

POLICY CHALLENGES IN MAPPING INTERNET INTERDOMAIN CONGESTION

kc claffy, David D. Clark, Steve Bauer, and Amogh Dhamdhere

ABSTRACT

Interconnection links connecting broadband access providers with their peers, transit providers and major content providers, are a potential point of discriminatory treatment and impairment of user experience. However, adequate data to shed light on this situation is lacking, and different actors can put forward opportunistic interpretations of data to support their points of view. In this article, we introduce a topology-aware model of interconnection to elucidate our own beliefs about how to measure interconnection links of access providers and how policy-makers should interpret the results. We use six case studies that show how our conceptual model can guide a critical analysis of what is or should be measured and reported, and how to soundly interpret these measurements.

Keywords: Internet measurement; Internet interconnection; Internet merger conditions; Internet performance; independent measurement expert

Even the best publicly available data about the global interconnection system that carries most of the world's communications traffic is incomplete and of unknown accuracy. There is no map of physical link locations, capacity, utilization, or interconnection arrangements. Recent public policy challenges have triggered the need for more transparency into the state of Internet interconnection. While concerns about interconnection might arise in any part of the Internet, the broadband access providers, who serve the public and provide retail access to the Internet, have attracted the most attention with respect to public policy issues. In particular, as a result of US telecommunications policy over the last 20 years, the Internet industry structure has evolved toward a state where carriage (infrastructure) networks increasingly own content and monetize transmission of that content on their networks, creating

kc claffy: CAIDA, USD

David D. Clark: CSAIL, MIT

Steve Bauer: CSAIL, MIT

Amogh Dhamdhere: CAIDA, UCSD, and Amazon

DOI: 10.5325/jinfopoli.10.2020.0001



JOURNAL OF INFORMATION POLICY, Volume 10, 2020

This work is licensed under Creative Commons Attribution CC-BY-NC-ND

naturally misaligned incentives regarding interconnection with other content providers. Specifically, as stated in the Federal Communication Commission's (FCC) 2015 Open Internet Order "broadband Internet access providers have the ability to use terms of interconnection to disadvantage edge providers and that consumers' ability to respond to unjust or unreasonable broadband provider practices are limited by switching costs."^{1,2}

Although regulators have gained some experience measuring the operating parameters of broadband Internet access links,³ and accommodated extensive debate to inform rule-making on what constitutes reasonable network management of these links,⁴ they (at least in the United States) have no experience with the measurement of Internet interconnections nor in analyzing its role in user quality of experience (QoE). And yet, links connecting access providers with their peers, transit providers and major content providers, are a potential point of discriminatory treatment and impairment of user experience. In the United States, the FCC has asserted regulatory authority over those links, although they have acknowledged they lack sufficient expertise to develop appropriate regulations thus far.⁵ Without a basis of knowledge that relates measurement to justified inferences about actual impairment, different actors can put forward opportunistic interpretations

1. In this article, we reference various orders of the FCC, including the orders related to network neutrality from 2010 to 2015. As we publish this article in 2020, the current FCC has taken steps to reverse these orders and remove itself from any regulation of issues such as Internet neutrality. We recognize that this event means that in the United States, the FCC is not likely at present to act on the issues that we discuss here. However, administrations change, and this article is not relevant just to the United States. The issues we raise here are relevant to any regulator concerned with the state of the Internet, and regulators in different countries operate with different scopes of authority.

2. Federal Communications Commission, "In the Matter of Protecting and Promoting the Open Internet."

3. FCC, "*Measuring Broadband America*"; Bauer, Clark, and Lehr, "The Evolution of Internet Congestion"; Bauer, Clark, and Lehr, "Understanding Broadband Speed Measurements."

4. Federal Communications Commission, "*Notice of Proposed Rulemaking*"; Federal Communications Commission, "*In the Matter of Protecting and Promoting the Open Internet: REPORT.*"

5. Quoting from the 2016 ruling, "As a result, the Commission concluded that it could regulate interconnection arrangements under Title II as a component of broadband service. It refrained, however, from applying the General Conduct Rule or any of the bright-line rules to interconnection arrangements because, given that it lack[ed] [a] background in practices addressing Internet traffic exchange, it would be premature to adopt prescriptive rules to address any problems that have arisen or may arise. Rather, it explained that interconnection disputes would be evaluated on a case-by-case basis under sections 201, 202, and 208 of the Communications Act." United States Court of Appeals for the D.C. Circuit, "*United States Telecom Association, et al., petitioners v. FCC,*" p. 52, citations omitted.

of data to support their points of view. The recent proliferation of performance-related data and claims leaves policymakers, researchers, and the general public with the tremendous challenge of interpreting it all.

In this context, the FCC has turned toward the research community for help with two measurement challenges: measurement of interconnection links and measurement of overall QoE for users accessing specific services. The FCC has expressed interest in augmenting its Measuring Broadband America (MBA) program with both types of measurements and is evaluating measurement methodologies for users accessing specific services. Unfortunately, this sort of measurement has not been a high priority for the academic research community nor its funding agencies.

There is no known way for a third party to remotely obtain direct measurements of basic parameters of an interconnection link, for example, capacity and utilization, and commercial concerns regarding sharing data or access to instrumentation belonging to the operators themselves generally prevent researchers from being able to validate measurements or methods. Thus, like the FCC, the research community cannot bring to this discussion much experience with this sort of measurement. Absent a regulatory requirement for mandatory sharing of such data, the research community is not in a position to offer concrete advice to regulators. But researchers can bring objectivity, insights into how to think about the problem, and suggestions for how to start gathering data to inform regulatory trajectories.

In this article, we take up this challenge in four parts. First, we provide background on the range of methods of harmful discrimination against interconnecting parties (“Approaches to Interconnection Discrimination and Policy Responses”). We then introduce a new topology-aware model (“Topology-Aware Hierarchical Model of Rich Interconnection”) that captures some of this complexity by recognizing hierarchical structure in interconnection architectures. We discuss the complexity of trying to interpret measurement (“Relating Performance Measures to QoE”) and explain how the model helps to explicate the ways that measurement data aggregated at different levels can serve different purposes from operational network management to regulatory oversight (“Different Objectives for Performance Analysis”). Aggregation of measurement results across different scopes can characterize how pervasive congestion is across a set of possible paths between two interconnected parties: is evidence of congestion observed only on a single link, or an aggregate set of links, in a single metropolitan region, or more broadly? Finally, we discuss measurements that do not focus on a single link (at a point of interconnection) but on

longer path segments, perhaps the complete path from the source to the destination (“Measurements of Path Segments”).

We do not intend this model to be predictive, that is, to allow forecasting of congestion or any other network dynamics, but rather as a *conceptual model*, to add clarity to a recent proliferation of data and claims, and to elucidate our own conclusions about how to define and measure interconnection-related performance problems, and the complexity of interpretation induced by different choices.

To concretely demonstrate the utility of our conceptual model, we apply it to the examination of six measurement projects (“Applying the Model to Specific Measurement Data Sets”) that different stakeholders have propounded to illuminate their view of the landscape of interconnection performance. These projects span data sets offered by access providers, edge providers, academic researchers, and mandated by the FCC. In the final case study, we describe our efforts as the independent measurement expert (IME) that worked with the FCC and AT&T in establishing a measurement methodology for reporting on the state of AT&T’s interconnection links. Finally, we offer some conclusions, implications, and recommendations for researchers and policymakers.

The key contributions of this article are:

- We construct a topology-aware model of interconnection that distinguishes itself from existing models by capturing nested aggregations of links between access providers and their interconnecting parties.
- We discuss the trade-offs and limitations of examining data at different aggregation granularities, and how to use them to evaluate the relative significance of performance impairment.
- We demonstrate the utility of our conceptual model by applying it to six case studies that show how it can guide informed decisions regarding interpreting, or mandating, measurements intended to reveal harmful (impairment-inducing) congestion at interconnection links.

Approaches to Interconnection Discrimination and Policy Responses

Network operators could impose several potentially harmful forms of discrimination in the context of interconnection; we describe five sorts of discrimination in this section. First, *differential treatment of packets* could occur at interconnection links. However, with today’s typical Internet

usage, content generally flows *toward* the access provider, in which case discrimination across the link would have to occur not on the access provider's router but on the upstream router operated by the interconnecting party, which makes this form of discriminatory treatment less likely. Once traffic is on the access provider's side of the interconnection link, from a regulatory perspective it would count as discrimination within the access Internet service provider (ISP), which would be prohibited under the various prior FCC orders unless it was justified as reasonable network management.⁶

A second form of discrimination uses *routing policy* internal to the access network: an access provider could engineer its network so traffic from different interconnecting parties traverses different links within the access ISP, and underprovision some of those links. Those underprovisioned links would afflict only the interconnecting parties using such links, without requiring any selective discrimination among packets passing over any link. However, an interconnecting party can detect this sort of congestion and has a strong incentive (especially if it is paying for interconnection) to do so and complain when it occurs.

A third and nontechnical approach is *price discrimination* across direct interconnections. Content from different directly connected content providers travels across different physical links, and the contracts for those links reflect terms that each pair of interconnecting parties may craft separately, perhaps with different costs for equal capacity. Since these agreements are almost always covered by nondisclosure agreements, there is no way for a given content provider (or a third-party observer) to tell whether two content providers seeking direct interconnection are receiving equivalent treatment. Since agreements may include complex business terms, including commitments to rates of capacity expansion over time or restrictions on routing, it might be difficult to compare treatments of different interconnecting parties, even if one could see the contracts.

Fourth, another nontechnical approach to discrimination is for an access provider to *limit the number of interconnections* with one

6. An episode of traffic differentiation outside the access provider occurred in 2014 when Cogent, a transit provider for Netflix among others, began to mark packets using the IP Type of Service (ToS) bits to classify their customers' traffic into wholesale and retail categories, and prioritized traffic belonging to retail customers (Kilmer, "M-Labs data and Cogent DSCP markings," see quote embedded in group discussion). In this case, an upstream (i.e., not broadband access) provider materially altered the performance of broadband users' traffic by employing traffic differentiation techniques.

content provider more than another. Fewer points of connection might disadvantage the content provider, increasing latency for traffic that flowed longer distances across the access ISP. But this approach would also increase transit costs internal to the access ISP, so it is often in the best interest of both parties to interconnect at many points. Indeed, an obligation of the opposite sort is often part of an interconnection agreement—the parties are required to interconnect at least some minimum number of points.

