

Online Broadband Test Tools

Lars-Petter Gunhildsberg Hansen, Ivar Conradi Østhus
and Poul Heegaard
Department of Telematics,
Norwegian University of Science and Technology

Abstract

The Norwegian Post and Telecommunications Authority (NPT) will launch an online tool for evaluation of the end-users' Internet connections based on active measurements between end-users and a measurement server located at the NIX. In this paper we study existing active measurement tools with respect to their features and precision. The experiments showed that challenges exists for some of the tools when the bandwidth is higher than 12 Mbit/s, when the bandwidth is not asymmetric, and when JavaScript or Flash is used.

1 Background

The majority of Norwegian households have access to a broadband Internet connection today [15]. The Internet service providers are continuously increasing the data rates of their broadband products in the fierce competition for customers. Bandwidth and price have become the most important selling pitch, and the operators are constantly trying to surpass each other. They all claim to deliver connections with a certain capacity, but the specifications are not uniform and often quite vague. "Up to" is a term that is commonly used by the Internet service providers when the capacity of broadband products are specified.

The usage of the private Internet connection has undergone a change as the typical contents communicated have developed from simple web pages to larger quantities of data and contents with real-time requirements. This development increases the users' demands related to the quality of the Internet connection. Customers are interested in getting the connection quality they pay for, and also be able to check

This paper was presented at the NIK-2009 conference; see <http://www.nik.no/>.

whether this is the case or not. This raises a need for reliable broadband test tools. There exists numerous online tools for broadband evaluation today, and this is a popular service provided by many different web pages.

The Norwegian Post and Telecommunications Authority (NPT) will develop and release an online tool for evaluation of the end-users' Internet connections. This work is inspired by the establishment of such a tool in Sweden where The Swedish Post and Telecom Agency (PTS) was participating. NTP's concern is that online tools are financed by advertisement and that commercial interests may affect how the test results are presented to the end-users, that the test servers usually are located in non-neutral locations producing biased test results and that the test tool principle and quality is unknown. Hence, NTP will launch a non-commercial tool in 2009 with measurement servers at neutral locations in the Internet. We will refer to this service as NPT Broadband Test Tool (NBTT).

The numerous tools that already exist for broadband evaluation are either based on active or passive measurement techniques [13]. Passive techniques only observe the real-traffic on the link, while active techniques actively insert artificial traffic to measure the link capacity. Both techniques have their strengths and weaknesses and in this paper we focus on tools based on active measurement techniques since the NBTT is an active measurement tool. Most tools focus on up- and download data rates measured on different protocol layers. Other aspects associated with a broadband connection are also important, and different users and applications have different requirements.

In this paper the main objective is to analyze and benchmark the existing tools. This paper presents the key findings from our master thesis [7]. Section 2 provides a brief overview over performance metrics and a selection of online and stand-alone tools that can be applied for assessment of these properties. In Section 3 benchmarking of the tools provides a better insight in how the tools perform in practice, and relate to how the planned NBTT will evaluate the broadband capacity. In Section 4 the paper is closed with a summary and a look into the future for potential further development of NTP's test tool.

2 Broadband Test Tools

The network performance metrics that are considered in this paper include; connectivity (ability to connect to ISP), delay (round trip time), jitter (variation in delay), loss ratio, throughput, utilization, available bandwidth (current) and capacity bandwidth (maximum in contract). The broadband test tool considered includes *online* tools (accessed through web) and *stand-alone* tools (installed).

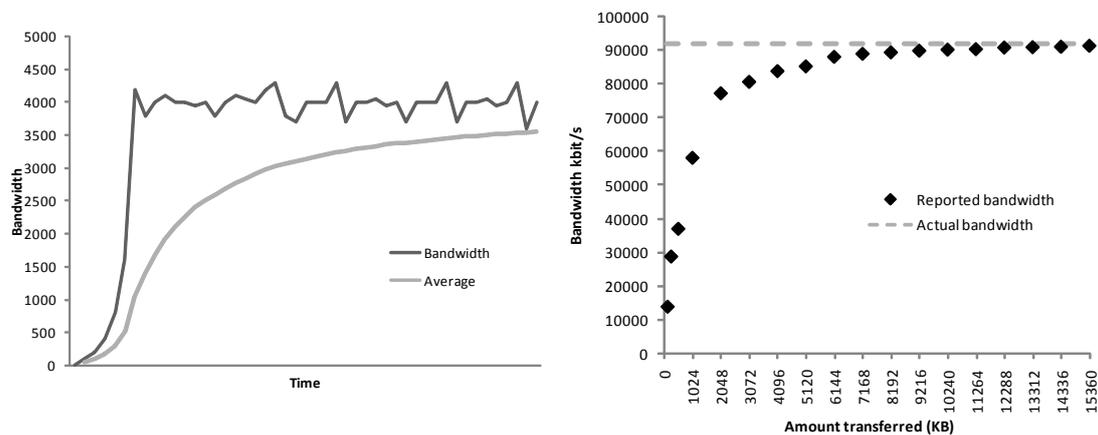
Table 1: Comparisons of online test tools

