

Are there limits to growth in data traffic?: On time use, data generation and speed

Mike Hazas, Janine Morley, Oliver Bates, Adrian Friday
Lancaster University
Bailrigg, LA1 4YW, United Kingdom
{m.hazas,j.morley,o.bates,a.friday}@lancaster.ac.uk

ABSTRACT

This discussion paper considers the nature of growth in data traffic across the Internet, as a basis for asking whether and how such growth might slow down or otherwise be limited. Over the last decade, data growth has been dramatic, and forecasts predict a similar ongoing pattern. Since this is associated with increasing electricity consumption, such a trend is significant to global efforts to reduce carbon emissions. In this paper, we selectively explore aspects of data growth that are linked to everyday practices and the way they draw upon and generate Internet data. We suggest that such growth does have some conceivable limits. However, the nature of ‘Internet use’ is changing and forms of growth are emerging that are more disconnected from human activity and time-use. This suggests that although there may well be limits, in principle, to some forms of growth, total data traffic seems likely to continue growing. This calls for careful attention to the nature of the trends involved, as a basis for intentionally building limits into this system before levels of Internet electricity demand becomes directly and more explicitly problematic.

CCS Concepts

•Social and professional topics → Sustainability; •Human-centered computing → Interaction design;

Keywords

Information infrastructures; social practice

1. INTRODUCTION

In a recent article in Low-Tech Magazine, de Decker [5] argues that *“there are no limits to growth when it comes to the Internet, except for the energy supply itself”* and so, uniquely, *“the energy use of the Internet can only stop growing when energy sources run out, unless we impose self-chosen limits”*. Other energy-demanding systems, he suggests, encounter inherent constraints which limit their growth:

for example, the weight and size of cars can only increase so much if they are still to remain compatible with infrastructures designed for them; infrastructures in which speed limits are imposed for safety reasons. Similarly, he suggests that some kind of ‘speed limit’ for the Internet is needed.

Putting to one side consideration of how such limits might be achieved, it is intriguing to examine the claim that the energy used by the Internet will continue to grow until the availability of energy itself becomes problematic, that is, unless some other kind of checks or limits to growth are imposed first. This is a rather radical, fascinating and, in so far as it is plausible, troubling claim.

Currently, the direct energy used to power the Internet and to produce, transport and dispose of its components (embodied energy) are relatively small contributors to global emissions in comparison to, say, global transport [13]. Current estimates suggest that operation of the Internet (powering devices, networks and data centres) amounts to around 5% of global electricity use; yet this is growing faster (at 7% per year) than total global electricity consumption (3% per year) [19]. In other words, the Internet is consuming an increasing portion of global electricity supply. In the context of changing energy systems that include greater renewable sources, and new forms of electricity demand such as electric vehicles, the growing portion of global electricity required to run the Internet may become increasingly significant in efforts to balance supply and demand, reduce carbon emissions and, as such, become potentially subject to energy-related limits as de Decker suggests. Yet given the radically distributed, and largely ‘invisible’ nature of this energy consumption, how large a portion of global electricity could this represent before such limits might be imposed? Some predictions suggest that production and use of information and communication technologies might grow to around 20% of global supply by 2030, or as much as 50% in a worst case scenario [2].

In this short discussion paper, we consider trajectories of growth in the energy that powers the Internet, not through modelling, but through a more conceptual exploration of some of the trends currently and historically associated with such growth. Specifically, we ask whether or not these trends may ultimately be limited in themselves, and thereby may help to slow and limit growth in the future, at least in principle. This is a complex question, so we frame and focus our exploration in two ways.

First, we focus on volumes data traffic as an indicator of operational energy use. This is by no means a direct or simple relationship but in broad terms *“the very substantial*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

LIMITS '16, June 08-10, 2016, Irvine, CA, USA

© 2016 ACM. ISBN 978-1-4503-4260-5/16/06...\$15.00

DOI: <http://dx.doi.org/10.1145/2926676.2926690>

gains from energy efficiency improvements have been more than offset by increased consumption of services” [12, p. 583]. Thus, as flows of data traffic have increased so too has energy use, albeit at a lesser rate.

