

Characterizing Residential Broadband Networks

Marcel Dischinger
MPI for Software Systems
mdischin@mpi-sws.mpg.de

Krishna P. Gummadi
MPI for Software Systems
gummadi@mpi-sws.mpg.de

Andreas Haeberlen
MPI for Software Systems, Rice University
ahae@mpi-sws.mpg.de

Stefan Saroiu
University of Toronto
stefan@cs.toronto.edu

A large and rapidly growing proportion of users connect to the Internet via residential broadband networks such as Digital Subscriber Lines (DSL) and cable. Residential networks are often the bottleneck in the last mile of today's Internet. Their characteristics critically affect Internet applications, including voice-over-IP, online games, and peer-to-peer content sharing/delivery systems. However, to date, few studies have investigated commercial broadband deployments, and rigorous measurement data that characterize these networks at scale are lacking.

In this paper, we present the first large-scale measurement study of major cable and DSL providers in North America and Europe. We describe and evaluate the measurement tools we developed for this purpose. Our study characterizes several properties of broadband networks, including link capacities, packet round-trip times and jitter, packet loss rates, queue lengths, and queue drop policies. Our analysis reveals important ways in which residential networks differ from how the Internet is conventionally thought to operate. We also discuss the implications of our findings for many emerging protocols and systems, including delay-based congestion control (e.g., PCP) and network coordinate systems (e.g., Vivaldi).

Categories and Subject Descriptors

C.2.2 [Computer Systems Organization]: Computer-Communication Networks—*Network Operations*; C.2.5 [Computer Systems Organization]: Computer-Communication Networks—*Local and Wide-Area Networks*; C.4 [Computer Systems Organization]: Performance of Systems

General Terms

Measurement, Performance, Experimentation

Keywords

Broadband access networks, DSL, cable, network measurement

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IMC'07, October 24-26, 2007, San Diego, California, USA.
Copyright 2007 ACM 978-1-59593-908-1/07/0010 ...\$5.00.

1. INTRODUCTION

Residential broadband networks such as Digital Subscriber Lines (DSL) and cable are increasingly being used to access the Internet. More than 158 million people use these networks worldwide [39], and this number is expected to rise to 477 million by 2011 [51]. In the United States alone, more than half of all Internet users connect via residential broadband networks [38]. In addition, many governments are adopting policies to promote ubiquitous broadband access [18, 48].

Residential broadband networks provide the critical 'last mile' access to the Internet infrastructure. It is widely thought that the bottlenecks in the performance of the Internet lie in its access networks [1]. So the reliability and the performance of Internet applications – including voice-over-IP (VoIP), video on demand, online games, and peer-to-peer content delivery systems – depend crucially on the characteristics of broadband access networks.

Despite the widespread deployment of residential broadband networks and their importance to emerging applications, they remain relatively unexplored by the academic community. Although many measurement studies have focused on the Internet's core [6, 26, 40] and academic/research edge networks [5, 35], rigorous measurement data that characterize residential network deployments at scale are lacking.

In the absence of systematic studies, knowledge about residential broadband networks is based on anecdotal evidence, hearsay, and marketing buzzwords. Although broadband networks are known to have very different characteristics from academic networks [5, 43], there have been no large-scale studies quantifying these differences. As a result, researchers today are left to second-guess how well protocols or systems evaluated in academic networks would work in the commercial Internet, where broadband networks are widely deployed.

One reason for the lack of large-scale measurement studies on residential networks is that researchers have limited access to broadband environments. Most academic institutions and research laboratories do not access the Internet over broadband. Even state-of-the-art research network testbeds, such as PlanetLab [41] and RON [2], have only a handful of broadband nodes. We overcame this problem by developing tools that can measure broadband networks remotely and without cooperation from end hosts connected to the broadband links.

In this paper, we present the first large-scale measurement study examining 1,894 broadband hosts from 11 major com-




Figure 1: A typical setup of (a) cable and (b) DSL access networks

mercial cable and DSL providers in North America and Europe. To conduct this study, we developed tools that enable us to measure a large number of remote broadband links. We performed a detailed characterization of an extensive set of properties of broadband links. Our analysis was driven by three questions:

1. What are the typical bandwidth, latency, and loss characteristics of residential broadband links?
2. How do the characteristics of broadband networks differ from those of academic or corporate networks?
3. What are the implications of broadband-network properties for future protocol and system designers?

Our study reveals important ways in which cable and DSL networks differ from the conventional wisdom about the Internet, accumulated from prior studies of academic networks. For example, many cable links show high variation in link bandwidths over short timescales. Packet transmissions over cable suffer high jitter as a result of cable’s time-slotted access policy. DSL links show large last-hop delays and considerable deployment of active queue management policies such as random early detection (RED). Both cable and DSL ISPs use traffic shaping and deploy massive queues that can delay packets for several hundred milliseconds.

Our findings have important implications for emerging protocols and systems. For instance, the high packet jitter in cable links can affect transport protocols that rely on round-trip time (RTT) measurements to detect congestion, such as TCP Vegas [9] and PCP [3]. Further, the large queue sizes found in cable and DSL ISPs can be detrimental to real-time applications such as VoIP when they are used concurrently with bandwidth-intensive applications such as BitTorrent [8].

The rest of the paper is organized as follows: Section 2 provides an overview of residential cable and DSL networks. Section 3 describes our measurement methodology, including the tools we built for gathering data over remote broadband network links. Section 4 then presents an in-depth analysis of our data set, characterizing the bandwidth, latency and loss properties of broadband networks. In Section 5, we discuss the implications of our findings for the designers of future protocols and systems. Section 6 presents related work, and Section 7 summarizes our conclusions.

2. BACKGROUND

Two types of broadband access networks are popular today: cable networks and DSL networks. In this section, we present a brief description of their architectures, and we point out differences to other access networks, such as corporate and academic networks.

2.1 Cable networks

Cable networks use the cable television infrastructure to connect home users to the Internet. In these networks, a master headend connects to several regional headends using fiber-optic cables. Each regional headend serves a set of customers (up to 2,000 homes). A single coaxial cable, carrying both television and data signals, links these customers to their headend.

DOCSIS [10] is the most common specification defining the interface requirements of cable modems. In DOCSIS, each cable modem (CM) exchanges data with a cable modem termination system (CMTS) located in a regional headend. In the downstream direction, the CMTS broadcasts data to all cable modems that are connected to it. The cable modems filter all received data and forward only the bytes destined for their customer’s host. In the upstream direction, the access channel is time-slotted – a cable modem must first reserve a time slot and wait until the CMTS grants the reservation. When the time slot has been granted, the cable modem can transmit data upstream. Figure 1(a) illustrates a typical setup of a cable access network.

There are several important differences between cable and other access networks. First, cable links typically have asymmetric bandwidths: their downstream bandwidth is much higher than their upstream bandwidth. Second, customers cannot use the full raw capacity of their cable links. Instead, cable networks use traffic shaping to restrict users from consuming more bandwidth than their contract stipulates. Although cable networks currently allow raw data rates of up to 40 Mbps, the contracts of individual customers specify much lower rates, between 128 Kbps and 10 Mbps. Further, some ISPs over-subscribe their cable access networks. In this case, the level of service experienced by customers can vary depending on the amount of competing network traffic.

Finally, cable modems can concatenate multiple upstream packets into a single transmission, which results in short bursts at high data rates. Thus, the upstream latencies can fluctuate heavily, depending on the allocation policy, and the amount of signaling and concatenation used by the CMTS.

2.2 Digital Subscriber Line networks

DSL access networks use existing telephone wiring to connect home users to the Internet [13]. Unlike cable customers, DSL customers do not share their access link. Each customer’s DSL modem uses a dedicated point-to-point connection to exchange data with a Digital Subscriber Line Access Multiplexer (DSLAM). The connection carries both data and telephone signals, which are encoded in different frequencies. On the customer side, a splitter separates the two signals and forwards the data signal to the DSL mo-

	DSL						Cable				
	Ameritech	BellSouth	BT Broadband	PacBell	Qwest	SWBell	Charter	Chello	Comcast	Road Runner	Rogers
Company	AT&T	AT&T	BT Group	AT&T	Qwest	AT&T	Charter Comm.	UPC	Comcast	TimeWarner	Rogers
Region	S+SW USA	SE USA	UK	S+SW USA	W USA	S+SW USA	USA	Netherlands	USA	USA	Canada
Hosts measured	113	155	173	158	97	397	114	120	118	301	148
Offered BWs (bps)	768K, 1.5M, 3M, 6M	768K, 1.5M, 3M, 6M	2-8M	768K, 1.5M, 3M, 6M	256K, 1.5M, 7M	768K, 1.5M, 3M, 6M	3M, 5M, 10M	384K, 1.5M, 3M, 6M, 8M	6M, 8M	5M, 8M	128K, 1M, 5M, 6M

Table 1: **Measured hosts:** We measured 1,894 broadband hosts from 11 major commercial cable and DSL providers in North America and Europe.

dem. Figure 1(b) illustrates a typical setup of a DSL access network.

