

van der Vorst, Tommy; Brennenraedts, Reg

**Conference Paper**

## Understanding the demand growth for digital connectivity

22nd Biennial Conference of the International Telecommunications Society (ITS): "Beyond the Boundaries: Challenges for Business, Policy and Society", Seoul, Korea, 24th-27th June, 2018

**Provided in Cooperation with:**

International Telecommunications Society (ITS)

*Suggested Citation:* van der Vorst, Tommy; Brennenraedts, Reg (2018) : Understanding the demand growth for digital connectivity, 22nd Biennial Conference of the International Telecommunications Society (ITS): "Beyond the Boundaries: Challenges for Business, Policy and Society", Seoul, Korea, 24th-27th June, 2018, International Telecommunications Society (ITS), Calgary

This Version is available at:

<https://hdl.handle.net/10419/190350>

**Standard-Nutzungsbedingungen:**

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

**Terms of use:**

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*

# Understanding the demand growth for digital connectivity

Tommy van der Vorst<sup>1</sup> & Reg Brennenraedts

## Abstract

While connectivity supply is growing exponentially, likely as a result of developments in the semiconductor industry, research on connectivity has mostly focused on the demand side. Such approach is however unable to account for the introduction of unforeseen services, which is also supply-driven. In this study we seek to validate the existence of a ‘residual’ of unexplained growth and quantify it as the difference between supply of connectivity and demand from existing service category. The hypothesis is confirmed: an increasing fraction of internet traffic volume expected at high levels of aggregation (i.e. an internet exchange) is unexplained by existing service category growth.

**Keywords:** digital connectivity, bandwidth, internet traffic, exponential growth

## Introduction

### *Research question*

In 1998 Jacob Nielsen introduced his law stating that a high-end user's connection speed grows by 50% per year [1] – a figure that he has since validated. Not only the speed, but also the cost of data transmission is improving exponentially; data from [2] shows a year-over-year decline of about 46% of the price per Mbit/s of transit connectivity between 1998 and 2015. Andrew Odlyzko even found that the growth in backbone traffic was around 100% per year around the year 2000. [3] Additionally several vendors of telecommunication equipment have reported exponential growth: Cisco [4] and Sandvine [5] are the two examples.

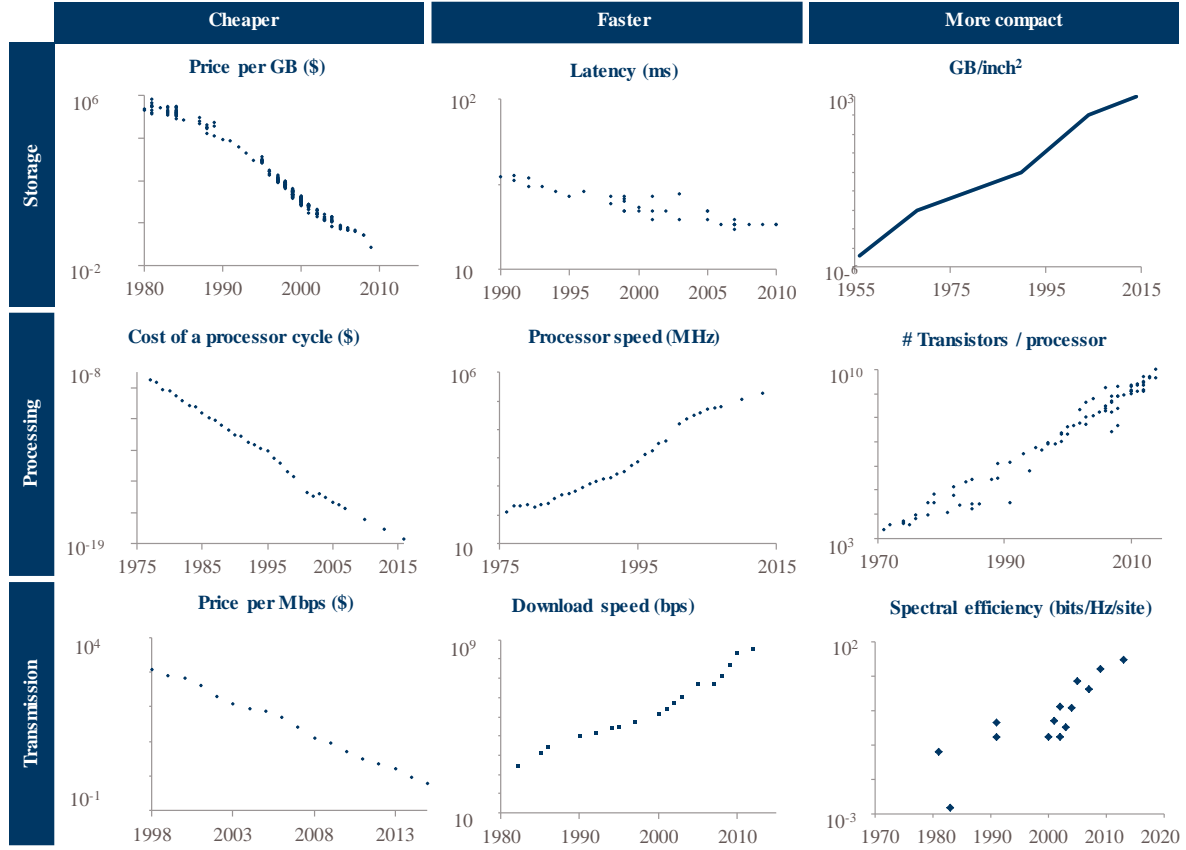
Exponential growth curves are not unique to bandwidth or even telecommunications in general. In many other fields of ICT exponential growths can be observed, with Moore's Law being the most well-known, stating that the number of transistors in dense integrated circuits will double every two years. [6] Figure 1 below depicts similar exponential trends

---

<sup>1</sup> Corresponding author. Dialogic *innovation & interaction*, Hooghiemstraplein 33, Utrecht, the Netherlands. [vandervorst@dialogic.nl](mailto:vandervorst@dialogic.nl).

in storage, processing and transmission for prices, speed and size. Note that these effects are likely all driven by the same set of developments in the semiconductor industry.

Figure 1 Exponential developments in information technology over the past decades [7]

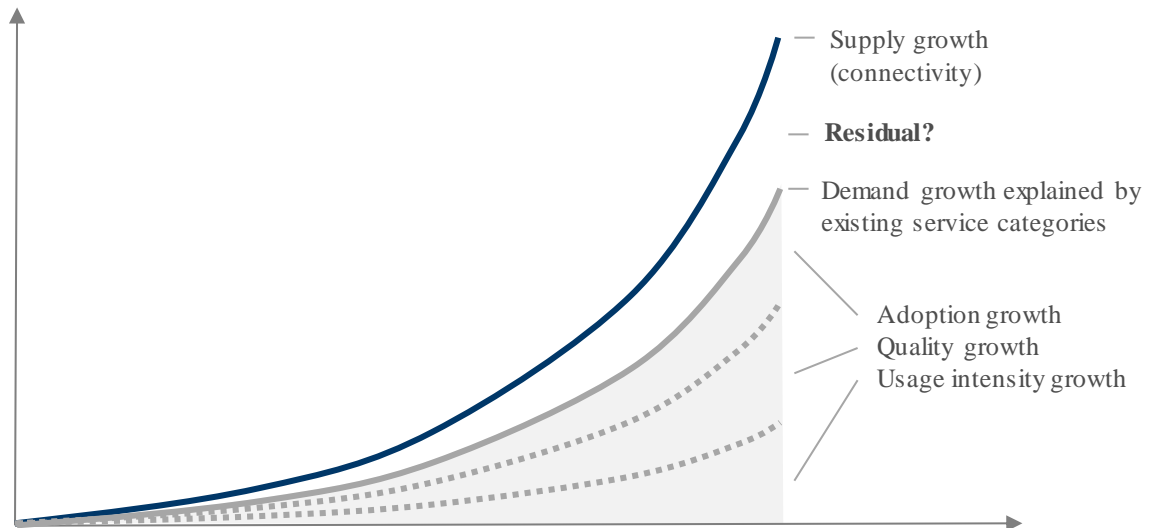


While connectivity supply is growing exponentially, likely as a result of developments in the semiconductor industry, academic research on connectivity has mostly focused on the demand side and its exponential growth. Such approach however is questionable at best considering the large role that technology plays in the exponential developments in connectivity supply, and the effect it has on the demand side. Understanding the interrelationship between the supply and demand side for connectivity is crucial in understanding, and also predicting, future demand for connectivity.

Growth of bandwidth capacity on the supply side appears to go hand-in-hand with growth on the demand side. Availability of bandwidth has allowed services to improve their quality and users to adopt services and increase their usage. Existing services are not able to fully account for demand growth. Historically there have been introductions of ‘revolutionary’ services too, which have suddenly increased the demand. Consider for instance Napster (1999), BitTorrent (2001), YouTube (2005) and Netflix (2007). Interestingly no research up to this point has focused on the exact mechanism of this growth; e.g. what amount of growth can be explained by growth (from adoption, quality increases, and

intensity increases) in existing service categories, and what remaining growth is expected to come from yet unknown service categories? Figure 2 shows this concept schematically.

Figure 2 Breaking down the demand growth for internet connectivity: is there a residual to be expected?



Assuming exponential growth in the ICT industry, both on the supply side (connectivity) as well as the demand side, we seek to identify the importance of this residual through the following main research question:

*Given exponential growth on the supply side, which drivers explain the growth of residential broadband connectivity demand from existing service categories, and to what extent is a residual expected?*

#### *Scope of this research*

In this study we focus on internet traffic via fixed networks for households of different applications in advanced economies in the period 2013-2022. Below we will go into more detail.

