

A Model for Green Design of Online News Media Services

Daniel Schien*
schien@cs.bris.ac.uk

Paul Shabajee*
paul.shabajee@bristol.ac.uk

Stephen G. Wood[‡]
stephen.g.wood@guardian.co.uk

Chris Preist*
cpreist@cs.bris.ac.uk

*Department of Computer Science, University of Bristol, UK

[‡]Center for Environmental Strategy, University of Surrey, UK and Guardian News Media, UK

ABSTRACT

The use of information and communication technology and the web-based products it provides is responsible for significant emissions of greenhouse gases. In order to enable the reduction of emissions during the design of such products, it is necessary to estimate as accurately as possible their carbon impact over the entire product system. In this work we describe a new method which combines models of energy consumption during the use of digital media with models of the behavior of the audience. We apply this method to conduct an assessment of the annual carbon emissions for the product suite of a major international news organization. We then demonstrate its use for green design by evaluating the impacts of five different interventions on the product suite. We find that carbon footprint of the online newspaper amounts to approximately 7700 tCO₂e per year, of which 75% are caused by the user devices. Among the evaluated scenarios a significant uptake of eReaders in favor of PCs has the greatest reduction potential. Our results also show that even a significant reduction of data volume on a web page would only result in small overall energy savings.

Categories and Subject Descriptors

H.3.4 [Systems and Software]: Performance evaluation (efficiency and effectiveness)

C.4 [Performance of Systems]: Modeling techniques, Performance attributes

General Terms

Measurement, Design, Economics, Human Factors.

Keywords

Carbon footprinting; digital media; sustainability; green software engineering

1. INTRODUCTION

The use of information and communication technology (ICT) and the software services it provides via the Internet is responsible for significant global energy consumption and the resultant emissions of greenhouse gases often quantified as product energy and

Copyright is held by the International World Wide Web Conference Committee (IW3C2). IW3C2 reserves the right to provide a hyperlink to the author's site if the Material is used in electronic media.
WWW 2013, May 13–17, 2013, Rio de Janeiro, Brazil.
ACM 978-1-4503-2035-1 /13/05.

carbon footprints. It has been estimated by Malmodin et al. [32] that in the year 2007, such usage was responsible for 710 tera watt-hours (TWh) of electricity consumption – 3.9% of global production, resulting in emissions of 447 million tones carbon dioxide equivalent (tCO₂e).

While many of companies that provide these services consider and optimize the energy efficiency of their data centers, this is often a small part of the total energy consumed compared to other parts of the product system, in particular third-party data centers, networks and user devices such as PCs and tablets (Figure 1).

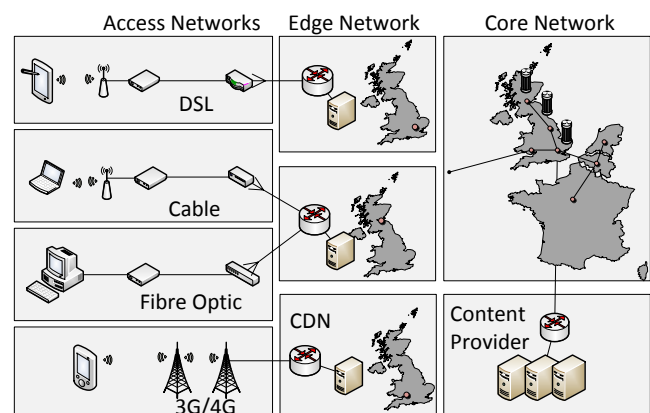


Figure 1 - Illustration of the service system. User devices connect with access network equipment via edge and core networks to the origin servers of the content provider and the CDN servers.

Design decisions by service providers regarding the architecture and use models of digital services can significantly influence their overall environmental impact. Such decisions may be taken explicitly with a view to reducing the overall energy and/or carbon impacts of their digital service, or may be decisions taken for other business reasons. In both cases, having an estimate of their likely resulting environmental impact will allow this factor to be considered alongside others when deciding whether to go ahead, or when considering which of several options to pursue.

The energy consumption by data centers and networks has recently received increasing attention by the engineering community. These efforts are directed to optimize the energy consumption in a single subsystem and can indirectly contribute to reductions of carbon footprints. However, in order to reduce the total energy consumption and avoid shifting burden between

subsystems, a model of the end-to-end energy consumption is needed. In particular, a quantification of potential savings in subsystems outside the operational control of the service provider is needed to effectively support efforts to reduce energy and carbon footprints.

1.1 Contributions

In this paper, we present an analysis of the suite of digital services offered by Guardian News and Media Ltd (GNM) including the guardian.co.uk website read on PCs, smartphones and tablets and determine the current operational energy and carbon footprint of the end-to-end delivery of these services – that is the energy use required to operate the services at the time of service use. Such delivery involves servers (both at GNM and third parties), network equipment and end-user access equipment. The GNM is an example of a complex media organization providing a mix of digital products and as such the results of this work can be generalized more broadly.

Existing approaches to assessing energy consumption and carbon emissions of digital services use a Life Cycle Assessment (LCA) methodology [20], and adopt a model of an average or prototypical service user. Such an approach provides an estimate of the overall footprint, but is relatively limited when exploring interventions on the service architectural design. We go beyond this state-of-the-art by combining LCA techniques based on a detailed product model with detailed parameters of the behavior of users of the services, and in so doing produce the most accurately modeled assessment of energy use of a digital service conducted to-date. We achieve this by synthesizing models from the engineering disciplines of networks and user devices into an end-to-end model of energy consumption.

We then consider potential design interventions in this system, and quantitatively estimate the change in emissions each enables. We consider six interventions and assess their relative impact on energy use. In doing so, we present the first quantification of carbon reductions of such interventions on a digital service. These interventions are illustrative, and the model can be used to assess other such interventions on a service.

2. RELATED WORK

Our research draws on work in the industrial ecology and computer systems engineering communities. From industrial ecology, we adopt and adapt life cycle modeling techniques used to identify energy use and carbon emissions from the creation and delivery of a given product. Using techniques from computer systems engineering, we integrate these LCAs with models of data flow and energy use across the Internet.

2.1 Carbon and Energy Footprinting of Digital Products

Research in industrial ecology has developed techniques for the environmental assessment of physical products and services, in particular for product life cycle assessment [26]. Recently, work applying these techniques to the carbon and energy footprinting of digital products has been conducted.

These existing studies differ in the level of detail to which they model the digital product systems they consider. Two alternative modeling approaches exist: bottom-up or top-down. A top-down model measures or estimates the total energy use of an entire subsystem, for example 'all data centers' or 'the internet', measures or estimates the total quantity of a given service type provided, for

example data transmitted, and divides the former by the latter to give the energy consumption per unit of service. Hence it treats a given subsystem as a black box, and does not model the usage of components within that subsystem by the digital product being assessed.

A bottom-up model on the other hand, includes a model of the subsystem and calculates the energy consumption by measuring or estimating the energy used by each component in delivering the digital product, and combining these figures to give a total. A bottom-up model of the energy consumption for Internet delivery, for example, sums the proportional energy consumption by each network device that plays a role in a typical route between two end points.