Finally, and the approach that has been the subject of media and policy attention in recent years, the parties responsible for an interconnection link could simply *fail to upgrade capacity* of that link, often due to business disputes, when evidence of congestion manifests. It requires the agreement of both parties with an interconnection to implement an upgrade to the capacity, and either party can refuse to cooperate. Thus far, the highest profile dispute about interconnection between access and content providers involved Netflix and some of its transit providers claiming that access providers were exercising this form of discrimination. When Netflix started delivering content streamed over the Internet in 2007, it first used its own servers in five locations in the United States. As traffic grew, Netflix began enlisting third-party content delivery networks (CDNs) such as Akamai and Limelight in 2008. In 2012, Netflix began moving away from third-party CDNs and started using transit providers and simultaneously started installing its own content servers (Netflix OpenConnect) across the Internet to interconnect directly with major access ISPs.⁷ Sometimes these negotiations for direct interconnection became contentious, with Netflix arguing that if it introduced its traffic into the access network at a point close to the consumers that requested it, Netflix was carrying most of the cost of the delivery and reducing the costs borne by the access provider, so the interconnection should be settlement-free, while major access ISPs asserted that these interconnections were commercial arrangements that required payment.⁸ Before negotiating direct interconnection with access ISPs, Netflix used its existing connections (where it was a customer) with its Tier 1 transit providers to transmit traffic to Netflix customers on access ISPs. Those links lacked sufficient capacity to keep up with the growth in Netflix traffic, leading to massive congestion on the Tier 1 interconnection links,

7. Netflix, “*Petition to Deny of Netflix, Inc.*”

8. Florence, “The Case Against ISP Tolls”; Khoury, “Comcast Response to Netflix”; Brodtkin, “Time Warner, net neutrality foes cry”; Brodtkin, “After Netflix pays Comcast, speeds improve 65%.”

which impaired both Netflix traffic and any other traffic unfortunate enough to be passing over the same interconnections.⁹ In the United States, the FCC had until this point avoided intervening in negotiation of commercial terms between large Internet players, but in this case, the FCC pressured the parties to resolve these conflicts quickly to relieve ongoing impairments. In other countries, interconnection disputes ended up in court: in France, the competition court ruled with respect to a Tier 1 provider (Cogent) delivering large volumes of content into France Telecom (FT) that FT had the right to demand payment from Cogent.¹⁰

The consumer complaints and media coverage about these interconnection disputes brought attention to the power of access providers to impose terms for interconnection and on the important role that interconnection plays in the stability and function of the Internet. In the 2015 Open Internet Order, the FCC pivoted from their 2010 position that excluded interconnection from their purview and explicitly asserted that their authority extends to the regulation of interconnection with access providers.¹¹ They did not impose any overall regulations on Internet interconnection while that order was in effect, but they did impose regulatory requirements as conditions on two large mergers involving access providers.¹²

The first such merger agreement was between AT&T and DirectTV in 2015.¹³ Responding to specific concerns in the merger review process that interconnection links could be a locus of unreasonable discrimination, the final agreement imposed a requirement on AT&T to report to the FCC the contractual terms of their interconnections with major peering and paid peering partners, as well as performance parameters of these interconnections (“Measuring the Interconnection Links of AT&T”).¹⁴ Similar

9. Luckie et al., “Challenges in Inferring Internet.”

10. Republique Francaise: Autorite de la concurrence, “*Internet Traffic—Peering Agreements.*”

11. Federal Communications Commission, “In the Matter of Protecting and Promoting the Open Internet: REPORT AND ORDER ON REMAND,” p. 295.

12. As the 2015 order (*ibid.*, para 294) notes, in the 2010 Open Internet Order, the Commission applied its Open Internet rules “only as far as the limits of a broadband provider’s control over the transmission of data to or from its broadband customers,” and excluded the exchange of traffic between networks from the scope of the rules. In the 2014 Open Internet NPRM, the Commission tentatively concluded that it should maintain this approach, but explicitly sought comment on suggestions that the Commission should expand the scope of the Open Internet rules to cover issues related to Internet traffic exchange. (See also footnote 1.)

13. Federal Communications Commission, “In the Matter of Applications of AT&T Inc. and DIRECTV For Consent to Assign or Transfer.”

14. The FCC was using this occasion as an opportunity to educate itself and to gain experience about what sort of data should actually be gathered, and how to interpret it.

concerns arose during the 2016 Time Warner/Charter merger, but instead of imposing measurement requirements, they imposed interconnection traffic volume reporting requirements, as well as constraints on the business relationships that the combined entity could negotiate with interconnecting parties.¹⁵

The most likely forms of discriminatory treatment we identify here include nontechnical approaches (e.g., price discrimination), but the visible disputes and ongoing concerns about congested interconnection links motivate the focus of this article on measurement of link congestion and actual achieved end-to-end performance.

Measurement of Internet interconnection performance is technically complex, for many reasons identified in the networking literature.¹⁶ In this work, we focus on three reasons with current policy implications. First, understanding performance of an interconnection link requires measuring several parameters, for example, utilization, loss rates, and variation in latency. As utilization reaches 100% (e.g., congestion begins to manifest), both latency and loss rate are affected. Operators that control interconnection links could measure such parameters for those links, although accurate assessment of these parameters may require cooperation of the operator at the other end of the link (“Measuring the Interconnection Links of AT&T”), and we are not aware of any commercial operators voluntarily cooperating to do so. Second, modern interconnection practices render

15. The Time Warner/Charter merger agreement includes the following reporting requirement (Federal Communications Commission, “In the Matter of Applications of Charter Communications, Inc., Time Warner Cable Inc., and Advance/Newhouse Partnership For Consent to Assign or Transfer Control of Licenses and Authorizations: MB Docket No. 15-149,” Appendix B.I): Information for each Interconnect Exchange Point, which shall include, as of the date that is the last day of the calendar quarter preceding the Report:

- Each Interconnection Party interconnected with the Company at that Interconnect Exchange Point.
- For each Interconnection Party, the aggregate link capacity between the Company and each Interconnection Party at that Interconnect Exchange Point.
- For each Interconnection Party, traffic exchange, in each direction, as measured by the 95th percentile method.
- For each port through which traffic is exchanged with an Interconnection Party, the percentage time within the reporting period that the port was over 75% capacity in the dominant direction.

The agreement also requires the new entity to provide settlement-free peering for 7 years, although it lays out a complex set of obligations on the parties in order to obtain these settlement-free arrangements, including the number of peering locations, rates of growth in traffic, and restrictions on routing and delivery of non-customer traffic. We revisit this issue in “Different Objectives for Performance Analysis.”

16. Luckie et al., “Challenges in Inferring Internet.”

it potentially necessary to measure multiple links, sometimes in different cities, in an integrated manner, to assess overall performance degradation (“Different Objectives for Performance Analysis”). Third, neither regulators nor researchers clearly understand how to relate variation in measured performance with impairment of the user’s QoE (“Relating Performance Measures to QoE”).¹⁷

Topology-Aware Hierarchical Model of Rich Interconnection

Researchers have explored increasingly refined economic models of Internet interconnection, including modeling paid peering, game-theoretic justifications of settlement-free peering, the business effects of transit versus peering relationships, transit pricing and provisioning of service tiers, simple pricing schemes that approximate complex revenue-maximizing pricing, network formation models, and more generally the evolution of the Internet from a game theoretic perspective.¹⁸ We offer a new conceptual model of interconnection (illustrated in Figures 1 and 2), which captures hierarchical aggregations of links between an access provider and

17. Quality of experience, or QoE, refers to a subjective characterization by users of their level of satisfaction using a particular application at a given moment. This term contrasts with Quality of Service (QoS) metrics such as throughput.

18. Shrimali and Kumar, “Can Bill-and-Keep Peering Be Mutually Beneficial?”; Shrimali and Kumar, “Paid Peering Among Internet Service Providers”; Shrimali and Kumar, “Bill-and-keep peering”; Dhamdhere, Dovrolis, and Francois, “A Value-based Framework for Internet Peering Agreements”; Baake and Wichmann, “On the Economics of Internet Peering”; Badasyan and Chakrabarti, “A simple game-theoretic analysis of peering and transit contracting”; Anshelevich, Shepherd, and Wilfong, “Strategic Network Formation Through Peering”; Anshelevich and Wilfong, “Network Formation and Routing by Strategic”; He and Walrand, “Pricing and revenue sharing strategies for internet”; Dai and Jordan, “Modeling ISP tier design”; Stanojevic, Laoutaris, and Rodriguez, “On economic heavy hitters: Shapley value analysis”; Valancius et al., “How Many Tiers?: Pricing in the Internet”; Shakkottai et al., “The Price of Simplicity”; Johari and Tsitsiklis, “Routing and peering in a competitive Internet”; Johari, Mannor, and Tsitsiklis, “A contract-based model for directed network formation”; Arcaute, Johari, and Mannor, “Network formation: Bilateral contracting”; Arcaute, Johari, and Mannor, “Local two-stage myopic dynamics for network”; Lodhi, Dhamdhere, and Dovrolis, “GENESIS: An Agent-based Model of Interdomain Network Formation”; Lodhi, Dhamdhere, and Dovrolis, “GENESIS-CBA: An Agent-based Model of Peer Evaluation and Selection”; Lodhi and Dovrolis, “A Network Formation Model for Internet”; Lodhi, Dhamdhere, and Dovrolis, “Open Peering by Internet Transit Providers: Peer?”; Chang and Jamin, “To Peer or not to Peer: Modeling the Evolution of the Internet”; Holme, Karlin, and Forrest, “An Integrated Model of Traffic, Geography and Economy”; Dhamdhere and Dovrolis, “The Internet is Flat: Modeling the Transition from a Transit”; Meirum, Mannor, and Orda, “Network Formation Games and the Internet Structure”; Ma, Lui, and Misra, “On the Evolution of the Internet Economic Ecosystem.”

their interconnecting parties. These nested aggregation levels are implicit in many measurements and often not explained in published assertions and data sets. We use this model to explicate two other important factors in understanding interconnection as it relates to performance impairments. First, interconnecting parties may take direct or indirect paths to reach access providers; indirect paths go through at least one peer or transit provider before reaching the access provider. Second, either party could impose performance bottlenecks on the other at some other location than the interconnection link itself, detection of which would require measurements of the overall path, or at least the right path segments. We first define the five levels of aggregation, then differentiate among direct paths, indirect paths, and potential paths between interconnecting parties.

Levels of Aggregation in Interconnection Relationships

Figure 1 illustrates the five aggregation granularities of our model. It shows two points of presence (PoPs, or physical locations) in New York, labeled NYC 0 and NYC 1, as well as locations in other cities. There is an access provider (in blue) providing service in a number of cities. This network is connected to other ISPs, which might be transit providers, or peers (such as other access providers).

Individual link: connection between two physical ports, either directly connected or through a switch (an Internet exchange). Individual links

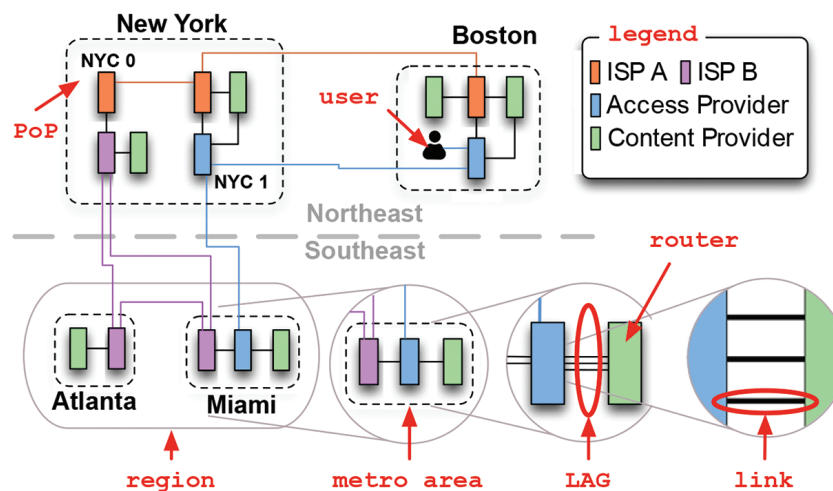


FIGURE 1 Hierarchical Aggregations of Links between Access Provider and Interconnecting Parties.

today are sometimes one (1 GB) but more usually ten (10 GB) gigabit connections. Links with 100 GB capacity are being used in some cases.

Link aggregation group (LAG): combination of physical ports and a more reliable and higher bandwidth connection between two parties, with a single router balancing traffic load across all links in a LAG. A LAG is often the lowest aggregation that network operators consider, as traffic management functionality on modern routers and switches renders a LAG essentially indistinguishable from a single link of aggregate capacity. More important from a measurement perspective, aggregate performance of the LAG is generally representative of performance over the individual links in the group.¹⁹ Currently, LAGs of large providers often consist of multiple 10-GB links. Multichassis LAGs implement node-level redundancy over a single logical link-level connection between two points.

Metro area group: aggregate of all link and LAG interconnections between two parties in a metro area. In contrast to individual links in a LAG, different LAGs connecting two parties in a metro area can have different traffic characteristics. Traffic is often not load-balanced equally across all interconnections in a metro area, so summary statistics about a metro area may not be representative of performance of all individual interconnections that compose the metro area group. Traffic flow across metro areas reflects routing policy as well as application and service considerations such as which cache serves a particular request. A LAG is part of a metro area if one end of the link exists in that metro area.²⁰

Region group: aggregate of all LAGs in multiple metro areas with geographic proximity between two parties, for example, Boston and NYC in a northeast region. This granularity matters because interconnection LAGs in one metro area can potentially substitute for those in a nearby metro area. For example, if LAGs in Boston are filled to capacity but available capacity exists in NYC, it may not be necessary to augment interconnection capacity in Boston. However, a path from the alternative metro area increases latency, and an increased hop count increases the potential that a problem will occur between user and server. The notion of *substitutability* is essential to reasoning about aggregation from a measurement perspective (“Different Objectives for Performance Analysis”).

