

[Leslie Haddon](#) and Peter Heinzmann

Implications of the variation in broadband speeds over time

Book section

Original citation:

Originally published in Haddon, Leslie and Heinzmann, P. (2011) *Implications of the variation in broadband speeds over time*. In: Haddon, Leslie, (ed.) *The Contemporary Internet: National and Cross-National European Studies. Participation in Broadband Society* (Vol. 3). Peter Lang, Frankfurt am Main, Germany, pp. 95-119. ISBN 9783631600986

© 2011 [Peter Lang](#)

This version available at: <http://eprints.lse.ac.uk/64817/>

Available in LSE Research Online: November 2018

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

This document is the author's submitted version of the book section. There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

The Experience of Broadband Speeds

Leslie Haddon and Peter Heinzmann

In Haddon, L. (Ed.) (2011) *The Contemporary Internet: National and Cross-National European Studies*, Peter Lang, Frankfurt, pp.95-119

Introduction

Many factors affect people's interest in acquiring and using new information and communication technologies (ICTs) and specifically new applications. Such factors include, for example, economic costs, network effects in the case of communications (i.e. how many other people use an ICT), the symbolic meanings associated with ICTs whereby it becomes fashionable to have some technologies, etc. This report, though, aims to reflect more on the features of technology itself, specifically its capability to support certain functionalities and, indeed, to do so in a way that makes them if not attractive, then at least acceptable. In particular it looks, historically, at various changes related to the experience of the speed of broadband¹ and how this has influenced users' experience of different applications and hence their different uses of the internet.

Some writers examining the topic of the digital divide have made the point that one should not just look at whether different people are internet users or not but also consider how the online experience is affected by factors such as their technical skills, their awareness of online possibilities, the support of their social networks – and the very technology through which they experience the online world (Bakardjieva, 2001; Haddon, 2004). It makes a difference, for example, if one is using older or newer equipment, with differences in processing power or one subscribes to ISPs with faster or slower connection speeds. Looking at the technologies and connections people use gives a more nuanced picture of whether some have more advantages over others, and how the affordances of

1 Since technologies are always changing, what counts as “broadband” also continues to evolve - indeed, most reports using the word “broadband do not provide a precise definition of the term. Today, the term broadband typically includes widely used internet connections 500 times faster than early internet dial-up, i.e. “narrowband”, technologies 10 years ago. However, the term broadband does not refer to either a certain speed or a specific service. In this chapter we understand by “broadband connections” internet connections with a download speed higher than 128kbit/s, with flat-rate (non-volume based) tariff charges for fixed landlines and volume based charges for mobile phones.

people's technologies affect the quality of experience they can have of the online world.

Applying this approach beyond the digital divide debates, one can simply say that certain usages that have now emerged would have been impossible, or less practical, or certainly less attractive, when connection speeds were slower. However, this is not just about looking at the changeover from narrowband² to broadband. From their introduction, although the offerings of different ISPs may have all been called broadband, the actual speeds of connections varied.

Let us consider some examples. Even with narrowband one could send images, but slower upload speeds meant that this involved spending some time sending those images to others or posting those images online (which could be both tedious as well blocking telephone lines for voice calls in many homes). Early broadband at around 600 kbit/s download and 100 kbit/s upload speed certainly made sending and receiving single images much quicker, but more speed was needed to make the practice of sending (and receiving) many images viable. For example, in interviews in Korea in 2005 some students were posting many photos of friends that they had taken with their camera phones on, say, a group excursion, where a 3 day trip might have produced something in the order of 200 pictures (Haddon and Kim, 2007).

Or take websites. While they have been around for many years, earlier versions contained mainly and almost exclusively text, compared to the websites of today, which have many different objects and sometimes extensive audio-visual material. At what point in time (and for whom, if we take a cross-national perspective) did the changing speed of broadband support practices such as posting many images and managing complex websites?

Meanwhile, one study of Italian youth discussed how they commonly share online music, TV and film material (Mascheroni, et al., 2008). Again, this has become more feasible as speeds have evolved. But what was the tipping point to allow this?

In calculating how people have experienced connection speeds and hence download and upload times, we note that the headline speeds - i.e. the top-speeds claimed by the ISPs who supplied those connections - provide just a starting point. Others have also demonstrated (e.g. Ofcom 2008, 2009) that the reality of speeds experienced by users for various reasons does not always match the potential. Hence, empirically, this report examines historical data about the actual average speeds experienced at different time points, and the pattern by

2 In this chapter "Narrowband" is understood to mean the dial-up connection to the internet with speeds of less than 128kbit/s, being charged on a time or volume basis.

which those speeds evolved, since this had a bearing upon what applications (and the hence the practices of users) could be supported or could become acceptable.

Lastly, averages conceal the variation in people's experiences. For example, the length of the connection line and congestion at certain peak times both influence speeds. Hence we have to consider what effect these have on transmission times in order to appreciate the range of experiences within countries at different points in time. Finally, while we can offer examples of single countries, transmission rates varied significantly in different European countries. So how much of a gap is there between the speed experienced by people in different countries and how has this gap changed over time? Has it become larger or smaller?

Thus the rest of the report proceeds through the following stages, providing:

- a) a first impression of the time that would be taken to upload and download files given different theoretical connection speeds
- b) an explanation of the methodology behind the next stages of the report.
- c) an illustration of how file sizes have grown over time for different applications and changes in streaming speeds for different applications.
- d) an historical overview of how long it took to transmit typical files in different years.
- e) a view of the difference between headline speeds (over the years) and the actual speeds people experienced, with implications for transmission times.
- f) a discussion of variations within countries and why these occur, with what implications.
- g) an overview of the differences in connections between countries, and how those differences have changed over time.

Methodological issues

Streaming and 'bursty' traffic

There are two basic types of information transfer in telecommunication and on the internet leading to streaming and 'bursty'³ traffic.

In the case of *streaming traffic* the receiver already consumes the first parts of information while other parts are still being sent by the sender. Streaming traffic involves the regular delivery of the pieces of information. For example the digitised audio signals that are common in conventional telephony require the regular delivery of one byte (8 bits) of information every 125 microseconds.

³ *Bursty traffic* refers to an intermittent pattern of data transmission.

This leads to a constant average data rate of 64 kbit/s (for the entire duration of the transmission). Another example is streaming audio or video information on the internet, as for example when people are listening to internet radio programs or when they are watching YouTube videos leading to an average data rate of 128 kbit/s and 900 kbit/s respectively. Here, compression is used and therefore the instant data rate may vary. But again, we have a rather constant data rate during the whole transmission.

