

The Energy Intensity of the Internet: Home and Access Networks

Vlad C. Coroama¹, Daniel Schien², Chris Preist², and Lorenz M. Hilty^{3,4,5}

¹ Measure-IT Research, Bucharest, Romania

vlad.coroama@measureit-research.eu

² Department of Computer Science, University of Bristol, UK

³ Department of Informatics, University of Zurich, Switzerland

⁴ Empa, Swiss Federal Laboratories for Materials Science and Technology,
St. Gallen, Switzerland

⁵ Centre for Sustainable Communications CESC, KTH Royal Institute of
Technology, Stockholm, Sweden

Abstract. Estimates of the energy intensity of the Internet diverge by several orders of magnitude. We present existing assessments and identify diverging definitions of the system boundary as the main reason for this large spread. The decision of whether or not to include end devices influences the result by 1-2 orders of magnitude. If end devices are excluded, customer premises equipment (CPE) and access networks have a dominant influence. Of less influence is the consideration of cooling equipment and other overhead, redundancy equipment, and the amplifiers in the optical fibers. We argue against the inclusion of end devices when assessing the energy intensity of the Internet, but in favor of including CPE, access networks, redundancy equipment, cooling and other overhead as well as optical fibers. We further show that the intensities of the metro and core network are best modeled as energy per data, while the intensity of CPE and access networks are best modeled as energy per time (i.e., power), making overall assessments challenging. The chapter concludes with a formula for the energy intensity of CPE and access networks. The formula is presented both in generic form as well as with concrete estimates of the average case to be used in quick assessments by practitioners. The following chapter develops a similar formula for the core and edge networks. Taken together, the two chapters provide an assessment method of the Internet's energy intensity that takes into account different modeling paradigms for different parts of the network.

Keywords: Internet, Energy Intensity, Energy Efficiency, Customer Premises Equipment, Access Network

1 Introduction

Information and Communication Technologies (ICT) are increasingly perceived as enablers of a reduction of anthropogenic greenhouse gas (GHG) emissions. Studies providing evidence for this enabling effect come not only from academia

This Accepted Author Manuscript is copyrighted by Springer. The final publication will be available via <http://link.springer.com/bookseries/11156> by end of August 2014. Suggested citation: Coroama, V.C., Schien, D., Preist, C., Hilty, L.M.: The Energy Intensity of the Internet: Home and Access Networks. In: Hilty, L.M., Aebischer, B. (eds.) ICT Innovations for Sustainability. Advances in Intelligent Systems and Computing 310, Springer International Publishing (2014, in press)

[1–3], but also from organizations as diverse as ICT industry associations [4], the European Commission [5], and the World Wildlife Fund [6]. Reductions are usually estimated based on quantitative scenarios [7], yielding an abatement potential which has to be adjusted downwards by an estimate of the ICT applications own footprint to calculate the net effect.

ICT can reduce energy consumption and related GHG emissions through three mechanisms: i) from optimization effects in domains such as smart engines, buildings, or logistics, ii) due to ICT-supported novel paradigms for the generation and distribution of electricity (i.e., smart grids), and iii) due to substitution effects in which information and communication services partly replace other more energy-intensive activities [4, 8].

Two problems must be solved to quantify the net effect of ICT applications in these cases: First, the energy savings induced by ICT must be assessed. As we argue in [9] and [10], this is methodologically challenging: The baseline scenario, among other factors, as it expands into the future, is inherently speculative. Moreover, allocation issues are raised by the fact that ICT typically does not induce efficiency on its own, but only in a suitable technological, political, or organizational context. Secondly, the energy consumption of the ICT solution involved must be determined. This is also technically challenging, and existing literature reports diverse results. The current and the subsequent chapter explore this issue, with a particular focus on Internet services.

The *energy intensity* of the Internet, expressed as energy consumed to transmit a given volume of data, is one of the most controversial issues. Existing studies of the Internet energy intensity give results ranging from 136 kWh/GB [11] down to 0.0064 kWh/GB [12], a factor of more than 20,000. Whether and to what extent it is more energy efficient to download a movie rather than buying the DVD, for example, or more sustainable to meet via videoconference instead of travelling to a face-to-face meeting are questions that cannot be satisfyingly answered with such diverging estimates of the substitute’s impact.

The Internet’s energy consumption, and the energy intensity of the Internet as a suitable metric, are the topics of both the current chapter and the next one. The current chapter presents a review of existing studies and provides explanations for the large spread in their results. The chapter then recommends a definition of the system boundary that is most useful for decision-making and concludes by zooming in on the peripheral parts of the network and assessing their contribution to the energy intensity of the Internet. The next chapter [13] will focus on the remaining components, i.e., the core of the Internet.

2 Definitions and System Boundaries

In the late 1960s and early 1970s, the term *Internet* literally indicated the interconnection of a small number of local area networks on university campuses; back then, “the Internet” comprised a few routers and cables. Today, the Internet is the vast and heterogeneous infrastructure connecting billions of computers worldwide using the TCP/IP family of communication protocols. The majority

of these computers belong to private users who connect to the Internet through their *Internet service provider* (ISP).

Figure 1 presents a high-level structure of the Internet. We comment only briefly on this structure here; more detailed discussions can be found in the following chapter in section 2.1, and in the network models presented by [14, 15]. Users' devices (such as desktop, laptop or tablet PCs and smartphones) connect through what is referred to as *customer premises equipment* (CPE), which are mainly WiFi routers and modems, to their ISP. The ISP bundles the data from several users in multiplexers. These vary depending on the subscription technology; for the widely used DSL connectivity they are *digital subscriber line access multiplexers* (DSLAM). Together with the cables connecting them to the CPE, these multiplexers constitute what is called the *access network*. After passing through an edge router, the traffic enters the metro and core parts of the Internet where routers with increasing capacities bundle the traffic. On the other side, the traffic is decomposed according to its destination; a large part is directed to data centers, while a smaller part is directed to other users (not represented in Figure 1).

