

Measuring Home Broadband Performance

By S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè

Abstract

We present the results from the first study of Internet access link performance measured directly from home routers. In conjunction with the Federal Communication Commission's study of broadband Internet access in the United States, we investigate the throughput and latency of network access links from about 4000 routers across eight ISPs. Our findings provide a snapshot of access network performance across the United States, offer insights on how access network performance should be measured and presented to users, and inform various ongoing efforts to evaluate the performance of access networks around the world.

1. INTRODUCTION

Of nearly two billion Internet users worldwide, about 500 million are residential broadband subscribers.¹⁰ Broadband penetration will continue to increase, with people relying on home connectivity for day-to-day and even critical activities. Accordingly, the Federal Communication Commission (FCC) is developing performance-testing metrics for access providers.^{4, 9, 22} Policymakers, home users, and Internet Service Providers (ISPs) need better ways to benchmark broadband Internet performance.

Unfortunately, benchmarking home Internet performance is not as simple as running one-time "speed tests." There exist countless tools to measure Internet performance.^{5, 16, 17, 20} Previous work has studied the typical download and upload rates of home access networks^{7, 14}; others have found that modems often have large buffers,¹⁴ and that DSL links often have high latency.¹⁵ These studies have shed some light on access link performance, but they have typically run 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 and, without continual measurements of the same access link, we cannot establish a baseline performance level or observe how performance varies over time.

We measure broadband Internet performance directly from the router that is connected to the user's ISP. Measuring the access link from the home router offers several advantages over conventional methods. First, the home router is typically always on. Second, because it connects the home network to the ISP's network (as shown in Figure 1), 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. Table 1 summarizes the challenges involved in measuring access ISP performance, and how performing measurements from the home router addresses each of them.

We collected data from two complementary deployments. First, we gathered data from routers in over 4200 homes across the United States and eight different ISPs from a deployment sponsored by the Federal Communications Commission and operated by SamKnows. Second, we collected data from 16 homes in the BISmark deployment, spanning three ISPs in Atlanta. The SamKnows deployment provides a large user base, as well as diversity in ISPs, service plans, and geographical locations. We can access BISmark routers remotely and run repeated experiments to investigate the effect of factors that we could not study in a larger deployment. For example, to study the effect of modem choice on performance, we installed different modems in the same home and conducted experiments in the same controlled setting. Both deployments run a comprehensive suite of measurement tools that periodically measure throughput, latency, packet loss, and jitter. We use active measurement data from both deployments from December 14, 2010 to January 14, 2011. Table 2 lists the ISPs that we study and the number of measured access links for each of them.

Figure 1. The home router sits directly behind the modem in the home network. It takes measurements both to the last mile router (first non-NAT IP hop on the path) and to wide area hosts.




Table 1. Confounding factors and how we address them.

Factor	How we address it
Wireless effects	Use a wired connection to the modem
Cross-traffic	Measure cross-traffic and avoid it/account for it
Router load	Use a well-provisioned router
Server location	Choose a nearby server
End-to-end path	Focus on characterizing the last mile
Router configuration	Test configuration in practice and controlled settings

A previous version of this paper was published in the *Proceedings of SIGCOMM '11* (Toronto, Ontario, Canada, Aug. 15–19, 2011).

We characterize access network throughput (Section 3) and latency (Section 4) from the SamKnows and BISmark deployments. We explain how our throughput measurements differ from common “speed tests” and also propose several metrics that capture different aspects of latency. When our measurements cannot fully explain the observed behavior, we model the access link and test our hypotheses using controlled experiments. We find that the most significant factors affecting throughput are the access technology, ISPs’ traffic shaping policies, and congestion during peak hours. On the other hand, the quality of the access link, modem buffering, and cross-traffic within the home affect latency the most.

This study offers insights into both access network performance and measurement methods for benchmarking home broadband performance. Our findings include the following:

- Access link throughput is more variable during peak hours.
- ISPs use different policies and traffic shaping behavior that can make it difficult to compare measurements across them.
- Home network equipment and infrastructure can affect performance. For example, buffering in a user’s modem varies across models and can affect the latency that a user experiences during an upload. We found “bufferbloat,” or excessive buffering that adversely affects network performance, in many DSL modems.
- There is no “best” ISP for all users. Some ISPs have better short-term throughput, while others may have better sustained throughput, lower latency, or generally more consistent performance. Different users may prefer different ISPs depending on their usage profiles and how those ISPs perform along performance dimensions that matter to them.

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 router; 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

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

ISP	Technology	SamKnows		BISmark
		Total	Active	Total
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	–

for a benchmark is clear, and the results from this study can serve as a principled foundation for such an effort.

