

# A Review of Top-Down Models of Internet Network Energy Intensity

Daniel Schien, Chris Preist  
Department of Computer Science  
University of Bristol  
Bristol, UK  
{schien,cpreist@cs.bris.ac.uk}

**Abstract**— Assessments of digital services wishing to take into account an energy footprint of the Internet typically require models of the energy intensity of the Internet. A common approach to the modeling of the energy intensity is to combine estimates of market surveys of installed network devices on a national or global scale and their related power consumption with the total data volume transported at the same scale. These are commonly referred to as top-down models. In order to be applicable, the system boundaries applied in the modeling of energy intensity need to match the system that delivers the particular digital service. The system boundaries applied in past studies vary considerably, resulting in increased uncertainty in the assessment results. In this text, we firstly reconsider the components of the network relevant for the delivery of general public digital services that should be included in an energy footprint. We then reinterpret the existing models in the light of those altered system boundaries. Secondly, we review existing assumptions about annual reduction of energy intensity and consider the effect of uncertainty in these on the overall estimates of energy intensity. We find that the energy intensity of the general Internet network in the year 2014 varies between 0.05 and 0.6 kWh/GB when top-down modeling is applied which is considerably lower than previous estimates had indicated.

**Index Terms**—Energy intensity, Internet, Networks, Energy Footprint

## I. Introduction

The continuing proliferation of digital services such as streaming videos, browsing websites or generally exchanging data over the Internet has attracted some attention to their environmental impact. This is often considered in terms of their direct and indirect impacts [1]: direct impacts being the negative environmental impacts of the infrastructure providing the service, and the indirect impacts being the (negative or positive) impacts of what the service enables to take place in the wider system. To understand the tradeoff between direct and indirect impacts requires a detailed modeling of the direct impacts of a service. This allows comparison to take place with alternatives, such as electronic vs paper-based news delivery

[2], CDs vs music streaming [3] and postal vs online software delivery [4]. It also allows alternative IT architectures for delivery to be compared, and so allowing ‘design for environment’ of digital services [5].

The direct environmental impact of digital services results from (i) the energy consumption by the electronic devices involved in the production, delivery and access to services and (ii) from resource use, energy consumption and waste emissions from the manufacture and disposal of the devices involved. Depending on the electronic device type and the environmental impact class considered, then either use phase energy consumption or impact from manufacturing can be the dominating factor. However, as a result of the complexity of information technology systems at all scales, significant uncertainty remains with regard to the direct impacts; both the impact from manufacturing of devices [6] and the use phase energy consumption [7]. Regarding the latter, significant uncertainty remains on the energy consumption by the network devices in the Internet, as Coroama and Hilty find [8]. They distinguish between three types of assessment approaches: bottom-up, top-down and model based. They find that top-down approaches estimate the use phase impact to be significantly higher than bottom-up or model based approaches have suggested. Two main reasons are provided as explanation for this disagreement: varying system boundaries and varying year of reference for assumptions.

Given the strong impact of the choice of system boundaries, an opportunity exists to reduce uncertainty in the past top-down models and contribute to a more accurate estimate of use phase impact of energy consumption.

In this paper, we review the most relevant top-down models and their system boundaries, and revise these models based on our analysis. The paper is structured as follows. We begin with a review of the main components that form the Internet and work out the system boundaries accordingly. In the following section, we review the most important top-down models and discuss discrepancy in the system boundaries. In section three, we propose changes to the system boundaries and alternate extrapolations of efficiency improvements over time that we then evaluate in order to present revised numerical estimates for energy intensity in section four. We close with a discussion of the consequences of these changes on existing results.

### A. Structure of the Internet Network

Energy and environmental footprints of digital services have mostly been studied for consumer products, as opposed to business-to-business services. In this context, a user device such as a PC or handheld device such as a tablet or laptop is connected via the Internet to a server in a data center as illustrated in Figure 1. The Internet is a so-called network of networks: several organizations operate independent, hierarchical tiers of networks that are connected with each other and pass traffic on, in order to establish connectivity between end-points.

The user operates customer premises network equipment (CPE) such as a wireless router or a modem that is connected to the network of the broadband provider. Residential network connectivity is predominantly provided by wired networks in the form of Digital Subscriber Line (DSL), Hybrid Fiber-Coaxial (“cable”) or a fiber optic variant (e.g. fiber to the curb – FTTC or fiber to the home - FTTH) between a residential node at the customer side and a terminal in the access network. This network link is often called local loop or last mile. The specific device type of the terminal varies with each type of network. In the case of DSL, a modem in the home modulates a phone and data signal on to the same line at one end and a Digital Subscriber Line Access Multiplexer at the other end

detail on the structure of metro and long haul networks, please refer to [9].

