Appendix A

Massachusetts Institute of Technology, *Measuring Internet Congestion: A Preliminary Report* (Jun. 18, 2014)

Measuring Internet congestion: A preliminary report

The goal of the research described here is to determine the location and extent of congestion in the core of the Internet. In particular, as a first focus, we are interested in whether the interconnections among ISPs, both customer-provider (*transit*) and ISP-ISP (*peering*) links, are subject to widespread congestion. This topic is relevant to the current debates about network neutrality because congestion is one of the reasons given for network management. Our goal is to bring a focus on the more important issues in this space.

Definition: Congestion occurs when a part of the Internet (a link, a server, and so on) does not have the capacity to service all the demand that is presented to it. Congestion is not in itself a bad thing: if the network has enough capacity so that congestion never occurred, this might well imply unjustified investment. The transport protocols of the Internet (for example TCP) are designed to detect and respond to congestion by slowing down so that the offered demand matches the capacity. In this way, the design of the Internet allows it to adapt gracefully to instantaneous fluctuations in demand.[1]

Economic causes of congestion: However, long-lived and recurring congestion can be taken as a signal of a real mismatch between capacity and demand. The underlying causes of such a mismatch are usually economic—either the very high cost of a specific link (e.g., a trans-oceanic link), inability to obtain new capacity in a cost-effective, timely manner (which may happen in cellular networks due to issues of cell siting, etc.), or a business disagreement about the terms under which shared capacity is to be paid for. The recent dispute between major access ISPs and Netflix over the provisioning of direct connections is a highly visible example of such a disagreement.

Results: For several of the major U.S. broadband providers, including Comcast, Time Warner, Verizon and Cox, we have measured all the interconnection links visible from our vantage points. The links we probe include interconnection points with providers of high volume content such as Netflix and Google (Youtube), and content delivery networks (CDNs) such as Akamai and Limelight. We have also looked at ISPs in other parts of the world, including BT (UK) and Free (France). Most of the congestion we find over the last several months can be traced to specific circumstances, in particular the adjustments and negotiations related to the delivery of Netflix content into the access networks. We see peering links carrying Netflix traffic that appear to be congested for 18 hours a day, and we also see all of this apparent congestion vanish essentially overnight as new interconnection links are put in place, presumptively as a result of the new business arrangement between Comcast and Netflix.

Congestion at interconnection points does not appear to be widespread. Apart from specific issues such as Netflix traffic, our measurements reveal only occasional points of congestion where ISPs interconnect. We typically see two or three links congested for a given ISP, perhaps for one or two hours a day, which is not surprising in even a well-engineered network, since traffic growth continues in general, and new capacity must be added from time to time as paths become overloaded. We see some congestion on costly links, such as trans-oceanic links.

---

[1] For a detailed discussion of congestion and its various definitions, see Bauer, Steven and Clark, David D. and Lehr, William, The Evolution of Internet Congestion (August 15, 2009). TPRC 2009. Available at SSRN: http://ssrn.com/abstract=1999830

Conclusions:

- *Our data does not reveal widespread congestion problem among the U.S providers.*
- *Most congestion we see can be attributed to recognized business issues, such as interconnection disputes involving Netflix. These issues are being resolved, if slowly.*
- *Congestion does not always arise over time, but can come and go essentially overnight as a result of network reconfiguration and decisions by content providers as to how to route content.*

Possible implications for policy:

Since ISPs are under no obligation to peer, a link would not be put in place unless it was of mutual benefit. It would seem odd that two ISPs would see it in their interest to exchange traffic but not to put in adequate capacity. While the issues with delivery of Netflix content are taking a while to resolve, it would appear that all parties are moving toward adequate resolution. *We would actually find it surprising if there were widespread congestion on peering links.*

It makes little sense to talk about "priority" video traffic if the link in question is not congested. For a link that is not chronically congested, the only arguable benefits of priority are to reduce jitter in jitter-sensitive traffic during instantaneous overloads or to mitigate the consequence of unexpected increases in traffic.