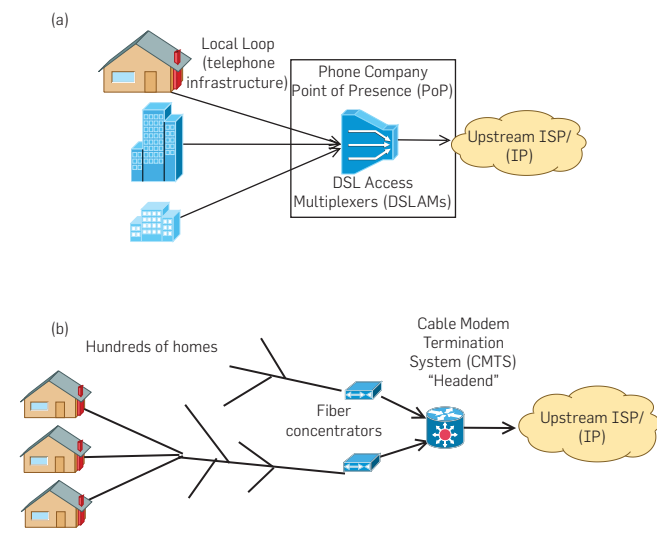
2. ACCESS TECHNOLOGIES

We describe the two most common access technologies from our deployments: Digital Subscriber Line (DSL) and cable. A few users in our deployments have Fiber-To-The-Node (FTTN), Fiber-To-The-Premises (FTTP), and WiMax, but 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 2a. 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 2b. 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 1¹¹ is 12 Mbits/s downstream and 1.8 Mbits/s upstream. The ADSL2+ specification¹² extends the capacity of ADSL links to at most 24 Mbits/s downstream and 3.5 Mbits/s upstream. Although the ADSL technology can theoretically achieve these speeds, many factors 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

Figure 2. Access network architectures (a) DSL and (b) cable.



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 it is received, which minimizes latency but prevents error correction from being applied across frames. Thus, ISPs typically configure fastpath 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),¹³ which specifies download rates up to 42.88 Mbits/s and upload rates up to 30.72 Mbits/s 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. An operator configures the service plan rate limit at the cable modem using a token bucket rate shaper. Many cable providers offer *PowerBoost*, which allows users to download (and, in some cases, upload) data at higher rates 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.

3. UNDERSTANDING THROUGHPUT

We first explore how different techniques 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. Finally, we explore the effects of ISP traffic shaping and the implications it has for throughput measurement.

3.1. How to measure and interpret throughput

Users 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 measuring it. For example, a run of www.speedtest.net in an author's home, where the service plan was 6 Mbits/s down and 512 Kbits/s up, reported a downlink speed of 4.4 Mbits/s and an uplink speed of 140 Kbits/s. Netalyzr¹⁴ reported 4.8 Mbits/s and 430 Kbits/s. Long-term measurements from the SamKnows router 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.

There is no standard way to measure throughput. 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 along a path with a congestion-aware protocol like TCP.³ The SamKnows routers measure bulk transfer capacity using

three parallel HTTP transfers; this approach increases the likelihood of saturating the access link. The software first executes a “warm-up” transfer until throughput is steady to ensure that the throughput measurements are not affected by TCP slow start. The following download tests use the same TCP connection to exploit the “warmed up” session. The tests last for about 30s; the software reports snapshots of how many bytes were transferred for every 5-s interval. BISmark measures throughput by performing an HTTP download and upload for 15s using a single-threaded TCP connection once every 30min, regardless of cross-traffic. To account for cross-traffic, the router counts bytes transferred by reading from `/proc/net/dev` and computes the “passive throughput” as the byte count after the HTTP transfer minus the byte count before the transfer, divided by the transfer time. Table 3 summarizes the throughput measurements collected by the two deployments. Although measuring throughput may seem straightforward, our results demonstrate the extent to which different measurement methods can produce different results and, hence, may result in different conclusions about an ISP's performance.

Throughput measurement techniques—even commonly accepted ones—can yield variable results. We compare throughput measurement techniques in two locations that have both the SamKnows and BISmark routers (we use only two locations due to the logistical difficulty of deploying both routers 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 Mbits/s down and 384 Kbits/s up). We normalize the throughput with the advertised service plan so that we can compare the service plans. Figure 3 shows a CDF of the normalized throughput reported by the four methods from Table 3. Each data point represents a single throughput measurement. A value of 1.0 on the *x*-axis indicates that the throughput matches the ISP's advertised rate; no method achieves that value. This shortfall could be caused by many factors, including the sync rate of the modem to the DSLAM, layer-2 framing overhead on the line, and overhead from the measurement

Table 3. SamKnows and BISmark throughput measurements.

