INFERRING PERSISTENT INTERDOMAIN CONGESTION

Amogh Dhamdhere

with David Clark, Alex Gamero-Garrido, Matthew Luckie, Ricky Mok, Gautam Akiwate, Kabir Gogia, Vaibhav Bajpai, Alex Snoeren, k Claffy

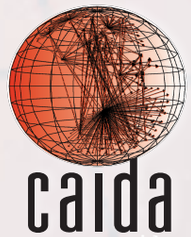

Problem: High Volume Content Strains Internet Technology and Economics

France Telecom Accused Of Holding YouTube Videos Hostage Unless It Gets More Money

from the *more-peering-disputes* dept

Netflix packets being dropped every day because Verizon wants more money

Verizon wants to be paid by consumers and Cogent, but Cogent refuses to pay.

Peak Time Congestion Hits Hardest for Broadband ISP Plusnet

Monday, April 23rd, 2018 (3:12 pm) - Score 14,382

Level 3 and Comcast Issue Statement

Jul 16, 2013

Broadband
by Karl Bode
Thu, Dec 21st 2017
10:44am

CenturyLink Pushed For Net Neutrality Repeal, Now Adorably Calls For FCC To Police Interconnection

from the *good-luck-with--that* dept

You'll probably recall that a few years ago, Netflix streams
users nationwide. Eventually, Netflix, Level3 and Cogent st
Netflix's end, but was occurring at peering points, where th
intentionally letting their networks congest by refusing to

el 3 and Comcast have resolved their prior interconnect
ails will not be released.

Verizon denies using net neutrality victory to sabotage Netflix, Amazon

BY BRIAN FUNG February 5 at 1:59 pm

Cogent Gearing for Another Peering Battle

Confirmed: Comcast and Netflix have signed a paid peering agreement

by Stacey Higginbotham FEB. 23, 2014 - 9:27 A

'Peering' Into AOL-MSN Outage

Netflix still sucks on AT&T, and now AT&T plans to offer Netflix clone

AT&T partners with an investment firm to buy and launch streaming services.

caida

Manifestation: Interdomain Congestion

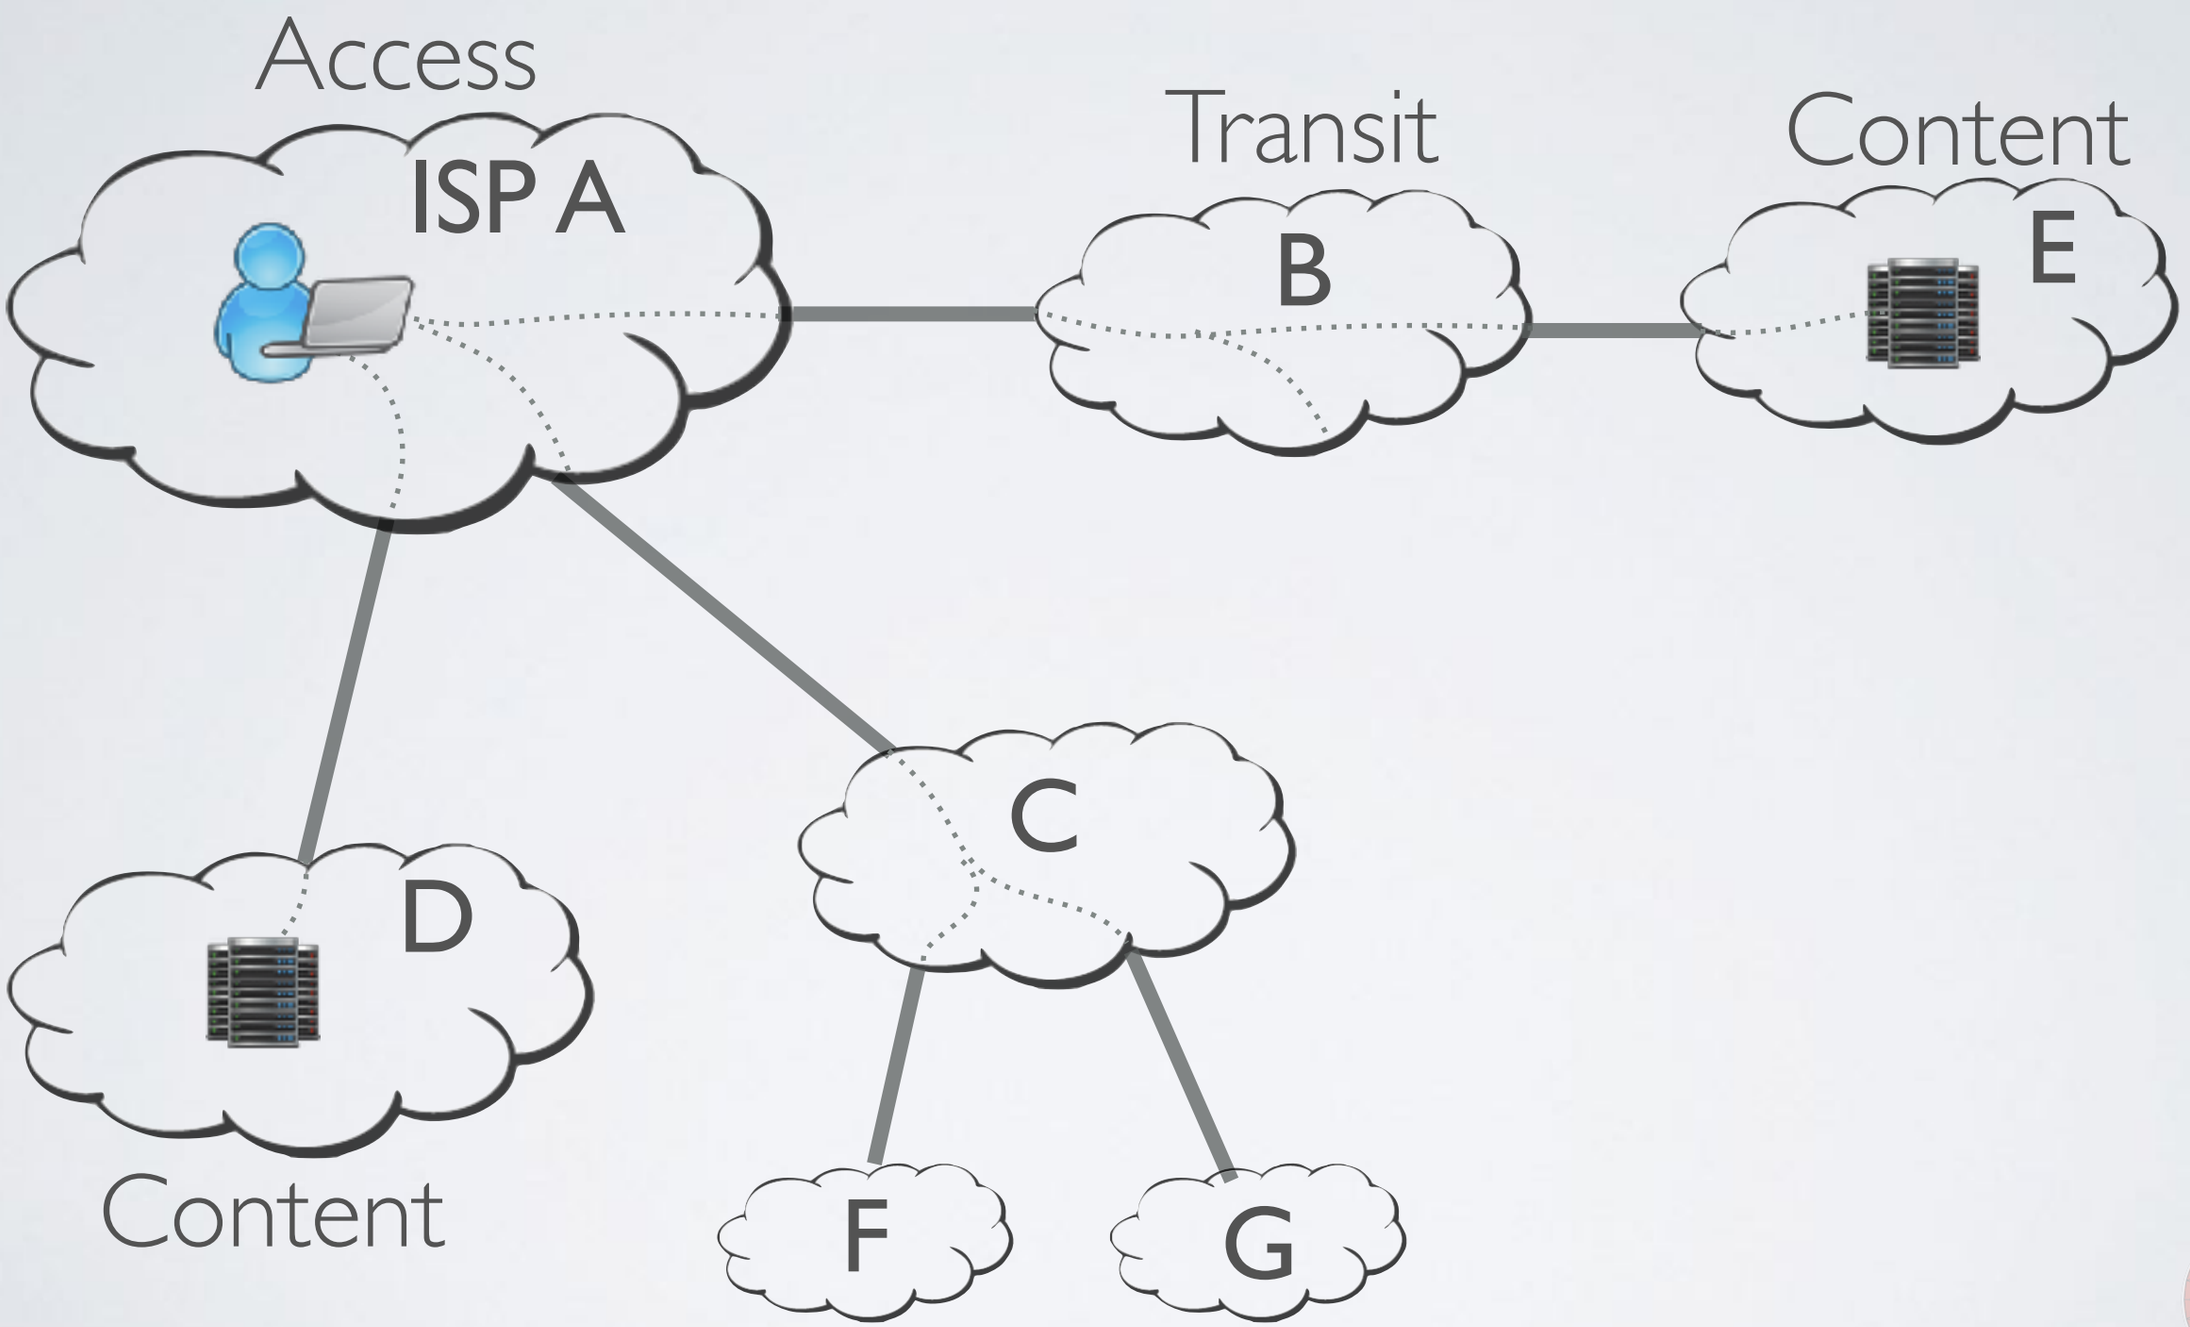

caida

Manifestation: Interdomain Congestion

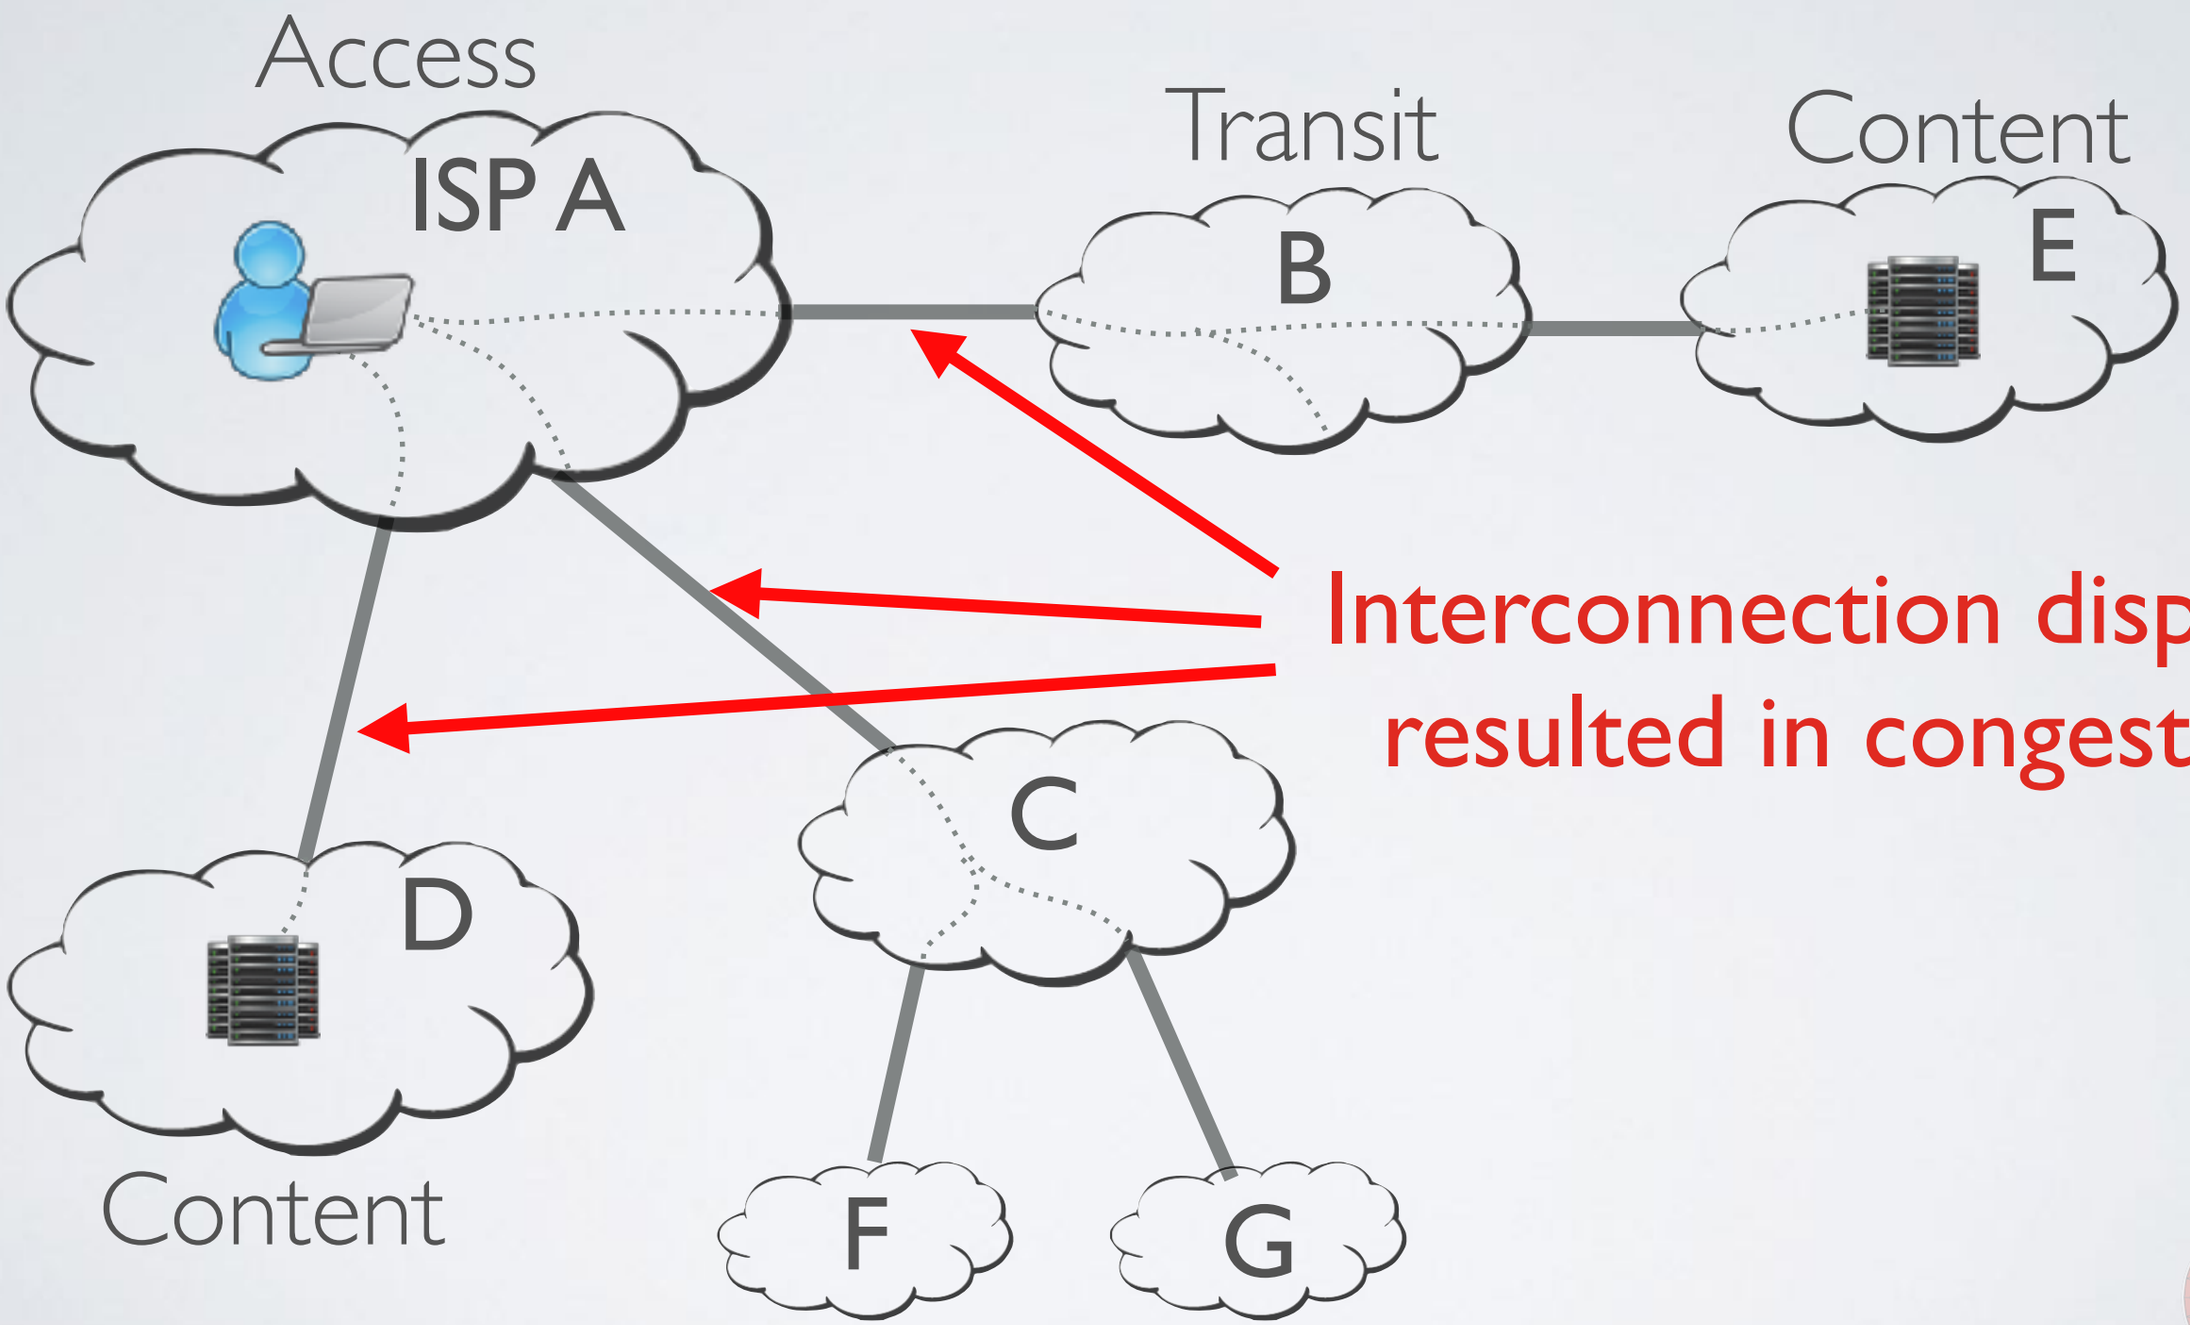

Manifestation: Interdomain Congestion

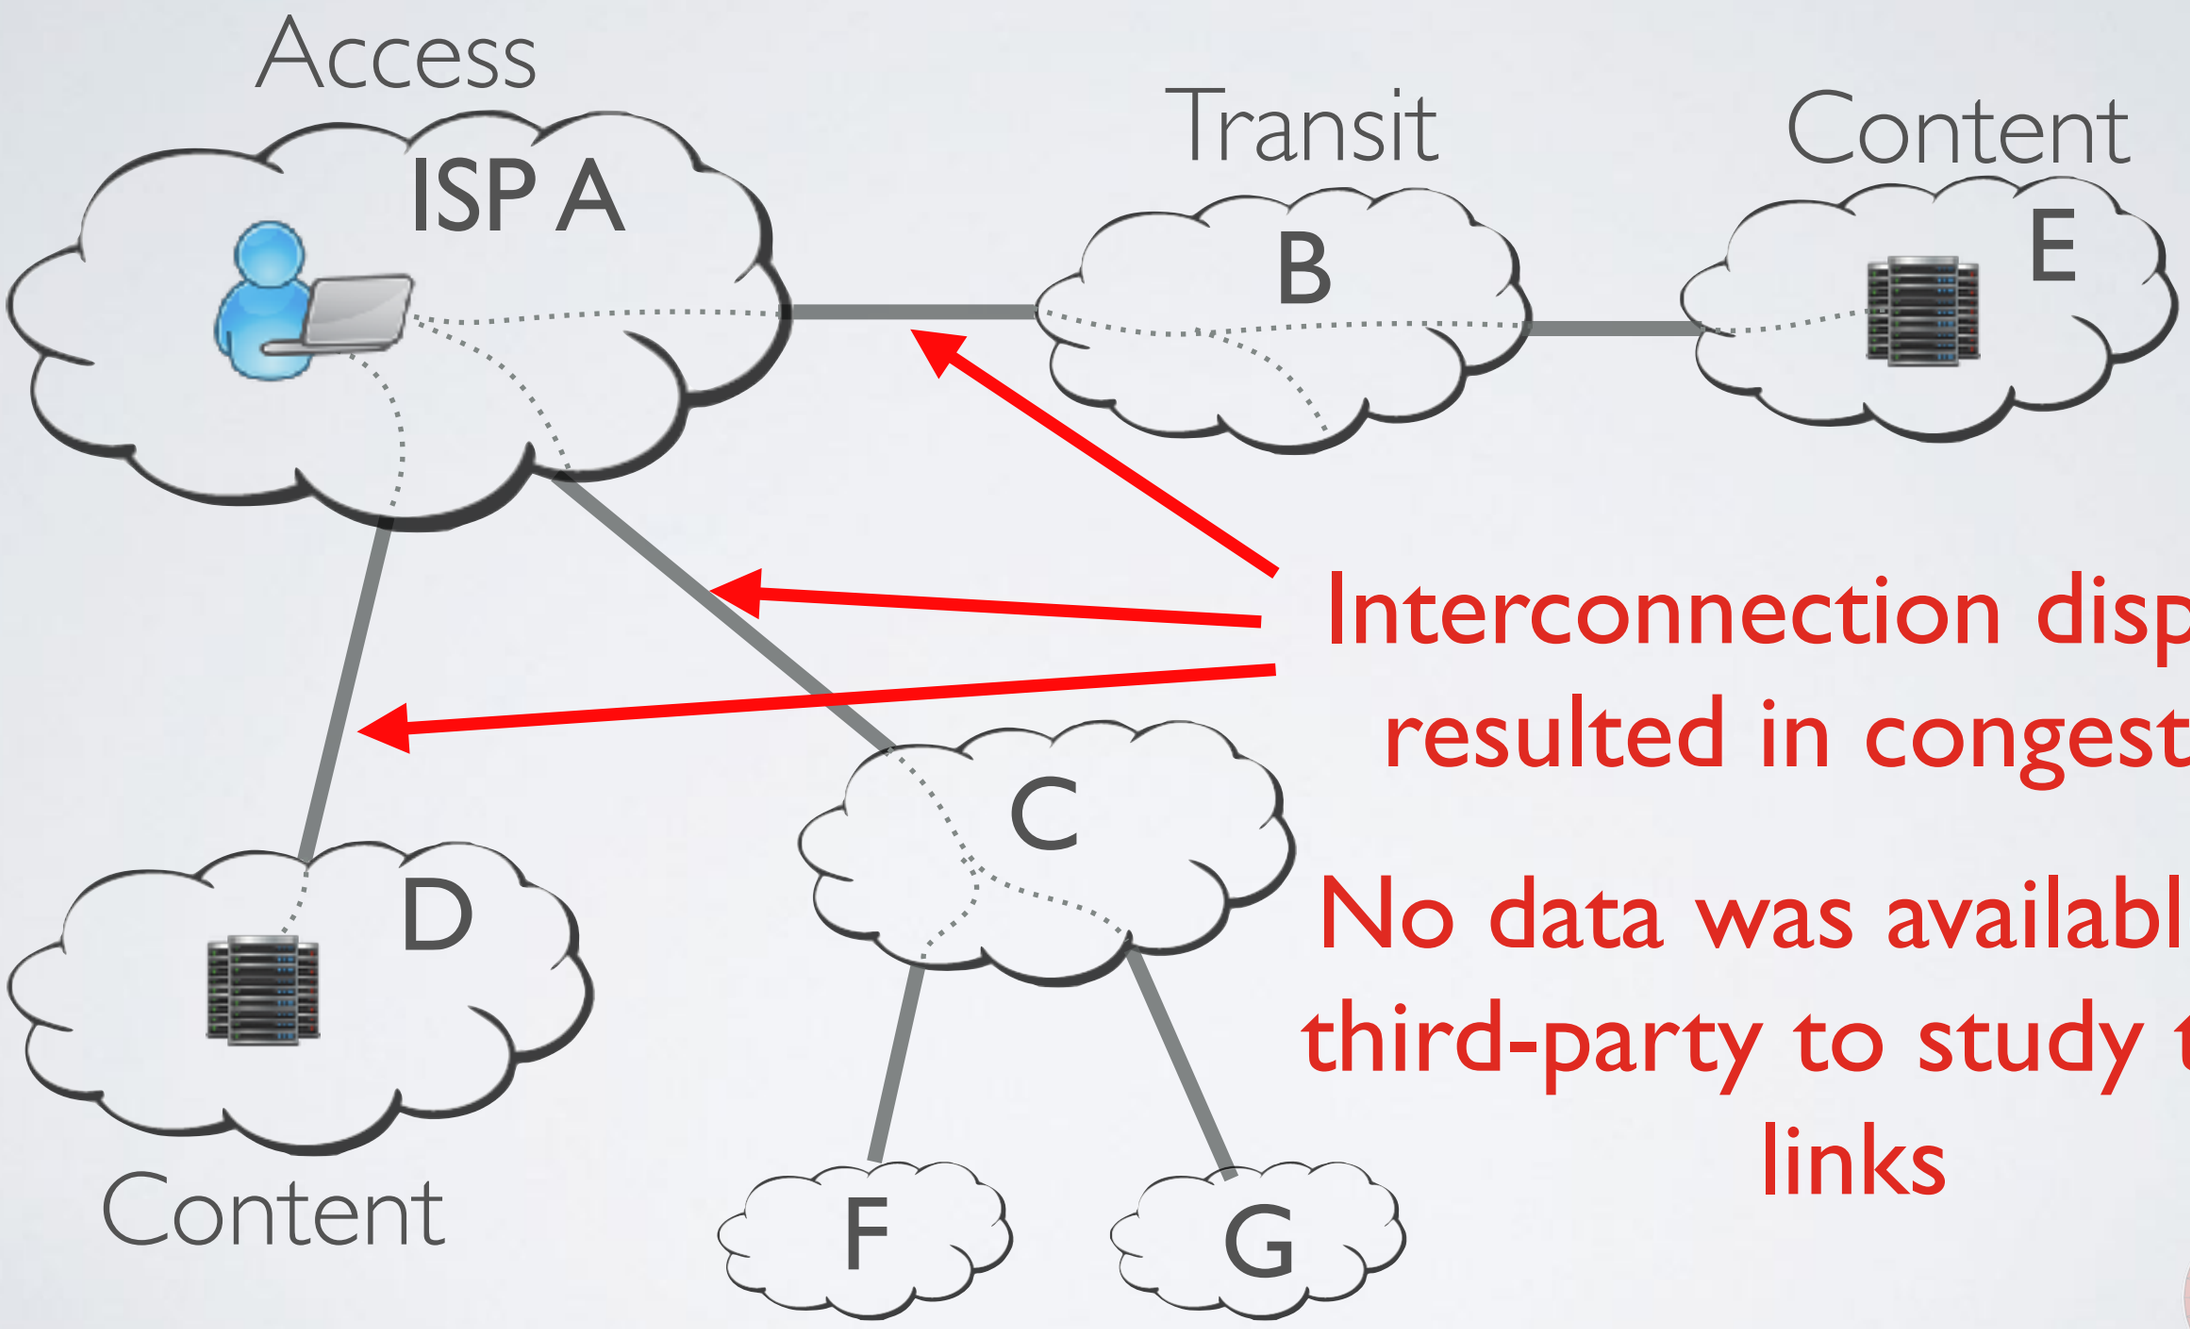

Interconnection disputes resulted in congestion

No data was available for third-party to study these links

caida

Consequences of the Problem

Consequences of the Problem

- Congestion on transit links **affects parties other than those involved in the dispute**

Consequences of the Problem

- Congestion on transit links **affects parties other than those involved in the dispute**
- **Limited data available to regulators and researchers** to increase transparency and empirical grounding of debate

Consequences of the Problem

- Congestion on transit links **affects parties other than those involved in the dispute**
- **Limited data available to regulators and researchers** to increase transparency and empirical grounding of debate
- Our goal: **third-party inference of congestion at interdomain interconnections**

Consequences of the Problem

- Congestion on transit links **affects parties other than those involved in the dispute**
- **Limited data available to regulators and researchers** to increase transparency and empirical grounding of debate
- Our goal: **third-party inference of congestion at interdomain interconnections**
- Scientific approach to achieving this goal involves challenges in **network inference, system development and data mining**

caida

Our Contributions

Our Contributions

1. **Methodology**: Operationalized a lightweight method for third-party inference of interdomain congestion, conducted thorough validation

Our Contributions

1. **Methodology**: Operationalized a lightweight method for third-party inference of interdomain congestion, conducted thorough validation

2. **System**: Built data collection and analysis platform to support the entire scientific workflow, and enable others to access and further study the data (ongoing)

Our Contributions

1. **Methodology**: Operationalized a lightweight method for third-party inference of interdomain congestion, conducted thorough validation

2. **System**: Built data collection and analysis platform to support the entire scientific workflow, and enable others to access and further study the data (ongoing)

3. **Observations**: Studied 8 large U.S. broadband providers from March 2016 to Dec 2017 (data collection ongoing)

Method: Time Series Latency Probes

(But first, an observation)

Peak-hour congestion fills up router buffers, resulting in elevated latency across an interdomain link

Method: Time Series Latency Probes

(But first, an observation)

Peak-hour congestion fills up router buffers, resulting in elevated latency across an interdomain link

How do we measure latency across an interdomain link?

caida

Time Series Latency Probes (TSLP)

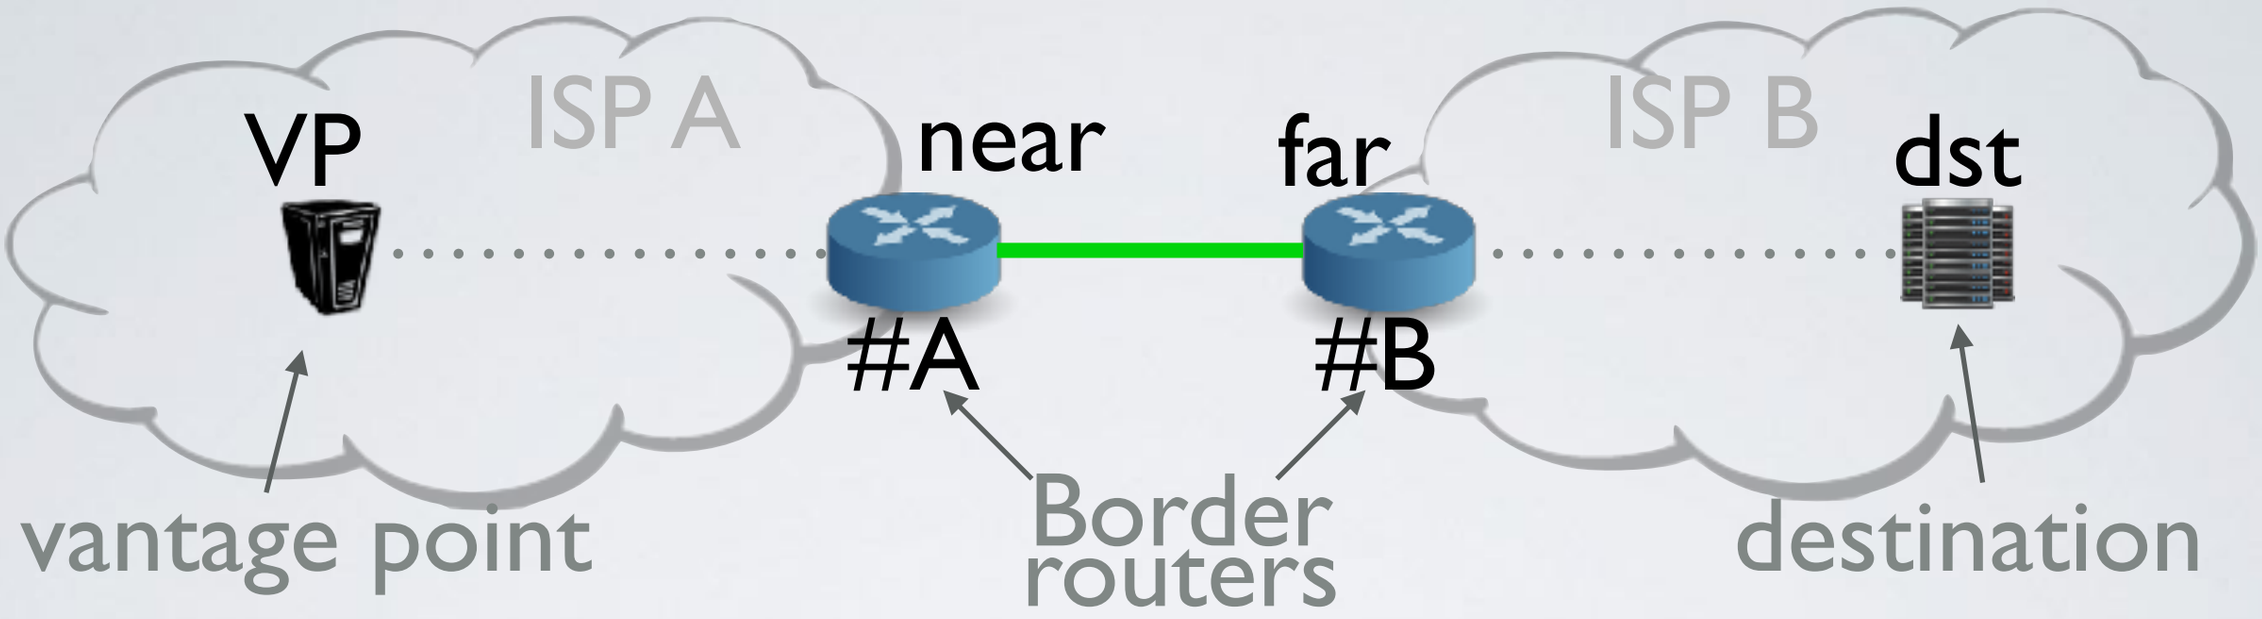

caida

Time Series Latency Probes (TSLP)

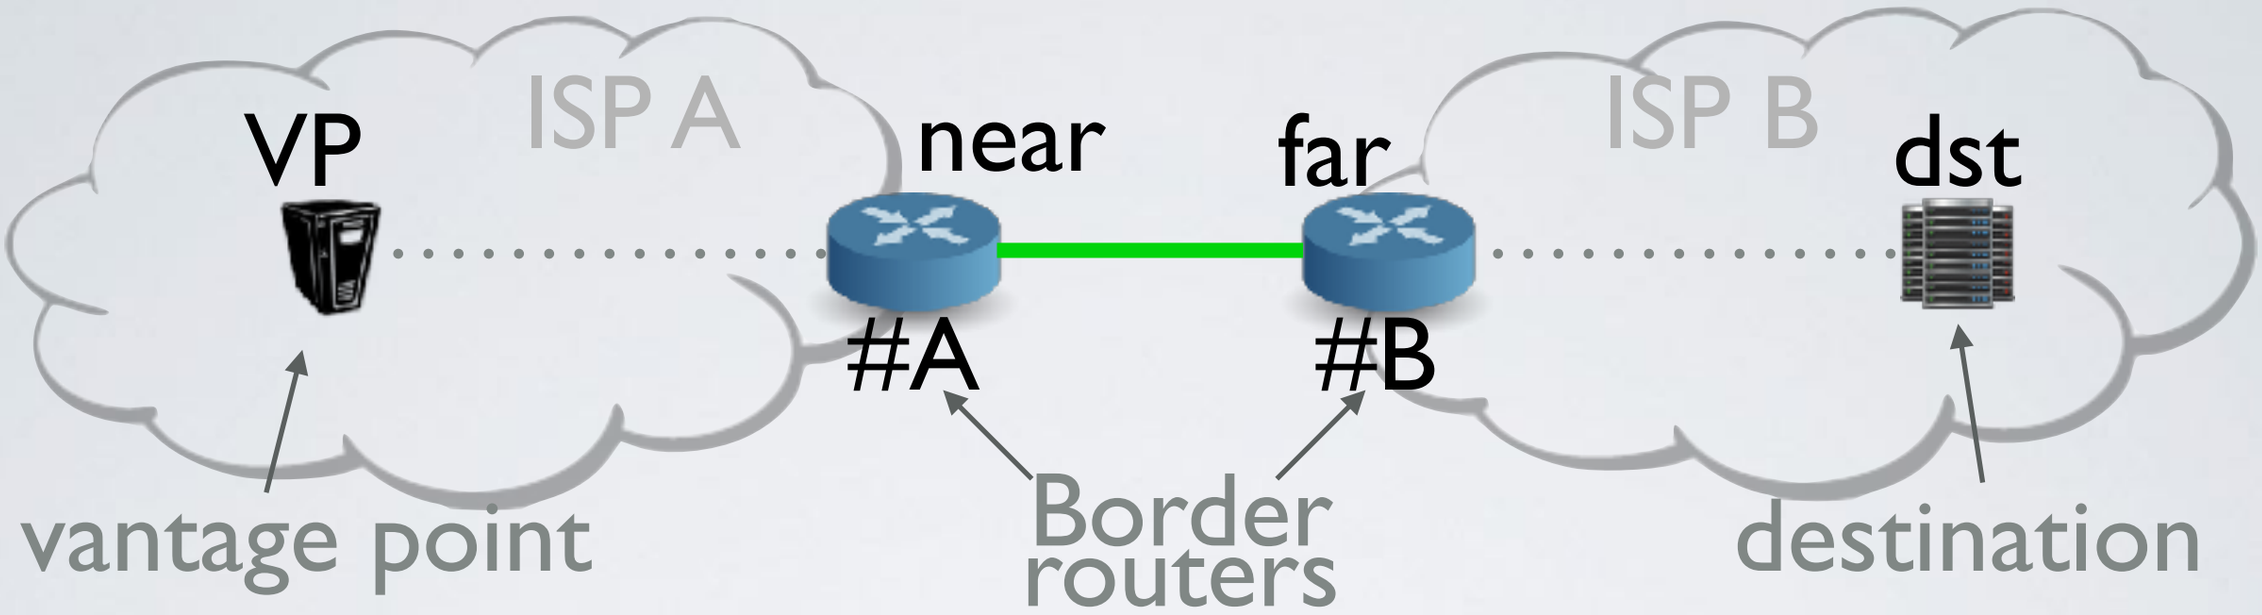

caida

Time Series Latency Probes (TSLP)

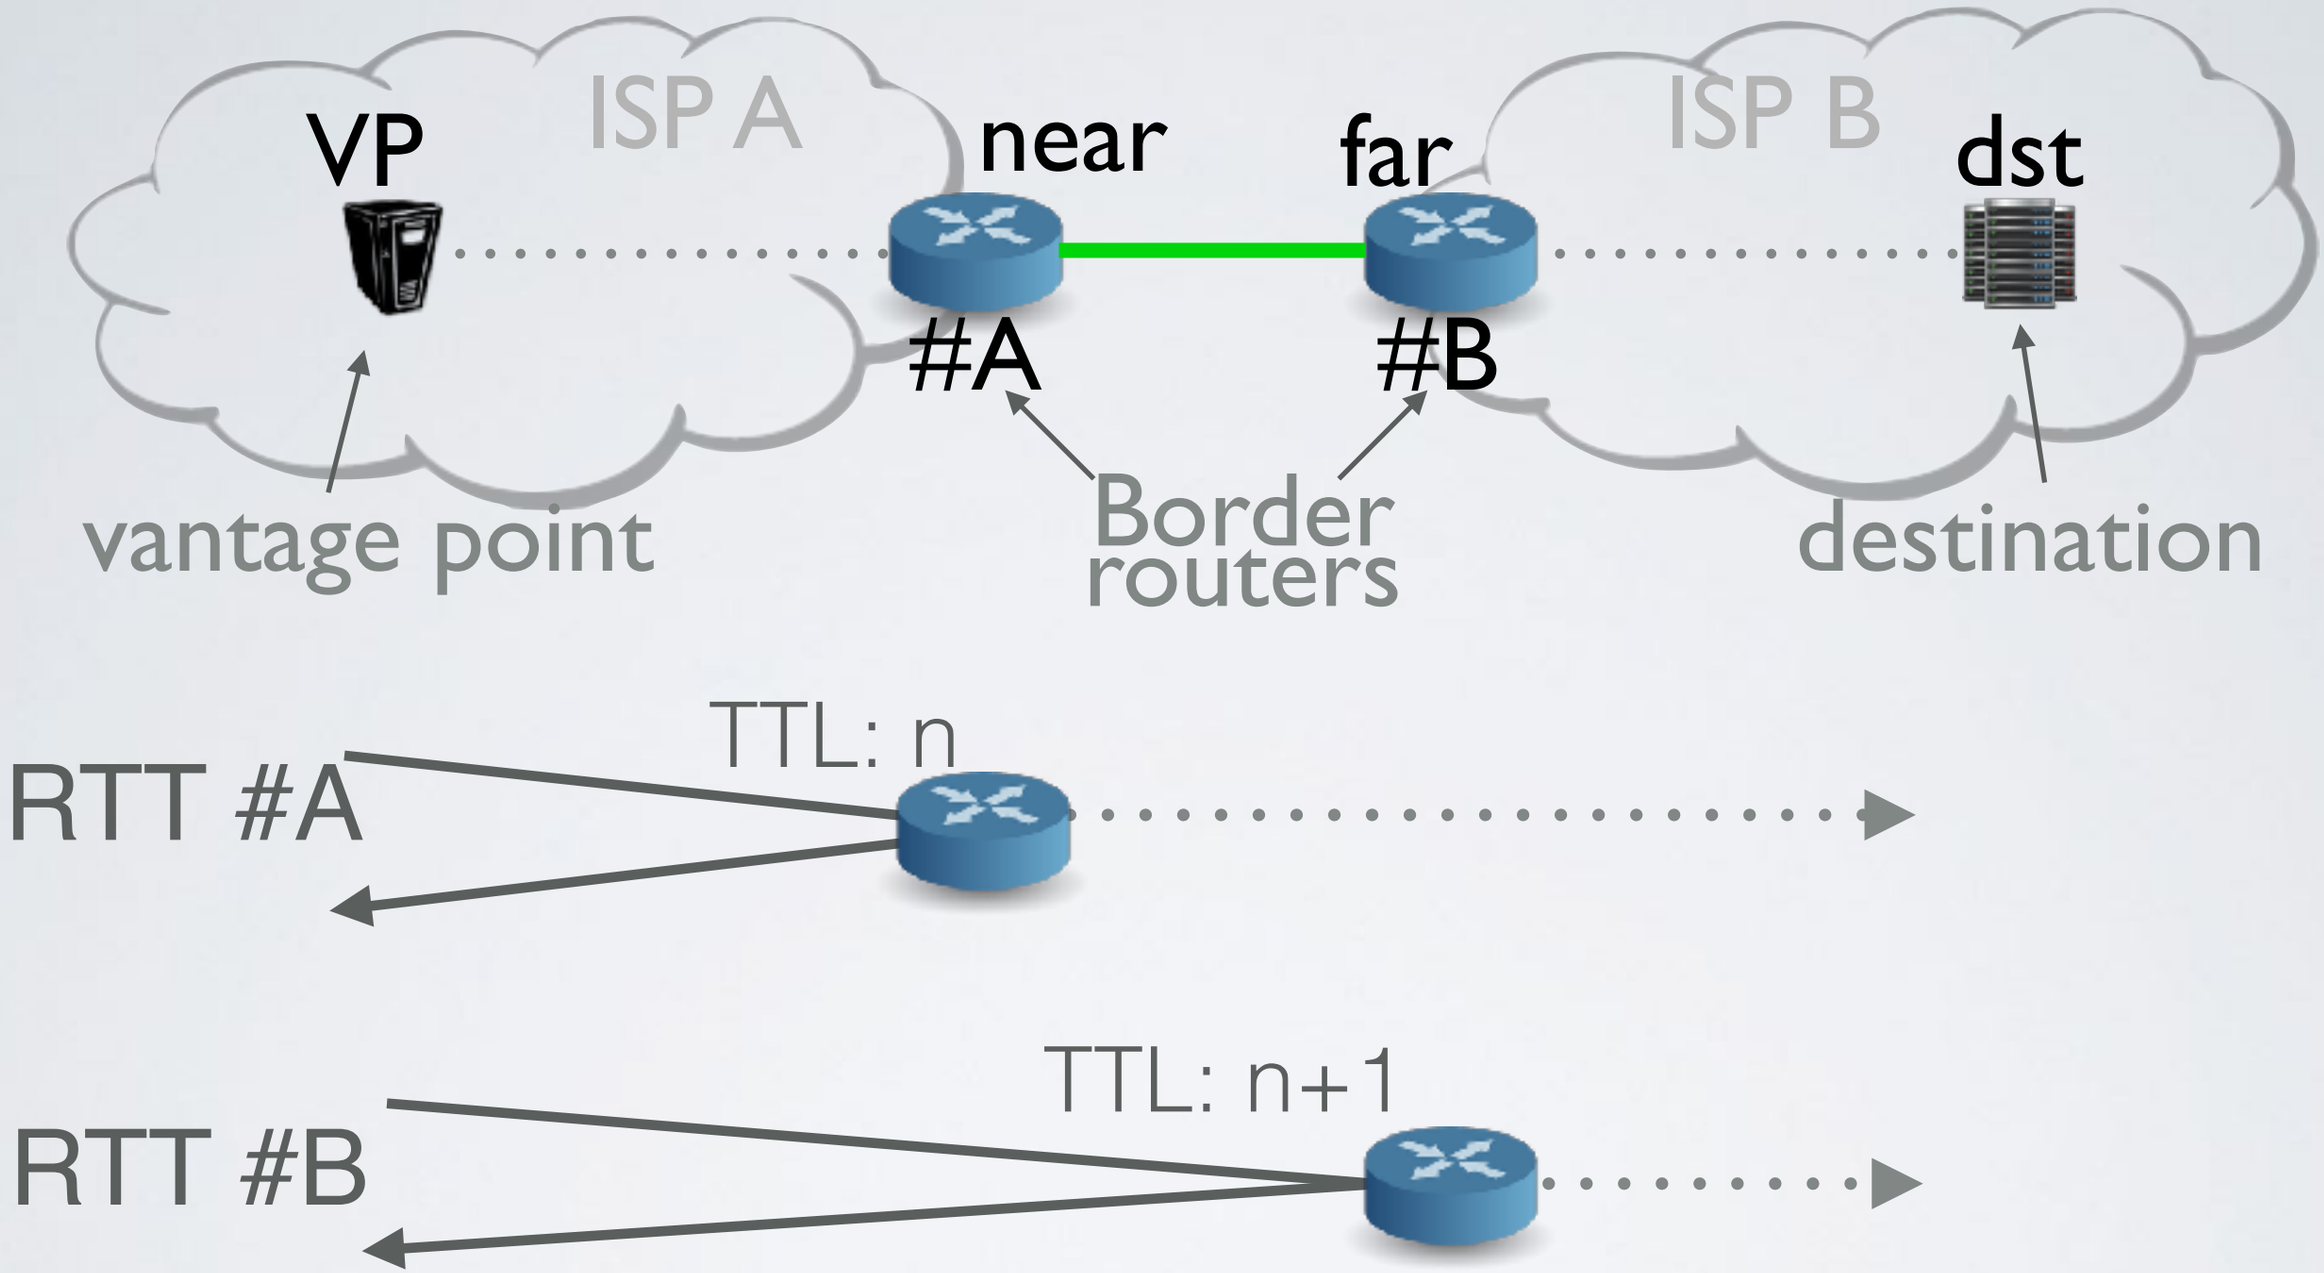

caida

Time Series Latency Probes (TSLP)