Second, we focus on the link between the use and generation of data and activities in which people are more and less directly involved. We explore the idea that, to date, human time and attention has been related to the growth in data traffic, in various ways. We ask whether this might represent some limits to ongoing growth, and whether other forms of growth are emerging and taking over.

We start by characterising the rate and composition of data growth. We then consider aspects of this growth related to activities and time spent ‘online’, and ask whether such trends will continue indefinitely or, at some point, represent some kind of check to growing data traffic. Whilst growing levels of access to the Internet, a greater range of ‘online’ activities and services migrating and emerging, and increases in time spent online continue to be relevant to the growth in data, there is, in principle, a limit to these forms of growth: the global population is finite (though clearly growing), and the hours in the day available to each person are also finite. We then turn to consider how the dynamics of data growth are changing: the relationship to attention and time-use is becoming less direct as data intensities of certain services increase and as ‘background’ connectivity and connected things become more prevalent.

2. GROWING DATA, GROWING ENERGY USE

Measures of actual Internet traffic volume and composition tend to be partial, limited to certain countries or narrow periods of time. Based on a “representative cross-section” of service providers, who volunteer to allow their aggregate data to be reported, Sandvine [17] report that the average monthly per-subscriber traffic volume on fixed lines like broadband and fibre increased by about 50% in North America, and 170% in Europe over 2013 and 2014 (Figure 1). According to Ofcom, the UK telecommunications regulator, UK home broadband data volumes also grew markedly: monthly average traffic volume rose from 17 GB in 2011, to 82 GB in 2015 [11].

Taking just one survey point, public aggregate input/output data statistics published for January each year 2002–2016 from the Amsterdam Internet Exchange (Figure 2), we see a clear pattern of year on year growth of between 20–140% (20–40% each year in the last 5 years). Crudely, this data demand has an equivalent direct energy cost as data is transmitted and processed. Despite step changes in energy efficiency as new technology is introduced, this could arguably be offset by innovations in the marketplace such as increased expectations around high-definition video (e.g. 1080p cameras on smartphones, 3DTV and 4K video). Holistically, this pattern of growth also causes increases in embodied energy (energy [13]) of the Internet as equipment needs to be replaced to cope with the continued increase in data demand, and its associated demands for processing and storage.

For mobile access, data volumes are generally smaller, but have consistently fast growth, doubling every few years according to Sandvine (Figure 1). In fact, Ericsson [6] report a 60% growth in global mobile cellular data in recent years, and the forecasts are for 45% or more annual growth through

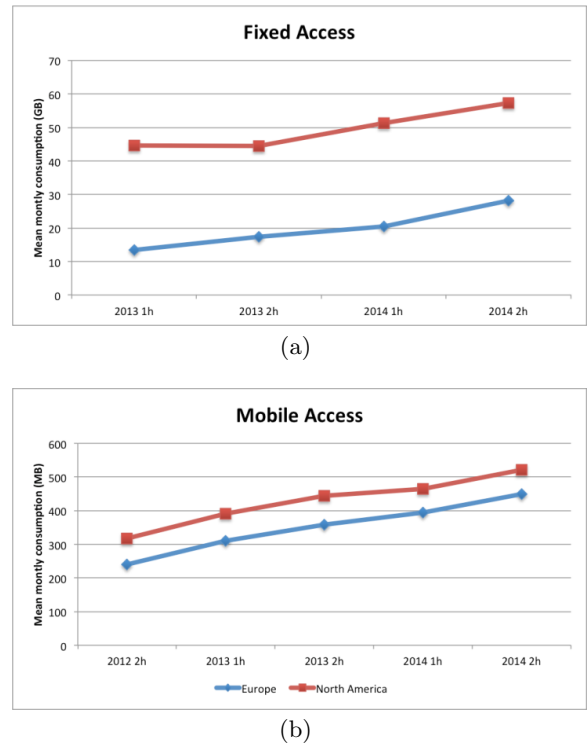


Figure 1: Growth in fixed and mobile traffic volumes. Data source: Sandvine reports 2012–2014.