Bursty traffic is irregular; it comes in bursts. Sometimes there is a need for communication at typically high data rate, sometimes not. Hence the average data rate during a so called ‘session’ - that is, an extended period of communication between two parties - is much lower than the peak data rate. In such bursty traffic situations the communication channel is typically shared by many users. While one user is not requesting information the channel is free for other users. This type of traffic is typical in computer communication networks as for example when people access web servers, in computer games, or when people do access remote computers.

Response time and data rate

In the case of bursty traffic the ‘response time’ a user experiences when requesting a piece of information is the most important technical performance indicator. The response time requirements depend on the ‘use case’ i.e. the type of application.

In the case of streaming traffic, however, it is the available data rate of the communication channel that is the most important technical performance indicator. If the streaming information belongs to a one way communication channel such as a radio broadcasting signal, then the response time - that is, the time it takes from the mouse click requesting the programme to the moment in time when the sound can be heard at the receiver side - is not very demanding (e.g. it can be up to a value measured in seconds).

But with streaming (e.g. audio signals belonging to a telephone call) the so called ‘round trip time’ must be less than 350 milliseconds in order to have a fluent conversation. When it is longer, the two conversation partners feel uncomfortable because of the delayed reactions of the other side. We are familiar with such longer round trip times from intercontinental telephone calls that are routed via satellite, where the time from sending out a phrase until the reaction from the conversation partner arrives is typically around half a second.

Depending on the use case, i.e. the application, there are three classes of bursty traffic to be distinguished: batch processing, human interactive and machine interactive.

In the case of ‘batch processing’, response times of tens of seconds to hours may be perfectly acceptable. When requesting the transmission of a ‘batch’ of

information, for example sending an e-mail or making a backup of files there is usually no need for instant response. However, in the so called ‘human interactive’ applications, people want an average response time of clearly less than one second. Moreover, this response time should be fairly stable - it should have a small variance. Human interactive applications include, for example, web browsing, database queries or remote access to computers. Finally, in the ‘machine interactive’ application we have computers interacting with each other.

The next step involves looking at what shapes response times. One factor that determines the response time is the available transmission speed (the data rate) between the source of the information (e.g. a Web server) and the destination or ‘sink’ for the information (e.g. a user’s client program). The other important factor is the volume of the piece of information requested.

In order to convey a first impression of what data volume and data rates can mean to our use of the internet, Table 1 shows the minimum response time when requesting typical ‘application objects’ (i.e. items we might want to send or receive online) on networks with different data rates. The application object data volume refers to data volume sizes used in 2009.

Object Type	Data Volume	Data Rate				
		33.1 kbit/s	600 kbit/s	2 Mbit/s	5 Mbit/s	20 Mbit/s
Web Page	500 kByte (4,000 kbit)	2 min	7 sec	2 sec	1 sec	0.2 sec
4 minute Music Track	4.6 MByte (36,800 kbit)	19 min	63 sec	19 sec	8 sec	2 sec
2 minutes Video Clip (352x288px)	20 MByte (158 Mbit)	1.3 h	4 min	1 min	32 sec	8 sec
143 minutes DivX Quality Movie (592x312 px)	940 MByte (7,876 Mbit)	3 days	3.6 h	66 min	26 min	7 min
143 minutes DVD Quality Movie (720x384 px)	2 Gbyte (15,938 Mbit)	6 days	7.4 h	2.2 h	1 h	13 min
143 minutes HD 720p Quality Movie (1280x688 px)	8 Gbyte (66,605 Mbit)	23 days	31 h	9 h	4 h	56 min
Blu-ray Disc (2 hours HD 1080p Full HD Quality Movie)	50 GByte (33,554 Mbit)	147 days	8 days	2.4 days	1 day	6 h

Table 1: Theoretical time taken to transmit different application objects⁴.

This is a straightforward table, showing that if we have files of different sizes on the left (presented in an approximate and more exact way⁵) and different data

4 Similar tables are also available in other reports, as in the one from the British regulator Ofcom 2008, p.14 (Ofcom, 2008).

rates (that is, connection speeds) on the top, then this is how long it will take to transmit the files - either to download or upload them. It is clear that for some file sizes, as in the 250kB webpage, one reaches the 'human interactive' stage very quickly, after 600kbit per second (kbit/s), where the user would probably find the speed acceptable. In contrast, for large file sizes, such as downloading DVD Quality Movies, the effect of faster speeds on download times at each stage in the evolution of broadband is more noticeable. In fact, at the very slow response times most people would not have bothered to download movies, video clips and probably even music. In those cases we are clearly in the 'batch processing' mode of traffic, or it might be even faster to physically transport the discs.

But these are all just typical 'example' values, even if they provide a rough guide. If one asks for the actual time it takes, for example, to get an actual video clip one must consider many more details. The size of the video clip file may vary considerably not just depending on the type of file but also on the year. The volume of a single web page is just one issue. But in 2009 most web pages consist of 10s to 100s of objects which may be fetched from multiple servers and therefore additional delay when accessing other servers may be the factor which mainly determines the overall response time. And we will demonstrate that all file types tend to get larger over the years. Finally, taking into account the actual speeds we will come up with figures and explanations which describe the users' actual experiences more realistically than the advertised speeds.

Next it is important to consider file sizes: these vary depending on the application and the file sizes for particular applications themselves change over time. In order to get some idea of current average file sizes as well as about historical evolution of file sizes, information about file sizes was assembled from a number of sources. Agrawal, et al (2007) have analysed metadata from over 10,000 file systems on Windows desktop computers at Microsoft Corporation over the time period 2000 to 2004. In addition, these data are supplemented by an analysis of the files on the file store of a small IT company with ten employees for the time period 1997 to 2009.

Lastly, one project within the COST298 action had the task of measuring broadband performance. This involved installing software on "reference servers" in different countries so that internet users could select a web page on their country's reference server and check their connection speeds. The resulting information was also kept centrally at the cnlab in Switzerland, so that over time

5 Note that sometimes a 'Byte' is abbreviated with a capital 'B' and a 'Bit' is abbreviated with a small "b". One Byte consists of 8 Bit.

a considerable number of measurement results were assembled. For the case of Switzerland it is estimated that about 10% of the internet users have performed such measurements.

This self-selecting group of people who check their speeds is by no means a random sample of users. For example, people who check their speeds probably tend to be more enthusiastic about the technical aspects of the internet and/or are more advanced users. They may well be more likely to pay more for faster internet speeds. Alternatively, people who experience speed problems may run more tests than people who are satisfied with their connection speed.

The point of the performance test platform was to compare the connections of different ISPs and to analyse regional differences in certain countries. The objective of this report is to see how speeds change over time, so it does not matter so much if the people running tests are a little ahead of the rest of the population – it just gives a broad indication of the wider trend and highlights some issues that others would also experience.

