


# INTERNET INTERDOMAIN CONGESTION

Matthew Luckie, Amogh Dhamdhere,  
Bradley Huffaker, Young Hyun

*mjl@caida.org*

Pls - **KC Claffy** (CAIDA)  
David Clark (MIT)



# WHAT IS CAIDA?

- The Cooperative Association for Internet Data Analysis (CAIDA)
  - Founded by PI and Director k claffy
  - Independent analysis and research group
  - Based at UC San Diego's Supercomputer Center (SDSC)
  - 15+ years experience in data collection, curation & research
  - Measurement infrastructure, tools, analysis, & data sharing
  - Research informing Internet science, technology, policy



<http://www.caida.org/>



# IN THE PRESS



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

from the *more-peering-disputes* dept

## 'Peering' Into AOL-MSN Outage

### Level 3 and Comcast Issue Statement

Jul 16, 2013

BROOMFIELD, Colo., July 16, 2013 – Level 3 and Comcast have resolved their prior interconnect dispute on mutually satisfactory terms. Details will not be released.

## Confirmed: Comcast and Netflix have signed a paid peering agreement

### 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.

### Cogent Gearing for Another Peering Battle

by [Stacey Higginbotham](#) FEB. 23, 2014 - 9:27 AM

## Verizon denies using net neutrality victory to sabotage Netflix, Amazon

BY **BRIAN FUNG** February 5 at 1:59 pm

## 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.




# BACKGROUND

- Modern peering disputes manifest as congested links
- Disputes among access, content, and transit providers
- Some content is carried over inadequate links between access and transit networks
- **Congestion on transit links affects everybody**, not just parties to the peering dispute



# INTERDOMAIN CONGESTION

- Steady flow of articles discussing interdomain congestion and peering disputes: focus is on individual links
- We are developing a method to characterize the extent of interdomain congestion
- **Our goals (1) atlas of interdomain links and their congestion state, (2) improve transparency, empirical grounding of debate**
- Work still in *early stage*: **seeking funding from NSF and Industry, as well as feedback**




# METHOD: TIME SERIES PING




Vantage Point

Border Routers on  
Interesting Link




*(repeat to obtain a time series)*



**More congestion on weekend than weekdays. Monday 11th was Veterans Day**




# RTT measurements of border routers






# TRAFFIC VS TIME SERIES PING



# TRAFFIC VS TIME SERIES PING



Increase in delay begins after 90% utilization.  
Our traffic data source doesn't have the demand  
that we hypothesize is behind the RTT level shift.






# CHALLENGE: VALIDATION

*(with a concrete example)*

- 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/>




# CHALLENGE: VALIDATION

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



“Ground Truth”

