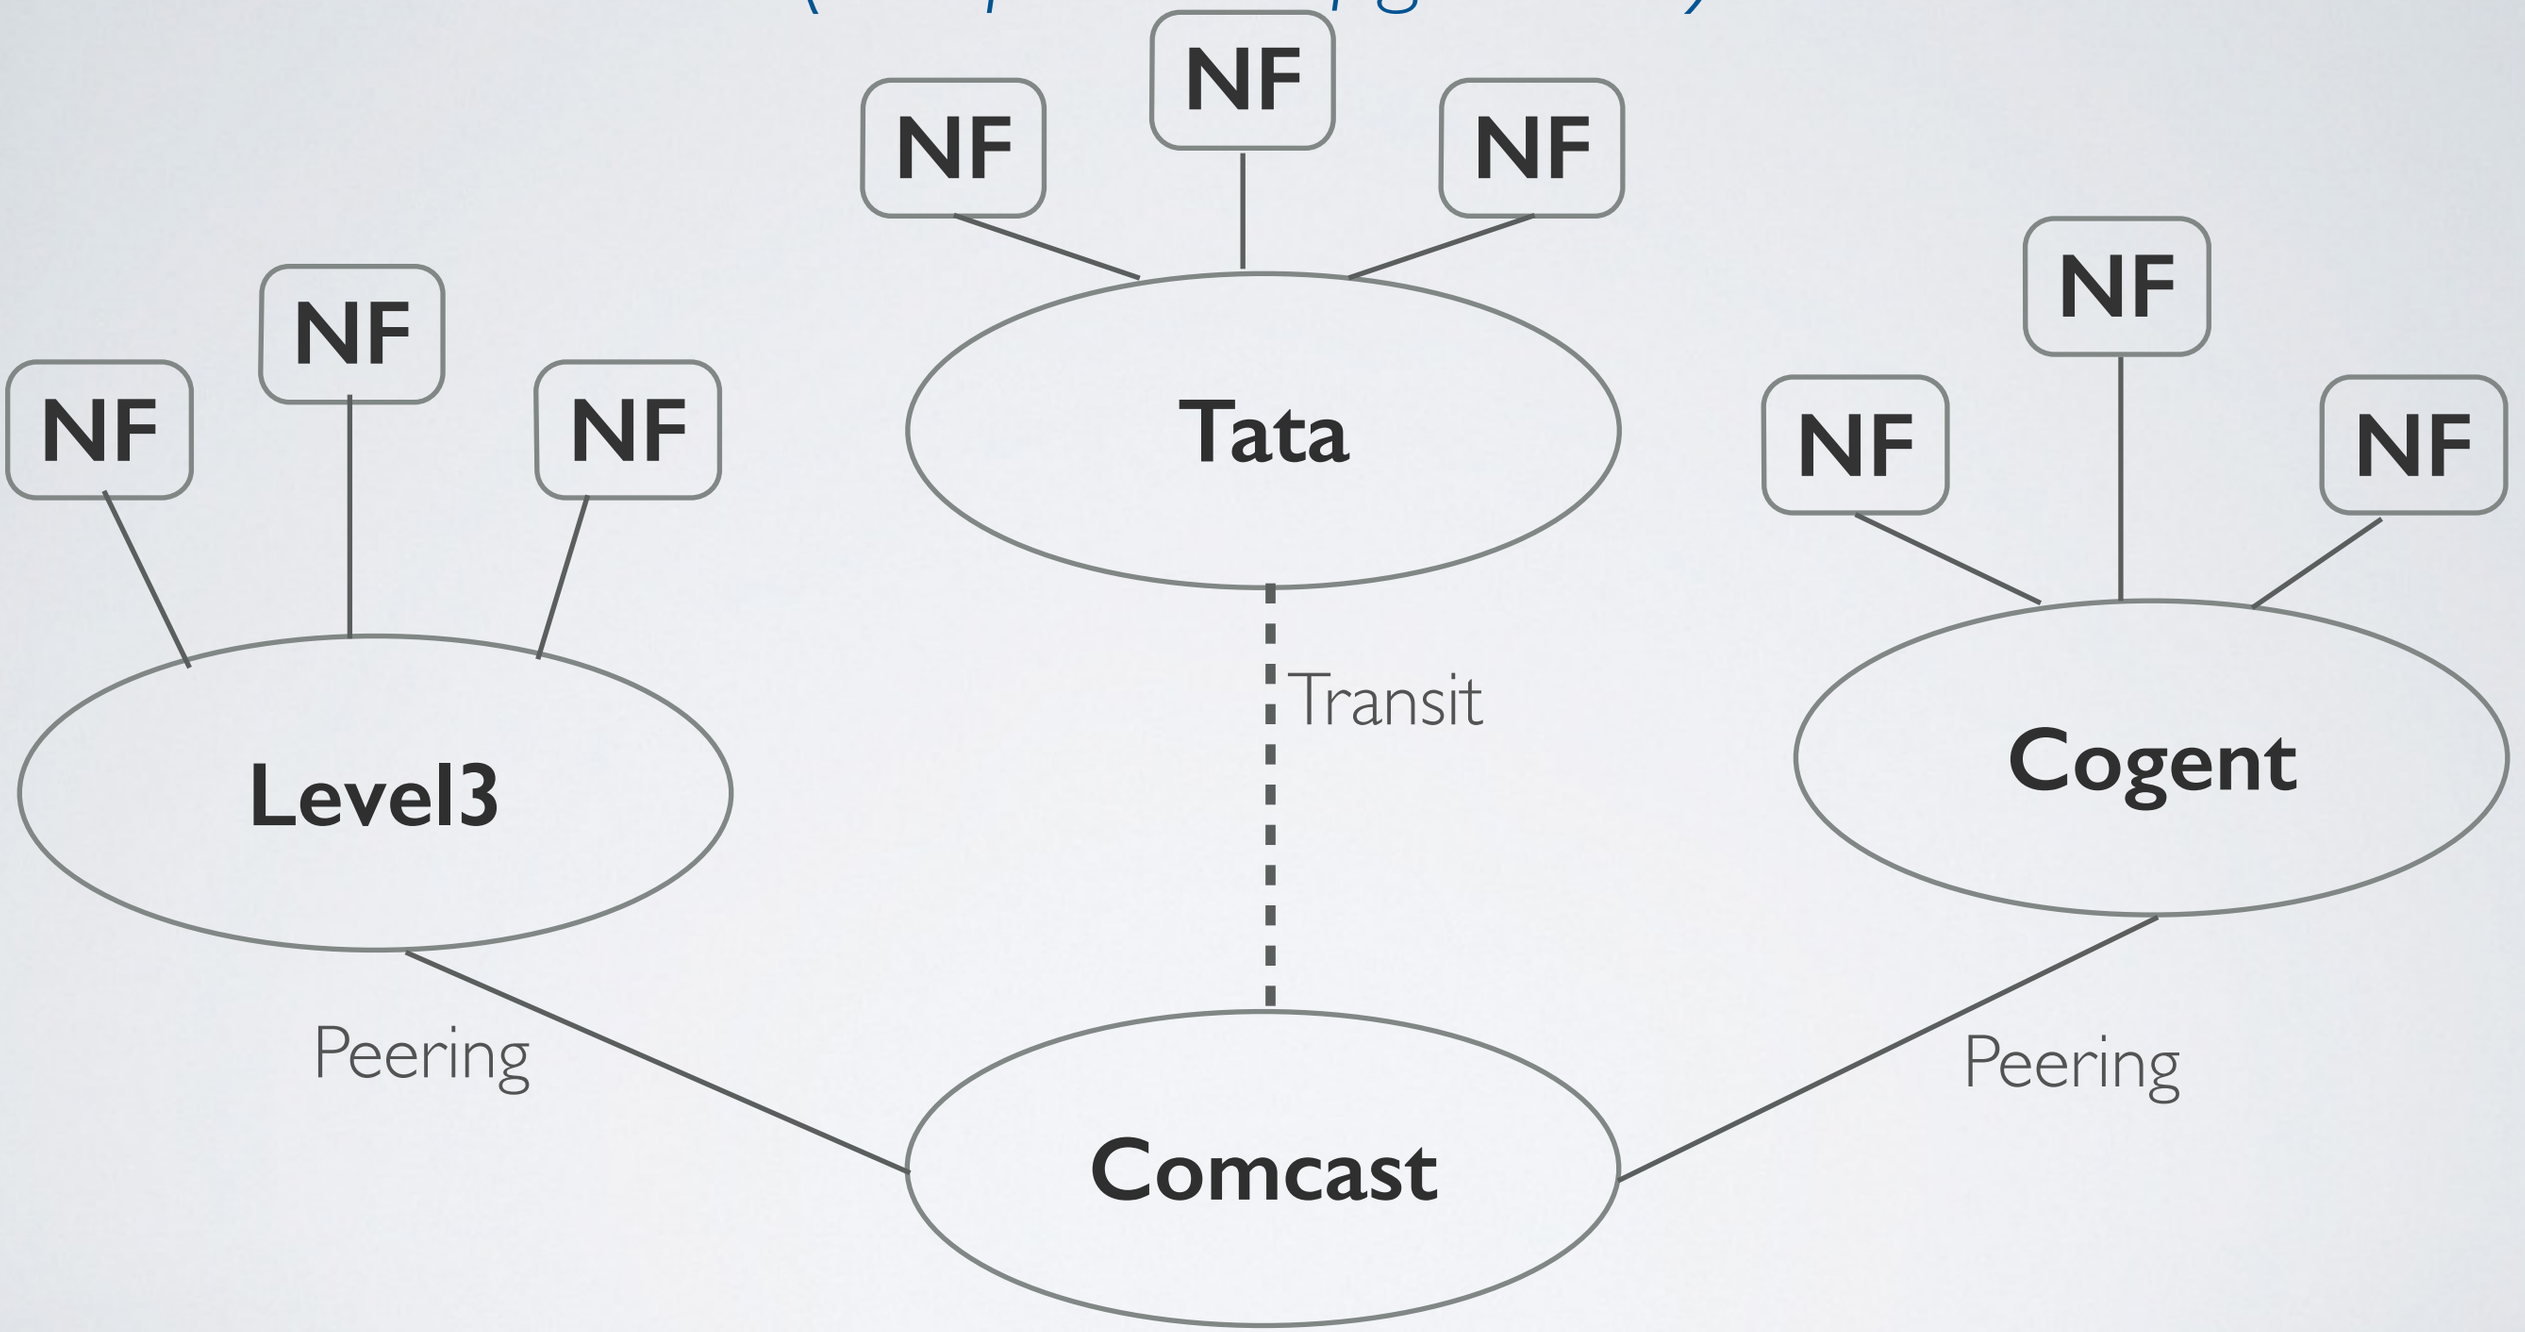

COGENT-COMCAST

