

# The utility argument – Making a case for broadband SLAs

Zachary S. Bischof<sup>†</sup>, Fabián E. Bustamante<sup>†</sup>, Rade Stanojevic<sup>\*</sup>

<sup>†</sup>Northwestern University    <sup>\*</sup>Qatar Computing Research Institute

## Abstract

Most residential broadband services are described in terms of their maximum potential throughput rate, often advertised as having speeds “up to  $X$  Mbps”. Though such promises are often met, they are fairly limited in scope and, unfortunately, there is no basis for an appeal if a customer were to receive compromised quality of service. While this ‘best effort’ model was sufficient in the early days, we argue that as broadband customers and their devices become more dependent on Internet connectivity, we will see an increased demand for more encompassing Service Level Agreements (SLA).

In this paper, we study the design space of broadband SLAs and explore some of the trade-offs between the level of strictness of SLAs and the cost of delivering them. We argue that certain SLAs could be offered almost immediately with minimal impact on retail prices, and that ISPs (or third parties) could accurately infer the risk of offering SLA to individual customers – with accuracy comparable to that in the car or credit insurance industry – and price the SLA service accordingly.

## 1 Introduction

In today’s broadband markets, service plans are typically described in terms of their maximum download throughput rate, often advertised as “up to  $X$  Mbps”. This advertised capacity, along with the associated monthly cost, are the two primary, and many times only, pieces of information available to consumers when comparing service providers. Such “constrained” service agreements place services using technologies as diverse as fiber, DSL, WiMAX or satellite on nearly equal grounds, and leave consumers without clear expectations given that, strictly speaking, any speed less than  $X$  would meet such a guarantee.

We argue that as Internet users and their devices become more dependent on connectivity and consistency, broadband will move from a loosely regulated luxury to a key utility. This in turn will usher in a growing demand for more encompassing, well-defined SLAs similar to those of other utilities, such as electricity and water.

We believe that the adoption of SLAs could benefit all players in the broadband market – service providers, customers, and regulators. From the ISP’s perspective, contracts with SLAs could allow them to better differentiate their retail services and fine-tune their contracts to the needs of particular classes of customers (e.g., a service for

gamers or business users).<sup>1</sup> For customers, SLAs could significantly simplify the process of comparing services offered by different providers, allowing customers to make more informed decisions. This could improve competition and potentially lower prices. Similarly, for regulators and policymakers, SLAs would provide a better way to gauge broadband infrastructure across communities and justify investments.

Despite these potential benefits, there are several challenges in defining SLAs for broadband services that range from identifying metrics and defining the appropriate SLA structures to engineering compliance monitoring.

SLAs must be designed so that they can be accurately and efficiently monitored and that they add value to providers and consumers, while limiting the risk of non-compliance. We expect broadband SLAs to be specified, as other network SLAs, in terms of transport-level performance assurances using Quality of Service metrics such as bandwidth, packet loss, delay and availability. While the relationship between such QoS metrics and users' experience with different applications is a topic of ongoing research, existing approaches rely on such QoS metrics as input to application specific models of QoE estimation (e.g., [5,15]).

An SLA could be seen as an insurance policy against the risk of not receiving the contracted level of service. Consequently, SLA-enhanced services would come with a price-tag for providers that depends on the structure of the SLA and degree of risk involved in the delivering the desired levels of service. Using four-years of data from the largest, publicly available dataset of broadband performance [10], we study the design space of broadband SLAs and demonstrate that certain SLAs could be offered almost immediately with minimal impact on retail prices and network investment.

In this paper, we make a case for broadband SLAs and follow a data-driven approach to explore some of these key challenges. We make the following contributions:

- We analyze different QoS metrics for use in SLA and define a set of broadband SLAs (§2). We find that, across all ISPs and access technologies, bandwidth is the most consistent of the studied performance metrics.
- We evaluate the relationship between SLA structure and the cost of supporting them with different access technologies (§3). We show that many of the studied ISPs could offer moderate SLAs with minimal impact on their existing business, but that SLAs with stringent constraints are much harder to deliver across the whole user-base.
- We show that ISPs (or third parties) could accurately infer the risk of offering SLA to individual customers – with accuracy comparable to that in the car or credit insurance industry – and price the SLA service accordingly (§4).

