

Broadband Internet Performance: A View From the Gateway

Srikanth Sundaresan
Georgia Tech
Atlanta, USA
srikanth@gatech.edu

Walter de Donato
University of Napoli Federico II
Napoli, Italy
walter.dedonato@unina.it

Nick Feamster
Georgia Tech
Atlanta, USA
feamster@cc.gatech.edu

Renata Teixeira
CNRS/UPMC Sorbonne Univ.
Paris, France
renata.teixeira@lip6.fr

Sam Crawford
SamKnows
London, UK
sam@samknows.com

Antonio Pescapè
University of Napoli Federico II
Napoli, Italy
pescapè@unina.it

ABSTRACT

We present the first study of network access link performance measured directly from home gateway devices. Policymakers, ISPs, and users are increasingly interested in studying the performance of Internet access links. Because of many confounding factors in a home network or on end hosts, however, thoroughly understanding access network performance requires deploying measurement infrastructure in users' homes as gateway devices. In conjunction with the Federal Communication Commission's study of broadband Internet access in the United States, we study the throughput and latency of network access links using longitudinal measurements from nearly 4,000 gateway devices across 8 ISPs from a deployment of over 4,200 devices. We study the performance users achieve and how various factors ranging from the user's choice of modem to the ISP's traffic shaping policies can affect performance. Our study yields many important findings about the characteristics of existing access networks. Our findings also provide insights into the ways that access network performance should be measured and presented to users, which can help inform ongoing broader efforts to benchmark the performance of access networks.

Categories and Subject Descriptors

C.2.3 [Computer-Communication Networks]: Network Operations—*Network Management*; C.2.3 [Computer-Communication Networks]: Network Operations—*Network Operations*

General Terms

Management, Measurement, Performance

Keywords

Access Networks, Broadband Networks, BISMark, Benchmarking

1. INTRODUCTION

Of nearly two billion Internet users worldwide, about 500 million are residential broadband subscribers [19]. Broadband penetration is likely to increase further, with people relying on home connectivity for day-to-day and even critical activities. Accordingly, the Federal Communication Commission (FCC) is actively

developing performance-testing metrics for access providers [5, 15, 35]. Policymakers, home users, and Internet Service Providers (ISPs) are in search for better ways to benchmark home broadband Internet performance.

Benchmarking home Internet performance, however, is not as simple as running one-time “speed tests”. There exist countless tools to measure Internet performance [7, 14, 29, 32]. Previous work has studied the typical download and upload rates of home access networks [12, 24]; others have found that modems often have large buffers [24], and that DSL links often have high latency [26]. These studies have shed some light on access-link performance, but they have typically run one-time measurements either from an end-host inside the home (from the “inside out”) or from a server on the wide-area Internet (from the “outside in”). Because these tools run from end-hosts, they cannot analyze the effects of confounding factors such as home network cross-traffic, the wireless network, or end-host configuration. Also, many of these tools run as one-time measurements. Without continual measurements of the same access link, these tools cannot establish a baseline performance level or observe how performance varies over time.

This paper measures and characterizes broadband Internet performance from home gateways. The home gateway connects the home network to the user's modem; taking measurements from this vantage point allows us to control the effects of many confounding factors, such as the home wireless network and load on the measurement host (Section 4). The home gateway is always on; it can conduct unobstructed measurements of the ISP's network and account for confounding factors in the home network. The drawback to measuring access performance from the gateway, of course, is that deploying gateways in many homes is incredibly difficult and expensive. Fortunately, we were able to take advantage of the ongoing FCC broadband study to have such a unique deployment.

We perform our measurements using two complementary deployments; the first is a large FCC-sponsored study, operated by SamKnows, that has installed gateways in over 4,200 homes across the United States, across many different ISPs. The second, BISMark, is deployed in 16 homes across three ISPs in Atlanta. The SamKnows deployment provides a large user base, as well as diversity in ISPs, service plans, and geographical locations. We designed BISMark to allow us to access the gateway remotely and run repeated experiments to investigate the effect of factors that we could not study in a larger “production” deployment. For example, to study the effect of modem choice on performance, we were able to install different modems in the same home and conduct experiments in a controlled setting. Both deployments run a comprehensive suite of measurement tools that periodically measure throughput, latency, packet loss, and jitter.

We characterize access network throughput (Section 5) and la-

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.

SIGCOMM'11, August 15–19, 2011, Toronto, Ontario, Canada.
Copyright 2011 ACM 978-1-4503-0797-0/11/08 ...\$10.00.

tency (Section 6) from the SamKnows and BISMark deployments. We explain how our throughput measurements differ from common “speed tests” and also propose several different latency metrics. When our measurements cannot fully explain the observed behavior, we model the access link and verify our hypotheses using controlled experiments. We find that the most significant sources of throughput variability are the access technology, ISPs’ traffic shaping policies, and congestion during peak hours. On the other hand, latency is mostly affected by the quality of the access link, modem buffering, and cross-traffic within the home.

This study offers many insights into both access network performance and the appropriate measurement methods for benchmarking home broadband performance. Our study has three high-level lessons, which we expand on in Section 7:

- ISPs use different policies and traffic shaping behavior that can make it difficult to compare measurements across ISPs.
- There is no “best” ISP for all users. Different users may prefer different ISPs depending on their usage profiles and how those ISPs perform along performance dimensions that matter to them.
- A user’s home network equipment and infrastructure can significantly affect performance.

As the first in-depth analysis of home access network performance, our study offers insights for users, ISPs, and policymakers. Users and ISPs can better understand the performance of the access link, as measured directly from the gateway; ultimately, such a deployment could help an ISP differentiate performance problems within the home from those on the access link. Our study also informs policy by illustrating that a diverse set of network metrics ultimately affect the performance that a user experiences. The need for a benchmark is clear, and the results from this study can serve as a principled foundation for such an effort.

2. RELATED WORK

This section presents related work; where appropriate, we compare our results to these previous studies in Sections 5 and 6.

From access ISPs. Previous work characterizes access networks using passive traffic measurements from DSL provider networks in Japan [8], France [33], and Europe [26]. These studies mostly focus on traffic patterns and application usage, but they also infer the round-trip time and throughput of residential users. Without active measurements or a vantage point within the home network, however, it is not possible to measure the actual performance that users receive from their ISPs, because user traffic does not always saturate the user’s access network connection. For example, Siekkinen *et al.* [33] show that applications (*e.g.*, peer-to-peer file sharing applications) often rate limit themselves, so performance observed through passive traffic analysis may reflect application rate limiting, as opposed to the performance of the access link.

From servers in the wide area. Other studies have characterized access network performance by probing access links from servers in the wide area [11, 12]. Active probing from a fixed set of servers can characterize many access links because each link can be measured from the same server. Unfortunately, because the server is often located far from the access network, the measurements may be inaccurate or inconsistent. Isolating the performance of the access network from the performance of the end-to-end path can be challenging, and dynamic IP addressing can make it difficult to determine whether repeated measurements of the same IP address are in fact measuring the same access link over time. A remote server

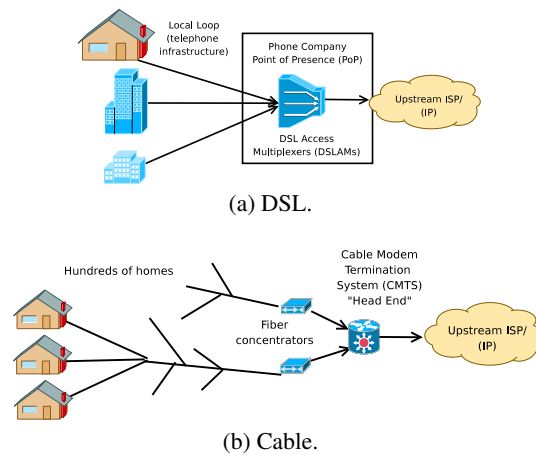


Figure 1: Access network architectures.

also cannot isolate confounding factors, such as whether the user’s own traffic is affecting the access-link performance.