Neither approach is intrinsically more accurate than the other. However, only a bottom-up model provides the level of detail needed to assess the impact of a particular design change on the energy footprint of the entire system.

The majority of existing studies use primarily a top-down approach. Taylor and Koomey in [46] develop the most widely referenced top-down model of the energy footprint of data transfer in the Internet, and quantify an estimate of efficiency improvements over time. Due to an overly wide choice of system boundaries (what is included within the 'product system'), they are likely to have overestimated the energy consumption for data transport in the Internet. Weber et al. [49] use an extrapolated value from Taylor and Koomey's model in a comparison of the environmental impact of different methods for delivering music. Moberg et al. [35] also use this, combined with a bottom-up model of the local delivery system, to compare the impact of a printed newspaper and reading news with an e-reader. Teehan et al. [47] also apply Taylor and Koomey's model in conjunction with user behavioral data a study of user behavior by Beauvisage's [6] to estimate the total energy consumption in the US for a variety of digital activities.

A bottom-up model of energy consumption by servers was used by Chandaria et al. to analyze the carbon footprint of digital services at the BBC [11]. Some modeling simplifications, for example, assuming full utilization at nominal throughput rates, mean that they significantly underestimate server energy use. They also assume the energy consumption by networks to be negligible. Williams and Tang use a more realistic bottom-up model for the network although they underestimate the utilization of servers in their assessment of browsing a web shop and downloading a large file [50].

Baliga et al. [5] provide a detailed bottom-up model to estimate energy use of transmitting data through the Internet but exclude servers and end devices, and do not consider carbon emissions.

Our modeling work described below primarily uses a bottom-up approach, and draws on these results, but goes beyond them in several ways. Most notably, all other studies use aggregate data and assumptions regarding an average or prototypical user. We instead use behavioral data derived from web analytics software to estimate far more accurately the spread of behavior and characteristics within the user population, and to calculate energy consumption across this population. This results in a more accurate estimate of the total energy use, and also gives detailed and flexible model which allows the testing of the impact of alternate business and design interventions.

Finally, assessment standards for reporting of emissions by ICT equipment are currently being devised by several international

organizations including ITU (International Telecom Union)[27] and GHG (Greenhouse Gas) Protocol [23]. These standards are not directly applicable in support of green design with the goal of reducing total emissions as they firstly, do not mandate the inclusion of scope 3 emissions and secondly suggest top-down modeling approaches which hide the detail required to guide design decision making.

For a more detailed discussion of these approaches, including analysis and critique of the assumptions and methods they use in comparison with our approach, see [42].

2.2 Green Engineering

Our work also draws from recent work on green software and system design in the engineering disciplines extending it through integration into an end-to-end model. It is this work which provides the white box perspective in a bottom-up model of a system.

Significant work already exists on green design of software. According to Nauman et al. [36] green software is that “whose direct and indirect negative impacts on economy, society, human beings, and environment that result from development, deployment, and usage of the software are minimal.” They further define green software engineering as “developing software products in a way, so that [...] the negative and positive impacts continuously assessed, documented, and used for a further optimization of the software product”. They propose a green software process in which the energy consumption caused from the consumption of a service by a user, its *use phase*, is identified as an impact category, yet they do not propose a method for its assessment.

Some work on the evaluation of green software design exists. Dick et al. in [16] evaluate the energy savings on web servers through several interventions such as reduced image resolution, yet they do not quantify the energy savings over the entire system. Simons in [43] measure the additional power consumption on a PC induced from flash content and find an increase in power consumption by 3.4%. Thiagarajan [48] perform a similar but far more sophisticated experiment on mobile phones and analyze the additional energy consumption induced from rendering individual web elements. They find that they can reduce the energy consumption of mobile phones by 30% by changing the JavaScript contents without impacting the user experience. Other work that does take the energy consumption of software into account was carried out by the human computer interaction community focusing on user devices, for example [34]. They also do not assess the savings across the whole product system but only the user device.

Besides models, several *metrics*, or measurement procedures, to quantify the energy consumption of ICT device have been proposed by both private organizations and academic researchers. The SPEC_power [45] initiative, for instance, provides a standardized protocol to measure the power consumption of servers. EnergyStar develops metrics for many different types energy consuming devices, including servers [19], displays [17] and personal computers (PCs) [13]. While metrics serve to quantify the energy consumption of a particular instance of a system, they cannot provide a *generic estimate* of energy consumption – a role that is fulfilled by models. Additionally, there is no metric which spans all system parts end-to-end. Such a metric is likely to remain infeasible in the near future because

subsystems operated by third-party organizations are not open to instrumentation for measurements of energy consumption.

Multiple models have been developed to estimate the energy consumption of all parts of the service system. Examples include user devices [28][14], servers [8], data centers [2] and clouds [22]. Such analytic models and simulations can precisely estimate the device energy consumption from the composition of utilization factors of the device components such as CPU, disc and memory but at present these approaches cannot be used in the evaluation of design interventions directly because such models, firstly, require calibration for each specific device model and additionally - and more importantly - do not include a model of the component utilization that is induced by a particular service. Additionally, these models are currently too complex to be of practical use during software development. We adopt a simplified variant of this approach to estimate the energy consumption on user devices.

Besides these approaches to quantify energy consumption, the vast amount of work on more efficient ICT system design (including of servers [21], data centers [31] and networks [7]) often quantitatively illustrate the potential energy savings which an engineering intervention would provide. However, these savings, firstly, are agnostic of services and do not take an end-to-end perspective and, secondly, are often presented in relative proportions and thus do not allow transfer between subsystems. In some cases, the academic prototype web services used to illustrate potential savings of energy lack the complexity of industrial counterparts and thus further limit the transferability of the results between services. One of the strengths of this work is that the parameters are calculated from measurements of an actual, globally operating news service.

3. METHODOLOGY

In line with current practice of energy and carbon analyses, we adopt methodological principles from Life Cycle Assessment, but adapt them to our specific purpose. Any such assessment must have a clearly defined goal, which helps guide the choice of scope, in particular, the system boundaries which determine what is included or excluded from the analysis.

Our goal is to provide an analysis of the current energy usage and associated greenhouse gas emissions (which we both refer to subsequently as *impacts*) resulting from the use of GNM’s digital product suite to its end customers. We wish to do this with sufficient detail to allow what-if analysis of alternative design and business decisions, and to estimate changes in energy and emissions associated with these.

We are interested in the overall impacts associated with this activity, rather than the carbon or energy footprint of an average webpage or an average user. We include within the boundaries all activities within and outside GNM associated with the delivery of the service to the end customer. Specifically, this includes activity by servers within GNM responsible for the dynamic generation of webpage content in response to a user request, activity by third-party servers elsewhere responsible for providing parts of this content (such as images or advertisements), activity within the internet to transfer content between the datacenters and the end users, and activity on end user devices for requesting and downloading the individual resources that constitute the service, and rendering and displaying it during consumption.