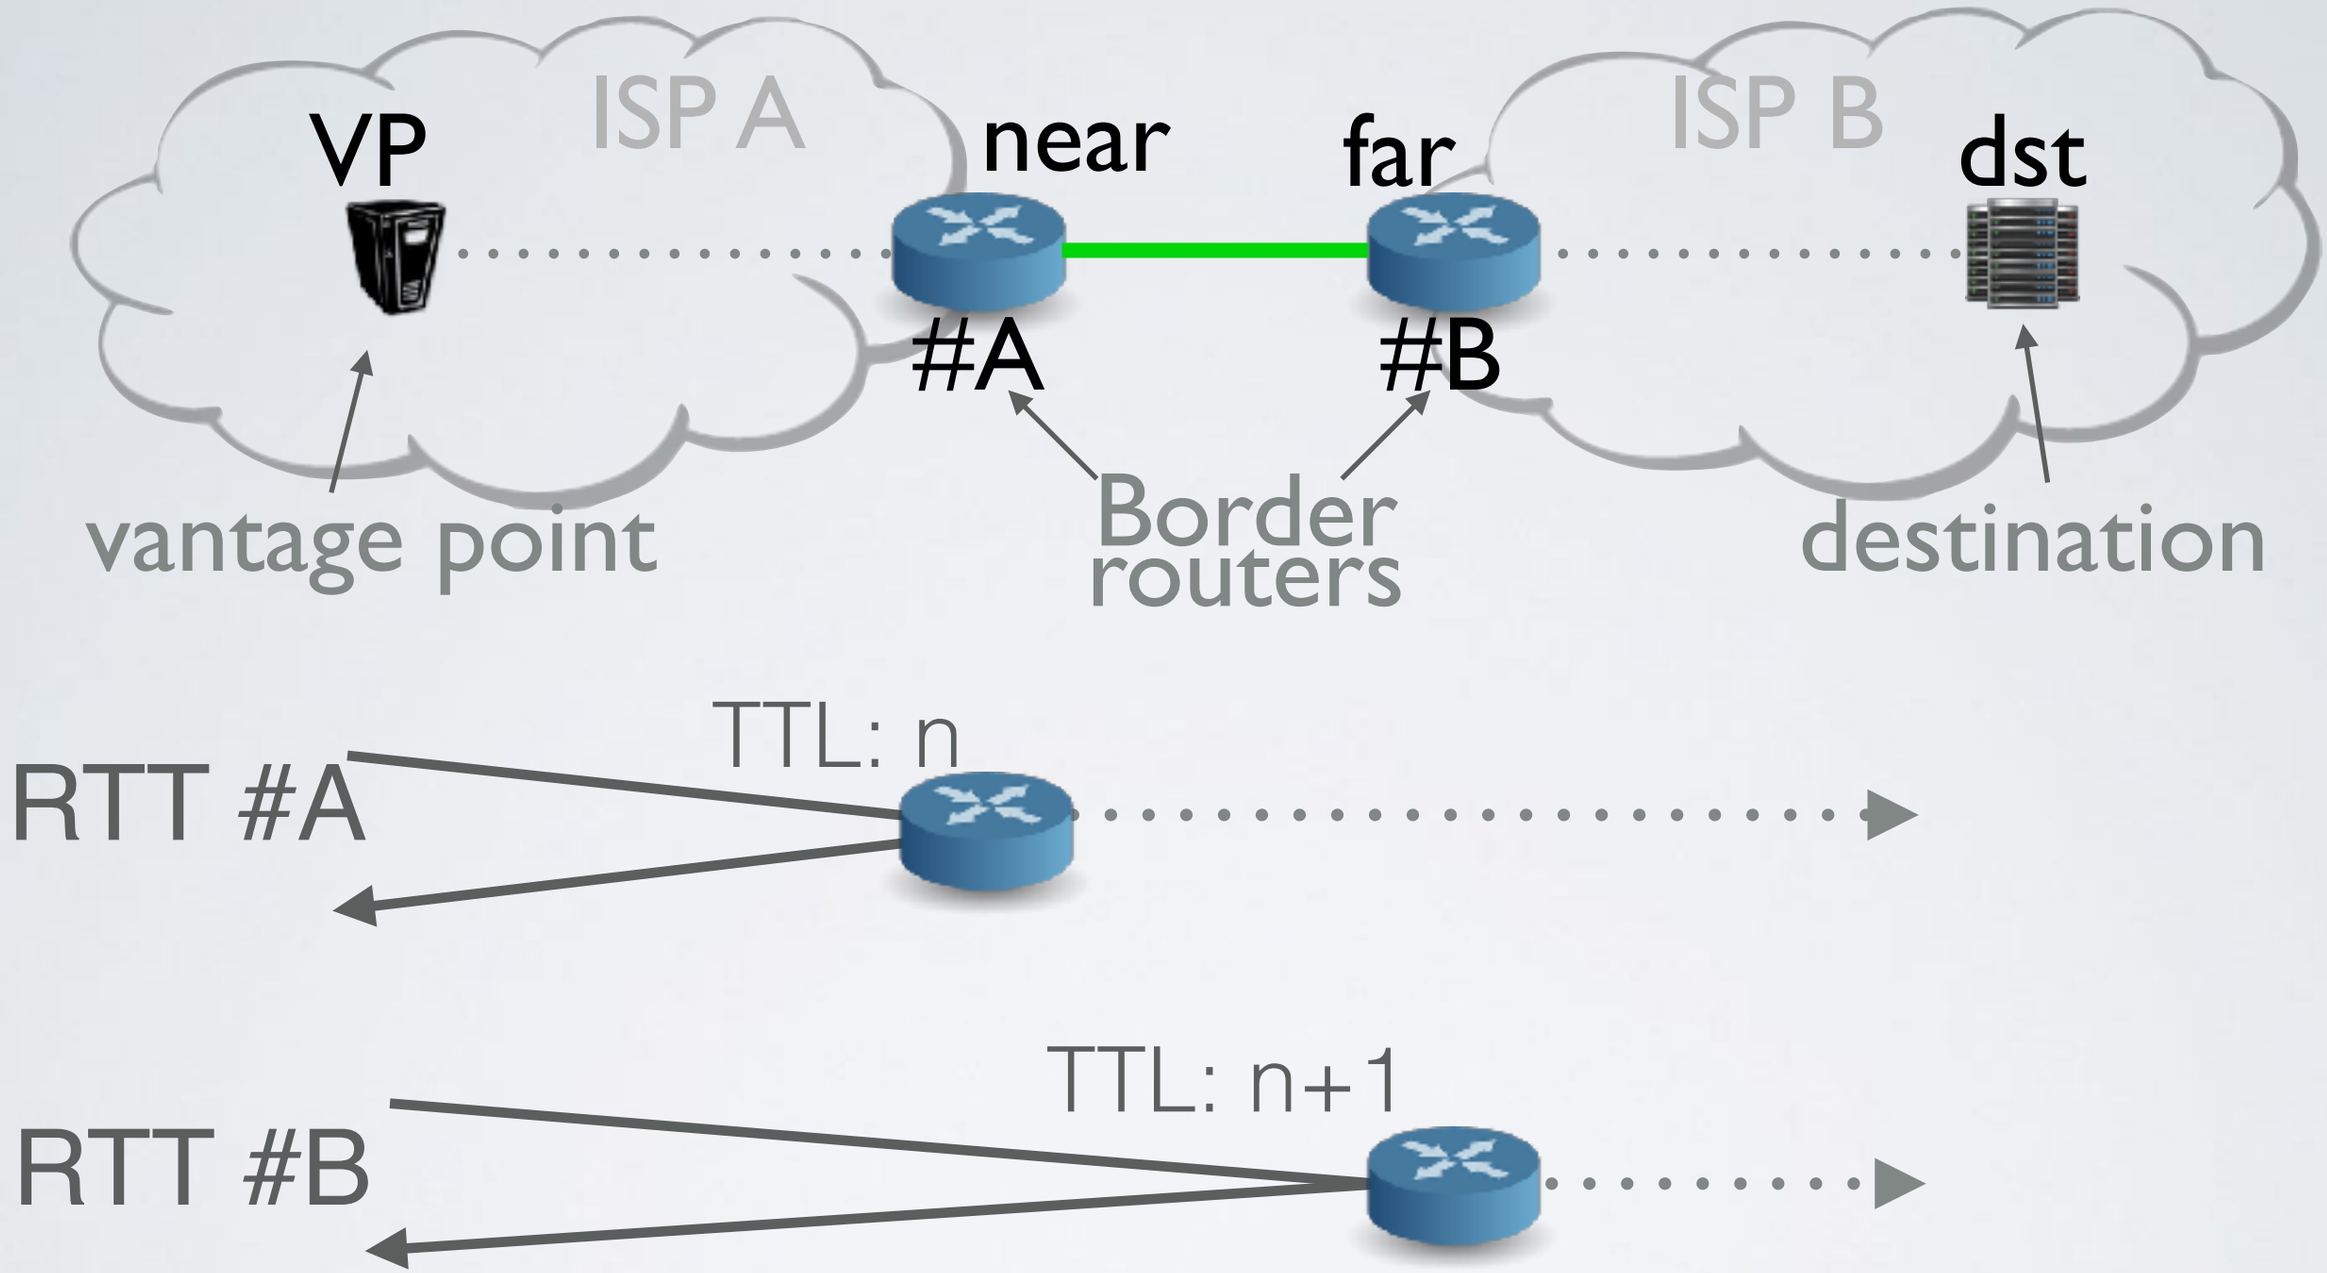

(repeat to obtain a time series)

caida

An Experiment with TSLP

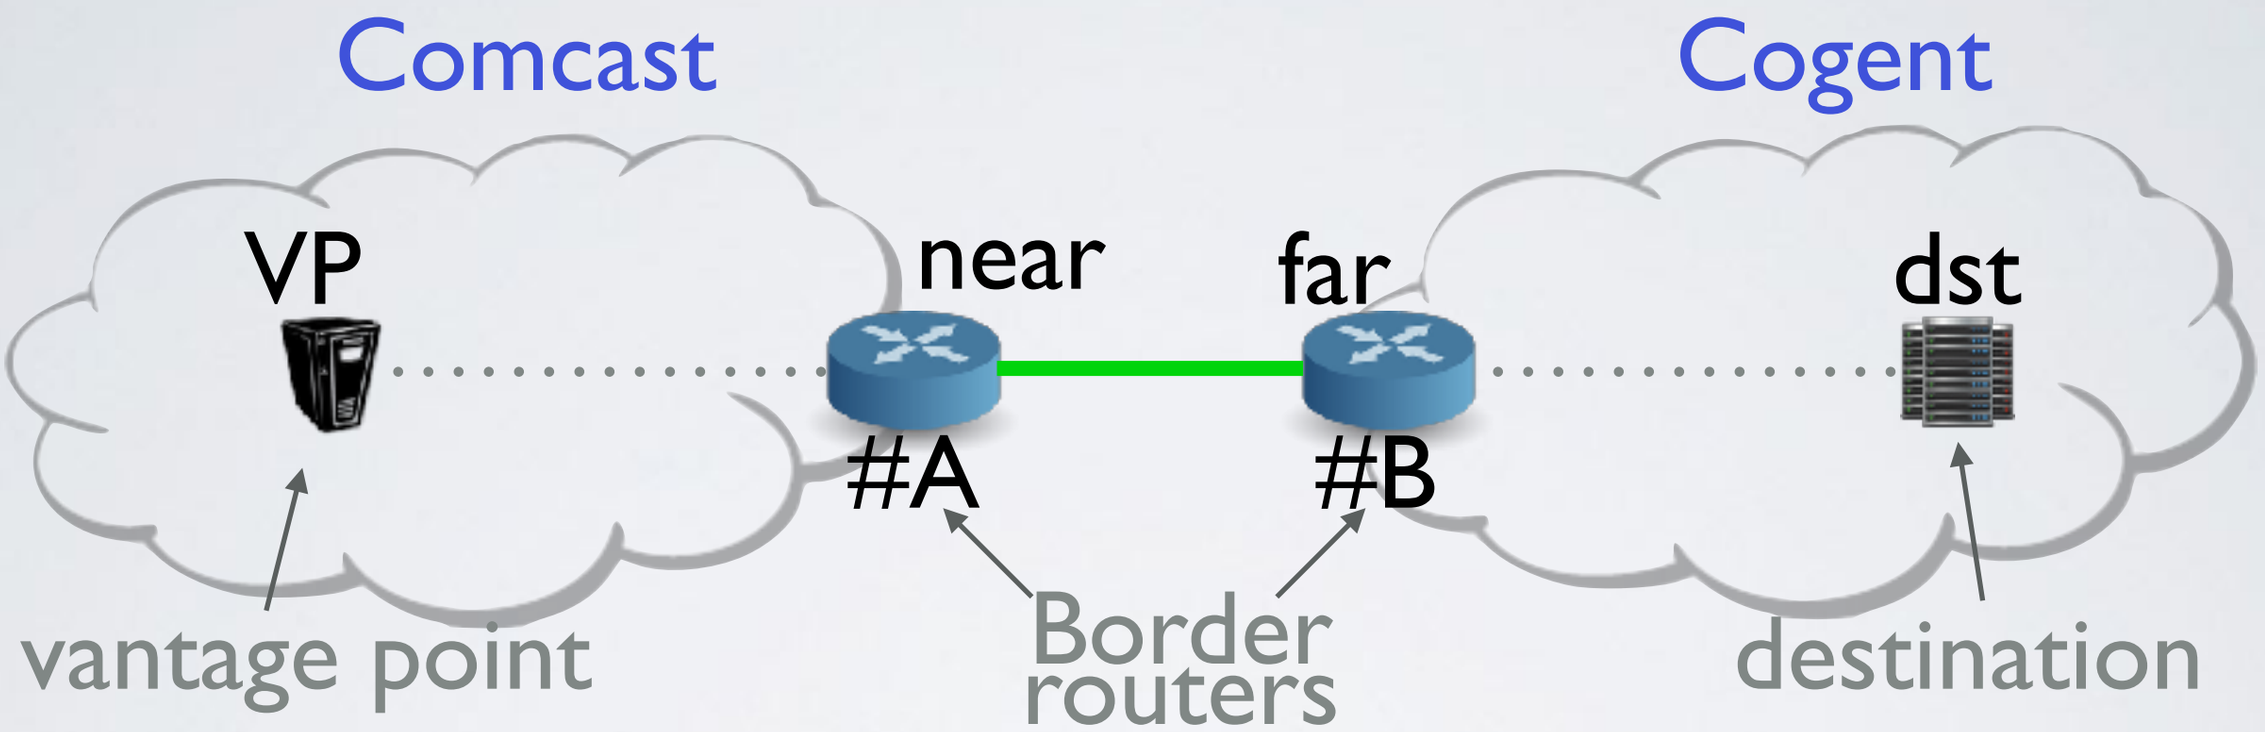

Measured interdomain link from Comcast to Cogent using VP in Comcast

caida

An Experiment with TSLP

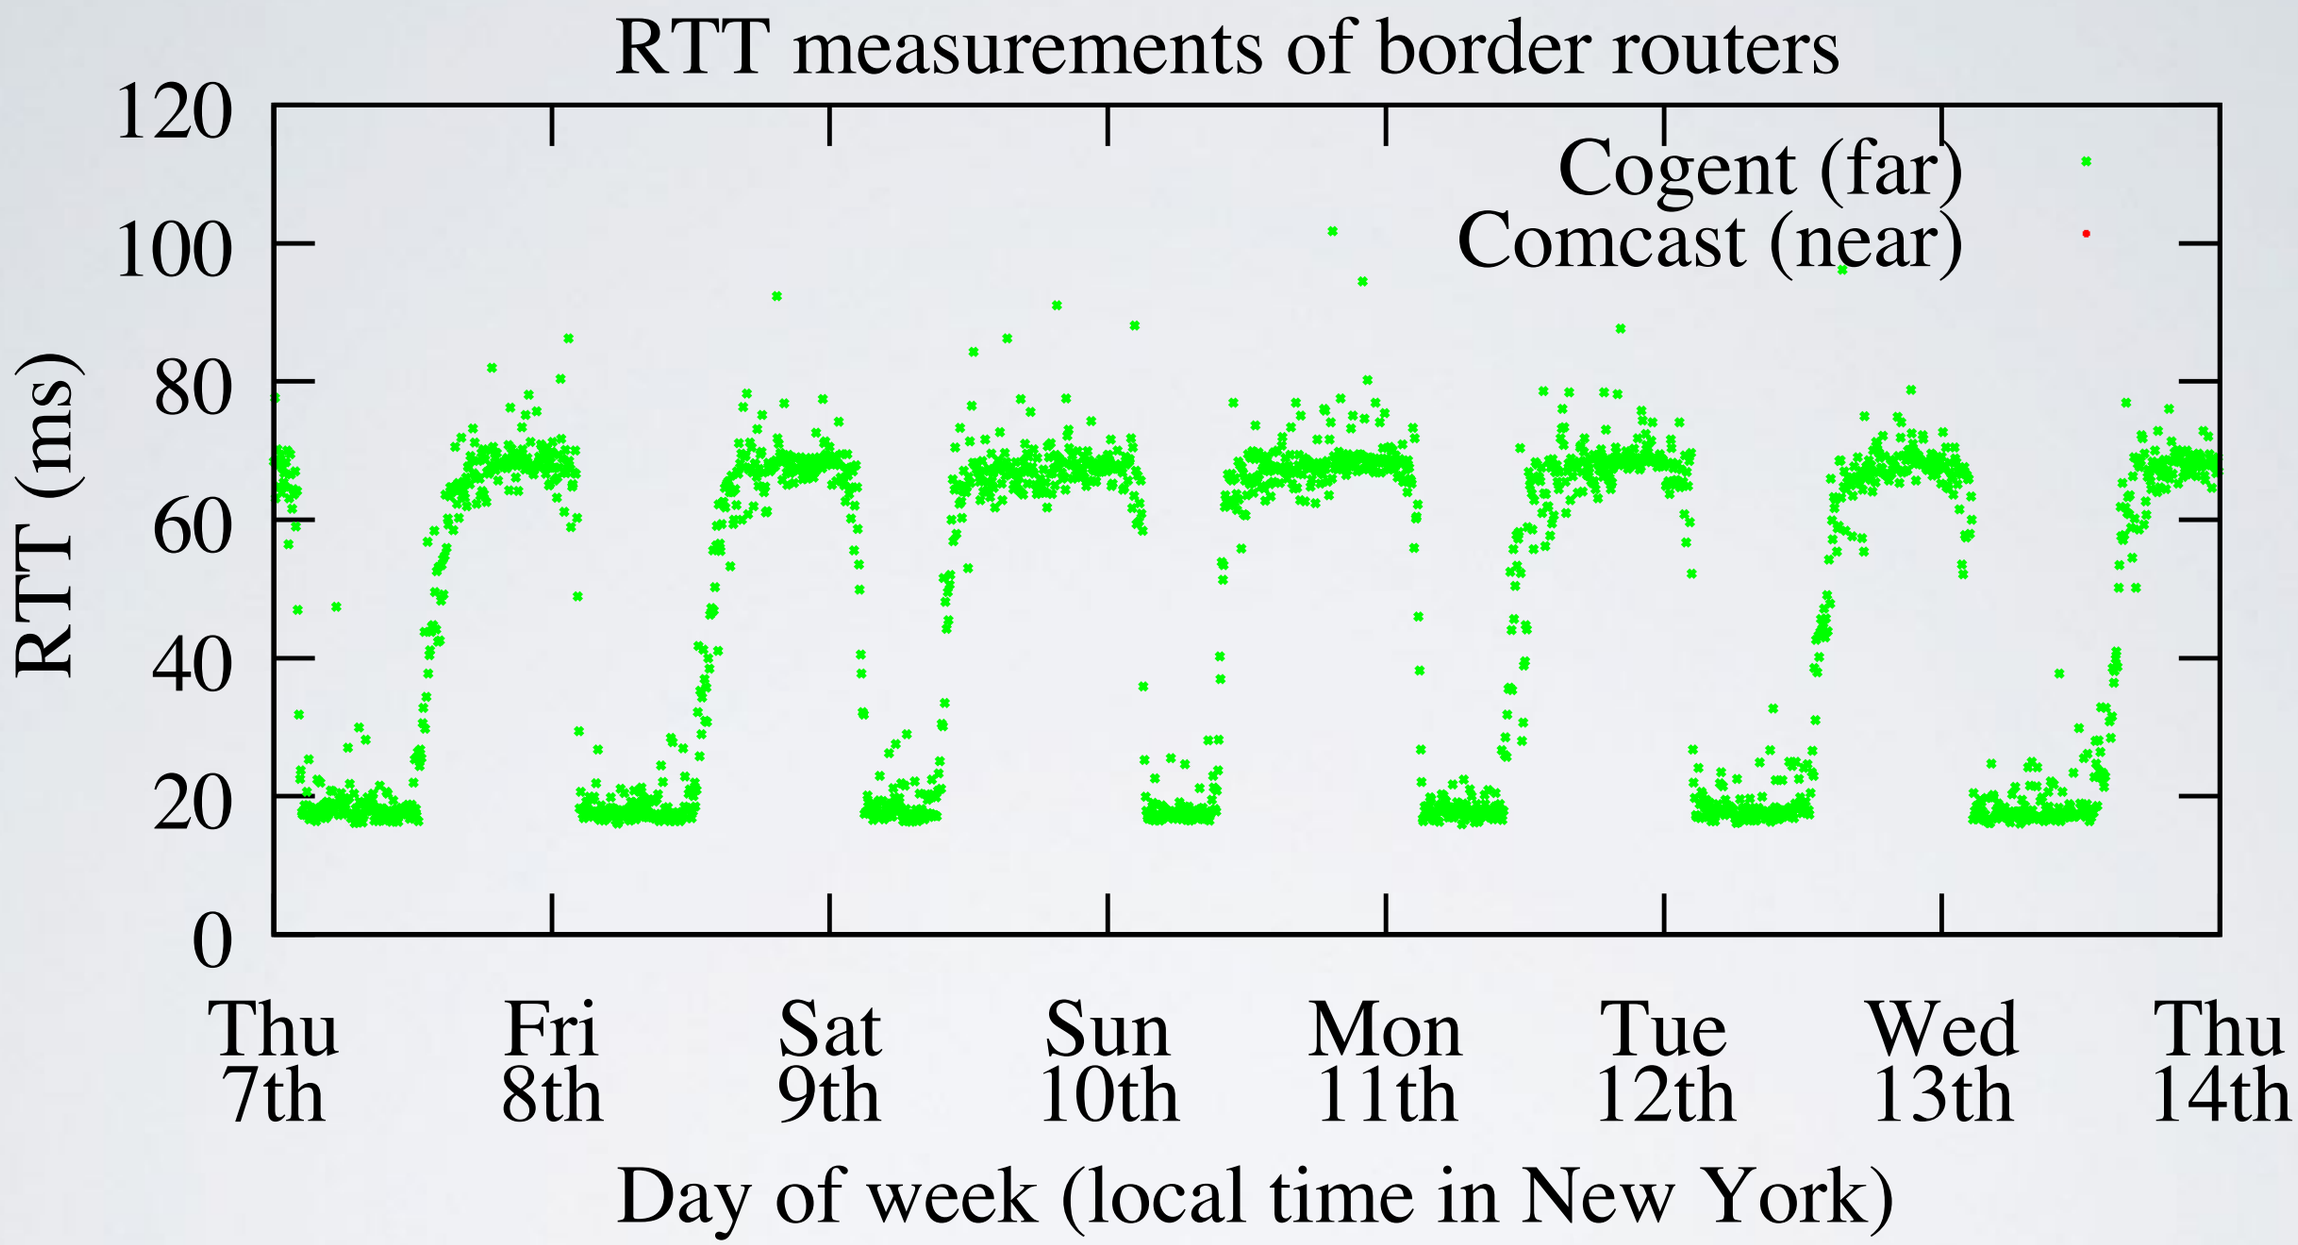

*Luckie, Dhamdhere, Clark, Huffaker, Claffy, "Challenges in Inferring Interdomain Congestion", IMC 2014

caida

An Experiment with TSLP

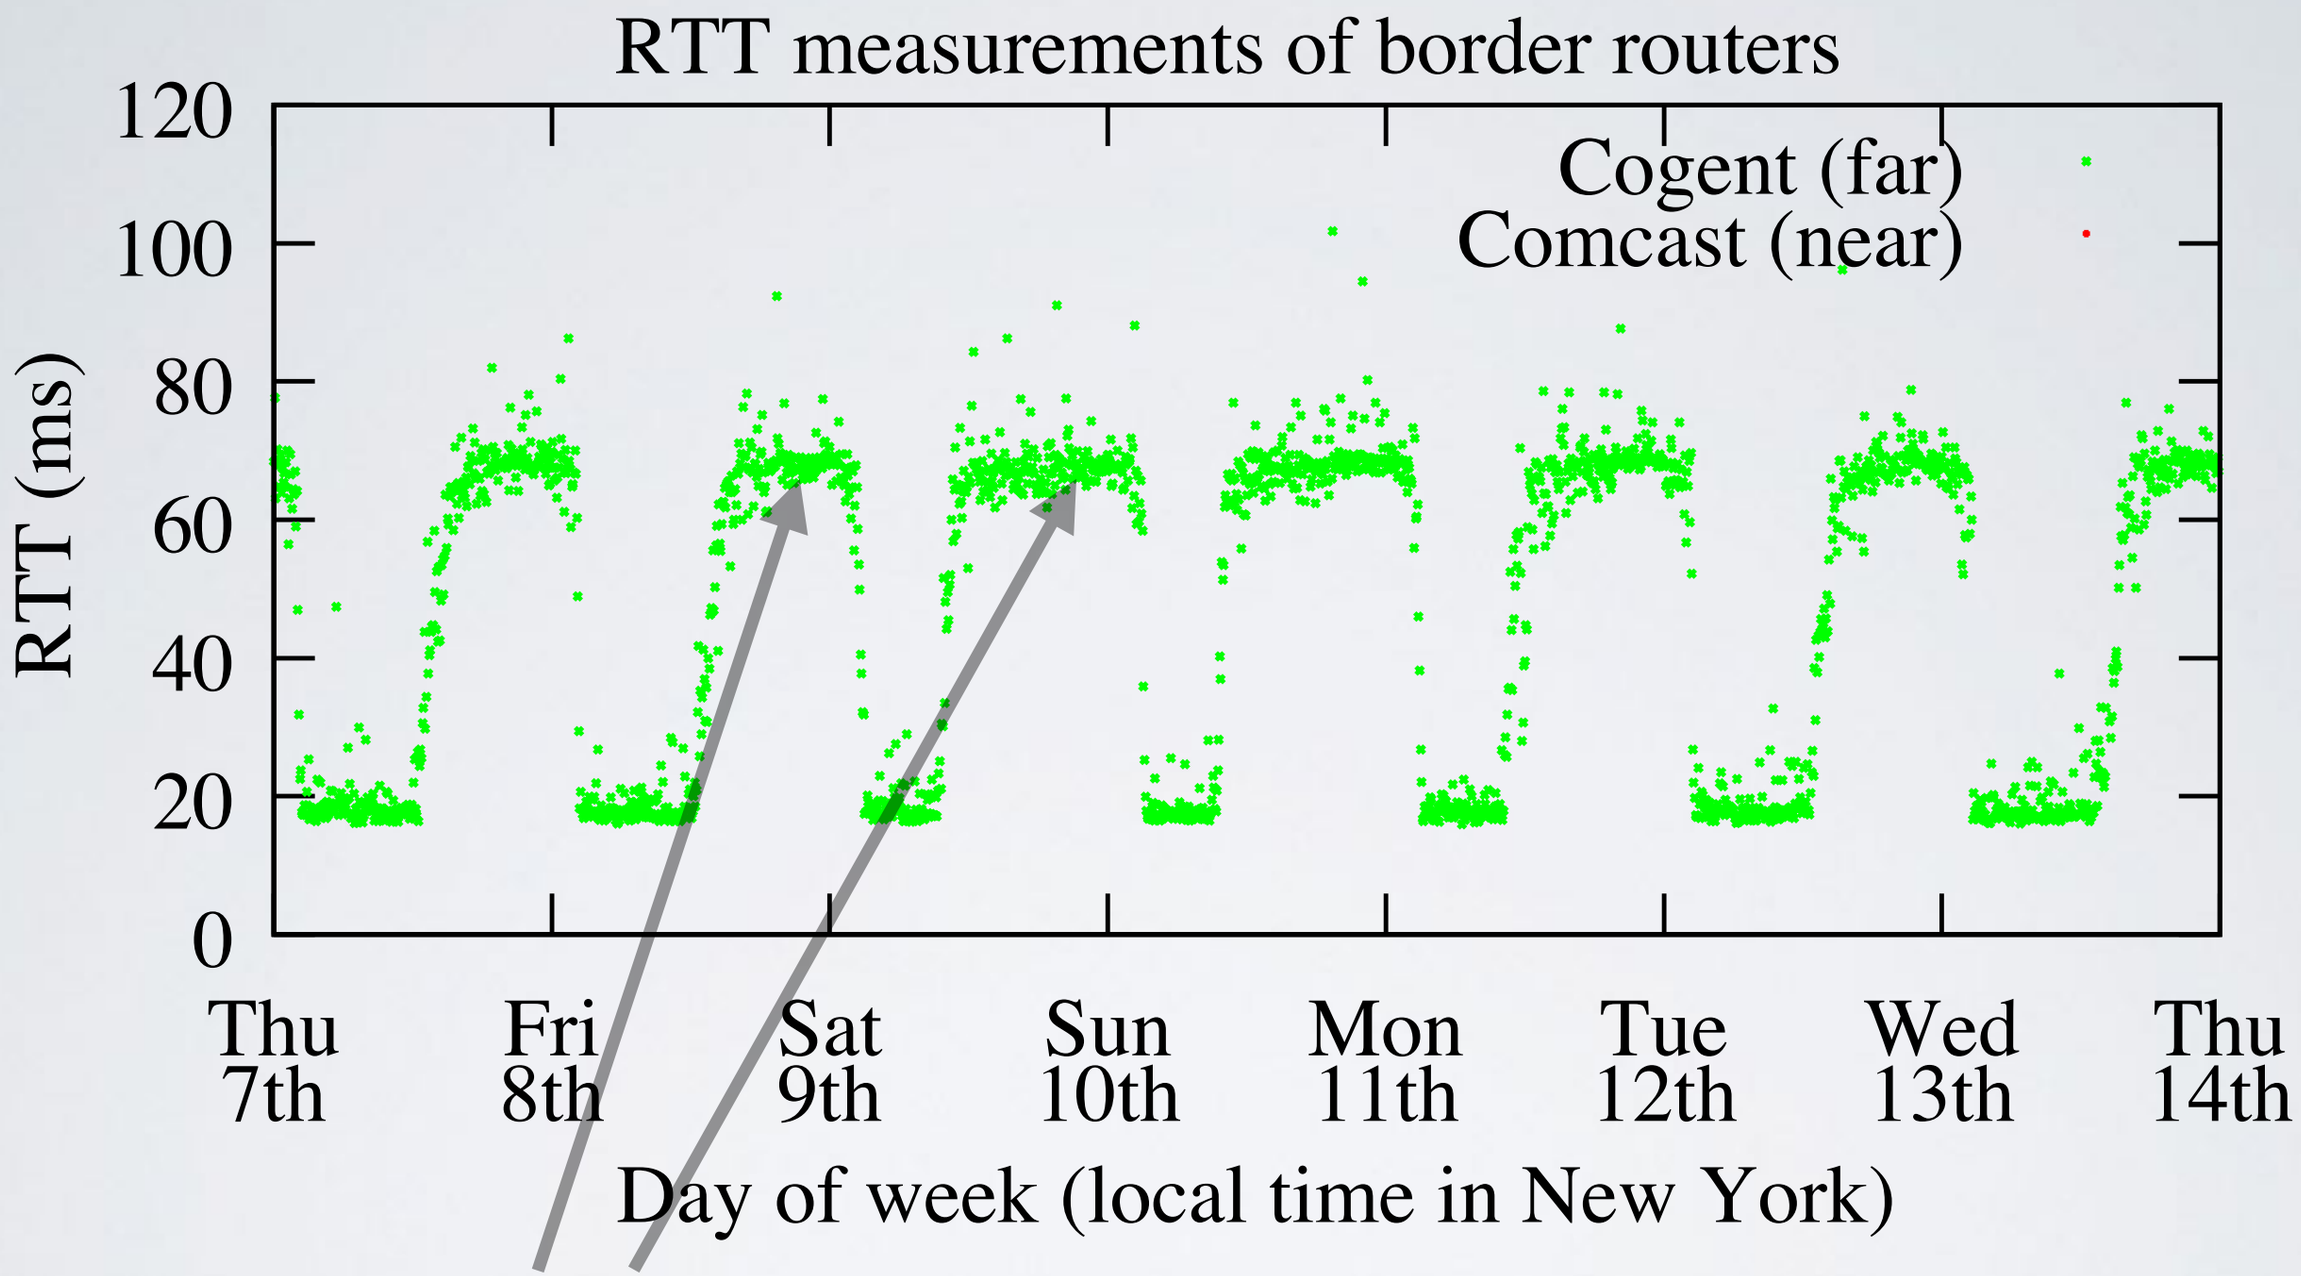

Diurnal elevation to far-side

*Luckie, Dhamdhere, Clark, Huffaker, Claffy, "Challenges in Inferring Interdomain Congestion", IMC 2014

caida

An Experiment with TSLP

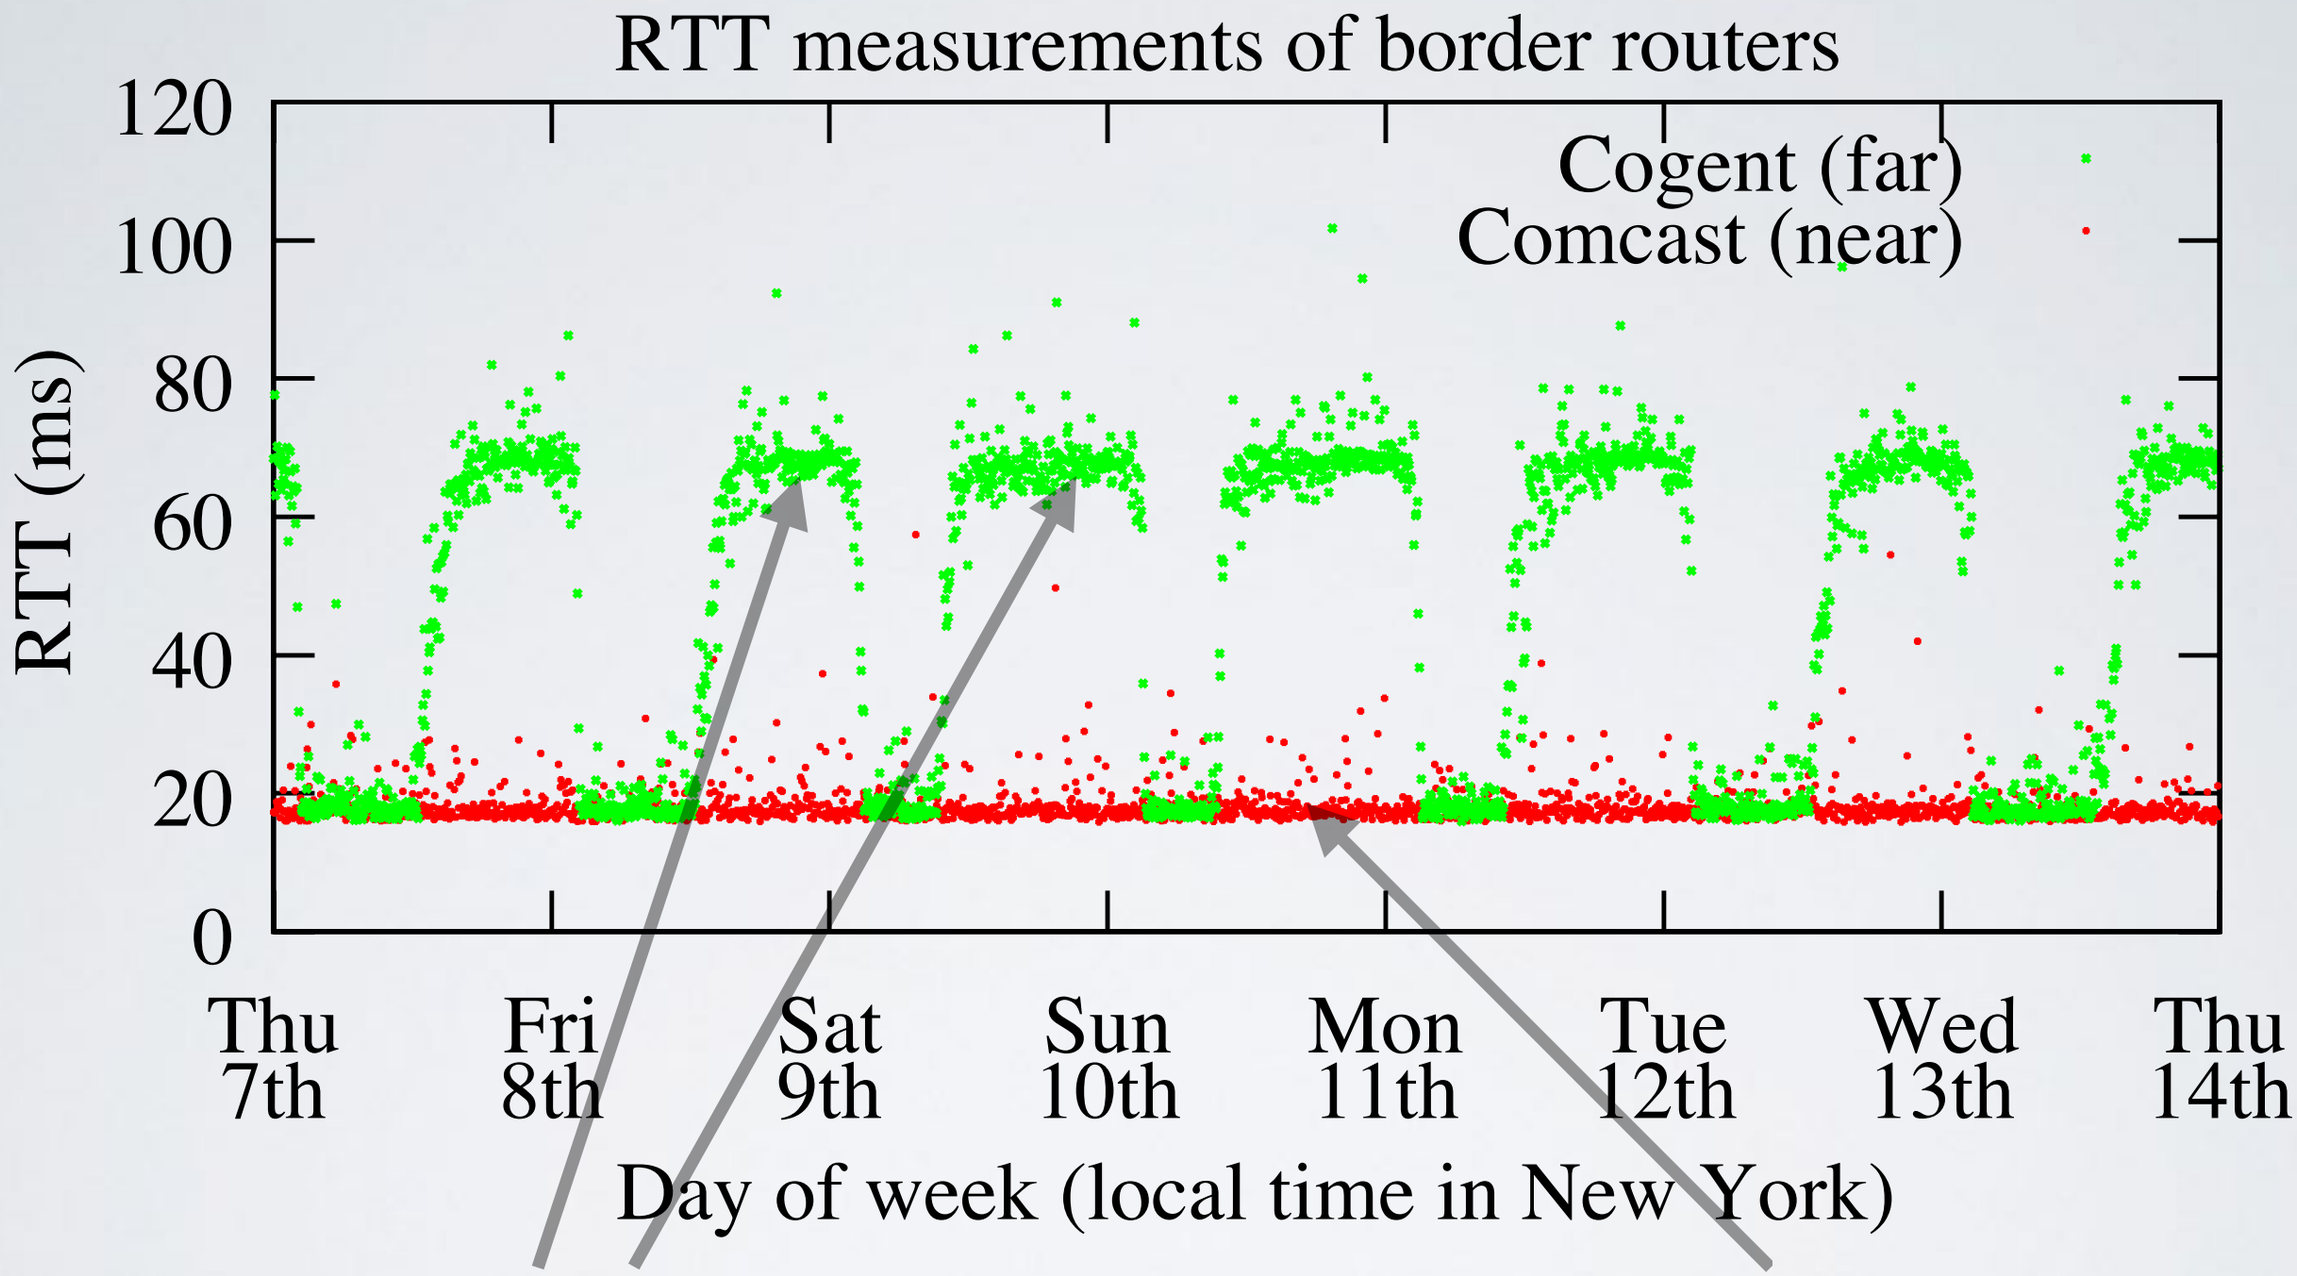

*Luckie, Dhamdhere, Clark, Huffaker, Claffy, "Challenges in Inferring Interdomain Congestion", IMC 2014

caida

Pieces of the Puzzle

caida

Pieces of the Puzzle

Pieces of the Puzzle

- **Interdomain link Identification:** Need to identify interdomain links to be able to probe them

caida

Pieces of the Puzzle

- **Interdomain link Identification:** Need to identify interdomain links to be able to probe them
- **Adaptive Probing:** Need to be adaptive to changes in the underlying topology and routing

Pieces of the Puzzle

- **Interdomain link Identification:** Need to identify interdomain links to be able to probe them
- **Adaptive Probing:** Need to be adaptive to changes in the underlying topology and routing
- **Identifying Congested Links:** Need time-series analysis techniques to find patterns in (noisy) data that indicate congestion

Pieces of the Puzzle

- **Interdomain link Identification:** Need to identify interdomain links to be able to probe them
- **Adaptive Probing:** Need to be adaptive to changes in the underlying topology and routing
- **Identifying Congested Links:** Need time-series analysis techniques to find patterns in (noisy) data that indicate congestion
- **Validation:** Need to validate inferences. Most peering agreements are covered by NDAs