In Switzerland the facility is used by about 2000 users per day carrying out around 10,000 tests per day. This is due to the media visibility of the project, which has not happened yet in other countries. Since there are more data about users in Switzerland and these have been collected over a longer period of time, the Swiss data have been used to demonstrate many of the processes at work, especially for the earlier years of broadband.

Where a comparative data is introduced this is also from the performance measures outlined above. Finally, in the discussion of variation in the experience of speeds within countries, UK data from the British regulator is used to supplement the main observations. We hope to further promote the performance test in various countries and hence to get much more users testing their connections in other countries than just Switzerland.

The time taken to transmit different file sizes

In the first section we take the first step of looking, empirically, at the current and historically changing file sizes and available data rates examining the implications for the time taken to download these files. At the moment we consider only “headline speeds” – i.e. what is the time taken if we assume the maximum speed of connection cited for particular years. In the next section we will then go on to look beyond this to consider factors affecting the actual speed users experience, and the implications for download times.

In Table 2 the current (2009) average file sizes for different activities and applications have been assembled from various sources. Table 2 shows also the time needed for the transmission of the different file types using the current widely used 5Mbit/s internet connection.

Object, Application	average file size 2009 [MB]	transmission duration at 5Mbit/s speed
Full disk backup	80,000	2,185 min
Video Clip (.avi format)	942	26 min
Windows Vista. Service Pack 2 (SP2), May 2009	348	10 min
Video Clip (.mov, .mp4, mpg format)	40	66 sec
Windows Media Player (wmp11)	25	40 sec
Video Clip (.wmv)	23	38 sec
Music Files (.mp3 format)	5.6	9 sec
File Hosting Rapidshare / Megaupload normal	5.0	8 sec
Fotos Camera	2.0	3 sec
Fotos Mobiles	0.8	1 sec
Web pages	0.5	1 sec

Table 2: File sizes for different objects / applications⁶

In Table 2 we can see that backing up an entire hard disk takes hours at 2009 speeds, whereas files that update software take minutes. Other transmissions need only seconds. Sometimes we have had to add several different possibilities, for example, in the case of video. This is because the resolution, the frame rate and the duration of videos clearly determines file size, and here we provide our best judgement as regards average file size. But the format also affects size; the file from a video clip of a given duration in the (uncompressed) .avi format is much larger than in the compressed .mpeg or .avi format. Music files are, understandably, smaller than video ones, and photos take up even less space, the file size between digital cameras and mobile phone photos being typically different because of the different resolution. File hosting refers to the files or pieces of files one can store on the internet instead of on one's PC, which is becoming an increasing popular way of sharing audio-visual material as opposed to exchanging peer-to-peer. Ultimately, Table 2 gives us a first idea of the time it takes to, usually download, different types of file, if we have the full speed of 5Mbits/s that the ISP claim to provide.

⁶ Statistics were taken in 2009 from files on the file servers of a small company.

The changing size of files and response times

While Table 2 showed the current average file sizes, these sizes have changed over the years – mostly they have increased. For example, the upper line of Figure 1 shows how home pages grew in size from June 2006 to December 2007 (17 months). This represents a growth of 34% between those dates, which means 24% a year – a growth rate of just under a quarter. This leads to larger files doubling in size in 3 years or growing tenfold over 10 years.




Figure 1: The changing size of home pages.

Taking an even longer term perspective, shows how *web video clips* distributed over the internet have on average become longer (in terms of duration) over a decade. While the median web video clip duration was 45 seconds in 1997, it had grown to 120 seconds in 2005 and 192.6 seconds in 2007 (Growth of the Average Top 1000 Home Page, 2008). Hence the video clip has grown fourfold in duration over 10 years. But at the same time, the image resolution and the frame rate have also increased, leading to an even higher video clip file size growth over the last 10 years.

Figure 2 shows how the average office file sizes - Word (designated by 'doc') Excel ('xls'), PowerPoint ('ppt') and the exe-file sizes - grew over the last

ten years⁷. For instance, PowerPoint files have grown about 20 times over the last 10 years, mainly due to the increasing use of images in PowerPoint charts⁸.




Figure 2: Average Office file size growth over time.

Modes of operation (user behaviour)

Obviously requesting a larger file will lead to larger response times than a smaller one. In the ‘user interactive’ mode people wait for the information they requested. When really working interactively e.g. remote computer access or surfing web pages users do need the information within sub-seconds. They may accept seconds for images or video clips appearing in web pages. But as soon as things take longer than tens of seconds, e.g. when downloading documents or applications, people tend to switch to other tasks rather than waiting for the information. Hence, in effect, they change to batch processing mode.

Note that when information is delivered in streaming mode (e.g. video and audio) users get the beginning of the information before the whole file has been

⁷ These data were taken from the file server of a small company with 10 employees.

⁸ Haddon and Heinzmann (2009) also chart the growth of Adobe Acrobat Reader, Windows Media Player and Skype, the voice over IP package. If we look at Windows service pack and Windows patch file sizes theoretically downloading a 150MB service pack over the network in 2002 would have taken almost 5 hours at the 64kbit/s speeds broadly used at that time. No wonder that the software was in practice distributed on CDRoms. Today, the 350MB to 450MB service packs can be downloaded within about 15 minutes thanks to the current 5Mbit/s speeds. Thus, use of the frequent online software updates and patching that it is common today would simply not have been practical at the download data rates available five years ago.

transferred; hence the response is faster than the time it would take to transmit the whole file.

To give an example of what this means in practice, at speeds of 5Mbit/sec a 2 minute video clip would take 32 seconds to download and a 143 minutes DivX quality movie would take 26 minutes to download. But playing the movie in streaming mode will show the first scenes within seconds. We will return to these examples as benchmarks in the later sections.

Symmetric and asymmetric communication

Traditionally most internet users are mainly consuming information, i.e. browsing web pages, downloading documents and programs, listening to audios and watching videos. Therefore ISP connection speeds are very asymmetric, having typically around 10 times smaller upload speeds than download speeds. For example, the upload speed for a line with 5Mbit/s is only 500kbit/s. Hence it would take about 3 seconds to download a single picture taken by a cameraphone but 32 seconds to upload it.

Around 2006 more and more internet users began producing and sharing information- writing information into forums or wikis, or uploading photos or videos to information sharing platforms like Flickr or YouTube. In this context the terms Web 2.0 or Read/Write Web and Prosumer⁹ were created. Until 2009 the so called 90-9-1 rule was almost certainly correct - i.e. the claim that in most online communities, 90% of users are lurkers who never contribute, 9% of users contribute a little, and 1% of users account for almost all the action. But the fame of information sharing platforms like YouTube, Flickr, Wikipedia or Twitter and the omnipresence of buzzwords like Blogging (e.g. Technorati), Social Bookmarking (e.g. Delicious), Social (e.g. Facebook, MySpace) or Business Networking (e.g. Xing) clearly point towards a growing amount of symmetric communication and hence more symmetric communication channel needs. Internet telephony (Skype) and video telephony applications provide further examples of such growing symmetric traffic.