We conclude the paper (§5 and 6) with a discussion of some of the key open issues and potential future directions for this work.

---

<sup>1</sup> Some ISPs already try this in coarser terms; e.g., Comcast's "What type of Internet connection is right for you?" <http://www.xfinity.com/resources/internet-connections.html>

SLA	Throughput (% of service)	Latency	Packet loss	Description
A	> 90%	< 50 ms	< 1%	Demanding applications (e.g., real-time gaming) Video streaming, telephony Web browsing, email
B	> 50%	< 150 ms	< 5%	
C	> 10%	< 250 ms	< 10%	

**Table 1.** Three examples of possible broadband SLAs.

## 2 Metrics for a broadband SLA

An SLA is a contract between a service provider and its customers that specifies what services the provider will support and what penalties it will pay upon violations. A meaningful SLA should (i) capture the needs of consumers, (ii) be feasible to support by most service providers today and (iii) be expressed in measurable terms that can be validated by both consumers and services providers.

To understand the need of broadband consumers, we must consider the requirements of commonly used network applications. Clearly, one would not expect that “broadband consumers” would be a homogeneous class in either the type of applications they value most or their expectations. For some consumers, being able to browse the web or read their email may be sufficient, and paying for a higher guaranteed throughput would not be a priority. Others may have higher performance requirements, wanting a connection that reliably allows them to stream HD video content or play online games with strict timing requirements.

Driven by these observations, current literature on the needs of different application classes (e.g., [7,6,27]) and our dataset [10], we drafted three potential broadband SLAs that cover a wide range of user requirements. Note that these are mere examples of possible SLAs, focused on the points relevant to our argument, and ignoring the specifics of a practical SLA, such as the form of reporting quality of service violations, the procedure to be invoked in case of violations or the exact cost model of violations.

Our basic SLAs (see Table 1) are stated in terms of throughput, latency and packet loss. Considering that subscription capacity is already advertised by ISP and varies across users, we structure SLAs in terms of the percentage of subscription speed available. For latency and packet loss, we adopt a simple “below-threshold” model. *SLA A* represents a service that should be able to fit the demands of users with very strict performance requirements for applications such as real time gaming. *SLA C* characterizes a service that could support simple applications, such as browsing the web or email. Finally, *SLA B* matches the middle-of-the-road services, capable of supporting most applications, such as video chat or video streaming, but with less than perfect performance for network-intensive applications.

Although they are somewhat arbitrary, the thresholds we use for our sample SLAs – from fractions of throughput to latency and loss rate – are based on existing literature and earlier studies of broadband services.

For service capacity, we selected 10% of capacity as a bottom threshold (*SLA C*) since the vast majority of users in our dataset had a connection much faster than 1 Mbps and that 100 kbps can support basic browsing and email requirements. We opted for 50% as a threshold for *SLA B* following a 2010 report from the UK Office of Communication (Ofcom) reporting that surveyed users received, on average, nearly half (46%) of their advertised speed [16]. Finally, for our highest SLA we opted for 90% as a

threshold to highlight providers that consistently deliver capacities close to their subscription speeds.

In terms of packet loss, previous work has shown that rates above 1% can have a negative impact on users' QoE while using gaming applications [7]. High loss rates can also affect other common services such as audio and video calls [6]. Xu et al. [27], for instance, shows that loss rates above 4% can significantly degrade iChat video calls and rates larger than 10% result in a sharp increase in packet retransmissions.

We selected thresholds for latency in a similar manner. Our least demanding SLA, *SLA C*, has a latency threshold of 250 ms, since larger latencies can significantly increase page loading time [3] and would likely have a negative impact on QoE. End-to-end latencies below approximately 150 ms, the threshold for *SLA B*, should be sufficient for Skype calls [21]. Last, our low threshold for *SLA A* is based on previous work showing that an increase of just 10 ms can yield an increase in page loading delays by hundreds of milliseconds [25].

### 3 Supporting SLA today


Building on these SLAs that would be meaningful to end-users, we now explore what sort of service guarantees it would be feasible for ISPs to provide to subscribers. We do this by looking at the performance and consistency of broadband services offered by US-based ISPs. We first describe the dataset on broadband services used throughout our analysis.

