Power Profiling the Internet Core: a Sensitivity Analysis

Aruna Prem Bianzino, Anand Kishore Raju, Dario Rossi
Telecom ParisTech, Paris, France
Email: {bianzino, raju, rossi}@telecom-paristech.fr

ABSTRACT

As in many scientific domains, the accuracy of results in green networking research depends on the accuracy of the underlying energy models adopted in the study. On the one hand, researchers are in agreement regarding the need for verifiability of such models. On the other hand, they have yet to agree on an effective validation methodology to quantify the reliability of the estimated results. Often, rather different estimates are gathered, via different models, for the same power figure – which holds especially true whenever large-scale networks, as opposed to individual devices, are considered.

In this chapter, we perform a careful sensitivity analysis of a power model for the Internet core: our results show that, no matter how carefully the data upon which the power-consumption model relies is chosen and cross-verified, the uncertainty of the overall results remains disappointingly high. We believe that part of the solution lays in a community-wide effort, to which we offer some initial guidance that could, if not solve the issue, at least greatly improve the current situation.

INTRODUCTION

Research on green ICT is evidently gaining momentum: a rich literature exists and continues to grow on energy-thrifty networks solutions and on network power profiling analysis. As a comprehensive overview of network energy-efficiency issues is out of the scope of this chapter, we refer the reader to the literature survey by Bianzino et al. (2011) or to earlier chapters of this book.

One of the most challenging aspects in green research is to gather a set of energy-related assumptions, such as device power profiles, which is both accurate for present-day devices and future-proof as well. Indeed, even considering present-day systems, it is often difficult to estimate various aspects of power measurements – as for example the power required by cooling in Katz (2009), or discrepancies between actual power drain and maximum drain reported by the equipment manufacturer in Juniper (2009a), and so on. Considering instead futuristic scenarios, it is clear that major advances originating from other areas are hardly foreseeable, but may deeply impact the overall results – as has been in recent years for voltage scaling (Weiser et al., 1994) from a hardware perspective, for tickless kernels (LessWatts Project, 2009) from a software standpoint, and for adaptive link rate (Christensen et al.,2010) from the communication networking front.

To further add complexity to the picture, there are many instances of inaccurate or grossly misinterpreted results published by energy analysts and media. Moreover, errors can easily propagate, as wrong numbers replicate themselves through direct references and citations, possibly under the camouflage due to manipulation. The above statement is especially true when the estimations involve rapidly changing technologies, requiring constant updates and time-to-time verifications. Developing methodologies and performing measurement is often challenging, e.g., because of precluded access to the real infrastructures, or proprietary data-centers, very large-scale networks spanning multiple domains, etc. Thus, unfortunately, gauging the exact nuance of *green* in the ICT context is inherently uncertain, hence prone to fallacies.

Evidence of existing fallacies in the literature are uncovered only sporadically, if ever. The best-known example is represented by a Forbes magazine article by Huber and Mills (1999), which claimed that PCs and networked devices were responsible for 8% of all electricity consumed across USA. Huber and Mills (1999) also projected a staggering growth up to 50% of all electricity usage in the following 10-20 years. These numbers were widely published as well as publicized by the media, but were later debunked by a study conducted at Lawrence Berkeley National Laboratory (LBNL) by Koomey et al. (2002), which resized the estimation to about 3%, for all office, telecommunications, and network equipment.

Along similar lines, more recently, an article published by Wissner-Gross (2009) on BBC News estimated that a couple of Google searches on a desktop computer produces about 14 g of CO_2, which is roughly the equivalent of boiling an electric kettle. These figures were later countered by Google (2009), claiming instead that a typical search produces only 0.2 g of CO_2. Relative overestimation in this case is in the order of 70×, which raises the question on how to disprove the erroneous figures, or at least whether the gap in such diverse estimations could be significantly narrowed down, e.g., by using more accurate input values for power models.

In our work, we investigate a similar issue, i.e., namely a disproportionate gap into power profiles of core IP technologies, and more precisely focus on the estimation of the Internet energy consumption. This is a challenging task, first because of the widespread extent of the Internet, followed by the number of assumptions that are needed to define such a complex model. The reminder of this chapter is organized as follows. First, we overview different work focusing on the issue of measuring the power rating of either a single device, or a set of networked devices. Afterward, we closely investigate one of such models, illustrating the challenges that lie in a reliable calibration of the model output. We then perform a careful sensitivity analysis of the model: special effort is made on collecting verifiable data which serves as underlying assumption and inputs, rather than exploring the widest possible spectrum of parameters. Similarly, our aim is more on gauging the precision in the estimation, i.e., finding the *boundaries* between which the actual value should be reasonably found, rather than yielding a single value, though accurate it may be. Our results show that, even when great care is taken in defining a scenario which is as realistic as possible, the uncertainty in the underlying data propagates to very large errors in the final result (up to 50-250%). Based on these findings, we then propose community-wide directions that could help in improving such a disappointing situation, prior to summarize the chapter.

BACKGROUND

Inconsistencies as those arising in the context of PCs between Huber and Mills (1999) and Koomey et al. (2002), or in the context of search queries between Wissner-Gross (2009) and Google (2009), are typical in the real world and are an accepted part of the results originating from mass media. On the other hand, scientific publications should be worthy of a higher degree of trust. To convince that the propagation of spurious scientific results can have rather serious implications, it is sufficient to consider the negative side effects in the case of nation-wide policy decisions made on the basis of inconsistent scientific results, as rightly pointed out by Koomey et al. (2002).

At least in part, some of these errors can be easily avoided: combination of careless reporting and reliance on secondary sources are very often solely responsible for the occurrence of errors. In several documents we overviewed, we found for instance: that the same reference was cited for different power-figures in different documents; that documents mistakenly reported the value of the figure they referenced; that two versions of the same document from the same author reported different values for the same figure; that the reported values were clearly affected by a typo; etc.

However, even when the methodology is technically sound and carefully proof-checked against the above flaws, a given degree of uncertainty always remains. To testify this intrinsic difficulty, we next overview the available literature on two different perspectives on energy measurement, namely considering (i) a single router and (ii) the whole Internet.

Table 1: Network energy cost

Context	Source	Energy Cost	
		Minimum	Maximum
Router (experimental)	Chabarek et al. (2008)	0.31 μJ/bit	
Router (experimental)	Qureshi et al. (2009)	0.43 μJ/bit	
Internet (microscopic)	Baliga et al. (2009)	2 μJ/bit	75 μJ/bit
Internet (macroscopic)	Gupta and Singh (2003)	16.0 mJ/bit	28.1 mJ/bit

We summarize the available estimations of ICT energy cost in Tab. 1, gathered from reliable and reputed scientific literature, namely: Baliga et al. (2009); Chabarek et al. (2008); Gupta and Singh (2003); Qureshi et al. (2009). It has to be noted that estimates in the table are not all directly comparable, as they are obtained using a different evaluation methodology, as they rely on different input data-sets, and as they possibly targets a different aspect of ICT energy consumption. Specifically, a microscopic modeling technique is adopted by Baliga et al. (2009), macroscopic modeling is instead exploited by Gupta and Singh (2003), while experimental measurements are employed by Chabarek et al. (2008) and Qureshi et al. (2009).

Power Models of IP Core Routers

As we previously introduced, power profiling a single device is a relatively simple task, at least with respect to the issue of profiling a complex set of devices such as the Internet. Yet, even in this case some difficulty may arise: indeed, there exist many classes of routers, suited for a given network segments from the access to the core, with rather different capabilities and, hence likely different power profiles. Moreover, due to the large plethora of devices, it is hard to gather measurements that are representative of all devices in a given class.

To date, only a few experimental studies of IP routers are available in the literature such as those performed by Chabarek et al. (2008), by Qureshi et al. (2009) and by Juniper (2009a). For instance, Chabarek et al. (2008) experimentally measures the energy cost of transmitting data through an IP router, considering two routers models (namely, Cisco GSR 12008 and Cisco 7507). The measurement campaign yields an average energy consumption of 770 W for 540 Kpkt/s, with packet size of 576 Bytes (mid-size), resulting in 0.31 μJ/bit. Considering the same router set, Qureshi et al. (2009) derived an energy footprint of 2 mJ per packet, that corresponds to 0.43 μJ/bit, again considering mid-size packets. While the above work is substantially in agreement, this partly follows from the coherence of the router dataset.