Parameter	Type	Prot.	Freq.	Comments
SamKnows				
Downstream throughput	Multi-threaded HTTP	TCP	2h	MLab, idle link
Upstream throughput	Multi-threaded HTTP	TCP	2h	MLab, idle link
BISmark				
Downstream throughput	Single-thread HTTP	TCP	30 min	curlget from Host/ proc/net/dev
	Passive throughput	N/A	30 min	
	Capacity	UDP	12h	ShaperProbe
Upstream throughput	Single-thread HTTP	TCP	30 min	curlput to Host/ proc/net/dev
	Passive throughput	N/A	30 min	
	Capacity	UDP	12h	ShaperProbe

techniques themselves. Multiple parallel TCP sessions nearly achieve the advertised throughput. UDP measurements also produce consistent measurements of throughput that are closer to the multi-threaded TCP measurement. A single-threaded TCP session may not achieve the same throughput, but accounting for cross-traffic with passive measurements does yield a better throughput estimate.

The behavior of single-threaded TCP measurements varies for different access links. We compare 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 4 shows that User 2 consistently sees nearly 20% higher throughput than User 1. One possible explanation for this difference is that the two users experience different loss rates: User 1 experiences four times more packet loss in both directions than User 2. The baseline latencies also differ by about 16ms (8ms vs. 24ms). We confirmed from the respective modem portals that User 1 has interleaving disabled and User 2 has interleaving enabled. Thus, User 2's connection recovers better from line noise. Single-threaded downloads suffer more from high packet loss rates than multi-threaded downloads; interleaving reduces the packet loss rate, and thus improves 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

Figure 3. Comparison of various methods of measuring throughput. (SamKnows and BISmark)





Figure 4. Users with the same service plan but different loss profiles see different performance. User 1 has higher loss and sees lower performance. (BISmark)



its accuracy. Multi-threaded TCP and the UDP capacity measurements measure the access link capacity more accurately and are more robust to loss.

3.2. Throughput performance

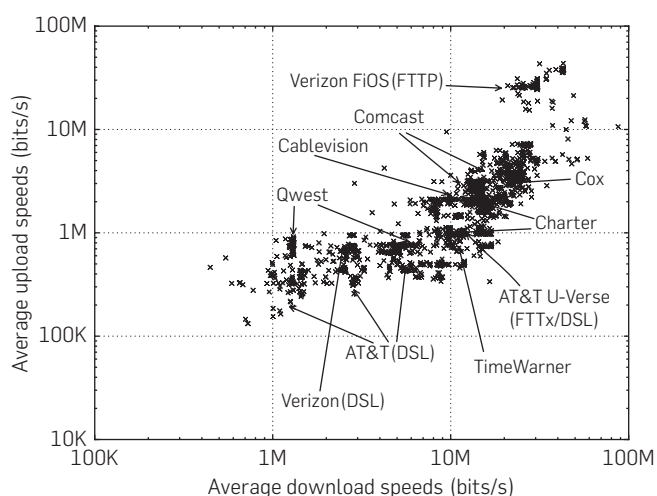
We investigate the throughput obtained by users in the SamKnows deployment and the consistency of these measurements.

What performance do users achieve? Figure 5 shows the average download and upload speeds for each router in the SamKnows dataset. Each point in the scatterplot shows the average performance that a single router in the deployment achieves. 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 previous^{7,14} work, although our dataset also includes Verizon FiOS (FTTP) users that clearly stand out and other recent service offerings (e.g., AT&T U-Verse). Although there are some noticeable clusters around various service plans, there appears to be considerable variation, even within a single service plan. We seek to characterize both the performance variations and their causes.

Do users achieve consistent performance? We analyze whether routers in the SamKnows deployment consistently achieve their peak performance 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. A higher ratio reflects that a router's upload and download rates more consistently achieve performance that is closest to the highest rate; lower values indicate that performance fluctuates.

Figure 6 shows the CDF of the *Avg/P95* metric across users from each ISP. Most access links achieve throughput close to their 95th percentile value. Certain ISPs (e.g., Cox and Cablevision) achieve an average download throughput that is significantly less than their 95th percentile. Upload

Figure 5. Average download rate versus the average upload rate obtained by individual users in the dataset. (SamKnows)



throughput is much more consistent, possibly because upload rates are typically much lower.