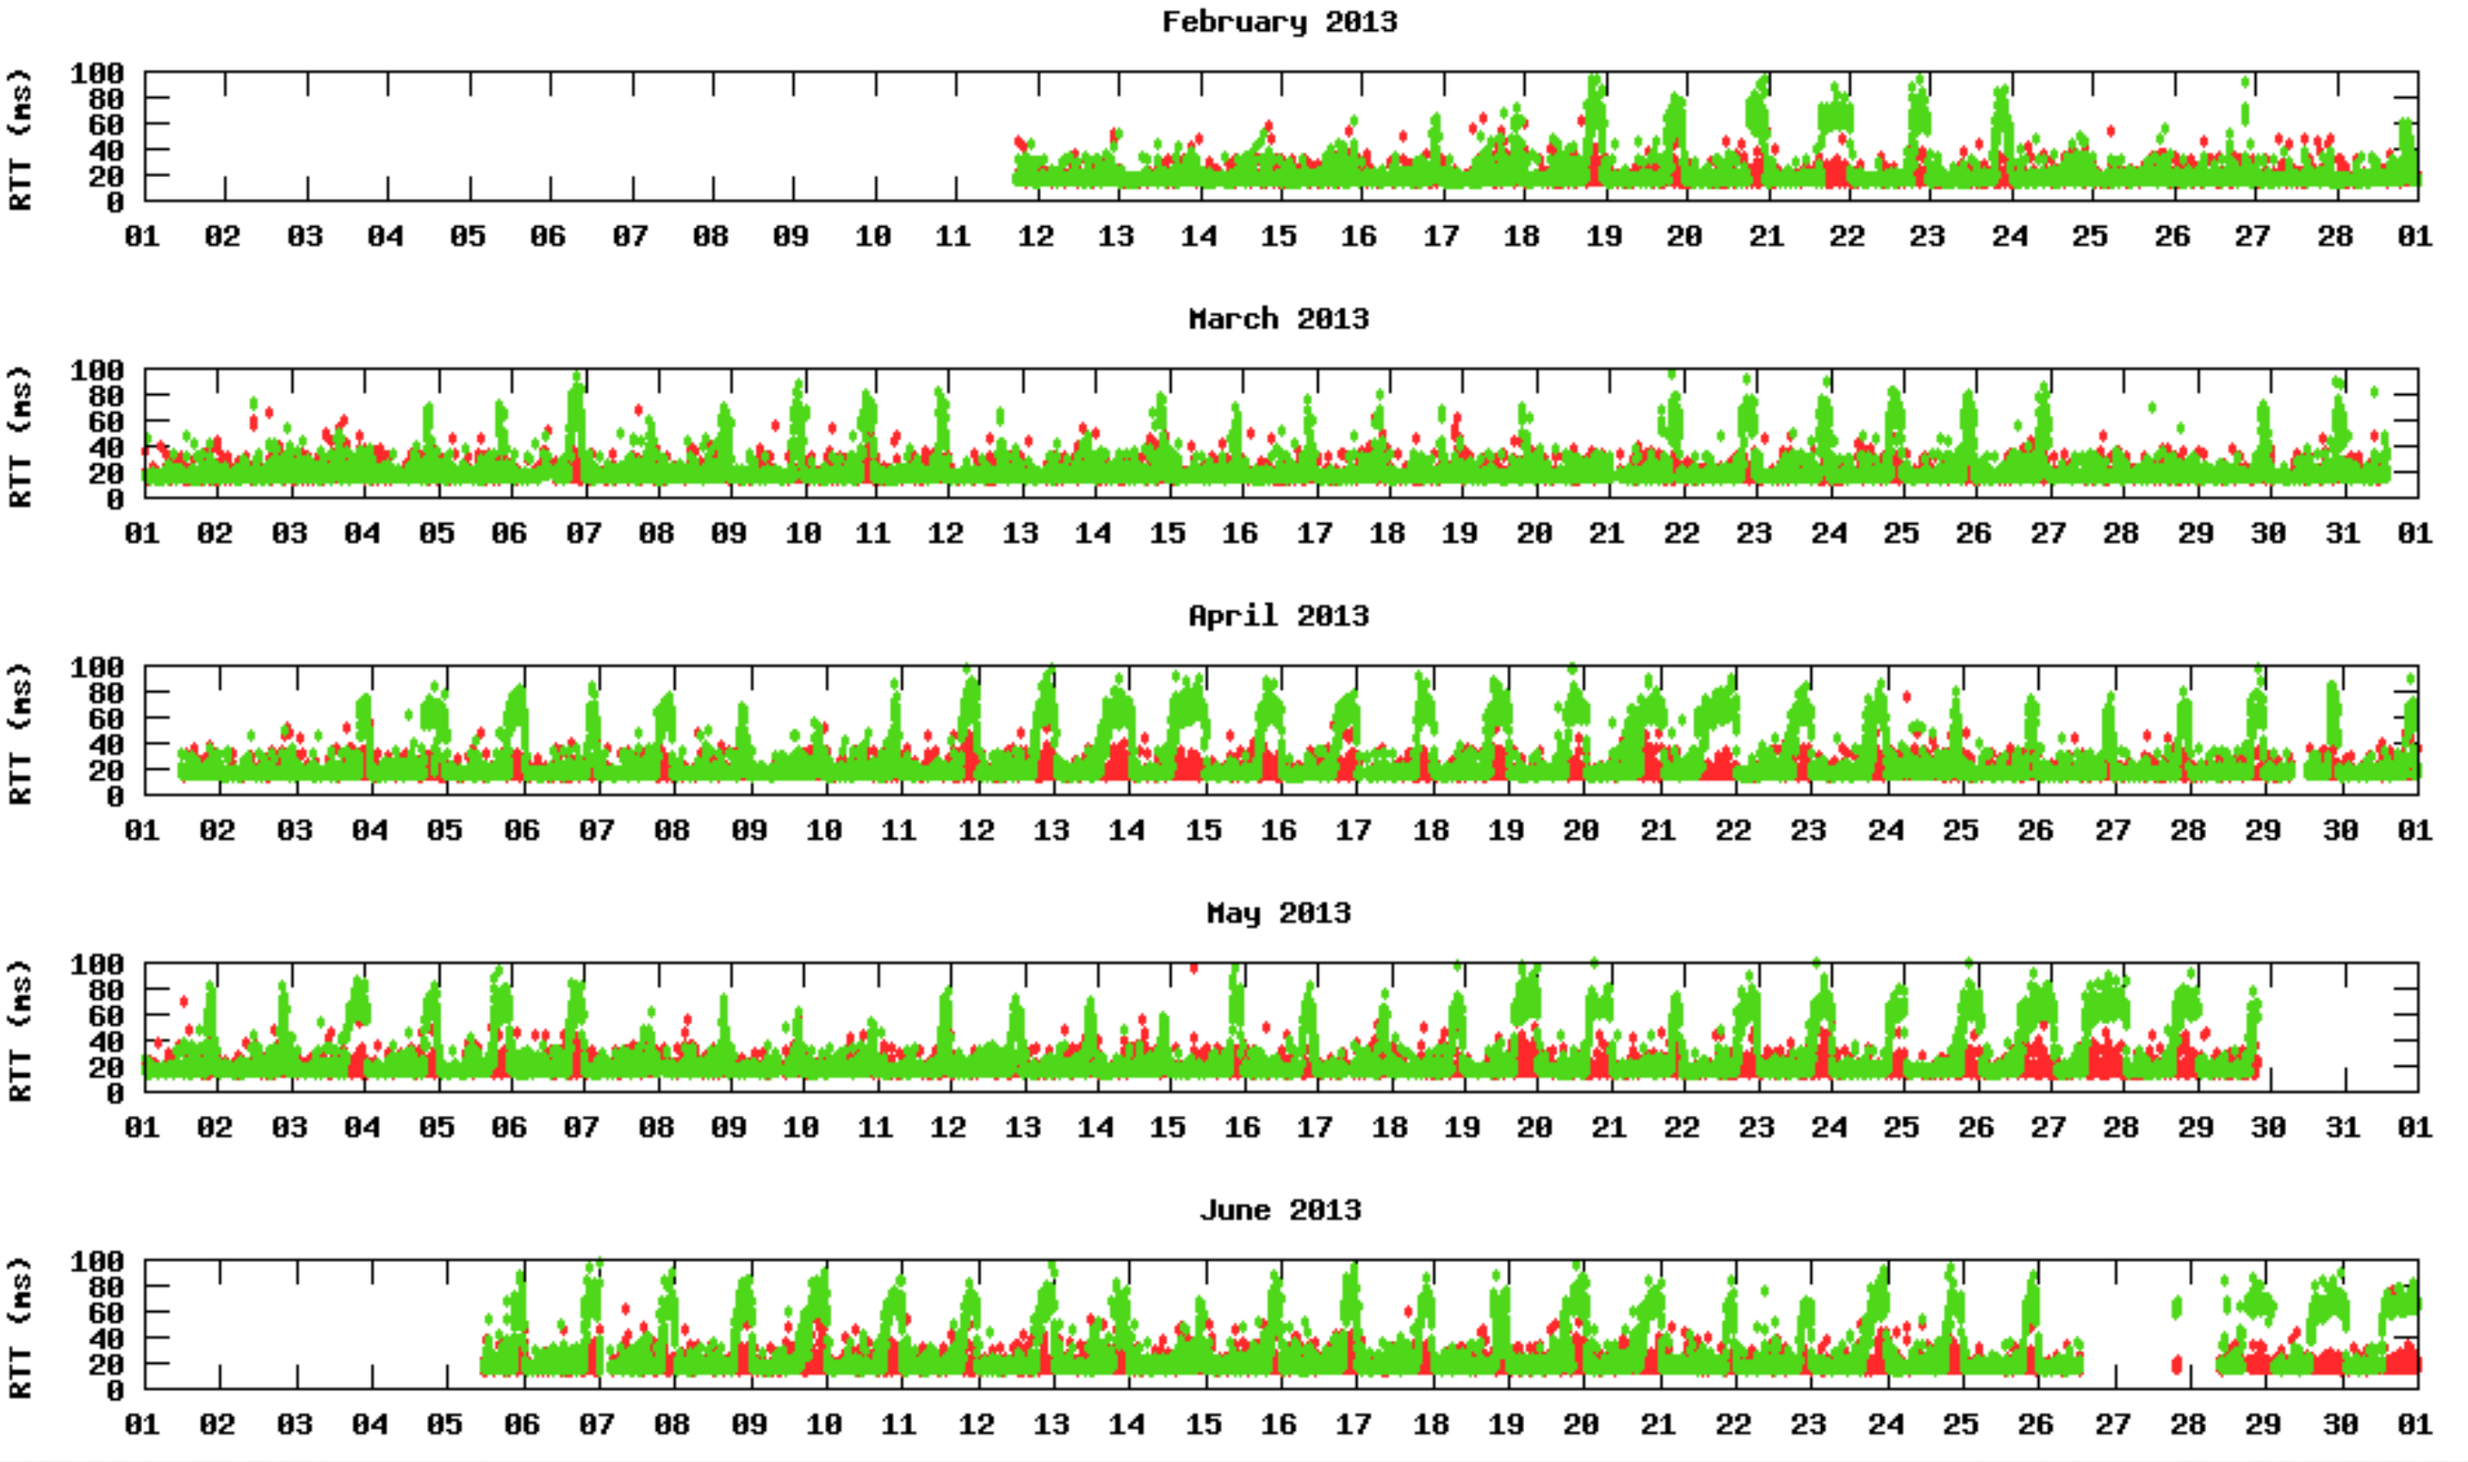

COGENT-COMCAST

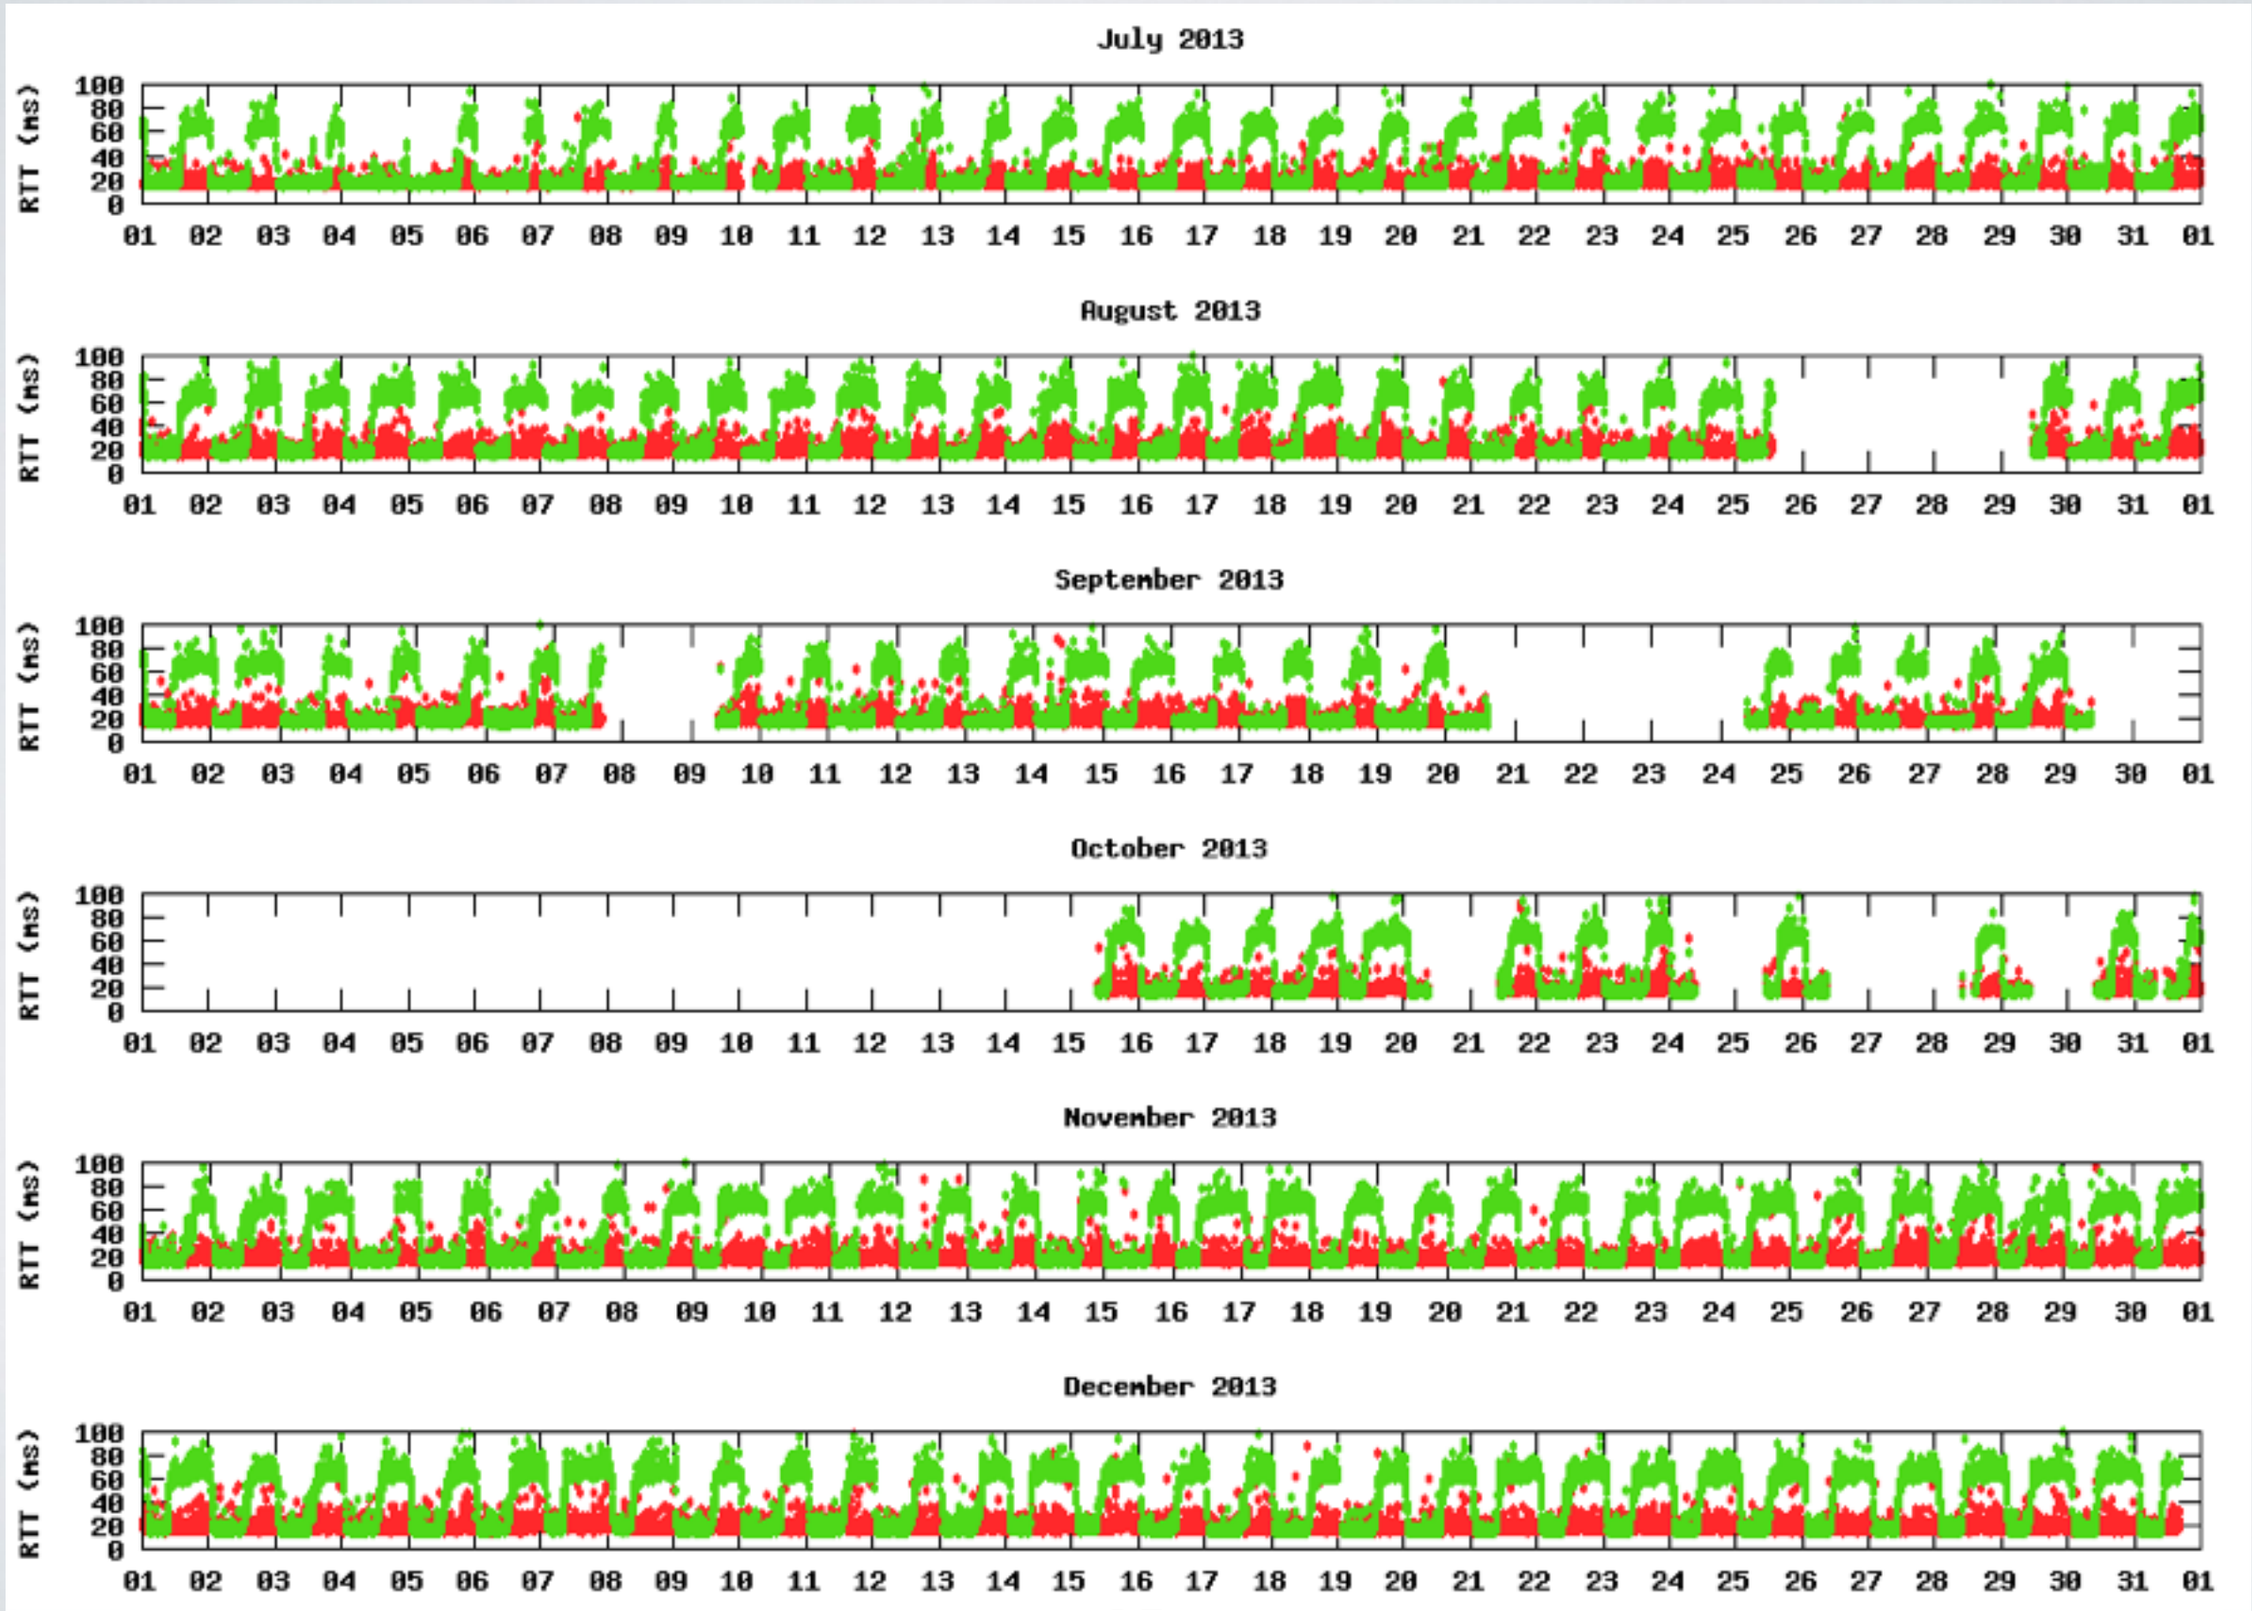

