

Modeling and Assessing Variability in Energy Consumption During the Use Stage of Online Multimedia Services

Daniel Schien, Paul Shabajee, Mike Yearworth, and Chris Preist

Keywords:

industrial ecology
information and
communications technology (ICT)
Internet
life cycle assessment (LCA)
Monte Carlo simulation
online publishing



Supporting information is available on the JIE Web site

Summary

In this study, we use an improved, more accurate model to analyze the energy footprint of content downloaded from a major online newspaper by means of various combinations of user devices and access networks. Our results indicate that previous analyses based on average figures for laptops or desktop personal computers predict national and global energy consumption values that are unrealistically high. Additionally, we identify the components that contribute most of the total energy consumption during the use stage of the life cycle of digital services. We find that, depending on the type of user device and access network employed, the data center where the news content originates consumes between 4% and 48% of the total energy consumption when news articles are read and between 2% and 11% when video content is viewed. Similarly, we find that user devices consume between 7% and 90% and 0.7% and 78% for articles and video content, respectively, depending on the type of user device and access network that is employed. Though increasing awareness of the energy consumption by data centers is justified, an analysis of our results shows that for individual users of the online newspaper we studied, energy use by user devices and the third-generation (3G) mobile network are usually bigger contributors to the service footprint than the datacenters. Analysis of our results also shows that data transfer of video content has a significant energy use on the 3G mobile network, but less so elsewhere. Hence, a strategy of reducing the resolution of video would reduce the energy footprint for individual users who are using mobile devices to access content by the 3G network.

Introduction

The climate-change impact of information and communications technology (ICT) has been studied by the academic community for some time, for example, most recently by Malmodin and colleagues (2010) and Weber and colleagues (2010a); it is also attracting increasing interest from the public (Greenpeace 2012). Attributional life cycle assessment (LCA), the quantifying of environmental impacts resulting from the creation, use, and disposal of a product or service, has played a key role in this analysis. A recent study by Weber (2012) of the sources and

extent of variability in LCAs of a server computer has identified the use-phase energy consumption to be the most uncertain. The observation that the variability in the use phase is very high has also been made for end-user devices (Beauvisage 2009).

The International Reference Life Cycle Data System (ILCD) Handbook distinguishes between *variance* as the degree of stochastic uncertainty in a single process within an LCA and *variability* as the single representation of multiple processes and systems with differing impacts (European Union European Union Joint Research [EU JRC]–Environment and Sustainability 2011). When faced with variability in a process

Address correspondence to: Daniel Schien, University of Bristol, MVB, Woodland Road, Bristol, BS8 1UB, UK. Email: daniel.schien@bristol.ac.uk, Web: www.cs.bris.ac.uk/home/schien/

© 2013 by Yale University
DOI: 10.1111/jiec.12065

Editor managing review: H. Scott Matthews

Volume 17, Number 6

flow, an LCA practitioner has a choice between a more detailed process model, which will require commensurately more acquisition of associated data, or a less detailed model, which takes an “average” or “prototypical” process and data set, concealing the underlying variability. The latter approach has the advantage of being easier, but may reduce the accuracy of the assessment and hide potential opportunities for improvement.

In this article, we present an analysis of energy use in the delivery and consumption of online digital news content per individual user, with particular reference to the variability in energy use that occurs in the delivery of the service. To do this, we have developed a model of digital service delivery that is significantly more detailed than earlier art, and gathered associated data from new primary and secondary sources. The model covers the use phase of the life cycle of digital services, including the dynamic creation and delivery of content from a distributed set of data centers, transmission of content through routers, switches, and cable (“the Internet”), delivery of content through the user access network, and consumption of content on user devices. Regarding the access network, we consider connectivity by digital subscriber line (DSL) modems in combination with on-premise wireless network (WiFi) and third-generation (3G) mobile networks. These are representative of the majority of access modalities in the Organisation for Economic Co-operation and Development (OECD) countries. In the United States, wired Internet access by cable is more popular than DSL. For the purposes of this text, the two are comparable, even though, on average, the power consumption for cable Internet access is higher than for DSL (Aleksić and Lovrić 2010). The model we have developed, and much of the associated data, is of general applicability to digital services. Our analysis is focused on a specific case study: the provision of the multimedia Web site by Guardian News & Media Limited (GNM). The majority of visits to GNM’s Web site come from the UK, but there also is a significant international audience. The functional unit we adopt in this article is 10 minutes of content browsing, and we explore the impact of variability in content type (video and Web page), end-user device (desktop, laptop, tablet, and smartphone), access network, geographical location, and browsing behavior on overall energy use. Regarding variability in browsing behavior, we consider the impact of varying the speed of changing between Web pages. We present results for a number of scenarios exploring the impact of this variability and variance within each scenario modeled, using a Monte Carlo approach, and identify how the significance of different components of delivery and consumption alters between scenarios.

This work contributes to state-of-the-art use stage energy and greenhouse gas (GHG) footprinting of ICT in a number of ways. First, the model of energy use by digital services we present is both more detailed than previous studies and more complete in its ability to capture and distinguish between different usage scenarios. Second, using the primary and secondary data we have gathered, we present the first detailed analysis of a diverse set of scenarios for digital media consumption, and thus update and complement existing studies of the environmental footprint of digital media. Third, we present novel methodolog-

ical advances in the modeling of digital services—specifically, the use of Monte Carlo simulation to draw from alternative subscenarios, rather than error estimates, and our discussion of the appropriateness of different allocation methods for different components of digital delivery. Allocation is an important consideration in this study because some of the equipment is used by multiple users and some of the equipment is used for multiple purposes by the same user, and the energy used by the equipment has to be reasonably apportioned to the functional unit of the study. Fourth, we present initial results that can be used for simplified analyses of digital services.

Our analysis of a large set of typical scenarios of digital online media consumption provides results that can be compared to several previous studies. The existing studies vary in the degree to which they apply *bottom-up* or *top-down* models for the majority of the life cycle processes. Our own model and others (Moberg et al. 2010; Chandaria et al. 2011; Williams and Tang 2011; Baliga et al. 2009) apply bottom-up models in as far as they calculate the energy consumption from the additive impacts for the most impactful processes per functional unit. A bottom-up model of the energy consumption for Internet delivery, as applied by Baliga and colleagues (2009), for example, sums the proportional energy consumption by each network device that plays a role in a typical route between two end points. On the other hand, in studies that mostly apply top-down models, such as Taylor and Koomey (2008) and Malmodin and colleagues (2012), the measured or estimated total aggregate impact of an entire (sub)system is related to the total number of functional units provided by the system. In Malmodin and colleagues (2012), for example, the total impact from consumption of electrical energy of a Swedish Internet network is related to the estimated data volume transmitted through this network by dividing the former by the latter to give the energy consumption per unit of service. The result of both top-down and bottom-up allocation is an average value for impact per functional unit, yet only the bottom-up model contains data about the elementary life cycle processes.

Our work draws from these studies, but goes beyond them in several ways. First, whereas they use aggregate data and assumptions regarding the average or typical user, we model the variability explicitly in a number of scenarios. Second, we use a Monte Carlo approach to account for both variance and variability within each scenario. Third, we present a principled approach to the challenge of allocation as applied to digital services and make use of it in our model.

To our knowledge, ours is also the first study that relates the energy footprint of 3G networks to a functional unit of a media service. Although Scharnhorst et al. (2006) and Stutz and colleagues (2006) both provide an LCA assessment of a 3G cellular wireless network, their functional unit is that of a year’s mobile service. At this level of aggregation, their results cannot be related to a single media service similar to our functional unit. Toffel and Horvath (2004), on the other hand, analyze the energy footprint of downloading newspaper content to a handheld reader by a mobile network, but do so for a two-generation (2G) wireless network. They reference the

total energy consumption of the wireless network based on an LCA study from 1999 and relate it, in top-down fashion, to the total number of subscribers in the network. The most widely referenced top-down model of the energy footprint of Internet data transfer is presented by Taylor and Koomey (2008) in a study of the impact of Web advertisements, and we compare its results to those derived from our model in the discussion. The findings by Taylor and Koomey (2008) are applied by Teehan and colleagues (2010) in a top-down model that is used to analyze the total energy consumption in the United States for a variety of tasks, assuming that user behavior in the United States is similar to survey data from France 2005–2006. They do not capture the wide variability in the energy footprint resulting from individual user behavior and variability in the power consumption of devices.