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 7a shows the Avg/P95 metric across the time of day. We obtain the average measurement reported by each user at that 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

Figure 6. Consistency of throughput performance: The average throughput of each user is normalized by the 95th percentile value obtained by that user. (SamKnows)


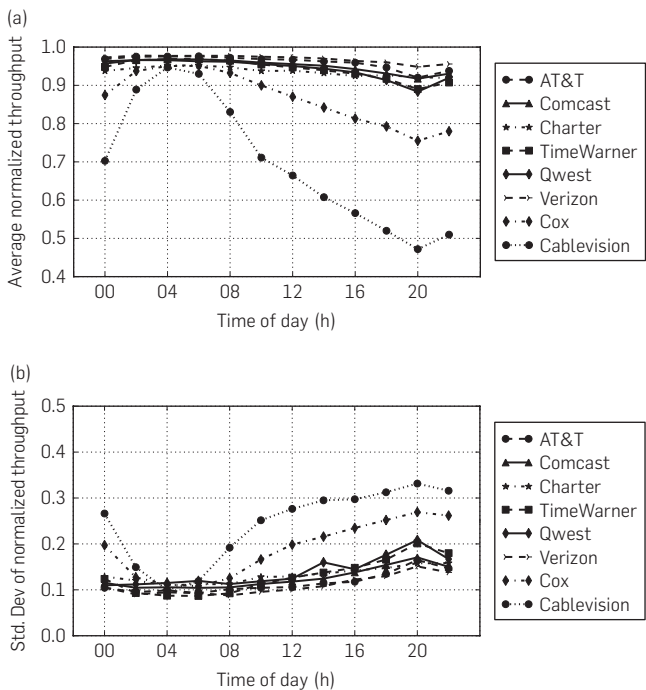


Figure 7. Time of day is significant: The average download throughput for Cablevision and Cox users drops significantly during the evening peak time. Throughput is also significantly more variable during peak time. (SamKnows) (a) Peak and worst performance differ by up to 40%. (b) The standard deviation of throughput measurements increases during peak hours, most significantly for ISPs that see lower throughputs at peak hours.



in the peak evening hours; Cox users experience a 20% performance reduction on average. This effect exists for other ISPs to a lesser extent, confirming prior findings.⁷ Without knowing the service plan for each user, we cannot, in general, say whether the decrease in performance 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). However, the FCC’s report,¹ which analyzes performance in comparison to advertised rates, confirms that Cablevision users do see average performance significantly less than the advertised rates. Figure 7b shows how the standard deviation of normalized throughput varies depending on the time of day. All ISPs experience more variable performance during peak hours. Although most ISPs do not suffer an increase in loss rates during peak hours, Cox does. ISPs that exhibit poor performance during peak hours may be under-provisioned; they may be experiencing congestion, or they may be explicitly throttling traffic during peak hours.

Takeaway: Throughput performance is more variable during peak hours. A one-time “speed test” measurement taken at the wrong time could likely report misleading numbers that do not have much bearing on performance over time.

3.3. 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 across different ISPs, time, and users. We also model how Comcast implements PowerBoost.

Which ISPs use PowerBoost, and how does it vary across ISPs? Each SamKnows throughput measurement lasts 30s, and each report is divided into six snapshots at roughly 5-s intervals for the duration of the test. This technique highlights the evolution of throughput performance over time. On a link that is subject to traffic shaping, 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. The normalized throughput on an unshaped link is close to one for all intervals. On the other hand, on an access link configured with PowerBoost, the throughput in the last 5s should be less than the throughput in the first 5s (assuming that PowerBoost lasts less than 30s). Figure 8 shows the progression of average throughput over all users in an ISP: the average normalized throughput decreases over time. Our data shows that most cable ISPs provide some level of PowerBoost for less than 30s, 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 30s consistently (in which case the throughput test would reflect only the effect of PowerBoost). The gradual decrease, rather than an abrupt decrease, could be because PowerBoost durations vary across users or because the ISP changes PowerBoost parameters based on network state. In the case of uploads, only Comcast and Cox seem to deploy PowerBoost; in these cases, we observed a difference in throughput of about 20%. DSL ISPs do not appear to implement PowerBoost.

Do different users see different PowerBoost effects? We investigate Comcast's use of PowerBoost using the BISmark testbed. According to Comcast,⁶ their implementation of PowerBoost provides higher throughput for the first 10MB of a download and the first 5MB 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, but the results are consistent across traffic generators and ports.

Figure 9 shows the downstream throughput profiles for four users, each identified by their modem type. Although the modem does not affect burst rates, it does have different amounts of buffering, which can affect latency. All four users experience PowerBoost effects, but, surprisingly, each user experiences a different traffic shaping profile: The user with a D-LINK modem sees a peak rate of about 21 Mbits/s for 3s, 18.5 Mbits/s for a further 10s, and a steady-state rate of 12.5 Mbits/s. The Motorola user sees a peak rate of 21 Mbits/s for about 8s. The D-LINK profile can be modeled as a cascaded token bucket filter with rates of 18.5 Mbits/s and 12.5 Mbits/s, and buffer sizes of 10MB and 1MB, respectively, with a capacity of 21 Mbits/s. Upload profiles vary across different users as well, although the shaping profiles seem to indicate that only a single token bucket is applied on the uplink.