	Speedometeret	Speedometeret JS	Bredbandskollen	Speedtest	MySpeed	Network Diagnosis Tool
Download rate	x	x	x	x	x	x
Upload rate	x		x	x	x	x
Round trip time (RTT)			x	x	x	x
Bulk transfer capacity	x	x			x	x
Achievable throughput			x	x		
HTTP-based transfer	x	x	x	x		
TCP-based transfer					x	x
Number of connections, download	1	1	2	2+	1	1
Number of connections, upload	1	n/a	2	2+	1	1
Server selected by user?	No	No	Yes	Yes	Yes	Yes
Java applet	x				x	x
Flash			x	x		
JavaScript		x				

Table 1 shows the online broadband test tools reviewed in this paper and a comparison of some of the key attributes. All the tests tools measure download and upload bandwidth and round trip time (RTT) except the Speedometeret. Speedometeret from ITavisen (aka Surfometeret @ Aller Internet) is a very simple bandwidth test which measures the download and upload rate at the application layer (HTTP). The actual throughput is underestimated since only the size of the transmitted file, and not the total amount of transmitted data is included. Time used to transfer a file over HTTP cannot be used to determine how the connection behaves when used with other types of applications like streaming, VoIP or other. Furthermore, the Speedometeret does not measure the *RTT*/latency, which is an important parameter even in web surfing performance evaluation. A user is most likely not satisfied with a connection with a large bandwidth but with poor responsiveness. Bredbandskollen and Speedtest also use HTTP. We have analyzed these tools with Wireshark [6] and found that they operate almost identically, except in the location of the test servers. MySpeed and Network Diagnosis Tool (NDT) both use TCP and measure the TCP goodput directly. This is less application dependent than HTTP (however, assumes TCP and not UDP) and will provide a slightly better accuracy in comparison with Speedometeret.

Common for all test tools is that the bandwidth estimates are based on transfer

of data/time. This assumes that the transfer instantly utilizes 100 % of the bandwidth, which is not the case in TCP, see Figure 1(a) that illustrates how a TCP connection acquires bandwidth over time as a result of TCP congestion mechanism (slow-start). In the slow-start period the bandwidth will be underestimated. The tools we have evaluated solve this in different ways. Speedometeret, Bredbandskollen and Speedtest do a pre-calculation of speed and capacity and select an appropriate transfer file size. In addition Bredbandskollen and Speedtest use a *sliding average* that efficiently eliminates the slow-start error by only including the last part of the download in the calculations. MySpeed and NDT operate on the TCP layer and have defined transfer times of 8 and 10 seconds, respectively. Both these solutions prevent the slow-start in messing up the result. The JavaScript version of Speedometeret only downloads a predefined picture file. The problem is that a fixed file size is applicable only for a specific range of bandwidth capacities. To illustrate the effects the file size may have on



(a) Effect on bandwidth estimation

(b) Variable data transfer volume

Figure 1: TCP's slow-start mechanism influences the bandwidth estimation

the reported bandwidth we used the Iperf [9] tool on computers connected to the network at NTNU. The Iperf client and server should in theory have a 100 Mbit/s Internet connection. We used an Iperf server in Oslo inside the UNINETT network. To measure the bandwidth between the server and client we first ran an Iperf test for 20 seconds. This reported a bandwidth of 92,029 Mbit/s. Then we ran Iperf tests with defined transfer sizes from 128 KB up to 15 MB. The results of these tests are plotted in figure 1(b). The bandwidth reported in the 20 seconds test, are shown with the dashed line. Our experiment shows that adjustment of download size in accordance with the bandwidth of the connection is essential.

Table 2: Comparison of stand-alone test tools.