#### 3.1 Dataset


We leverage the largest, publicly available dataset of broadband performance collected through the FCC's Measuring Broadband America effort [10]. Since 2011, the US FCC has been working with SamKnows to distribute home gateways ("whiteboxes") to broadband customers that conduct and report network measurements. These devices have collected increasingly rich data, including metrics such as latency, throughput and page loading time for a number of popular websites. A full description of all the tests performed and data collected is available in the FCC's technical appendix [9]. This data has been mostly used by the FCC to create periodic reports on the state of broadband access in the United States.

We employ the full four years of measurements made available in order to quantify network performance in terms of latency, packet loss, and download/upload throughput. For this, we used three different measurement tables (out of eleven present) from the dataset for our analysis: (1) UDP pings, (2) HTTP GETs, which measure download throughput, and (3) HTTP POSTs, which measure upload throughput.

The UDP pings run continuously, measuring round-trip time to two or three measurement servers. These servers are hosted by either M-Lab or within the user's provider's network. Over the course of each hour, the gateway will send up to 600 probes to each measurement server at regular intervals, less if the link is under heavy use for part of the hour. Each gateway reports hourly statistical summaries of the latency measurements (mean, min, max, and standard deviation) as well as the number of missing responses. We use the average latency to the nearest server (in terms of latency) to summarize the latency during that hour. We also use the number of missing responses to calculate the packet loss rate over the course of each hour.



**Fig. 1.** CDF of measured download throughput rates as a fraction of the subscriber’s service capacity.



**Fig. 2.** CCDF of the fraction of a user’s download throughput measurements per month that are below a percentage of the subscription capacity. Each gray vertical line represents a frequency of (from left to right) once a month, once a week, once a day, and once every other hour.

As mentioned above, the HTTP GET and POST tables record the measured download and upload throughput rate, respectively. Similar to the latency measurements, throughput measurements are typically done to two different target servers. However, throughput measurements are run once every other hour, alternating between measuring upload and download throughput rates.

We combined these performance measurements with user metadata, which includes the user’s provider, service tier (i.e., subscription speed), service technology (e.g., cable or DSL), and geographic region. This allows us to group measurements by ISP, compare the achieved throughput as a fraction of the subscription capacity and differentiate between subscribers of the same ISP with different access technologies (e.g., Verizon customers with DSL or fiber).

### 3.2 Throughput

We first analyze ISPs’ download and upload throughput. A challenge in comparing performance across providers and services, is that users do not have the same subscription speeds; individual ISPs typically offer a number of service capacities and the stated capacities of such offerings vary from one ISP to another. In order to directly compare the consistency of performance, we first normalize throughput measurements by the speed that each user should be receiving. For this, we use the reported download and upload subscription rate included as part of the FCC dataset, as described in Section 3.1.

**Throughput distribution.** Figure 1 shows a CDF of each normalized download throughput measurement from subscribers of four services: AT&T’s DSL service, Clearwire, Comcast, and Frontier’s fiber service. Of the services we studied, Frontier’s fiber

had the most consistent throughput rates, both in terms of the fraction of probes that measured at least 90% of the subscription speed and the variations between measurements. Although measurements were unlikely to achieve download rates significantly higher than their subscription speed, 96% of measurements were above 90% of the subscription speed.

For Comcast (cable), measurements were slightly less likely to reach 90% of the subscription speed (about 91%). However, download throughput measurements were often much higher than the user’s subscription speed – the median measurement on Comcast’s network was 135% of the subscription speed. We observed a similar trend for most cable broadband providers, as well as Verizon’s fiber service.

Download throughput measurements from subscribers of AT&T’s DSL service were fairly consistent (i.e., showing little variation). However, in contrast to cable and fiber services, they rarely exceeded the subscription speed, with less than 10% of measurements at or above the subscription speed. Nearly half (48%) of measurements were below 90% of the subscription speed. Other DSL providers showed a similar trend. Of the ISPs in our study, Clearwire had the largest fraction of measurements (73%) below 90% of the subscription speed.

**Variation over time.** Looking only at Figure 1, it is still unclear how much performance can vary for an individual subscriber over the course of a month. To capture this, we aggregated all measurements that were conducted from the same vantage point and run during the same month, which we refer to as a “user-month”. For each user-month, we calculate the fraction of measurements that were below a threshold of 10%, 25%, 50%, 75%, and 90% of the subscription speed.