There are two important differences between DSL networks and other access networks. First, like cable networks, DSL networks often have asymmetric bandwidths; their downstream bandwidth is higher than their upstream bandwidth. Second, the maximum data transmission rate falls with increasing distance from the DSLAM. To boost the data rates, DSL relies on advanced signal processing and error correction algorithms, which can lead to high packet propagation delays. Consequently, the properties of DSL access links vary depending on the length or the quality of the wiring between a modem and its DSLAM.

3. MEASUREMENT METHODOLOGY

The goal of our study was to perform a rigorous characterization of broadband access networks. For this, we measured their link bandwidths, latencies, and loss rates. We also characterized the properties of broadband queues, including queue sizes and packet drop policies. Finally, we examined a physical property specific to the cable transmission medium: the time-slotted access policy of the upstream channel. We measured the effects of this access policy on latency and jitter. Because broadband access links are asymmetric, we measured the properties of the upstream and downstream directions separately.

For our measurements to be generally applicable, the study needed to be performed at large scale. Previous studies of broadband [14, 32, 33] used measurement tools that required cooperation from the remote broadband hosts. Such a methodology restricts the scale of the measurement study. Instead, we developed a different methodology for conducting large-scale detailed broadband measurements. Our approach requires minimal cooperation from the remote hosts, allowing our measurements to scale to thousands of broadband links.

Remote hosts need to cooperate only in two simple ways. First, they must respond to ICMP echo request packets with ICMP echo responses. Second, they must send TCP reset (RST) packets when they receive TCP acknowledgments (ACK) that do not belong to an open TCP connection. Both responses are mandated by the corresponding protocols, and previous work shows they are supported widely [23].

At a high level, our technique is simple – we probe the broadband link with packet trains of different rates, using packets of various types and sizes. We use the responses received to infer a broad range of characteristics, both down-

stream and upstream. This approach requires support from only one endpoint of an Internet path, but obtaining accurate measurements is more challenging than with tools that require support from both endpoints or with tools that have been explicitly designed to measure one specific property [17, 29, 33].

In the remainder of this section, we present our measurement methodology in more detail. We describe how we selected broadband hosts from different ISPs. We list the types of probe trains used to gather data. And we describe how we inferred the characteristics of the broadband links. Finally, we present how we validated the assumptions of our methodology, and we discuss potential concerns and limitations of our tools.

3.1 Selecting residential broadband hosts

We used techniques similar to those described in [23] to select 1,894 broadband hosts from 11 major cable and DSL providers in North America and Europe. We identified IP address ranges of popular residential ISPs from IP-to-DNS mappings (e.g., BellSouth’s DNS names are `ads1-*.bellsouth.net`), and we scanned for IP addresses responding to our probes.

Table 1 summarizes high-level information about the ISPs we measured. Our study includes five out of the top ten largest broadband ISPs in the U.S. [27]¹, the largest cable provider in Canada [28], the second-largest cable provider in the Netherlands [50], and the largest DSL provider in the U.K. [42]. From each ISP, we chose approximately 100 hosts randomly and measured them.

Table 1 also shows the bandwidths advertised by ISPs on their web sites. Although a range of speeds is available, all advertised bandwidths are lower than 10 Mbps. We took advantage of this property by using 10 Mbps probe streams to saturate these broadband links and their routers.

3.2 Probe trains to measure broadband links

We used five types of probe packet trains to measure each broadband link. Each probe train was sent from well-connected hosts located in four academic networks (Figure 2). The academic networks used are dispersed geograph-

¹During the recent consolidation of the U.S. telecom industry, many large ISPs merged with each other. Four of the eleven ISPs we measured are owned today by AT&T, a single company. However, our measurements show that their networks have very different characteristics. For the purposes of this study, we treat them as independent ISPs.




Figure 2: Experimental setup

ically – three in North America (in the south, northwest, and northeast) and one in Europe. We also probed the last-hop router before each broadband link. We used `traceroute` to discover these routers.

We sent probe trains at different rates. We refer to our high-rate probe trains as floods, and to our low-rate probe trains as trickles. All floods were sent at 10 Mbps to saturate the broadband links. Consequently, packet floods measure the network under congestion. By contrast, all packet trickles were sent at a rate of a few tens of Kbps, so they characterize the broadband network under normal operational conditions.

We limited the packet floods to at most 10 s, whereas we allowed trickles to last from several hours to several days. To capture diurnal variations of network properties, we repeated the floods every half hour for one week.

Asymmetric large-TCP flood: We sent large (1,488-byte²) TCP ACK packets, and the remote host responded with small (~40-byte) TCP RST packets. The ACK packets saturated the downstream links and router queues, but the responses, being smaller and fewer, did not saturate the upstream links or queues.

Symmetric large-ICMP flood: We sent large (1,488-byte) ICMP echo request (PING) packets, and the remote host responded with ICMP echo response packets of the same size. This packet train saturated the links and router queues in both downstream and upstream directions.

Symmetric small-TCP flood: We sent small (40-byte to 100-byte) TCP ACK packets, and the remote host responded with small (~40-byte) TCP RST packets. Like the symmetric large-probe flood, this packet train saturated the network in both downstream and upstream directions but with much smaller packets.

Symmetric large-ICMP trickle: We sent large (1,488-byte) ICMP echo request packets spaced at large intervals randomly chosen between 10 ms and 30 ms, and the remote host responded with ICMP echo response packets of the same size. Unlike the above probe trains, this packet train did not saturate the downstream or upstream links.

Symmetric small-TCP trickle: We sent small (40-byte) TCP ACK packets spaced at large random intervals between 10 ms and 30 ms, and the remote host responded with small (~40-byte) TCP RST packets. This packet train did not saturate the downstream or upstream links.

²We used 1,488-byte probes because some DSL links running PPPoE or PPPoA have an MTU of less than 1,500 bytes.

3.3 Measured broadband link properties

Our measurements rely on a simplifying assumption: that the broadband access link is the only bottleneck along the Internet path between our measurement hosts and the remote broadband hosts (Figure 2). We validate this assumption in the next section. This section describes how we measured the properties of the broadband links based on this assumption.

Link bandwidth: To estimate the allocated downstream bandwidth, we calculated the fraction of answered probes in the large-TCP flood, which saturates the downstream link only. For example, we estimate the downstream bandwidth of a link to be 6 Mbps when 60% of packets in our 10 Mbps large-TCP flood are answered. We used the same technique to estimate upstream bandwidths from the symmetric large-ICMP flood. The behavior of the large-ICMP flood is driven by the bandwidth of the slower link, which for cable and DSL is the upstream link.

Our techniques yield incorrect estimates in the presence of cross-traffic. We use IPID-based techniques described in [23] to identify and eliminate all measurement probes affected by cross-traffic.

Packet latencies and jitter: We characterized three types of packet delays and their variation (jitter) for each link: queuing delay, propagation delay, and transmission delay.

We estimated the maximum possible queuing delays (or queue lengths) by calculating the variation in RTTs of packets in our floods. To determine downstream queue lengths, we calculated the difference between the 95th percentile highest RTTs and minimum RTTs of packets in the large-TCP flood, which overflows only the downstream router queues. A similar calculation for the large-ICMP flood, which overflows queues in both directions, estimated the sum of downstream and upstream queue lengths. We subtracted the downstream queue length from this estimate to obtain the length of the upstream queue.

To study propagation delays of broadband links, we estimated their last-hop delays. We calculated last-hop delay as the difference between the latencies of small-TCP trickle probes to the broadband host and to its last-hop router. By comparing the last-hop delays for different packet sizes, we were able to infer the transmission delays in broadband links. We discuss transmission delays in more detail in Section 4.2.