Takeaway: Many cable ISPs implement PowerBoost, which could distort speedtest-like measurements. In particular, any throughput measurement that lasts less than 35s will mainly capture the effects of PowerBoost. While some people may be interested only in short-term burst rates, others may be more interested in long-term rates. Any

Figure 8. The average throughput during the measurement decreases for the ISPs that enable PowerBoost. (SamKnows)





Figure 9. The level and duration of the burstiness are different for users with different modems, suggesting different shaping mechanisms or parameters. (BISmark)



throughput benchmark should aim to characterize both burst rates and steady-state throughput rates.

4. 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.

4.1. How (and why) to measure latency

Latency not only affects the throughput that users achieve, but it also affects the performance that users perceive, since it affects everything from DNS lookup time to the time to set up a TCP connection. Although measuring latency appears straightforward, arriving at the appropriate metric is a subtle challenge because our goal is to isolate the performance of the access link 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 IP hop inside the ISP's network. The last-mile latency captures the latency of the access link, which could affect gaming or short downloads. We measure last-mile latency in both of our deployments. 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 end-to-end latency under load. Table 4 summarizes the latency measurements we collect.

We investigate the effect of last-mile latency on download times for popular Web pages. Figure 10 shows the download time for `www.facebook.com` and how it varies by both the user's throughput and baseline last-mile latency. Figure 10a 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

Table 4. SamKnows and BISmark latency and loss tests.

Parameter	Type	Prot.	Freq.	Comments
SamKnows				
Latency	End-to-end	UDP	600 pkts/h	MLab
	End-to-end	ICMP	5 pkts/h	MLab
	Last mile	ICMP	5 pkts/h	First IP hop
	Upstream load	ICMP	2 h	During upload
	Downstream load	ICMP	2 h	During download
Loss	End-to-end	UDP	600 pkts/h	MLab
Jitter	Bidirectional	UDP	1 h	500 pkts/30s
Web GET	HTTP	TCP	1 h	Alexa sites
BISmark				
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

125KB. As expected, the download times decrease as throughput increases; interestingly, there is negligible improvement beyond a rate of 6 Mbits/s. Figure 10b 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 10ms to 40ms. The pronounced effect of latency on download time for Web objects underscores the influence of baseline latency.

4.2. 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, because some ISPs have layer-two DSLAMs that are not visible in traceroute. The possibility of measuring slightly further than the DSLAM or CMTS should not materially affect our results, 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? Last-mile latency is generally quite high, varying from about 10ms to nearly 40ms (ranging from 40–80% of the end-to-end path latency). Last-mile latency is also highly variable. One might expect that variance would be lower for DSL ISPs, 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

Figure 10. Effect of downstream throughput and baseline latency on fetch time from facebook.com. (SamKnows) (a) Fetch time stabilizes above 6 Mbits/s. (b) Baseline latency affects fetch times.