MANIC System

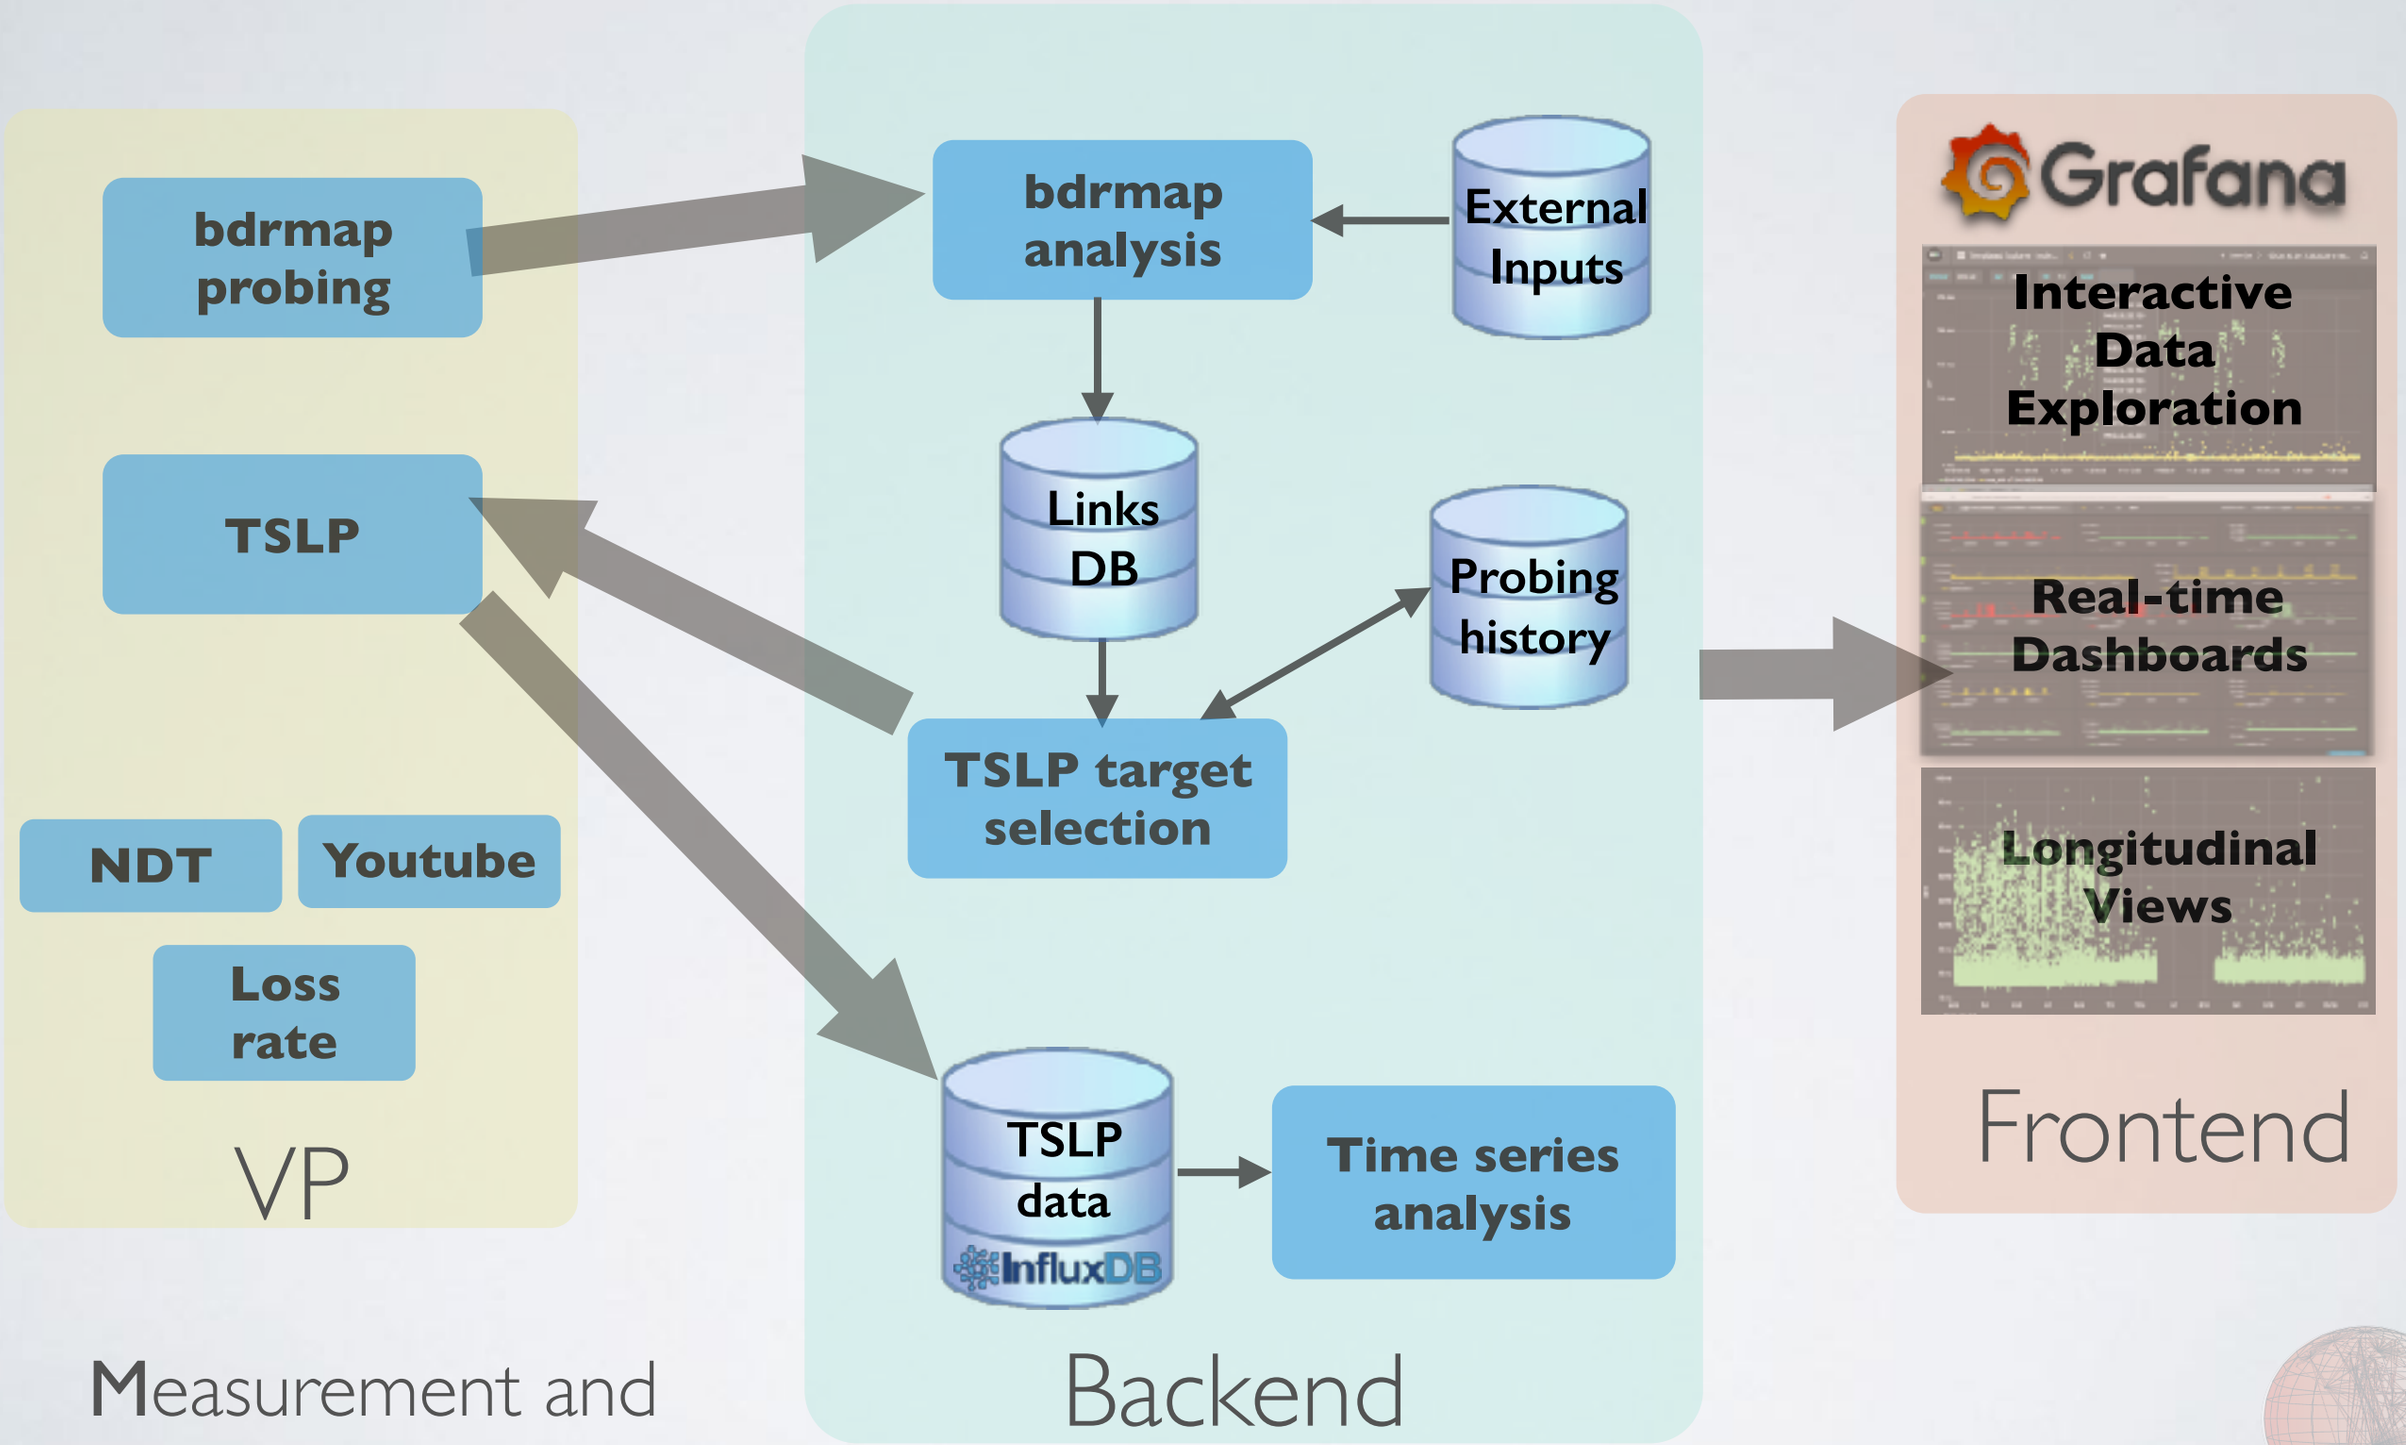

Measurement and ANalysis of Internet Congestion

caida

MANIC System

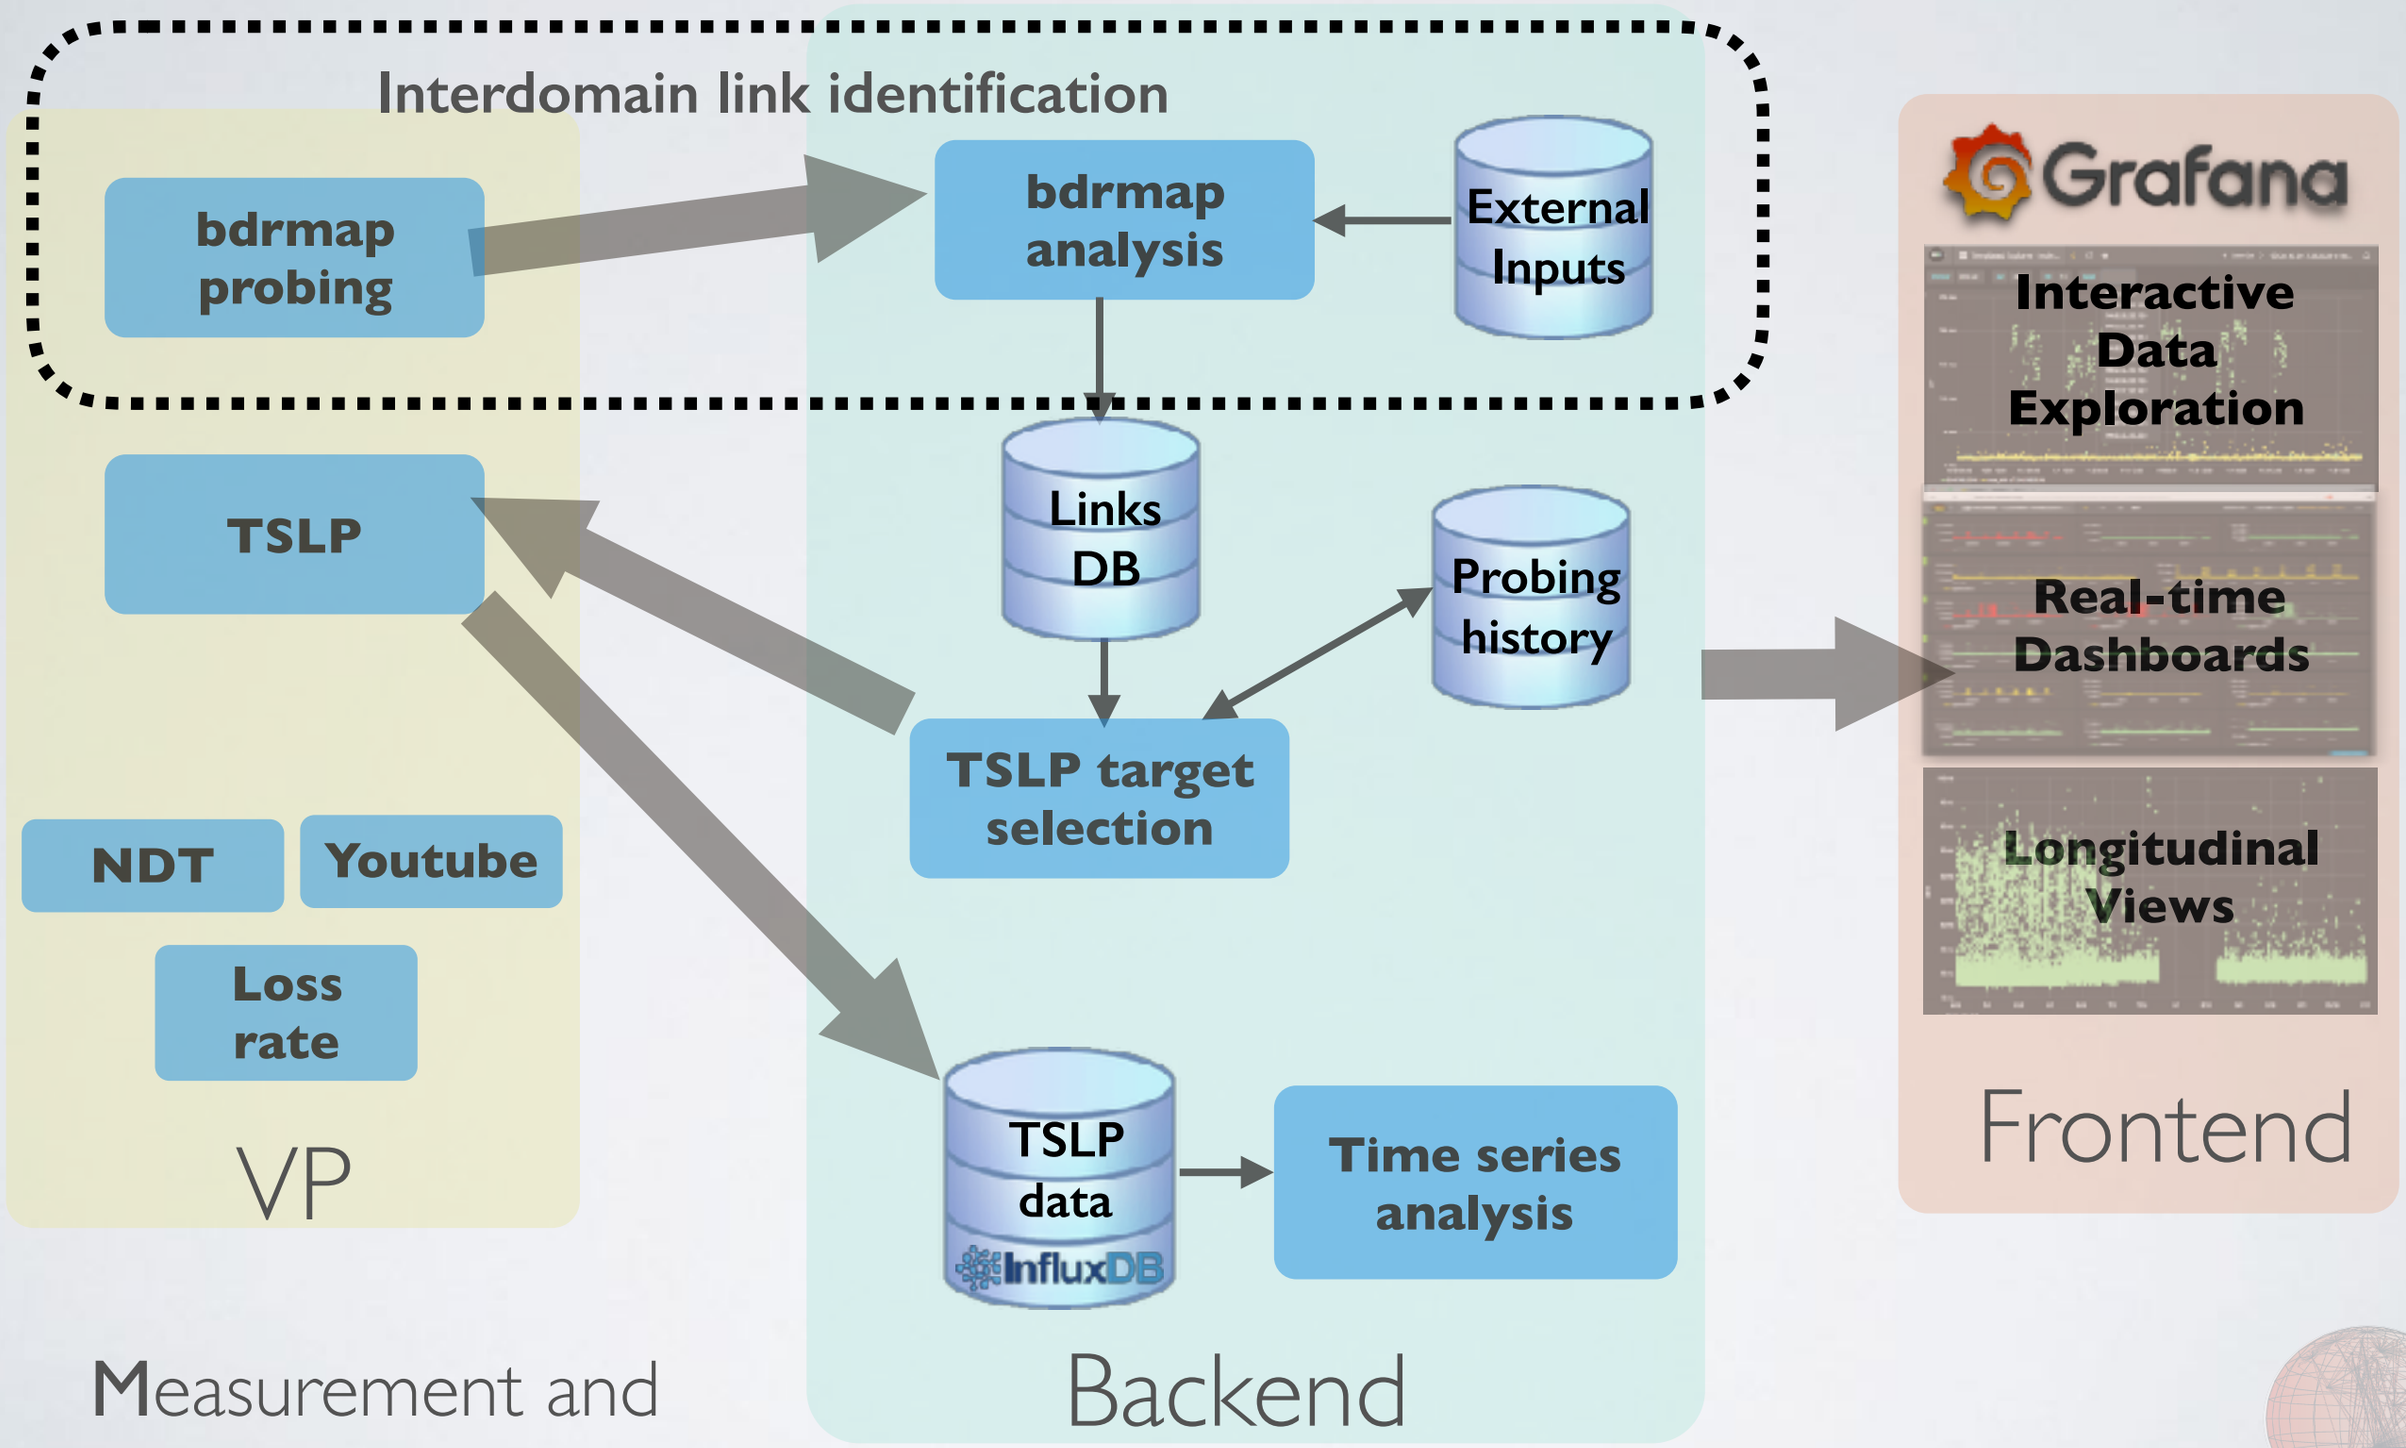

Measurement and ANalysis of Internet Congestion

caida

MANIC System

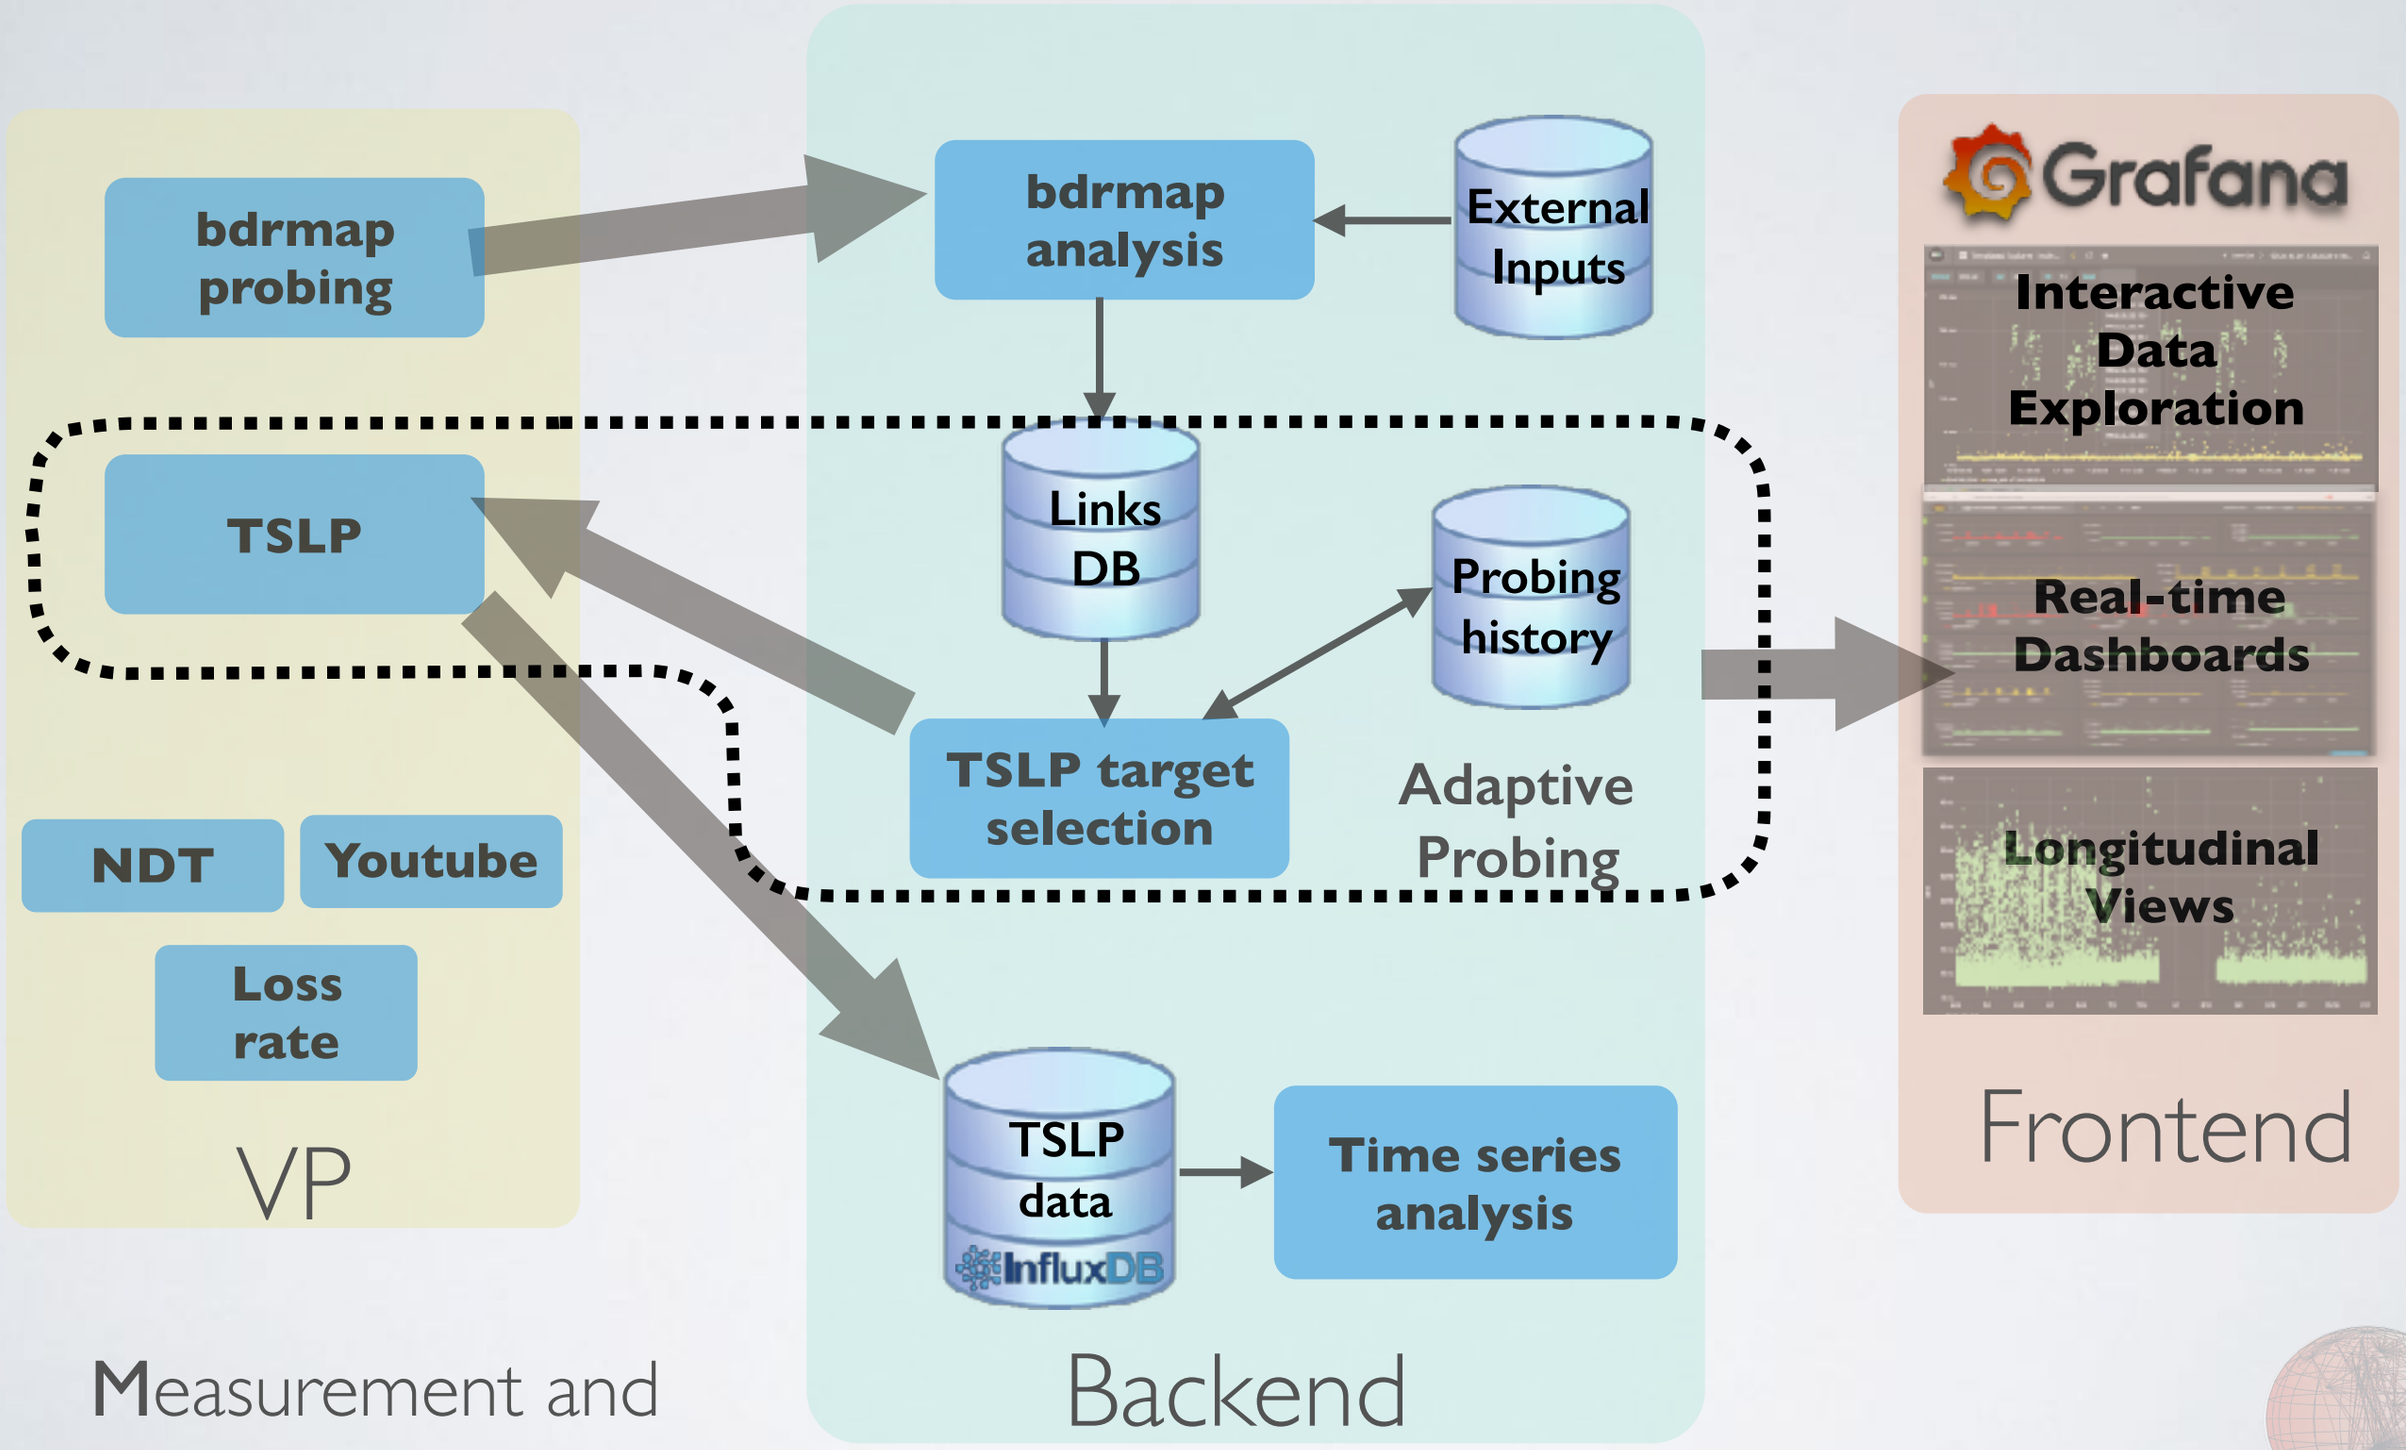

Measurement and ANalysis of Internet Congestion

caida

MANIC System

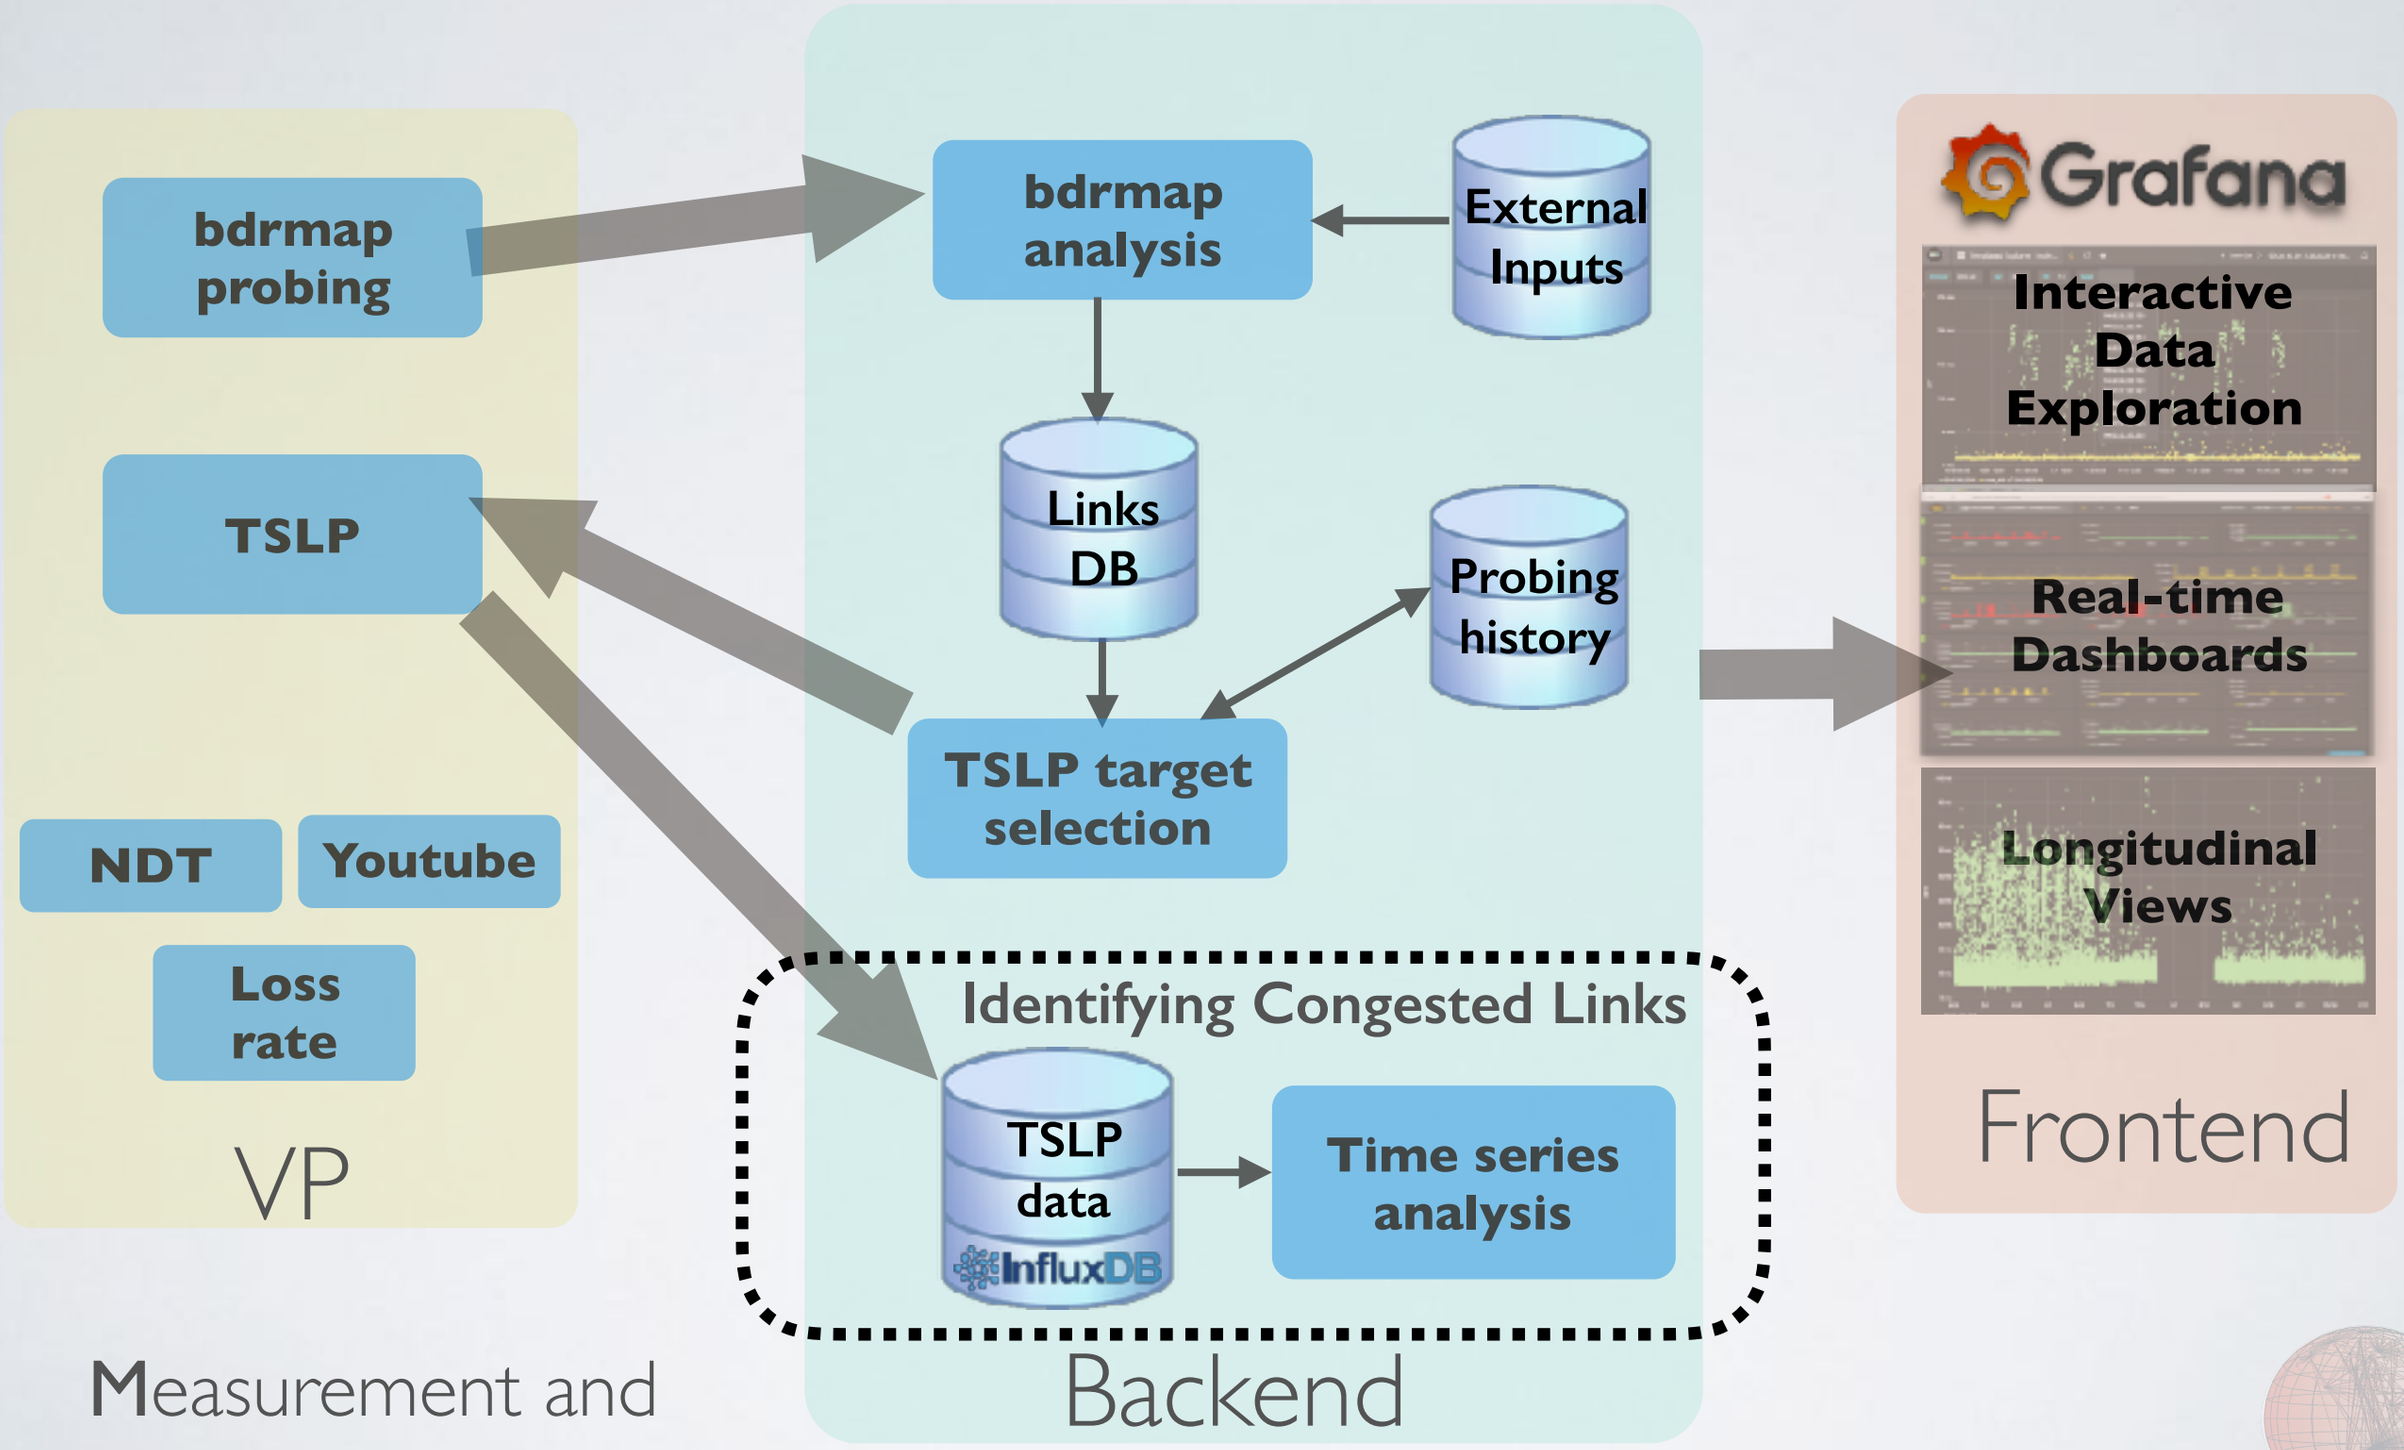

Measurement and ANalysis of Internet Congestion

caida

MANIC System

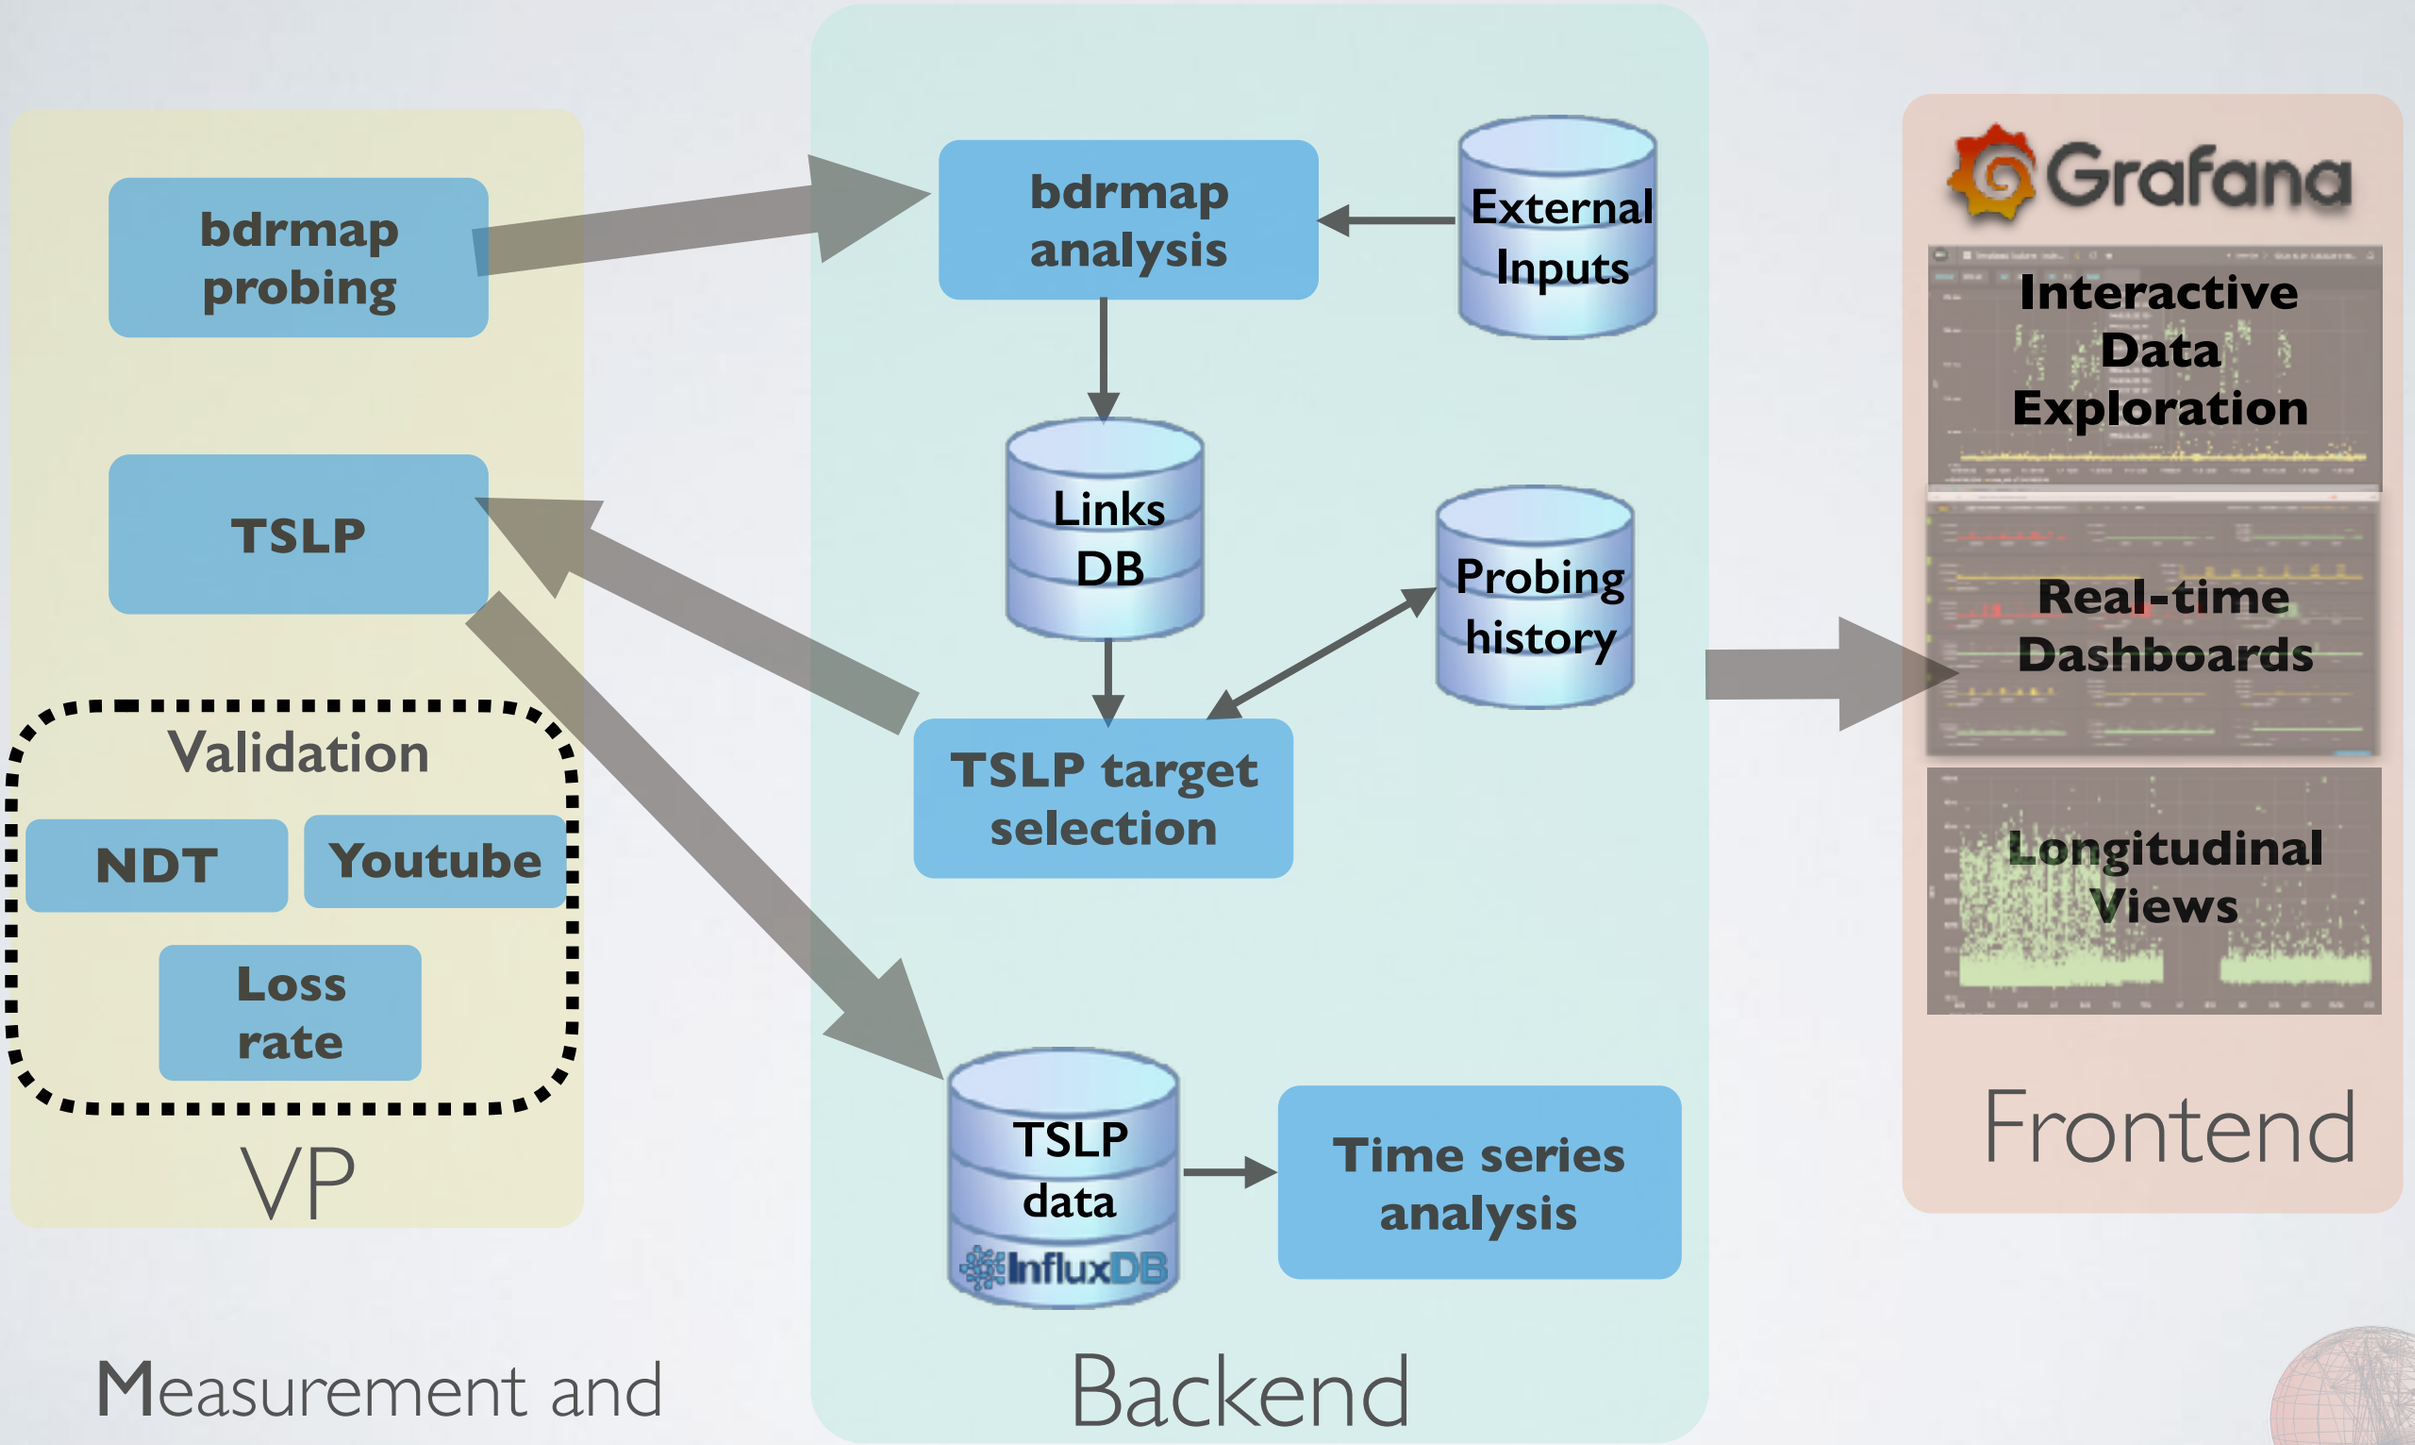

Measurement and ANalysis of Internet Congestion