Historical evolution of transmission speed and file size

Clearly in the past, when transmission data rates were slower, it would have taken longer to download the current versions of the various files listed above. On the other hand, we were simply not sending such large files in the past.

⁹ The buzzword 'Prosumer' has multiple conflicting meanings. Here it is meant as the combination of information producer and consumer (as introduced by Alvin Toffler in 1980. In the business sector the prosumer is the market segment of professional consumers.

Figure 1 and Figure 2 illustrate how, for different types of applications, the files were much smaller then. So what we can now ask is, given speeds at different time points but also the fact that there were different sized files at those time points, has the response time changed – are people able to download these items faster?

In Table 3: Historical data showing file size and response times for downloads.

on downloading we can see that the trend across most files was that file sizes grew larger, but download speeds increased by an even greater factor so that the actual times for downloading all these files was reduced. Even in 1997 web pages with file sizes typical at that time could be downloaded fairly quickly and so subsequent faster times would hardly have been noticeable to the user. Those same users would have noticed changes to downloading photos and PC files after 2001, but then these times all remain very quick, and so at this point the difference is hardly noticeable.

Year	Download Data Rate [kbit/s]	Web Pages		Office Files (Powerpoint)		Photos		Music Song		Software (e.g. Adobe Reader)	
		File Size	Response Time	File Size	Response Time	File Size	Response Time	File Size	Response Time	File Size	Response Time
1997	34 kbit/s	20 kB	5 sec	0.4 MB	2 min	0.8 MB	3 min	1.0 MB	4 min	4.5 MB	18 min
2001	64 kbit/s	60 kB	8 sec	0.5 MB	1 min	0.8 MB	2 min	2.0 MB	4 min	8.4 MB	18 min
2005	600 kbit/s	150 kB	2 sec	0.8 MB	11 sec	1.5 MB	20 sec	3.0 MB	41 sec	16.3 MB	4 min
2009	5000 kbit/s	500 kB	1 sec	1.6 MB	3 sec	3.6 MB	6 sec	7.0 MB	11 sec	26.6 MB	1 min

Table 3: Historical data showing file size and response times for downloads.

Table 4, dealing with uploading speeds, indicates that the change of upload transmission speeds was less than for downloading, but the overall picture remains the same – file size may have increased but upload times have decreased.

Year	upload Data Rate [kbit/s]	Office Files (Powerpoint)		Photos		Music Song	
		File Size	Response Time	File Size	Response Time	File Size	Response Time
1997	34 kbit/s	0.4 MB	2 min	0.8 MB	3 min	1.0 MB	4 min
2001	64 kbit/s	0.5 MB	1 min	0.8 MB	2 min	2.0 MB	4 min
2005	100 kbit/s	0.8 MB	1 min	1.5 MB	2 min	3.0 MB	4 min
2009	500 kbit/s	1.6 MB	26 sec	3.6 MB	1 min	7.0 MB	2 min

Table 4: Historical data about file size and response times for uploads.

Table 5 shows the typical streaming speeds for different applications.

Streaming data from an online source to one’s computer, as in watching a live video or listening to a live internet radio programme, is only possible if the network supports the required streaming speed. One may adapt the speed

requirements by selecting lower quality sound or videos, e.g. lower image resolution. However, if the available transmission data rate is not sufficient, the programme is not just delayed - it cannot be viewed at all. In that case the programme might be downloaded and viewed later. Hence the user has to switch to batch processing mode.

Date Introduced	Streaming Media Type	Typical (average) Data Rate [kbit/s]
	Audio	
1998	Music MP3	128
2002	Internet Radio	128
2005	Internet telephony (Skype)	128
2005	Audio Podcasting	64
	Video	
2003	Video Podcasting	495
2005	YouTube Videos, standard quality, 320x240	300
2008, March	YouTube videos, medium quality, 480x360	700
2007, June	YouTube videos, mobile quality, 176x144	300
2008, Nov	YouTube videos, HQ (720p)	3,500
2009, Nov	YouTube videos, HD (1080p), 1920x1180	8,000

Table 5: Streaming speeds for different streaming applications

We can take the case of YouTube to show changes over time. Table 5 shows that the videos that one could watch streamed from YouTube were always, in a sense, ‘viewable’ at the broadband connection speeds available at different time points. But the quality did vary. In fact, YouTube now offers high definition to take into account the fact that some viewers now have very high connection speeds, but users with lower speeds select videos at a lower quality.

The experience of broadband speeds

As noted earlier there is a difference between the advertised connection speeds claimed by ISPs (also known as the ‘headline speed’ or the ‘target speed’) and the real speeds that people get in practice (also known as ‘effective speed’). Over the years ISPs have regularly increased the connection speed, as shown in Figure 3 for Switzerland where internet access is mainly provided by ISPs working with Digital Subscriber Line (DSL) and Cablenet Technologies.

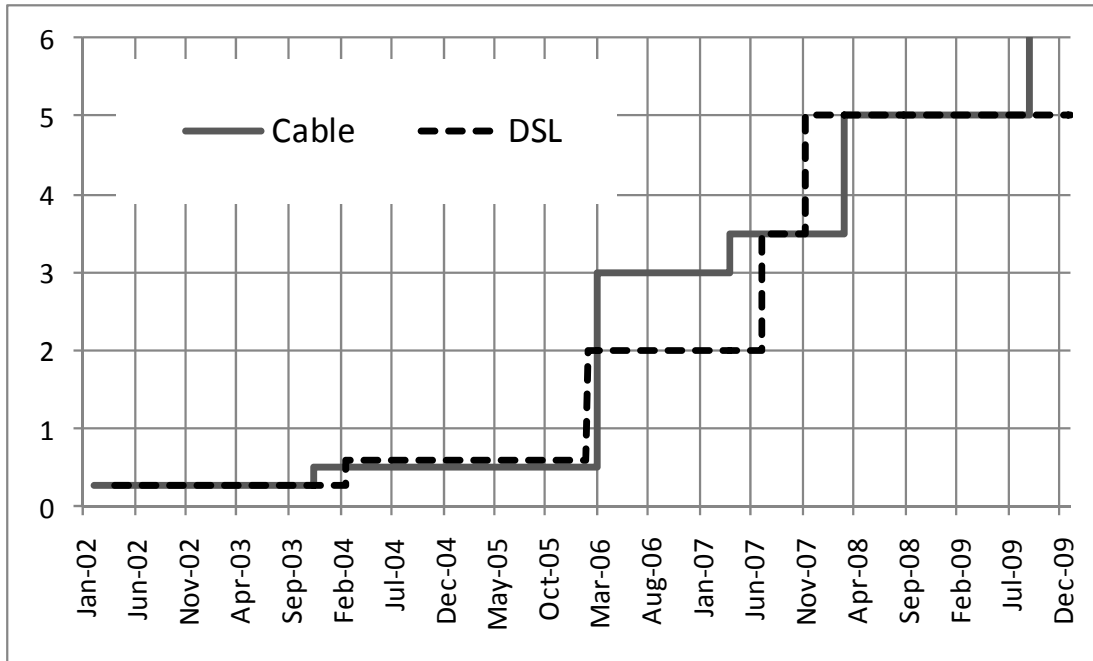


Figure 3: Evolution of advertised transmission data rate [Mbit/s] for the most widely used internet connection products.