TATA, LEVEL3, COGENT - COMCAST

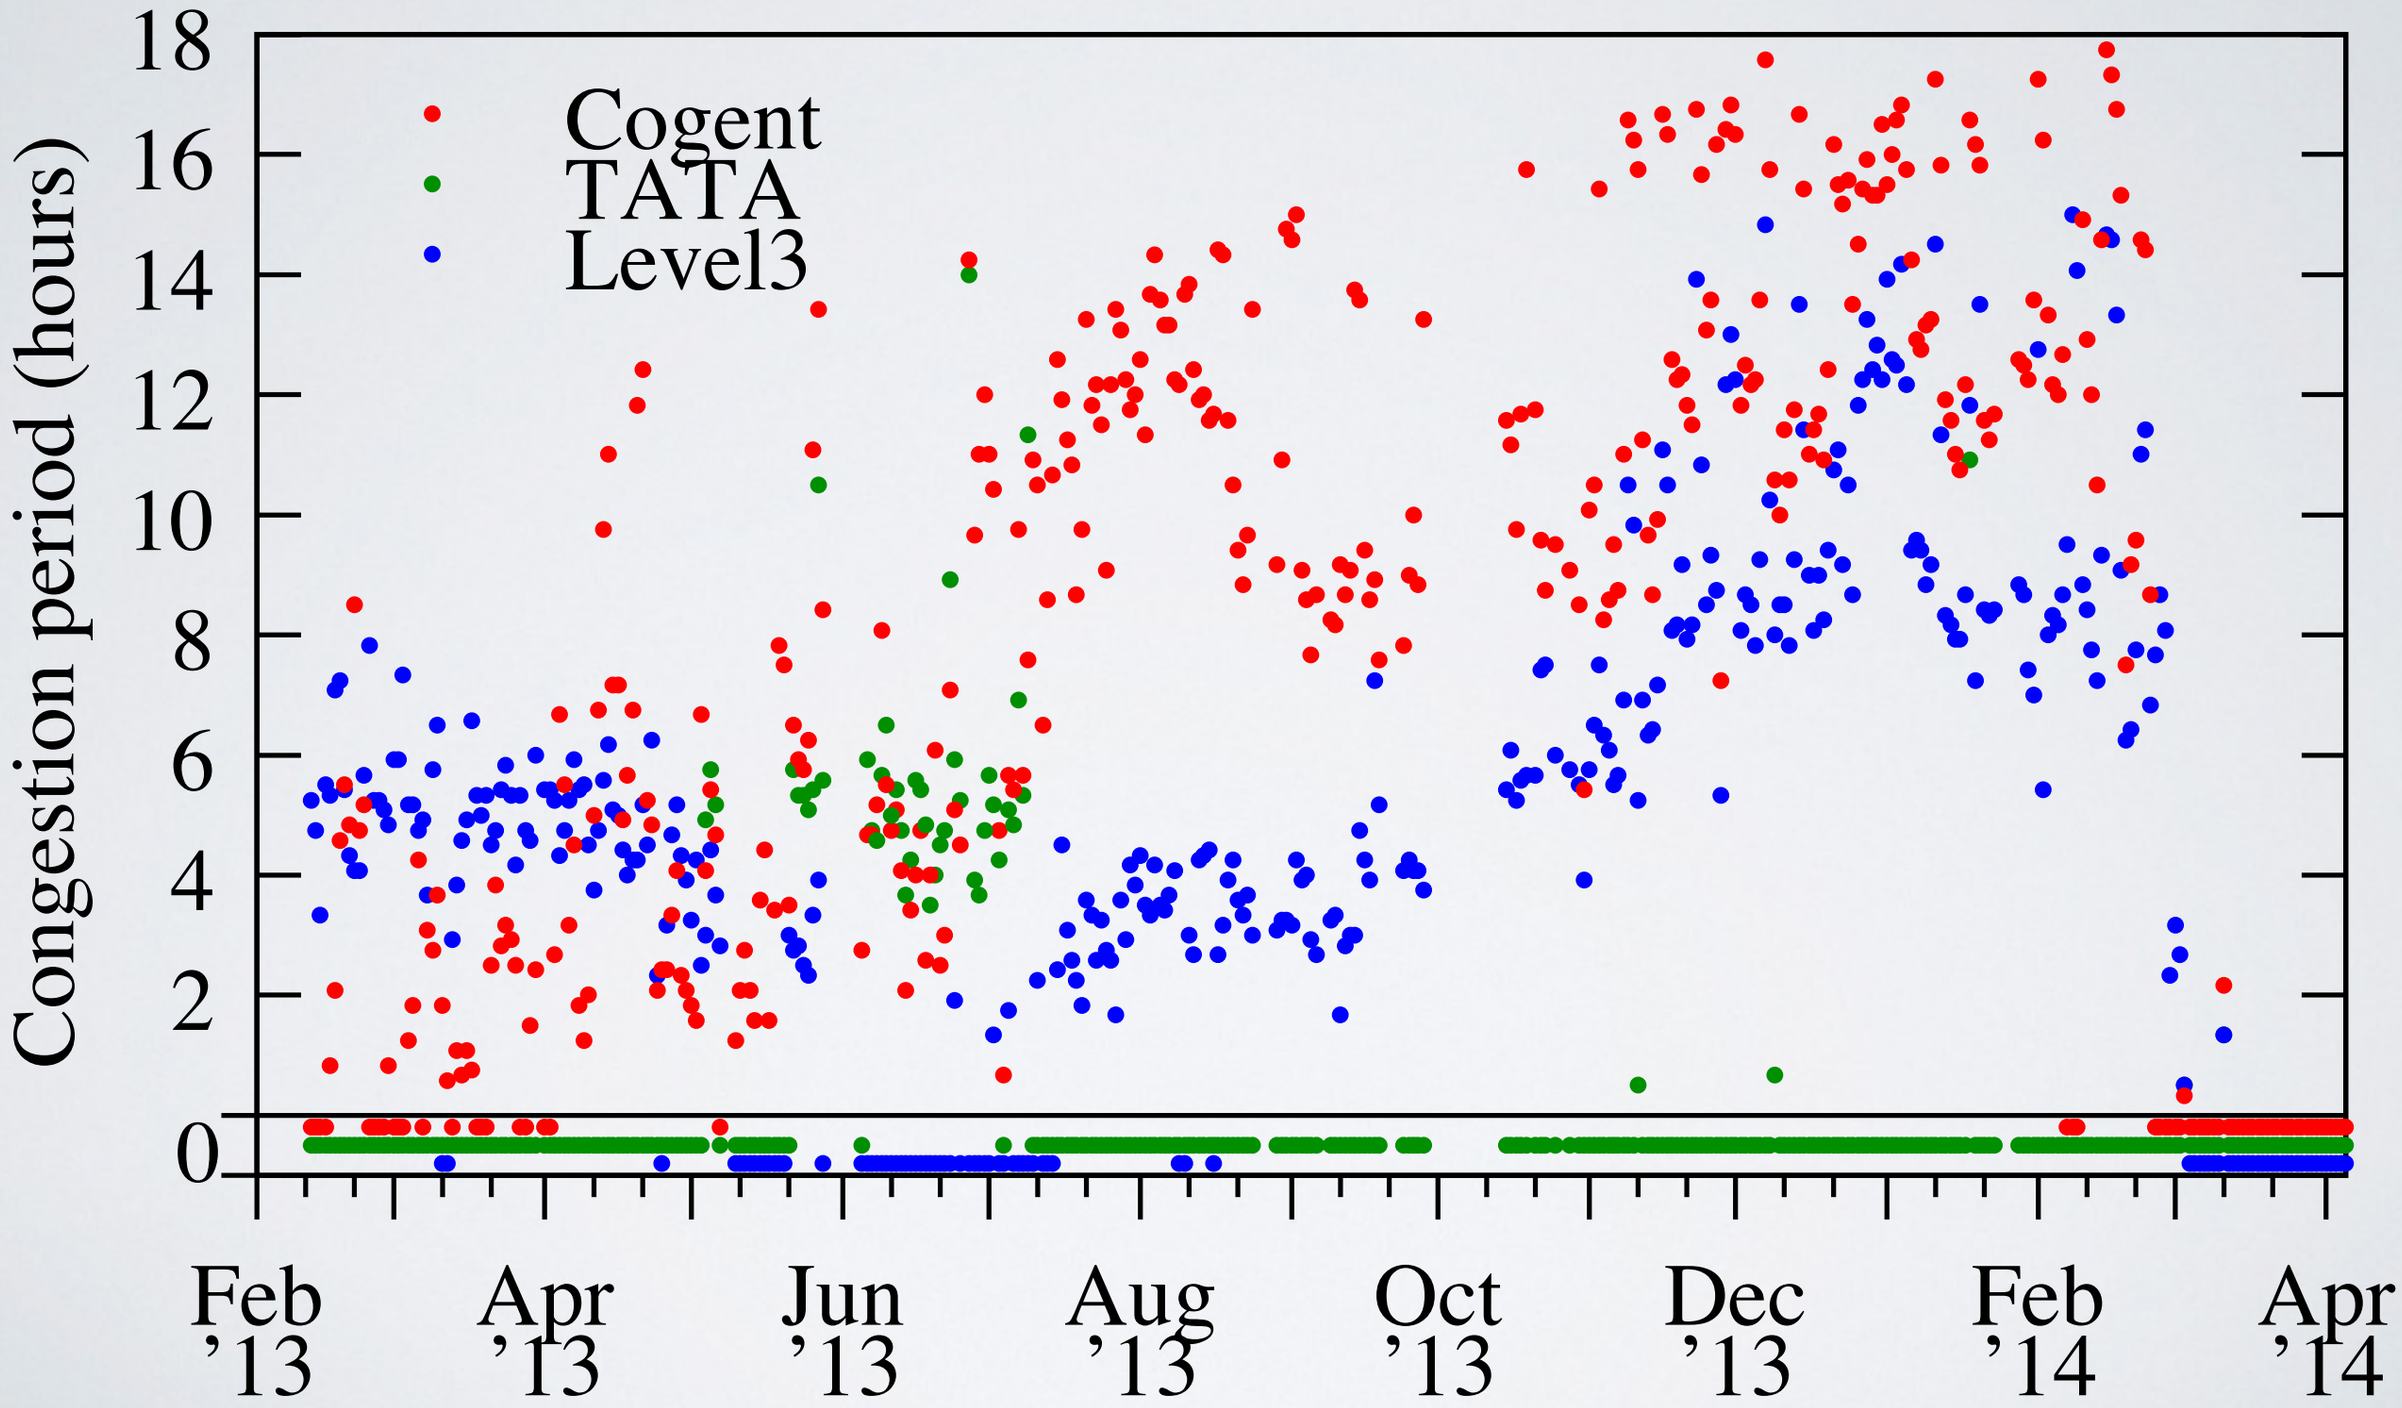

CONGESTION TRENDS

- Two interpretations
  - ability of content providers to shift traffic "firehose" (from Level3 to TATA in June 2013)

