Policy Implications of Third-Party Measurement of Interdomain Congestion on the Internet

David D. Clark[†], Amogh Dhamdhere, Matthew Luckie[*], and kc claffy[1]

[†]MIT CSAIL
CAIDA/UC San Diego
[*]University of Waikato

February 13, 2019

Abstract

Internet engineering, science, and public policy communities have significant interest in understanding the extent and scope of, as well as potential consumer harm induced by, persistent interdomain congestion on the Internet. In recent work (Dhamdhere et al., 2018), we developed and implemented a lightweight active measurement method and system to measure evidence of congestion on thousands of interconnection links between broadband access ISPs and major interconnecting parties, including directly connected content providers. This method provides empirical grounding for discussions of interconnection congestion, without requiring direct access to interconnection links.

We first review the previous work to provide context. We then present new techniques for visualizing the data in ways we believe are conducive to policy analysis, e.g, of infrastructure resilience, performance metrics, and potential harm to consumers of persistently under-provisioned interconnection links. We present data in ways that allow us to compare different access providers, and show how congestion varies over time. We focus on seven large U.S. broadband access networks, but there is nothing U.S.-specific about the methods we use. Finally, we describe policy-relevant limitations and implications of the work.

1 Introduction

Internet engineering, science, and public policy communities have significant interest in understanding the extent and scope of, as well as potential consumer harm induced by, persistent interdomain congestion on the Internet. In recent work (Dhamdhere et al., 2018), we developed and implemented a lightweight active measurement method and system to measure evidence of congestion on thousands of interconnection links between broadband access ISPs and major interconnecting parties, including directly connected content providers. This method provides empirical grounding for discussions of interconnection congestion, without requiring direct access to interconnection links.

This work is based in part on an earlier work: Inferring Persistent Interdomain Congestion, presented at SIGCOMM '18, , Budapest, Hungary, © ACM, August 20–25, 2018. This work was partly funded by NSF grants CNS-1413905 and CNS-1414177.

In this work, we focus on policy implications of the previous work. We first review the previous work to provide context. We then present new techniques for visualizing the data in ways we believe are conducive to policy analysis, e.g, of infrastructure resilience, performance metrics, and potential harm to consumers of persistently under-provisioned interconnection links. We present data in ways that allow us to compare different access providers, and show how congestion varies over time. We focus on seven large U.S. broadband access networks, but there is nothing U.S.-specific about the methods we use. Finally, we describe policy-relevant limitations and implications of the work.

2 Background

In its strictest definition—demand exceeds capacity of a resource—congestion is a widespread phenomenon on the Internet, and is central to the proper functioning of TCP. The sender in a TCP connection induces congestion as a means to ascertain its most appropriate sending rate. It repeatedly increases its packet sending rate until it triggers a transient episode of congestion, which in turn triggers a lost packet. The sender detects this event, and responds by slowing its sending rate. These transient episodes of congestion do not normally degrade the user experience (QoE). Other common occurrences, including traffic management transitions, router operating system overheads, network configuration errors, flash crowds (e.g., software releases), and malicious attacks can induce isolated episodes of network congestion. These are inevitable and inherent aspects of packet switched public IP-based networks, and are not the focus of this work.

Our interest is in persistent congestion due to a long-term mismatch between installed capacity and actual traffic, particularly at points of interconnection between networks. This more persistent form of congestion manifests as an increase in the latency of packet delivery due to the time that the packet waits in a router buffer, and potentially dropped packets, which may be an impairment itself but also causes a sender to slow its sending rate, and poses a risk to the user quality of experience (QoE).

Several recent peering disputes covered in the press (Engebretson, 2010, 2013; Verizon, 2013; Brodkin, 2013b; Andrews and Higginbotham, 2013; Buckley, 2013; Brodkin, 2013a) shed light on various claims by involved parties regarding the causes of congestion and poor performance. However, there is a dearth of publicly available data that can shed light on interconnection issues. In isolated cases, providers have volunteered limited, anonymized information that reveals congestion at interconnection points (Taylor, 2014). Some content providers release end-to-end performance data (e.g., video quality reports) that provide aggregate information about performance between themselves and access providers (Google Video Quality Report, 2016; Netflix, 2017). However, end-to-end statistics cannot accurately map congestion to specific interconnection points, and the coverage of crowd-sourced measurements is low (Sundaresan et al., 2017). In 2016, seven access providers volunteered anonymized and aggregated statistics of interdomain link utilization, but the level of aggregation prevented inference of congestion on individual links (Feamster, 2016).

The U.S. broadband ecosystem (access and transit ISPs as well as edge providers) is at a daunting crossroads—without clear regulatory oversight, and without transparency into performance of critical components of the infrastructure. Although of regulatory interest in the U.S. for years (Federal Communications Commission, 2015b), there has thus far been no comprehensive system for lightweight measurement from the edge that is capable of inferring this type of congestion between networks *at a link-level granularity*. Our recently reported effort (Dhamdhere et al., 2018) to prototype such a system, which motivates new policy questions related to interpretation and application of this data.

In brief, our system identifies all the interdomain links visible from a vantage point (VP) in a given access ISP (Luckie et al., 2016), and uses latency measurements to infer evidence of congestion on those links (Luckie et al., 2014). Since March 2016, we have collected measurements from 86 vantage points in 47 ISPs. For every link we measure, we infer whether that link shows evidence of congestion, and for each day estimate the duration for which congestion was present. A detailed description of the method, validation, and design and implementation of a prototype implementation appears in in (Dhamdhere et al., 2018). A primary discovery of that work, further confirmed here, is that for the period of time over which we gathered measurements, we did not find evidence of widespread persistent congestion on interdomain links between U.S. access ISPs and directly connected peer, transit, and content providers. However, we did find evidence of congestion between specific parties, in some cases quite severe, e.g., exceeding half the day for many days.

3 Related Work

Bauer et al. (2009) presented an overview of the evolution of Internet congestion, arguing that recent increases in edge traffic demands and changes in user expectations have forced network operators to use non-TCP congestion control mechanisms, such as volume-based limits and active traffic management. At the same time, the Internet's interconnection structure has evolved (Gill et al., 2008; Dhamdhere and Dovrolis, 2010), most notably the proliferation of Internet Exchanges (IXes) as anchor points in the mesh of interconnection (Ager et al., 2012). Labovitz et al. (2010) reported that the majority of interdomain traffic was exchanged directly between content and consumer networks, and by 2013 Sandvine reported that for its monitored networks, Netflix accounted for a third of peak downstream traffic (Sandvine, 2013). These changes have resulted in heated disputes (e.g. Engebretson (2010, 2013); Verizon (2013); Brodkin (2013b); Andrews and Higginbotham (2013); Buckley (2013); Brodkin (2013a)) over the provision and management of interdomain links, as no single operator has complete control of the link's operation. Work on characterizing interdomain congestion in the U.S. is sparse, largely due to the proprietary nature of the data required to validate inferences. Different stakeholders publish reports, aggregated data, or end-to-end throughput measurements, but none of these approaches allow insight into individual interconnections. In 2014 and 2015, the Measurement Lab (M-Lab) consortium released reports that inferred congestion at interconnection points between large content and access networks in the U.S. using crowd-sourced Network Diagnostic Tests (NDT) by end users (M-Lab Research Team, 2014; Anderson, 2015). But these studies did not include path information, without which it is challenging to conclude that observed congestion is at the interconnection (it could be internal to ASes). Furthermore, the crowd-sourced nature of NDT tests used in these studies make it difficult to conclude that observed diurnal variations are due to congestion as opposed to variations due to a different testing sample (Sundaresan, Lee, Deng, Feng, and Dhamdhere, 2017). Sundaresan, Dhamdhere, Allman, and claffy (2017) proposed a technique that uses TCP connection statistics to determine if a TCP flow experienced *self-induced* or *external* congestion. While the technique provides more information than a regular speed test about what limits a TCP flow, it still cannot localize *where* the bottleneck lies.

Deng and Kuzmanovic (2008) introduced a method for detecting interdomain congestion using latency measurements. Because their *pong* tool requires cooperative endpoints, i.e., a system on the other side of an interdomain link, the approach is challenging to deploy at scale. Our scheme Dhamdhere et al. (2018) also uses measurement of latency, but avoids the necessity of having two cooperative end-points, so it can be deployed at a much larger scale.