Figure 3 illustrates the historical evolution of advertised internet access data rates for the most widely used broadband product in Switzerland. We will call this product, which costs around 35 Euros a month, the ‘standard product’. In Switzerland, over the last 8 years the access speed provided by the standard product roughly doubled every two years, and this was at constant price. Standard product users were upgraded automatically when the higher connection speed became available. The ISPs upgraded these access speeds in order to increase or keep their customer base - i.e. for product management and marketing reasons. This is illustrated by the fact that the two large internet providers in Switzerland typically followed each other’s speed increases within months. It is worth noting that, in addition, there were new products that regularly became available which had significantly faster speeds but at higher cost. However, in practice even when the cost of these products was only 10% higher cost than the standard one, only a minority of users upgraded to the significantly higher speeds.

Figure 4 shows the average speed experienced by users in Switzerland measuring their connection in comparison to the advertised speed. The historical evolution of these speeds is shown for the most widely used (standard) product and for the fastest available product at a given time. As advertised speeds continuously get faster (see Figure 3) it becomes more difficult to achieve those advertised speeds. Figure 4 illustrates that the effective speed of the standard product was about 95% and 72% of the advertised speed in January 2005 for DSL and Cablenet, respectively. At that time the standard product was at 1,200

kbit/s download data rate. However, in July 2009, when the standard product was at 5 Mbit/s the effective speed was only at 82% and 71% of the advertised speed for DSL and Cablenet, respectively.




Figure 4: Evolution of measured (i.e. effective) connection speed in comparison to the advertised speed for fastest and standard (most widely used) products in Switzerland.

If we start to look at speed in more detail, the experience of speeds becomes even more nuanced.

First of all one has to define the endpoints of the transmission i.e. the network connections. It makes a difference whether we look at speed for a download from the local server of one’s own ISP compared to a download from a remote server somewhere else in the world. For accessing the ISP’s local server we just need the so called ‘access network’ i.e. the network between the subscribers’ homes and the internet service provider’s backbone network. In access networks with Digital Subscriber Line (DSL) technology, the length of the telephone lines might be the limiting factor. In cable networks we might have local overload situations because many users share the same cable. For accessing a server somewhere in the world – perhaps on another continent – the traffic has to cross international lines going from the provider’s backbone network to the final destination somewhere in the world. This traffic is typically routed over different network paths for different providers, implying different speed limitation for different providers. Finally the speed limiting factor might

not be the transmission but the capacity of the server providing the information or the client software consuming the information.

One important factor for the copper wire connection of DSL (as opposed to cable connections), is the length of the line from the exchange.

Figure 5 shows how the maximum achievable DSL data rate decreases with increasing distance from the exchange. ADSL 1 is the older, slower scheme providing speed of up to 8Mbit/s download, and ADSL2+ is the newer, faster scheme with speed up to 24Mbit/s for very short distances. In both cases, the speed declines for homes located further from the exchange (more so for ADSL2+).

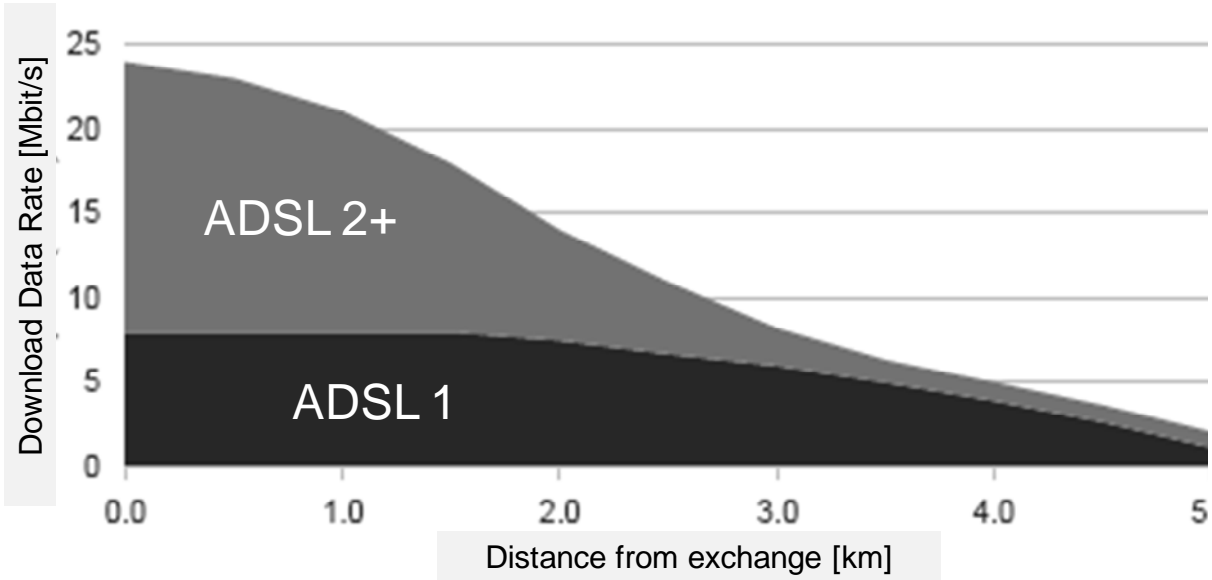


Figure 5: Theoretical maximum DSL speeds by length of line from exchange to premises (Reference: Ofcom 2009, p.11).

The British regulator Ofcom collected data from 1,621 panelists who had a broadband monitoring unit connected to their router in the 30 days from October 23rd to November 22nd, 2008. This analysis shows that consumers with the most popular broadband headline speed package (‘up to’ 8Mbit/s) only received an average actual download speed of 3.6Mbit/s (45% of headline speed).

In Switzerland similar data is collected for the most popular Swiss broadband headline speed package (‘up to’ 5Mbit/s) by letting users run tests in their internet browsers.