The destination of the client connection is usually a server computer (except for peer-to-peer traffic which is destined to another user client) that most typically is located in data center facilities. The organization that commissions the servers will additionally operate a local area network (LAN) at campus or building level.

### B. Definition of Energy Intensity

In order to estimate the energy footprint of a digital service, a share of the energy consumption by the devices in the Internet must be apportioned to the digital service, which is commonly referred to as allocation. For public network devices, allocation by transferred data volume over time is useful as it forms the limiting factor for the use of the device [7]. Energy intensity, most commonly used in previous studies, denotes the energy consumption per bit of data transfer and is usually stated in (Kilo)watt-hours per Gigabyte or Joule per bit. During an assessment the data volume for a particular service can be multiplied with the energy intensity to give the energy footprint of the service.

Two modeling approaches have been used in the past to

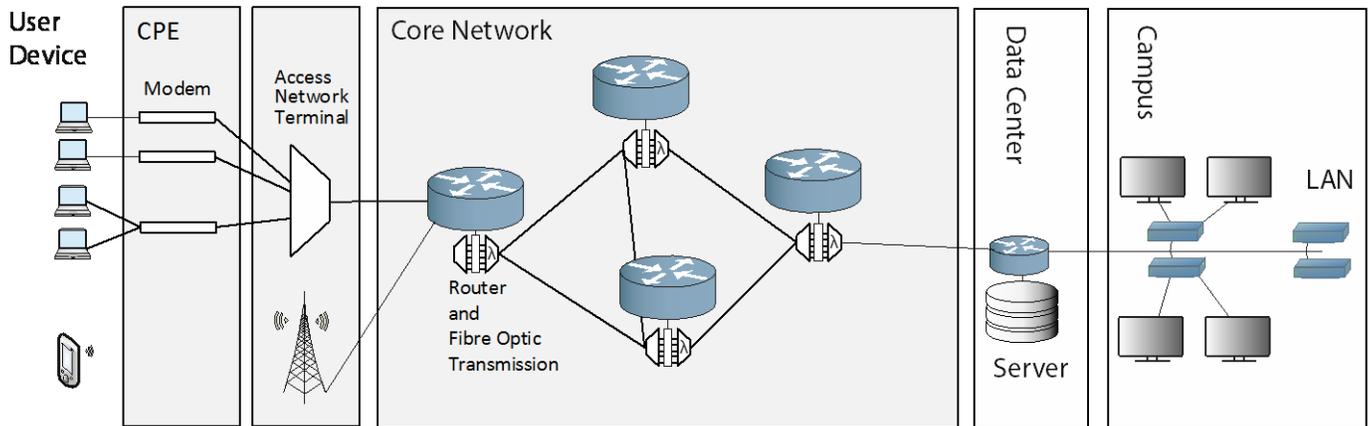


Fig. 1 – Structure of the network device types in layers of the Internet as distinguished in top-down models.

separates the phone and the data signal and passes the latter on towards the public Internet. Similarly, a cable modem combines TV and data signal in the home and is separated into individual streams by the Cable Modem Termination System (CMTS). The terminal systems form the access network. Here data traffic from multiple users is aggregated into a single line that connects to the next higher network tier, the core network, usually distinguished between metro and long haul networks. Mobile networks (3G and LTE) provide connection via a base station that connects to the higher network tiers. In the core network routers direct traffic between these multiple, tiered networks towards the destination. While the purpose of the routers is to determine paths through the Internet, the physical connection between devices is established by fiber optic links, also called the transmission system. These commonly use dense wavelength division multiplexing (DWDM). For more

estimate the energy intensity: bottom-up and top-down. The bottom-up approach estimates energy intensity per network device class (such as a WAN router or transmission links) as the ratio of power consumption over data throughput and sums this up over all network devices in the end-to-end connection.

A top down model, instead, relates energy consumption of the entire population of network devices to the total data volume. It assumes average values of power draw. Top-down models explicitly set the aggregate data volume transferred while bottom-up models treat data volume implicitly by modeling data flow. The accuracy of either approach, that is, how close the estimate energy intensity is to the “real” value can be characterized by the accuracy of the assumptions used in either model. In case of the bottom-up model, this mainly refers to assumptions of average power consumption and data throughput per device class and the number of such devices in

the end-to-end connection. In the case of the top-down model, the total energy consumption of all network devices and the total data volume in the network are estimated and affect the accuracy of the energy intensity. Top-down models, in principle, represent the entire population and thus should be based on a large sample size. This can add to the robustness of their results.