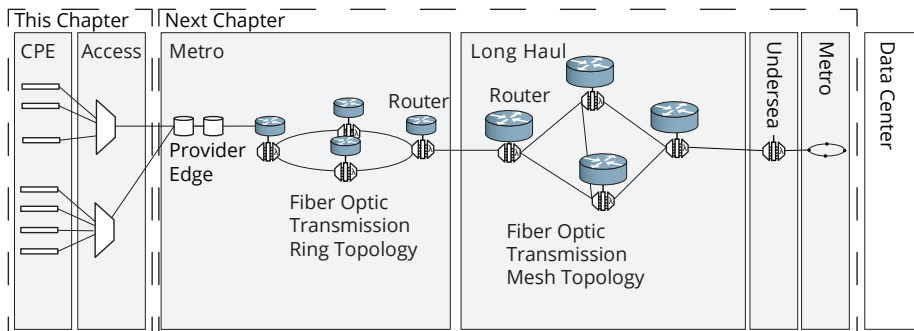


Fig. 1. Model of the Internet structure.

First attempts to understand the energy consumption of this distributed and heterogeneous power-consumption system were undertaken a decade ago, in 2003-2004. Starting from statistical data and studies on the ICT equipment in use, both [16] and [11] estimated the yearly power consumption of the Internet in the US. Dividing this value by an estimate of the US Internet traffic for that year resulted in estimates of the energy intensity of the Internet, i.e., the energy consumed throughout the Internet per amount of data transferred.

The two studies differed in their definition of what constitutes the Internet. While [16] considered only networking equipment (i.e., Ethernet hubs, LAN and WAN switches in offices and buildings, together with the routers of the Internet), [11] also took into account devices in data centers such as servers or data storage (see Figure 1). As we discuss in more detail later, this definitional discrepancy

across studies persists until today, and is one of the main causes for the large spread of the published results.

Before going into detail, we note that we discuss estimates for direct energy consumption in the form of electricity only. The energy supply chain (containing the supply of primary energy, power plants transforming primary energy sources to electricity, and grids bringing them to the consuming devices) is excluded from the system under study. We also exclude the “grey” energy embedded in ICT hardware, although the material flows caused by producing and disposing of hardware are significant (see [17] and the chapter by Hischier et al. [18] later in this volume).

All studies that will be discussed apply an average allocation rule, distributing the equipment energy consumption evenly among the total traffic volume over a certain period of time. We found no study focusing on marginal instead of average effects.

3 Existing Assessments and Methods Used

The two assessments mentioned above, [11] and [16], both used the same methodology: dividing an estimate of the overall US Internet energy consumption by the estimated total Internet traffic in the US. Since then, several more studies used the same approach, while other studies deployed different methodologies. According to the basic methodological approach they use, existing studies can be classified in two classes:

Top-down. According to [19], *top-down* analyses are based on two estimates: 1) the overall energy demand of either the entire Internet or a part of it (e.g., a country or a continent), and 2) the total Internet traffic of that region.

Bottom-up. By contrast, *bottom-up* approaches model parts of the Internet (i.e., deployed number of devices of each type) based on network design principles. Such a model combined with manufacturers’ consumption data on typical network equipment leads to an estimate of the overall energy consumption [15], which is then related to an estimate of the corresponding data traffic. Bottom-up studies can also use direct observations made in one or more case studies. These provide primary data on one or more of the following: deployed equipment, network topology and routing, power consumption of specific devices, and data volume passing through specific devices. Case studies present energy intensity values for specific cases, typically followed by a discussion of how the results can be generalized.

3.1 Top-Down Assessments

The earliest top-down assessments were introduced above: Gupta et al. published their results in 2003 [16], and Koomey et al. in 2004 [11]. The two studies used

the same statistical inventory data on ICT equipment in commercial buildings in the US for the year 2000 [20]. Despite building on the same inventory, they make different assumptions as to which devices belong to the Internet and thus yield distinct consumption results.

The assessment by Gupta et al. [16], taking into account Internet routers, WAN and LAN switches as well as Ethernet hubs, yields a yearly energy consumption of 6.05 terawatt-hours per year (TWh/a) for all networking devices in the US. In assessing the energy intensity of “the Internet” in the US, the study narrows the focus even more. It leaves campus devices (i.e., LAN switches and hubs) outside the calculation, and considers only the consumption of WAN switches and Internet routers, estimated at just 1.25 TWh/a in 2000.

Koomey et al. [11], on the other hand, consider not only the campus networking devices, but also data center devices – servers (10.2 TWh/a) and data storage (1.5 TWh/a) – as well as uninterruptible power supplies (5.8 TWh/a), leading to a total of 23.65 TWh/a. Furthermore, they multiply this result by 2 to account for overhead such as cooling and ventilation, leading to a total of 47 TWh/a, 37 times higher than the value in [16].

The two studies also use an identical source to estimate the US Internet traffic: data from the Minnesota Internet Traffic Studies (MINTS) by Andrew Odlyzko and colleagues. These data, published e.g. in [21], estimate a US Internet traffic in 2000 of 20,000-35,000 TB/month. Using the same traffic data, and consumption data different by a factor of 37, one could expect from the two studies energy intensity results differing by the same factor 37. This, however, is not the case. Koomey et al. [11] use the lower end of the Internet traffic data from [21] but complement it with traffic on other public data networks and private lines, leading to a total traffic of 348,000 terabyte per year (TB/a). The study thus yields an energy intensity of 136 kilowatt-hours per Gigabyte (kWh/GB) [11]. As for the other study, we can only speculate that Gupta et al. [16] misinterpreted the data in [21] as yearly instead of monthly traffic values. The study thus calculates with the range of 20,000-35,000 TB/a, which leads to an energy intensity of 0.128-0.225 Joule/Byte, or 38-67 kWh/GB. The correct energy intensity for the system boundaries used by [16], however, should have been twelve times lower because the yearly traffic estimate was actually twelve times higher. The corrected values for [16] are 3.2-5.6 kWh/GB (see Table 1).