Figure 6 illustrates the download data rate as measured by 15,408 different consumers between April 1st and June 30, 2009. Again due to the line length, many customers only get the so called ‘fall back speeds’, illustrated by the numbers in the chart (Figure 6). Looking at these measurements in more detail, it transpires that 87% of the customers get at least 2.25Mbit/s (45% of headline

speed) and 61% of the customers get at least 4,000kbit/s (80% of the headline speed). This looks better than in the UK study, but still only one third of customers really do get the advertised download data rate of 5,000kbit/s.

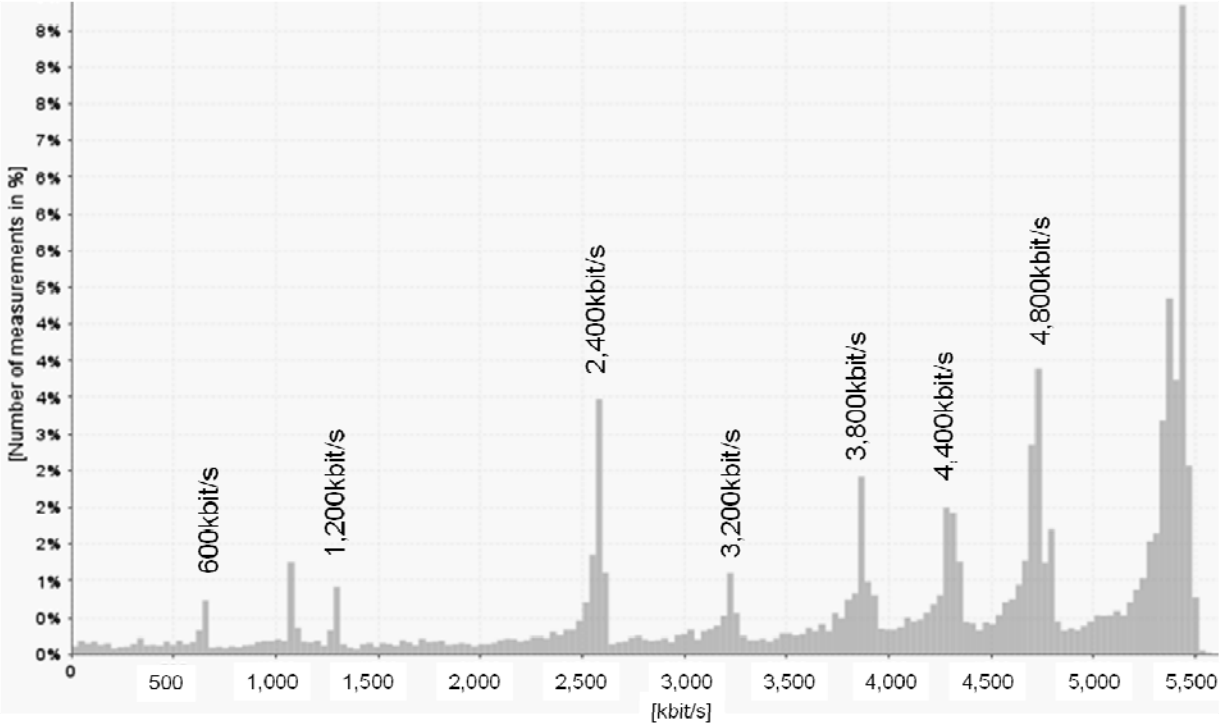


Figure 6: DSL: Experienced download data rate (in Switzerland).

The DSL line length impact has ramifications for regional experiences within a country: typically larger urban areas have shorter line lengths and hence higher speeds (Ofcom, 2009) p.28. It also has implications for the urban-rural divide: Ofcom found speeds were 15% higher in urban areas because of typically shorter line lengths than in rural areas (Ofcom, 2009, p.33). Not surprisingly, that same study showed that rural users were more dissatisfied with their speed. That means, to take the early example, if it takes 32 seconds to download a 2 minutes video clip in urban areas where 5Mbit/s are available it would take 4 minutes in rural ones where only 600 kbit/s download speeds are available. This is a very noticeable difference, and might well make the rural user think twice about whether it was worthwhile to download.

Line length is not the only important factor affecting speeds. The most important factor for the Cablenet internet connections is the number and activity of subscribers who are sharing the same Cablenet infrastructure. Cablenet operators try to keep the number of concurrent users in a so called ‘cell’ sufficiently low. But this is always not possible in situations where there is a fast growth of subscribers or where there is a rapid growth in network use.

Previous studies of both telephone traffic and early internet use underlined the fact that people communicate or go online at certain times rather than others, and this is mainly for social reasons (De Gournay and Smoreda, 2001; Lelong and Beaudouin, 2001 – both discussed in Haddon, 2004). For example, both types of traffic can be influenced by the timing of work and of school. It can also be influenced by social commitments (where there are norms about having dinner together as a family) or rules within the household (e.g. concerning when and for what purpose lone children at home can go online, for example). The point is that one of the factors affecting internet speeds is the number of users online at any one time, and so if social constraints and commitments mean that more of us go online at certain times rather than others. This creates overload situations (known as ‘congestion’ or ‘contention’) at those peak times that may make especially cablenet internet connections slower.

This is exactly what is shown in Figure 7, where 25Mbit/s cablenet speeds start to decline after working and school hours at certain locations - i.e. cells - while speed is fairly stable in other locations. This data from the Swiss speed tests illustrates how the effective end user data rate is influenced by concurrent users: Speed declines considerably into the evening, only starting to rise again after 10pm¹⁰.

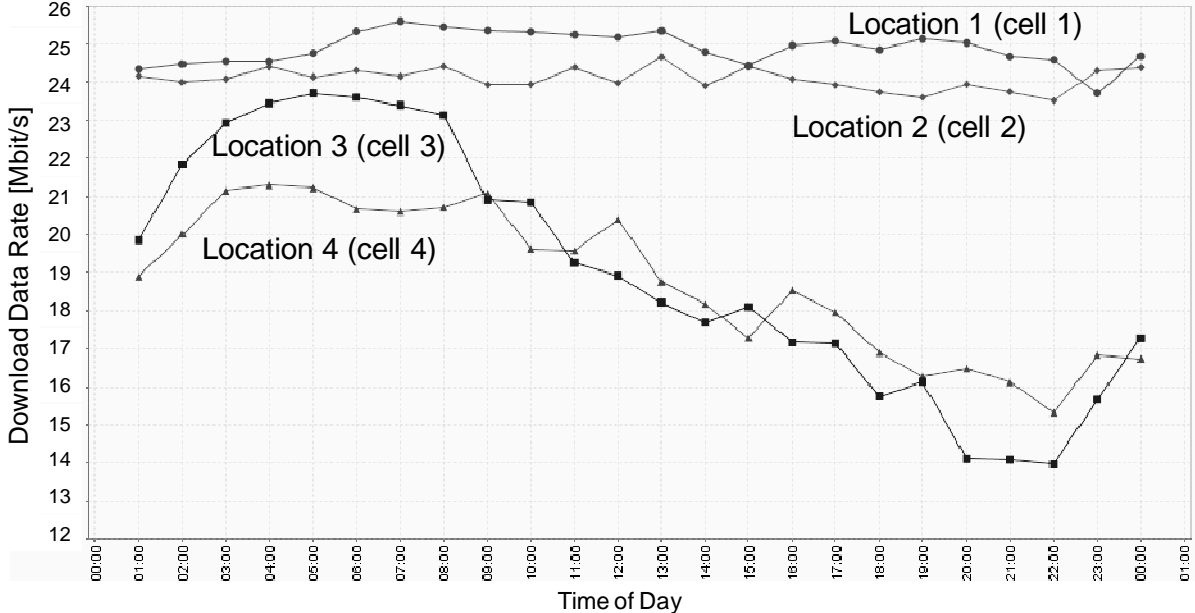


Figure 7: Cablenet: Experienced download speed by time of day (in Switzerland).

