Apple-to-Apple: A Framework Analysis for Energy-Efficiency in Networks

Aruna Prem Bianzino, Anand Kishore Raju, Dario Rossi
Institut TELECOM, TELECOM ParisTech, CNRS LTCI UMR 5141, Paris, France
bianzino, raju, drossi@telecom-paristech.fr

ABSTRACT

Research on energy-efficiency in and through communication networks has already gained the attention of a broad research community. Specifically, we consider efforts towards improving environmental sustainability by making networks energy-aware. One of the first, and perhaps most important step in this direction is establishing a comprehensive methodology for measuring and reporting the energy consumption of the network. In this work, we compare and contrast various energy-related metrics used in the recent literature, by means of a taxonomy definition, as well as through relevant case studies. We believe this to be a first necessary step towards the definition of a common framework for the performance evaluation of energy-aware networks.

1. INTRODUCTION

Information and Communication Technologies (ICT) offer opportunity to rationalize the energy-usage and reduce the corresponding carbon emissions for non-ICT process, e.g., by reducing physical travel and enhancing on-line connectivity. At the same time, while ICTs undoubtedly represent an important cornerstone for the development of a carbon-lean economy, there are possible issues that can hamper this process. Indeed, this potential will meet a key road block in terms of a rapid and widespread deployments in ICT (in the form of networking equipments and data-centers).These deployments in turn translate into an increased energy consumption, and in the consequent escalation of carbon emissions. As such, in recent years, research effort devoted towards embedding energy-awareness in ICTs has proliferated, and is usually referred to as "greening" of the network.

A number of proposals exists which aims at offering Telecom Operators and Internet Service Providers (ISPs) the ability to optimize their energy-efficiency, hence regulating both Total Cost of Ownership (TCO) and the related carbon footprint. At the same time, we point out that the effectiveness of this optimization process is strongly dependent on the accuracy of both the modeling and the input data (e.g., power profile of individual devices). Till now, the green networking community lacks these very definitions of effectiveness: indeed, the current literature proposes many heterogeneous metrics to qualify and quantify the energy savings [4–7, 10, 11, 13, 15–18, 21–23]. Due to this heterogeneity, comparison of competing solutions may be more favorable to one or to another approach, depending on the metrics used to express the results of the benchmark. Hence, it becomes fundamental to define a coherent framework for the evaluation of networks, which should be able to fully characterize any possible trade-off between energy consumption and system performance.

Our work should be considered as a first step in this direction. It includes overview and comparison of different metrics that have been proposed in the literature so far. We exemplify the use of these metrics considering both (i) a simple toy case, so as to give a litmus test of their usage, as well as (ii) a realistic in-depth case study, where we use these metrics to compare two routing devices (namely, Cisco CRS-1 and Juniper T1600).

2. A TAXONOMY OF GREEN NETWORKING METRICS

A comprehensive evaluation of energy savings and system performance in a real-world networking scenario is a challenging task, as many technological and non-technological aspects must be taken into account. Diverse issues like fault tolerance and transmission technologies (e.g., Ethernet, optical, wireless), can be difficult to homogeneously quantify, which complicates the task of developing such a common framework. Non-technical issues like verifiability, accessibility, and representivity of a framework makes the process even more complex.

By a common framework, we imply first of all the use of coherent and widely accepted *metrics* for the performance evaluation of an energy-aware network. More precisely, a framework can be defined as a performance indicator which must be: *simple* enough to be understood; *accurate* enough to withstand scrutiny; and both *usable* and *relevant* enough to be an effective agent of change. As research in green networking has proliferated, so has the number of proposals and studies, resulting in a profusion of heterogeneous energy-related metrics. As such, efforts from research community are needed to improve the current state-of-the-art evaluation practices in green networking research. In other

Table 1: Summary and Taxonomy of Energy-Related Metrics

Country-Level	Uses	Description
EPI [4]	Rankings	Environmental Performance Index, benchmarks the country's environmental performance.
ESI [5]	Rankings	Environmental Sustainability Index, indexes the country's environmental sustainability.
EVI [6]	Indicators	Environmental Vulnerability Index, characterizes various environmental issues.
HPI [10]	Rankings	Happy Planet Index, reveals the ecological efficiency of human well-being.