- We only focus on internet traffic. This includes all the traffic from and to subscribers via their internet service provider. However, other services that are part of a triple play package typically use virtual networks are excluded from the scope. This typically means that video (often traditional linear TV) that is sent to set-up boxes or devices with comparable functionality is not seen as internet traffic. This holds for both TV via DVB-C and IP-TV. Obviously video that uses the internet (e.g. Netflix, YouTube) is seen as a part of internet traffic. For telephony roughly the same holds. Telephony services that are not routed over the internet ('over the top') are excluded. These traditional services typically use a separate

virtual network. Services that use the internet (e.g. Skype, Facetime) are included in internet traffic.

- We focus on fixed networks. In advanced economies these typically consist of access networks based on fibre, DSL and/or coax.
- We focus on households. We define households as subscribers that use the consumer proposition of ISPs. Obviously, many households consist of more than one person. Therefore we see the CPE as the nexus for measurement.
- We focus on advanced economies in this study. We apply mainly apply data from Western Europe, but also use data from North America to triangulate data or to fill gaps in the data.
- We apply the timeframe 2013-2022. We chose to apply a rather small time window of ten years. We do this because the highly dynamic nature of the internet economy.

### *Limitations of this research*

An obvious limitation of our model that is developed, is that using it outside the scope addressed above severely limits its validity. And although we will present measurements that form a highly consistent picture looking backwards, predicting the future is quite a different story. Finally, a fundamental issue with predictive quantitative models is that it is generally impossible to model the so-called ‘black swans’, that is, unlikely events that nevertheless can have substantial impact on the outcome. [8] A few examples of such ‘black swans’ in the context of our study are the following:

- Blocking of certain services or content. Several governments have already decided that certain services or websites should not be accessible at all (e.g. file sharing websites) or only with an explicit opt-in (e.g. porn sites). Although such blocking can usually be circumvented by technically skilled users, the majority of users will simply be unable to gain access.
- Separation or decentralization of the internet. Several countries have installed virtual ‘Chinese walls’ that disallow many types of foreign services.
- In the future, triple play providers could redesign their services and stop using virtual networks for IP-TV but direct this traffic via the internet line. If we apply the system describe above, this will lead in a spectacular overnight increase in internet traffic. The other way around, we can also speculate that ISPs start delivering certain services not via the Internet, but via a specific virtual network.

### *Challenges of this research*

This study has several challenges, of which the following two are the most relevant:

- Analyses on data usage per subscriber have shown that the distribution of average usage per subscriber per day does not resemble anything like a normal distribution. A small set of top users account for a vast amount of traffic, while the remaining set of less active users account for a relatively small amount of traffic. Using averages would result in a suboptimal fit of the model. In order to tackle this issue, we apply different users groups. These groups are defined based on literature and data analyses.
- Many measurements have a relative low resolution; typically for example we see measurements on traffic volumes over the course of a month, rather than Mbit/s specified over smaller time frames. We have developed a model that combines measurements at the aggregate level (which cover a larger population but are less specific) with measurements at lower levels (which cover a much smaller sample, but are more specific). We also attempt to break down aggregate figures to individual user groups and service categories.

### *Assumptions of this research*

As any forecasting model, ours too requires making a set of assumptions:

- First of all, we make the assumption that the relevant trends can be modelled. We assume that certain developments in the past can be extrapolated.
- Second, we make the assumption that the adoption of services follows an S-curve. This is consistent with standard literature on the diffusion of innovations. [9]
- Third, we assume that some services are generic and will result in full adoption in the long run. But we will also assume that some services are specific by nature and will only be used by a subset of the total population.
- Fourth, we assume that households in advanced economies face no restrictions in the supply of bandwidth. If we look at data from the European Commission, we see that in 2016 over 99% of the households in the EU have access to broadband. [10] The availability of NGA broadband (with at least a sustained rate of 30 Mbit/s download) is over 90% in most Western European countries. [10]

## **Theoretical framework**

### *Previous research*

Previous research addresses (some of the) applications that are considered drivers for residential broadband connectivity demand, e.g. [3], [5] and [11]. These papers each apply

a different unit of analysis. Some papers only discuss the traffic on internet exchanges, some focus on backbones [12] and other on the level of access networks. In our perspective, analyses in the lowest level (i.e. subscribers) is preferable, since it also allows many of the analyses on higher levels (e.g. backhaul, backbone, et cetera). Data on higher levels severely limits the possibility for many analyses on the micro-level.

Among policy makers there has been some debate on the need for symmetric access networks that provide the same upload and download bandwidth. Traditionally, access networks are not symmetric and they provide more bandwidth for download than for upload. As the argument goes, consumers consume and don't produce: *nomen est omen*. On the other hand, some policy makers stress the need for symmetric access networks stating that consumers are not able to produce since the current networks hinder them. The European Commission stated that "*not only download speeds are important in that context, but higher symmetry (much higher upload speeds) and lower latency may be required for innovative services and applications.*" [13] Also, if we for example look at some recent ITU-standards, we find specifications symmetric access networks.<sup>2</sup>

There are dozens of papers discussing the optimization of streaming video over the internet. [14] Video uses a lot of resources of the internet ecosystem. [15] This stems from the simple fact that streaming video (1) requires a relatively large amount of bandwidth and (2) consumers are not inclined to accept to wait for content. The fact that humans have a brain that is highly focussed on visual input could be the most important driver for this. In fact, some scholars state that half of our brain focuses on vision. [16] We are simply made for processing very high amounts of visual data. While computers even in very early stages beat humans in simple calculations, it took computers decades to match human levels of visual pattern recognition. Humans finds it much easier to recognize the face of their mother than to calculate  $782 \times 624$ .

#### *Availability of data on this research topic*

If we look at the data on this research topic, we typically find many research that uses measurements in the core of the network. For example they use the traffic over an Internet Exchange or via back bones networks. A limited amount of research actually uses data on end users. In these papers, we typically find that there is a high level of aggregation. In other words, all the data from a large set of (known) clients is integrated in one set and end-users cannot be distinguished anymore. In many other publications only one (type of) service is addressed; for example the use of video of eHealth or the impact of smart grids on networks.

---

<sup>2</sup> See for example ITU-T G.9807.1

For this research we have acquired measurements of traffic from and to CPEs in households. We have these data on actual the level on individual households and we have this data for two ISPs. The data is split to different types of services. Please note that the data is not privacy sensitive in any way; the only data available per household was a (non-identifying) number, and the volume of traffic divided over different types and over time. This data is allows us to make new analysis on this research topic.

### *Hypotheses of this research*

Based on the existing literature and the functionality of the model, we come to three hypotheses:

- H1: Video will remain the primary driver for bandwidth demand growth.
- H2: Upload will grow in importance for the next few years.
- H3: Existing types of services will not be able to fully account for the future growth in bandwidth: there is a residual, to be filled by ‘unforeseen’ services.

## **Methodology**

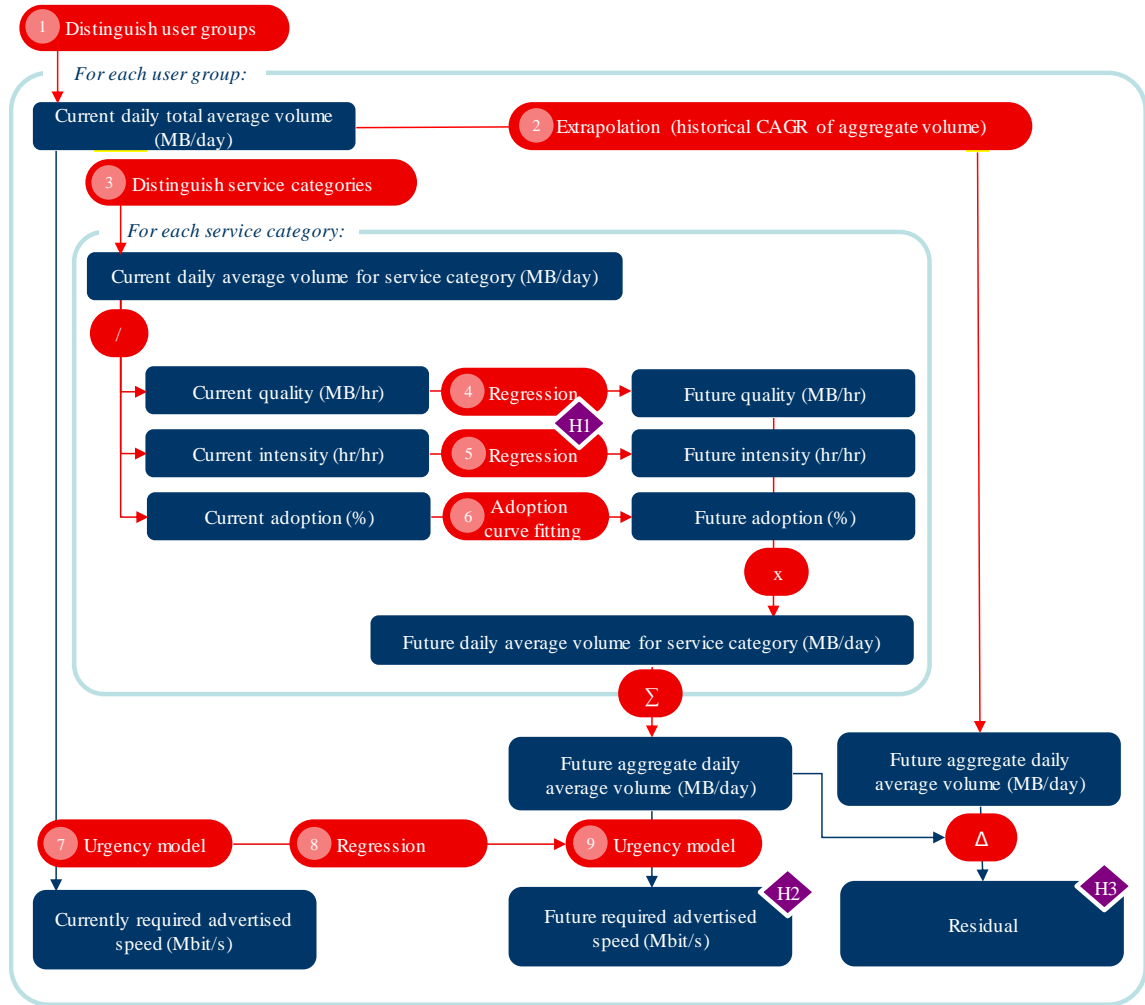
### *Overview*

In order to be able to answer the research questions, a model was developed allowing estimation of future demand for bandwidth. The model combines various aspects of broadband demand (derived in smaller ‘sub models’). Figure 3 below provides an overview of the overall research setup. The model is a refinement of a model used in our earlier studies on the phenomenon: [17] and [18].

