Powerboost

Steven Bauer
MIT
32 Vassar St, G-814
Cambridge MA 02139
bauer@mit.edu

David Clark
MIT
32 Vassar St, G-814
Cambridge MA 02139
ddc@csail.mit.edu

William Lehr
MIT
32 Vassar St, G-814
Cambridge MA 02139
wlehr@mit.edu

ABSTRACT
This paper examines a feature present in some cable broadband networks that can enable, for multiple megabytes of data, higher data rates than what can be sustained over the long term. Commonly referred to as "Powerboost," this feature represents a strategy for sharing unused link capacity among users of cable broadband networks. We explain how Powerboost works under current implementations, consider how it may impact the experiences of broadband users, and examine the challenges that Powerboost poses for the design of performance metrics that may be used to evaluate the service quality of broadband ISPs. We present sample measurement data for a Powerboost enabled broadband connection and discuss how such measurements might be reported in the current large scale broadband measurement study led by the FCC.

Categories and Subject Descriptors
C.2.5 **[Computer-Communications Networks]**: Local and wide-Area Networks – *access schemes.*

General Terms
Measurement, Performance.

Keywords
Broadband access, Token bucket, Applications.

1. INTRODUCTION
As noted in [13], the traditional model of a network link as having a well-defined bit rate and first-in-first-out (FIFO) scheduling of traffic does not hold in broadband cable networks. Instead, a cable network has a variety of complex mechanisms for coordinating access to links that are shared by numerous users [13]. The peak data rate specified for the services that most broadband users subscribe to is usually substantially less than the total link capacity. One mechanism that commonly implements these subscription rates is a token bucket [7], which specifies the mean rate (or committed rate) in bits per second as well as the maximum burst size in bytes. Bursts of packets can be transmitted at rates higher than the committed rate.

HomeNets'11, August 15, 2011, Toronto, Ontario, Canada.

As we discuss, Powerboost will change the behavior of some applications. The presence of Powerboost will also affect how access speed is measured: depending upon the time interval over which one observes traffic different bit rates can be measured. This leads [13] to draw a distinction between the following rates in broadband access networks:

1) **Raw bit rate of the channel**. These are rates that can be measured over sub-second time intervals. This may be measured by packet pairs or even short packet trains. For DOCSIS 2.0 compliant cable networks these can be as high as 38Mbps.
2) **Bit rate for sustained transfers.** These are rates that can be measured over minute or multiple minute time intervals**.** This will correspond roughly to the mean rate specified for the token bucket mechanism governing the traffic for a user.

Our paper examines an additional rate that has become important with the advent of Powerboost. It is the rate that can be measured during the initial *n*-megabytes of traffic for a user who has not been saturating their portion of the access link for some period of time. The value of *n* is a configurable parameter selected by the broadband provider and may be different in the downstream and upstream direction.

3) **Bit rate for *n*-megabyte transfers.** The time intervals over which this can be measured depend upon *n*, the subscription rate of the user, as well as factors such as the amount of traffic from other users sharing the same link. For typical connection speeds today, this time interval is on the order of seconds to tens of seconds and can result in significantly higher data rates for the transfer of megabytes of data.

This paper attempts to document as much about Powerboost as we could discover including its history, motivation, and implementation. With this background we then discuss some potential effects of Powerboost on broadband measurement methodologies. In particular we focus on the effect of Powerboost on broadband "speed" measurements, a topic that is particularly important to both the policy and research communities [3]. We conclude with sample measurement results from a Powerboost enabled broadband connection and a discussion detailing how such measurements might be reported in the current broadband measurement study led by the FCC [19].

2. POWERBOOST
Powerboost was first deployed in broadband networks by Comcast in June of 2006 [8]. The Comcast press release announcing Powerboost described the technology and benefits in the following manner:

Table 1: Sample of cable broadband companies that implement Powerboost/Speedboost technologies and their respective download and upload byte limits if documented. This information was gathered from a search of various broadband companies' websites in March 2011 and thus may be incomplete. If we could not find an advertised Powerboost size, we simply list "Available". Speeds achieved with Powerboost will depend upon the subscription speed tier and traffic from other users sharing the same link.

