Assessing Broadband Reliability: Measurement and Policy Challenges

William Lehr[1]
Steven Bauer[2]
Mikko Heikkinen[3]
David Clark[4]

Massachusetts Institute of Technology[5]

Abstract

With broadband availability now approaching universal service in the United States, the focus on broadband metrics is shifting toward the assessment of broadband performance.[6] The initial focus has been on the measurement and reporting of the speeds achieved by broadband services.[7] As we develop more comprehensive and, hopefully, more meaningful frameworks for characterizing, measuring, and evaluating broadband performance,[8] one obvious measure of merit will be *reliability*. In a survey, consumers identified concerns about reliability as second only to speed in importance.[9] Moreover, as our collective reliance on broadband as critical socio-economic infrastructure increases, and as a growing range of services move into the Internet, it is reasonable to expect that demand for reliable broadband service will increase. At the same time, the growth of streaming and interactive rich multimedia, the convergence of fixed and mobile services, and the evolution of the Internet services industry pose new challenges for developing appropriate reliability metrics. While the measurement of broadband availability or speeds and the appropriate interpretation of measurement data are not simple tasks, evaluating reliability is inherently much more complex and difficult.

---

[1] Corresponding author: wlehr@mit.edu, 617-258-0630.

[2] bauer@mit.edu

[3] mikko.heikkinen@aalto.fi. During this work, Mr. Heikkinen was a Visiting Student at MIT CSAIL. He is also affiliated with the Aalto University School of Electrical Engineering and the Helsinki Institute for Information Technology HIIT. His work was supported by FICNIA and TEKES.

[4] ddc@csail.mit.edu

[5] In completing this work, Drs. Lehr, Bauer, and Clark would like to acknowledge support from NSF Awards 1040020, 1040023, the MIT Communications Futures Program and the MITAS Project. All opinions expressed herein are those of the authors alone.

[6] See Bauer et al. (2009a) for a discussion of the need for better broadband data metrics.

[7] See Bauer et al. (2010) for a discussion of broadband speed measurements.

[8] On June 1, 2010, the FCC announced plans for deploying a new broadband measurement infrastructure by the UK-based firm SamKnows (see http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-298525A1.pdf).

[9] See Figure 5.2 "Main reason for dissatisfaction with ISP" from Ofcom (2009).

Building on our earlier work on Internet metrics, we focus here on the challenge of assessing broadband reliability.[10] There is a long history of work on reliability in other industries and within various parts of the networking community. However, how that history of thought, measurements, and lessons learned applies to broadband reliability has not previously been explored in much depth. Topics we consider are different potential definitions of reliability, how to measure broadband reliability (metrics), and some of the major policy questions such as what, if any, reporting requirements should be for various types of outages at the current time. We also discuss the economics of reliability to better understand and anticipate how reliability concerns may impact our collective thinking about issues such as fixed/mobile substitution, network neutrality, and quality-of-service guarantees. Finally, we discuss some of the measurement efforts currently underway, including the FCC/SamKnows study, with an eye toward what these may reveal about broadband reliability.

1. Introduction

Earlier empirical research directed at characterizing and measuring the performance of broadband service markets focused on questions of broadband availability,[11] usage/adoption patterns,[12] pricing,[13] economic impacts,[14] and more recently, aspects of broadband service quality.[15] For example, in Bauer et al. (2010), we discussed some of the challenges associated with measuring the speed of broadband. In this paper, we turn to considering the challenges of assessing the reliability of consumer broadband services. The motivations for this inquiry are several, but as we explain, assessing reliability poses a much more difficult problem.

First, consumers care about the reliability of their broadband services. While the number one concern is generally "speed," a close second is "reliability."[16] As access speeds continue to increase, broadband reliability looks like it may become the primary frustration point for consumers.

[10] See, for example, Bauer et al. (2009a, 2009b, 2010), Lehr & Grieco (2009), and Lehr et al. (2008).

[11] See, for example, Downes & Greenstein (2002), Lehr & Gillett (1999), NTIA (1999), Flamm (2004), or Horrigan (2010).

[12] See, for example, the Pew Internet & American Life Project (http://www.pewinternet.org/) or Horrigan and Smith (2007).

[13] See, for example, Atkinson et al. (2008).

[14] See, for example, Gillett et al. (2006), Czernich et al. (2011), or OECD (2011).

[15] See Bauer et al. (2010).

[16] See Figure 5.2 "Main reason for dissatisfaction with ISP" from Ofcom (2009).

Second, broadband access to the Internet is increasingly viewed as basic infrastructure for our society and economy.[17] As such, there is a public policy responsibility to ensure affordable services are available to all citizens. Earlier, policymakers recognized the need for ensuring universal access to basic telephony services. Today, they are engaged in refocusing universal service policies toward promoting broadband Internet access.[18] Either implicitly or explicitly, any public subsidy program for promoting universal access to broadband will need to embed a notion of what constitutes an acceptable minimal level of service. Some notion of reliability will necessarily be embedded in any such definition.

Third, the Internet ecosystem is continuing to evolve. Important phenomena like the growth of mobile broadband,[19] the increased share of traffic associated with streaming media services like Netflix and YouTube,[20] and increased demand for cloud-based services are contributing to the complexity of assessing the performance of Internet services, including broadband. There is growing recognition that the Internet protocols

[17] President Obama has reaffirmed this position. Speaking for his administration, Susan Crawford commented in a speech on May 14, 2009 that "Broadband is the new essential infrastructure" (see http://www.broadcastingcable.com/article/232506-President_Obama_Focused_On_Broadband.php). Similar positions have been adopted in Europe, where the European Commission has concluded that "widespread and affordable broadband access is essential to realize the potential of the Information Society" (see http://ec.europa.eu/information_society/eeurope/2005/all_about/broadband/index_en.htm); in Australia, where a government report concludes that "ubiquitous, multi-megabit broadband will underpin Australia's future economic and social prosperity" (see http://www.dcita.gov.au/communications_for_consumers/internet/broadband_blueprint/broadband_blueprint_html_version/chapter_one_broadband_as_critical_infrastructure); in Japan, where the Japanese have joined with regional partners to "enable all people in Asia to gain access to broadband platforms" by 2010 (see http://www.dosite.jp/asia-bb/en/pdf/abp005.pdf).

[18] The U.S. National Broadband Plan (NBP) recommends "reform current universal service mechanisms to support deployment of broadband and voice in high-cost areas; and ensure that low-income Americans can afford broadband; and in addition, support efforts to boost adoption and utilization." (see page XI in FCC, "Connecting America: the National Broadband Plan," Federal Communications Commission, Washington, DC, March 16, 2010). Following up on this mandate, the FCC issued its Notice of Proposed Rule-making on Universal Service Reform in February 2011 (see FCC, "Notice of Proposed Rulemaking and Further Notice of Proposed Rulemaking (NPRM), in the Matter of Connect America Fund (WC Docket No. 10-90), a National Broadband Plan for Our Future (GN Docket No. 09-51), Establishing Just and Reasonable Rates for Local Exchange Carriers (WC Docket No. 07-135), High-Cost Universal Service Support (WC Docket No. 05-337), Developing an Unified Intercarrier Compensation Regime (CC Docket No. 01-92), Federal-State Joint Board on Universal Service (CC Docket No. 96-45), and Lifeline and Link-up (WC Docket No. 03-109), before the Federal Communications Commission, Washington DC, February 9, 2011).

[19] See, for example, http://www.oecd.org/document/4/0,3746,en_2649_34225_42800196_1_1_1_1,00.html.

[20] See, for example, http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360_ns827_Networking_Solutions_White_Paper.html.

need to be upgraded to enable better security, support for rich policy-based routing, and for more diverse traffic types.[21] In this more complex matrix of Internet services and needs, identifying appropriate notions and metrics for reliability will pose difficult challenges for industry, policymakers, and end users alike.

In the U.S., the Federal Communications Commission (FCC) launched a *Notice of Inquiry* in April 2011 to gather stakeholder perceptions and data on Internet reliability, to follow up on the National Broadband Plan's recommendation that the FCC "engage in an exploration of the reliability and resiliency standards being applied to broadband networks in order to ascertain what action, if any, the Commission should take to bolster the reliability of broadband infrastructures."[22] This NOI follows up on earlier inquiries regarding the ability of our infrastructure to support continuous communications in the face of manmade and natural disasters like the 9/11 terrorist attacks and Hurricane Katrina; and seeks to initiate the discussion regarding how best to assess the reliability and resiliency of Internet services. This paper contributes to that discussion.

There is a wide variation in what one may mean by "reliability." Being precise about what one is talking about is a necessary first step toward having an informed policy debate. In Section 2, we present several reasonable perspectives on what might be meant by "reliability" and review some of the large engineering and policy literature on the subject.

In Section 3, we discuss the different economic roles that notions of reliability may play. This helps provide a contextual basis for evaluating reliability and its relevance to public policy. We illustrate and explore some of the challenges that reliability poses for policymakers by considering several topical issues, including the convergence of fixed

[21] With support from the National Science Foundation, a number of research teams are working on new Internet architectures that are intended to address these issues (see http://www.nets-fia.net/, and for a list of analogous projects underway in Europe, see http://ec.europa.eu/information_society/activities/foi/research/fiaprojects/index_en.htm). The rise of Distributed Denial of Service (DDoS) attacks, identity theft, and malware pose serious threats to the open Internet, at the same time that a growing share of the population are becoming ever-more dependent on the Internet in their daily economic and social lives. The challenge is to improve trust and security mechanisms, while protecting privacy. At the same time, the growth of services like VoIP, peer-to-peer traffic, and streaming media are increasing the need for more than best effort packet transport support. Finally, the growing complexity of Internet interconnections and intradomain and interdomain traffic routing, coincident with the rise of content delivery networks (CDNs) fuels commercial interest in more granular, policy-based routing than is supported by the legacy BGP routing protocol.

[22] See FCC, *Notice of Inquiry*, In the Matter of Reliability and Continuity of Communications Networks, Including Broadband Technologies (PS Docket No. 11-60), Effects on Broadband Communications Networks of Damage or Failure of Network Equipment or Severe Overload (PS Docket No. 10-92), and Independent Panel Reviewing the Impact of Hurricane Katrina on Communications Networks (EB Docket No. 06-119), Before the Federal Communications Commission, Washington DC, April 7, 2011 (available at: http://transition.fcc.gov/Daily_Releases/Daily_Business/2011/db0407/FCC-11-55A1.pdf).

and mobile service regulation, network neutrality, and enforcement of service level guarantees (SLAs).

In the preceding sections, our perspective is abstract and quasi-theoretical. In Section 4, we turn to the more practical challenge of measuring and assessing the reliability of consumer broadband services and discuss current measurement efforts.

Section 5 concludes with thoughts about future directions for research. As with our earlier work on measurement of broadband speed, we conclude that reliability is a multifaceted notion for which there is no single best definition. The best measure will depend on the context and no single measure will adequately capture all that will be considered important in assessing broadband performance, and the quality and reliability of other services and components in the Internet ecosystem.

Reliability will be an important focal issue in markets for Internet interconnection and broadband service, in the formulation of end-user perceptions of service quality, and for Internet policymaking. Given its critical role, better data and better collective understanding of that data and the metrics used in its interpretation will be important for properly evaluating broadband reliability.