An update for Koomey et al. [11] was published a few years later. The new assessment by Taylor and Koomey [22] referred to the year 2006. Estimating again the US Internet energy consumption and using three existing estimates of the US Internet traffic per year, the new study yielded as result the range 8.8-24.3 kWh/GB [22]. This 2006 estimate was yet again updated for the year 2008 in an article by Weber et al. [23]. For this period, the authors assumed that total Internet traffic increased by 50% per year, and that total Internet electricity use grew at a yearly rate of 14%, which had been the average global growth rate of data center electricity use between 2000 and 2005. These assumptions resulted in an average Internet electricity intensity of about 7 kWh/GB for 2008 [23].

The study by Lanzisera et al. [24] is another well-known top-down estimate. The analysis only includes networking equipment, excluding not only end devices but also the transmission lines. Estimating the total of both the US and the world networking equipment stock for 2008, the power of each device and their individual usage patterns, the article computed an annual electricity consumption of 18 TWh for all networking equipment in the US and of 50.8 TWh for the world. The study did not relate this consumption to traffic values to compute the Internet energy intensity. To make the result comparable with other studies, we divide it by an estimate for Internet data traffic for 2008 in order to calculate the energy intensity. According to Cisco’s “Visual Networking Index,” “global IP traffic grew 45 percent during 2009 to reach an annual run rate of 176 exabytes per year” [25]. We therefore assume a traffic volume of 121 exabytes (EB) for 2008. Using this value as a worldwide traffic estimate for 2008 yields an energy intensity of 0.39 kWh/GB for the world average.

3.2 Bottom-Up Assessments

The model-based approach has been used by Kerry Hinton’s research group at the University of Melbourne [12, 14, 15, 26, 27], as well as in [28] and [29]. Some of these studies are not directly comparable to the results of top-down assessments because they have different focuses such as analyzing only a part of the Internet transmission (e.g., [27]) or analyzing the Internet power consumption per subscriber and not per amount of data [26]. A few of these results may, nevertheless, be adapted to be made compatible with studies on Internet energy intensity. As we present in detail in [30], the very first assessment from the Melbourne group [26] yields an Internet energy intensity of 0.91-2.52 kWh/GB, depending on the estimate of worldwide Internet traffic used.

[14] provides a direct estimate of the energy intensity of Internet data transmission: 75 micro-Joule per bit ($\mu\text{J}/\text{bit}$), equal to 0.179 kWh/GB, at the access rates typical of 2008. As the authors point out, their result represents a lower bound or optimistic estimate in terms of energy consumption, because the model assumes only state-of-the-art equipment and ignores the fact that less energy-efficient legacy network equipment is still in use. They further state that they expect this energy intensity to drop in the near future to 2-4 $\mu\text{J}/\text{bit}$ with increasing access rates.

[12] puts forward a value of 2.7 $\mu\text{J}/\text{bit}$. This value corresponds to 0.0064 kWh/GB and represents the lowest value published thus far. This study, however, aimed to compare the energy demand of traditional computing with that of cloud computing. As the energy consumption of the access network is largely independent of the traffic and would leave the result untouched, the authors legitimately ignored it: “The access network does not influence the comparison between conventional computing and cloud computing. Therefore, it is omitted from consideration and is not included in our calculations of energy consumption” [12]. While this assumption stands to reason within the scope of the study, it can lead to misinterpretation if taken out of context, as we will discuss in the next section.

Finally, Schien et al. [29] used a network model to analyze the download of the UK newspaper “The Guardian,” as well as the download of a 640 second video from the Guardian’s video section. The newspaper’s html homepage was located on a server within the UK, while video and images were outsourced to a Content Distribution Network (CDN) and mirrored on several continents within the CDN’s network. Downloads from clients in Oceania, North America, and Europe were studied. Results showed that because of the CDN architecture, geographical distance played only a minor role; the energy intensities of the downloads from different continents were similar. For both the homepage and the video, the intensity was 8-9 Joules per megabit (J/mbit), which corresponds to approximately 0.02 kWh/GB. The study, while considering access, metro, and core parts of the Internet, did not account for the CPE.

In [31], this work was extended to include CPE and end devices, and to explore uncertainty and variability in assessments of digital services. In contrast to this chapter and to [30], which try to represent the existing variability in previous assessments, [31] estimates how uncertainty in energy intensity affects the overall result. Combining earlier results in a triangular distribution, the study arrived at a mean energy intensity for metro and core networks of 0.038 kWh/GB. For access network and CPE together, and excluding end devices, it provided a mean energy estimate of 0.019 kWh/GB.

In [32], Coroamă and Hilty present an assessment of a 40 megabit per second (Mbps) videoconferencing transmission of the case study introduced in [33]. For a system boundary that included network devices and optical fibers but no end devices, and making pessimistic assumptions in terms of energy consumption where specific data were not available, the study yielded an energy intensity of 0.2 kWh/GB for 2009. As we argued in [32], many characteristics of the study (such as an above-average number of hops) justify considering its result above-average in terms of energy intensity. This implies that the case-study result, when generalized, should be considered an upper bound for the average energy intensity.

The setting of this case study was such that no CPE or access network was distinguishable. The conference was held between a large conference center in Switzerland and a university campus in Japan. Both sites were directly connected to the metro network in the same way that Figure 1 depicts data centers. Yet, the edge routers on each side of the connection behaved similarly as CPE and the access network behave in the typical setting: they had a low load far from their capacity, and with an energy consumption that had to be allocated entirely or to a large extent to the case study.

4 Factors Influencing the Results

Results of the surveyed studies span a very wide range: from the 136 kWh/GB of [11] down to the 0.006 kWh/GB from [12], there is a spread of four orders of magnitude. In this section, we show how the distinct assumptions about the system boundary and further factors affected the results of the individual studies.