The FCC has struggled with this level of opacity of the infrastructure, and tried to take its own steps to address it. In the context of the 2014 AT&T-DirectTV merger, AT&T and the FCC appointed a team from CAIDA and MIT as Independent Measurement Experts (IME) to define a methodology to measure AT&T's interconnections as a condition of the merger Claffy et al. (2016a). While the IME team was allowed to view data from AT&T's interconnects under NDA for the purpose of developing their method, that data will not be publicly available. The FCC also supports the Measure Broadband America platform, a mesh of ~10,000 vantage points (operated by SamKnows, Inc.) that perform active measurements of access (last mile) performance. The FCC has thus far not used this platform to tackle measurement of interconnection performance, although has expressed interest to us in lightweight technology to support such measurement.

The methods we use in this paper enable fine-grained link-level congestion inferences, which can be aggregated to higher levels of granularity such as region and provider-wide. We hope that our systematic validation of the method yields an opportunity for the FCC and other regulators, as well as third-parties, to augment existing measurement fabrics with a lightweight method for capturing potential interconnection performance issues.

4 System and Methods

The measurement system includes several components:

- Probes deployed across the Internet that serve as Vantage Points (VPs) to run our probing: the CAIDA *Archipelago (Ark) Measurement Infrastructure.*
- The *bdrmap* algorithm to infer the interconnection links.
- The probing method used to detect congestion on those links: *TSLP*.
- The method used to analyze the raw probe data, infer congestion and estimate its duration: the *Autocorrelation* algorithm.
- A back-end database system for storage and retrieval of the raw data.
- Tools for visualization of the raw data and congestion inferences.
- Supplemental probing methods used as part of validation of our inferences.

Figure 1 provides an overview of the measurement system we built to operationalize the TSLP method and execute it at scale.

4.1 The Archipelago infrastructure

The Archipelago measurement nodes used for this study are small computers (e.g., Raspberry Pis) installed in homes of residential broadband customers around the world. They support scientific experiments beyond the probing described here, as part of the larger CAIDA research agenda. We refer to these devices as vantage points (VPs). Since March 2016, our system has used 86 VPs in 47 networks in 24 countries. Due to the volunteer-based nature of Ark VP hosting, there is churn in the set of usable VPs. As of December 2017, our measurements spanned 63 VPs in 39 networks in 22 countries. In §5 we focus on results from 28 VPs in 7 broadband access networks in the United States.

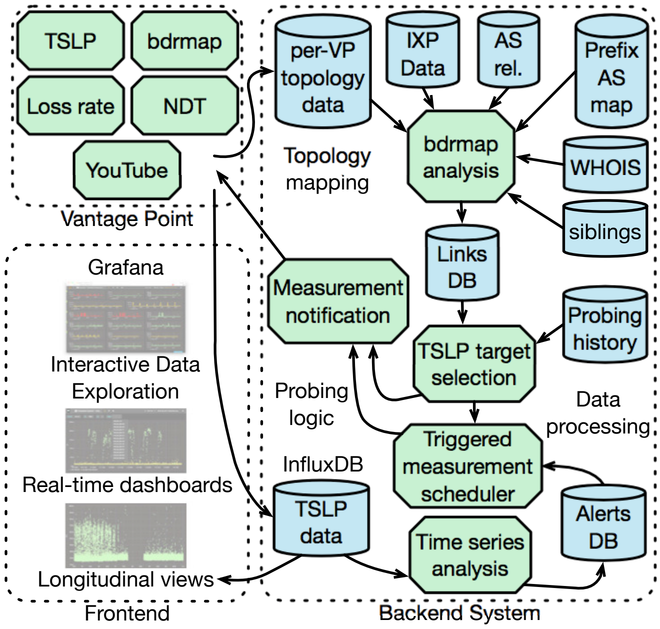

Figure 1: System for interdomain link discovery, active measurements, and congestion inference. *(Figure borrowed from Dhamdhere et al. (2018))*

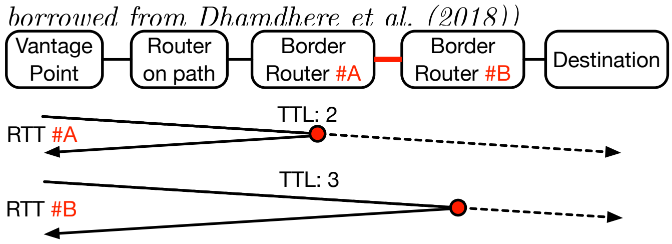

Figure 2: The Time-series latency probing (TSLP) method sends TTL-limited packets expiring at border routers #A and #B to measure link delay patterns.

4.2 Time-Series Latency Probing method

The scheme we use to detect evidence of congestion is called Time Series Latency Probes (TSLP) (Luckie et al., 2014). TSLP builds on a basic intuition: if the offered load at a link approaches (or exceeds) capacity, then packets are buffered, leading to an increase in measured latency through that link. Given a vantage point (VP) inside a network, we send a special packet (an ICMP packet) to measure latencies over time to the two ends of an identified interdomain link (*near* and *far* end, BR $\#A$ and BR $\#B$ in Figure 2). If the latency to the far end of the link is elevated but that to the near end is not, then a possible cause of the increased latency is congestion at the interdomain link.

The TSLP approach is appealing because of its lightweight nature and the feasibility of implementing measurements from the edge of the network, without cooperation from operators or direct access to border routers (unlike Dinu and Ng (2011)).

4.3 *bdrmap*: Identifying links to probe

Each VP runs *bdrmap* (Luckie et al., 2016) to infer the interdomain links between the host network and neighbor networks visible from the VP. The goal of *bdrmap* is to find every link visible from that VP between the border routers of the host network and every neighbor network. At a high-level, the *bdrmap* algorithm proceeds as follows:

- *bdrmap* attempts to find all the links that exit the source network (the network to which the VP is attached). Using data from the Border Gateway Protocol (BGP—the interdomain routing protocol of the Internet), *bdrmap* accumulates a list of all IPv4 address blocks reachable from this source network, and does a traceroute to an address from each block to see what path the packet takes out of the source network. (In August 2018 the global Internet routing table has about 725k address blocks.) A large broadband access ISP, might have several hundred such paths.

  Bdrmap uses as input a prefix-to-AS mapping constructed from public BGP data (RouteViews (University of Oregon, 2017) and RIPE RIS (RIPE, 2017)), a set of AS-relationships from CAIDA's AS-relationship algorithm (CAIDA, 2017; Luckie et al., 2013), a list of IXP prefixes curated based on data from Packet Clearing House (PCH) (2017) and peeringDB (2017), WHOIS data from RIR delegation files (ARIN, 2017; RIPE NCC, 2017; LACNIC, 2017; APNIC, 2017; AfriNIC, 2017), and a list of sibling ASes of the network hosting the VP.

- For each such path, *bdrmap* identifies up to three of those addresses to use as destinations for further probing. We prefer destinations in the address space of the neighbor network; however, this is not possible for all inferred interdomain links.

- *bdrmap* then heuristically determines which hop along the path to those destinations is the immediately adjacent interdomain link. The simple approach, which is prone to errors, is to use an Internet (BGP) routing table to look up the autonomous system (AS) number originating the address blocks containing the IP addresses returned by the traceroute, and infer that the interdomain link is the one where the AS number changes from that of the source network. Unfortunately, operational IP address assignment practices hinders the accuracy of this approach. Sometimes IP address assigned to the interfaces at the two ends of an interdomain link may be from an address block belonging to one of the two ASes. Depending on which AS it is, the actual interdomain link could be one hop before the link where the AS number mapping changes. Navigating the complexities associated with assignment of IP addresses to router interfaces led to the development and validation of the subtle and complex *bdrmap* algorithm, described in detail in Luckie et al. (2016).

4.4 Inferring congestion: the Autocorrelation method

From a policy perspective, recurring congestion, e.g., for days or weeks, is more significant than a single episode of congestion. Further, we argue that when we find recurring episodes of elevated latency, we can claim that these are indeed a manifestation of congestion, while an isolated event might arise from some other cause, such as a routing change or temporary configuration error. One cannot find a recurring pattern by looking at each day in isolation. *Autocorrelation* finds patterns of similarity between elements of a time series separated by a repeating interval, in this case 24 hours. Our autocorrelation scheme looks at the days in 50-day samples to find evidence of correlation across days; specifically, the algorithm looks for multi-day repetition of elevated delays at the same times of day that indicates congestion driven by diurnal demand.