We view the growing concern with issues of Internet reliability as justified, and recognize the fear that the Internet may not be adequately robust or resilient in keeping with its new role as critical infrastructure. However, we urge policymakers to proceed cautiously. There is a pressing need for better data and metrics. It would be easy to proceed with overly ambitious attempts to specify reporting requirements or minimum reliability requirements that may prove overly burdensome, resulting in more harm than any problem they may seek to resolve. We remember numerous discussions in the 1980s and 1990s in which the Internet's lack of service quality guarantees (a byproduct, in part, of best effort packet delivery) and what was commonly referred to as "carrier-grade" reliability were decried as evidence of the Internet's inferiority relative to traditional telephone networks. The commercial success of the Internet, its robustness in adapting and scaling to meet new communications challenges, and its resiliency in supporting communications across very diverse environments and user needs should remind us to have a healthy skepticism of any new claims of impending doom for the Internet.

2. Defining broadband reliability

The Merriam-Webster dictionary defines "reliability" as the "extent to which an experiment, test, or measuring procedure yields the same results on repeated trials" and identifies "trustability" and "dependability" as synonyms.[23] This definition captures both the objective notion of a "reliable system" as one with consistent and predictable behavior; and in the identification of synonyms, the normative notion that one can trust (place confidence in) a reliable system not to fail, where a failure implicitly implies causing harm either to oneself or others.

[23] See http://www.merriam-webster.com/dictionary/reliability (accessed 8/7/2011).

Reliability engineering is its own sub discipline within the engineering sciences, and there is an extensive technical literature that focuses on the theory, practice, and empirical assessment of the reliability of engineered systems.[24] At the most general and abstract level, the concept of service "reliability" engages two distinct elements: 1) a service definition which includes implicitly or explicitly a notion of what it means to be “available” or “in service”;[25] and 2) a statement about predictability of that service definition. A common metric for reliability is “availability” which is the amount of time that a system is expected to be in-service. It is often expressed as a statistical time measure (e.g., Mean Time To Failure, MTTF) or the percent of time over some period that the system is available for service.

In telephony systems, we are accustomed to hearing about "five nines reliability" which is usually taken to mean that components of the telephone network such as the switching fabric are available "99.999%" of the time, or equivalently, experience down time of less than 5.26 minutes per year.[26] This is normally assumed to be a very high-level of reliability that is appropriate for a service as important as the telephone network. In days past, it was used to contrast the highly reliable telephone network that most consumers take for granted and the "best-effort" Internet or mobile telephone service for which less consistent performance was often expected or experienced.

However, the use of reliability metrics and their application to different networking environments is often extremely crude and misleading.[27] For example, while it might be true that central office switches were designed for better than 99.999% availability, this did not mean that an individual homeowner’s service was expected to meet such a high standard of availability. In fact, the Internet or mobile communication services are not inherently less reliable than legacy telephone networks, but rather differently reliable. Whether that matters depends on one’s perspective. For example, a mobile phone is “available” to make calls in more locations than a fixed telephone; and the Internet protocols support interoperability, facilitating higher “availability” of communications across heterogeneous platforms.

---

[24] As one very rough gauge, consider that the IEEE Xplore digital library (http://ieeexplore.ieee.org/) has 17,813 publications (out of a total of 2.97 million) with “Reliability” included in the title; and Amazon.com’s bookstore has 393 books in its Computer & Internet category with “Reliability” in the title (search conducted August 7, 2011).

[25] Peter Neuman commented that “a system without requirements cannot fail; it merely presents surprises” (see Neuman, 2004).

[26] 5.26 minutes/year = (365 days/yr * 24 hrs/day * 60 min/hour)*(1-0.99999). Analysts are not always precise as to whether the 99.999% availability applies to each switch individually, or to the fabric of switches that provide redundant (back-up) routing capabilities.

[27] For example, an industry white paper notes that the “five-nines benchmark is a common way to quantify ‘carrier class’ — one of the most widely used — and misused — adjectives in the telecom industry. With no real established measurement standards, vendors can claim products are carrier-class without having to detail under what operating conditions and circumstances the products were measured,” (see Tellabs, 2007).

When one considers complex systems, assessing reliability also becomes complex because the number of potential failure modes may expand, and may not even be enumerable *ex ante*.[28] Indeed, Internet architects interested in improving the reliability of the Internet in the face of arbitrary threats strive for what is referred to as "Byzantine" robustness.[29] The failure of some subset of components may result in a deterioration of service, but not a complete lack of availability. Defining what constitutes a sufficient level of service disruption to constitute a failure is not always easy. To loosely paraphrase the philosophical riddle, "if an outage occurs that no one notices, does it count?"

Usually, it is possible to increase reliability (reduce the probability of failure) by investing more resources (e.g., in redundant facilities and in hardening of components). When one introduces the notion of optimal investment, it becomes natural to consider the relative harm that is expected to result from an outage, which is likely to depend on the number of users impacted and the harm suffered per user. Thus, a loss of service to a city would generally be thought to constitute a more costly outage than the loss of service to a few subscribers; and a long duration service outage more costly than a short duration outage. Hence, in addition to MTTF, reliability engineers are also interested in the Mean Time To Restore (MTTR),[30] which when coupled with information about the relative costs of outages, recovery, and failure avoidance options provides the basic elements for a cost-benefit analysis of good reliability management. Such an analysis often implies that there is an optimal level of reliability, and improvements beyond that optimal level do not promise benefits that justify the additional costs required. Indeed, in days past when rate of return regulation was common, there was concern that telephone companies might seek to gold plate their networks by over-investing in capacity and network reliability.[31] In many environments, consumers might prefer to accept lower quality or less reliable service in return for new functionality (e.g., mobile phones over fixed line phones) or a lower price.

Finally, sometimes investments that are intended to improve the reliability of a system from one perspective may introduce new types of failure modes, and thereby threaten reliability from another. For example, the move to software control for networks,

---

[28] Even if it is possible to model the failure modes, computing network reliability may be NP-hard (see, Ball 1980).

[29] See, for example, Castro & Liskov (2002) or Avramopoulus et al. (2004) for papers discussing Byzantine Fault Tolerant (BFT) algorithms. BFT algorithms are robust to arbitrary failures of some number of sub-components in a system. The "Byzantine" refers to a classic agreement coordination problem in which a group of geographically distributed Byzantine generals are trying to organize a simultaneous attack and their communications may fail and some of the generals may be traitors, and thus, the goal may fail to be realized in multiple, arbitrary ways.

[30] The M and R in the abbreviation MTTR may have different meanings such as Minimum, Mean or Maximum; and Recovery, Repair, Respond, or Restore (see http://en.wikipedia.org/wiki/MTTR_(disambiguation), accessed 8/8/2011).

[31] See Averch & Johnson (1962), and for how this related to telecommunication network reliability, see Lehr (1995).

represented by the implementation of Signaling System 7 (SS7) in the public switched telecommunications network, increased the ability of the network to detect, route around, and recover from network faults (e.g., associated with a link or switch failure). While this helped increase the reliability and resiliency of portions of the network (component reliability), the chance that a software bug could crash SS7 and bring down a much larger part of the network was introduced. Centralized network control that makes it faster to respond in a coordinated way to network problems may introduce a single point of failure. Similarly, transitioning to faster, higher capacity routing and backbone transmission links means that a backhoe that cuts a fiber may impact a much higher volume of traffic than in an earlier era with lower capacity networks.

2.1. A review of the engineering literature

All of these issues arise in the context of thinking about the reliability of traditional telecommunication networks, and more recently, the Internet. Some of the key points emphasized include the need to define what constitutes an outage worthy of note, focusing on the duration, scope, and severity of an outage; and the need to collect and report data on user experiences (e.g., satisfaction surveys), indicators of network management practices (e.g., installation and repair intervals), and data on the network's ability to anticipate, withstand, and recover from a variety of threat scenarios. [32]

There is also a large literature focusing of a range of specific problems that may degrade the performance of the Internet, and hence, adversely impact the reliability of the Internet.[33] The focus of this literature is generally less on characterizing outages and their impacts, and more on performance measurement and identifying and proposing fixes for various problems. There are also a number of ongoing research projects that are engaged in collecting and analyzing performance-related data relevant to questions about broadband reliability. These include projects at CAIDA,[34] Georgia Institute of Technology,[35] and MIT.[36] Also noteworthy is the collaborative Measurement Lab

[32] For defining and measuring outages in the public switched telephone network (PSTN), see for example FCC (2009a & 2009b), Kaâniche & Kanoun (1996), Nojo & Watanabe (1993), Sanso et al. (1991), Snow et al. (2000), and Spragins et al. (1986). For reliability concepts applicable to services on the Internet, see for example Brewer (2001), Carlier et al. (1997), FCC (2010), Keralapura et al. (2004), and Tortorella (2005). Also, there is a wealth of materials available from the Network Reliability and Interoperability Council (www.nric.org).

[33] Such literature can be found, for example, from the Proceedings of the ACM SIGCOMM (http://www.sigcomm.org/), The Internet Measurement Conference (IMC, http://www.sigcomm.org/events/imc-conference/), and The Passive and Active Measurement Conference (PAM).

[34] See, for example, The Cooperative Association for Internet Analysis (CAIDA, http://www.caida.org/home/) for pointers to a wealth of interesting work. Specifically, see Huffaker et al. (2001).

[35] See http://projectbismark.net

[36] See http://mitas.csail.mit.edu

project[37] which provides an open, distributed server platform for researchers to deploy measurement tools, particularly ones that provide the public with information about their broadband connections.

A large body of work has defined measures for reliability of performance: Alves et al. (2002) and Prasad et al. (2003) discussed the complexities related to bandwidth estimation. Bovy et al. (2002) characterized end-to-end delay components; Dischinger et al. (2007) described their variation (i.e., jitter). Jiang and Dovrolis (2002) explained two methods for passive estimation of TCP round-trip time (RTT). Zhang et al. (2002) studied the characteristics and origins of Internet flow rates. Sommers et al. (2007) described how to accurately and efficiently monitor these measures.

Optimizing route selection and avoiding overloads is a key challenge in increasing reliability of IP connectivity. Labovitz et al. (2001) and Rexford et al. (2002) examined Border Gateway Protocol (BGP) routing stability. Markopoulou et al. (2004) analyzed routing updates in the Sprint backbone to classify and characterize failures affecting IP connectivity. Nguyen et al. (2009) suggested ways of minimizing probing cost for detecting interface failures. Wang et al. (2006) demonstrated how traffic engineering algorithms could improve reliability of IP connectivity in various scenarios Another approach is to deploy application-layer overlay networks where nodes across routing domains cooperatively relay traffic according to specific criteria (Andersen et al., 2001).

Many users are unable to diagnose their home network for faults before blaming other parties (Dong & Dulay, 2011). Maier et al. (2009) demonstrated that local architecture could be a significant bottleneck for the performance of a consumer's broadband connection. Additionally, user equipment is increasingly responsible for faults (Song et al., 2011).

Out of the possible "core" services of the Internet, Voice over IP (VoIP) has stood out as a popular service for reliability analysis: Cole and Rosenbluth (2001) identified delay, network packet loss and the decoder's de-jitter buffer packet loss as the relevant transport level measures to describe the quality of VoIP connections. Markopoulou et al. (2003) showed a large variety of VoIP behavior across backbone networks. Kushman et al. (2007) found BGP routing updates hindering VoIP usability as much as network congestion. The Domain Name System (DNS), essentially a core service of the Internet, is often suspected for faults, but Pang et al. (2004) found it to be highly reliable for most end-users.

2.2. Towards a practical Taxonomy for Broadband Service Metrics

The preceding discussion highlights the challenges of assessing the reliability of network services. No single definition or metric will suffice to capture the variety of ways in which the service may be degraded or suffer an "outage." Different definitions of what

[37] See http://www.measurementlab.net

constitutes an outage will impose different requirements on how outages might be detected, measured, and evaluated.