Table 1 summarizes the characteristics and the results of the studies presented. Special emphasis is given to their system boundaries. Our analysis revealed the following factors to be the most important influences on the result:

- The reference year of the study,
- the inclusion of data center devices within the system boundary,
- the inclusion of customer premises equipment and access network, and
- the inclusion of overheads such as cooling and redundancy equipment.

Each influencing factor is discussed in the following subsections.

Table 1. Estimates of the energy demand of Internet transmissions. The columns below “system boundary” show which parts of the Internet and of the end devices (as introduced in Figure 1) were accounted for by the individual studies. *CPE* is the customer premises equipment. *Core* stands for the metro and long haul Internet together – all studies consider both of them. *Fibers* are the optical transmission lines, and *DCs* stands for the equipment in data centers.

Study	Method	System boundary					Data on Energy intensity	
		CPE	Access	Core	Links	DCs	[year]	[kWh/GB]
[16]	Top-down	—	—	X	—	—	2000	38 – 67
[16], corrected	Top-down	—	—	X	—	—	2000	3.2 – 5.6
[11]	Top-down	—	—	X	X	X	2000	136
[22]	Top-down	—	X	X	X	X	2006	8.8 – 24.3
[23]	Top-down	—	X	X	X	X	2008	7
[24]	Top-down	X	X	X	—	—	2008	0.39
[26]	Model	—	X	X	X	—	2007	0.7 – 2.1
[14]	Model	X	X	X	X	—	2008	> 0.179
[12]	Model	X	—	X	X	—	2011	0.006
[29]	Model	—	X	X	X	—	2009	0.02
[31]	Meta-Analysis	X	X	X	X	—	2009	0.057
[32]	Case study	X	X	X	X	—	2009	< 0.2

Reference Year. An important part of the large differences can be explained by the year of reference of the individual studies, ranging from 2000 [16, 11] to 2011 [12]. The ICT sector is characterized by fast innovation cycles, and the equipment is becoming ever more energy efficient, needing less energy per amount of data being processed or transmitted. Taking [22]’s estimate that the energy intensity of the Internet decreases by 30% each year, this technological progress alone leads to a reduction by a factor of 50 over the period of 11 years covered by the studies.

System Boundary: Data Centers. The most important determining factor is the system boundary, in particular, whether or not data center devices (i.e.,

data storage and server-type devices running in server rooms or data centers) are viewed as part of the Internet. This decision has a large impact on the result: As shown in Table 1, for 2008, which is referred to by several studies, the two studies not including data centers result in energy intensities of 0.39 [24] and 0.179 kWh/GB [14] – factors of 18 and 39 below the 7 kWh/GB of the study that includes data centers [23], respectively.

The original statistical data from [20] (used by both [11] and [16]) supports this observation. While the consumption of core Internet devices was estimated at just 1.25 TWh/a, the consumption of storage devices and servers together was estimated at 11.7 TWh/a.

System Boundary: CPE and Access Networks. From the equipment inventories in [16] and [11] (presented in Table 1 of each paper), it is clear they do not consider customer premises equipment or access networks. This stands to reason: both rely on 2000 inventory data for ICT in US *offices*, not in homes. And the Internet had not yet exploded, reaching every home as it does today. In 2000, ISPs, access networks, and customers' modems and routers were not nearly as prevalent or important as they later quickly became. [22] and [23], on the other hand, explicitly include access networks and exclude CPE. Most of the other studies account for both, except [27] and [29] which do not include CPE in their assessment and [12] which does not consider the access network.

Not considering CPE or access networks, however, has a great impact on the result. When the otherwise dominating data centers are excluded from the calculation, numerous studies point out that CPE and access networks dominate the energy intensity of the Internet over core and metro networks. Figure 2, for example, shows the cumulated power consumption of the case study presented in [32]. The peripheral parts of the network clearly dominate the overall power consumption.

Both [14] and [15] also point out that for the peak access rates typical of 2008 and 2010, respectively, access networks dominate the power consumption of the Internet. Core and metro networks would become the dominating part only if access rates were to grow extensively and reach peak access rates in excess of 100 Mbps (for the energy efficiency of 2010 networking technology). For networking technology that becomes more efficient along the foreseeable trends, typical peak access rates would have to be over 1 Gbps for the core and metro networks to surpass the power consumption of the access networks. For the moment, as typical access rates are much lower, access networks use more energy than the core of the Internet and their exclusion profoundly changes the result.

The consumption of customer premises equipment is within the same order of magnitude as the access network but, as a general tendency, is slightly higher: according to [14], the access network needs 2.8W of power, while the CPE needs between 4-10W, depending on the technology used. [31] calculate with 2W per subscriber for the DSLAM of the access network and with 10 W for the CPE (5W for the modem and another 5W for the WiFi router). As with the access

network, the decision of whether or not to include the CPE decisively influences the result.

