Schien, D., & Preist, C. (2014). Approaches to energy intensity of the internet. *IEEE Communications Magazine*, 52(11), 130-137.
<https://doi.org/10.1109/MCOM.2014.6957153>

Peer reviewed version

Link to published version (if available):
[10.1109/MCOM.2014.6957153](https://doi.org/10.1109/MCOM.2014.6957153)

[Link to publication record in Explore Bristol Research](#)
PDF-document

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Approaches to Energy Intensity of the Internet

Dan Schien and Chris Preist, University of Bristol

Abstract

With more and more activities taking place online concern over the environmental impact of digital services has drawn attention to the energy intensity of the network. Estimating the network energy intensity has been the subject of research for some time but results have differed widely, thus weakening the robustness of any conclusions drawn from assessments. A review of past studies shows two separate communities at work, applying different methods and assumptions. In this text we consider the approaches of top-down and bottom-up modeling. Top-down models have in the past usually given higher estimates of energy intensity than bottom-up models. We find that among the main reasons for the difference are varying system boundaries and assumptions on the number and energy efficiency of routers and optical transmission equipment. Through application of consistent system boundaries around the metro and core network and excluding access networks and customer equipment we reduce the difference between the energy intensity estimates of the alternative approaches. Additionally, we review the varying assumptions in existing bottom-up models and combine them in a meta-model. Through Monte-Carlo simulation over the distributions behind the varying assumptions we provide a more robust estimate of approximate energy efficiency for networks of 0.02 kWh/GB that can be used in the environmental impact assessment of digital services.

Introduction

The continuing growth of digital services such as streaming videos, browsing websites or generally exchanging data over the Internet has drawn some attention to their environmental impact, which is either indirect, referring to the potentially beneficial impact of changes that digital services induce in the wider society and economy, or direct, resulting from manufacturing and energy consumption of devices. An understanding of the tradeoffs between potential benefits and the negative direct impacts enables consumers, businesses and policy makers to take environmental impact into account.

Sustainability practitioners working for the businesses providing digital services (such as online news or online video) are experts in taking an end-to-end perspective and model all environmental impacts during the life cycle of a product but lack the resources and expertise to create detailed models of each subsystem under consideration, such as the network. Instead, they require guidelines and off-the shelf models.

Despite some progress, efforts such as ICT sector guidance service chapter to the Greenhouse Gas Protocol or the ITU L 1410 "Methodology for the assessment of the environmental impact of information and communication technology goods, networks and services" currently lack such models. Hence practitioners adopt results from past studies without detailed analysis of underlying assumptions. In the case of energy usage by the Internet, this can be particularly problematic because the great variation of figures used means that the selection of one rather than another can dramatically effect the conclusions of an assessment.

In this paper, we present a meta-analysis of past studies of energy consumption in the network. While this text provides an estimate of energy intensity of only edge, metro and core networks, a complete assessment of a digital service needs to take all network parts into account, including the customer premise equipment, wired access networks, wireless access networks and metro and long haul networks as shown in Figure 1.

Figure 1 – End-to-end model of the network. Servers in data centers are providing digital services to user devices in the home via the network including customer premise equipment (CPE) connecting to user devices in the home, access network, metro, long haul/core network, undersea cable transport.

Two Communities: Industrial Ecology and Network Research

Life cycle assessments of digital services usually use a measure of the energy intensity of their network usage to determine their allocation of energy. This is normally stated in Joules per bit (J/b) or Kilowatt-hours per Gigabyte (kWh/GB). It is calculated as a share of the network energy consumption relative to the data volume transported or, equivalently, of the power consumed per bandwidth sustained. Given the energy intensity, the energy footprint of a service is estimated as the product of energy intensity and the data volume of the service; per single unit of service (such as one minute of video stream) or for the entire audience (i.e. all videos streamed per year).

Energy intensity has been estimated using two different approaches: top-down and bottom-up. Each approach relates energy consumption to data traffic but differs in the kind of input data it applies and enables its use in different applications. The distinction between top-down and bottom-up modeling approaches can be found in several domains where an overall property being calculated is also present on the level of model components.