Besides, the measurement in Chabarek et al. (2008); Qureshi et al. (2009) focus on *access routers*, configured with limited line-cards (four 1 Gbps ports) and limited switching capabilities (2.5 Gbps). While these routers are ideal candidates for modeling the devices that are currently deployed at the edge of a network (e.g., 4 Gbps can roughly handle a residential population of about 5,000 ADSL2+ users with 20 Mbps downlink capacity using an oversubscription factor of 25), they clearly fail to represent aggregation and core IP routers. Lack of sufficient experimental measurements and data availability for energy consumption of core routers in the literature further forces us to adopt their *nameplate consumption*, i.e., the power rating advertised by the manufacturer, which can be gathered directly from the vendor catalog. However, this choice may not be accurate, as manufacturers in their data-sheets often *overestimate* the advertised power budget, so as to reduce the occurrence of power breakdown events, as pointed out by the ECR Initiative (2009). Hence, an accurate power profiling of edge and core IP routers requires to take into account several aspects, which we will attempt at doing in the following sections.

Power Models of the Internet

If already power-profiling a single device poses some difficulties, power-profiling the whole Internet is clearly a much more challenging issue, due to the large number and heterogeneity of the devices involved. So far, only a few researchers have taken up the challenge, namely Gupta and Singh (2003) and Baliga et al. (2009), following two complementary approaches.

While the main focus of Gupta and Singh (2003) is the adaptation of Ethernet link rate during low traffic period, so to reduce the LAN power budget, their work initially provide an estimate of Internet power profile to motivate the interest of their proposed approach. In more detail, Gupta and Singh (2003) use a *macroscopic* approach to evaluate the energy cost of the data transmission in public Internet, which is expressed as the ratio of the total energy consumed by ICT devices during 2002 (as estimated by Roth et al. (2002)) over the total Internet traffic transferred in the same year (as reported by Schulzrinne (2010)). Hence, to calibrate the model not much can be done other than verifying (and possibly correcting) the estimates of Roth et al. (2002) and Schulzrinne (2010).

An orthogonal approach is instead adopted by Baliga et al. (2009), whose main focus is to design a power-profile model for a large scale network (such as the Internet), comprising several access technologies carried over an optical core. The model conceived by Baliga et al. (2009) is *microscopic* in the sense that it starts considering all network components at an individual level, integrating thus the power contribution of each component in the final estimation. The model, other than being very detailed from a technological standpoint, is also *conservative* by design, in order to avoid raising false alarms as done previously by Huber and Mills (1999) and by Wissner-Gross (2009). The model proposes an estimated range of energy cost for the data transmission across the network, which is evaluated to be as low as 2μJ/bit, and as high as 75μJ/bit depending on the access rates.

Clearly, as it can be seen from Tab. 1, different perspectives (e.g., single device vs Internet) lead to rather different power figures, which is expected. At the same time, even when the same perspective is considered, different approaches (e.g., microscopic vs macroscopic) may again lead to significant differences, that possibly amounts to several orders of magnitude. Notice that, even though technical development in the time dimension (i.e., improvement of device power efficiency) has surely contributed to the reduction of the energy cost per bit between 2003 and 2009, it cannot however account for the 10,000 fold improvement over the same time lapse (rather, a 10- to 100-fold energy efficient improvement over a 15-years time window is forecasted by Baliga et al., 2009).

MODELING THE INTERNET POWER PROFILE

With all evidences pointing towards the large gap that exists in the estimation of the Internet power profile, this raises the need for further investigation: as the macroscopic approach is however rather inflexible, we therefore decide to perform a careful sensitivity analysis of the microscopic model proposed by Baliga et al. (2009).

Model Motivation

Before we describe the model, we need however to illustrate the relevance of our choice. We decide to focus on a power-consumption model of the Internet, since its large-scale deployment makes it a very challenging task, which is also far less studied than, e.g., the power-modeling of a single device. We also choose to focus on a specific portion of the network, namely the Internet core. While it could be argued that focusing on absolute power figure estimates of a specific network segment is rather a sterile exercise, this issue is instead very relevant for the following reasons.

First of all, we consider the absolute estimation of network core energy consumption as an intermediate step towards a more interesting problem: i.e., the relative assessment of component-wise network energy consumption. Such a relative viewpoint is extremely important, as it allows to identify the network segment which has the larger energy footprint: in turn, this would pinpoint where large green optimization gain could occur, thus providing a stepping stone towards further research efforts.

Second, although for a long time the access portion has been considered as the biggest contributor to the overall Internet energy consumption, however some researchers expect this trend is to change in future. For instance a Deutsche Telekom study conducted by Lange (2009) forecasts that by year 2017, the energy consumed by the network core will be equal to that of the network access. The study also suggested a stunning 300% rise in power rating of the network core in coming decade, mainly due to the IP/MPLS layers. Yet, as another recent study by Bolla et al. (2011) suggests that the core network consumption will instead play a minor role with respect to the other network segments, this issue needs further attention.

Third, we also point out that the access-segment is much more heterogeneous due to the large number of available access technologies: hence, this larger variability translates not only into a larger cross-validation effort, but also into a higher chance of making the very similar mistakes that were earlier outlined.

We thus undertake a sensitivity analysis of the Internet core power profile as a first, necessary step toward a more comprehensive assessment of the whole network power profile. As starting point of our work, we consider the model proposed in Baliga et al. (2009). To the best of our knowledge, this work represents the only network power profiling model available in the literature defined at a *microscopic* level – i.e., it considers individual elements (e.g., router, access devices, etc.) to extrapolate overall figures concerning the whole network. This makes it a very interesting, rich and detailed model, that, as opposed to macroscopic modeling, further allows to pinpoint the contribution of several parameters in the overall power figures.

Model Overview

We point out that the energy model proposed in Baliga et al. (2009) aims at *conservatively* estimating the power-per-customer P_{core} expenditure by the Internet infrastructure. Since the interconnection structure of a core router is highly dependent on the nature of the traffic, the authors evaluate the power rating of the network core by taking into account the total switching capacity required to support the traffic and the average number of routers through which traffic transits on its end-to-end path. The component that accounts for the core segment is then expressed as:

$$P_{core} = (H+1)\frac{A}{C_{router}}\gamma \; P_{router} \qquad (1)$$

where H represents the number of core hops traversed by the traffic, A represents the minimum access rate for public Internet (considering a fixed oversubscription rate of 25), and C_{router} and P_{router} represent the capacity and power rating of the specific routing device considered. Finally, a multiplicative factor γ=8 is used to account for cooling, redundancy and future growth (2×2×2).

The model expresses the power-per-customer as a function of the desired minimum access rate A, and fixes all the other parameters to some reasonable value available from data sheets or empirical studies. Specifically, the authors fix the number of core hops to H=10 from measurements by Van Mieghem (2006), and consider a Cisco CSR-1 core router with P_{router}=10,900 W and C_{router}=640 $Gbps$, using data reported by Cisco (2004).

The model (1) defined in Baliga et al. (2009) is subject to a number of assumptions, which we aim to cross-validate in the following section. Irrespectively of the fact that the assumptions are indeed reasonable and are well motivated, they nevertheless constitute a simplification of a much more complex reality. For instance, the value of H may have changed since the hop count measurement performed by Van Mieghem (2006). Also, Cisco is perhaps the most popular core router manufacturer, but other manufacturers also exist (such as Juniper, Alcatel-Lucent and Huawei, to mention a few). Moreover, as the model intends to propose a conservative lower bound estimation of the power profile, the choice of parameters is also crucial in determining the correctness and validity of the original aim.

Notice that our purpose is not to invalidate, debunk or confute the findings in Baliga et al. (2009). Rather, taking (1) as starting point, we aim at quantifying the possible extent of errors in the power model. Indeed, notice that in (1), individual errors in the estimation of any factor γ, H, P_{router} linearly propagate as an error in the estimation of P_{core}. However, several errors may occur at the same time, their effects are multiplied when the model is considered. Thus, it is imperative to assess the *actual boundaries* within which the Internet power-per-user P_{core} can be estimated.

While for some factors such as H, P_{router} several estimations exist and can be cross-checked, other factors may be harder to estimate: for instance, γ depends on choices (e.g., cooling, future growth, etc.) made by the facility managers and telecom operators, that are clearly not publicly available and thus difficult to verify. In the following, we will dig further all the relevant parameters described so far.