In the following three sub-sections, we propose three distinct types of reliability metrics that we expect to be useful for assessing broadband services. These metrics are:

1) Reliability of performance
2) Reliability of connectivity
3) Reliability of core services

2.2.1. Reliability of performance

By "reliability of performance," we mean that the quality-of-service (QoS) provided by broadband meets or exceeds some target level of performance over some specified timeframe. There are multiple QoS measures that one might be concerned with, included the speed (data rate), latency, jitter, or bit error rate. One might be interested in some composite (weighted index) of multiple of such measures.

While one might presume that a service which offered a more consistent experience of broadband speeds might be perceived to be more reliable, and hence preferable, this need not be the case. The Internet owes much of its robustness to the fact that it is highly adaptive to variable performance, and for traffic degrading gracefully in the face of congestion, especially traffic that is relatively insensitive to variable delays (e.g., email traffic more so than VoIP traffic). Because it is generally understood that the performance of the Internet may be instantaneously variable, applications have been designed both to take advantage of this, and to mask this variability from the user's perception.[38] Thus, on the one hand, to measure the behavior of the underlying network without taking into account the compensations provided by the application may give an insufficient understanding of the user-perceived reliability, but on the other hand, the user experience provided by each application may differ.

We demonstrated in our earlier paper (Bauer et al. 2010) that even measuring speed is difficult. We argued that while there is no single best measure that is appropriate for all circumstances, there are many ways to measure broadband speeds so as to generate misleading inferences.[39] Similar challenges and problems will assuredly arise as the measurements of performance reliability evolve.

---

[38] For example, Bauer et al. (2011) discuss how the "Powerboost" service offered by a number of cable broadband providers contributes to broadband performance.

[39] For example, as we explained in Bauer et al. (2010), a number of the speed measurement organizations used a single TCP session to measure broadband access speeds. This resulted in a significant share of the test results estimating rates that were effectively throttled by TCP's receive window rather than because they encountered any rate limiting congestion or traffic management associated with the broadband access service. Similarly, users who try and measure broadband speeds from a computer on a home WiFi network may be rate limited by their home network or the activity of other users in the home, rather than by the broadband access service. Measurements based on the FCC's new broadband measurement infrastructure uses specialized edge and network boxes with software provided by SamKnows. It seeks to measure more directly

2.2.2. Reliability of connectivity

An even more basic notion of broadband service reliability might focus on Internet connectivity.[40] If I sit down at my computer and find web pages loading too slowly or streaming video performance degraded, I may conclude that my broadband service is "available," but not providing a reliable (acceptable) level of performance. If a user cannot obtain a network IP address from their service provider[41] or connect to websites on the Internet, then most users would probably regard their service as not available (but whether the problem is with the broadband service or something in the user's home configuration remains to be determined). Often, one of the first things to try if you are experiencing Internet connectivity problems is to open a terminal window and "ping" some sampling of websites to see if it is possible to reach those websites. If you receive no return packets from your attempt to ping your Internet Service Provider's (ISP) Domain Name Service (DNS) service, or popular websites like Google, Yahoo, or Facebook, you might conclude that you are experiencing a broadband service outage.[42]

Because many popular websites are hosted on multiple servers, not all users may experience the same problems at the same time, and not all users may be interested in reaching the same sites at the same time. Any particular server may be down for either a brief or long time. The failure of a single server or even subset of servers need not constitute an outage of broadband service. Or, pings to target servers may be returned but only with unacceptable delays. There is some inherent arbitrariness to the definition of what constitutes a sufficient loss of connectivity to be classified as an outage.

As with the earlier discussion of performance metrics, whether one detects an outage may depend on whether one is actively monitoring the connection. One might take the view "no harm, no foul," or that one is due the service contracted for whether one is using it or not (and, of course, that depends on one's interpretation of what the service contract entitles one to). Consumer broadband service commitments by service providers are notoriously vague with respect to their performance guarantees. This lack of specificity

---

the performance of broadband access services. It has reported significantly higher data rates being achieved than were suggested by the earlier measurements (see, FCC 2011).

[40] Sundaresan et al. (2011) propose defining "availability as the faction of the time that home users have IP connectivity to their access ISP." This definition may suggest that connectivity is limited to packets traversing the access link, but that would be a rather limited notion. For many, connectivity would imply being able to get packets to/from off-net locations, but determining what fraction of which targets need to be unavailable and for how long to qualify as an outage is complex.

[41] For example the Dynamic Host Configuration Protocol (DHCP) may not be working between the customer and the service provider.

[42] One must note that the firewall infrastructure of some web sites obstructs ping efforts, leading to the possible false conclusion that a web site is unavailable, even though it may be available. Details such as this confound simple attempts to design tests of availability.

need not be perceived as a problem since providing firm minimum performance commitments may only be feasible at a much higher cost.[43]

2.2.3. Reliability of core services

A third notion of reliability that seems likely to be of interest may focus on the availability of certain core services. A broadband user, application provider, or ISP may be interested in whether certain core applications or functionality of the Internet is operating properly. The potential list of target services or functionalities and the metrics for assessing whether they are operating properly is long. Several of the obvious candidates include services like DNS, email, and the Web. Each of these services may fail in multiple ways (e.g., no connectivity, diminished performance, localized or system-wide loss of availability, et cetera).

As with the detection of QoS, performance or connectivity outages, a natural way to measure the availability of such services would be to actively invoke the critical functionality from a number of source and destination nodes and report the percentage of times when the appropriate (expected) response is not received. When the number of test results exceeds some threshold, an outage is recorded. The duration and severity (scope, number of users) of the outage should also be recorded.

With each of these metrics, active measurements may be supplemented by passive measurement or alerts. For example, if Google detects that its servers are down, then it might announce or record that event, even in the absence of active measurements by broadband service nodes.[44] Other critical sub-components for broadband service could likewise record events when they go out of service and thereby contribute to the detection of broadband service outages. However, since many aspects of the Internet depend on redundancy to enhance reliability, the failure of a component does not necessarily imply a user-perceived availability failure. While the tracking of component failures is important (if only to assist in the assessment of how much redundancy is needed to reach an overall level of reliability), it may be difficult to reason from measurements of component failure to assertions about user-perceived failures, because these will depend, among other things, on the design of the redundancy scheme itself. So direct measurement of the overall service may be the best way to understand its achieved reliability.

Of course, an appropriate measurement framework will have to balance the costs of improved detection and reliability assessment with the potential direct and indirect economic cost/benefit of such improved reliability assessment. In the next section, we

---

[43] In Bauer et al. (2009b), we discuss how the Internet's congestion management for shared infrastructure contributes to overall efficiency. The fact that service may degrade occasionally because of congestion should not be interpreted as bad engineering.

[44] Like many other providers, we can be sure Google certainly tracks and records such information, but does not generally announce this data publicly. However, it published a relevant study a few years ago (see Pinheiro et al., 2007).

consider some of the economic and policy challenges posed by better measurements, standards, and reporting of reliability.

3. Economic and policy challenges

Reliability measures and/or standards may serve multiple economic purposes. In the following, we discuss several of the ways in which they may play a role, and point to some of the relevant economics literature for exploring these roles more deeply.

First, the realization of an outage is a payoff-relevant event. The actual experience of an outage may result in real costs and transfers being incurred, including insurance claims, incremental costs to restore service, or otherwise respond to any outages. These costs, their relative magnitudes, and their distribution produce direct and indirect economic consequences that will reverberate through markets and society. For example, cyber-attacks that bring down eCommerce websites can result in significant economic harm as measured by the stock price of affected firms.[45] These losses arise from lost sales, lost brand image, and costs associated with recovering from the disruptions. There are few good estimates of the total costs of various types of Internet outages, and estimates vary widely.[46]

Second, the prospect of outages may induce strategic behavior among stakeholders. The expectation of an outage implies a contingent payoff-relevant event, and in anticipation of those, stakeholders may be induced to invest in insurance or other preventative measures. The expectations of outages and their payoff implications may influence the selection of technologies, product-market strategies, and other aspects of business decision-making. These expectations may be reflected in contract terms (e.g., in Service Level Agreements or SLAs, which may define minimum performance, penalties, or other remedies) and market norms.

---

[45] For example, Garg et al. (2003) studies 22 events from 1996-2002 and found the average event resulted in a loss of $918 million in shareholder value, reflecting a drop in price of 5.6% over three days. Cashell et al. (2004) found stock-price impacts of 1-5%, translating into shareholder losses of $50-$200 million. They report anecdotal evidence from security firms estimating global losses from all overt attacks of $226 billion. A report by Anthony et al. (2006) looked at the impact on share prices of website outages and found adverse impacts of 3-4% associated with the days around the event.

[46] For example, Dynes et al. (2006) attempted to estimate the macroeconomic impact of several outage scenarios that might disrupt the Internet communication capabilities of U.S. firms. They estimated a wide-range of potential impacts. For example, for a Gulf oil refiner representing 10% of U.S. supply, they estimated an impact of $0 for the loss of basic Internet, but a loss of $405 million associated with a 10-day outage of the SCADA safety network. Wilson (2003) estimated the impact of enterprise network downtime for a series of case studies across multiple industries to range from $2k to $97k per hour-long outage. And, the Ponemon Institute (2011) estimated the cost of data center outages at $505k per incident.

Expectations about reliability may be based on empirical observations (past system performance) or theory (provable statements about system performance or models). Those expectations may be common or asymmetric (different parties may have different information or reach differing conclusions based on the same information). The expectations may be subject to varying degrees of uncertainty, and the information may be complete or incomplete (participants may not even know how to parameterize the uncertainty). Fully understanding how expectations may shape behavior is complicated by the fact that we may never actually observe the harms being realized by the anticipated outage.[47]

Third, markets for reliability may suffer from a variety of imperfections. For example, we may expect buyers to base their selection of service provider on reliability claims. If consumers cannot distinguish high-reliability from low-reliability services there may be a "Lemons" problem.[48] Alternatively, "reliability" claims or contracts may induce strategic behavior to address potential adverse selection or moral hazard problems.[49] Imagine if ISPs sought to make credible commitments about their reliability (e.g., to address a potential Lemons problem) by committing to significant rebates or penalties if outages occur.[50] Such promises might attract the reliability equivalent of "patent trolls," that is, users with a higher-than-expected probability to be able to detect or experience outages that would trigger the compensating payoffs (adverse selection).

From a public policy perspective, we may assume that outages are (by definition) bad (something to be avoided), but we need to consider how bad (costs of an outage) in order to understand what efforts are warranted to avoid outages, and how to assign responsibility for undertaking the requisite investments. If the markets for broadband services are deemed to be effectively competitive, then we may be able to rely on market forces to induce an efficient reliability equilibrium. However, as noted above, there are a number of potential market failures that might lead us to doubt the market's ability to induce the socially optimal reliability behavior. Unfortunately, recognizing the potential for a problem is not the same as demonstrating that the problem poses a significant risk, or if it does, that there is a viable policy intervention that will improve outcomes.

---

[47] That is, market behavior is impacted (a real effect), while the motivation for that behavior is not observed (it is because of the potential for some event to happen that does not actually happen). Thus, I continue to invest in fire insurance for my home, even though my house is unlikely to ever burn down.

[48] See Akerlof (1970).

[49] See Holmstrom (1979) for a discussion of moral hazard and Abbring et al. (2003) for a discussion of the difference between moral hazard and adverse selection. Both are problems that arise as a consequence of contracts in the presence of asymmetric information. Moral hazard relates to potential for an agent to change his behavior after a contract is selected in a way that harms the principal; whereas adverse selection relates to the choice of contract by agents in such a way that it results in harm to the principal.