Packet loss: We estimated typical packet loss rates in broadband networks by calculating the fraction of lost packets in the small-TCP trickle. To detect packet loss due to queue management policies, such as RED, we examined how the loss rate varies with the latencies of the packets. We discuss the details of RED detection in Section 4.3.

3.4 Validating our assumptions

Next, we discuss five important concerns about our methodology:

1. To be accurate, our probes must traverse the entire Internet path reaching the broadband host and not be answered by an intermediate router. *Do our measurements reflect accurately the properties of broadband access links?*
2. We assumed that the broadband links are the bottlenecks in the measured Internet paths. *How often are broadband links the bottlenecks along the measured Internet paths?*




Figure 3: **The broadband link is the bottleneck:** Comparison of the paths to the residential broadband hosts and their corresponding last-hop routers. The former include the broadband link, while the latter do not. The two sets of paths have very different characteristics, which validates our assumption that broadband links are the bottlenecks along the Internet paths to broadband hosts.

3. We assumed broadband hosts respond to all probes without any delay. In practice, end hosts could drop or rate limit their responses. *How often do broadband hosts delay or drop response packets?*
4. Our probes can be interpreted as port scans or attacks. *What are the best practices we used in our measurements?*
5. Large-scale Internet studies suffer from limitations and shortcomings. *What are some of the limitations of our study?*

3.4.1 Do our measurements reflect accurately the properties of broadband access links?

We ran controlled experiments using five broadband hosts (two cable and three DSL) under our control, located in North America and Europe. These experiments were performed on a small scale because they required end-host cooperation. Although we hoped to recruit more volunteers, the effort required to setup our experiments made it difficult to convince users to perform them. Our experiments require root access and manual changes to the modems' firewalls.

First, we checked whether the probe packets were being sent over the broadband link or whether they were being answered by a router in the middle of the network. We found that in all cases the probes were being responded to by the NAT-enabled modems in the customers' premises. By configuring the modems to forward any arriving probe packets to end hosts, we were able to receive the probes at our end hosts (Figure 2). Note that the probes must cross the broadband link to reach the modems.

Second, we checked whether the NAT-enabled modems affected the measurements by delaying or rate-limiting their responses. We gathered two traces for each link: one when the modem responded to the probes, and another when the modem forwarded all probes to the broadband hosts. We configured the broadband hosts to respond to the probes without any delay (less than 100 μ s) or rate-limiting. We compared the two traces with respect to latencies and losses of probes and responses. The two traces matched closely in all cases, suggesting that the modems do not adversely affect our measurements.

Finally, we verified the accuracy of our bandwidth and queue length measurements. We compared the measured bandwidths of the access links with the rate speeds advertised by their ISPs. We found that these bandwidths matched very closely – the average difference in downstream bandwidths was less than 3%. To validate our queue length

estimates, we used our access to the end hosts to measure the upstream and downstream queue lengths separately and accurately. The measurements matched the estimated queue lengths very well. The close match suggests that both our bandwidth and queue length measurements are accurate.

3.4.2 How often are broadband links the bottlenecks along the measured Internet paths?

Our methodology assumes the broadband link is the bottleneck on the Internet path measured. Because our probes are sent from well-connected academic hosts, the broadband links are likely to be the bottlenecks in these paths. To validate this assumption, we sent a large-TCP flood probe train to the broadband host and another train to its last-hop router. Comparing these two probe trains revealed that the broadband links are in fact the bottlenecks.

Figure 3 compares the available bandwidth, the RTT increases, and the packet loss rate of the two traces for 1,173 randomly selected broadband hosts. Most paths to the last-hop routers achieved the full 10 Mbps throughput, experiencing almost no losses or RTT fluctuations. By contrast, the paths including the broadband link had much lower throughput, considerable RTT increases, and high packet loss. This suggests that these variations are caused by the last hop (i.e., the broadband link).

3.4.3 How often do broadband hosts delay or drop response packets?

Our methodology assumes broadband hosts respond to probes without any delay. Several factors could prevent hosts from responding to some or all of our probes. For example, a firewall may block certain types of probes, such as PINGs. Some routers add a delay between the arrival of a probe and the departure of the response [21]. Also, a host with limited processing power might delay or drop packets arriving at high rates.

We removed all hosts that did not respond to our probes. We also removed the broadband hosts that rate-limited their probe responses. We identified such hosts by checking for large loss episodes occurring periodically.

Finally, we performed the following experiment to check whether our probe trains were too aggressive for the processing power of some hosts. We sent probe trains at 10 Mbps but with varying packet sizes. Although the trains consumed the same bandwidth, their packet sending rates were different. We checked whether hosts experienced higher losses at faster sending rates. A higher loss rate suggests that an end host cannot process packets at fast rates. We checked how

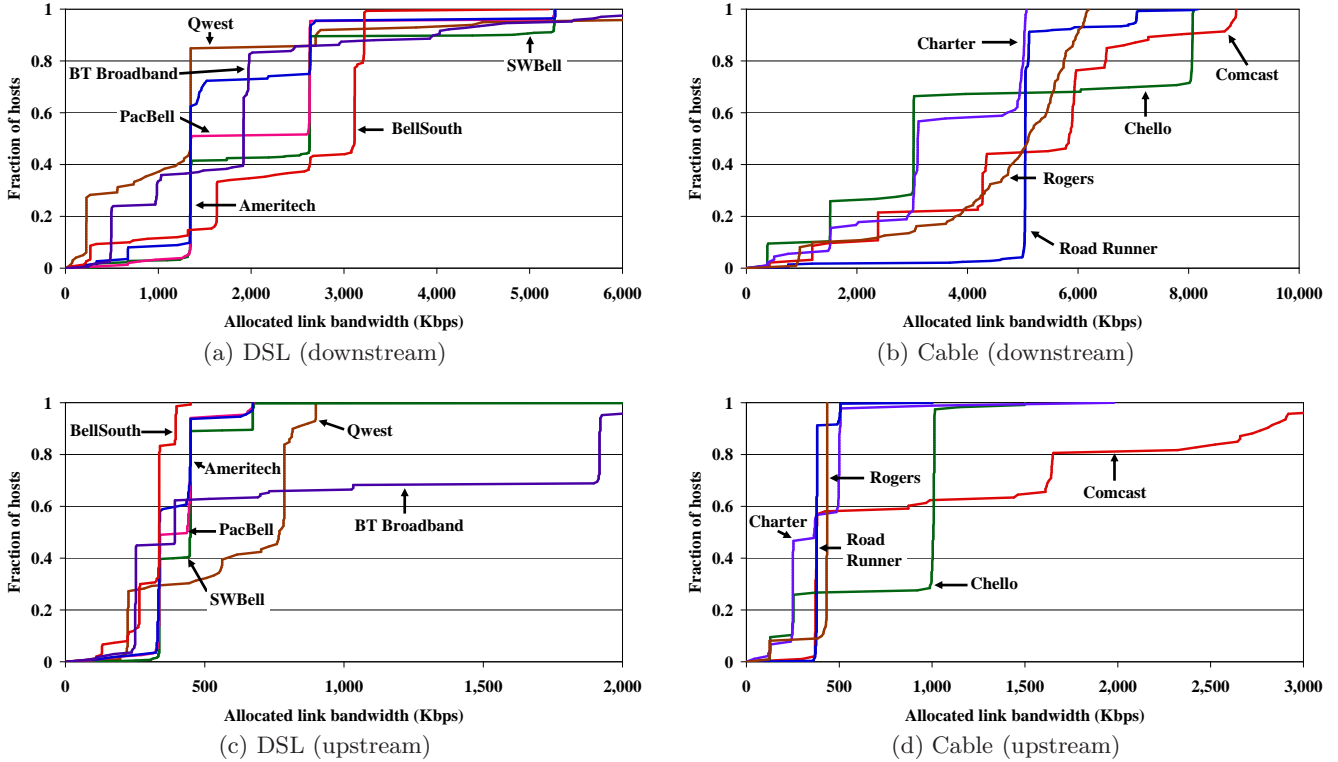


Figure 4: **Allocated downstream and upstream link bandwidths:** Most ISPs offer upstream bandwidths of 500 Kbps or less, even when the downstream bandwidths exceed 5 Mbps.

losses vary with packet sending rates for all broadband hosts in our study. The loss rates remained constant for over 99% of the hosts in our study, suggesting that the end hosts have sufficient processing power to handle our probing rates.