A top-down model, for example [1], estimates the total energy use of an entire subsystem, for example 'all data centers' or 'the internet', measures or estimates the total quantity of a given service type provided, for example data transmitted, and divides the former by the latter to give the energy consumption per unit of service. Hence regarding energy consumption, it treats a given subsystem as a black box. Top-down models can evaluate change on the level of aggregate variables: total network traffic, average energy consumption per device class. Since top-down models are parameterized with market data, they are accessible to non-experts in network technology and are open to external validation.

For energy footprinting of digital services, the most influential top-down models (specifically, [1], [2]) were developed by researchers from the inter-disciplinary industrial ecology community – although this categorization is loose as academic communities are not clearly separated, and individual researchers publish in a variety of venues and collaborate across boundaries.

One of the defining goals of studies from this community is to quantify energy and material flows in industrial systems in order to increase the sustainability of industrial systems by understanding their relationships on a technological, social, economic and environmental level. A particular focus is taking whole-systems

perspective: investigating the dynamic relationship over these levels in order to prevent shifting the burden from one part of the system to another. For example, the shift towards distributed services provided by central servers through low power clients might result in greater energy consumption by the network.

Top-down models are conducive to this whole-systems perspective. Energy consumption in top-down models is usually estimated from market sales per device class and corresponding average power consumption to give a total over all considered device classes. By comparing this total with other macro scale energy statistics, it is easy to sanity-check them. And by treating the modeled system partly as a black box, they also don't require detailed knowledge on the network architecture. On the other hand, they cannot be used to evaluate changes to part of the network but only on trends of changing total network traffic or total energy consumption.

A bottom-up model (for example [3]), in contrast, calculates the overall energy intensity from the sum of the energy intensity of the subsystem components – usually the physical devices in the network. These models have also been referred to as transactional models as they allocate energy consumption to the transaction of data from end to end. As they represent energy intensity on the level of the system components they are more flexible than top-down models to evaluate change: they can be used to evaluate the impact of modifications to the system architecture and its components.

Such models thus require detailed knowledge of the operation and design of networks. Although this knowledge is already held by network operators and thus, in principle, it is possible to represent each individual device in a bottom-up model, in practice network operators do not publicly disclose this information for business

reasons. Instead, bottom-up models have been build based on implicit assumptions around the typical architecture of networks and they are thus more difficult to validate.

It is no surprise that bottom-up models originate from the network research community, which investigates the design of networks and considers a number of metrics, including energy consumption. Network researchers have quantified energy consumption to estimate environmental impact from carbon emissions and, more frequently, to address network operator costs. If environmental impact was being assessed then the interpretation of results typically focused on directing research in network design.

Bottom-up models facilitate the evaluation of alternative design choices if they represent the network components that are to be altered. The scope of an investigation thus affects the level of detail by which the network is modeled. For example, if the goal of a study is to evaluate savings from optical switching then these devices must be explicitly modeled, if not then fiber optic components might be modeled in less detail with average values for energy consumption and capacity. At the same time the level of detail of modeling is naturally constrained by the simultaneously increasing complexity of the model, which is particularly relevant for end-to-end models.

Both modeling approaches introduce significant sources of uncertainty: Accuracy of a top-down model depends on the assumed total energy consumption and data volume. Accuracy of bottom-up models depends on how closely the assumed network architecture mirrors real network deployments. The energy elasticity of network devices, that is the ratio between marginal change of the utilization of a device and the

resulting marginal change in energy consumption, cannot be taken into account by top-down models. However, the current generation of network devices has very low energy elasticity, thus not limiting the potential accuracy of top-down models, as [4] find.

Yet, both approaches arrive at significantly different estimates for the energy intensity. In [5] a review of top-down and bottom-up models by Coroama and Hilty finds that top-down studies consistently arrive at higher estimates for energy intensity than bottom-up studies, differing by 4 orders of magnitude. Reasons for this discrepancy are given as being varying years of reference, and varying system boundaries sometimes including user devices, data centers, and customer premise equipment. However, they make no attempt to study the discrepancy by normalizing the system boundaries and thus leave open the question if the differences can be resolved. Their Table 1 and Figure 1 list estimates of energy intensity of the past studies along a seemingly inverse exponential curve between 136 kWh/GB and 0.006 kWh/GB. We note that none of the three top-down studies with the highest energy intensity value includes end-user devices or optical fibers. Thus, even among studies with a similar year of reference and not including user devices a variation between one and two orders of magnitude remains. Specifically, the top-down model [2] and the bottom-up model [3], both with data for 2008 and both not including end user devices or customer premise network equipment, arrive at estimates of 7 kWh/GB and 0.006 kWh/GB. Hence the explanation for this discrepancy given by Coroama and Hilty is only partial, and further analysis is necessary to provide network energy intensity values for life cycle assessment practitioners.