Corporate-Level	Established as	Description
ISO/TC207 [11]	Workgroups	Standardization in the field of environmental management tools and systems.
GHG Protocol [7]	Standards	GHG Protocol Corporate Standard provides standards and guidance for companies and other organizations preparing a GHG emissions inventory.

Facility-Level	Range	Description
PUE [20]	1 to ∞	Power Usage Effectiveness, is defined as the ratio of data center power to IT power draw.
DCiE [20]	0 to 100%	Data-center infrastructure Efficiency, is a reciprocal of PUE times 100.
W/ft^2 [25]	0 to ∞	Over-estimates the area actually devoted to the servers.
DCP [20]	1 to ∞	Data-center Productivity, accounts only for the "useful work done" by the data-center.

Equipment-Level	Units	Description
ECR [13, 21]	Watt/Gbps	$ECR = E_f/T_f$. Aggregated energy consumption normalized to capacity.
EER [21]	Gbps/Watt	$EER = 1/ECR$. For convenience, defined under the name of EER.
ECRW [13]	Watt/Gbps	$ECRW = (\alpha E_f + \beta E_h + \gamma E_i)/T_f$, with (α,β,γ)=(0.35,0.4,0.25). A weighted metric considering energy consumption at different loads.
EPI [23]	Percentage	$EPI = 100(M-I)/M$. Based on consumption at idle (I) and maximum (M) workload.
Power Per User [18]	Watt/user	Power consumed by each subscriber in the public Internet.
TEER [16]	Gbps/Watt	Telecommunications Energy Efficiency Ratio. Ratio of "useful work" to power consumption.
TEEER [17]	$-\log \frac{watt}{Gbps}$	Telecommunications Equipment Energy Efficiency Rating. Log of the power/throughput ratio.
CCR [22]	rad (Dimensionless)	Consumer level energy metrics, $CCR = E/\Sigma A(j)$, where E is the power rating of a consumer network device, A is the energy allowance per function (DSL, Wi-Fi, etc).
Watts per circuit [15]	Watt	Used for point-to-point (Virtual Leased Line) Ethernet-Line services.
Watts per MAC [15]	Watt	Used for multipoint (MAC Address) Ethernet-LAN services.

words, a common ground is needed for defining a coherent set of energy-efficiency metrics, along with establishing a community wide consensus over their usage and applicability.

When considering the energy metrics present in the literature, we can broadly classify them into four categories, on the basis of their respective domain of application. The resulting categories are: (i) *Equipment-level metrics*, which account for the lowest level, by evaluating the energy-efficiency rating of an individual piece of ICT equipment; (ii) *Facility-level metrics* relates instead to higher-level systems, where equipment is interconnected (e.g., data-centers, Internet exchange points, ISP networks, etc.); (iii) *Corporate-level metrics*, mostly obscure at a user level, represent the framework where corporations abide by and implement programs in order to meet their social responsibility as good corporate citizens; and (iv) *Country-level metrics*, which evaluate and benchmark the relative environmental sustainability of a country as a whole, of which ICT represent a non-negligible part. The resulting taxonomy, together with the metrics definition, is reported in Tab. 1.

The extent of Tab. 1 well illustrates the aforementioned need of a common evaluation framework. Moreover, this need can be felt even when focusing on specific levels: e.g., in the case of data-center research, which is a relatively mature research field, we see that although several metrics have already been proposed (Facility-level metrics in Tab. 1) nevertheless only some of them are widely used. The risk is that, without a coherent community-wide effort, the situation could further degenerate into an inorganic set of unpopular, heterogeneous and non-comparable metrics. On the other hand, having a clearly agreed and widely accepted set of energy metrics can be beneficial for both the research community (e.g., to promote cross comparison studies) and the industry as well (e.g., to guide manufacturers and service providers in making informed decisions regarding infrastructure deployments and purchases).

Due to space constraints, in this work we focus on equipment-level metrics, which have recently enjoyed a consistent proliferation across the literature: a cursory viewpoint might only consider proposing a new energy metric, while ignoring its comprehensiveness, compatibility and redundancy with other prevailing metrics. Let us consider first a number of technology-agnostic metrics such as ECR [13,21], EER [21], TEER [16] and TEEER [17] (which roughly corresponds to the DCP facility-level metric): these metrics aim at weighting the "energy expenditure" on the ground of the "work". At the same time, the definition of "work" is not unique and can instead be rather flexible (e.g., CPU operations, capacity, throughput, etc.), while this can have some important consequences (e.g., intuitively, expressing the work in terms of capacity (C) or throughput (Th) can have a strong bias on the actual results).