From inside home networks. The Grenouille project in France [1] measures the performance of access links using a monitoring agent that runs from a user’s machine inside the home network. Neti@Home [23] and BSense [4] also use this approach, although these projects have fewer users than Grenouille. PeerMetric [25] measured P2P performance from about 25 end hosts. Installing software at the end-host measures the access network from the user’s perspective and can also gather continuous measurements of the same access link. Han *et al.* [18] measured access network performance from a laptop that searched for open wireless networks. This approach is convenient because it does not require user intervention, but it does not scale to a large number of access networks, cannot collect continuous measurements, and offers no insights into the specifics of the home network configuration.

Other studies have performed “one-time” measurements of access-link performance. These studies typically help users troubleshoot performance problems by asking the users to run tests from a Web site and running analysis based on these tests. Netalyzer [29] measures the performance of commonly used protocols using a Java applet that is launched from the client’s browser. Network Diagnostic Tool (NDT) [7] and Network Path and Application Diagnostics (NPAD) [14] send active probes to detect issues with client performance. Glasnost performs active measurements to determine whether the user’s ISP is actively blocking BitTorrent traffic [17]. Users typically run these tools only once (or, at most, a few times), so the resulting datasets cannot capture a longitudinal view of the performance of any single access link. In addition, any technique that measures performance from a device inside the home can be affected by factors such as load on the host or features of the home network (*e.g.*, cross-traffic, wireless signal strength). Finally, none of these studies measure the access link directly from the home network gateway.

3. ACCESS NETWORKS: BACKGROUND

We describe the two most common access technologies from our deployments: Digital Subscriber Line (DSL) and cable. Then, we explain how a user’s choice of service plan and local configuration can affect performance. Although a few users in our deployments have fiber-to-the-node (FTTN), fiber-to-the-premises (FTTP), and WiMax, we do not have enough users to analyze these technologies.

DSL networks use telephone lines; subscribers have dedicated lines between their own DSL modems and the closest DSL Access Multiplexer (DSLAM). The DSLAM multiplexes data between the access modems and upstream networks, as shown in Figure 1a. The most common type of DSL access is asymmetric (ADSL), which provides different upload and download rates. In cable access networks, groups of users send data over a shared medium (typically coaxial cable); at a regional *headend*, a Cable Modem Termination System (CMTS) receives these signals and converts them to Ethernet, as shown in Figure 1b. The physical connection between a customer’s home and the DSLAM or the CMTS is often referred to as the *local loop* or *last mile*. Users buy a service plan from a provider that typically offers some *maximum* capacity in both the upload and download directions.

ADSL capacity. The ITU-T standardization body establishes that the achievable rate for ADSL 1 [20] is 12 Mbps downstream and 1.8 Mbps upstream. The ADSL2+ specification [21] extends the capacity of ADSL links to at most 24 Mbps download and 3.5 Mbps upload. Although the ADSL technology is theoretically able to reach these speeds, there are many factors that limit the capacity in practice. An ADSL modem negotiates the operational rate with the DSLAM (often called the *sync rate*); this rate depends on the quality of the local loop, which is mainly determined by the distance to the DSLAM from the user’s home and noise on the line. The maximum IP link capacity is lower than the sync rate because of the overhead of underlying protocols. The best service plan that an ADSL provider advertises usually represents the rate that customers can achieve if they have a good connection to the DSLAM. Providers also offer service plans with lower rates and can rate-limit a customer’s traffic at the DSLAM.

Modem configuration can also affect performance. ADSL users or providers configure their modems to operate in either *fastpath* or *interleaved* mode. In *fastpath* mode, data is exchanged between the DSL modem and the DSLAM in the same order that they are received, which minimizes latency but prevents error correction from being applied across frames. Thus, ISPs typically configure fast-path only if the line has a low bit error rate. *Interleaving* increases robustness to line noise at the cost of increased latency by splitting data from each frame into multiple segments and interleaving those segments with one another before transmitting them.

Cable capacity. In cable networks, the most widely deployed version of the standard is Data Over Cable Service Interface Specification version 2 (DOCSIS 2.0) [22], which specifies download rates up to 42.88 Mbps and upload rates up to 30.72 Mbps in the United States. The latest standard, DOCSIS 3.0, allows for hundreds of megabits per second by bundling multiple channels. Cable providers often offer service plans with lower rates. The service plan rate limit is configured at the cable modem and is typically implemented using a token bucket rate shaper. Many cable providers offer *PowerBoost*, which allows users to download (and, in some cases, upload) at rates that are higher than the contracted ones, for an initial part of a transfer. The actual rate that a cable user receives will vary with the network utilization of other users connecting to the same headend. The CMTS controls the rate at which cable modems transmit. For instance, Comcast describes that when a CMTS’s port becomes congested, it ensures fairness by scheduling heavy users on a lower priority queue [3].

4. MEASUREMENT INFRASTRUCTURE

We describe the measurement infrastructure that we deployed and the datasets that we collected. We first motivate the need for

Factor	How we address it
Wireless Effects	Use a wired connection to modem.
Cross Traffic	Measure cross traffic and avoid it/account for it.
Load on gateway	Use a well-provisioned gateway.
Location of server	Choose a nearby server.
End-to-end path	Focus on characterizing the last mile.
Gateway configuration	Test configuration in practice and controlled settings.

Table 1: Confounding factors and how we address them.




Figure 2: Our gateway device sits directly behind the modem in the home network. They take measurements both to the *last mile router* (first non-NAT IP hop on the path) and to wide area hosts.

deploying measurement infrastructure directly at the gateway; then, we describe the *SamKnows* and *BISMark* (Broadband Internet Service benchMark) gateway deployments.

4.1 Why a Gateway?

Deploying measurements at gateway devices offers the following advantages over the other techniques discussed in Section 2:

- *Direct measurement* of the ISP’s access link: the gateway sits behind the modem; between the access link and all other devices at the home network as shown in Figure2. This allows us to isolate the effect of confounding factors such as wireless effects and cross traffic.
- *Continual/longitudinal measurements*, which allow us to meaningfully characterize performance of ISPs for individual users.
- *The ability to instrument a single home with different hardware and configurations*, which allows us to explore the effects of multiple factors on performance. In some deployments, we were even able to swap modems to study their effect on performance, holding all other conditions about the network setup equal.

Table 1 summarizes the challenges involved in conducting such a study, and how deploying gateways solves them. We now describe the two gateway deployments in our study.

4.2 Gateway Deployments

Our study uses two independent gateway deployments. The first, the FCC/SamKnows gateway deployment, collected data from over 4,200 users across different ISPs in the United States, as of January 2011. This deployment currently has over 10,000 users. Our goal in using the measurements from this deployment is to achieve *breadth*: we aim to classify a large set of users across a diverse set of ISPs and geographical locations. The second, the BISMark deployment, collects measurements from a smaller, focused group of users from different ISPs and service plans in Atlanta. Our goal with the measurements from this deployment is to achieve *depth*: this platform allows us to take measurements with detailed knowledge of how every gateway is deployed; we can also take repeated measurements and conduct specific experiments from the same deployment with different settings and configurations.