Figure 2 shows, for AT&T, Comcast and Frontier fiber subscribers, how frequently measurements during the same month measured below a particular threshold. The vertical gray lines represent a particular frequency of throughput measurements being below a given threshold (from left to right): once a month, once a week, once a day, and once every other hour.

In the case of AT&T, shown in Figure 2a, during 47% of the user-months, subscribers got less than 90% of their service capacity at least once every other hour (the right-most vertical line). In contrast, for Comcast subscribers, shown in Figure 2b, only about 9% of user-months measured less than 90% of the subscription speed at the same frequency. Comcast users were also less likely to receive less than 50% of their subscription speed. Frontier’s fiber was even less likely to have degradations in download throughput every other hour; less than 3% of months of Frontier measurements saw throughput rates below 90% of the subscription speed.



In general, the distributions of upload (Fig. 3) and download throughput measurements shown similar trends. The most obvious difference was that upload measurements from Clearwire subscribers were noticeably higher, more consistent, and much closer to the subscription speed. For each ISP in Figure 3, the median measurement was at least 90% of the subscription speed.

### 3.3 Latency

As mentioned in Section 3.1, we use the average latency (measured over an hour) to the nearest measurement server as an estimate of the subscriber’s latency. Figure 4 shows a CDF of the hourly average latency for five ISPs. Cablevision, with 96% of



**Fig. 3.** CDF of upload throughput rates as a fraction of the subscriber's service capacity.



**Fig. 4.** CDF of latency measurements to servers. **Fig. 5.** CCDF of hourly loss rates to servers.

hourly averages below 20 ms, showed the lowest latencies of all ISPs in our dataset and appeared to consistently meet even the most demanding SLA. Other fiber, DSL, and cable ISPs had slightly higher latencies, but were fairly consistent, with at least 90% of average latency reports for each provider being less than 70 ms. AT&T, with 95% of measurements below 57 ms had the lowest latencies of all DSL providers, but the overall distribution was higher than most cable providers.

Latency measurements from Clearwire subscribers were noticeably higher, with a median of approximately 90 ms. Satellite providers had the highest latency measurements, consistently above 600 ms, as a result of the fundamental limitations of the technology.


### 3.4 Packet loss

Using the number of UDP pings that succeeded and failed to the target measurement server, we calculated the percentage of packets lost over each hour. Figure 5 shows the CCDF of the hourly packet loss rates for four ISPs. On average, fiber providers tended to have lower loss rates and had the lowest frequency of high loss. More specifically, Verizon had the lowest frequency of loss rates above 1%, occurring during only 0.82% of hours. Comcast (not in the figure) and TimeWarner had the lowest frequency for cable providers, with loss rates above 1% occurring in approximately 1.5% of hours. Satellite providers had the highest frequency of loss rates above 1%, occurring during over 26% of hours.

### 3.5 Applying an SLA

In Section 2, we defined SLAs in measurable terms with thresholds that would be meaningful to users' Quality of Experience. Building on our characterization, we now explore how effectively today's ISPs could meet our proposed set of SLAs.

There are a number of ways that a broadband SLA could be structured in terms of how users are compensated for periods of poor performance. As an example, we looked



**Fig. 6.** CDF of the percent of hours per month in violation of each SLA for four providers.

at how some broadband ISPs structure the agreements that they offered to businesses. In the case of Comcast [8], business class subscribers are compensated once the network become unavailable for more than four hours in a single month. For each hour of downtime after the first four, customers are reimbursed  $1/30$  of the monthly subscription price.<sup>2</sup> We believe that general broadband service plans could have a similar structure. For example, the SLA could state that the network may be unavailable for up to two hours per day (or about 8.33% of hours in a month). This would allow ISPs to schedule downtime for maintenance and provide a guarantee for subscribers that their service will not be down for days at a time (or that they will be compensated if it is).

However, our focus in this paper is not on the structure of compensation for SLA violations. Instead, we look at how well the ISPs in the FCC’s dataset are able to meet the SLAs defined in Section 2, and whether it would be at all feasible to provide guarantees of service.

Figure 6 summarizes the total number of SLA violations per month for four example ISPs. AT&T, shown in Figure 6a struggles to meet the requirements of *SLA A* but is able to meet *SLA B* during 90% of hours per month for 73% of users and meets *SLA C* during 90% of hours for 82% of users.