3.4.4 What are the best practices we adopted?

Performing active measurements on the Internet raises important usage concerns. Although it is difficult to address and eliminate all such concerns, we adopted a set of precautions to mitigate these concerns. We restricted our high rate probe trains to no more than 10 s each. We also embedded a custom message in each of our probe packets, which described the experiment and included a contact email address. To date, we have not received any complaints.

Another cause for concern was that users with a per-byte payment model end up paying for our unsolicited traffic. To mitigate this concern, we only measured hosts in ISPs that offer flat-rate payment plans, and we restricted the total amount of data sent to any single broadband host over our entire study.

3.4.5 What are some of the limitations of our study?

Two important limitations affect our measurements. First, we studied only major cable and DSL ISPs in North America and Europe. Our conclusions are unlikely to generalize to high-speed fiber-based broadband ISPs, such as those in Japan or South Korea [12]. Second, we removed all hosts that did not respond to our probes or that were rate-limited, which could introduce some unknown bias.

4. CHARACTERIZING BROADBAND LINKS

In this section, we analyze the data gathered from sending probe packet trains to a large number of residential broadband hosts in several major ISPs (see Table 1). We examine three important characteristics of broadband networks, namely link bandwidths, packet latencies, and packet loss. Analyzing these properties is important because they affect the performance of protocols and systems running over broadband.

4.1 Allocated link bandwidth

Allocated link bandwidth refers to the bandwidth reserved by a provider to a single broadband user. In cable networks, allocated link bandwidth is the portion of the shared link’s capacity assigned to an individual user, whereas in DSL networks it is the ISP’s cap on a user’s traffic rate. Characterizing allocated link bandwidths in broadband networks helps to predict the maximum throughput any transport protocol (such as TCP Reno or TCP Vegas) or application (such as BitTorrent) can achieve. As described in Section 3.3, our probe streams measured allocated bandwidths by saturating the broadband links.

4.1.1 What are the allocated link bandwidths?

Figures 4(a) and (b) show the cumulative distributions of downstream link bandwidths for the different DSL and cable ISPs. For many ISPs, the distributions jump sharply at distinct bandwidth levels, such as 256 Kbps, 384 Kbps, 512 Kbps, and 1 Mbps. Only two cable ISPs (Rogers in

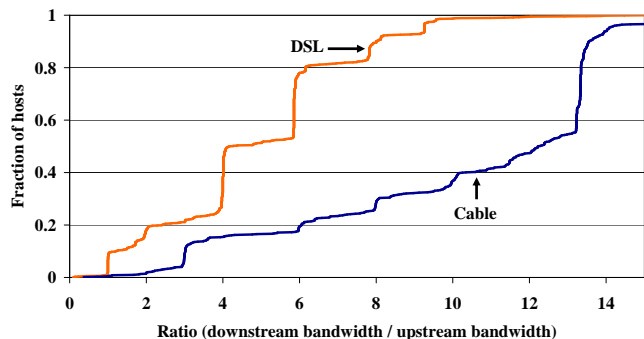


Figure 5: **The ratio of downstream to upstream link bandwidths:** The gap between downstream and upstream bandwidths is much wider for cable networks than for DSL networks.

Canada and Comcast in the United States) allocate bandwidths distributed along a continuous spectrum.

By comparing these measured allocated bandwidths to the advertised link speeds from Table 1, we can confirm and quantify some commonly held opinions. We find that most DSL ISPs have bandwidth rates corresponding to those advertised. By contrast, major cable ISPs, such as Comcast and Rogers, show rates different from those advertised (both higher and lower). We consider that this discrepancy is due to the nature of the two technologies – cable is a shared medium, whereas DSL is not. Our data also shows that many cable ISPs have significantly higher downstream bandwidths than DSL.

Figures 4(c) and (d) show the cumulative distributions of upstream link bandwidths. Upstream bandwidths are strikingly different from downstream bandwidths – with the exception of a few ISPs, most upstream bandwidths are lower than 500 Kbps, even when their downstream bandwidths exceed 5 Mbps. To examine this difference, we plotted the ratio of downstream to upstream link bandwidths in Figure 5. Most DSL hosts have much smaller ratios than cable hosts, because compared to cable, DSL hosts have lower downstream but similar upstream bandwidths. For over half of the cable hosts, the downstream bandwidths exceed upstream bandwidths by a factor of more than 10.

The highly asymmetric nature of bandwidths does not align well with the requirements of emerging peer-to-peer systems [8, 24], whose workloads tend to be symmetric. Despite all the excitement surrounding user-driven content generation and distribution, residential networks continue to be predominantly optimized for client-server workloads.

4.1.2 How stable are the allocated link bandwidths?

Next, we studied the short-term and long-term stability of link bandwidths. Understanding the stability of link properties is useful for designing network protocols that can quickly adapt to changing link conditions.

We examined the stability of the allocated link bandwidths over the 10 s duration of our packet floods. For this, we divided the 10 s into 100 ms intervals (the RTT of a typical Internet path), we estimated the bandwidth within each interval, and we compared the different estimates across intervals. Figure 6 shows how bandwidths for a PacBell link (DSL) and a Rogers link (cable) vary over time. Whereas




Figure 6: **Stable and unstable link bandwidths:** The allocated link bandwidth is stable for the PacBell DSL host. For the Rogers cable host, the access link bandwidth varies greatly over time.




Figure 7: **Fraction of hosts with ‘stable’ downstream link bandwidths:** Most DSL links show stable bandwidths, whereas most cable links do not. The results for upstream bandwidth are similar.

the PacBell link shows stable bandwidth, the Rogers link weaves above and below its average bandwidth of 3 Mbps.

Figure 7 shows the fraction of DSL and cable links that exhibit stable bandwidths in the downstream direction. We classified a link as stable if at least 90% of the 100 ms intervals show a bandwidth estimate within 10% of the average bandwidth. Although most DSL ISPs show stable link bandwidths, we found that most cable ISPs have bandwidths that vary significantly even within the short 10 s duration of our probes. We also found that upstream bandwidths have unstable short-term characteristics. We have omitted these results because of space constraints.

This large short-term variation in cable bandwidths poses new challenges to transport protocol designers. Traditionally, transport protocols have been developed to achieve stable throughput and to avoid reacting to short-term events (on timescales less than one RTT) [19]. However, when running in a cable network environment, protocols need to adjust quickly to rapidly changing link conditions. Slow reacting protocols might not achieve good throughput in cable networks.

We now turn our focus to the long-term diurnal stability of link bandwidths. We took measurements of the upstream and downstream bandwidths every half an hour for one week

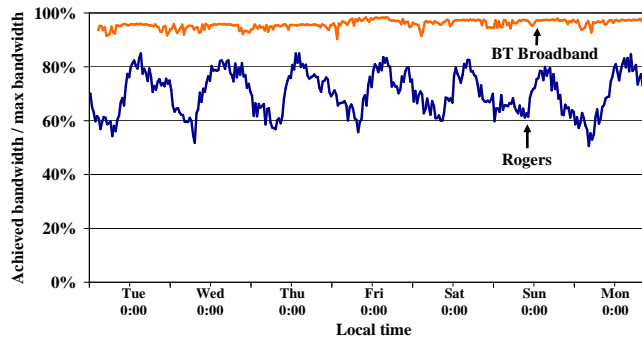


Figure 8: **Long-term link bandwidth stability:** Whereas BT Broadband has stable bandwidths over time, Rogers’s allocated link bandwidths show diurnal patterns.

from 70 randomly chosen hosts from each ISP.³ Figure 8 shows the diurnal variation in bandwidths for one DSL ISP (BT Broadband) and one cable ISP (Rogers). Each curve shows the bandwidth variation averaged across all measured links within one ISP. To account for links with different bandwidths, we normalize each link’s bandwidth by using the maximum measured bandwidth of that link during the entire measurement period.

We found that most ISPs have high long-term stability (not shown). As the curve for BT Broadband illustrates, their bandwidths do not vary with the time of the day. By contrast, a small number of ISPs, such as Rogers, show a clear diurnal trend in link bandwidths. Rogers’s end hosts see significantly lower bandwidths (almost a 25% reduction) in the evening (between 4 PM and 7 PM) than in the early morning (between 1 AM and 5 AM). In the upstream direction, we find stable bandwidths (not shown) for all ISPs, including Rogers. These findings seem to contradict the popular idea that competing traffic affects the bandwidths of broadband hosts. For most ISPs, we found little evidence that competing traffic affects link bandwidths during the day.