Our autocorrelation algorithm aggregates the raw TSLP measurements into 15-minute intervals. It then selects the minimum measurement in each 15-minute interval to filter outliers. It tests for elevated latencies to the near side (which indicates possible congestion within the access network) and exclude those intervals so that the analysis can focus on the interconnection link itself.

For each 15-minute interval of the day, the scheme looks across all 50 days in the current sample and counts the number of days for which the RTTs in that interval were above an empirically determined threshold. The more days that contribute elevated latency in the same interval, the more likely it is that the event that triggered the elevated RTTs is a recurring one. The algorithm then looks for clusters of intervals where a minimum number of days had elevated RTTs. The current algorithm assumes, and infers, the single most likely cluster as the *congestion interval*, the period of the day in which recurring congestion occurred. (As future work we are considering the potential for multiple congested intervals per day.) The algorithm then separately analyzes each day, and counts the number of 15-minute intervals in the congestion interval in which the RTT for that day is elevated above our threshold. The algorithm uses the number of elevated 15-minute intervals as the estimate of congestion on that day. For example, if a link had one elevated 15-minute interval on a day, the algorithm infers a congestion level of 1.04% (1/96). We manually inspect the results of the algorithm in cases where it asserts evidence of congestion, to confirm that the assertion is appropriate.

The final stage of the scheme merges estimates from all VPs that observe a given interdomain link to derive an overall inference of congestion on that link. The scheme then identifies multiple links at the same location that appear to be part of the same *link group*. ISPs install multiple links in a given location when they need more capacity than a single link can carry. These are usually managed jointly (traffic is sent across them in a balanced way), and the resulting set of links is called a *link group* or a *link aggregation group*. It is useful to identify links that appear to be part of the same link group and deal with them as a single unit with larger capacity.

To identify links that could be part of a link group, the algorithm must first determine that they are in the same location. One approach would be to take the IP address of each link and look it up in some geolocation database, but these are sometimes not accurate for the IP addresses associated with routers. We infer that two links are in the same location if the minimum latency to the two links is the same, measured from every monitor that is probing the link. Another hint that two links are part of the same link group is that the IP addresses are often chosen from the same address range.

The strongest signal that two or more links are part of the same link group occurs if the links show congestion. Since operators roughly equally balance traffic across a link group, links that are part of a link group will show similar patterns of congestion. For each link in a potential link group, the algorithm looks at the duration of congestion for each day, and compare those values for every day for which it has data. Our system combines links into a link group only if the patterns of congestion are sufficiently similar.

Since the algorithm only combines links into link groups if the congestion on the links is similar, we can report congestion on link groups with no loss of information. In Section 5, we report levels of congestion on link groups rather than individual links, to make the results more compact.

Depending on VP deployment density, the *bdrmap* algorithm may not discover all link groups from a given access network to a given interconnected party. ISPs normally route traffic headed for any interconnected party to the closest point of interconnection. If there is no VP near a particular interconnection, our *bdrmap* probing may not discover it. The completeness of our visibility depends on the density of VPs that we are able to deploy, which varies across networks.

4.5 Back-end data base

The backend system maintains the current TSLP probing state and manages time-series data – millions of measurements over thousands of links. We selected an open source system called InfluxDB (2017) for this task. InfluxDB allows for efficient retrieval of raw probe data, optionally aggregated into time intervals and pre-processed in a number of ways. We are developing and will make public a set of query APIs for the data stored in InfluxDB to allow exploration of the data by those who are prepared to write code.

4.6 Data visualization

We selected the Grafana (2017) front end to provide interactive visualization of our data via a web browser. Grafana is integrated into InfluxDB, and provides an easy way to explore the raw data without having to write software.

5 U.S. Interdomain Congestion

Our probing system generates a large volume of data. Using Comcast (one of the largest ISPs in our dataset) as an example, and AS relationships (approximating business relationships between ISPs) inferred by CAIDA's AS-relationship algorithm (Luckie et al., 2013), our system discovered links with 1353 customers, 108 inferred peers (due to limitations of the current algorithm, some of these are actually customers), and 2 transit providers. We focus on measurements collected from March 2016 to June 2018, from U.S. access providers to major U.S. peers and content providers. There were 34 ASes in this reduced set for Comcast. Further limiting our analysis to links we observed for at least seven days yielded a total of 973 interdomain links. The population of links varies with our dynamic, but generally growing, visibility of interdomain links. In December 2017, there were 345 visible links to this reduced set of peers.

The congestion numbers we report derive from our *autocorrelation* analysis method. We report overall levels of congestion between two interconnected parties, and at the level of a *link group*. Autocorrelation and merging of links into link groups is described in §4.4. For each day for each link group we observe, which we refer to as a *day-link*, the inference algorithm classifies the day as *congested* or *uncongested*. If *congested*, the algorithm computes the duration of the congestion episode as a percentage of the day; a metric we call the *day-link congestion percentage*.

5.1 Overview of congested day-links

Table 1 reports the percentage of day-links showing evidence of congestion between each of the selected access providers (AP) and each interconnected network (transit, content provider or peer). This table reveals that different interconnections show different congestion profiles. Our system classified as congested 98% of the observed day-links between CenturyLink and Google, but only 4% of those between Cox and Google. Comcast and Google had 20% of day-links classified as congested, whereas that number was only 4% for Comcast and Netflix. The prominence of Google in Table 1 is consistent with efforts to actively maintain high peering link utilization (Yap et al., 2017).

Interconnected Network	Comcast	Verizon	CenturyLink	AT&T	Cox	TWC	Charter
TATA	66/44	9/4	12/5	49/35	—	14/12	—
NTT	48/3	2/0.6	Z/Z	43/18	18/9	—	—
XO	37/13	1/0.3	8/3	43/21	—	22/10	33/5
GOOGLE	20/14	77/63	98/93	9/5	4/1	—	2/2
CENTURYLINK	22/8	1./0.2	Z/Z	30/11	—	—	—
LEVEL3	9/3	4/1	4/3	5/3	37/32	5/3	Z/Z
AT&T	3/1	Z/Z	23/3	—	—	—	—
C&W	14/9	11/3	9/7	—	Z/Z	3/2	—
NETFLIX	4/2	4/3	11/9	5/3	23/21	3/2	5/4
LIMELIGHT	12/1	1/0.3	1./0.7	0.3/Z	2/2	6/3	3/1
TW	9/2	—	—	2/0.6	—	—	0.7/0.7
COMCAST	—	0.4/0.3	22/4	—	1/1	8/2	Z/Z
YOUTUBE	6/5	6/5	—	—	1/0.7	—	—
MFNX	7/1	1/0.1	7/4	4/0.7	4/1	0.8/0.1	—
AKAMAI	5/0.2	2/0.2	5/0.2	4/1	3/1	6/3	4/4
FACEBOOK	3/1	2/0.4	0.8/0.04	3/0.4	3/1	2/1	2/Z
HURRICANE	0.7/0.06	2/0.5	4/1	3/0.4	9/7	1/1	Z/Z
GTT	4/1	0.8/Z	Z/Z	1/0.6	Z/Z	1/0.4	4/2
AOL	2/0.7	0.8/Z	2/0.2	—	2/1	2/2	—
COGENT	0.6/0.3	0.9/0.4	2/0.4	2/1	Z/Z	3/2	4/3
CABLEVISION	1/0.4	—	—	—	—	—	—
SAVVIS	3/0.5	Z/Z	0.7/0.3	—	8/5	0.7/0.7	2./2
AMAZON	0.7/0.05	4/3	1/0.5	0.4/0.08	3/3	0.2/0.1	1/0.2
APPLE	2/0.4	0.9/0.8	0.4/0.4	0.4/0.2	2/0.8	Z/Z	0.2/Z
MICROSOFT	0.3/Z	1/1	0.4/Z	2/0.1	21/0.9	0.5/0.4	3/0.7
SPRINT	0.7/0.2	Z/Z	2/0.4	0.7/0.2	—	—	—
VERIZON	0.5/0.3	Z/Z	Z/Z	0.1/0.1	—	0.4/0.2	—
COX	0.2/0.1	—	—	Z/Z	—	—	—

Table 1: Two metrics of congested day-links for each pair of providers: percentage of day-links where any congestion was inferred, and percentage of day-links where congestion persisted for more than 4% of the day (about an hour). The table is ordered based on the average congestion to all of the included interconnected networks. Note that some interconnected networks exhibit frequently congested day-links with some access providers. We selected this 4% threshold to capture significantly persistent congestion; it is a tunable parameter. Z: congested day-links < 0.01%. —: No observations.