### C. System Boundaries for Digital Media Assessments

In past assessments of digital services, energy consumption of user devices, the network and data center infrastructure was accounted for separately. Besides incommensurable user device modalities (e.g. the use e-readers for magazine reading as opposed to PCs for file download) and differing allocation approaches [7], a reason is the practicality of modeling network energy consumption. Because publicly available primary data on the network is scarce and its structure complicated (compared to user devices), assessments have usually sourced ready-made estimates of network energy intensity from secondary sources. For similar reasons, data centers are treated separately. Additionally, in some cases primary data center energy consumption is available for a digital service assessment. Given that the Internet is a network of networks each of which is operated by an independent organization, the collection of primary data is especially difficult.

Meanwhile, it is useful to account for service use over the various types of access network (DSL, Cable, FTTx, 3G, LTE) separately from the public Internet [5]. Thus, the energy intensity of the public Internet includes the routers and transmission devices in the metro and long haul network. Finally, network devices in campuses are not involved in the connection between the consumer of a digital service and the data center providing it and it should not be included in the energy intensity. For supply chain organization of the digital services, they should be accounted for as part of the building infrastructure operation.

Last but not least, an inventory of server totals in operation based on market data will include devices that are used for the delivery of digital services but also devices used for general business information processing. Energy consumption for these servers should be accounted for separately from the direct energy demand during consumption of digital services.

## II. Literature Review

### A. Inventory of Energy Consumption

As was mentioned above, the accuracy of the top-down model depends on the estimate of total energy consumption and total data volume. At present, no comprehensive inventory of network devices currently in operation exists, thus past estimates of the total energy consumption by network devices were based on market sales data for a sample of devices per device class, to which average estimates of device type power consumption was added. The robustness of such an extrapolation generally depends on its statistical viability and is subject to selection bias. Typically, the greater the sample size, the more robust the study will be.

The most comprehensive estimate of the total energy consumption by network devices was performed by Roth and colleagues [10] for the US. Estimating the energy consumption by office and telecommunications equipment in commercial buildings, the authors estimate total energy consumption for several device classes, including PCs, Server Computers, Printers and network equipment by multiplying sales data per device type and the respective annual energy consumption, taking into account varying power consumption per usage mode (such as stand-by idle and in-use). They provide estimates for CMTS, WAN switches, routers, hubs and LAN switches and also fiber transmission terminals. Stock estimates are based on market surveys and expert judgment.

Uncertainty around the estimates of installed device stock and power consumption increases with the age of the study. The study was published in 2002, while some of the referenced market surveys date back to 1998. Secondly, the sample size of power measurements for some devices is small. For the estimate of power consumption they combine 6 sample devices of entry-level and Internet Service Provider (ISP) class devices into an estimate based on expert opinion of market proliferation.

However, it remains the most comprehensive account that sources market studies and breaks up energy consumption per device class. A more recent estimate of the US and worldwide energy consumption by data network equipment by Lanzisera and colleagues [11] provides a segmented estimate of routers and switches but does not include optical transmission equipment. Also, their segmentation is not identical with that applied by Roth and colleagues. Nonetheless, if the categories WAN switches and Routers from Roth et al. are aggregated and extrapolated to 2008, with a 14% compound annual growth rate (CAGR), a total energy annual consumption of 3.57TWh results. That is remarkably close to sum of the “Modular Core Switches & 10G Switches” and “Large Routers” and “Small and Medium Routers” in the account by Lanzisera et al. of the same year with 3TWh.

### B. Top-Down Models

In the review by Coroama and Hilty [8] five top-down studies are listed. Out of these, the study by Lanzisera [11] is not a top-down model of energy intensity by itself as it does not include an estimate of network traffic volumes. A second study [12] does not present a model of energy intensity but only the final numerical value for energy intensity and worldwide energy consumption, the calculation of which is not explained and only an exemplary source, dating to 2003, for its input is given. Coroama and Hilty combine the study’s power consumption value with an estimate of Internet traffic from Cisco’s VNI (visual networking index) and yield a value for energy intensity of 1.8kWh/GB for 2008. Given the lack of transparency on the contents of the inventory data and the uncertainty regarding the reference year, we have chosen to ignore this study and others that are similarly coarse, such as [13].

The remaining three studies combined the estimate of the annual energy consumption by Roth and colleagues together with an estimate of the data volume transferred during the same

time to give an estimate of energy intensity [14], [15], [16]. Of these, the most frequently cited in energy footprinting was [16] by Taylor and Koomey 2008; for example in [2], [3]. This is likely to be the result of a significantly more detailed explanation and discussion of its assumption. The study provides an estimate for the energy efficiency of the Internet in order to assess the carbon emissions of a typical ad campaign. Identical to [14], Taylor and Koomey apply estimates of energy consumption in the network by Roth et al. but in contrast to the former also include servers and storage devices. The model is conceptually similar to an earlier study by the same authors [15] from 2004. The model differs in regards to the source of estimates for data traffic and server energy consumption. While [16] applies estimates for server energy consumption from a more recent 2007 EPA report [17], the 2004 study applies estimates given in Roth's report. In both models the estimates of network device energy consumption are from Roth et al.