4.1.3 Is there evidence of traffic shaping?

Traffic shaping is likely to be one of the factors leading to the bandwidth instability encountered in broadband networks. Some ISPs allow an initial burst of bandwidth that is often many times greater than the advertised bandwidth. For example, Comcast’s PowerBoost feature [15] doubles the customer’s allocated bandwidth for a short time. This provisioning reduces the download times of relatively small files, such as MP3s. Other ISPs throttle the bandwidth allocated to long running transfers to discourage the heavy hitters from consuming a disproportionate share of the bandwidth.

Because our probe floods were limited to 10 s, we could only detect the traffic shaping associated with short-duration flows. To do this, we performed the following experiment. We used our packet streams to compute the allocated link bandwidth of each 100 ms interval. To detect the presence of traffic shaping, we checked for a consistent and significant drop in bandwidth after some initial period. Figure 9




Figure 9: **Traffic-shaped downstream:** The bandwidth of this link is initially 2.5 Mbps, but it drops to 1.5 Mbps after one second.

shows an example link from Ameritech DSL, whose bandwidth drops from 2.5 Mbps to 1.5 Mbps (its long-term rate) after the first second.

We found similar downstream traffic-shaping techniques used by three ISPs – Ameritech, Comcast, and Chello. 11% of the Ameritech links, 26% of the Comcast links, and 67% of the Chello links provide an initial burst of bandwidth to speed up short transfers. The burst rates are typically more than 1 Mbps above the long-term bandwidth. However, in many cases, we were unable to quantify precisely the burst rates because they exceeded the rate of our probe train. In the upstream direction, we found no evidence of traffic shaping or bandwidth throttling of our probe stream. The short duration of our probe trains (10 s) could have prevented us from detecting upstream traffic shaping.

History-based bandwidth prediction is a popular technique used in several transport protocols [3, 9, 19] and content distribution systems [25]. Although our traffic-shaping analysis is preliminary, it suggests that using past bandwidth estimates to predict future bandwidth conditions might not work well over broadband links.

4.2 Packet latencies

We analyzed each of the three components of packet latencies: propagation delays, transmission delays, and queueing delays.

4.2.1 Do broadband links have large propagation delays?

A link’s propagation delay is the time elapsed between sending a bit at one end and receiving it at the other end. On one hand, broadband propagation delays could be short because the links themselves are short. On the other hand, sophisticated signal processing and error correction algorithms could increase broadband propagation delays.

Our methodology prevents us from directly measuring the propagation delay of a broadband access link. Instead, we were able to estimate the round-trip delay of the last-hop of the path between our measurement hosts and the broadband hosts. This last-hop delay roughly approximates the sum of downstream and upstream broadband propagation delays.

To do this, we sent small-TCP trickle probes to both the broadband host and its last-hop router. The trickle consisted of several hundred widely spaced small probes and their responses. We calculated the last-hop RTT by subtracting the

³To minimize DHCP effects, we discarded any host that went offline (i.e., did not respond to probes) during this period. We also excluded measurements when we detected cross-traffic.




Figure 10: **Last-hop delay and jitter in cable and DSL networks:** DSL shows higher last-hop RTTs than cable, while cable exhibits higher jitter than DSL.

minimum RTT to the last-hop router from the minimum RTT to the broadband host. We used the minimum RTT estimates to avoid transient jitter as a result of queueing at intermediate routers.

Figure 10 shows our results for last-hop RTTs for cable and DSL networks. DSL hosts exhibit considerably higher propagation delays than cable hosts. 75% of all DSL hosts have last-hop delays larger than 10 ms, while 15% have propagation delays larger than 20 ms. These delays are surprising because many last-hop routers are located in the same city as their end hosts.⁴

Figure 10 also shows the jitter in our latency measurements. We used the RTTs of the small-TCP trickle to estimate the jitter of the broadband link. We calculated jitter by subtracting the 10th percentile RTT from the 90th percentile highest RTT. Compared to cable, DSL links have higher last-hop delays but lower jitter. We believe that the characteristics of the upstream cable links are responsible for these differences. We examine this next.

4.2.2 How do cable’s time-slotted policies affect transmission delays?

Transmission delay refers to time elapsed between a router starting to transmit a packet and ending its transmission. It is usually calculated by dividing the packet length by the link bandwidth. However, cable links use a reservation policy to transmit packets in the upstream direction. This policy can cause additional delays to a packet’s transmission. We examined the effects of such transmission policies under both low and high network loads.

First, we studied transmission delays under low network loads. We used the large-ICMP trickle to calculate the last-hop delays, similar to the experiment conducted in the previous section. We compared these last-hop large-packet delays to the last-hop small-packet delays measured in the earlier experiment. The differences in the last-hop delays between large (1,488-byte) and small (100-byte) packets are mostly due to the additional transmission delays incurred by sending larger packets.

Figure 11 shows the difference in transmission delays between large and small packets for cable and DSL hosts. We found that the transmission delays for DSL are large, on the same order of magnitude as their propagation delays, shown

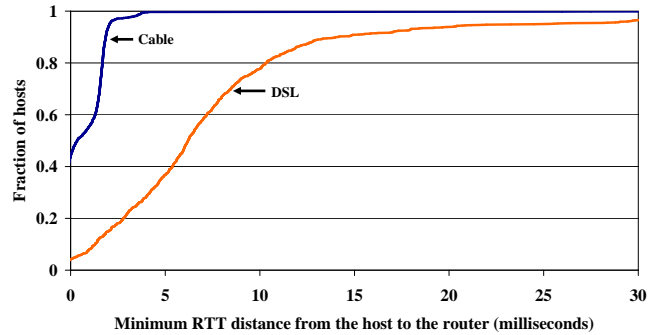


Figure 11: **Difference in transmission delays between large and small packets:** DSL shows longer transmission delays than cable.

in Figure 10. By contrast, the transmission delays for cable are surprisingly low – 99% of hosts show an increase of less than 1 ms to send an extra 1,388 bytes.

We believe that the time-slotted nature of cable links is responsible for these short transmission delays. All our probes, both large and small, experience similar waiting times for a time slot. When a slot has been granted, packets are transmitted at the full link speed (10.24 Mbps according to the DOCSIS 1.0 specification). This matches our data very well; our measured transmission delays correspond to an upstream link speed of 11 Mbps.


Next, we examined transmission delays under high network load. In this case, packets have to wait longer to reserve a time slot. When the reservation is granted, multiple waiting packets can be concatenated (see Section 2.1) and sent in a single burst. Although concatenation reduces the overhead of scheduling many small packets, such as TCP ACKs, it introduces a systematic jitter, which we refer to as the concatenation jitter.

We used the small-TCP flood to examine the effects of concatenation because it saturates the upstream link with a large number of small packets, which are well suited for concatenation. We clustered probe responses received in very close succession (separated by less than 100 μ s) as part of a single bursty transmission, and we calculated the number of packets in the largest cluster. Because there is no known concatenation feature for DSL, we expected these links to show only minimal burst sizes.


Figure 12(a) shows the extent of packet concatenation in DSL and cable ISPs. As expected, DSL links show only very short bursts, whereas 50% of cable links can concatenate 19 packets or more in a single burst. We used the link’s speed and the number of packets in a burst to estimate a lower bound on the amount of concatenation jitter when the link is saturated. Figure 12(b) shows the results. Whereas the mean concatenation jitter for cable networks is about 5 ms, many links experience 10 ms or more of jitter due to concatenation.

In cable networks, the concatenation jitter under high network load can be higher than the end-to-end jitter over the entire path under normal load (shown in Figure 10). The presence of high jitter in cable networks has important consequences for protocols such as TCP Vegas [9] and PCP [3], which interpret changes in RTT as a sign of incipient congestion. High jitter could cause these protocols to enter congestion avoidance too early, leading to poor performance.

⁴We inferred the locations of hosts and routers from their DNS names as suggested in [47].

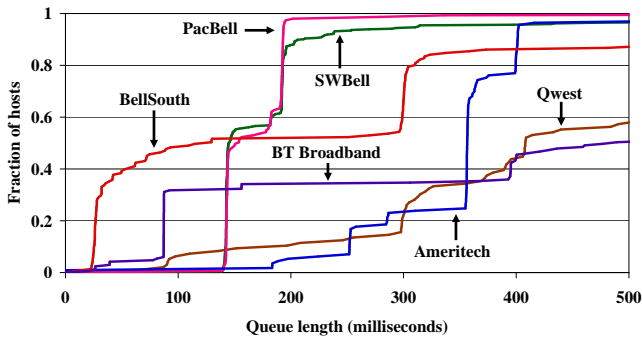


(a) Maximum number of packets per burst




(b) Lower bound estimate of concatenation jitter


Figure 12: **Cable links show high RTT variation:** In addition to a high level of basic jitter, cable modems can send small packets in a single burst and thus cause additional jitter.