	Ping	Traceroute	Iperf	Netperf	TPTEST 5.02	Pathload	Abget
Round trip time	x	x		x	x		
Loss	x		x	x	x		
Jitter	x		x		x		
Bulk transfer capacity			x	x	x		
Achievable bandwidth			x		x		
Available bandwidth						x	x
Request/Response performance				x			
TCP			x	x	x		x
UDP			x	x	x	x	
ICMP	x	x			x		

In addition to the online tools, we have studied a selection of active stand-alone measurement tools. They are all analyzed with Wireshark. The evaluated stand-alone tools, except Abget, are all using either TCP, UDP or ICMP. Table 2 summarizes the stand-alone tools. The performance metrics and supported protocols are listed. Neither of the tools assess all aspects of a broadband connection.

Ping [14] is an excellent tool for checking the bidirectional connectivity from one host to another. Ping measure the *RTT* sending *ICMP* packets. In a controlled environment, ping has reported high precision in order of 0.1 ms in Linux[1], but ICMP might get blocked or delayed by the operator, or experience that the end host is not responding to these messages in a timely manner. Ping is also applied for loss ratio estimation. Ping records the number of packets sent, and number of received responses, where the difference is the experienced loss. Ping sends ICMP packets regularly and will not necessary reflect the loss experienced by the users both due to difference in traffic pattern and protocol type. Traceroute [17] uses ping repeatedly with increasing Time To Live (TTL) and is a tool for finding the path traversed by packets through the network. TPTEST uses ping to measure the *RTT*, jitter and loss rate and presents the results in its own graphical user interface, which is the same as ping without the option of specifying any parameters, such as message size or the number of probe packets.

Iperf and Netperf [12] measure the loss probability by UDP. This is a better than ICMP because *UDP* is often used to transmit loss- and time-sensitive information. Both Iperf and Netperf allow the user to define the content of the messages,

which makes it possible to use these tools to mimic application layer protocols. Iperf, Netperf and TPTEST are all tools which can be used to measure the Bulk Transfer Capacity (BTC), i.e. a paths ability to transfer large files. While Iperf and Netperf are applications which allows (or requires) the user to set-up its own server, TPTEST depends on the pre-established TPTEST-servers. For advanced network debugging and performance tuning it is essential to be able to define both hosts in the end-to-end path being evaluated. The server location is important and will affect the obtained results. A user located in Norway running a *BTC* test to a server located in Sweden will measure the capacity for the whole path.

Iperf and Netperf can be applied to measure the *achievable bandwidth* (i.e maximum capacity over TCP or UDP) through either (i) use of multiple *TCP* sessions at the same time to give the application a higher share of the path-capacity, or (ii) use of *UDP* for transport. *Available bandwidth* can be measured with tools like abget and pathload. The motivation for these tools is to estimate a paths' available bandwidth, and to be as little intrusive as possible. The available bandwidth is useful for applications over UDP to adjust their sending rate. Pathload is a fast and accurate available bandwidth estimation tool [10]. In Pathload the sender transmits periodic *UDP* to the receiver at a certain rate. Abget has another approach where the client is able to estimate the available bandwidth to "any" web server on the Internet holding a file with a suitable size. The idea is to make available bandwidth estimation as easy as ping and traceroute. Abget is able to do this because this tool modifies the *TCP* behavior and utilize *ACKs*, but requires root privileges and is dependent on libraries that only are available on UNIX based systems.

3 Evaluation of Tools

Several tools were introduced in previous section. In this section benchmarking of different online broadband test tools is conducted and cross-validated against Iperf-measurements towards a server located at the Norwegian Interchange Point (*NIX*). The test setup is shown in Figure 2. The tests consist of independent experiments repeated ten times. The Test Computer connects through a Bandwidth Limit Server, which operates as a network bridge that acts as a packet rate limiter. This enables us to emulate *ADSL* broadband connections and to change the bandwidth capacity. A few of the existing rate limiter tools were tested, including Traffic Shaper XP Client [3], NetLimiter [16], Trickle [4] and Wonder Shaper [8]. Unfortunately, we found them to be inaccurate, unstable and to deliver unpredictable results. Their common problem is that they operate in user mode, which add extra variable delay and CPU scheduling problems. For