Moreover, even whether a common definition of work would be agreed, we notice that simple relationships exist among the different metrics: as further exacerbated in Tab. 2, these metrics provide the same information, expressed in a slightly different quantitative fashion. The question still remains whether this redundancy is actually needed (e.g., as a logarithmic scale may be more readable in the case of a metric spanning several orders of magnitude) or whether it may instead represent a source of confusion.

On the other hand, the above stated range of metrics (ECR,

Table 2: Apple-to-apple metric comparison

	TEER	TEEER	ECR	EER
TEER	–	e^{TEEER}	$\frac{Th}{C \cdot ECR}$	$\frac{Th \cdot EER}{C}$
TEEER	$\log TEER$	–	$\log \frac{Th}{C \cdot ECR}$	$\log \frac{Th \cdot EER}{C}$
ECR	$\frac{Th}{C \cdot TEER}$	$\frac{Th}{C} e^{-TEEER}$	–	$\frac{1}{EER}$
EER	$\frac{C \cdot TEER}{Th}$	$\frac{C}{Th} e^{TEEER}$	$\frac{1}{ECR}$	–

EER, TEER and TEEER) is unable to capture some relevant properties of the system, such as the degree of proportionality between the energy consumption and the level of utilization [19]. In particular, EPI [23] evaluates energy proportionality of a device or a system, on the basis of energy consumption at idle and maximum workloads respectively. Similarly, ECRW [13] can be considered as a *weighted* extension of the ECR metric: it quantifies the energy consumption of a device by taking different workloads into account (at full, half and idle load). It is to be noted that, using weights, the notion of proportionality can be incorporated (to some extent) into other metrics as well. Nevertheless, this solution raises further degrees of freedom concerning the choice of the weights (α, β, γ) in ECRW (and contrasting ECRW values gathered with different weights does not result in an apple-to-apple comparison).

The toy-case depicted in Fig. 1 is helpful to better understand the differences among a few relevant metrics (namely, EPI, ECR, ECRW). We consider three different power-profiles of a device: a *fully proportional* model (in which the energy consumption is directly proportional to the utilization), an *energy-agnostic* model (in which the power drain is constant and irrespective of the utilization level) and a *realistic* model (which includes both fixed and proportional consumption components). For sake of simplicity, both maximum energy consumption and capacity are taken unitary (i.e., 1 Watt and 1 bps respectively, which results in highly unrealistic *absolute* values, that can however be *relatively* compared). It is here evident how different metrics highlight different aspects of a power consumption profile: EPI considers only the slope, ECR only the efficiency at maximum load, whereas ECRW weights both factors by sampling power efficiency at different loads (notice that given $(\alpha, \beta, \gamma) = (0.35, 0.4, 0.25)$, a $ECRW = 0.55\,ECR$ denotes a fully proportional system while $ECRW = ECR$ denotes an energy-agnostic system).

Finally, the technology-dependence of metrics should be considered in the ICT context. Indeed, metrics that depend on the type of services offered, like power-per-{subscriber, port, VLL, circuit, MAC address} [15, 18], all express relevant information from an ISP perspective. It is also fairly difficult to coalesce all the above metrics: e.g., while power-per-subscriber and power-per-port fits well the case of a single-purpose PSTN network, ISPs providing triple-play services would find these conventional energy metrics insufficient for accurate evaluations. As a consequence, for telecom standards, the area of sustainability metrics has become a new battlefield where evaluation metrics are highly volatile mainly due to rapidly evolving services offered by the ISPs.

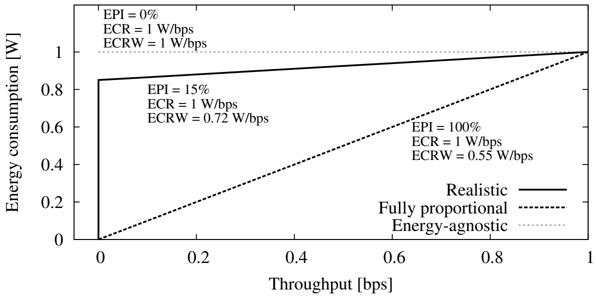

Figure 1: Litmus test of EPI, ECR, ECRW energy-related metrics over a toy-case.