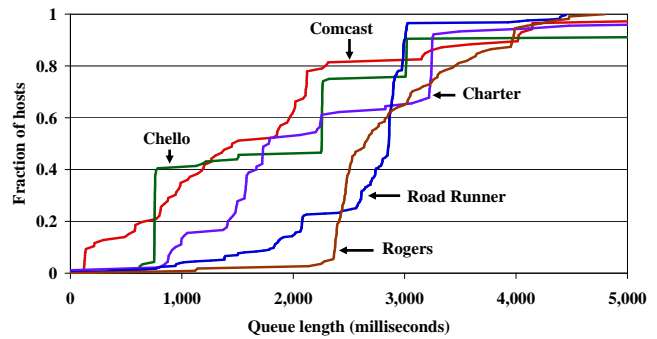
(a) DSL (downstream)



(b) Cable (downstream)



(c) DSL (upstream)



(d) Cable (upstream)

Figure 13: **Downstream and upstream queue length in milliseconds:** Some downstream queue lengths follow the recommendation for voice calls (150 ms), but most are significantly longer. The upstream queue length can be massive, especially for cable links.

4.2.3 How large are broadband queuing delays?

Sizing router queues is a popular area of research (e.g., [4]). A common rule of thumb (attributed to [49]) suggests that router queues' lengths should be equal to the RTT of an average flow through the link. Larger queues lead to needlessly high queuing delays in the network. We investigated how well this conventional wisdom holds in broadband environments.

We measured queue lengths in milliseconds by calculating the RTT variation of our probe streams' packets. To estimate downstream queue lengths, we used large-TCP flood probe trains, which saturate the downstream but not the up-

stream link. We calculated the difference between the minimum RTT and the 95th percentile highest RTT. To estimate upstream queue lengths, we first measured the difference between the minimum RTT and the 95th percentile highest RTTs of large-ICMP flood probe trains. This difference corresponds to the sum of downstream and upstream queue lengths. We then subtracted the estimate of the downstream queue length to obtain the length of the upstream queue.

Figures 13(a) and 13(b) show the cumulative distributions of downstream queue lengths for different cable and DSL providers. Across most cable ISPs and two DSL ISPs (PacBell and SWBell), the curves show a sharp rise at 130 ms. This value is consistent with that recommended by the ITU

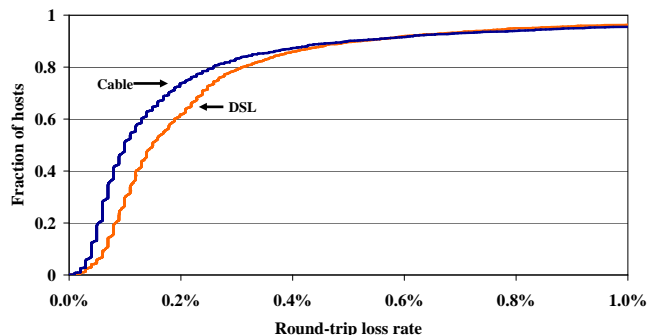


Figure 14: **Observed round-trip loss rate for residential broadband paths:** DSL and cable paths show similar loss rates. 95% of all DSL and cable hosts have loss rates of less than 1%.

G.114 standard for maximum end-to-end latency in a network running interactive traffic – 150 ms. Nevertheless, these queue lengths are significantly higher than a typical flow’s average delay, which ranges between 50 and 75 ms within North America or Europe. By contrast, we observed queuing delays of up to 2 s for a significant number of Comcast and Qwest hosts and up to 6 s for some BT Broadband hosts (not shown). Our findings show diverse queue configurations for broadband links, with most hosts exhibiting queue lengths significantly higher than 130 ms.

Figures 13(c) and 13(d) show the cumulative distributions of upstream queue lengths for the different cable and DSL providers. Compared with downstream queues, the lengths of upstream queues are very large. Most DSL links exhibit queues of 600 ms or higher, and many cable links allow their upstream queues to grow to several seconds. Although some of the upstream queues’ build-up results from the low upstream link bandwidths, the excessive lengths will negatively affect interactive traffic like VoIP whenever users upload content, such as when using BitTorrent.

4.3 Packet loss

In this section, we characterize packet loss in residential broadband networks. We contrast our results with the commonly held idea that broadband networks have high packet loss rates. Our tools cannot measure the access links’ loss rates. Instead, we examined the packet loss rates of the Internet paths between our well-connected measurement hosts and the broadband hosts. Because the broadband access links are part of these Internet paths, our measured loss rates provide an upper bound on the broadband links’ loss rates.

4.3.1 Do broadband links see high packet loss?

We used the small-TCP trickle probe trains to calculate the loss rates along the round-trip paths to remote broadband hosts. We sent widely spaced trickle probes at a very low rate for a week, and we measured the fraction of probes for which the broadband hosts did not respond. This includes losses on both the upstream and the downstream paths, and it measures the loss rate under normal operating conditions of the network. Note that the loss rate we measured might differ from the loss rate that application traffic (e.g., TCP flows) saturating broadband links would suffer.

Figure 14 presents our results. We found that both cable and DSL have remarkably low packet loss rates. The

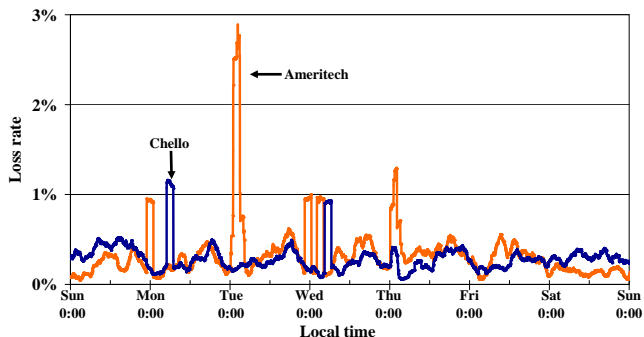


Figure 15: **Packet loss over time:** The loss rate is generally low and shows heavy diurnal variations with intermittent spikes. Note that this graph includes both upstream and downstream losses; the time axis shows local time (EST for Ameritech and CET for Chello).

loss rate is below 1% for more than 95% of all DSL and cable paths. Overall, we found that the packet loss rates for broadband access networks are similar to those observed in academic network environments [11, 40].

We also examined how loss rates varied over the course of the week. Figure 15 shows our measurements for two typical providers: a DSL ISP (Ameritech) and a cable ISP (Chello). The horizontal axis shows the local time for the ISPs. The loss rates shown along the vertical axis are averaged over intervals of 120 minutes. We found that loss rates exhibit diurnal patterns with occasional spikes. Both ISPs follow similar diurnal patterns, showing lower loss rates in the early morning than in the evening.

4.3.2 Do ISPs use active queue management?

When packets are sent very quickly, they begin to fill up queues, and the routers must eventually drop some of the packets. The most common queue management policy is tail-drop – i.e., all packets arriving after the queue is full are discarded. More active queue management policies, such as RED [20], proactively drop packets using probabilistic schemes when the queue starts to fill up but before the queue is full. Active queue management has been extensively studied, but relatively little is known about the extent to which it is deployed in practice.

We performed the following experiment to infer whether the broadband ISPs are using active queue management policies. We used the small-TCP flood to overflow both downstream and upstream links, and we used IPIDs to distinguish between losses occurring upstream and those occurring downstream [36]. For each successfully received response, we recorded the RTT, and we calculated the average loss rate over a sliding window of 40 packets. We examined the correlation between the loss rates and the corresponding RTTs. On the basis of this correlation, we can infer whether routers use tail-drop or more active queue management policies. A tail-drop policy will result in a steep increase in loss rate when the queue is full (i.e., for a large RTT value); if an active queue management policy such as RED is used, then the loss rate will increase proportionally to the RTT after a certain threshold.