Bandwidth demand is expected to differ greatly between user groups and service categories. We therefore choose to model different categories and user groups separately. The basic idea of the model is that by distinguishing user groups, service categories within these groups, and growth drivers for each of these services (adoption, intensity and quality), the demand growth can be explained to a high degree. By making extrapolations at this level and then summing the results, the model will be able to predict future demand from existing services. This estimate can then be compared with an estimation based on an extrapolation of aggregate demand (over the whole group or over the whole population). The difference between these two estimates signifies the ‘residual’: the growth that may be caused by yet unknown service categories.



Figure 3 Overview of the model used



### Estimation per service category and user group

Given a certain service category as well as a user group, we estimate the household demand in terms of daily volume, expressed in the amount of data (as MB) demanded on average over the course of a day. We model this demand as an integration of the following components:

- The service *quality*. This relates to the data volume the service needs to transmit to the user in order to provide the service at a certain level of quality. Over time, we expect services to increase their service quality (e.g. an online video service will increase the video resolution, a music service will increase the bit rate of the audio served, et cetera). The service quality is expressed as the data volume (as MB) required to provide the service over the course of a certain time interval (an hour).

- The usage *intensity*. The demand for bandwidth can vary still after adoption; a household may choose to use the service (category) for a longer time period per day, or the service may be used by multiple persons simultaneously (as is common with online video streaming services, who differentiate their subscriptions by the number of simultaneous streams allowed). The usage intensity is defined as the number of hours per hours in the day a service is used. As we are also interested in the concentration of the usage over the course of a day, we also estimate a concentration metric.
- The service *adoption among relevant users*. After introduction of a service (category), households will (according to innovation theory) gradually adopt the service. The service adoption is defined as the percentage of the households that have adopted the service category (by using at least one service in the category).

In order to estimate future demand, we employ regression models to predict future quality, intensity and adoption. For the quality aspect, we assume that (for the foreseeable future) quality will be increasing at the same pace as it has in the previous years, and that trade-offs (with respect to transmission vs. computation vs. storage) remain constant (this assumption is validated separately).

## *Operationalization*

### *Literature defining service categories and user groups*

Rogers [9] studies the diffusion of innovations (and ideas in general) among a society. Rogers distinguishes five phases in the adoption process, which are linked to different social groups who sequentially adopt and ‘spread’ the innovation. Rogers also provides a general ‘rule of thumb’ regarding the distribution of these groups. The first to adopt are *innovators* (2.5% of the population), followed by *early adopters* (13,5%), *early majority* (34%), *late majority* (34%) and finally *laggards* (16%). The groups can also be linked to lifecycle phases for a product or service: from introduction to growth, maturity, saturation, to decline.

With respect to service categories, we strived to cluster the categories commonly found in literature, e.g. [5] [19] and [20]. To estimate the traffic demand for these concrete services, we used a top-down as well as a bottom-up approach. For some service groups, specific literature was available to estimate this parameter, while for other, no concrete literature sources were found. In that case, we applied a bottom-up approach, meaning that we estimated the traffic for this service based on the traffic needed for one single action, multiplied by estimates for the amount of actions per day and the number of users.

### ***Traffic measurements and literature for current demand***

In this study we are interested in the amount of traffic flowing between CPEs and the first endpoint of the ISP (e.g. the traffic over the *access network*). In general there are three ways of measuring such traffic: *user centric* (measure the traffic as it is generated and received at the user side), *network centric* (measure the traffic as it flows over the network) and *service provider centric* (measure the traffic as it arrives/leaves the service provider's network). Obviously the user-centric approach is the most reliable, as it allows for the least amount of interference. An example of this approach is the one taken by SamKnows, which places a box measuring internet performance at the consumer's premise. [21] It is however also the most invasive method, requiring the installation of monitoring software or hardware at the end-user side. Network-centric measurement methods are generally easier to deploy, and can provide detailed measurements when actually performed close to the user (see e.g. [22]). In many cases however ISPs are reluctant to cooperate due to legal prohibitions with respect to deep packet inspection as well as end user privacy.

An important caveat with respect to traffic measurements is that part of the traffic measured may be involuntary, e.g. not the result of an actual (functional) demand of an end user. The most impactful appears to be mechanisms that update software in the background (e.g. [23]) as well as botnet traffic [24]. Service-centric approaches may have to deal with filtering larger amounts of involuntary traffic originating from non-residential users (e.g. bots) [25].

In order to mitigate the issues described above, we chose to combine various methods for measuring and triangulate between different measurement locations. In this study we relied on data sets provided to us by ISPs (network-centric measurements), generally at an aggregate level (one Dutch ISP, 2013) as well as aggregates at the individual subscriber level (one smaller Dutch ISP). Literature sources used include [22], Cisco's Visual Networking Index [19] as well as Sandvine's Global Internet Phenomena reports (e.g. [5]). In order to distinguish the total traffic volumes by service, ISPs provided us with figures on the amount of traffic to be deducted for non-internet traffic (e.g. IPTV provided by the ISP) as well as a distinction of traffic over different destination/source ASes. This data was combined with insights from literature.

### ***Estimation of growth factors***

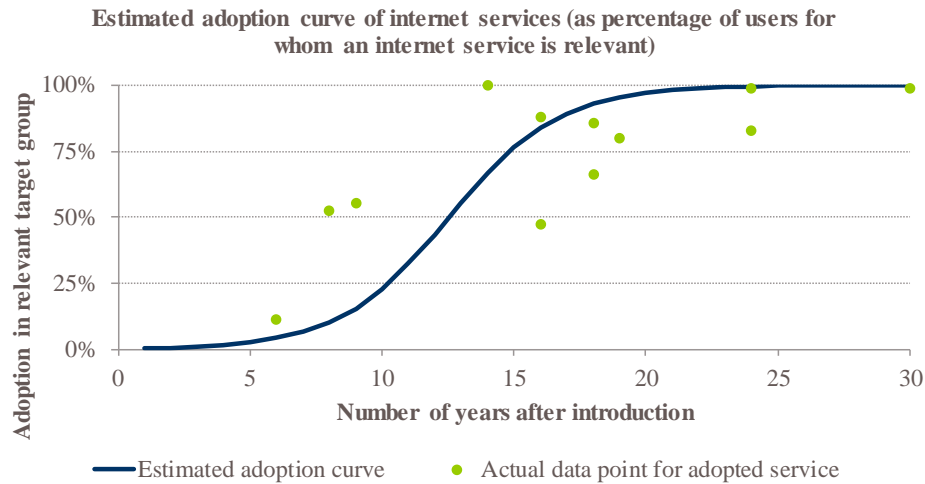
For each service category we estimate a growth factor based on three components: (1) growth from adoption, (2) growth by increased intensity of use and (3) growth from increased quality of the service.

### Adoption growth

With respect to the adoption component, we follow innovation literature. Rogers [9] concludes that due to the underlying mechanisms through which innovation spreads, the adoption rate for an innovation rises slowly after introduction, is at its maximum when a certain critical mass is reached, and then decreases again as maximum adoption is reached. The cumulative penetration over time can, according to literature, be modelled as an S-curve. For this study we assume that this holds for internet-based consumer-oriented services as well.

As we are interested in modelling the future adoption for a service category, we attempt to fit an S-curve for internet-based consumer-oriented services. A similar method is described in [26]. The fitted S-curve is then used to predict future adoption of a service, given its current position along the curve (innovation phase as discussed earlier) as well as its age (which together define the speed at which the diffusion has happened up to the measuring date). Historical adoption data were obtained from Eurostat [27]. Adoption is capped at 100%<sup>3</sup> Figure 4 shows the resulting fitted S-curve.

Figure 4 Best-fitting adoption curve based on adoption data for the service categories under study



Notice that there is significant spread between services, especially in the middle phases of adoption (mainstream phase). Many of the service categories are however in their later

---

<sup>3</sup> This is not trivial: it is possible for a service to be adopted more than 100% depending on the definition of adoption. Our definition does not include intensity of use and therefore cannot be above 100% (e.g. a household that holds two subscriptions for a video on demand service, e.g. because it is used on two devices simultaneously, will still count as one).

stages, where the fitted curve fits the data relatively better. The fitted overall S-curve can be described as follows:

$$a(t) = \left( \frac{1}{1 + e^{-12 \cdot \frac{t}{T} - 6}} * \frac{M}{100\%} \right) * 100\%$$

In the above formula  $t$  is the number of years after introduction of a service,  $T$  is the fitted total adoption time,  $M$  is the maximum adoption level (assumed to be 100%).