The wireless provider in our dataset, Clearwire (Fig. 6b), face difficulties in meeting *SLA A*, as the average latencies were almost always higher than 50 ms. This appears to be a result of the underlying technology and many cellular providers are unable to meet this latency requirement [18]. Interestingly, Clearwire actually did a better job of meeting *SLA C* than AT&T, with 94% of users meeting *SLA C* performance during at least 90% of hours in a month.

Both Comcast and Verizon’s fiber service did a relatively good job of meeting the requirements of all three SLAs. Comcast was able to meet *SLA A* during 90% of hours in a month for 75% of users while Verizon was able to do the same for 83% of fiber subscribers. Both were able to provide both *SLA B* and *C* during 90% of hours for at least 90% of users.

To summarize our findings in this section, moderate SLAs (those which require SLA compliance up to 90% of time) are feasible nowadays and could be offered by many ISPs with minimal effect on their current business. However, stricter SLAs (those which require SLA compliance 99% of the time or more) would be much more challenging to offer across the whole user base. In the following section, we examine how difficult

<sup>2</sup> This effectively means that if the service was ‘unavailable’ for 34 hours in a month (approximately 5% of the month) the user gets the monthly subscription for free.



it would be to assess the individual risk of breaking and SLA, a central challenge in personalized SLA offerings.

## 4 Personalized SLAs

As we noted in a previous section, SLA can be seen as an insurance policy against poor broadband experience, which may in turn have financial consequences in case of broken SLA. In this section we study if SLAs could be tailored for each end-user individually.

The key question we try to answer is whether the provider could infer the likelihood of delivering the SLA. For instance, it is possible that certain user characteristics are correlated with the quality of service the user receives and hence the SLA provider may choose to price the service (premium in insurance terms) according to the risk of not delivering promised SLA to this set of users. With a good understanding of how likely it is to break the SLA the insurer (either a third party or the broadband provider itself) can fine tune the SLA parameters and the premium<sup>3</sup> (in \$ per month) in order to improve user satisfaction with the service and ensure the profitability of the SLA service.

We train a simple model to examine the predictability of the service of individual subscribers complying to an SLA based on several simple user features available to us: (1) access technology, (2) base latency (to the nearest measurement server), (3) aggregate usage (in bytes per month) and (4) city population (a proxy of urban/rural residence). More advanced models, using a range of additional demographic and technological features, would likely improve the prediction accuracy, yet such analysis is out of scope of this short study and is left for future work.

We use supervised learning for estimating the likelihood of breaking the SLA, for the three SLA types described in Table 1 with 95% time threshold (i.e., the users' performance complies with the SLA 95% of the time). This is basically a binary classification task, where we use four user features described above to predict whether the user complies with SLA or not. The features are extracted on 4038 active users in October and November 2014. The categorical feature describing access technology is projected to a binary vector (of length 4) encoding the access technology of every user.

We experimented with several classification methods including L2-regularized logistic regression, gradient boosting trees and random forests. We report the results from random forests which showed slightly better performance although the performance of all methods were comparable. The hyper-parameters were optimized using a grid-search over a validation set extracted from the training set. We use fourfold cross validation to predict the chance of breaking SLA. The features are extracted in October 2014 and the (binary) SLA compliance is extracted for November 2014.

We use *Area Under Curve Receiver Operating Characteristic (AUCROC)*, a standard metric for measuring the performance of the binary classifiers[13]. The ROC curve as well as the AUCROC are reported in Figure 7 for the three SLAs from Table 1.

The AUCROC for all three SLAs: A, B and C, is similar and is around 0.8. Such AUCROC is comparable to the precision of classifiers build from demographic user

---

<sup>3</sup> The key cost for the ISP selling an SLA is the loss of revenue when the SLA is broken. Hence the stricter SLA the higher expected cost for the ISP which may be passed down to the end-user in the form of higher premium/monthly subscription.