Provider	*Down*	*Up*
Comcast	10 MB	5 MB
Time Warner	10 MB	No Powerboost
Cox	18-22 MB	Available
Charter	Available	No Powerboost
Rogers	10MB	No Powerboost
Shaw	Available	No Powerboost

PowerBoost technology utilizes available capacity already built into Comcast's advanced fiber network to provide customers with extra bursts of download speed – up to 12Mbps and 16Mbps, respectively, on Comcast's 6 and 8 Meg services.

The PowerBoost speed enhancer makes it faster than ever to download software, games, music, photos, and videos. For example, with PowerBoost, Comcast's 6Meg customers can download three MP3 songs (approximately 10MB) in a quick 6.6 seconds – up to eight times faster than the 53 seconds required with a 1.5Mbps DSL connection.

This suggests that Powerboost was viewed as both a way to improve the user experience for activities that benefit from higher download rates and as a way to differentiate the cable broadband service from its broadband competitors.

While initially Powerboost only boosted speeds for traffic in the download direction, Powerboost for the upload direction was added in 2007 [9]. Other broadband cable providers have since deployed Powerboost as well. Some of the providers that implement Powerboost (as of March 2011) are listed in table 1.

As can be seen in the table, Powerboost is deployed in various configurations across network providers. Comcast and Cox employ Powerboost in both the downstream and upstream direction, Time Warner Cable and Charter employ Powerboost only in the downstream direction. Powerboost is not always available for every package that a provider offers. Both Cox and Shaw, for instance, note that Powerboost is not available for their lowest tier service. Some companies adopt a different name for the technology in their marketing literature. Rogers, for instance, uses the name 'Speedboost.' Also see [14] for a measurement based study and methodology for detecting Powerboost and the parameters in use by various broadband providers.

2.1 Powerboost implementation

Powerboost may be implemented in a variety of ways. A patent [10] initially filed by Comcast in 2003 and granted in 2005 describes multiple implementation designs. These include one that dynamically adjusts the configuration of a token bucket based upon traffic demand, a design that regulates the rate in one direction based upon traffic flow in the opposite direction, a design that regulates both packet and data rates, and a design that varies the classification of packets assigned by the traffic regulator.

Based upon our discussions with network operators and equipment vendors, it appears that Powerboost is often implemented by composing existing features of the cable modem termination systems (CMTSs) and cable modems (CMs). This leverages various configurations of token buckets for each subscriber.

In the next two subsections, we examine two variants of Powerboost, *Uncapped Powerboost* and *Capped Powerboost*. While Powerboost is often advertised to users in terms like "Powerboost provides speeds up to 15Mbps for the first 10MB," for much of the deployment history of Powerboost, no such actual upper rate limit existed. A user's traffic that was eligible for Powerboost could make use of whatever spare capacity (if any) was available on the link. On a subscriber line with service that is characterized as having a 12 Mbps peak data rate with Powerboost rates up to 15Mbps, we have regularly measured speeds in excess of 20 Mbps during the Powerboost window, and on at least on occasion we have received the full capacity of the channel.

In some networks a Capped Powerboost is now being deployed. This cap limits the rate that can be achieved during the Powerboost window. These caps may have been prompted by the much higher data rates available in DOCSIS 3.0 networks. Without capping the Powerboost peak rates, users may not have as much incentive to subscribe to higher tiers of service that DOCSIS 3.0 networks enable.

As we show below, Powerboost can be modeled as a shaper with different token bucket configurations. Traffic that does not conform to the token bucket specification is deferred and rate shaped up to a limit defined by the available buffer space. Until recently the amount of buffering for best effort traffic was static and fairly large.[1] Recent revisions of the DOCSIS specification have introduced a buffer control parameter that can be used to match the amount of buffering to the subscription speed tier and Powerboost configuration.[2]

2.1.1 Uncapped Powerboost

The behavior of Uncapped Powerboost can be specified in terms of a token bucket with rate MSTR (Maximum Sustained Traffic Rate[3]) and token bucket size PBS (Powerboost Bucket Size) e.g. 10 MB in some of the previous examples. Uncapped Powerboost eligible traffic can fill the spare capacity left over after all users have had their MSTR rate's satisfied. If a link is congested and there is no spare capacity, Powerboost has no effect. The share of a congested channel that each user receives will be determined by the TCP congestion dynamics of all concurrent traffic flows.