This article is structured as follows: In the next section, we present a conceptual system model and identify suitable approaches for allocating energy usage to the functional unit for individual subsystems. This section is followed by a description of the most significant model parameters. In the subsequent section, we present the results of the analysis. In particular, we demonstrate variance and variability between scenarios of device, service type, and access network combinations. We close with a discussion and conclusion in the final section.

Models

Broadly, four categories of devices are involved during the use phase of a digital service. First, data centers consisting mainly of servers, but also of networking and storage infrastructure, provide the service content. In the case of GNM, this content is split between origin servers belonging to the organization itself and a number of third-party server data centers. These third-party servers either provide additional content, such as advertising, or belong to *content delivery networks* (CDNs), which cache content in different regions around the world to improve service performance. Second, the devices that make up the *edge and core networks* of the Internet, also called the Internet network, transport content from its sources and the third-party data centers to the end user. Third, the *shared access network* (sometimes in conjunction with a customer premise access network), which is DSL, 3G, or DSL in combination with WiFi, links the user's device with the Internet network. The fourth and final category is the user's device itself.

Figure 1 depicts the devices and energy flows that are the subject of this article. Not included in our assessment is the impact of energy embodied in the devices (the energy during manufacturing and transport as well as the end-of-life stage of devices), energy required for software development activities, and energy used during editorial work. Our collaboration with GNM allows us to use primary data for many, but not all, processes.

Our functional unit is 10 minutes of browsing, during which we assume the user issues one or more requests for content. Each such request involves opening an individual uniform resource locator (URL) with a Web browser. The energy consumption

for each individual request is the sum of the consumption by the four subsystems in the delivery model. The energy footprint for the functional unit is the sum of the energy consumption of all requests issued during the time of the functional unit. Not included in this energy footprint is the energy consumption of other life cycle phases, notably the manufacturing of the devices.

In the remainder of this section, we will present the model in detail, starting with the methodology for allocating energy used by the shared information technology (IT) infrastructure: the origin servers of the content provider; third-party servers; and the Internet. We will then look at the energy consumption of the network connection between the servers and the end user. Finally, we will discuss the energy consumption of the end-user equipment itself.

Allocation Approaches for Digital Products

A key methodological decision in LCAs, which can significantly impact the assessment result, is made during the *allocation* of environmental burden, which is defined as the act of “partitioning the input or output flows of a process or a product system between the product system under study and one or more other product systems” (ISO 2006, 4). In this section, we consider alternative allocation approaches possible for digital products and discuss their appropriateness in different situations. In the case of digital products, allocation is necessary for two reasons. First, equipment use may be shared between multiple users, such as a content server providing Web pages to many people, or a DSL multiplexer providing broadband connections to a number of households. Second, equipment may be used for multiple services, such as a physical server running multiple virtual machines, or a household laptop providing Web access as well as e-mail, playing music, and many other applications. In this study, allocation of energy use between multiple services running on an end-user device, such as the household laptop in this example, was not conducted; all the energy use for the device was assumed to be used for accessing video content or nonvideo content during the 10-minute functional unit.

The ILCD Handbook distinguishes between two approaches to allocation. The preferred approach, *physical causality*, allocates burden based on the share of some physical (or other) flow that is directly related to the environmental burden generated (EU JRC–Environment and Sustainability 2011). The second approach is to use some other relationship, such as economic activity.

We consider three approaches to allocation for digital services:

1. **Data flow:** In the case of digital services, there is no clear “physical” flow to study, but there is a flow of data. Allocation can take place based on the share of data passing through an energy-using device. This approach is adopted by Lee and colleagues (2011) and Baliga and colleagues (2009), which considered both energy consumption by Internet routers and home access network devices.

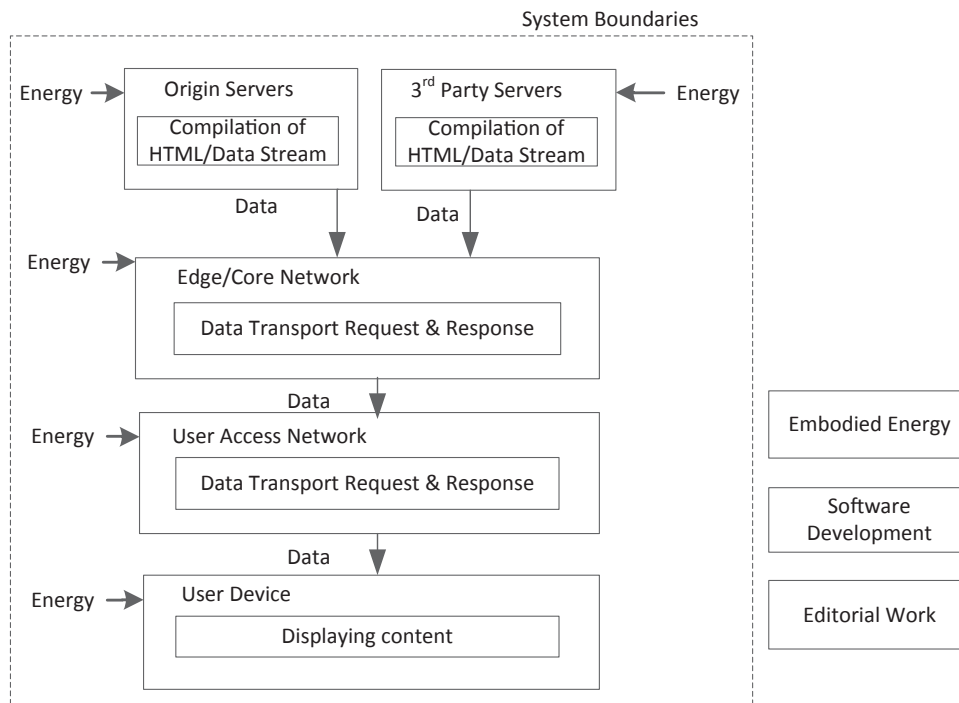


Figure 1 System boundaries. HTML = HyperText Markup Language.

2. Number of users: If a device is shared between services offered to a number of users, the energy can be allocated equally to each.
3. Service “attention” time: If a user is using a given device for a number of services, energy usage by the device can be allocated based on the amount of time the user spends using the different services.

None of these cases correspond to the preferred approach of physical causality because energy usage of devices does not vary directly proportionally to data flow, users or services being used. We discuss this in more detail for each device in subsequent sections. To determine which approach to use, we adopt a principle of allocating based on which of these is the *limiting factor* to device usage—namely, the factor which, if increased, would first limit or degrade the quality of service. In the case of most network devices, this is usually bandwidth. Similarly, video, audio, or image content servers, such as those used by CDNs, conduct relatively little computation and are limited by their capacity to transmit data at speed. In the case of a DSL multiplexer, on the other hand, which is used to provide access to the Internet for a number of premises, the limiting factor is the number of connections it can provide and therefore the number of users it supports. Also, in the case of a Web server, the limiting factor is the computational power required to construct pages, rather than the speed of data it outputs. Finally, in the case of a user device, the limiting factor is usually—but not always—the user’s attention: The device could easily run more applications, but the user would only be able to make use of a limited number at a time.

In addition to the limiting factor, we allocate along that dimension, which, if changed—given current levels of typical utilization—would result in the most significant change in the energy consumption. For user devices, for example, one such choice is between data volume received and time of service consumption. In the case of most online multimedia services, and, in particular, online news, a reduction or increase in the device operation time will result in a much greater change in the energy consumption of the service than a change in the data volume transferred. In summary, we allocate by duration of use for user devices and network devices inside the customer premises and by data volume for third-party servers and publicly shared network devices. In the origin data center, we allocate by number of services. In the “Allocation Approach” section of the supporting information available on the Journal’s Web site, we provide a detailed formalization of our model and allocation approach for each system part.