[50] Apparently, some ISPs have done precisely that: see http://www.hkbn.net/2010/eng/en_service1_1c2.html.

Varian (2001) illustrates some of the complexity of that may arise if one considers the issue of free riding on system reliability under a number of obvious contexts. He considers three prototypical situations: 1) Total effort: the reliability of the system depends on the effort of all; 2) Weakest link: the reliability of the system depends on the minimum effort; and 3) Best shot: the reliability depends on the maximum effort. In his stylized model, he shows how the level of effort exerted typically falls short of the socially optimal level of effort, but the behavior and outcomes are materially different in each situation; hence, we should expect that the appropriate remedy would depend critically on the reliability model that is applicable. There is an obvious interplay here among different sorts of "remedies." Varian's different situations can be mapped to specific mechanism design (e.g., the design of redundancy schemes), architecturally defined tradeoffs (e.g., the balance of responsibility between the transport layer and the application), or to operational obligations imposed on different actors.

Moreover, if the service becomes sufficiently highly reliable so the probability of observing an outage becomes very small, then providers of sub-components may seek to free ride on the reliability investments by other component providers, thereby undermining the mechanisms on which system reliability depends (a form of moral hazard). This is potentially a hard problem to address, and points to one reason why concerns over the reliability of the core components of the Internet (e.g., the backbone transport and routing infrastructure, for which it is relevant to talk about "carrier grade five-nines or better reliability") and the reliability of consumer broadband services need to be clearly distinguished.

As noted earlier, the potential harm from an outage that significantly impaired the Internet's core backbone would have a much higher total cost than the disruption of service to a subset of consumer broadband customers. The former would be expected to impact the experiences of a much larger number of consumers, and in an effort to avoid what might be viewed as potentially catastrophic failures, we would expect a much higher level of preventative investment to protect against such failures. Indeed, we might view significant outages of a large portion of the Internet or its core operating elements as potentially so harmful that it is socially desirable to limit the probability of any such event to be extremely rare. This takes us into the realm of catastrophe economics, which raises a host of public policy difficulties. Kahneman and Tversky (1979) have suggested how we might tend to overweight the probability of rare events (e.g., risk of a shark attack or a lightning strike) in decision-making, thereby overinvesting in protection against such rare events. On the other hand, if the potential harms are sufficiently bad, it may be more likely that we would under-invest in protection and over-invest in responses.[51] Generically, as individuals and collectively, we are poorly prepared to make socially optimal decisions when it comes to anticipating so-called "Black Swan" events.[52]

---

[51] See Posner (2004) and Noll (1996).

[52] In considering the threat of an asteroid hitting the earth, we need to balance the small probability of such an event occurring with the very large negative payoff if it did occur (potential to annihilate all life on the planet). In assessing this risk, we need to figure out how much effort we should expend today in designing counter-measures (e.g., lasers that might divert such an asteroid) relative to the risk we seek to avoid. See Chichilinsky & Eisenberger (2010) for a

When it comes to protecting us against catastrophic network failures, we cannot rely on ex post enforcement and the monitoring of outages to provide adequate incentives for network reliability. In such cases, we need to rely more on our ability to audit the architecture, processes, and reliability-investment behavior of service providers to ensure that adequate provisions for reliability have been made.[53]

With respect to the reliability of consumer broadband services, the problem is different. In that context, we may expect to be able to monitor the actual experience of individual users, collecting measurements over large samples that are amenable to statistical characterization, just as today we measure and report on the data rates experienced by users. While the collection and analysis of such large sample data is inherently complex (as we discuss in the next section), the fact that we can formulate solid empirical data on which to measure service reliability renders this sort of performance monitoring a fundamentally different business and policy challenge than is posed by the challenge of protecting against catastrophic failures. In the case of monitoring the performance of consumer broadband, the predictable statistical regularity of observing faults means that we can rely on competitive markets to send the appropriate signals to providers and users to push them toward the efficient equilibrium.[54]

In such markets, the challenge of regulating reliability may be viewed as equivalent to the challenge of regulating service quality. There is a large literature on the economics of service quality regulation.[55] Although most economists generally believe that market competition will enhance incentives to invest in quality, this presumes that quality differences are observable by consumers (which is not easy as the discussion in the next section demonstrates) and that markets are not imperfectly competitive.[56] Ultimately, the economics literature suggests that that the relationship between the optimal level of quality and market structure is inherently ambiguous.[57]

---

discussion of this problem. Or, for a more popular discussion of challenges "Black Swans" pose for descision-making, see Taleb (2007).

[53] A key component of the Nebula research project (see note 21 *supra*) is to design a highly robust core routing architecture that can be distributed across multiple routers and afford better than five-nines reliability.

[54] This assumes that broadband services are competitive, but the relative competitiveness of broadband services is a separate question we wish to abstract from in this paper.

[55] See, for example, Sappington (2005).

[56] If broadband services are not competitive, then there may be strategic motivations for a provider with market power to seek to manipulate reliability or market perceptions of reliability (e.g., by manipulating the quality of interconnections) to raise rivals costs. This raises a host of additional problems that we ignore here.

[57] For example, Spence (1975) shows that a monopolist may over- or under-invest in quality relative to the social optimum, depending on the willingness-to-pay of the marginal consumer relative to the inframarginal consumer. Auriol (1998) explains how competition may result in lower quality if quality is not verifiable. Sappington (2005) explains how the intensification of price pressure associated with competition may reduce ex ante incentives to invest in quality, and

In the presence of market imperfections resulting in quality of service issues, policymakers have a number of options, but all of these are second-best solutions (in part, because of the economic costs inherent in regulation). Generically, because quality of service is hard for regulators to observe or verify (relative, for example, to other service characteristics like price or quantity of output) and because it is inherently multidimensional, it poses a greater challenge for regulators. Designing optimal incentives for multidimensional goals is problematic because too much focus on one dimension may result in excessive investment in that dimension, and too little in other dimensions.[58] For example, regulators may establish minimum quality of service requirements (MQR) and impose penalties or offer financial incentives for meeting those standards. While such strategies may help, they may also have perverse impacts. For example, by increasing the quality of services offered by lower quality producers, MQRs may reduce incentives of higher quality producers to invest in quality (Sappington, 2005). Or, the threat of regulatory penalties may result in excessive investment in quality of service because firms fear not only the financial penalties from being found at fault, but potentially even more, the loss of brand image.[59]

Also, quantity or price regulations also may have perverse effects for incentives to provide quality. The transition from rate of return to price cap regulation of telecommunications providers was motivated, in part, by the desire to provide operators with enhanced incentives to improve economic efficiency, but this regulatory transition had mixed impacts on service quality.[60]

In the next three sub-sections, we briefly consider how reliability considerations may impact several current policy issues.

3.1. Role of reliability in the Network Neutrality Debate

One interpretation of Network Neutrality (NN) is that it implies restrictions on the ability of ISPs to actively manage traffic. One response to efforts by ISPs to engage in discriminatory traffic management is to try and hide information about the traffic being carried, to make it difficult for ISPs to selectively manage that traffic.[61] Moreover, efforts to hide traffic characteristics (through encryption, random port hopping, or onion routing)

---

enhance ex ante incentives to try and soften price competition by offering differential quality – both of which may reduce total welfare.

[58] See Holmstrom & Milgrom (1991).

[59] See Price et al. (2002) for an empirical study of quality of service regulation of utilities that finds such an effect.

[60] Uri (2003) found a significant decline in service quality in the US following the transition to incentive regulation, based on a review of data from 1991 to 2000. Ai et al. (2004) found mixed results, with service quality for some measures improving and for others declining.

[61] See Lehr et al. (2007).

that may be viewed as user-employed defenses against non-neutral treatment may increase the complexity of detecting and recovering from certain outages. The result of either protecting NN by regulation (rules that limit ability of ISPs to actively manage or discriminate among traffic) or seeking to enforce it by end-user actions (e.g., seeking to obscure traffic features that might be used to actively manage the traffic) may be, perversely, to harm broadband reliability. However, we also recognize the risk that an ISP interested in violating principals of network neutrality might seek to hide behind claims that their behavior is motivated by reliability or security concerns. Reliability concerns should be seen as a part of a balanced NN policy, not as an excuse for discriminatory traffic management.

For example, active management is one of the key ways that ISPs seek to anticipate, detect, and recover from outages. For example, one view of what Comcast was trying to do when it interrupted selected BitTorrent sessions (the action that caused Comcast to run afoul of the FCC) was that it was trying to protect the performance of VoIP services that might otherwise have been crowded out by BitTorrent sessions. Similarly, efforts by an ISP to block malware or defend against Distributed Denial of Service attacks may also be viewed as "reliability-enhancing", but overly restrictive interpretation of NN rules might interfere with such responses.

The interplay between network neutrality and reliability is exacerbated by the fact, noted above, that different applications have different requirements for the reliability of the underlying service, and take different steps to mask variation in the underlying service from the user-visible experience. To the extent that different applications benefit from (and can exploit) different sorts of behavior in the underlying network, it is not in the best interest of the user that all packets be treated the same by the network.

In the original design of quality of service (QoS) mechanisms for the Internet, the mechanisms to request a particular service from the network (the Type of Service bits in the IP header) are distinct from any indication of what application is running (e.g. the well-known port number in the TCP header). This distinction was intended to allow an application to select the service it required without necessarily revealing which application it was. As the public Internet currently operates, the ISPs do not provide service variation based on the Type of Service bits, but as the conversation about reliability (and NN) become more nuanced, their motivations may change.

Video over the Internet provides an interesting example of the interplay between application design, network design, and the user perception of reliability. First, when a user watches video over the Internet, they become highly sensitive to any user-visible "glitch" that occurs during their watching experience. Hence, the user perception of potential broadband reliability issues may increase if they shift a greater share of their TV watching to over broadband (and thus, have occasion to monitor performance more often). Moreover, because different schemes for delivery of video use different techniques to deal with variation in speed and latency (as well as transient outages), different schemes for delivery of video may reveal different aspects of underlying network behavior. There is an obvious tension between application designer and network designer as to "who should do the work" of making the user experience acceptable, a tension with technical, economic, and regulatory components.

3.2. Mobile v. Fixed reliability

Certainly one of the most important forces impacting the Internet ecosystem is the growth of mobile broadband and the convergence of mobile and fixed services. This raises important questions for broadband policies ranging from universal service to network neutrality to interconnection.

Today, there are significant differences in how mobile and fixed telecommunication services are regulated. As mobile broadband services continue to improve, we will need to ask ourselves whether such asymmetric treatment should be sustained.

Lehr and Chapin (2010) explain why we should expect to see significant differences persist between mobile and fixed broadband services, and identify fundamental differences in the reliability of the two services as a key distinguishing feature for the foreseeable future. Because of the fundamental technical differences between mobile (wireless) and fixed networks, they are vulnerable and resilient in different ways. And, as noted earlier, they are available in different ways. In disaster scenarios, there is the chance that one or the other network might continue to operate while the other fails, and so the overall reliability of communications is likely enhanced by having both.