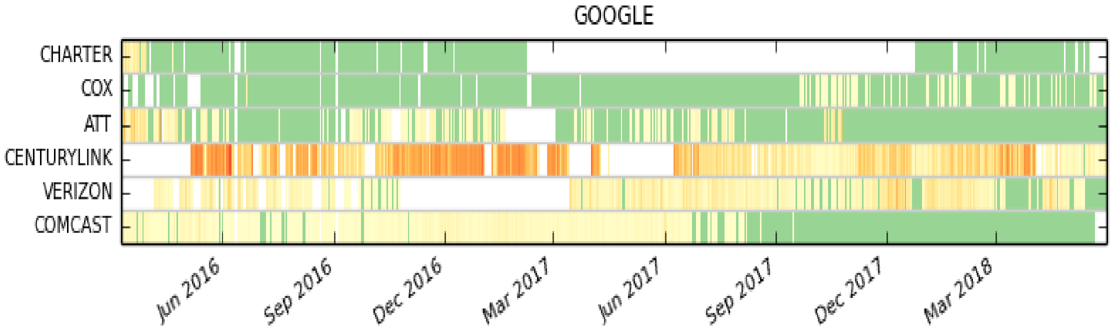

Figure 3: Congestion over time from set of U.S. access ISPs to Google, averaging across link-groups our VPs detect. CenturyLink shows episodes of substantial congestion; Comcast shows mild congestion ending in the summer of 2017. The long gap in the data for Charter is due to a failure of the vantage point. The gap in the data for Verizon reflects a period when our monitors in Verizon detected no link groups to Google, which might reflect an error in *bdrmap*, or actual interconnection conditions.

5.2 Temporal evolution of congestion

Table 1 provides no insight into the dynamics of congestion over time, and little visibility into the duration of congestion. But congestion on interconnection links changes over time. A link group that is congested an hour a day is probably not causing as much of an impairment as a link group that is congested 50% of the day, but both of these would contribute to the computed percentage of congested day-links. In this section, we provide a view into how evidence of congestion, as inferred by our system, evolves over time.

As an example, figure 3 shows the pattern of congestion from each of our observed networks across their direct interconnections (again, the ones that we observed) to Google. The graph plots the overall congestion for each day, averaging the congestion on all the interconnecting link groups, and weighting each group by the number of links it contains. The color for each day indicates the *duration* of congestion: days in green had no congestion across any of the links: for days with congestion the colors range from light yellow (short-term congestion) to deep red (congestion all day). Comcast experienced mild congestion ending around August 2017. CenturyLink experienced episodes of more substantial congestion. Cox and Charter experienced few periods of inferred congestion. (As previously mentioned, averaging may mask behavior of specific links. Other approaches to heatmaps may be worth exploring.)

Figure 4 plots inferences for connections to Tata, a tier 1 transit provider. The links from Verizon and CenturyLink show only occasional periods of mild congestion, AT&T has episodes of very substantial congestion, Time-Warner had congestion that ended near the end of 2016, while Comcast shows a continued level of mild congestion.

A superficial comparison of figure 4 and the data in table 1 may be confusing. According to that table, 66% of the day-links from Comcast to Tata were congested, but figure 4 shows some congestion for almost every day. This results because the data in figure 4 computes the average congestion for each day across *all* the link groups from Comcast to Tata, so congestion on one group will cause the average to be greater than zero, even if other groups on that day showed no congestion. Such aggregation can mask important phenomenon happening at the link-level granularity (Claffy et al., 2016b), Figure 6 shows the link groups from Comcast to Tata. (Link groups typically reflect the set of links in a

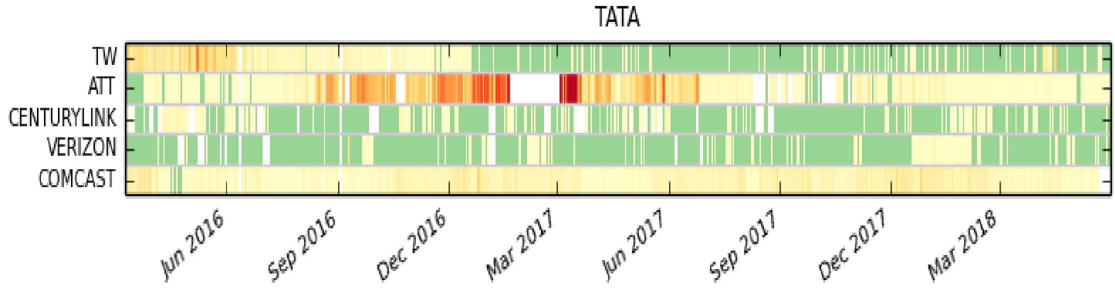

Figure 4: Congestion over time from the set of U.S. access ISPs to Tata. AT&T shows episodes of substantial congestion; Comcast shows a steady level of mild congestion.

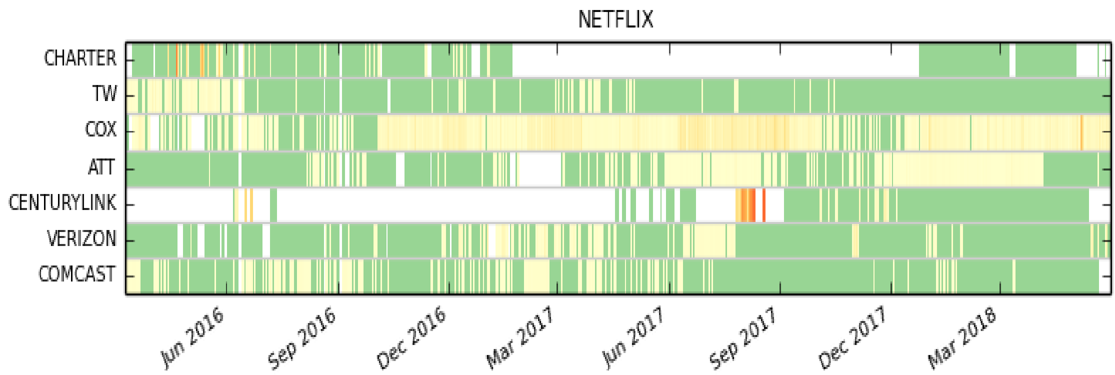

Figure 5: Congestion over time from U.S. access ISPs to Netflix. With the exception of Cox, there is little evidence of persistent congestion. CenturyLink had one interval of substantial congestion.

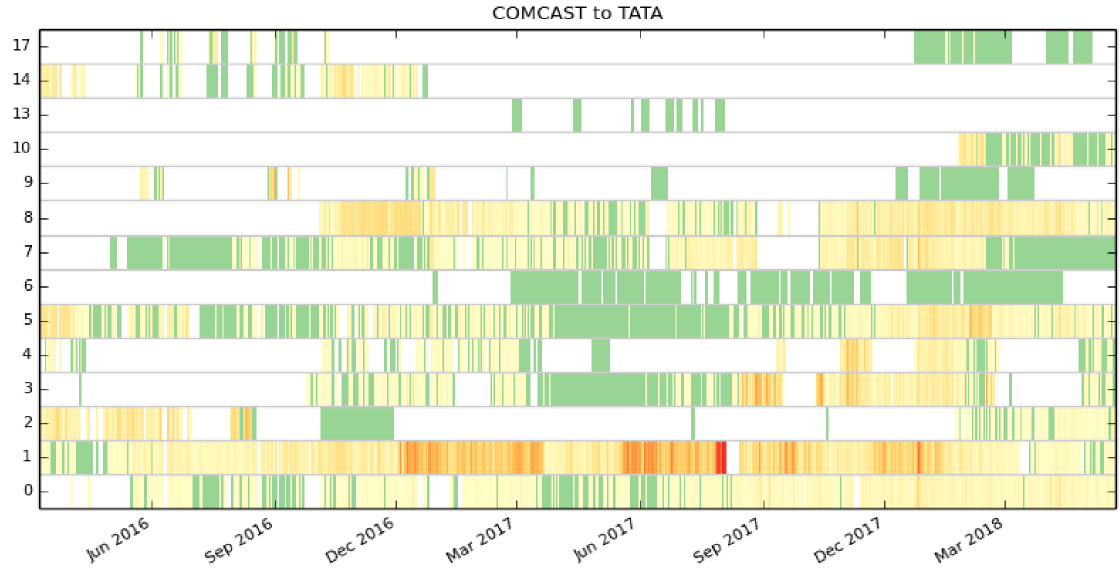

Figure 6: Congestion on different link groups from Comcast to Tata, where we detected activity for at least one month. Group 1 has a higher level of congestion compared to the other groups; evidence of congestion is not consistent across groups.

given interconnection point, e.g., a city.) Many of them only show mild congestion, with intervals where there is no congestion. However, one group (group 1) shows more significant congestion, with almost all days having at least some congestion. All this structure is masked with an aggregated view of the data, as when we report, as in Table 1 that 66% of the day-links from Comcast to Tata were congested to some degree.