Model Parameters

In the following section, we present a parameterization that allows calculating an energy footprint for a media service provided by GNM. The main model parameters are power consumption values, throughput capacity and operational use time of servers, and network and user devices. We will look at each parameter in turn and discuss possible sources for values and variability.

First, we consider the GNM data center. GNM operates virtualized blade servers that are arranged in a tiered array of 25 blades, each of which has a power consumption varying

between 140 and 300 watts (W) (Beckett and Bradfield 2011) from base to peak load. Apportioned to the relative number of monthly visitors, this number of servers is similar to that reported for the German online magazine, stern.de (heise Verlag 2011). The average number of pages served per second from all servers ranges from approximately 40 to 200 with a trough during what are the early morning hours in the UK (Wood 2012). Also, from internal data, we know the load of the servers typically varies between 15% and 30%, which is in agreement with typical utilization of data centers as presented by Barroso and Hölzle (2007). Network and storage equipment in the data centers are often shared between independent parts of organizations, which necessitates allocation decisions. In order to assess the impact from varying the allocation of this equipment, we sample this contribution from a triangular distribution between a minimum of 10%, a maximum of 30%, and a mode value of 15% overhead each. For the additional overhead from cooling and power transformation PUE_{origin} , we apply a distribution, based on the values from Bertoldi (2010), between 1.25 and 2.86 (see figure S12 in the supporting information on the Web for more details). Data transfer is measured in bits (b) or bytes (B) (1 byte = 8 bits). During assessments of digital products and services, it is common to estimate the energy consumption by networks and servers from the energy intensity of data, which is expressed in Joules per bit (J/b) or the equivalent kilowatt-hours per gigabyte (kWh/GB). In the existing literature, values of energy have been presented in units of Joules (J) and watt-hours (Wh) and energy intensity in J/b, Joules per megabit (J/Mb), or kWh/GB. In order to make it easier to compare our results with past studies, we provide numeric results for energy intensity in this text in J/Mb and kWh/GB. GNM commissions several CDNs, of which Akamai delivers the largest data volume, for pages without video content. For September 2010, Akamai reports a quarterly GHG footprint for sustained bandwidth of serving data of 1.2 kilograms carbon dioxide equivalent (kg CO₂-eq)/megabit per second (Mbps) (Akamai 2010). Based on this value and an assumed average carbon intensity of electricity supplied to their globally distributed servers of 0.62 kg CO₂-eq/kWh, we estimate the energy consumption per data volume transferred to be 2.0×10^{-3} kWh/GB (0.89 J/Mb). Details for arriving at this estimate, including a rationale for the lower and upper bound, are provided in the section titled “Energy Intensity Akamai CDN Servers” in the supporting information on the Web. Our estimated value is a 15th of the value reported by Google for YouTube servers, which is 2.96×10^{-2} kWh/GB (13.3 J/Mb), assuming a bit rate of 900 kilobits per second (kbps) (Google 2011) and more than two times higher than the value of 8.89×10^{-4} kWh/GB (0.4 J/Mb) assumed by Chandaria and colleagues (2011) for CDN servers streaming video data. Given the discrepancy between those values, we do not make additional assumptions about the possible variation of energy efficiency from changing utilization at different times of day.

Our network model for edge and core networks, including fiber-optic equipment, is defined in equation (5) in the “Allocation Approach” section in the supporting information on the

Web. It is a bottom-up model that needs to be parameterized with the number of hops in the connection between a server and the user. The alternative modeling approach of top-down modeling, in the case of the core network, estimates the energy efficiency of a single service by taking the total energy consumption of the entire network and apportioning it to all services delivered. Both approaches arrive at different values, and we discuss the discrepancy between top-down and bottom-up models in more detail below. During the Monte Carlo simulation, we evaluate the impact of these different assumptions by sampling from a triangular distribution. The minimum and mode values of 0.010 and 0.023 kWh/GB (4.5 and 10.5 J/Mb) are based on a bottom-up model, which we present in detail elsewhere (Schien et al. 2012). For the maximum value, we apply a value of 0.080 kWh/GB (36 J/Mb) based on a top-down model (Malmodin et al. 2012).

The power consumption per subscriber by the DSL access multiplexer (DSLAM) was assumed to be approximately 2 W by several studies (Aleksić and Lovrić 2010; Lee et al. 2011; Baliga et al. 2009). In addition to that earlier research, our own measurements also find that the power consumption of DSL modems is approximately 5 W. Separate WiFi routers have a similar power consumption, as measurements for the new Energy Star rating of small network equipment indicate (Energy Star 2012a). We assume a home setup of a single wireless router and single DSL modem with both being actively used for the same time as the end-user devices.

The energy efficiency of cellular wireless networks varies strongly with the allocation of the energy consumption for cell subscription required to receive calls. In our Monte Carlo simulation, we apply a triangular distribution with a maximum value of 0.73 kWh/GB (328 J/Mb), which is based on uniform allocation of total power consumption to data packets. For the mode and minimum of the distribution, we apply values that are based on the allocation of the instantaneous power consumption of the base station and the data rate per subscriber. We apply 0.12 kWh/GB (54 J/Mb) as the mode and 0.030 kWh/GB (13 J/Mb) as the minimum, which are based on a power consumption of 460 W per subscriber and a data rate of 11 and 45 Mbps for High Speed Packet Access, a 3G cellular network evolution from Deruyck and colleagues (2010). The minimum and mode value also include an overhead of 1.3 to account for the energy consumption of the remaining parts of the cellular network in addition to the mobile base station. This is based on the yearly energy consumption of 4.177 gigawatt-hours (GWh) for the whole network of the German mobile network operator, Vodafone, for operation of 224,000 base stations (Vodafone 2011), resulting in an allocated average power draw of 2,129 W per base station, which is approximately 30% more than the average nominal power consumption of base stations operated by Vodafone Germany (Zwemke 2012). We assume that the energy efficiency of mobile networks is similar between countries of the OECD, although no systematic study exists.

We distinguish between the following classes of end-user devices: smartphones; tablets; laptops; and desktop computers. For the laptop and desktop computers (including monitors),

distributions of power consumption are based on data from the Energy Star measurements (Energy Star 2011, 2012a). These extensive lists contain power measurements of several thousand energy-efficient devices that were awarded the Energy Star rating. They do not represent the relative popularity of these devices.

On top of the power consumption in active idle mode (non-standby), the execution of computer programs introduces a dynamic portion of power consumption relative to their utilization of device components, such as the central processing unit, network interfaces, or disk drives. The relative and absolute magnitude of this dynamic power consumption depends on several parameters: for example, the specific device and—in the case of browsing online news—the amount of JavaScript embedded in a page or the algorithm used to decode a video data stream to images (codec) used. Yet, systematic research of the influence of these parameters does not exist. We conducted a scoping experiment on a single, modern Energy Star-rated laptop and found no statistically significant variation from idle power when browsing text; hence, we set $\alpha = 1$ in equation (10) in the “Allocation Approach” section in the supporting information on the Web. In the consumption of video, however, the same experimental setup found a significant increase in power consumption. In our model, we assume the power consumption of devices increases by 15% when video is viewed; hence, we set $\alpha = 1.15$, which is similar to values reported in Somavat and colleagues (2010).

Based on GNM data, we apply an empirical distribution of the duration that users spend reading or watching the news. The distributions can be found in the “Sampled Distributions” section in the supporting information on the Web in figures S9 and S10. The distribution of news reading is heavy tailed and has its average at 1.5 minutes. Video content is being watched for approximately two minutes, on average.