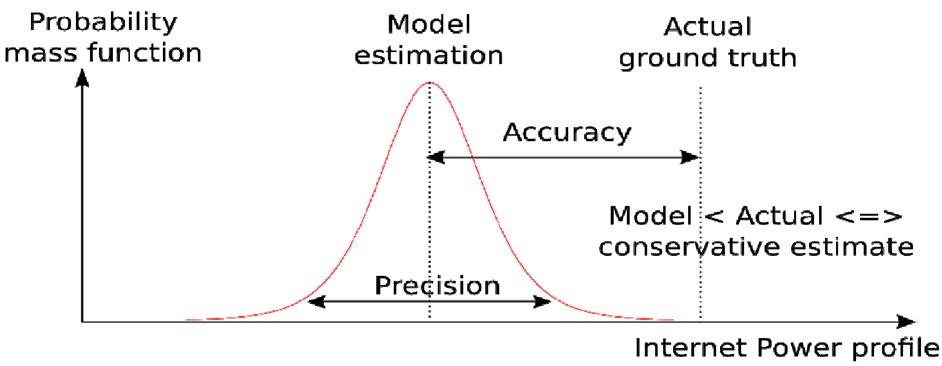

Figure 1: Pictorial representation of the Internet power profile sensitivity analysis

SENSITIVITY ANALYSIS OF INTERNET POWER PROFILE

In this section, we perform a sensitivity analysis of (1) considering several relevant factors that were earlier outlined. From a high level point of view, we first consider each factor in isolation, and then assess their joint effect. As far as each factor is concerned, we are more interested in carefully evaluating the boundaries between which it can reasonably vary, more than just arbitrarily selecting a carefully tuned fixed value to refine the estimation. This way, our sensitivity analysis will not merely yield a refinement of the estimation *accuracy* of Baliga et al. (2009), but also assess the achievable *precision* in the estimate as represented in Figure 1. Notice that significant effort is devoted in Baliga et al. (2009) to precisely assess the time dimension in the power profile evolution: in more detail, authors speculate on power profile evolution during a window of 15 years (2008-2023), considering the forecasted growth of different parameters (e.g., such as the Internet traffic volume growth, trends in router capacity, etc.). In this work, we instead adopt a different viewpoint and, focusing our attention to a specific time snapshot, we perform an assessment of many different parameters that make the Internet power profile estimation challenging.

Specifically, (1) models the changing topology of the Internet in terms of the average number H of IP core routers that packets traverse in their end-to-end journey from a client host to a server in the Cloud. In this section, we closely examine the effect of the Internet topology, by surveying different experimental measurement work that have targeted the Internet topology study, to refine and cross-check the model input parameter H. Then, notice that (1) expresses the impact of core router devices assuming an (i) homogeneous Cisco router population, (ii) whose power drain is furthermore modeled according to vendor nameplate rating. In this section, we therefore address both these issue, by considering (i) an heterogeneous Cisco vs Juniper population of routers, taking special care in defining a relevant breakdown parameter α for the router population, and (ii) refining the power model to address the gap between nameplate rating and the actual power rating, expressed by the factor β. Besides, (1) also accounts for a number of additional factors, such as cooling, redundancy and future growth, that are lumped in the factor γ, and that we also consider later on in this section.

Summarizing, this section reports a thorough evaluation of the impact of different parameters on the accuracy and precision of the model (1). These factors are first inspected in isolation, as follows:

- Topological assumptions
  - Number of IP core hops (H)
- IP router assumptions
  - Breakdown of an heterogeneous router population (α)
  - Nameplate vs actual power rating (β)
- Further assumptions
  - Cooling, redundancy and future growth (γ)

Then, we report on the *joint* effect of all the above factors, refining the estimate and evaluating its boundaries.

Impact of Core Hops Number

As the first parameter of the sensitivity analysis, we consider a plausible range of values for H in the model. In general, the average hopcount in the network core depend on two factors: on the one hand, we have users, with their preferences, the locality and popularity of the content they seek, and the traffic they generate; on the other hand, we have the Autonomous System (AS) network structure, and policies to confine the traffic in the proximity at possibly different layers of the networking stack (e.g., caching, CDN, application-layer proximity-aware peer selection of P2P overlays).

Due to the complexity in modeling the above factors, active or passive measurement campaigns are typically used to experimentally assess the distance traveled by network traffic. In order to establish a plausible range for the parameter *H*, we conducted a survey on available research studies focusing on hop count measurements. We summarize our findings in Tab. 2, ordered on the basis of the time-line at which the measurement campaign was effectuated: notice that the different work samples a snapshot of the Internet at different years, spanning over the decade 1998-2009. To be consistent with Baliga et al. (2009), we also assume that it takes 3 to 6 hops for data to reach the core from the access: thus, Tab. 2 suggests the boundaries for the variable *H* as $6 \le H \le 16$.

Table 2: Average Hop Count for end-to-end Internet data transfer

Source	Year	Source	Dest.	Hops	Core	Tools
Fei et al. (1998)	1999	US US US	US APAC Europe	13 21 26	9 17 22	Traceroute
Van Mieghem et al. (2001)	2001	Europe Europe Europe	Europe APAC US	14.5 19.5 15	11 16 12	Traceroute
Donnet et al. (2006)	2004	N.A	N.A	17	12	Traceroute
Van Mieghem (2006), adopted by Baliga et al. (2009)	2005	N.A	N.A	13	10	N.A
Valancius et al. (2009)	2009	AS Country	AS Country	9 14	6 10	DipZoom

Notice that the number of core hops actually diminishes over time: indeed, recent studies all consistently report a lower hop count number with respect to the previous ones. This is the result of two different root causes. On the one hand, older studies considered inter-continental distances between US-APAC-Europe as done by Fei et al. (1998) or between Europe-APAC as done by Van Mieghem and Begtašević (2001), while more recent studies are confined in a single continent or AS – hence, part of this shortening may be due to an implicit bias in the measurement and thus may be only apparent.

On the other hand, the lower number of hops of recent studies such as Valancius et al. (2009) can be explained in terms of an evolution of the Internet structure at the AS level. As suggested by Haddadi et al. (2010), the Internet core is shrinking due to an increase in peering between Tier-2 ASs: as they avoid to transiting through Tier-1 ASs, this results in an overall reduction in the path length for the Internet core.

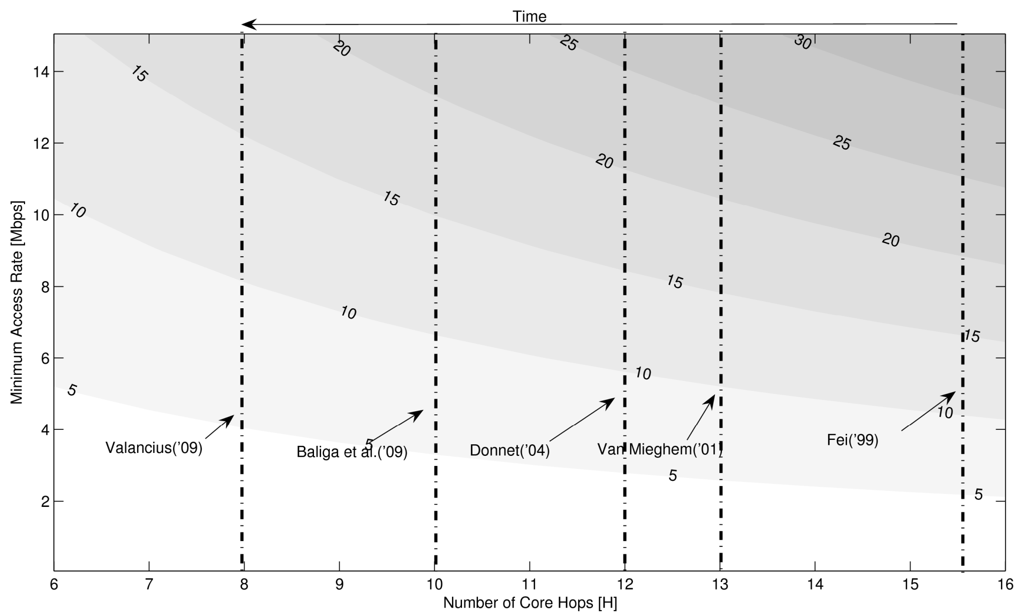

Figure 2: Contour plot of Power-per-user (Watts) consumed by the Internet core for different hop count *H* and minimum access rates (x-axis) *A* values (y-axis)

Fig. 2 shows the results of the topology sensitivity analysis, performed on the energy model in response to variations of parameters *H* and *A*, depicted as a contour plot. Notice that in the last decade, for a given power-per-user the decrease in hop counts corresponds to an increase in minimum access rate *A* for the same power level (or, equivalently, that the power-per-user needed to sustain a given minimum access rate *A* has decreased). For instance, according to the results, in the year 1999, 10 W per user were spent to deliver a minimum access rate of 4 Mbps, whereas the shrink of the network 10 years later allows the network operator to sustainably offer twice the capacity within same power budget (all other parameters being unchanged).