As our aim is to provide feedback to software developers, system architects and product managers regarding energy use, we explicitly exclude from our system two impact areas which would be included in a full product carbon assessment for reporting or comparison with alternative products, for a printed newspaper. Firstly, we exclude impacts associated with creation of the product – both journalism for the content (which would be shared with the newspaper), and IT development of the products. Both are straightforward to calculate, and already included in the Guardian News and Media sustainability report [24]. Secondly, we exclude the impacts associated with some share of the manufacture of the IT equipment used. While it is conceptually relatively straightforward in estimation, data availability is poor. Furthermore, these two factors are not impacted by the design of the service delivery.

3.1 Model Structure

To ensure the level of detail necessary to assess design interventions, our model is made up of two parts: a system model and a set of parameter values. Firstly, we use a fully parameterizable model of the service system which can estimate the energy use for an individual user of a digital service based on that presented in [29], which we summarize here. The delivery of a digital service is divided into the following subsystems: user access device; customer premise network devices, access network, edge network, core network, servers/data centers, including both origin data center of the service provider and third-party, such as Content Delivery Networks, ad networks or analytics (see Figure 1). The parameters of the model include user access device (e.g. phone, laptop, tablet, desktop), access network technology (e.g. 3G network, Digital Subscriber Line, Cable Modem, local WiFi, corporate LAN), service choice (web page, video), geographical location and duration of using the service. The model consists of a set of equations and associated data which can be used to determine energy use across the system for a given choice of user access parameters. These equations are presented in detail in [29], and summarized in Figure 2. It illustrates in abstract the structure of the model. For a specific user group, the energy footprint is calculated as the sum of the energy consumption by servers, edge and core networks, access networks (only one variety) and user devices. For example, to calculate the energy consumption by wired access networks, we multiply the power consumption with the duration of the service consumption.

The equation parameters are divided into constants, subsystem variables and service variables. Constants are values which vary only slowly as the internet evolves, such as the energy efficiency of delivery of data across the 3G network. Subsystem variables are those which are also fixed, but are determined by user choice of subsystem, such as the power of the user device or home network setup. Finally, the service variables are those which are determined by the usage of the service itself, such as how much data is transferred and how long the user interacts with it.

Secondly, we segment the user base of the GNM into groups according to service and subsystem variables. This determines for a given (actual or hypothetical) scenario, the number of users from a specific country accessing the service with a specific combination of user and access network devices and includes the number of requests these user groups make to a certain content segment as well as the average duration they spend consuming this content. This population model is generated from detailed data on the behavior of users which is derived from GNM's web analytics tool (Omniure SiteCatalyst). One example of such a

group is the set of all users based in the UK who browse the Guardian web site on a specific tablet model via WiFi and DSL together with the average time they spend reading it.

To determine the spread of parameters in the population, we use access data for the month of March 2012, taken as a representative month. It was chosen because nothing unusual, such as the launch of a new digital product was occurring and it does not include significant holiday periods. In order to enable and support comparison with other studies and reporting we scale up all results to give annual equivalent values. Conceptually, we apply the energy model to the parameters of each individual user access, and sum these to determine the overall energy used by the population in accessing the service. In practice, for efficiency purposes, we create batches of similar users within each subsystem to calculate the overall energy used by that subsystem to provide service to the batch of similar users.

As with any such model, there is uncertainty about the data values used. We run a Monte Carlo analysis over 10000 iterations to handle this: we use a spread of values for each equation parameter and randomly sample from this spread repeatedly to obtain a distribution of outcomes.

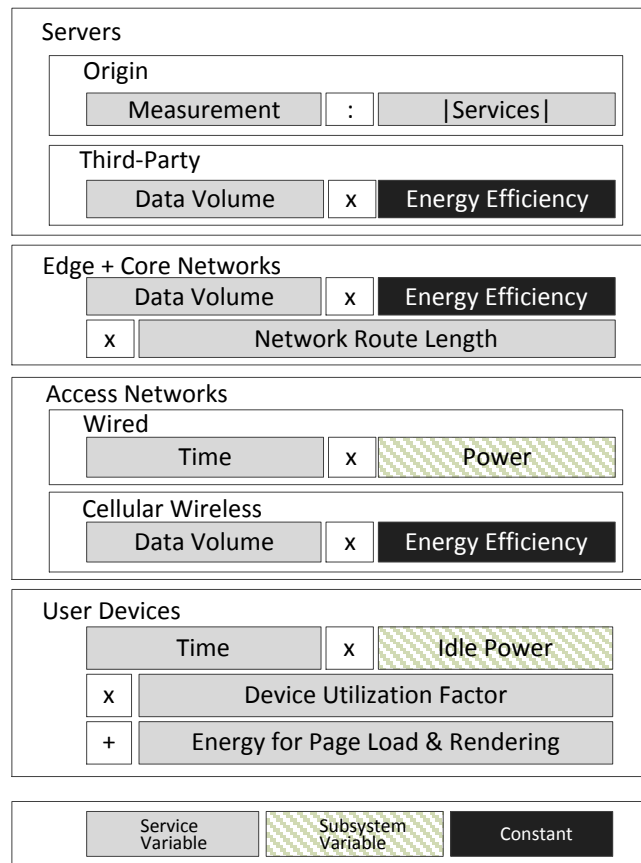


Figure 2- Abstract illustration of the model structure used to calculate the energy consumption per user group. The energy footprint for the service is calculated as the sum over the energy consumption of the relevant subsystems involved in the delivery and consumption.

Having used the data from the GNM's user population to estimate the current actual energy use and associated carbon emissions from the digital service, we then explore alternative design modifications. We do this by adjusting associated parameters in

the model and/or user population, followed by a re-running the simulation.

We now turn to each of the subsystems in the model, and discuss how we determine both the equation parameters, and the spread of user behaviors, for each of these. Note, though that the illustration in Figure 2 does not represent the full set of parameters for clarity. For additional detail on the system model, please refer to [42].

3.2 Servers

The servers at GNM are responsible for dynamically generating web page HTML skeletons and text, including comments etc. In the origin datacenter at the GNM the energy consumption of each server and the supporting networking infrastructure and storage devices can be measured directly. We allocate the total monthly energy consumption uniformly between all page requests served during the measurement interval. We include an overhead for cooling and power transformation infrastructure in form of the power utilization effectiveness (PUE) which is a measure of the portion of electrical energy used in computation compared to that used for cooling and power conversion.

Third-party servers are responsible for providing additional content to fill the HTML template. Content Delivery Networks (CDNs) hold image and video caches in servers around the world, allowing data to be provided more locally and therefore faster. Ad Servers provide advertising content. The energy consumption of these servers cannot be measured directly as they are operated outside of the control of the content provider. We estimate the energy consumption of data served to be 3.8 watt-hours per gigabyte (Wh/GB), based on public reports by one of GNM’s CDNs – Akamai [1]. We estimate the total energy consumption by multiplying the energy efficiency with the total data volume transferred in the connection to retrieve a specific resource from the CDN server. During the Monte Carlo simulation we use a triangular distribution with a lower and upper bound for this parameter of 0.89 Wh/GB and 29.56 Wh/GB based on estimates by Chandaria [11] and Google [25] respectively. We apply this distribution to all third-party providers.