Method

We have conducted simulations of a number of different scenarios of users accessing GNM digital services to explore the impact of variability on the energy footprint. Each scenario has a specific user device, access network technology, and service type associated with it. The service type can be either a HyperText Markup Language (HTML) Web page, including text, images, and gif animations, or HTML with embedded video content. The functional unit for either service is 10 minutes of browsing. The average duration spent per Web site is 90 seconds for reading text and 121 seconds for videos. The precise value of the duration per page is randomly sampled during each iteration of the Monte Carlo simulation. We simulate the most popular user-device technologies: smartphones; tablets; laptops; and desktop personal computers (PCs). We exclude some exotic combinations of local-access network types and user devices, such as the combination of wired connection of a phone or a tablet to a DSL modem, but we include the simulation of smartphones connecting to a WiFi. We only simulate

mobile network access for phones, tablets, and laptops. We do not consider mobile access with a laptop in combination with an external screen.

For each scenario, we conduct a Monte Carlo simulation of 1 million runs. This figure was determined by experiment to ensure convergence of average energy consumption of each subsystem (servers, networks, and user devices) to within 1% overall for the same scenario. Each run draws from distributions based on both variance and variability within a given scenario. Variance is handled in the usual way, by using distributions around a mean value based on data quality factors and correlation between different secondary data sources. We give details of the distributions used in the Supporting Information on the Web. Our approach to handling variability is novel; instead of a statistical distribution around a mean, we make random draws from a representative population of discrete observed values. For example, in the case of time taken to view a given Web page or video, we make draws from a distribution generated from actual usage data provided by GNM. Similarly, the geographical location of the end user and the time of day accessed draw from distributions of GNM primary data. In the case of end-user device power consumption, we make draws from a population of potential device models, each with an associated power consumption. Again, details are provided in the Supporting Information on the Web, except where commercially confidential. In all cases of variability represented in this way, our model allows values to be fixed to give results for a specific subscenario—for example, modeling a user accessing the service from Boston with an iPhone at 6:00 p.m. GMT.

Results

We now present results of our simulations for the scenarios we explored. We start with the presentation of average absolute values of energy consumption, broken down according to contributions by different system components. We then show which components affect the total allocated individual power consumption most and explore this influence in more depth.

Figure 2 shows the average energy consumption for the different scenarios. Error bars indicate the 25th and 75th percentile of the total energy used for each device/access network scenario's sample distribution for ten minutes of viewing either text or video content. The charts in this figure share the same vertical scale. Energy consumption varies widely between the scenarios, highlighting the need to take the particular combination of device types, local access networks, and service type into account and precisely estimate their share of use by an audience when assessing digital products or services.

Average energy consumption for consuming video is higher than the energy for consuming text. Not surprisingly, in the case of reading articles, the scenarios with a desktop computer arrive at the highest total energy consumption, with an average of 96 kilojoules (kJ) for ten minutes of reading. The least amount of energy is consumed when reading articles for ten minutes with a smartphone over a 3G cellular wireless connection




Figure 2 Average energy consumption in kilojoules (kJ) for ten minutes of news consumption by system components for selected combinations of access network and user device type. Shared access network is either the 3G mobile network base station or a digital subscriber line (DSL) access multiplexer (DSLAM). WiFi = on-premise wireless networks (customer premise access network together with a DSL modem); 3G = third-generation mobile networks; PC = personal computer.

with a 9-kJ energy consumption. In the case of consuming ten minutes of video content, the scenarios with a cellular wireless connection consume the most energy, with the 3G network alone contributing 121 kJ, which is 83% of the average total energy consumption for this scenario for smartphones and tablets. Energy consumption by the Internet network is smaller, but—at 15 kJ—not insignificant. Energy consumption of scenarios for consuming video with hand-held devices (smartphones and tablets) is dominated by the Internet network when accessing content by WiFi and by the 3G network when video is consumed over the 3G network. Energy consumption for consuming text with hand-held devices is dominated by both the origin data centers and type of access network being used to access content. In contrast, except for viewing video over the 3G network on a laptop, energy consumption of scenarios with laptops and desktops is primarily dependent on user-device power consumption. Of particular significance is the energy consumption by cellular wireless networks in the case of viewing video, which outweighs all other subsystems. Total energy consumption varies substantially between scenarios. The data center where the news content originates consumes between 4% and 48% of the total energy consumption when news articles are read and between 2% and 11% when video content is viewed. Similarly, user devices consume between 7% and 90% and 0.7% and 78% for articles and video content, respectively. The full numeric values are presented in the “Numerical Results” section of the supporting information on the Web in tables S1 and S2.

Figure 3 shows histograms of selected sample distributions within the 2.5th and 97.5th percentile. The y-axis in this figure is the number of times a given energy consumption was calculated during the Monte Carlo simulation. Considering the nonvideo scenarios, those with a smartphone show a much smaller degree of variability, compared to those with PCs. This is mainly a result of the larger variability in the power consumption of PCs, in comparison to smartphones or tablets. The histograms also show that there is a clear distinction in energy use between lap- and desktop computers, with laptops only using more energy than desktops in 2.13% and 2.42% of the scenario samples for text and video, respectively.

Following Weber (2012), we use a Spearman rank analysis over several scenarios, varying the access network type and the service type, to determine how different parameters of the model affect the final result. A Spearman rank analysis generates coefficient values (ρ) between +1 and -1, with +1 indicating perfect direct correlation, -1 indicating perfect indirect correlation, or anticorrelation, and 0 indicating no correlation.

Figure 4 shows the values of those correlation ranks between scenarios of consuming video or text, depending on the local access network type. The graphs only show the variables with an absolute value in at least one scenario of greater than 0.05. These are mainly the aggregate variables that represent the energy consumption of subsystems. The remaining variables are listed with parameter values in the section on “Model Parameter Values” in the supporting information on the Web. The markers




Figure 3 Histograms of total energy consumption in kilojoules (kJ) by all subsystems for selected combinations of access network and user-device types and consumption of ten minutes of news without video content and consumption of 10 ten of video content. PC = personal computer; UMTS = universal telephone mobile system (third-generation mobile network); WiFi = on-premise wireless network.

in 4a show the Spearman rank values for scenarios with user devices applying DSL plus WiFi access, namely, smartphones, tablets, laptops, and desktops, whereas the markers in figure 4b show the correlation values for user devices applying 3G mobile access, namely, smartphones, tablets, and laptops. For example, in the scenario of consuming video over WiFi in figure 4a, the Spearman rank for the correlation between device power and total energy consumption in the lowest energy-consuming smartphone is 0.04 and in the scenario of the most energy-consuming desktop it is 0.92. Importantly, for video services consumed on hand-held devices, total energy consumption depends most strongly on the access network, rather than the user device. Also, for the video scenarios, the access network is much more relevant to total energy usage than it is for the text scenarios. Not surprisingly, the correlation between server utilization, expressed by pages per second, and the total energy consumption is higher in the ranks for 3G mobile access than in those for WiFi, because the former do not include desktop scenarios. Negatively correlated coefficients indicate inverse correlation of components; for example, the lower the utilization of the origin servers (indicated in the row “Page Request per Second All Users” referring to the total number of page requests hitting the origin servers, as opposed to the number of page requests issued by the individual user, which is expressed by the variable “Duration Reading/Watching Video”), the higher the total energy consumption. Also, the duration spent viewing each of the possibly multiple Web pages visited during the ten minutes of the functional unit appears negatively correlated with the total energy footprint because it is inversely proportional to the number of repeated page requests submitted within the ten minutes (indicated with negative values in the column “Duration” in figure 4a and 4b). When connecting with 3G, the shared access network has a stronger impact on the total power consumption than the home networking equipment when connecting with WiFi. In the case of watching video, data volume is directly dependent on duration of service consumption. When consuming

text only, data volume has much less impact on total energy consumption.