2.1.2 Capped Powerboost

The behavior of the Capped Powerboost can be specified with an additional token bucket. This token bucket limits the peak sending rate for Powerboost eligible packets to the rate P (Peak Traffic Rate[4]) where P is greater than MSTR. In the DOCSIS

[1] See [12] for problems caused by overly large buffers.

[2] See section "C.2.2.5.11 Buffer Control" in [5].

[3] See section "C.2.2.5.2 Maximum Sustained Traffic Rate" in [5].

[4] See section "C.2.2.5.10 Peak Traffic Rate" in [5].

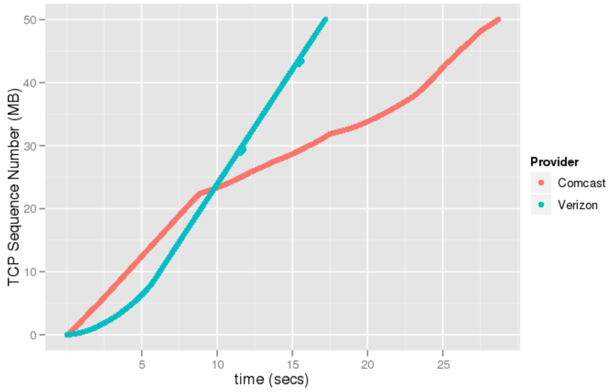

Figure 1: TCP sequence number plot comparing the download of a 50MB file on a Comcast 12 Mbps tier with Powerboost to a Verizon 25 Mbps FiOS tier. An artificial delay of 10 ms was added to the Comcast line using Linux netem so that both connections faced a ~20ms RTT to the server hosting the file.

specification document the token bucket size for the peak rate control is specified as a constant 1522 [5]. Multiple small packets therefore might be transmitted consecutively at the raw line rate, but full sized packets will be limited to the peak rate P. The two token buckets are arranged so that sufficient tokens from both buckets are required to send a packet.

2.1.3 Scheduling

Everything we have described up to this point is how traffic is shaped by Powerboost. Since each subscriber to a broadband network has their own set of token buckets in the downstream and, if available, the upstream direction, the important next question is how each user's traffic is scheduled along with traffic from other subscribers and sources. The answer to this depends upon the scheduler implementation of the shared link.

Along with best effort traffic governed by the preceding token bucket mechanisms, cable systems often will carry packets associated with other DOCSIS service flows providing video or telephony services. These flows may have higher priority access to the DOCSIS channel. But even with just best effort traffic, the exact scheduling of subscriber traffic in both the upstream and downstream direction can be complex. See Chapter 7 and Chapter 8 in [5] for an overview. The potentially complex dynamics introduced by the scheduler of a DOCSIS channel will generally perturb timing of traffic on a sub-second time scale. Longer term structures are likely to be the result of the shapers described above, cross traffic, or congestion.

2.1.4 Important other implementation details

We note the following miscellaneous details that address common questions we have faced while discussing Powerboost with colleagues:

- Powerboost is application agnostic and has no relationship to deep packet inspection (DPI) technologies.
- Powerboost is per user connection, not per TCP flow. All simultaneous TCP flows from a subscriber's cable modem will share the same Powerboost token buckets. This includes multiple users in the subscriber's household.
- Upstream Powerboost is independent of downstream Powerboost. A user can deplete Powerboost tokens in one direction without affecting the other direction.
- It has been possible to implement Powerboost since DOCSIS 1.1.
- As far as we can determine users can not strategically game Powerboost.
- There are no known security issues with Powerboost.

2.2 Powerboost formulas

In this section, we present formulas for calculating Powerboost relevant information.

Given a rate R where R > MSTR, the time duration D of Powerboost can be calculated as D = PBS/(R-MSTR). The Powerboost token bucket is refilling at rate MSTR even as it is being drained by traffic at rate R. If the maximum sustained traffic rate is 12 Mbps and the Powerboost Bucket Size is 10 MB, an R of 15 Mbps would imply a Powerboost duration of approximately 27 seconds. Higher bit rates R during the Powerboost period would imply a shorter duration D of time to complete the transfer. Also note that different speed tiers with different MSTRs will have different Powerboost durations.