To the extent we move toward thinking about fixed and mobile broadband services as substitutes that should be subject to symmetric regulation, there will be obvious challenges to designing a set of metrics and strategies for measurement and reporting that will be comparable across mobile and fixed services. For example, the channelization of frequencies in wired services means the local variation in performance is more limited than what is experienced in the wireless frequency domain (where issues the mobility of end nodes, multipath, out-of-band interference, line-of-sight, or rain fade) mean that performance assessment is inherently more dynamic.[62]

More generally, we may expect that end-user concerns about what constitutes appropriate reliability for mobile and fixed broadband services to differ. For example, users may care more about being able to connect to the best available wireless network connection wherever they are, rather than to the network of a particular provider. With advances in

---

[62] That is, the reliability of wireless transmission is highly dependent on the location of the transmitters and receivers and the conditions in the local RF environment, in ways that are inherently more difficult to isolate than is the case with wired communications (see Lehr & Chapin 2010 for further discussion). This implies that if you wanted to evaluate the reliability of a wireless service you may need many more sensors at many more locations than you would in a wired network to adequately model expected performance (i.e., the reliability assessment should include not just the locations where the transmitters and receivers are currently, but where they are or are expected to be over time).

wireless technology,[63] the potential for users to switch among wireless networks and platforms has the potential to offer new ways of provisioning Internet services and managing the user's reliability experience (e.g., end-user controlled infrastructure).[64] It will be increasingly important to be able to think about the reliability of service afforded by the mosaic of overlapping wireless infrastructures, each of which is likely to be robust in different, but potentially complementary ways.

3.3. Interconnection SLAs

The markets for Internet interconnection are growing increasingly complex. The convergence of mobile and fixed services discussed above is one reason, but there are a number of others (see Clark et al, 2011). ISP interconnection agreements today are not regulated (whereas legacy telephone interconnection agreements were heavily regulated). When the markets for interconnection were relatively simple and bifurcated into fairly standardized transit and settlement-free peering, it was relatively easy to figure out what was going on. Today, a growing share of ISP traffic is being routed on the basis of interconnection agreements that are increasingly diverse in their terms and conditions. Moreover, ISPs are interested in extending their ability to offer managed services like Virtual Private Networks (VPNs) across provider boundaries. This requires new types of interprovider Service Level Agreements (SLAs).

For example, the design of the large content delivery networks (CDNs) that are used today to deliver commercial content (for example real-time entertainment) to consumers illustrate the complex interplay of different system aspects in achieving overall reliability. A large access ISP today may have many 10's of interconnection agreements, and many 100's of actual interconnection points. CDNs implement sophisticated algorithms to select the source and interconnection path for different content downloads, and by changing the chosen source, they may be able to avoid "or compensate for" a failure as significant as the loss of a major interconnection point in the Internet.

Reliability concerns are also important in the product descriptions of broadband service offerings and the service agreements negotiated with large enterprise customers. Collectively, we may refer to all of these agreements as SLAs.[65] These SLAs already include a number of terms and conditions related to reliability. For, example service descriptions often note that availability commitments provide allowances for maintenance downtime; providers may tout the availability of customer support services

---

[63] Motivated in part by the scarcity of radio frequency spectrum, trends in wireless technology (including cognitive radios, smart antenna systems, and dynamic spectrum management) are enabling the unbundling of applications and wireless technologies. This unbundling makes it possible for applications to move across wireless network platforms (and frequency bands), as for example, when users switch between using WiFi and 3G networks for broadband wireless Internet access.

[64] Enabling such highly mobile architectures is a goal of the MobilityFirst research project (see note 21 *supra*).

[65] See Lehr & McKnight (2000) for an early discussion of consumer-grade SLAs for broadband.

such as 24-hour hot lines or offer special repair commitments that are related to outage detection and recovery; and, SLAs may already specify availability commitments and remedies in the event those commitments are not meant. With the growth of cloud services, we may expect service descriptions and negotiated contracts to address performance issues such as accessibility to data centers, diverse physical routing, and stricter latency bounds. If we are correct in our prediction that reliability concerns will become increasingly salient for mass market and enterprise customers alike, we can expect the terms and conditions related to reliability to play a larger role in broadband service descriptions and marketing, and in the SLAs negotiated with large enterprise customers and between ISPs.

Although SLAs already include a range of reliability-related terms and conditions, they are not standardized.[66] This contributes to the complexity and heterogeneity of interconnection agreements and broadband product offerings, which poses a challenge to policymakers.[67] The interpretation of these terms and their application is not always clear, potentially even to the parties to the agreement. Most consumers do not read the fine print of broadband service descriptions, and SLAs for large enterprise customers or between ISPs are complex legal documents that are open to variable interpretation.[68]

If arriving at suitable contracts to address reliability concerns proves intractable, then the perceived QoS benefits of keeping traffic on-net (relative to off-net) will persist; whereas if efficient intercarrier reliability assurance contracts are negotiated, this might open up new market opportunities for Internet services.

4. Current data, metrics, and tools for assessing broadband reliability

In this section we look at currently available data and metrics on broadband reliability for consumer broadband. In general, we are at the beginning stages of measuring broadband reliability and the user experience of it. Over the coming years, we expect that the data,

---

[66] One study of current peering policies identified a number of reliability-related terms, including 24/7 network operations center, commitments to prompt response to outages, geographic diversity of connection points, notification of outages, and access to tools to provide internal visibility. However, as we note, there was great variation in requirements across agreements. See http://drpeering.net/white-papers/Peering-Policies/A-Study-of-28-Peering-Policies.html.

[67] For example, because higher availability commitments, including special provisions like support for diverse routing or maximum time to restore commitments, are costly, the specific terms and conditions may be expected to vary in price as well as options across SLAs. This diversity may compound difficulties in assessing overall system reliability. While this complexity and diversity may make policymakers' jobs more difficult and may complicate reliability planning, it does not mean that the complexity and diversity are bad – we are only pointing out some of the challenges that arise as a consequence of the Internet's success.

[68] Because reliability terms and conditions address rare events, the provisions may not be invoked often and the processes to be followed in the event of an outage may not be well understood or easily verified.

metrics, and tools for understanding and quantifying broadband reliability will evolve considerably.

Currently, information on broadband reliability comes from three main sources: 1) survey results of broadband users; 2) tools for ad-hoc measurements, self-reporting and online discussion of reliability; and 3) dedicated broadband measurement infrastructures such as those employed by FCC and Ofcom.[69] Each of these sources provides useful, if imperfect, insight into the reliability of broadband services.

4.1. Survey results of broadband users

Survey results of broadband users indicate the majority of users are satisfied with their broadband service. The following Figures 1-3 represents survey results for broadband users in the UK.[70] While broadband users in the United States or other countries could potentially have very different broadband experiences, we have not seen evidence that suggests broadband users elsewhere are less satisfied with their broadband connections.[71] Most users (at least 80%) in the UK survey reported being "very satisfied" or "fairly satisfied" with their broadband connectivity (see Figure 1).

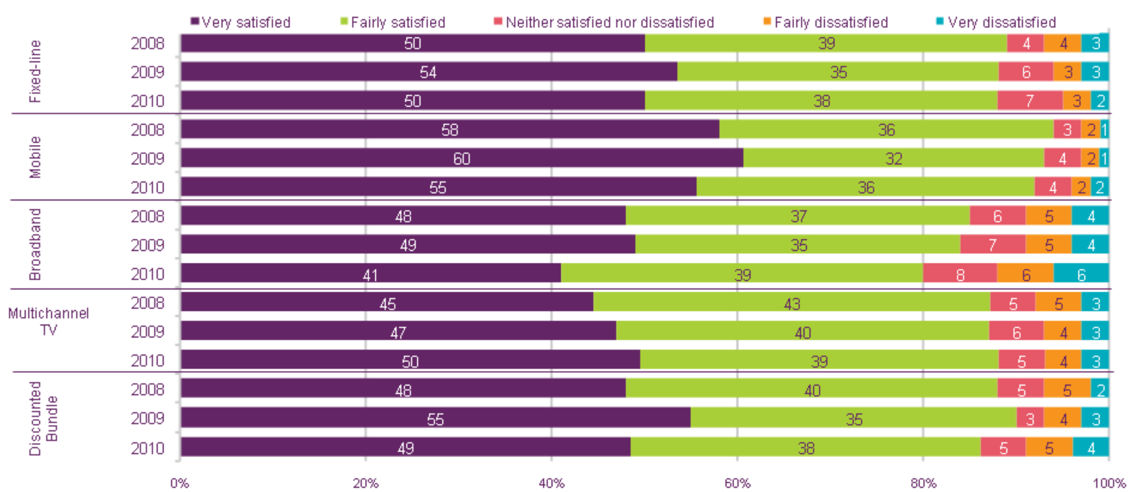

QL5/QM5/QI5/QT5/QB5 - How satisfied are you with the overall service provided by [SERVICE PROVIDER]?
Source: Ofcom decision making survey carried out by Saville Rossiter-Base in July to August 2008, 2009 and 2010
Base: All adults aged 16+ who are the decision maker and express an opinion on fixed line (2008, 771) (2009, 660) (2010, 617), mobile (2008, 1251) (2009, 1205) (2010, 1189), broadband** (2008, 302) (2009, 275) (2010, 222), multichannel TV (2008, 797) (2009, 760) (2010, 768), discounted bundled services (2008, 534) (2009, 631) (2010, 566). *Don't know responses have been excluded from the base. **Base for broadband in 2010 represents those with fixed broadband rather than fixed or mobile broadband as in previous years. Too few interviews were conducted with those with mobile broadband to report these separately. Trend data may be affected by this change.

Figure 1: Satisfaction with overall services from communications supplier, over time (Source: Ofcom, 2010)

---

[69] Both the FCC and Ofcom partnered with Samknows (http://samknows.com) to build, deploy, and run large-scale measurements systems.

[70] See Ofcom (2010).

[71] While likely out of date, the last survey in the United States that we found which asked about broadband reliability was in 2002. In that year, "90% of broadband users rate the quality of their Internet connection at home as "excellent" or "good", with 40% saying it is "excellent" (see http://www.pewinternet.org/Reports/2002/The-Broadband-Difference-How-online-behavior-changes-with-highspeed-Internet-connections/Main-Report/Part-3.aspx).

When asked specifically about their level of satisfaction with the reliability of their broadband connection, 8% overall reported being dissatisfied with their broadband reliability (see Figure 2). Younger groups are more likely to be dissatisfied than older groups. While the data is insufficient to know for sure, we speculate that younger generations are more likely to be dissatisfied since they tend to use the network more intensely both in terms of more hours of usage (and hence more likely to notice temporary outages) and in terms of applications (streaming audio and video, gaming, and other services where disruption in service is more noticeable).

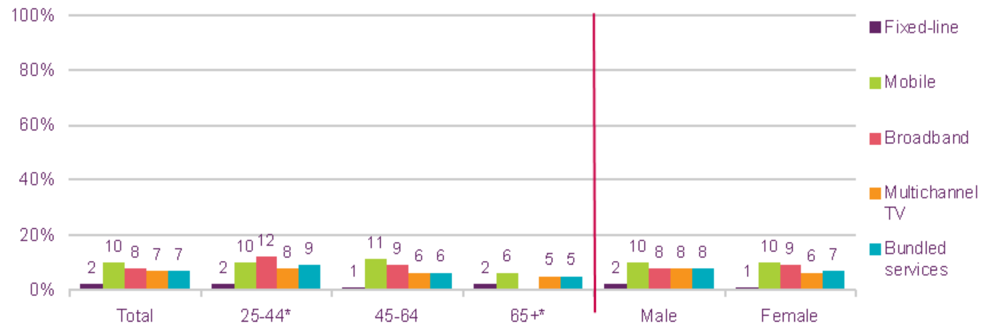

QL5B/QI5B/QT5B/QB5B - How satisfied are you with the reliability of your service from [SERVICE PROVIDER]?/ QM5B – How satisfied are you with the reception or ease of accessing the [SERVICE PROVIDER] network?
Source: Ofcom decision making survey carried out by Saville Rossiter-Base in July to August 2010
Base: All adults aged 16+ who are the decision maker and express an opinion on fixed line** (2010, 620), mobile** (2010, 1195), broadband** (2010, 220), multichannel TV** (2010, 765), bundled services (2010, 567). *Caution: Low base. Base for 16-24 years olds and 75+ too small to analyse. Base for 65+ for broadband too small to analyse. Don't know responses have been excluded from the base. **NB Base amended in 2010 to exclude those who receive this service along with another service from the same supplier without receiving a discount. Base for broadband in 2010 represents those with fixed broadband rather than fixed or mobile broadband as in previous years. Too few interviews were conducted with those with mobile broadband to report these separately. Trend data may be affected by these changes.