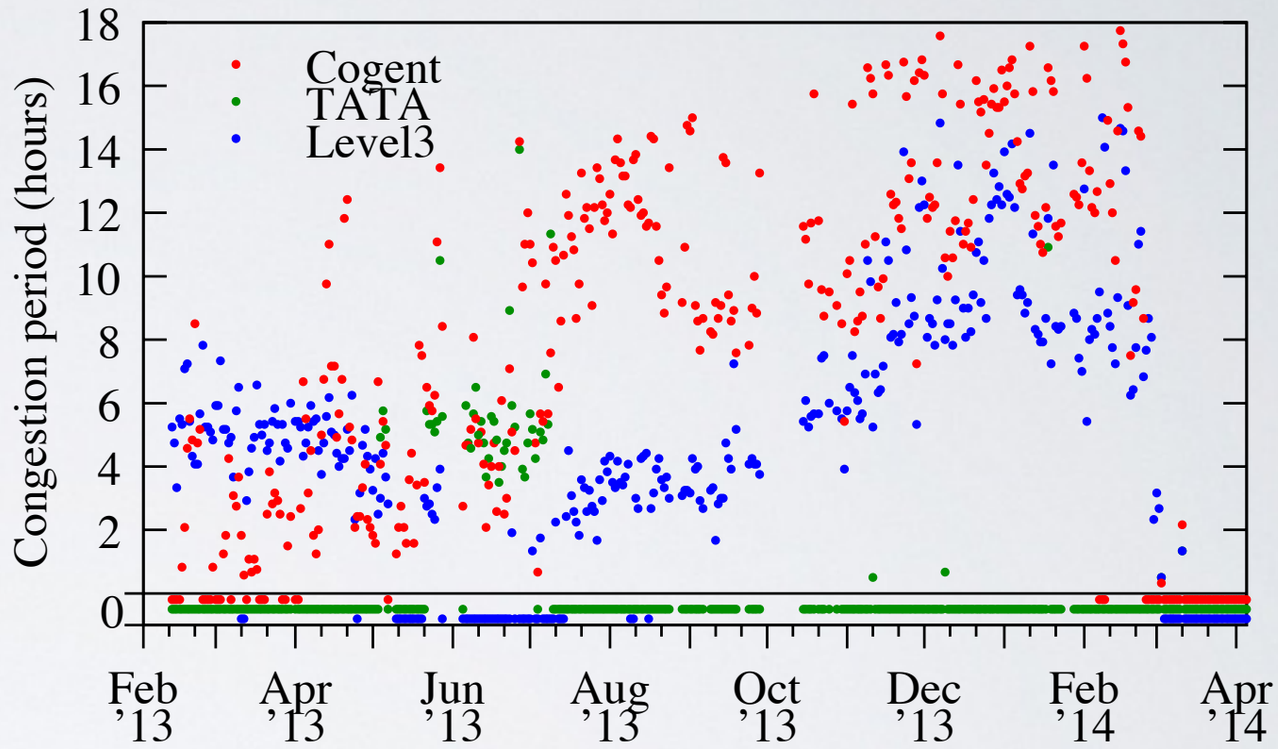

  - demonstrates year-long, worsening, congestion patterns until Netflix / Comcast peering agreement)

NETFLIX-COMCAST

(after the deal)