Provider-wide: aggregate of all direct LAGs between two parties. This includes LAGs in potentially diverse geographic locations, for example, Boston, Chicago, and San Diego.

19. IEEE, “802.1AX-2014 – IEEE Standard for Local and metropolitan.”

20. Normally, both ends of large-volume interconnection LAGs are in one metro area.

In “Different Objectives for Performance Analysis,” we discuss the utility of measurements at different levels of aggregation for different purposes.

Direct, Indirect, and Potential Connections

In addition to capturing structure in the set of interconnection *links* between two organizations, our model recognizes structure in the set of *paths* between two organizations. Interconnecting parties often have both direct and indirect paths to each other. Figure 2a depicts a party (green) directly connecting to an access provider (blue), and Figure 2b depicts alternative indirect paths through other networks (orange) to that access provider. These alternative paths might be through a transit network that connects to both the content and access providers, or through a CDN like Akamai. The diverse set of paths connecting two parties may have different performance and economic characteristics. Although end users do not care about what path traffic takes unless it impacts their perceived performance, it can make a huge economic difference to the interconnecting parties.

In addition to direct and indirect paths utilized between two networks, other *potential* indirect paths may exist—paths that are potentially available from a business and routing perspective, but not actively used by default. A nuance of defining this superset of potential paths is knowing what action is necessary to make a path available for use, for example, is it available on reasonable terms? In general, a third party has no way of either identifying or measuring the set of potential paths.²¹

21. Notably, in the Time Warner/Charter merger, the order’s restriction on settlement-free peering obligations specifically takes into account the existence of indirect paths: “In the event that the Interconnection Party begins conveying data to or from New Charter that was previously conveyed to or from New Charter by a third party, the parties shall account for this additional data transfer as the Interconnection Party’s own for the purposes of measuring growth rates during subsequent measuring periods. The parties shall not count in the growth rate any portion of that incremental traffic that was previously being delivered to New Charter by third parties.” Furthermore, the interconnecting party cannot use the settlement-free interconnection to deliver traffic unless the source of the traffic is a customer, and the agreement explicitly excludes the option of a customer that purchased a path only to Charter. Federal Communications Commission, “In the Matter of Applications of Charter Communications, Inc., Time Warner Cable Inc., and Advance/Newhouse Partnership For Consent to Assign or Transfer Control of Licenses,” para 456, Appendix B.III.

Applying the Conceptual Model to Analyze Performance Measurements

We apply our model to help understand the implications of performance and impairment measurements for different actors in the ecosystem. We argue that data at different aggregation granularities, where each granularity reveals some characteristics of interconnection between two networks while obscuring others, will be useful in different contexts. We first consider what performance measures are relevant to assessing the quality of a link. We then discuss the different purposes for which data might be used—operational network management, business relationship, and regulatory oversight. In “Measurements of Path Segments” we consider additional measurements that can contribute to a more complete picture of performance: measures along a path segment that include several links (including end-to-end measurements) and higher-level measurements related to QoE. The background in this section will prepare us to use the

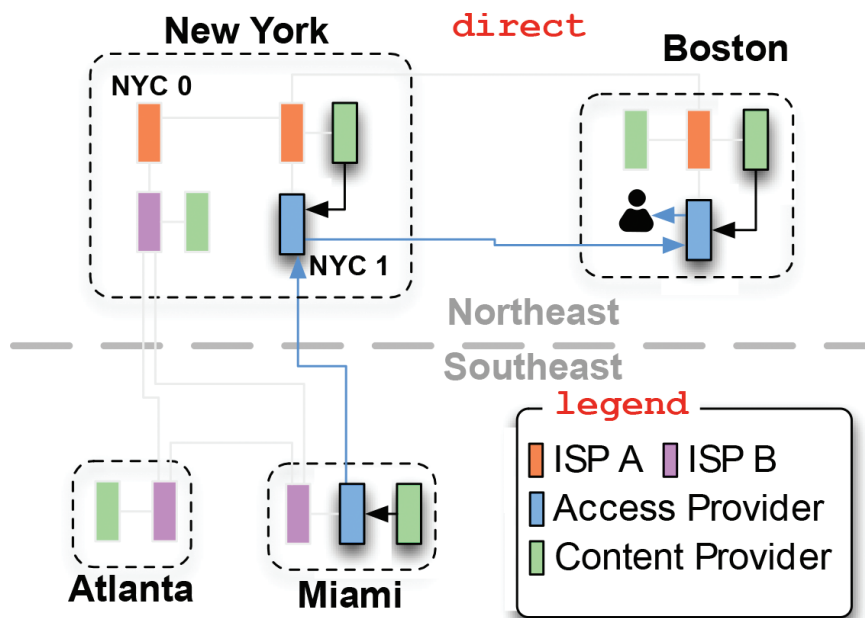


FIGURE 2A Direct Paths between Interconnecting Parties. The content provider (green) is directly connected to the access provider (blue).

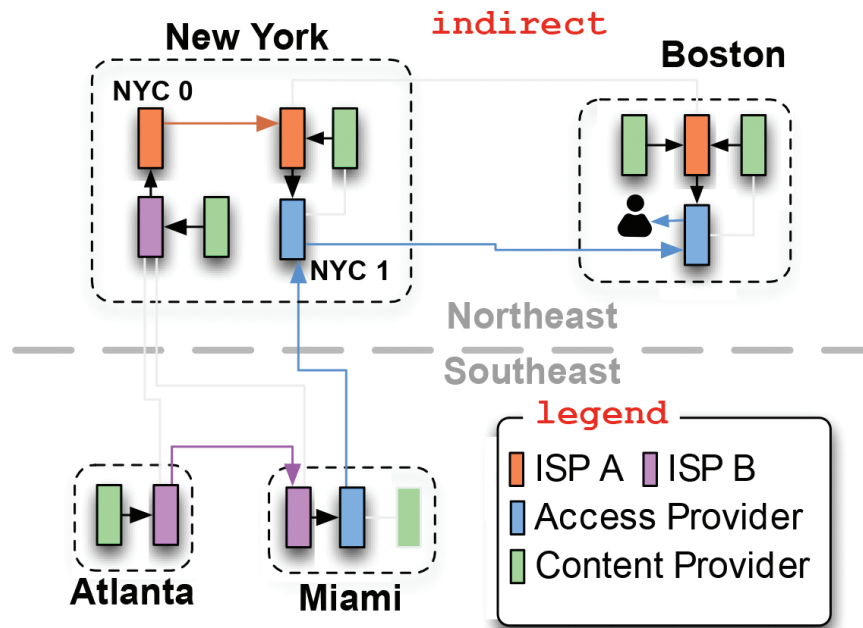


FIGURE 2B Indirect Paths between Interconnecting Parties. In this case the content provider (green) is not connected directly to the access provider (blue) but to intermediate ISPs (orange and lavender) which in turn connected to the access provider.

model to critically examine six concrete case studies in “Approaches to Interconnection Discrimination and Policy Responses” to gain insight into what each approach can and cannot say about interconnection.

Relating Performance Measures to QoE

We begin with a caveat that applies to all aggregation granularities. Measurement of *link utilization* is a common way for an operator to assess link (or LAG) behavior, but interpreting utilization measurements and its relation to congestion and impairment is not always straightforward. In the past, when ISPs typically interconnected with a single LAG and there were few options for alternative routes, a single LAG showing persistent plateaus of essentially full utilization (such as illustrated in Figure 3a) implied congestion: unserved demand between the two interconnecting parties and impairment of flows crossing the LAG. Since Internet traffic typically shows a diurnal variation, a plateau suggested that the LAG was fully utilized well before (and after) its peak demand and was thus underprovisioned, losing some packets, delaying others, and reducing




FIGURE 3 Utilization (Left) and Loss (Right) Plots for a Level-3 100-Gbps Interconnect with an Unnamed Broadband Access Provider in Dallas for One Week Ending April 3, 2014. The significant loss confirms that the high utilization plotted for the same interval induced impairment for consumers of traffic over this link. (The loss rate indicated is around 1K packets/second. What is an acceptable loss rate at interconnection links is an important question, but not one we discuss here.)

throughput of flows. In other words, measurements from a single LAG served as strong evidence of significant impairment.

Today's sophisticated traffic engineering and interconnection practices preclude this conclusion without additional information, particularly in the case of aggregates between a content provider and an access provider. Business terms, loads on servers, and internal LAG capacities may induce an operator to fully load one LAG before starting to load another LAG beyond a certain level. Thus, if multiple LAGs in an aggregate (e.g., a metro area) connect two providers, assessing congestion in that area requires consideration of all of them. If the edge provider can control content source selection, one cannot verify congestion without additional measurement, such as the packet loss ("discards") plotted on the right of Figure 3.

A second issue with performance measurements is how to relate them to the higher-level question of whether the observed measures actually relate to any degradation of the user experience.

The relationship between congestion and degradation in actual experience is complex. Understanding a little about the dynamics of the predominant Internet transport protocol, Transmission Control Protocol (TCP), provides some insight into the complexity. When excess traffic arrives at the ingress to a link, a queue of packets forms, and the holding time of the packets in this queue adds to the delay across the link. This variation, usually called *jitter*, is a signal of potential congestion and may impair latency-sensitive applications such as real-time communications and multiplayer games. It normally does not impair video streaming. If the queue becomes full and excess traffic continues to arrive, the router will drop some arriving traffic. Dropped packets can impair service, but routers drop packets for other reasons, and applications on the sending end will typically detect and retransmit dropped packets. More significantly, TCP

(which most Internet applications use as a transport protocol) treats a lost packet as a signal of congestion and reduces the sending rate. Dropped packets are thus a necessary part of the transport control mechanism that regulates the sending rate of traffic sources to match the capacity of the link. Reductions in throughput can be a severe impairment to applications such as streaming content, but of almost no concern to a delay-tolerant application such as email. In general, while dropped packets at the ingress to a link may signal congestion, there is no way to measure the link to determine how much excess traffic an endpoint could send if the link had more capacity, since the content sources control the sending rate, not the link ingress.

In addition to normal TCP congestion control behavior, applications may use even more sophisticated ways to adapt to an indication of a congested link. Large CDNs or content networks typically interconnect to broadband access networks at multiple points and can engineer server placement and routing policy to avoid congested links. Alternatively, content providers may adapt to a signal of congestion along a path to a recipient by degrading the content encoding to fit into a lower data rate, for example, from high definition to standard definition video encoding. This adaptation is impossible to detect with active measurement.²² These factors prevent the use of measurements of individual links to prove impairment of the user experience.

In relating measures of quality of service (QoS) (measurable performance metrics) to QoE, operators and researchers tend to make intuitive assumptions. If an aggregate, for example, a metro area, shows persistent congestion for many hours a day, users are likely experiencing negative consequences. In other cases, the argument is harder to make. Imagine that a content provider reports that over a 24-hour period, overall traffic into a metro area shows periods where the throughput of individual flows drops by 5%. First, the congestion is not necessarily causing this drop; users could be downloading different sorts of content, or using different devices, which changes the mix of content coding.

However, if the drop is due to congestion, would users perceive a drop in QoE? The answer is application specific. Applications such as real-time communication (VoIP, teleconferencing) and multiplayer games are

22. One could use passive monitoring at the right link to collect evidence of adaptive coding, or try to directly assess the impairment of the QoE assuming a user can distinguish encoding qualities.

sensitive to jitter. Streaming content (audio and video) applications are less affected by jitter, since they buffer some content at the receiver to smooth out variation in arrival time, but a reduction in bandwidth may require the source to reduce the encoding quality, which often affects QoE. Jitter or reduced throughput (unless severe) has less effect on interactive applications like web browsing and negligibly affects background traffic. Although considerable research literature exists in the space of QoE, it is not easy to translate these results into operational criteria for asserting impairment.²³ This sort of work has not been thus far of high priority to the research community, nor to its funding agencies. In the United States, the FCC has recognized the need for this sort of basic research and co-sponsored (with the National Science Foundation) a workshop on QoE, one goal of which was to explore the relationship between observed operational measures of network performance and impairments to QoE.²⁴

Different Objectives for Performance Analysis

Different actors in the ecosystem have different needs and motivations for performance analysis, and may find data at different aggregates most useful.