An orthogonal view of congestion results from looking at different interconnections from a given access network, rather than interconnects to a given party from multiple access networks. Figure 7 shows the daily average congestion on the interconnections from Comcast to select parties. Comcast shows recurring congestion to Tata and NTT, congestion for part of the measurement period to Google, as well as congestion to XO. The patterns clearly evolve over time. Figures 8 and 9 show the same interconnections for AT&T and Verizon. AT&T, in common with Comcast, shows persistent congestion to Tata and NTT, but lower levels of congestion to Google. The vertical gap early in 2017 was due to failures of our vantage points in that network. Verizon has little persistent congestion to any of the traditional Tier 1 peers, but does show congestion to Google.

6 Other Measures of Congestion

Recurring diurnal patterns of elevated latency appear to be a legitimate signal of congestion, but they could also arise from other phenomenon, such as a load-dependent variation in scheduling the process in the router that responds to the probe packet. Given the politically loaded nature of interdomain congestion inferences, Dhamdhere et al. (2018) details three approaches used to cross-validate the inference method against other measurement methods, specifically of packet loss, throughput, and video streaming performance. They also used direct feedback from network operators to validate specific congestion inferences. We briefly summarize the points of this validation effort that are relevant to our analysis of policy-related strengths and limitations of this work (§7.5).

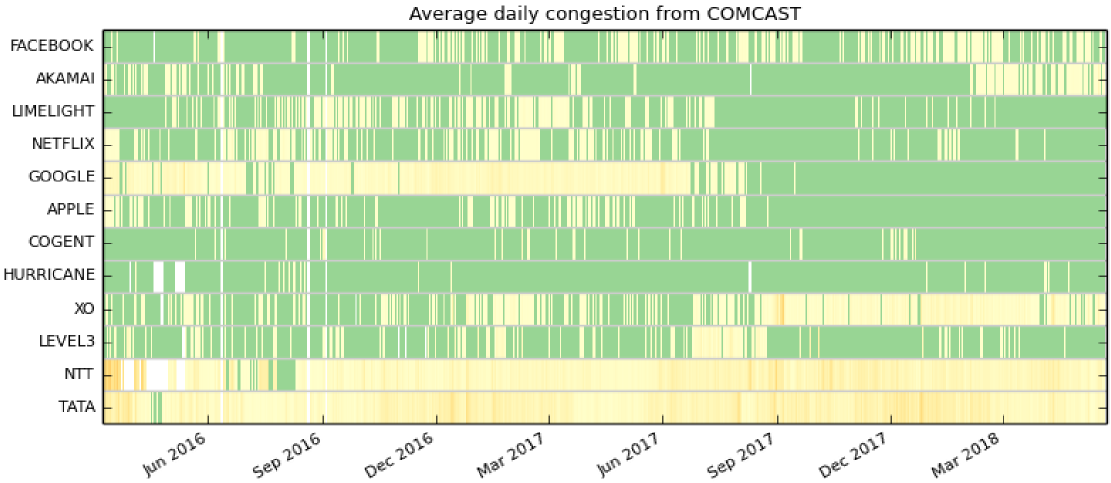

Figure 7: Congestion over time from Comcast to select interconnecting parties.

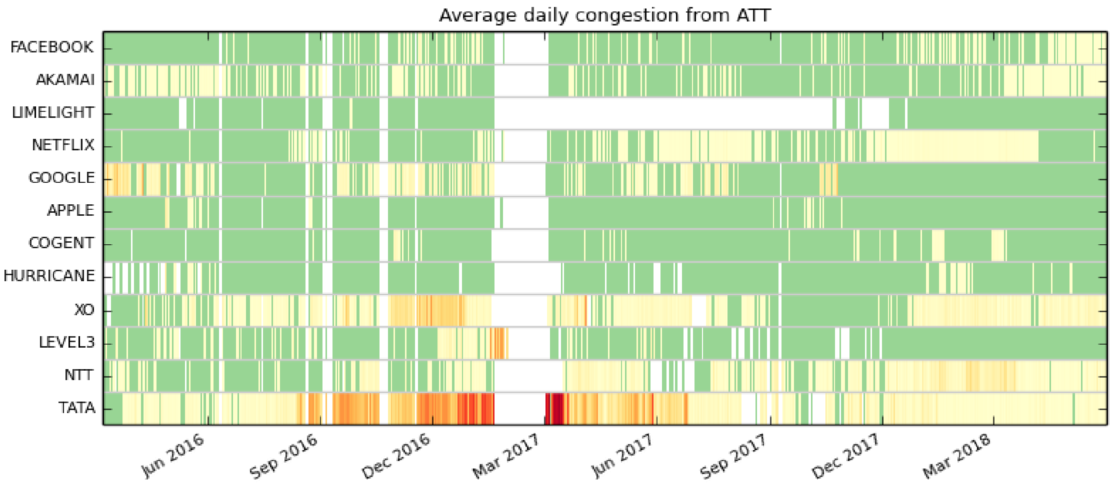

Figure 8: Congestion over time from AT&T to select interconnecting parties.

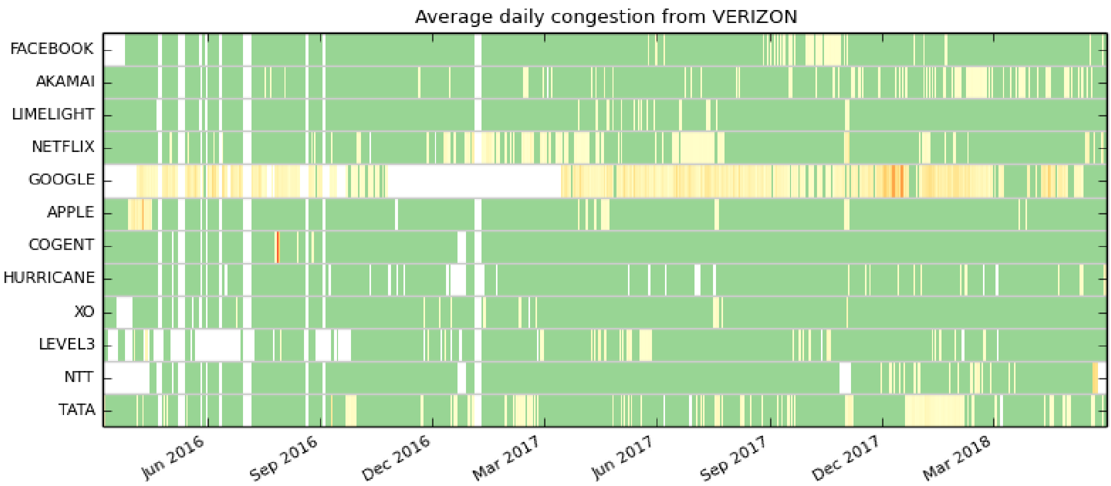

Figure 9: Congestion over time from Verizon to select interconnecting parties.