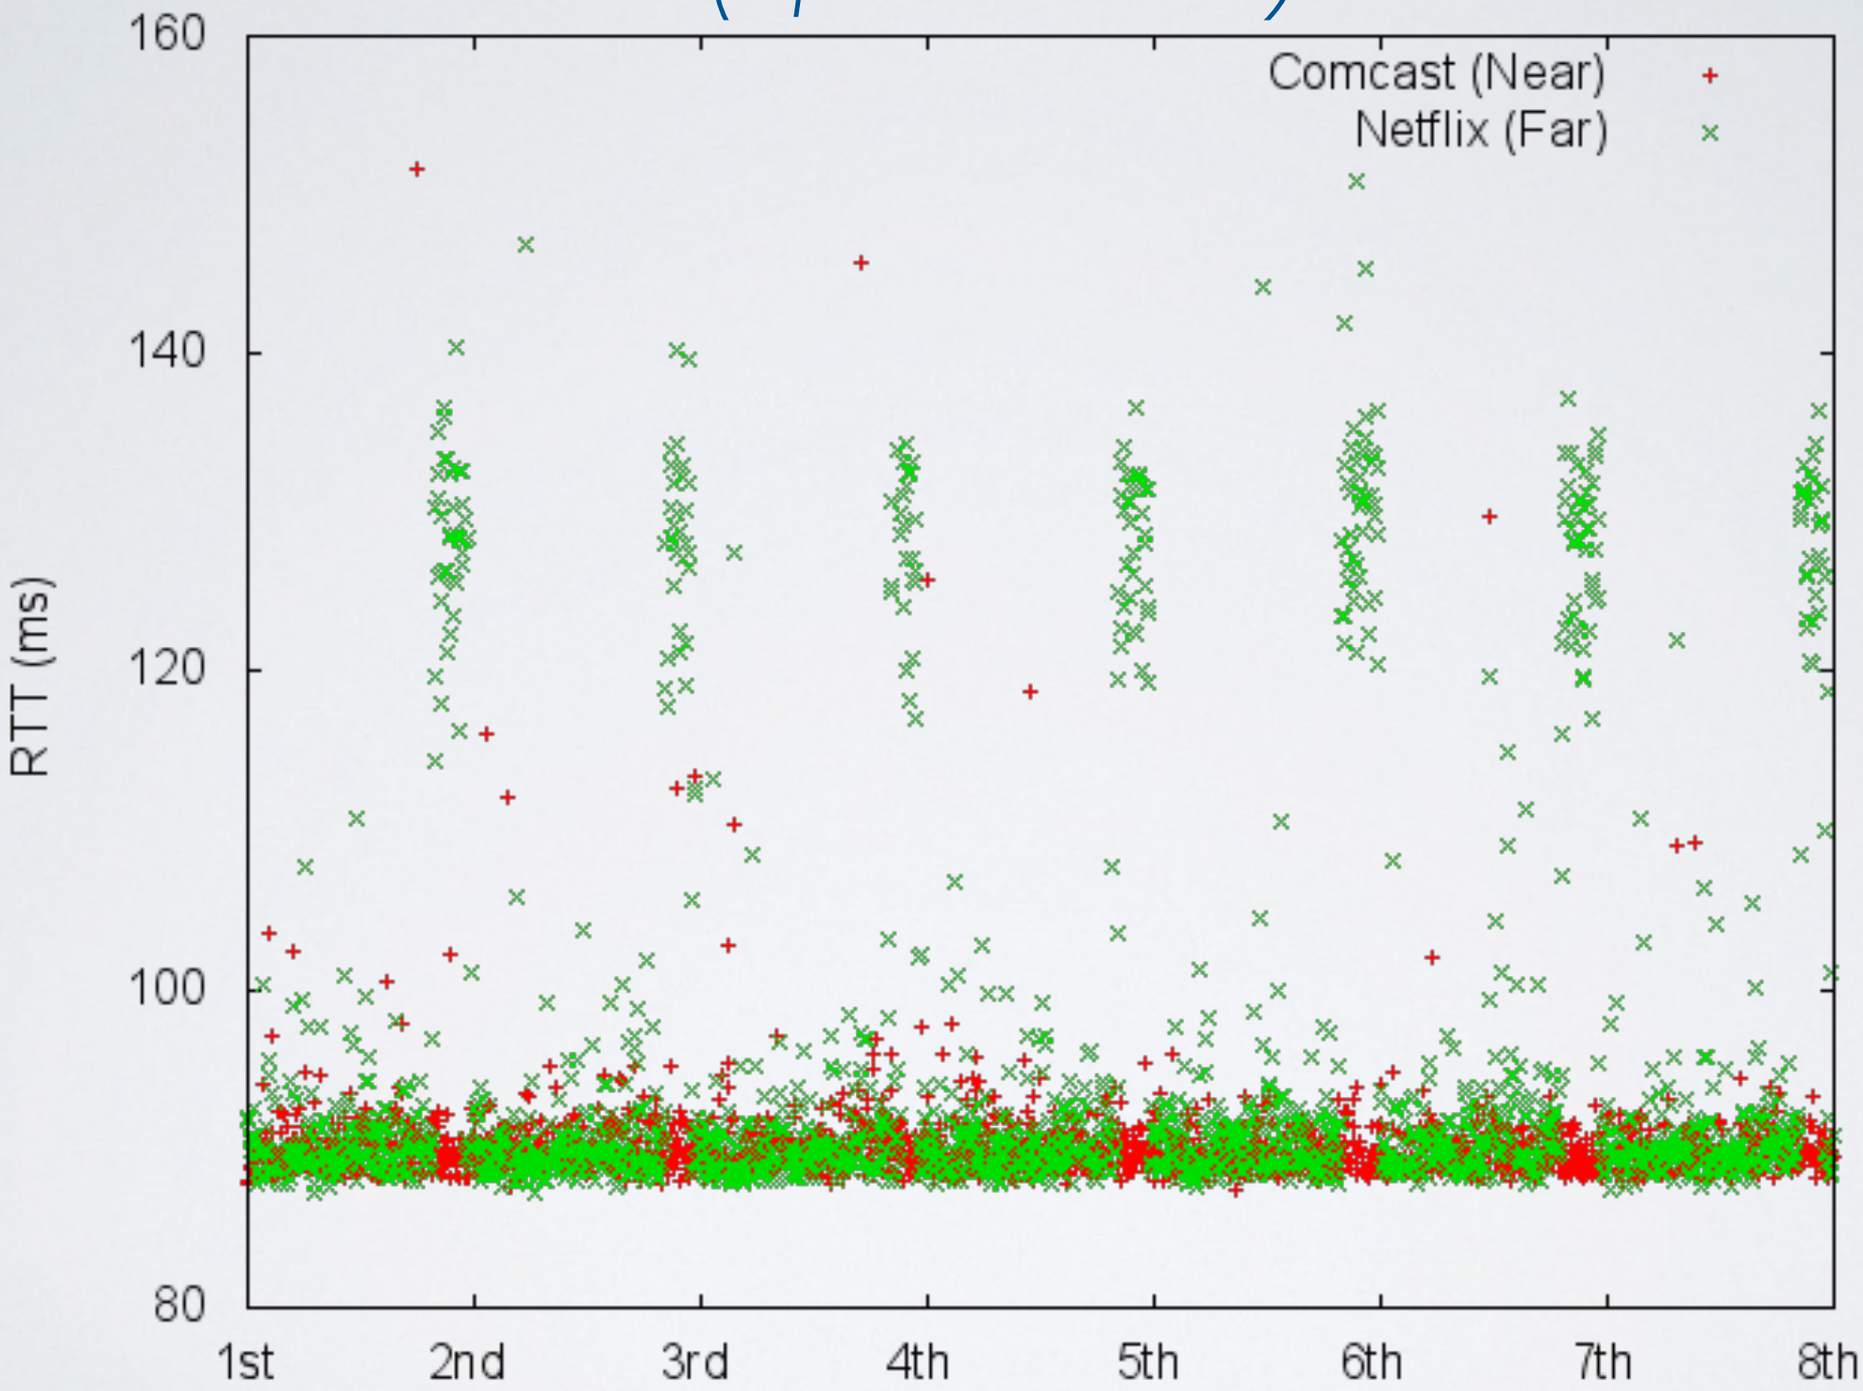

Week in April 2014, local time in San Jose

LEVEL3

(we happen to have a good view of Level3-Dallas)

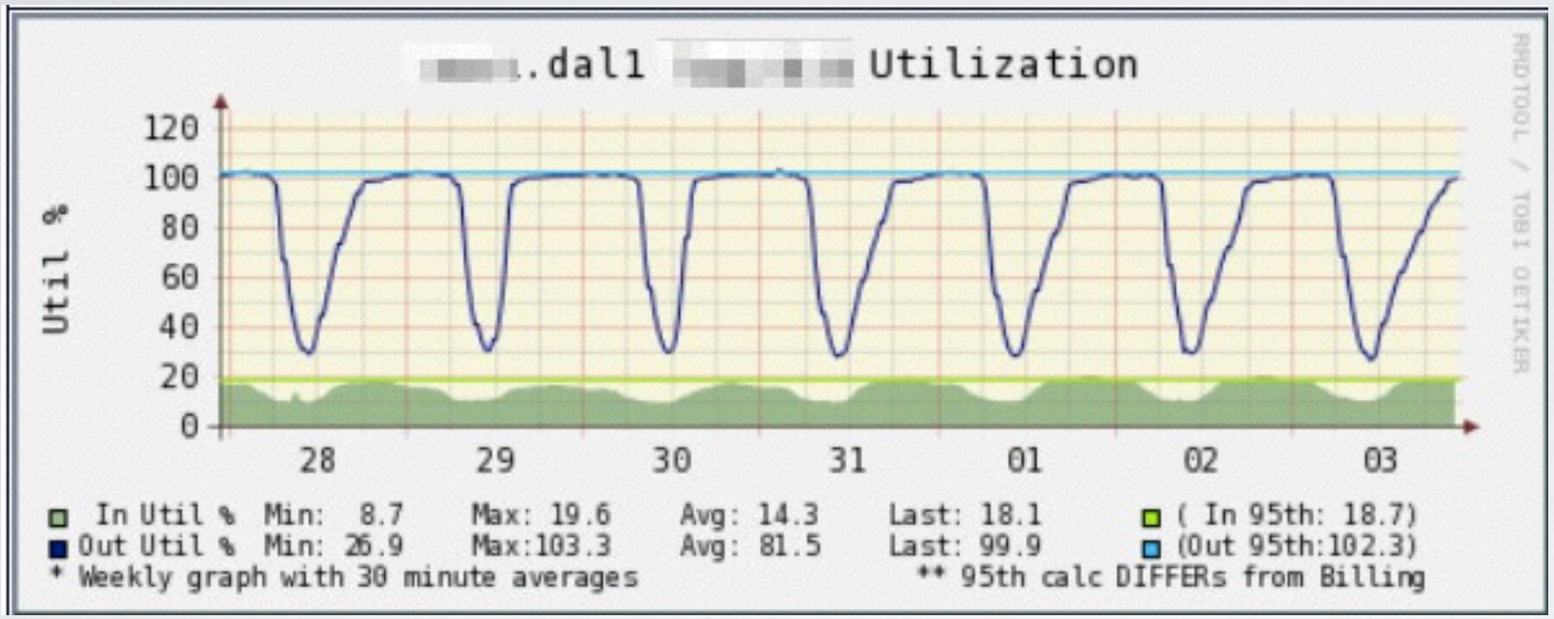

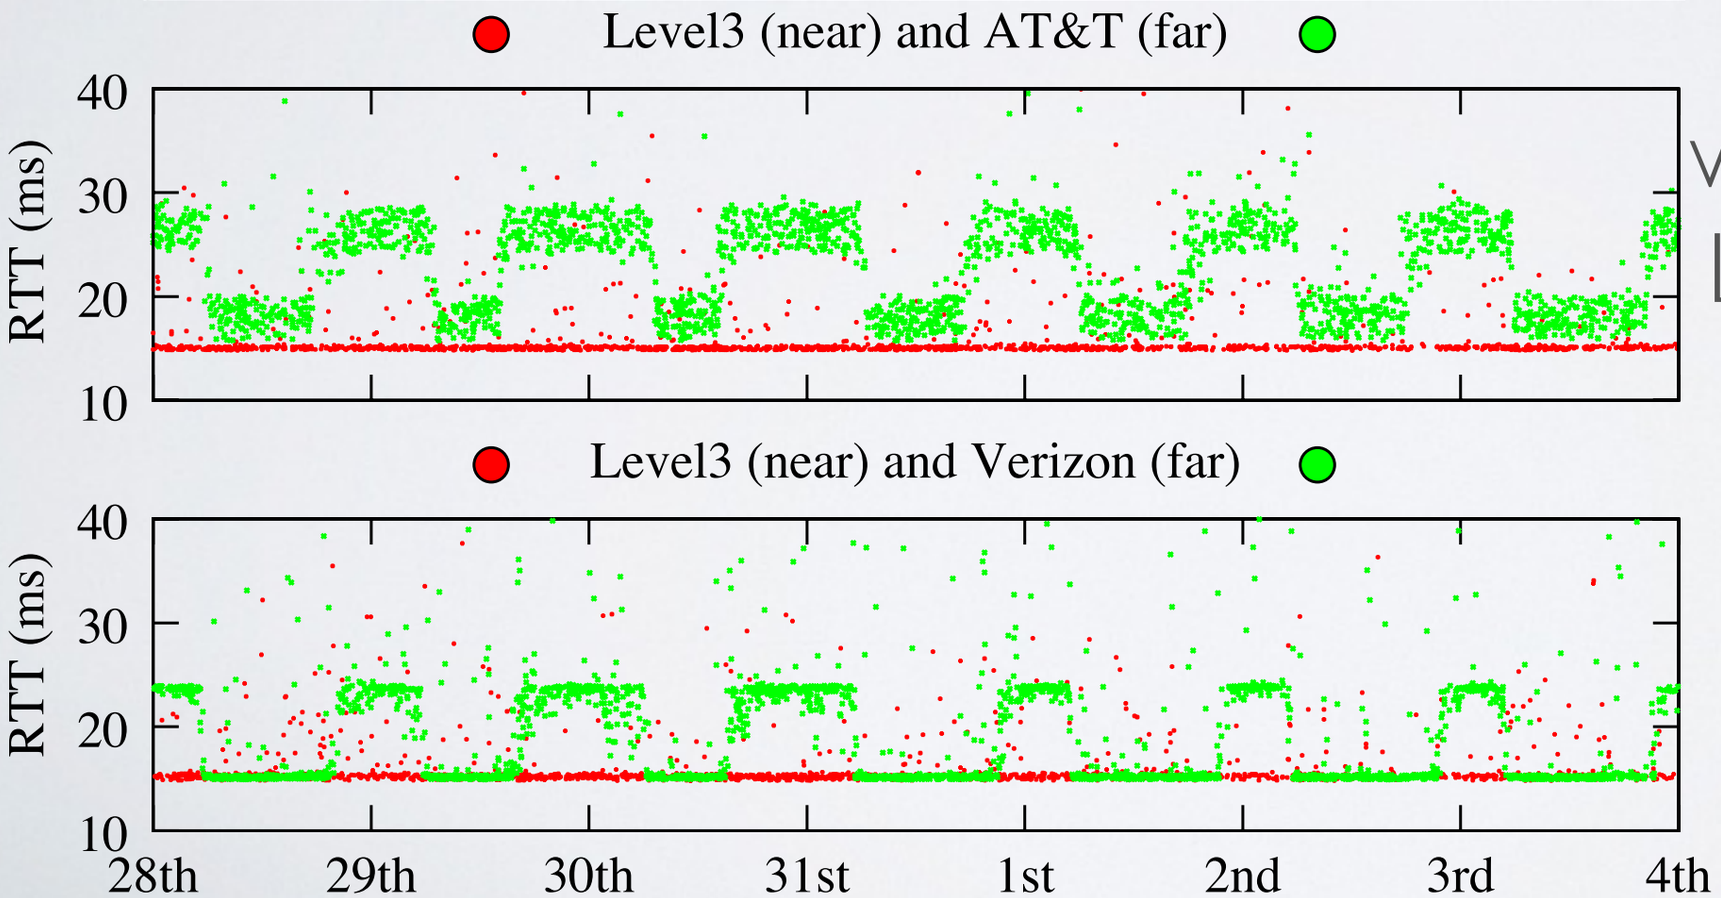

"Ground Truth"

