

INTRODUCTION TO ENERGY EFFICIENCY IN LARGE-SCALE DISTRIBUTED SYSTEMS

Jean-Marc Pierson¹ and Helmut Hlavacs²

¹*IRIT, University of Toulouse, France*

²*Faculty of Computer Science, University of Vienna, Austria*

1.1 ENERGY CONSUMPTION STATUS

The demand for research in energy efficiency in large-scale systems is supported by several incentives [1–3], including financial incentives by government or institutions to energy efficient industries/companies [4–5]. Indeed, studies such as [6] reported already in 2006 that the information technology (IT) consumption accounts for 5% to 10% of the growing global electricity demand and for a mere 2% of the energy while data centers alone account for 14% of the information and communication technology (ICT) footprint. It was projected that by 2020, the energy demand of data centers will represent 18% of the ICT footprint, the carbon footprint rising at an annual 7% pace, doubling between 2007 and 2020 [7]. The study of Koomey [8] in 2011 highlights that the rise of energy consumption is not as bad as expected in 2007: between 2005 and 2010, the electricity demand for data centers increased by (only) about 56% worldwide instead of the projected doubling and even as low as 36% in the United States. Altogether the electricity used worldwide for operating data centers in 2010 accounted for about 1.3% of total electricity use.

The past 5 years have witnessed the increase of research focusing especially in energy reduction. While being a major concern in embedded systems since decades, the

Large-Scale Distributed Systems and Energy Efficiency: A Holistic View, First Edition.

Edited by Jean-Marc Pierson.

© 2015 John Wiley & Sons, Inc. Published 2015 by John Wiley & Sons, Inc.

problem is quite new in the large-scale infrastructures where performances have been for long the sole parameters to optimize. The motivation comes from two complementary concerns: first, the electrical cost of running such infrastructure is equivalent nowadays to the purchase costs of the equipment during a 4-year usage [9]. Second, electricity providers are not always able to deliver the needed power to run the machines, capping the amount of electricity delivered to one particular client.

Modern usage of ITs relies on the existence of large data centers, high-performance infrastructures, and performance networks, core and mobile networks.

Cloud computing is one of the major evolutions in IT in the past decade. It mainly relies on data centers, some hosting thousands of servers. In 2010, Google was hosting already 900,000 servers (almost 1 million must be the case today, and is estimated to be even more than 1.5 million). In 2013, Microsoft's CEO Steve Ballmer claimed hosting more than 1 million servers. Amazon is guessed to have about the same number of servers. For 1 million servers, at about 200 W per server, plus something like 50 W for cooling and electricity distribution losses, it represents a total power consumption of 250 MW, which is likely 2 TWh/year. However, in [8], it is shown that less than 1% of electricity used by data centers worldwide was attributable to Google's data center operations: The big players are often cited as examples, but they represent only a few percentage of the problem. When they exhibit better energy efficiency, it must be remembered that most other companies have less advances and the average is far from these big players.

While, traditionally, supercomputers have been mainly compared by their raw performance measured now in PFlops (petaflops), they are now also assessed based on their energy efficiency. The ranking of supercomputers by their energy efficiency places great emphasis on their energy consumption through the number of GFlops (gigaflops) they can achieve per watt. For instance, the Tianhe-2 machine, the leading one of the top performance list (Top500¹), delivers a computing power of over 33 PFlops and shows an energy efficiency of 1.9 GFlops/W; while the CINECA machine, which tops the green list (Green500 list²) with an energy efficiency of 3.9 GFlops/W, delivers a low computing power of less than 2 PFlops. Nevertheless, it can be noted that supercomputers are getting greener, or more exactly, their energy efficiency is continuously increasing, while their energy consumption itself is nevertheless growing. Despite this trend, it will be difficult to achieve exascale computing for 20 MW by 2020, the limit given by the US Department of Energy (DoE).

Network operators are among the most power-consuming players. Telecom Italia [10] estimated that its consumption represented 1% of the Italian total power consumption in 2011 (compared to 0.7% in 2008). Similarly, British Telecom estimates 0.7% to be its share of electricity usage in the United Kingdom (2.3 TWh), same as that of NTT in Japan. These numbers account not only for networks but also for associated infrastructures to operate them. For instance, for Telecom Italia, 65% of electricity is consumed in the networks (wired and mobile) and 10% by their data centers. However, these numbers do not account for the equipment at final clients. In France, a study from

¹top500.org.

²green500.org.

IDATE [11] shows that the total electricity consumption of the telecom is 8.5 TWh in 2012 (for 6.7 TWh in 2008). The share is 40% for wired and mobile networks, 6% for data centers, 24% for the ADSL boxes at client places, and, finally, 18% for the fix and mobile phones themselves. We can notice that the total energy consumption of the internet boxes at clients' home is estimated to be 3.3 TWh in 2012 (40 millions of boxes).