Operational network management: network operators are concerned with failures of network elements, as well as changing traffic patterns over time that may lead to overloads or underutilization of links, which trigger traffic engineering (at one time scale) and changes in network provisioning (at a longer time scale). To detect failures, it is necessary to look at the performance of individual links. Normally, if a link completely fails, some relevant equipment will raise an alert, but links often degrade before they fail outright. To understand traffic patterns and how they are changing, it is probably most useful to track performance at the level of a LAG. Capacity can be added to a LAG by adding another link, and on a longer time scale new LAGs can be provisioned to carry new traffic patterns.

Business relationships and negotiations: Business agreements between interconnected parties often include requirements to interconnect at multiple geographically distinct locations. Capacity commitments would thus

23. Moller and Raake, "Quality of Experience: Advanced Concepts, Applications and Methods."

24. The report of the first FCC/NSF workshop is available at Bustamante, Clark, and Feamster, "Workshop on Tracking Quality of Experience in the Internet."

normally be at the metro level of aggregation, which in many cases will be the same as LAG aggregation. However, there are circumstances where there are multiple LAGs in the same metro area. If there are multiple LAGs in a metro area, interconnecting parties will want to confirm that traffic is routed over them in a balanced way (e.g., that the two parties have good traffic engineering practices), but contract compliance will most likely hinge on adequacy of total metro capacity.

Regulatory agencies: Regulatory agencies may have a number of distinct missions. Responsibilities vary by country, but can include determining unreasonable business practices and imposing remedies, responding to harms to consumers, acting to improve deployment of broadband in unserved areas, and public safety. Visibility into network performance seems necessary as part of those first two objectives, although additional data (such as pricing data and terms of contracts) may also be important. We do not discuss here the full range of what data might be useful, nor under what authority different regulators collect data; we focus on what visibility regulators need into network performance, again with a focus on interconnection.

Consistent with a long history of telephony regulation, one might hypothesize that the more people affected by an impairment, the sooner it will merit regulatory attention; other factors held constant. Thus, persistently observed performance problems across multiple metro-level areas might warrant regulatory attention, while observations of brief impairments at few links or LAG are more likely operational problems of less regulatory interest. This assumption would suggest that if a network is required to report performance data to a regulator, it does so at the metro level, across all metro areas, with reporting of LAG-level data only to verify that the metro-level data is representative.

However, this level of data reporting (depending, of course, on the time granularity) will impose a burden both on the reporting operator and on the regulator that must analyze this data. This raises the question of whether data at a larger granularity would be adequate.

Regional aggregation and substitutability: We identified aggregation at a regional level as a level of reporting higher than metro. The relevant question is how to define a region, in order to be useful and representative for different purposes. For a region to be scoped in a useful way, the elements that make up the region should be *substitutable*—equivalently suited to purpose. But substitutability can be defined in several ways, again

related to technical factors, business considerations, and regulatory scope. Each presents trade-offs.

A technical justification for region size likely depends upon the nature of the typical applications in use. For highly latency-sensitive gaming content, a region of substitutable links would be smaller than for less latency-sensitive entertainment traffic, due to speed of light limitations on transmission speed. Returning to Figure 1, the additional latency between New York and Boston is small enough that for almost all applications, delivering traffic to a Boston customer over an interconnection in New York would not degrade the user experience, so interconnections in New York and Boston would be substitutable based on technical considerations.

However, this scoping might not be applicable in the context of business relationships. In our discussion earlier of the dispute between Netflix and access providers such as Comcast, Netflix proposed that it be allowed settlement-free interconnection if it delivered its content into the access network at an interconnection point close to the eventual consumer. An access provider might argue that the cost of carrying content from New York to Boston was material, and therefore the content provider had not complied with the “close enough” requirement if it just delivered traffic destined for Boston in New York. From this business perspective, the links in a region the size of New England would not be valid substitutes, even if consumers saw no degradation in the performance of the applications.

A regulator with a certain scope of authority (e.g., a state regulator in the United States) would not benefit from reporting at a regional size larger than that scope. The Massachusetts telecommunications regulator would not be able to act on data that reported on the overall New England region. It might require service providers to report at a state-level region. The question in this case is whether that level of aggregation masks important observations. But fully loaded LAGs in one metro area in a region where other metro areas have unloaded LAGS to the same party may not be of concern, depending on substitutability of the LAGs, and whether traffic is actually being routed over those underloaded LAGs. Unfortunately, regulators currently have no way to measure LAG substitutability or how traffic is actually being routed.

The risks of aggregation: In any aggregate that contains a number of LAGs (as would be the case with aggregation at the regional level), it is highly likely that there will be underutilized links. Consider, for example, a 10-GB link that becomes slightly congested at peak traffic times. In

practice, capacity is usually added in units of 10 GB. That is, to upgrade the capacity of a 10-GB link an operator would make a LAG by adding a second 10-GB link, with a resulting total capacity of 20 GB. If the original 10-GB link was only slightly congested, this new LAG with twice the capacity will be hardly more than 50% loaded at peak. Averaging such a link into a mix of links (or LAGs) some of which are congested may result in an overall value that does not suggest congestion.

In fact, it is not obvious how to aggregate LAGS of different capacity to provide the most representative metric for the resulting region. How does one describe the overall state of a region with a 100-GB link that is not congested and three 10-GB links that are substantially congested? One answer is that if the links are substitutable, then the excess traffic on the 10-GB links could be routed over the 100-GB links, so the overall region need not have any congestion. In order to validate this assertion, it is not enough to ask if the links are substitutable based on technical parameters, but as well whether the business agreements describing those links would actually allow the excess traffic to be redirected over those links.

When receiving third-party data in support of some argument, the regulator must carefully consider whether its level of presented aggregation (over time and in space) is potentially masking material events. One sort of dispute that might arise before a regulator would be one party to an interconnection providing measurements of congested LAGs as evidence of a failure of the other party to address the business issue, and the other party providing measurement at a regional aggregation to argue that there are uncongested paths available to carry the excess traffic. When a regulator finds itself in this sort of situation, the regulator should require not only performance data on the various LAGs in the region, but disclosure of the relevant business agreements, to confirm that the links in the region are actually valid substitutes.

Provider wide-data: This level of aggregation considers all direct links between a provider and a single other partner network. Level 3 posted the plots in Figure 3 in a 2014 essay that also claimed that similar patterns on most interconnect points between two ISPs had persisted for over a year.²⁵ In the context of our model, Level 3 aggregated this measurement over a single LAG, but asserted that all other LAGs to the same interconnection partner were behaving similarly. A regulator would probably want to confirm that claim with data from other individual LAGs. Impairments might manifest as persistent problems across all paths connecting two networks

25. Taylor, "Observations of an Internet Middleman."

or as patterns of path-specific periodic impairment. While the aggregate of “all direct interconnection LAGs” is easy to define, it can be hard to measure for parties other than the two directly connecting networks, due to the challenge of discovering all LAGs in the set. In “Applying the Model to Specific Measurement Data Sets” we turn to some concrete case studies that hopefully demonstrate how our model can help guide appropriate conclusions.

Measurements of Path Segments

Our model provides a structured way to describe and interpret interconnection measurements. To this point, we have considered measures of an interconnection link or aggregate. However, sources of impairment between networks may not manifest as congestion on the interconnection link itself, but at some other bottleneck in the path, as we described in the discussion of routing policy internal to an access network in “Approaches to Interconnection Discrimination and Policy Responses.” Measurement across a longer path segment, up to and including end-to-end measurements from sender to receiver, may be essential for assessing performance between two networks.

Measurement of a partial path segment (i.e., including several links) may be helpful in localizing impairment or excluding certain segments as a source of impairment, but it may not be possible to position test probes so as to measure the right partial path segments. More commonly reported are end-to-end measurements that record overall performance of a transfer, such as achieved throughput. How these reported results aggregate individual end-to-end measurements is crucial for making sound inferences from them.

The edge content provider is well-positioned to record end-to-end performance measures such as throughput of an individual flow transfer, but has flexibility in how to aggregate these measures for reporting purposes, with respect to both the source and the destination. A content provider would not normally report transfer speeds to a single destination, but to a set of destinations that correspond to one of our aggregates: a metro area, a region, or overall for an access provider. They can also choose to report all transfers *from* a single content source (not typical), or for all sources in a region, or from all sources belonging to the content provider. An aggregated report of end-to-end throughput from any content source into a metro or region of an access ISP may or may not include both direct and indirect paths. The mix of traffic flowing on the direct and indirect paths may vary over time, likely a function of how much capacity is available

on direct paths. Direct and indirect paths may have different performance characteristics, and aggregating them together may give an overall result that is not indicative of either class of path. Using our previous terminology, aggregated end-to-end measures may lump together paths that are not *substitutable*, and thus obscure important information about specific paths.

If the path segment is the entire end-to-end path, then another challenge for interpretation is that it may reflect impairments in the user home or problems in the edge provider network or services, both of which vary over time, rather than anything related to the interconnection.

Finally, these measurements of actual user performance can only measure actively used paths. Access providers have argued that an analysis should consider both paths that edge providers are using to send traffic as well as alternative (potential) paths, since edge providers may choose not to avail themselves of all available paths. On the other hand, some unused alternative paths may not be practical for business or topology reasons, which third parties cannot determine on their own. Under these circumstances, regulators are justifiably skeptical of expansive claims that all interconnection links of an access provider represent an appropriate scope of analysis. Our advice would be that any aggregated reporting of end-to-end performance should include explicit descriptions of the level of aggregation on both the source (typically content sources) and the destination (which might be a metro or regional aggregation). The description should make clear the set of interconnection paths that the aggregated traffic used. Without this level of detail, it is too easy to draw unwarranted inferences from the results.

An alternative to end-to-end measurement by the content provider is measurement from the client end that attempts to estimate end-to-end performance. As an example, in late 2013, the FCC began to expand the range of measurements in their Measuring Broadband America program, to better understand the implications of end-to-end measurement. They were exploring how to test the performance of video services like YouTube and Netflix.²⁶ This effort expanded from a preliminary pilot phase in 2014 to a wider roll out throughout 2015, but is still remarkably tentative.²⁷ The FCC notes that: “the video streaming tests developed by SamKnows and the FCC in collaboration with content providers like Netflix, YouTube and Hulu are not intended to compare the performance of the carriers, but

26. Miller, “Collaborative Meeting Report.”

27. Samknows, 2014 “Collaborative Meeting Presentation”; Samknows, 2015 “Collaborative Meeting Presentation.”

rather to develop a methodology study.”²⁸ To date, the FCC has released no results from this study, in part due to concerns from stakeholders about the accuracy of the methodology and results.²⁹ In July 2016, the FCC announced a new CDN test to measure the download throughput of small objects hosted on the following CDNs: Apple, Akamai, Microsoft, Google, Cloudflare, and Amazon. They have not released any results from this test either, due to similar concerns. The slow pace and lack of resolution in this project is an indication of both the difficulty of sound measurement and its contentious context. It is not easy for the research community to participate actively in the design of these experiments, nor participate in the discussion about the interpretation of the data, due to concerns about control over disclosure of early results. Researchers also do not have access to any suitable experimental platform to carry out similar experiments at a suitable scale, nor any assurance that content providers would cooperate in validating measurement methods.

Applying the Model to Specific Measurement Data Sets

We consider six case studies that span data sets offered by seven US broadband access providers, edge providers Google and Netflix, academic researchers, and the IME for the AT&T/DirecTV merger.³⁰ Each example covers large content and service providers interconnecting with large access

28. Razdan, “Collaborative Meeting Report.”

29. Quoting from the cable lobbying organization’s letter to the FCC: “With regard to the Netflix streaming tests in particular, the ISP Representatives questioned whether the proposed testing would accurately measure the performance that a consumer actually experiences in streaming a Netflix video. We expressed our concern that the testing of Netflix streaming currently under way uses synthetic 25MB binary files instead of actual video files that are delivered to Netflix customers. The ISP Representatives stressed that the testing should replicate the real-life consumer experience of streaming a video, and that therefore the testing should randomly access actual video files from the same servers that deliver videos to Netflix customers. The ISP Representatives also stressed the importance of requiring that participating streaming services sign a Code of Conduct to ensure that there is no gaming of the testing process, similar to the Code of Conduct that all of the participants in the fixed-line MBA Collaborative signed.” National Cable and Telecommunications Association, “Letter to the FCC Re: Broadband Performance Measurement.”