3.3 Networks

The network used to transfer data between the various servers and the end user devices can be divided into core Internet and more local *edge* and *access* networks. The equipment involved is spread throughout the world among many parties, and so direct energy measurements are not available. Using a combination of industry and academic data we have built a model which, given a traceroute between two IP addresses, can estimate the energy required to transfer a given quantity of data through the core and edge networks [41]. It estimates the likely number of routers and repeaters of different kinds, and uses data regarding their power consumption to estimate this value. We found that the number of network hops in the route between two devices grows proportional with their geographical distance. In particular, this was evident in the connection to the origin servers. We also found that there was little variation in the route lengths to CDN servers and concluded the relative effectiveness in the CDNs in serving the data intensive sections (image, audio, video) from relatively close to the end user. Given that the service provider can very quickly change between CDNs or decide to serve all data locally, we decided to model the energy consumption by core and edge networks by the number of network hops in routes between geographical regions of user location and the average energy consumption per hop. Following the assumptions in above text

regarding utilization and energy efficiency of core and edge routers we estimate an average energy efficiency per hop of 1.42 Wh/GB, including a share of electricity for optical equipment. The energy consumption for data transport for a user group is calculated as the product of the proportional volume of each type of data, the average route length to each type of server and the energy efficiency per hop. The average route lengths between continents with the majority of the GNM audience and types of servers are listed in Table 1. These include core and edge network devices. The traceroute servers from which the measurements were made are located in a mix of academic networks, privately hosted servers and ISP’s looking glass servers. Thus, these routes include some access network equipment. The average route length is 13.

Table 1 - Average route lengths between regions of GNM readership and server types. 'Other' includes ad and analytics servers and mainly located in the US. Origin servers are located in the UK, CDN relatively local to customers.

	Origin	CDN	Other
Europe	14	9	12
North America	17	8	13
Oceania	21	8	14

During the Monte Carlo simulation, we vary the energy coefficient over a distribution. As an upper bound we assume a value from recent work by Malmodin et al. [33] who found the average energy efficiency of a major Swedish network to be 80 Wh/GB. In order to estimate a per-hop value on this basis, an assumption about the average route lengths in the Swedish network is required. Given that the average route length in our global measurements is 12, on average, routes in the Swedish network are likely to be shorter. We assume a value of 10 hops which results in an energy efficiency of 8 Wh/GB per hop. Given the fivefold difference between both values of per-hop energy efficiency we use those as minimum and maximum values in a triangular distribution during the Monte Carlo simulation and set the mode to the mean of the two at 4.71Wh/GB.

The energy consumption by the access networks is far more variable. It depends on networking equipment deployed locally to the user, which is more diverse. Broadly, we can divide it into home/small office, institutional/workplace/campus networks, and mobile (3G) access. We categorize users into one of these three categories using user domain data from the web analytics tools. Certain domains are known to be mobile, workplace or domestic/small to medium enterprise Internet Service Providers (ISPs) and can be straightforwardly categorized. To determine unknown domains, we look at the user access patterns from these and compare them with the access patterns of known domains.

Figure 3 gives a smoothed and normalized distribution of page views for a typical Sunday and Monday in March 2012 – this retains the broad qualitative shapes of the distributions but relative numbers of page views and fine detail have been removed for reasons of commercial confidentiality. Page views from three different domains are being displayed: a popular, largely domestic ISP, a typical workplace domain, and a typical academic domain from the UK (ac.uk). The workplace domains have significant daytime peaks around lunchtime, while the domestic domains peak in the evening.

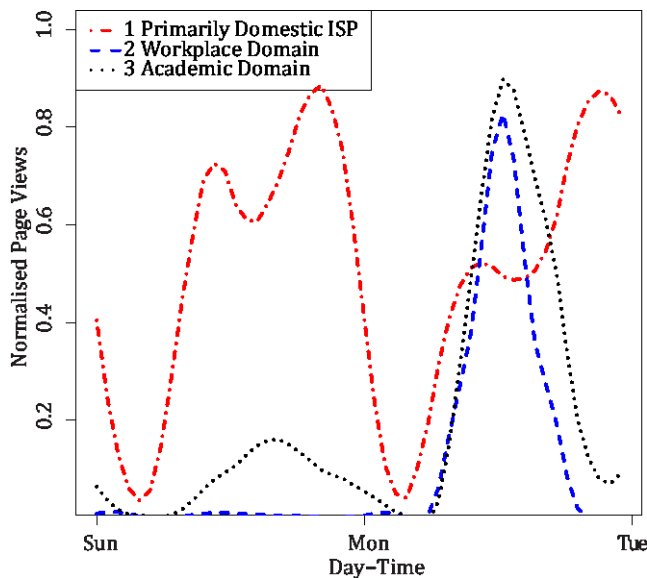


Figure 3 - Variation of number of page views during Sunday and Monday from domestic, commercial and academic ISPs.

Based on our analysis, we adopt a heuristic of classifying an unknown domain based on the ratio of peak access rate in evening to peak access rate in office hours. If the ratio is $<20\%$, we classify a domain as workplace, if it is $>200\%$ we classify it as domestic, else we classify it as mixed. Mixed domains are assumed to have a mix of workplace to domestic access based on the overall proportion of accesses in known domains. We explored other approaches, such as the ratio of weekend to weekday daytime traffic, but found this to be the most effective. Based on this distinction, we estimate the energy consumption by wired networks for each batch of users as the product of the service use time and the power consumption per connected user device of each type of access network.

The most common types of access technology for home and small office are ADSL, Cable and Fiber-optic LAN. It is not possible to work out which of these options is being used by a given user based on site analytics data. However, the analytics data does allow us to determine which country users are from, and apportion usage to each technology type based on the relative share according to data from the Organisation for Economic Co-operation and Development (OECD). For instance, the UK has 78% ADSL, 20% Cable and 2% Fiber/LAN while the USA has 36% ADSL, 56% Cable and 7% Fiber/LAN [37].

We assume that those components of Cable, DSL or fiber-optic access network equipment which are shared between multiple subscribers - most commonly those in a neighborhood - consume 19, 2 and 4W per subscriber respectively, derived from Aleksić & Lovrić [3]. Inside of homes, the power consumption varies depending on whether WiFi is deployed, and if so, whether as part of a modem. Based on measurements by Energy Star [44] the power consumption of cable and DSL modems is typically 7W, increasing to 11W if they include a WiFi router. We assume that about 85% of all households use WiFi based on statistics by Ofcom [38]. We follow Lanzisera [30] who assume WiFi routers are built into DSL modems in 80% of subscribers and cable modems in 20% of all subscribers. Remaining WiFi usage is assumed to be a separate device.

Having considered domestic network access, we now turn to workplace campus networks and 3G mobile access. There is relatively little data on energy use in offices. The most widely cited study in is by Roth [40] from 2002 and is now so dated that we believe the use of this data is not justified. In our model, we assume a power figure of 8 W per user, based on averaging results of studies of the LBNL campus [30] and the Stanford Computer Science Department [29].