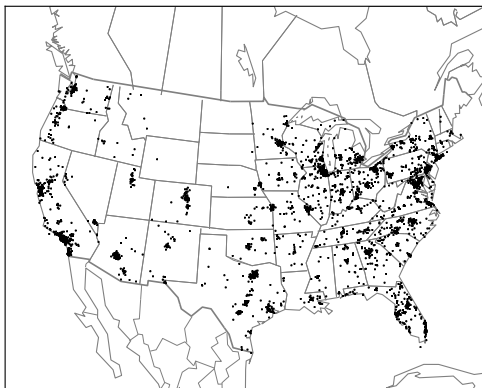


Figure 3: SamKnows deployment: 4,200 devices, 16 ISPs.

Gateway deployments entail significant challenges concerning the resource constraints of the gateway platform and the need to remotely maintain and manage the devices (especially because these devices are deployed in homes of “real users”); we omit discussion of these logistical challenges due to lack of space and instead focus on the details of the platforms and the measurements we collect.

4.2.1 SamKnows

SamKnows specializes in performance evaluation of access networks; it has studied access ISP performance in the United Kingdom and has now contracted with the FCC for a similar study in the United States. SamKnows deployed gateways in each participant’s home either directly behind the home user’s router or behind the home wireless router; the devices can be updated and managed remotely. The gateway is a Netgear WNR3500L RangeMax Wireless-N Gigabit router with a 480 MHz MIPS processor, 8 MB of flash storage, and 64 MB of RAM. We use active measurement data from the SamKnows study from December 14, 2010 to January 14, 2011. This dataset comprises measurements from 4,200 devices that are deployed across sixteen different ISPs and hundreds of cities in the United States. The volunteers for the study were recruited through <http://www.testmyisp.com>. Figure 3 shows a map of the deployment.

Table 2 lists the ISPs that we study, the number of gateways deployed in them, and the number of gateways that report more than 100 throughput measurements. Gateways are rolled out in phases. These devices perform measurements less aggressively when users are sending a lot of traffic. Therefore, not all gateways report data for the entire duration of the study. When we report averages and 95th percentile values for some metric, we only consider gateways that have reported more than 100 measurements for that metric. We also only consider the eight ISPs with the most gateways.

Table 3 shows the active measurements that we use from the SamKnows deployment; some of these (e.g., *last mile latency*) were inspired from our experience running them on BISMarks. The gateways conduct upstream and downstream measurements to servers hosted at Measurement Lab [27] about once every two hours.

There are many ways to measure throughput, though there is no standard method. Bauer *et al.* list several notions of “broadband speed”: *capacity* is the total carrying capacity of the link; and the *bulk transfer capacity* is the amount of data that can be transferred

ISP	Technology	SamKnows		BISMarks Total
		Total	Active	
Comcast	Cable	864	560	4
AT&T	DSL/FTTN	787	335	10
TimeWarner	Cable	690	381	-
Verizon	DSL/FTTP	551	256	-
Cox	Cable	381	161	-
Qwest	DSL/FTTN	265	117	-
Charter	Cable	187	51	-
Cablevision	Cable	104	53	-

Table 2: The SamKnows and BISMarks deployments. *Active* deployments are those that report more than 100 download throughput measurements over the course of our study.

along a path with a congestion-aware protocol like TCP. In Section 5.1, we evaluate several methods for measuring these metrics.

The SamKnows gateways measure bulk transfer capacity using an HTTP client that spawns three parallel threads; this approach increases the likelihood of saturating the access link. The software first executes a “warmup” transfer until throughput is steady to ensure that the throughput measurements are not affected by TCP slow start. The download tests that follows use the same TCP connection to exploit the “warmed up” session. The tests last for about 30 seconds; the software reports snapshots of how many bytes were transferred for every five-second interval.

The gateways also measure different aspects of latency: (1) end-to-end latency; (2) latency to the first IP hop inside the ISP (*last mile latency*); and (3) latency coinciding with an upload or download (*latency under load*). They measure end-to-end latency in two ways: (1) Using a UDP client that sends about six hundred packets an hour to the servers and measures latency and packet loss, and (2) using ICMP ping to the same set of servers at the rate of five packets per hour. To measure latency under load, the gateway measures end-to-end latency during both the upload and the download measurements. They also measure jitter based on RFC 5481 [28] and the time to download the home page of ten popular websites. Before any test begins, the measurement software checks whether cross traffic on the outgoing interface exceeds 64 Kbits/s down or 32 Kbits/s up; if traffic exceeds this threshold, it aborts the test.

4.2.2 BISMarks

BISMarks comprises gateways in the home, a centralized management and data collection server, and several measurement servers. The gateway performs passive and active measurements and anonymizes the results before sending them back to the central repository for further analysis. This gateway also periodically “phones home” to allow the central repository to communicate back through network address translators, to update network configurations and to install software updates. The gateway is based on the NOX Box [31], a small-form-factor computer resembling an off-the-shelf home router/gateway. The NOX Box hardware is an ALIX 2D13 6-inch by 6-inch single board computer with a 500MHz AMD Geode processor, 256 MB of RAM and at least 2 GB of flash memory. The Nox Box runs Debian Linux.

Table 3 lists the measurements that BISMarks collects.¹ We collect throughput, latency, packet loss, and jitter measurements.

BISMarks measures *bulk transfer capacity* by performing an HTTP download and upload for 15 seconds using a single-threaded TCP connection once every 30 minutes, regardless of cross traffic. We do this to have more readings, and to account for cross-traffic, we count bytes transferred by reading directly from `/proc/net/dev`, and compute the “passive throughput” as the

¹The data is available at <http://projectbismark.net/>.

Parameter	Type	Prot.	Freq.	Comments
SamKnows: 4,200 devices, 16 ISPs				
Latency	End-to-end	UDP	600 pkts/hr	MLab
	End-to-end	ICMP	5 pkts/hr	MLab
	Last-mile	ICMP	5 pkts/hr	First IP hop
	Upstream load	ICMP	2 hours	During upload
	Downstream load	ICMP	2 hours	During download
Loss	End-to-end	UDP	600 pkts/hr	MLab
Downstream Throughput	Multi-threaded HTTP	TCP	2 hours	MLab, idle link
Upstream Throughput	Multi-threaded HTTP	TCP	2 hours	MLab, idle link
Jitter	Bi-directional	UDP	1 hour	500pkts/30sec
Web GET	HTTP	TCP	1 hour	Alexa sites
BISMark: 17 devices, 3 ISPs				
Latency	End-to-end	ICMP	5 min	Host
	Last-mile	ICMP	5 min	First IP hop
	Upstream load	ICMP	30 min	During upload
	Downstream load	ICMP	30 min	During download
Packet loss	End-to-end	UDP	15 min	D-ITG
Jitter	End-to-end	UDP	15 min	D-ITG
Downstream Throughput	Single-thread HTTP	TCP	30 min	curlget to Host
	Passive throughput	N/A	30 min	/proc/net/dev
Upstream Throughput	Capacity	UDP	12 hrs	ShaperProbe
	Single-thread HTTP	TCP	30 min	curlput to Host
Passive throughput	Capacity	N/A	30 min	/proc/net/dev
	Capacity	UDP	12 hrs	ShaperProbe

Table 3: Active measurements periodically collected by the SamKnows and BISMark deployments.

byte count after the HTTP transfer minus the byte count before the transfer, divided by the transfer time. This yields the combined throughput of the HTTP transfer and the cross traffic. To measure *capacity*, we run ShaperProbe [32] once every twelve hours to measure UDP capacity. The gateways measure end-to-end latency to a nearby wide-area host, last-mile latency, and latency-under load to the last-mile router. They also measure packet loss and jitter using the D-ITG tool [6]. The gateways perform each measurement at the frequency presented in Table 3 regardless of cross traffic. All measurements are synchronized to avoid overlapping towards the same measurement server.