The duration of Powerboost, of course, is not what matters. It is rather the amount of traffic that can be sent at a higher rate R that is important. The total bytes B that can be transferred with a Powerboost Bucket Size PBS and refill rate MSTR and average rate R will be B = MSTR(PBS/(MSTR-CIR)).

The max refill time M before the Powerboost token bucket is refilled can be calculated as M = PBS/MSTR. Assuming a PBS again of 10MB and a MSTR of 12 Mbps, the refill time M will be approximately 7 seconds.

2.3 Impact of Powerboost on traffic

The impact of Powerboost on traffic can be seen in the TCP sequence number plot of Figure 1 where a 50 MB file is downloaded on both a Comcast 12 Mbps tier with Powerboost and a Verizon FiOS 25 Mbps tier. The speed of the transfer can be inferred from the slope of each line. A steeper slope indicates a faster transfer.

In this test, up to 23 MB of data was downloaded more quickly on the Comcast connection. After Powerboost depleted, the Verizon connection downloaded the rest of the 50 MB file approximately 10 seconds faster. So which network is faster depends upon the file size.

Note that this is a comparison of single TCP transfers. It is generally representative of the tests we have taken from these two connections and a limited sample of other Powerboost enabled connections. The speeds measured during a Powerboost interval depends upon the amount of other competing traffic for the channel. We do not yet know if the Powerboost speeds we can measure in the Boston area are representative of Powerboost speeds in other locations. The FCC/SamKnows study discussed below will eventually shed some light on this question.

The benefit of Powerboost is that users whose average rate has been below their CIR can benefit from available capacity. This raises the question of why, if there is spare capacity, it is not made available to all users all the time i.e. by simply raising everyone's tier of service. The answer to this is that Powerboost enables providers to target the spare capacity for use by subscribers who do not put a sustained load on the network. This may be considered by providers to be a more desirable distribution of capacity [2].

2.4 Impact of Powerboost on applications

While the impact of Powerboost can be demonstrated with traffic traces, in this section we explore whether or not common classes of applications and services actually benefit from it. We argue that Powerboost does provide a benefit, but as we have noted elsewhere [3], speed is only one dimension of a complicated mixture of factors that can impact application performance. Speed may not be the dominant factor for many classes of applications. Furthermore no application or service actually requires Powerboost, as application and service designers build systems that operate over a wide spectrum of broadband connection capabilities. Further, networks that provide Powerboost may not always have the excess capacity to realize it, depending on the instantaneous level of user demand[5]

From the statistics that Google gathers from their web crawls, 90% of all web pages had less than 663KB of total content. Web pages often consist of content loaded from multiple hosts. 90% of web pages loaded 179KB or less from each host [15]. With these object sizes, the majority of web connections finish before TCP exits its slow start phase. [11] Thus the committed rate or Powerboost rates are not dominant contributors to the total transfer times. However, higher Powerboost rates still do have an impact. A formula for calculating the latency of a transfer completing during slow start without losses [6,11], is:

$$\left\lceil \log_{\gamma}\left(\frac{s(\gamma-1)}{init\_cwnd}\right)+1\right\rceil * RTT + \frac{S}{C}$$

where S is transfer size, C is bottleneck link-rate, γ is 2 or 1.5 depending upon where delayed TCP acknowledgements are employed, and init_cwnd is typically 3 or 4. Plugging some sample numbers into this equation with γ=1.5, init_cwnd=3, RTT=25ms, and S=179KB, one finds a total transfer time of 422ms for C = 12Mbps, 398ms for C=15Mbps, and 361ms for C=24Mbps for a total difference from the highest to the lowest of about 60ms.[6]

Some of the factors that cause most web browsing to be so driven by latency are (potentially) being addressed with larger initial congestion windows [11], DNS prefetching [16], improved rendering and Javascript engines [17] that mitigate the local processing times, and proposed new protocols such as SPDY [18]. To the degree that these are deployed, and web pages continue to increase in size, the impact of Powerboost for web browsing may increase.

