2011 .se Health Status

Network neutrality

The influence by Service Providers on Internet traffic

Table of Contents

1 Introduction

In 2011, .SE carried out surveys within the framework of a project titled "IP Surveys," to determine whether Internet suppliers in Sweden influence traffic by various means. The report compiled backgrounds, analyses and conclusions from these surveys.

The full report is available in Swedish; this is an abridged edition.

1.1 Purpose of the report

The purpose of this report and the underlying tests is twofold: partly to investigate whether it is possible to detect the influence on Internet traffic and partly to provide an impression of whether, and if so, how, Internet traffic in Sweden is being influenced in ways that may not be entirely known or accepted by the general public. The impression is not comprehensive - it is only of sufficient scope to determine whether there is a need for more exhaustive studies and analyses.

Furthermore, the background that the report endeavors to provide on the topic is not entirely scientific. It is relatively comprehensive and aims to provide those who do not work in the trade with some insight into existing history, the opportunities and challenges that Internet suppliers face with a fast-changing Internet and how they may have been manifested in the management of online traffic and various types of business models.

Finally, the report is intended to be a means for advancing the debate on the topic. This is not the most accessible subject and to a certain extent, the industry and experts themselves are to blame for this due to their lack of clear and straightforward terminology. One of the most difficult aspects for outsiders to comprehend is when the same term is used in different contexts or when different words are used for the same thing.

1.2 Delimitation

The report does not aim to describe openness or neutrality in physical infrastructures that are not IP-based, such as optic-fiber leasing, channelization, frequency ranges or the like. This may be more of an issue concerning the regulation of infrastructure.

Nor does the report claim to deal with descriptions of brands or trademarks. There have been on-and-off discussions about whether regulation is required, for example, what may be referred to as "Internet" and "Internet access." However, this issue, regardless of its urgency, is not addressed in this report.

Nor does the report shed light on "openness" - we will not even attempt to describe people's perceptions of what this means.

2 Summary

We believe that our surveys confirm that Swedish Internet Service Providers, for both fixed and mobile connections, take measures to influence traffic. However, it is not always easy to determine exactly what those measures are. It is somewhat akin to trying to understand what goes on inside a watch without opening the watch casing.

Certain aspects can be determined while others will probably remain unknown, as long as the Swedish Post and Telecom Authority (PTS) do not force Internet Service Providers to disclose details of their operations by means of the supplementary tools provided to them by the regulations on disclosure requirements under the Electronic Communications Act.

The aim of the project has NOT been to create a perfect platform, but rather to take a closer look at what is really going on through limited spot checks. Now that this has been undertaken, we can draw a number of conclusions about what we could do as a next step. What these next steps may entail is reported in section 10, but we consider automated surveys that include more Internet Service Providers to be among the opportunities for improvement.

3 About “network neutrality”

For many years, the Internet was neutral when it came to the various applications that utilize it. TCP was the basic protocol of the Internet. This protocol creates connections that may be likened to a telephone call between several computers on a network. TCP has built-in mechanisms for discovering if a connection is overloaded, and if so, limit how it tries to transmit data to increase fairness.

In the mid 1990s, the http protocol for world-wide web traffic took over domination of the Internet from the classic protocols: telnet (terminal communication), smtp (e-mail) and ftp (file transfer). The http protocol soon represented an absolute majority in all Internet traffic.

Fast-forwarding through history, various file-sharing protocols then went on to replace http on the throne as largest user of Internet capacity.

When Youtube launched its operations in 2005, the status quo was once again displaced. If everyone was not already involved in file sharing, it seemed this was at least about to become the major service of the broader masses. By then the “Internet bubble” had burst and Internet Service Providers (ISPs) were no longer interested in spending vast sums of money to expand their networks. Instead, the focus turned to how YouTube actually generated earnings – on networks for which ISPs were bearing the cost, networks which ISPs often established with a business plan promising profit “sometime in the future,” when they become the biggest provider of all. That was when they considered charging for supplementary services that were intended for sale to their own customer base.

However, thanks to YouTube and other similar phenomena we were apparently about to be deprived of our privileges!