2021 [4, 6].

Several forms of growth are reflected in the overall increases in mobile traffic volume: a) per-subscriber demand, which is b) compounded by the increasing number of tablet and smartphone handsets connected to mobile cellular data services, which is c) related to, but not entirely explained by, the proportion of the global population who have some form of Internet access.

On the face of it, the latter form of growth has some limits, even though these may be many years away: broadband services currently only reach around 30% of homes in most developed countries, and only 43% of the world population is using the Internet [8]. There are biophysical limits to human perception and attention, which ostensibly could place an upper bound on the fidelity of digital media we deem ‘sufficient’. Some limits are indeed already visible: in economically developed nations, mobile phones are reaching saturation. For example, the number of mobile subscriptions already exceeds the population in the US, Scandinavia and Australia. And while by 2020 over 3.1 billion mobile subscribers will have access to LTE and 4G communications networks worldwide, this falls far short of world population.

3. CURRENT DATA GROWTH: DEVICES PRACTICES AND TIME-USE

Alongside questions of population and ‘saturation’, it is important to consider what all this Internet traffic is for - what activities it seems to be intertwined with. As Røpke points out, “People are practitioners who indirectly, through the performance of various practices, draw on resources” [15].

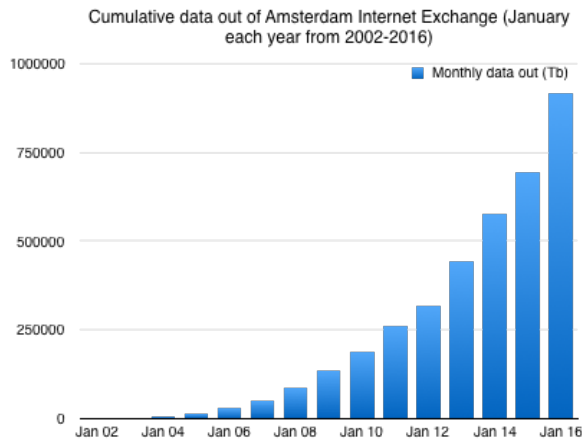


Figure 2: Monthly data from Amsterdam Internet Exchange aggregated for each January 2002-2016.

Communications enabled devices increasingly augment everyday practices that previously would have not have required or even benefited from, Internet connectivity. In our recent study of the communication impacts of applications of mobile devices in everyday life [9], we found practices such as exercise regimes to be augmented by social media applications; and migration of once offline or broadcast viewing and listening of TV and radio to media rich communication and streaming media platforms.

The personal ownership of these devices, coupled with the convenience of accessing the Internet and their mobility, enable communication to fill and even expand small pockets of ‘dead time’. Along with increased multitasking supported by digital technologies “enabled by the partial decoupling of many practices from previous time and space constraints through the use of ICT, contribute to a more densely packed everyday life” [16, p. 356].

In other words, because Internet connectivity is being integrated into a huge variety of practices which may take place throughout the day, time ‘spent online’ attending to digital services is growing. Indeed, in 2005 in the UK an average of 9.9 hours was spent online in a “typical week” across home, work and elsewhere. By the end of 2014 this average had grown to 20.5 hours per week [10, p. 28].

Yet, not all ‘online’ activities (those that require an Internet connection) are equal. Firstly, significant differences emerge between traffic types on domestic and mobile Internet networks. Drawing on Sandvine, the category of “real-time entertainment”, which is primarily video, accounts for a large portion of total Internet traffic, at over a third of peak period traffic. This is followed by web browsing, which might support a wide variety of activities. Social networking is also significant, particularly on mobile networks. Secondly, these activities ‘take-up’ time in different ways, with some clearly taking more time than others.

3.1 Watching

Some of the growth in time which is reportedly spent online can be attributed to a daily increase in the number of Internet users who watch TV or films online (10%–27% between 2007–2014) [10]. The number of Internet users who watch short video content (i.e. video clips) has almost doubled in the same amount of time, growing from 21% to 39%




Figure 3: Composition of peak period traffic in Europe in the second half of 2014. Redrawn from Sandvine.

in 2014.