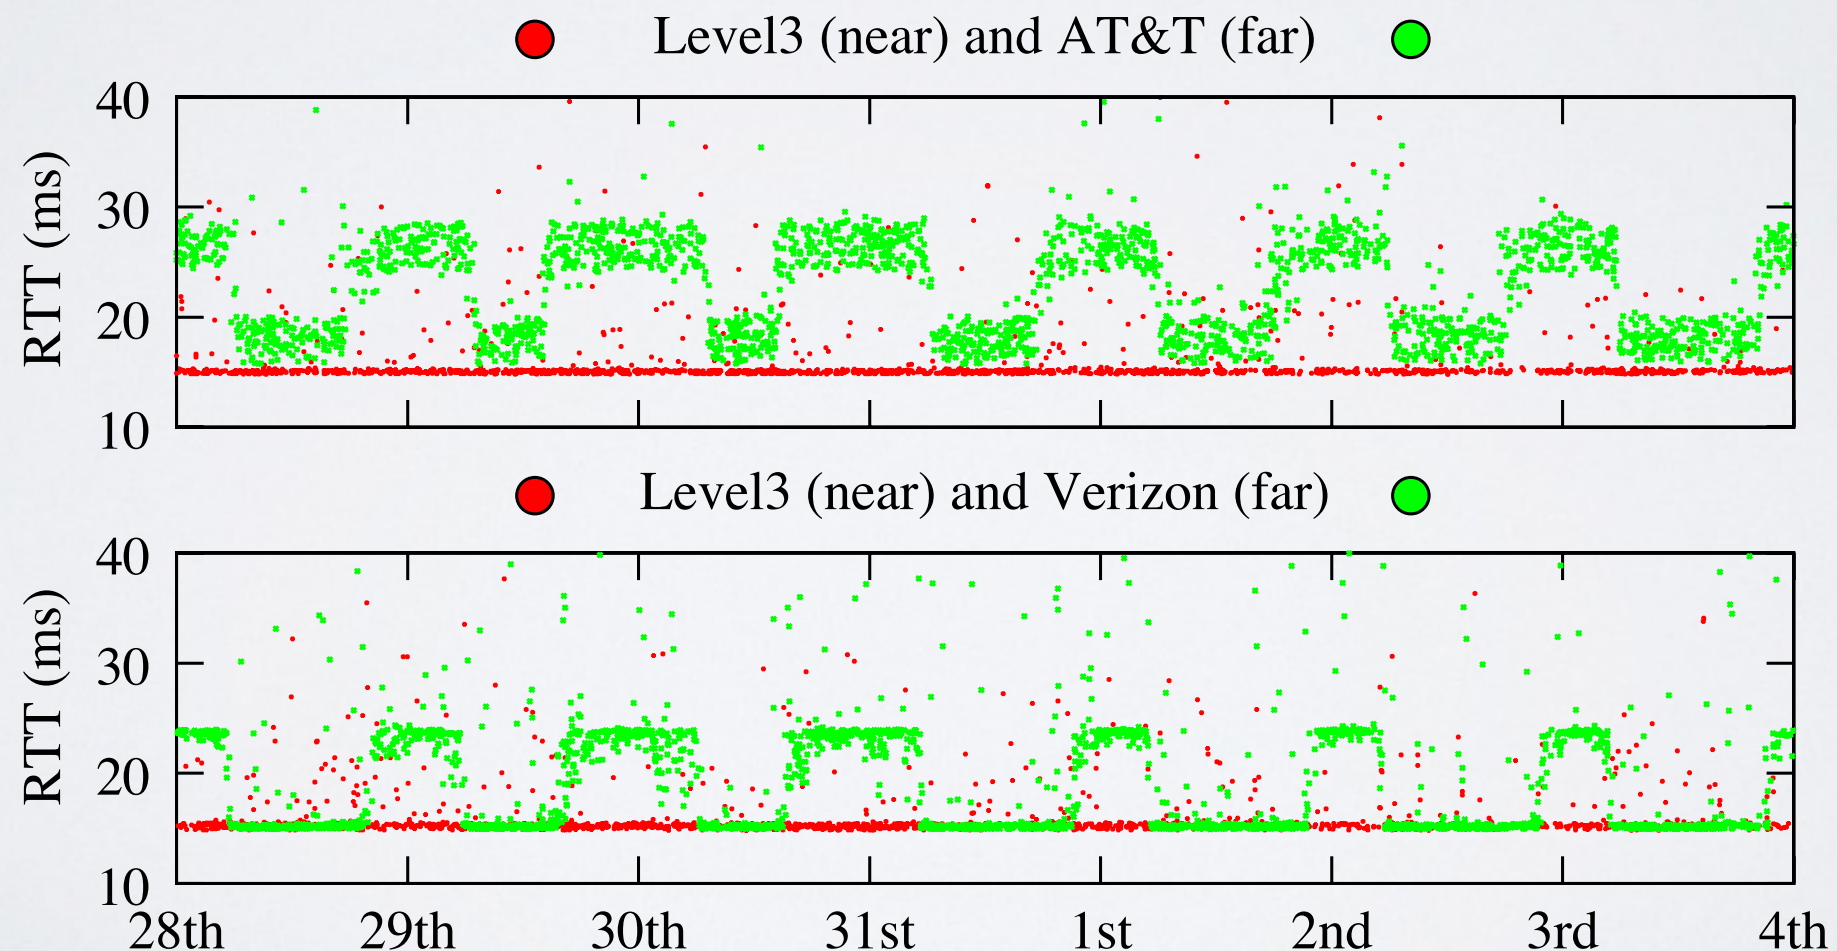
We believe both AT&T and Verizon are congested with Level3 in Dallas