30. Center for Information Technology Policy, “Interconnection Measurement Project”; Google, “Google Video Quality Report”; Netflix, “Fast.com”; Clark et al., “Policy Implications of Third-Party Measurement of Interdomain Congestion on the Internet”; claffy et al., “First Amended Report of AT&T Independent Measurement Expert: Reporting requirements and measurement methods.”

providers in many different locations. We chose these examples to show how our model and approach described previously can facilitate understanding of what is being measured, and how to soundly interpret it. The first three projects focus on end-to-end path measurement data: Google's Video Quality Report; Netflix's ISP Speed Index; and Measurement Lab's (M-lab) interconnection study. The second three projects focus on specific interconnection link measurements: the CAIDA/MIT interconnection measurement project; Princeton's Center for Information Technology Policy (CITP) interconnection measurement project; and the measurements proposed by the IME for the AT&T/DirecTV merger conditions.

Google Video Quality Reports

Google's Video Quality Report includes data derived from end-to-end measures, aggregated per ISP per city, of throughput of YouTube streams.³¹ This data set includes for each YouTube video request: timestamp, access network, estimated geographical region (e.g., country, metro), total bytes transferred to the client, time at which receiver acknowledges receipt of all bytes.³² The reported data shows demand as a function of time of day, averaged over 30 days, so it is possible to speculate from the data about performance at peak times versus off-peak times.

This chosen aggregation granularity for reporting removes visibility into what paths the flows are using to reach their destination. Internally Google knows in detail how they route traffic to the users for any given measurement, for example, whether it crossed a direct interconnection link with the broadband access provider or via an indirect path through a third party. Unlike Google's *peering.google.com* peering portal, which provides to an interconnected party capacity and traffic volume statistics on direct interconnection links to that party, our understanding from discussion with Google is that these Video Quality Reports aggregate measures over both direct and indirect paths. Exactly how they aggregate these measures is not clear. Figure 4 illustrates a video streaming quality report for AT&T Digital Subscriber Line (DSL) service in State College, PA, reporting the fractions at different times Google is sending the video at high definition and standard definition.³³ A reduced fraction of high-definition flows,

31. Google, "Google Video Quality Report."

32. Google, "Google Video Quality Report—Methodology."

33. We note that we had to look at many cases to find plots worthy of discussion. Almost all the plots we examined showed no degradation at peak times, suggesting that much of the infrastructure they measure has adequate capacity.

particularly during peak hours, and when other comparable providers and comparable regions do not experience a drop, is potential evidence of an impairment. However, what this picture implies is not exactly clear. DSL is a slow-speed service, so individual users might not have the capacity to watch video in high definition. But the limited capacity of a DSL link does not change as total load on the system rises at peak time. The increased fraction of video at standard definition during the morning hours could be due to a number of reasons. Figure 5 shows another daily distribution of high and standard definition, in this case for Verizon in Cambridge, MA. This plot shows a similar pattern in the morning hours as the plot from State College, but in this case it is clear that the problem is not a lack of overall capacity, since in this plot there is an afternoon peak that is substantially higher than the morning peak, during which almost all traffic is in high definition. Perhaps these systems have users with different usage patterns in the morning and afternoon. There is no way to tell—pictures at this level of aggregation often raise more questions than answers.




FIGURE 4 Google Video Quality Reports Summarize Video Performance at the Metro or State Level, in this Case for ATT DSL Service in State College, PA (Sept 17, 2019).




FIGURE 5 Google Video Quality Report for Verizon Service in Cambridge, MA (Sept 17, 2019). Google classified this service as FIOS, but FIOS is not offered in Cambridge, so this data is either geolocated to the wrong city, or more likely is Verizon DSL service.

These published statistics allow one to examine data at a metro, regional (which Google defines by state in the United States), and entire provider level of aggregation in the United States. For instance, one can look at the performance of Comcast in a given city (such as Boston), across an entire state (Massachusetts), or across the whole United States. The web interface facilitates comparisons between broadband providers at these three granularities, which provide evidence that Google is capable of delivering high definition video content to users in some networks even if impairments are evident in others.

Another limitation of this data, and all end-to-end performance measurement, is its inability to localize the source of any impairment to the access provider itself. The measurements also capture, but do not distinguish, impairments in the home network, which may manifest more often as access speeds increase.³⁴ On the other hand, large and geographically distributed end-to-end measures of similar users across multiple ISPs mitigate this concern. Examining diurnal patterns can also establish some baseline of impairments, although some portion of the variation detected during peak hours are also likely due to changes in user behavior, for example a larger mix of devices accessing content.

In summary, this data set should capture impairments at the metro, regional (state-based) or provider-wide level of aggregation, but may also imply that a broadband provider is responsible for an impairment that is not under its control. For the regulator, data of this sort may be a suggestive starting point, depending on what question is being asked, so long as the provider is clear about exactly how the aggregation has been done—over what sources, whether it includes indirect paths, and so on.

Netflix ISP Speed Index

Netflix reports per-ISP values of their *speed index*, an aggregate summary of end-to-end download speed across all paths from Netflix servers to a given access provider's customers (Figure 6).³⁵ They do not report any time-of-day data, as Google does, but just longer-term trends. This plot shows significant improvement of this metric for many ISPs in late summer 2014, which appears to be the result of resolution of noticeable earlier performance degradation to customers of those ISPs. After that time, the

34. Sundaresan, Feamster, and Teixeira, "Home Network or Access Link?"

35. Netflix, "Netflix Speed Index."

plots show steady improvement in average speeds, with two clusters of performance: the higher cluster being cable and fiber providers, and the lower cluster DSL providers. However, it is not clear what inference to draw from the steady increase. The performance of different ISPs over time seems to vary in lock-step, which suggests that the variation month to month is not due to changes within the ISPs, but some other cause—perhaps a change in the mix of videos that are downloaded by the Netflix customers.

In our model, this data aggregation includes content served directly from Netflix caches in access provider networks, direct interconnections, as well as indirect paths. Why is Netflix aggregating statistics at this overall provider-wide granularity, rather than metro or regional levels of aggregation as Google does? One plausible explanation is that Netflix is balancing the desire to establish a public record of access provider performance with concerns that they do not want to discourage adoption of their services by suggesting certain regions have performance problems. This public provider-wide aggregate prevents inferences about regional interconnection issues. Also, the end-to-end measures here suffer the same risks of the Google's Video Quality Report, in terms of potentially masking issues happening other places in the network. For the regulator, data at this high level of aggregation may not serve to answer many questions. Since the approach to aggregation is different from that of (for example) Google, there is no effective way to compare Netflix data and Google data to see if the two content providers are experiencing similar treatment. In general,

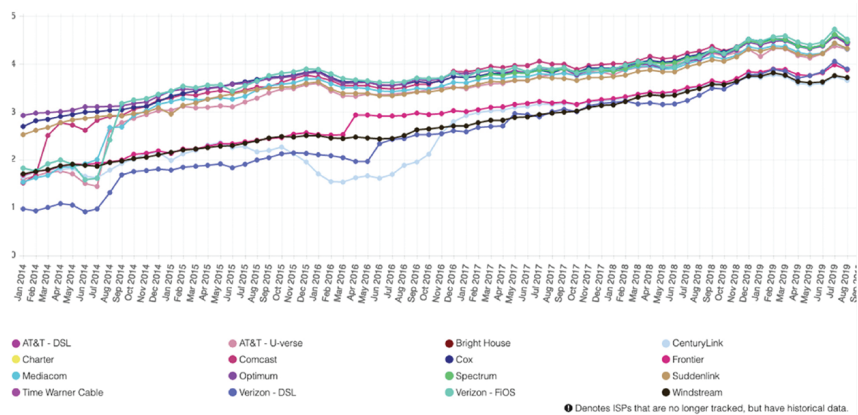


FIGURE 6 Netflix ISP Speed Index is an Aggregate Summary of Performance Across All Possible Paths by Which Netflix Serves traffic to a Given Access Provider's Customers. Both direct and indirect paths across all regions and metro areas.

without an agreement across firms to report data in a similar way, it will be impractical for a regulator to compare how different interconnecting parties are being treated, based on data provided by those firms.

Measurement Lab

Google's Measurement Lab (M-lab) operates a set of servers against which clients can measure throughput, traffic shaping, and traffic differentiation. Network Diagnostic Tool (NDT) is a popular test hosted on the M-lab infrastructure that clients use to measure achievable throughput of their Internet connection. When a client initiates the test, the M-lab backend directs the client to an available server geographically near the client. The client and the server then conduct throughput tests in both directions. A report from M-lab used this end-to-end measurement data, that is, the achieved throughput in NDT tests from a server hosted in a certain AS (say S) to clients in an access AS (say A), to infer performance degradations on the path from S to A.³⁶ The methodology looked for significant differences between peak and off-peak throughput for a server–client AS pair to infer peak-hour congestion on paths from that server to the client AS. The report went a step further, attributing observed performance degradation between S and A to the interconnection between S and A in the region of the server S, for example, the interconnection between Cogent and Comcast in Los Angeles.

We can use our model to critically consider three assumptions underlying this inference method. First, similar to the previous two case studies, this method assumes that congestion is more likely to exist at the interconnection points between networks rather than within a network. Second, also an unspoken assumption of the previous two case studies is that the server and client AS directly connect, thus any observed interconnection congestion exists on that direct link. It is possible an individual test could execute over an indirect path, in which case reported measurements reflect a combination of *direct* and *indirect* paths between server and client ASes. Third, the method assumes that M-lab's server selection algorithm works well enough that clients in a certain metro region are directed to servers in that same region, and the NDT test thus reflects performance from S to A in that metro region. The M-lab report did not examine performance at a finer or coarser granularity than metro region, that is, at the LAG level.

36. M-Lab Research Team, "ISP Interconnection and its Impact on Consumer Internet Performance."

M-lab publishes all NDT test data, along with path data in the form of traceroutes from servers to clients. This data could enable performance evaluation of specific LAGs, although we have found that due to the way the NDT tests sample individual LAGs, only a small set of LAGs admit statistically significant inferences. One could also aggregate the data to a provider-level view by aggregating all tests from servers in an AS *S* to clients in an AS *A*. One concern with such aggregation is that some metro regions may be under- or over-represented based on deployment of server-side infrastructure.

CAIDA/MIT Study

The CAIDA/MIT project developed methods to detect interconnection links and evidence of persistent congestion on these links, for example, recurring patterns of increased latency.³⁷ The interconnection discovery phase uses vantage points (VPs) inside a network to perform an extensive active topology discovery process that infers all interconnections of that network visible from each VP. By widely distributing many active probes in an access network, they assert that they can find essentially all points of interconnection with other connecting parties. They also developed the time-series latency probing (TSLP) method, a method to infer congestion at these discovered interdomain interconnections. The VPs send probes toward the near and far end of each discovered interdomain link to obtain two time series of latencies. The presence of diurnal patterns in latency to the far end of the link but not to the near end of the link signals evidence of congestion at the interconnection. The method generates raw measurements on each interconnection LAG discovered from the VP, and they examine the resulting latency data (using heuristic geolocation) at the metro, region, or provider level. The TSLP method does not measure the utilization or capacity of LAGs; it uses active measurement to reveal only whether individual LAGs show latency-based evidence of congestion.