Low-quality user experience can arise from many sources other than problems in the core of the Internet: The fact that we see little congestion in the core of the network does not mean that we believe users will always have a perfect experience when transferring data. There are many issues that can impair the quality of the user experience, including overloaded servers, problems with home networks (often wireless issues), and issues surrounding the last mile access link. We anticipate further research to explore the real causes of impaired user experience. Access links are traditionally priced based on peak speed, and that speed is often the limiting factor in the rate of a long-lived data transfer. To the extent there is little congestion in the core of the network*, the bottleneck most flows will encounter is either at the content server or the "last mile" access link to the customer*.

An uncongested core benefits small edge providers. For smaller providers of content and applications, who would normally reach their customers across the Internet either by using a third-party content delivery platform or by using the paths provided by peering and transit links, the lack of widespread congestion means they have adequate ways to reach their customers. During disputes they are harmed in passing by the congested interconnection links.

A note on our research method: Our method takes advantage of the fact that in most parts of the Internet, demand shows a strong diurnal pattern, with peak demand during the evening, and low demand in the early morning. Congested links are thus normally congested for only part of a day, corresponding to times of heavier load. We probe the links over time and look for different responses during congested periods. This approach allows us to find links that show recurring congestion over a number of days, but does not reveal episodes of short-term congestion, which can contribute to jitter.

The work described here results from a collaboration between the Information Policy Project at MIT and the Cooperative Association for Internet Data Analysis (CAIDA) at UCSD. For further information, contact David Clark at MIT (ddc@csail.mit.edu) or kc claffy at UCSD (kc@caida.org). Visit the MIT CSAIL Information Policy Project for more details on the IPP. http://ipp.mit.edu/
Visit the CAIDA web site at www.caida.org for more information on their range of research.

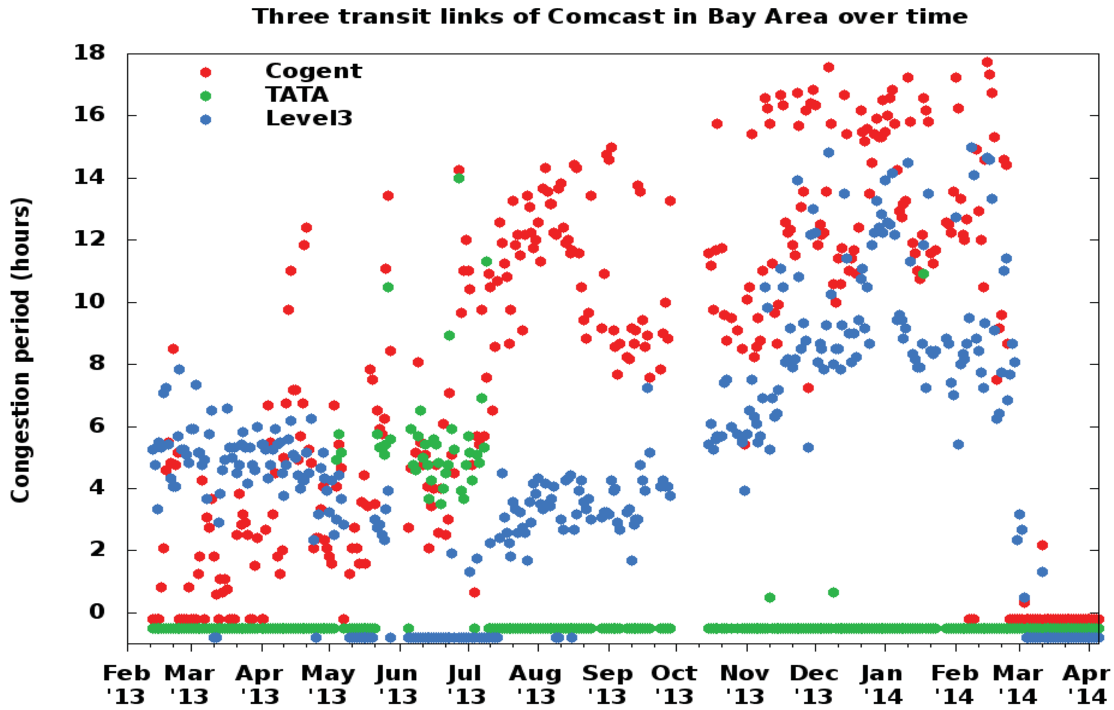

Inferred congestion duration for specific links connecting three major networks to Comcast. By February 2014, the Cogent and Level3 links appear to be congested up to 18 and 14 hours per day, respectively. Around the time Netflix and Comcast signed a peering agreement in February 2014, evidence of congestion on those links disappeared.