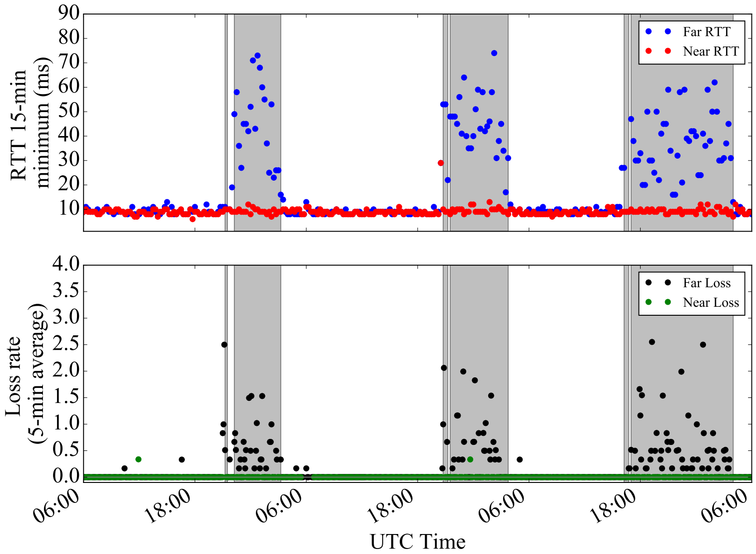

Figure 10: Time series of TSLP latency (top) and packet loss percentage (bottom) for an interdomain link between Verizon and Google on 7–9 December 2017. Periods inferred as congested are shaded in gray. *Plot taken from Dhamdhere et al. (2018).*

Loss rate. Correlation between periods of elevated latency and periods of elevated packet loss is a strong indication that the underlying phenomenon we are measuring is congestion. As with TSLP latency measurements, the loss measurement module of the system sends TTL-limited ICMP echo probes once per second toward both the near and far ends of interdomain links, set to expire at the target interfaces.[1] Figure 10 (borrowed from Dhamdhere et al. (2018)) shows a well-correlated example of TSLP latency and loss rate measurements for a link between Verizon and Google in December 2017. The latency pattern to the far end was elevated during peak hours every day, and our autocorrelation method flagged this link as congested during the periods shaded gray. The lower panel of Figure 10 shows the loss rate computed over 5-minute intervals during congested and uncongested intervals. Two important takeaways from this exercise were that a) loss rate to the far end was higher during congested periods than during uncongested periods and b) loss rate to the far end was higher than that to the near end during congested periods. (We do not always observe such correlation; § 9 discusses why traffic engineering techniques may induce cases of elevated latencies with no observable losses.)

Correlation with NDT throughput. Dhamdhere et al. (2018) also ran controlled experiments of the Network Diagnostic Tool (NDT) to measure upload and download TCP throughput to M-Lab NDT servers (M-Lab, 2017). Drops in throughput during periods inferred as congested provided additional validation. But several limitations impede the use of NDT measurements for broader validation: end-to-end measurements lack the ability to localize congestion to the interdomain link; lack of pervasive NDT server deployment limits the overlap with observed interconnections; and running throughput measurements frequently can congest the VPs home network.

[1] For statistical significance, the loss probing must be a much higher rate than for TSLP, which generates 1-3 data points per link in a 5-minute window.

Correlation with YouTube performance. Finally, Dhamdhere et al. (2018) investigated correlation of interdomain congestion with application performance metrics by studying YouTube video streaming, using both archived and new measurements taken with the `YouTube-test` tool (Ahsan et al., 2015). This tool emulates streaming playback by buffering and decoding video data, and requires a cooperating server in Google's network. After establishing a network connection to the YouTube server, the `YouTube-test` client starts filling its video buffer by downloading data using all available bandwidth. In steady state, traffic shows an ON-OFF pattern where during each ON-period, the client downloads a burst of packets from the video cache (Rao et al., 2011; Ghobadi et al., 2012). Figure 11a shows the CDFs of this ON-period throughput. Reduced peak (ON) throughput during intervals when TSLP detects elevated latency (the red line) strengthens the evidence for (but note, still does not prove) interdomain congestion. As with NDT, this test cannot localize congestion to a given link along the path; heavy server load or some other factor could reduce throughput.

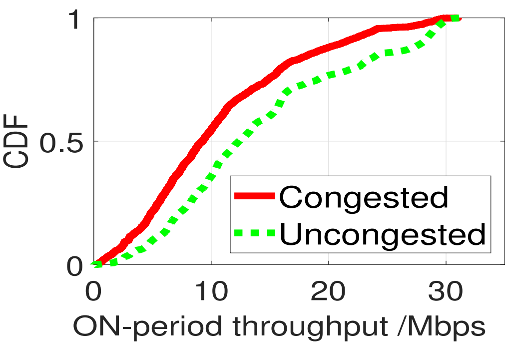

(a) YouTube streaming performance experiences lower throughput during periods inferred as congested using TSLP.

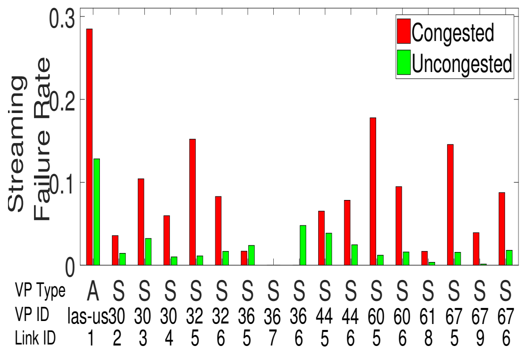

(b) Median streaming failure rates for Ark (A) and SamKnows (S) VPs. Failure rates higher during periods inferred as congested.

Figure 11: Validation using `YouTube-test` tool (Ahsan et al., 2015) on Ark and SamKnows vantage points to a cooperating YouTube server. *(Figures taken from Dhamdhere et al. (2018)).*

The download test also measures streaming failures, such as failing to download the next video segment quickly enough during streaming, which manifests as a re-buffering event on a viewer's player. Figure 11b shows the fraction of downloads with at least one streaming failure, as measured from Ark probes and SamKnows probes. It shows that except for one SamKnows VP, failure rates were generally higher during congested periods. For the Ark VP, almost 30% of the tests failed during congested periods. This measure does relate to perceived QoE, and indicates that in some cases, the degree of congestion was sufficient to cause an impairment.

7 Limitations of Current Approach

Dhamdhere et al. (2018) discuss several technological limitations to their approach to third-party measurement of interconnection performance. These include the possibility that observations are distorted by behavior of the router, the fact that *bdrmap* cannot necessarily identify all interconnection links (or link groups) between a network where we have VPs and a given interconnected network, and the possibility of asymmetric paths, where the probe packet to the far side of the interconnect leaves by one link and returns by another. Below, we discuss some limitations that have direct bearing on policy applications of this work, and future work that may overcome them.

7.1 Limited coverage path visibility

Our current system limits congestion measurements to interdomain links identified by *bdrmap*, which identifies only immediate neighbors of the network hosting a VP. We have not yet tried to measure any aspect of the path beyond the interconnection link. In other words, if there is congestion within the network of the far-side interconnected party, we do not detect it. Recently, Marder and Smith (2016) introduced the MAP-IT tool to identify interdomain links in a set of collected traceroutes. In the future we plan to investigate whether the combination of *bdrmap* and *MAP-IT* can enable measurement of interdomain links farther than one AS hop away from the network hosting our VP.

In the meantime, the current method can detect episodes of elevated latency on the path to the *near side* of the interconnection link, which might reflect congestion within the access network where our VP is located, or perhaps some local overload of the link that connects our VP to the network. We have not yet investigated this data for possible inferences about the access network's performance.

7.2 Quality of Experience (QoE) inferences

One risk of this sort of third-party measurement and reporting of evidence of congestion is that someone will incorrectly conclude that some actor or set of actors is doing something wrong. Even the sort of persistent congestion we observe does not necessarily cause, and this method cannot measure, impairment to the quality of experience (QoE) for users. Video content providers may respond to congestion detected along a path by downgrading to a lower quality of encoding, which users may or may not perceive as a QoE degradation. For example, some links from Google to access providers exhibited latency-based evidence of congestion for significant periods of the day, which correlated with a reduced peak sending rate in limited cross-validation (§6). But neither of these metrics imply a degradation in QoE. The only measured evidence where the consequence of congestion likely impairs QoE is the increased probability of a "re-buffering event" (figure 11b). Definitive assessments of QoE degradation require experiments in measurements of user satisfaction. In our ongoing work, we are building a system to crowd-source QoE measurements such as video streaming performance and web page loads (similar to the Eyeorg platform (Varvello et al., 2016)), and correlate those with congestion inferences from TSLP.