In [16] the annual direct energy demand of the Internet is estimated as 19.3TWh. Separate from that and not factored in to the estimate of the energy intensity, the annual direct energy demand of the telephone system is estimated as 3.8TWh.

The energy consumption of network devices is then extrapolated from the original values dating from the year 2000 to 2006 by applying an estimate of 14% of annual increase of aggregate energy consumption of "network equipment in server rooms, localized data centers, mid-tier data centers, and enterprise-class data centers" given in the 2007 EPA report to network devices, 13% CAGR to servers and 20% to storage devices.

Additionally, in the 2008 study, the traffic estimates by Andrew Odlyzko from 2001, used in the 2004 text, are replaced with more recent upper and a lower bound estimates from Minnesota Internet Traffic Studies (MINTS) from 2006. Such estimates are based on periodical measurements of data volume throughput on central and representative network nodes that is then extrapolated to monthly or annual time periods and scaled to the size of the remaining network. These measurements do not capture traffic that remains inside business LANs.

In both studies, the aggregate energy consumption is doubled to account for cooling and power transformation overheads, which was not applied by [14]. The resulting estimate of energy intensity of the Internet is 135.9 kWh/GB for the 2004 study and between 8.8 and 15.7 kWh/GB for the 2008 study, respectively.

In their account of direct energy demand of the Internet (Table 3 in [16]), they include estimates of total annual energy consumption of servers, data storage, and as network devices, hubs, routers, LAN switches and WAN switches from Roth and colleagues. Since the traffic estimate does not include LAN networks, this is necessarily an over estimate for the direct energy demand of the Internet.

As confirmed in conversation with Jon Koomey, the categories of WAN switches and routers in Roth's data include an inventory of ISP network devices, however they also include LAN switches and hubs that do not form part of core networks. It is impossible to retrospectively factor out the

lower class LAN routers from the ISP class devices. Not included in the estimate of Internet energy demand are fiber-optic devices.

As mentioned above, the estimates by Taylor and Koomey's had significant influence to assessments of digital services. For example, [3] apply these values in a comparative assessment of music downloads over physical shipping of CDs but update estimates of Internet traffic and energy consumption by extrapolation. Compound annual growth rate of Internet data volume was assumed to be 50% based on MINTS estimates while for the growth rate of network energy consumption the value from data center network equipment of 14% in the EPA report was applied (13% for servers, 20% for storage) resulting in a lower and upper estimate of 9.2 and 5.1 kWh/GB. By combining the EPA baseline estimate for servers and storage for the year 2000 with traffic estimates by Odlyzko the authors derive an average improvement rate of 30%.

### III. Methodology

#### A. Updating System Boundaries

In order to reduce the envelope of over estimation included in Taylor and Koomey's estimate, we evaluate the direct energy demand under a change of system boundaries. Firstly, we exclude servers and storage devices, as we believe they should be accounted for separately from the network devices, as explained above. Secondly, we exclude the categories of hubs (1.6TWh/a) and LAN network devices (3.3TWh/a) as the traffic estimate does not include them and they are not directly used in the delivery of digital services, but may be used to support the business processes that facilitate the digital services. However, we include the transmission devices (1.8TWh/a) from the category of telecom systems as these also transport data.

If we change system boundaries to exclude servers and storage and only include the categories routers (1.1TWh/a), WAN switches (0.15TWh/a) plus additionally the public Telephone networks equipment, the resulting value for the energy intensity in 2000 is 3.05TWh, compared to 6.15TWh given the original system boundaries. For 2006, the extrapolated value of annual network device energy consumption with altered system boundaries under a 14% CAGR is 3.9TWh, compared to the original 13.4TWh.

#### B. Efficiency improvements

Besides the choice of system boundaries, uncertainty enters into assessments through necessary extrapolation of increases in energy efficiency. Given the few publicly available estimates of total operated device energy consumption, and their significant age, device efficiency improvements and increases of traffic demand over time need to be taken into account. To that end, baseline estimates of total energy consumption as well as network traffic are extrapolated towards a reference year.

For example, Taylor and Koomey, in [16] extrapolate energy consumption of network equipment from the year 2000 values provided by Roth and colleagues to 2006 at 14% p.a. As frequently updated estimates for US network traffic are scarce as well, Weber and colleagues in [3] similarly extrapolated

reference values of monthly Internet data volume by the Minnesota Internet Traffic Studies [18] (450 - 800 PB per month in 2006) at a growth rate of 50% p.a the year 2008. Comparing their original 2006-estimates of Internet energy intensity, based on the extrapolated year-2000 annual device energy consumption from Roth et al. with those newly extrapolated values for 2008 they arrived at a combined efficiency improvement rate of 30%. This continuous extrapolation, although necessary in the absence of more recent estimates, amplifies the uncertainty already present in the original input data.