Since (1) is linearly dependent on *H*, errors in the estimation of *H* linearly propagates as errors in the estimation of P_{core}. With respect to the *H*=10 value chosen in Baliga et al. (2009), experiments in Tab. 2 give us the lower and upper bounds for the variable *H*: overall, these work suggest that P_{core} could be *overestimated* by 40% (*H*=6), as well as *underestimated* by 60% (*H*=16). At the same time, we point out that, if as Haddadi et al. (2010) suggest the trend for *H* is to shrink over time, then the conservative power estimate of Baliga et al. (2009) could be instead overestimated.

Besides, notice that even though the number of core hops is not directly related to power ratings, nevertheless it significantly impacts the considered power model. This observation highlight the need for *reliable input data* to quantitatively gather accurate, reliable and authoritative output power figures– an important issue on which we will come back in the solution and recommendations section.

Impact of Heterogeneous Router Population

Next, we refine assumptions concerning the type and number of router devices in the Internet core. Authors in Baliga et al. (2009) assume all core routers to be Cisco CRS-1. Furthermore, as far as the CSR-1 power profile is concerned, they assume a *nameplate consumption*, which can be adopted directly from the vendor catalog.

Both these assumptions have a strong impact on the estimation of the core network energy consumption. Considering market shares, we highlight that assuming a homogeneous router population is unrealistic given reports such as Current Analysis (2010) that study the market share of different equipment vendors. Furthermore, as ECR Initiative (2009) points out, manufacturers data-sheets often overestimate the power budget to reduce the occurrence of power breakdown events. Notice that the latter observation is especially critical in case the model of Baliga et al. (2009) is used to gather a conservative lower bound.

Therefore, in this subsection, we address the above issues by considering Juniper T1600 as an alternative to Cisco CRS-1 core routers. We therefore consider a *heterogeneous router population,* based on the Cisco vs Juniper market share estimate reported by Current Analysis (2010). Moreover, we refine the power model of the two considered routers, by taking into account the differences between their *nameplate vs actual ratings* adopted by experimental measurements published by manufacturers study such as Juniper (2009a).

Router Power Profile

In this section, building on Bianzino et al. (2010), we report a detailed power profile of the Cisco CRS-1 and Juniper T1600 core routers devices. In order to build a valid model for different configurations of the Cisco CRS-1 and Juniper T1600, it is necessary to consider the architectural and operational details of such devices. Notice that, for the time being, we focus on the interconnection and configuration issues of the device, excluding external factors (like cooling, power-redundancy and console systems), that are instead addressed later in this chapter. Our choice falls on these two router models since, as estimated by Current Analysis (2010), these account for the most significant fraction of the core IP router population.

Generally speaking, systems can be configured in *Single-Chassis* or *Multi-Chassis* mode. A basic Single-Chassis system is composed of several linecards (LC), such as Modular Service Cards (MSC) or Physical Interface Cards (PIC) linecards. A Multi-Chassis system includes instead multiple LCs, which are interconnected by one or more dedicated switching fabric (SF) chassis. Interconnection through SFs allow multi-chassis systems to scale up the aggregated system capacity. For example, as a Cisco CRS-1 switching fabric can interconnect up to 9 linecard shelves, a CRS-1 Multi-Chassis System can support an array of 72 linecard shelves interconnected by eight switching fabrics. In the following, we first develop separate power profiles models for the CRS-1 and T1600 routers configuration; then, we further develop a single unifying model for both router families. We point out that, while our numerical examples only refer to T1600 and CRS-1, the methodology outlined here applies in principle to *any* router (as the Single- vs Multi-Chassis alternatives comprise all the possible router configurations) for which power ratings of the Chassis, LC and SF sub-component are known.

Tab.3 summarizes the power rating of the individual sub-components of the CRS-1 and T1600 systems. Notice that, for the time being, we consider nameplate ratings reported from vendors, which are typically overestimated (i.e., well above the actual power values) thereby ensuring a safe system operation by reducing the odds that power-breakdowns occur. Later on, we will drop the nameplate rating assumption to further refine the estimations. Typically, the manufacturers directly provide the power rating of LC and SF sub-components, as reported in Tab. 3. Conversely, the $P_{chassis}$ power rating of an empty chassis has to be determined by subtracting LC ratings from a "typical" router configuration. In more detail, in Tab. 3. the power consumption of 16 line cards has to be subtracted from the power consumption of the typical 640Gbps configuration (considered by Baliga et al., 2009), in order to gather the power $P_{chassis}$ of an empty chassis.

Table 3: Power footprint of individual components for Cisco CRS-1 and Juniper T1600 routers

Label	**Functionality**	**Power (W)**	**Source**
LC	Linecard (CRS-1) Linecard (T1600)	500 66	Cisco (2007) Juniper (2009c)
SF	Switching Fabric (CRS-1) Switching Fabric (T1600)	9100 12750	Cisco (2011) Juniper (2009b)
Typical	CRS-1 Single chassis 16LC, 640Gbps T1600 Single chassis 16LC, 640Gbps	10900 8350	Cisco (2011) Juniper (2009a)

A Cisco CRS-1 uses OC-768c/STM-256c linecards Cisco (2007), which can support a 40 Gbps throughput for a power of 500 W. Conversely, a Juniper T1600 equipped with a *Tx Matrix Plus* Juniper (2009b) switching fabric can interconnect 16 T1600 chassis into a single routing entity. T1600 uses OC-768c/STM-256c PIC linecards Juniper (2009c), which also provide a capacity of 40 Gbps for a power of 66 W.

Total power P_{total} of a Multi-Chassis system can be calculated by summing the power rating of each component: namely, the power $P_{chassis}$ of an empty chassis (i.e., without active linecards), plus the power P_{LC} of an active linecard installed in a linecard shelf, plus the power P_{SF} of the switching fabric used to interconnect the Multi-Chassis system.

From Tab. 3, we find the power rating of the OC-768c/STM-256c linecard (P_{LC}=500 W) and the switching fabric (P_{SF}=9,100 W) for a CRS-1. Since the power rating of the chassis $P_{chassis}$ is not publicly available, we derive it with the previously outlined methodology. To do so, we consider a fully equipped Cisco CRS-1 with a single chassis and 16 linecards, which Cisco (2011) reports to consume P_{total}=10,900 W. Ripping off 16 active linecards, each consuming 500 W, gives us a conservative upper bound of $P_{chassis}$=2,920 W for the empty CRS-1 chassis.

The above architectural details allow us to derive a model for the total power rating P_{total} of any configuration of a Cisco CRS-1 Multi-Chassis system:

$$P_{total}^{CRS-1}(n_{LC}) = n_{LC}P_{chassis} + 16\,n_{LC}P_{LC} + \left\lceil \frac{n_{LC}-1}{8} \right\rceil P_{SF} \qquad (2)$$

where $i \in [1,72] \subset N$ corresponds to the number of linecard chassis installed in the CRS-1 Multi-Chassis System. Notice that, for the sake of simplicity, our model assumes that once a linecard chassis is installed, it is fully utilized (i.e., it consumes $P_{chassis}$, plus P_{LC} for each of the 16 cards it supports). Also notice that a variable number $\left\lceil \frac{n_{LC}-1}{8} \right\rceil$ of SF chassis are needed to support n_{LC} linecard chassis. More precisely, the model states that any eight slot is occupied by a SF element needed for the interconnection, and the SF is not needed in case a single linecard chassis is used. Notice that this factor may change for other Multi-Chassis systems, notably depending on supported slots per chassis.

Using a similar profiling technique, we can derive the generic power model for the Juniper T1600 Multi-chassis System. From Tab. 3, we get $P_{LC} = 66\,\mathrm{W}$ and P_{SF}=12,750 W for the Juniper T1600. Interestingly, notice that while the Juniper SF is much more power-hungry than the Cisco counterpart, the opposite happens concerning linecards (which is due to Short Reach interfaces, that consume much less power): overall, it is thus hard to guess the global system power footprint, which furthermore depends on the (unknown) $P_{chassis}$. As before, we gather $P_{chassis}$ by ripping off the 16 installed linecards (each consuming 66 W) from a system equipped with a single-chassis that Juniper (2009a) reports to have P_{total}=8,350 W. Finally, we have:

$$P_{total}^{T1600}(n_{LC}) = n_{LC}P_{chassis} + 16\,n_{LC}P_{LC} + I_{(n_{LC}>1)}P_{SF} \qquad (3)$$

where I represents the identity function (i.e., $I_{(x>y)}$ equals to 1 when $x>y$ and 0 otherwise), $n_{LC} \in [1,16] \subset N$ corresponds to the number of linecard chassis installed in the T1600: notice that, unlike in the case of CRS-1, a T1600 Multi-Chassis system only support a SF which delimits its scalability to 16 linecard chassis. It is to be noted that the SF is needed only when more than a single linecard chassis is in use (i.e., $n_{LC}>1$).