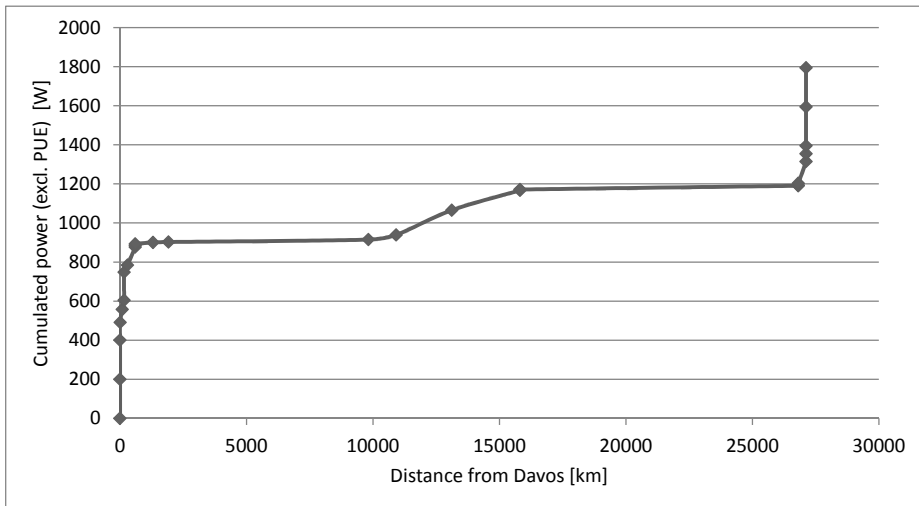


Fig. 2. Cumulated power demand for the videoconference case study [33] along a distance of 27,000km. The power demand of local network components, albeit relatively small, is allocated to a relevant extent to the case study’s traffic due to the low average load of the local components. These local components therefore clearly dominate the overall power consumption of the transmission. The large core Internet nodes and the transoceanic “data highways,” while utilizing a relatively large amount of power, typically have a switching/transporting capacity on the order of hundreds of GB/s, or even TB/s. Their contribution to the overall demand amounts to less than 1 Watt/Mbps in the case study. By contrast, the power demand of transcontinental links does contribute significantly to the overall consumption, due to their relatively low bandwidths.

Overheads: Cooling and Redundancy Equipment. The facilities (rooms, buildings) hosting ICT networking equipment and datacenters induce a power overhead due to non ICT-related consumption such as cooling or lighting. The measure widely used to account for this consumption is the *Power Usage Effectiveness* (PUE). The PUE is computed as a facility’s total power divided by the power needed to run the ICT equipment only [34]. As the former includes the latter, the PUE is larger than, or equal to, 1. The closer the PUE is to 1, the less power is “wasted” for activities other than information processing. The average PUE for datacenters nowadays is slightly lower than 2 [35] with a decreasing tendency; for Internet routers it was around 1.7 in 2009 [36]. Most studies do consider a PUE between 1.78 [29] and 2 [12, 14, 26, 32] in their calculations. No-

table exceptions are [16] and [24] that do not account for the PUE, i.e., they present results one would obtain for a $PUE = 1$.

Another overhead is induced by redundancy devices. Referring to routers, [31] notes that “devices are operated with at least dual redundancy in order to cope with failure.” [14], too, notes that a minimum of two uplinks are used for redundancy in metro and core networks. Several of the bottom-up studies (both model-based and case studies) account for the redundancy equipment via a redundancy factor of 2 [12, 14, 29, 32, 31].

5 Discussion

As the system boundary plays such a decisive influence on the end result, this section discusses possible best practices in defining the system boundary of the Internet for energy-related studies. It then shows how the access network and CPE differ from the more central parts of the network, and concludes with an assessment of the energy intensity of the former.

5.1 System Boundary

End Devices. [30] argues extensively against the inclusion of end devices (both user end devices such as personal or laptop computers, as well as servers and storage in data centers) within the system boundary for the energy intensity of the Internet. As we have shown above for data centers, including such devices can dramatically change the results. This, however, is inadequate not only as a matter of semantics; The concept of the Internet by definition does not include end devices but only the infrastructure connecting them. There are also practical arguments against including end devices, since that would yield results ill-suited and potentially misleading on most questions:

- The consumption in access, metro and core networks is largely independent of both the end devices and the application generating the traffic. Meaningful averages can thus be defined and estimated. By contrast, end devices such as desktop computers, smartphones or web servers have very different power demands. Different applications imply different sets of end devices: web browsing a server at a data center and a desktop, laptop or tablet computer; peer-to-peer file exchange two computers; and high-end videoconferencing two large LCD screens in combination with codec devices.
- Moreover, even for identical devices at both ends, distinct applications can induce very different consumption levels per amount of transferred data. While a peer-to-peer file exchange, for example, can use a bandwidth of several Mbps, a Skype voice call gets by with a bandwidth of only 60 kbps. Assuming exclusive usage of the two client devices, the low-bandwidth case induces a much higher energy consumption *per bit* at the terminal nodes due to the low utilization. This could lead to the seriously misleading conclusion that “the Internet” uses more energy per amount of data for applications

with lower bandwidth demand, in this case for a Skype voice call, than for a highly demanding file exchange.

Under these circumstances, with different devices and different applications inducing varying consumption in the end devices, it is unclear how these consumption levels could be aggregated into meaningful averages. It seems more meaningful always to assess network energy and the energy of end devices separately, and to add them up when needed – for example, for the assessment of the energy needs of a specific service [32].

CPE and Access Networks. The consumption of customer premises equipment and access networks, on the other hand, should always be considered. Unlike end devices, these devices have no stand-alone meaning. They only exist to connect end devices and thus semantically belong to the Internet.

Moreover, as shown in several studies [14, 15, 32], CPE and access network dominate the energy consumption of the Internet over metro and core network. Although this might not be true for services with a very high bandwidth usage (see next chapter [13]) and might permanently change in the future with increasing access rates [15], it is certainly the rule for the moment. It is advisable always to include these factors, even if they make an equal contribution in the cases under scrutiny. Although dropping such factors may simplify comparisons such as in [12], it will also change the absolute values which, if taken out of context, may lead to misunderstandings.

Overhead: PUE and Redundancy. The cooling and other types of overheads included in the PUE, as well as the redundancy equipment, provide support or safety functions in the Internet. Hence, their consumption should be accounted for.

As one of their few advantages over bottom-up assessments, top-down studies inherently include the redundancy equipment, as this is included in the stock inventories these are based on. Bottom-up analyses must account for them explicitly – as mentioned above, a factor of 2 for redundancy and a factor of 1.7-2 for the PUE are the values most studies use.

It must be noted, however, that both PUE and redundancy only apply to access, edge, metro and core equipment. CPE have no redundancy nor are they cooled. Whether redundancy and overheads are considered thus has a relatively low influence on end results that include CPE, and there is a risk of overestimation if one includes redundancy and PUE for the CPE as well (as we did in [32]).