10 The pattern is different for the weekends (or vacation periods, such as school holidays) but the graph here is clearly strongly influenced by the figure for patterns for weekdays outside of vacations.

How much difference does this all make to user’s experiences? It depends. In certain situations (i.e. cells) there is no decrease at all because there is enough network capacity available to serve all concurrent users. In other situations speed decrease of up to 50% may be noticed.

At the very fast speed of the example in Figure 7 there is only a small response time difference when downloading a 2 minutes video clip. If somebody downloads at peak load time it takes 11 seconds i.e. only 5 seconds more than it would take during the low load period.

Overload may also occur in the backbone or at the server side. In the Ofcom example from Figure 8 it is shown, that there may also be a time of day dependency in the case of DSL access. In Figure 8 we notice a dip in speeds after 5pm and an increase after 11pm, where the peak hour of 5-6pm on Sunday can be over 30% slower than the average speeds during the off-peak hours between 4-7am. This dip is likely to be the result of contention within ISP networks and the broader internet or at the server side, meaning that speeds are also degraded as multiple users share the same bandwidth and server resources (Ofcom, 2009, p. 28).




Figure 8: Average download speeds for panellists by hour of day and day of week (UK).

What does this all mean to end users? In the Ofcom example from Figure 8, if they tried to download a 2 minutes video clip it would take 60 seconds at peak time instead of 38 seconds during low load periods.

Measurements in Switzerland outline speed differences that occur when connecting from an ISP backbone location to international web servers. The average download data rate from local servers (i.e. servers in Switzerland) is 82 Mbit/s, which is close to the maximum possible connection speed of the test system. In contrast, the download data rate from servers in Asia is a mere 2 Mbit/s, i.e. more than 40 times slower than the one to local servers. Looking at other continents we get 36Mbit from servers in Europe, 16 Mbit/s from servers in the USA and 4 Mbit/s from servers tested in Africa. Hence, it not only means

slower speeds from international servers, but considerable variation depending upon which part of the world one is dealing with. These differences are mainly due to differences in the response time when connecting to these servers, which leads to TCP/IP protocol speed limitations. The TCP/IP protocol uses a so called receive window size parameter to control the message flow. In widely used (e.g. Windows XP) operating system's default configurations this parameter is fixed and too small for large response time and high data rate values. With novel operating systems (e.g. new Linux distributions and Windows Vista / Windows 7) or with specific TCP/IP parameter tuning, the client and the server do adapt their receive window size and therefore can achieve higher speeds at high response time.

The cross-national dimension

Figure 9 shows how the average download data rates evolved since 2002 in Austria, Germany, Switzerland, France and Italy based on end user measurements. The average effective download data rate increased from approximately 500 kbit/s in 2002 to 3,500 kbit/s in 2009 in Austria and Italy. In Switzerland and France it increased to around 7000 kbit/s, the same as in Germany, which had already higher data rates in 2002. The statistical relevance of these data is not so high because there are not so many measurements in some countries. However, the general increase of the experienced (effective) download data rate is obvious.

These speed differences are due to a combination of speed limits on some paths and on round trip times which lead to speed limitations imposed by the transmission protocol i.e. by TCP/IP. The speed furthermore differs by ISP due to the different paths used to access these reference servers.

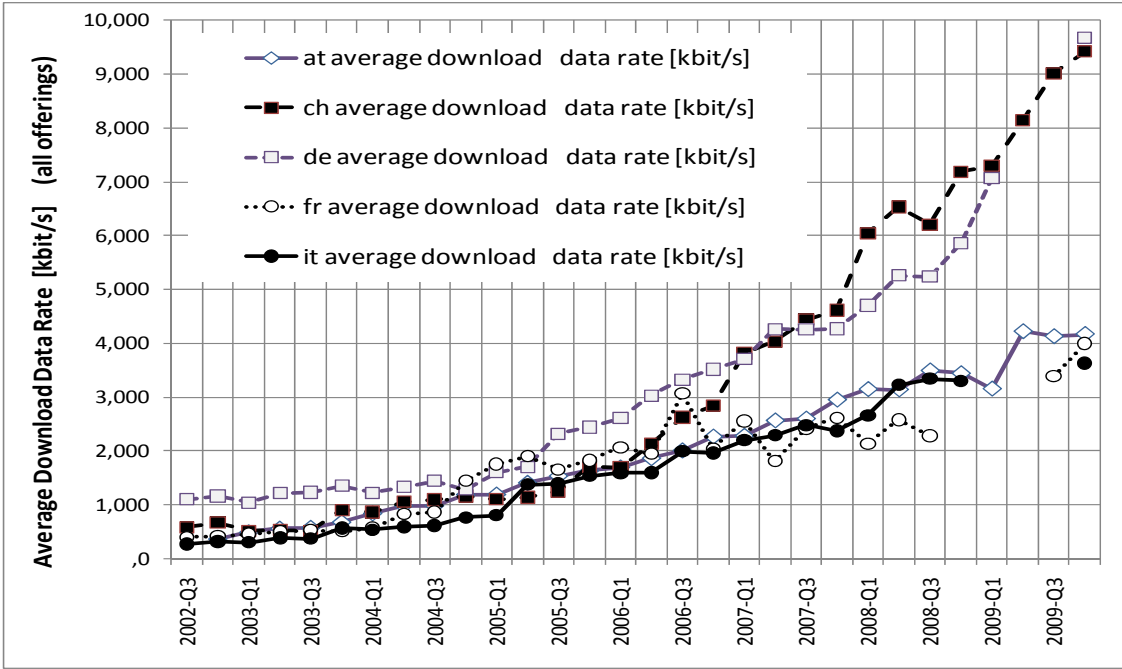


Figure 9: Evolution of experienced download rates by country.