Ericsson’s February 2016 report sums up why mobile video is traffic is growing at such a rate, due to: larger screens enabling better quality streams; video content increasingly appearing as part of other applications (e.g. social networking, news, advertising); growth of video streaming (50–70% of video traffic for some mobile networks is YouTube); growth in uptake of video on demand; changes in where and where video is consumed and faster infrastructure.

3.2 Other online services

Whilst the number of people who are using online services to consume media content is growing, it is worth noting that the amount of time spent using these services is also growing. In the US, the average time spent per day paying attention to digital activities (including digital video, social networks, digital radio, Facebook, Pandora) has increased from 226–364 minutes per day between 2011–2015 [20]. The time spent online watching digital video has grown from 39–115 minutes, with growth also observed for social networks (71–104 minutes), and digital radio (53–65 minutes) [20].

3.3 Concurrent uses of the Internet

In 2014 additional connected devices are being used to go online compared to 2009, including Smart TVs, e-book readers and wearable devices, with all technologies (computers, smartphones, tablets, games consoles) other than “portable media players” increasing in their use for going online [10, Fig. 34]. It’s worth mentioning that multitasking that takes place whilst watching TV is performed by 53% of UK adults [14, p. xi]. This multitasking is encouraged by “living room connected devices” (e.g. devices in the living room that are connected to the Internet) seen to “blur the line between passive and active entertainment” [14, p. 4]. Whilst we are currently unaware of the overlap between watching on living room connected devices and primary viewing devices (e.g. Smart TVs, laptops, consoles) it is worth acknowledging that living room connected devices are likely to increase time spent online as they are predicted to increase from 114 million to 267 million units shipped worldwide by 2017 [14, p. 5].

4. FUTURE DATA DEMAND: INTENSITY, BACKGROUND AND NON-HUMANS

So far we have talked about ‘attention-connected data demand’, that is, exchanges in data that occur as the (more or less) direct result of what people do by paying attention to online services, reflecting the presence and significance of media rich communication in everyday practices. As can be appreciated from Figure 3, much of the data traffic during peak times, currently seems to be associated with such activities (e.g. entertainment, web browsing, social networking). Although such traffic may continue to grow as a result of increases in the data-intensity of these activities, such as through higher definition video content or through new entertainment services requiring complex forms of cloud computing such as virtual reality and cloud gaming, there is some limit to the time that is available to ‘invest’ in them, and thus some ‘check’ to growth in these types of data demand.

Yet not all traffic is so directly associated with attention-linked services, and thus with patterns of time-use. One already automated form of demand on the Internet is through software updates. These are currently about 6% of download traffic, or perhaps up to 10% if computer game downloads and updates are included (on marketplaces such as Steam, the PlayStation Store, and Xbox Live) [17].

In addition, activities are at times accompanied by ‘unintentional’ or ‘background’ data demand. In recent studies of mobile device use, we found unexpectedly high levels of communication between apps and the cloud when specific applications were not in active use, for instance ensuring that applications and operating system services were up to date (900Mb or 5% of their overall traffic for one participant’s iPad), backup of application data and digital photos to the cloud (2.25Gb/week or 71% of data demand cloud syncing of photos and videos in one case). These levels of data traffic, were both unobserved, uncoordinated, and largely unmanaged by or even unmanageable for most of our participants [3, ch. 6], [9]. While these studies are arguably with small and isolated populations, the very unremarkable nature of these findings coupled with the vast numbers of similar devices in common use, suggests a very significant data demand that is difficult to manage and limit at scale. As the dominant paradigm for software development, tool chains, and indeed the very business models that drive the mobile eco-system push the design of software as thin clients to powerful cloud backend services, we will doubtless witness these kinds of data patterns repeated many-fold. Most significant perhaps is the ongoing development of machine to machine communication enabled by the so called ‘Internet of Things’ (IoT). What about the data and energy impacts of smart ‘things’, in homes, workplaces and civic infrastructures? This introduces another form of growth that is more dissociated from the limits associated growth in direct forms of Internet ‘use’ that have to date been so significant. Connecting ‘things’ or ‘machines’ to the Internet changes the connection between what people do and exchanges of data. This communication will occur transparently, without observation or interaction, and potentially without limit. At the time of writing, the existing 6.4bn connected IoT devices is only slightly less than world population (86%), but market predictions suggest this will reach 21bn devices by 2020—roughly three times world population estimates. Some predictions put machine to machine communication as 45% of the whole Internet traffic by 2022 [7], with additional re-