The impact of Powerboost is somewhat limited for web browsing because the content transferred (in terms of MB) is typically small. This is why the target of Powerboost was to improve the performance of larger MB transfers for such content as music, photos, games, email attachments, etc. Today one might also list apps for mobile devices, podcasts, and ebooks all downloaded over the home network. In these usage cases, the volume of content and transfer duration are long enough for TCP to ramp up its sending rate to a level that can benefit from Powerboost and also short enough that Powerboost can make a discernable impact on the overall transfer time. Even though Powerboost runs out after a certain amount, its effect on the total transfer times can still be significant if the Powerboost duration is a non-trivial portion of the total transfer time. Obviously, if the file transferred is a 400MB operating system update, then Powerboost's impact on the average data rate and total duration of the transfer will be much less significant.

Online gaming is one class of applications traditionally thought of as almost exclusively concerned with latency. But even here, there are numerous points where download rates matter. In the case of browser based Flash games the initial download of the application can benefit from Powerboost rates. For non-browser based games, transferring the game map or virtual world can also be accelerated by Powerboost. While Powerboost may only benefit the ‘startup’ costs of gaming sessions, these may have a significant impact on users' appreciation of the gaming experience.

The last application category we consider is video. Again at first glance this is a category where latency, jitter, and reliability are considered the dominant drivers of application performance. The typical transfer rates for video, 6 Mbps or less in 2011, are well below the level where Powerboost would have any apparent effect. However, many video players download a burst of video in the initial seconds to rapidly fill a playback buffer [1]. This initial burst, as well as re-buffering that occurs if there is any interruption in service, can be accelerated by Powerboost. This may contribute to a better user viewing experience, one better protected from dropouts resulting from temporary congestion or network disruptions.

Powerboost can potentially cause problems for some applications if an estimate of the long term achievable throughput is derived from a Powerboost enabled time window. For instance, a video conference application might select a codec for video that is non-sustainable if Powerboost is not in effect (particularly in the upload direction.) Most applications however are capable of adapting to varying network conditions.

3. BROADBAND MEASUREMENT AND POWERBOOST

For the past nine months we have been one of the academic groups participating in the FCC broadband measurement project which has now deployed approximately 10,000 SamKnows [19] measurements boxes on broadband networks throughout the United States. As part of this project, we have been studying and offering feedback based upon our analysis of data collected from a small number of SamKnows boxes that we have deployed for testing purposes in the Boston area.

Part of our motivation for writing this paper stemmed from studying the data from these boxes, particularly as the test suite employed has evolved. We have come to realize that Powerboost can have some subtle, non-obvious effects on broadband measurements. Understanding these effects is important for the FCC study, but more generally documenting the lessons learned is important for the research community as it increasingly deploys its own novel broadband measurements on networks with Powerboost.

[5] The availability of Powerboost at any particular moment depends on the degree of demand on the system and the resulting level of congestion. This paper does not attempt to measure the variation in congestion across these networks; our limited data is only intended to demonstrate the character of Powerboost. The larger SamKnows data will give insights into the level and variation in congestion on these networks.

[6] Seemingly small differences in latency can influence user behavior [4].

Table 2: Average speeds for all bi-hourly download and upload tests during January 2011 as measured by FCC/SamKnows measurement boxes on a Comcast (12/2 Mbps) connection and a Verizon (25/25 Mbps) connection. Speeds are first calculated for each individual test over the time intervals indicated and are then averaged with other tests measuring the same interval. The speeds in the "+ warm_up" columns include the byte count and time from the warm-up period of each test. We include tests to all target servers.

	Comcast (Mbps)				Verizon FiOS (Mbps)			
	Download	+ Warm-up	Upload	+ Warm-up	Download	+ Warm-up	Upload	+ Warm-up
Warm-up	28.25		2.36		26.79		25.72	
0-5 secs	23.80	24.34	1.95	2.18	30.58	30.16	27.02	26.36
0-10 secs	18.30	18.95	1.95	2.11	30.52	30.31	27.00	26.57
0-15 secs	16.42	16.96	1.95	2.08	30.50	30.36	27.01	26.68
0-20 secs	15.48	15.93	1.95	2.05	30.49	30.38	27.01	26.74
0-25 secs	14.92	15.29	1.95	2.03	30.49	30.40	27.01	26.79
0-30 secs	14.54	14.86	1.95	2.03	30.49	30.42	27.03	26.84
25-30 secs	12.64		2.00		30.52		27.12	

3.1 General impact

In §2 we noted that the max refill time M before the Powerboost token bucket is refilled could be calculated as M = PBS/MSTR. This can be important when scheduling broadband tests because tests that are conducted too close together in time may have a first test that measures during the Powerboost interval and a second test that measures the non-Powerboost periods. This may invalidate an assumption of independent tests.