Improvements in device energy efficiency commonly referred to as Moore's law and that follow the famous rule of doubling capacity every 18 months (improvement rate of 58% p.a.) - and energy efficiency every 1.57 years [19] - are mainly the result of increasing density from miniaturization of structures in CMOS logic. Network device (in particular routers but also fiber optic technology), however, cannot benefit from this development to the same degree as computers.

Estimates of energy efficiency improvement rates for network devices are rare. Baliga and colleagues in [20] include a prognosis of future energy consumption in their study which they base on Neilson [21]. Neilson in turn combine trend data of annual router capacity increments from [22] and a spreadsheet from Nick McKeon with trend data on power consumption of routers (for which they unfortunately provide no reference) to estimate the annual improvements of energy efficiency per capacity. Neilson finds that the energy efficiency increases by 20% per year (22.4% when evaluating the data from McKeon). More recently, [23] reference Tamm [24] when stating that the rate of improvement has slowed down to about 10% per year for network equipment.

Tamm et al. perform a systematic analysis of annual efficiency improvements of network devices. First, they find that CMOS logic in telecom devices only improves by ~26% per year as opposed to ~41% of Moore's law for CMOS in general. Secondly, they find that the real savings are smaller than that because other system parts including fans and power supplies improve much slower. Tamm and colleagues don't provide a figure of 10% explicitly but a reconstruction of their Figure 7 yields an improvement rate of 12.5%. They also state that fiber optic equipment improves to the same degree.

Finally, an update of the original data by McKeon with more recent data on Cisco and Juniper core routers from [25] fits the trend line well ( $R^2$  increased from 0.90981 to 0.92022) and indicates overall capacity increments of 54% p.a. The subsequent regression analysis of these values yields an improvement rate of energy efficiency of only 18% (Fig. 2).

In order to realise these improvements on devices as savings in the operation of networks, operators need to replace their network devices by more efficient ones. However, as for example Baliga and colleagues noted the real reduction of average energy intensity of networks might be lower than the 12 to 18 percent of devices themselves. A point that is supported by the relative age of models in the empirical data from [26].

Besides these improvement of the network devices energy efficiency of networks is also improved by reduced energy consumption elsewhere. As network devices are housed in the datacentres, ongoing efforts to reduce PUE overheads are beneficial here, too. Although these developments might be slower given that many network devices are housed in so-called Points of Presence as opposed to data centers for the operation of server computers.

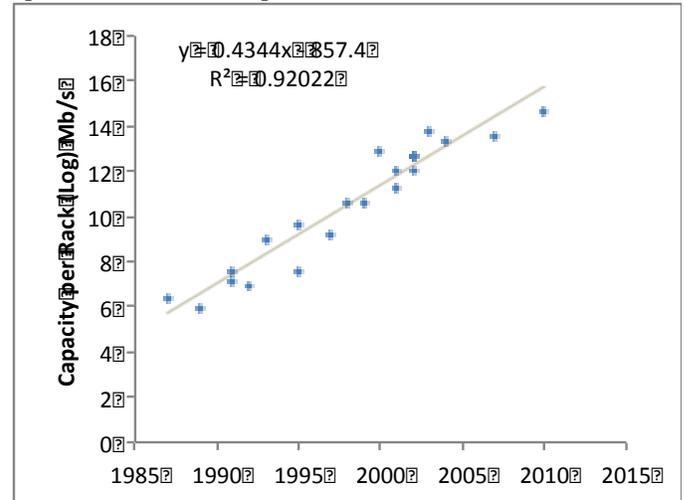


Fig. 2 – Trend of router per-rack capacity increments over time. The resulting annual increment is 54% p.a.

Given relatively low technology improvements rates then the single largest source for improvements that remains is better utilization of equipment since Weber and colleagues assume the PUE to remain constant at 2. Among the drivers for better utilization could be increased competition among operators due to lower margins from dropping cost of bandwidth over the years, which results in a greater pressure to reduce operational expenditure from energy consumption. If the device improvement rate only range between 12 to 18% these improvements must have been substantially higher than 30%. Given that these are one-off savings, improvement rates would have to drop for later years.

Without a new inventory of network device energy consumption, the improvement rate is difficult to corroborate. In order to evaluate the significance of a lower energy intensity reduction rate, we evaluate the resulting energy intensity for 10, 20 and 30% improvement rate.