Figure 2: Dissatisfaction with reliability of service, by age and gender (Source: Ofcom, 2010)

Surveys of broadband users indicate that reliability is one of the main sources of dissatisfaction with broadband service. These surveys tend not to define reliability in the questions asked to panelists or in the summary reports. Hence survey respondents and readers of the resulting reports likely adopt different interpretations of reliability ranging from "failure to perform as expected" to "total failure of network connectivity." When users were asked specifically about their experiences with different types of problems while accessing the network, they noted problems associated with both speed and connection problems (see Figure 3).

Are there any problems using (X) to access the Internet?	Fixed line (WiFi)	Fixed line (no WiFi)	Laptop via dongle/ USB at home	Laptop via dongle/ USB out of home	Mobile phone at home	Mobile phone out of home
Speed of connection is too slow	18%	24%	22%	34%	27%	22%
The internet connection is unreliable	7%	7%	12%	13%	12%	12%
Poor coverage – it's hard to get a connection	3%	2%	7%	9%	6%	12%

Source : Q.43 Are there any problems using (main device at Q.14) to access the Internet?
Base: 2,001 All respondents

Figure 3: Main problems experienced when accessing the Internet (Source: Ofcom, 2010)

One caveat for all discussions of reliability based upon user-reported issues is that many users are likely attributing any failure of a network dependent activity to their broadband provider. Many problems users detect are likely not the responsibility of the broadband provider. Problems can arise along any part of the network that is critical to the user activity and experience. In particular, the computers, access points, wires, and other components inside a user's home are known to be a significant source of failures.[72] The "server side" and other networks also can experience failures. There is a lack of good data on the relative importance of different problems, for example, how often are broadband subscriber problems due to problems with the user's home configuration or with applications or content being accessed versus what share of subscriber problems are attributable to problems in the access provider's or some other ISP's network. To our knowledge, what fraction of failures is attributable to broadband providers has not been quantified.

4.2. Tools for user-generated measurements and online discussion of reliability

While many of the tools that are employed to monitor enterprise or service IT infrastructure (such as *Nagois* and *Cacti*[73]) could be applied to monitor residential broadband connections, few users bother monitoring their home broadband connections.[74] Instead of systematically monitoring, these sophisticated users report

[72] Indeed, one of the authors of this paper has experienced a number of failures that he initially suspected were the fault of his broadband provider that ultimately were traced to sources such as faulty in-house wiring and connectors and faulty wireless access points. See also Dong & Dulay (2011).

[73] See http://en.wikipedia.org/wiki/Comparison_of_network_monitoring_systems (accessed 08/10/2011).

[74] We base this claim upon informal surveys of colleagues in the networking research community.

investigating faults and failures on an ad hoc basis as they arise. They self-diagnosis their own broadband problems using tools such as *traceroute* (to trace paths), *ping* (to test basic connectivity), *dig* (to test and debug DNS problems), and speed tests tools (to test throughput).[75] Based upon these investigations, they build up informal mental models of the reliability and causes of fault on their broadband connections.

There are relatively few tools for casual users to assess the long-term reliability of their broadband connection. There are a services, such as those offered by dslreports.com,[76] and tools, such as Smokeping,[77] which will monitor a broadband connection and produce email and web based reports on the status of a broadband connection (see Figure 4 for one part of such a report). Both of these tools operate by periodically pinging the broadband users IP address.[78]

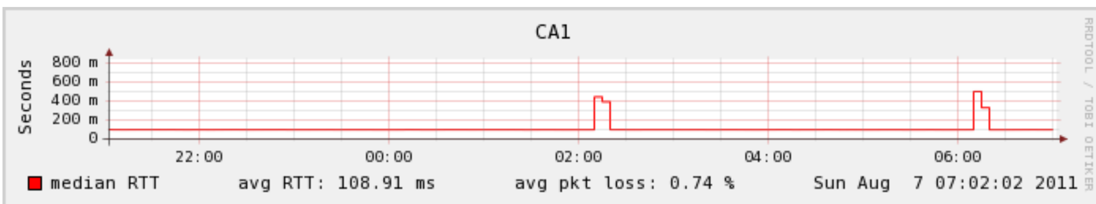

Figure 4: 24 hour ping test from dslreports.com Smokeping service. Chart shows the ping time across 24 hours for a single broadband connection in MA measured from a server in CA.[79]

All of the above tools and services provide a view about the reliability of a single user's own connection. However, this does not tell the user about the more general reliability of connections (performance or failure wise) of other users or other networks.

To get a wider perspective on reliability one can turn to various online forums that discuss broadband providers.[80] These forums tend to discuss a wide range of topics, but discussion of network problems is common. Discussion of current outages increasingly takes place on various real-time communication platforms like Twitter. Searching Twitter for a broadband provider's name plus "outage" or "problem" or "down" tends to return numerous hits particularly when an outage or problem affects a large number of users.

---

[75] For further discussion of these tools, see Bauer et al. (2010).

[76] On August 8, 2011, the dslreports.com service is monitoring 1560 static IPs and 946 dynamic IPs (http://www.dslreports.com/schedule). The dslreports.com also runs another service based upon the Smokeping tool (http://www.dslreports.com/smokeping).

[77] Smokeping (http://oss.oetiker.ch/smokeping/) is an open source tool that users can employ to run one type of broadband connectivity test based upon period pings.

[78] The Smokeping tool sends 20 ICMP Echo pings every 300 seconds.

[79] Source: http://www.dslreports.com/r3/smokeping.cgi?target=network.bb494e0966ee7afcc7112eb2fbb3064f&r=802

[80] See http://www.broadbandreports.com/forums/all

While some of these conversations are just airing user frustrations, they also serve a real purpose in validating the problem and checking its scope. If other users are not experiencing the same problem, a user may need to report the issue himself.

The fact that people do coordinate and air their frustrations in real time on these services illustrates an interesting phenomenon. In a number of cases, users utilize one platform (say mobile) to report problem on another (say fixed). This illustrates the resiliency afforded by having access to diverse infrastructure and a potential benefit of multi-homing.

4.3. Dedicated Infrastructure

The previous sections have covered data on reliability that is either limited in the time duration (e.g., occasional failure reports or snapshots from occasional surveys) or data that covers a topologically limited region of the network (e.g., a single user's connection). Generating a consistent record of measurements and data that can be used to gauge the reliability of broadband connections requires dedicated infrastructure. These measurements and data can come from two different sources: 1) Dedicated measurement infrastructure and 2) data from content and application providers who produce reliability relevant statistics as a by-product of analyzing the data they collect to assess and monitor their own networks.

4.3.1. Dedicated measurement infrastructure

Both Ofcom in the UK and the FCC in the United States have run large-scale measurements of broadband networks utilizing dedicated measurements boxes and infrastructure provided by Samknows. In the FCC version of the study, over 7,000 broadband subscribers have been running specially built access points which run numerous tests every hour of every day.[81] See Table 1 for a list of tests in the FCC study.

While most of the published analysis of this data has so far focused on characterizing the average behavior, almost any of the data produced by these tests could be analyzed in terms relevant to a discussion of reliability. Data from these tests could be used in various formulations of the reliability metrics discussed in Section 2.2. For instance, a measure of the reliability of performance might involve characterizing the bottom 5% of tests results or characterizing the percentage of tests where performance fell below some applicable threshold would be another.

[81] See FCC (2011).

Table 1: List of tests and their primary measures in the FCC broadband measurement study[82]

Test	Primary measure(s)
Download speed	Throughput in Megabits per second (Mbps) utilizing three concurrent TCP connections
Upload speed	Throughput in Mbps utilizing three concurrent TCP connection
Web browsing	The total time taken to fetch a page and all of its resources from a popular website
UDP latency	Average round trip time of a series of randomly transmitted UDP packets distributed over a long timeframe
UDP packet loss	Fraction of UDP packets lost from the UDP latency test
Video streaming	The initial time to buffer, the number of buffer under-runs and the total time for buffer delay
Voice over IP	Upstream packet loss, downstream packet loss, upstream jitter, downstream jitter, round trip latency
DNS resolution	The time taken for the ISP's recursive DNS resolver to return an A record for a popular website domain name
DNS failures	The percentage of DNS requests performed in the DNS resolution test that failed
ICMP latency	The round trip time of five regularly spaced ICMP packet
ICMP packet loss	The percentage of packets lost in the ICMP latency test
Latency under load	The average round trip time for a series of regularly spaced ICMP packets sent during downstream/upstream sustained test
Availability	The total time the connection was deemed unavailable for any purpose, which could include a network fault or unavailability of a measurement point
Consumption	A simple record of the total bytes downloaded and uploaded by the router

Initially, the availability test appears to be directly applicable to a discussion of connection reliability:[83]

> *This test measured the availability of the network connection from the Whitebox to multiple target test nodes by sending and receiving TCP segments to a receiving server located on each test node.*
>
> *The Whitebox established long-lived TCP connections to the server on each test node, periodically sending TCP packets containing a timestamp in microseconds.*
>
> *The server echoed back the same data to the Whitebox and if it failed to respond or the connection was reset via TCP RST or FIN then the Whitebox would attempt to re-establish the connection. If the Whitebox was unable to re-establish the connection to all three servers simultaneously, it was inferred that Internet connectivity was at fault, and the test recorded a failure locally, along with a timestamp to record the time of failure.*

A naive analysis of this data however would be misleading. Our initial analysis of the data turned up time intervals where the availability test indicated failures but other tests ran successfully. One source of errors was that the target test nodes, while generally very

---

[82] Source: Technical Appendix to FCC (2011), p. 21-22 (available at: http://transition.fcc.gov/cgb/measuringbroadbandreport/technical_appendix/Technical_Appendix_Full.pdf).

[83] Source: *ibid.*, p. 26.

reliable, may have occasionally experienced minor problems. Also it is suspected that in some cases gateways in the home may have been periodically dropping the long lived TCP flows currently being used for the availability test.

Samknows notes in a footnote that "[t]he measurement of availability provided a check on how often tests could not be run and was used as a quality metric overall, but was not used in analysis of broadband performance."[84] In particular, Samknows "removed measurements [for all tests] for any unit that exhibited greater than or equal to 10% failures [on the availability test] in a particular one hour period (the purpose was to remove periods where units were unable to reach the Internet)."[85]

All of the tests noted in the table above record a failure and success value. In our exploratory analysis of the data from these boxes, these at first initially looked like potential source of data on broadband reliability. However again, a failure is recorded under conditions that include factors beyond just problems with the broadband access network. Note of this is to say that no meaningful analysis is possible, just that simple approaches to analyzing the data would be inadequate.

If one takes a view of reliability as consistency, or time above a given threshold, the Samknows data is a rich source of information. Figure 5 shows a time of day plot of performance for each of the broadband providers in the FCC study. The one standout, Cablevision, has a considerable performance drop during peak usage hours in the evening. Cablevision's response[86] was that they have "consistently won top ratings in much broader and more extensive consumer surveys conducted by J.D. Power & Associates, PC Magazine and others." While this may be a typical marketing response, we agree that the connection between different metrics and actual user experience and preferences is generally complex.