We believe both AT&T and Verizon were congested with Level3 in April 2014

Anonymous link is probably AT&T (based on duration of level shift)

LEVEL3

- Netflix signed paid-peering deals with Verizon and AT&T in summer 2014

- Congestion signals on the Level3-AT&T and Level3-Verizon links disappeared

GOOGLE-FREE

(another dispute in the news)

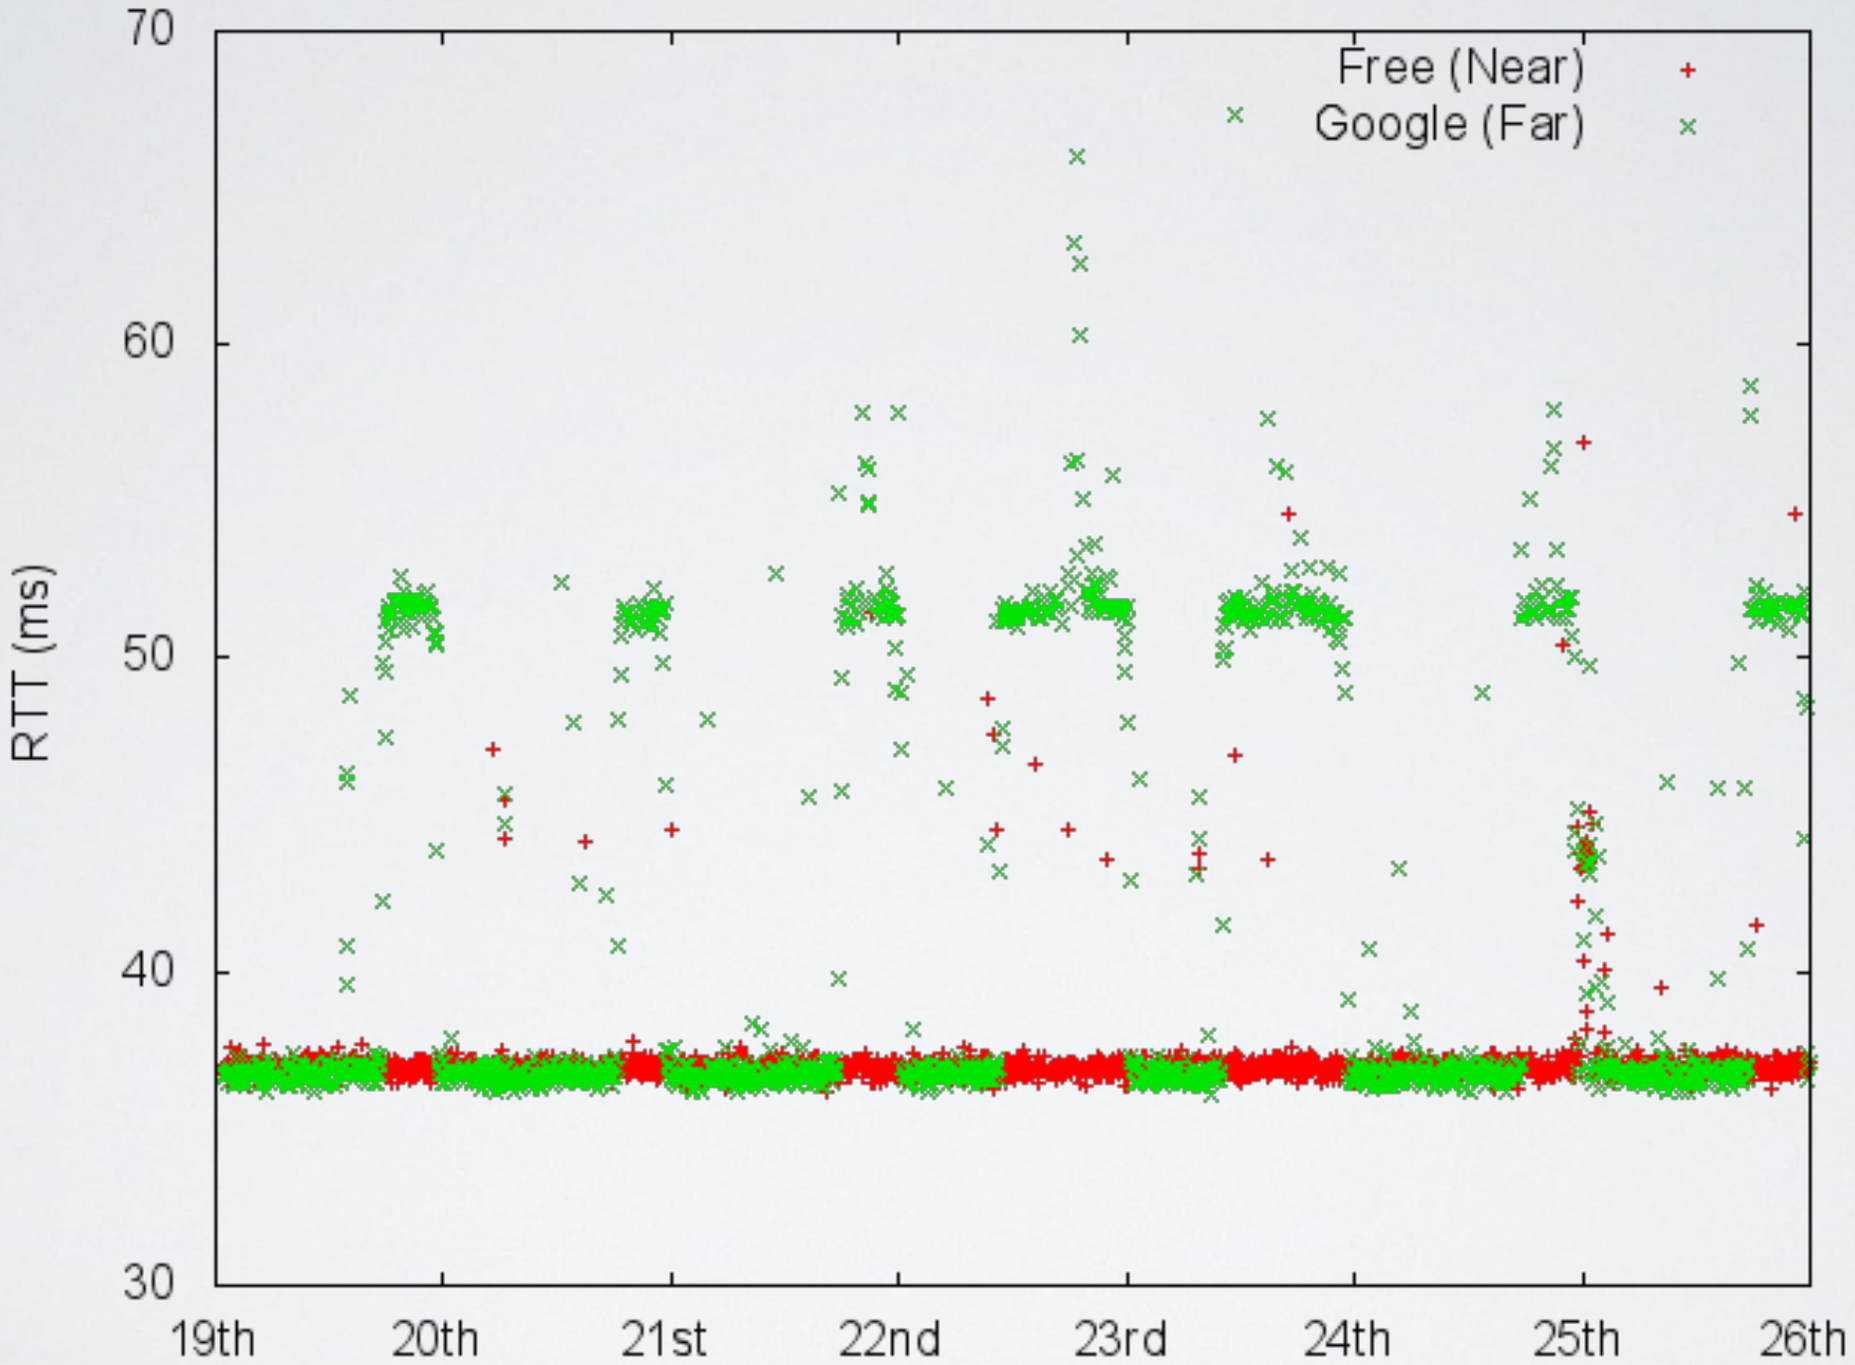

Week in March 2014, local time in France