7.3 Link capacity inferences

While our method can measure the duration of congestion on a link, it cannot measure the capacity of a link, or whether the offered load on the congested link is 101% of what the link could carry, or 200%. A link that is congested because the offered load is 101% of what it can carry will cause the average sender to slow its sending rate by 1%. That might not matter at all. If the offered load exceeds the capacity by a factor of two, then the average (TCP) sender will cut its sending rate to half of what it would otherwise have been. That might matter a lot. We can speculate, based on experience talking to operators and observing traffic patterns, that for a traditional peering link, which carries whatever traffic arrives over it, it is highly unlikely that a congested link would be overloaded by only 1%, especially if that link is congested for a significant period of the day. Traffic levels typically display a somewhat sinusoidal pattern. If a link is congested for a significant period, one would expect the offered traffic at the peak to be well above the link capacity. But it is important to note that there is no current method for a third-party to accurately infer the capacity of a remote link on the Internet.

7.4 Traffic engineering inferences

Some content providers with dedicated direct connections to an access ISP use sophisticated dynamic traffic engineering techniques to carefully control the volume traffic they send over those links. In addition to using content-encoding methods that adapt to congestion in real time, content providers can also shift requested content over an alternative link with lower utilization, including indirect links that connect to an intermediate transit provider. As mentioned in §6, some interconnection links show recurring episodes of increased latency, but no increase in packet losses. Extremely precise traffic engineering of link load could lead to this outcome, with minimal to no impairment of QoE.

7.5 Economics of interconnection

Even if a content provider can manage an interconnection link to be just on the edge of congestion, that content provider must deliver requested content somehow. In a carefully engineered set of interconnection links, why would one link be provisioned to have excess capacity when others do not? One answer might be that the cost of installing that link was lower than the cost of the congested link. The answer is a business and economic one, not a technical one. We have no visibility into that sort of data. To illustrate this point, our data shows that the most congested links into most U.S. access networks where we have VPs are the links to the upstream transit providers of those networks, i.e., ISPs like Tata and NTT. If content providers use those links to handle peaks in traffic on their direct interconnection links, it is not surprising that those links would remain congested. An access ISP must pay for this upstream transit capacity, and has little incentive to increase that capacity if much of the traffic across that link is from content providers that would otherwise be cost-sharing with the access network for adequate interconnection capacity.

8 Policy implications

Considering all the caveats associated with this work, we pose two questions that merit further community debate: what role should such third-party measurement of interconnection play in the future policy landscape, and do the lessons we learn generalize to (or echo) other aspects of communications policy?

8.1 Enlightening policy debates

Our original motivation for this work was an increase in heated peering disputes between powerful players in the U.S., which raised questions about intentional degradation of performance as a business strategy to obtain (or avoid) interconnection fees (Norton, 2010). The prevalence of these public disputes dropped around the time of the FCC's 2015 Open Internet Order, in which the FCC asserted authority over interconnection, sending a signal to industry to resolve disputes or trigger regulatory oversight. (We cannot prove causation there; commercial players may have realized independent motivations to resolve peering disputes.) Furthermore, the text in the Order related to interconnection was among the most vague; the FCC explicitly recognized that they lacked sufficient understanding of interconnection (Federal Communications Commission, 2015b). In part to close this gap in understanding, during the next merger between an access and content provider (AT&T and DirecTV (Federal Communications Commission, 2015a)), the FCC imposed interconnection measurement and reporting conditions, for 4 years, under NDA agreements (Claffy et al., 2016). Like other sources of interconnection data, this data tells a partial story, but in this case, a secret one (Claffy et al., 2016a).

Thus, the most important contribution of this work is addressing this decades-long gap in a third-party's ability to study peering disputes in an open, objective, scientifically validated way. Our measurements reveal ongoing indications of persistently congested transit links, which – regardless of cause – implies clear motivation for large players to engage in direct peering negotiations. Especially in today's deregulatory political climate, we consider such measurement to be the most promising strategy for incentivizing transparent and accountable ISP behavior.

8.2 Doing science in cooperation with a competitive world

While ISPs have direct access to some of the metrics researchers try to measure, including interconnection congestion, in general they have knowledge only of their part of the Internet. The scale and complexity of the Internet, as well as the range of actors, leads to emergent behaviors that derive from non-technical inputs such as business decisions, which are not easily predicted. An example where this matters is understanding QoE. There is a potential virtuous cycle in which better QoE could increase user engagement which could yield new revenue opportunities. A less virtuous cycle is a cycle of blame, where frustrated users call someone at random and complain.

But there is a tension with respect to measurement done by the research community, or by any third party. What if some method or results are wrong? Wrong measurement does not just lead folks in the research community astray, although that certainly happens. It can lead others astray as well. The FCC's Measuring Broadband America (MBA) program originated in part from an impression at the FCC that consumers were not getting the speeds that were advertised by providers, and this impression arose in part from some faulty assumptions about how to analyze commercial data about access speed (Sevick and Wetzel, 2010).

Performance measurement challenges are only increasing with gigabit wired bandwidths and wireless network constraints. As providers try to compete on performance, or regulators consider imposing performance metrics and performance transparency, the research community could help by evaluating or calibrating tools used to measure high speed performance, to see if they are measuring the right thing (Bauer et al., 2010). But it order to do so, cooperation to validate methods and inferences is essential.

We also caveat that regulatory climates change, and nobody wants misguided regulation based on a misunderstanding of what is actually happening. The history of the Internet has relatively little regulation, and this was (in our view) essential to its early success. But regulatory attention to the Internet is generally increasing, and as with other sectors, regulation is generally mostly reactive, i.e., crisis-driven. The policy community uses the term *policy window* to describe an opportunity for a policy intervention in response to a crisis. Industry may reject the word "opportunity", and rather view regulatory intervention as the second half of a two-part disaster. But one thing the research community can do is to think ahead, and try to think through potential policy windows and the sorts of responses that would not be malformed and unnecessarily disruptive. This is another motivation for the operational community to make sure the scientific research community is informed about the real state of the Internet, and to facilitate unbiased measurement that can inform policy.

9 Next Steps: Reproducibility and Extension

To encourage reproducibility of our results and to facilitate community use of our raw data and tools, we are sharing our analysis scripts, and the underlying datasets via an interactive visualization interface[2] and query API. Our data management system allows interactive data exploration, near real-time views of interdomain links, and longitudinal views. We are working on enhancing the capabilities of this platform, so interested parties can issue more complicated queries about Internet topology and performance, eventually leveraging infrastructure geolocation capabilities. We also want to create functionality for operators to inject validation, or explanations

We hope to expand our measurement system to use the FCC's Measuring Broadband America (MBA) infrastructure (Federal Communications Commission, 2017) consisting of thousands of home routers. By doing so, we hope to obtain visibility into virtually all router-level interdomain links in the U.S. Improving the quality of data, as well as expanding to thousands of additional FCC MBA VPs will bring new challenges in system scale and data processing. Before we take on such a challenge, we seek feedback from the telecommunications policy community regarding how such a system should, or is likely, to be used, and what functionality and access control it should support.

This work was partly funded by NSF grants CNS-1413905 and CNS-1414177.

References

(2017). Grafana: The Open Platform for Beautiful Analytics and Monitoring. `https://grafana.com/`.

(2017). InfluxDB. `https://www.influxdata.com/`.

(2017). PeeringDB. `http://www.peeringdb.com`.

AfriNIC (2017). AfriNIC delegaged files. `ftp://ftp.afrinic.net/`.

Ager, B., N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger (2012). Anatomy of a Large European IXP. In *Proceedings of ACM SIGCOMM.*

Ahsan, S., V. Bajpai, J. Ott, and J. Schönwälder (2015). Measuring YouTube from dual-stacked hosts. In *Proceedings of the Passive and Active Measurement Conference (PAM).*

Anderson, C. (2015, June). New Opportunities for Test Deployment and Continued Analysis of Interconnection Performance. `http://www.measurementlab.net/blog/interconnection_and_measurement_update`.

Andrews, R. and S. Higginbotham (2013). YouTube sucks on French ISP Free, and French regulators want to know why. `http://gigaom.com/2013/01/02/youtube-sucks-on-french-isp-free-french-regulators-want-to-know-why/`.