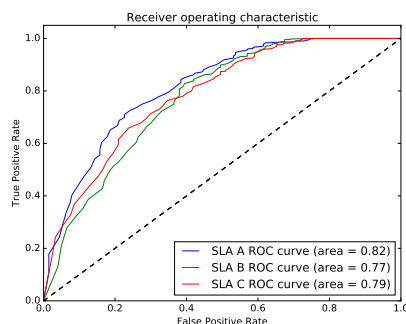


Fig. 7. Area under ROC for Random Forest classifiers.

information in other insurance products such as cars and credit ratings [17]. This accuracy of prediction for SLA compliance suggests that it would be possible to offer personalized SLAs with a price which accurately matches the likelihood of breaking the SLA.

City population and population density are loosely correlated with last-mile performance, especially in the context of DSL [1]. However, city population appears to have minor predictive power for inferring the likelihood of SLA violations in our dataset.

## 5 Discussion

Recent efforts [24,4,19,20] have attempted to address the lack of detailed evaluations of ISPs. Annual reports published by the FCC in the US and Ofcom in the UK have studied whether or not ISPs are providing the capacities promised to users. The recent Net Neutrality ruling from the FCC [11,12] discussed the issue of how service plans are described to subscribers. One part of the ruling states that ISPs must disclose reasonable estimates of performance metrics, including both latency and packet loss. Unfortunately, what exactly is a “reasonable” estimate of these metrics is somewhat unclear. Additionally, providing the estimates alone does not offer any protection for consumers that may experience seriously degraded performance.

This work points to a number of interesting research directions that are crucial for implementing broadband SLAs. For example, perhaps the largest roadblock to adoption of broadband SLAs is the lack of infrastructure for monitoring performance and reporting SLA violations. One potential avenue to explore would be the deployment of a system, such as SLAM [23], on home gateways or modems that could monitor SLA compliance. These devices could be distributed by the SLA provider (either the ISP or a third party). The design of a reliable processes for the automatic generation and filing of SLA violation reports and reporting, to both the subscriber and the ISP, is another interesting research direction.

There is also a need to consider factors beyond throughput, latency, and packet loss. For example, high packet delay variations could impact user quality of experience. Furthermore, recent peering disputes between content providers and broadband access providers [14,22] highlight the importance of measuring congestion on a provider’s peering links and its potential impact on performance. Poor quality of experience while

streaming via Netflix or making Skype calls would not be captured by the measurements used in this paper if this is caused by congestion at the edge of the provider's network.

Another aspect we have not explored is the design of SLAs that both fit what a user's needs and what they can afford, an area we have explored in past work [2]. For example, an SLA that promises to provide lower latency, from 25 ms to 15 ms, could come at a hefty price for the ISP and yet provide little value to subscribers. Additionally, the availability of other services that are typically hosted by the ISP, such as DNS or email, may be more important to some users than a guaranteed throughput rate.

Previous work has suggest that consumers could benefit from improvements in how service offerings are described to customers [26] and shown that the relationship between QoS metrics (as those we used in our definitions of SLA) and users' experience with different applications is an open research problem. Nevertheless, all existing approaches we are aware of rely on such QoS metrics as input to application specific models of QoE estimation (e.g., [5,15]).

## 6 Conclusion

This work is partially motivated by the FCC's recent classification of broadband as a utility. We believe that this is a natural course for broadband Internet, as it progresses from a luxury to a key utility and, in some countries, considered a basic human right. The growing understanding of broadband connectivity as a utility will, in turn, usher in a demand for more encompassing, well-defined SLAs. The introduction of SLAs could enable broadband operators to personalize the service offerings down to the individual customer and improve their efficiency and overall user satisfaction. Broadband SLAs could also facilitate transparent competition, ultimately benefiting both consumers and service providers. In this paper, we explored the possibility of implementing broadband SLAs and demonstrated that certain SLAs could be offered almost immediately with small impact on the retail prices and network investment. We showed that ISPs (or third parties) could accurately infer the risk of offering SLA to individual customers, with accuracy comparable to that in other insurance markets, and price SLA services accordingly.

## Acknowledgments

We thank our shepherd Monia Ghobadi and the anonymous reviewers for their invaluable feedback. This work was supported in part by the National Science Foundation through Award CNS 1218287.