Fibers. Whether the power of the amplifiers along the fiber cables is considered (as it is in most studies) or not (as in [16, 24]), also has only a marginal impact on the result. The relatively high power of dozens of kilowatts of transoceanic fibers gets divided by such a large amount of traffic that the contribution per

amount of data becomes negligible [32]. The consumption along fibers only becomes relevant for fibers with a low load such as the US transcontinental links in [32]. Such a case leads to a considerable allocation to relatively low amounts of traffic. Often, however, this is not the case. [14] argues that “the core optical transport (wavelength division multiplexed links) accounts for only a small fraction of the total energy consumed by the Internet.” But even when relevant, the consumption along the fibers was still smaller than the consumption of the access network [32].

5.2 The Challenge of Defining the Energy Intensity of the Internet

A basic methodological problem with the energy intensity of the Internet is that for some devices the energy consumption scales with the traffic volume and for other devices it scales along different dimensions, especially time of usage [31, 30]. The former is true of most networking devices in the metro and core network [31], while the energy consumption of devices in the access network and CPE (as well as end devices, which have been excluded from the analysis) usually scales with the time of usage and is largely traffic-independent [27].

To account for these differences, [31] recommends allocating the energy consumption of a device based on an approach that takes into account the limiting factor of the device – i.e., the factor that, if increased, would first limit the quality of service. If the limiting factor is in practice very hard to reach, the approach allocates energy according to the dimension that, if changed, results in the most significant change in energy consumption [31]. For the overall network energy intensity, [30] thus suggests using a combined approach that encompasses intensities defined as both energy per data and energy per time (i.e., power) where appropriate. We will elaborate on this approach in subsection 5.3.

Existing studies, however, do not differentiate between categories of devices when defining a metric for the energy intensity. Instead, each study defines the energy intensity for all devices along the same dimension: Most studies define the network energy intensity as energy per amount of data [16, 11, 22, 23, 29, 32]. With a partial focus on CPE and access networks, [31] defines the energy intensity for those devices as energy per time. Some studies with a focus on access networks, noting that access networks’ devices are always on and their consumption is thus both traffic- and time-independent, define the energy intensity as energy per subscriber [26, 27]. Finally, top-down studies can avoid the problem altogether by computing only the overall energy of the Internet and not relating it to any other dimension for a measurement of energy intensity (e.g., [24]).

We conclude this chapter by putting forward a formula for the computation of the energy intensity of the access network and customer premises equipment. The next chapter in this volume [13] addresses the metro and core networks, and proposes a formula for the energy intensity of those parts of the Internet. That energy intensity is defined as energy per data. Summing these two leads to the first formula for the energy intensity of the Internet that combines an energy per time component with an energy per data component and thus models reality more closely than previous work.

5.3 The Energy Intensity of CPE and Access Networks

In this subsection, we develop the formula for the energy intensity of CPE and access networks. We build on the analysis of [31], which analyzes the energy consumption of online multimedia services. For the consumption in access networks and CPE, that article puts forward the following formula (formula 9 in the article):

$$E_{AN} = t_s \left(\frac{P_{CPE}}{N_{CPE}} + \frac{P_{TU}}{N_{TU}} PUE_{Net} \right) \quad (1)$$

where

- E_{AN} is the energy consumption in the access network (including CPE in that article’s terminology) for the consumption of a given service,
- t_s the time of service consumption,
- P_{CPE} and N_{CPE} the power of all CPE taken together and the number of users connected to them, respectively,
- P_{TU} and N_{TU} the power of the access network devices and the number of users connected to them, respectively, and
- PUE_{Net} the PUE of the DSLAM, which typically requires cooling.

We start from this formula to define the energy intensity per unit time (i.e., power) of customer premises equipment and access networks. For more clarity, we consider access network and CPE separately and add them back together in the end. As we are interested in the energy intensity and not the total amount of energy, the time factor t_s disappears. Additionally, the energy intensity of the Internet is the average value for one Internet communication and thus always includes exactly one set of typical customer premises equipment ($N_{CPE} = 1$) – what “typical” means in this context will be addressed shortly.

With these observations, and renaming P_{TU} , N_{TU} , and PUE_{Net} to P_{AN} , N_{AN} , and pue_{AN} for more consistency, the intensities of the access network and the CPE, i_{AN} and i_{CPE} , are:

$$i_{AN} = \frac{P_{AN}}{N_{AN}} pue_{AN} \quad (2)$$

$$i_{CPE} = P_{CPE} \quad (3)$$

The trivial formula 3, however, ignores one important aspect that was not considered in [31] either: the energy used by CPE while idle, i.e., while not providing any service. This energy has to be somehow distributed among the services provided during a certain period [37]. We choose to distribute the idle energy consumption among the services provided over a given period of time proportionally to the time those services are active.

For the entire cycle over which meaningful averages can be built (i.e., a day, a week, a year), we define

- t_{On} the total time in which the equipment is on,

- $t_{U_{se}}$ the total time in which the CPE is in use, i.e., in which it is used for data transmission, and
- t_{Idle} the total idle time, when the CPE is on but not used ($t_{Idle} = t_{On} - t_{U_{se}}$).

The CPE consumes power for the time t_{On} but only provides services during $t_{U_{se}}$. Distributing the entire power on the services provided during $t_{U_{se}}$ needs the factor $t_{On}/t_{U_{se}}$ for extrapolation. With this, formula 3 becomes

$$i_{CPE} = \frac{t_{On}}{t_{U_{se}}} P_{CPE}, \quad (4)$$

which, because $t_{On} = t_{U_{se}} + t_{Idle}$, can also be written as

$$i_{CPE} = \left(1 + \frac{t_{Idle}}{t_{U_{se}}}\right) P_{CPE} \quad (5)$$

The energy intensity of the access network and CPE taken together is

$$i_{CPE\&AN} = \left(1 + \frac{t_{Idle}}{t_{U_{se}}}\right) P_{CPE} + \frac{P_{AN}}{N_{AN}} pue_{AN} = \frac{t_{On}}{t_{U_{se}}} P_{CPE} + \frac{P_{AN}}{N_{AN}} pue_{AN} \quad (6)$$

While such a generic formula is interesting from a theoretical point of view, for added practical relevance, the formula should ideally be parametrized with meaningful average values. These values should be based on the equipment and usage patterns considered typical and will change over time.