In a notable move in autumn 2005, SBC Communications Inc (AT&T) announced that they intended to charge Youtube, Yahoo, MSN and others for providing services that were dependent on SBC’s access connections and Internet backbone. However, it seems that this idea was never feasible.

The debate may not have begun there, but that was probably when the general public first became aware of the existence of such concepts. The debate has since been loudest in the US, but the matter has also been discussed in the EU within the framework of what is known as the data communications directive.

Since the possibility of obtaining payments from service suppliers seems to be farfetched, ISPs have instead looked into implementing, and to a certain extent chosen to implement, other solutions for resolving the growing need for capacity, often erroneously referred to as “bandwidth,” among users.

4 Why do Internet Service Providers want to limit traffic?

There are several reasons for why an ISP (this refers to any party who provides Internet connections regardless of what it is called: broadband, Internet, mobile broadband, surf, etc.) would seek to limit certain types of traffic:

- Certain types of traffic are costly for an ISP – for example, the capacity to major providers of Internet connections for shared traffic that provides access to the Internet as a whole, or proprietary transatlantic connections.
- Certain online resources are limited, for example, radio spectrums for 3G/4G services.
- The ISP wants to create worse conditions for a nonproprietary competing service and by so doing, increase its profits. This may also involve preventing a customer from using a free Internet-based service, for example, IP telephony and forcing them to use services that are subject to a charge by the ISP.
- The ISP wants to protect its network from DOS attacks which may be both intentional and unintentional.
- The ISP wants to prevent customers who use relatively small capacities from being run over by heavy consumption customers.

In other words, there are many different underlying motives for limiting traffic. Both are financially profitable: to enable the network to deliver a service that is equal for everyone and to protect it from harm.

The limitations may be implemented on various parts of the network and thus affect different numbers of customers in different ways:

- Limit the amount of traffic per customer, for example, by setting limits on the customer's access connection.
- Limit all customers, for example by setting the limit for particular types of traffic on an interconnect or connection to an Internet exchange point.
- Limit customers in a certain location, for example, a certain sector in a radio tower.

Where the limitation is placed may in turn lead to various types of influences in terms of the number of customers, but modern equipment that is available for example, on a transit connection may distinguish and solely limit a singular customer's data traffic.

5 What is traffic management?

Traffic management is a general term for how network traffic is managed, based on information other than purely destination-based routing, which is the default in IP-based networks.

“Other information” may include the degree of network load. If a link is congested, it may be reasonable to send certain traffic via alternative, if not equally good, routes to prevent data loss.

A relatively new variety of traffic management is to limit how much of a particular type of data an individual customer may send. In this regard, it is most common for providers to limit the most substantial types of traffic, such as, file sharing or video. This can be accomplished by several means:

By limiting only the traffic flow or protocols that generate considerable traffic (for example, BitTorrent); other protocols are not affected, but the total volume of traffic is reduced. In certain cases, the limitation is implemented in full, meaning that the traffic is blocked in its entirety.

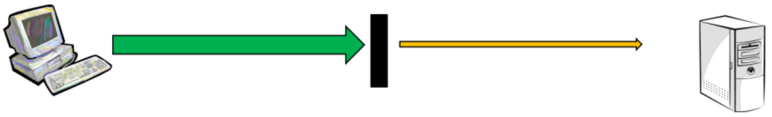

Limit!

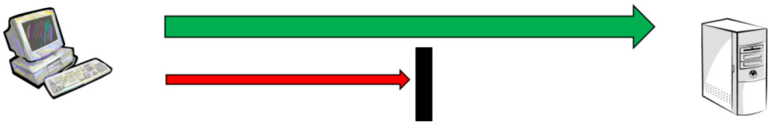

Block!

By prioritizing traffic to ensure that, for example, IP telephony or IP-TV always has sufficient space.

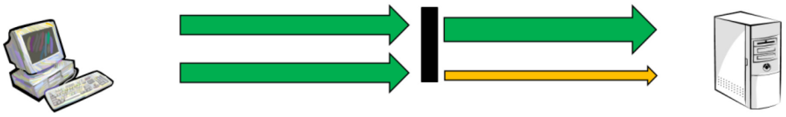

Prioritize!