caida

MANIC System

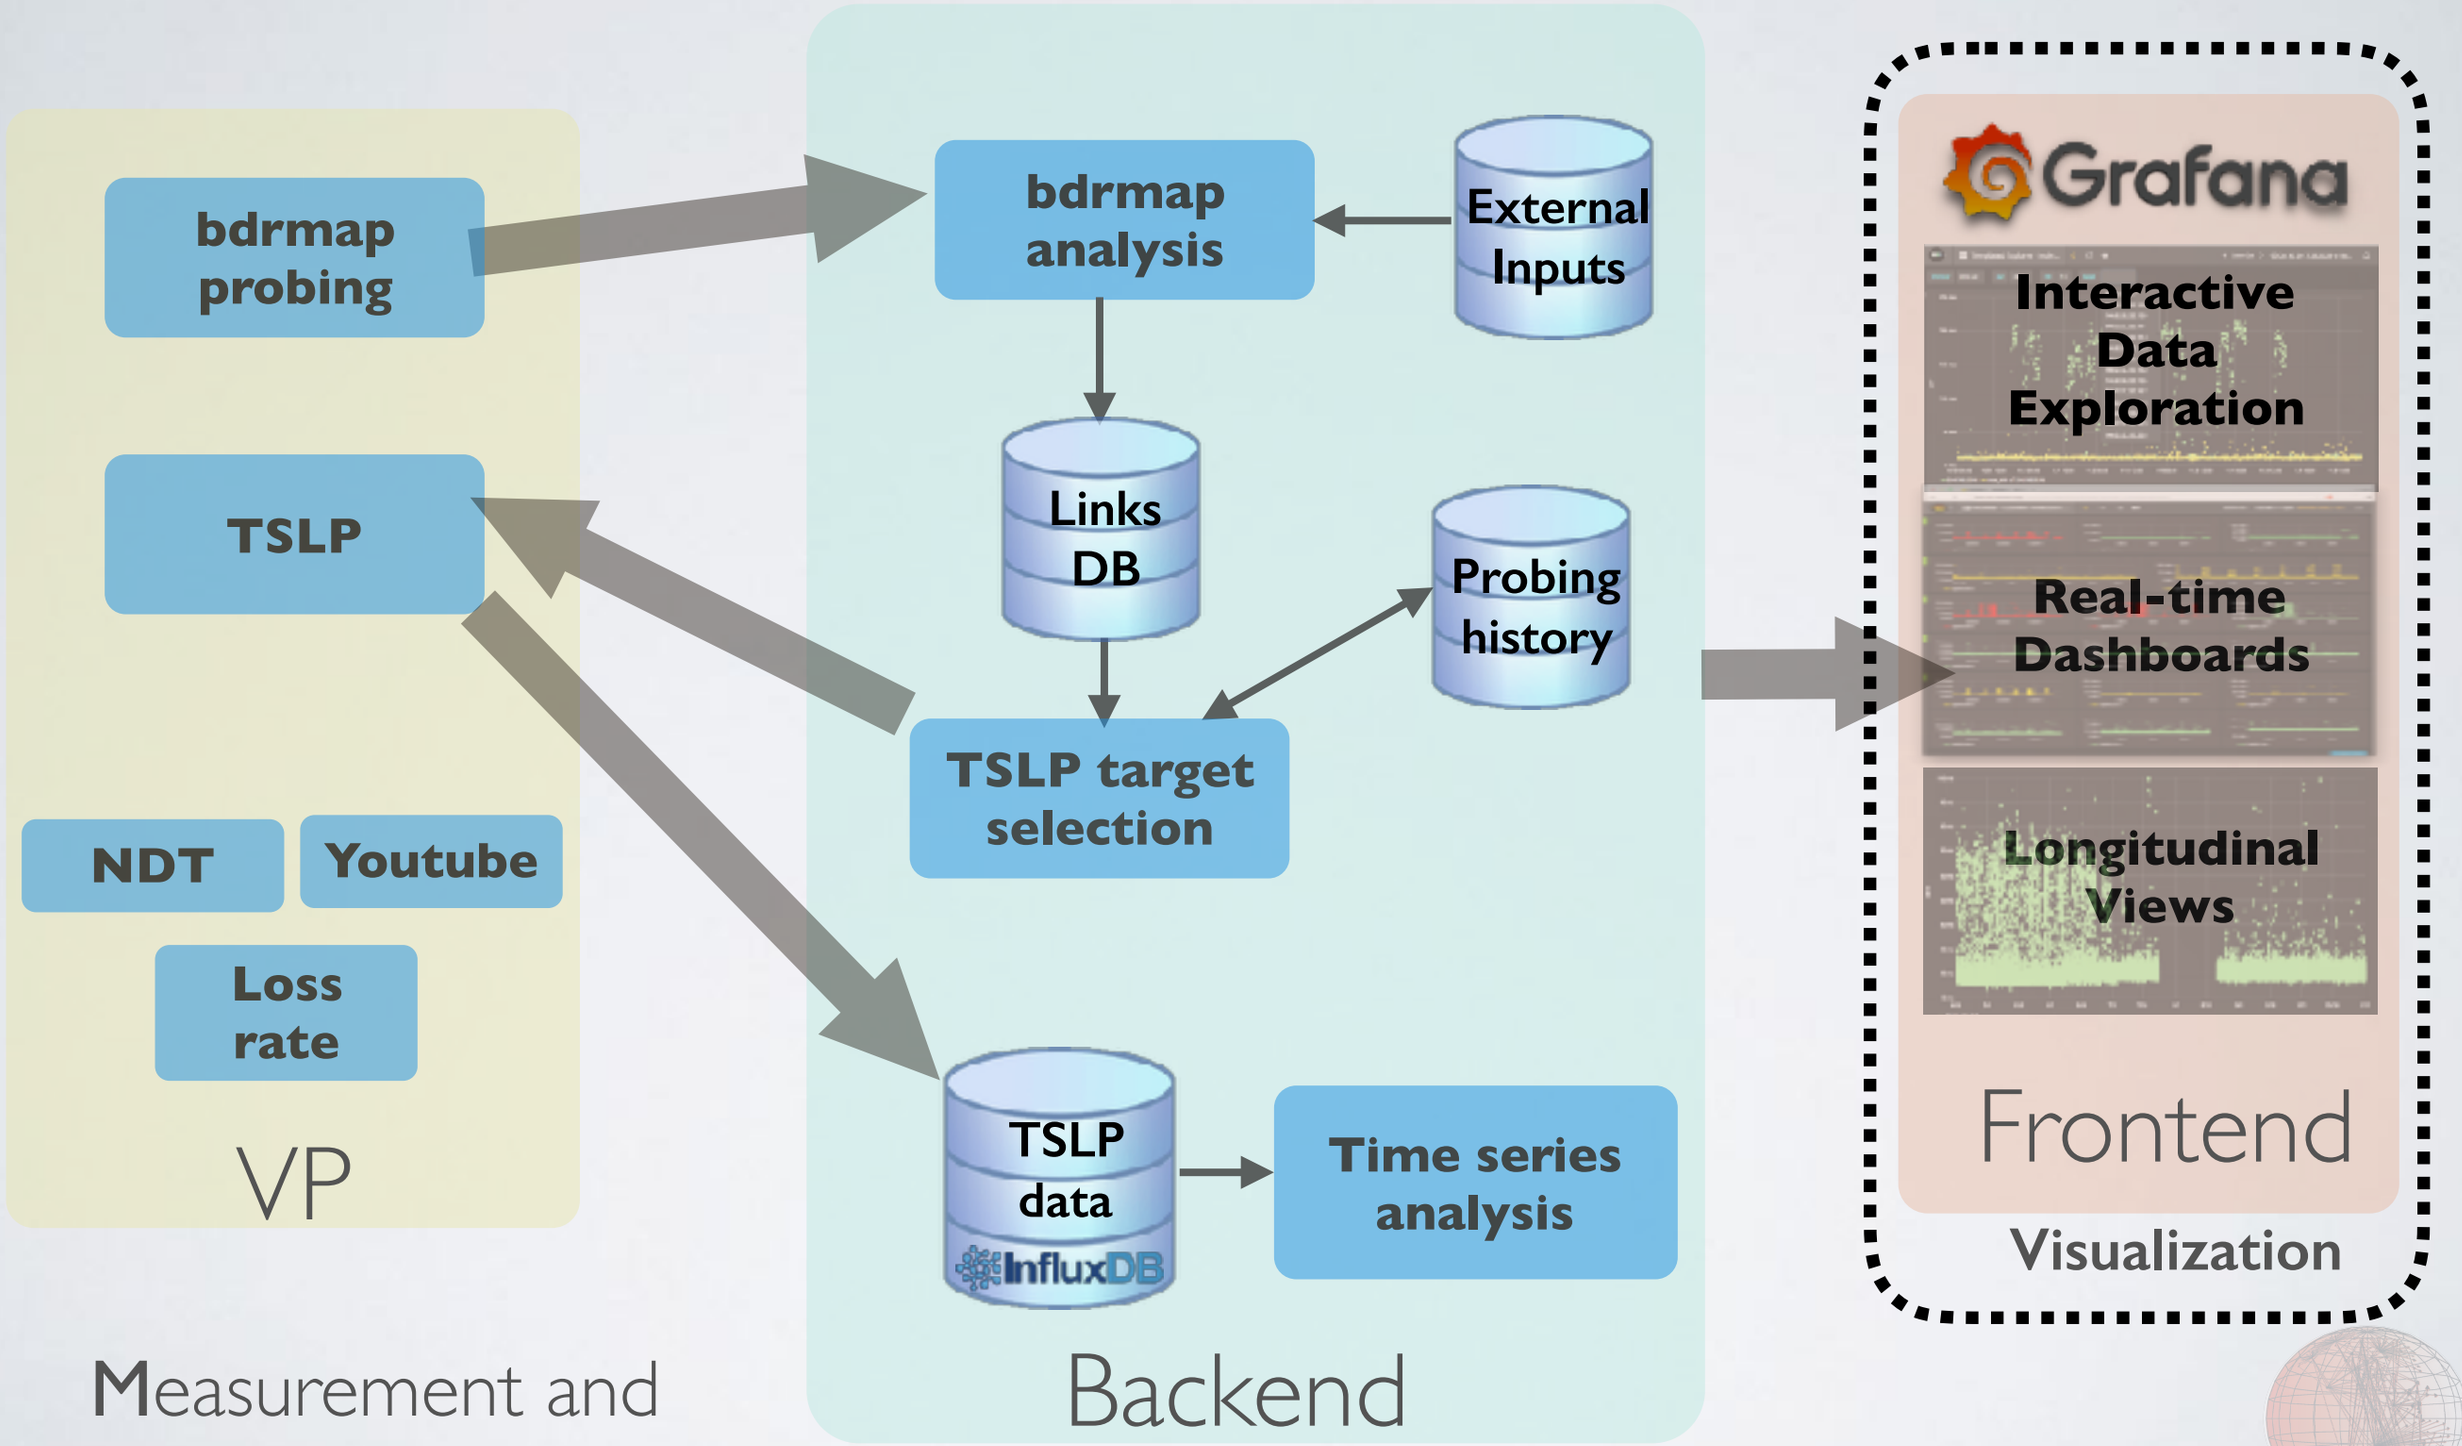

Measurement and ANalysis of Internet Congestion

caida

Identifying Interdomain Links

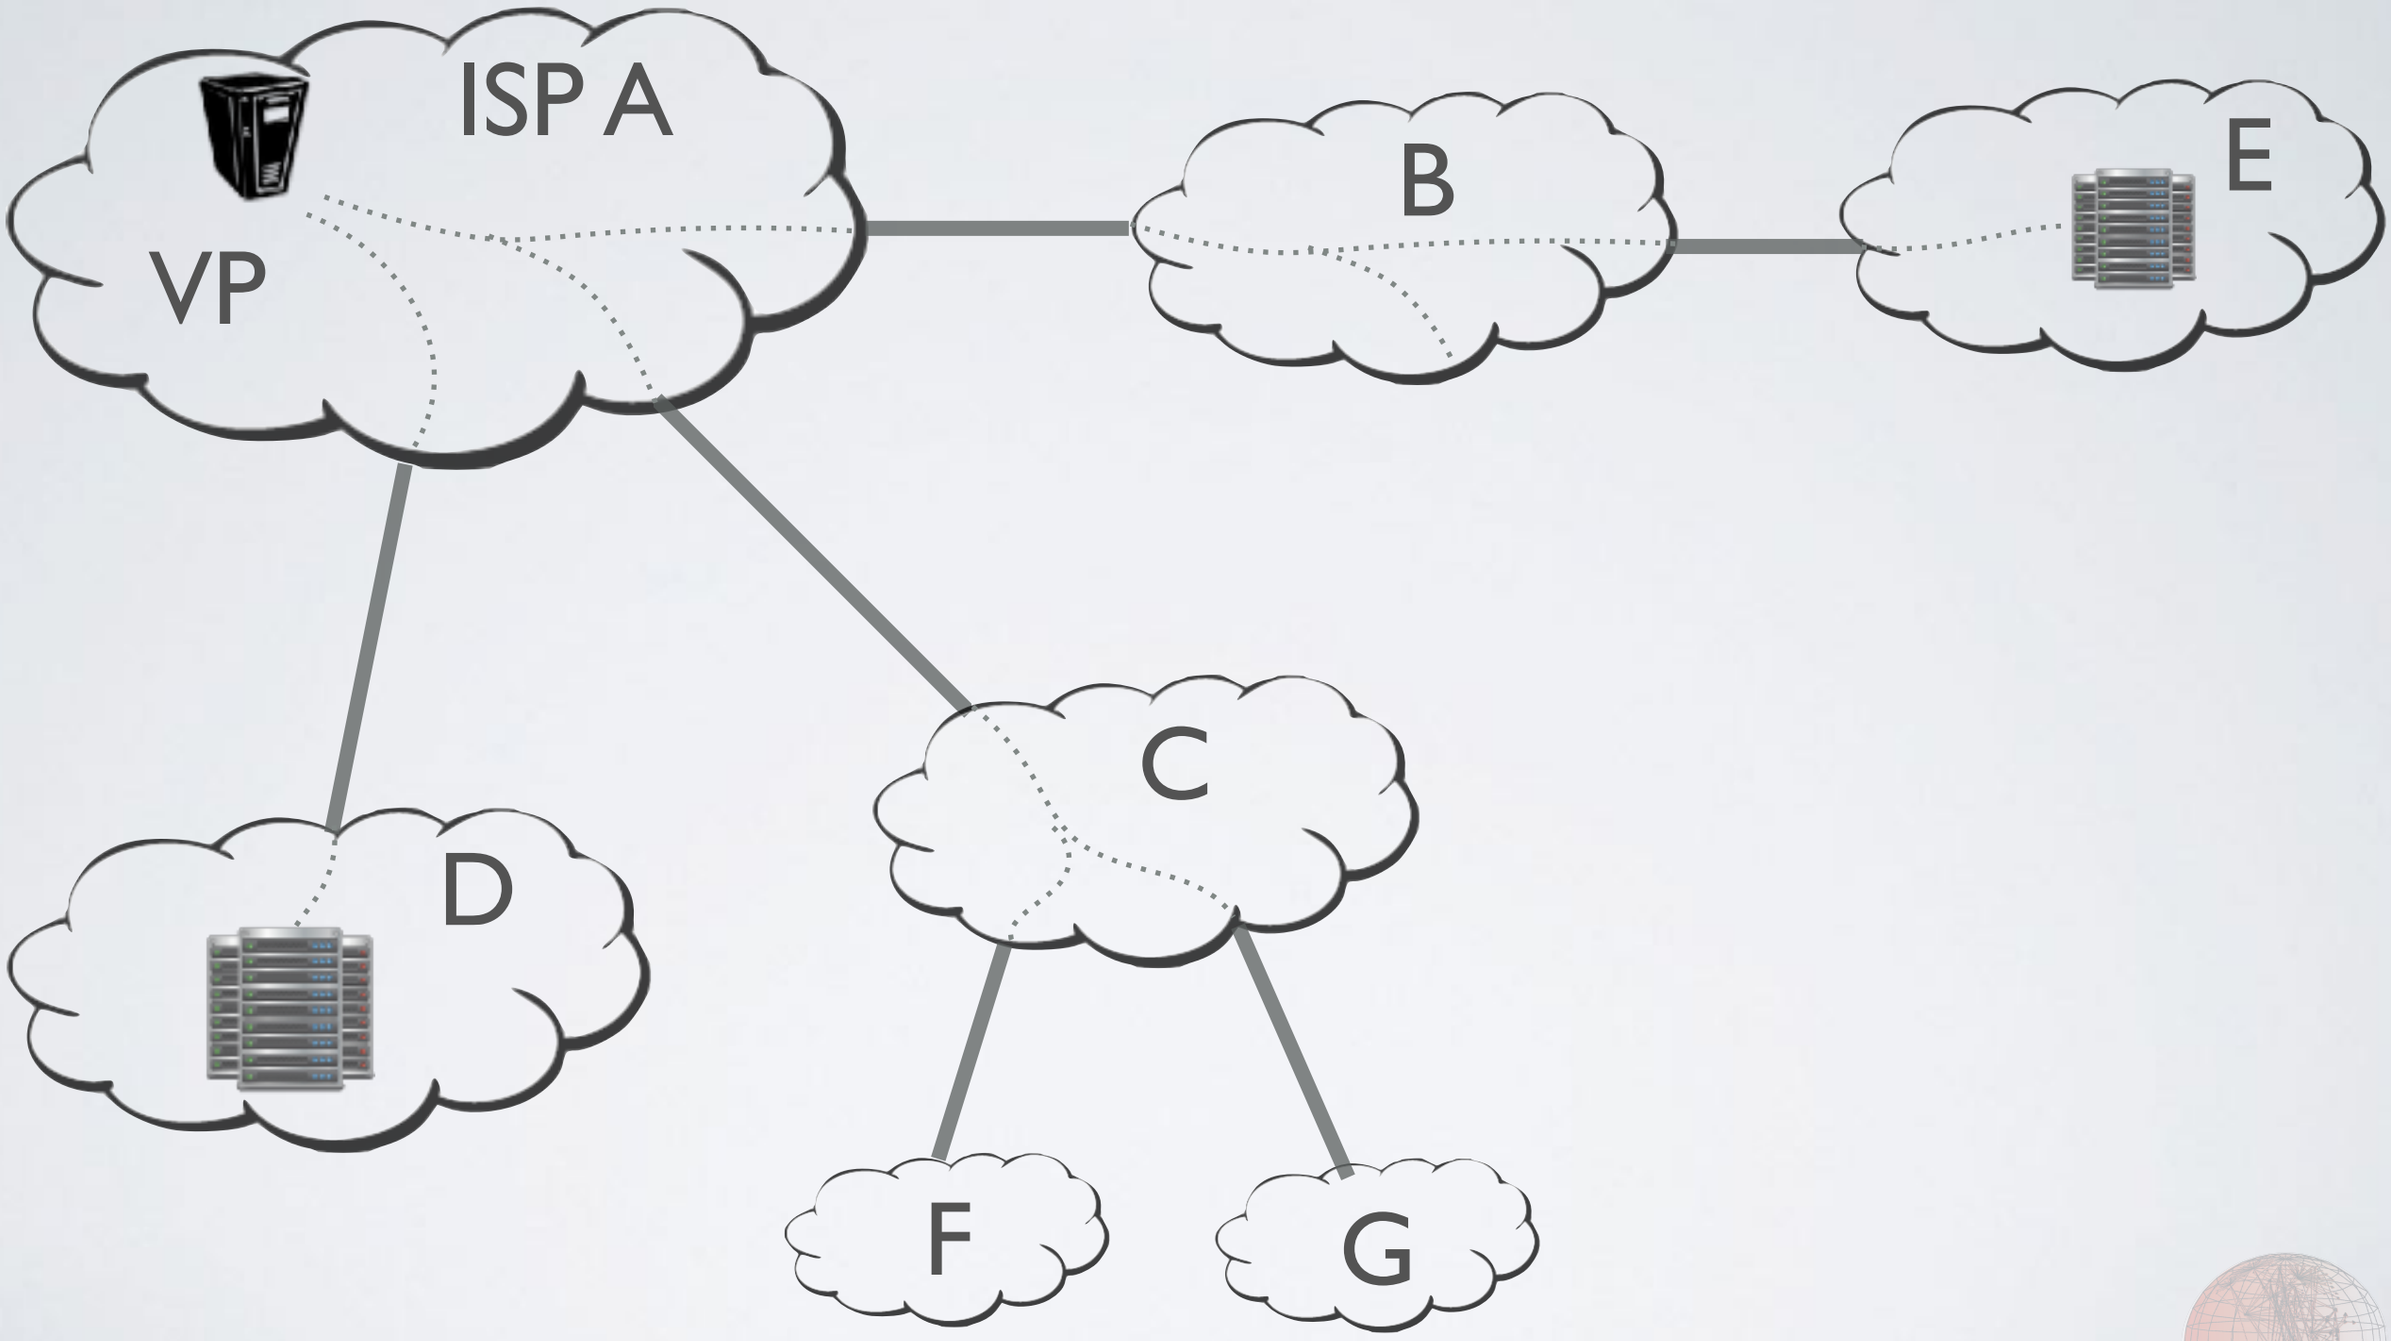

caida

Identifying Interdomain Links

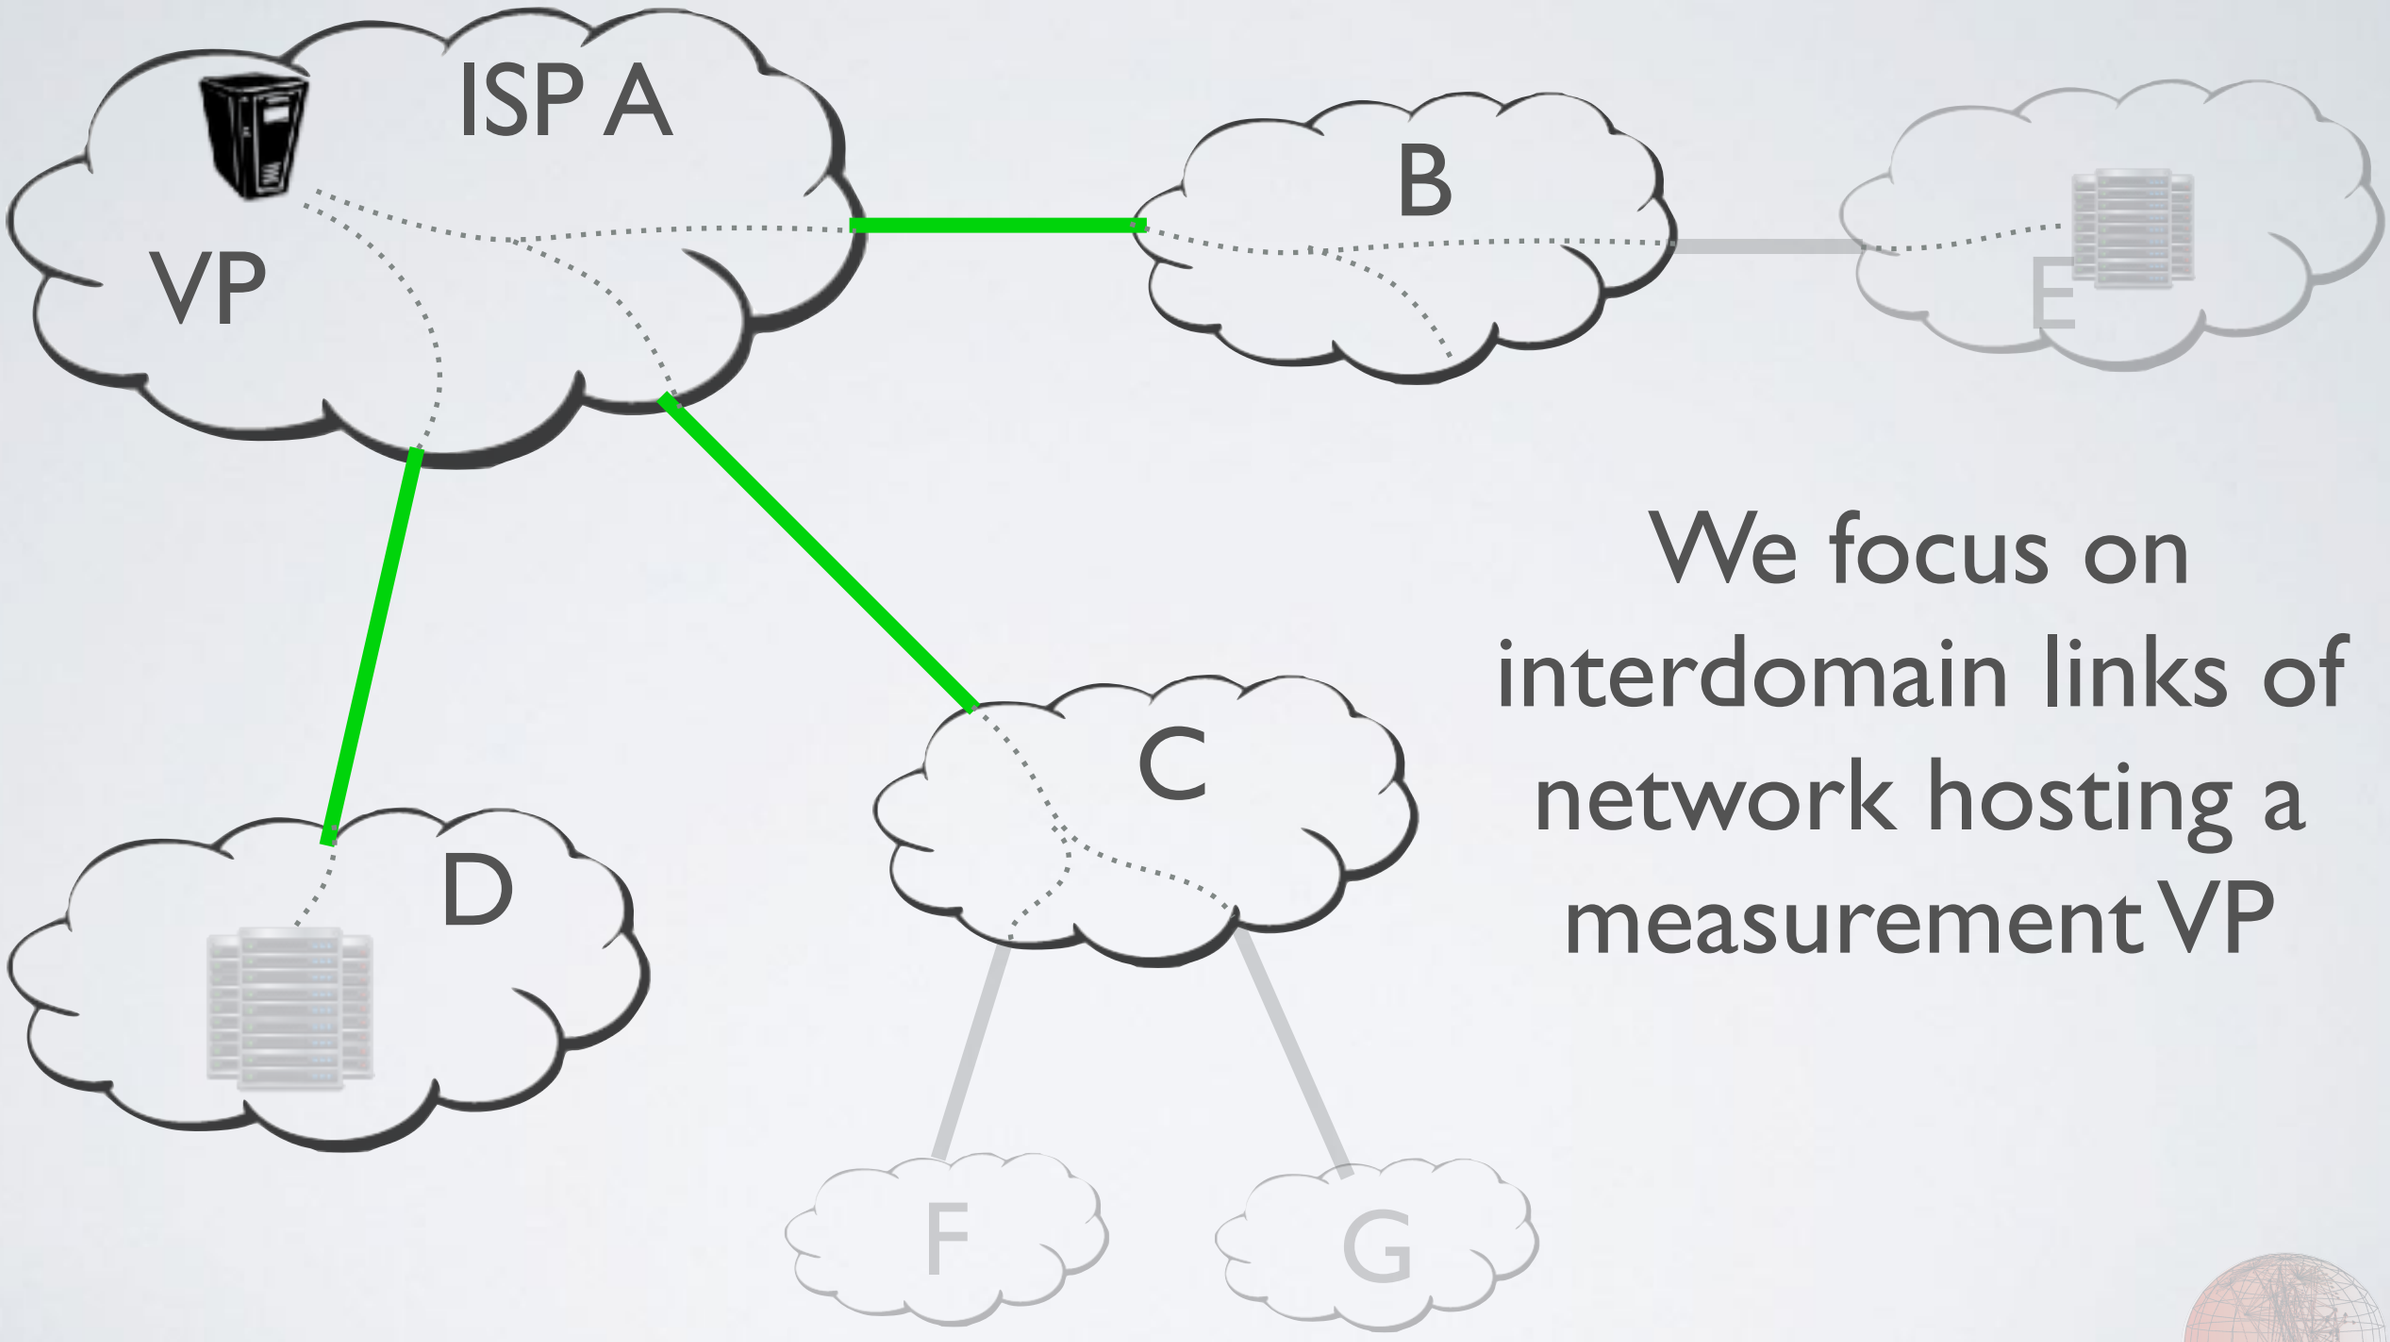

We focus on interdomain links of network hosting a measurement VP

caida

Identifying Interdomain Links

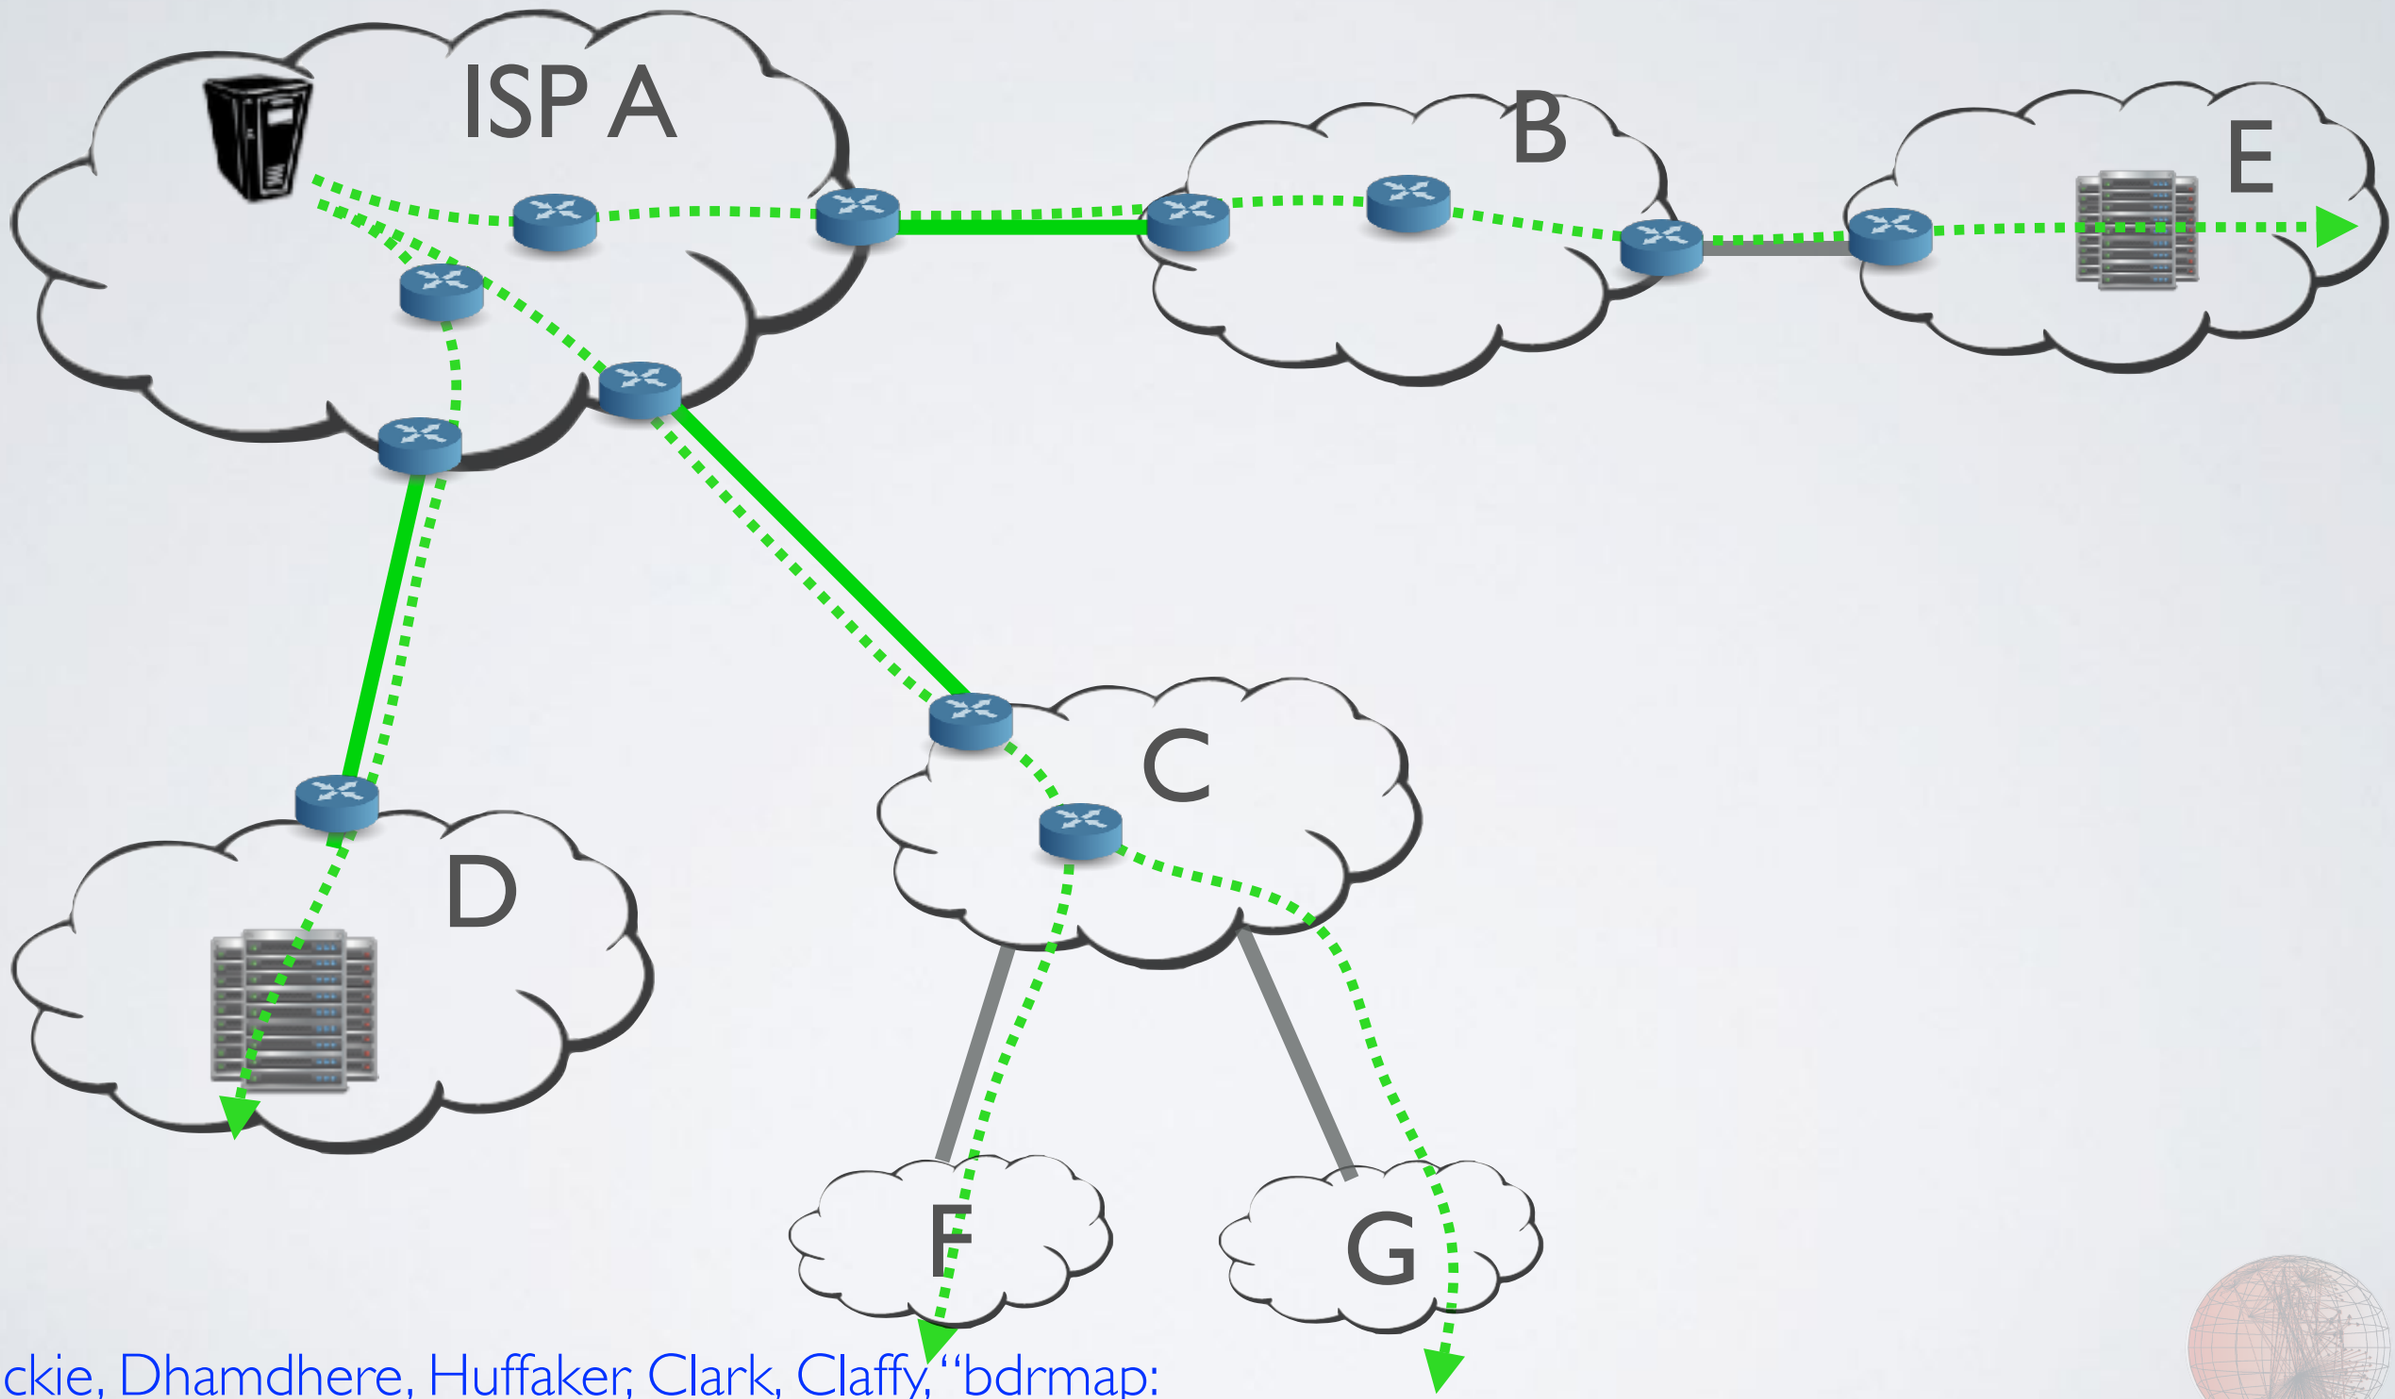

*Luckie, Dhamdhere, Huffaker, Clark, Claffy, "bdrmap: Inference of Borders Between IP Networks", IMC 2016