---

[84] Source: *ibid*., p. 22.

[85] Source: *ibid*., p. 31.

[86] See http://www.reuters.com/article/2011/08/02/idUS358118963520110802

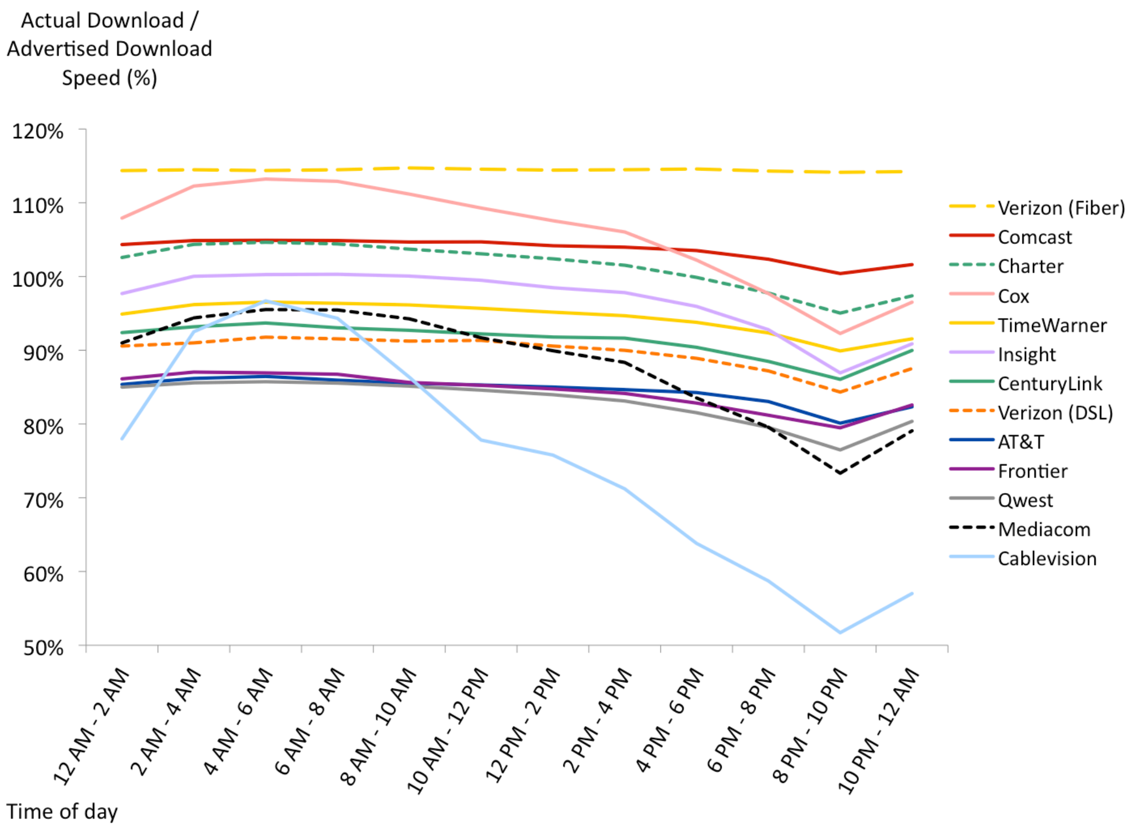

Figure 5: Actual Download/Advertised Download Speed (%) (Source: FCC, 2011).

4.4. Data from content and service providers

The other entities with both the infrastructure and the incentive to report on performance and reliability are large content, application and service providers. A number of noteworthy providers reporting on the experience of their users include YouTube's my_speed[87] and Netflix's characterization of its video performance on broadband networks[88]. We expect this will be an increasing trend where we will increasingly see content provider specific metrics of both performance and reliability. While measurements of "speed" came first from content providers, some metrics of reliability will likely follow. Blizzard and their World of Warcraft servers for instance could report connection drops where players are dropped from the game for having too large a delay or lack of a timely response. These faults should not naively be attributed to broadband providers. Home networks and other end-to-end impairments are frequent sources of problems as well.

---

[87] See http://www.youtube.com/my_speed

[88] See http://techblog.netflix.com/2011/01/netflix-performance-on-top-isp-networks.html

5. Conclusions

The science and practice of Internet and broadband performance measurement and assessment is evolving rapidly. This is motivated by the recognition that broadband Internet access constitutes basic infrastructure, and the fact that consumers, industry participants, and policymakers need good data in order to make informed decisions about purchasing, using, investing in, and regulating Internet services.

As part of the National Broadband Plan in the US,[89] and similar plans abroad, efforts are under way to assess the availability and quality of broadband services available to consumers. Over the past year, much of the debate has focused on issues such as the speeds actually provided by Internet service providers, and how those compare with advertised rates. In this paper, we address the challenge of assessing the reliability of broadband services. As we explain, this poses a much more difficult challenge than speed measurement.

There are distinctly different policy challenges associated with ensuring the "five-nines" reliability of core Internet backbone components and with monitoring the reliability of consumer broadband services. The former challenge confronts the need to provide strong incentives to adequately provision against what would be regarded as very high cost failures, and consequently, events that should – by design – be very rare (beyond five-nines reliability). The latter problem, which forms the focus of this analysis, is different and addressable, in part, by appropriate measurement strategies.

In order to design appropriate reliability metrics, it is important to recognize that reliability is inherently multidimensional (like speed but even more so). Multiple reliability concepts and metrics are needed to appropriately evaluate consumer broadband. After reviewing some of the technical literature on Internet reliability, we propose a tripartite grouping of reliability metrics: 1) Performance metrics, which characterize the probability that the service will meet or exceed a target level (e.g., average data rate); 2) Connectivity metrics, which measure the ability to connect to the Internet; and 3) Core service availability metrics, which measure the availability of certain core Internet services (e.g., email, DNS). Each category of metrics poses different sets of challenges.

To provide context for understanding the economic and policy challenges associated with regulating reliability, we discuss the multiple economic roles that reliability plays in markets and in governing the behavior of participants. The failure of systems imposes real costs, and the desire to avoid those costs influences behavior in wide-ranging ways. There are a number of reasons why we should be concerned that competitive markets may fail to deliver the socially optimal level of reliability. While there are a range of policy interventions that may be considered (including the specification of minimum

[89] See FCC, "Connecting America: the National Broadband Plan," Federal Communications Commission, Washington, DC, March 16, 2010.

reliability requirements, process standards, penalties and financial incentives, and reporting responsibilities), none are perfect and all have their own sets of challenges.

We also provide a preliminary review of some recent and on-going empirical work addressing reliability-related measurements. Our goal in discussing the multiple definitions for reliability, reviewing the literature and economics of reliability, and current measurement efforts, is to contribute to framing what we expect will be an important and vigorous public discussion.

6. References

Abbring, J., P. Chiappori, J. Heckman and J. Pinquet (2003), "Adverse selection and moral hazard in insurance: can dynamic data help to distinguish?" *Journal of the European Economic Association* 1 (2/3), pp. 512-21.

Ai, C., S. Martinez and D. Sappington (2004), "Incentive regulation and telecommunications service quality," *Journal of Regulatory Economics* 26 (3), pp. 263-85.

Akerlof, G. (1970), "The market for lemons: quality uncertainty and the market mechanism," *The Quarterly Journal of Economics* 84 (3), pp. 488-500.

Alves, M., L. Corsello, D. Karrenberg, C. Ögüt, M. Santcroos, R. Sojka, H. Uijterwaal, and R. Wilhelm (2002), "New measurements with the RIPE NCC test traffic measurements setup," Passive and Active Measurement Workshop (PAM), Fort Collins, CO, March 25-27, 2002.

Andersen, D., H. Balakrishnan, F. Kaashoek, and R. Morris (2001), "Resilient overlay networks," 18th ACM Symposium on Operating Systems Principles (SOSP), Banff, Canada, October 21-24, 2001.

Anthony, J., W. Choi and S. Grabski (2006), "Market reaction to e-commerce impairments evidenced by website outages," *International Journal of Accounting Information Systems* 7 (2), pp. 60-78.

Atkinson, R., D. Correa and J. Hedlund (2008), "Explaining international broadband leadership," The Information Technology and Innovation Foundation (ITIF), Washington DC.

Auriol, E. (1998), "Deregulation and quality," *International Journal of Industrial Organization* 16 (2), pp. 169-94.

Averch, H., and L. Johnson (1962), "Behavior of the firm under regulatory constraint," *American Economic Review* 52 (5), pp. 1052-1069.

Avramopoulus, I., H. Kobayashi, R. Wang, and A. Krishnamurthy (2004), "Highly secure and efficient routing," IEEE INFOCOM 2004, Hong Kong, China, March 7-11, 2004.

Ball, M. (1980), "Complexity of network reliability computations," *Networks* 10 (2), pp. 153-165.

Bauer, S., D. Clark, and W. Lehr (2009a), "Broadband microfoundations: the need for traffic data," Beyond Broadband Access, Washington DC, September 2009.

Bauer, S., D. Clark, and W. Lehr (2009b), "Evolution of Internet congestion," 37th Research Conference on Communication, Information and Internet Policy (TPRC), George Mason University School of Law, Arlington, VA, September 25-27, 2009.

Bauer, S., D. Clark, and W. Lehr (2010), "Understanding broadband speed measurements," 38th Research Conference on Communication, Information and Internet Policy (TPRC), George Mason University School of Law, Arlington, VA, October 1-3, 2010.

Bauer, S., D. Clark, and W. Lehr (2011), "Powerboost," HomeNets'11, Toronto, Ontario, Canada, August 15, 2011.

Bovy, C. J., H. T. Mertodimedjo, G. Hooghiemstra, H. Uijterwaal, and P. Van Mieghem (2002), "Analysis of end-to-end delay measurements in Internet," Passive and Active Measurement Workshop (PAM), Fort Collins, CO, March 25-27, 2002.

Brewer, E. (2001), "Lessons from giant-scale services," *IEEE Internet Computing*, Jul-Aug, pp. 46-55.

Cashell, B., W. Jackson, M. Jickling and B. Webel (2004), "Economic impact of cyber-attacks," CRS Report for Congress, Library of Congress, Order Code RL32331, April 1, 2004.

Carlier, J., Y. Li, J. Lutton, (1997), "Reliability evaluation of large telecommunication networks," *Discrete Applied Mathematics* 76 (1-3), pp. 61-80.

Castro, M. and B. Liskov (2002), "Proactive Byzantine fault tolerance and proactive recovery," *ACM Transactions on Computer Systems* 20 (4), pp. 398-461.

Chichilinsky, G. and P. Eisenberger (2010), "Asteroids: assessing catastrophic risks," *Journal of Probability and Statistics*, doi:10.1155/2010/954750.

Clark, D., W. Lehr, and S. Bauer (2011), "Interconnection in the Internet: the policy challenge," 39th Research Conference on Communication, Information and Internet Policy, George Mason University, Arlington, VA, September 23-25, 2011.

Cole, R. G., and J. H. Rosenbluth (2001), "Voice over IP performance monitoring," *SIGCOMM Comput. Commun. Rev.* 31 (2), pp. 9-24.

Czernich, N., O. Falck, T. Kretschmer and L. Woessmann (2011), "Broadband infrastructure and economic growth," *The Economic Journal* 121 (552), pp. 505-32.

Dischinger, M., A. Haeberlen, K. P. Gummadi, and S. Saroiu (2007), "Characterizing residential broadband networks," Proc. Internet Measurement Conference (IMC), pp. 43-56.

Dong, C., and N. Dulay (2011), "Argumentation-based fault diagnosis for home networks," HomeNets'11, Toronto, Ontario, Canada, August 15, 2011.