In order to understand why bottom-up models arrive at significantly lower estimates than top-down variants and to reduce the overall uncertainty, we review system boundaries of past models, reconstruct the models within system boundaries around on the edge and core network. We find that even with normalized system boundaries the bottom-up models arrive at varying results due to differing assumptions on bottom-up model parameters regarding the number and energy intensity of routers and fiber optical equipment. Without further qualification, it must be assumed that these models represent the real variability of existing network deployments. Based on this assumption, we then construct a bottom-up meta-model and parameterize it with distributions to represent the varying assumption in existing models. By means of a Monte Carlo simulation we then generate a distribution of the overall energy intensity from which we suggest a new authoritative value for use by sustainability practitioners in assessments.

This text thus makes the following contributions:

- Provide a comparison of energy intensity estimates of top-down and bottom-up models within appropriate system boundaries for edge and core networks
- A review of the most robust bottom-up models of energy intensity and a normalization within common boundaries
- A distribution of the energy intensity with a single average value together with a confidence interval based on a principled approach.

Top-Down models

The most influential top-down model comes from the industrial ecology community estimating the energy intensity of the US Internet for the year 2006 (Taylor and

Koomey [1]). In this study, the annual direct energy demand of the Internet is estimated as 19.3 TWh based on sales data of device types for the year 2000 in [6] and then extrapolated to a value of 42.3 TWh for the year 2006. A PUE value of 2 is applied additionally. The total energy consumption is then divided by an upper and lower bound estimate of annual network traffic of 5.4 to 9.6 Exabyte to give a resulting energy intensity of 9 to 16 kWh per GB. The authors state that some assumptions were conservative and thus the results constitute an overestimate. A later study [2] then applied an annual rate of reduction of 30% to the average between to high and low estimate of energy intensity estimates of the year 2006 to account for increasing efficiency of devices and arrive at a value of 7 kWh/GB for 2008. Extrapolated to 2014, this would result in a mean energy intensity of 0.84 kWh/GB, a value that is considerably higher than most of the bottom-up estimates listed in the review [5].

To allow more accurate comparison of this model with bottom-up models of the core network, we rework this study to model the core network alone. We use the same data set [6] and methodology but change the system boundaries. Referring back to the end-to-end model of the network in Figure 1, it is necessary to include fiber optic equipment and edge and core routers and switches but to exclude servers and data storage as well as campus network equipment such as office floor level hubs and small switches. While this is mostly straightforward, the 'router' category includes both high-end core routers and small office-level models. Given that the router category is the largest position in the inventory, this results in a significant overestimate. Although the two remaining top-down models in [5] are not based on Roth's inventory, these cannot be used to triangulate the portion of core routers from all routers as one is equally focused on campus networks and the other only provides an aggregate result for network device energy consumption.

In **Error! Reference source not found.** we list the inventory categories from [6] as used in [1] and in our reworking to focus on the core network alone. The resulting estimate of total annual energy consumption in the updated inventory is 3.9 TWh/year compared to 42.3 TWh/year in the original estimate. If the assumed 30% annual improvement rate by Weber and colleagues in [2] is applied to this estimate, the resulting energy intensity would be 0.55 kWh/GB for 2009 (down from 7 kWh/GB) and 0.07 kWh/GB for 2014 (down from 0.84 kWh/GB). However, this annual improvement rate was calculated relative to the observed growth of energy consumption by data center network equipment from 2000 to 2006, and might be too high for carrier network equipment. Kilper et al. in [7] refer to Tamm et al. in [8] for an estimation of annual improvements of telecom equipment of 10%. Although Tamm et al. do not provide the value of 10% explicitly but a reconstruction of their Figure 7 results in a value of 12.5%. At this lower annual improvement rate the resulting average energy intensity for the top-down model would be 0.39 kWh/GB for 2014.