This project makes the data available on a per-link and per-LAG basis, so a researcher or a regulator could in principle combine data at various levels of aggregation, depending on what question was being considered. One could use the data at different aggregations to infer that “all LAGs connecting providers *A* and *B* in a certain region appear congested” or that “4 out of 5 LAGs connecting providers *A* and *B* in a certain region appear

37. Dhamdhere et al., “Inferring persistent interdomain congestion”; David Clark et al., “Policy Implications of Third-Party Measurement of Interdomain.”

congested.” In principle, one could aggregate data at the LAG level in any way that addresses a specific question. The main challenge in aggregation based on geography is accurately geolocating interconnection links to specific cities or metro regions, which is known to be less accurate for core router infrastructure than edge devices.³⁸ A further challenge is that a VP inside an access network may not observe interconnection LAGs in geographically distant regions, depending on the nature of routing between the access network and its interconnection partner. A full provider-level view from an access provider to an interconnection partner may require a dense deployment of VPs inside the access network.³⁹

An ongoing project at CAIDA is focused on making this and other data related to Internet topology and performance easier to use.⁴⁰

Princeton’s CITP Interconnection Measurement Project

Seven broadband access providers, which service over half of US broadband subscribers, cooperated to provide Princeton’s CITP with aggregated utilization data for interconnection links of broadband providers.⁴¹ It contains data on nearly all of the paid peering, settlement-free peering, and ISP-paid transit links of Bright House Networks, Comcast, Cox, Mediacom, Midco, Suddenlink, and Time Warner Cable. Each broadband ISP submitted the following data for each five-minute interval (See Figure 7): timestamp; region (which maps to *metro* in our model) representing an aggregated link group; anonymized interconnecting party; total ingress bytes; total egress bytes; and capacity of the aggregated link group. We try to map these terms to our model of interconnection, but note that although the data *collection* granularity maps to terms in our model (if we consider each metro area its own region), the *reporting* granularity does not match any aggregation granularity in our model.⁴²

First, Feamster states that “to protect the confidentiality of information pertaining to usage on specific interconnects, the data is aggregated into

38. Huffaker, Fomenkov, and Claffy, “Geocompare: a comparison of public and commercial.”

39. That project found 45 distinct interconnection router-level links from a major broadband access provider into Level 3, the largest of the Tier 1 providers. Because different points are announced in different regions of that access network, it took 19 observation points to find all of those interconnection points.

40. See <http://www.caida.org/funding/dibbs-panda/>.

41. Center for Information Technology Policy, “Interconnection Measurement Project.”

42. Feamster, “Revealing Utilization at Internet Interconnection Points.”




FIGURE 7 Center for Information Technology Policy's Interconnection Measurement Project Architecture. (Borrowed from Feamster, p. 4.)

a single link group per geographic region.”⁴³ In other words, the broadband providers are aggregating capacity and utilization data to a metro level before sharing it with CITP. In defense of this aggregation, Feamster notes that “we can assume a relatively uniform load balance of inbound traffic flows for a link aggregation group.”⁴⁴ But content providers do not necessarily balance load equally across multiple LAGs in a metro area. From the data CITP shares publicly, it is unclear whether many of the LAGs in the CITP data set have multiple constituent LAGs. If they do, this assumption is not valid.

More problematic is that although each broadband provider anonymized the partner network for each data point they shared, they also required Non-Disclosure Agreements that further limited the analysis and reporting to even more heavily aggregated forms of this data. In particular, the participating ISPs required aggregation of capacity and utilization levels either *across all regions* or *across multiple broadband networks* to anonymize any given interconnecting party. Feamster acknowledges that such aggregation does not enable understanding of impairments between any pair of networks:

The paper states:

In the public dataset, it is possible to assess the overall utilization in some region across all ISPs and partner networks, but not for any individual interconnection point in a region. Similarly, it is possible to see the aggregate utilization for any of the participating ISPs, but not for a specific region or neighbor ISP. As a result, the aggregates make it difficult to drill down into the utilization between any pair of networks, either as a whole or for any particular region. As a result, it is not possible to conclude that

43. Ibid, p. 4.

44. Ibid, p. 5.

*no interconnection links experience high utilization. Because the public data shows utilization across each ISP, we can conclude that each ISP has spare capacity—although we cannot conclude that it has spare capacity in each region or on any individual port.*⁴⁵

We use the terminology of our model to explain how this excerpt is unnecessarily misleading. First, the intentionally chosen aggregates do not just make it “difficult” to drill down into issues between any pair of networks—they make it impossible; as Feamster has previously clarified, obfuscating that information was a requirement of the NDAs between the access and content providers. Furthermore, the conclusion that each ISP has spare capacity refers to that ISP’s connections with *all* of its interconnecting parties, which is not a meaningful aggregation granularity, since, for example, Netflix cannot use Google’s links. Indeed, we note that the most interesting observation in the public data that CITP posted is that a nonnegligible number of interconnection links experience nearly 100% utilization!⁴⁶

Feamster proceeds to aggregate the utilization of interconnect groups for *all* interconnections from *all access ISPs* to *all interconnecting parties* across *all* regions, as if they were substitutable, to conclude that even if there are heavily congested individual LAGs (which he cannot disclose), there is some LAG somewhere in the highly aggregated data set that has spare capacity that a partner could be using but is choosing not to use. He draws similar conclusions from aggregating and presenting utilization statistics *per region* but *across all ISPs*. The problem with such broad aggregations, across all regions *or* across all interconnecting parties in a region, is that the reported capacity could include links that are unlikely to be available to all interconnecting content providers. (Again, Netflix links are probably not a substitute for delivering Google content.) We consider these aggregation granularities to hide exactly the information a regulator needs to see to assess performance problems with interconnection between parties. He again acknowledges the limits of these aggregations:

Certain answers remain obscured, such as whether a particular partner network is experiencing persistent congestion, or whether particular types of connections (e.g., paid peering) are experiencing more or less congestion.

45. Ibid, p.7.

46. This page provides an anonymized view of each aggregated LAG: <http://interconnection.citp.princeton.edu/project/viewsby-interconnect/>; CITP, Interconnection Measurement Project.

but concludes that the reported granularities:

*reveal a general picture of (1) all ISPs having spare capacity in aggregate across interconnects; (2) most interconnect capacity in aggregate showing spare capacity at peak. Both of these conclusions reveal significantly more than we have known to date.*⁴⁷

Unfortunately, we do not believe these measurements reveal anything new that can address any disagreements related to Internet interconnection that have appeared in the research literature, media coverage, or FCC filings.

Measuring the Interconnection Links of AT&T

During the AT&T/DirecTV merger proceedings, publicly filed objections from edge providers and their representatives focused on interconnection as a locus of harmful discrimination and strongly advised the FCC impose conditions related to interconnections on any approved merger. In response to these concerns, the merger order imposed a requirement for AT&T to provide to the FCC the business terms for all their significant interconnecting parties and to report key performance metrics of those interconnections, for a duration of four years beyond the merger approval date. This information will educate the FCC about the character of interconnection and whether the contracts suggested unreasonable discrimination among AT&T's interconnecting parties. Since objectors focused only on interconnection as a point of possible discriminatory behavior, the merger order limited the scope of measurement to interconnection links, and no other part of the path from senders to receivers.

To define exactly how AT&T would gather and report these measurements, the FCC merger agreement called for the joint appointment of an IME. AT&T and the FCC selected CAIDA, at the University of California San Diego, to serve as this IME. The full report of their methodology and the supporting filing with additional justification for some of methodology are available as FCC filings and on CAIDA's web site.⁴⁸ The FCC specified a

47. Feamster, "Revealing Utilization," p. 9.

48. claffy et al., "First Amended Report of AT&T"; claffy et al., "Report of AT&T Independent Measurement Expert: Background and supporting."

range of required measurements; here we examine how those measurements fit within our conceptual model.⁴⁹

Capacity and utilization. Consistent with the reasoning in “Different Objectives for Performance Analysis,” the IME required that AT&T report on LAGs, and as well metro- and provider-wide aggregations, but not on individual links. To account for factors that might make this data ambiguous, including possible use of the link to carry non-Internet traffic, for example, carrier VoIP, IPTV transport, or other such specialized (or “non-BIAS”) traffic. The IME required that AT&T disclose to the FCC if any such sharing is taking place. The IME emphasized that it might be necessary to seek the cooperation of the interconnecting party to fully characterize how the interconnecting party is managing different sorts of traffic.⁵⁰

Packet Loss rate. Since utilization data alone cannot confirm the presence of congestion (“Relating Performance Measures to QoE”), the FCC

49. In full, the merger agreement placed the following requirements on the IME (Federal Communications Commission, “In the Matter of Applications of AT&T Inc. and DIRECTV For Consent,” Appendix B.V.2.c.iv), and required the specification of the following measurements:

...the Company, in consultation with the Independent Measurement Expert, will submit for approval by the Commission’s Office of General Counsel, in consultation with the Wireline Competition Bureau and the Chief Technologist, a report describing the Independent Measurement Expert’s proposed methodology for the measurement of the performance metrics described herein. Such report shall also be submitted to the Independent Compliance Officer. The proposed methodology should, at a minimum, address the following criteria:

- (1) *Identification of Internet Interconnection Points, including the identity of the interconnecting parties and the location and capacity of each interconnection point.*
- (2) *Identification of a disclosure exemption threshold for a de minimis volume of traffic exchanged between the Company and interconnecting parties.*
- (3) *A definition of “Latency,” which shall include the disclosure of the probability distribution.*
- (4) *A definition of “Packet Loss”.*
- (5) *Time of measurements, which shall, at a minimum, include an identified window within peak usage periods.*
- (6) *For any performance metric contingent upon an interconnecting party’s participation in the selected measurement methodology, a process for waiving the disclosure of that metric at points of interconnection where the interconnecting party declines to participate.*
- (7) *Frequency and duration of measurements.*
- (8) *Any devices used for measurement.*
- (9) *End points of measurements.*
- (10) *Placement of any devices.*
- (11) *Frequency of disclosures.*

50. claffy et al., “Report of AT&T Independent Measurement Expert: Background and Supporting.”

required the reporting of packet losses and variation in latency (jitter), presumptively an indication of packet queues due to congestion. The IME recommended several way of measuring packet loss. The first is to use the network management capability on the routers at the endpoints of the LAG. Routers track how many packets they drop at the ingress to a link.⁵¹ But the FCC's regulatory concern was primarily with traffic flowing *into* AT&T, which means that the packet loss of interest occurs on the other end of the link, on the router belonging to the interconnecting party. In the context of these merger conditions, the FCC could compel AT&T to report, but the IME could only *ask* that the interconnecting party provides this information, and those parties have legitimate reasons for reluctance to cooperate with this request. Most obviously, if the link is congested in the incoming direction, the interconnecting party may not want to reveal it. Additionally, the router may drop packets that are malformed, classified as malicious (e.g., part of a DDoS attack), mis-routed, and so on. A dropped packet does not necessarily mean congestion, and if the link is only occasionally congested, other causes of drops may distort the overall measure. Given these factors, the IME required that AT&T use a second approach to measure losses (and latency, as discussed below) on the interconnecting links, which is to send a probe packet across the link with the goal of triggering a response packet. The final methodology required using both approaches to the extent possible, which provides a means to determine whether the two yield similar answers, and (if not) the circumstances under which the two differ, for example, Barford and Sommers.⁵²

Latency and jitter. In addition to loss rates, the FCC required (that the IME specify a method for) AT&T to report on variation of latency across the LAG. While some commercial routers support a measurement of queue length, operators do not generally use it, and the IME had no way to calibrate it. So the final methodology recommended two approaches for active measurement, both of which depend on measuring the time between the sending of the probe and the return of the reply, which estimates the round-trip latency. Variation in this measurement is evidence of jitter. Limitations of the canonical probing protocol (ICMP) motivated the requirement for another more accurate protocol (Two-Way Active Measurement Protocol [TWAMP]) for interconnecting parties willing

51. These loss counters may be incomplete, that is, reporting only drops the router properly records.

52. Barford and Sommers, "Comparing Probe- and Router-Based Packet-Loss Measurement."

to cooperate to support it.⁵³ TWAMP tries to remove the uncertainty associated with a simple ICMP probe using a more complex set of timings. While TWAMP may be more accurate, it is not widely deployed in the Internet, and the operational and research community has little experience with validating the accuracy of TWAMP probes. The interconnecting party would have to install a TWAMP responder at the far end of the link, which there is no way to mandate in the context of the merger order. All the IME (and AT&T) could *ask* the interconnecting party is if they will install and activate a TWAMP responder. The fallback is the less accurate ICMP probing.