Based on our findings, and the power profiling of two families of Core IP routers, the power model for a generic router, expressed as a function of the number n_{LC} of linecard chassis installed and of the number n_{slot} of slots per chassis that can be interconnected by a single SF, is then:

$$P_{total}(n_{\mathrm{LC}}, n_{slot}) = n_{LC}(P_{chassis} + 16P_{LC}) + \left\lceil \frac{n_{LC}-1}{n_{slot}} \right\rceil P_{SF} \qquad (4)$$

It is to be noted that that (4) degenerates in (2) and (3) when parameters are properly set (i.e., respectively, n_{slot}=8 and $n_{LC} \in [1,72]$ for the Cisco CRS-1, n_{slot}=16 and $n_{LC} \in [1,16]$ for the Juniper T1600).

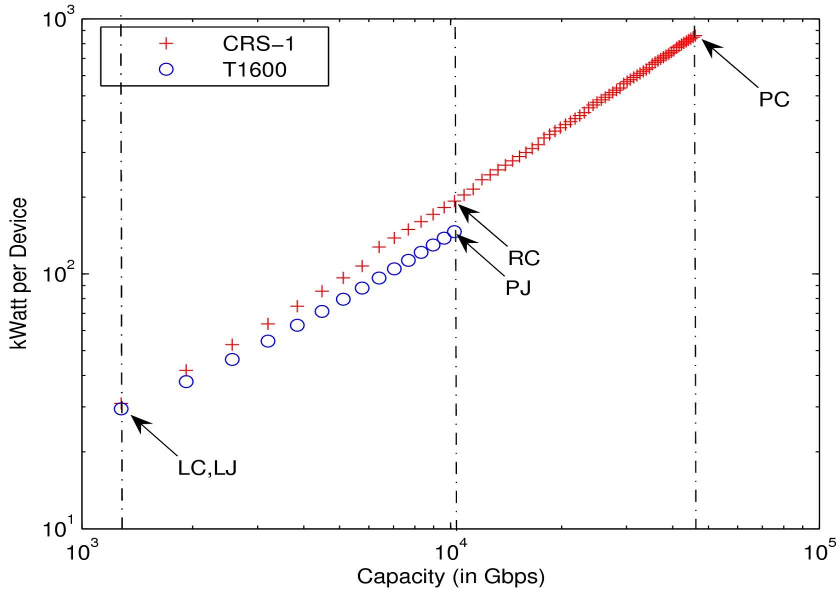

(a) Total power Vs Capacity

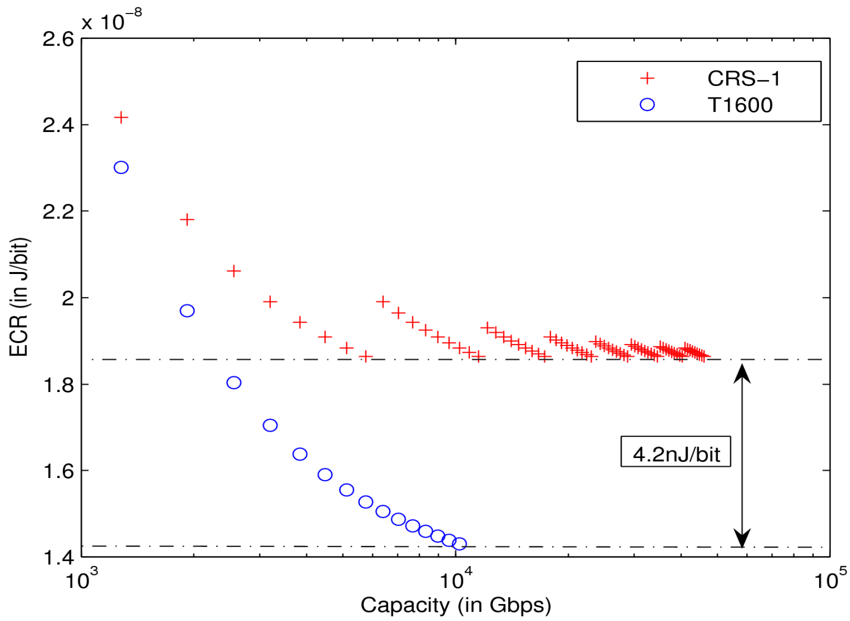

(b) Energy per bit Vs Capacity

Figure 3: Comparison of Cisco CRS-1 and Juniper T1600:
(a) Watts-per-device and (b) ECR metric as a function of the router aggregated capacity.

The resulting power profiles, for both Cisco CRS-1 and Juniper T1600, are shown in Fig. 3, which depicts the router power drain as a function of the aggregated system capacity achieved under different configurations. Notice that in both cases the system capacity can be expressed as

$C_{total}(n_{LC}) = 16 n_{LC} C_{LC}$, where C_{LC} is the capacity of a single linecard. Also, the range of capacities depicted in Fig. 3 exceeds significantly the 640 Gbps capacity considered in Baliga et al. (2009) , corresponding to a 16 line cards single-shelf configuration.

Fig. 3(a) reports the raw power profile of the device, i.e., the total power required by CRS-1 and T1600 at a given capacity is reported with crosses and circles respectively. As expected, the power grows roughly linearly with the capacity. Notice that while differences are small at lower capacities (since Juniper lower cards consumption is offset by the higher chassis consumption), differences in the curves increases at higher capacities (when the number of line cards is larger, hence their contribution more important).

In the figure, we mark with vertical bars a few reference cases: $L_C = L_J$ corresponding to the lowest capacity for both Cisco and Juniper, while P_C and P_J corresponding to the peak capacity configurations (it can be seen that Cisco CRS-1 is able to achieve a higher capacity as it allows the use of multiple switching fabrics).

To allow a fair system comparison (i.e., for an equal amount of work done), we consider a further reference $R_C = P_J$ for Cisco, corresponding to Juniper capacity peak. In this reference case, we can compare both systems on the basis of the amount of power needed to offer the $R_C = P_J$ capacity: as Juniper power curve is lower than the one of Cisco in Fig. 3(a) it is easy to understand that Juniper achieves in this case better performance. This can better be seen in Fig. 3(b), which reports the Energy Consumption Rating (ECR) metrics by normalizing the power drain over the achieved capacity, thus expressing the energy cost for the device to process (i.e., to route) a single bit of information. It can be seen that the ECR metric in the case of Cisco exhibits a non-monotonous behavior. Recall that, every 8th slot needs to be occupied by a switching fabric, which has a higher power rating with respect to a linecard: this yields a spike in the power profile, which (although present) could not be spotted in Fig. 3(a) due to the logarithmic y-axis scale. Notice that in $R_C = P_J$ the maximum energy saving of Juniper T1600 with respect to Cisco CRS-1 amounts to about 4.2 nJ/bit.

Nameplate vs Actual Rating

As it follows from the above analysis, and is confirmed Juniper (2009a) a T1600 configured to have a 640 Gbps throughput has a nameplate rating of 8,352 W. Notice that this is nearly 25% less than the Cisco CRS-1 originally considered by Baliga et al. (2009).

However, power profile not only varies between different routers, or different configurations of the same router. Indeed, the nameplate rating is known to overestimate the theoretical energy consumption based on power ratings of individual components, which is done to reduce the risk of power outage. For instance, Juniper reports in Juniper (2009a) that the aggregated rating for the T1600, based on power ratings of individual components, would amount to 7,008 W, i.e., about 16% less than its nameplate rating.

Furthermore, since power rating of individual components is possibly subjected to the very same conservative overestimation principle, the actual power drain in normal conditions may be even lower. For example, Juniper (2009a) reports results of an experimental study of T1600 power rating under realistic throughput. The study measures the actual power rating of a T1600 sending 64 Byte packets at three different loads, namely 0%, 50% and 100% of the capacity. In the above conditions, the measured power amounts to 5,640 W: thus, the use of nameplate ratings may result in an overestimation of the actual power rating by more than 30%.

Notice that similar observations are also found in independent research. According to Fan et al. (2007), nameplate rating of IT devices is well in excess of the actual running load by a factor of at least 1/3, for more than 40% of the servers. Despite there are no publicly available measurements for Cisco CSR-1 router (e.g., we recall that Chabarek et al., 2008 performs experimental measurement for Cisco GSR 12008 and 7507 access routers), however, in reason of Juniper (2009a) and Fan et al. (2007), it may be argued that similar considerations should hold as well. As such, original projections evaluated through equation (1) may be overestimated by a further 30% due to nameplate rating.