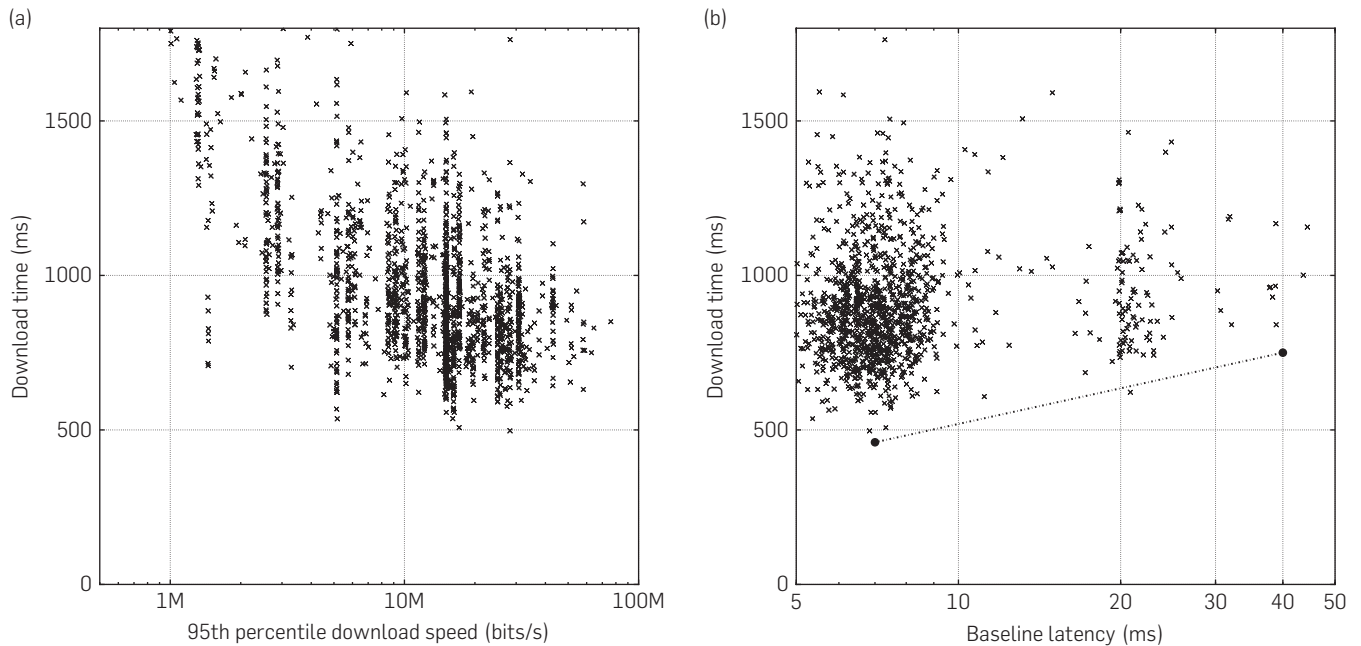



Figure 11. The baseline last-mile latency for each user is computed as the 10th percentile of the last-mile latency. Most users see latencies less than 10ms, but there are a significant number of users with the last-mile latency greater than 10ms. (SamKnows)



fraction of the mean. To understand this variance, we divide different users in each ISP according to their baseline latency, as shown in Figure 11. Most users of cable ISPs are in the 0–10ms interval. On the other hand, a significant proportion of DSL users have baseline last-mile latencies more than 20ms, with some users seeing last-mile latencies as high as 50 to 60ms. Based on discussions with network operators, we believe DSL ISPs may be enabling an interleaved local loop for these users. We also analyzed loss rate and jitter. We found that the average loss rates are small, but variance is high for all ISPs, suggesting bursty loss. Jitter has similar characteristics; although the average jitter is low, the variation is high, especially on the upstream connection.

Takeaway: Cable users tend to have lower last-mile latency and jitter, while for DSL users it may vary significantly based on physical factors such as distance to the DSLAM or line quality.

4.3. Latency under load

We now turn our attention to latency under load, a characteristic that is often affected by “bufferbloat,” or excessive buffering in network devices.¹⁹

Problem: Bufferbloat. Buffers on DSL and cable modems are too large. Buffering affects latency during periods when the access link is loaded; during these periods, packets can see substantial delays as they queue in the buffer. The capacity of the uplink also affects the latency that buffering introduces. For a given buffer size, the queuing delay will be lower on access links with higher capacities because the draining rate 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.

How widespread is bufferbloat? Figure 12 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 plot confirms that bufferbloat affects users across all ISPs, albeit to different extents. The factor of increase when the uplink is saturated is much higher than when the downlink 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 is often better provisioned than the downlink, so downstream traffic experiences less buffering in the modem. The high variability in the latency under load can be partly explained by the variety in service plans; for example, 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 BISmark. We ran tests on the best AT&T DSL plan (6 Mbits/s down; 512 Kbits/s up). We first started ICMP ping to the last

mile hop. After 30s, we flooded the uplink (at 1 Mbits/s for AT&T and at 10 Mbits/s for Comcast, using `iperf`'s UDP measurement). After 60s, we stopped `iperf`, but let ping continue for another 30s. The ping measurements before and after the `iperf` test established the baseline latency.

Figure 13 shows the latency under load for three different DSL modems. In all cases, the latency skyrockets when flooding begins and plateaus when the buffer is saturated. This latency plateau indicates the size of the buffer, since we know the uplink draining rate. Surprisingly, we observed more than an order of magnitude of difference in buffering in different modems. The 2Wire modem introduces the lowest worst-case latency of 800ms, the Motorola modem about 1600ms, while the Westell introduces more than 10s of latency! Comcast users experienced as much as 350ms of latency under load. Because modems are usually the same across service plans, we expect that latency under load may be even worse for users with slower plans (and, hence, slower drain rates).

We perform experiments in Emulab⁸ to model modem buffering; the topology has two end-hosts and one router. We configured a token bucket filter using `tc` with the buffer size as 512 Kbits/s times the maximum latency that the modem introduces. This calculation yields 640KB for Westell, 100KB for Motorola, and 55KB for 2Wire. This simple setup almost perfectly captures the latency profile that the actual modems exhibit. 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 20KB do little for throughput, but cause a linear increase in latency under load. Our

Figure 12. Latency under load: the factor by which latency goes up when the upstream or the downstream is busy translates to significant real latencies, often in the order of seconds. (SamKnows)