We use the same curve for each service, but vary the position along the curve for each service individually, depending on each service's age and maturity. The curve is adapted to each user group as well (the period of adoption as well as the start date of the adoption are shifted). The year-over-year growth rate of adoption of each service was calculated from the difference in adoption as estimated by the curve during the analysed time period.

The S-curves fitted above concerns the full population. As we are interested in adoption within each user group as well, we use a separate S-curve for each user group; the weighted sum of these S-curves approximates the fitted S-curve for the full population.

The adoption curves for the individual group each use the same formula, but have a different value for  $T$ . Following innovation literature we also add a delay component, which indicates the year in which the group starts adopting the innovation. Finally we should account for the fact that the groups each represent different fractions of the total user group:

$$a_n(t) = \left( \frac{1}{1 + e^{-12 \cdot \frac{t-d_n}{T_n} - 6}} * \frac{M_n}{100\%} \right) * w_n$$

In the above,  $T_n$  refers to the total adoption time for group  $n$ ,  $M_n$  refers to the maximum adoption within group  $n$ , and  $w_n$  represents the relative size of group  $n$  to the total population (The weights of all groups should sum to 1).

The group curves, when added together, should then closely approximate the aggregate adoption curve, e.g. the error term to be minimized (while also matching innovation literature on the characteristics of each group) is:

$$\text{argmin}(w_n, d_n, T_n, T, M) \sum_{t=0}^T \left[ a_{all}(t) - \sum_{n=0}^N a_n(t) \right]$$

In this formula,  $N$  is the number of groups,  $n$  is the group number,  $a_n$  is the adoption curve for each group,  $a_{all}$  is the adoption curve for all users. The final values were fitted and chosen in such a way that they match innovation literature as well as provide a good fit

to the aggregate adoption curve when summed together. Table 1 shows the final fitted values for each group.

*Table 1 Adoption curve parameters for the different user groups*

Adoption of services	All inter- net users	Laggards	Main	Innovators	Power users
Start adoption after (years)	0,0	15,0	9,0	4,0	0,0
Time to 100% adoption (years)	25,1	10,1	6,0	5,0	4,0

### *Quality and intensity growth*

Estimates on quality and intensity growth are obtained from literature for each service category separately. Sources include [28] (on streaming video) and [29] (peer-to-peer file sharing)., among others. Of particular interest (due to its large relative share in traffic volume) is online video. At this moment, HD is the *de facto* standard for the resolution of streaming video at this moment. This requires a bit rate of approximately 4 Mbit/s for compressed video. In the future 4K will increase its market share. Expectations for Western Europe are between 40% of the households in 2020 [30] and 20% in 2023 [28]. The effective bitrate of compressed 4K video is around 15 Mbit/s.

Note that in many measurements, adoption and intensity of use are provided as an aggregate figure (e.g. they provide adoption figures that may exceed 100%, to reflect the usage of a service by more than one person simultaneously in a household, for instance). In cases where we were unable to obtain separate measures, we used such totals.

### *Mapping traffic volume to minimum required provisioned speeds*

Data from ISPs was used to model the relationship between traffic volume and the resulting minimum sufficient required provisioned speed. First of all, we gathered the aggregate traffic volumes for individual FttH connections on the network over a period of 24 hours. We then calculated the concentration of this traffic over time.

Knowing the distribution of traffic volume over the course of a day, we calculate the minimum sufficient bandwidth by looking at the busiest time interval, and dividing the traffic volume during that interval by the length of the interval. Comparing these results to the actual maximum speeds of the connections (which, in the FttH case, matched the provisioned speeds) leads to a translation formula. In order to translate demand volume to demand speed, we first calculate the speed it would take to transfer the traffic volume in a given period of time, as follows:

$$s(v, u) = \frac{v}{u}$$

In the above,  $s$  is the required bandwidth (in Mbit/s),  $v$  is the traffic volume (in Mbit) and  $u$  is the time period. We then find the value for  $u$  that best matches the current speed-to-volume ratio by comparing against current offerings from ISPs. In other words, we calibrate the average ‘urgency’ of traffic against the currently prevalent ratios of advertised speeds versus volume.

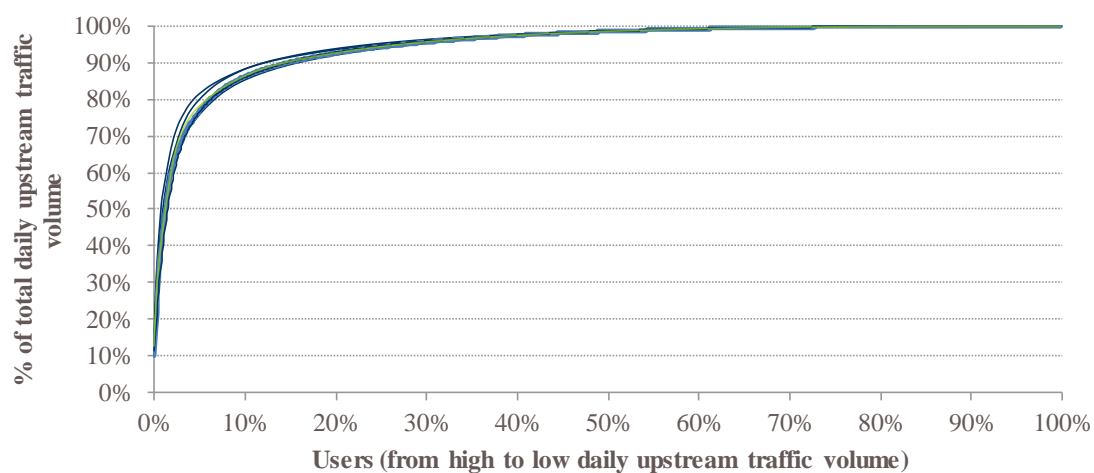
The mapping between volume and required bandwidth may change over time (e.g. traffic may become more or less ‘urgent’). In order to account for this effect we take multiple measures and extrapolate any changes over time.

## Results

### *User groups*

Measurements at a Dutch ISP’s network largely confirm the suspicion that the distribution of traffic volume consumption over users is highly skewed. Figure 5 shows the distribution of upstream traffic over the user base, as measured by a Dutch ISP on its network over the course of seven days in 2013 for a random sample of 1000 subscribers. Notice that about 2% of the users is responsible for 60% of all upstream traffic already.

*Figure 5 The distribution of upload traffic over users, based on measurements by a Dutch ISP in 2013 (the different lines indicate different measurement days over the course of a week).*



The results displayed above largely fit with Roger’s phases of innovation diffusion. While showing a great disparity between innovators and mainstream data, the results do not provide conclusive evidence to justify a separation between early and late majority users. For this study we therefore settle for four user groups: power users (2%), innovators (20%), mainstream users (60%) and laggards (20%).

### *Service categories*

From literature and interviews with experts, we have identified several service categories:

- **Consultative web browsing:** all activities related to obtaining information. Usually, this is done through the World Wide Web (WWW).
- **E-mail, conversational applications, social media and 'Web 2.0':** all online services where sharing user-generated content with other users plays a central role. This category includes social networks like Facebook and Twitter, as well as review websites and market places (eBay and Airbnb).
- **Remote backup:** services that allow consumers to periodically make a back-up of their files (in some cases restricted to certain types of files, e.g. photos or documents), only meant to be retrieved in an emergency. Examples are Carbonite, Tarsnap and Backblaze.
- **Online video and music (streaming and peer-to-peer):** any service that provides non-linear video streaming (users can start watching what they want at any time). The best known services are YouTube and (more recently) Netflix. Additionally any form of online music streaming or download service. This includes online music shops (Apple iTunes Store) as well as streaming services (Spotify, Deezer). It also includes service that allows files to be shared between end-users from their home computer(s). Older examples are Napster or Kazaa and a contemporary one is BitTorrent.
- **Remote workplace and work file access:** secured access to documents and other services normally only accessible in the workplace. This is a typical business application.
- **File download:** any type of download not already included in other categories. Examples are app/software downloads and updates and OS downloads and updates such as Windows Update and Steam.
- **Online gaming:** service that allows the end-user to play games against or with other end-users through the internet. An example is the Blizzard server, facilitating games like World of Warcraft, Diablo and Starcraft.
- **Personal cloud storage and file synchronization:** a service that allows files to be stored at a centralized location, from where it is accessible from a plurality of device types, only requiring an internet connection and certain software. An example is Dropbox. This service could also be used for business applications.
- **Other services:** any service that does not fall into one of the categories defined above. An example could be the offloading of mobile traffic by means of a femtocell or Wi-Fi.



An important aspect to consider when comparing downstream and upstream traffic demand is the fact that downstream traffic actually requires an upstream traffic flow as well; most applications (those based on TCP) will require ‘acknowledgements’ to be returned for each downstream packet in order to know whether delivery succeeded. As we are reasoning from the services themselves (and only the *useful* traffic volume), we need to account for this overhead as well. From discussions with ISPs, we arrived at an estimated 10% overhead traffic in one direction to allow traffic in the other direction to flow (e.g. if there is 100 MB of downstream traffic, we add 10 MB of upstream ‘acknowledgement’ traffic, and also the other way around).

### Traffic measurements

Traffic levels at the Amsterdam Internet Exchange (AMS-IX) provide insight in the total demand for internet traffic at the aggregate level, albeit at a different location in the network. Figure 6 shows the development of volume of the traffic exchanged over AMS-IX.