Table 1 Inventory of system components in the top-down model [1] in the original system boundaries and our rework.

	2006 Electricity [TWh/year]	
Equipment Type	Original Inventory [1]	Reworked Inventory
Servers	24.5	
Data Storage	4.4	
WAN Switches	0.3	0.3

Routers	2.4	2.4
LAN Switches	7.2	
Hub	3.5	
Transmission Networks		1.2
Sum	42.3	3.9

More recently, another top-down model for the Swedish core network by Teliasonera estimated its energy efficiency as 0.08 kWh/GB for the year 2010 [9] which is lower but not entirely dissimilar to our reworked values. This study is supported by confidential data from Teliasonera and thus important for corroboration but does not provide enough detail in order to compare and explain differences to other studies.

Bottom-Up Models

Bottom Up models of end-to-end network energy intensity combine a network architecture for access, metro and core network layers, with a specific parameterization of device energy intensity values.

Any variation in the overall energy intensity results from different assumptions on the route length of metro and core networks as well as the energy intensity of router and fiber optic equipment, which varies between specific device types and models and with device age. Additionally, overheads for building infrastructure, expressed as power usage effectiveness (PUE), redundancy, and overcapacity increase the overall network energy intensity.

One of the first end-to-end models, by Baliga and colleagues [3], estimated power draw per user of the optical Internet as a function of bandwidth in the access network for several access network technologies and has been referenced in assessments of digital services several times. Given that user bandwidth was estimated from statistical average values, the overall estimate can be converted equivalently to energy consumption per bit. More recent formulations of the model by the same authors have maintained parameterization and architecture largely unaltered.

They describe a reference end-to-end architecture for the network including customer premise equipment (CPE), access, edge, metro, long-haul networks and undersea cables as well as an additional IPTV network.

Baliga and colleagues only published the energy efficiency result including the access network, which depends on the access rate. Our reproduction of their model resulted in a value for the energy consumption of the core Internet – not including the access network but including undersea traffic – of 2.66 J/Mb which equals 0.0059kWh/GB. As the authors acknowledge in their text, the model provides an underestimate of the energy efficiency of the Internet.

Another notable end-to-end bottom-up model by Kilper et al. [7] evaluate how the power consumption of optical networks is likely to change until 2020 and take a mix of different types of services into account. They provide a layered network path model to estimate energy consumption for services using a specific network topology, for example Peer to Peer vs. video, by summing up the energy consumption of each layer traversed. However, unlike Baliga, they do not include undersea cables and associated terminals in their model.

In total, our reproduction of their model yields an energy consumption of 3.28 J/Mb (0.0073 kWh/GB) for a path that includes one leg of edge, metro and long haul network, which despite absence of a leg of undersea cable, is higher but of comparable magnitude to Baliga's values. The year of reference for equipment efficiency values in both studies is 2008 [7],[5].

The distribution of energy intensity over the subsystems is substantially different in the two models. Although Kilper et.al. agree with Baliga et.al. that the core layer is more impactful than the edge, their model indicates that the fiber optic devices contribute to a much greater degree.

Figure 2 – Composition of energy intensity by edge and metro routers, core routers and fiber optic transport in Baliga et al. [3] and Kilper et al. [7]. The fiber optic transport includes overland and subsea cables in Baliga et al.

Although we consider these studies to be the most robust end-to-end models of energy intensity, there are many other excellent models of energy consumption in

networks. However, these are usually modeling energy consumption on a network scale (as opposed to end-to-end) or evaluate relative changes without providing absolute values, and thus are not applicable to our needs. [12] provides the most detailed model of the optical layer, and so provides a valuable source of individual parameters and to corroborate assumptions in the meta model we present below.