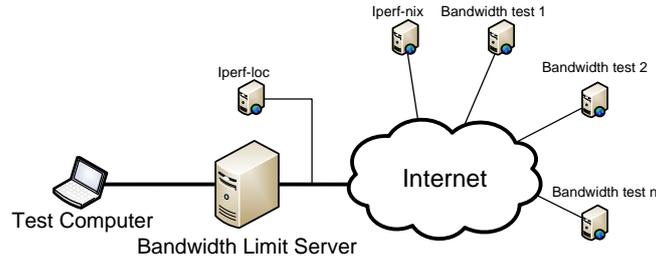


Figure 2: Functional overview of the Test Setup.

our testing, we need a rate limiter that allows us to limit the total traffic rate to a specified rate, whether it is one flow or multiple simultaneous flows. So, we configured a Linux server as a network bridge, see Figure 3, and used the Linux based tool *tc* [2]. This is a tool used to configure the Traffic Control module in the Linux kernel. This tool enables us to shape the traffic to a specified rate in a random early detection (RED) queue discipline [5]. The server uses a queue in both directions¹, which will add some extra delay to the sent and received packets.

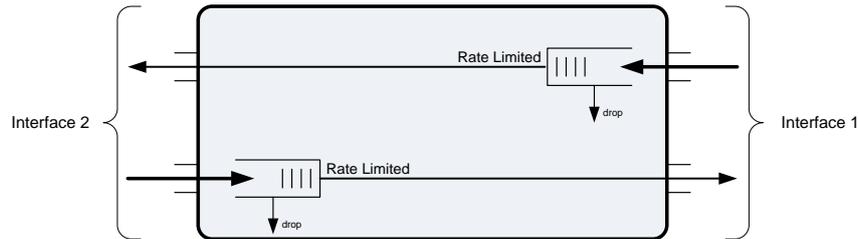
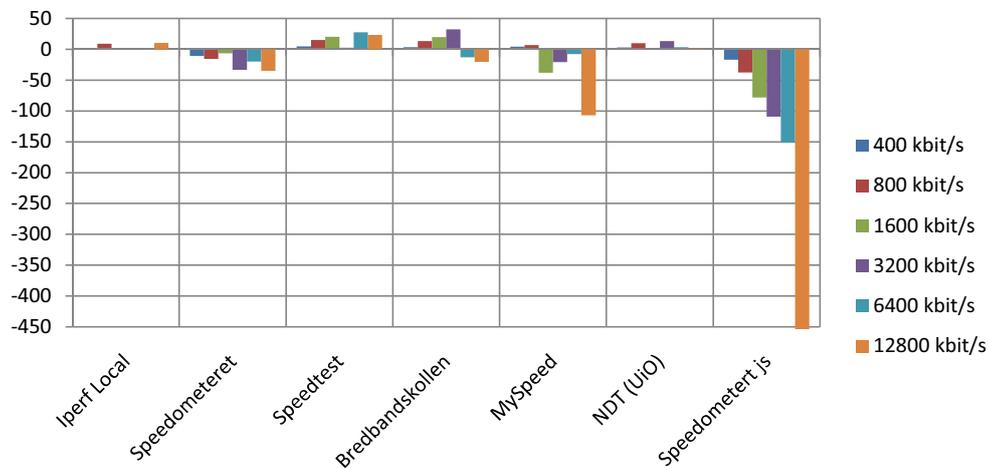


Figure 3: Network interfaces of the Bandwidth Limit Server.

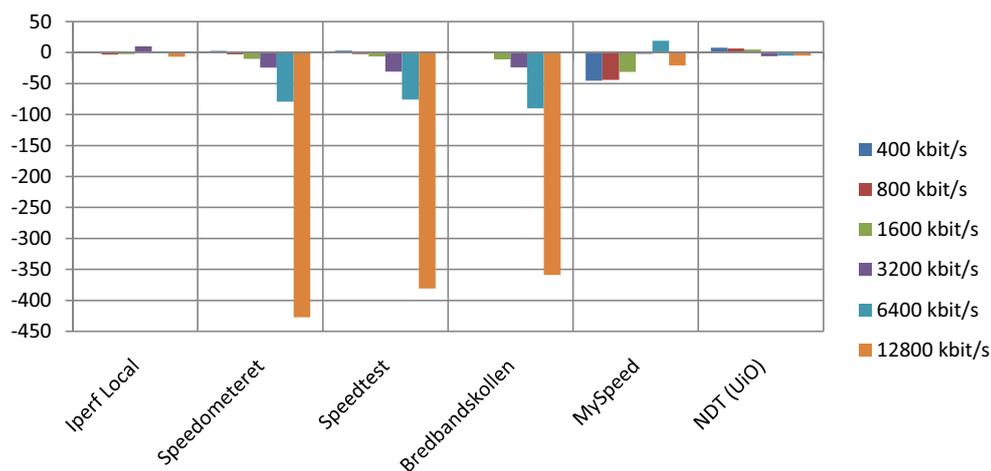
The test results are affected by many factors in the test setup. The Internet has many competing users producing *cross traffic* that will cause temporal resource contention. If the *rate limit* technique is inaccurate this could affect the results and might lead to a wrong conclusion. The *queue discipline* is *RED* which makes it easier for *TCP* flows to adapt to the available rate. The test tools perform measurements on different *protocol layers* which make the cross validation inaccurate. The Test Computer and its network connection might be busy with parallel processing tasks or network traffic. We have used passive monitoring to ensure that we know when this happens. The location of the *broadband test server* will influence the accuracy. We experienced this problem with the MySpeed server.