3. CASE STUDY: PROFILING ROUTER ENERGY CONSUMPTION

In this section, as a realistic case study, we address a detailed power-profile comparison of widely used network equipment. Specifically, using the previously defined metrics, we characterize the power consumption for different configurations of the Cisco CRS-1 and Juniper T1600 core routers.

Before quantifying the metrics, it is necessary to consider the architectural and operational details of the devices. Notice that, in what follows, we focus on the configuration of the device, excluding external variants like cooling, power-redundancy and console systems, etc. Generally speaking, systems can be configured in *Single-Chassis* (as the Juniper T1600) and *Multi-Chassis* (as the Cisco CRS-1) 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 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.

Tab.3 summarizes the power consumption of the individual sub-components of the CRS-1 and T1600 systems. Notice that we report nameplate ratings from vendors, which are typically overestimated (i.e., well above the actual power consumption values) thereby ensuring a safe system operation by reducing the odds that power-breakdowns occur. CRS-1 uses OC-768c/STM-256c linecards [2], which can support a 40 Gbps throughput for a power consumption of 500 W. Conversely, a Juniper T1600 equipped with a *Tx Matrix Plus* [12] switching fabric can interconnect 16 T1600 chassis into a single routing entity. T1600 uses OC-768c/STM-256c PIC linecards [14], which also provide a capacity of 40 Gbps for a power consumption of 66 W.

3.1 Profiling the Power Consumption

Total power consumption P_{total} of a Multi-Chassis system can be calculated by summing the power consumption of each component: namely, the power consumption $P_{chassis}$ of

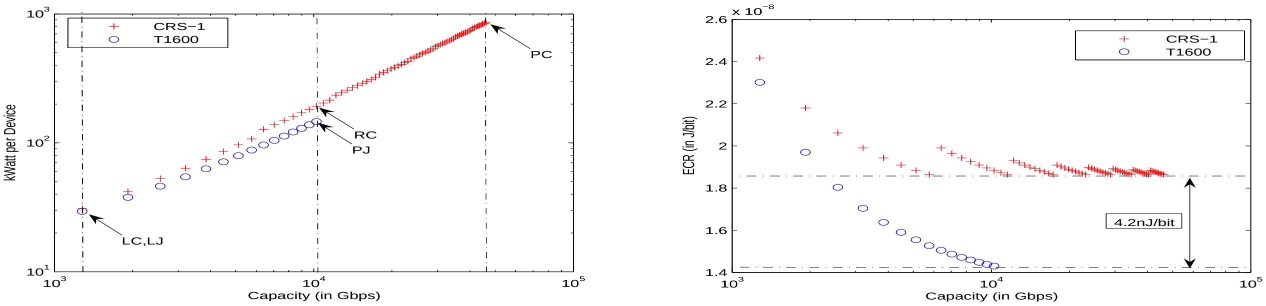

(a) Total power Vs Capacity

(b) Energy per bit Vs Capacity

Figure 2: Comparison of Cisco CRS-1 and Juniper T1600: Watts-per-device and ECR metrics as a function of the router aggregated capacity.

Table 3: Power footprint of individual components

Equipment	Functionality	Power (kW)
OC-768c/STM-256c	Linecard (CRS-1)	0.5 [2,24]
OC-768c/STM-256c PIC	Linecard (T1600)	0.066 [14]
Switching Fabric	SF (CRS-1)	9.1 [1]
Tx Matrix Plus	SF (T1600)	12.75 [12]

an empty chassis (i.e., without active linecards), plus the power consumption P_{LC} of an active linecard installed in a linecard shelf, plus the power consumption P_{SF} of the switching fabric used to interconnect the Multi-Chassis system.

From Tab. 3, we find the power consumption of the OC-768c/STM-256c linecard ($P_{LC} = 0.5$ kW) and the switching fabric ($P_{SF} = 9.1$ kW) for a CRS-1. However, the power rating of the chassis $P_{chassis}$ is not publicly available. To derive $P_{chassis}$ we consider a fully equipped Cisco CRS-1 with a single chassis and 16 linecard chassis which is reported in [1,18] to have $P_{total} = 10.9$ kW. Ripping off 16 active linecards (each consuming 500 Watts), gives us a conservative upper bound of $P_{chassis} = 2.92$ kW for the empty CRS-1 chassis.