caida

Identifying Interdomain Links

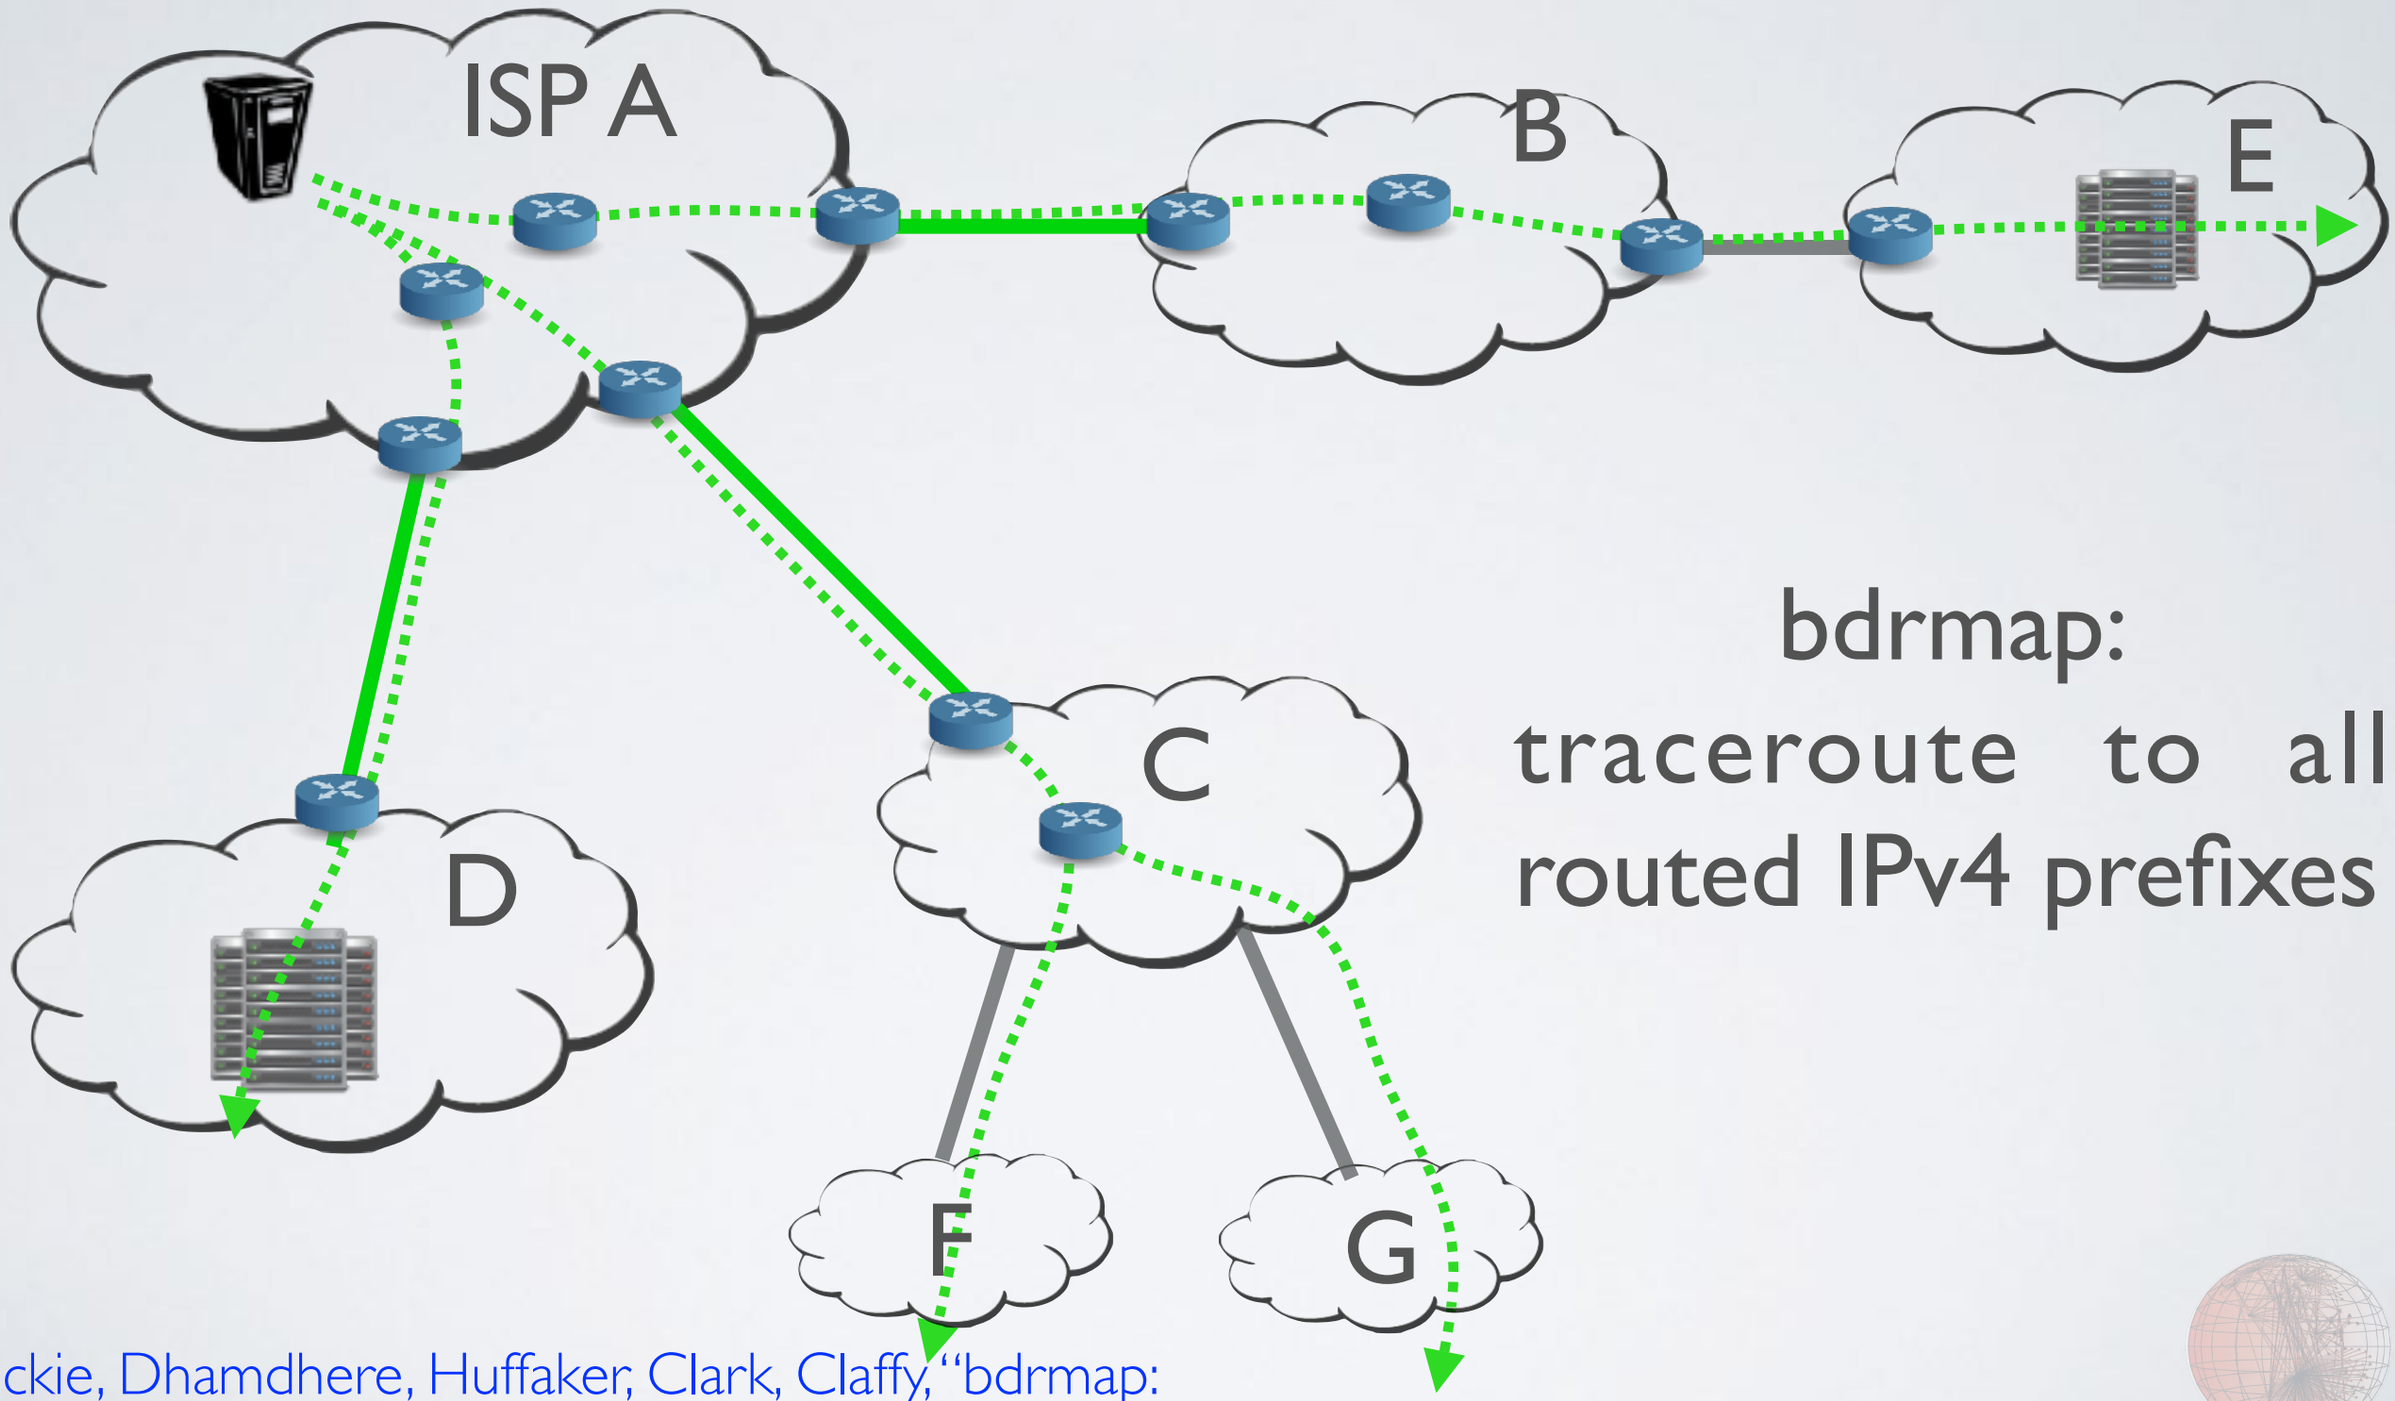

bdrmap:
traceroute to all routed IPv4 prefixes

*Luckie, Dhamdhere, Huffaker, Clark, Claffy, "bdrmap: Inference of Borders Between IP Networks", IMC 2016

caida

Identifying Interdomain Links

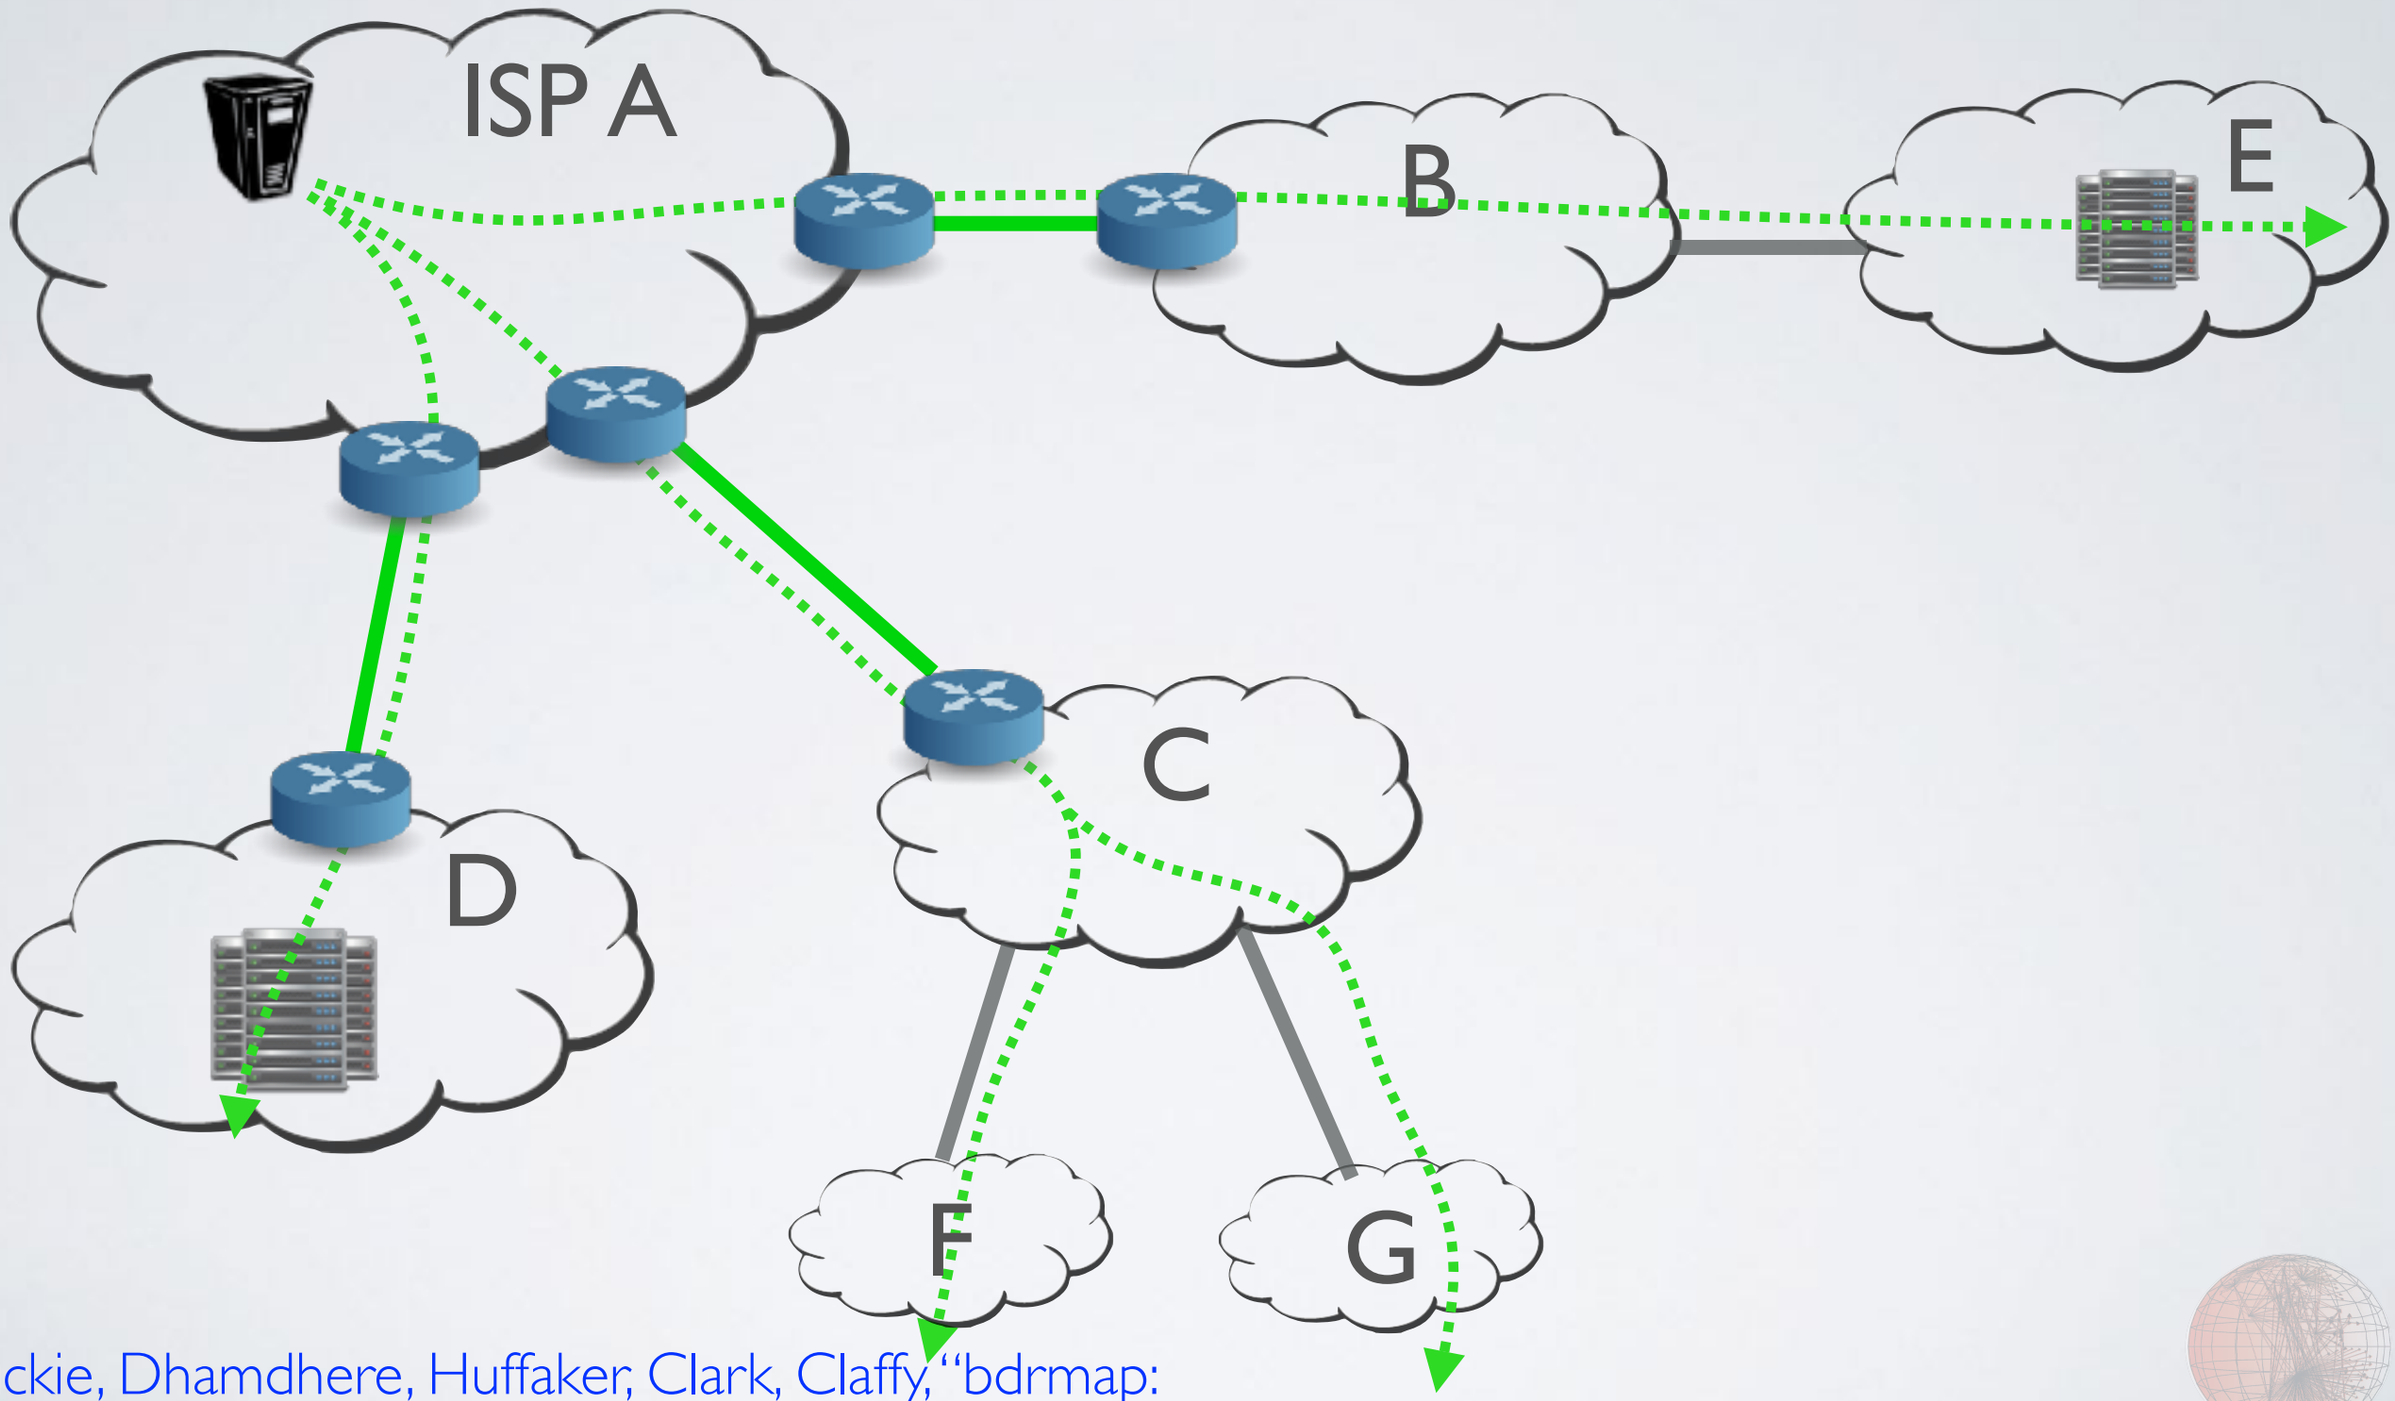

*Luckie, Dhamdhere, Huffaker, Clark, Claffy, "bdrmap: Inference of Borders Between IP Networks", IMC 2016

caida

Identifying Interdomain Links

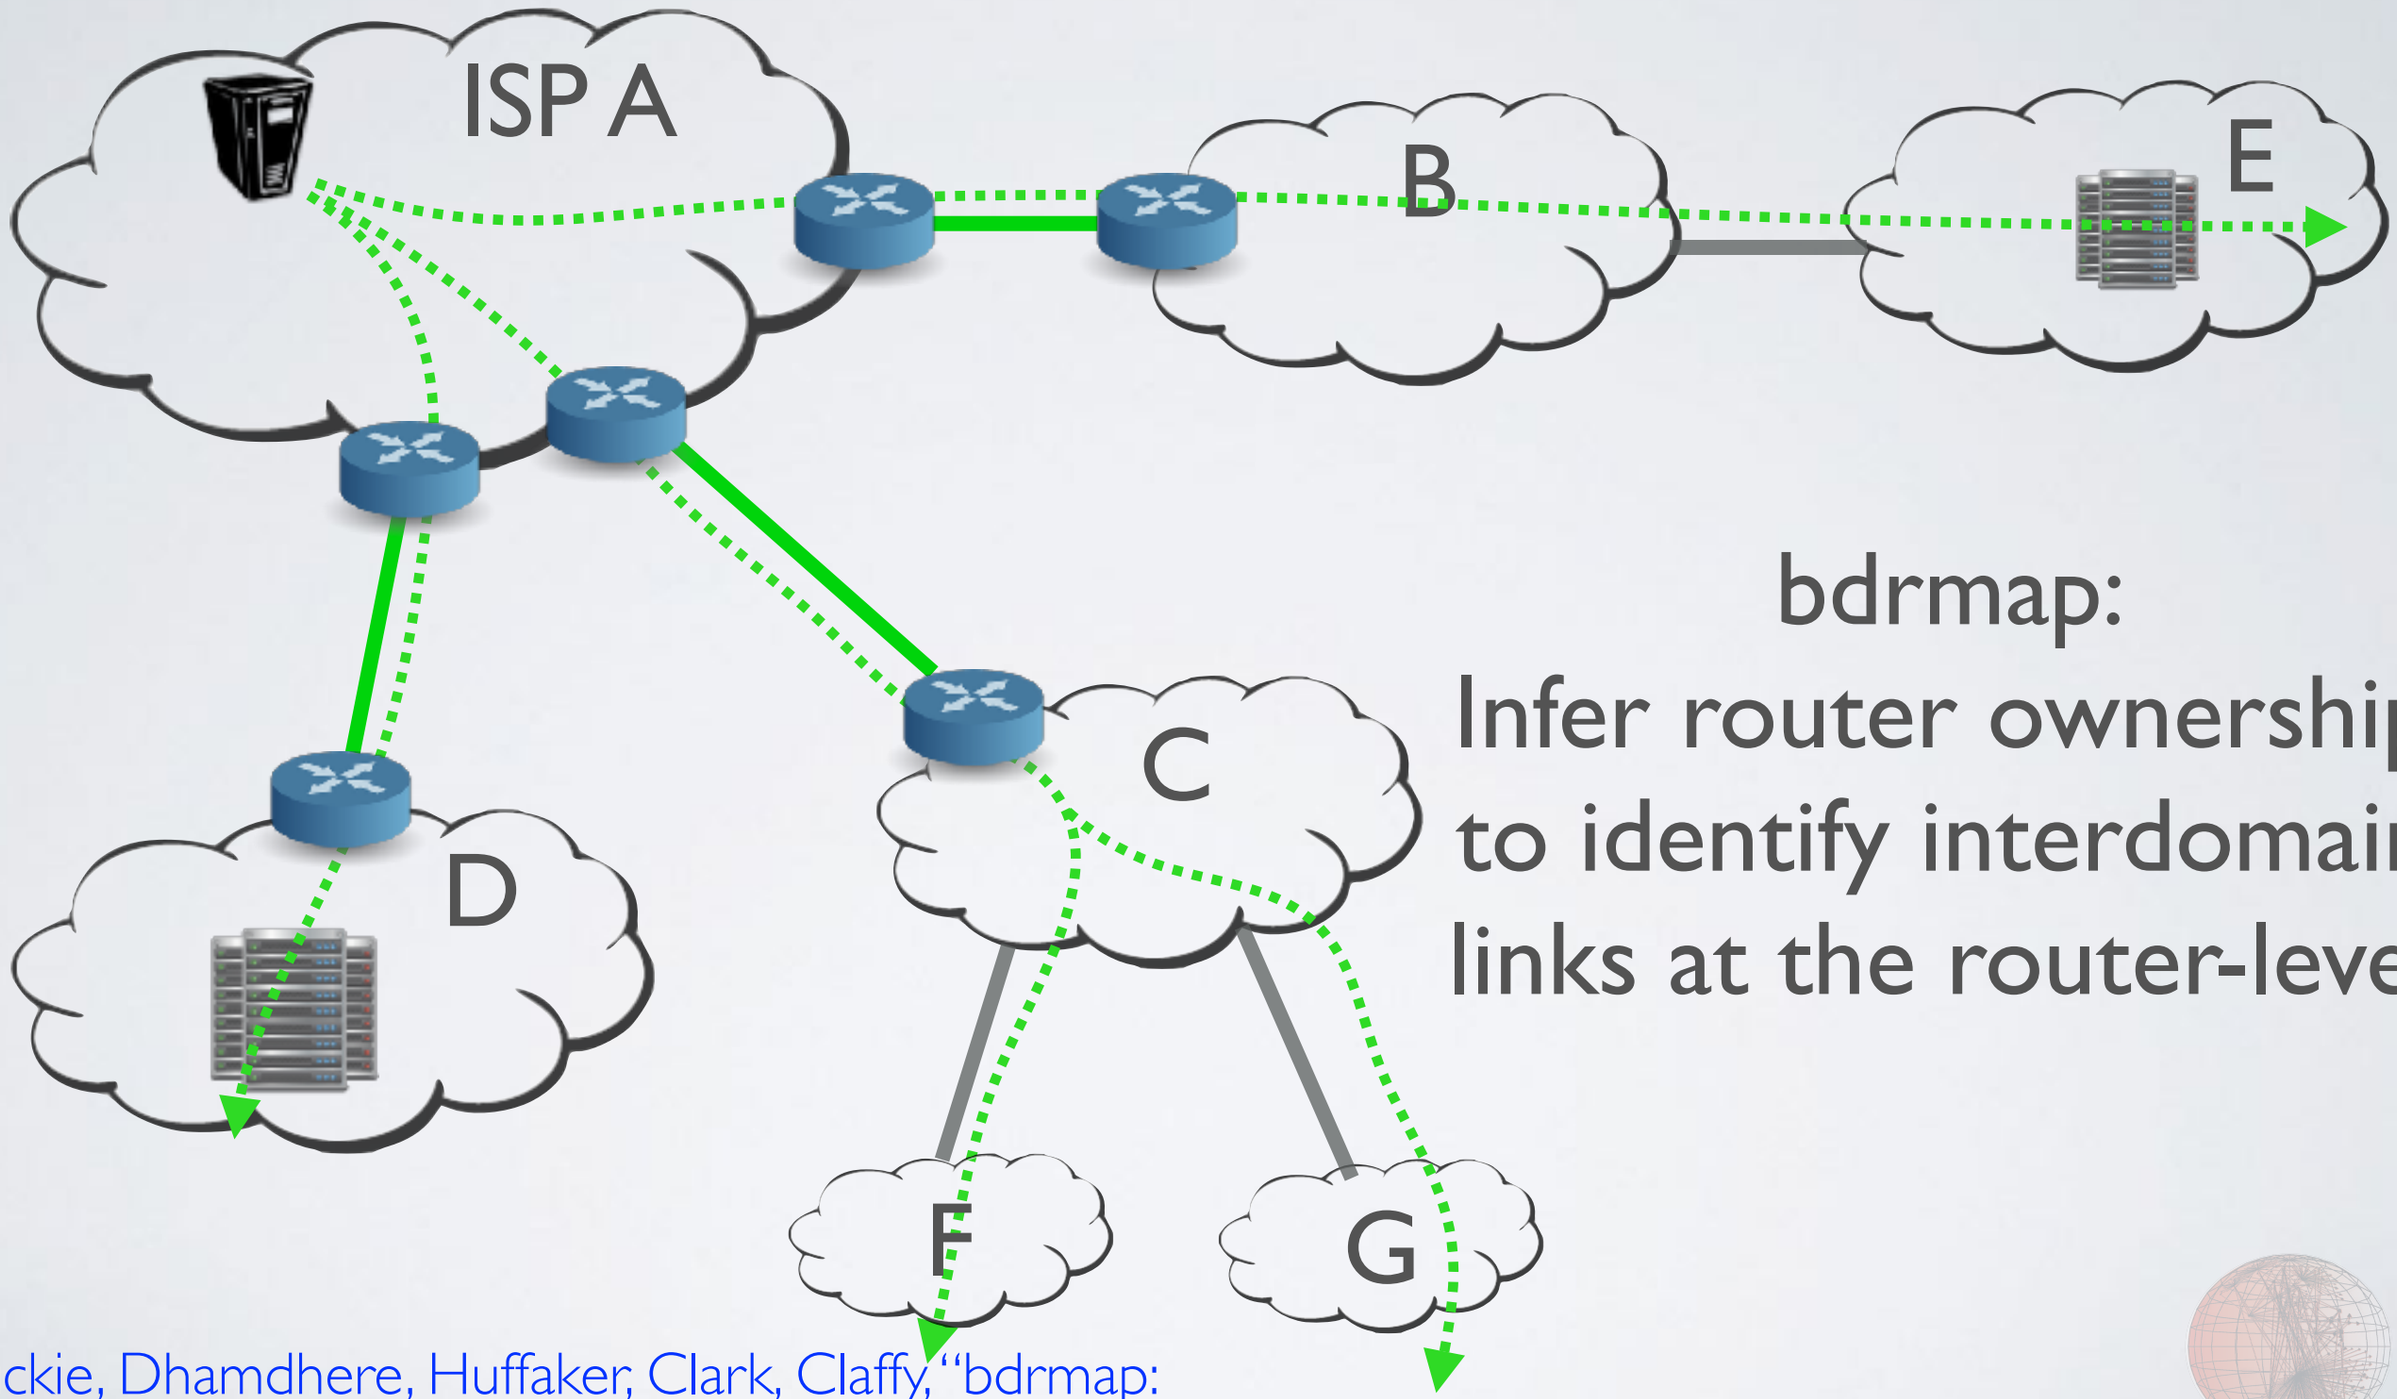

bdrmap:
Infer router ownership to identify interdomain links at the router-level

*Luckie, Dhamdhere, Huffaker, Clark, Claffy, "bdrmap: Inference of Borders Between IP Networks", IMC 2016

caida

Identifying Congested Links

- Focus on **persistently congested** links
- Look for periods of elevated latency that correlate across days (**autocorrelation method**)