Figure 16 shows how the loss rates increase with the RTT for two broadband hosts, one in PacBell and one in SWBell. For the PacBell host, the loss rate increases steeply around




Figure 16: **Tail-drop and active queue management:** When a tail-drop queue overflows, the loss rate increases sharply. If the loss rate increases proportionally to the queue length after a threshold, then this suggests that active queue management (probably RED) is being used.

an RTT of 850 ms, which suggests that a tail-drop queue is used. The loss rate for the SWBell host shows a different trend; after 500 ms, it increases almost linearly with the RTT before stabilizing at around 85%. This behavior matches the description of the RED active queue management policy.

To quantify the extent of RED deployment in broadband networks, we tested whether the increases in RTT and loss rates are strongly correlated. If the correlation coefficient is high (≥ 0.9) beyond a threshold loss rate of 5%, we conclude that the link may be using RED as its drop policy. We did not calculate the correlation coefficient for low loss rates (below 5%) because these losses might be sporadic and not representative of the broadband router’s queue policy.

We found that 26.2% of the DSL hosts show a RED-style drop policy on their upstream queues. The three providers owned by AT&T (i.e., Ameritech, BellSouth, and PacBell) exhibit deployment rates between 50.3% and 60.5%, whereas all other DSL providers’ deployment rates are below 23.0%. The partial deployment of RED-style policies within ISPs could be due to heterogeneity in the ISPs’ equipment. We did not detect RED in any of the cable ISPs measured.

4.4 Summary

We have presented an in-depth characterization of the properties of residential broadband networks. Our analysis reveals important ways in which these networks differ from academic networks, and it quantifies these differences. We summarize our key findings below:

Allocated link bandwidths: Our results show that downstream bandwidths exceed upstream bandwidths by more than a factor of 10 for some ISPs. In contrast to popular belief, for most ISPs, the measured bandwidths matched well with the advertised rates at all times of day, and we found little evidence of competing traffic affecting their links. Although link bandwidths remain stable over the long term, they show high variation in the short term, especially for cable hosts. For some ISPs, link bandwidths change abruptly as a result of traffic shaping.

Packet latencies: Many DSL hosts show large (≥ 10 ms) last-hop propagation delays. Cable hosts suffer higher jitter than DSL hosts as a result of time-slotted packet transmission policies on their upstream links. Packet concatenation on the upstream links can add another 5 – 10 ms of jitter in cable links.

All ISPs deploy queues that are several times larger than their bandwidth-delay products. Whereas downstream queues can delay packets by more than 100 ms, the upstream queueing delays can exceed several hundreds of milliseconds and, at times, a few seconds.

Packet loss: Both DSL and cable ISPs exhibit surprisingly low packet loss. We also found that many DSL hosts use active queue management policies (e.g., RED) when dropping packets.

5. IMPLICATIONS OF OUR FINDINGS

We consider that our observations about broadband networks’ characteristics can help researchers to understand how well existing protocols and systems work in the commercial Internet. Our findings offer useful insights for the designers of future applications. To illustrate this, we briefly discuss the potential implications of our measurements for three popular Internet-scale systems.

Transport Control Protocols: Our bandwidth and latency findings have several implications for transport protocol designs. For example, protocols such as TCP Vegas [9] and PCP [3] use RTT measurements to detect incipient congestion. In the presence of the high jitter found in our measurements, this mechanism might trigger congestion avoidance too early. Bandwidth-probing techniques, such as packet-pair [31], could return incorrect results in the presence of traffic shaping or packet concatenation. This could be detrimental to transport protocols that rely on probing to adjust their transfer rates, such as PCP.

Network coordinate and location systems: Many IP-to-geolocation mapping tools [22, 52] use latency measurements to determine a host’s location. The large propagation delays and high jitter found in broadband networks are likely to seriously interfere with the accuracy of these systems.

Similarly, network coordinate systems [16, 37] use latency estimates to assign a set of coordinates to their participating hosts. A recent study [34] found that network coordinate systems do not perform well when deployed in BitTorrent networks, because RTTs between nodes can vary by up to four orders of magnitude. Our measurements explain and provide insights into these findings: BitTorrent networks typically include many residential links, which have very large RTT variations as a result of their long queues. BitTorrent traffic compounds these variations because it tends to fill up the queues.

Interactive and real-time applications: Recently, the popularity of VoIP and online games has grown considerably. Our data shows that real-time applications will be negatively affected by the broadband links’ large queueing delays. Because queueing delays increase in the presence of competing traffic, these time-sensitive applications are likely to experience degraded service when they are used concurrently with bandwidth-intense applications, such as BitTorrent.

6. RELATED WORK

There is a large body of previous measurement work characterizing Internet paths other than broadband. Paxson [40] studied network packet dynamics among a fixed set of Internet hosts located primarily in academic institutions. More recently, several studies have examined the characteristics of the network paths connecting the PlanetLab testbed [5, 43].

Although our paper uses similar measurement techniques, the network environment we study is different.

Compared with other parts of the Internet, broadband access networks have received relatively little attention. Claypool et al. [14] performed a measurement study of access networks' queue sizes using 47 volunteering broadband hosts. They found that the median queue size was 350 ms in DSL networks and 150 ms in cable networks, and they showed in simulation that large queue sizes are detrimental to network traffic from interactive applications. Our results are consistent with these earlier findings, but they are based on a set of hosts that is more than two orders of magnitude larger. Similarly, Jehaes et al. [30] observed a large increase in round-trip delays over saturated broadband links. Their experiments were limited to one DSL and one cable link.

Some recent studies have examined the traffic generated by residential customers in Japan [12] and France [46]. These results complement ours, because we examined the properties of the networks themselves rather than the properties of traffic traversing them. A comprehensive view of residential networks requires a good understanding of both.

Lakshminarayanan and Padmanabhan [32] performed a network measurement study from 25 broadband hosts to different Internet hosts, covering several application-level metrics such as TCP throughput and latency. Our study confirms some of their findings but at a much larger scale. In their later work, Lakshminarayanan et al. [33] outlined pitfalls in measuring link capacities of cable and DSL networks by using existing bandwidth estimation tools. In particular, the accuracy of these tools is greatly influenced by the rate regulation schemes used in cable and DSL networks. Our measurement methodology does not suffer from such inaccuracies because it relies on saturating the links for a short duration.

Many previous studies have measured network properties of hosts participating in file-sharing peer-to-peer systems [7, 44, 45]. Because a large fraction of peers in these systems access the Internet over cable and DSL networks, all these studies indirectly include measurements of broadband access networks. However, these results cannot be compared directly with ours because the focus of these studies is primarily on application-level performance and not on the link level characteristics of broadband networks.

7. CONCLUSIONS

In this paper, we presented the first large-scale measurement study of major cable and DSL providers in North America and Europe. Our study characterized several important characteristics of broadband networks, including available link capacities, packet transmission policies, jitter, packet drop policies, and queue lengths. Our analysis revealed important ways in which residential networks differ from the conventional wisdom about the Internet. We also discussed the implications of our findings for many emerging protocols and systems, such as delay-based congestion control (e.g., PCP) and network coordinate systems (e.g., Vivaldi).

8. ACKNOWLEDGMENTS

We thank Steve Gribble and Dan Sandler for generously hosting measurement servers for our experiments. We are grateful to our shepherd Guillaume Urvoy-Keller, Peter Druschel, and our anonymous reviewers for providing detailed and helpful feedback on this paper.