In January 2011, the BISMark deployment had 14 gateways across AT&T (DSL) and Comcast (Cable), and two more on Clear (WiMax). We recruited users from our research lab and other graduate students in the computer science department; some users were recruited using a recruiting firm.² We only study AT&T and Comcast using BISMark. The AT&T users form the most diverse set of users in the current deployment, with five distinct service plans. We use data from the same period as the SamKnows study.

5. UNDERSTANDING THROUGHPUT

We study throughput measurements from both the SamKnows and BISMark deployments. We first explore how different mechanisms for measuring throughput can generate different results and offer guidelines on how to interpret them. We then investigate the throughput users achieve on different access links, the consistency of throughput obtained by users, and the factors that affect it. Finally, we explore the effects of ISP traffic shaping and the implications it holds for throughput measurement.

5.1 Interpreting Throughput Measurements

Users of access networks are often interested in the throughput that they receive on uploads or downloads, yet the notion of “throughput” can vary depending on how, when, and who is mea-

²All of our methods have been approved by Georgia Tech’s institutional review board.




Figure 4: Comparison of various methods of measuring throughput. (SamKnows and BISMark)

asuring it. For example, a sample run of `www.speedtest.net` in an author’s home, where the service plan was 6Mbits/s down and 512Kbits/s up, reported a downlink speed of 4.4 Mbits/s and an uplink speed of 140 Kbits/s. Netalyzr reported 4.8 Mbits/s and 430 Kbits/s. Long-term measurements (from the SamKnows gateway deployed in that author’s home) paint a different picture: the user achieves 5.6 Mbits/s down and 452 Kbits/s up. Both `www.speedtest.net` and Netalyzr measurements reflect transient network conditions, as well as other confounding factors. Users cannot complain to their ISPs based solely on these measurements. Although measuring throughput may seem straightforward, our results in this section demonstrate the extent to which different measurement methods can produce different results and, hence, may result in different conclusions about the ISP’s performance.

We compare several methods for measuring upstream and downstream throughput from Table 3. We normalize the values of throughput by the service plan rates advertised by the ISP so that we can compare throughput across access links where users have different service plans.

Throughput measurement techniques—even commonly accepted ones—can yield variable results. We perform comparisons of throughput measurement techniques in two locations that have deployed both the SamKnows and BISMark gateways (we are restricted to two due to the logistical difficulty in deploying both gateways in the same location). In both cases, the ISP is AT&T, but the service plans are different (6 Mbits/s down and 512 Kbits/s up; and 3 Mbit/s down and 384 Kbits/s up). Figure 4 shows a CDF of the normalized throughput reported by the four methods we presented in Table 3. Each data point in the distribution represents a single throughput measurement by a client. A value of 1.0 on the x-axis indicates that the throughput matches the ISP’s advertised rate. None of the four methods achieve that value. This could be due to many factors: the sync rate of the modem to the DSLAM; layer-2 framing overhead on the line; or overhead from the measurement techniques themselves. The throughput achieved by multiple parallel TCP sessions comes closer to achieving the advertised throughput. UDP measurements (obtained from ShaperProbe) also produce consistent measurements of throughput that are closer to the multi-threaded TCP measurement. A single-threaded TCP session may not be able to achieve the same throughput, but accounting for cross traffic with passive measurements can provide a better estimate of the actual achieved throughput.

The behavior of single-threaded TCP measurements varies for different access links. We compare the passive throughput for two BISMark users with the *same ISP and service plan* (AT&T; 3 Mbits/s down, 384 Kbits/s up) who live only a few blocks apart. Figure 5 shows that User 2 consistently sees nearly 20% more throughput—much closer to the advertised rate—than User 1. One possible explanation for this difference is the loss rates experienced by these two users; User 1 suffers more loss than User 2




Figure 5: Users with the same service plan but different loss profiles see different performance. User 1 has higher loss and sees lower performance. (BISMark)




Figure 6: Average download rate versus the average upload rate obtained by individual users in the dataset. (SamKnows)

(0.78% vs. 0.20% on the downlink and 0.24% vs. 0.06% on the uplink). Their baseline latencies differ by about 16 milliseconds (8 ms vs. 24 ms). We confirmed from the respective modem portals that User 1 has interleaving disabled and that User 2 has interleaving enabled. Therefore, User 2 is able to recover from noisy access links that cause packet corruption or losses. Single-threaded downloads are more adversely affected by the loss rate on the access link than multi-threaded downloads (even when accounting for cross traffic); reducing the loss rate (*e.g.*, by interleaving) can improve the performance of a single-threaded download. For the rest of the paper, we consider only multi-threaded TCP throughput.

Takeaway: Different throughput measurement techniques capture different aspects of throughput. A single-threaded TCP session is sensitive to packet loss. Augmenting this measurement with passive usage measurements improves its accuracy. Multi-threaded TCP and the UDP capacity measurements measure the access link capacity more accurately and are more robust to loss.

5.2 Throughput Performance

We investigate the throughput obtained by users in the SamKnows deployment. We then study the consistency of their performance.

What performance do users achieve? Figure 6 shows the average download and upload speeds obtained by each user in the SamKnows dataset. Each point in the scatter plot shows the average performance obtained by a single user in the deployment. Clusters of points in the plot reveal common service plans of different ISPs,


identified in the plot by labels. In general, these results agree with the findings from both Netalyzr [24] and Dischinger *et al.* [12], although our dataset also contains Verizon FiOS (FTTP) users that clearly stand out, as well as other more recent service offerings (*e.g.*, AT&T U-Verse). Although the statistics do show some noticeable clusters around various service plans, there appears to be considerable variation even within a single service plan. We seek to understand and characterize both the performance variations and their causes. We do not yet have access to the service plan information of each user, so we focus on how and why throughput performance varies, rather than whether the measured values actually match the rate corresponding to the service plan.

Do users achieve consistent performance? We analyze how consistently users in the SamKnows achieve their peak performance deployment using the $Avg/P95$ metric, which we define as the ratio of the average upload or download throughput obtained by a user to the 95th percentile of the upload or download throughput value obtained by the same user. Higher values for these ratios reflect that users' upload and download rates that are more consistently close to the highest rates that they achieve; lower values indicate that user performance fluctuates.

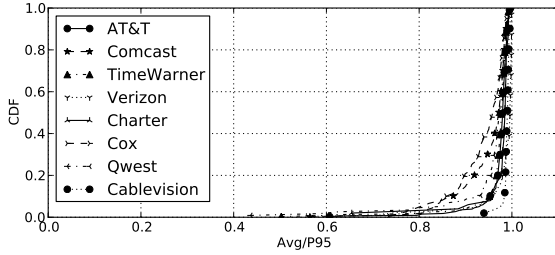
Figure 7a shows the CDF of the $Avg/P95$ metric for each user; Figure 7b shows the same metric for uploads. Most users obtain throughput that is close to their 95th percentile value. Users of certain ISPs (*e.g.*, Cox, Cablevision) experience average download throughput that is significantly less than their 95th percentile. (Both ISPs have more than 50 active users in our data set; see Table 2). Upload throughput performance is more consistent across ISPs. The big difference between download rates and upload rates for popular service plans could account for the fact that upstream rates are more consistent than downstream rates. We also studied the Median/ $P95$ performance; which is similar to $Avg/P95$, and so we do not show them. Our results suggest that upload and download throughput are more consistent than they were when Dischinger *et al.* performed a similar study few years ago [12], especially for some cable providers.