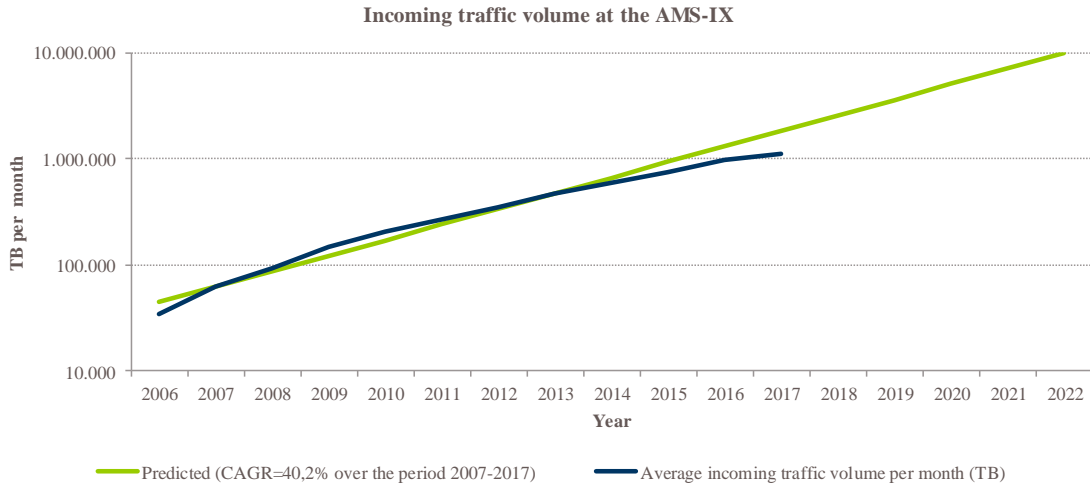


Figure 6 Measured and predicted growth of the average monthly incoming<sup>4</sup> traffic volume at the AMS-IX [31]

The compound annual growth rate (CAGR) for the traffic volume is estimated as follows:

$$CAGR = \left( (y - y_0)^{\frac{1}{t-t_0}} - 1 \right) * 100\%$$

Here,  $y_0$  is the value in the starting year ( $t_0$ ), and  $y$  is the observed value after  $t$  years have passed. Volume is extrapolated from a base value and the CAGR, as follows:

<sup>4</sup> As an Internet Exchange primarily exchanges traffic between networks and does not consume or produce traffic, the total volume of incoming traffic is in theory equal to the total volume of outgoing traffic. In practice, the totals differ slightly due to technical reasons (packet loss, disruptions, et cetera).

$$y(t) = y_0 * \left(1 + \frac{CAGR}{100\%}\right)^{t-t_0}$$

Over the period 2004-2014, a compound annual growth rate of 34% can be observed. Using data over the period 2005-2016, the growth rate observed is 29%. For 2022, we predict that the aggregate traffic flowing through AMS-IX will be 5x that of the volume flowing today, or between 4 and 6 exabyte (million terabyte) per month, on average. Note that the growth figures at the aggregate level are not comparable to growth figures at the consumer level. A significant portion of traffic from and to consumers in the Netherlands does not flow over the AMS-IX. In addition, traffic over the AMS-IX also includes peer-traffic between networks that are not necessarily ‘eyeball’ consumer networks.

Traffic measurements were obtained from various sources, and at various locations and periods. Table 2 shows an overview of the different sources employed, as well as the total traffic volume calculated to a monthly total.

*Table 2. Average total monthly volume of traffic per household, according to various sources.*

Source	Period	Location	Total traffic volume (Mbyte/month)
Sandvine	2013 H1	Western Europe	13,400
Sandvine	2013 H2	Western Europe	17,400
ISP A	2013-11	The Netherlands	7,466
ISP B	2013-11	The Netherlands	2,655
Cisco VNI	2014	Western Europe	38,800
ISP A	2015-Q1	The Netherlands	55,955
ISP A	2016 Q1	The Netherlands	77,924
ISP B	2016-Q1	The Netherlands	76,986
ISP B	2016-Q2	The Netherlands	75,315
Cisco VNI 2015	2015 (predicted)	Western Europe	46,661
Cisco VNI 2015	2016 (predicted)	Western Europe	56,115
Cisco VNI 2015	2017 (predicted)	Western Europe	67,484
Cisco VNI 2015	2018 (predicted)	Western Europe	81,157
Cisco VNI 2018	2016	Western Europe	27,000
Cisco VNI 2018	2021 (predicted)	Western Europe	78,000

The sources indicated in Table 2 as well as the AMS-IX data provide varying estimates for the growth rates of aggregate traffic volume. We suspect many of the differences to be related to differences in the definition of geographical scope as well as what is counted as ‘internet traffic’. Weighing the sources above we arrive at an estimated CAGR of 40.5% for downstream and 44.1% for upstream traffic volume.

## Adoption of services

Figure 7 Estimated adoption of services by households

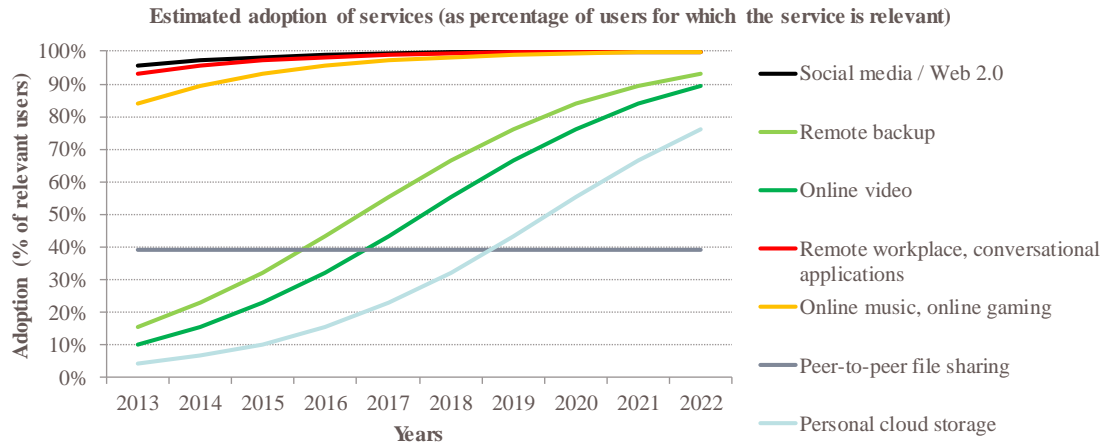


Figure 7 shows the resulting estimation for the adoption curves for the different service categories. Notice that peer-to-peer filesharing was excluded, and is modelled as a flat line instead. Adoption of peer to peer services is primarily influenced by the risk posed by viruses or spyware and penalties for downloading or sharing illegal content. In addition, the introduction of legal alternatives (and their adoption) has a negative effect on the adoption of peer-to-peer services. Indeed, if legal alternatives become successful, the adoption of peer-to-peer may decrease before it has reached its maximum adoption potential. Whether a ‘churn’ from illegal peer-to-peer towards legal distribution is occurring or will occur in the future is still a topic of debate (see e.g. [32]).

Several studies confirm our hypothesis that there will no further uptake of peer to peer file sharing services, and some even find a decrease. Studies by market research firm GfK show a 40% penetration of Video on Demand Services among internet users in the 13+ years group in Q4 2015. [33] [34] Sandvine predicts a negative adoption rate in their 2014-2019 forecast for the Western European region. [35] A Norwegian study confirms the strong repressing effect of ‘all-you-can-stream’ services on banning digital piracy. [36]

## Future demand volume

Integrating and summing all data collected up to this point allows for the estimation of future traffic demand. Figure 8 shows the estimated demand for upstream traffic volume over the period 2013-2022, broken down by service category.