9. REFERENCES

- [1] A. Akella, S. Seshan, and A. Shaikh. An empirical evaluation of wide-area Internet bottlenecks. In *Proceedings of IMC'03*, Oct 2003.
- [2] D. G. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris. Experience with an Evolving Overlay Network Testbed. *ACM Computer Communication Review*, 33(3):13–19, July 2003.
- [3] T. Anderson, A. Collins, A. Krishnamurthy, and J. Zahorjan. PCP: Efficient endpoint congestion control. In *Proceedings of NSDI'06*, San Jose, CA, May 2006.
- [4] G. Appenzeller, I. Keslassy, and N. McKeown. Sizing router buffers. In *Proceedings of ACM SIGCOMM'04*, Aug 2004.
- [5] S. Banerjee, T. G. Griffin, and M. Pias. The interdomain connectivity of PlanetLab nodes. In *Proceedings of the 5th Passive and Active Measurement Conference (PAM)*, Apr 2004.
- [6] J. Bellardo and S. Savage. Measuring packet reordering. In *Proceedings of the 2nd ACM SIGCOMM Internet Measurement Workshop (IMW)*, Marseille, France, Nov 2002.
- [7] R. Bhagwan, S. Savage, and G. M. Voelker. Understanding availability. In *Proceedings of 1st International Workshop on Peer-to-Peer Systems (IPTPS)*, Boston, MA, March 2002.
- [8] Bittorrent. <http://bittorrent.com>.
- [9] L. S. Brakmo and L. Peterson. TCP Vegas: End to end congestion avoidance on a global Internet. *IEEE Journal on Selected Areas in Communication*, 13(8):1465–1480, Oct 1995.
- [10] CableLabs. DOCSIS 1.1 interface specification, 2006. <http://www.cablemodem.com/specifications/specifications11.html>.
- [11] B. Chandra, M. Dahlin, L. Gao, and A. Nayate. End-to-end WAN service availability. In *Proceedings of the 3rd USITS*, 2001.
- [12] K. Cho, K. Fukuda, H. Esaki, and A. Kato. The impact and implications of the growth in residential user-to-user traffic. In *Proceedings of SIGCOMM'06*, Sep 2006.
- [13] Cisco. Internetworking technology handbook, 2006.
- [14] M. Claypool, R. Kinicki, M. Li, J. Nichols, and H. Wu. Inferring queue sizes in access networks by active measurement. In *Proceedings of the 4th Passive and Active Measurement Workshop (PAM)*, Antibes Juan-les-Pins, France, Apr 2004.
- [15] Comcast PowerBoost press release, Jun 2006. <http://www.cmcsk.com/phoenix.zhtml?c=147565&p=irol-newsArticle&ID=890297>.
- [16] F. Dabek, R. Cox, F. Kaashoek, and R. Morris. Vivaldi: A decentralized network coordinate system. In *Proceedings of SIGCOMM'04*, Aug 2004.
- [17] C. Dovrolis, P. Ramanathan, and D. Moore. Packet dispersion techniques and a capacity estimation methodology. *IEEE/ACM Transactions on Networking*, Dec 2004.
- [18] Federal Communications Commission ComCom, Switzerland. Broadband in the universal service, Sep 2006.
- [19] S. Floyd, M. Handley, J. Padhye, and J. Widmer. Equation-based congestion control for unicast applications. In *Proceedings of SIGCOMM'00*, Aug 2000.
- [20] S. Floyd and V. Jacobson. Random early detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking*, 1(4):397–413, Aug 1993.
- [21] R. Govindan and V. Paxson. Estimating router ICMP generation delays. In *Proceedings of Passive and Active Measurement (PAM)*, Fort Collins, CO, USA, 2002.

- [22] B. Gueye, A. Ziviani, M. Crovella, and S. Fdida. Constraint-based geolocation of Internet hosts. *IEEE/ACM Transactions on Networking*, 14(6):1219–1232, 2006.
- [23] A. Haeberlen, M. Dischinger, K. P. Gummadi, and S. Saroiu. Monarch: A tool to emulate transport protocol flows over the Internet at large. In *Proceedings of IMC'06*, Rio de Janeiro, Brazil, Oct 2006.
- [24] Y. hua Chu, A. Ganjam, T. S. E. Ng, S. G. Rao, K. Sripanidkulchai, J. Zhan, and H. Zhang. Early experience with an Internet broadcast system based on overlay multicast. In *USENIX Annual Technical Conference*, Jun 2004.
- [25] Y. hua Chu, S. G. Rao, S. Seshan, and H. Zhang. Enabling conferencing applications on the Internet using an overlay multicast architecture. In *Proceedings of SIGCOMM'01*, Aug 2001.
- [26] G. Iannaccone, C. nee Chuah, R. Mortier, S. Bhattacharyya, and C. Diot. Analysis of link failures in an IP backbone. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurement (IMW)*, 2002.
- [27] ISP Planet. Top 24 U.S. ISPs by subscriber: Q3 2005, 2006. <http://www.isp-planet.com/research/rankings/usa.html>.
- [28] ISP Planet. Top seven ISPs in Canada by subscriber: Q4 2006, 2006. http://www.isp-planet.com/research/rankings/2006/canada_q42006.html.
- [29] M. Jain and C. Dovrolis. End-to-end available bandwidth: Measurement methodology, dynamics, and relation with TCP throughput. *IEEE/ACM Transactions on Networking*, Aug 2003.
- [30] T. Jehaes, D. D. Vleeschauwer, T. Coppens, B. V. Doorselaer, E. Deckers, W. Naudts, K. Spruyt, and R. Smets. Access network delay in networked games. In *Proceedings of the 2nd workshop on Network and system support for games (NetGames)*, 2003.
- [31] S. Keshav. A control-theoretic approach to flow control. In *Proceedings of SIGCOMM'91*, Zurich, Switzerland, Sept 1991.
- [32] K. Lakshminarayanan and V. N. Padmanabhan. Some findings on the network performance of broadband hosts. In *Proceedings of IMC'03*, Miami, FL, USA, Oct 2003.
- [33] K. Lakshminarayanan, V. N. Padmanabhan, and J. Padhye. Bandwidth estimation in broadband access networks. In *Proceedings of IMC'04*, Taormina, Italy, Oct 2004.
- [34] J. Ledlie, P. Gardner, and M. Seltzer. Network coordinates in the wild. In *Proceedings of NSDI'07*, Apr 2007.
- [35] S.-J. Lee, P. Sharma, S. Banerjee, S. Basu, and R. Fonseca. Measuring bandwidth between PlanetLab nodes. In *Proceedings of the Passive and Active Measurement Workshop (PAM)*, March 2005.
- [36] R. Mahajan, N. Spring, D. Wetherall, and T. Anderson. User-level Internet path diagnosis. In *Proceedings of the 19th Symposium on Operating Systems Principles (SOSP)*, Bolton Landing, NY, USA, Oct 2003.
- [37] T. Ng and H. Zhang. Predicting Internet network distance with coordinates-based approaches. In *Proc. INFOCOM 2002*, New York, NY, USA, June 2002.
- [38] Nielsen/NetRatings. U.S. broadband connections reach critical mass, 2004. http://www.nielsen-netratings.com/pr/pr_040818.pdf.
- [39] OECD broadband statistics, Dec 2005. <http://www.oecd.org/sti/ict/broadband>.
- [40] V. Paxson. End-to-end routing behavior in the Internet. *IEEE/ACM Transactions on Networking*, 5(5):601–615, Oct 1997.
- [41] PlanetLab. <http://www.planet-lab.org/>.
- [42] Point Topic Ltd. UK broadband market monitor, 2006. <http://point-topic.com/content/ukplus/email%20archive/wukbbmm7050602.htm>.
- [43] H. Pucha, Y. C. Hu, and Z. M. Mao. On the impact of research network based testbeds on wide-area experiments. In *Proceedings of IMC'06*, Oct 2006.
- [44] S. Saroiu, K. P. Gummadi, and S. D. Gribble. A measurement study of peer-to-peer file sharing systems. In *Proceedings of the Multimedia Computing and Networking (MMCN)*, San Jose, CA, Jan 2002.
- [45] S. Sen and J. Wang. Analyzing peer-to-peer traffic across large networks. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet Measurement (IMW)*, Marseille, France, Nov 2002.
- [46] M. Siekkinen, D. Collange, G. Urvoy-Keller, and E. Biersack. Performance limitations of ADSL users: A case study. In *Proceedings of the 8th Passive and Active Measurement Conference (PAM)*, Apr 2007.
- [47] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson. Measuring ISP topologies with Rocketfuel. *IEEE/ACM Transactions on Networking*, 12(1):2–16, Feb 2004.
- [48] United Kingdom e-Minister and e-Envoy. UK online: The broadband future, 2001.
- [49] C. Villamizar and C. Song. High performance TCP in ANSNET. *ACM Computer Communication Review*, 24(5):45–60, 1994.
- [50] Wikipedia. Broadband Internet access in the Netherlands, 2006. http://en.wikipedia.org/w/index.php?title=Broadband_Internet_access_in_the_Netherlands&oldid=152236688.
- [51] Windsor Oaks Group, LLC. Market outlook report, May 2005. http://www.broadbandtrends.com/Report%20Summary/2006/BBT_GlobalBBOutlook2006_061110_TOC.pdf.
- [52] B. Wong, I. Stoyanov, and E. G. Sirer. Octant: A comprehensive framework for the geolocalization of Internet hosts. In *Proceedings of NSDI'07*, Apr 2007.