The share of power consumption in servers is evolving continuously, because of the improvement in electronics for individual components. Processors (central processing unit, CPU) and memory account together for about 54% of the total consumption, with a rough 37% share for CPU and 17% for memory while the other components are consuming less: Peripheral Component Interconnect (PCI) slots (23%), motherboard (12%), disk (6%), and fans (5%) [12] (see Chapter 2 for details). When graphics processing unit (GPU) are present, they can represent up to a tremendous 50% share of the total consumption. It is, therefore, not surprising that most of the efforts have been put on reducing the power consumption of processors and memory. However, despite the urge for proportional computing already demonstrated in 2007 [13], the current servers are not consuming proportionally to their usage. This makes a lot of work trying to switch off components (consolidation in clouds) or using them at lower speed and capacity (dynamic voltage frequency scaling (DVFS) for CPU, Low Power Idle (LPI) for network cards, disk spin down for hard disk) very valuable. It should be noted that the situation is improving: 5 years ago, a server was consuming as much as 50% of its peak power when idle. Now this drops to 20%, and the peak power itself is decreasing. One can wonder if the works based on the nonproportionality of power consumption will still be interesting in the future. We believe that the delay is still long enough to see achievements in this dimension and also that the aforementioned researches may be used in conjunction and transferred at lower levels (at the components architecture), to allow for actual proportional computing.

One must not forget also the impact of cooling in the global consumption, especially in data centers or large-scale networking equipment rooms. The power usage effectiveness (PUE), promoted since 2007 as a criteria for assessing the power efficiencies of infrastructures, is the ratio between the global power usage to the power usage for IT. While it was common to have a PUE of 2 or more (meaning that as much electricity was used for the infrastructure – mainly cooling and distribution losses) the state-of-the-art values are now at about 1.5 or 1.6. Still 50–60% of power is used for cooling IT equipment. However, as outlined earlier, many data centers do not operate with state-of-the-art solutions, and their PUE are more likely to be at about 1.8 or 1.9 [8].

Energy concerns have been integrated in many works at the different levels of the IT stack: hardware, network, middleware, and software levels in large-scale distributed systems, being high-performance computing (HPC), clouds, or networks.

In the following, we exhibit some actions undertaken at these levels, in particular in the scope of an European-funded initiative.

1.2 TARGET OF THE BOOK

The focus and context of the book is on large-scale distributed systems. We will not study embedded systems in this book. Also we are not investigating hardware-specific

optimization for energy saving. Instead, we focus in this work on energy-efficient computation and communication in large-scale distributed systems. These systems consist of thousands of heterogeneous elements that communicate via heterogeneous networks and provide different memory, storage, and processing capabilities. Examples for very large-scale distributed systems are computational and data grids, data centers, clouds, core and sensor networks, and so on.

The target audiences of the book are manifold: from IT and environmental researchers to operators of large-scale systems, up to small and medium sized enterprises (SMEs) and startups willing to understand the global picture and the state of the art in the field. It helps in building strategies and understanding upcoming developments in the rapid field of energy efficiency to speedup transfer of technologies to industries [14].

1.3 THE COST ACTION IC0804

This section introduces the European Cooperation in Science and Technology (COST) Action IC0804. The COST Action instrument is a 4-year funding scheme in European research framework aimed at helping the development of networks of researchers.³ The funding goes mainly to four main objectives: (i) to organize meetings on a specific field with experts of the members' countries and to foster closer cooperation; (ii) to help exchanges of researchers through short-term scientific missions (STSM), for a duration of 1 week to several weeks ; (iii) to develop the young researchers' skills in the field through the organization of annual training schools; and (iv) to shape and rethink the research agenda on the field. One important information at this stage is that COST Actions do not fund research directly: individual projects still rely on European and national funding agencies.

1.3.1 Birth of the Action

The COST Action IC0804⁴ finds its root in 2008. It started in May 2009 and finished in May 2013. The idea of the Action started with eco-awareness during the previous years among colleagues around the world. This topic appeared in both national and international research agendas and scientific interests.

First, the ecological concerns came more and more on the agenda of many funding institutions, raising the foundations for some research projects in various countries and then at the European level. At the European level, the Environmental Research themes were still not investigating the issues related to the electricity consumption in the IT.⁵ In FP7 (Framework Program 7), the Energy theme⁶ is concerned with the energy consumption and production. The Environment theme of the FP7⁷ focuses its actions on

³www.cost.eu.

⁴http://www.cost.eu/COST_Actions/ict/Actions/IC0804

⁵ec.europa.eu/research/environment/themes/themes_en.htm.

⁶cordis.europa.eu/fp7/energy/.

⁷cordis.europa.eu/fp7/environment/.

environmental issues and climate change but nothing is done on the ICT side. For the ICT theme, the COST Action fitted into Challenge 1 on pervasive and trusted network and service infrastructure and in Challenge 6 on mobility, environmental sustainability, and energy efficiency.⁸ Altogether, these programs raised a number of projects with at least eco-awareness. The Action IC0804 served as a means to coordinate focused initiatives. As for examples, the Virtual Home Environment (VHE) project, a specific subproject of the European Network of Excellence (NoE) Euro-FGI (Design and Engineering of the Future Generation Internet) has been carried out on energy-efficient home networking. In the FP7 call, the AIM⁹ and the BE-AWARE¹⁰ projects were among the first ones to investigate complementary approaches to energy awareness.