By limiting all traffic for a particular customer for an extended period if the customer exceeds a certain volume of data. (Subscription with limited transfer, for example, per month):

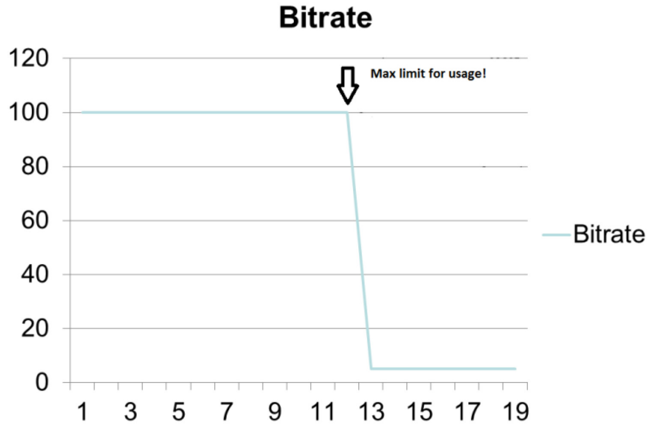

Maximum data volume exceeded

The limitation of a particular traffic flow/protocol may be implemented by more or less sophisticated means:

- Softer "shaping" by means of delays or similar measures.
- Harder "rate limitation," which involves discarding packets that exceed the maximum limit.
- By limiting the number of new connections that may be created or send interfering packets to make the sender or receiver believe that it is time to stop sending data and to terminate the connection.

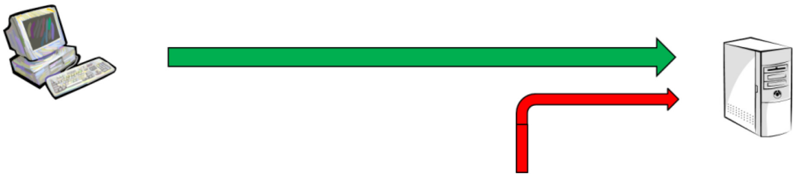

Interfere!

6 Measuring whether traffic is influenced

Measuring traffic on the Internet is no simple feat. It is neither economical nor practical to build a network that enables every end customer to use their full capacity to connect to all destinations simultaneously. Accordingly, a network is built based on statistics and financial factors. Statistics indicate the amount of load in various sections of the network over time and the differences depend on the day of the week and time of day. Financial factors determine whether to increase the number connections in anticipation of periods when full capacity will be exceeded, or how much and how long a connection may be congested before it is expanded.

In addition, traffic flows between two locations on the Internet through one, two or more ISPs, and each has a different policy for network expansion.

In other words, there is no specific location on the Internet from which measurements can be made to determine how much data can be sent or received for a particular form of traffic, since the measurements will likely be influenced by all of the different sections involved.

Nor can the assumption be made that specific measurements could be repeated with the same results, since loads may change over time depending on the behavior of other users.

However, measurements can be made according to a model whereby two or more types of traffic are sent simultaneously and compared with each other. While this model certainly does not provide any clear answer to the total amount that can be transferred (which may vary over time), it may indicate whether there are differences between the different types of traffic. To increase the reliability of the measurements, they should be of sufficient length, repeated and performed during different times of the day and year. Measurements generating results that excessively deviate from each other should also be processed in a balanced manner.

In other words, a certain amount of caution must be observed when interpreting the results, particularly if the percentages of relative differences between various types of traffic are minor.

7 Measurements and results

In autumn 2010, .SE decided to implement a project that included the following objectives:

- To perform a number of manual measurements.
- To analyze results from the measurements.
- To discuss the results with ISPs, technical experts and other stakeholders.
- To prepare a report that presents the results.
- To produce a glossary for the topic.
- To serve as documentation for future decisions regarding whether it is pertinent to develop a service that performs automated measurements.

The implemented measurements were performed by means of Internet access through:

- a number of ADSL subscriptions,
- fiber-optic broadband,
- cable television broadband, and
- a number of mobile broadband subscriptions.

The server and method of measurement chosen was Glasnost from the Max Planck Institute for Software Systems. Software was installed on two servers located with different ISPs to enable comparisons to be made with the various servers from the same measuring control points. The tests performed by Glasnost are in line with the criteria indicated in the section "How can you detect whether ..."