Discussion

Analysis

In this section, we discuss our results in the context of previous work modeling the use-phase energy consumption of digital media. We compare the quantitative results with those of other investigators, where there is overlap of the models, and explore the reasons for differences. The energy per bit varies between cellular wireless and wired access network connections such as DSL. For comparison, we also calculate the energy intensity per data volume for DSL and WiFi access networks, although, for this subsystem, energy is allocated by time in our model. The energy consumption of edge and core networks is calculated as the product of energy intensity and data volume. Our average energy intensity for the edge and core network (the Internet network) for either text or video is 0.038 kWh/GB (17 J/Mb), which is the average value of a triangular distribution with the minimum, mode, and maximum values, as defined above, equal to 0.010, 0.023, and 0.080 kWh/GB (4.5, 10.5, and 36 J/Mb). According to the same approach, our value for average energy intensity of 3G wireless networks is 0.293 kWh/GB (132 J/Mb). Energy consumption by shared and customer premise equipment for DSL (DSLAM, DSL modem, and WiFi router) is calculated as the product of power consumption of devices and time of use. A value for the effective energy intensity can be calculated as the energy consumption over the total data volume transferred. Energy intensity of data over DSL and WiFi differs between consuming text or video: 0.019 kWh/GB (8.49 J/Mb) for video over DSL and WiFi, compared to an average total consumption of 0.733 kWh/GB (330 J/Mb) for text. Williams and Tang (2011) allocate power consumption only for the duration of the data transfer, resulting in energy consumption

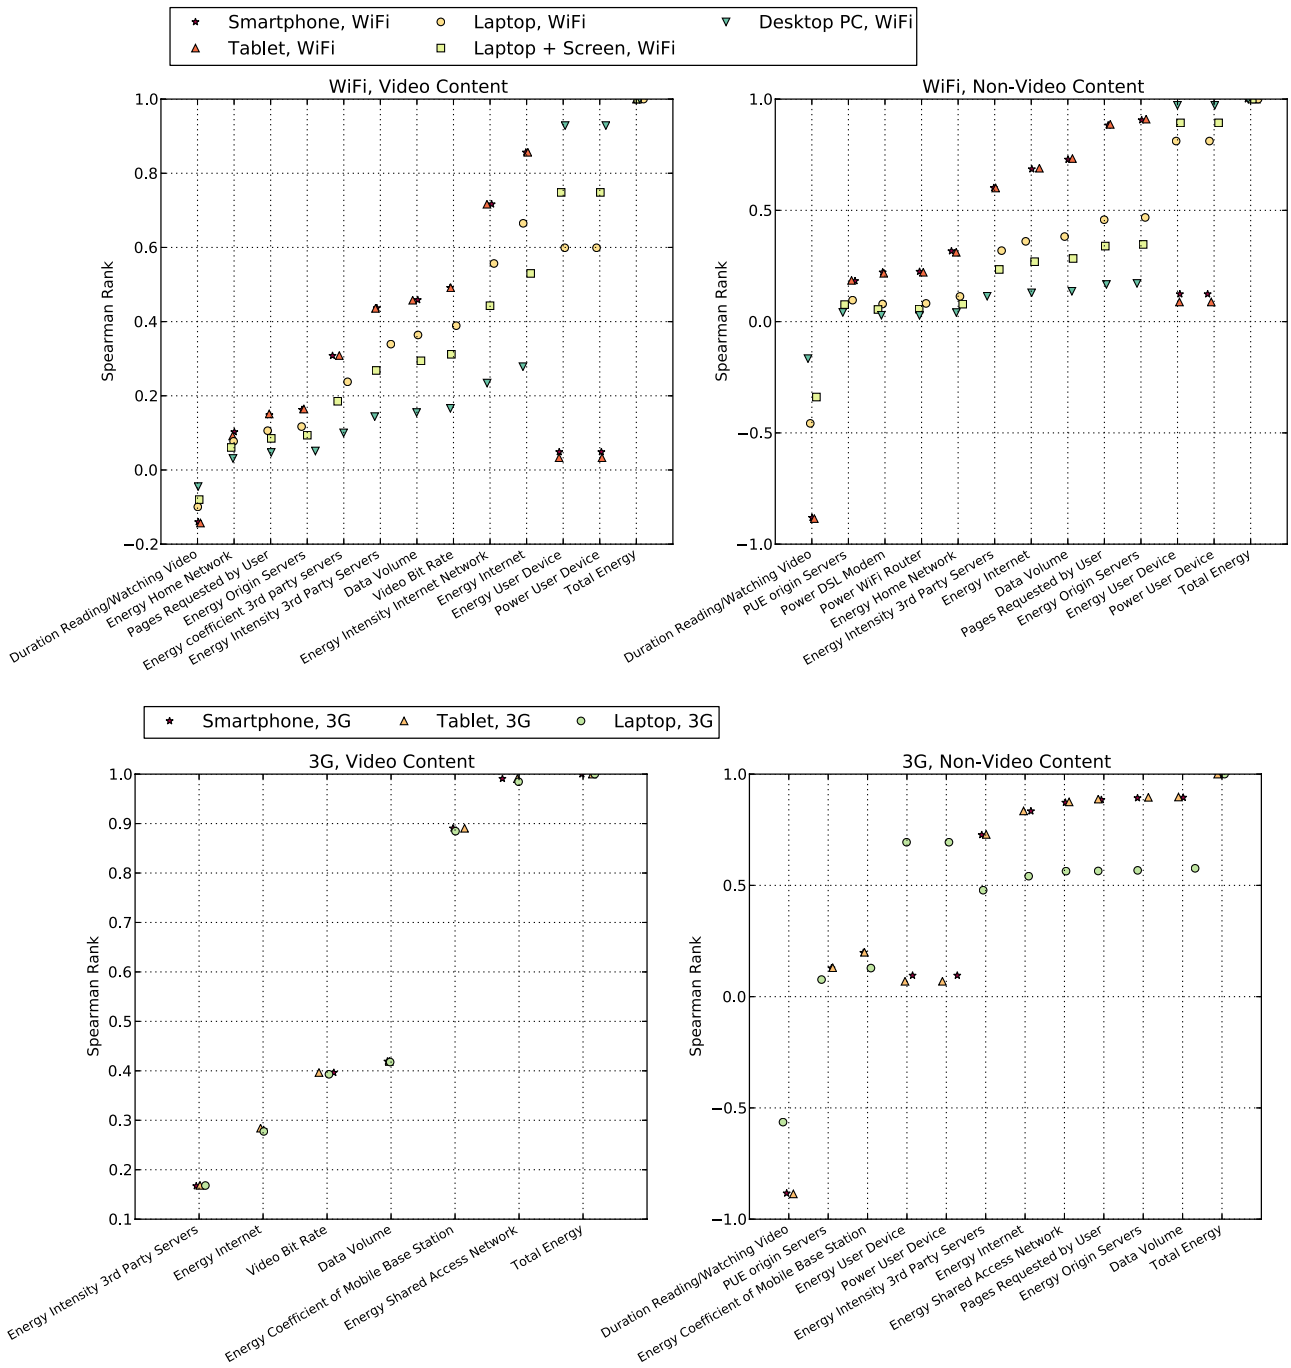


Figure 4 Spearman rank correlation values between most impactful model variables and total energy consumption for (a) scenarios with smartphones, tablets, and laptops for connections by digital subscriber line (DSL) and on-premise wireless networks (WiFi) and (b) scenarios with smartphones, tablets, and laptops connecting by third-generation mobile networks (3G). The Internet consists of transmission of content through routers, switches, and cable. Home network refers to DSL plus WiFi.

per bit for wired connections of approximately 0.009 kWh/GB (4 J/Mb), which is approximately one fourth of our results of 0.038 kWh/GB (17 J/Mb). For the servers, on the other hand, Williams and Tang (2011) arrive at a much higher energy footprint per user for browsing Web pages by assuming that a server is occupied during 50% of the duration the user spends reading a page, whereas we use primary data from GNM showing

that Web servers complete page requests in subsecond time intervals. Figure 5 compares results for energy per bit on the Internet and access network in two of our scenarios with the results from the earlier works. In this figure, we include only the energy consumption per bit by the Internet and the shared access networks, which we assume to be independent from the type of data transferred (text or video). Baliga and colleagues




Figure 5 Comparison of average energy consumption for data transfer by Internet and access network for bottom-up studies. WiFi = on-premise wireless networks; 3G = third-generation mobile networks; J = Joule; Mb = megabits.