We noticed this effect when the order of tests happened to be changed, bringing two download speed tests back to back. The first test to an on-net server had depleted the download Powerboost, resulting in the second test to an off-net server measuring significantly lower speeds than had been recorded previously. Failure to note the impact of Powerboost might lead one to falsely infer a much larger gap between on-net and off-net data rates.

In a related vain, user traffic can also deplete the Powerboost interval potentially making test results dependent upon how closely they occurred in time to user activity. The Samkows test box addresses this issue by deferring its broadband tests if the user is actively using the broadband connection and it also monitors traffic both prior to and during testing and defers or cancels tests if traffic in either direction exceeds a threshold.

This is not to say that all tests should be conducted while Powerboost is in effect. Indeed, if there were differences in broadband metrics after Powerboost had ended, they would be interesting to quantify. But having some level of confidence that the two contexts can be separated would be desirable.

3.2 Speed measurements

While there has been less debate in the FCC study over what should be measured, how to conduct the measurements and what to report remain contentious issues. This is particularly true when it comes to speed measurements.

A speed test in the FCC/SamKnows study employs three TCP connections to a dedicated test server. There is a variable length warm-up period that attempts to drive each TCP through its initial slow start to a rate that fully saturates the link. After the warm-up period, the test runs for 30 seconds and records the average throughput every 5 seconds for the cumulative amount of traffic sent so far. More details are described in the SamKnows methodology document [19].

This is different than a test that might measure what users or applications achieve in a bulk transfer test. The test is targeted at accessing what rates are achievable when the access network can be fully loaded with multiple already accelerated TCP connections. Thus discounting TCP startup effects and various congestion response dynamics of TCP was important. One motivation for this methodology was to have a test that enabled comparisons of advertised speeds to provisioned speeds.

3.3 Sample of speed measurement results

Table 2 presents the average speeds for all bi-hourly tests during the month of January 2011 for two of our SamKnows test boxes deployed in the Boston area. The point of this table is to demonstrate the potential effect of both Powerboost and the warm-up interval employed by this particular speed test, not compare these two providers. The question is what speeds are fair (we assume their accuracy) to report for each provider.

For the download tests, the samples from the 25-30 second test interval are consistent with sustained rates we have measured using other tools. Comcast measured 12.6 Mbps and Verizon 30.5 Mbps. We view these as important numbers to report.

Establishing a download speed for an initial interval however is harder. For Comcast, the average speed recorded during the warm-up period was 28.25 Mbps which lasted an average of 738ms and transferred an average of 2.4MB of data. On Comcast the average for the other test intervals is also higher if one includes the warm-up interval. For the 0-5 second interval, the speed is 23.80 Mbps without the warm-up interval, but 24.34 Mbps with it included. This means the warm-up period actually results in lower measured speeds for Powerboost connections. Instead of allowing TCP time to ramp up, the warm-up period is depleting the Powerboost token bucket.

The warm-up period for the Verizon tests records markedly lower speeds (at 26.79 Mbps) than the other test intervals (~30.5 Mbps). Thus the warm-up period plays its designed role for these tests of allowing TCP time to ramp up its sending rate. (During the

Verizon warm-up period and average of 1.9MB of data was transferred in 604ms on average.)

In the upstream direction, we find a similar issue. In the Verizon upload tests the warm-up period records 25.72 Mbps while all other test periods record ~27.0 Mbps. Again on the Comcast network the warm-up period appears to consume the available Powerboost token bucket as the highest average speeds are recorded during that time period (2.36 Mbps). Including the time and bytes of the warm-up period in calculating the other intervals also results in larger measured speeds (e.g 2.18 Mbps compared to 1.95 Mbps for the 0-5 second interval.)

This suggests that fair treatment of different providers might require different treatments of the warm-up period. While it helps some providers, it results in underreporting speed for providers which employ Powerboost.