liance on data upload [18] and processing in the cloud [1].

While IoT devices are presumably assumed to be low communications footprint devices such as smart meters, or smart home thermostats that communicate sporadic values to backend services; others like self-driving cars, and remotely monitored wireless cameras and wearable medical devices will be highly data intensive, and require matching facilities in communications networks and data centres to offer timely and responsive communication, computation and storage. Some see IoT as a potentially ‘net negative’ contributor to energy demand due to the efficiency gains made through better measurement and control. We merely observe, that these new facilities have the potential to be adopted into everyday practice, to raise expectations and lower the burden of creating additional data and electricity demand and have an initial and ongoing embodied energy cost, by way of rebound effect—in exactly the same way that smart devices, media and broadcast are already doing. What seems clear, is that the impact of IoT is unknown or unknowable at this point. Algorithms that create/consume data can carry on doing so, regardless of who is paying attention to them, and even less clear are the limits and braking functions that will ensure that such systems will operate within sustainable limits.

5. CONCLUSION

This returns us to the question of whether there is an inherent limit to Internet traffic growth (other than energy)? On the one hand, yes. Some of the reasons for current and past growth will themselves be limited at some point in the future; there is after all a finite, if growing, population on the planet and only a certain number of devices that one might *explicitly* use to access Internet services. Even if we assume that the current access patterns in the most ‘connected’ societies are replicated worldwide; the full panoply of second screen, higher definition, and on-demand over 4G and higher, there is plausibly a limit to the hours in the day and content that can be actively engaged with. On the other hand, ‘the data under the hood’ is growing unwatched and unabated. The automated updates, cloud syncing, off-loading of storage and computation to the cloud, that are an increasing feature of the design of applications, and endemic to the tools and pervading technological culture that is bringing these about.

Further, the Internet of Things is set to trigger a whirlwind of investment and connected infrastructure growth that has the massive potential to grow operational electricity use and emergy of the Internet. Despite sometimes questionable benefits and motivations, the IoT is currently under construction, in many different ways. This raises key questions as to whether, and what kinds, of limits there may be to potentially self-generating cycles of data generation, processing and circulation within such an Internet. If such cycles are largely automated and operate at remove from the time-limits associated with human activity then, potentially, any ‘inherent’ limits to the growth in Internet traffic will fall a long way short in the future.

Yet this Internet of Things is still in-the-making, and such limits (or lack of them) are not yet ‘inherent’. Thus, as the Internet continues to develop, de Decker’s proposal of some kind of speed limit, that might be built into this system, shaping its development for years to come, is an important proposition to consider; especially in comparison to an alter-

native perspective of making retrospective reductions in Internet traffic in the future. It is far from clear how such limits could and should be formulated and enforced: should ‘unlimited’ data tariffs be replaced with volume quotas or differential pricing for services of various ‘importance’? Should micro-payment schemes incentivise use of more bandwidth frugal and offline media?

It is clear, however, that the dynamics of Internet traffic growth are changing. In many ways the technical capacities of data infrastructures, such as broadband networks, continue to limit to the data flows and services that they also make possible. Specifically, many broadband and mobile contracts are limited and traffic is managed. In many parts of the world, including areas in developed countries, these infrastructural capacities are increasingly experienced as tangible limitations to accessing a range of services that are taken for granted elsewhere. To the extent that such limits are experienced directly as limits to activity, then there are important issues of social equity that cannot be ignored. But where differences in data demand are less ‘visible’ and less connected to explicit forms of ‘using the Internet’ other possibilities for limiting data intensities might be explored in ways that do not impact and consolidate existing inequalities of access.