(2009) estimate a value of 0.0074 kWh/GB (3.32 J/Mb), which is slightly lower than our minimum assumption of 0.01 kWh/GB (4.5 J/Mb) for the sum of edge and core routers and fiber-optic equipment. Whereas they assume 100% utilization of Internet routers and we assume between 12% and 25% (TeleGeography 2005), our measurements of hop count per route (average of 12) is lower than their assumed value of 14. The difference regarding the shared-access network power consumption is the result of a different allocation model. They allocate by throughput capacity; we allocate by time. We argue, in the “Models” section above, why we believe time to be a more appropriate metric. Moberg and colleagues (2010) do not take into account energy consumption by servers. Idle energy consumption is then apportioned relative to duration of service use.

Chandaria and colleagues (2011) do not take into account energy consumption of the Internet in their calculations. Their result for the wired-access network is 0.024 kWh/GB (11 J/Mb), and ours is 0.019 kWh/GB (8.49 J/Mb). This similarity is accidental. They take into account the idle power consumption of the DSL modem based on the assumption that it is used for 10.75 hours per day and idle for the remaining time and allocate it to the active use time similarly to Moberg and colleagues (2010). We, on the other hand, include a wireless network router besides a modem (both 5 W), but do not account for idle power consumption. The reason why we do not include idle power consumption of user equipment in this assessment is the current lack of systematic studies of this important factor to energy consumption at user premises. This problem is further compounded by allocation questions of idle power consumption.

Even though every new generation of mobile networks brought a decrease of the energy consumption per bit of data, the total power consumption of base stations increases with their total throughput capacity (Manner et al. 2010). This, together with higher bandwidth usage by mobile services (Cisco Newsroom 2012), means their relevance will grow. Our assumptions regarding the energy efficiency of mobile data transfer

overlap with those by Toffel and Horvath (2004). They relate the total energy consumption of a 2G mobile network to the total number of subscribers in the network and determine a power draw of 840 W per subscriber. This average power draw is then applied to the transmission of data, which is assumed to endure 60 seconds over a 56-kbps modem. The resulting energy consumption of 33.33 kWh/GB (15,000 J/Mb) is two orders of magnitude higher than our average values of 0.293 kWh/GB (132 J/Mb). This discrepancy mainly results from outdated values for the utilization of mobile networks and from using the energy footprint of voice service to calculate the footprint of a data service.

In their top-down study, Taylor and Koomey (2008) find the energy footprint per data volume to range between 9 and 16 kWh/GB (4,050 to 7,200 J/Mb). This figure has been referenced and updated by several other studies extrapolating using a trend identified by Taylor and Koomey (2008). Weber and colleagues (2010a) use this in a comparison of the environmental impact of different methods for delivering music and assume a value of 5 to 7 kWh/GB (2,250 to 3,150 J/Mb). Preist and Shabajee (2010) estimate an upper bound on future global energy use for the provision of media services and extrapolate to a value of 4 kWh/GB (1,800 J/Mb) from Weber and colleague’s value. In order to compare this value to the results derived from our model, it is useful to consider the results separately for servers and the network in the way that Moberg and colleagues (2010) performed their calculation. For the data transport, they also apply Taylor and Koomey’s (2008) values, excluding the contribution of servers, to give a value of 3 kWh/GB (1,350 J/Mb). Taylor and Koomey (2008) take the energy consumption values from a study by Roth and colleagues, which accounts for network components used in a commercial context (Roth et al. 2002). Roth and colleague’s inventory is now severely outdated, but in order to compare Taylor’s values with ours, it is necessary to analyze this data in more detail. They distinguish between several device types, among which only the wide area network (WAN) switches and routers map to our model of the public Internet. They calculate the energy consumption on total shipments of network devices, which include Internet service providers (ISPs), commercial intranet, and household deployments, and, accordingly, their results are likely to overestimate the energy consumption when applied to calculate the power consumption of the public Internet. Assuming that three device categories (hubs, routers, and WAN switches) contain the devices that we consider the public Internet, then the energy consumption of the Internet would only account for 14% of Taylor and Koomey’s (2008) values. Applied to the values extrapolated by Preist and Shabajee (2010), this would result in an energy footprint for the Internet of 0.52 kWh/GB (230 J/Mb), which is roughly 14 times higher than our value of 0.038 kWh/GB (17 J/Mb). Chiefly among the reasons for this discrepancy are a potential overestimation on the side of Roth and colleagues (2002), an underestimation of the network traffic in the Internet, or a severe underestimation of the number of devices and their energy consumption in our bottom-up model.

Data Quality

Data on GNM server energy consumption and duration of service use was provided as primary data by GNM, so it is of high quality. For energy consumption by third-party servers, we use a figure estimated from publicly available emissions data from Akamai, one of the largest content delivery networks for the media industry, and use this for all third parties. Other CDNs that mainly provide static files are likely to operate their servers at similar energy efficiency. For servers that provide custom content, such as advertising content providers and data analytics servers, this energy-efficiency value is likely to be a significant underestimate. Our model of the Internet distinguishes between edge and core routers. For each class of router, we have a number of data points from manufacturers' specifications and peer-reviewed literature, which we use to generate a mean value and statistical distribution. Our model of the shared wired access network uses a similar approach. It omits certain equipment, which is operated by some, but not, all ISPs—for example virtual private network (VPN) connection servers between the access network provider and the Internet ISP or Remote Authentication Dial-In User Service servers—because of a lack of publically available data. Although we believe that this is acceptable as a lower bound for the shared access network power consumption and that inclusion would increase the portion of the shared access network without significantly altering the overall results of our assessment, further research would benefit from transparency of ISPs in this area. Data on energy use of end-user devices comes from Energy Star and so can be considered primary data. The relative quality of the different data points was used in determining the range of variance of parameter distributions used within the Monte Carlo simulation described previously.

Implications and Applications

Our work highlights the importance of allocation techniques that are in accord with the technical functionality and usage of the system under study, and this is particularly challenging in the area of IT systems. The choice of an allocation technique can have a significant impact on the results of the LCA. Our work makes a contribution to the debate of how best to do this, although we do not claim that we have provided the definitive answer. In particular, we allocate all energy of a user device to one function—namely, browsing a Web site—while the user is carrying this out, even though the system could be carrying out other functions simultaneously. For example, it may be playing music. Also, it is likely providing instantaneous availability of services, such as e-mail, Internet telephony, or instant messaging chat. The question of how best to allocate user-device energy between these requires further work. Further, a user device has periods when it is consuming energy on standby, or is on but not providing any active functionality. How best to allocate the energy used during these periods between the various functionalities it provides is also a question meriting further exploration.

Beyond the scope of this article, it is relatively straightforward to extend our analysis to cover GHG emissions in a specific location resulting from consumption of online content. The model identifies the different locations where electricity consumption takes place in the use phase of a service. This can be combined with national and regional GHG emission figures, where they exist, to give a more precise estimate than would be possible using a single global or national figure. Our work can also be extended to allocate energy and GHG emissions associated with manufacturing the equipment used to provide the digital services. This is obviously an important part of the overall footprint and should be accounted for when making comparisons with alternative delivery methods of news content, such as paper based.

The global IT system is responsible for the consumption of 3.9% of all electricity generated (Malmodin et al. 2010). A significant amount of effort has been put into reducing energy use of individual components—such as laptops and data centers—motivated by eco-efficiency and cost savings. Although this is valuable, it does not address the energy and environmental consequences of design decisions taken by the various parties involved in providing services across the Internet. The complexity of the business ecosystem involved in such services means that a design decision by one can have energy (and therefore environmental and cost) implications on many others. Similarly, choices by the end user have effects throughout the system, and those choices are influenced by the service provider. The energy model presented in this article is detailed enough to allow assessment of the implications of such decisions and choices. It allows the systemic approach that characterizes industrial ecology to be applied to the IT business ecosystem in a number of ways.