Since the measurements were performed manually, it has not been possible to perform them on a major scale (recurring and at various times of the day), instead, these tests were more of a random-sampling. However, these are entirely in line with the objectives of the project.

8 Measurement results

The following is a sample compilation of excerpted results from the measurements that were performed as of July 1, 2011 up to and including August 19, 2011.

8.1 Compilation of measurement results per provider and measuring control point

Measurements were performed on two servers at different locations in the network. The results per ISP are: measuring control point 1 and downloading, measuring control point 1 and uploading, measuring control point 2 and downloading and measuring control point 2 and uploading.

Each measurement is represented by a pair of bars (blue-red) and the protocols and ports that were used are stated below the bars – the first protocol is the blue bar and the second is the red. The measurements were performed in the order left to right in each diagram with a brief pause (approximately 30 seconds) between each measurement and new protocol.

We have noted that something seems to be interfering with the first measurement for flash videos since it essentially always ends up lower than the measurements that follow. We will examine whether this is due to a fault in the survey software, the measurement procedures used or some other reason.

8.2 ISP C

8.2.1 Measuring control point 1

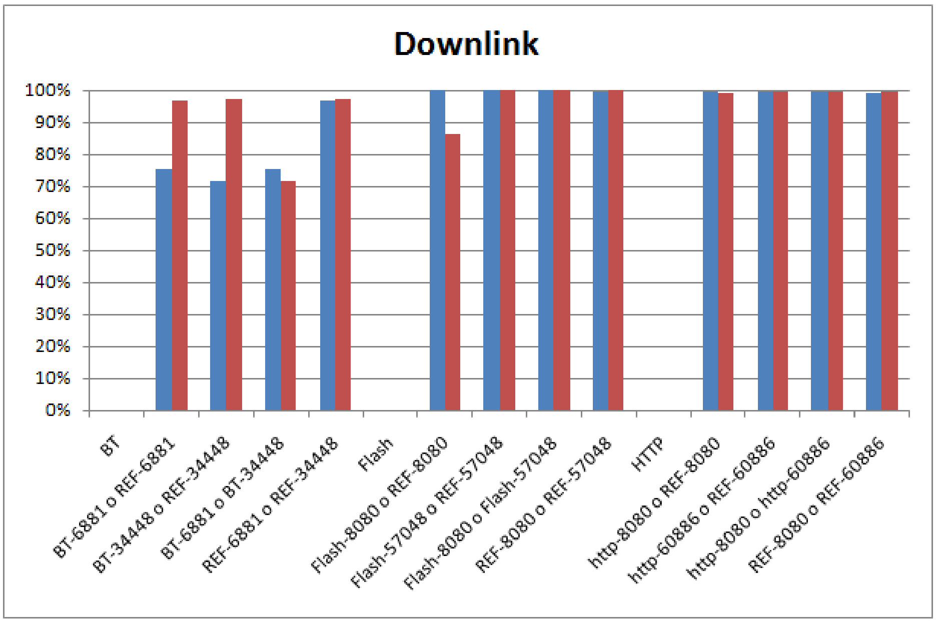

8.2.2 Measuring control point 2

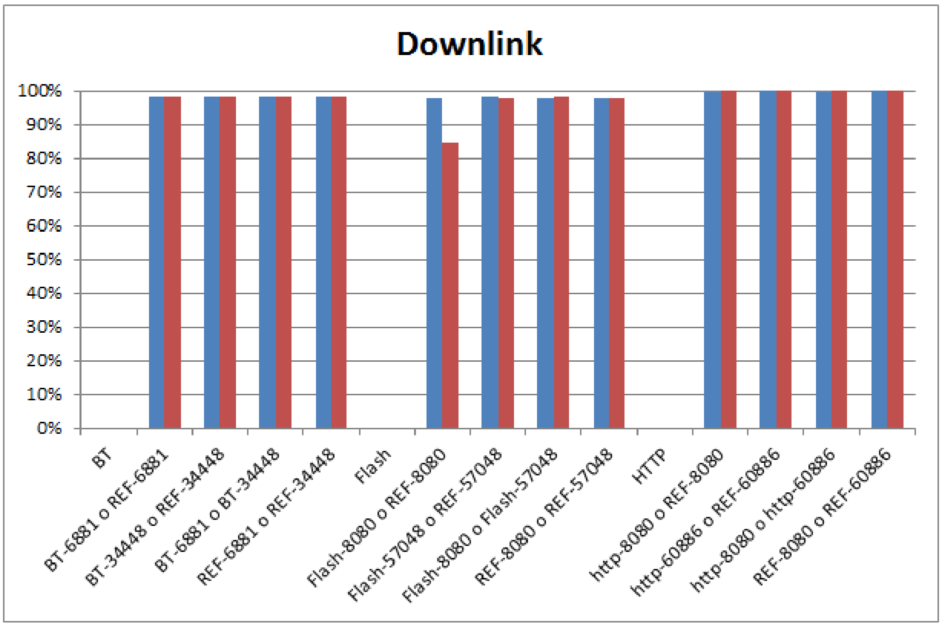

8.2.3 Comments on the results from ISP C

The BitTorrent traffic seems to specifically be 30 percent lower than the reference protocols for measuring control point 1. The results are more balanced for measuring control point 2. The traffic between ISP C and measuring control point 1 passed through a third provider (transit), while the traffic to measuring control point 2 took a direct route. One assumption is that ISP C is limiting

BitTorrent traffic since traffic that passes through a transit provider has to be paid for.