Models of the power consumption per user of wireless cellular networks still vary widely. Our model is based on third generation networks. Based on our estimates in [42] we assume an average value of 293 Wh/GB. During the Monte Carlo simulation we apply a triangular distribution with a lower bound value of 63 Wh/GB based on a lean component-level model of the most efficient HSPA variant by Deruyck et al. [15]. As an upper bound we apply a value of 729 Wh/GB from the high-use scenario in the system-level model of LTE networks in [18].

3.4 User Devices

The final source of energy consumption is the device used by the user to access the service. The amount of energy consumed will depend both on what the device is, and also what service it is accessing, more specifically web browsing or video content. The service is known from the site visit data. To determine the user device, we use data collected by the web analytics system from the User-Agent string provided by the client and embedded JavaScript in the HTML pages. This provides the operating system of the device, and the screen resolution. This information is sufficient to identify tablets and smartphones devices and there is little variation in power consumption between models. In the case of laptops and desktops with monitors, however, further analysis is required. To do this, we constructed a database containing screen resolution data for laptops and monitors, using information from online stores and other sites. If a screen resolution is only used in the manufacturing of laptop displays, we assume the access is a laptop, and if it only appears as a monitor, we assume it is a desktop in combination with a monitor. If either is possible, we distribute the alternatives across the population of all users with similar parameters according to the relative proportion of laptops and monitors with this resolution in the database.

Our model estimates the energy consumption by user devices as the sum of the base power consumption at active idle and a dynamic portion induced from consuming a service. For web pages from the GNM the dynamic portion is approximated from the sum of the energy for loading and initial rendering. For smartphones, Thiagarajan [48] provide detailed measurements with several different pages and find that the average energy consumption for loading and rendering 20 joules (J). We assume that these findings are representative for tablets, as well. For laptops and desktops, we assume this is 50 J based on our own scoping experiments with a modern EnergyStar-rated laptop. The dynamic portion of power consumption for watching video is assumed to be 15 percent of base power consumption based on the results by [10] for smartphones and our own experiments with a laptop.

To estimate the power consumption of laptop, desktop and monitors, we take power consumption of models as measured by EnergyStar [13], for the one hundred most popular models on Amazon, and calculate the average. This yielded figures of 114W for desktops, 24W for monitors and 27W for laptops. We assume

the base power consumption in active idle of smartphones is 1W [10],[48] for smartphones, 3W for tablets [4].

3.5 From Energy to Carbon Footprint

Our model presented so far allows determining the energy footprint of a given user. In order to derive the corresponding carbon footprint we combine this with data about average carbon emissions per unit of energy (kgCO₂e/kWh), known as an emissions factor. The emissions factor can vary significantly between countries, as the mix of different energy generation technologies varies. For a given user, the site analytics software can infer their country of origin using their IP address and a geolocation database. We assume electricity use by the end device and access network takes place in this country. We cannot locate the edge and core network components involved with this degree of precision, or most of the CDN servers involved in a given interaction, and so we use the average global emissions factor for OECD countries. Finally, the GNM datacenter is known to be located in the UK, so uses the UK emissions factor. Together, these can be used to convert the various energy consumption figures in the model to provide an overall carbon footprint.

4. RESULTS

Based on our model we calculate the combined end-to-end energy consumption of the service system over on year for the audience of the GNM online news service. We present the annual equivalent values for the operational energy and carbon footprint (in metric tonnes) separately for each subsystem as shown in Table 2.

From our analysis by far the most impactful part of the media subsystem is the user devices which accounts for 74% of all carbon emissions which result from the generation of the consumed electricity. The second most impactful subsystem is the access network equipment which account for 22% of all carbon emissions. Compared to that, the impact of servers and the network is relatively small, however not insignificant, at 3.4%.

The 25th and 75th percentile of the resulting distribution of energy consumption by third-party servers and networks from the Monte Carlo simulation are 7.3% lower and 8.5% higher than the average.

Table 2. Electrical Energy and Carbon Footprint

	<i>Energy</i> [MWh]	<i>Carbon</i> [tCO ₂ e]	<i>% of</i> <i>CO₂e</i>
Origin Data Centre	369	199	3%
Shared Networks	111	60	1%
Third Party Servers	29	15	0.2%
Access Networks	3049	1681	22%
Users	10475	5736	75%
Sum	14033	7693	100%

5. ASSESSMENT OF DESIGN INTERVENTIONS

Having used our model to calculate current emissions from GNM digital products, we now turn to assessing six potential interventions on the product. The interventions presented below are a representative sample of interventions we have identified as being of interest to business strategists and sustainability

professionals based on interviews and at GNM. Interventions may take place for reasons of business strategy, product improvement, sustainability or a combination of those. Our model is detailed enough to assess the impact of such interventions on carbon emissions, allowing this factor to be included in decisions about whether to go ahead or not.

5.1 Reducing the Data Volume of Web Pages

Preist and Shabajee [39] have identified web pages as a potential source of “digital waste” – transportation of data which is of no value to the end user – suggesting that large web pages could be reduced in size to reduce emissions. We consider a reduction in data volume of 30% in the 1000 most popular pages on GNM. This could be achieved through a combination of reduction of JavaScript (responsible, according to our analysis, for 15-25% of data volume of a web page) through code optimization, and a reduction in the number, size and/or resolution of images. (Note that our original model already accounts for existing caching by browsers, as it uses actual data transferred by CDNs rather than original page size. Caching typically reduces data transfer by an average of 25% of the original page size.)

We assume that the structure of pages remains unchanged. Hence this intervention will not alter energy consumption by the origin servers, responsible for the generation of the HTML template and text content of the page. It will result in energy savings in the CDN servers and core/edge network, and mobile networks due to reduced data transmission.

According to our model, the intervention to reduce the data volume of the most popular web pages by 30% would result in a total reduction of 4,132 kgCO₂e, or 0.05% of the overall footprint. This comes from a reduction of 3,332 kgCO₂e (5%) of network emissions, 682 kgCO₂e (4.4%) of third-party server emissions and savings of 118 kgCO₂e in wireless access networks.

5.2 Simplifying Page Rendering

Separately from savings in data volume from optimized JavaScript contents, we also want to evaluate potential reductions on the user devices. In [48] Thiagarajan et al. measure the savings from simplifying the JavaScript contents in a Wikipedia page. They find that a single optimization that does not affect the user experience can realize 30% savings in the energy consumption of loading and rendering the page. We assume this change does not affect the base power consumption.

We evaluate the potential reductions of energy and carbon footprints under the assumption that an optimization could reduce the energy consumption for loading and initial rendering a page by 30% relative to the baseline assumptions between 20 and 50 joules as stated in section 3.4.

Assuming that in this scenario all other system parts are unaffected, we expect a reduction of the carbon footprint by the user devices of 9,344 kgCO₂e (0.16%) or 0.12% of the total carbon footprint.

5.3 Reducing Video Resolution

GNM offers significant amounts of video on its website, and one design option we consider is a reduction in the video quality. The typical bit rate is 1100 kbps and the resolution is 360p. A reduction of the bit rate by circa 50% can be achieved with a change of the resolution to 240p.