If one reports the 25-30 second interval, the warm-up period has no effect. But for other intervals in the download (and analogously for upload) speeds for each interval might better be calculated as:

Speed=Max(Download, Download + Warm_up)

For these two test boxes this would result in an initial period speed of 24.34 Mbps for Comcast and an initial period speed of 30.85 Mbps for Verizon. Taking differences between successive Comcast readings allows us to calculate the speeds during 5 second intervals: once Powerboost is ended (essentially at the 5 second mark), the speeds are stable to less than 1% between 12.64 and 12.69 mb/s, virtually equal to the observed speed of 12.64 in the 25-30 second range.

In the FCC broadband study, and other similar studies, we would like to see reports that detail performance during initial periods of a speed test and the last 5 or 10 seconds of the test. These would shed light on both Powerboost speeds (if present) and sustained speeds. Reporting some mix of the two speeds e.g. 0-20 seconds would be the harder to interpret as it is not representative of either Powerboost or the sustained rate.

4. CONCLUSION

Shaping of traffic is often thought of as something that decreases the sending rate relative to the status quo. Powerboost is, in a sense, a positive shaper which enables traffic to be sent at a rate that is higher than the normal rate. We have argued in this paper that this shaping does benefit broadband traffic and the user experience. So along with other broadband metrics, we hope to see more measurements of both sustained rates and rates during intervals when Powerboost is in effect; both are important.

5. ACKNOWLEDGMENTS

We acknowledge the support of participants in the MIT MITAS and CFP research programs. Along with other members, these research programs include cable, DSL, and wireless network operators. SamKnows has also been very helpful providing both access to data, test boxes, and answers to our endless list of questions. Any opinions expressed and errors are solely our responsibility.

REFERENCES

[1] S. Akhshabi, A Begen, and C. Dovrolis. An experimental evaluation of rate-adaptation algorithms in adaptive streaming over HTTP. In Proceedings of the second annual ACM conference on Multimedia systems (MMSys '11). ACM, New York, NY, USA, 2011.

[2] S. Bauer, D. Clark, W. Lehr, "The Evolution of Internet Congestion,", 37th Research Conference on Communication, Information and Internet Policy. September 2009.

[3] S. Bauer, D. Clark, W. Lehr, “"Understanding broadband speed measurements," 38th Research Conference on Communication, Information and Internet Policy. Oct 2010.

[4] J. Brutlag, “Speed Matters for Google Web Search,” Google. June 2009.

[5] Cablelabs, SP-MULPIv3.0: MAC and Upper Layer Protocols Interface Specification, Feb 10, 2011.

[6] N. Cardwell, S. Savage, and T. Anderson. Modeling TCP Latency. In Proceedings of IEEE Infocom, 2000.

[7] Cisco, Policing and Shaping Overview , Cisco IOS Quality of Service Solutions Configuration Guide, Release 12.2.

[8] Comcast Powerboost press release, Jun 2006. http://www.comcast.com/About/PressRelease/PressReleaseDetail.ashx?PRID=65

[9] Comcast Powerboost press release, Aug 2007. http://comcastcalifornia.mediaroom.com/index.php?s=43&item=170

[10] Comcast, “Method and packet-level device for traffic regulation in a data network,” Patent 7289447, Oct 30, 2007.

[11] N. Dukkipati, T. Refice, et al “An argument for increasing TCP's initial congestion window” SIGCOMM Comput. Commun. Rev. 40, 3 (June 2010).

[12] J. Gettys, Bufferbloat: Dark Buffers in the Internet, IEEE Internet Computer, May/June 2011.

[13] K. Lakshminarayanan, V. Padmanabhan, and J. Padhye. 2004. “Bandwidth estimation in broadband access networks” (IMC '04).

[14] P. Kanuparthy and C. Dovrolis , End-to-end Detection of ISP Traffic Shaping using Active and Passive Methods, George Tech Technical Report, May 2011..

[15] Google, Web metrics: Size and number of resources. http://code.google.com/speed/articles/web-metrics.html

[16] Google, DNS Prefetching. http://www.chromium.org/developers/design-documents/dns-prefetching

[17] Google, V8 Javascript Engine, http://code.google.com/p/v8/

[18] Google, SPDY: An experimental protocol for a faster web. http://www.chromium.org/spdy/spdy-whitepaper

[19] Samknows, FCC Study, Test Methodology Document. (Public Release Pending).