First, such a model can be used to support real-time feedback to a user about the energy and climate impacts of their online behavior, as proposed by Weber and colleagues (2010b). We discuss how distributed systems technology can be used to support this in Schien and colleagues (2011). Though this may be of interest to some users, we do not see this as likely to lead to significant energy reduction without action by the service providers. Service providers can use our model to assess the effect of possible user trends on energy use by their service and use this to consider which trends to encourage and which to discourage. For example, in the case of consuming news for 10 minutes analyzed here, the Spearman rank analysis shows that different variables dominate the overall energy footprint depending on which service type, user device, and access network are chosen. When accessing nonvideo services by WiFi, the user device dominates in the case of lap- and desktop devices, but when accessing nonvideo services by WiFi with smartphones or tablets, the origin servers and per-page duration of reading are most significant. In contrast, when accessing video content over cellular wireless networks, the most significant factor is the wireless access network. When reading text over a cellular wireless network, the user device is most important when using a laptop and the data volume and origin servers when using a smartphone or a tablet.

Figure 2 shows that except for viewing video over a 3G network, scenarios with low-power user devices, such as smartphones and tablets, result in lower overall energy consumption, compared to scenarios with lap- or desktops. This suggests that encouraging a move to smartphone and tablet access will, in many cases, decrease overall energy consumption. This can be done through the provision of applications that enable enhanced experiences on such devices, provided that such a move does not stimulate additional purchases or an increased upgrade rate of such devices.

Second, such a model can be used to assess the impact of decisions by designers of a digital service on the energy use of that service across the IT system and propose design modifications that result in reduced energy use. For example, our analysis shows that data transfer of video content has a significant energy use on the 3G mobile network, but is less when transferred over other networks. Hence, a possible design intervention would reduce the resolution of video automatically when the service provider detects that the service is being delivered over 3G, but leave high resolution at other times. Such an intervention, if widely adopted among video service providers, could significantly reduce load on the 3G mobile network and hence associated energy use, environmental impacts, and costs. Beyond the analysis for GNM's Web site that is presented in this article, our approach can be used to evaluate other IT design and architectural decisions from an energy perspective. For example, Apple's iCloud music-match service fingerprints songs of a user's music collection locally and adds the identified songs to the cloud library from the existing cloud repository and thus avoids redundantly uploading terabytes of music files (Schien 2012). Another intervention that can be evaluated with the model is increasing outsourcing of data from the servers of a host, such as GNM to the CDNs, who can serve content more efficiently and benefit from economies of scale at the same time as reducing bandwidth in the core network, realizing additional energy savings.

More broadly, such a detailed model can be applied to questions of "virtual industrial symbiosis." Certain Internet architectures used by service providers, such as the peer-to-peer architecture used by the Spotify music streaming service, use "waste" compute cycles on customer machines to deliver content on other machines. The prime motivation of such architectures is cost reduction (by avoiding energy and infrastructure) at the service provider. Our model could be extended to allow assessment of such architectures to determine whether they do reduce energy consumption across the system or simply move the energy burden away from the service provider.

With the increasing pervasiveness of digital technology, sophistication of online services, energy consumption by IT, and complexity of the business and technical systems that deliver them, it is necessary to go beyond local optimization of energy use and environmental impacts and adopt a systemic perspective to mitigation, as advocated by industrial ecology. By providing a model of digital services detailed enough to explore the impact of design interventions on energy use across the system, we facilitate the adoption of such a perspective.

Conclusion

The use of aggregate figures and assumptions about typical user behavior may be adequate for environmental accounting and reporting purposes, yet it can conceal insights into the impact of variability on an energy footprint that can be used for a number of other purposes. As proposed by Weber and colleagues (2010b), a more detailed model can be used to support real-time feedback to a user about the energy and climate impact of their behavior. We discuss how distributed systems technology can be used to support this in Schien and colleagues (2011). Such a model can also be used to support the environmental strategy of an organization wishing to reduce the footprint of its digital services. It can be used to assess different interventions for their potential impact and support "design for environment" of digital products.

For example, our results show that data transfer of video content has a significant energy use on the 3G mobile network, but less so elsewhere. Hence, a strategy of reducing the resolution of video would be appropriate for mobile devices, but unnecessary for other devices. If the browsing time of users is assumed constant, the model also shows that the duration of time spent on a page is inversely correlated with energy consumption, particularly if that page is of text or images, rather than video, because when the user is looking at multiple pages, the delivery of each adds to energy consumption. This suggests that focusing on the design of Web service and content to enable users to easily get to content that is most of interest to them, and ensuring it is of sufficiently high quality that they want to stay with it, is beneficial in terms of both the energy footprint and as a business strategy.

Recently, data-center energy consumption has received heightened public attention, for example, by Cook and Van Horn (2011). Though increasing awareness of this issue is justified, our analysis, together with that of others, shows that in most cases, energy use by user devices and the 3G mobile network are bigger contributors to the service footprint. In the case of online news, it is only when accessing text by a mobile device on the 3G network that energy consumed by data centers is dominant (this might differ for other more computationally intensive services). Data centers are assuming the role that plastic bags have for supermarkets, receiving attention disproportionate to their relative contribution of environmental burden, compared to other parts of the retail business. It is important that the analysis of the impacts of ICT, and the means to mitigate these, takes a view of the entire system.

Acknowledgments

This work was conducted as part of the SYMPACT project, funded by the RCUK Digital Economy and Energy programs (EPSRC EP/I000151/1). The authors acknowledge the contribution of time and data made by Guardian News & Media, particularly Matthew Malthouse, Christopher Hodgson, Jo Confino, and Stephen Wood. For her suggestions

and for proofreading of the article, the authors also thank Elaine Massung.