Finally, a study from the industrial ecology community by Coroama et al. [10] provides an estimate of the energy intensity of transporting the video signal of a virtual conference on a network path between Switzerland and Japan of 0.2 kWh/GB, which is higher than our reworked top-down estimate. This study is unique in that network operators provided specific values for power consumption, utilization and capacity of routers and fiber optic transmission equipment, which is relevant to validate other bottom-up models from the network research community. As the authors acknowledge, the study investigates a worst-case scenario given the unusually long distance of the video channel spanning three continents from Europe across the US to Japan and 27000 km of distance.

Meta-Model

The highest estimates of energy intensity for edge and core network by the bottom-up models in the previous section 0.2 kWh/GB is 33 times higher than the lowest of 0.0059 kWh/GB. Although the route between Davos and Nagoya analyzed by Coroama and Hilty is untypically long, this difference alone cannot explain the variance. Given that there is no clear underlying reason in the models to favor one over the other, the difference partly represents the actual variability found in real network deployments and the actual average energy intensity of edge and core networks is to be found along the spectrum defined by the variability of underlying parameters. We now combine

different data and structural assumptions within the models discussed so far to allow us to calculate an average value for the energy intensity of the core internet. From the studies discussed above, we adopt the most detailed and robust model for each layer under consideration. Similarly to [7] we model each layer as composed of a number of nodes (IP + fiber optic devices) and the energy intensity per layer is the sum of the energy intensity of all nodes in a layer multiplied by factors for overcapacity, PUE and redundancy. These intensity values are then added over all devices that comprise a layer. The overall energy intensity is the sum of the edge, metro and core layers. We follow [11] in explicitly modeling the components of the optical layer (OTN switches, transponders, line amplifiers, regenerators) and we follow [3] in the modeling of the undersea transport. We apply an energy efficiency improvement rate of 12.5% per annum (taken from [8]) on deployed network devices, to normalize all data to a reference year of 2014. The model is available in code and with typeset documentation online¹.

Given the combined structure and parameterization of the model we then perform a Monte Carlo simulation to give us a distribution of the overall energy estimate including a mean value to represent the average case of core networks in general.

The parameterization of the model is provided in Table 2. The distributions for routers are calculated by resampling from a Gaussian kernel density estimated distribution, using the data in Table 3. For PUE and redundancy, we apply the same single value of 2 that was assumed by all studies.

¹ <http://nbviewer.ipython.org/gist/dschien/1859c0f525473211f66f>

Table 2 – Model parameters including parameter name, the type of distribution applied in the Monte Carlo simulation, and the distribution parameters. Parameters are grouped by applied distribution type: uniform, triangular, choice. For a uniform distribution min and max denote the boundary values. A choice denotes discrete values. A point estimate refers to a single value. Device energy intensity values are listed with their extrapolated value for 2014 based on a 12.5% percent annual improvement rate and the original value in brackets. References to the original sources are listed at the bottom of the table and are referred to by superscript indices.

Model Variable				Mean (2014)	Unit
Single data points					
Overcapacity edge layer: 10^2 , PUE: $2^{1,2,3,4}$, Redundancy: $2^{1,2,3,4}$					
Triangular Distributions	Min	Mode	Max	Mean	Unit
Energy intensity optic amplifiers	$0.03 (0.065)^2$	$0.21 (0.27)^3$	Proprietary (Proprietary) ⁴	Proprietary	J/Gb
Energy intensity optical switch	$0.04 (0.05)^3$	$0.35 (0.46)^2$	$1.42 (1.85)^2$	0.60	J/Gb
Number core hops	3^2	6^4	10^1	6.33	-
Number metro hops	3^1	4^2	12^4	6.33	-
Total distance	7500^1	7500^2	8217^4	7739	km
Uniform Distributions	Lower	Upper		Mean	Unit
Energy intensity regenerator	$7.66 (10)^3$	Proprietary (Proprietary) ⁴		Proprietary	J/Gb
Distance sea cable	6000	12000^4		9000	km
Energy intensity edge switch	$2 (4.46)^1$	$3.59 (8)^2$		2.80	J/Gb
Energy intensity OTN switch	$1.57 (3.5)^2$	$2.6 (3.4)^3$		2.09	J/Gb
Energy intensity per km undersea cable	0.021^8	0.066^9		300	J/Gb/km
Energy intensity transponder	$2.11 (4.7)^2$	$3.83 (5)^3$		2.97	J/Gb
Overcapacity core layer	$2^{1,2}$	4^4		3	-
Overcapacity metro layer	5^2	10^6		7.50	-
Span optical amplifier	80^2	100^1		90	km
Undersea traffic share	0.1^{10}	0.5^{10}		0.3	-
¹ [3], ² [7], ³ [11], ⁴ [10], ⁵ Exclude regenerators from model, ⁶ Half of the utilization in the SWITCH research network reported in [10], ⁷ Based on the total power consumption and number of optical amplifiers and regenerators for the Internet2 core network from conversation with the authors of [10]. ⁸ Based on formula 15 in [3]. ⁹ Based on the average energy intensity per km in [10]. ¹⁰ Portion of undersea traffic varies with location of the user and service.					