#### IV. Results

Given the inventory data of annual network device energy consumption by Roth and colleagues and altering the system boundaries for included device types as described above, the estimate of annual energy consumption by the Internet is 7.8TWh, including an overhead factor of 2 for infrastructure services. This is less than a tenth of the original value of 84.6TWh by Taylor and Koomey. Accordingly, energy intensity for the year 2006 is reduced from the original values between 8.8 and 15.7 kWh/GB in [16] to between 0.81 and 1.44 kWh/GB for the scenario of high and low MINTS network traffic volume, respectively.

For the year 2008, assuming 30% annual reduction of energy intensity, values of 0.4 to 0.71 kWh/GB with an average of 0.55 kWh/GB would result, compared to the values by Weber and colleagues in [3] of 5.15 to 9.15 with an average of 7.15 kWh/GB. This is displayed in Fig. 3 while the numerical values are presented in Table 1.

However, if the rate of annual reduction of energy intensity is assumed to be lower than 30%, the difference is less dramatic. In Fig. 4 the reduction of energy intensity over time is plotted, assuming 10, 20 and 30 percent reduction p.a. for the MINTS high and low scenarios of data volume. The numerical values are listed in Table 2. According to these estimates, for the year 2014, the original estimates by Weber and colleagues would yield an average energy intensity of 0.84kWh/GB which is only a third higher than the value assumed in worst-case MINTS low traffic scenario with 10 percent p.a. reduction of energy intensity.

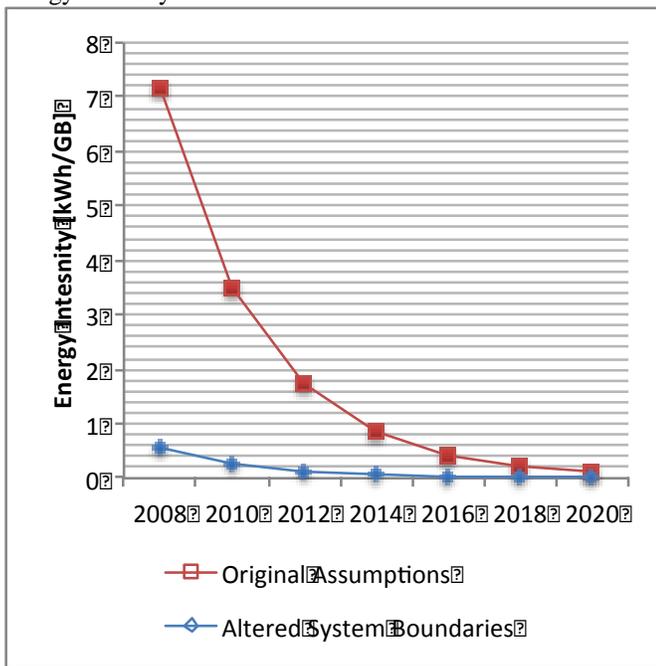


Fig. 3 – Energy Intensity [kWh/GB] change over time with the original system boundaries in [16] and [3] and the proposed changes.

Table 1 - Energy Intensity [kWh/GB] change over time with the original system boundaries in [16] and [3] and the proposed changes.

	Altered	Original
2006	1.13	12.24
2008	0.55	7.15
2010	0.27	3.50
2012	0.13	1.72
2014	0.07	0.84
2016	0.03	0.41
2018	0.02	0.20
2020	0.01	0.10

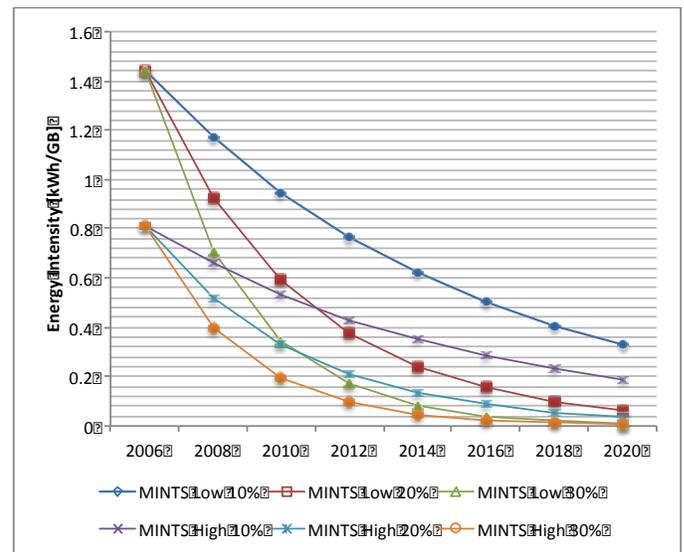


Fig. 4 – Top-Down modeled energy intensity [kWh/GB] change over time with altered system boundaries to Taylor and Koomey’s study [16] and lower and upper estimates of data volume according to MINTS.