## References

1. Z. Bischof, F. Bustamante, and N. Feamster. (The Importance of) Being Connected: On the Reliability of Broadband Internet Access. Technical Report NU-EECS-16-01, Northwestern University, 2016.
2. Z. S. Bischof, F. E. Bustamante, and R. Stanojevic. Need, want, can afford – broadband markets and the behavior of users. In *Proc. of IMC*, November 2014.
3. Z. S. Bischof, J. S. Otto, and F. E. Bustamante. Up, down and around the stack: ISP characterization from network intensive applications. In *Proc. of W-MUST*, 2012.

4. Z. S. Bischof, J. S. Otto, M. A. Sánchez, J. P. Rula, D. R. Choffnes, and F. E. Bustamante. Crowdsourcing ISP characterization to the network edge. In *Proc. of W-MUST*, 2011.
5. P. Casas, B. Gardlo, R. Schatz, and M. Melia. An educated guess on QoE in operational networks through large-scale measurements. In *Proc. of SIGCOMM Workshop Internet-QoE*, August 2016.
6. K.-T. Chen, C.-Y. Huang, P. Huang, and C.-L. Lei. Quantifying Skype user satisfaction. In *Proc. of ACM SIGCOMM*, 2006.
7. K.-T. Chen, P. Huang, and C.-L. Lei. How sensitive are online gamers to network quality? *Commun. ACM*, 49(11), Nov. 2006.
8. Comcast Business Class. Service level agreement. <http://business.comcast.com/pdfs/cbc-trunks-sla-110922.pdf>.
9. FCC. 2013 measuring broadband America February report. <http://data.fcc.gov/download/measuring-broadband-america/2013/Technical-Appendix-feb-2013.pdf>.
10. FCC. Measuring Broadband America. <http://www.fcc.gov/measuring-broadband-america>.
11. FCC. In the matter of preserving the Open Internet broadband industry practices, December 2010.
12. FCC. In the matter of protecting and promoting the Open Internet, February 2015.
13. W. Green. *Econometric Analysis*. Prentice Hall, 2003.
14. S. Higginbotham. Why the consumer is still held hostage in peering disputes. <http://bit.ly/1KbBBh1>, July 2014.
15. A. Nikravesh, D. K. Hong, Q. A. Chen, H. V. Madhyastha, and Z. M. Mao. QoE inference without application control. In *Proc. of SIGCOMM Workshop Internet-QoE*, August 2016.
16. Office of Communication (Ofcom). UK fixed broadband speeds, November/December 2010. Technical report, London, UK, March 2011.
17. J. S. Pedro, D. Proserpio, and N. Oliver. Mobiscore: Towards universal credit scoring from mobile phone data. 2015.
18. J. P. Rula and F. E. Bustamante. Behind the curtain: Cellular DNS and content replica selection. In *Proc. of IMC*, 2014.
19. SamKnows. Samknows & the FCC American broadband performance measurement. [http://www.samknows.com/broadband/fcc\\_and\\_samknows](http://www.samknows.com/broadband/fcc_and_samknows), June.
20. M. A. Sánchez, J. S. Otto, Z. S. Bischof, D. R. Choffnes, F. E. Bustamante, B. Krishnamurthy, and W. Willinger. Dasu: Pushing experiments to the Internet's edge. In *Proc. of USENIX NSDI*, 2013.
21. Skype. Plan network requirements for skype for business. <https://technet.microsoft.com/en-us/library/Gg425841.aspx>.
22. J. E. Solsman. Cogent: Comcast forced netflix with clever traffic clogging. <http://cnet.co/113aDw1>, May 2014.
23. J. Sommers, P. Barford, N. Duffield, and A. Ron. Accurate and efficient SLA compliance monitoring. In *Proc. of ACM SIGCOMM*, 2007.
24. S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè. Broadband Internet performance: A view from the gateway. In *Proc. of ACM SIGCOMM*, 2011.
25. S. Sundaresan, N. Feamster, R. Teixeira, and N. Magharei. Measuring and mitigating web performance bottlenecks in broadband access networks. In *Proc. of IMC*, October 2013.
26. S. Sundaresan, N. Feamster, R. Teixeira, A. Tang, W. K. Edwards, R. E. Grinter, M. Chetty, and W. de Donato. Helping users shop for ISPs with Internet nutrition labels. In *Proc. of HomeNets*, 2011.
27. Y. Xu, C. Yu, J. Li, and Y. Liu. Video telephony for end-consumers: measurement study of Google+, iChat, and Skype. In *Proc. of IMC*, 2012.