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

$$P_{total} = iP_{chassis} + 16iP_{LC} + \left\lceil \frac{i-1}{8} \right\rceil P_{SF} \qquad (1)$$

where $i \in [2,72] \subset N$ corresponds to the number of linecard chassis installed in the CRS-1 Multi-Chassis System. Notice that our model assumes that once a linecard chassis is installed, it is fully utilized (i.e., it consumes $P_{chassis}$, plus iP_{LC} for each of the 16 cards it supports). Also notice that a variable number $\lceil \frac{i-1}{8} \rceil$ of SF chassis are needed to support i linecard chassis (more precisely, any eight slot is occupied by a SF element needed for the interconnection).

Using a similar profiling technique, we can derive the generic power model for the Juniper T1600 Multi-chassis System. From Tab. 3, we get $P_{linecard} = 0.066$ kW and $P_{SF} = 12.75$ kW for the Juniper T1600. Interestingly, notice that while the Juniper SF is much more power-hungry than the Cisco's one, the opposite happens concerning linecards (which is due to Short Reach interfaces, which consume much less power): overall, 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 [3] reports to have $P_{total} = 8.35$ kW. Finally, we have:

$$P_{total}(i) = iP_{chassis} + 16iP_{LC} + \mathbb{I}_{i>0}P_{SF} \qquad (2)$$

where $i \in [2,16] \subset N$ corresponds to the number of linecard chassis installed in the T1600: notice that, unlike in the CRS-1 case, 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., $\mathbb{I}_{i>0}$).

The resulting power profiles, for both Cisco CRS-1 and Juniper T1600, are shown in Fig. 2(a) as a function of the aggregated system capacity achieved under different configurations. Notice that in both cases the system capacity can be expressed as $C_{total}(i) = 16iC_{LC}$, where C_{LC} is the capacity of a single linecard. Fig. 2(a) reports the raw power consumption of the devices (i.e., Watt-per-device metric in Tab. 1); the total power required by CRS-1 and T1600 at a given capacity is reported with crosses and circles respectively. As expected, the power consumption grows roughly linearly with the capacity.

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: in the case in which *capacity* (i.e., the absolute amount of work done) is the primary metric to compare the above systems, Cisco is the clear winner. 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: it is easy to understand that Juniper achieves better performance. This can better be seen in Fig. 2(b), which reports the ECR metrics by normalizing the power consumption over the achieved capacity, thus expressing the energy cost for the device to process (i.e., route) a single bit of information. It can be seen that the ECR metric in the

Table 4: Comparison summary of T1600 vs CRS-1

Metric	Juniper T1600		Cisco CRS-1	
	L_J	P_J	L_C	R_C
TEER [*Gbps/Watt*]	0.04	0.07	0.04	0.05
TEEER $[-\log \frac{Watt}{bps}]$	7.63	7.84	7.61	7.72
ECR [*Watt/Gbps*]	23.3	14.3	24.4	18.9
EER [*Gbps/Watt*]	0.04	0.07	0.04	0.05

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 consumption with respect to a linecard: this yields a spike in the power consumption, which (although present) could not be spotted in Fig. 2(a) due to the logarithmic y-axis scale. Notice that in $R_C = P_J$ the energy saving of Juniper T1600 with respect to Cisco CRS-1 amounts to about 4.2 nJ/bit. Cisco Visual Index [8] estimates a yearly Internet traffic of about 885.29 Exabits. Considering this data traverses at least one core router, the annual energy saving can be as high as 1.03 TWh, which corresponds to approximately 700,000 Metric Tons of Carbon Dioxide Equivalent [9].

Finally, for the sake of completeness and to provide a litmus test for this case study as well, we quantify the metrics presented in Tab. 2. The resulting figures are reported in Tab. 4. It is to be noted that computation of TEER and TEEER assumes maximum device utilization. It can be seen that EER and TEER provide valuable but redundant information; moreover, ECR conveys the same quantitative information with perhaps a more intuitive unit of measure. Conversely TEEER, due to the use of a logarithmic scale, is useful for contrasting only very diverse values of measured efficiency.