¹Also NextGenTel uses queues in both directions for their ADSL subscribers.

Problems with *poor cables* and *unreliable switches* might cause problems like high loss probability and varying bandwidth. To evaluate the different online test tools



(a) Download



(b) Upload

Figure 4: Rate deviation from Iperf-nix in kbit/s.

we connected our Test Computer to the Internet and executed each broadband test-tools *ten* times at different rates. By running each tool multiple times at each rate we ensure that transient capacity changes on the Internet should not influence the final result too much. The tested rates are 400 kbit/s, 800 kbit/s, 1600 kbit/s, 3200 kbit/s, 6400 kbit/s and 12800 kbit/s in both upload and download direction. To find the actual *TCP* goodput we have established two Iperf-test servers. One server is located locally (Iperf-loc), on the same switch as the Bandwidth Limit

Server, while the other, Iperf-nix, is located near *NIX* as illustrated in figure 2. The Iperf-nix server gives a natural location of a bandwidth test-tool, at/or nearby a neutral exchange point on the Internet. We use this server to calibrate our results and all other tools are compared with Iperf-nix. The Iperf-loc is used to verify the results reported by Iperf-nix and ensures that nothing else disturbs our results.

Figure 4 shows how much each tool deviates on average from Iperf-nix in kbit/s. Most tools underestimate the bandwidth, except Bredbandskollen and Speedtest, which seem to overestimate the download rate². When we increase the rates to 6400 kbit/s and above, we see that Speedometeret, Speedtest and Bredbandskollen really start to underestimate the upload rate. We think this is because these tools are designed for asymmetric bandwidths where they assume that the upload rate is much smaller than the download rate and too little data is sent. Speedometeret completes its upload test at the highest rate in just 0.6 seconds which is an indication that a too small amount of data is used. The figure also illustrates that Speedometeret generally measures almost the correct value with a little underestimation. We think this is because Speedometeret uses *HTTP* to perform the bandwidth test without accounting for the overhead, and thus actually measure the *HTTP* payload throughput. MySpeed underestimates the highest rate. We suspect that this is a result of the server location which is far away from the Norwegian backbone, and we have to pass links of 2.5 Gbit/s to reach the MySpeed server in Ålesund.

Many more tests have been conducted and more results details of the test plans can be found in [7]. Important findings include:

1. Speedometeret JavaScript - JavaScript timing issues make the JavaScript version of Speedometeret less accurate than the Java based version. This has clearly been verified by the download result in Figure4(a) and more results in [7]. The deviation in kbit/s measured compared with Iperf-nix gets worse as the download rate increases.
2. High Speed Download Testing - As more and more subscribers get higher and higher download rates, it is becoming more important that the broadband evaluation tools can handle these increased rates. This has been tested for 20 Mbit/s, 40 Mbit/s and 60 Mbit/s in Figure 5. The tools handle the high rates fairly well except Speedometeret which underestimates the 60 Mbit/s connection with 12 %. MySpeed cannot be applied at high rates, e.g., it reported 22 Mbit/s on a 100 Mbit/s link. Bredbandskollen underestimates

²To verify that Speedtest and Bredbandskollen actually gets a higher bandwidth we performed bandwidth measurements against Iperf-nix, Speedtest and Bredbandskollen at our private ADSL-connection (NextGenTel). We found the same "overestimation" in these measurements. This verifies that our Bandwidth Limit Server works like it should.

the bandwidth with 0,86 % at the highest rate probably because the server is in Sweden.