The distribution of hops in metro and core networks is based on the assumptions in [3], [7], [10]. In [10] 6 hops are located in the core network and 12 hops are in the

Swiss and Japanese research networks SWITCH and NICT with 7 and 5 hops respectively which we use as the high estimate for the metro network.

Results from the bottom-up model are strongly influenced by assumptions of the network utilization or overcapacity. This refers to the difference between maximum capacity, which serves as the basis for the calculation of the devices' energy intensity, and the actual use of capacity. [3] assume no overcapacity for edge and metro, which is an idealization that we ignore. In [10] utilization on routers and links in core network combined is 26.3%, excluding the undersea cables and terminals, resulting in an overhead coefficient of 4. [10] also provide utilization values for the SWITCH research network of 5% which we exclude because it is likely to be lower than commercial networks.

Table 3 Energy Intensity of Metro and Core Routers as provided in previous studies and extrapolated to 2014 based on an improvement rate of 12.5% per year.

Router Model	Source	Energy Intensity [J/Gb]	Year of Reference	Energy Intensity 2014 [J/Gb]
Metro Routers				
Cisco 12816	[3]	25.75	2008	11.56
Cisco 7603	[10]	25.00	2009	12.82
Cisco 7606	[10]	16.04	2009	8.23
Cisco 7613	[3]	38.33	2008	17.20
Cisco 10008	[3]	137.50	2008	61.71
Hitachi GS4000 320E	[10]	12.50	2009	6.41

Hitachi GS4000 160E	[10]	10.00	2009	5.13
Cisco 6513	[3]	8.36	2008	3.75
Cisco 6513	[10]	40.00	2009	20.52
Cisco 6509	[10]	40.00	2009	20.52
Juniper MX960	[10]	16.20	2009	8.31
Mean Energy Intensity [J/Gb]		33.61		16.01
Core Routers				
Juniper T1600	[10]	34.48	2009	17.69
Juniper T640	[10]	17.47	2009	8.96
Juniper T320	[10]	16.20	2009	8.31
Cisco CRS - 1	[3]	17.03	2008	7.64
Cisco CRS - 3	[11]	10.00	2012	7.66
Generic	[7]	12.60	2008	5.65
Mean Energy Intensity [J/Gb]		17.96		9.32

Results

The resulting distribution from the Monte Carlo simulation is displayed as a box and whisker plot in Figure 3 showing the total energy intensity as well as the contribution of the edge, metro, and core layers and undersea segments.

The mean energy intensity for the year 2014 is 0.02 kWh/GB with 25th, 75th percentile of 0.0144 and 0.023 and a median of 0.018 kWh/GB, respectively.

Figure 3 – Box and Whisker plot of energy intensity from the Monte Carlo simulation of the meta-model showing the total energy intensity for the edge and core networks and the energy intensity for individual layers. The red line indicates the median values, the vertical edges of the boxes mark the 1st (lower edge) and 3rd (upper edge) quartile. The blue dots indicate the mean. The horizontal black lines indicate 1.5x the inter quartile range (IQR), the distance between the 1st and 3rd quartile. Outliers outside of the IQR are marked as crosses.