Figure 8 Estimated future demand volume (all users) – downstream traffic

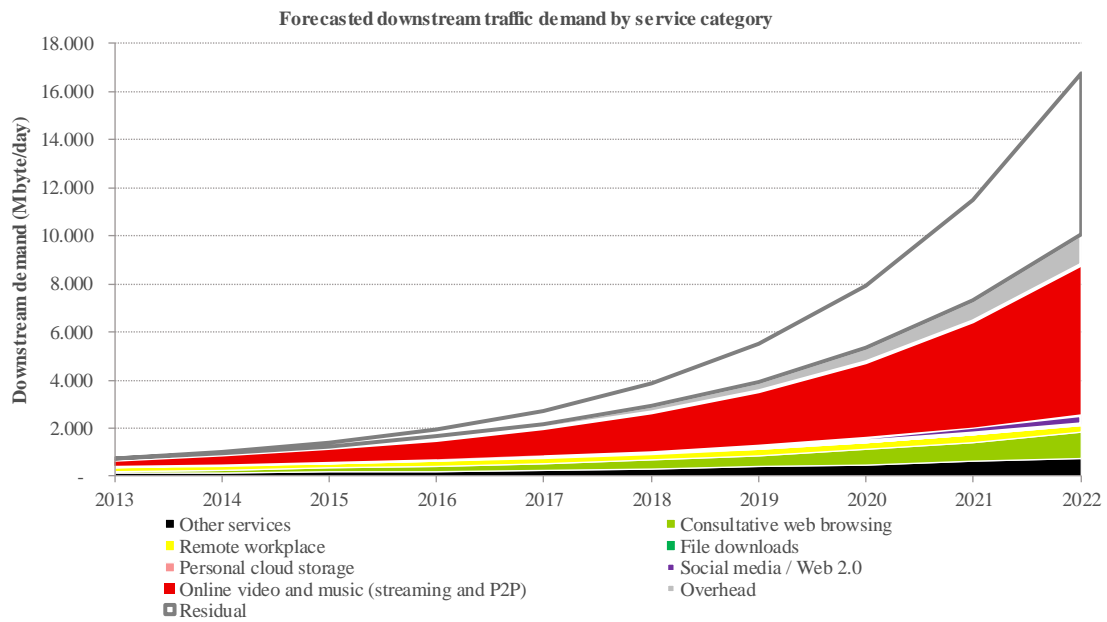
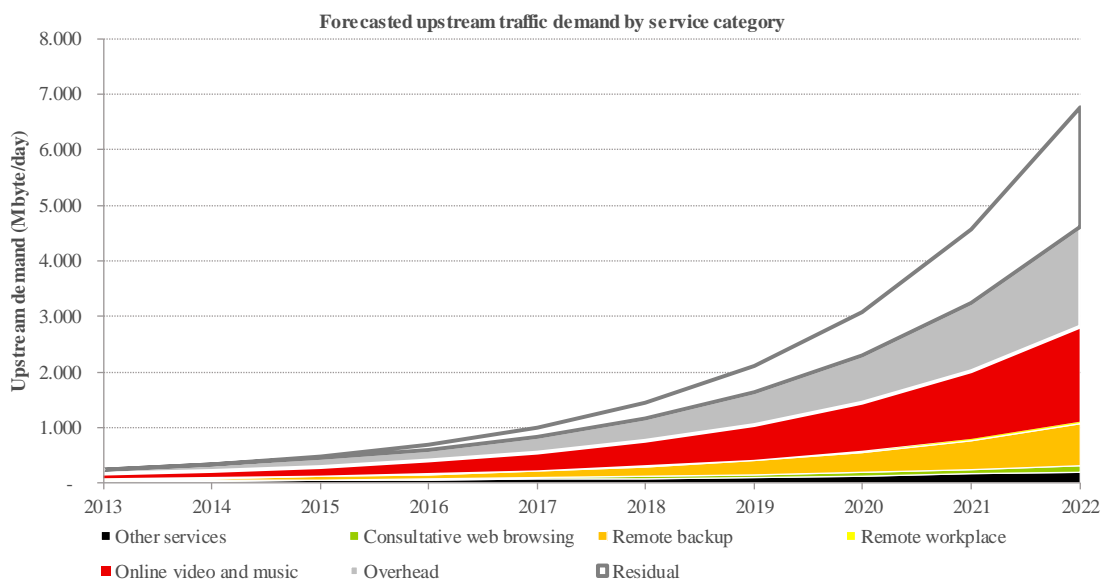


Figure 9 Estimated future demand volume (all users) – upstream traffic



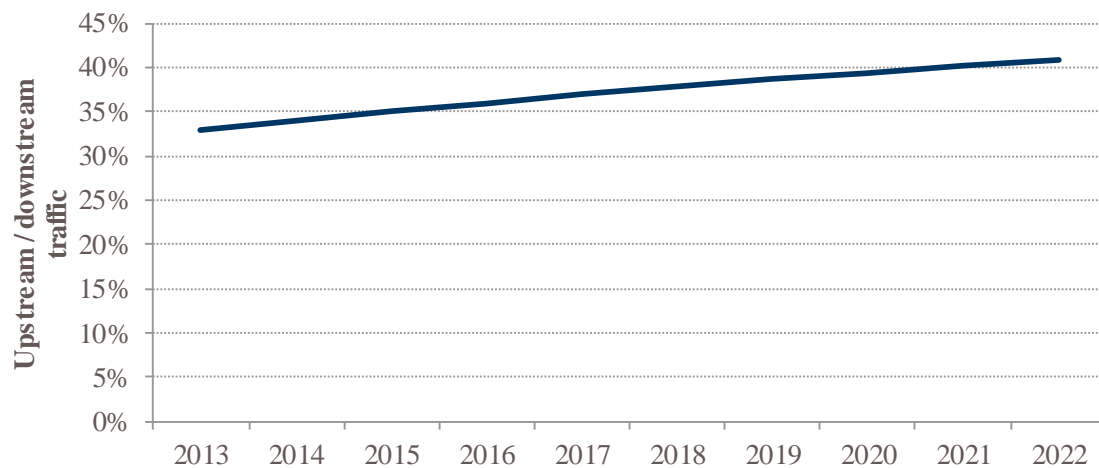
Unsurprisingly, online video and music are the major driver of downstream traffic growth. The growth of online video consumption is primarily due to intensity growth (i.e. the move towards HD and higher resolutions) but also growth from adoption by lagging users (primarily driven by the introduction of legal video streaming services such as Netflix).

There is also a major role for future revolutionary services in the downstream direction, which we expect to cause 40% of the total traffic volume demand by 2022. The total downstream demand for 2020 is estimated at 16.6 Gbyte per day, per household.

### *Upstream vs. downstream*

Figure 10 shows the forecasted upstream traffic volume divided by the downstream traffic volume. Interestingly, upstream traffic is expected to grow faster than downstream traffic. Nevertheless, even in 2022, downstream traffic volume is expected to still be more than double the upstream traffic volume.

*Figure 10 Forecasted volume ratio between upstream and downstream traffic*



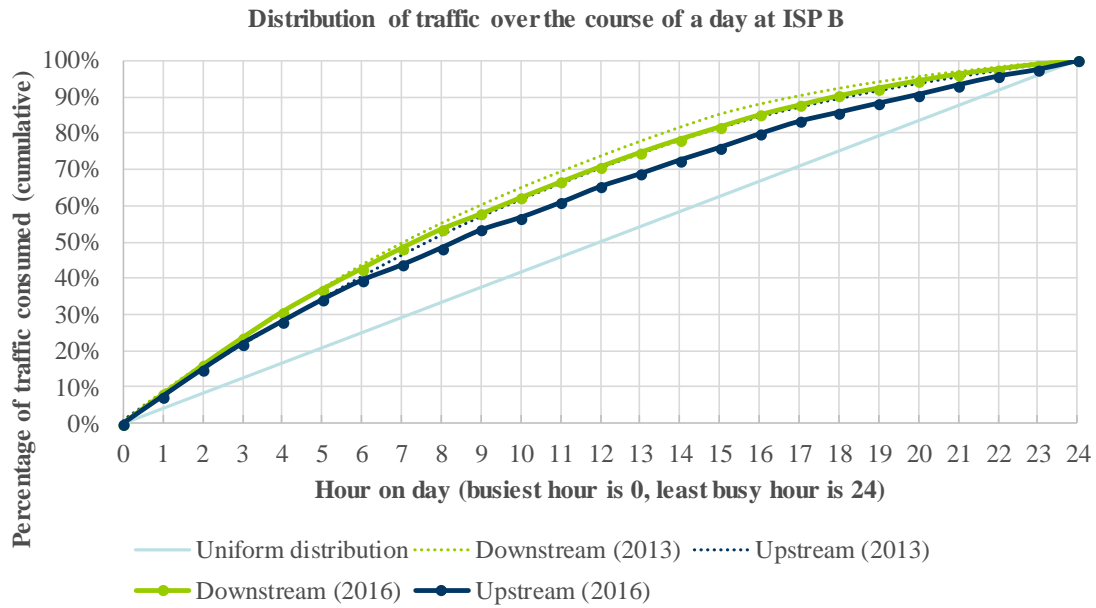
### *Traffic volume to minimum required provisioned speed*

Finally we look at the minimum required provisioned speeds that are required given the demanded traffic volumes. Figure 11 shows the measured distribution of traffic over the course of a day for connections on a Dutch FttH network, measured in 2013 as well as 2016, for upstream and downstream traffic separately.<sup>5</sup>

---

<sup>5</sup> This result was also replicated for measurements on another ISP's network.

Figure 11 Cumulative distribution of traffic volume over the course of a day [17]



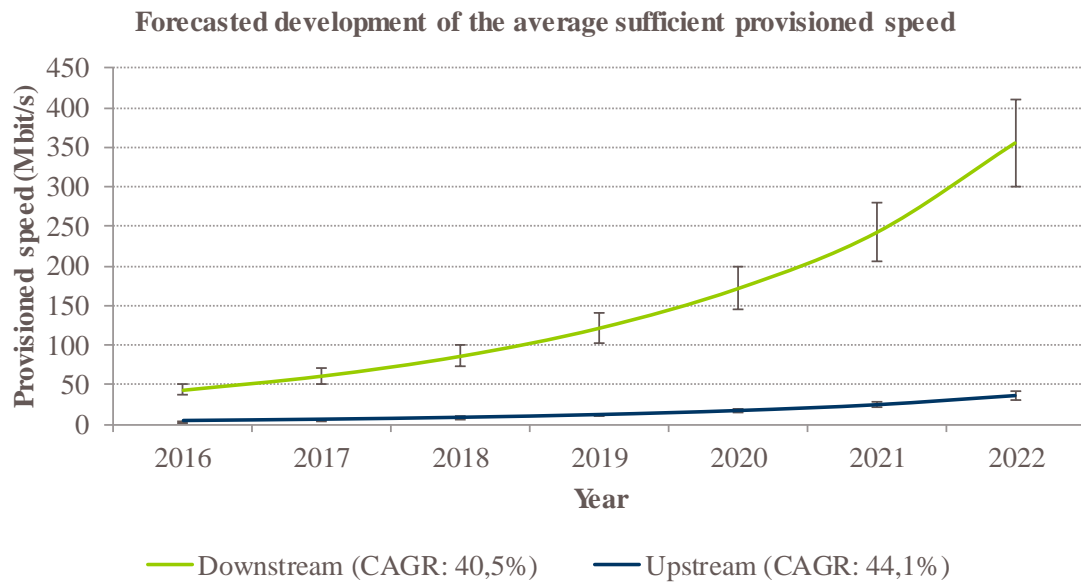
From Figure 11 it appears that the distribution of traffic volume over the course of a day has changed between 2013 and 2016. In [17] we drew the conclusion that downstream traffic has become more evenly spread over the course of a day, leading to lower minimum required bandwidth at similar traffic volumes. The same holds for upstream traffic, although less strongly so; this means that the ratio of upload to download bandwidth still increases, but at a slower rate than the upload to download volume ratio.