References

- Akamai. 2010. CDP 2010 investor information request: Akamai Technologies Inc. www.akamai.com/dl/akamai/cdp_akamaitechnologies_2010.pdf. Accessed October 2012.
- Aleksić, S. and A. Lovrić. 2010. Power efficiency in wired access networks. *e & i Elektrotechnik und Informationstechnik* 127(11): 321–326.
- Baliga, J., R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker. 2009. Energy Consumption in Optical IP Networks. *Journal of Lightwave Technology* 27(13): 2391–2403.
- Barroso, L. A. and U. Hözlle. 2007. The case for energy-proportional computing. *IEEE Computer* 40(12): 33–37.
- Beauvisage, T. 2009. Computer usage in daily life. In *Proceedings of the 27th International Conference on Human Factors in Computing Systems*, pp. 575–584. Boston, MA, USA: Association for Computing Machinery.
- Beckett, J. and R. Bradfield. 2011. Power efficiency comparison of enterprise-class blade servers and enclosures. www.dell.com/downloads/global/products/pedge/en/BladePowerStudyWhitePaper_08112010_final.pdf. Accessed October 2012.
- Bertoldi, P. 2010. The European Programme for Energy Efficiency in Data Centres: The Code of Conduct—Presentation Slides to the Joint Research Centre—JRC—European Commission. www.itu.int/dms_pub/itu-t/oth/3C/02/T3C020000390001PDFE.pdf. Accessed October 2012.
- Chandaria, J., J. Hunter, and A. Williams. 2011. The carbon footprint of watching television, comparing digital terrestrial television with video-on-demand. In *Sustainable Systems and Technology (ISSST), 2011 IEEE International Symposium On*, pp. 1–6. Chicago, IL, USA: Institute of Electrical and Electronics Engineers.
- Cisco Newsroom. 2012. Cisco visual networking index. <http://newsroom.cisco.com/press-release-content?type=webcontent&articleId=668380>. Accessed October 2012.
- Cook, G. and J. Van Horn. 2011. How dirty is your data? Greenpeace. www.greenpeace.org/international/en/publications/reports/How-dirty-is-your-data/. Accessed October 2012.
- Deruyck, M., W. Vereecken, E. Tanghe, W. Joseph, M. Pickavet, and L. Martens. 2010. Comparison of power consumption of mobile WiMAX, HSPA and LTE access networks. In *6th Conference on Telecommunication Techno-Economics, 2010. CTTE 2010*, pp. 1–7. Ghent: IEEE.
- EU JRC (European Union Joint Research Centre)—Environment and Sustainability. 2011. International Reference Life Cycle Data System (ILCD) Handbook. <http://lct.jrc.ec.europa.eu/assessment/publications>. Accessed October 2012.
- Energy Star. 2011. Computers. www.energystar.gov/index.cfm?fuseaction=find_a_product.showProductGroup&pgw_code=CO. Accessed October 2012.
- Energy Star. 2012a. Small network equipment. www.energystar.gov/index.cfm?c=new_specs.small_network equip. Accessed October 2012.
- Energy Star. 2012b. Displays. www.energystar.gov/index.cfm?fuseaction=find_a_product.showProductGroup&pgw_code=MO. Accessed October 2012.
- Google. 2011. Google data center efficiency measurements. www.google.com/about/datacenters/inside/efficiency/index.html. Accessed October 2012.
- Greenpeace. 2012. How clean is your cloud? www.greenpeace.org/international/en/publications/Campaign-reports/Climate-Reports/How-Clean-is-Your-Cloud/. Accessed October 2012.
- ISO (International Organization for Standardization). 2006. ISO 14044: Environmental management—Life cycle assessment—Requirements and guidelines. Geneva: International Organization for Standardization.
- Lee, U., I. Rimac, D. Kilper, and V. Hilt. 2011. Toward energy-efficient content dissemination. *Network, IEEE* 25(2): 14–19.
- Malmodin, J., D. Lundén, M. Nilsson, and G. Andersson. 2012. LCA of data transmission and IP core networks. In *Electronics Goes Green 2012+*, pp. 1–6, Berlin: Fraunhofer Verlag.
- Malmodin, J., Å Moberg, D. Lundén, G. Finnveden, and N. Lövehagen. 2010. Greenhouse gas emissions and operational electricity use in the ICT and entertainment & media sectors. *Journal of Industrial Ecology* 14(5): 770–790.
- Manner, J., M. Luoma, J. Ott, and J. Hämäläinen. 2010. Mobile networks unplugged. In *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking—e-Energy '10*, p. 71. New York, NY, USA: ACM Press.
- Moberg, Å., M. Johansson, G. Finnveden, and A. Jonsson. 2010. Printed and Tablet E-paper newspaper from an environmental perspective—A screening life cycle assessment. *Environmental Impact Assessment Review* 30(3): 177–191.
- Preist, C. and P. Shabajee. 2010. Energy use in the media cloud. In *2nd IEEE International Conference on Cloud Computing Technology and Science*. Indianapolis, IN, USA: Institute of Electrical and Electronics Engineers.
- Roth, K. W., F. Goldstein, and J. Kleinman. 2002. Energy consumption by office and telecommunications equipment in commercial buildings. Vol. I: Energy consumption baseline. Reference no. 72895-00. Prepared by Arthur D. Little for U.S. Department of Energy. Washington, DC: A.D. Little.
- Scharnhorst, W., L. M. Hilty, and O. Jolliet. 2006. Life cycle assessment of second generation (2G) and third generation (3G) mobile phone networks. *Environment International* 32(5): 656–675.
- Schien, D. 2012. How big are the power savings through iTunes match? *The Guardian*. www.guardian.co.uk/sustainable-business/itunes-icloud-energy-consumption. Accessed October 2012.
- Schien, D., C. Preist, P. Shabajee, and M. Yearworth. 2011. Approaches to the dynamic energy footprinting of online media. In *EnviroInfo 2011 Innovations in Sharing Environmental Observation and Information Proceedings of the 25th International Conference on Informatics for Environmental Protection*. Aachen, Germany: Shaker Verlag.
- Schien, D., C. Preist, M. Yearworth, and P. Shabajee. 2012. Impact of location on the energy footprint of digital media. In *IEEE International Symposium on Sustainable Systems and Technology (IEEE ISSST 2012)*. Boston, MA, USA: Institute of Electrical and Electronics Engineers.
- Somavat, P., S. Jadhav, and V. Namboodiri. 2010. Accounting for the energy consumption of personal computing including portable devices. *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking—e-Energy '10*, p. 141. Passau, Germany: Association for Computing Machinery.
- Stutz, M., M. F. Emmenegger, R. Frischknecht, M. Guggisberg, R. Witschi, and T. Otto. 2006. Life cycle assessment of the mobile

- communication system UMTS: Towards eco-efficient systems. *The International Journal of Life Cycle Assessment* 11(4): 141–146.
- Taylor, C. and J. Koomey. 2008. *Estimating energy use and greenhouse gas emissions of Internet advertising*. Working paper for IMC2. <http://evanmills.lbl.gov/commentary/docs/carbonemissions.pdf>. Accessed October 2012.
- Teehan, P., M. Kandlikar, and H. Dowlatabadi. 2010. Estimating the changing environmental impacts of ICT-based tasks: A top-down approach. In *2010 IEEE International Symposium on Sustainable Systems and Technology (ISSST)*. Arlington, VA, USA: Institute of Electrical and Electronics Engineers.
- TeleGeography. 2005. International Internet statistics. ITU ICT Indicators Meeting. www.itu.int/dms_pub/itu-d/md/02/isap2b.1.1/c/D02-ISAP2B.1.1-C-0025!!PDF-E.pdf. Accessed October 2012.
- Toffel, M. W. and A. Horvath. 2004. Environmental implications of wireless technologies: News delivery and business meetings. *Environmental Science Technology* 38(11): 2961–2970.
- heise Verlag. 2011. iX-12-2011. www.heise-shop.de/heise-zeitschriften-verlag/ix-12-2011_pid_16800353.html?hsc=IK_Zeitschriften_Modul_2011. Accessed October 2012.
- Vodafone. 2011. Sustainability report. www.vodafone.com/content/dam/vodafone/about/sustainability/reports/2010-11_vodafone_sustainability_report.pdf. Accessed October 2012.
- Weber, C. L., J. G. Koomey, and H. S. Matthews. 2010a. The energy and climate change implications of different music delivery methods. *Journal of Industrial Ecology* 14(5): 754–769.
- Weber, C. L., E. A. Olivetti, and E. D. Williams. 2010b. Data and methodological needs to assess uncertainty in the carbon footprint of ICT products. In *2010 IEEE International Symposium On Sustainable Systems and Technology (ISSST)*, p. 2009. Arlington, VA, USA: Institute of Electrical and Electronics Engineers.
- Weber, C. L. 2012. Uncertainty and variability in product carbon footprinting. *Journal of Industrial Ecology* 16(2): 203–211.
- Williams, D. R. and Y. Tang. 2011. A methodology to model the energy and greenhouse gas emissions of electronic software distributions. *Environmental Science and Technology* 46(2): 1087–1095.
- Wood, S. G. 2012. Personal communication with Stephen G. Wood, Research Engineer, Sustainability. Guardian News & Media, London, 2012.
- Zwemke, T. 2012. Personal communication with Thomas Zwemke, Mobilfunk und Umwelt/TRPE-N. Vodafone D2 GmbH, Langenhagen, Germany, 2012.

About the Authors

Daniel Schien is a research associate in the Faculty of Engineering at the School of Computer Science, **Paul Shabajee** is a research fellow in the Faculty of Engineering at the School of Computer Science, **Mike Yearworth** is a reader in engineering systems in the Department of Civil Engineering, and **Chris Priest** is a reader in sustainability and computer Systems, all at the University of Bristol, Bristol, UK.

Supporting Information

Additional supporting information may be found in the online version of this article at the publisher's web site:

Supporting Information S1: This supporting information contains detailed information on methodological assumptions and input data sets used in the article, as well as additional findings.