Router Market Share

Involving market surveys as a data source is a topic of contention in the research community. Often, in practice, a direct approach for gathering data is difficult to follow (e.g., due to limited access to facilities, unavailability of public data, restricted access to the experimental data-sheets from the vendors, etc.). In such cases, researchers and analysts are compelled to rely on indirect data sources, such as those provided by market surveys. At the same time, special care must be taken on the correctness of the input: typically indeed, the data is non-verifiable (as it is used when no other data is available), and the data-gathering methodology is seldom unclear (e.g., possibly biased by commercial interests).

While we are aware of the above fallacies, in cases where no direct data source is available, we argue that it may be preferable to exploit data gathered in such surveys, rather than rely on arbitrary assumptions. For example, in order to keep the model simple, Baliga et al. (2009) assume the Internet core router population to be *entirely* made up by Cisco CRS-1 core routers. However, as estimated in Current Analysis (2010), Cisco devices may represent only about 60% of the core network market. Moreover, as reported in the previous section, nameplate ratings suggest Juniper T1600 power to be 25% lower than Cisco CRS-1. Thus, heterogeneity of the router population constitutes a further source of uncertainty in the estimate, and considering a homogeneous Cisco CRS-1 router population may result in a further overestimation of P_{core}.

Impact of Exogenous Factors

Here, we briefly consider all those factors that the model (1) lumps in the γ constant: namely (i) cooling, (ii) redundancy and (iii) future growth. The choice made by Baliga et al. (2009) to fix $\gamma = 2 \times 2 \times 2$ bares indeed additional discussion. Reliable sources for the selection of these parameters are hard to find, as they are deeply tied to ISP planning (e.g., redundancy) and investment (e.g., future growth) strategies. At the same time, an educated guess can assist us in refining the value of ϕ.

For example, claiming a factor of two for redundancy corresponds to the 1-to-1 protection upper-bound, which is unlikely to be applied in practice, while it would be more reasonable to assume a 1-to-N protection policy with $N>1$. If an ISP used a $N=2$ policy (i.e., a backup node/link shared by two nodes/links), considering $N=1$ would result in a 25% overestimation (more generally, considering 1-to-1 protection when 1-to-N is in place, yields to an overestimation factor that grows with N as $1-\frac{N+1}{2N}$).

As far as cooling is concerned, U.S. Environmental Protection Agency (EPA) (2007) state that state-of-the-art data-center power usage efficiency (PUE) could significantly ameliorate by employing energy-efficient power and cooling technologies (such as liquid cooling and combined heat-and-power energy generation solutions): hence, the present PUE of 1.9 could reach in the near future a PUE of 1.2 as suggested by Google (2009).

It follows already that when only minimal conservative corrections are adopted (i.e., improved PUE estimate, 1-to-2 protection, no correction for future growth), this would result in $\gamma'=1.2\times1.75\times2=4.2$, i.e., almost a factor of $2\times$ reduction with respect to the original estimate.

Joint Impact

Finally, we assess the joint impact of the above factors, so to estimate bounds for values of (1) and for reference purposes, we compare the bounds with the values gathered by Baliga et al. (2009) in terms of P_{core} relative error. We reformulate (1) to take into account the manufacturer market share by means of the weight α. Thus, we consider two families of routers, both configured to attain the same aggregated capacity C_{router}. Each family $F\in\{$Cisco, Juniper$\}$ is characterized by a given nameplate power profile P^F_{router}, and we further consider that the actual power drain is lower than the nameplate rating, i.e., $\tilde{P}^F_{router}=\beta P^F_{router}$ with $\beta\in[0,1]$. From our previous observations, it follows that a careful evaluation would consider $(\alpha,\beta)=(0.6,0.7)$. For the sake of simplicity, we fix the minimum per-user access rate to A=4 Mbps as considered by Baliga et al. (2009).

$$\tilde{P}_{core} = (H+1)\frac{A}{C_{router}}\gamma\ \beta\ \left(\alpha P^{CRS-1}_{router} + (1-\alpha)P^{T1600}_{router}\right) \qquad (5)$$

In Fig. 4, we report the model (5) evaluated for:

- two different values for the router power-model: namely, the experimental $\beta=0.7$ and the nameplate $\beta=1$ adopted by Baliga et al. (2009) as a reference;
- three path length H values, corresponding to the lower H=6 and upper H=16 bounds, and to the reference value H=10 adopted in Baliga et al. (2009);
- three different market share values, considering $\alpha=1$ as in Baliga et al. (2009), $\alpha=0.6$ as a more realistic estimate, and $\alpha=0$ as a further reference;
- two values of the lumping constant γ: namely, the full bar height correspond to $\gamma=8$, while horizontal thick line in each bar corresponds to $\gamma'=4.2$.

Lower and upper bounds of the power-per-user statistics $\tilde{P}_{core}(H,\alpha,\beta)$ are gathered for parameter setting of Lo=(6,0.6,0.7) and Hi=(16,1.0,1.0), while the reference case by Baliga et al. (2009) corresponds to Ref=(10,1.0,1.0). Considering for the time being the case $\gamma=8$, from Fig. 4 it is easy to see that, in absolute terms, the estimated power-per-user varies from a minimum of

$\tilde{P}_{core}(Lo)$=2.33 W, to a maximum of $\tilde{P}_{core}(Hi)$=9.26 W, with the reference by Baliga et al. (2009) $\tilde{P}_{core}(Ref)$=6.00 W closer to the latter. In relative terms, it is important to notice that the reference *Ref* value exceeds by about 2.5 times the more realistic lower bound, while the distance from the upper bound is only about 50%. In case the lumping factor γ'=4.2 is taken into account on top of (H,α,β), the ratio between the reference and the lower bound grows to almost a factor of 5×.

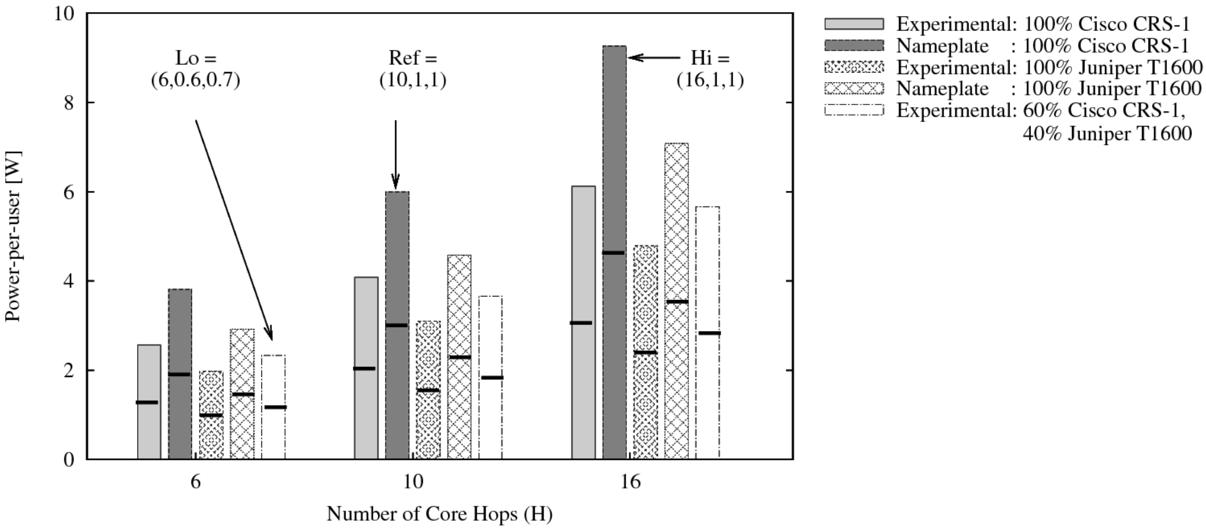

Figure 4: Sensitivity analysis of the power-per-user in the Internet core. Bars correspond to different parameters (H,α,β) settings; the full bar height further correspond to the original setting γ=8 of the lumping constant, while the thick line reports the estimate for γ'=4.2.

We again point out that our aim is not to disprove, confute or weaken the importance of Baliga et al. (2009), as our sensitivity analysis would not have been possible without their effort. Rather, the results we gather seem to suggest that further work is still needed to gather reliable, and verifiable, power figures of the Internet.

SOLUTIONS AND RECOMMENDATIONS

Overall, our results suggest that power profiling the Internet remains a challenging task: the original estimation of Baliga et al. (2009) is not easy to refine and, despite the care one can put in its evaluation, results have a deceiving precision. Indeed, when we compare the gap between the *Hi* and *Lo* boundaries remain very large: to reduce this uncertainty, further effort is needed in terms of methodological approach, data gathering, and comparative power modeling.