caida

Autocorrelation method

Day 1

Day 2

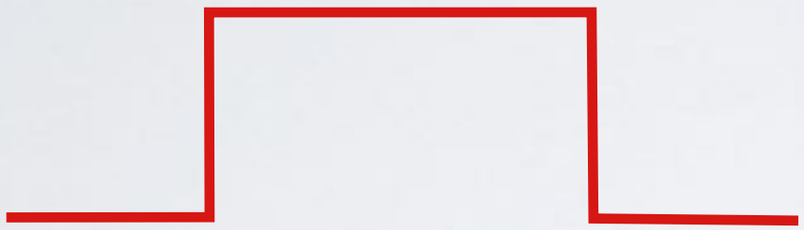

Day 3

Day 50

Far-side latency

caida

Autocorrelation method

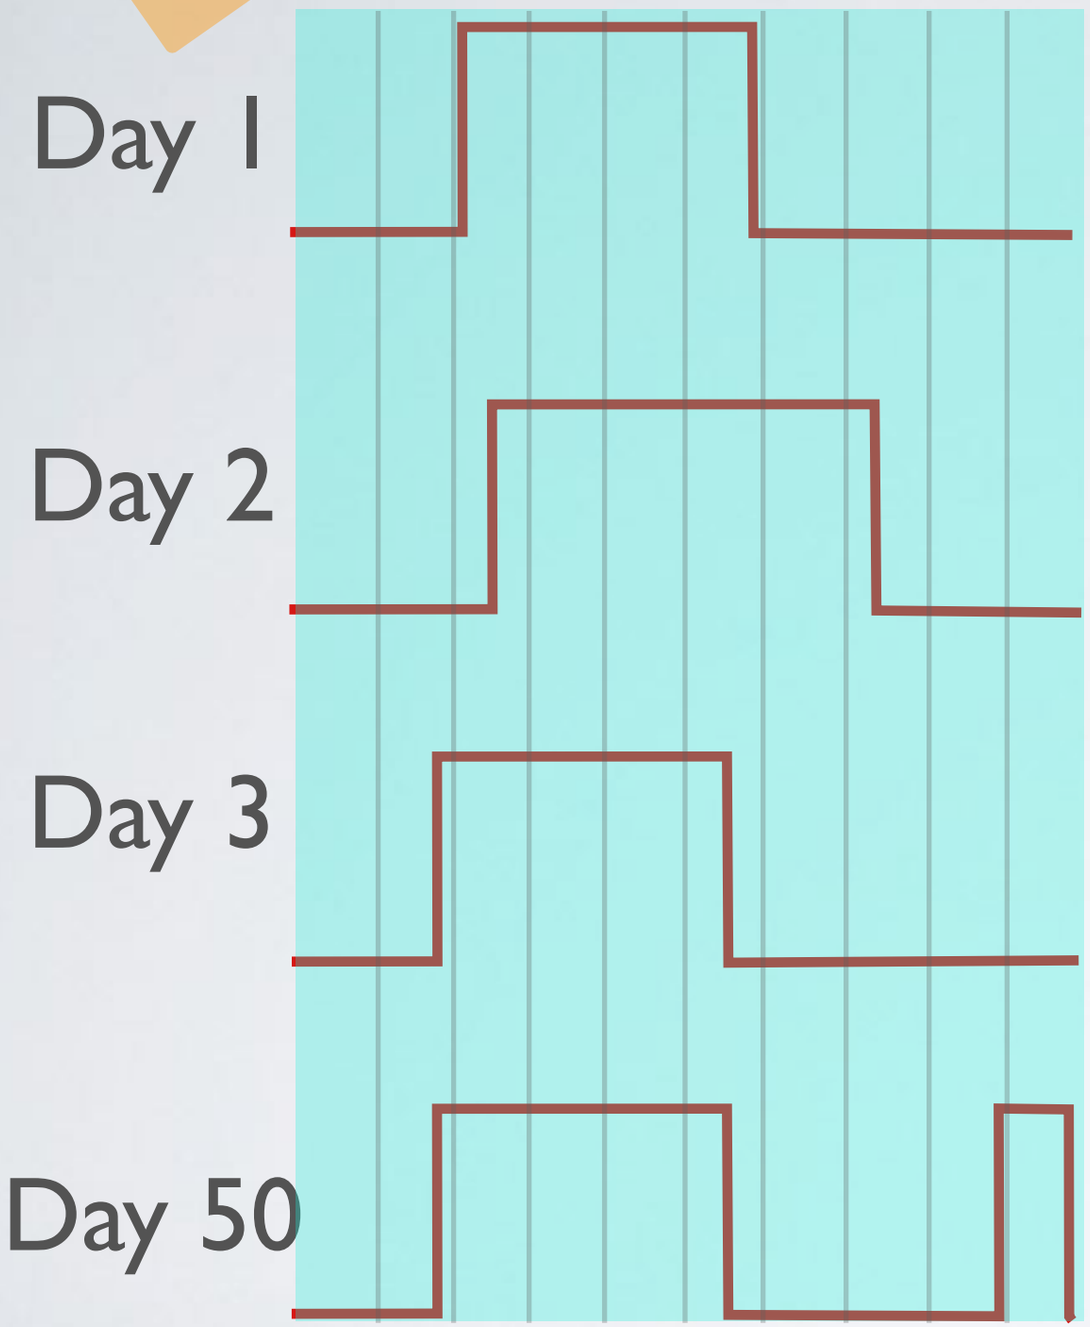

caida

Autocorrelation method

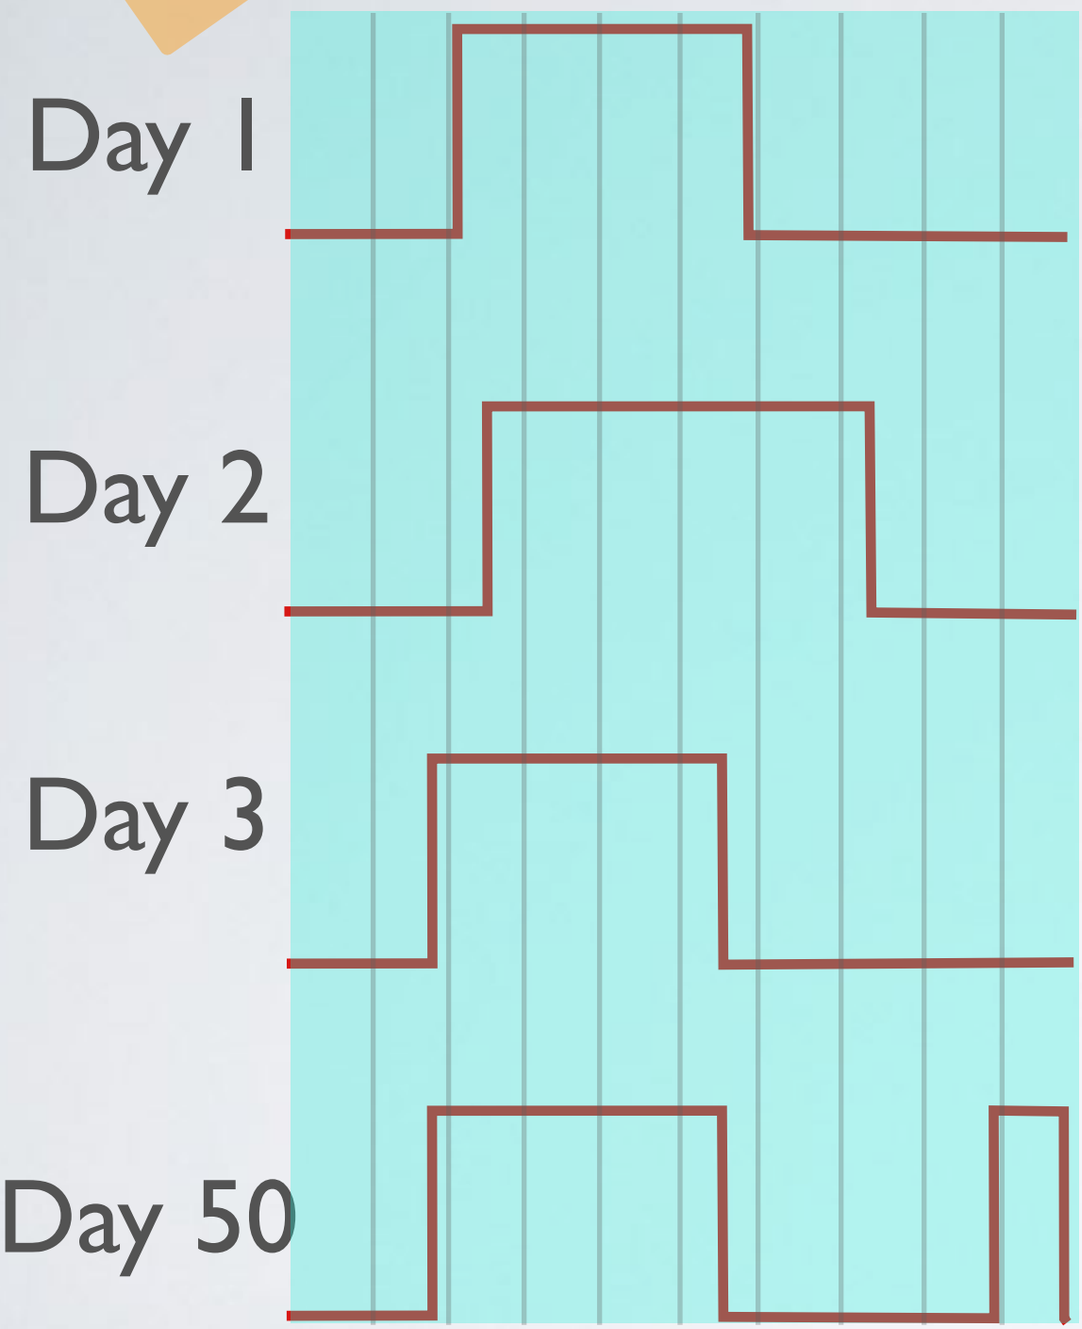

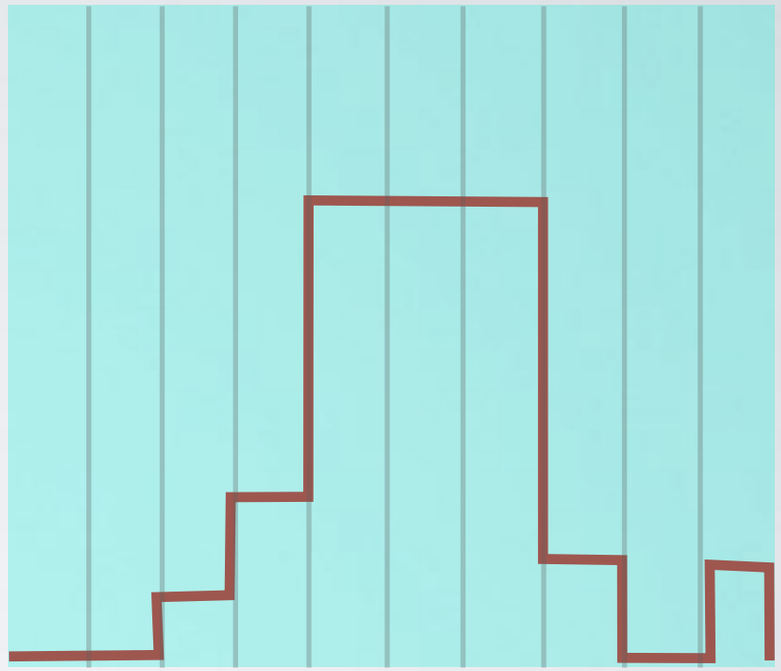

Sum number of days showing elevation in each 15-min interval

caida

Autocorrelation method

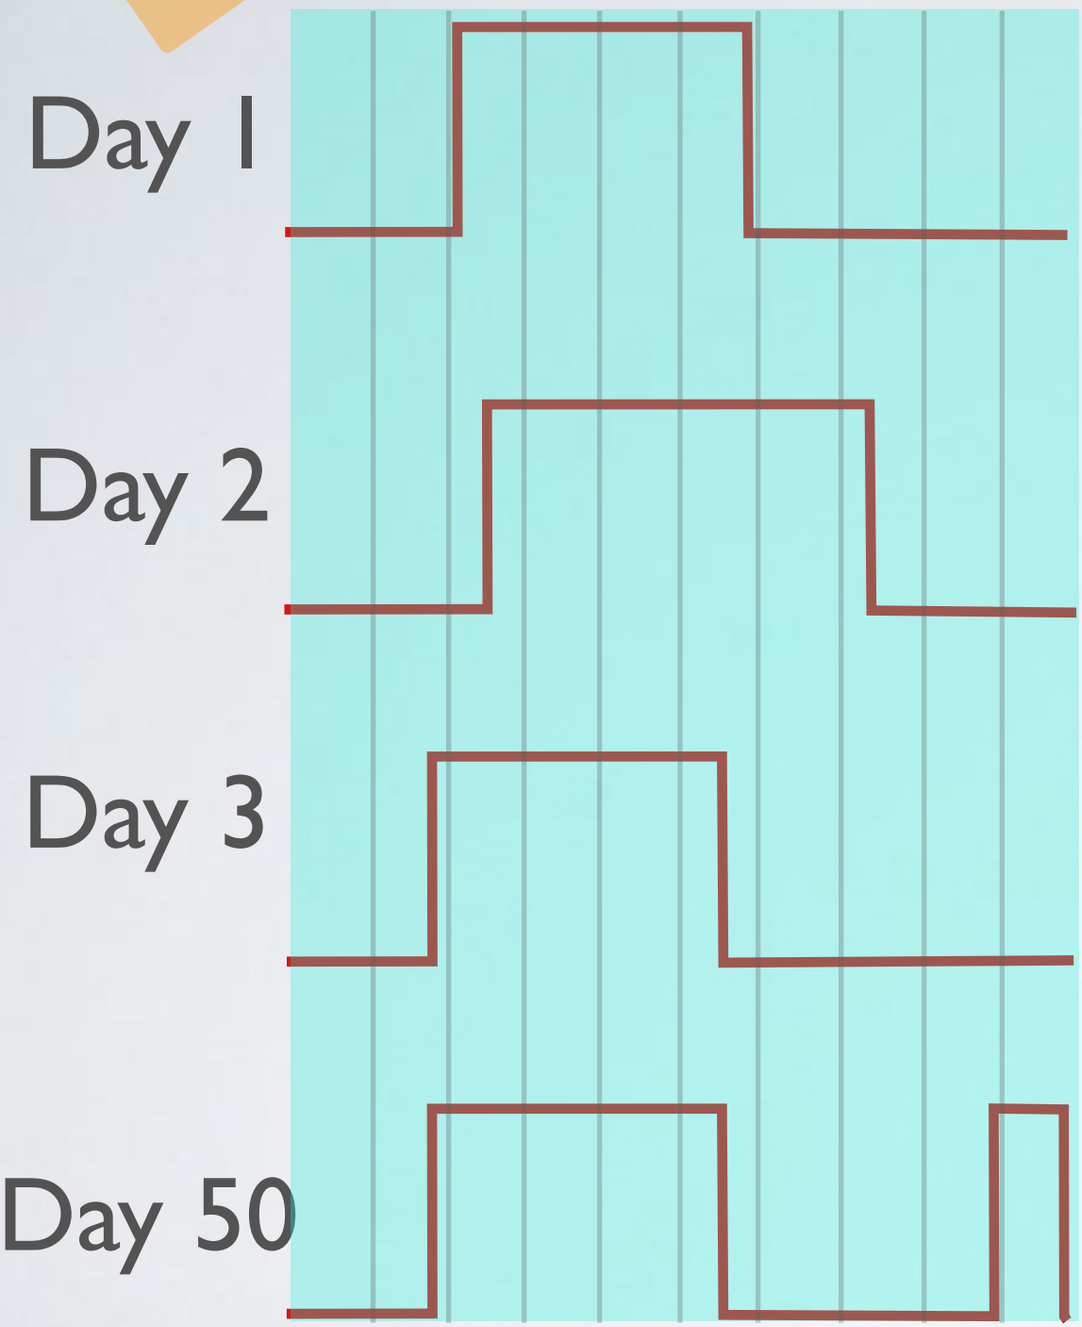

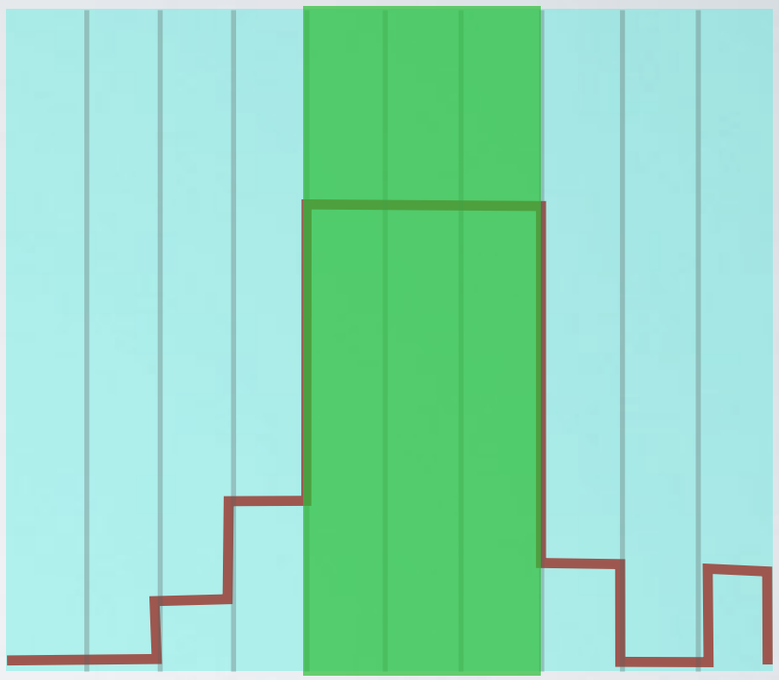

Infer an interval of recurring congestion

caida

Autocorrelation method

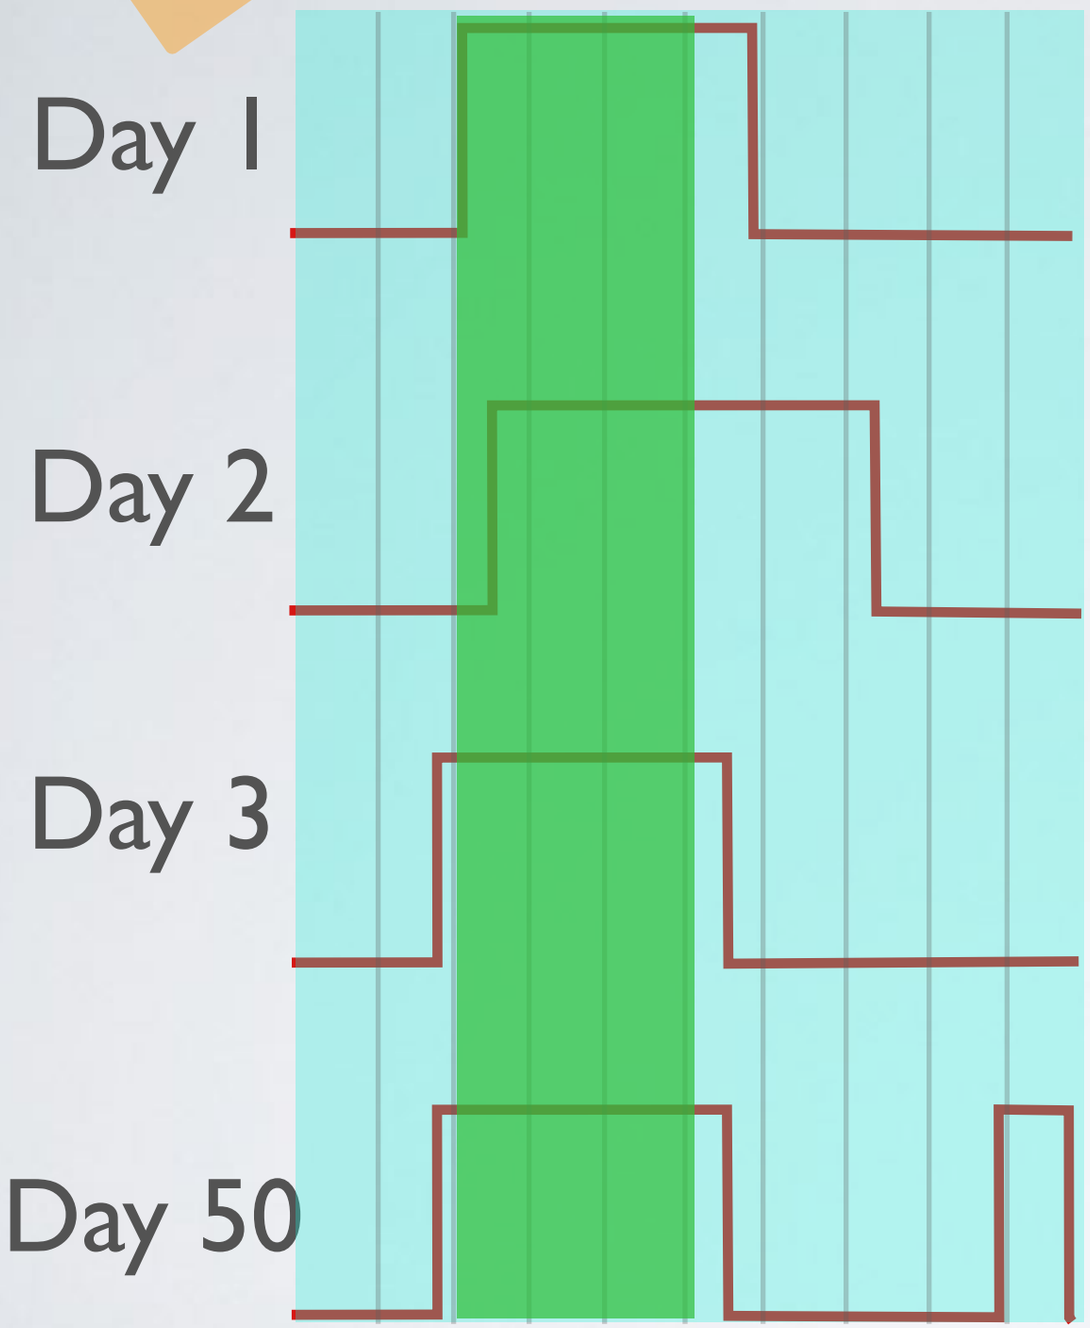

Far-side latency

Map inferred interval onto each day

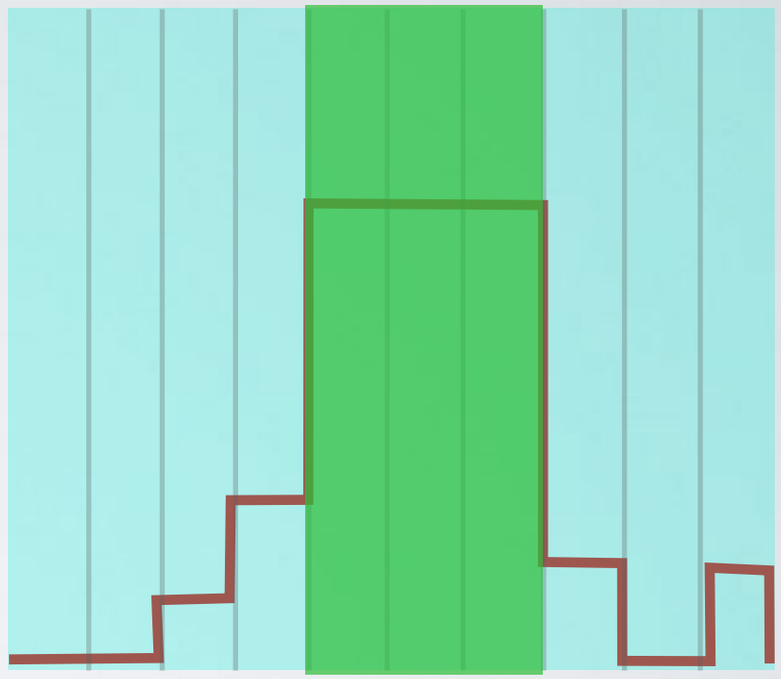

caida

Autocorrelation method

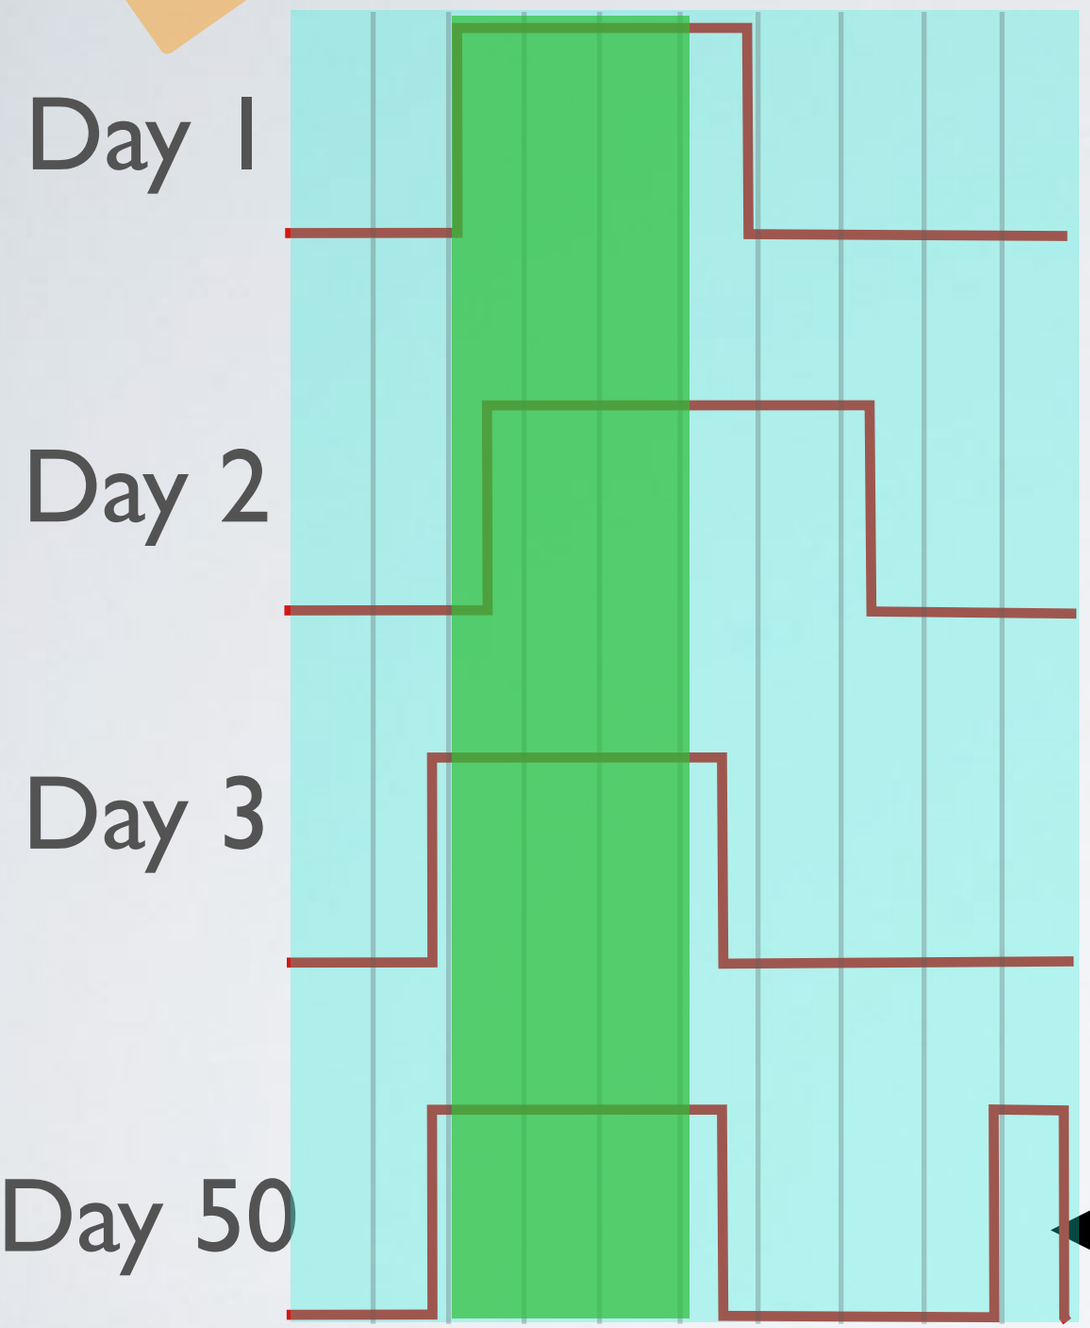

Far-side latency

Days

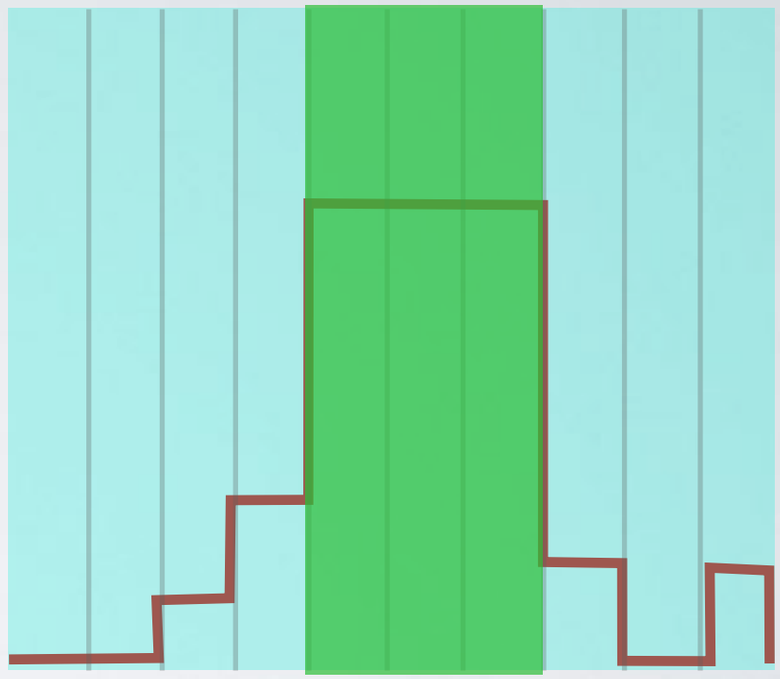

Map inferred interval onto each day

Reject elevated intervals that do not overlap with recurring interval

caida

Heavy Emphasis on Validation

(Critical due to political nature of inferences)

Heavy Emphasis on Validation

(Critical due to political nature of inferences)

- Validation of inference method
  - Correlation with loss
  - Correlation with throughput
  - Correlation with YouTube streaming

Heavy Emphasis on Validation

(Critical due to political nature of inferences)

- Validation of inference method
  - Correlation with loss
  - Correlation with throughput
  - Correlation with YouTube streaming
- Validation of specific inferences
  - Ground truth from operators

caida

Heavy Emphasis on Validation

(Critical due to political nature of inferences)

- Validation of inference method
  - **Correlation with loss**
  - **Correlation with throughput**
  - Correlation with YouTube streaming
- Validation of specific inferences
  - **Ground truth from operators**

caida

Heavy Emphasis on Validation
(Critical due to political nature of inferences)

- Validation of inference method
  - **Correlation with loss**
  - **Correlation with throughput**
  - Correlation with YouTube streaming
- Validation of specific inferences
  - **Ground truth from operators**

Brief overview in this talk
See paper for full details

caida

Does TSLP inference correlate with loss?

TSLP latency

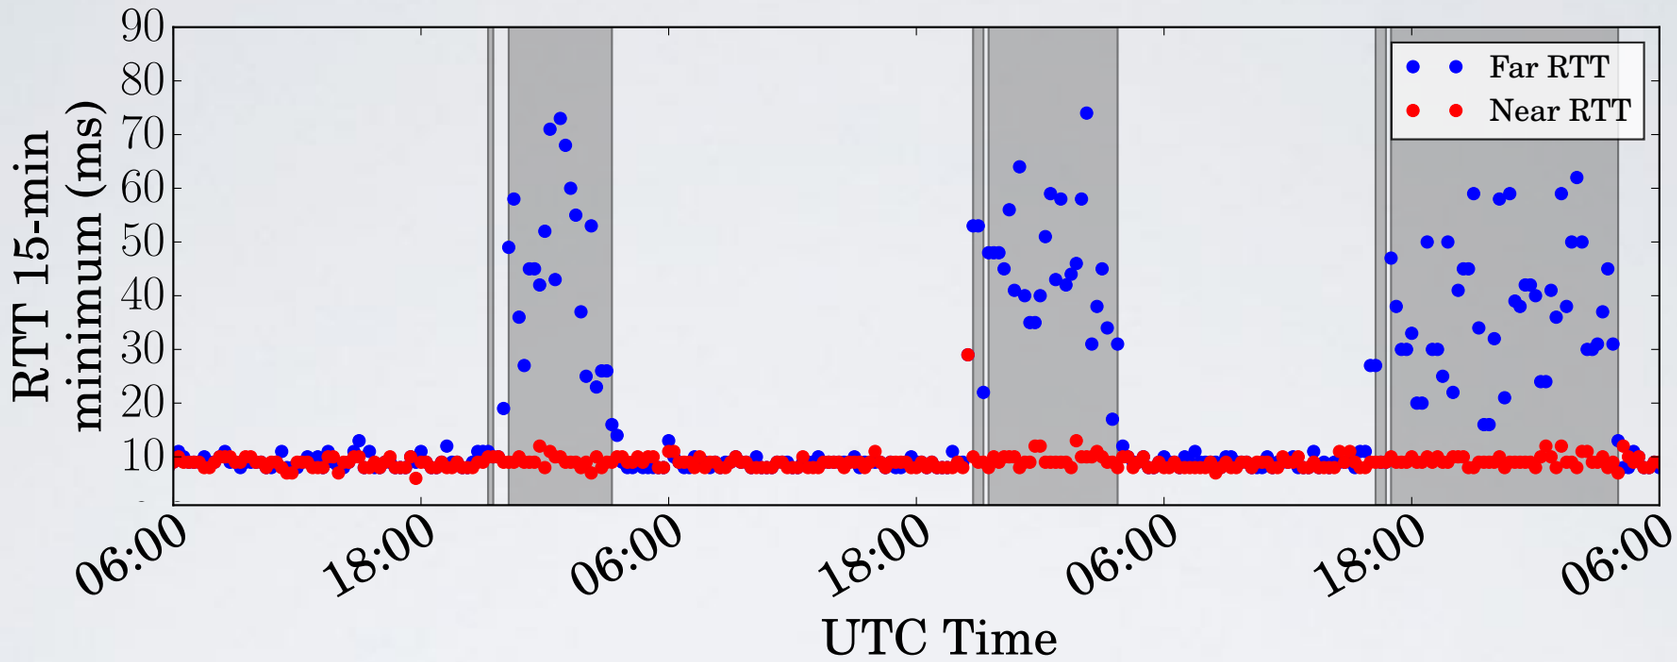

caida

Does TSLP inference correlate with loss?

TSLP latency

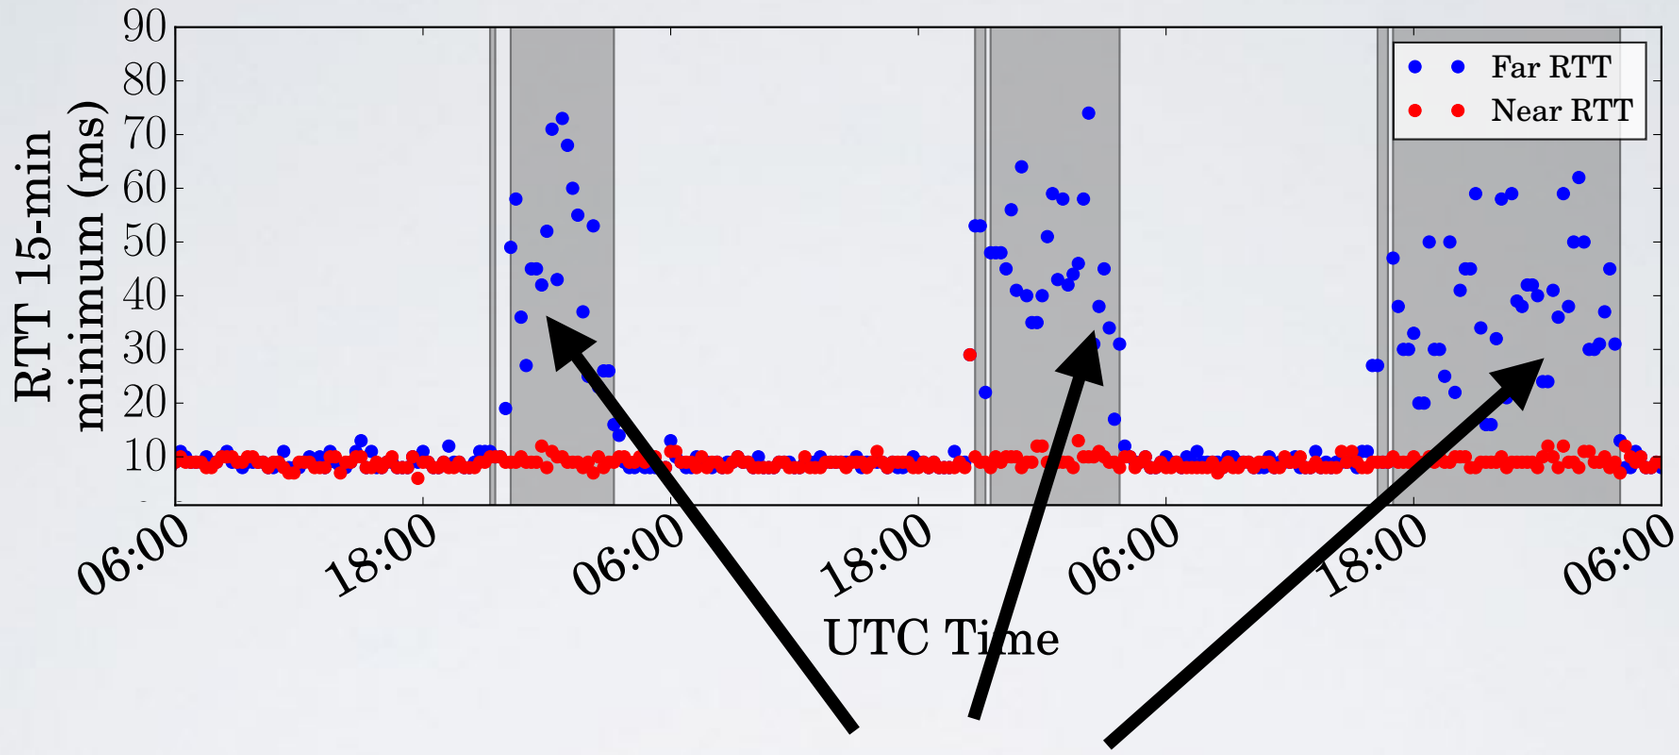

Diurnal latency elevation to far side indicates congestion

caida

Does TSLP inference correlate with loss?

TSLP latency

Loss Rate

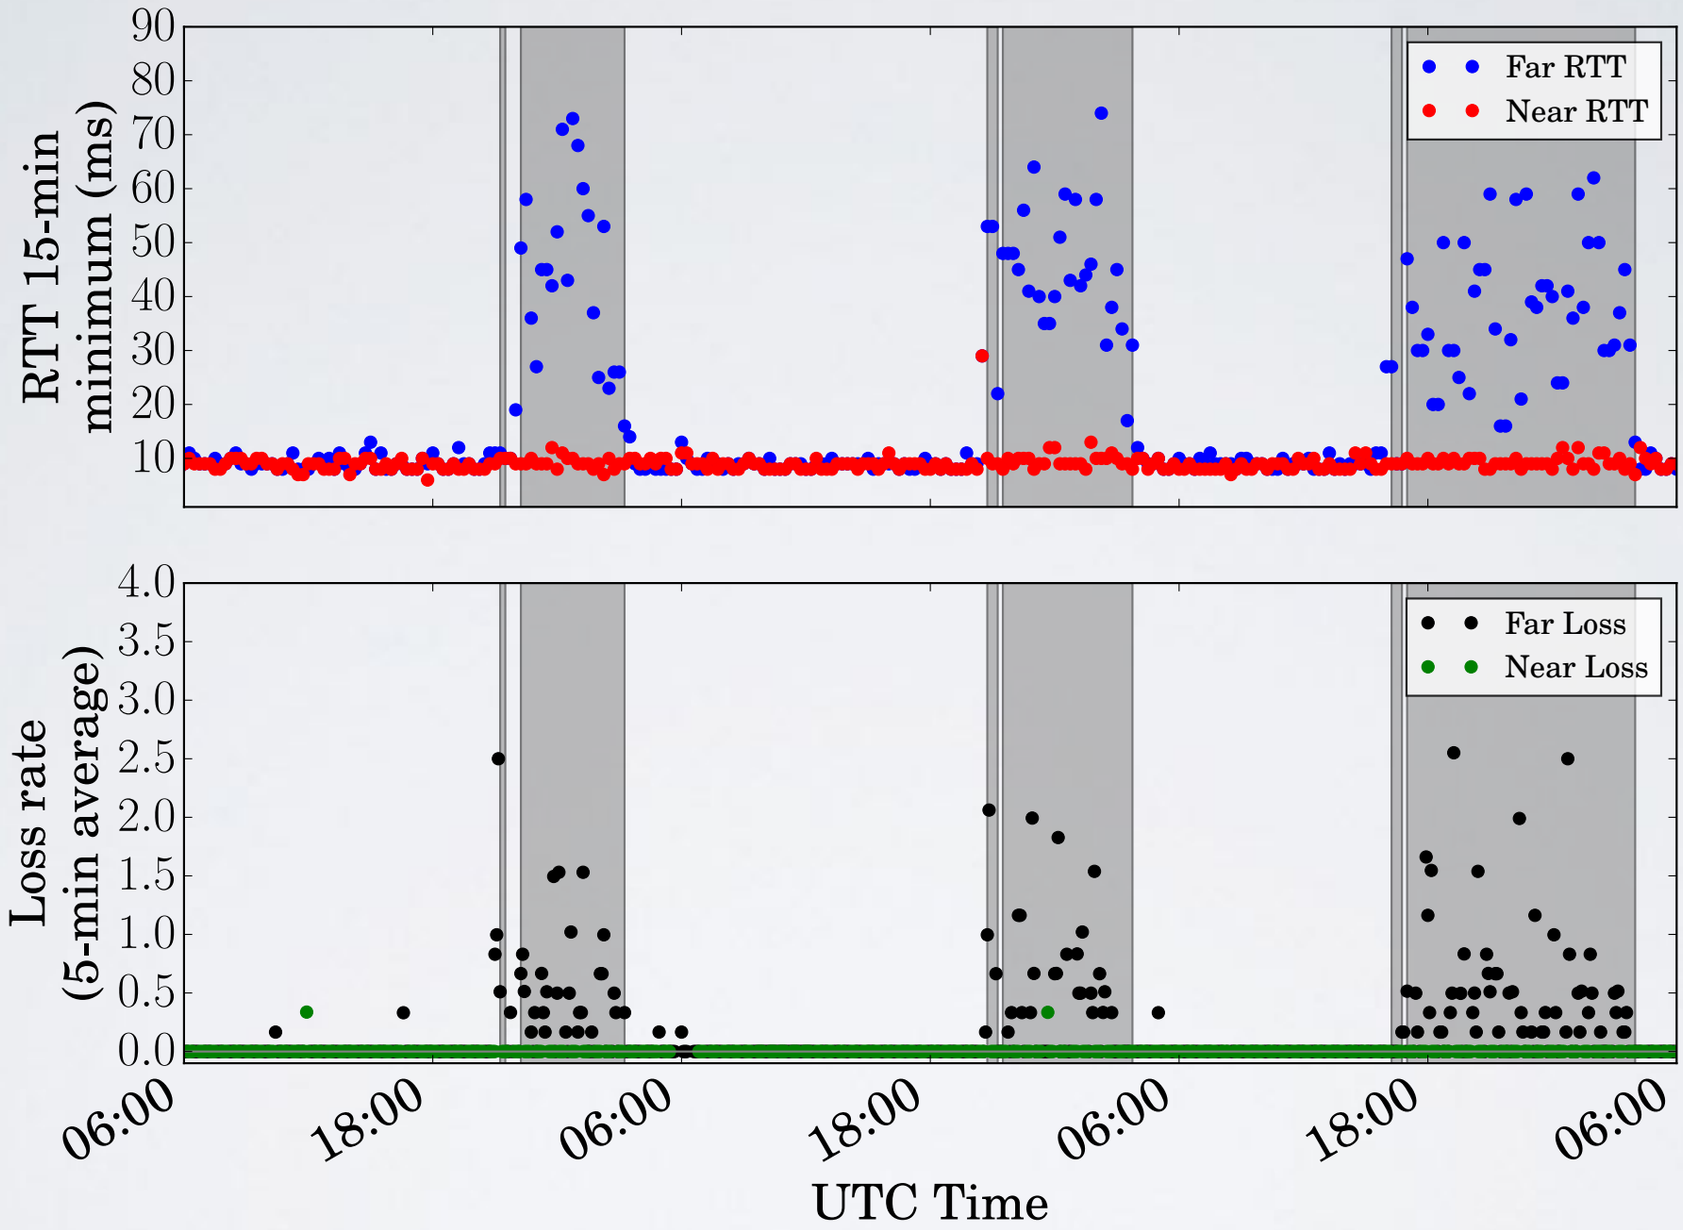

caida

Does TSLP inference correlate with loss?

TSLP latency

Loss Rate

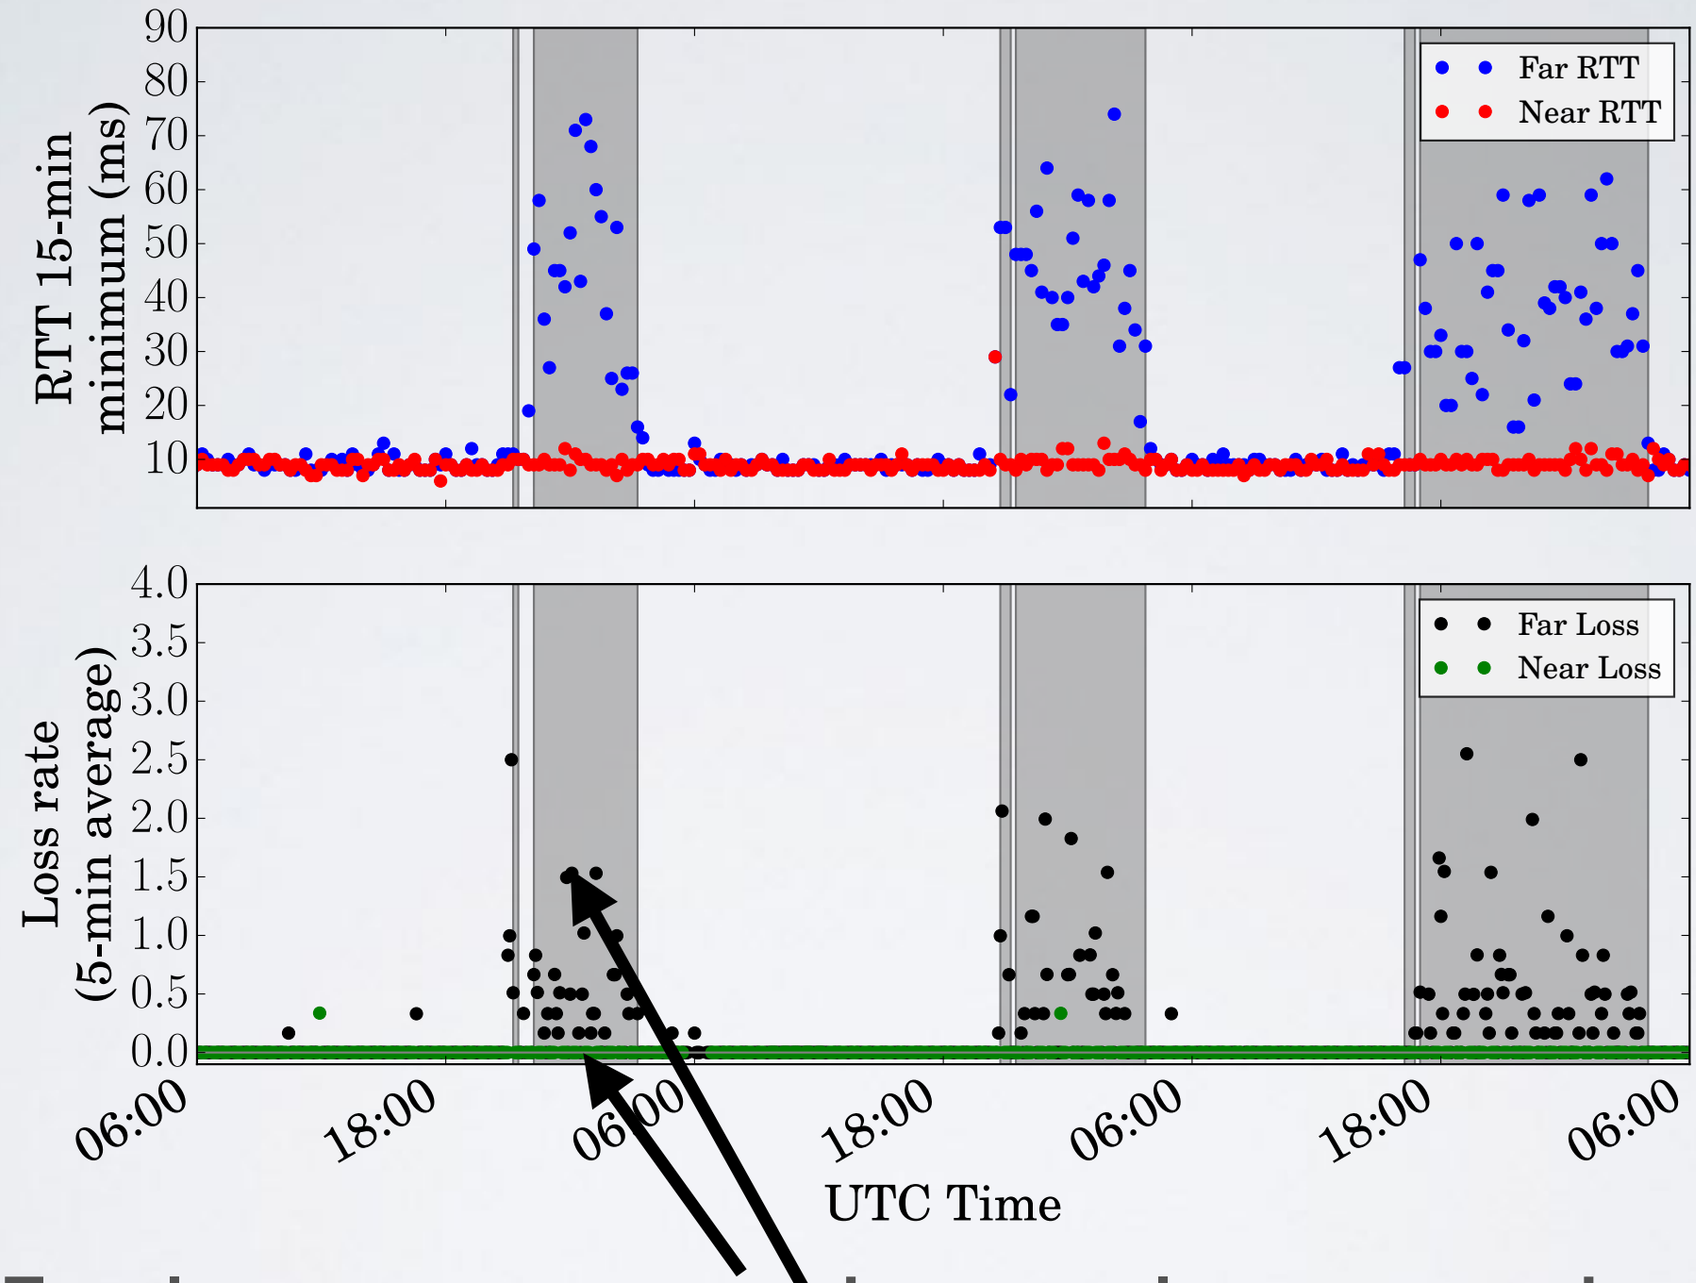

Far loss rate exceeds near loss rate during periods of inferred congestion

caida

Does TSLP inference correlate with loss?

TSLP latency

Loss Rate

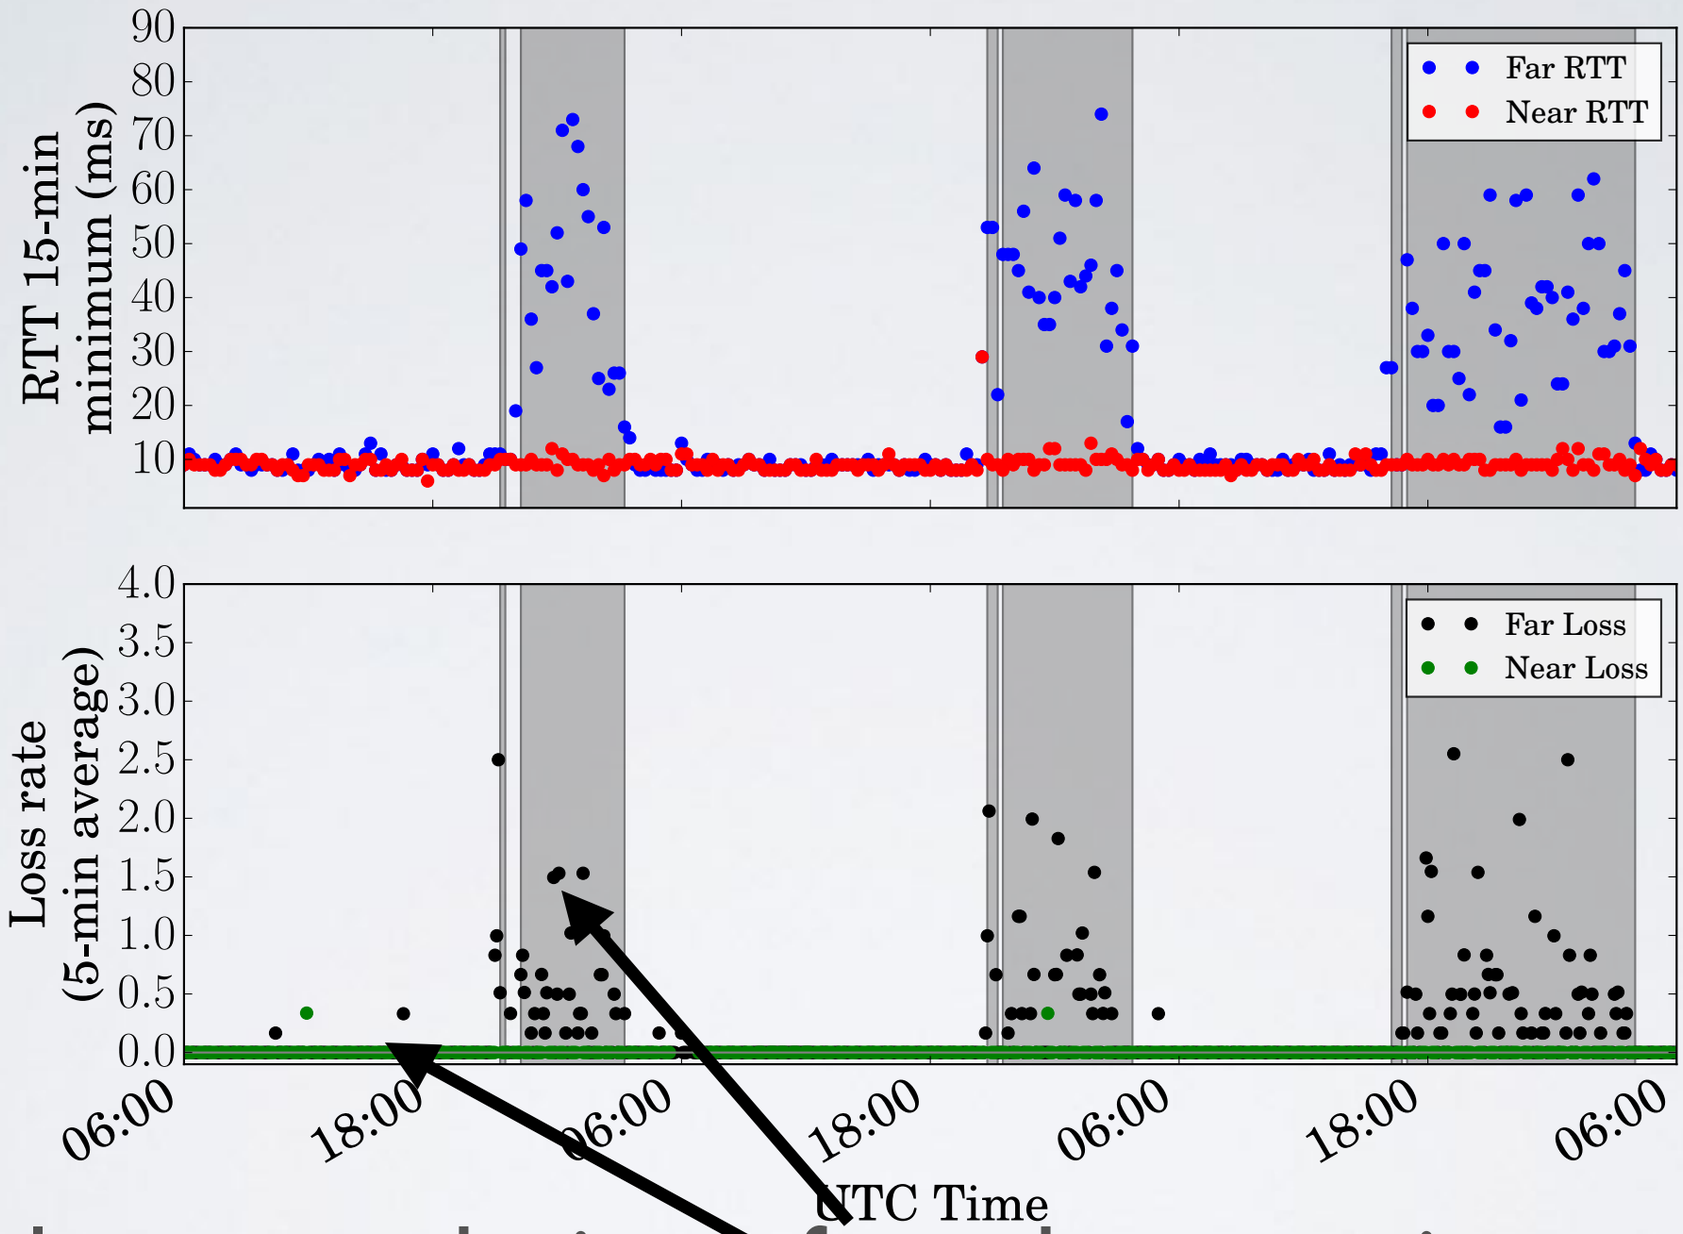

Far loss rate during inferred congestion exceeds far loss rate during non-congested period

caida

Does TSLP inference correlate with loss?

TSLP latency

Loss Rate

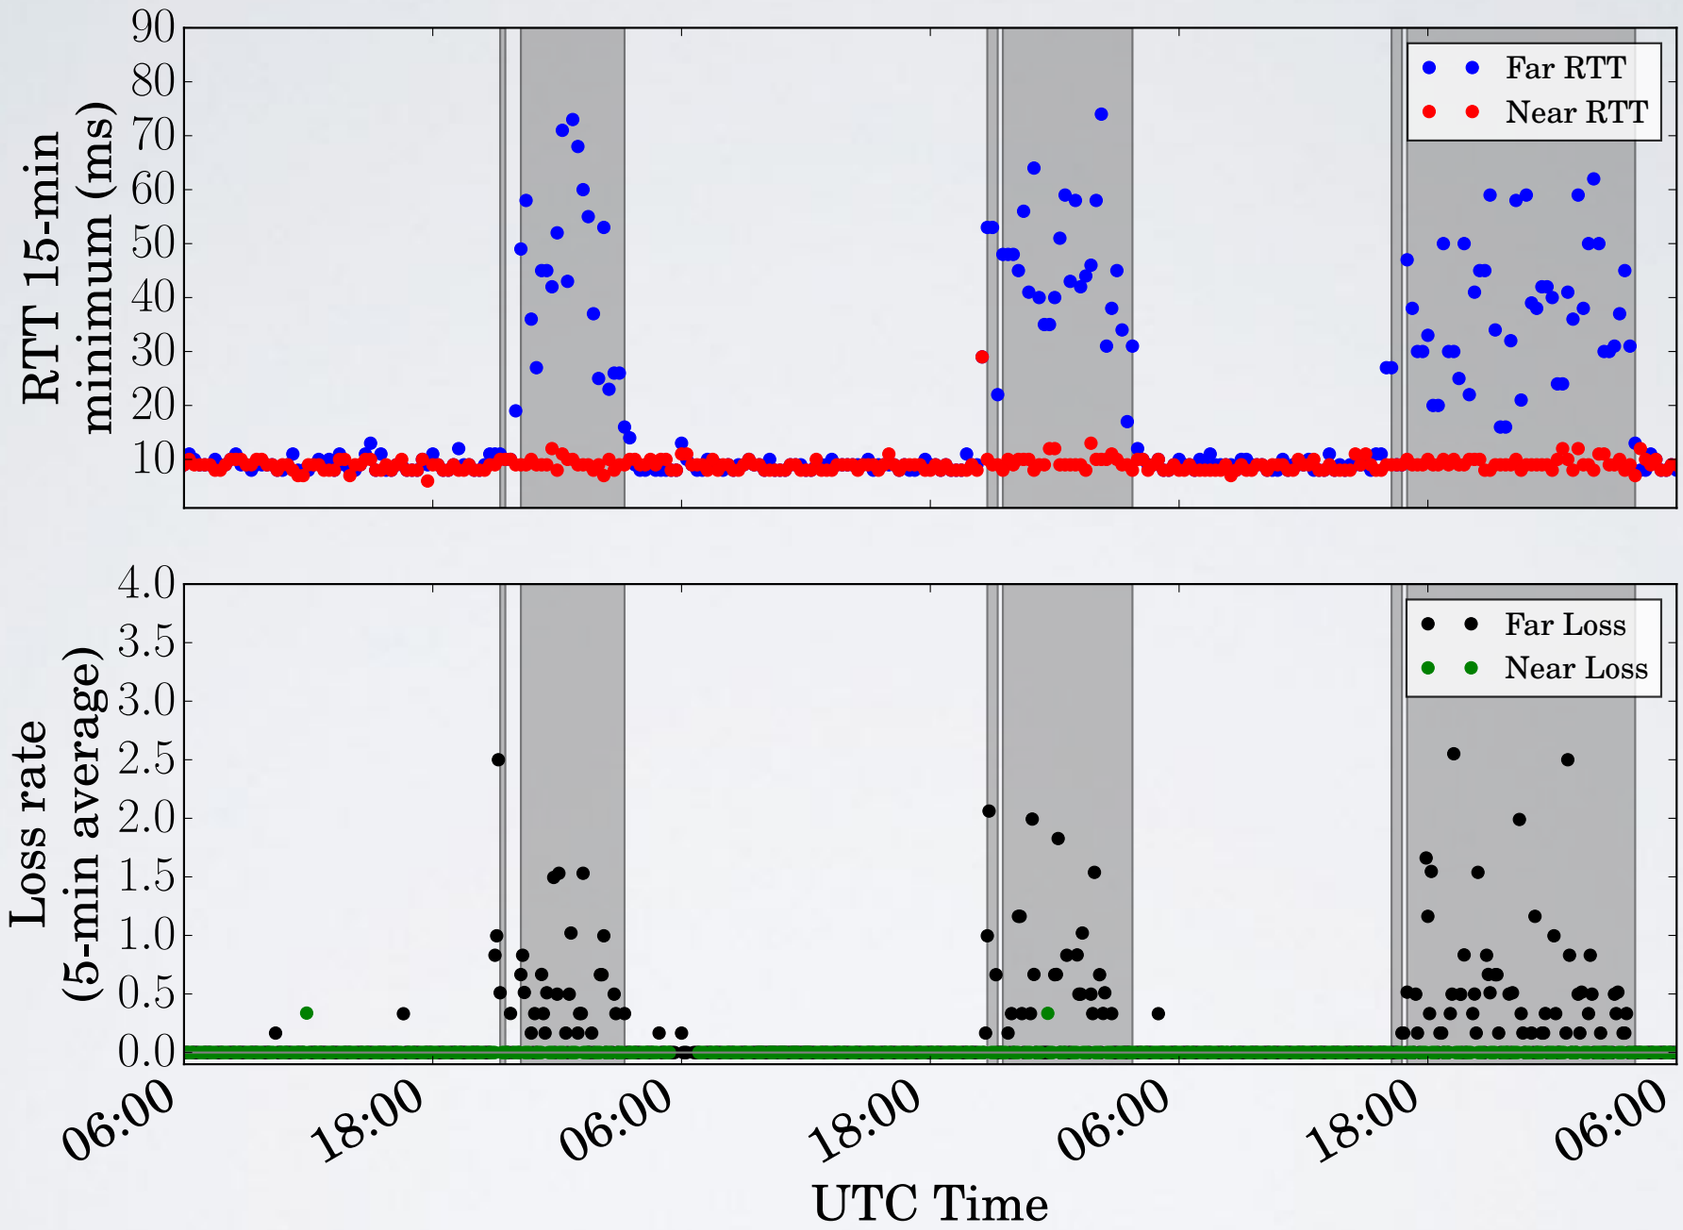

In 81% of cases, far loss rate was higher during congested periods and higher than near loss rate

caida

Does TSLP inference correlate with throughput?

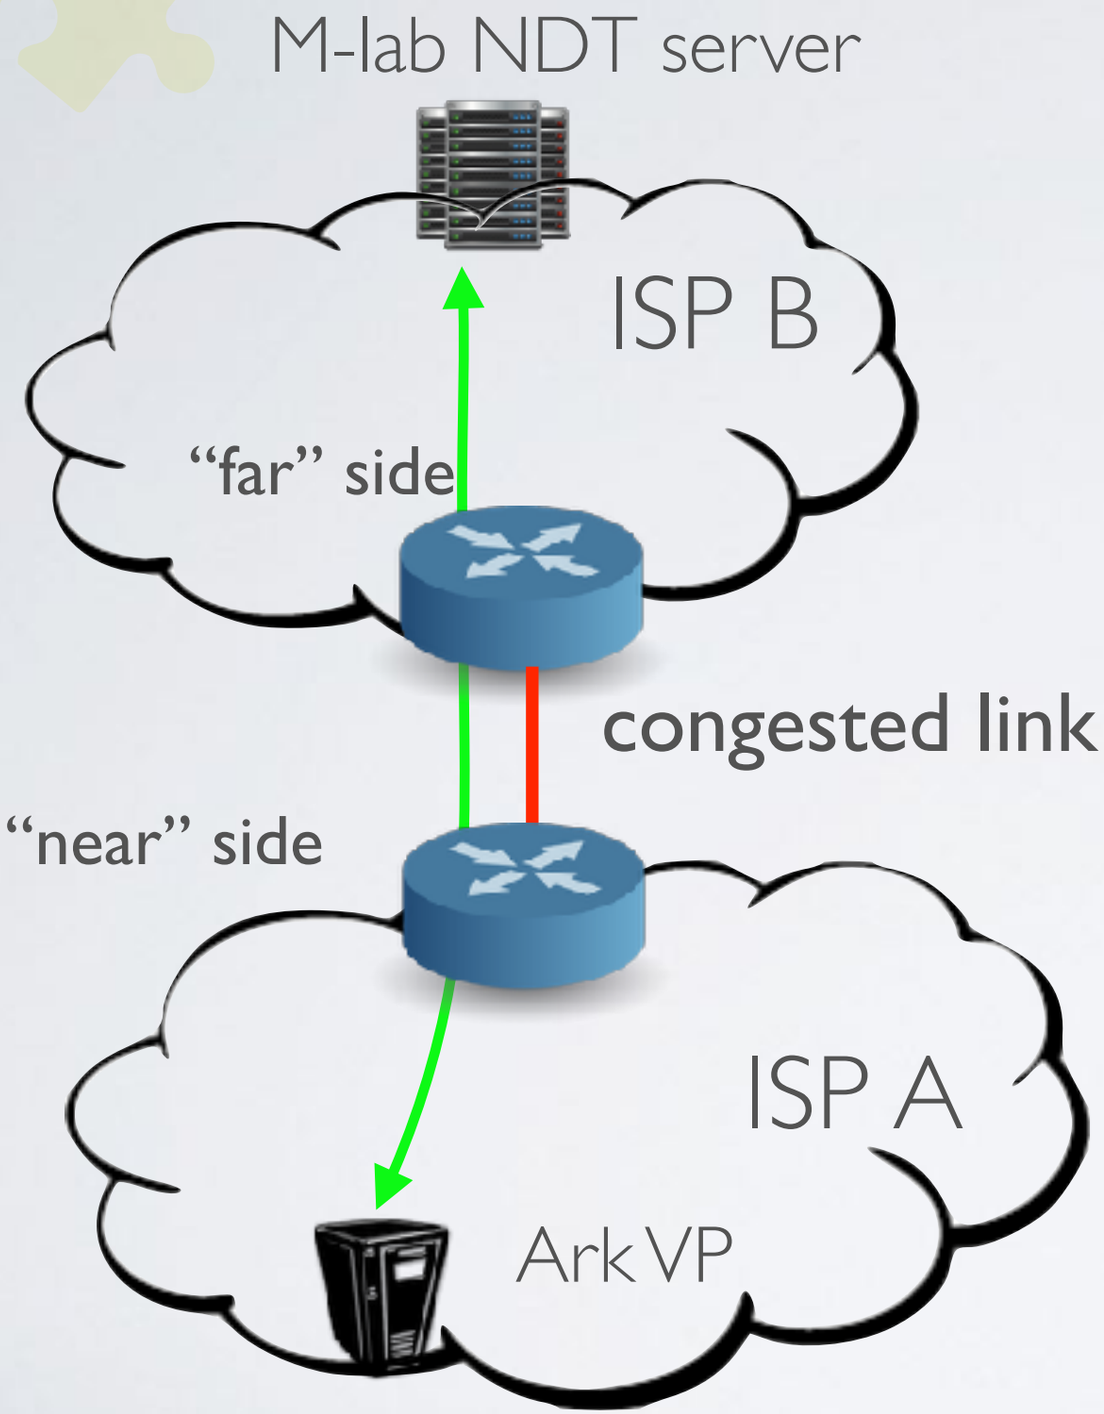

Approach: throughput measurements from Ark VP to M-lab NDT server traversing congested interdomain link

caida

Does TSLP inference correlate with throughput?