[2]Measurement and Analysis of Interdomain Congestion: https://manic.caida.org

APNIC (2017). APNIC delegated files. ftp://ftp.apnic.net/pub/stats/apnic/.

ARIN (2017). ARIN delegated files. ftp://ftp.arin.net/pub/stats/arin/.

Bauer, S., D. Clark, and W. Lehr (2009). The Evolution of Internet Congestion. In *Research Conference on Communication, Information and Internet Policy (TPRC).*

Bauer, S., D. Clark, and W. Lehr (2010). Understanding Broadband Speed Measurements. In *TPRC.*

Brodkin, J. (2013a). Time Warner, Net Neutrality Foes Cry Foul Over Netflix Super HD Demands. http://arstechnica.com/business/2013/01/timewarner-net-neutrality-foes-cry-foul-netflix-requirements-for-super-hd/.

Brodkin, J. (2013b). Why YouTube Buffers: The Secret Deals that Make-and-break Online Video. http://arstechnica.com/information-technology/2013/07/why-youtube-buffers-the-secret-deals-that-make-and-break-online-video/.

Buckley, S. (2013). France Telecom and Google Entangled in Peering Fight. http://www.fiercetelecom.com/story/france-telecom-and-google-entangled-peering-fight/2013-01-07.

CAIDA (2017). AS Relationships. http://www.caida.org/data/as-relationships/.

Claffy, K., D. D. Clark, S. Bauer, and A. Dhamdhere (2016a). Policy Challenges in Mapping Internet Interdomain Congestion. In *Research Conference on Communication, Information and Internet Policy (TPRC).*

Claffy, K., D. D. Clark, S. Bauer, and A. D. Dhamdhere (2016b). Policy challenges in mapping internet interdomain congestion. In *Research Conference on Communication, Information and Internet Policy (TPRC).* Available at SSRN:https://ssrn.com/abstract=2756868.

Claffy, K., A. Dhamdhere, D. Clark, and S. Bauer (2016, August). First amended report of at&t independent measurement expert: Reporting requirements and measurement methods. https://www.caida.org/publications/papers/2016/att_ime_first_amended_report/.

Deng, L. and A. Kuzmanovic (2008, October). Monitoring Persistently Congested Internet Links. In *Proceedings of the International Conference on Network Protocols (ICNP).*

Dhamdhere, A., D. D. Clark, A. Gamero-Garrido, M. Luckie, R. K. P. Mok, G. Akiwate, K. Gogia, V. Bajpai, A. C. Snoeren, and kc claffy (2018, August). Inferring persistent interdomain congestion. In *Proceedings of SigComm '18.*

Dhamdhere, A. and C. Dovrolis (2010). The Internet is Flat: Modeling the Transition From a Transit Hierarchy to a Peering Mesh. In *Proceedings of ACM CoNEXT.*

Dinu, F. and T. S. E. Ng (2011). Inferring a Network Congestion Map with Zero Traffic Overhead.

Engebretson, J. (2010, November). Level 3/Comcast Dispute Revives Eyeball vs. Content Debate. http://www.telecompetitor.com/level-3comcast-dispute-revives-eyeball-vs-content-debate.

Engebretson, J. (2013, July). Behind the Level 3-Comcast peering settlement. http://www.telecompetitor.com/behind-the-level-3-comcast-peering-settlement/.

Feamster, N. (2016). Revealing Utilization at Internet Interconnection Points. In *Research Conference on Communication, Information and Internet Policy (TPRC)*.

Federal Communications Commission (2015a, July). In the matter of applications of at&t inc. and directv for consent to assign or transfer control of licenses and authorizations: Mb docket no. 14-90. https://apps.fcc.gov/edocs_public/attachmatch/FCC-15-94A1.pdf.

Federal Communications Commission (2015b, March). In the Matter of Protecting and Promoting the Open Internet: Report and Order on Remand, Declaratory Ruling, and Order: GN Docket No. 14-28. https://apps.fcc.gov/edocs_public/attachmatch/FCC-15-24A1.pdf.

Federal Communications Commission (2017). Measuring Broadband America. https://www.fcc.gov/general/measuring-broadband-america.

Ghobadi, M., Y. Cheng, A. Jain, and M. Mathis (2012). Trickle: Rate limiting YouTube video streaming. In *Proceedings of USENIX ATC*.

Gill, P., M. Arlitt, Z. Li, and A. Mahanti (2008). The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse? In *Proceedings of the Passive and Active Network Measurement Conference (PAM)*.

Google Video Quality Report (2016). https://www.google.com/get/videoqualityreport/.

Labovitz, C., S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian (2010, Aug). Internet inter-domain traffic. In *Proceedings of ACM SIGCOMM*.

LACNIC (2017). LACNIC delegated files. ftp://ftp.lacnic.net/pub/stats/lacnic/.

Luckie, M., A. Dhamdhere, D. Clark, B. Huffaker, and k. claffy (2014, Nov). Challenges in Inferring Internet Interdomain Congestion. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*.

Luckie, M., A. Dhamdhere, B. Huffaker, D. Clark, and k. claffy (2016, Nov). bdrmap: Inference of Borders Between IP Networks. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*.

Luckie, M., B. Huffaker, k. claffy, A. Dhamdhere, and V. Giotsas (2013, Oct). AS Relationships, Customer Cones, and Validation. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*.

M-Lab (2017). NDT (network diagnostic tool) - m-lab. https://www.measurementlab.net/tests/ndt/.

M-Lab Research Team (2014). ISP Interconnection and its Impact on Consumer Internet Performance - A Measurement Lab Consortium Technical Report. http://www.measurementlab.net/publications.

Marder, A. and J. M. Smith (2016). Map-it: Multipass accurate passive inferences from traceroute. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC).*

Netflix (2017). ISP Speed Index. `https://ispspeedindex.netflix.com/country/us/`.

Norton, W. (2010, August). The Art of Peering: The Peering Playbook. `http://drpeering.net/white-papers/Art-Of-Peering-The-Peering-Playbook.html`.

Packet Clearing House (2017). Full Exchange Point Dataset. `https://prefix.pch.net/applications/ixpdir/menu_download.php`.

Rao, A., Y. sup Lim, C. Barakat, A. Legout, D. Towsley, and W. Dabbous (2011). Network characteristics of video streaming traffic. In *Proceedings of ACM CoNEXT.*

RIPE (2017). Ripe ris. `http://www.ripe.net/ris/`.

RIPE NCC (2017). RIPE NCC delegated files. `ftp://ftp.ripe.net/ripe/stats/`.

Sandvine (2013). Global Internet Phenomenon Report, 1H13.

Sevick and Wetzel (2010). FCC Relying on Faulty ISP Performance Data from comScore.

Sundaresan, S., A. Dhamdhere, M. Allman, and k. claffy (2017, Nov). TCP Congestion Signatures. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC).*

Sundaresan, S., D. Lee, X. Deng, Y. Feng, and A. Dhamdhere (2017, Nov). Challenges in Inferring Internet Congestion Using Throughput Measurements. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC).*

Taylor, M. (2014, May). Observations of an Internet Middleman. `http://blog.level3.com/global-connectivity/observations-internet-middleman/`.

University of Oregon (2017). RouteViews. `http://www.routeviews.org/`.

Varvello, M., J. Blackburn, D. Naylor, and K. Papagiannaki (2016). EYEORG: A Platform For Crowdsourcing Web Quality Of Experience Measurements. In *Proceedings of ACM CoNEXT.*

Verizon (2013, June). Unbalanced Peering, and the Real Story Behind the Verizon/Cogent Dispute. `http://publicpolicy.verizon.com/blog/entry/unbalanced-peering-and-the-real-story-behind-the-verizon-cogent-dispute`.

Yap, K.-K., M. Motiwala, J. Rahe, S. Padgett, M. Holliman, G. Baldus, M. Hines, T. Kim, A. Narayanan, A. Jain, V. Lin, C. Rice, B. Rogan, A. Singh, B. Tanaka, M. Verma, P. Sood, M. Tariq, M. Tierney, D. Trumic, V. Valancius, C. Ying, M. Kallahalla, B. Koley, and A. Vahdat (2017). Taking the edge off with espresso: Scale, reliability and programmability for global internet peering. In *Proceedings of ACM SIGCOMM.*