#### *Future demand for bandwidth*

Translating the forecast traffic volumes to minimum sufficient provisioned speeds leads to the forecast shown in Figure 12.

The forecast shows that an average subscription will have a sufficient provisioned downstream speed of about 355 Mbit/s in 2020 and an average sufficient provisioned upstream speed of 37 Mbit/s. This estimate is only valid assuming that the current advertised speeds are a reasonable indication of the speed of a ‘sufficient’ connection. In addition, it is assumed that the urgency of traffic will not change. The error bars in Figure 12 show the speeds required if urgency changes by 20% (i.e. traffic needs to be transferred in 20% more or less time than it currently is).

Figure 12 Forecast of average sufficient provisioned speeds



### Bandwidth demand by user group

Table 3 shows the forecast demanded bandwidth by user group. Large differences exist between the user groups. Since the intensity is expected to grow equally among the user groups, the differences exclusively arise from the adoption of these services. The graph shows that especially laggards are expected to adopt services on a large scale, in particular remote backup, conversational applications, online video and to a lesser extent online music. Personal cloud storage is on the other hand expected to be adopted by innovators on a large scale. This is because currently the service is assumed to be used mainly by power users and will subsequently dissipate to the innovators group.

Table 3 Forecast minimum sufficient provisioned speeds by user group

		2013	2014	2015	2016	2017	2018	2019	2020	2021	2022
Power users	Up	24,9	34,8	48,8	69	96	136	191	269	380	550
	Down	147,8	198,7	267,5	360	486	656	887	1200	1625	2267
* Note power users:		The estimations for the sufficient provisioned speeds for power users are based on a different method in which traffic for peer-to-peer is modelled to be supply-driven rather than demand-driven. This means that the power users will always maximally utilize the provisioned bandwidth.									
Innovators	Up	3	4	6	9	12	17	25	35	50	72
	Down	46	61	82	111	149	201	272	367	496	694
Mainstream users	Up	0	0	1	1	1	2	2	3	4	6
	Down	7	9	12	17	22	31	42	57	77	109
Laggards	Up	0	0	0	0	0	0	1	1	1	2
	Down	1	1	1	2	3	4	5	7	9	17
All users	Up	1	2	3	4	6	8	12	17	25	37
	Down	16	22	31	44	62	87	122	172	243	356

## Discussion

### *Conclusions*

#### *Future bandwidth demand*

In this study, we have developed a method for estimating households' future bandwidth demand. The future bandwidth demand for existing services can be estimated by combining data on current usage of services with projections of the adoption and growth rates of service usage intensity. Future revolutionary services are expected to play a major role in the growth of bandwidth demand. Such services are however difficult to foresee. We have developed a 'next best' estimation model for these services, where we modelled a probability distribution of the impact of revolutions and their expected occurrence frequency.

The method developed in this study was applied to data on residential subscriptions in Western Europe. We predicted the compound annual growth rate (CAGR) of upstream and downstream traffic demand to be 44% and 40% respectively. While demand in 2013 is on average 15.3 Mbit/s downstream and 1.6 Mbit/s upstream, in 2020 demand is expected to increase to 165.4 Mbit/s downstream and 20.1 Mbit/s upstream. Large differences can be found between the types of services and the user groups. Power users, constituting 2% of the total users, will require 1,155 Mbit/s downstream and 315 Mbit/s upstream by 2020, whereas the laggards will only need 6.6 Mbit/s downstream and 0.8 Mbit/s upstream by that time.

#### *Hypotheses of this research*

Our first hypothesis is that growth from existing service categories (be it from growth related to increased adoption, quality or intensity) would not be able to account for the full growth in bandwidth. The hypothesis appears to be confirmed. An increasing fraction of internet traffic volume expected at high levels of aggregation (i.e. an internet exchange) is unexplained by our model, which we assume to cover all currently relevant service categories. Note that the observed gap's size may even be underestimated if (as expected) internet traffic will become more decentralized and closer to edges (e.g. due to *edge computing* as well as the increased reliance on content delivery networks).

The second hypothesis concerned the question whether only video will remain the primary driver for bandwidth demand growth. Even compared to the residual, online video will remain the primary driver for residential internet traffic in the foreseeable period.

Thirdly, we stated our hypothesis that the demand for upload traffic will grow in importance over the next few years. This hypothesis too is confirmed: our model predicts that the volume of upstream traffic demanded will grow to 41% of the downstream demanded volume, up from an estimated 38% in 2018.



### *The residual*

In this study we compare growth that can be explained by existing services with the average growth, and arrive at the conclusion that not all growth can be explained by existing services. There exists a ‘residual’ which is assumed to be growth caused by the introduction of yet-unknown services. In the past, the introduction of new services that were previously not envisioned, but went on to drive a large fraction of bandwidth demand has happened many times, such as has happened with the introduction of Napster (1999), BitTorrent (2001), YouTube (2005) and Netflix (2007). While the nature of ‘unforeseeable services’ does not allow us to identify these up front, we can still make assumptions about how often such services are introduced, and what their expected impact could be on demand.

For the purposes of this study, we calculated the expected probability of the development of future services and their expected impact on demand growth. Table 4 is an excerpt of the way we propose to model such ‘revolutions’. Each row represents a particular kind of revolution: a ‘once every year’ revolution occurs frequently, but has a negligible impact on growth; the ‘once every fifty years’ kind of revolution is rare, but has enormous impact. For each type, we calculated the expected number of times such a revolution will occur in the seven-year period. We subsequently modelled the impact of each kind of revolution inversely to the occurrence frequency: high-frequency revolutions have a low impact, and low-frequency revolutions a higher impact. The impact is measured as an increase in yearly traffic in proportion to existing traffic (e.g. an impact of 5% indicates that each year, traffic will grow by an additional 5%). We chose to assign 50% growth and 200% impact to the rarest type of revolutionary event (‘once every 50 years’) and scaled the impact to the other types using a quadratic and cubic interpolation formula, respectively.

*Table 4. Example for modelling the expected growth in bandwidth from ‘revolutionary services’*

Revolutionary service type	Next seven years		Estimated impact on traffic demand		Expected growth in seven year period	
	Chance of at least one introduction	Expected number of services	Min	Max	Min	Max
Once every 1 year	100%	7,00	0,00%	0,00%	0,0%	0,0%
Once every 2 years	99%	3,50	0,00%	0,00%	0,0%	0,0%
...	...	...	...	...	...	...
Once every 49 years	13%	0,14	23,06%	55,34%	3,3%	7,9%
Once every 50 years	13%	0,14	25,00%	60,00%	3,5%	8,4%

Figure 13 gives a graphical overview of the distribution of the probability of revolutionary events happening and the corresponding impact of such an event on bandwidth demand.

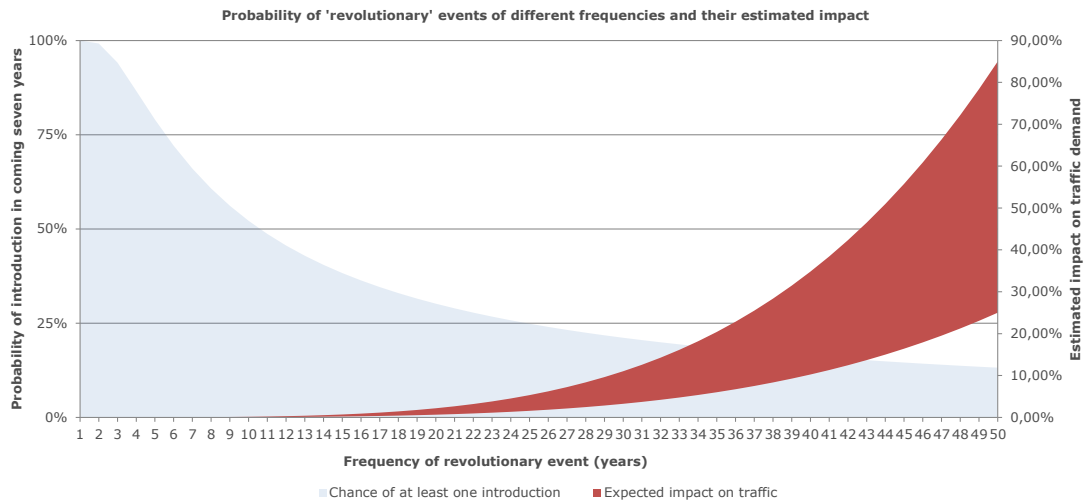


Figure 13. Probability of 'revolutionary' events of different frequencies and their estimated impact

By summing the expected growth percentages, it is possible to calculate a compound annual growth rate over the seven-year period. We estimate the year-over-year growth from 'revolutionary' events in the coming seven years will be between 5,5% and 11,1%. The percentages are highly dependent on the estimated impact of the various types of events.