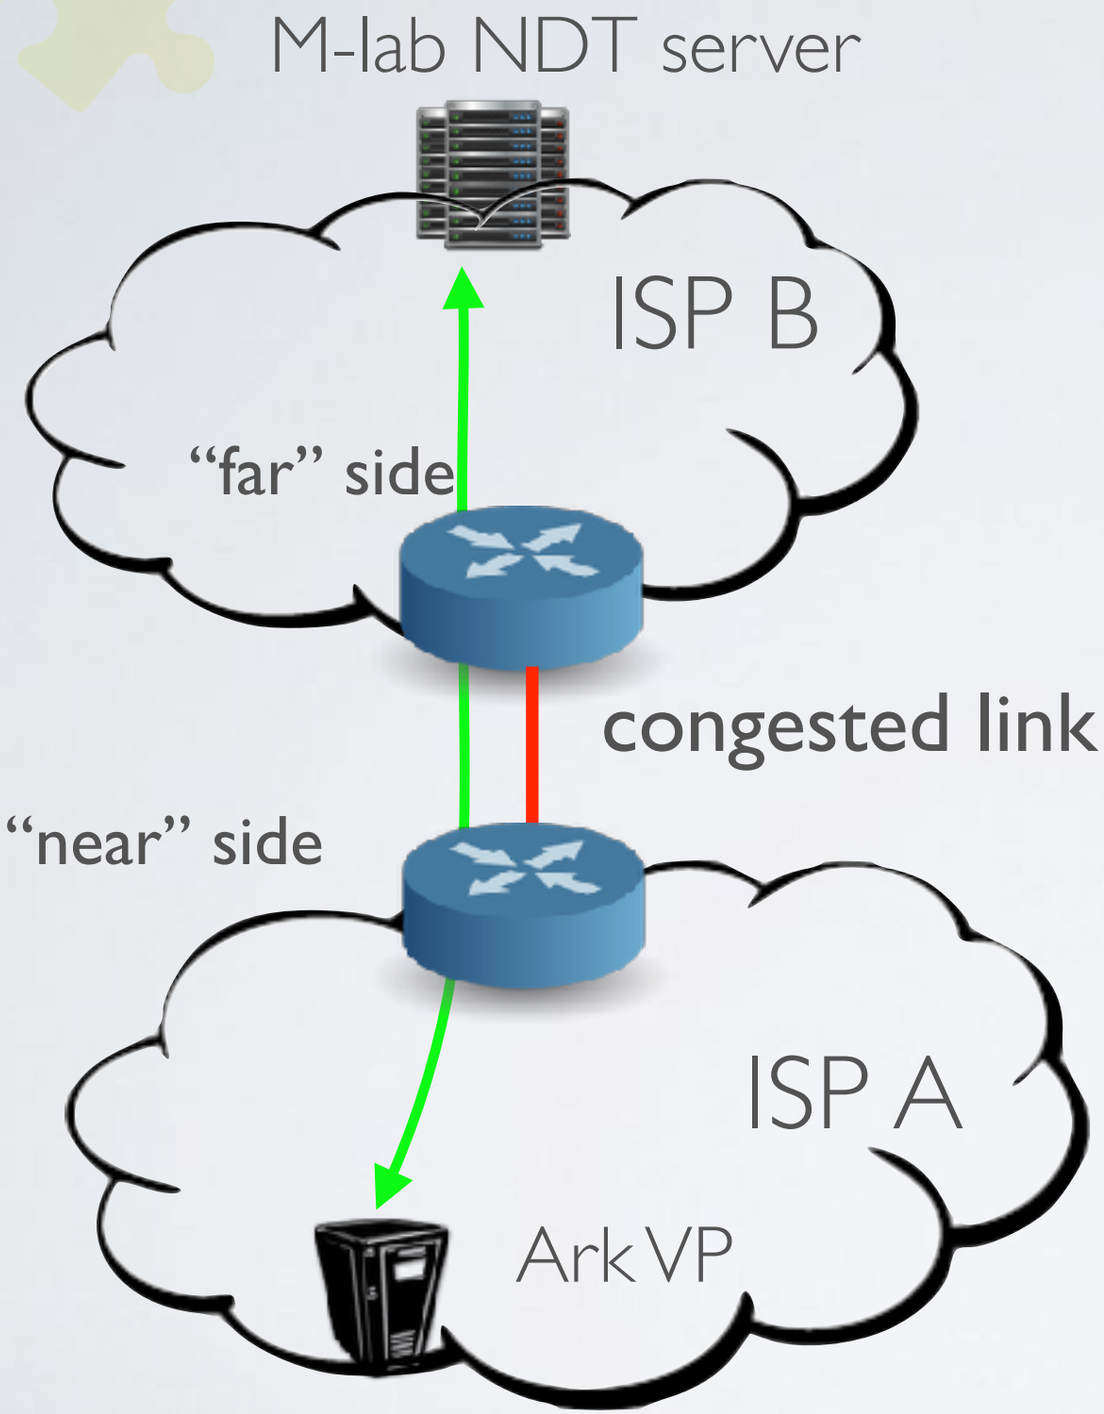

Approach: throughput measurements from Ark VP to M-lab NDT server traversing congested interdomain link

Challenge: difficult to find NDT servers that cover specifically observed interconnections

caida

Does TSLP inference correlate with throughput?

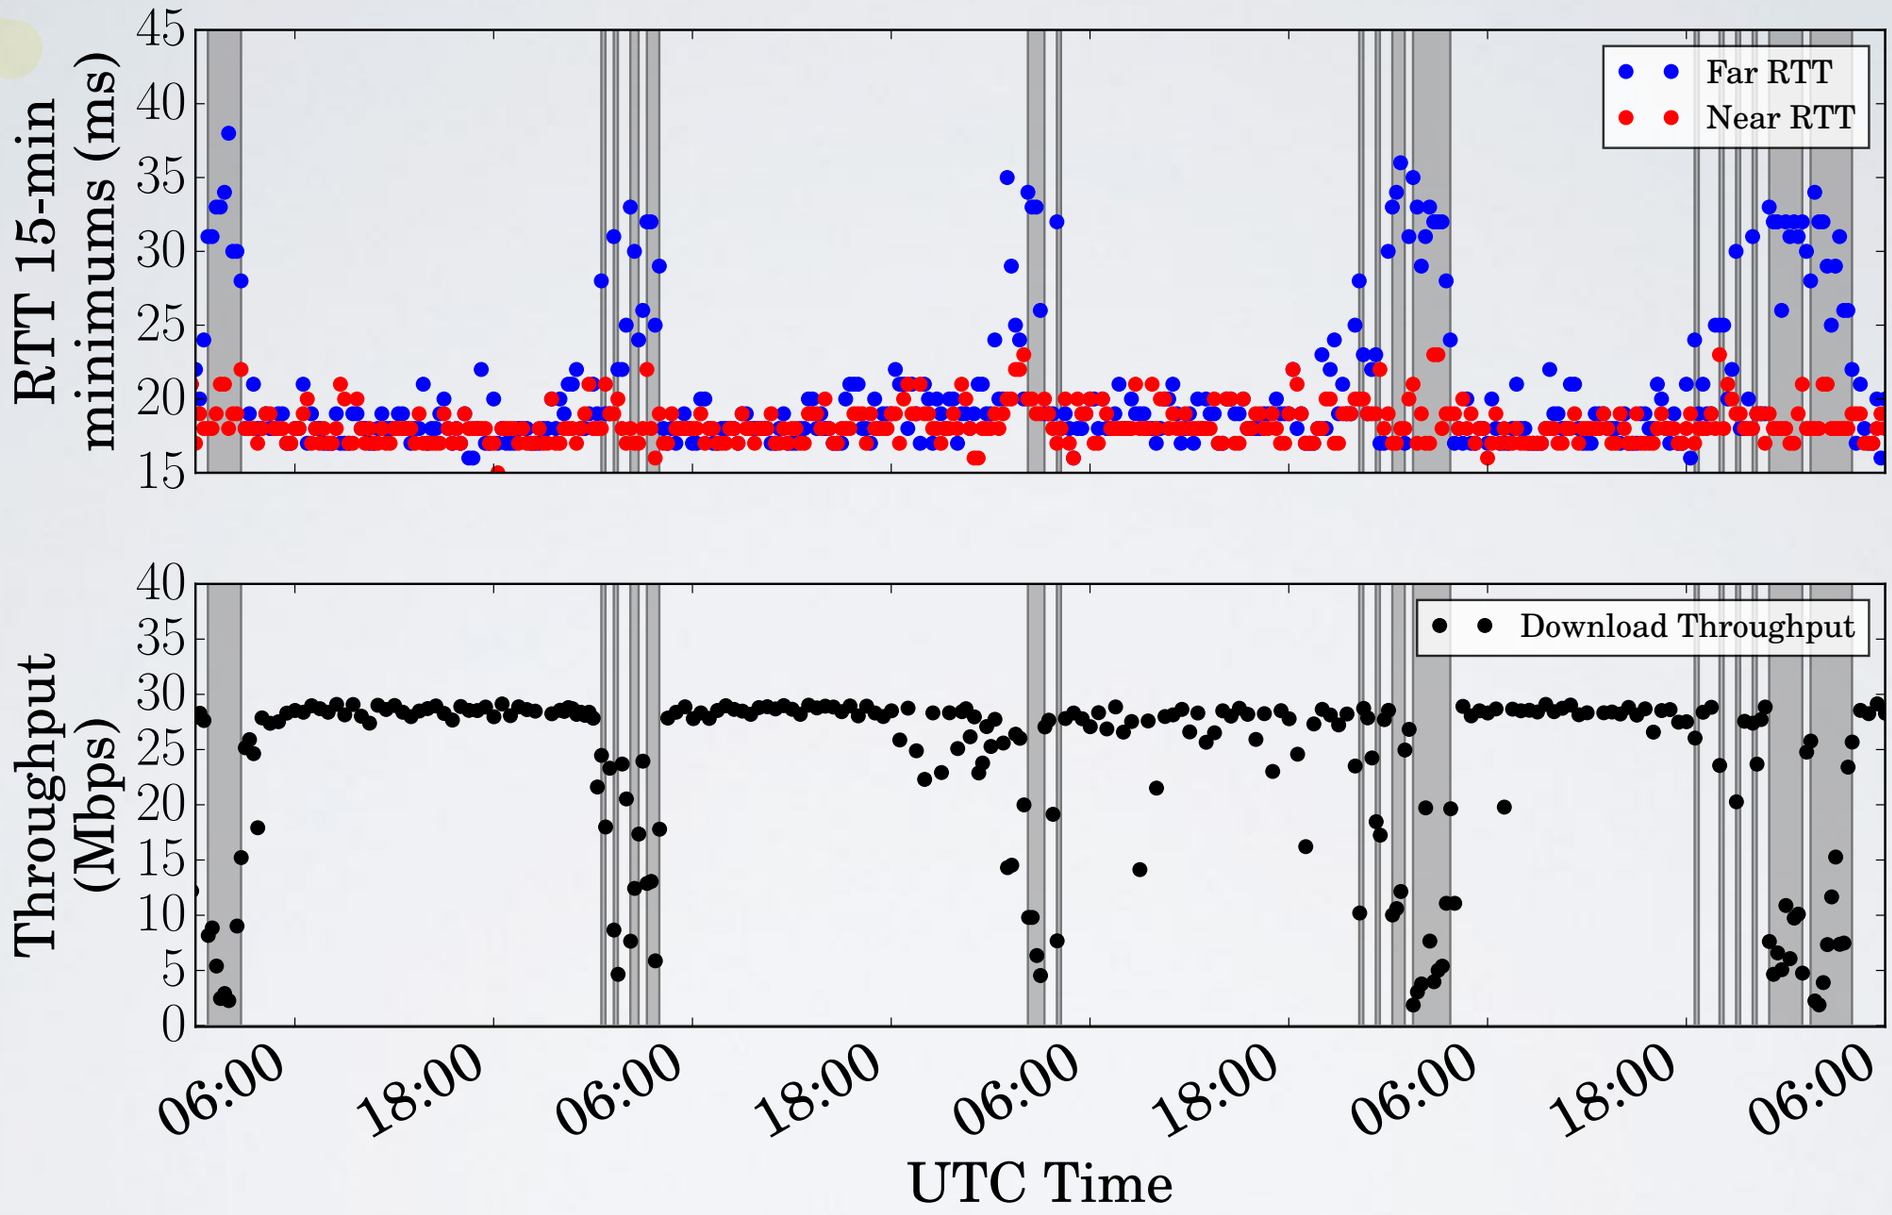

Does TSLP inference correlate with throughput?

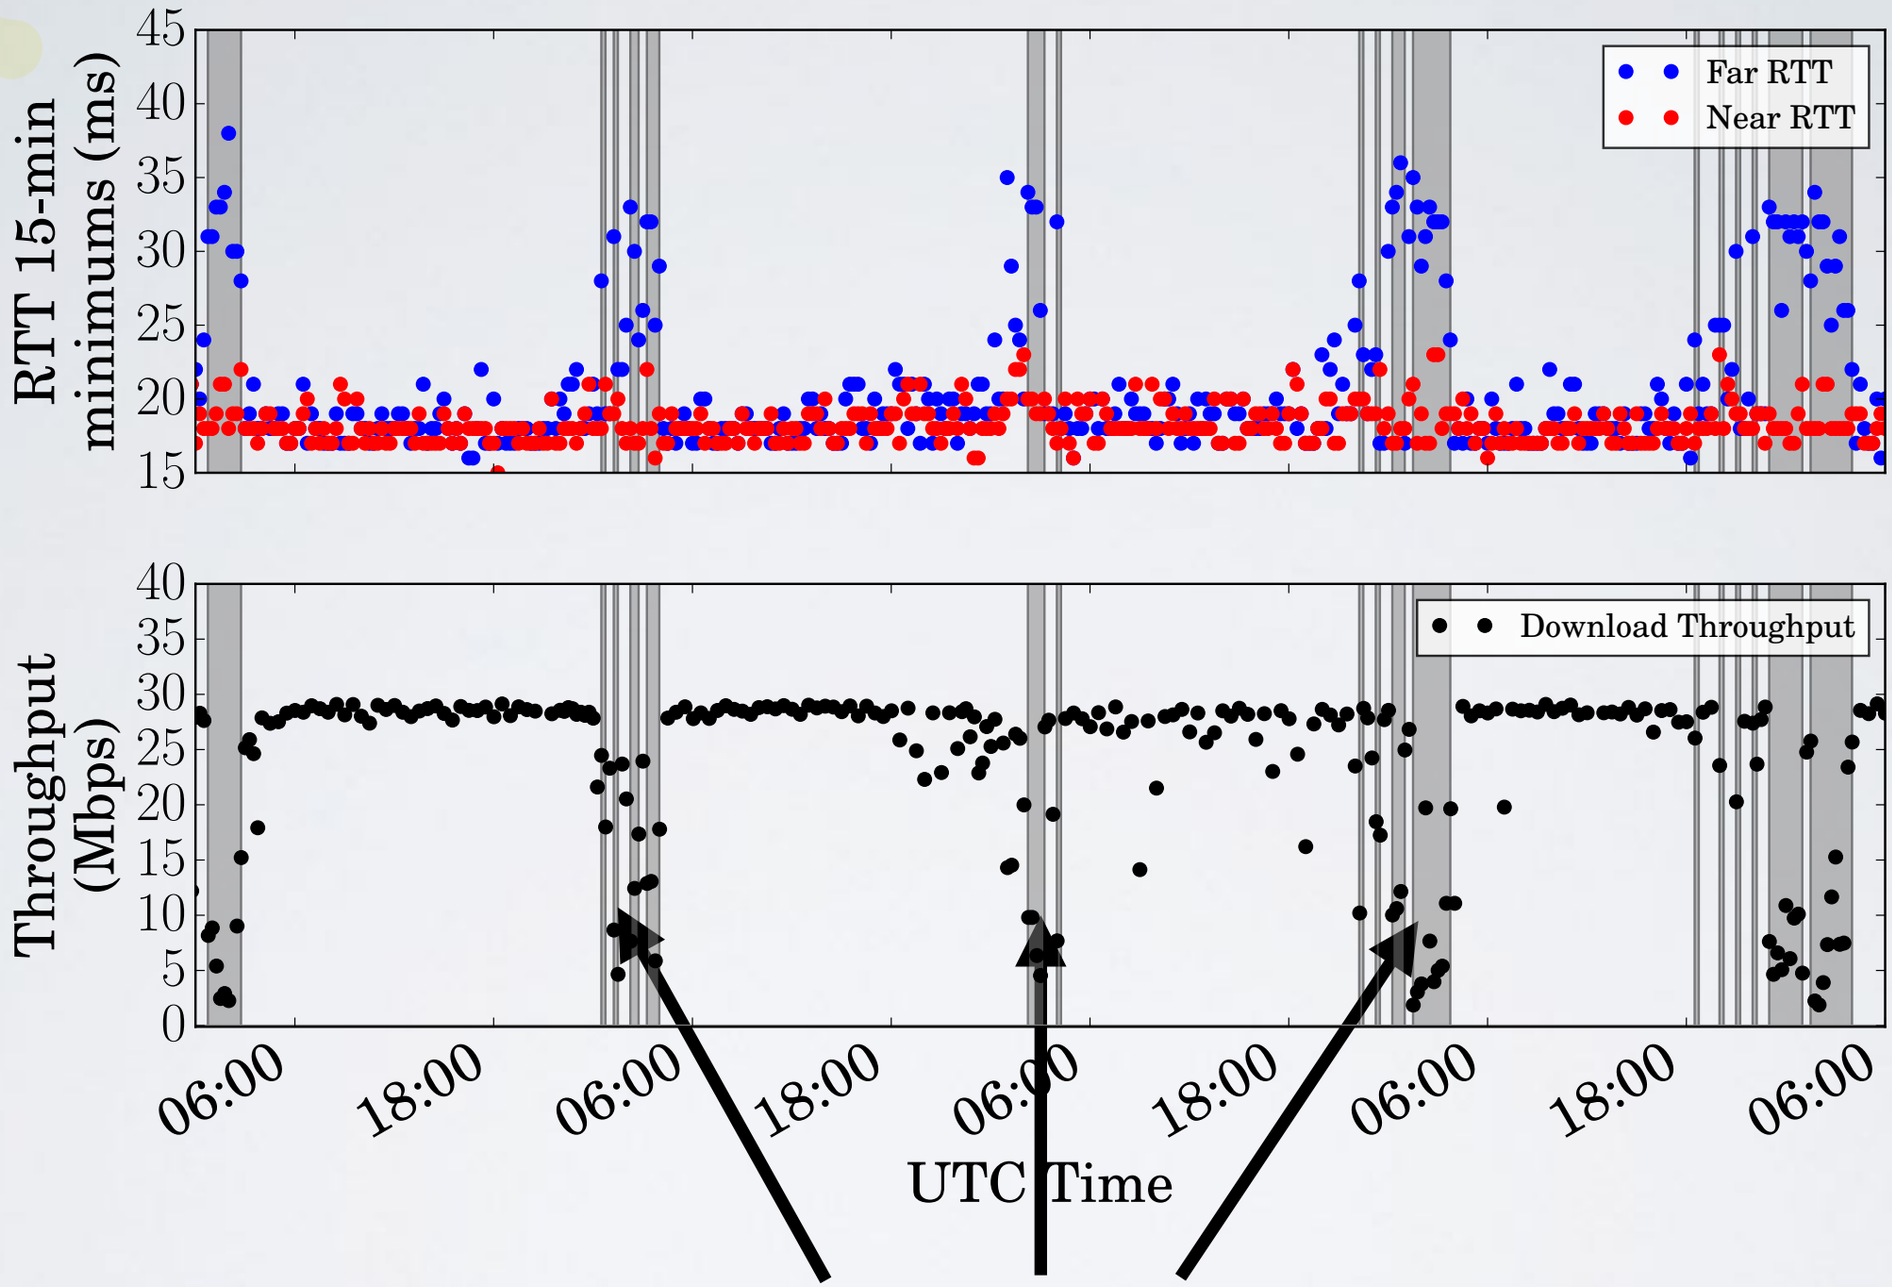

Lower throughput during periods inferred congested

caida

Does TSLP inference correlate with throughput?

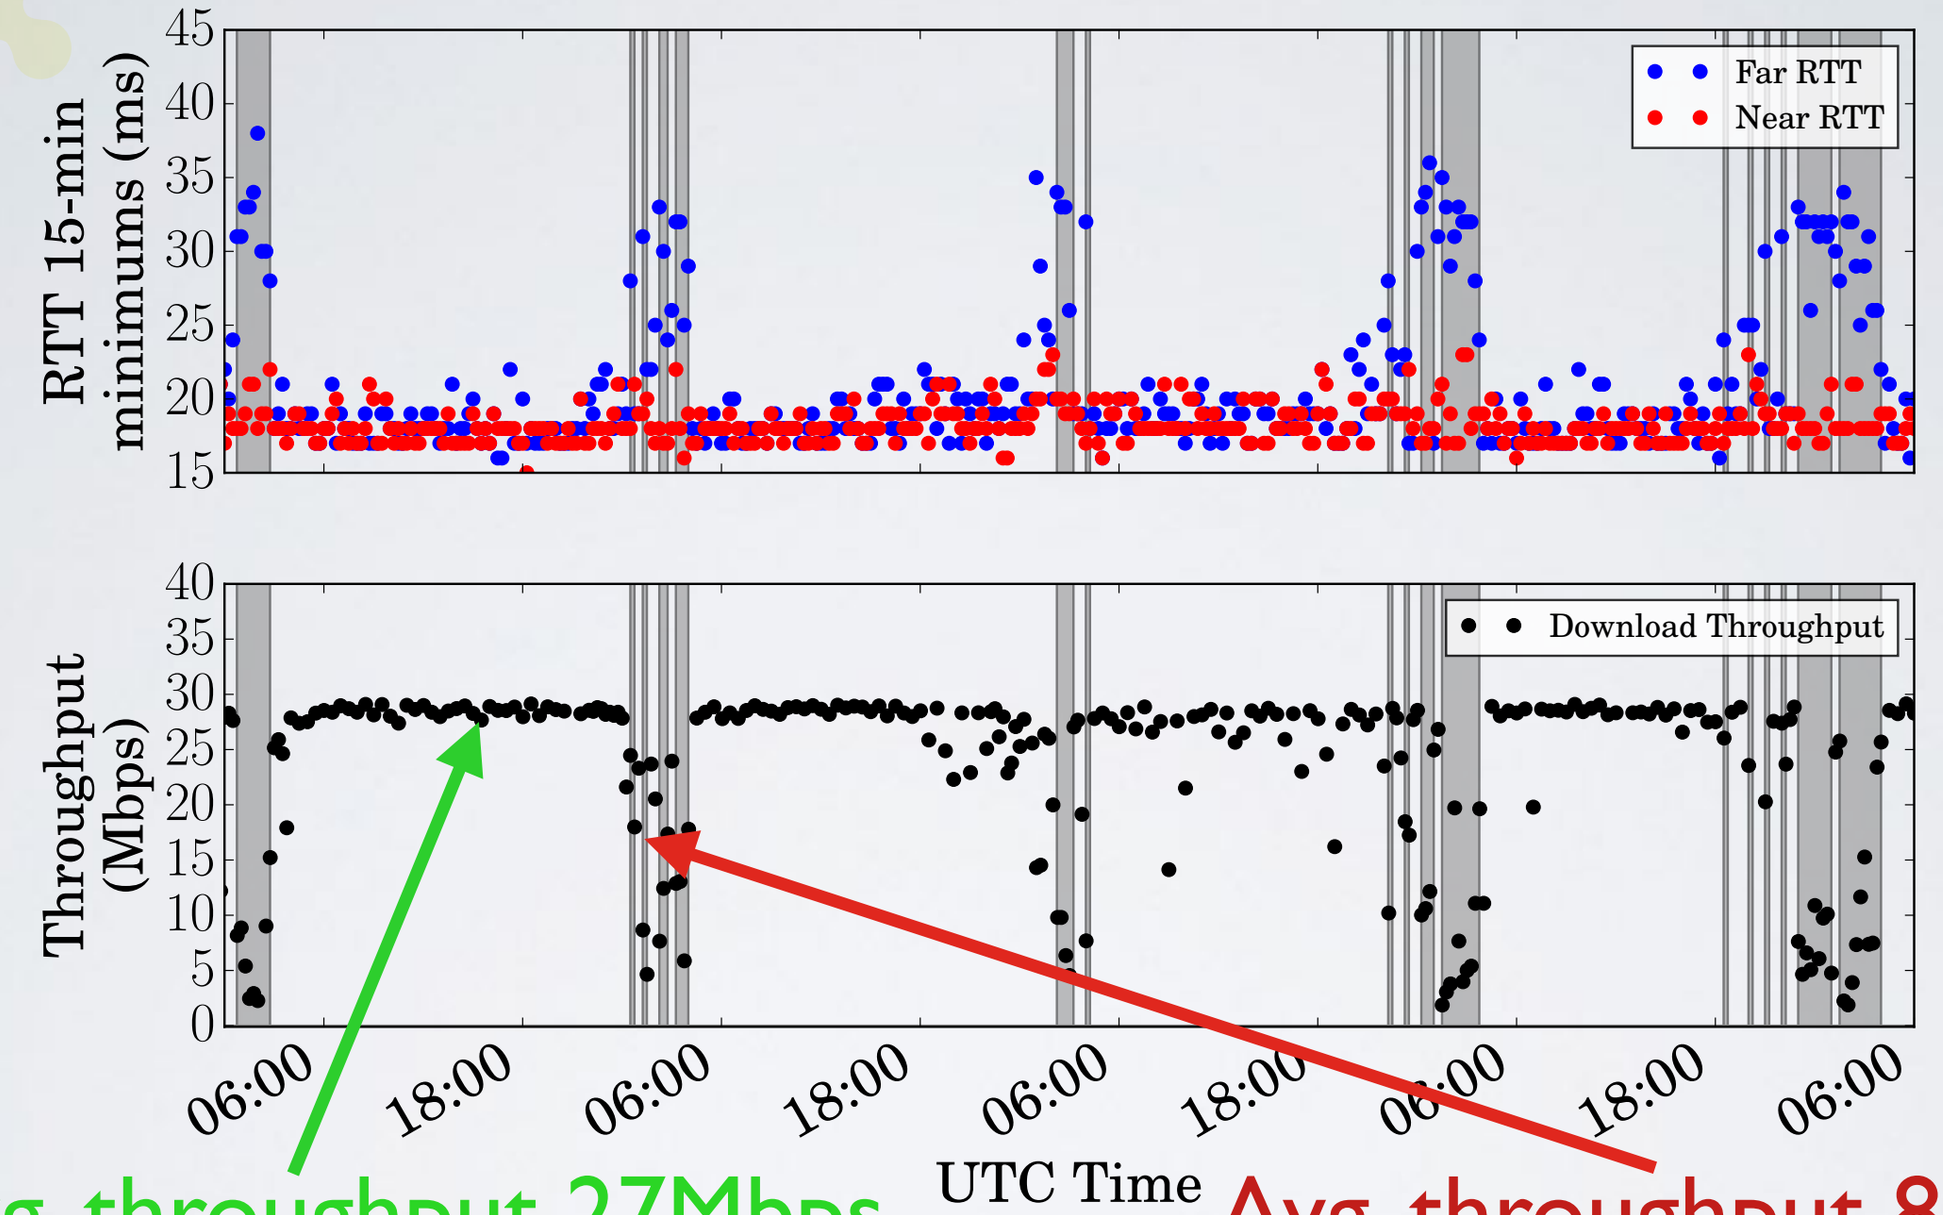

Avg. throughput 27Mbps during uncongested periods

Avg. throughput 8Mbps during congested periods

caida

Validation: Operator Feedback

caida

Validation: Operator Feedback

- Validated our inferences with operators from two large U.S. access ISPs

Validation: Operator Feedback

- Validated our inferences with operators from two large U.S. access ISPs
- ISP A: 7 links (all inferred congested)
- ISP B: 20 links (10 inferred congested, 10 uncongested)

Validation: Operator Feedback

- Validated our inferences with operators from two large U.S. access ISPs
- ISP A: 7 links (all inferred congested)
- ISP B: 20 links (10 inferred congested, 10 uncongested)
- **Our inferences were correct in each case: no false positives or false negatives**

caida

Longitudinal Study

- Collecting data since March 2016
- Focused on interdomain links of 8 large access ISPs in the U.S. to their transit providers and peers from Mar 2016 to Dec 2017
- Driving questions:
  - How prevalent is interdomain congestion?
  - Which transit/content providers are most often congested to access providers?
  - Can we characterize trends over time?

What Did We Find (so far)?

- No evidence of widespread (pervasive) congestion between Mar 2016 and Dec 2017
- Small fraction of peers of the 8 studied access providers showed evidence of congestion
- Certain transit providers (e.g., TATA) and content providers (e.g. Google) most often showed evidence of congestion
- Interesting dynamics of interdomain congestion