Downes, T. and S. Greenstein (2002), "Universal access and local Internet markets in the US," *Research Policy* 31, pp. 1035-1052.

Dynes, S., E. Andrijcic and M. Johnson (2006), "Costs to the US economy of information infrastructure failures: estimates from field studies and economic data," Fifth Workshop on the Economics of Information Security, Cambridge, UK, June 26-28, 2006.

FCC (2009a), "Quality of service of incumbent local exchange carriers," Federal Communications Commission, Washington DC, December 2009
(available at: http://www.fcc.gov/wcb/stats).

FCC (2009b), "Code of Federal Regulations, Title 47 Telecommunication, Chapter 1 Federal Communications Commission, Part 4 Disruptions to Communications," pp. 675-681, Revised as of October 1, 2009.

FCC (2010), "Broadband performance," OBI Technical Paper No. 4, Federal Communications Commission, Washington DC
(available at: http://www.broadband.gov/plan/broadband-working-reports-technical-papers.html).

FCC (2011), "Measuring broadband America: a report on consumer wireline broadband performance in the U.S.," Office of Engineering and Technology, Federal Communications Commission, August 2011
(available at: http://www.fcc.gov/measuring-broadband-america).

Flamm, K. (2004), "The role of economics, demographics, and state policy in broadband competition: an exploratory study," Telecommunications Policy Research Conference, Arlington, VA, October 2, 2004.

Garg, A., J. Curtis, and H. Halper (2003), "Quantifying the financial impact of IT security rreaches," *Information Management & Computer Security* 11(2), pp. 74-83.

Gillett, S., W. Lehr, C. Osorio and M. Sirbu (2006), "Measuring the economic impact of broadband deployment," Report to US Department of Commerce, Economic Development Association, Washington DC.

Holmstrom, B. (1979), "Moral hazard and observability," *Bell Journal of Economics* 10 (1), pp. 74-91.

Holmstrom, B., and P. Milgrom (1991), "Multi-Task Principal-Agent Analysis: Incentive Contracts, Asset Ownership and Job Design," *Journal of Law, Economics, and Organization* 7(1): 26–52.

Horrigan, J. (2010), "Broadband adoption and use in America," OBI Working Paper Series No. 1, Federal Communications Commission, Washington, DC.

Horrigan, J., and A. Smith (2007), "Home broadband adoption 2007," Pew Internet & American Life Project
(available at: http://www.pewinternet.org/PPF/r/217/report_display.asp).

Huffaker, B., M. Fomenkov, D. Moore, and K. Claffy (2001), "Macroscopic analyses of the infrastructure: measurement and visualization of Internet connectivity and performance"
(available at: http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.28.4969).

Jiang, H., and C. Dovrolis (2002), "Passive estimation of TCP round-trip times," *SIGCOMM Comput. Commun. Rev.* 32 (3), pp. 75-88.

Kaaniche, M., and M. Kanoun (1996), "Reliability of a commercial telecommunications system," Proc. Software Reliability Engineering, pp. 207-212.

Kahneman, D. and & A. Tversky (1979), "Prospect theory: an analysis of decision under risk," *Econometrica* 47, pp. 263-291.

Keralapura, R., C. Chuah, G. Iannaccone, S. Bhattacharyya (2004), "Service availability: a new approach to characterize IP backbone topologies," Proc. IEEE International Workshop on Quality of Service (IWQOS 2004), pp. 232-241.

Kushman, N., S. Kandula, and D. Katabi (2007), "Can you hear me now: it must be BGP," *SIGCOMM Comput. Commun. Rev.* 37 (2), pp. 75-84.

Labovitz, C., A. Ahuja, A. Bose, and F. Jahanian (2001), "Delayed Internet routing convergence," *IEEE/ACM Transactions on Networking* 9 (3), pp. 293-306.

Lehr, W. and J. Chapin (2010), "On the convergence of wired and wireless access network architectures," *Information Economics and Policy* 22 (1), pp. 33-41.

Lehr, W. and T. Grieco (2009), "Measurement and assessment of broadband availability," Report prepared for John Adams Innovation Institute of the Massachusetts Technology Collaborative, January 2009.

Lehr, W., T. Grieco, and G. Woo (2008), "Broadband metrics best practices: review and assessment," Report prepared for Massachusetts Technology Collaborative/John Adams Innovation Institute (MTC/JAII), February 2008.

Lehr, W., M. Sirbu, J. Peha, and S. Gillett (2007), "Scenarios for the network neutrality arms race," *International Journal of Communication* 1 (available at: http://ijoc.org/ojs/index.php/ijoc/article/view/164/89).

Lehr, W. and L. McKnight (2000), "A Broadband Access Market Framework: Towards Consumer Service Level Agreements," Twenty-Eighth Annual Telecommunications Policy Research Conference, September 25-27, 2000, Alexandria, VA

Lehr, W. and S. Gillett (1999), "Availability of broadband Internet access: empirical evidence," Twenty-Seventh Annual Telecommunications Policy Research Conference, September 25-27, 1999, Alexandria, VA.

Lehr, W. (ed.) (1995), *Quality and Reliability of Telecommunication Infrastructure*, Lawrence Erlbaum Associates, Hillsdale, NJ.

Maier, G., A. Feldmann, V. Paxson, and M. Allman (2009), "On dominant characteristics of residential broadband Internet traffic," Proc. Internet Measurement Conference (IMC), pp. 90-102.

Markopoulou, A. P., F. A. Tobagi, and M. J. Karam (2003), "Assessing the quality of voice communications over Internet backbones," *IEEE/ACM Transactions on Networking* 11 (5), pp. 747-760.

Markopoulou, A., G. Iannaccone, S. Bhattacharyya, C. Chuah, and C. Diot (2004), "Characterization of failures in an IP backbone network," IEEE INFOCOM 2004, Hong Kong, China, March 7-11, 2004.

Neuman, P. (2004), "Principled assuredly trustworthy composite architectures," Final Report Contract number N66001-01-C-8040, DARPA Order No. M132, SRI Project P11459, San Diego CA, December 28, 2004 (available at: http://www.csl.sri.com/users/neumann/chats4.html).

Nguyen, H. X., R. Teixeira, P. Thiran, and C. Diot (2009), "Minimizing probing cost for detecting interface failures: algorithms and scalability analysis," Proc. IEEE INFOCOM 2009, pp. 1386-1394.

Nojo, S., and H. Watanabe (1993), "Incorporating reliability specifications in the design of telecommunication networks," *IEEE Communications Magazine* 31 (6), pp. 40-43.

Noll, R. (1996), "The complex politics of catastrophe economics," *Journal of Risk and Uncertainty* 12, pp. 141-46.

NTIA (1999), "Falling through the net: defining the digital divide: a report on the telecommunications and information technology gap in America," National Telecommunications and Information Administration (NTIA), Washington DC.

OECD (2011), "Economic impact of Internet/broadband technologies," DSTI/ICCP/IE(2011)1/REV1, Working Party on the Information Economy, Directorate for Science, Technology and Industry, OECD, Paris, May 30, 2011.

Ofcom (2009), "UK broadband speeds 2009: consumers' experience of fixed-line broadband performance," July 28, 2009
(available at: http://stakeholders.ofcom.org.uk/binaries/research/telecoms-research/broadbandspeeds.pdf).

Ofcom (2010), "The consumer experience 2010," December 7, 2010
(available at: http://stakeholders.ofcom.org.uk/binaries/research/consumer-experience/tce-10/consumer-esperience-10.pdf).

Pang, J., J. Hendricks, A. Akella, R. De Prisco, B. Maggs, and S. Seshan (2004), "Availability, usage, and deployment characteristics of the domain name system," Proc. Internet Measurement Conference (IMC), pp. 1-14.

Pinheiro, E., W. Weber, and L. Barroso (2007), "Failure trends in large disk drive population," Proceedings of the 5th USENIX Conference on File and Storage Technologies (FAST-07), February 2007.

Ponemon Institute (2011), "Calculating the cost of data center outages," February 1, 2011
(available at: http://www.emersonnetworkpower.com/en-US/Brands/Liebert/Documents/White Papers/sl-24659.pdf).

Posner, R. (2004), *Catastrophe: Risk and Response*, Oxford, UK: Oxford University Press.

Prasad, R., C. Dovrolis, M. Murray, and K. C. Claffy (2003), "Bandwidth estimation: metrics, measurement techniques, and tools," *IEEE Network* 17 (6), pp. 27-35.

Price, C., B. Brigham and L. Fitzgerald (2002), "Service quality in regulated monopolies," University of East Anglia Centre for Competition and Regulation, CCR Working Paper 02-4, August 2002.

Rexford, J., J. Wang, Z. Xiao, and Y. Zhang (2002), "BGP routing stability of popular destinations," Proc. ACM SIGCOMM Workshop on Internet Measurement, pp. 197-202.

Sanso, B., F. Soumis, and M. Gendreau (1991), "On the evaluation of telecommunications network reliability using routing models," *IEEE Transactions on Communications* 39 (10), pp. 1494-1501.

Sappington, D. (2005) "Regulating service quality: a survey," *Journal of Regulatory Economics* 27 (2), pp. 123-54.

Snow, A., U. Varshney, A. Malloy (2000), "Reliability and survivability of wireless and mobile networks," *Computer* 33 (7), pp. 49-55.

Sommers, J., P. Barford, N. Duffield, and A. Ron (2007), "Accurate and efficient SLA compliance monitoring," *SIGCOMM Comput. Commun. Rev.* 37 (4), pp.109-120.

Song, H. H., Z. Ge, A. Mahimkar, J. Wang, J. Yates, and Y. Zhang (2011), "Analyzing IPTV set-top-box crashes," HomeNets'11, Toronto, Ontario, Canada, August 15, 2011.

Spence, A. (1975), "Monopoly, quality, and regulation," *The Bell Journal of Economics* 6 (2), pp. 417-29.

Spragins, J., J. Sinclair, Y. Kang, H. Jafari (1986), "Current telecommunication network reliability models: a critical assessment," *IEEE Journal on Selected Areas in Communications* 4 (7), pp. 1168-1173.

Sundaresan, S., N. Feamster, R. Teixeira, A. Tang, W. Edwards, R. Grinter, M. Chetty, W. de Donato (2011), "Helping users shop for ISPs with Internet nutrition labels," HomeNets'11, Toronto, Ontario, Canada, August 15, 2011.

Taleb, N. (2007), *The Black Swan: The Impact of Highly Improbable Events*, New York: Random House.

Tellabs (2007), "Beyond five nines: the new age of 'uber reliability'"
(available at: http://www.tellabs.com/resources/papers/tlab_uberreliability.pdf).

Tortorella, M. (2005), "Service reliability theory and engineering," *Quality Technology of Quantitative Management* 2 (1), pp. 1-37.

Uri, N. (2003), "Service quality effects of incentive regulation on access service in telecommunications in the United States," *European Journal of Law and Economics* 16 (3), pp. 369-90.

Varian, H (2001), "System Reliability and Free Riding"
(available at: http://people.ischool.berkeley.edu/~hal/Papers/2004/reliability).

Wang, H., H. Xie, L. Qiu, Y. R. Yang, Y. Zhang, and A. Greenberg (2006), "COPE: traffic engineering in dynamic networks," Proc. ACM SIGCOMM'06, pp. 99-110.

Wilson, J. (2003), “The Cost of Enterprise Downtime,” Infonetics Research.

Zhang, Y., L. Breslau, V. Paxson, and S. Shenker (2002), “On the characteristics and origins of Internet flow rates,” *SIGCOMM Comput. Commun. Rev.* 32 (4), pp. 309-322.