The reduced data volume will ease load on the data center of the video CDN and the core and access networks similar to the

reduction in page volume but to a greater degree as the data volume per page view is higher. Arguably, the energy consumption of user devices is likely to also be reduced, because less work may be done in rendering lower resolution video, yet a systematic analysis of this effect was out of scope of this analysis and we refrain from speculation and assume no change in their power consumption.

This resulted in a total reduction of 18,595 kgCO₂e, or 0.24% of the overall footprint. This comes from a reduction of 9,572 kgCO₂e (14.5%) of core network emissions and 6,341 kgCO₂e (0.37%) reductions in access networks, and 2,682 kgCO₂e (17.3%) in third-party data center emissions. The savings in the energy footprint of the mobile network are equivalent to 47% compared to the corresponding baseline value.

5.4 Disabling CDN Servers

In [41] we found that CDN servers are very effective in reducing the route lengths for data transport to GNM customers. In this scenario we evaluate the impact on energy consumption if delivery of all GNM content, including those parts that are currently served by CDNs, was performed centrally from the origin servers in London. To this effect we modify our model to assume the route length for connections to CDNs is identical to the route length to the origin servers: 14 hops on average from Europe, 17 hops from North America and 21 hops for connections from Australia. We assume that the energy efficiency of servers remains unchanged.

Such a change would result in an increase of the energy footprint of the core and edge networks by 38.04% or 25,130 kgCO₂e (0.3% reductions from total).

5.5 Increasing uptake of E-Readers

GNM provide a number of products targeted at specific devices. We now consider a scenario where a product targeted at passive-display eReaders is actively promoted through a combination of attractive pricing models, exclusive content and user-centric design. In this scenario, we assume this results in an increase in popularity of these devices to 10% of all page views and would displace consumption on PCs to the same degree and uniformly distributed between desktop and laptop. We assume that page volume remains constant. As a result, server and network activity is unchanged, but energy consumption by end-user devices is.

This resulted in a total reduction of 751,761 kgCO₂e (9.77% of the overall footprint).

5.6 Increasing Consumption of Video

We now consider a scenario where the product designers increase the quantity of video content on the site, and actively promote it through means such as high profile banners on the home page, aiming to increase the consumption of video by 10%. We assume that the video quality, and therefore the bit rate, is unchanged.

This will affect the energy consumption of all subsystems: the origin data center will serve additional page requests containing the video, the CDNs will serve additional content, the core and mobile networks will carry increased traffic, and the user devices will be active for the additional period while watching the video.

This resulted in a total increase of 13,304 kgCO₂e, or 0.17% of the overall footprint. This comes from an increase of 8,447 kgCO₂e (0.15%) in user device emissions, 1,984 kgCO₂e (3%) in network emissions, 2,216 kgCO₂e (0.13%) increase in emissions from access networks - including a small increase in emissions

from mobile networks of 38 kgCO₂e and 551 kgCO₂e (3.56%) in third-party data center emissions. The additional page views would contribute to additional emissions in the origin data center by 103 kgCO₂e (0.05%).

The results of all scenarios are shown graphically in Figure 4.

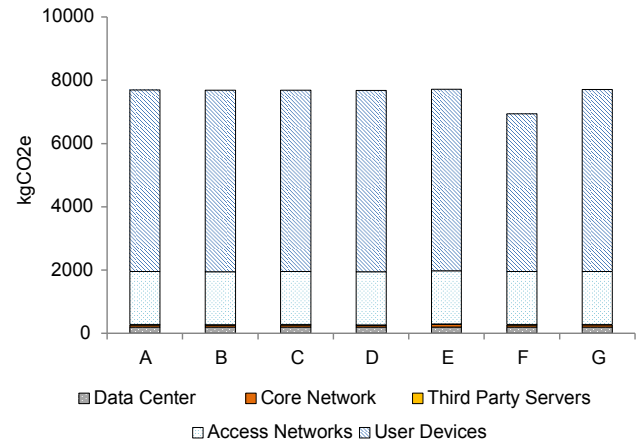


Figure 4 Total Annual Carbon Emissions by Subsystem for Scenarios: A – Baseline Results, B – Reducing the Data Volume of Web Pages, C – Simplifying Page Rendering, D – Reducing Video Resolution, E – Disabling CDN Servers, F – Increasing uptake of E-Readers, G - Increasing Consumption of Video

6. DISCUSSION AND FURTHER WORK

In the previous sections we have presented a methodology for the assessment of carbon emissions by digital services. It combines a detailed model of energy use by subsystems involved in the delivery of a service to a given user with a model of a diversity of behavior in the user population. In this way, it goes beyond the current state of the art which has focused on aggregate measurements and models of an average user. We also provide more detailed and accurate data on energy use by specific subcomponents in the system than has been used in prior studies. We have applied the model to give an accurate assessment of carbon emissions resulting from the delivery of GNM's digital product suite. Unlike prior coarser-grained models of carbon emissions from digital service use, it can be used to assess the impact of design changes on carbon emissions. We have demonstrated this by assessing the impact of six potential interventions on GNM services.

6.1 Intervention Scenarios

Based on the evaluations of interventions presented above we can draw three conclusions about the use of the model for the sustainable design of online news:

Firstly, at current levels of consumptions the displacement of PCs by lower power user devices has by far the highest potential to reduce the total carbon emissions - assuming that the PCs are decommissioned, in sleep mode or turned off instead. Secondly, interventions which significantly reduce the transferred data volume in networks can contribute to a substantial reduction of emission in networks and data centers, although to a much lower degree.

Thirdly, despite mobile networks having a much higher energy consumption per transferred data volume and thus are a risk for

the future sustainability of digital media consumption, at the present moment the emissions from use of GNM products over mobile networks are relatively small compared to the emissions of other subsystems, including wired access networks in particular. Thus, the rapid increase of access via mobile networks should not be the central concern of the GNMs sustainability agenda at this time.

6.2 Generality

More broadly the results presented are specific to the product suite and user community of GNM. However, the findings are likely to broadly apply to similar news and media sites and product suites, such as CNN or BBC. The methodology is more broadly applicable than this, and given the availability of appropriate data could be used to assess design interventions on a more video intensive digital product such as YouTube or a social networking service such as Facebook. We would expect the set of potential interventions, and their relative effectiveness, to change with the nature of the digital product.

6.3 Change over Time

One limitation of the work we have presented is that it assumes a steady state of the digital product, based on a snapshot at a given time, when assessing the impact of design interventions. However, it is clear that digital products are in a state of flux, and the web analytics data confirms that the products at GNM are no exception. Changes are resulting both from the uptake of specific new GNM products entering the market, and also because of broader trends such as the increased uptake of tablets. Through the use of web analytics systems our method increases accuracy above others based on annual data. Additionally, close integration of analytic suites significantly reduces the effort to update assessment results in accord to the evolution of the system. Thus, our methodology can be used to closely follow and project the impact of such trends on emissions over an extended period of time. Furthermore, rather than using a snapshot as baseline for assessing design interventions, such a projection can be used as a 'business as usual' baseline to give a longer term assessment of the impact of design interventions against such trends. Additionally, the choice of product design by GNM is co-evolving with such trends: a product may arise in response to a trend, but also will influence the uptake or otherwise of such a trend. Our methodology can support including environmental impacts in the choice of which trends to encourage and which to discourage.