Table 2 Top-Down modeled energy intensity [kWh/GB] with altered system boundaries to Taylor and Koomey’s study [16] and lower and upper estimates of data volume according to MINTS

	10%	20%	30%
MINTS Low			
2006	1.444	1.444	1.444
2008	1.170	0.924	0.708
2010	0.948	0.592	0.347
2012	0.768	0.379	0.170
2014	0.622	0.242	0.083
2016	0.504	0.155	0.041
2018	0.408	0.099	0.020
2020	0.330	0.064	0.010
MINTS High			
2006	0.813	0.813	0.813
2008	0.658	0.520	0.398
2010	0.533	0.333	0.195
2012	0.432	0.213	0.096
2014	0.350	0.136	0.047
2016	0.283	0.087	0.023
2018	0.229	0.056	0.011
2020	0.186	0.036	0.006

## V. Discussion

As consequence of the suggested change to the system boundaries that should be applied when delineating the

Internet, the baseline for the total energy consumption was reduced to a tenth of the original value. In all scenarios of energy intensity extending from this, the estimates of energy intensity were significantly lower than the values originally suggested.

Two comparative assessments based on the original estimates [2], [3] found that the digital service had a smaller environmental footprint than the physical service. The proposed reduced energy intensity would not invalidate these conclusions, but rather would strengthen them. However, the original estimates are popular and have influenced intuition and back-of-the-envelope calculations about energy footprints of digital services. Conclusions drawn from these may need revisiting and revising.

Our results present multiple scenarios of annual reduction rate and baseline network traffic. Instead of providing the average value as a single representative for energy intensity we chose to present an envelope in which the current top-down estimate should be located. Future assessments applying network energy intensity in an environmental assessment of a digital service can thus take this uncertainty into account during the interpretation of their findings.

The size of the envelope is increasing the greater the extrapolated time span. For the year 2014 the extrapolated estimates under the 30% reduction rate are one order of magnitude less than those achieved under a 10% reduction rate. This improvement rate is thus the greatest source of uncertainty.

The uncertainty in the initial assumption of baseline data volume propagates towards the extrapolated values - the lower initial value from MINTS was 50% smaller than the upper estimate of data volume.

From the changed system boundaries, the greatest reduction in annual energy consumption is from servers and data storage (60 and 8% of initial total) followed by LAN switches and hubs (17% and 8% of initial total). As we have argued above, we believe that servers are best accounted for specifically in the context of each digital service assessment and that LAN device should be covered outside of the direct energy demand. In order to reduce remaining uncertainty, new detailed inventories are needed.

The estimate is based on data for the US. In absence of representative data, it is reasonable to apply this value globally. More research is needed to understand how the structure and the energy efficiency of the Internet differs among regions on a global scale.

## VI. Conclusion

We have reviewed top-down estimates of energy intensity of the Internet network and updated the system boundaries and reference year in order to support assessments of digital services most effectively.

We have found that estimates of Internet energy intensity vary significantly depending on assumptions of traffic volume in networks and the chosen system boundaries but most importantly, on the assumed rate of annual efficiency

improvements. For the year 2014 this review suggests an envelope of energy intensity between 0.05 and 0.6 kWh/GB.

When used to estimate the use-phase energy footprint of digital services as the product of the energy intensity and the data volume transferred, this revised value will yield more robust estimates for the overall energy consumption.

## Acknowledgment

Thanks to Jonathan Koomey for very helpful email discussions. This work was funded in part by UK Research Council Digital Economy grant SYMPACT (grant number EP/I000151/1) as well as partly by UK EPSRC Large-Scale Complex IT Systems Initiative (grant EP/F001096/1).