Our community is well placed to help shape these debates and possible futures. By measuring and understanding the holistic value and impacts of such systems, and understanding the varied ways in which they change, can we help to bring about a future where these systems operate within similar and sufficient limits?

6. ACKNOWLEDGMENTS

We would like to give our thanks to our colleague Kelly Widdicks, for her compilation and analysis of the Sandvine reports from 2012 to 2015.

7. REFERENCES

- [1] A. Al-Fuqaha, M. Guizani, M. Mohammadi, M. Aledhari, and M. Ayyash. Internet of things: A survey on enabling technologies, protocols, and applications. *Communications Surveys & Tutorials, IEEE*, 17(4):2347–2376, 2015.
- [2] A. S. Andrae and T. Edler. On global electricity usage of communication technology: trends to 2030. *Challenges*, 6(1):117–157, 2015.
- [3] O. E. G. Bates. *The impacts of domestic media and ICT: a study of digital technology, energy consumption, data demand and everyday practice*. PhD thesis, 2015.
- [4] Cisco. Cisco visual networking index: Forecast and methodology, 2012–2017. Technical report, May 29, 2013.
- [5] K. de Decker. Why we need a speed limit for the internet. Technical report, Low-tech Magazine, 2015.
- [6] Ericsson. Ericsson mobility report: On the pulse of the networked society (mwc edition). Technical report, Ericsson, Feb 2015.
- [7] J. Gantz and D. Reinsel. The digital universe in 2020: Big data, bigger digital shadows, and biggest growth in the far east. *IDC iView: IDC Analyze the future*, 2007:1–16, 2012.
- [8] ITU. The world in 2015. Technical report, International Telecommunications Union, 2015.
- [9] C. Lord, M. Hazas, A. K. Clear, O. Bates, J. Morley, and A. Friday. Demand in my pocket: mobile devices and the data connectivity marshalled in support of everyday practice. In *Proc. CHI*, 2015.
- [10] Ofcom. Adults’ media use and attitudes report. Technical report, Ofcom, 2014.
- [11] Ofcom. Communications market report. Technical report, Ofcom, 2014.
- [12] C. Preist and P. Shabajee. Energy use in the media cloud: Behaviour change, or technofix? In *Cloud Computing Technology and Science (CloudCom), 2010 IEEE Second International Conference on*, pages 581–586. IEEE, 2010.
- [13] B. Raghavan and J. Ma. The energy and emergy of the internet. In *Proceedings of the 10th ACM Workshop on Hot Topics in Networks, HotNets-X*, pages 9:1–9:6, New York, NY, USA, 2011. ACM.
- [14] N. Robinson, J. Freeman, J. Gaspers, V. Horvath, T. Hellgren, and A. Hull. Living room connected devices. 2014.
- [15] I. Røpke. Theories of practice —new inspiration for ecological economic studies on consumption. *Ecological Economics*, 68(10):2490–2497, 8 2009.
- [16] I. Røpke and T. H. Christensen. Energy impacts of ICT – insights from an everyday life perspective. *Telematics and Informatics*, 29(4), 2012.
- [17] Sandvine. Global internet phenomenon report. Technical report, Sandvine Incorporated ULC, 1h 2014.
- [18] M. Z. Shafiq, L. Ji, A. X. Liu, J. Pang, and J. Wang. A first look at cellular machine-to-machine traffic: large scale measurement and characterization. *ACM SIGMETRICS Performance Evaluation Review*, 40(1):65–76, 2012.
- [19] W. Van Heddeghem, S. Lambert, B. Lannoo, D. Colle, M. Pickavet, and P. Demeester. Trends in worldwide ict electricity consumption from 2007 to 2012. *Computer Communications*, 50:64–76, 2014.
- [20] A. Walgrove. The explosive growth of online video, in 5 charts. Technical report, 2015.