6.4 Limitations

When using an LCA approach, it is important to address issues of data quality and data uncertainty. Comprehensive discussion of this with regard to the underlying energy model is beyond the scope of this paper, but addressed in [42]. Here, we consider these issues with regard to the user behavioral data, which is used to augment the energy model in the methodology presented in this paper. We extract the user data from site user analytics software. There are a range of well-known issues with reliability of such systems, see for example [12]. In our context one specific potential issue with this is the use of cookies and JavaScript to conduct analytics of specific users, meaning that user-based analysis is not reliable if cookies are rejected. We mostly overcome this as our analysis is 'per access', and so a user rejecting cookies will simply appear as multiple accesses. The one exception is in determining time on page. We also place an upper bound of 30mins for time-on-page, to cap cases where a browser

is left open on a page, but the user has finished using it and is doing something else.

6.5 Further improving the model

While our model is finer grained than those that have gone before, further improvements would increase its accuracy and allow finer-grained design decisions to be assessed. Firstly, the modeling of the relationship between service type and power use on the end device is relatively simplistic. More sophisticated models of power use by end user devices, based on utilization of CPU, memory, disk and IO could be used to refine this. Secondly, the granularity of data collection by the user analytic software also places limitations on what can be assessed. For example, one design intervention we identified as being of interest but were unable to assess is the effect of 'bounces' - rapid visits to a web page because on arrival the user discovers it isn't really what they want, and is therefore another example of digital waste. This was because the user data did not distinguish rapid (2-3 sec) visits from 15 second visits, where a person may have received some value from the page.

6.6 Extending the model

The model we have presented in the paper is focused on the energy used to deliver a digital service. We have focused on this because it is what is most directly influenced by service design decisions. The scope of the model could be extended to include other carbon emissions that are indirectly associated with the delivery of the service. These would include (a) some share of the emissions associated with the manufacture of IT kit used to deliver and use the service. (b) some share of the energy used by user equipment at times when it is idle or on standby (i.e. not providing any service); (c) some share of energy used (IT equipment, office heating, etc.) by the developers, content providers and maintainers of the service. Such an extension is an area of further work for this model. In particular, questions of accurate allocation (the decision of how to share such emissions between the many services which IT equipment provides) require both further theoretical work and further study of user behavior. Specifically, this will involve analysis of how much time users spend on different services, and how much time devices stay in energy using idle or sleep states

6.8 Short Term, Longer Term and Systemic Impacts of Service Use

In the methodology presented in this paper, in line with standard practice in carbon footprinting, we have adopted an attributional approach. This means sharing out all impacts between the services involved. The disadvantage with this approach is that it does not distinguish between reductions which are immediately realizable and those that are not. For example, using a low power device instead of a higher power device results in an immediately realizable reduction in energy use. Reducing data traffic through a 3G network, on the other hand, may not, as a base station may use nearly as much energy to serve a lower quantity of data. However, such a change does lead to a longer term reduction in impacts if being part of a bigger trend - in that it will decrease the pressure to add new equipment.

There are more subtle long-term systemic trends associated with IT energy use which are notoriously difficult to model. For example, as observed by Blevis [9], the provision of digital services can contribute to the uptake of more digital devices. When making a decision of whether, for example, to promote

services on low-power eReaders, organizations such as GNM need to consider the trade-off between the efficiency gains from using such a device against the increase in manufacturing emissions resulting from any increased device purchases motivated by the new service. However, it is very difficult to estimate how much increase in demand for eReaders could be attributed to the new GNM service, and would require modeling and study using econometric techniques. Another broader issue for investigation is the extent to which new services on low power devices displace activity on other higher power devices, and to what extent it results in additional device usage. Our analysis of the eReader intervention assumes that the GNM product does not stimulate new product purchases, and time spent on the eReader correspondingly reduces time using a laptop or desktop. Further work on user behavior is needed to determine if this assumption is appropriate, and if not how best to model these systemic effects.

7. CONCLUSIONS

In this paper, we have presented a methodology that combines user behavior analysis with life cycle analysis to provide the most detailed and extensible carbon and energy footprint of a digital product suite produced to date. The carbon footprinting methodology we have presented is the first that is detailed enough to be able to assess the impact of alternative design decisions in digital products. We demonstrate this by applying it to digital services deployed by GNM, and assess the impact of the current services and potential changes resulting from six alternate interventions. The methodology, and many of the data points we have gathered, can be applied directly to many other digital services with similar delivery architectures, and could be extended to cover services with more complex architectures such as P2P.

The methodology and models we have developed have application in several areas. Firstly, they can be used to support and inform design decisions made by service providers. Secondly, they can be used as an educational tool to make software developers more aware of how design decisions they make impact energy use of the final product. Thirdly, they can be used to guide research in energy-efficient and green design. In particular, the results of our study emphasize the value of work focused on reduction of energy use by services on end user devices.

8. 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 would like to acknowledge the contribution of time and data made by Guardian News and Media, particularly Matthew Malthouse, Christopher Hodgson and Jo Confino.

9. REFERENCES

- [1] Akamai Environmental Sustainability Initiative Report 2011: 2012. http://www.akamai.com/dl/sustainability/Environmental_Sustainability.pdf. Accessed: 2012-11-22.
- [2] Aksanli, B. et al. 2012. Using Datacenter Simulation to Evaluate Green Energy Integration. *Computer*.
- [3] Aleksić, S. and Lovrić, A. 2010. Power efficiency in wired access networks. *e & i Elektrotechnik und Informationstechnik*. 127, 11 (2010), 321–326.
- [4] Apple - Environment - Reports: 2011. <http://www.apple.com/environment/reports/>. Accessed: 2011-11-22.
- [5] Baliga, J. et al. 2009. Energy Consumption in Optical IP Networks. *Journal of Lightwave Technology*. 27, 13 (Jul. 2009), 2391–2403.
- [6] Beauvisage, T. 2009. Computer Usage in Daily Life. *Proceedings of the 27th international conference on Human factors in computing systems* (Boston, MA, 2009), 575–584.
- [7] Bianzino, A.P. et al. 2012. A Survey of Green Networking Research. *Communications Surveys & Tutorials, IEEE*. 14, 1 (2012), 3–20.
- [8] Bircher, W.L. and John, L.K. 2012. Complete System Power Estimation Using Processor Performance Events. *IEEE Transactions on Computers*. 61, 4 (Apr. 2012), 563–577.
- [9] Blevins, E. 2007. Sustainable interaction design: invention & disposal, renewal & reuse. *Proceedings of the SIGCHI conference on Human factors in computing systems* (2007), 503–512.
- [10] Carroll, A. and Heiser, G. 2010. An Analysis of Power Consumption in a Smartphone. *USENIX annual technical conference* (Berkeley, CA, 2010), 21–21.
- [11] Chandaria, J. et al. 2011. The carbon footprint of watching television, comparing digital terrestrial television with video-on-demand. *Sustainable Systems and Technology (ISSST), 2011 IEEE International Symposium on* (2011), 1–6.
- [12] Clifton, B. 2012. *Advanced Web Metrics with Google Analytics*. John Wiley & Sons.
- [13] Computers: 2011. http://www.energystar.gov/index.cfm?fuseaction=find_a_product.showProductGroup&pgw_code=CO. Accessed: 2011-11-22.
- [14] Da Costa, G. and Hlavacs, H. 2010. Methodology of measurement for energy consumption of applications. *Energy Efficient Grids, Clouds and Clusters Workshop (co-located with Grid)(E2GC2 2010)* (2010).
- [15] Deruyck, M. et al. 2010. Comparison of power consumption of mobile WiMAX, HSPA and LTE access networks. *6th Conference on Telecommunication Techno-Economics, 2010. CTTE 2010* (2010), 1–7.
- [16] Dick, M. et al. 2012. Green Web Engineering - Measurements and Findings. *EnviroInfo* (2012).
- [17] Displays: 2012. http://www.energystar.gov/index.cfm?fuseaction=find_a_product.showProductGroup&pgw_code=MO. Accessed: 2012-05-24.
- [18] EARTH Project 2010. *Energy efficiency analysis of the reference systems, areas of improvements and target breakdown*.
- [19] Enterprise Servers: 2011. http://www.energystar.gov/index.cfm?fuseaction=find_a_product.showProductGroup&pgw_code=DC. Accessed: 2011-07-21.