8.3 ISP F

8.3.1 ISP F measuring control point 1

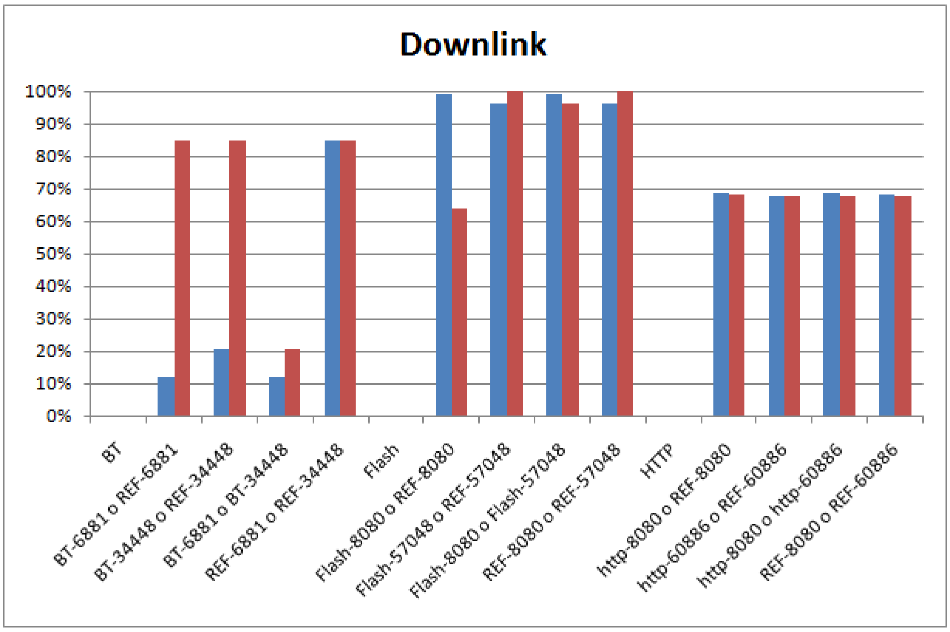

8.3.2 ISP F measuring control point 2

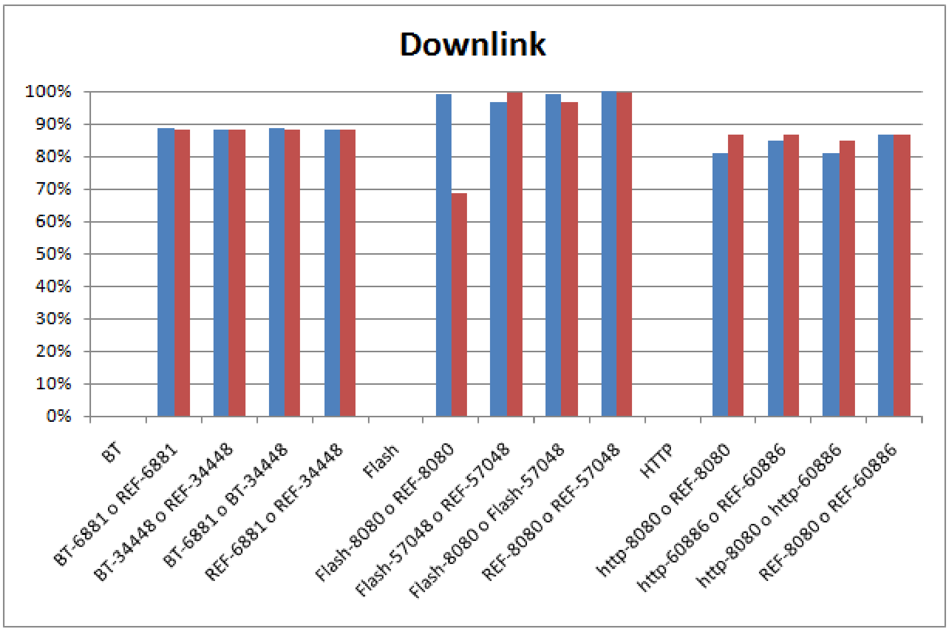

8.3.3 Comments on the results from ISP F

The BitTorrent traffic with measuring control point 1 is clearly lower than the maximum limit – and in comparison with the reference protocol. We are also noting that this is not due to the port but to the protocol. Thus, equipment