INTERNET INTERDOMAIN CONGESTION

Matthew Luckie, **Amogh Dhamdhere,** Bradley Huffaker, Young Hyun, KC Claffy (CAIDA/UC San Diego) Steve Bauer, David Clark (MIT)

amogh@caida.org



THE INTERDOMAIN INTERNET

- The Internet consists of ~45000 independently operated and managed networks (ASes)
- Broadly two types of relationships between networks: customer-provider or (settlement-free) peering
- Settlement-free peering based loosely on a notion of balance of trade



BACKGROUND

- Historically, peering disputes were between large transit provider networks and were motivated by unbalanced traffic
- One network would de-peer the other to encourage a resolution
- Numerous examples: BBN vs. Exodus Communications (1998), PSINet vs. Cable and Wireless (2001), AOL vs. Cogent (2002), France Telecom vs. Cogent (2005), Level3 vs. Cogent (2005), Sprint vs. Cogent (2008), Level3 vs. Comcast (2010), France Telecom and Free vs. Google (2012)

MORE RECENTLY ...

- Peering disputes appear to manifest as performance degradation of high-bandwidth applications (particularly video) as seen from large broadband access providers
- When videos buffer, people get angry..



Netflix is slow lately on FIOS

06-11-2013 12:53 PM

For the past few weeks I am intermittently getting only 1-2 bars on the bandwidth meter for the XBOX Netflix app, which I think means it is ~1-2Mbps.

Options -

Reply

I have a 35/35Mbps plan with FIOS, so it should be getting HD all the time. (it used to until recently).

Are others having this problem lately? Is it a FIOS issue or is Netflix downgrading their service for any customers not on OpenConnect ISPs?

Comcast is definitely throttling Netflix, and it's infuriating

First, let's talk about the worst company in America

Comcast is one of the most hated companies in America, yet inexplicably, also one of the most successful. Two nationwide surveys done by the American Customer Satisfaction Index in 2004 and 2007 showed that <u>"Comcast had the worst customer satisfaction rating of any company or government agency in the country, including the Internal Revenue Service."</u> Wow, just let that sink in for a second. People would rather get audited by the IRS than call up Comcast to deal with a faulty cable modem. Also, they are consistently given terrible ratings by every <u>consumer advocate organization</u> across the board.

Dave's Blog

My Collection of Hobbies, Code and Other Ramblings

Blog Archives

FEB 5TH, 2014

Verizon Using Recent Net Neutrality Victory to Wage War Against Netflix

UPDATE: The team over at Speedchecker Ltd has created a speedtest oriented around this issue. They are going to be collecting data and presenting findings if they get enough data:

http://netneutralitytest.com/

I usually don't post articles about current affairs. However, a recent series of events has inspired me to write about this.

Towards the end of January, the president of our company – <u>iScan Online, Inc.</u>, was complaining that our service was experiencing major slowdowns. I investigated the issue, but I couldn't find anything wrong with our production environment. We were stumped.

One evening I also noticed a slowdown while using our service from my house. I realized that the one thing in common between me and our president was that we both had FiOS internet service from Verizon.







No VPN for me, thank you

The Washington post PostTV Politics Opinions Local Sports National World Busines



Archives Authors

2

Verizon denies using net neutrality victory to sabotage Netflix, Amazon

BY BRIAN FUNG S February 5 at 1:59 pm





Today, Verizon on their policy blog said that we "don't have the full story" and that this is a problem on Cogent's part. Here is what they wrote:

Cogent is not compliant with one of the basic and long-standing requirements for most settlement-free peering arrangements: that traffic between the providers be roughly in balance. When the traffic loads are not symmetric, the provider with the heavier load typically pays the other for transit (see our ex <u>parte filing[PDF]</u> from the 2010 Comcast/Level 3 spat for more info on peering and transit agreements). This isn't a story about Netflix, or about Verizon "letting" anybody's traffic deteriorate. This is a fairly boring story about a bandwidth provider that is unhappy that they are out of balance and will have to make alternative arrangements for capacity enhancements, just like any other interconnecting ISP.