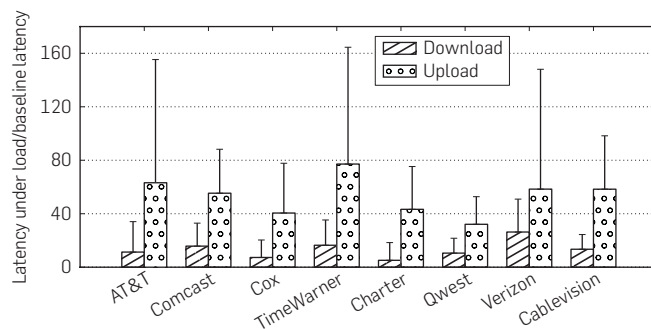



Figure 13. Different buffer sizes across modems lead to wide disparities in latencies when the upstream link is busy. (BISmark)



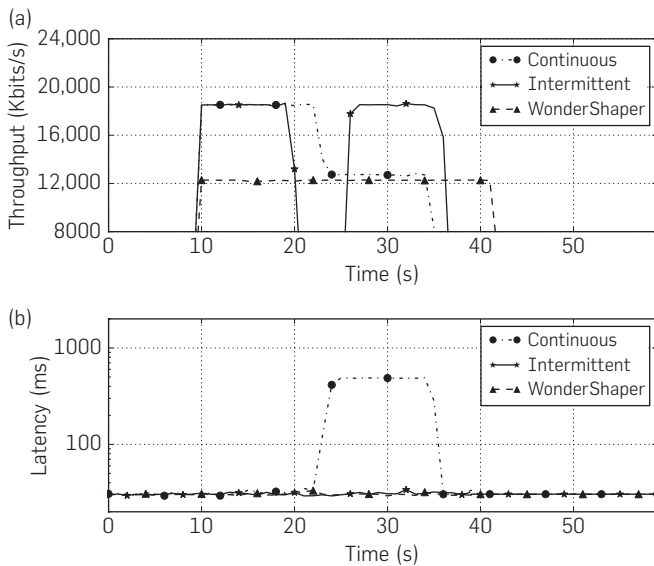
experiments confirm that buffer sizes in all three modems are too large for the uplink.

Can data transfer be modified to improve latency under load?

There has been recent work done on active queue management techniques to mitigate the effect of large buffers.¹⁹ In this section, we explore how large “bulk” flows and delay-sensitive flows can coexist without interfering with one another. We compare the effects of a 50MB download on a G.711 VoIP call in three different conditions: (1) not applying any traffic control, (2) sending intermittent traffic at capacity on 10.8s ON and 5.3s OFF cycle, and (3) shaping using the WonderShaper²³ tool. Figure 14 shows the result of this experiment. Without traffic control, the transfer takes 25.3s; immediately after the PowerBoost period, however, the VoIP call starts suffering high latency and loss until the end of the transfer. In the second scenario, traffic is sent in pulses, and the download takes 26.9s. In the third case, traffic is sent at just under the long-term rate and the download takes 32.2s. Sending intermittent traffic and shaping the traffic with Wondershaper do not increase latency much, because they do not ever fully deplete the tokens. The appropriate ON/OFF periods for intermittent transfers depend on the token bucket parameters^a and the size of the file to be transferred. Both approaches achieve similar long-term rates but yield significant latency benefit. These approaches require the user to properly tune the traffic shaping parameters according to the access link.

Takeaway: Modem buffers are too large. The smallest buffers we see induce nearly 1-s latency under load for AT&T

Figure 14. It is possible to maintain low latency by modifying data transfer behavior. (BISmark) (a) Throughput (b) Latency



^a If 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.

and 300ms for Comcast. Buffering degrades both interactivity and throughput. Transferring data in shorter bursts or shaping traffic using tools like WonderShaper can mitigate buffering problems.

5. 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 *understanding access network performance requires continual measurement from the home router*. Existing speed tests and end-to-end latency measurements do not reflect access network performance over an extended period of time, and they neglect confounding factors within the home network. Our study of broadband networks 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 performance¹⁸ 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 dimensions, 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 load time. Similarly, a gamer might require 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”²²). Our recent work proposes some first steps in this direction.²¹

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 Walter Johnston at the FCC for help and access to the data from the SamKnows study. This project is supported by the National Science Foundation through awards CNS-1059350, CNS-0643974, a generous Google Focused Research Award, the European Community's Seventh Framework Programme (FP7/2007-2013) no. 258378 (FIGARO), and the ANR project C'MON.