- [20] European Commission - Joint Research Centre - Institute for Environment and Sustainability 2010. International Reference Life Cycle Data System (ILCD) Handbook - General guide for Life Cycle Assessment - Detailed guidance. Publications Office of the European Union.
- [21] Frachtenberg, E. et al. 2011. High-efficiency server design. *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis on - SC '11* (New York, New York, USA, 2011), 1.
- [22] Garg, S.K. and Buyya, R. 2011. NetworkCloudSim: Modelling Parallel Applications in Cloud Simulations. *2011 Fourth IEEE International Conference on Utility and Cloud Computing* (Dec. 2011), 105–113.
- [23] GHG Protocol Product Life Cycle Accounting and Reporting Standard ICT Sector Guidance: 2012. <http://www.ghgprotocol.org/feature/ghg-protocol-product-life-cycle-accounting-and-reporting-standard-ict-sector-guidance>. Accessed: 2012-05-30.
- [24] GNM Sustainability Report 2011: 2011. <http://www.guardian.co.uk/sustainability/series/sustainability-report-2011>. Accessed: 2012-09-19.
- [25] Google data center efficiency measurements: 2011. <http://www.google.com/about/datacenters/inside/efficiency/index.html>. Accessed: 2012-07-06.
- [26] International Organization For Standardization 2006. ISO 14040: Environmental management — Life cycle assessment — Principles and framework. ISO copyright office.
- [27] ITU 2012. *L.1410 - Methodology for environmental impact assessment of information and communication technologies (ICT) goods, networks and services*.
- [28] Joulemeter: 2012. <http://research.microsoft.com/en-us/downloads/fe9e10c5-5c5b-450c-a674-daf55565f794/>. Accessed: 2012-01-24.
- [29] Kazandjieva, M. et al. 2012. Green Enterprise Computing Data: Assumptions and Realities. *Proceedings of the Third International Green Computing Conference* (2012).
- [30] Lanzisera, S. et al. 2010. Data network equipment energy use and savings potential in buildings. *ACEEE Summer Study on Energy Efficiency in Buildings* (2010), 1–14.
- [31] Lin, M. et al. 2011. Dynamic right-sizing for power-proportional data centers. *IEEE INFOCOM* (2011).
- [32] Malmodin, J. et al. 2010. Greenhouse Gas Emissions and Operational Electricity Use in the ICT and Entertainment & Media Sectors. *Journal of Industrial Ecology*. 14, 5 (Oct. 2010), 770–790.
- [33] Malmodin, J. et al. 2012. LCA of data transmission and IP core networks. *Electronics Goes Green 2012+* (2012), 1–6.
- [34] McLachlan, R. and Brewster, S. 2012. Towards new widgets to reduce PC power consumption. *Proceedings of the 2012 ACM annual conference extended abstracts on Human Factors in Computing Systems Extended Abstracts - CHI EA '12* (New York, New York, USA, 2012), 2153.
- [35] Moberg, Å. et al. 2010. Printed and tablet e-paper newspaper from an environmental perspective — A screening life cycle assessment. *Environmental Impact Assessment Review*. 30, 3 (Apr. 2010), 177–191.
- [36] Naumann, S. et al. 2011. The GREENSOFT Model: A reference model for green and sustainable software and its engineering. *Sustainable Computing: Informatics and Systems*. 1, 4 (Dec. 2011), 294–304.
- [37] OECD Broadband Portal: 2011. <http://goo.gl/Q30jB>. Accessed: 2012-08-27.
- [38] Ofcom 2012. *Communications Market Report 2012*.
- [39] Preist, C. and Shabajee, P. 2010. Energy Use in the Media Cloud. *2nd IEEE International Conference on Cloud Computing Technology and Science* (Indianapolis, USA, 2010).
- [40] Roth, K.W. et al. 2002. Energy Consumption by Commercial Office and Telecommunication Equipment. *2002 ACEEE Summer Study on Energy Efficiency in Buildings* (2002).
- [41] Schien, D. et al. 2012. Impact of Location on the Energy Footprint of Digital Media. *IEEE International Symposium on Sustainable Systems and Technology (IEEE ISSST 2012)* (Boston, MA, 2012).
- [42] Schien, D. et al. 2013. Modeling and Assessing Variability in Use Phase Energy of Online Multimedia Services. *Journal of Industrial Ecology*. [to appear] (2013).
- [43] Simons, R. and Pras, A. 2010. *The Hidden Energy Cost of Web Advertising*. Citeseer.
- [44] Small Network Equipment: 2012. http://www.energystar.gov/index.cfm?c=new_specs.small_network equip. Accessed: 2012-10-09.
- [45] SPEC power_ssj2008: 2011. http://www.spec.org/power_ssj2008/. Accessed: 2011-02-09.
- [46] Taylor, C. and Koomey, J.G. 2008. Estimating Energy Use and Greenhouse Gas Emissions of Internet Advertising. *Network*. Working paper for IMC2.
- [47] Teehan, P. et al. 2010. Estimating the Changing Environmental Impacts of ICT-Based Tasks : A Top-Down Approach. *Sustainable Systems and Technology (ISSST), 2010 IEEE International Symposium on* (2010).
- [48] Thiagarajan, N. et al. 2012. Who killed my battery?: analyzing mobile browser energy consumption. *WWW* (2012), 41–50.
- [49] Weber, C.L. et al. 2009. *The energy and climate change impacts of different music delivery methods*.
- [50] Williams, D.R. and Tang, Y. 2012. A Methodology to Model the Environmental Impacts of Electronic Software Distributions. *Environmental Science and Technology*. 46, 2 (2012), 1087–1095.