is being utilized somewhere along the route to look inside the packets to identify the protocol used, through what is known as deep packet inspection.

8.4 3G provider B

8.4.1 3G provider B measuring control point 1

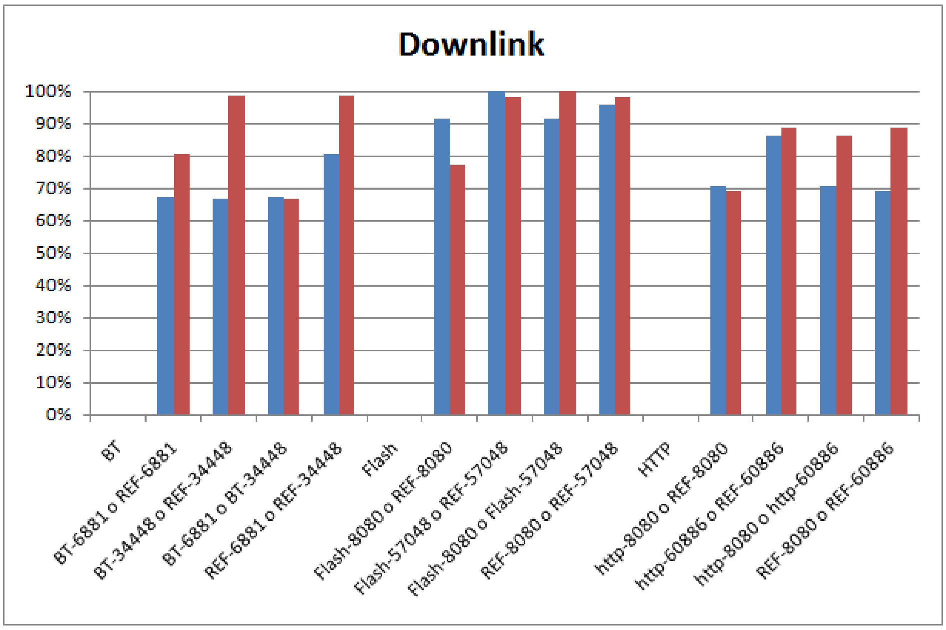

8.4.2 Comments on the results from 3G provider B

As far as we can determine, there seems to be a specific limitation of BitTorrent and port 8080, regardless of the protocols used.