Why is performance sometimes inconsistent? One possible explanation for inconsistent download performance is that the access link may exhibit different performance characteristics depending on time of day. Figure 8a shows the $Avg/P95$ metric across the time of day. We obtain the average measurement reported by each user at that particular time of day and normalize it with the 95th percentile value of that user over all reports. Cablevision users see, on average, a 40% drop in performance from early morning and evening time (when users are likely to be home). For Cox, this number is about 20%. As the figure shows, this effect exists for other ISPs to a lesser extent, confirming prior findings [12]. Because we do not know the service plan for each user, we cannot say whether the decrease in performance for Cox and Cablevision represents a drop below the service plans for those users (*e.g.*, these users might see rates higher than their plan during off-peak hours). Figure 8b shows how the standard deviation of normalized throughput varies depending on the time of day. Performance variability increases for *all* ISPs during peak hours. Figure 8c shows the loss behavior for different times of day; although most ISPs do not see an increase in loss rates during peak hours, Cox does. This behavior suggests that some access ISPs may be under-provisioned; those ISPs for which users experience poor performance during peak hours may be experiencing congestion, or they may be explicitly throttling user traffic during peak hours.

Takeaway: Although there is no significant decrease in performance during peak hours, there is significant variation. A one-time



(a) Download throughput is mostly consistent, with some exceptions.



(b) Upload throughput is consistent across ISPs.

Figure 7: Consistency of throughput performance: The average throughput of each user is normalized by the 95th percentile value obtained by that user. (SamKnows)


“speed test” measurement taken at the wrong time could likely report misleading numbers that do not have much bearing on the long-term performance.

5.3 Effect of Traffic Shaping on Throughput


ISPs shape traffic in different ways, which makes it difficult to compare measurements across ISPs, and sometimes even across users within the same ISP. We study the effect of PowerBoost 3 across different ISPs, time, and users. We also explore how Comcast implements PowerBoost.

Which ISPs use PowerBoost, and how does it vary across ISPs?


The SamKnows deployment performs throughput measurements once every two hours; each measurement lasts 30 seconds, and each report is divided into six snapshots at roughly 5-second intervals for the duration of the 30-second test (Section 4). This measurement approach allows us to see the progress of each throughput measurement over time; if PowerBoost is applied, then the throughput during the last snapshot will be less than the throughput during the first. For each report, we normalize the throughput in each period by the throughput reported for the first period. Without PowerBoost, we would expect that the normalized ratio would be close to one for all intervals. On the other hand, with PowerBoost, we expect the throughput in the last five seconds to be less than the throughput in the first five seconds (assuming that PowerBoost lasts less than 30 seconds, the duration of the test). Figure 9 shows the average progression of throughput over all users in an ISP: the average normalized throughput decreases steadily. We conclude that most cable ISPs provide some level of PowerBoost for less than 30 seconds, at a rate of about 50% more than the normal rate. Cablevision’s line is flat; this suggests that either it does not provide PowerBoost, or it lasts well over 30 seconds consistently, in which case the throughput test would see only the PowerBoost effect. The



(a) The biggest difference between peak and worst performance is about 40%.



(b) The standard deviation of throughput measurements increases during peak hours, most significantly for ISPs that see lower throughputs at peak hours.



(c) Loss increases during peak hours for Cox. Other ISPs do not see this effect as much.

Figure 8: Time of day is significant: The average download throughput for Cablevision and Cox users drops significantly during the evening peak time. Throughput is also significantly more variable during peak time. (SamKnows)

gradual decrease, rather than an abrupt decrease, could be because PowerBoost durations vary across users or that the ISP changes PowerBoost parameters based on network state. From a similar analysis for uploads (not shown), we saw that only Comcast and Cox seem to provide PowerBoost for uploads; we observed a decrease in throughput of about 20%. Dischinger *et al.* [12] also reported PowerBoost effects, and we also see that it is widespread among cable ISPs. For the DSL ISPs (not shown), the lines are flat.

Takeaway: Many cable ISPs implement PowerBoost, which could distort speedtest-like measurements. While some people may be only interested in short-term burst rates, others may be more interested in long-term rates. Any throughput benchmark should aim to characterize both burst rates and steady-state throughput rates.

Do different users see different PowerBoost effects? Using BIS-Mark, we study Comcast’s use of PowerBoost in depth. According to Comcast [9], their implementation of PowerBoost provides




Figure 9: The average throughput obtained during the course of the measurement goes down significantly for the ISPs that enable PowerBoost. (SamKnows)

higher throughput for the first 10 MBytes of a download and the first 5 MBytes of an upload. We measure the shaped throughput for download and upload at the receiver using `tcpdump`. Because our tests are intrusive, we conducted them only a few times; however the results do not vary with choice of traffic generators or ports.

Figure 10 shows the observed throughput for four users for both download and uploads. All four users see PowerBoost effects, but, surprisingly, we see many different profiles even in such a small subset of users. Figure 10a shows download profiles for each user (identified by the modem they use; while the modem doesn't have an effect on burst rates, it does have an effect on buffering latencies as we show in Section 6). The user with a D-LINK modem sees a peak rate of about 21 Mbits/s for 3 seconds, 18.5 Mbits/s for a further ten seconds, and a steady-state rate of 12.5 Mbits/s. The Motorola user sees a peak rate of 21 Mbits/s for about 8 seconds. The PowerBoost technology [10] provides token buckets working on both packet and data rates; it also allows for dynamic bucket sizes. The D-LINK profile can be modeled as a cascaded filter with rates of 18.5 Mbits/s and 12.5 Mbits/s, and buffer sizes of 10MBytes and 1Mbyte respectively, with the line capacity being 21Mbits/s. We see varying profiles for uploads as well, although we only see evidence of single token buckets (Figure 10b). The D-LINK user sees about 7 Mbits/s for 8 seconds, Scientific Atlanta and Thomson users see about 4 Mbits/s for 20 seconds, and the Motorola user sees about 3.5Mbits/s for nearly 35 seconds. Because our results do not vary with respect to the packet size, we conclude that Comcast does not currently apply buckets based on packet rates.

Takeaway: Depending on how throughput measurements are conducted and how long they last, the measurements across users may vary considerably. Specifically, any speedtest measurement that lasts less than 35 seconds will *only* capture the effects of PowerBoost in some cases, and any short-term throughput measurement may be biased by PowerBoost rates.


6. UNDERSTANDING LATENCY

We show how latency can drastically affect performance, even on ISP service plans with high throughput. We then study how various factors ranging from the user's modem to ISP traffic shaping policies can affect latency.


6.1 How (and Why) to Measure Latency

Latency affects the performance that users experience. It not only affects the throughput that users achieve, it also affects *perceived* performance: on a connection with high latency, various operations ranging from resolving DNS queries to rendering content may simply take longer.

Although latency appears to be a straightforward characteristic to measure, arriving at the appropriate metric is a subtle challenge because our goal is to isolate the performance of the access link



(a) PowerBoost download behavior for 4 users.



(b) PowerBoost upload behavior for 4 users.