Looking at these historical data in Figure 9, some of the gaps in speed (here download speed) have grown wider. The Swiss and Germans, recently joined very dramatically by the French, have increasingly higher actual speeds than the Austrians and Italians – who have to wait longer when downloading standard international files, like the Skype one cited earlier. Returning to our example of downloading the 143 minutes DivX Quality Movie (i.e. 940 MBytes), if the average Swiss, German or French user enjoys a 7Mbit/s download data rate in January 2009 whilst the average Italian or Austrian user gets only 3.5Mbit/s, it will take the latter over twice as long to download the same file, 36 min instead of 16 min.

Figure 10 uses the COST298 data to compare the distribution of speed in the five countries Austria, Germany, Switzerland, France and Italy in 2009. At the point when the measurements were taken, the highest headline speed offered (called the nominal data rate on the graph) was 20,000 kbits/s (or 20 Mbit/s). The ‘data points’ at the bottom refer to the numbers in the various national samples, i.e. the number of times people checked their speeds with the speed test applications. The number of measurements carried out is higher since in practice most check their connection speed three to five times.




Figure 10: Distribution of experienced download speed in different countries.

Hence, 10% of this, (the 10 on the horizontal axis), was 2 Mbits/s, and 65% of the Austrians, French and Italians had at least this speed, 75% of Germans had a least this speed and 85% of the Swiss had a least this speed. Country

differences exist at this end of the scale, but as we look at even slower speeds, going to the left in the graph, country differences diminish since most people in more of the countries can have at least these speeds. If we now go to the right in the graph and look at 50% of the nominal data rate, 10 Mbit/s, only 5% of Austrians and Italians have this speed or greater than it, whilst the figure for France is just over 15%, and for Germany and Switzerland 30%. This conveys some idea of the level of country variation and reminds us that before we consider averages the reality is that we are always talking about proportions of the internet users with different speeds.

Conclusions

In general, this chapter has been an interdisciplinary exercise, unpacking the technical considerations to be considered and translating this into terms that connect with some of the more social science based approaches and questions about the nature of internet use. Moreover, it has attempted to achieve this through comparative analyses both within and between countries, an approach that could be used in subsequent research.

The first substantive point to observe is that the historical data showing the growth in broadband speed and file sizes illustrate the interdependence of the content offered and the available download speed. As soon as higher speeds are broadly available the application designers begin to offer better quality, e.g. resolutions in the packages involving image, and video files of longer duration. On the other hand, availability of high volume content then in turn drives the evolution of access speed. The implication is that this is an on-going process. In other words, speed was not only an issue in the early years of the internet but it will always be an issue because there will always be some new applications 'eating up' all bandwidth.

Certainly the online distribution of software packages, updates and service packs would not be possible at the speeds we had ten years ago. Some things have been made possible, or made more practical, by the faster internet. But while on an everyday basis many of the tasks people conduct can be achieved in reasonable time, history suggests that applications will be developed that in the future require more of the public to upgrade to yet faster connections. The currently much discussed application service-providing solutions (aka "cloud computing") require speeds of at least 15Mbit/s to allow for sub-second response time when working with current PowerPoint files of about 1.6MB file size.

These histories also provide us with a good idea of what will happen in the mobile world in the next few years to come. The speeds that we had on the fixed internet connections ten years ago are now common on mobile phone connections. It is expected that the evolution of mobile phone data rates will be

even faster than it was on the fixed internet. But the interdependence of content offered and available download speeds will be very similar.

The next question addressed in this report is that of what these changes in speed and file size mean to the user, that is, how are they experienced? This required some technical clarification, since the average size (at any one time) of different types of file, types of video, and types of photo will vary. Nevertheless, the tables provide some sense of the history of when (and at what connection speeds) an increase in speed was noticeable to a user and for what types of application. This becomes important because while social science research regularly asks people (such as the Italian youth cited earlier) what they do online, the technical analysis here conveys a sense of how much time it requires to achieve different goals (including the quality of what can be viewed) – for example, the time taken for peers to acquire, exchange or post online audio-visual material.

The comparative analysis within countries takes this a stage further. When we hear that a certain speed is available for companies or people in certain areas, not all those people have it. The analysis then goes on to show how much where you live can make a difference to the speed experienced, particularly relevant for rural vs. urban differences. Moreover, the time of day when you go online also makes a difference, and in some cases their other commitments means that certain people may only go online at more congested times. All these factors highlight the ways in which the experience of using the internet can vary, certainly in terms of the time taken and hence attractiveness of particular actions.

Finally, the comparative analysis conducted cross-nationally illustrates not just the gaps between countries (in terms of the average experience of speeds) but also how these gaps change over time, with certain gaps expanding or narrowing, sometimes slowly, sometimes more dramatically. As before, translating these into the time that users in different countries might take to achieve certain goals casts some light on why practices are more attractive in some countries more than in others. While other social factors influence patterns of use, the technical constraints noted in this chapter can also play a part.

Bibliography

Agrawal, N. Bolosky, W.J., Douceur J.R., & Lorch, J.R., 2007. A five-year study of file-system metadata. *ACM Transactions on Storage*, 3 (3), Article 9.

Bakardjieva, M., 2001. Becoming a domestic internet user. *e-Usages*. Paris, France 12-14 June 2001. ADERA: Bordeaux

De Gournay & Smoreda, Z., 2001. La sociabilité téléphonique et son ancrage spatio-temporel. *e-Usages*. Paris 12-14 June 2001. ADERA: Bordeaux.

Growth of the Average Top 1000 Home Page, 2008

Haddon, L., 2004. *Information and communication technologies in everyday life: a concise introduction and research guide*. Oxford: Berg.

Haddon, L. & Heinzmann, P., 2009. *The Experience of Broadband Speeds*. A report for COST298. [Online] Available at: <http://www.cost298.org/uploadi/editor/1264711174Changing%20broadband%20speeds%20report.doc> [Accessed 1 March 2010].

Haddon, L. & Kim, S-D., 2007. Mobile phones and web-based social networking - emerging practices in Korea with Cyworld. *The Journal of the Communications Network*, 6 (1), January-March.

Lelong, B & Beaudouin, V., 2001. 'Usages d'internet, nouveaux terminaux et hauts debits: premier bilan après quatre années d'expérimentations. *e-Usages*. Paris 12-14 June 2001. ADERA: Bordeaux.

Ofcom 2008. UK broadband speeds 2008. Consumer experience of broadband performance: initial findings. [Online] Available at: http://www.ofcom.org.uk/research/telecoms/reports/bbspeed_jan09/ [Accessed 19 July 2009].

Ofcom 2009. Accessing the internet at home: a quantitative and qualitative study among people without the internet at home. [Online] Available at: <http://www.ofcom.org.uk/research/telecoms/reports/bbresearch/bbathome.pdf> [Accessed 19 July 2009].