Solution to this problem however requires *a coordinated and community-wide effort*: we believe the creation and maintenance of a *green repository* of energy-related figures would be a very helpful step to foster future research. Clearly, the repository would not only be useful for the specific problem considered in this work (i.e., Internet power profiling), but would also serve as valuable input for any class of green research work. For instance, data stored in the repository may include specific power models for different devices (e.g., router, switches, set-top-boxes, etc.) as well as any data not directly related, but still relevant for green benchmarking (e.g., the IP hop count metric or market share information considered in this work, or a set of standard workloads and traces to perform Adaptive Link Rate studies, or sets of topologies and traffic matrices for energy-aware routing performance evaluation, etc.).

Indeed, there is a wealth of information already available in the literature: at the same time, the fact that figures are widely distributed across many white papers and surveys, makes it difficult to both *gain access* to the data itself, and to *cross-check* it as well. Conversely, the process of building a centralized repository would also expose data to frequent cross-checking, thereby increasing its consistency. Hopefully, the repository could become a trusted, up-to-date and authoritative centerpiece of the green networking community: in this case, the use of a common evaluation framework consisting of a shared set of standard benchmarks and metrics (as the one overviewed in Bianzino et al.,. 2010), input data, and good practices would further promote the *cross-comparison* among different green-research studies as a beneficial side effect.

Such a repository would, e.g., allow us to further refine and elaborate the investigation performed in this work. Despite the effort we put in this sensitivity analysis, the estimation of some of parameters still have a degree of uncertainty. For example, while the assumption of the reduction factor β=0.7 with respect to the nameplate rating is based on multiple sources (Juniper, 2009a and Fan et al., 2007), nevertheless further independent assessments of β could further validate (or disprove) this assumption. For instance, it may happen that manufacturers will be even more conservative in their data sheets so that future values of β can be expected to be $\beta<0.7$. Similar consideration applies to all other parameters we considered in the sensitivity analysis.

Per definition, *independent validation* requires multiple research laboratories, universities, institutions, etc. to be involved in the process - so that collaborative work performed by the scientific community can help counter and pinpoint the occurrence of individual errors in green networking research. We hope that, equipped with powerful community tools such as a “green repository”, future work can further refine the sensitivity analysis carried out in this chapter.

CONCLUSIONS

In this chapter, we consider the challenging issue of power profiling the Internet: by taking the model proposed in Baliga et al. (2009) as a starting point, we focus on the consumption of the core segment, and perform a sensitivity analysis of the parameters affecting the model (such as path length, router power profile, market share, cooling, redundancy, etc.).

Two main conclusions can be derived from the analysis. First, it seems that although the original aim of Baliga et al. (2009) is to provide a lower-bound of Internet power profile, this may not have been entirely achieved: under this light, this chapter can be interpreted as a cautionary tale on the perils of placing too much emphasis on specific figures of Internet-wide power rating.

In fact, the most important and troublesome finding is that, even considering a very simple model of a specific part of the whole Internet, gathering reliable and accurate figures is extremely challenging: indeed, no matter how carefully the model data is selected and cross-verified, the boundaries within which the actual figures may lay remain unreasonably wide. This also means that, so far, the scientific community is unable to verify the gain brought at planetary scale by green research efforts, or to reliably quantify the energy consumed by each network segment (useful to pinpoint, over time, which network requires further research effort).

As such, we advocate the necessity for a coordinated, community-wide effort in order to, if not solve the issue, at least greatly improve the current situation by the use of commonly agreed assumptions, practices, input data, metrics and methods.

ACKNOWLEDGEMENTS

This work was funded by Eureka Celtic project TIGER2.

REFERENCES

Baliga, J., Ayre, R., Hinton, K., Sorin, W. V., and Tucker, R. S. (2009). Energy consumption in optical ip networks. *Journal of Lightwave Technology*, 27(13):2391–2403.

Barroso, L. and Hölzle, U. (2007). The Case for Energy-Proportional Computing. *IEEE Computer*, 40(12):33 – 37.

Bianzino, A. P., C.Chaudet, Rougier, J., and Rossi, D. (2011). A survey of green networking research. *IEEE Communications Surveys & Tutorials, to appear*.

Bianzino, A. P., Raju, A., and Rossi, D. (2010). Apple-to-apple: A common framework for energy-efficiency in networks. In *ACM SIGMETRICS, GreenMetrics workshop*, New York, USA.

Bolla, R., Bruschi, R., Christensen, K., Cucchietti, F., Davoli, F., , and Singh, S. (2011). The Potential Impact of Green Technologies in Next Generation Wireline Networks – Is There Room for Energy Savings Optimization? *IEEE Communication Magazine*.

Chabarek, J., Sommers, J., Barford, P., Estan, C., Tsiang, D., and Wright, S. (2008). Power Awareness in Network Design and Routing. In *IEEE INFOCOM 2008*, pages 457–465, Phoenix, USA.

Christensen, K., Reviriego, P., Nordman, B., Bennett, M., Mostowfi, M., and Maestro, J. A. (2010). IEEE 802.3az: The Road to Energy Efficient Ethernet. *IEEE Communication Magazine*, 48(11):50–56.

Cisco (2004). CRS-1 Modular Services Card. *http://www.cisco.com/en/US/prod/collateral/routers/ps5763/ps5862/product_data_sheet09186a008022d5ee.pdf*.

Cisco (2011). CSR-1 Carrier Routing System. *http://www.cisco.com/en/US/prod/collateral/routers/ps5763/prod_brochure0900aecd800f8118.pdf*.

Cisco (February 2007). CRS-1 Carrier Routing System Gigabit Ethernet Physical Layer Interface Module Installation Note. *http://www.cisco.com/en/US/docs/routers/crs/crs1/plim/installation/guide/eth6437.pdf*

Current Analysis (2010). Product comparison: Carrier infrastructure / core ip routing. *http://www.currentanalysis.com/teb/samples/ci-coreiproutcompare.asp*.

Donnet, B., Raould, P., and Friedman, T. (2006). Efficient Route Tracing from a Single Source. *Technical Report* arXiv:cs/0605133v1.

ECR Initiative (2009). Network and Telecom Equipment Energy and Performance Assessment. *http://www.ecrinitiative.org/pdfs/ECR_1_0_4.pdf.*

Fan, X., Weber, W.-D., and Barroso, L. A. (2007). Power provisioning for a warehouse-sized computer. In *34th International Symposium on Computer Architecture (ISCA'07)*, pages 13–23, San Diego, USA.

Fei, A., Pei, G., Liu, R., and Zhang, L. (1998). Measurements On Delay And Hop-Count Of The Internet. In *IEEE GLOBECOM'98 Mini-Conference*, Sydney, Australia.

Google (2009). Data center efficiency measurements. *http://www.google.com/corporate/datacenter/efficiency-measurements.html*

Gupta, M. and Singh, S. (2003). Greening of the Internet. In *ACM SIGCOMM'03*, pages 19 – 26, New York, USA.

Haddadi, H., Fay, D., Uhlig, S., Moore, A., Mortier, R., and Jamakovic, A. (2010). Mixing Biases: Structural Changes in the AS Topology Evolution. In *Passive and Active Measurements (PAM'10)*, Zurich, Switzerland.

Huber, P. and Mills, M. (1999). Dig more coal, the PCs are Coming. *Forbes Magazine*, 31.

Juniper (2009a). Energy Efficiency for the Network Equipment: Two Steps Beyond Greenwashing. *http://www.juniper.net/us/en/local/pdf/whitepapers/2000284-en.pdf.*

Juniper (2009b). Product data sheet. Juniper T Series Core Routers: T320, T640, T1600, TX Matrix, and TX Matrix Plus. *http://www.juniper.net/us/en/local/pdf/datasheets/1000051-en.pdf.*

Juniper (2009c). T1600 Routing Node PIC Guide. *http://www.juniper.net/techpubs/en_US/release-independent/junos/information-products/topic-collections/hardGreenwaware/t-series/t1600/pics/t1600-pic.pdf*

Katz, R. (2009). Tech titans building boom. *Spectrum, IEEE*, 46(2):40–54.

Koomey, J. G., Calwell, C., Thornton, J., Brown, R. E., Eto, J. H., Webber, C., Cullicott, C., and Laitner, S. (2002). Sorry, wrong number: The Use and Misuse of Numerical Facts in Analysis and Media Reporting of Energy Issues. *Annual Review of Energy and the Environment*, 27:119–158.

Lange, C. (2009). Energy-related aspects in backbone networks. In *European Conference and Exhibition on Optical Communication (ECOC'09)*, Wien, AU.

Lebeck, A. R., Fan, X., Zeng, H., and Ellis, C. (2000). Power Aware Page Allocation. *ACM SIGOPS Operating Systems Review*, 34(5):105–116.

LessWatts Project (2009). http://www.lesswatts.org/projects/tickless/.