In Figure 4 we compare the energy intensity values of the studies discussed so far with that of the reworked top-down model (0.39 kWh/GB) and the bottom-up meta-model (0.02 kWh/GB). Though the discrepancy is substantially reduced from the original estimates they cannot be compared like for like due to the inclusion of campus level routers in the top-down estimate which highlights an important area for further research. Other reasons that will contribute to the discrepancy will be (i) the age of the underlying data behind the top-down estimate means that the margin of error of the projection forwards is high; (ii) bottom-up models tend to be leaner, and will miss some deliberate redundancy or spare equipment. The top-down model by

Malmodin et al. [9] with 0.08 kWh/GB arrives at a value that is only four times higher than the meta-model result and thus provides partial corroboration. Unfortunately, it does not provide a detailed account of the model inventory to investigate what specific assumptions differed or were identical.

Figure 4 - Energy intensity estimates from top-down and bottom-up models with similar system boundaries on a log scale.

Discussion and Conclusion

The estimate of resulting energy intensity presented in the previous section is based on normalizing boundaries and statistically combining assumptions from previous studies, which in turn were based on measurements and experience.

The overall estimate can be used in sustainability assessments to estimate the network energy consumption which can be attributed to specific digital services. For example, a content service provider such as the BBC could apply this energy intensity to estimate the network energy consumption to be attributed to the downloading one hour of HD video on the BBC iPlayer service. The associated file is approximately 1 GB

in size, and so would be attributed 20Wh of the energy consumed by the core and edge networks. More detailed analyses of this kind can be used to explore the impact on energy consumption of alternative deployment architectures for digital services. [12].

A complete model of energy consumption involved in the delivery of digital services must also model access networks – the use of home, campus and mobile networks to access a service. The energy consumption of these is significant. As we argue in [12], usage characteristics of such equipment means that energy intensity is not an appropriate metric, and other allocation approaches are needed. This is discussed further in [13], and models are proposed.

More broadly, energy intensity of the network constitutes an example of an assessment of environmental impact of an industrial system. For these to be reliable, more input from the engineering community is required. Models by the industrial ecology community tend to provide over-estimates, in order to err on the safe side while engineering models tend to provide under-estimates, for example by abstracting from legacy systems where not needed.

Further research is necessary to provide a transparent inventory for top-down models that specifically identifies service provider network routers from campus network routers.

Both the community of network researchers and that of industrial ecology can contribute in order that the energy intensity values which are continuously used by practitioners are reliable and accurate.

Acknowledgement

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).

Biographies

Daniel Schien (daniel.schien@bristol.ac.uk) received his ‘Diplom’ (MSc equivalent) in computer science from TU Berlin, Germany. He is currently working toward a Ph.D. degree in computer science at the University of Bristol, UK. His research interests include environmental assessments of ICT and smart city solutions.

Chris Preist (Chris.Preist@bristol.ac.uk) is a reader in sustainability and computing systems at the University of Bristol, UK.

References

- [1] C. Taylor and J. G. Koomey, “Estimating Energy Use and Greenhouse Gas Emissions of Internet Advertising,” Working paper for IMC2, 2008.
- [2] C. L. Weber, J. G. Koomey, and H. S. Matthews, “The energy and climate change impacts of different music delivery methods,” 2009.
- [3] 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.

- [4] C. a Chan, A. F. Gyga, E. Wong, C. a Leckie, A. Nirmalathas, and D. C. Kilper, "Methodologies for Assessing the Use-Phase Power Consumption and Greenhouse Gas Emissions of Telecommunications Network Services.," *Environ. Sci. Technol.*, Dec. 2012.
- [5] 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.
- [6] K. W. Roth, F. Goldstein, and J. Kleinman, "Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings Volume I : Energy Consumption Baseline," 2002.
- [7] D. C. Kilper, 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.
- [8] 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.
- [9] J. Malmudin, D. Lundén, Å. Moberg, G. Andersson, and M. Nilsson, "Life Cycle Assessment of ICT," *J. Ind. Ecol.*, May 2014.
- [10] 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.

- [11] 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.
- [12] 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.
- [13] V. C. Coroama, D. Schien, C. Preist, and L. M. Hilty, "The Energy Intensity of the Internet: Home and Access Networks," in *ICT Innovations for Sustainability*, vol. 310, L. M. Hilty and B. Aebischer, Eds. Cham: Springer International Publishing, 2015, pp. 137–155.