For the access network, the fraction P_{AN}/N_{AN} represents the energy intensity per subscriber. For ADSL2+, this was assessed as 3.4W by [14] and 2W by [31]. As [31] is the more recent analysis, we use its value. We further assume a PUE of 2, following the majority of studies.

For the power of the customer premises equipment (P_{CPE}), a few older studies consider only modems [14], while the more recent studies consider both modems and WiFi routers [24, 31, 37]. We follow [24] who assumes that only a few users use a modem without a WiFi router and that their number is comparable to those with multiple WiFi routers or WiFi repeaters. This is equivalent to assuming that 100% of users use both a modem and a WiFi router, either as two separate devices or integrated into an *Integrated Access Device* (IAD). Taking into account the US distribution of IADs versus modems with WiFi routers, [24] puts forward 7.1W as average CPE consumption for DSL and 9.5W for cable. The November 2013 Energy Star requirements for small network equipment [38] call for a base power of at most 5.5W for ADSL and 6.1W for cable IADs, respectively, allowing another 0.8W for fast Ethernet and WiFi, and 0.5W for the telephony functionality of DSL modems. Considering these numbers as well, both ADSL and cable IADs are required to be just below 7W. Allowing for the slightly higher consumption of two separate devices as well as for legacy equipment, we use $P_{CPE} = 8W$.

Finding data for the idle and usage times of modems is far more challenging and is by far the greatest source of uncertainty. A 2011 BBC study [39] found

that set-top boxes (the modems which deliver both cable TV and Internet connectivity) are on for 15.57 hours/day, but it did not address their usage time. A 2007 study found that in Europe, DSL modems are idle for 20 hours/day and in use for the remaining 4 hours/day [40]. Although this study is older, we use its assumptions. The study distinguishes between the on and idle state of modems, and we feel that the on-time of 24 hours/day from [40] better reflects reality than the 15.57 hours/day from the 2011 study [39]. With these assumptions, $t_{On}/t_{Use} = 6$.

With all these specific values, formula 6 leads to a “currently typical” value for the energy intensity (per time) of the access network and CPE of

$$i_{CPE\&AN} = 6 * 8W + 2W * 2 = 52W \quad (7)$$

The average value from formula 7 already includes the PUE of the access network and the idle consumption of the CPE, and can thus be used for quick assessments of energy consumptions in access networks and by consumer premises equipment. As mentioned above, however, the formula has a relatively large degree of uncertainty, especially due to the uncertainty of the idle time of the CPE.

6 Conclusion

We have shown that the energy intensity of customer premises equipment and access networks has to be assessed differently from the intensity of metro and core networks. We proposed a formula for the intensity of the former, both generically and parametrized with typical data for 2014. The next chapter [13] complements this work with a formula for metro and core networks. Taken together, the two chapters provide an assessment method for the Internet energy intensity that appropriately uses different allocation approaches for different parts of the network. Parametrized with typical values for 2014, this method can be used by practitioners for quick assessments of various Internet-based services.

References

1. Laitner, J.A.: Semiconductors and information technologies: The power of productivity. *Journal of Industrial Ecology* **14**(5) (2010) 692–695
2. Mattern, F., Staake, T., Weiss, M.: ICT for green – how computers can help us to conserve energy. In: *Proceedings of e-Energy 2010*, Passau, Germany, ACM (April 2010) 1–10
3. Hilty, L.M., Ruddy, T.: Sustainable development and ICT interpreted in a natural science context: the resulting research questions for the social sciences. *Information, Communication & Society* **13**(1) 7–22
4. GeSI: Smart 2020: Enabling the low carbon economy in the information age. A report by the climate group on behalf of the global esustainability initiative (gesi) (2008)

5. European Commission: Addressing the challenge of energy efficiency through information and communication technologies. Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions 241 final, Brussels (2008)
6. Pamlin, D., Pahlman, S.: Outline for the first global it strategy for co2 reductions. Technical report, World Wildlife Fund (2008)
7. Erdmann, L., Hilty, L.M.: Scenario analysis: Exploring the macroeconomic impacts of information and communication technologies on greenhouse gas emissions. *Journal of Industrial Ecology* **14**(5) (2010) 826–843
8. Hilty, L.M.: *Information Technology and Sustainability. Essays on the Relationship between ICT and Sustainable Development.* Books on Demand, Norderstedt (2008)
9. Hilty, L.M., Coroama, V.C., de Eicker, M.O., Ruddy, T., Mller, E.: The role of ICT in energy consumption and energy efficiency. Technical Report ICT-ENSURE: European ICT Environmental Sustainability Research; call identifier FP7-ICT-2007-2 (2009)
10. Coroama, V.C., Hilty, L.M.: Energy consumed vs. energy saved by ICT – a closer look. In Wohlgenuth, V., Page, B., Voigt, K., eds.: *Environmental Informatics and Industrial Environmental Protection: Concepts, Methods and Tools.* (2009) 353–361
11. Koomey, J., Chong, H., Loh, W., Nordman, B., Blazek, M.: Network electricity use associated with wireless personal digital assistants. *Journal of Infrastructure Systems* **10**(3) (2004) 131–137
12. Baliga, J., Ayre, R.W., Hinton, K., Tucker, R.S.: Green cloud computing: Balancing energy in processing, storage, and transport. *Proceedings of the IEEE* **99**(1) (January 2011) 149–167
13. Schien, D., Coroama, V.C., Hilty, L.M., Preist, C.: The energy intensity of the internet: Edge and core networks. In Hilty, L.M., Aebischer, B., eds.: *ICT Innovations for Sustainability.* Springer, Heidelberg, Germany (2014)
14. Baliga, J., Ayre, R., Hinton, K., Sorin, W.V., Tucker, R.S.: Energy consumption in optical IP networks. *Journal of Lightwave Technology* **27**(13) (2009) 2391–2403
15. Hinton, K., Baliga, J., Feng, M.Z., Ayre, R.W., Tucker, R.S.: Power consumption and energy efficiency in the Internet. *IEEE Network* **25**(2) (March 2011) 6–12
16. Gupta, M., Singh, S.: Greening of the Internet. In: *Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM '03).* SIGCOMM '03, New York, NY, USA, ACM (2003) 19–26
17. Teehan, P., Kandlikar, M.: Comparing embodied greenhouse gas emissions of modern computing and electronics products. *Environmental Science & Technology* **47**(9) (2013) 3997–4003
18. Hischier, R., Coroama, V.C., Schien, D., Achachlouei, M.A.: Grey energy and environmental impacts of ICT hardware. In Hilty, L.M., Aebischer, B., eds.: *ICT Innovations for Sustainability.* Springer, Heidelberg, Germany (2014)
19. Chan, C.A., Gyax, A.F., Wong, E., Leckie, C.A., Nirmalathas, A., Kilper, D.C.: Methodologies for assessing the use-phase power consumption and greenhouse gas emissions of telecommunications network services. *Environmental Science & Technology* **47**(1) (2012) 485–492
20. Roth, K.W., Goldstein, F., Kleinman, J.: Energy consumption by office and telecommunications equipment in commercial buildings – volume i: Energy consumption baseline. Technical report, US Department of Energy, Office of Energy Efficiency and Renewable Energy, Washington DC (2002)