Applying our model for aggregating the AT&T measurements.

As measurement at different aggregations reveals different behavior, the IME required AT&T to report utilization at three aggregations: individual LAGs, metro, and overall for each interconnecting party. The IME required reporting on individual LAGs so that the FCC can detect an imbalance among LAGs in an aggregation, which might have implications for substitutability of those LAGs. The requirement for reporting at three levels of aggregation was first to allow for the possibility that measurements at one level may suggest a problem with data aggregated at another level. But more fundamentally, the goal is to learn what these levels of aggregation reveal, given that they are based on the same underlying LAG-level data. Defining a region raised enough ambiguities that it was omitted from the requirement. Aiming for a fine-grained view of the variation of demand, the methodology required that AT&T gather raw data in five-minute intervals and to plot key parameters (across all three aggregations) as well as summary statistics. The IME did not require that AT&T report either loss or jitter at higher levels of aggregation, since they knew of no way to aggregate these values in a way that yields meaningful results. This challenge is a future research question.

Summary and Policy Recommendations

We have used our conceptual model to position a number of current data gathering efforts and to frame the justification for the measurements the IME required AT&T to report to the FCC. Our goal is not to recommend a specific set of measurements, but to assist policymakers in making

53. claffy et al., "First Amended Report of AT&T Independent Measurement Expert"; Hedayat et al., A Two-Way Active Measurement Protocol (TWAMP): RFC 5357."

informed decisions regarding how to interpret measurements intended to reveal significant performance impairments. Here we summarize lessons learned in the process of defining the AT&T/DirecTV methodology and in the process of defining the model we offer in this article.

Measurement is political, and often adversarial. If parties are in a dispute, they may favor measurement that reflects well on them and poorly on another disputant. Measurements that reveal the location of an impairment may be in the best interest of one party but not the other. ISPs are only likely to share data that will reflect well on them.

Measuring individual interconnection points does not tell a complete story. Content providers have control over how they source content into an access provider network such as AT&T (including potentially over indirect links), and so a coherent view requires examination of the larger aggregates in our model, possibly including traffic flow over indirect paths, to see if available capacity can meet demand. A challenge of obtaining a more expansive view is that a third party may not be able to ascertain which transit links might be available to content providers as indirect paths, taking into account both technical and business considerations. Unanonymized network-level traffic flow data is required to confirm that a transit connection really is being used by a given content provider.

Reporting for a given pair of interconnecting parties at two aggregation granularities—metro and region—are generally useful in determining whether material impairments are occurring that might warrant regulatory attention. But a caveat applies. Specifically, if data from finer aggregation granularities does not lend credibility to the claim that the contained elements are substitutable, looking *only* at an aggregate may mask exactly what one is trying to use the data to ascertain. Particularly with respect to content providers, which may connect to an access provider with many LAGs in a metro area (and connect in many metro areas), the metro level of aggregation is useful, assuming substitutability of component links. The largest levels of aggregation (such as all links between two providers) are rarely a useful scope of regulatory analysis. It is unlikely that links in such a large scope will be substitutable, due to variation in the lower level elements that compose the aggregate.

Path measurements provide a useful and complementary view of performance across multiple ISPs. But caveats apply here too. Absence of impairments in a single interconnection link does not imply absence of impairments in the end-to-end path. Obtaining a reasonably complete picture will require consideration of both link and path (or path segment) measurements. The end-to-end measurements from Netflix and Google

(“Google Video Quality Reports” and “Netflix ISP Speed Index”) provide another measure of whether overall capacity is meeting demand, and whether there is actual variation in throughput. However, the specific decisions about how those data were aggregated make the conclusions ambiguous. End-to-end measurement will reveal behavior that most directly maps to possible impairment, but it does not reveal where impairment is arising, nor indicate if there is any actual degradation in user-perceived QoE. Mapping such data to QoE impairment requires agreement on application-specific thresholds for impairment, an open research question.

An accurate picture of packet loss and latency across interconnection links requires cooperation of interconnecting parties with counterincentives to cooperate. Content providers, in particular, navigate a complex set of issues surrounding interconnection, starting with commercial contracts. Using the AT&T case as an example, AT&T described the content providers as customers of AT&T, not traditional peers. Further, the FCC compelled AT&T to reveal under a protective order the contracts the interconnecting parties signed with AT&T. If a direct interconnection link is apparently congested, it does not necessarily mean AT&T is preventing the interconnecting party from obtaining wanted capacity. An alternative explanation is that the interconnecting party has, for economic reasons, chosen not to purchase additional capacity. Content providers may have many paths to deliver their content, but if their interconnection links become congested, it might not reflect well on them, especially if they had recently lodged complaints about their ability to interconnect, and the FCC now had the terms of their interconnection agreement in hand. These vested interests may motivate a preference for reporting larger aggregates (metro or provider-wide, including indirect links) and reporting measures that are not scoped to interconnection links specifically, but to longer path segments, or end to end.

Each stakeholder brings a unique contribution to the overall picture of interconnection conditions. The research community has not devised third-party measurement tools to remotely measure the capacity or utilization of an interconnection link. But neither do the ISPs have full visibility into behavior. An ISP observing a link may measure utilization and congestion but it cannot easily measure how the senders have chosen to deal with this congestion: how much they have slowed their sending, or whether they have changed the encoding of the traffic. Third-party observers (or providers of higher-level applications and content) may be better positioned than operators to measure the overall QoE of an activity by end-to-end measurement, but they are not well positioned

to assess substitutability of LAGs. An ISP can only see part of the path and may have excellent visibility into that part, but cannot see all of it. An end-to-end measurement has imperfect visibility into the path but can detect if there is an overall problem.

The FCC should pursue wider visibility of what is learned in the face of protective orders. It is important for the larger community to learn about the utility and effectiveness of these various measurement methods. However the measurement methodology defined by the IME is public, the gathered data is not. The merger order requires that AT&T share data with the FCC under a protective order, and with the IME itself only in the beginning of the reporting process, for the purpose of resolving flaws in the measurement methods.⁵⁴ Some operational issues that triggered the reporting requirement, such as overloaded LAGs, may not arise during the limited period in which the IME is reviewing the measurement method and resulting reporting. Validation of a method to detect an issue is not possible until and unless the condition of interest arises.

Furthermore, while the data shared with the FCC will inform regulatory policies and the FCC's understanding of access network interconnection

54. From the contract [<http://www.caida.org/funding/att-interconnection/>]:

(c) CAIDA and AT&T jointly will review the first report that AT&T must submit to the FCC on Internet interconnection performance metrics resulting from the Methodology (the "Metrics Report").

- (i) AT&T acknowledges and understands that in order for CAIDA to fulfill its obligations as part of this review process, including to assert confidence as to the validity of the Methodology, CAIDA must have reasonable access to certain underlying data necessary to validate the Methodology. Any method for measurement must be tested and evaluated in practice, and CAIDA must be materially involved in this activity.
- (ii) If either CAIDA or AT&T believes that there is an issue with the performance metrics contained in the first Metrics Report, CAIDA will (A) propose reasonable adjustments to the Methodology to resolve the issue(s) that are satisfactory to the FCC, and (B) consult with AT&T on AT&T's explanation of the issue(s) and the proposed adjustments to the Methodology to the FCC and the ICO. CAIDA and AT&T will repeat this process until there has been a Metrics Report that CAIDA believes contains appropriate performance metrics.

If, after CAIDA's repeated, good faith attempts to propose reasonable adjustments to the Methodology to resolve issue(s), the FCC fails to approve the proposed Methodology, CAIDA reserves the right to terminate this Agreement upon prior written notice to AT&T and the FCC and indicating the reasons therewith. If such termination occurs: (1) This agreement shall terminate and CAIDA shall no longer have the right to access and use the Protected Information (as described in 5(b)), (2) CAIDA is relieved of any obligation or penalty under this Agreement; and (3) AT&T shall reimburse CAIDA for all services performed and reimbursable expenses incurred up to that point of termination.

links, some sanitized public view of this data might dispel many concerns about the health of the interconnection ecosystem. The merger agreement contained a high-level acknowledgment of the goal of public release of the data in some form.⁵⁵ Because of the structure of the merger agreement, only the FCC is in a position to make this happen. We urge that this effort be undertaken.

One potentially productive role for the academic research community is as a partner to the regulator, which is only possible if the regulator can facilitate access to data for validation and deployment of measurement infrastructure in support of policy needs. Sustained funding and support for such infrastructure is an on-going challenge, and no country (to our knowledge) has stable sources of government funding specifically for measurement infrastructure to enable Internet research by independent third parties, nor to fund objective analysis of data obtained from such infrastructure. Lacking such a capability, the regulators must assume that actors will choose to gather and report numbers in a form that represents the interest of those actors.

BIBLIOGRAPHY

- Anshelevich, E., and Gordon Wilfong. "Network Formation and Routing by Strategic Agents Using Local Contracts." In WINE, 386–93, Shanghai, 2008.
- Anshelevich, Elliot, Bruce Shepherd, and Gordon Wilfong. "Strategic Network Formation through Peering and Service Agreements." In Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science, Berkeley, CA, 2006.
- Arcaute, Esteban, Ramesh Johari, and Shie Mannor. "Local two-stage myopic dynamics for network formation games." In Papadimitriou, Christos, Shuzhong Zhang, eds. *Internet and Network Economics*, 263–77. Berlin, Heidelberg: Springer-Verlag, 2008.
- . "Network Formation: Bilateral Contracting and Myopic Dynamics." In Deng, Xiaotie, Fan Chung Graham, eds. *Internet and Network Economics*, 191–207. Berlin, Heidelberg: Springer-Verlag, 2007.

55. From the merger order:

This condition will enable the monitoring of the combined entity's future interconnection agreement's terms to determine whether the combined entity is using such agreements to deny or impede access to its networks in ways that limit competition from third-party online video content providers. In addition, this condition requires the combined entity to work with an independent measurement expert to report certain Internet interconnection performance metrics, and to the extent possible, make such metrics publicly available. Federal Communications Commission, "In the Matter of Applications of AT&T Inc. and DIRECTV For Consent to Assign or Transfer Control of Licenses," P. 148.