Figure 10: The level and duration of the burstiness is different for users with different modems, suggesting different shaping mechanisms or parameters. (BISMark)

from the performance of the end-to-end path. End-to-end latency between endpoints is a common metric in network measurement, but it reflects the delay that a user experiences along a wide-area path. We use two metrics that are more appropriate for access networks.

The first metric is the *last-mile latency*, which is the latency to the first hop inside the ISP's network. This metric captures the latency of the access link, which could affect gaming or short downloads. We measure last-mile latency in both of our deployments. As we show in this section, the last-mile latency is often a dominant factor in determining the end-user performance. The second metric we define is *latency under load*, which is the latency that a user experiences during an upload or download (*i.e.*, when the link is saturated in either direction). For BISMark, we measure the last-mile latency under load; on the SamKnows platform, we measure latency under load on the end-to-end path.

To investigate the effect of latency on performance, we measured how the time to download popular Web pages varies for users with different throughput and latency. Figure 11 shows the download time for `www.facebook.com` and how it varies by both the user's throughput and baseline last-mile latency. Figure 11a plots the 95th percentile of each user's downstream throughput versus the average time it takes to download all objects from `www.facebook.com`. The average size of the download is 125 KByte. As expected, the download times decrease as throughput increases; interestingly, there is negligible improvement beyond a rate of 6 Mbits/s. Figure 11b plots download time against the baseline latency for all users whose downstream throughput (95th percentile) exceeds 6 Mbits/s. Minimum download times increase by about 50% when baseline latencies increase from 10 ms to 40 ms. The fact that this effect is so pronounced, even for small downloads, underscores importance of baseline latency.

We investigate the effects of cable and DSL access-link technologies on last-mile latency, packet loss, and jitter. We also explore how different DSL modem configurations, such as whether the modem has interleaving enabled, affects last-mile latency and loss. Finally, we study the effect of modem hardware on perfor-



(a) Fetch time stabilizes above 6Mbits/s. (b) Latency affects fetch times.

Figure 11: Effect of downstream throughput and baseline latency on fetch time from `facebook.com`. (SamKnows)

mance. Specifically, we investigate how oversized modem buffers that has recently received much attention from both operators and users [16]—affects interactivity and throughput.

ISP	Last mile latency		Loss	
	Average	Std. dev	Avg(%)	Std. dev
AT&T	25.23	33.47	0.48%	3.59
Comcast	10.36	14.49	0.27%	2.79
TimeWarner	11.87	25.18	0.33%	3.09
Verizon	12.41	20.60	0.51%	4.07
Charter	11.87	11.80	0.43%	3.29
Cox	13.88	28.02	1.11%	8.35
Qwest	39.42	32.27	0.33%	3.38
Cablevision	10.21	7.52	0.33%	3.14

Table 4: Last-mile latency and variation is significant; Variation in loss is high, suggesting bursty losses. (SamKnows)

ISP	Downstream		Upstream	
	Average	Std. dev	Average	Std. dev
AT&T	1.85	7.63	3.02	12.92
Comcast	1.15	6.37	3.24	6.60
TimeWarner	1.68	3.35	3.67	12.52
Verizon	1.71	5.01	1.97	4.82
Charter	1.17	1.66	2.66	7.48
Cox	1.18	1.89	4.27	7.10
Qwest	3.04	12.59	2.16	10.95
Cablevision	1.69	3.52	2.25	1.18

Table 5: Downstream jitter is quite low, however upstream jitter is significant. (SamKnows)

6.2 Last-Mile Latency

We obtain the last-mile latency by running `traceroute` to a wide-area destination and extracting the first IP address along the path that is not a NAT address. Note that we are measuring the latency to the first network-layer hop, which may not in fact be the DSLAM or the CMTS, since some ISPs have layer-two DSLAMs that are not visible in `traceroute`. This should not be problematic, since the latency between hops inside an ISP is typically much smaller than the last-mile latency.

How does access technology affect last-mile latency? Table 4 shows the average last-mile latency experienced by users in the ISPs included in our study. Last-mile latency is generally quite

high, varying from about 10 ms to nearly 40 ms (ranging from 40 – –80% of the end-to-end path latency). Variance is also high. One might expect that variance would be lower for DSL, since it is not a shared medium like cable. Surprisingly, the opposite is true: AT&T and Verizon have high variance compared to the mean. Qwest also has high variance, though it is a smaller fraction of the mean. To understand this variance, we divide different users in each ISP according to their baseline latency, as shown in Figure 12 Most users of cable ISPs are in the 0–10 ms interval. On the other hand, a significant proportion of DSL users have baseline last-mile latencies more than 20 ms, with some users seeing last-mile latencies as high as 50 to 60 ms. Based on discussions with network operators, we believe DSL companies may be enabling an interleaved local loop for these users.

Table 4 shows loss rates for users across ISPs. The average loss is small, but variance is high for all ISPs, suggesting bursty loss. Jitter has similar characteristics, as shown in Table 5; while the average jitter is low, the variation is high, especially on the upstream, also suggesting burstiness.

How does interleaving affect last-mile latency? ISPs enable interleaving for three main reasons: (1) the user is far from the DSLAM; (2) the user has a poor quality link to the DSLAM; or (3) the user subscribes to “triple play” services. An interleaved last-mile data path increases robustness to line noise at the cost of higher latency. The cost varies between two to four times the baseline latency.

Takeaway: Cable providers in general have lower last-mile latency and jitter. Baseline latencies for DSL users may vary significantly based on physical factors such as distance to the DSLAM or line quality.

6.3 Latency Under Load

We turn our attention to a problem that has gathered much interest recently because of its performance implications: modem buffering under load conditions [16]. We confirm that excessive buffering is a widespread problem afflicting most ISPs (and the equipment they provide). We profile different modems to study how the problem affects each of them. We also see the possible effect of ISP policies such as active queue and buffer management on latency and loss. Finally we explore exploiting shaping mechanisms such as PowerBoost might help mitigate the problem.

Problem: Oversized buffers. Buffers on DSL and cable modems are too large. Buffers do perform an important role: they absorb bursty traffic and enable smooth outflow at the configured rate [24]. Buffering only affects latency during periods when the access link is loaded, but during such periods, packets can see substantial delays as they queue up in the buffer. The capacity of the uplink also affects the latency introduced by buffering. Given a fixed buffer size, queuing delay will be lower for access links with higher capacities because the draining rate for such buffers is higher. We study the effect of buffering on access links by measuring latency when the access link is saturated, under the assumption that the last-mile is the bottleneck. We also present a simple model for modem buffering and use emulation to verify its accuracy.

How widespread are oversized buffers? Figure 13 shows the average ratios of latency under load to baseline latency for each user across different ISPs for the SamKnows data. The histogram shows the latencies when the uplink and the downlink are saturated separately. This figure confirms that oversized buffers affect users across all ISPs, though in differing intensity. The factor of increase when the uplink is saturated is much higher than when the downlink




Figure 12: The baseline last mile latency for each user is computed as the 10th percentile of the last mile latency. Most users see latencies less than 10 ms, but there are a significant number of users with the last mile latency greater than 10 ms. (SamKnows)

is saturated. One plausible explanation is that the downlink usually has more capacity than the uplink, so buffering on the ISP side is lower. The home network (at least 10 Mbits/s) is also probably better provisioned than the downlink, so there is minimal buffering in the modem for downstream traffic. The high variability in the latency under load can be partly explained by the variety in service plans; for instance, AT&T offers plans ranging from 768 Kbits/s to 6 Mbits/s for DSL and up to 18 Mbits/s for UVerse and from 128 Kbits/s to more than 1 Mbit/s for upstream. In contrast, Comcast offers fewer service plans, which makes it easier to design a device that works well for all service plans.