## References

- [1] L. Erdmann and L. M. Hilty, "Scenario Analysis Exploring the Macroeconomic Impacts of Information and Communication Technologies on Greenhouse Gas Emissions," *J. Ind. Ecol.*, vol. 14, no. 5, p. 826+, 2010.
- [2] Å. Moberg, M. Johansson, G. Finnveden, and A. Jonsson, "Printed and tablet e-paper newspaper from an environmental perspective — A screening life cycle assessment," *Environ. Impact Assess. Rev.*, vol. 30, no. 3, pp. 177–191, Apr. 2010.
- [3] C. L. Weber, J. G. Koomey, and H. S. Matthews, "The Energy and Climate Change Implications of Different Music Delivery Methods," *J. Ind. Ecol.*, vol. 14, no. 5, pp. 754–769, 2010.
- [4] D. R. Williams, Y. Tang, and Daniel R. Williams, "A Methodology to Model the Energy and Greenhouse Gas Emissions of Electronic Software Distributions," *Environ. Sci. Technol.*, vol. 46, no. 2, pp. 1087–1095, 2011.
- [5] D. Schien, P. Shabajee, C. Preist, and S. G. Wood, "A Model for Green Design of Online News Media Services," in *WWW2013*, 2013, pp. 1111–1121.
- [6] P. Teehan and M. Kandlikar, "Comparing Embodied Greenhouse Gas Emissions of Modern Computing and Electronics Products," *Environ. Sci. Technol.*, Apr. 2013.
- [7] D. Schien, P. Shabajee, M. Yearworth, and C. Preist, "Modeling and Assessing Variability in Energy Consumption During the Use Stage of Online Multimedia Services," *J. Ind. Ecol.*, vol. 17, no. 6, pp. 800–813, Dec. 2013.
- [8] V. C. Coroama and L. M. Hilty, "Assessing Internet energy intensity: A review of methods and results," *Environ. Impact Assess. Rev.*, vol. 45, pp. 63–68, Feb. 2014.
- [9] E. Iannone, *Telecommunication Networks*. CRC Press, 2011, p. 918.
- [10] K. W. Roth, F. Goldstein, and J. Kleinman, "Energy Consumption by Office and Telecommunications

- Equipment in Commercial Buildings Volume I: Energy Consumption Baseline,” 2002.
- [11] S. Lanzisera, B. Nordman, and R. E. Brown, “Data network equipment energy use and savings potential in buildings,” in *ACEEE Summer Study on Energy Efficiency in Buildings*, 2010, pp. 1–14.
- [12] M. Pickavet, W. Vereecken, S. Demeyer, P. Audenaert, B. Vermeulen, C. Develder, D. Colle, B. Dhoedt, and P. Demeester, “Worldwide energy needs for ICT: The rise of power-aware networking,” *2008 2nd Int. Symp. Adv. Networks Telecommun. Syst.*, no. December, pp. 1–3, Dec. 2008.
- [13] B. Raghavan and J. Ma, “The Energy and Emergy of the Internet,” in *Hotnets*, 2011.
- [14] M. Gupta and S. Singh, “Greening of the internet,” *Proc. 2003 Conf. Appl. Technol. Archit. Protoc. Comput. Commun. - SIGCOMM '03*, p. 19, 2003.
- [15] J. Koomey, H. Chong, W. Loh, B. Nordman, and M. Blazek, “Network Electricity Use Associated with Wireless Personal Digital Assistants,” *J. Infrastruct. Syst.*, vol. 10, no. September, pp. 131–137, 2004.
- [16] C. Taylor and J. G. Koomey, “Estimating Energy Use and Greenhouse Gas Emissions of Internet Advertising,” *Network*. Working paper for IMC2, 2008.
- [17] U.S. Environmental Protection Agency, “Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431,” 2007.
- [18] Minnesota Internet Traffic Studies, “Internet Growth Trends & Moore’s Law,” 2008. [Online]. Available: [www.dtc.umn.edu/mints/igrowth.html](http://www.dtc.umn.edu/mints/igrowth.html). [Accessed: 08-Mar-2014].
- [19] J. G. Koomey, S. Berard, M. Sanchez, and H. Wong, “Assessing trends in the electrical efficiency of computation over time,” *IEEE Ann. Hist. Comput.*, 2009.
- [20] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker, “Energy Consumption in Optical IP Networks,” *J. Light. Technol.*, vol. 27, no. 13, pp. 2391–2403, Jul. 2009.
- [21] D. T. Neilson, “Photonics for switching and routing,” *IEEE J. Sel. Top. Quantum Electron.*, vol. 12, pp. 669–678, Jul. 2006.
- [22] I. Keslassy, S. Chuang, K. Yu, and D. Miller, “Scaling Internet routers using optics (Report TR03-HPNG-080101),” 2003.
- [23] D. C. Kilper, S. Member, G. Atkinson, S. K. Korotky, S. Goyal, P. Vetter, D. Suvakovic, and O. Blume, “Power Trends in Communication Networks,” *Sel. Top. Quantum Electron. IEEE J.*, vol. 17, no. 2, pp. 275–284, 2011.
- [24] O. Tamm and C. Hermsmeyer, “Eco-sustainable system and network architectures for future transport networks,” *Bell Labs Tech.*, vol. 14, no. 4, pp. 311–327, 2010.
- [25] W. Van Heddeghem, F. Idzikowski, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester, “Power consumption modeling in optical multilayer networks,” *Photonic Netw. Commun.*, vol. 32, no. January, Jan. 2012.
- [26] V. C. Coroama, L. M. Hilty, E. Heiri, and F. M. Horn, “The Direct Energy Demand of Internet Data Flows,” *J. Ind. Ecol.*, vol. 17, no. 5, pp. 680–688, Jul. 2013.