- Baake, Pio, and Thorsten Wichmann. "On the Economics of Internet Peering," *NETNOMICS* 1, no. 1 (1999): 89–105. [In English] doi:10.1023/A:1011449721395.
- Badasyan, Narine, and Subhadip Chakrabarti. "A Simple Game-Theoretic Analysis of Peering and Transit Contracting among Internet Service Providers," *Telecommunications Policy* 32, no. 1 (2008): 4–18.
- Barford, Paul, and Joel Sommers. "Comparing Probe- and Router-Based Packet-Loss Measurement." *IEEE Internet Computing* 8(5) (2004): 50–56.
- Bauer, Steven, David Clark, and William Lehr. "The Evolution of Internet Congestion." In Proceedings of the 37th Research Conference on Communication, Information and Internet Policy (TPRC), Arlington, VA, 2009.
- . "Understanding Broadband Speed Measurements." In TPRC, Arlington, VA, 2010.
- Brodkin, Jon. "After Netflix Pays Comcast, Speeds Improve 65%," *Ars Technica*, April 14, 2014, <https://arstechnica.com/information-technology/2014/04/after-netflix-pays-comcast-speeds-improve-65/>.
- . "Time Warner, Net Neutrality Foes Cry Foul Over Netflix Super HD Demands," *Ars Technica*, January 17, 2013, <https://arstechnica.com/information-technology/2013/01/timewarner-net-neutrality-foes-cry-foul-netflix-requirements-for-super-hd/>.
- Bustamante, Fabian, David Clark, and Nick Feamster. "Workshop on Tracking Quality of Experience in the Internet: Summary and Outcomes." *ACM SigComm Computer Communications Review*, 37(1) (2017): 55–60.
- Center for Information Technology Policy. "Interconnection Measurement Project," April 2016. <http://interconnection.citp.princeton.edu/project/views-by-interconnect/>.
- Chang, Hyunseok and Sugih Jamin. "To Peer or Not to Peer: Modeling the Evolution of the Internet AS-Level Topology." In IEEE INFOCOM, Barcelona, Spain, 2006.
- claffy, kc, Amogh Dhamdhere, David Clark, and Steven Bauer. First Amended Report of AT&T Independent Measurement Expert: Reporting Requirements and Measurement Methods," August 2016, <https://ecfsapi.fcc.gov/file/108042516812991/MB%20Dkt%2014-90%20AT%26T%20Inc.%20First%20Amended%20IME%20Report%20ECFS.PDF>.
- . "Report of AT&T Independent Measurement Expert: Background and Supporting Arguments for Measurement and Reporting Requirements," August 2016, https://www.caida.org/publications/papers/2016/att_ime_justification/.
- Clark, David, Steven Bauer, William Lehr, kc claffy, Amogh Dhamdhere, Bradley Huffaker, and Matthew Luckie. "Measurement and Analysis of Internet Interconnection and Congestion." In Telecommunications Policy Research Conference (TPRC), September 2014.
- Clark, David, Amogh Dhamdhere, Matthew Luckie, and kc claffy. "Policy Implications of Third-Party Measurement of Interdomain Congestion on the Internet." In TPRC 46th Research Conference on Communication, Information and Internet Policy, August 15, 2018. <https://ssrn.com/abstract=3141671> or <http://dx.doi.org/10.2139/ssrn.3141671>.
- Dai, Wei, and Scott Jordan. "Modeling ISP Tier Design." In 25th International Teletraffic Congress (ITC), 1–9, Shanghai, China, 2013. doi:10.1109/ITC.2013.6662957.
- Dhamdhere, Amogh, David Clark, Alexander Gamero-Garrido, Matthew Luckie, Ricky K. P. Mok, Gautam Akiwate, Kabir Gogia, Vaibhav Bajpai, Alex C. Snoeren and kc claffy. "Inferring Persistent Interdomain Congestion." In Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication (SIGCOMM '18), ACM, New York, NY, 1–15, 2018. <https://doi.org/10.1145/3230543.3230549>.
- Dhamdhere, Amogh, Constantine Dovrolis, and Pierre Francois. "A Value-Based Framework for Internet Peering Agreements." In International Teletraffic Congress (ITC), Amsterdam, the Netherlands, October 2010.
- Dhamdhere, Amogh, and Constantine Dovrolis. "The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh." In Proceedings of the 6th

- International Conference (21:1–21:12. Co-NEXT '10), ACM, Philadelphia, PA, , 2010. doi:10.1145/1921168.1921196. <http://doi.acm.org/10.1145/1921168.1921196>.
- FCC. “Measuring Broadband America,” 2015. <http://www.fcc.gov/measuring-broadband-america>.
- Feamster, Nick. “Revealing Utilization at Internet Interconnection Points.” arXiv e-prints, March 2018, <https://arxiv.org/abs/1603.03656>.
- Federal Communications Commission. “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,” July 2015. https://apps.fcc.gov/edocs_public/attachmatch/FCC-15-94A1.pdf.
- Federal Communications Commission. In the Matter of Applications of Charter Communications, Inc., Time Warner Cable Inc., and Advance/Newhouse Partnership For Consent to Assign or Transfer Control of Licenses and Authorizations: MB Docket No. 15-149,” March 2016, https://transition.fcc.gov/Daily_Releases/Daily_Business/2016/db0510/FCC-16-59A1.pdf.
- . “In the Matter of Protecting and Promoting the Open Internet: REPORT and Order on Remand, Declaratory Ruling, and Order, GN Docket No. 14–28,” March 2015, https://apps.fcc.gov/edocs_public/attachmatch/FCC-15-24A1.pdf.
- . “Notice of Proposed Rulemaking,” May 2014, <https://www.fcc.gov/document/protecting-and-promoting-open-internet-nprm>.
- Florence, Ken. “The Case against ISP Tolls,” 2014, <https://web.archive.org/web/20150915103905/http://blog.netflix.com/2014/04/the-case-against-isp-tolls.html?m=1>.
- Google. “Google Video Quality Report - Methodology,” 2019. <https://www.google.com/get/videoqualityreport/m/methodology.html>.
- Google. “Google Video Quality Report,” 2019. <https://www.google.com/get/videoqualityreport/>.
- He, Linhai, and Jean Walrand. “Pricing and Revenue Sharing Strategies for Internet Service Providers.” In Proceedings IEEE-INFOCOM 24th Annual Joint Conference of the IEEE Computer and Communications Societies, vol. 1, 205–216, Miami, FL, 2005.
- Hedayat, K., R. Krzanowski, A. Morton, K. Yum, and J. Babiarz. *A Two-Way Active Measurement Protocol (TWAMP): RFC 5357*, 2008. Internet Engineering Task Force.
- Holme, Petter, Josh Karlin, and Stephanie Forrest. “An Integrated Model of Traffic, Geography and Economy in the Internet.” In *SIGCOMM Computer Communication Review*, New York, NY, vol. 38, no. 3, July 2008, 5–16. doi:10.1145/1384609.1384611. <http://doi.acm.org/10.1145/1384609.1384611>.
- Huffaker, B., Marina Fomenkov, and kc claffy. *Geocompare: A Comparison of Public and Commercial Geolocation Databases* (Technical Report). San Diego, CA: Cooperative Association for Internet Data Analysis (CAIDA), May 2011.
- IEEE. 802.1AX-2014 – IEEE Standard for Local and metropolitan area networks – Link Aggregation, 2014.
- Khoury, Jennifer. “Comcast Response to Netflix,” 2014, <https://corporate.comcast.com/comcast-voices/comcast-response-to-netflix>.
- Kilmer, Hank. “M-Labs Data and Cogent DSCP Markings,” 2014, https://groups.google.com/a/measurementlab.net/forum/#!msg/discuss/vcQnaZJO6nQ/ltfi_3Aif9gJ.
- Lodhi, A., Amogh Dhamdhere, and Constantine Dovrolis. “GENESIS: An Agent-based Model of Interdomain Network Formation, Traffic Flow and Economics.” In Proceedings IEEE INFOCOM, Orlando, FL, 2012.
- . “GENESIS-CBA: An Agent-Based Model of Peer Evaluation and Selection in the Internet Interdomain Network,” *Complex Adaptive Systems Modeling* 1 (2013): 10.
- . “Open Peering by Internet Transit Providers: Peer Preference or Peer Pressure?” In Proceedings IEEE INFOCOM, Toronto, ON, Canada, 2014.

- Lodhi, Aemen, and Constantine Dovrolis. "A Network Formation Model for Internet Transit Relations." In Proceedings of the 2010 Workshop on Economics of Networks, Systems, and Computation. Vancouver, BC, Canada, ACM, 2010.
- Luckie, M., Amogh Dhamdhere, David Clark, Bradley Huffaker, and kc claffy. "Challenges in Inferring Internet Interdomain Congestion." In Internet Measurement Conference (IMC), Vancouver, BC, Canada, November 2014, 15–22.
- Ma, Richard T. B., John C. S. Lui, and Vishal Misra. "On the Evolution of the Internet Economic Ecosystem." In Proceedings of the International Conference on World Wide Web, Rio de Janeiro, Brazil, 2013.
- Meirom, E. A., Shie Mannor, and Ariel Orda. "Network Formation Games and the Internet Structure." *ArXiv e-prints*, July 2013. arXiv: 1307.4102[cs.GT].
- Miller, James. "Collaborative Meeting Report", *FCC Filing*, April 2015, <https://ecfsapi.fcc.gov/file/60001043350.pdf>.
- M-Lab Research Team. "ISP Interconnection and its Impact on Consumer Internet Performance—A Measurement Lab Consortium Technical Report," 2019, <http://www.measurementlab.net/publications>.
- Moller, Sebastian, and Alexander Raake, eds. *Quality of Experience: Advanced Concepts, Applications and Methods*. Basel, Switzerland: Springer, 2014.
- National Cable and Telecommunications Association. "Letter to the FCC Re: Broadband Performance Measurement, GN Docket No. 12–262". March 2015, <https://ecfsapi.fcc.gov/file/60001040966.pdf>.
- Netflix. Fast.com, 2016, <https://netflixtechblog.com/building-fast-com-4857feof8adb>.
- . "Petition to Deny of Netflix, Inc.," *FCC Filing*, August 2014, <https://ecfsapi.fcc.gov/file/7521819696.pdf>.
- Ramesh, Johari, Shie Mannor, and John N Tsitsiklis. "A Contract-Based Model for Directed Network Formation," *Games and Economic Behavior* 56, no. 2 (2006): 201–24.
- Ramesh, Johari and John N. Tsitsiklis. *Routing and Peering in a Competitive Internet*. LIDS Publication 2570, Cambridge, MA: Massachusetts Institute of Technology, 2003.
- Razdan, Rajender. "Collaborative Meeting Report," *FCC Filing*, January 2016, <https://ecfsapi.fcc.gov/file/60001409201.pdf>.
- Republique Francaise: Autorite de la concurrence. *Internet Traffic – Peering Agreements*. http://www.autoritedelaconcurrence.fr/user/standard.php?id_rub=418&cid_article=1971, 2012.
- Samknows. "Collaborative Meeting Presentation," *FCC Filing*, July 2014, <https://ecfsapi.fcc.gov/file/7521751768.pdf>.
- . "Collaborative Meeting Presentation," *FCC Filing*, April 2015, <https://ecfsapi.fcc.gov/file/60001043352.pdf>.
- Shakkottai, S., R. Srikant, Asuman Ozdaglar, and Daron Acemoglu. "The Price of Simplicity," *IEEE Journal on Selected Areas in Communications* 26, no. 7 (September 2008): 1269–76. doi:10.1109/JSAC.2008.080923. <http://dx.doi.org/10.1109/JSAC.2008.080923>.
- Shrimali, Gireesh, and Sunil Kumar. "Bill-and-Keep Peering," *Telecommunications Policy* 32, no. 1 (2008): 19–32.
- . "Can Bill-and-Keep Peering Be Mutually Beneficial?" In Deng, Xiaotie, Yinyu Ye, eds., *Internet and Network Economics*, 738–47. Hong Kong: Springer, 2005.
- . "Paid Peering among Internet Service Providers." In Proceeding from the 2006 Workshop on Game Theory for Communications and Networks, GameNets '06, Pisa, Italy, ACM, 2006. doi:10.1145/1190195.1190207. <http://doi.acm.org/10.1145/1190195.1190207>.
- Stanojevic, Rade, Nikolaos Laoutaris, and Pablo Rodriguez. "On Economic Heavy Hitters: Shapley Value Analysis of 95th-Percentile Pricing." In Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement, ACM, 75–80, Melbourne Australia, 2010.

- Sundaresan, Srikanth, Nick Feamster, and Renata Teixeira. "Home Network or Access Link? Locating Last-Mile Downstream Throughput Bottlenecks." In *Passive and Active Measurement (PAM)*, 2016, <http://www1.icsi.berkeley.edu/~srikanth/docs/hoa-pam-2016.pdf>.
- Taylor, Mark. "Observations of an Internet Middleman," 2014, <https://web.archive.org/web/20141108012025/http://blog.level3.com/open-internet/observations-internet-middleman/>.
- United States Court of Appeals for the D.C. Circuit. United States Telecom Association, et al., petitioners v. FCC, June 2016. [https://www.cadc.uscourts.gov/internet/opinions.nsf/3F95E49183E6F8AF85257FD200505A3A/\\$file/15-1063-1619173.pdf](https://www.cadc.uscourts.gov/internet/opinions.nsf/3F95E49183E6F8AF85257FD200505A3A/$file/15-1063-1619173.pdf).
- Valancius, Vytautas, Cristian Lumezanu, Nick Feamster, Ramesh Johari, and Vijay V. Vazirani. "How Many Tiers? Pricing in the Internet Transit Market." In Proceedings of the ACM SIGCOMM 2011 Conference, Toronto, ON, Canada, ACM, 194–205, 2011. doi:10.1145/2018436.2018459. <http://doi.acm.org/10.1145/2018436.2018459>.