Mandviwalla, M. and Tzeng, N.-F. (2006). Energy-Efficient Scheme for Multiprocessor-Based Router Linecards. In *International Symposium on Applications on Internet (SAINT'06)*, pages 156–163, Washington, DC, USA. IEEE Computer Society.

Qureshi, A., Weber, R., Balakrishnan, H., Guttag, J., and Maggs, B. (2009). Cutting the Electric Bill for Internet-Scale Systems. In *ACM SIGCOMM'09*, Barcelona, Spain.

Roth, K. W., Goldstein, F., and Kleinman, J. (2002). Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings Volume I: Energy Consumption Baseline. Technical Report Reference No. 72895-00, Arthur D. Little, Inc.

Schulzrinne, H. (2010). Long-term traffic statistics. *http://www.cs.columbia.edu/hgs/Internet/traffic.html*

U.S. Environmental Protection Agency (EPA) (2007). Report to congress on server and data center energy efficiency.

http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf.

Valancius, V., Laoutaris, N., Massoulié, L., Diot, C., and Rodriguez, P. (2009). Greening the Internet with Nano Data Centers. In *ACM CoNEXT'09*, Rome, Italy.

Van Mieghem, P. (2006). *Performance analysis of communications networks and systems*. Cambridge Univ Pr.

Van Mieghem, P. and Begtašević, F. (2001). Measurements of the Hopcount in Internet. In *In Passive and Active Measurement (PAM'01)*, pages 23–24, Amsterdam, Nederland.

Weiser, M., Welch, B., Demers, A., and Shenker, S. (1994). Scheduling for Reduced CPU Energy. In *Proceedings of the 1st USENIX Conference on Operating Systems Design and Implementation (OSDI 1994)*, Monterey, California.

Wissner-Gross, A. (2009). 'Carbon cost' of Google revealed. BBC News, *http://news.bbc.co.uk/2/hi/7823387.stm*.

Zhai, B., Blaauw, D., Sylvester, D., and Flautner, K. (2004). Theoretical and Practical Limits of Dynamic Voltage Scaling. In *Proceedings of the 41st Annual ACM Design Automation Conference (DAC 2004)*, pages 868–873, San Diego, California, USA.

ADDITIONAL READING SECTION

Bianzino, A.P., Chaudet, F., Rossi, D., Rougier, J.L. and Moretti, S., (2011) The Green-Game: Striking a Balance between QoS and Energy Saving, In *22nd International Teletraffic Congress (ITC),* September, 2011.

Bianzino, A.P., Chaudet, F., Rossi, D. and Rougier, J.L., (2010) Energy-Awareness in Network Dimensioning: a Fixed Charge Network Flow Formulation. *In ACM SIGCOMM eEnergy'10,* Extended Abstract, Passau, Germany, April 13 - 15 2010.

Bianzino, A.P., Chaudet, C., Larroca, F., Rossi, D. and Rougier, J.L., (2010) Energy-Aware Routing: a Reality Check . *In IEEE Globecom Workshop on Green Communications (GreenComm3),* Miami, FL, USA, December 6-10 2010.

Bo, Y., Binqiang, W., Zhigang, S. and Yi, D., (2011) A Green Parallel Forwarding and Switching Architecture for Green Network, *IEEE/ACM International Conference on Green Computing and Communications (GreenCom'11)*. Chengdu, China, August 4-5, 2011

Bolla, R., Bruschi, R., Carrega, A. and Davoli, F., (2011) Green Network Technologies and the Art of Trading-off, In *IEEE INFOCOM Workshop on Green Communications and Networking,* April 15, 2011

Chen D., Henis, E., Kat, R.R., Sotnikov, D., Cappiello, C., Mello Ferreira, A., Pernici, B., Vitali, M., Jiang, T., Liu, J., and Kipp, A. (2011) Usage Centric Green Performance Indicators, *In ACM SIGMETRIC Workshop on Green Networking (GreenMetrics)* San Jose, CA, USA

Cianfrani, A., Listanti, M., Eramo, V. and Polverini, M., (2011) An OSPF Enhancement for energy saving in IP Network, *In IEEE INFOCOM Workshop on Green Communications and Networking,* April 15, 2011

Feng, M. Z., Hinton, K., Ayre, R. and Tucker, R. S. (2010) Reducing NGN Energy Consumption with IP/SDH/WDM, *In ACM SIGCOMM eEnergy,* Passau, Germany

Fisher, W., Suchara, M., Rexford, J. (2011) Greening Backbone Networks: Reducing Energy Consumption by Shutting Off Cables in Bundled Links, *In ACM SIGCOMM Workshop on Green Networking,* New Delhi, India, August 2011

Hylick, A., Sohan, R., Rice, A., and Jones, B. (2008). An Analysis of Hard Drive Energy Consumption. In *Proceedings of the IEEE International Symposium on Modeling, Analysis and Simulation of Computers and Telecommunication Systems, 2008. (MASCOTS 2008)*, pages 1–10, Baltimore, Maryland, USA.

Garroppo, R.G., Giordano, S., Nencioni, G., Pagano, M., (2011) Energy Aware Routing based on Energy Characterization of Devices: Solutions and Analysis, *In 4th International Workshop on Green Communications (GreenComm4),* Kyoto, Japan, 9 June 2011

Kansal, A. and Zhao, F. (2008). Fine-grained Energy Profiling for Power-aware Application Design. *ACM SIGMETRICS Performance Evaluation Review*, 36(2):26–31.

Lebeck, A. R., Fan, X., Zeng, H., and Ellis, C. (2000). Power Aware Page Allocation. *ACM SIGOPS Operating Systems Review*, 34(5):105–116.

Lewis, A., Ghosh, S., and Tzeng, N.-F. (2008). Run-time Energy Consumption Estimation Based on Workload in Server Systems. In *Proceedings of the USENIX Workshop on Power Aware Computing and Systems (HotPower), held at the Symposium on Operating Systems Design and Implementation (OSDI 2008)*, San Diego, California, USA.

Mahadevan, P., Sharma, P., Banerjee, S., and Ranganathan, P. (2009). A Power Benchmarking Framework for Network Devices. In *IFIP Networking*, Aachen, Germany.

Poess, M., Nambiar, R. O., Vaid, K., Stephens, J. M., Huppler, K. and Haines, E. (2010) Energy Benchmarks: A Detailed Analysis, *In ACM SIGCOMM eEnergy,* Passau, Germany

Raju, A., Bianzino, A.P., and Rossi, D., (2011) Greening the Internet: Measuring Web Power Consumption. *IEEE ITProfessional*, January 2011.Rivoire, S., Ranganathan, P., and Kozyrakis, C. (2008). A Comparison of High-Level Full-System Power Models. In *Proceedings of the USENIX Workshop on Power Aware Computing and Systems (HotPower), held at the Symposium on Operating Systems Design and Implementation (OSDI 2008)*, San Diego, California, USA.

Rivoire, S., Shah, M. A., Ranganatban, P., Kozyrakis, C., and Meza, J. (2007a). Models and Metrics to Enable Energy-Efficiency Optimizations. *IEEE Computer*, 40(12):39–48.

Rivoire, S., Shah, M. A., Ranganathan, P., and Kozyrakis, C. (2007b). JouleSort: A Balanced Energy-Efficiency Benchmark. In *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data (SIGMOD)*, pages 365–376, Beijing, China.

Rossi, D., (2010) Energy-Aware Networks: a Tutorial on Green Networking . *In IEEE Networks,* Warsow, September 2010. http://www.enst.fr/~drossi/paper/rossi10networks.pdf

Sivaraman, V., Vishwanath, A., Zhao and Z., Russell, C.L., (2011) Profiling Per-Packet and Per-Byte Energy Consumption in the NetFPGA Gigabit Router, *In IEEE INFOCOM Workshop on Green Communications and Networking,* April 15, 2011

Tzanakaki, A. et al. (2011) Energy Efficiency in Integrated IT and Optical Network Infrastructures: The GEYSERS Approach. *In IEEE INFOCOM Workshop on Green Communications and Networking,* April 15, 2011

Vahdat, A., Lebeck, A., and Ellis, C. S. (2000). Every Joule is Precious: the Case for Revisiting Operating System Design for Energy Efficiency. In *Proceedings of the 9th ACM SIGOPS European Workshop*, pages 31–36, Kolding, Denmark.

Vasic, N. and Kostic, D., (2010) Energy-Aware Traffic Engineering, *In ACM SIGCOMM eEnergy,* Passau, Germany, April 2010

Wang, L., Lu, R., Li, Q., Zheng X. and Zhang H., (2011) Energy Efficient Design for Multi-shelf IP over WDM Networks. *In IEEE INFOCOM Workshop on Green Communications and Networking,* April 15, 2011