Asa Mathat

Netflix Says Verizon Isn't Slowing Down Its Streams

February 11, 2014, 6:42 AM PST

By Peter Kafka





BACKGROUND

- Modern peering disputes among access, content, and transit providers manifest as congested links
- 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: End-users complain, and this usually makes it to the news
- But data about the location of congested links is sparse and anecdotal



INTERDOMAIN CONGESTION

- This research project aims to characterize the extent of interdomain congestion
- Our goals (1) Methods to detect and localize congestion, (2) Map of interdomain links and their congestion state, (3) Data to improve transparency, empirical grounding of debate
- Trying to infer which network actors are responsible, or the incentives for their behavior is not our focus
- This is early work: we are still developing the method, and seeking feedback/validation

MEASURING LINK CONGESTION

- When links become congested, queues fill up
- As queues fill, network delay and loss rate increases
- Delay increase is related to queue size of congested router







Vantage Point

Border Routers on Interesting Link





Vantage Point

Border Routers on Interesting Link











(repeat to obtain a time series)





Diurnal trend in RTT to far end No trend in RTT to near end





November 2013: more congestion on weekend than weekdays. Monday 11th was Veterans Day







CHALLENGE: REVERSE PATH

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



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



CHALLENGE: PARALLEL LINKS

 Some interdomain connections consist of many parallel links



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

 Should we probe all parallel links? If all links are equally loaded, then we could probe only one representative link

CHALLENGE: WHICH IS THE INTERDOMAIN LINK?



- Inter-AS link is between routers R1 and R2
- If A I, A5, B3 seen in traceroute, we would infer interdomain link between R2 and R3. Inference depends on how the interfaces are numbered and the direction of our trace
- Possible approach: Reverse DNS, or probe all 3 IPs



OTHER CHALLENGES

- 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
- Large number of links, huge volume of data. Need methods to find patterns in time-series ping traces



SYSTEM DESIGN

- CAIDA Ark boxes at interesting network locations continuously probe the entire routed space to find interconnection links
- Time-series ping for each discovered interdomain link: optimize probing to minimize number of packets sent and still get RTTs from near and far end of each link
- Adaptive probing: continuously update the set of links to probe based on current routing state
- Find "interesting" links: Level-shift detection and frequency-domain analysis to find diurnal trends



EXAMPLES

 Link between large access network and transit network serving popular streaming video service shows persistent congestion





EXAMPLES

- Many users of a certain access ISP have complained of poor performance for a popular video-sharing website
- The ISP's peering link with that content provider shows signs of congestion



UNCONGESTED LINKS

- Equally important is to find links that are NOT congested
- Most links we have measured so far do not appear congested
- Example: Yahoo! peering with Cox Communications





HISTORICALTRENDS

- The time-series ping method is equivalent to running frequent traceroutes towards a destination behind the target link
- At CAIDA we do a lot of traceroutes: A set of ~70 Ark monitors continuously probe all /24 prefixes in the IPv4 Internet
- We have historical reverse DNS lookups and alias resolution runs
- Can we find congestion signals in historical data we have collected?

HISTORICALTRENDS

- Some high-profile interdomain links are seen (somewhat) consistently and sampled at a high frequency
- Example: From an Ark box at a large access provider (AP), we were able to sample links between the AP and 3 transit providers (TPs) serving a popular video streaming service



HISTORICALTRENDS

- Some high-profile interdomain links are seen (somewhat) consistently and sampled at a high frequency
- Example: From an Ark box at a large access provider (AP), we were able to sample links between the AP and 3 transit providers (TPs) serving a popular video streaming service



AP - TP A



28



CONGESTIONTRENDS

(three interconnection links of an access network over time)



SUMMARY

- Our end goal: a lightweight and easily deployed method to find link congestion patterns
- Still very early work, need feedback:
 - validation of congestion signal from network operators
 - what data should we be collecting and keeping?
 - Improvements to the method
- We can use more vantage points at residential locations. Please let me know if you can host an Ark monitor

THANKYOU! Questions?