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

# CHALLENGE: VALIDATION

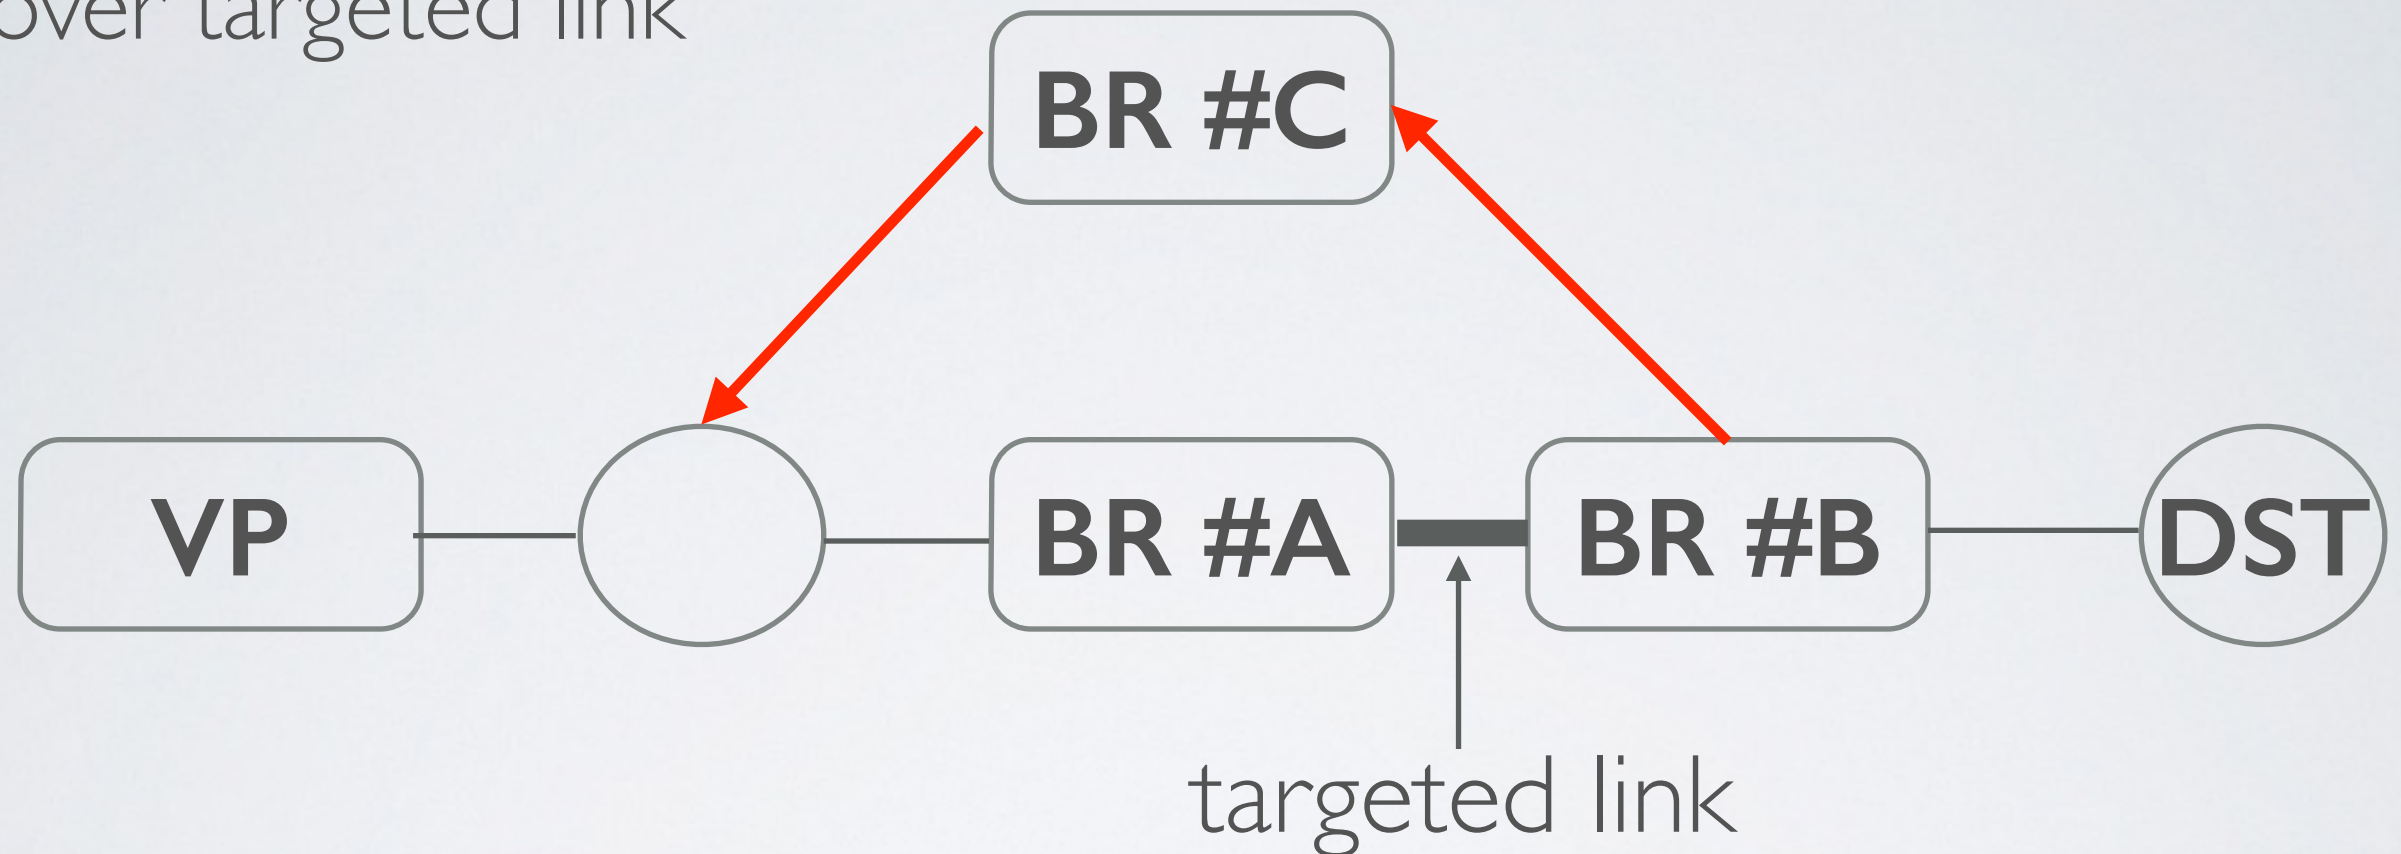
- Duration of level shift is different between AT&T and Verizon, particularly after the 1st.
- But it could also be a congested virtual output queue internal to the router (i.e. not a congested link)
- Would really like some feedback, offline.






# 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.

# OTHER CHALLENGES

- Building the system!
- Which interface IP addresses represent border routers?
- Interdomain interconnections come and go
- Need to adapt to routed paths that change over time
- Not trivial to determine direction of congestion
- ICMP responses may queue differently from other traffic





# CURRENT STATUS

- First pass at a system to infer interdomain links of an attached network and collect RTT time series to make congestion inference
- Deployments in various access networks (and other network types, see <http://www.caida.org/projects/ark/>)
- We continue to deploy ark nodes, using Raspberry Pi hardware, in homes of our friends (or friends of friends)


Power Supply →

Raspberry Pi  
(512MB ram, 8GB flash)  
smaller, more powerful  
than typical wireless routers




# CONGESTION TRENDS

*(three transit links of Comcast in Bay Area over time)*



# 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)





# 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
- Seeking industry funding and feedback to support research
- We view this as a long term project, similar to other long term CAIDA projects

*Email: [mjl@caida.org](mailto:mjl@caida.org), [kc@caida.org](mailto:kc@caida.org)*