References

- Measuring Broadband America: A Report on Consumer Wireline Broadband Performance in the U.S. http://www.fcc.gov/cgb/measuringbroadbandreport/Measuring_US_-_Main_Report_Full.pdf.
- 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>, *Ars Technica* (Oct. 2010).
- Bauer, S., Clark, D., Lehr, W. Understanding broadband speed measurements. In *38th Research Conference on Communication, Information and Internet Policy* (Arlington, VA, 2010).
- Bode, K. FCC: One Million Speedtests and Counting. <http://www.dslreports.com/shownews/FCC-One-Million-Speedtests-And-Counting-109440>, July 2010.
- Carlson, R. Network Diagnostic Tool. <http://e2epi.internet2.edu/ndt/>.
- Comcast FAQ. <http://customer.comcast.com/Pages/FAQViewer.aspx?Guid=024f23d4-c316-4a58-89f6-f5f3f5dbdcf6>, Oct. 2007.
- Dischinger, M., Haeberlen, A., Gummadi, K.P., Saroiu, S. Characterizing residential broadband networks. In *Proceedings of ACM SIGCOMM Internet Measurement Conference* (San Diego, CA, USA, Oct. 2007).
- Emulab. <http://www.emulab.net/>, 2006.
- National Broadband Plan. <http://www.broadband.gov/>.
- Internet World Stats. <http://www.internetworldstats.com/dsl.htm>.
- Asymmetric Digital Subscriber Line Transceivers. ITU-T G.992.1, 1999.
- Asymmetric Digital Subscriber Line (ADSL) Transceivers - Extended Bandwidth ADSL2 (ADSL2Plus). ITU-T G.992.5, 2003.
- Data-over-Cable Service Interface Specifications: Radio-Frequency Interface Specification. ITU-T J.112 2004.
- Kreibich, C., Weaver, N., Nechaev, B., Paxson, V. Netalyzr: Illuminating the edge network. In *Proceedings of Internet Measurement Conference* (Melbourne, Australia, Nov. 2010).
- Maier, G., Feldmann, A., Paxson, V., Allman, M. On dominant characteristics of residential broadband internet traffic. In *ACM Internet Measurement Conference* (2009).
- Mathis, M. et al. Network Path and Application Diagnosis. <http://www.psc.edu/networking/projects/pathdiag/>.
- Netalyzr. <http://netalyzr.icsi.berkeley.edu/>.
- NetFlix Performance on Top ISP Networks. <http://techblog.netflix.com/2011/01/netflix-performance-on-top-isp-networks.html>, Jan. 2011.
- Nichols, K., Jacobson, V. Controlling queue delay. *Queue* 10, 5 (May 2012), 20:20–20:34.
- ShaperProbe. <http://www.cc.gatech.edu/~partha/diffprobe/shaperprobe.html>.
- Sundaresan, S., Feamster, N., Teixeira, R., Tang, A., Edwards, K., Grinter, R., Chetty, M., de Donato, W. Helping users shop for ISPs with internet nutrition labels. In *ACM SIGCOMM Workshop on Home Networks* (2011).
- Vorhaus, D. A New Way to Measure Broadband in America. <http://blog.broadband.gov/?entryId=359987>, Apr. 2010.
- WonderShaper. <http://lartc.org/wondershaper/>, 2002.

Srikanth Sundaresan (srikanth@gatech.edu), Georgia Tech, Atlanta, GA.

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

Nick Feamster (feamster@cs.umd.edu), University of Maryland College Park, MD.

Renata Teixeira (renata.teixeira@lip6.fr), CNRS/UPMC Sorbonne University, Paris, France.

Sam Crawford (sam@samknows.com), SamKnows, London, UK.

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

© 2012 ACM 0001-0782/12/09 \$15.00

World-Renowned Journals from ACM

ACM publishes over 50 magazines and journals that cover an array of established as well as emerging areas of the computing field. IT professionals worldwide depend on ACM's publications to keep them abreast of the latest technological developments and industry news in a timely, comprehensive manner of the highest quality and integrity. For a complete listing of ACM's leading magazines & journals, including our renowned Transaction Series, please visit the ACM publications homepage: www.acm.org/pubs.

ACM Transactions on Interactive Intelligent Systems



ACM Transactions on Interactive Intelligent Systems (TIIS). This quarterly journal publishes papers on research encompassing the design, realization, or evaluation of interactive systems incorporating some form of machine intelligence.

ACM Transactions on Computation Theory



ACM Transactions on Computation Theory (ToCT). This quarterly peer-reviewed journal has an emphasis on computational complexity, foundations of cryptography and other computation-based topics in theoretical computer science.

PLEASE CONTACT ACM MEMBER SERVICES TO PLACE AN ORDER
 Phone: 1.800.342.6626 (U.S. and Canada)
 +1.212.626.0500 (Global)
 Fax: +1.212.944.1318
 (Hours: 8:30am–4:30pm, Eastern Time)
 Email: acmhelp@acm.org
 Mail: ACM Member Services
 General Post Office
 PO Box 30777
 New York, NY 10087-0777 USA



Association for Computing Machinery

Advancing Computing as a Science & Profession

www.acm.org/pubs