How does modem buffering affect latency under load? To study the effects of modem buffers on latency under load, we conduct tests on AT&T and Comcast modems using BISMarks. We ran tests on the best AT&T DSL (6 Mbits/s down; 512 Kbits/s up) and Comcast (12.5 Mbits/s down; 2 Mbits/s up) plans. We perform the fol-




Figure 13: Latency under load: the factor by which baseline latency goes up when the upstream or the downstream is busy. The high ratios translate to significant real latencies, often in the order of seconds. (SamKnows)


lowing experiment: we start ICMP ping (at the rate of 10 pkts/s for Comcast and 2 pkts/s for AT&T as some modems were blocking higher rates) to the last mile hop. After 30 seconds, we flood the uplink (at 1 Mbits/s for AT&T and at 10 Mbits/s for Comcast using iperf UDP). After 60 seconds, we stop iperf, but let ping continue for another 30 seconds. The ping measurements 30 seconds on either side of the iperf test establishes baseline latency. The Motorola and the 2Wire modems were brand new, while the Westell modem is about 5 years old, and was in place at the home where we conducted the experiment. We also saw the same Westell modem in two other homes in the BISMarks deployment.

Figure 14a shows the latency under load for the three modems. In all cases, the latency increases dramatically at the start of the


flooding and plateaus when the buffer is saturated. The delay experienced by packets at this stage indicates the size of the buffer, since we know the uplink draining rate. Surprisingly, we see more than an order of magnitude of difference between modems. The 2Wire modem has the lowest worst case latency, of 800 ms. Motorola's is about 1600 ms, while the Westell has a worst case latency of more than 10 seconds. Because modems are usually the same across service plans, we expect that this problem may be even worse for users with slower plans.

To model the effects of modem buffering, we emulated this setup in Emulab [13] with a 2 end-host, 1-router graph. We configured a token bucket filter using tc. We compute the buffer size to be: 512 Kbits/s × max(latency of modem), which yields a size of 640 Kbytes for Westell, 100 Kbytes for Motorola, and 55 Kbytes for 2Wire. This simple setup almost perfectly captures the latency profile that the actual modems exhibit. Figure 14b shows the emulated latencies. Interestingly, we observed little difference in throughput for the three buffer sizes. We also emulated other buffer sizes. For a 512 Kbits/s uplink, we observed that the modem buffers exceeding 20 Kbytes do little for throughput, but cause a linear increase in latency under load. Thus, the buffer sizes in all three modems are too large for the uplink.

How does PowerBoost traffic shaping affect latency under load? To understand latency under load for cable users, we study the Comcast users from BISMarks. All of the modems we study have buffers that induce less than one second of delay, but these users see surprising latency under load profiles due to traffic shaping. Figures 15a and 15b show the latency under load for two Comcast users. The other two Comcast users (with the Scientific Atlanta and the Motorola modems) had latency profiles similar to the user with the Thomson modem, so we do not show them. The difference in the two latency profiles is interesting; the D-LINK user sees a jump in latency when the flooding begins and about 8 seconds later, another increase in latency. The Thomson user sees an initial increase in latency when flooding starts but then a decrease in latency after about 20 seconds. The first effect is consistent with buffering and PowerBoost. Packets see lower latencies during PowerBoost because, for a fixed buffer, the latency is inversely proportional to the draining rate. The increase in latency due to PowerBoost (from 200 ms to 700 ms) is proportional to the decrease in the draining rate (from 7 Mbits/s to 2 Mbits/s, as shown in Figure 10b). The decrease in latency for the Thomson user cannot be explained in the same way. Figure 15 shows the average loss rates alongside the latencies for the two users; interestingly for the user with the Thomson modem, the loss rate is low for about 20 seconds after the link is saturated, but there is a sharp hike in loss corresponding



(a) *Empirical measurements* of modem buffering. Different modems have different buffer sizes, leading to wide disparities in observed latencies when the upstream link is busy. (BISMark)



(b) *Emulated modems* with token bucket filters. We see similar latency progression. Emulated buffer sizes have minimal effect on throughput.

Figure 14: Buffering in AT&T modems. There is little benefit to the buffering seen in most modems.


to the drop in latency. This behavior may correspond to dynamic buffer sizing, as discussed in Section 5.3.

Can data transfer be modified to improve latency under load?


We explore whether a user can modify their data transfer behavior so that large “bulk” flows and delay-sensitive flows can co-exist without interfering with one another. We compare the impact of a 50 MByte download on a G.711 VoIP call in three different conditions: (1) not applying any traffic control, (2) intermittent traffic at capacity on 10.8 seconds ON and 5.3 seconds OFF cycle, and (3) shaping using the WonderShaper [36] approach. Figure 16 shows the result of this experiment. In (1), the transfer takes 25.3 seconds; however, just after the PowerBoost period, the VoIP call starts suffering high latency and loss until the end of the transfer. In (2), traffic is sent in pulses, and the download takes 26.9 seconds. In (3), traffic is sent at just under the long term rate and the download takes 32.2 seconds. Both (2) and (3) do not increase latency significantly, this is because they do not deplete the tokens at any time, and therefore cause no queuing. In approach (2), the ON/OFF periods can be configured depending on the token bucket parameters,³ and the size of the file to be transferred. Both approaches achieve similar long-term rates but yield significant latency benefit. The drawback is that any approach that exploits this behavior would need to know the shaping parameters.

Takeaway: Modem buffers are too large. Even the smallest buffers we see induce nearly one-second latency under load for AT&T and 300 ms for Comcast. Buffering is detrimental to both interactivity and throughput. Modifying data transfer behavior using short bursts or tools like WonderShaper might help mitigate the problem in the short term.

³If ρ_r is the rate we want to reserve for real-time applications, and ρ_t the token rate, the condition to be satisfied is: $(\rho_b + \rho_r - \rho_t) \times \tau_{on} \leq \tau_{off} \times (\rho_t - \rho_r)$, where ρ_b is the sending rate during the pulse, and τ_{on} and τ_{off} are the ON and the OFF times, respectively.



(a) Comcast user with D-LINK modem.



(b) Comcast user with RCA Thomson modem.

Figure 15: Possible effect of active buffer management: Loss rates increase when the latency drops. (BISMark)

7. LESSONS LEARNED


We conclude with some high-level lessons and suggestions for future research directions. One significant takeaway for users, policymakers, ISPs, and researchers is that *continual measurements, directly from home network gateways are crucial for understanding the details of home access network performance*. Existing “speed test” downloads and end-to-end latency measurements do not often reflect access network performance over an extended period of time, and they neglect various confounding factors on the host and within the home. Our ability to execute measurements directly, both from a small set of gateways where we can control the network conditions and measurements (BISMark) and a larger, more representative set of gateways across the United States (SamKnows), yields several lessons:

Lesson 1 (One Measurement Does Not Fit All) *Different ISPs use different policies and traffic shaping behaviors that make it difficult to compare measurements across ISPs.*


There is no single number that characterizes performance, or even throughput. Certain ISP practices such as PowerBoost can distort benchmarking measurements; ISPs might even design their networks so that widely used performance tests yield good performance. Developing a benchmarking suite for ISP performance that users can understand (*e.g.*, in terms of the applications they use) is critical; the measurements we develop in this paper may be a good starting point for that. Along these lines, more work is needed to understand the performance of specific applications, such as how video streaming performance compares across ISPs. The Netflix study on ISP streaming performance [30] is a good start, but more such performance benchmarks are needed.