See paper for details

caida

Percent of congested day-links over time

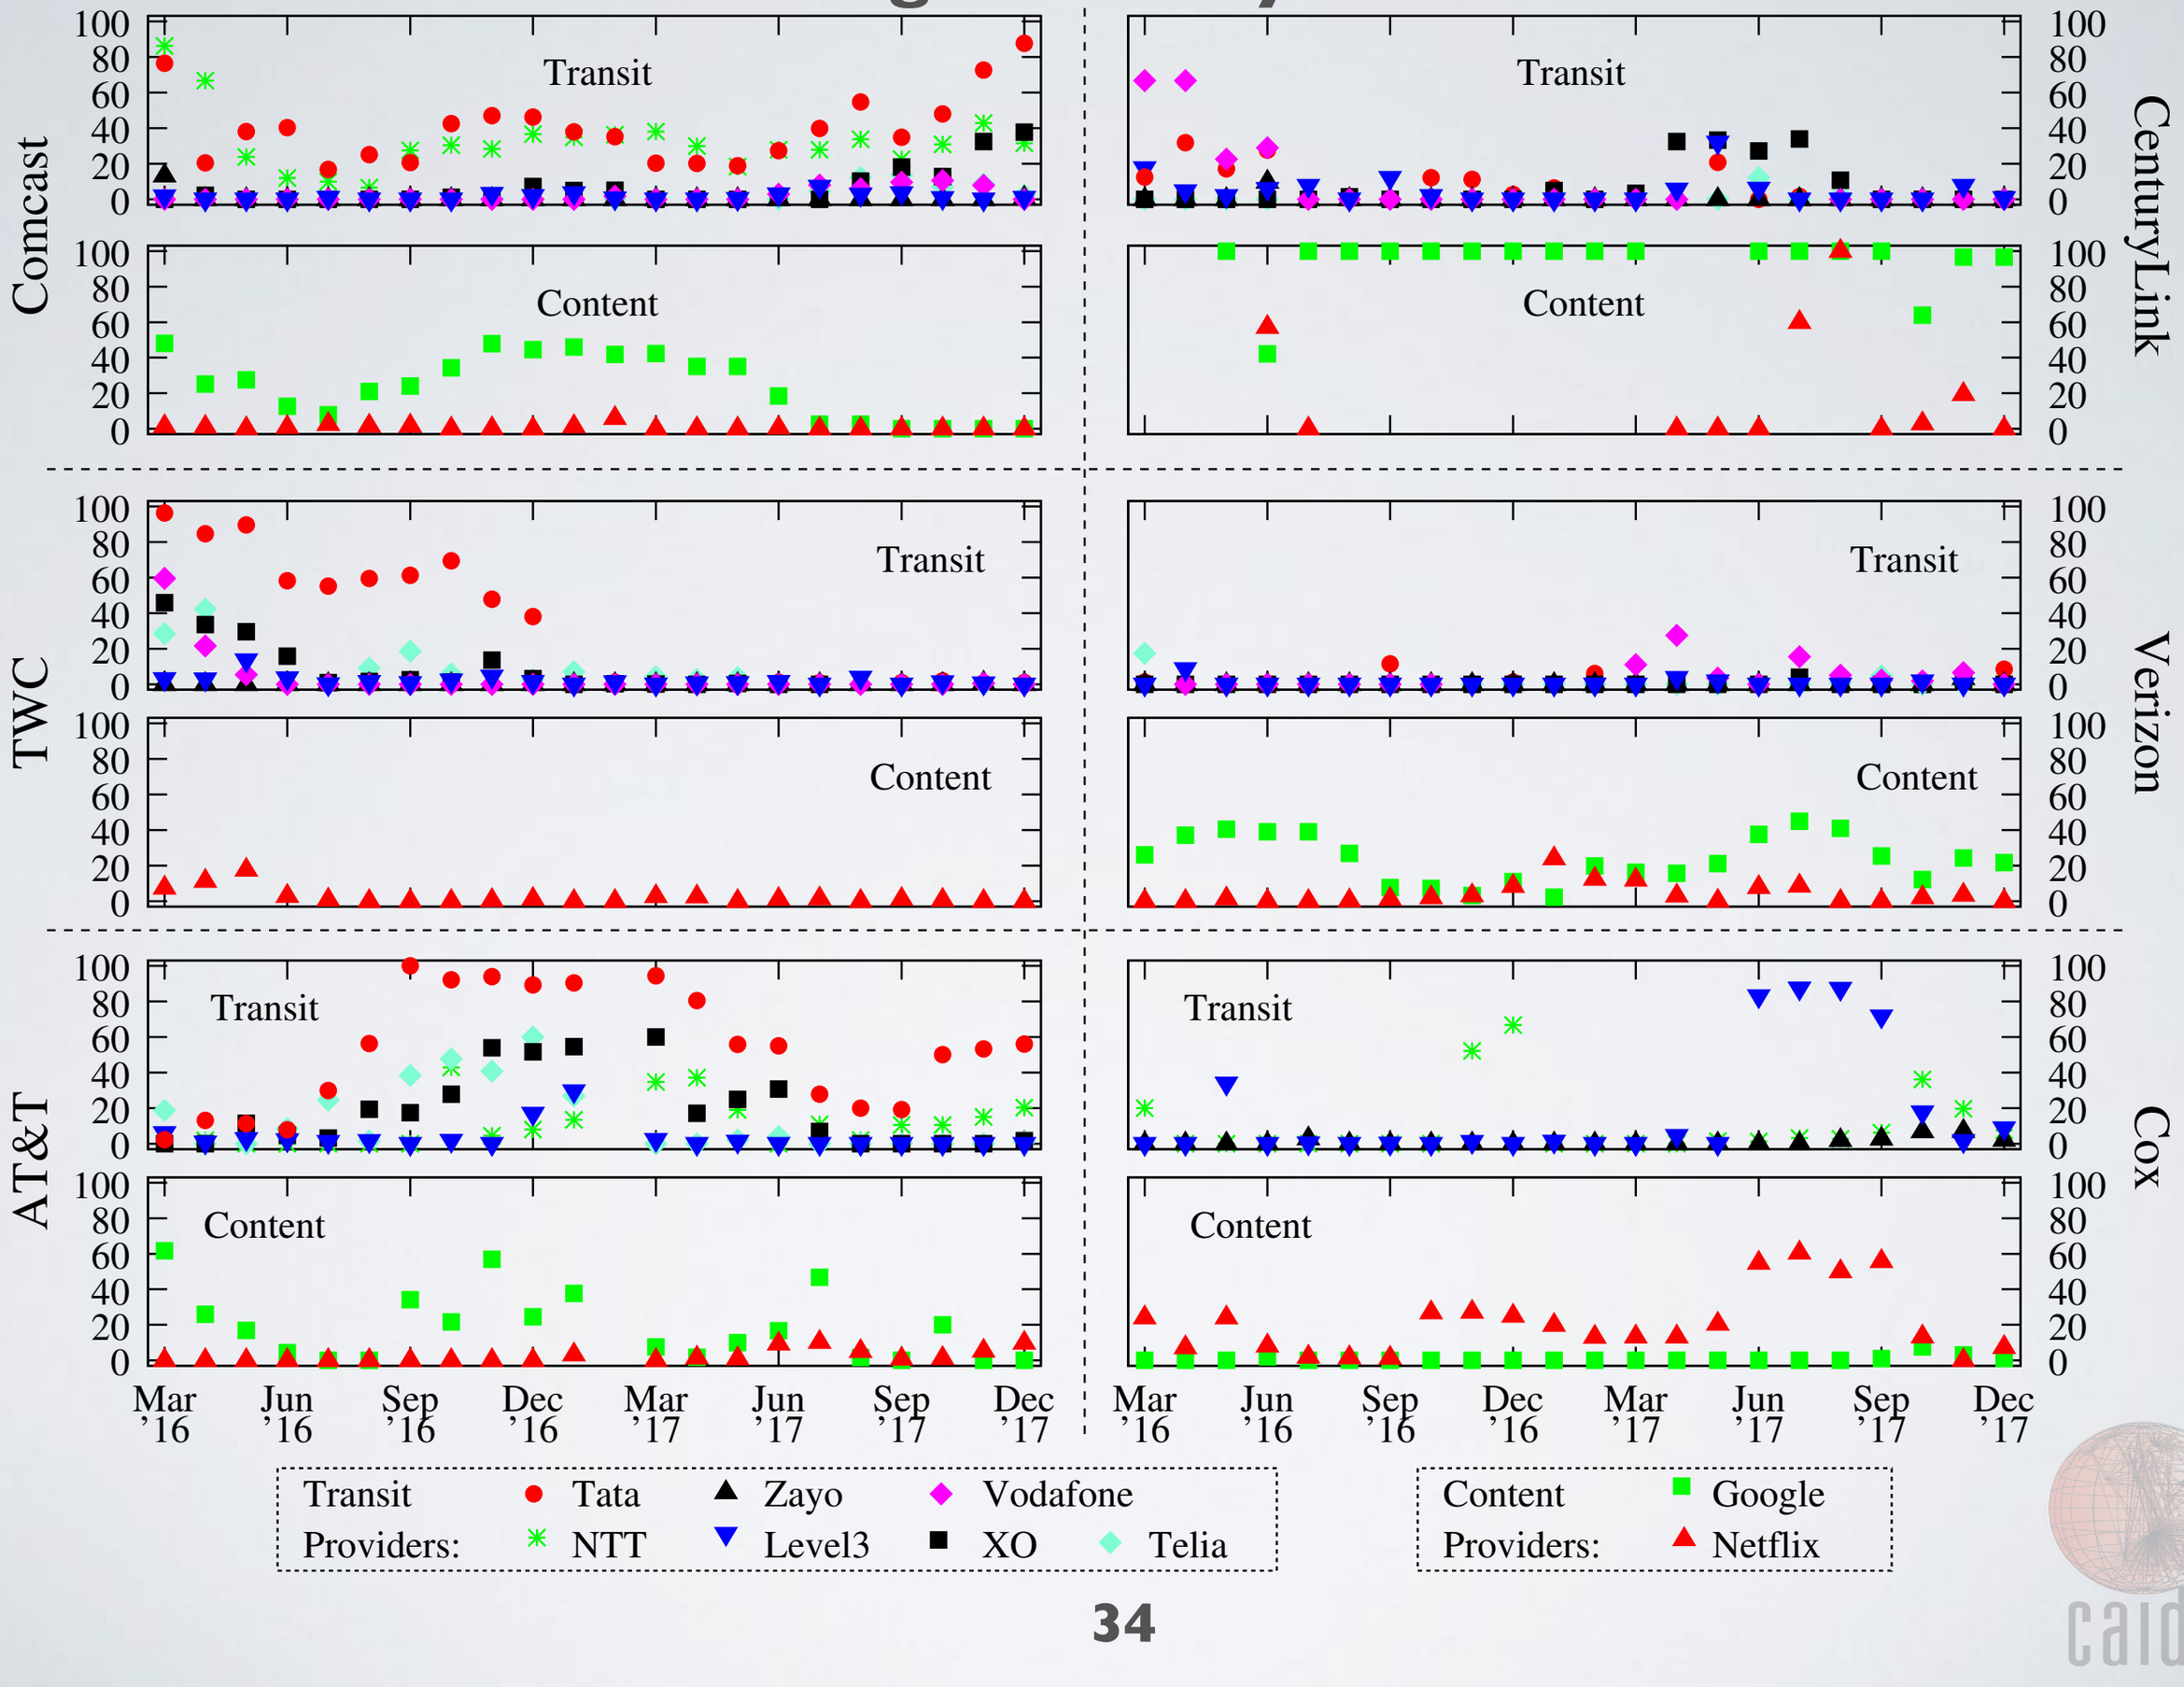

caida

Percent of congested day-links over time

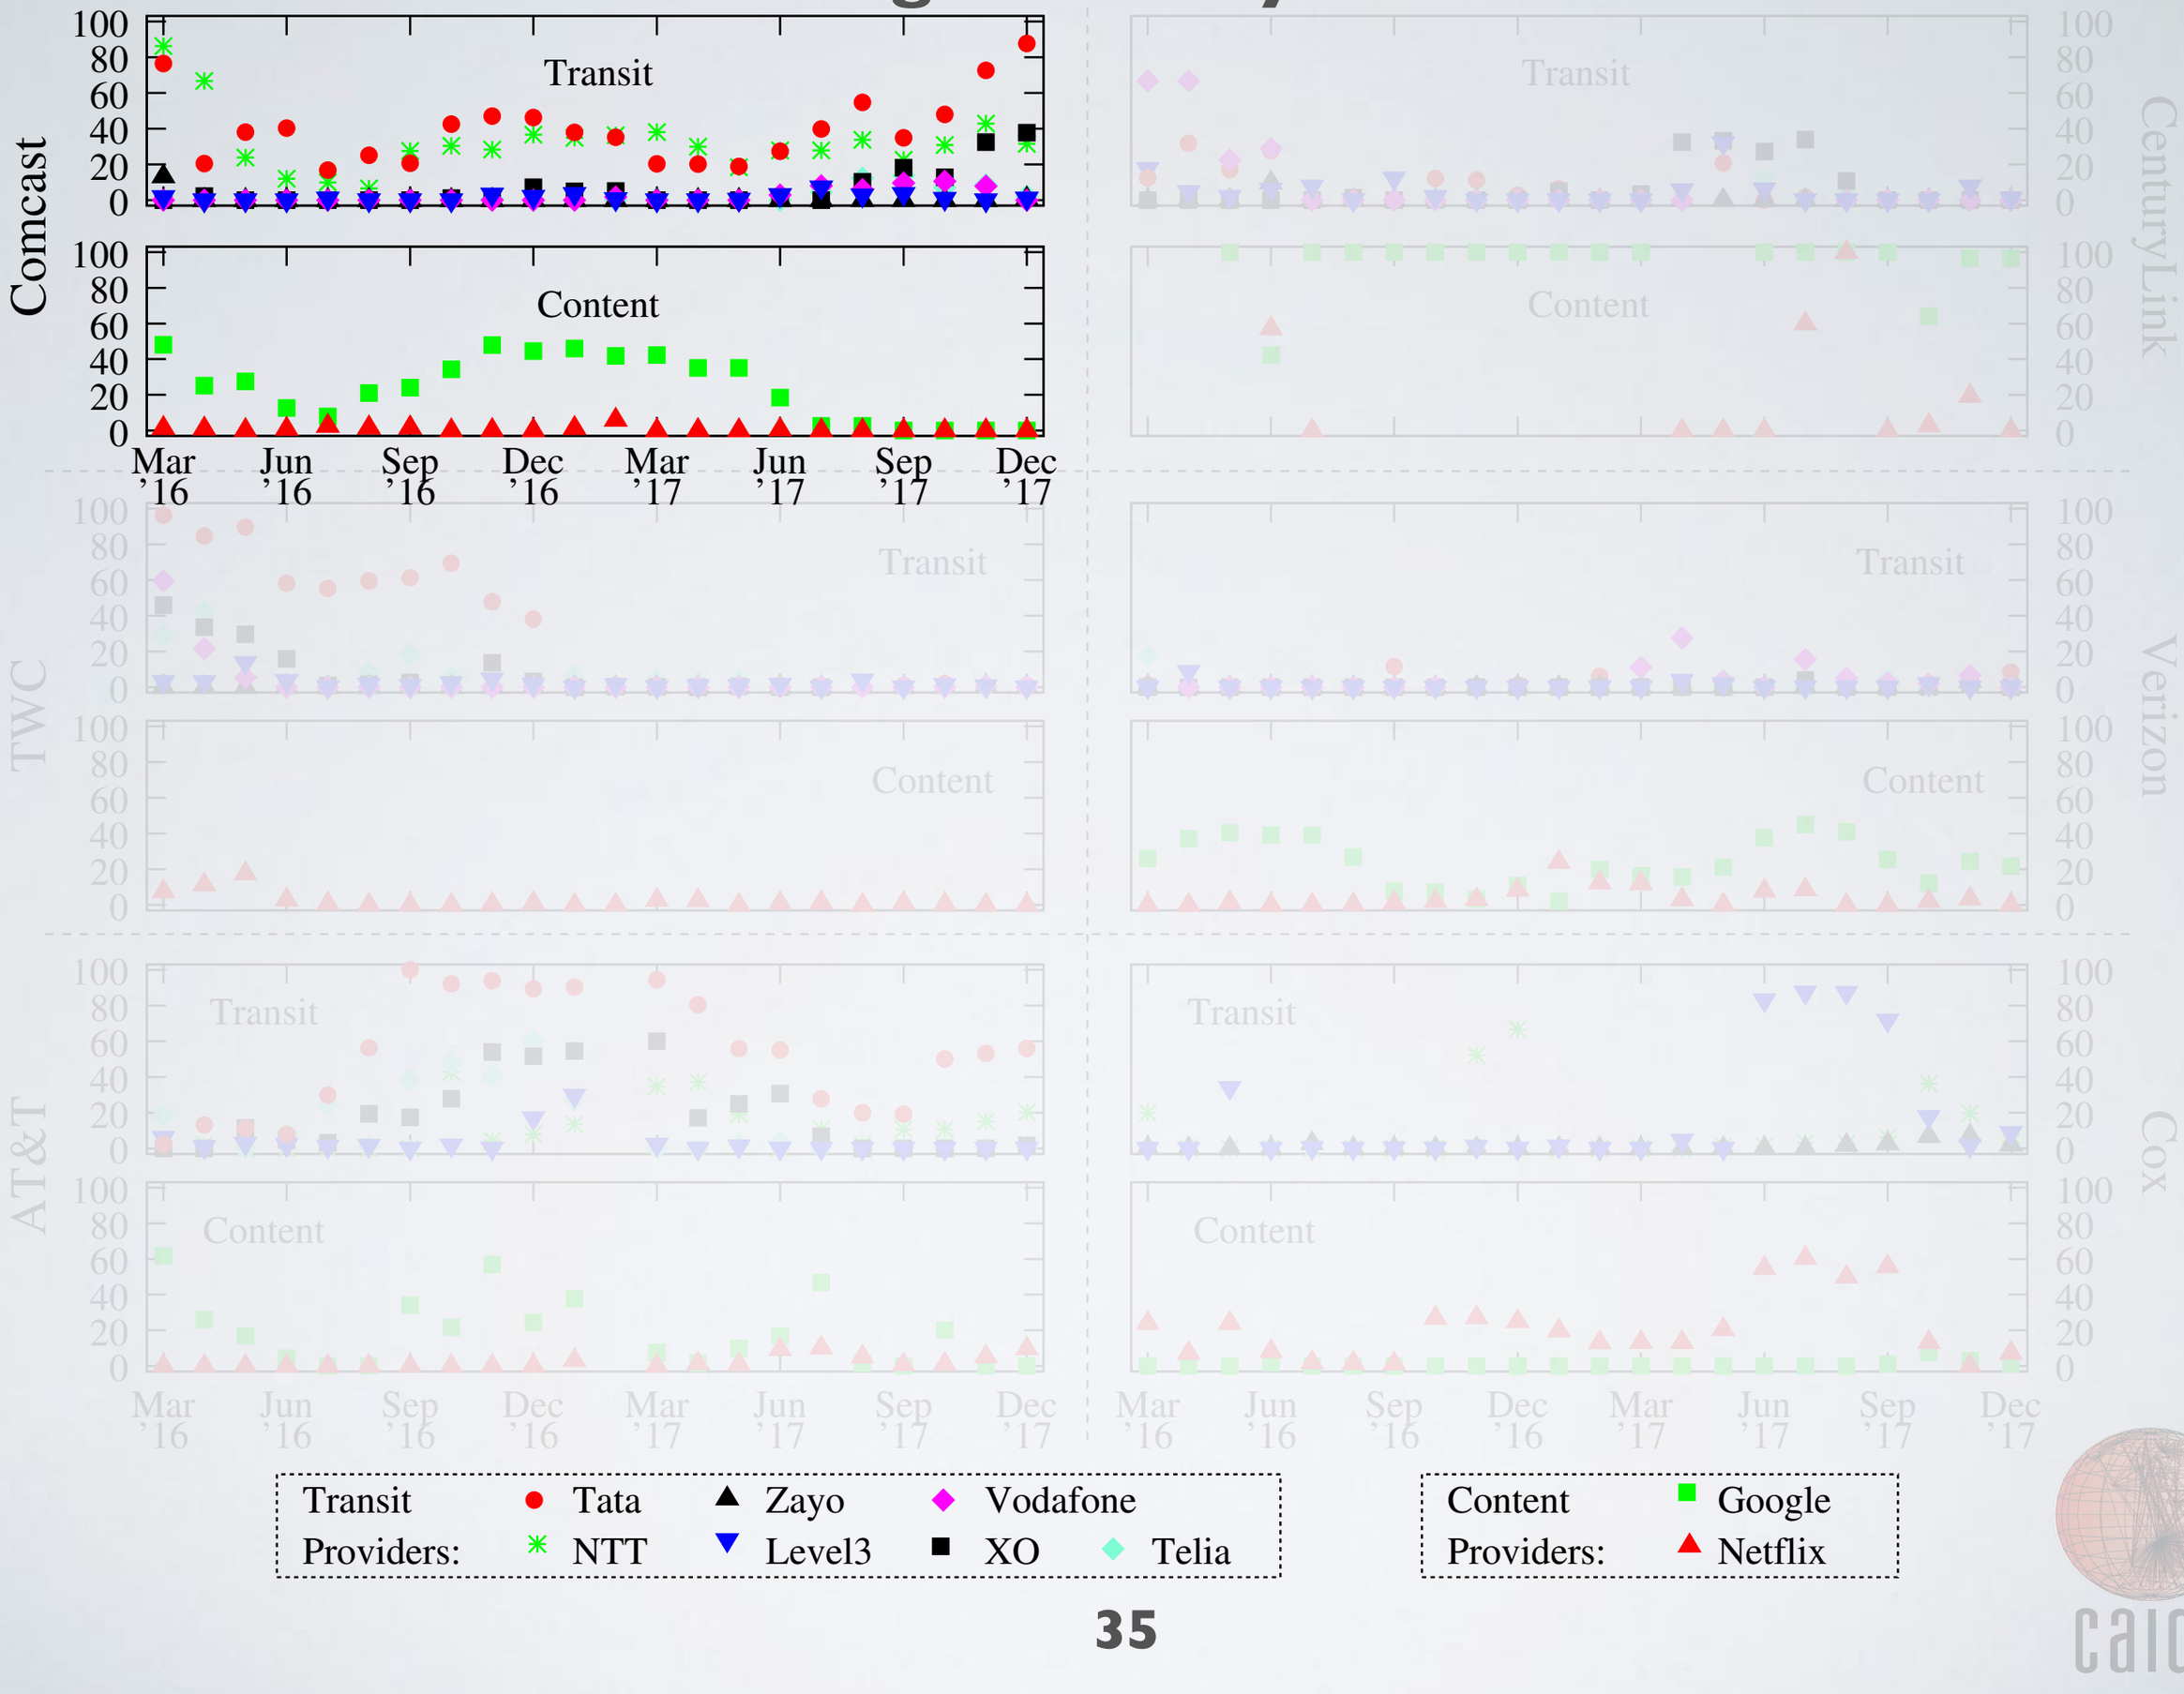

Percent of congested day-links over time

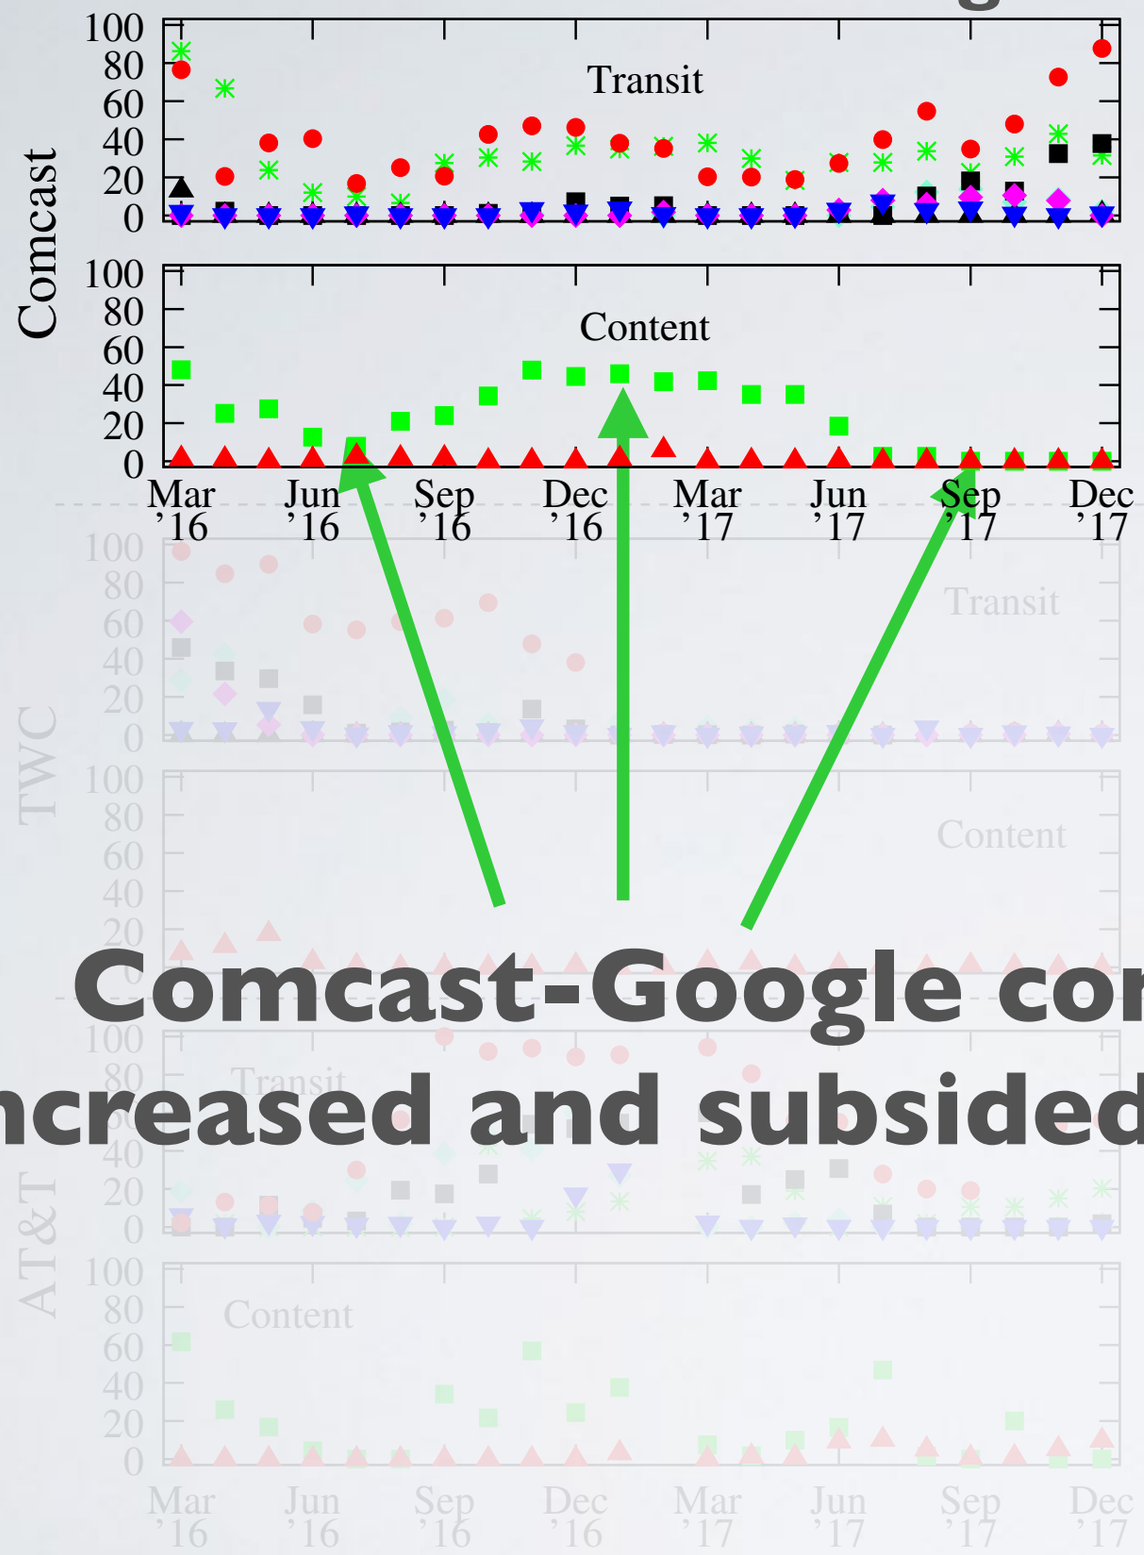

Comcast-Google congestion increased and subsided over time

Transit Providers: Tata, NTT, Zayo, Level3, Vodafone, XO, Telia

Content Providers: Google, Netflix

caida

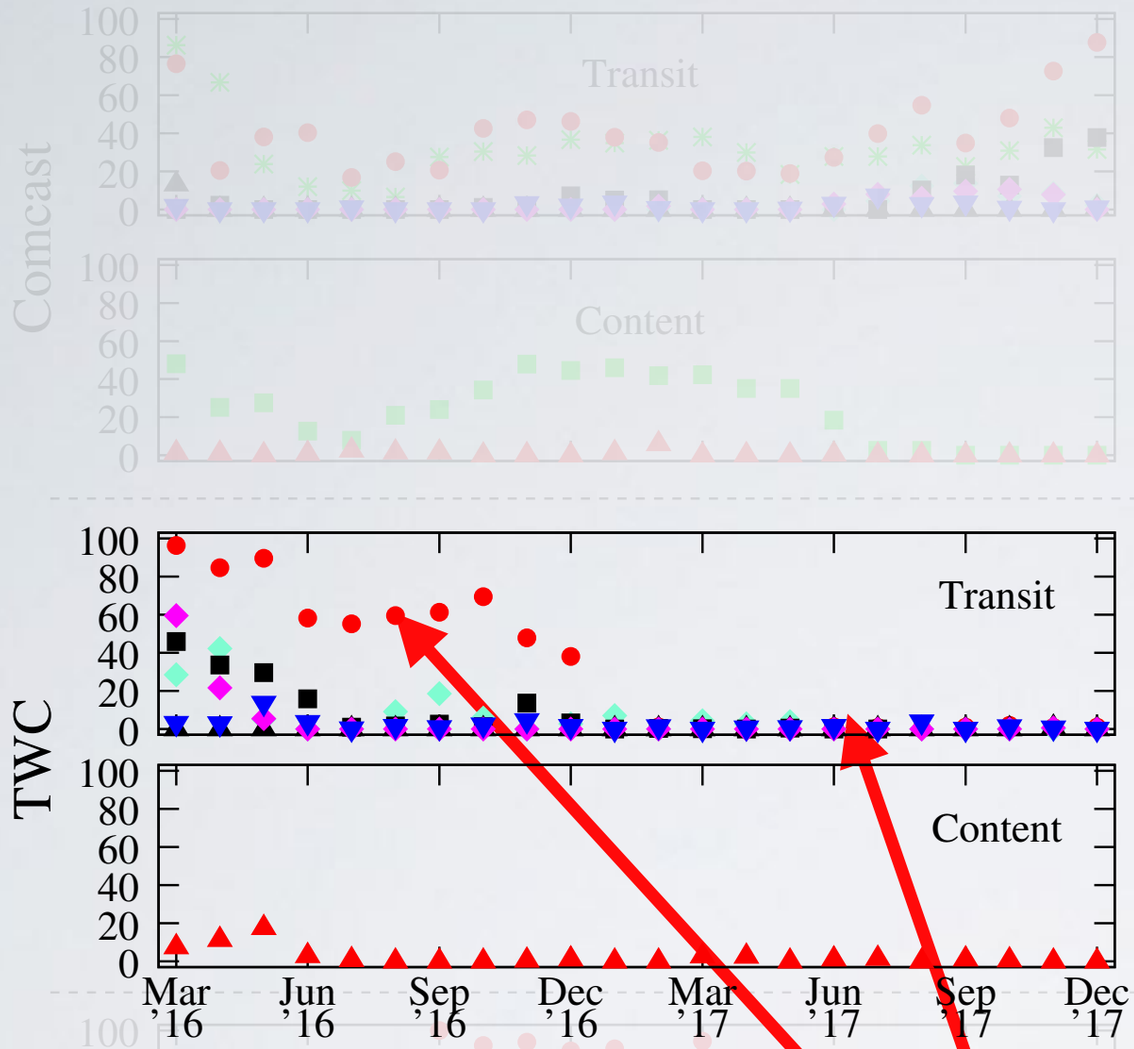

Time Warner - TATA congestion subsided in December 2016

Transit Providers: Tata, Zayo, Vodafone, NTT, Level3, XO, Telia

Content Providers: Google, Netflix

Percent of congested day-links over time

AT&T - TATA congestion increased and subsided over time

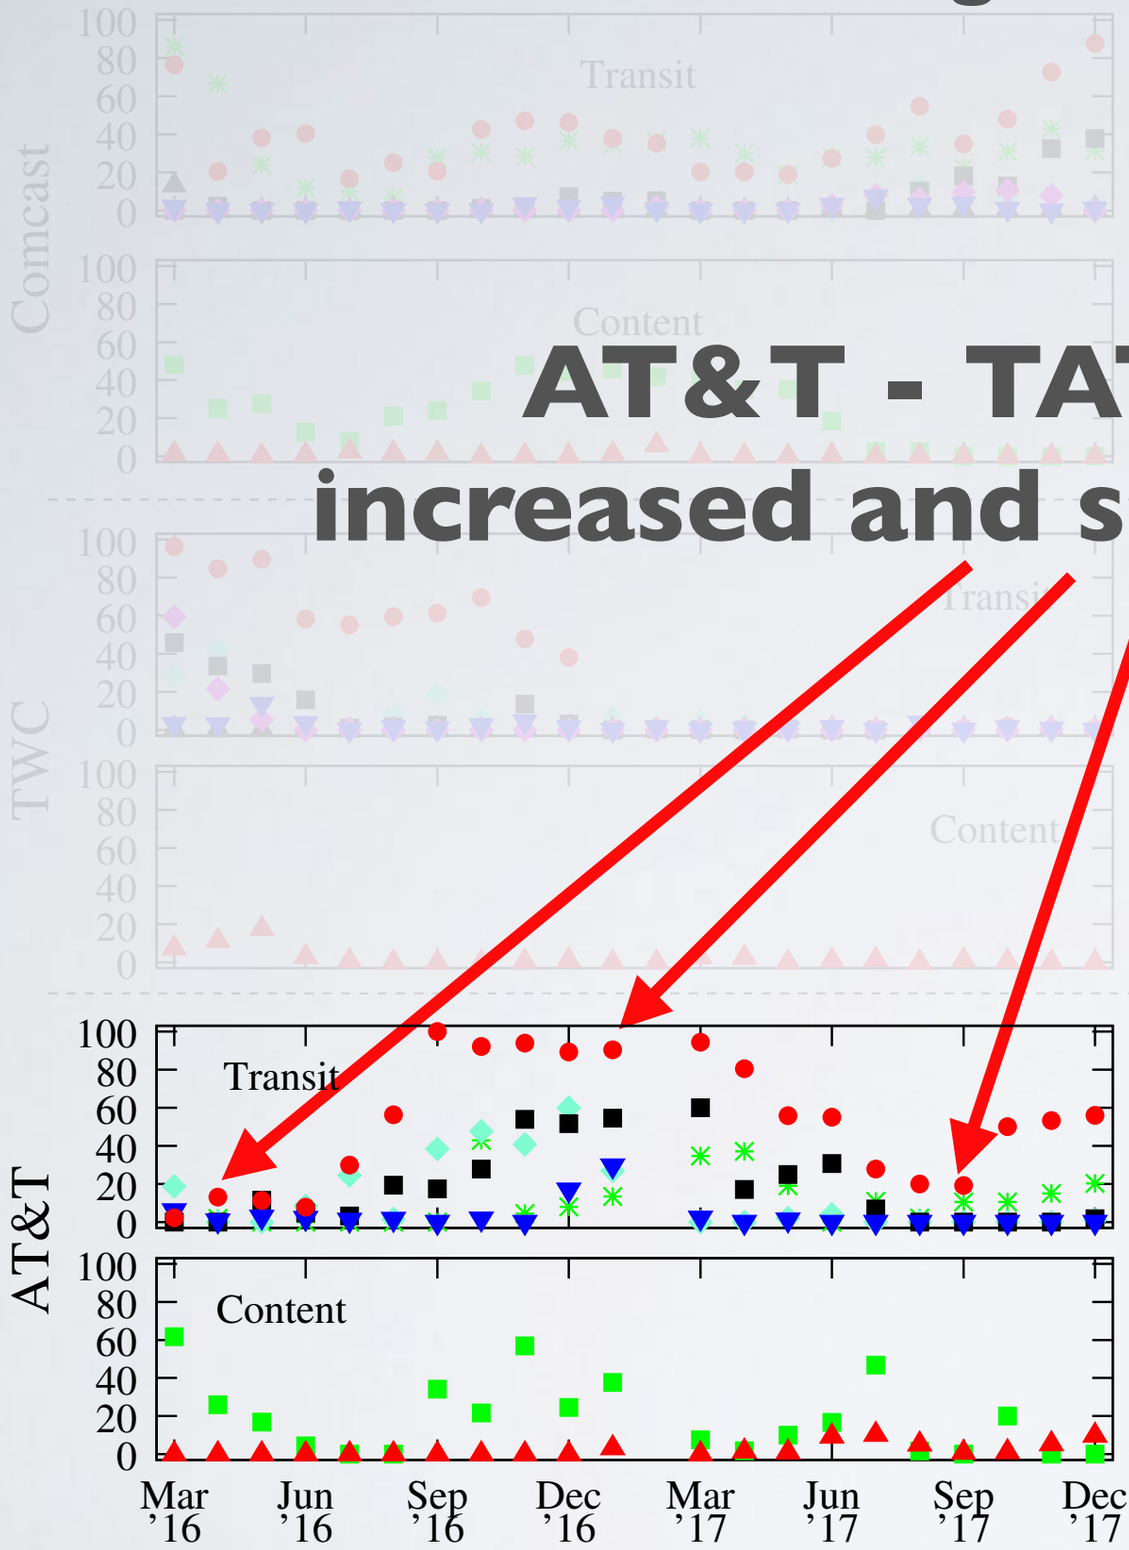

Transit Providers: Tata, NTT, Zayo, Level3, Vodafone, XO, Telia

Content Providers: Google, Netflix

Interdomain congestion evolves over time.

Need for ongoing measurements!

caida

Public Access to Data

- We are publicly releasing our data via an interactive visualization system (based on Grafana)
- And API access to the time series data (based on InfluxDB)

Interactive Visualization

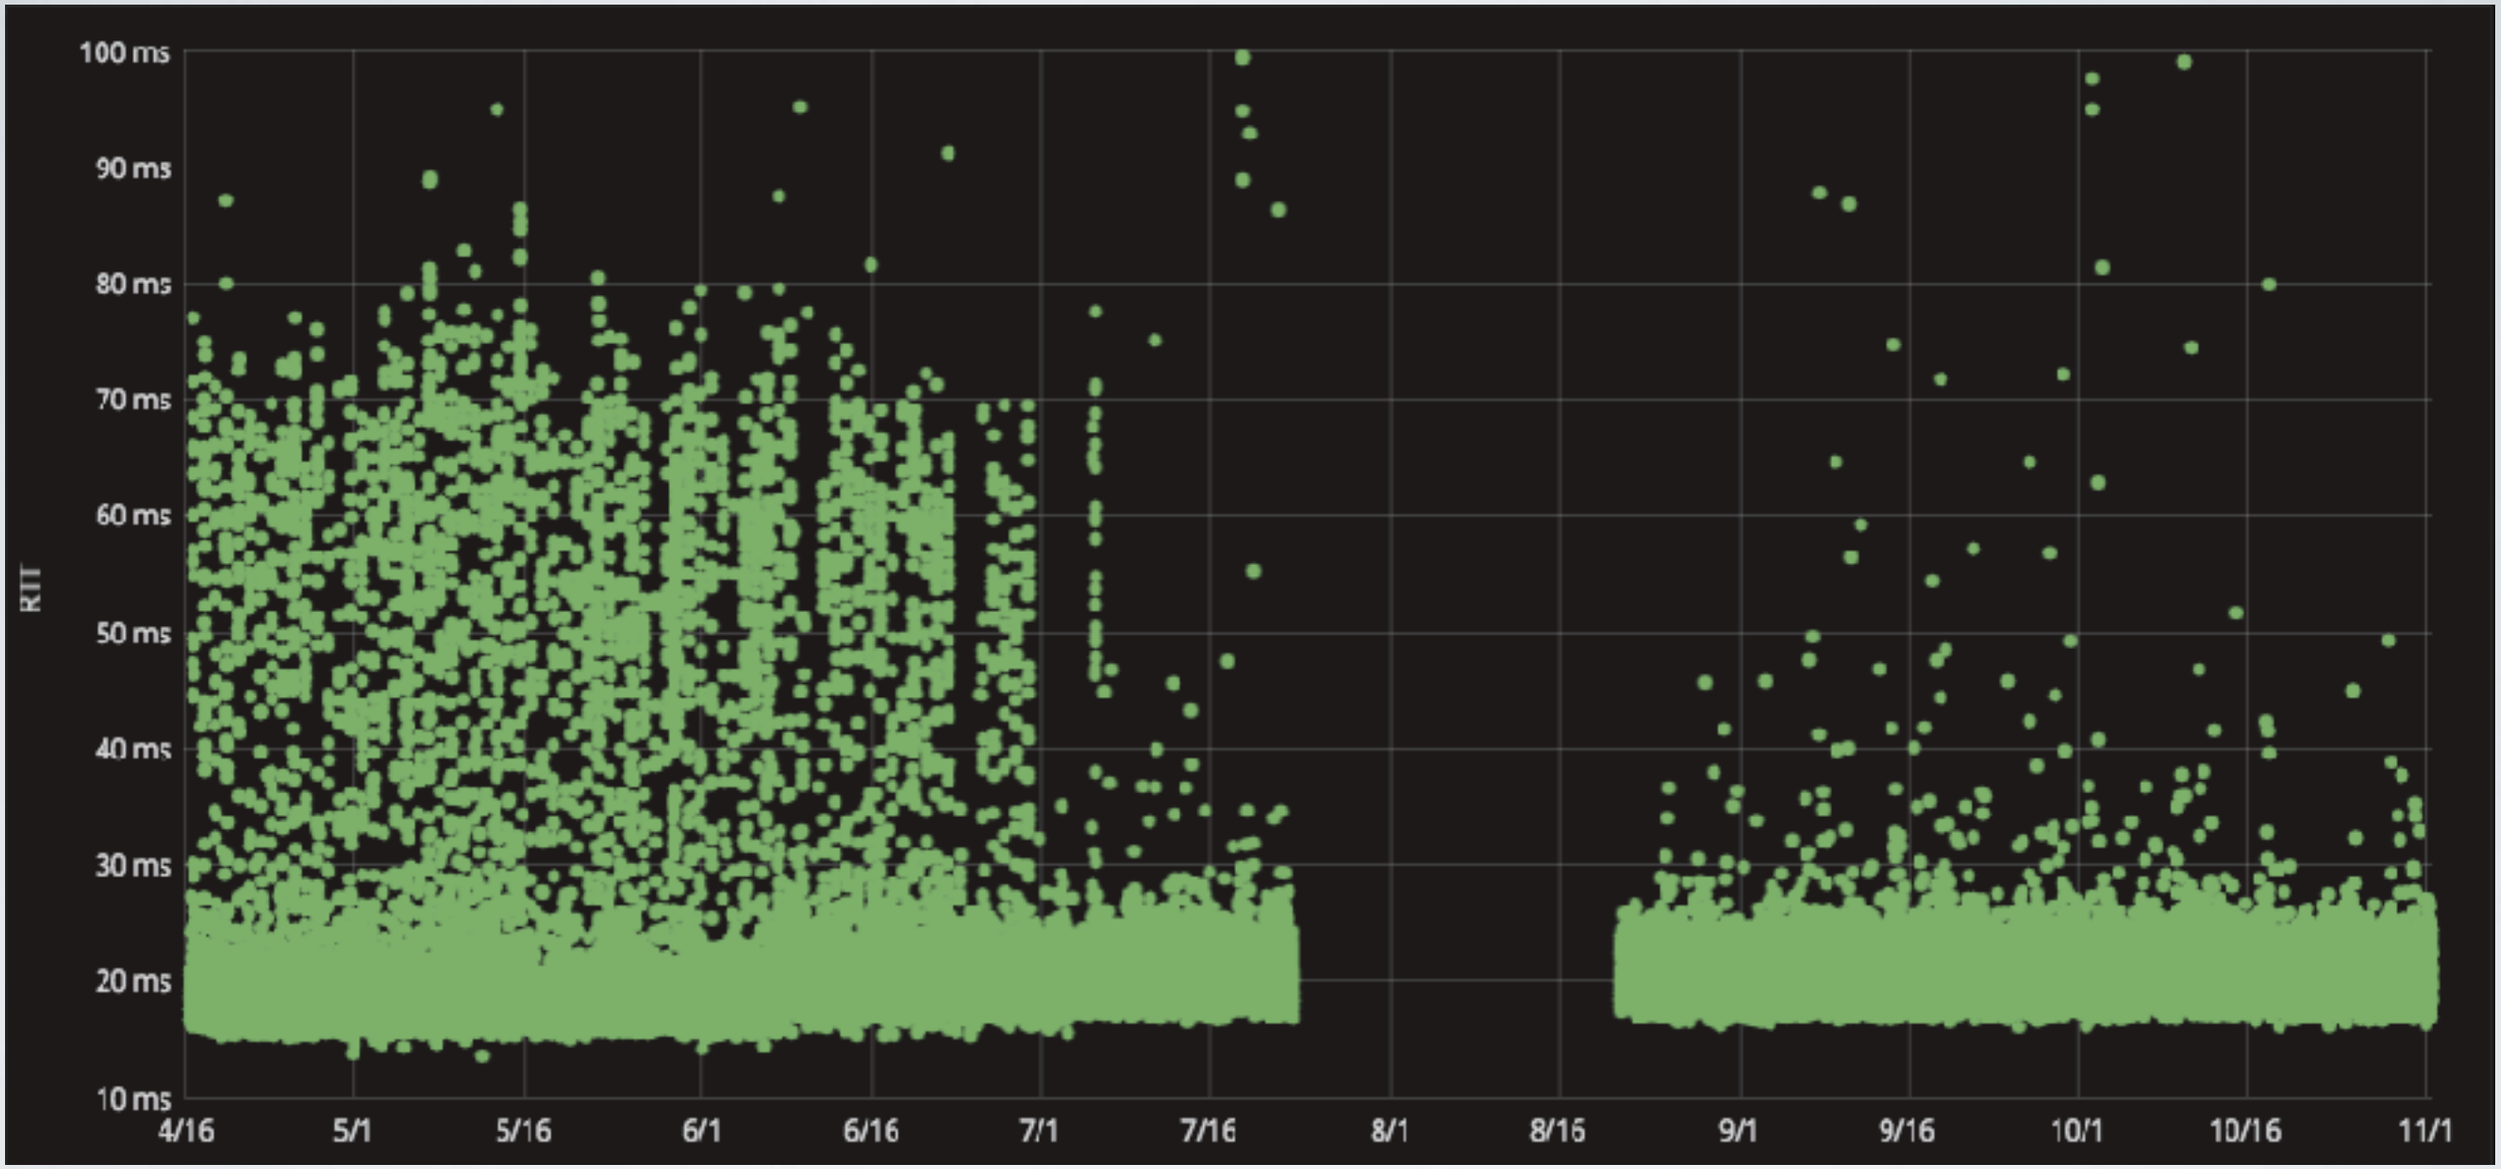

Longitudinal view of a single link, April - November 2017

caida

Interactive Visualization

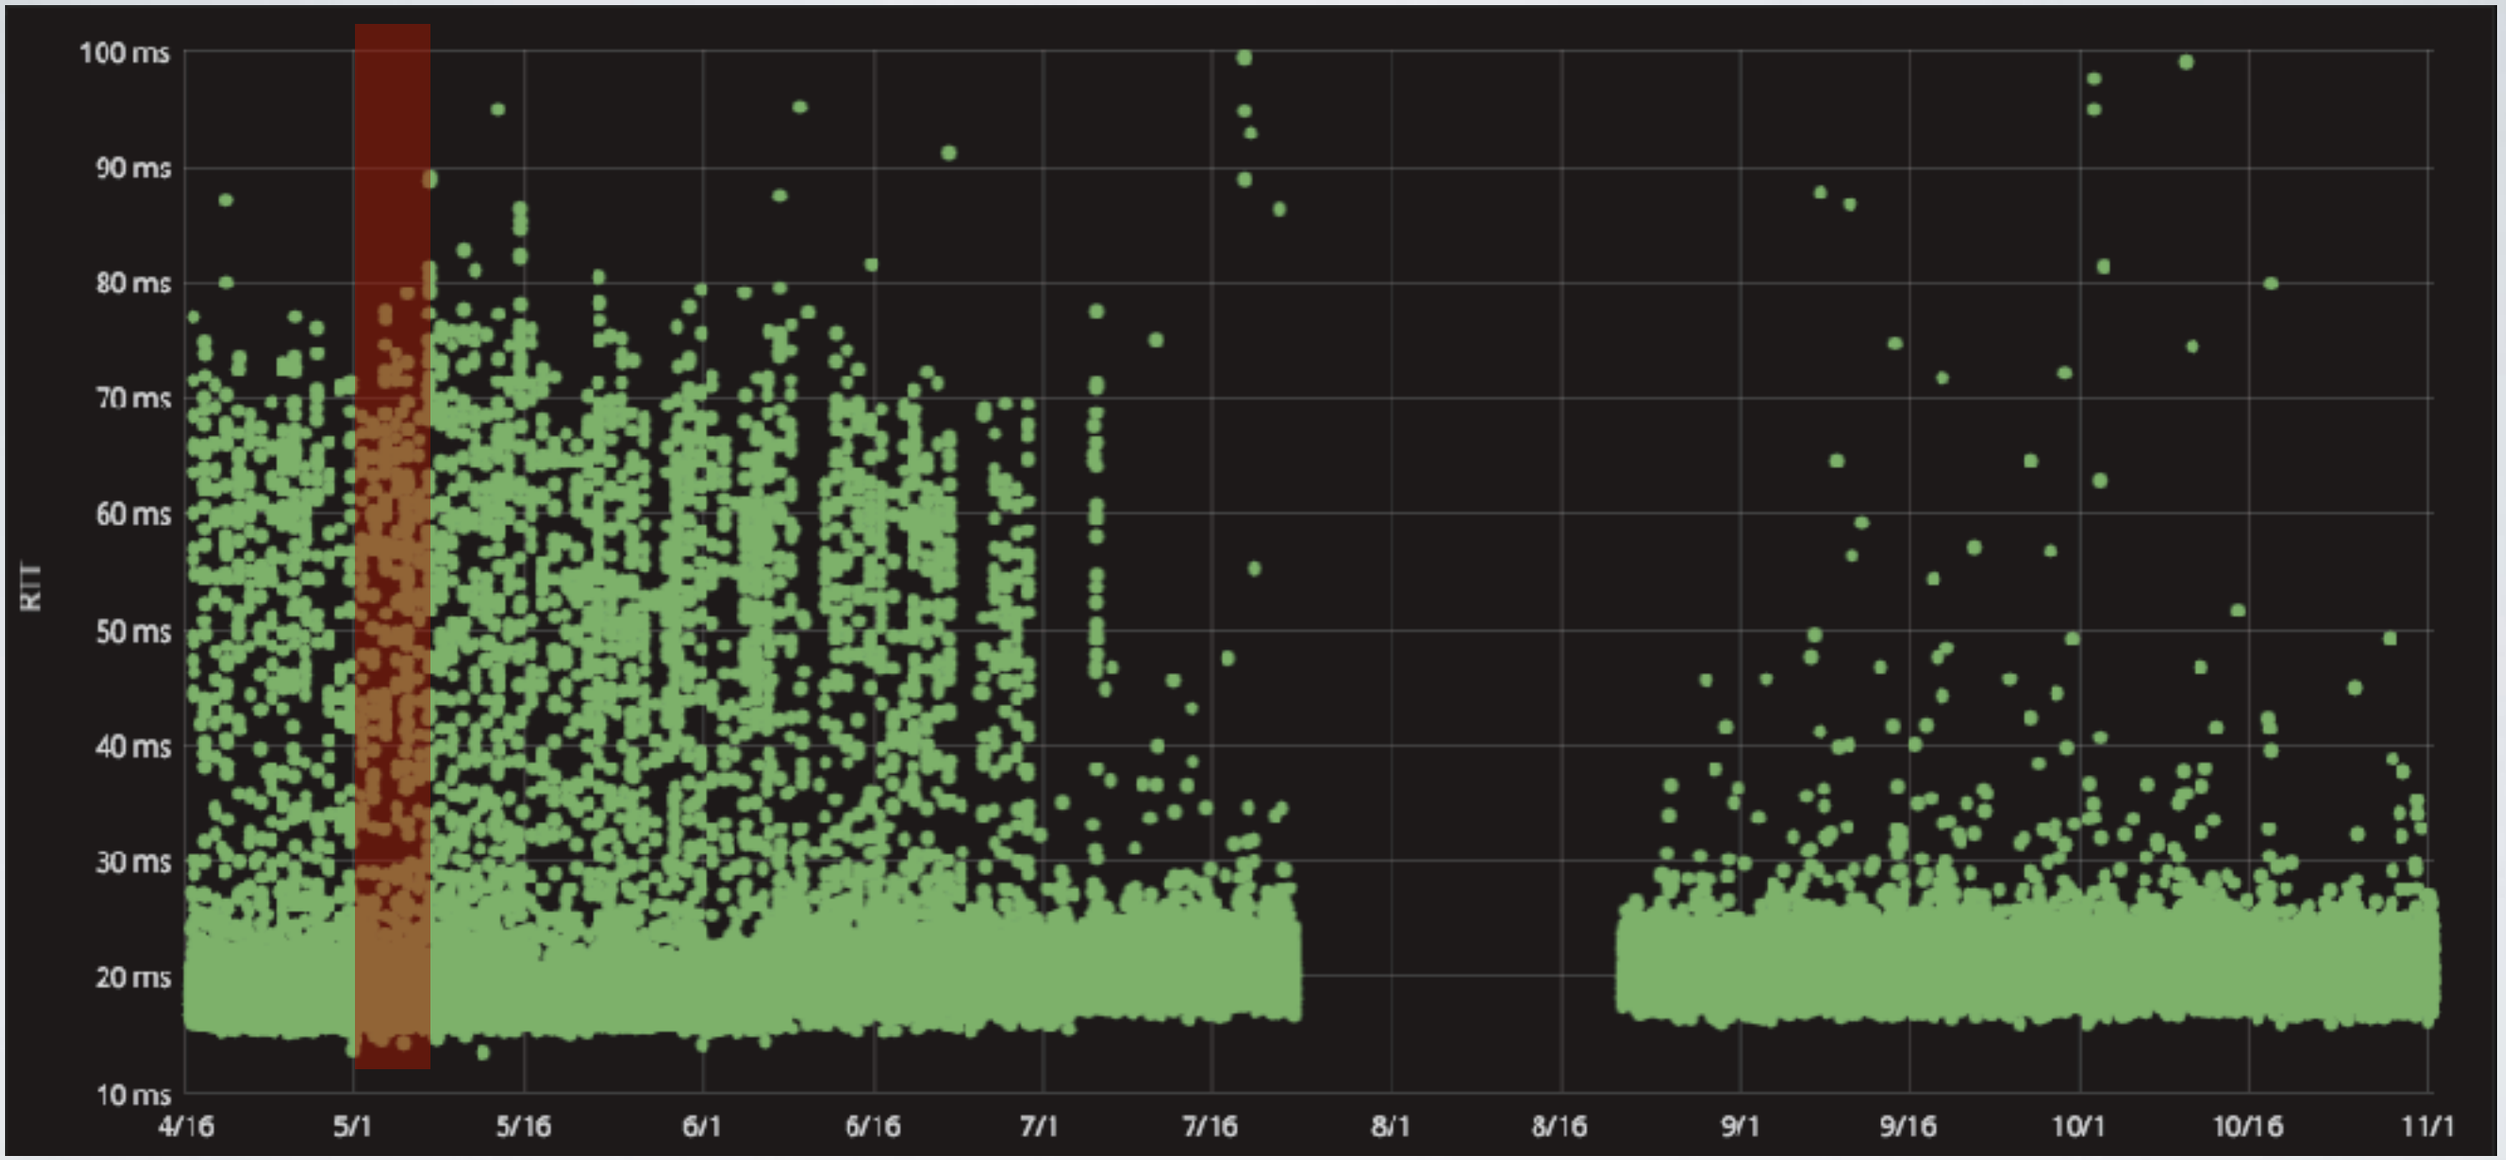

Zooming in for more detail

caida

Interactive Visualization

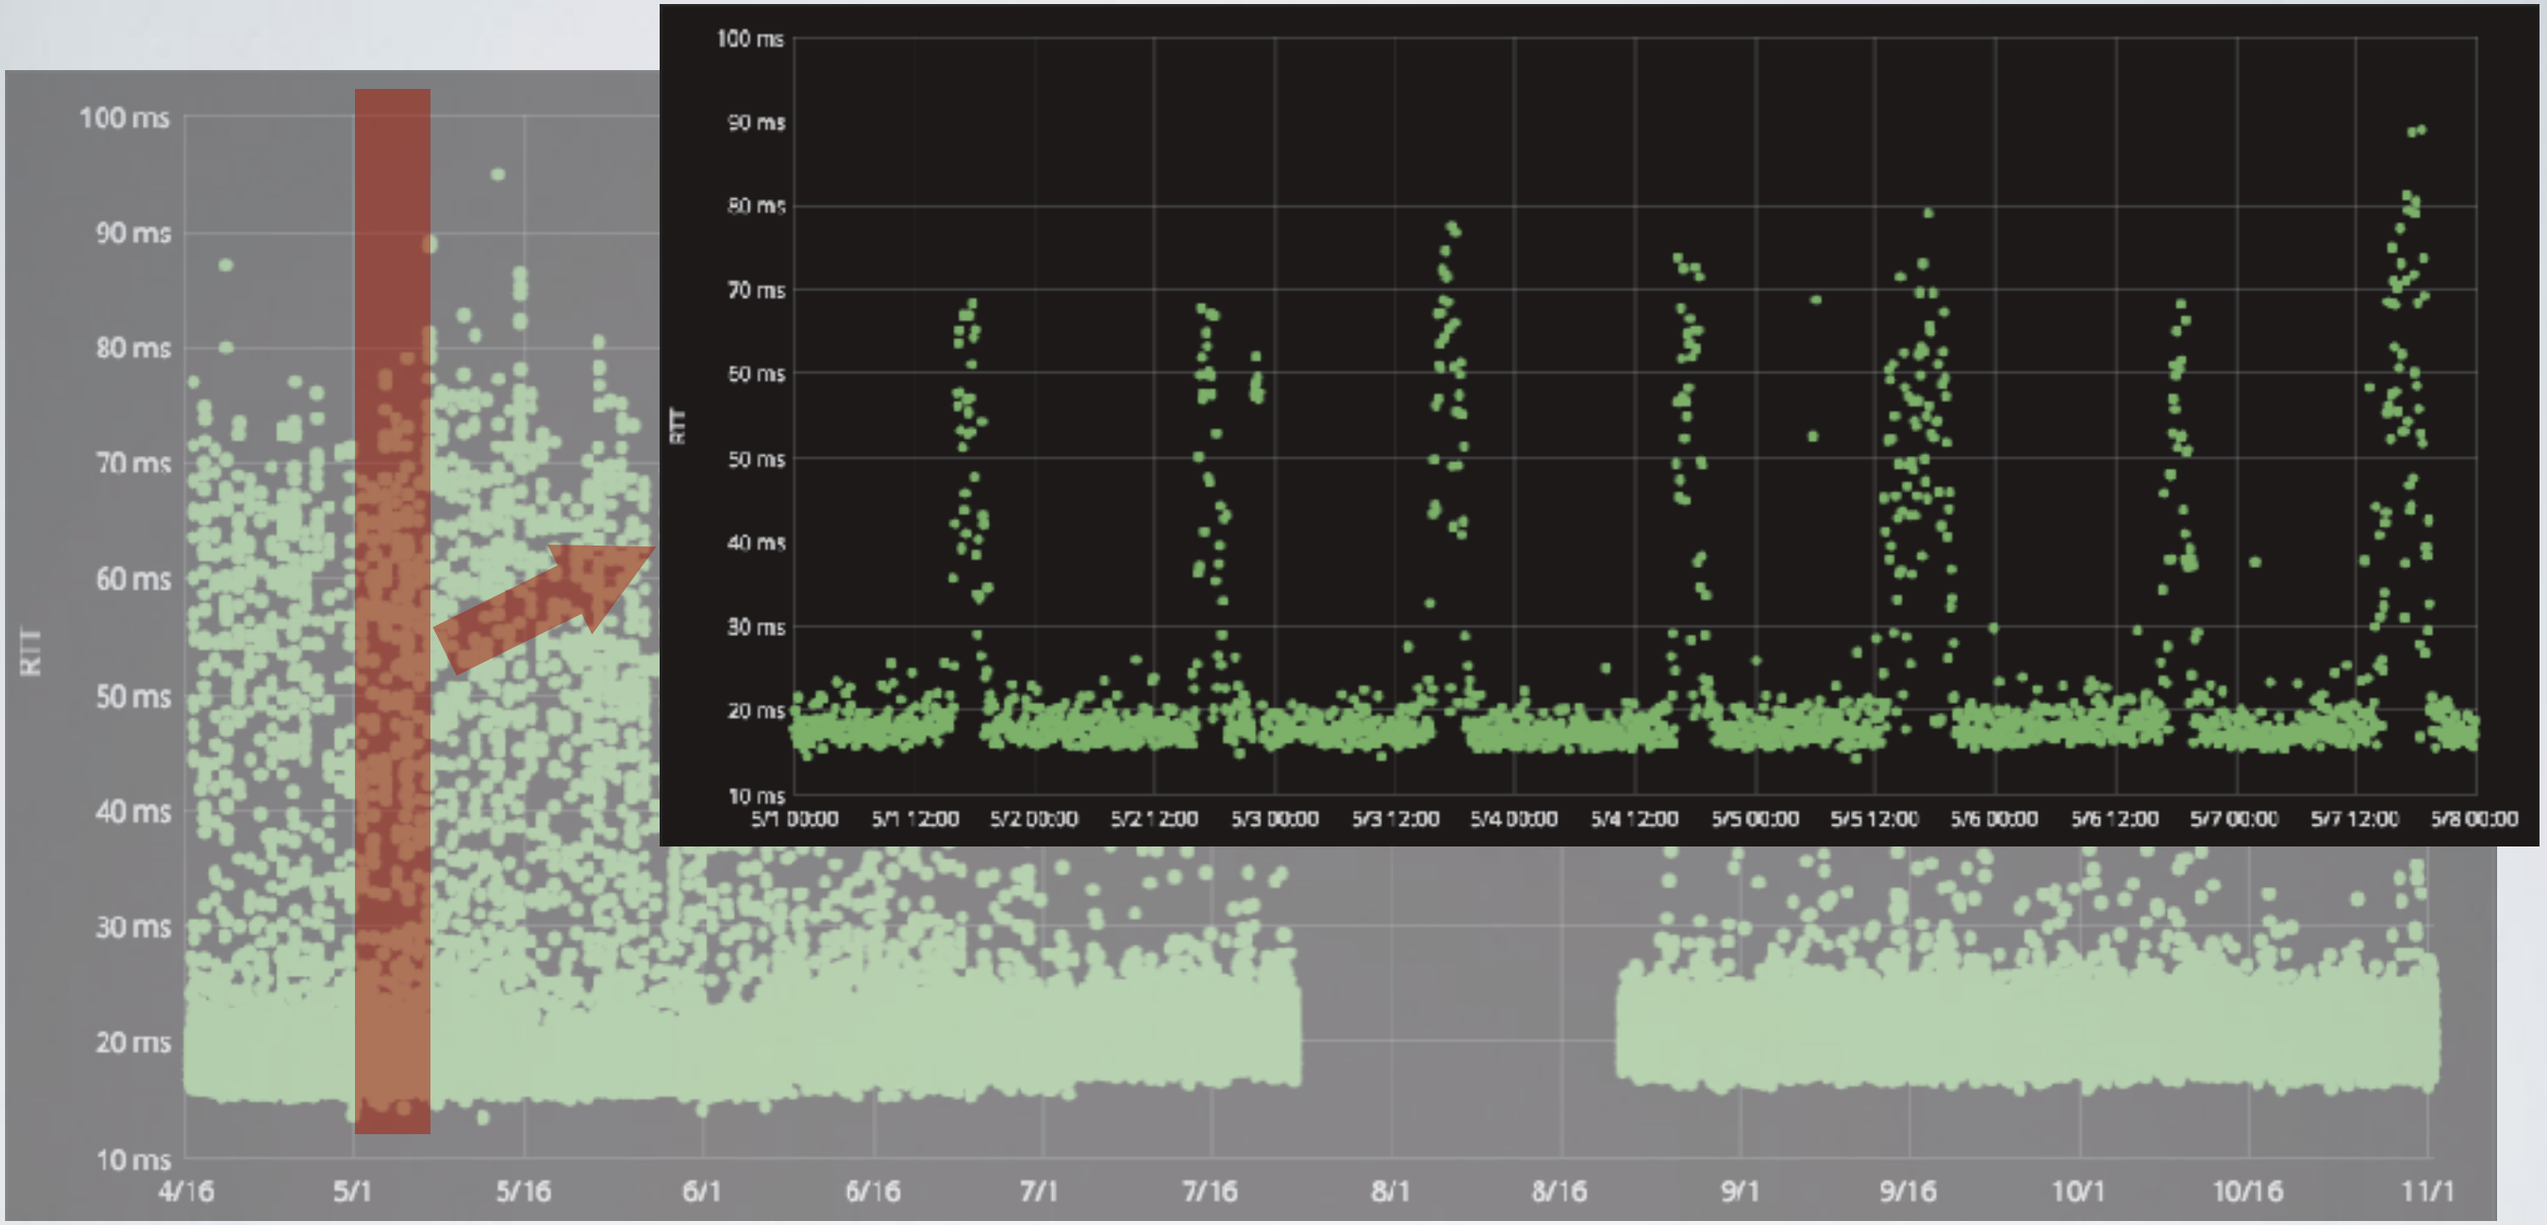

Zooming in for more detail

caida

Takeaways

- We have developed a lightweight method and system to provide third-party visibility into interdomain congestion
- We hope that our data can provide empirical grounding to debates over interconnection performance
- Contact us for access to the data: manic-info@caida.org

Host a Measurement VP!

We are always looking for volunteers to host VPs!

Contact us:

manic-info@caida.org

caida

Thanks!
Questions?

manic-info@caida.org