21. Coffman, K.G., Odlyzko, A.M.: Growth of the Internet. In Kaminow, I., Li, T., eds.: *Systems and Impairments. Volume IV B of Optical Fiber Telecommunications*. Academic Press (2002) 17–56
22. Taylor, C., Koomey, J.: Estimating energy use and greenhouse gas emissions of Internet advertising. Technical report, *IMC²* (2008)
23. Weber, C.L., Koomey, J.G., Matthews, H.S.: The energy and climate change implications of different music delivery methods. *Journal of Industrial Ecology* **14**(5) (2010) 754–769
24. Lanzisera, S., Nordman, B., Brown, R.E.: Data network equipment energy use and savings potential in buildings. *Energy Efficiency* **5**(2) (2012) 149–162
25. Cisco: Hyperconnectivity and the approaching zettabyte era. Technical report, Cisco, San Jose, CA, USA (2010)
26. Baliga, J., Hinton, K., Tucker, R.S.: Energy consumption of the Internet. In: *Proceedings of the Joint International Conferences on Optical Internet and the 32nd Australian Conference on Optical Fibre Technology, COIN-ACOFT*. (2007)
27. Baliga, J., Ayre, R., Sorin, W.V., Hinton, K., Tucker, R.S.: Energy consumption in access networks. In: *Proceedings of the Optical Fiber Communication Conference and The National Fiber Optic Engineers Conference (OFC/NFOEC)*, Optical Society of America (2008)
28. Kilper, D.C., Atkinson, G., Korotky, S.K., Goyal, S., Vetter, P., Suvakovic, D., Blume, O.: Power trends in communication networks. *IEEE Journal of Selected Topics in Quantum Electronics* **17**(2) (2011) 275–284
29. Schien, D., Preist, C., Yearworth, M., Shabajee, P.: Impact of location on the energy footprint of digital media. In: *Proceedings of the 2012 IEEE International Symposium on Sustainable Systems and Technology (ISSST)*, IEEE (2012) 1–6
30. Coroama, V.C., Hilty, L.M.: Assessing Internet energy intensity: A review of methods and results. *Environmental Impact Assessment Review* **45** (February 2014) 63–68
31. Schien, D., Shabajee, P., Yearworth, M., Preist, C.: Modeling and assessing variability in energy consumption during the use stage of online multimedia services. *Journal of Industrial Ecology* **17**(6) (2013) 800–813
32. Coroama, V.C., Hilty, L.M., Heiri, E., Horn, F.M.: The direct energy demand of Internet data flows. *Journal of Industrial Ecology* **17**(5) (2013) 680–688
33. Coroama, V.C., Hilty, L.M., Birtel, M.: Effects of Internet-based multiple-site conferences on greenhouse gas emissions. *Telematics and Informatics* **29**(4) (November 2012) 362–374
34. Rawson, A., Pflueger, J., Cader, T.: Green grid data center power efficiency metrics: PUE and DCIE. Technical Report White Paper #6, The Green Grid (2008)
35. Stansberry, M., Kudritzki, J.: 2012 Data Center industry survey. Technical report, Uptime Institute (2012)
36. Moth, J., Norris, M.: GÉANT3 carbon footprint. In: *Terena Networking Conference*. (2010)
37. Schien, D., Shabajee, P., Wood, S.G., Preist, C.: A model for green design of online news media services. In: *Proceedings of the 22nd International Conference on World Wide Web (WWW'13)*, Republic and Canton of Geneva, Switzerland, International World Wide Web Conferences Steering Committee (2013) 1111–1122
38. EPA: Energy star program requirements for small network equipment. Technical report, US Environmental Protection Agency (November 2013)
39. Chandaria, J., Hunter, J., Williams, A.: A comparison of the carbon footprint of digital terrestrial television with video-on-demand. Technical Report WHP 189, British Broadcasting Corporation, London, UK (March 2011)

40. Nissen, N.F.: EuP preparatory study lot 6: Standby and off-mode losses – task 3: Consumer behaviour and local infrastructure. Technical Report TREN/D1/40, Fraunhofer Institute for Reliability and Microintegration, IZM, Berlin, Berlin, Germany (October 2007)