Lesson 2 (One ISP Does Not Fit All) *There is no “best” ISP for all users. Different users may prefer different ISPs depending on their usage profiles and how those ISPs perform along performance dimensions that matter to them.*

Different ISPs may be “better” along different performance dimen-



(a) Throughput.



(b) Latency.

Figure 16: It is possible to maintain low latency by modifying data transfer behavior. (BISMark)

sions, and the service plan that a user buys is only part of the picture. For example, we saw that, above a certain throughput, latency is the dominant factor in determining Web page loading time. Similarly, a gamer might be interested in low latency or jitter, while an avid file swapper may be more interested in high throughput. An imminent technical and usability challenge is to summarize access network performance data so that users can make informed choices about the service plans that are most appropriate for them (akin to a “performance nutrition label” [2]). Our recent work proposes some first steps in this direction [34].

Lesson 3 (Home Network Equipment Matters) *A user’s home network infrastructure can significantly affect performance.*

Modems can introduce latency variations that are orders of magnitude more than the variations introduced by the ISP. Other effects inside the home that we have not yet studied, such as the wireless network, may also ultimately affect the user’s experience. More research is needed to understand the characteristics of traffic inside the home and how it affects performance.

Acknowledgments

We thank the participants in the SamKnows and BISMark studies, and to Walter Johnston at the FCC for help and access to the data from the SamKnows study. We are grateful to Dave Täht for his continuing support of the BISMark platform. We also thank Sam Burnett and Hyojoon Kim for constructive comments on paper drafts. This project is supported by the the National Science Foundation through awards CNS-1059350, CNS-0643974, a generous Google Focus Grant, the European Community’s Seventh Framework Programme (FP7/2007-2013) no. 258378 (FIGARO), and the ANR project C’MON.

REFERENCES

[1] Grenouille. <http://www.grenouille.com/>.
 [2] Does broadband need its own government nutrition label? <http://arstechnica.com/tech-policy/news/2009/10/>

does-broadband-needs-its-own-government-nutrition-label.ars, Oct. 2010. Ars Technica.

[3] C. Bastian, T. Klieber, J. Livingood, J. Mills, and R. Woundy. *Comcast’s protocol-agnostic congestion management system*. Internet Engineering Task Force, Dec. 2010. RFC 6057.

[4] G. Bernardi and M. K. Marina. Bsense: a system for enabling automated broadband census: short paper. In *Proc. of the 4th ACM Workshop on Networked Systems for Developing Regions (NSDR ’10), June 2010.*, 2010.

[5] K. Bode. FCC: One Million Speedtests and Counting. <http://www.dslreports.com/shownews/FCC-One-Million-Speedtests-And-Counting-109440>, July 2010.

[6] A. Botta, A. Dainotti and A. Pescapé. Multi-protocol and multi-platform traffic generation and measurement. IEEE INFOCOM, Demo session, May 2007.

[7] R. Carlson. Network Diagnostic Tool. <http://e2epi.internet2.edu/ndt/>.

[8] K. Cho, K. Fukuda, H. Esaki, and A. Kato. The impact and implications of the growth in residential user-to-user traffic. In *ACM SIGCOMM 2006*, 2006.

[9] Comcast FAQ. <http://customer.comcast.com/Pages/FAQViewer.aspx?Guid=024f23d4-c316-4a58-89f6-f5f3f5dbdcf6>, Oct. 2007.

[10] R. Compton, C.L. Woundy and J. Leddy. Method and packet-level device for traffic regulation in a data network. U.S. Patent 7,289,447 B2, Oct. 2007.

[11] D. Croce, T. En-Najjary, G. Urvoy-Keller, and E. Biersack. Capacity Estimation of ADSL links. In *CoNEXT*, 2008.

[12] M. Dischinger, A. Haeberlen, K. P. Gummadi, and S. Saroiu. Characterizing residential broadband networks. In *Proc. ACM SIGCOMM Internet Measurement Conference*, San Diego, CA, USA, Oct. 2007.

[13] Emulab. <http://www.emulab.net/>, 2006.

[14] M. M. et al. Network Path and Application Diagnosis. <http://www.psc.edu/networking/projects/pathdiag/>.

[15] National Broadband Plan. <http://www.broadband.gov/>.

[16] J. Gettys. Bufferbloat. <http://www.bufferbloat.net/>.

[17] Glasnost: Bringing Transparency to the Internet. <http://broadband.mpi-sws.mpg.de/transparency>.

[18] D. Han, A. Agarwala, D. G. Andersen, M. Kaminsky, K. Papagiannaki, and S. Seshan. Mark-and-sweep: Getting the inside scoop on neighborhood networks. In *Proc. Internet Measurement Conference*, Vouliagmeni, Greece, Oct. 2008.

[19] Internet World Stats. <http://www.internetworldstats.com/dsl.htm>.

[20] Asymmetric Digital Subscriber Line Transceivers. ITU-T G.992.1, 1999.

[21] Asymmetric Digital Subscriber Line (ADSL) Transceivers - Extended Bandwidth ADSL2 (ADSL2Plus). ITU-T G.992.5, 2003.

[22] Data-over-cable service interface specifications: Radio-frequency interface specification. ITU-T J.112, 2004.

[23] C. R. S. Jr. and G. F. Riley. Neti@home: A distributed approach to collecting end-to-end network performance measurements. In *the Passive and Active Measurement Conference (PAM)*, 2004.

[24] C. Kreibich, N. Weaver, B. Nechaev, and V. Paxson. Netalyzer: Illuminating the edge network. In *Proc. Internet Measurement Conference*, Melbourne, Australia, Nov. 2010.

[25] K. Lakshminarayanan and V. N. Padmanabhan. Some findings on the network performance of broadband hosts. In *Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*, IMC ’03, pages 45–50, New York, NY, USA, 2003. ACM.

[26] G. Maier, A. Feldmann, V. Paxson, and M. Allman. On dominant characteristics of residential broadband internet traffic. In *ACM Internet Measurement Conference*, 2009.

[27] Measurement Lab. <http://measurementlab.net>.

[28] A. Morton and B. Claise. *Packet Delay Variation Applicability Statement*. Internet Engineering Task Force, Mar. 2009. RFC 5481.

[29] Netalyzer. <http://netalyzer.icsi.berkeley.edu/>.

[30] Netflix Performance on Top ISP Networks. <http://techblog.netflix.com/2011/01/netflix-performance-on-top-isp-networks.html>, Jan. 2011.

[31] NOX Box. <http://noxrepo.org/manual/noxbox.html>.

[32] ShaperProbe. <http://www.cc.gatech.edu/~partha/diffprobe/shaperprobe.html>.

[33] M. Siekkinen, D. Collange, G. Urvoy-Keller, and E. Biersack. Performance limitations of ADSL users: A case study. In *the Passive and Active Measurement Conference (PAM)*, 2007.

[34] S. Sundaresan, N. Feamster, R. Teixeira, A. Tang, K. Edwards, R. Grinter, M. Chetty, and W. de Donato. Helping users shop for ISPs with internet nutrition labels. In *ACM SIGCOMM Workshop on Home Networks*, 2011.

[35] D. Vorhaus. A New Way to Measure Broadband in America. <http://blog.broadband.gov/?entryId=359987>, Apr. 2010.

[36] WonderShaper. <http://lartc.org/wondershaper/>, 2002.