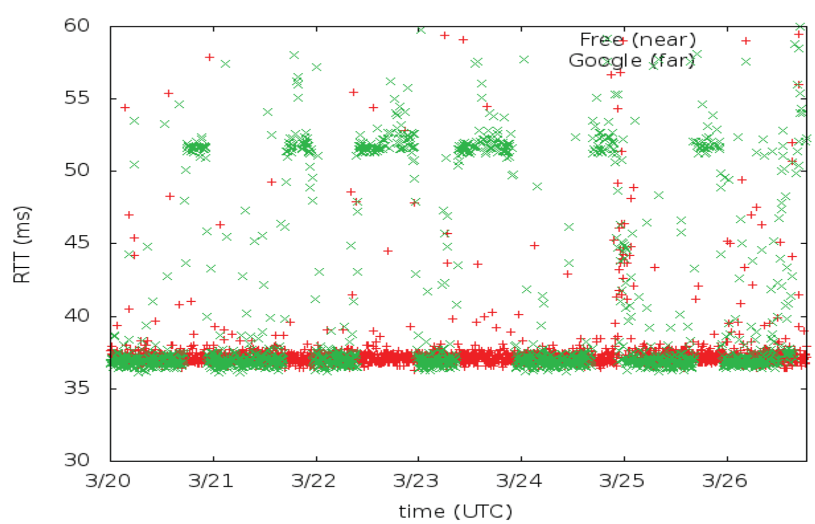

Inspired by customer reports of poor performance of Youtube videos on Free's network, we used our Ark monitor in Free's network to measure its direct links to Google with TSLP. This link appears congested for 4-6 hours at peak time on weekdays, and more than 12 hours on the weekends (March 22nd and 23rd).

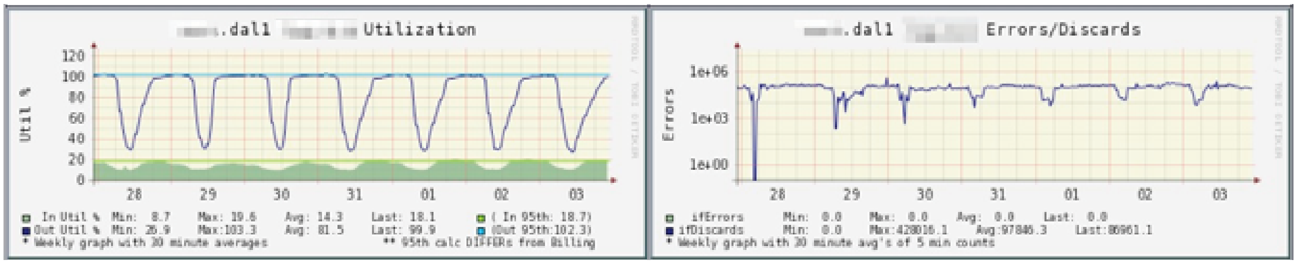

In a recent Level 3 blog, they gave some information (and some graphs) that provide evidence of congestion at a point of interconnection with an access ISP.
--http://blog.level3.com/global-connectivity/observations-internet-middleman/

Their blog provides data as to where the congested links were (Dallas). Using our tool, we can find links from Level 3 to two different ISPs in that location, ATT and Verizon, that show evidence of the same congestion as reported by Level 3.

Data analysis and graphing (except for the Level 3 blog data) by Matthew Luckie and Amogh Dhamdhere, CAIDA, UCSD