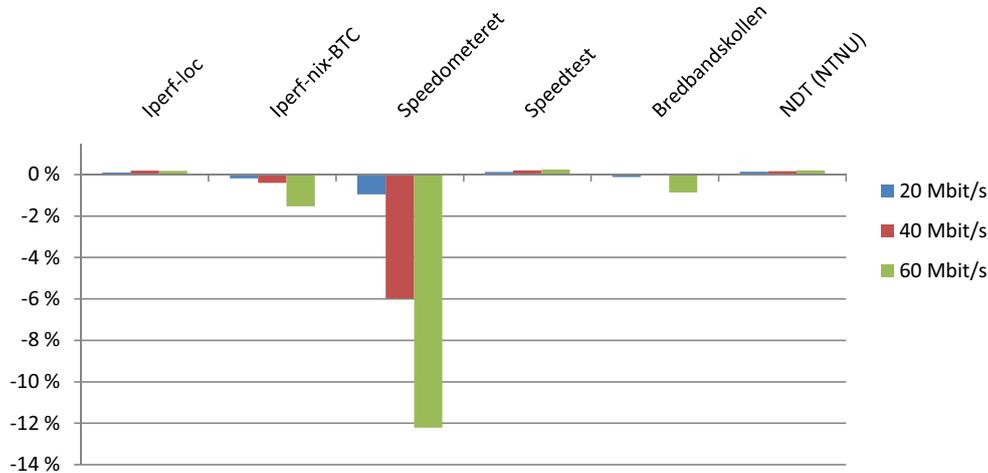


Figure 5: Percentage deviation from Iperf-nix download rate.

- Flash Upload Bandwidth Issues - During our testing we discovered that the Mozilla Flash plug-in has some performance issues with high speed upload rates. The problem seems to be related to the Windows Flash plug-in³ used by Opera and Firefox. Figure 6 shows the measured upload rates for different browsers and operating systems when our Test Computer was connected to a 100 Mbit/s access and executed the Bredbandskollen test tool. The upload rates are significantly lower for Firefox and Opera on the Windows platform, than for Internet Explorer in Windows and Firefox in Linux. This has been confirmed in a parallel study by Artur Janc master thesis [11]. This limitation with the Flash plug-in must be taken in to consideration when designing a broadband test tool that should support high speed clients.

4 Closing Remarks

We have in this paper tested different tools end-users can use to evaluate their broadband access. Generally we have seen that most of the tools are able to measure the download rate rather accurately up to 12,8 Mbit/s. When we tested the tools at higher download rates (20 Mbit/s, 40 Mbit/s and 60 Mbit/s) we found that Speedtest and Bredbandskollen handled these rates well. The evaluation also showed that server location is important. MySpeed server located in Ålesund failed to measure bandwidths above 20 Mbit/s. The evaluation of Speedometeret

³Both Shockwave Flash v9.0 and v10.0 r22 (Mozilla Version) had the same problems

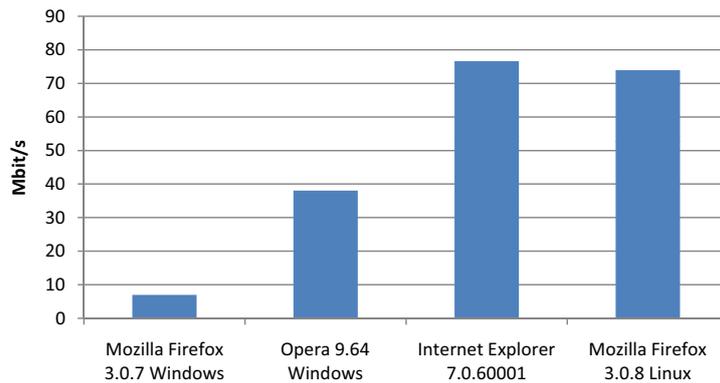


Figure 6: Obtained upload rate with Flash in various browsers.

JavaScript showed that JavaScript is not a good choice for measuring bandwidth. We also discovered upload bandwidth issues with the Windows version of the Flash plug-in for Mozilla for high-speed upload rates. The NPT plans to use the same test engine for NBTT as used by Bredbandskollen and Speedtest, and we can thus expect that the obtained results for this service will be close to the results obtained by these tools. Overall we think most of the available tools perform well for common bandwidth rates today. For higher rates, especially upload rates, there is still some room for improvements.