8.5 3G provider C

8.5.1 Measuring control point 1

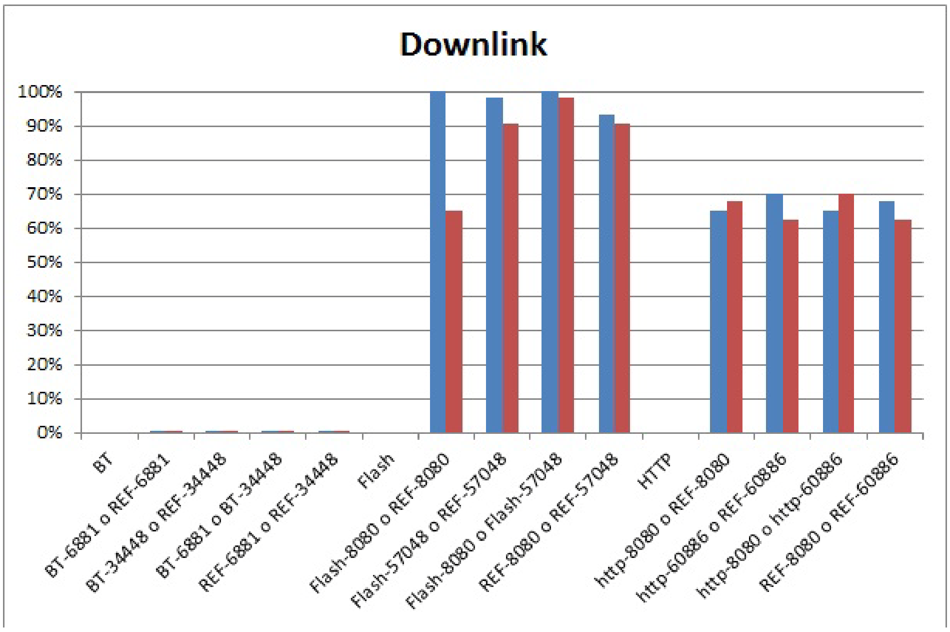

8.5.2 Comments on the results from 3G provider C

BitTorrent seems to be considerably limited for both measuring control points.

8.6 3G provider D

8.6.1 Measuring control point 1

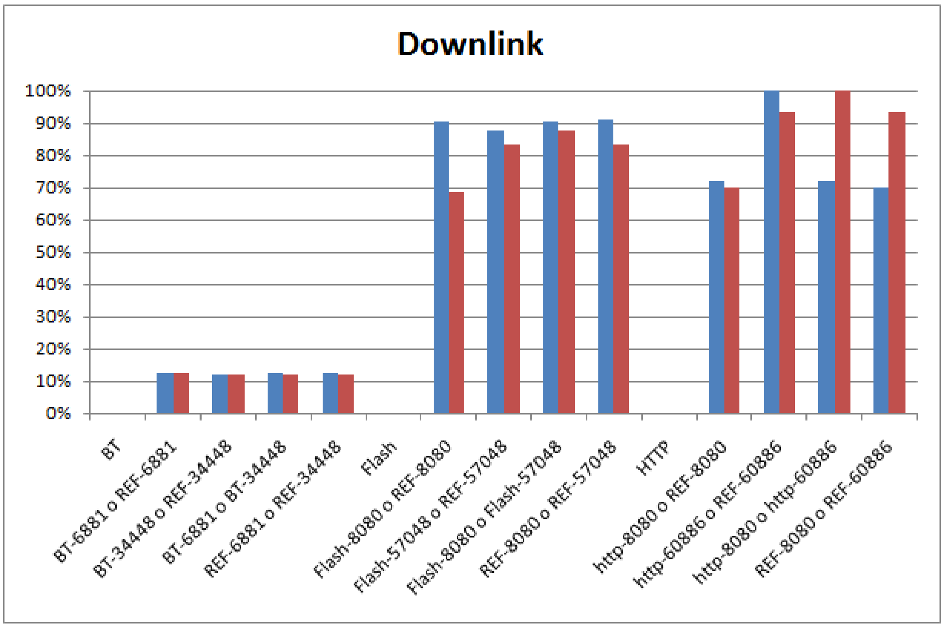

8.6.2 Comments on the results from 3G provider D

BitTorrent seems to be choked to approximately 10 percent of the maximum for measuring control point 1.