### *Will bandwidth growth continue exponentially?*

In this research, we included modelled demand for future revolutionary services. Future revolutionary services are services that do not exist at this moment, but are expected to come about in the coming years. We modelled their traffic by means of a probability distribution of the impact of revolutions and their expected outcome frequency. The future revolutionary services are expected to constitute 26% of the upstream traffic by 2020. One typical example of a driver for this traffic is a surge in the number of connected devices and accompanying services in a household.

A fundamental question remains whether, and if how long, exponential growth will continue. Several scholars have expressed their doubts (e.g. [37]) at various points in time, referring to physical limits (typically in relation to semiconductor developments) or the limits to information processing capabilities of us humans. Knudsen et al. estimate the upper bound of this capability between 50 Gbit/s and 2,800 Tbit/s. [38].

Obviously, extrapolation can should only be used on the short term. If one extrapolates data for a too long time frame, strange conclusion will be drawn. We found a research that extrapolates the growth of the internet and the growth in HDD and concludes that in 2050 the complete internet can be stored on one customer-of-the-shelve HDD [39].

Another classic example from 1894 is: “*In 50 years, every street in London will be buried under nine feet of manure*”. Who could have foreseen that in decades the combustion engine would replace manure with exhaust fumes?

### *Directions for future research*

The modelling approach seems robust and capable of predicting the future demand for fixed household internet bandwidth demand. It however assumes a quite ‘traditional’ model where consumers primarily use their home connection. Current and future developments in mobile networks (e.g. small cells, Wi-Fi and LTE interworking) as well as consolidation and integration of network owners will likely blur the line between wireless and fixed networks. In the future, we expect connectivity to be wireless-first for end users, and not necessarily tied to home connections. Future research should therefore focus on finding ways to change the unit of analysis from ‘home connection’ to the user itself.

Additionally we expect there to be more *tiers* in access networks, as certain applications (e.g. those related to smart mobility) will make use of edge computing. This makes the performance profile of an access network more opaque and greatly dependent on the architecture of the application. Future research should focus on connecting the tiered connectivity with the trade-off between computation and storage, which drives adoption of edge computing.

## References

- [1] Nielsen, J. (1998). *Nielsen's Law of Internet Bandwidth* [[www.nngroup.com](http://www.nngroup.com)] Nielsen Norman Group.
- [2] Norton, W.B. (2015). *Internet Transit Prices - Historical and Projected* [[drpeering.net](http://drpeering.net)] DrPeering.net.
- [3] Odlyzko, A.M. (2003). *Internet traffic growth: Sources and implications* [[www.dtc.umn.edu](http://www.dtc.umn.edu)] vol. II, Orlando, Florida: SPIE.
- [4] Cisco (2017). *Cisco Visual Networking Index: Forecast and Methodology, 2016–2021* [[www.cisco.com](http://www.cisco.com)]
- [5] Sandvine (2016). *Global Internet Phenomena: Latin America & North America* [[www.sandvine.com](http://www.sandvine.com)] Sandvine.
- [6] Moore, G. (1965). *Cramming more components onto integrated circuits* [[drive.google.com](http://drive.google.com)] vol. 38,
- [7] Brennenraedts, R., Vankan, A., te Velde, R., Minne, B., Veldkamp, J., and Kaashoek, B. (2014). *De impact van ICT op de Nederlandse economie* [[www.dialogic.nl](http://www.dialogic.nl)] Utrecht: Dialogic.
- [8] Taleb, N.N. (2007). *The Black Swan. The Impact of the Highly Improbable*
- [9] Rogers, E.M. (2010). *Diffusion of Innovations* Simon and Schuster.
- [10] IHS Markit; Point Topic (2016). *Broadband Coverage in Europe 2016. Mapping progress towards the coverage objectives of the Digital Agenda* [[ec.europa.eu](http://ec.europa.eu)] London: European Commission, DG Communications Networks, Content and Technology.
- [11] Cho, K., Fukuda, K., Esaki, H., and Kato, A. (2006). *The impact and implications of the growth in residential user-to-user traffic* [[www.fukuda-lab.org](http://www.fukuda-lab.org)] Pisa, Italy,
- [12] Gerber, A., and Doverspike, R. (2011). *Traffic Types and Growth in Backbone Networks* [[www.osapublishing.org](http://www.osapublishing.org)]
- [13] European Commission (2010). *European Broadband: investing in digitally driven growth* [[www.europarl.europa.eu](http://www.europarl.europa.eu)] Brussels: European Commission.
- [14] Wu, D., Hou, Y.T., Zhu, W., Zhang, Y., and Peha, J.M. (2001). *Streaming Video over the Internet: Approaches and Directions* [[repository.cmu.edu](http://repository.cmu.edu)] vol. 11,
- [15] van der Merwe, J., Sen, S., and Kalmanek, C. (2002). *Streaming Video Traffic : Characterization and Network Impact* [[www.cs.utah.edu](http://www.cs.utah.edu)]
- [16] MIT News (1996). *MIT Research - Brain Processing of Visual Information* [[news.mit.edu](http://news.mit.edu)] Cambridge: Massachusetts Institute of Technology.
- [17] van der Vorst, T., Brennenraedts, R., Driesse, M., and Bekkers, R. (2016). *Beyond Fast – How the speed of residential internet access will develop between now and 2022* Utrecht, the Netherlands: Dialogic.
- [18] van der Vorst, T., Brennenraedts, R., van Kerkhof, D., and Bekkers, R. (2014). *Fast forward: how the speed of the internet will develop between now and 2020* Utrecht, the Netherlands: Dialogic.
- [19] Cisco (2018). *VNI Forecast Highlights Tool* [[www.cisco.com](http://www.cisco.com)] Cisco.
- [20] International Telecommunications Union (2017). *Measuring the information society report 2017* [[www.itu.int](http://www.itu.int)] Geneva, Switzerland: ITU.
- [21] Sam Knows (2018). *The Whitebox: Keeping an eye on your internet performance.* [[www.samknows.com](http://www.samknows.com)]

- [22]Maier, G., Feldmann, A., Paxson, V., and Allman, M. (2009). *On Dominant Characteristics of Residential Broadband Internet Traffic* [[www.icir.net](http://www.icir.net)] Chicago: ACM.
- [23]Crozier, R. (2013). *Networks hammered by Apple iOS7 release* [[www.itnews.com.au](http://www.itnews.com.au)] ITNews.
- [24]Akamai, (2018). *state of the internet] / security. Carrier insights report (Spring 2018)* [[www.akamai.com](http://www.akamai.com)] Akamai.
- [25](2013). *Report: Bot traffic is up to 61.5% of all website traffic* [[www.incapsula.com](http://www.incapsula.com)] Redwood Shores, CA: Imperva Incapsula.
- [26]Van Buskirk, R. (2012). *An Adoption Curve Fitting Method for Estimating Market Efficiency Improvement and Acceleration* [[www.researchgate.net](http://www.researchgate.net)]
- [27]Eurostat (2017). *Digital economy and society statistics - households and individuals* [[ec.europa.eu](http://ec.europa.eu)] Eurostat.
- [28]Kenny, R., and Broughton, T. (2013). *Domestic demand for bandwidth: An approach to forecasting requirements for the period 2013-2023* [[www.broadbanduk.org](http://www.broadbanduk.org)] Broadband Stakeholder Group.
- [29]Karagiannis, T., Broido, A., Brownlee, N., Claffy, K., and Faloutsos, M. (2004). *File-sharing in the Internet: A characterization of P2P traffic in the backbone.* [[citeseerx.ist.psu.edu](http://citeseerx.ist.psu.edu)]
- [30]Strategy Analytics (2016). *1 in 8 homes in North America to own an Ultra HD TV by end 2016* [[www.strategyanalytics.com](http://www.strategyanalytics.com)] Strategy Analytics.
- [31]AMS-IX (2001-2018). *Historical monthly traffic volume* [[ams-ix.net](http://ams-ix.net)] Amsterdam: AMS-IX.
- [32]Poort, J., Leenheer, J., van der Ham, J., and Dumitru, C. (2013). *Baywatch: two approaches to measure the effects of blocking access to The Pirate Bay* [[hdl.handle.net](http://hdl.handle.net)] Amsterdam/Tilburg: Institute for Information Law, University of Amsterdam/CentERdata, Tilburg University/System and Network Engineering, University of Amsterdam.
- [33]Telecompaper (2015). *GfK: een derde gebruikers betaalt voor streaming muziek* [[www.telecompaper.com](http://www.telecompaper.com)]
- [34]GfK (2015). *Ruim 5 miljoen Nederlanders gebruiken Video on Demand diensten* [[www.gfk.com](http://www.gfk.com)]
- [35]Sandvine (2014). *Global Internet Phenomena Report, 1H 2014* [[www.sandvine.com](http://www.sandvine.com)] Sandvine.
- [36]Ingham, T. (2015). *Music piracy has been 'virtually eliminated' in Norway* [[www.musicbusinessworldwide.com](http://www.musicbusinessworldwide.com)] MusicBusiness Worldwide.
- [37]Ayres, R. (2006). *Turning point: The end of exponential growth?* [[pure.iiasa.ac.at](http://pure.iiasa.ac.at)] vol. 73, pp. 1188-1203.
- [38]Knudsen, T.P., Patel, A., Pedersen, J.M., and Madsen, O.B. (2006). *A Method for Upper Bounding Long Term Growth of Network Access Speed* [[www.iiisci.org](http://www.iiisci.org)] vol. 4, pp. 8-11.
- [39]Trotman, A., and Zhang, J. (2013). *Future Web Growth and its Consequences for Web Search Architectures* [[arxiv.org](http://arxiv.org)]