The planned NBTT will be based on the measurement engine used by both Bredbandskollen and Speedtest. Because of this we can expect that NBTT will perform equivalent to these services. Because NBTT is planned to be located at optimal locations in the Norwegian Internet infrastructure we should also assume that the results will be slightly more accurate and the service will provide less variations in the obtained results. Since NBTT will be developed in Flash, we should also be aware that NBTT will have the same high speed upload bandwidth issues demonstrated for Bredbandskollen shown in figure 6.

The NBTT will start with only one test server but should be extended to multiple servers in order to account for differences due to ISP domains and geographically distributed users. Multiple servers are also required to test network neutrality. It is also recommended that the tool should be able to acquire details about price and bandwidth guarantees directly from the host ISP when a user conduct the bandwidth test. Furthermore, it is found to be useful to define user profiles that describe the services and its performance attributes to be able to specify and customize test pattern in the future development of NBTT.

References

- [1] Patrik Arlos and Markus Fiedler. Accuracy Evaluation of Ping and J-OWAMP. *Swedish National Computer Networking Workshop*, 2006.
- [2] Martin A. Brown. Traffic Control HOWTO, Version 1.0.2. Internet, Oct 2006. http://tldp.org/HOWTO/html_single/Traffic-Control-HOWTO/ Last visited: 23.03.2009.
- [3] Bandwidth Controller. Traffic Shaper XP. Internet, 2007. <http://bandwidthcontroller.com/> Last visited: 24.03.2009.
- [4] Marius Aamodt Eriksen. Trickle. Internet, 2007. monkey.org/~marius/trickle/ Last visited: 24.03.2009.
- [5] Sally Floyd and Van Jacobson. Random Early Detection Gateways for Congestion Avoidance. *IEEE/ACM TRANSACTIONS ON NETWORKING*. VOL I . NO 1., Aug 1993.
- [6] Wireshark Foundation. Wireshark. Internet, 2009. <http://www.wireshark.org/> Last visited: 12.05.2009.
- [7] Lars-Petter Gunhildsberg Hansen and Ivar Conradi Østhus. NPT Online Broadband Test Tool. Master's thesis, Norwegian University of Science and Technology, Trondheim, Norway, June 2009. <http://daim.idi.ntnu.no/show.php?id=4619&type=masteroppgave>.
- [8] Bert Hubert. The Wonder Shaper. Internet, 2002. lartc.org/wondershaper/ Last visited: 24.03.2009.
- [9] Iperf. Internet. <http://iperf.sourceforge.net/> Last visited: 22.05.2009.
- [10] Manish Jain and Constantinos Dovrolis. End-to-End Available Bandwidth: Measurement Methodology, Dynamics, and Relation with TCP Throughput. *SIGCOMM Comput. Commun. Rev.*, 32(4):295–308, 2002.
- [11] Artur Janc. Network Performance Evaluation within the Web Browser Sandbox. Master's thesis, WORCESTER POLYTECHNIC INSTITUTE, Worcester, MA, USA, January 2009. www.wpi.edu/Pubs/ETD/Available/etd-011909-150148/unrestricted/artur-janc-msc-thesis.pdf.
- [12] Rick Jones. Netperf Manual. Internet, 1995. <http://www.netperf.org/svn/netperf2/tags/netperf-2.4.3/doc/netperf.html> Last visited: 22.05.2009.
- [13] Tønnes Brekne, Marius Clemetsen, Poul Heegaard, Tone Ingvaldsen, and Brynjar Viken. State of the Art in Performance Monitoring and Measurements. Technical Report R & D R15/2002, Telenor, 2002. www.telenor.com/rd/pub/rep02/R_15.pdf.
- [14] Ping. wikipedia, 2009. <http://en.wikipedia.org/wiki/Ping> Last visited: 26.06.2009.
- [15] Statistisk sentralbyrå. Andel med ulike typer internett-abonnement, etter husholdningstype, husholdningsinntekt, kjønn, alder, utdanning og arbeidssituasjon. 2. kvartal 2008. prosent. Internet, 2009. <http://www.ssb.no/emner/10/03/ikthus/tab-2008-09-18-02.html> Last visited: 25.06.2009.
- [16] Locktime Software. NetLimiter - Ultimate Bandwidth Shaper. Internet, 2008. <http://netlimiter.com/> Last visited: 12.05.2009.
- [17] traceroute. wikipedia, 2009. <http://en.wikipedia.org/wiki/Traceroute> Last visited: 26.06.2009.