4. CONCLUSION

Proper evaluation of energy-efficiency in green networking is a serious matter that the green networking community shall face soon: on the one hand, power profiles (i.e., the input data) are extremely volatile, due to the fast pace at which technology is evolving; on the other hand, the sprouting of different heterogeneous key performance indicators (i.e., the output data) may pose further challenges. In this work, we overview and compare the most common metrics that can be found in the literature. For illustration purposes, we then consider two case studies: a simple toy-case apt to show qualitative differences between the metrics, and a more realistic case considering widely popular core router devices. We point out that, due to their intrinsic difference and relevance, it is likely no single metric can represent the whole state of the system. Nevertheless, choosing a metric rather than another may yield significantly different evaluation results: thus, we believe that a large consensus needs to be reached as soon as possible on a reduced set of well defined performance indicators to promote cross-comparison.

Acknowledgment

This work was funded by Eureka Celtic TIGER2 project.

5. REFERENCES

[1] Cisco Carrier Routing System. www.cisco.com/en/US/prod/collateral/routers/ps5763/prod_brochure0900aecd800f8118.pdf.

[2] Cisco CRS-1 Carrier Routing System Packet-over-SONET/SDH Physical Layer Interface Module Instalation Note. www.ciscosystems.com/en/US/docs/routers/crs/crs1/plim/installation/guide/pos6436.pdf.

[3] Energy Efficiency for the Network Equipment: Two Steps Beyond Greenwashing. www.juniper.net/us/en/local/pdf/whitepapers/2000284-en.pdf.

[4] Environmental Performance Index (EPI). epi.yale.edu/.

[5] Environmental Sustainability Index (ESI). www.yale.edu/esi/.

[6] Environmental Vulnerability Index (EVI). www.vulnerabilityindex.net/.

[7] GHG Protocol Corporate Standard. www.ghgprotocol.org/standards/corporate-standard.

[8] Global IP Traffic Forecast and Methodology, 2006-2011. www.hbtf.org/files/cisco_IPforecast.pdf.

[9] Greenhouse Gas Equivalencies Calculator. www.epa.gov/RDEE/energy-resources/calculator.html.

[10] Happy Planet Index (HPI). www.happyplanetindex.org/engage/.

[11] ISO/TC207. www.tc207.org/About207.asp.

[12] Juniper Networks T Series Core Routers. www.juniper.net/us/en/local/pdf/datasheets/1000051-en.pdf.

[13] Network and Telecom Equipment - Energy and Performance Assessment. www.ecrinitiative.org/pdfs/ECR_1_0_4.pdf.

[14] T1600 Core Router PIC Guide. www.juniper.net/techpubs/en_US/release-independent/junos/information-products/topic-collections/hardware/t-series/t1600/pics/t1600-pic.pdf.

[15] The Greening of Telecom at The Edge & Metro. archive.ericsson.net/service/internet/picov/get?DocNo=25/28701-FGB101504&Lang=EN&HighestFree=Y.

[16] Transmittal of ATIS Energy Efficiencies Specifications. www.itu.int/md/dologin_md.asp?lang=en&id=T09-FG.ICT-IL-0003!!MSW-E.

[17] Verizon NEBS Compliance: TEEER Metric Quantification. www.verizonnebs.com/TPRs/VZ-TPR-9207.pdf.

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

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

[20] C. Belady. Green Grid Data Center Power Efficiency Metrics: PUE and DCiE, 2008.

[21] L. Ceuppens, A. Sardella, and D. Kharitonov. Power Saving Strategies and Technologies in Network Equipment Opportunities and Challenges, Risk and Rewards. *IEEE/IPSJ*, 2008.

[22] D. Kharitonov. Energy Efficiency for ICT Problem Space, Metrics, Gaps, 2008.

[23] P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan. A Power Benchmarking Framework for Network Devices. Aachen, Germany, May 2009.

[24] M. Mandviwalla and N.-F. Tzeng. Energy-Efficient Scheme for Multiprocessor-Based Router Linecards. In *SAINT '06*, Washington, USA, 2006.

[25] J. Mitchell-Jackson, J. G. Koomey, B. Nordman, and M. Blazek. Data center power requirements: measurements from Silicon Valley. *Energy*, 2003.