Second, conferences specialized in distributed systems, clusters, or grids saw the emergence of research papers in the field of energy efficiency and energy awareness (topics that have finally become part of their calls for research papers). When done separately, the researches are likely to contradict each other, leading to the idea for a common forum and exchange place. Energy has thus become a fixed topic in the distributed systems community research. Also, researchers working outside the field of energy efficiency and savings may want to gain more information about the energy consumption of their distributed systems, especially of large scale or cluster likes. They want to be guided in some good practices so as to use alternative eco-aware solutions if the performance and costs of these are satisfying. The common fear is the critical loss of performance while trying to reduce the ecological footprint. The COST Action aimed to show that efficient energy aware solutions exist. The work carried out in the course of the Action, summarized in the rest of the book, shows that this is achievable.

1.3.2 Development of the Action

The COST Action IC0804 investigated energy efficiency in large-scale distributed systems. These systems include, among others, wired and wireless networks, HPC, cloud and desktop computing, and smart grids. The scientific outcome of the Action have been numerous: a brochure on hardware leverages for energy reduction; some methodologies to monitor and model the energy consumption of hardware and software components under the constraints imposed by power meters (accuracy, frequency); concrete techniques and their analyses for energy savings (from using Power On/Power Off, DVFS, idle states of components, to algorithmic and rethinking of algorithms, and management of the platforms); outline of the need for self-management of infrastructures; trade-off analysis, including not only the electricity price but also other costs (CO₂, quality of experience (QoE), human resources, management efforts of organizations, etc.).

1.3.2.1 Numbers and Facts about the Action. At the end of the Action duration, 23 European countries and 7 non-European institutions were members of the Action. About 150 researchers participated on the various Action activities, coming from academy and industry (including IBM, Microsoft, Intel, Yahoo, Ericsson, and EDF and

⁸cordis.europa.eu/fp7/ict/programme/challenge6_en.html.

⁹<http://www.ict-aim.eu/>.

¹⁰<http://www.energyawareness.eu/beaware/>.

also various SMEs and startups originated during the 4 years). Twenty-four STSMs have been organized, and altogether more than 76 joint publications have been released.

1.3.2.2 Concrete Actions Toward Eco-Friendliness. In this section, we highlight some of the results that can be awarded to the Action. This section is not designed to provide a summary of the research results obtained but rather gives the opportunity to have a glimpse of a selection of them. All the works cited here are collaborative works involving a minimum of two institutions from two different countries. Some of them will be detailed in the upcoming chapters, while the reader can refer to the individual research papers for the others.

- *Hardware.* Energy efficiency is directly related to the efficiency of individual components. Therefore, it is very important to follow the actual development of technologies. Tremendous efforts are spent on improving electronics, and more energy-efficient new technologies are coming out regularly. Still, energy proportional computing [4404806] is not in place, and it is still worth investigating how to use the different possible status for a component (power off, standby, idle, etc.) and for every component: CPU, memory, network, disk, and accelerators (including GPU and cells, for instance). Hardware improves at a fast pace. Future hardwares are moving towards dedicated pieces of hardware instead of the current “able to do all, perfect for nothing” approach. Future is going towards high-scale dedicated infrastructure such as data mining (such as Google and Facebook) for which newly dedicated hardware will be available. From a software point of view, this means that in the future, new possibilities will arise as decisions will be more complex due to software requirements.
- *Monitoring.* Concerning the monitoring of energy consumption, we have witnessed a change from ad-hoc solutions at the beginning of the Action to more industrial solutions at the end, including the raise of the ACPI 4.0 norm, defining that in the future, components will be able to tell about their energy consumption individually. We also saw that the accuracy and the frequency of measurements play a large role in the evaluation of the solutions: Power meters have to be carefully studied and/or chosen before starting any serious experiments. Currently, mostly ad-hoc solutions are provided for monitoring data centers, wired network devices, and wireless networks using external measurement devices to get energy readings and derive models from these readings. Differences in hardware manufacturing require different models. Thus, more information from manufacturers would fasten the model creation step as well as classification of devices in terms of energy consumption characteristics. Often, the required information about measurements is missing in distributed monitoring, which makes it hard to compare results.

During the course of the Action, a joint methodology has been put in place between two teams (France and Spain), which proposed an analysis and evaluation of different external and internal power monitoring devices to validate the accuracy of the equipment in terms of power dispersion and energy consumption. This experimental study is completed by some results for a variety of benchmarks

that exercise intensively the main components (CPU, Memory, hard disk drives (HDDs), and network interface controllers (NICs) of the target platforms. This study has been the first one to carefully deal with the monitoring of energy consumption of computing servers and to correct several wrong assumptions on this topic: internal wattmeters do not register neither an equal energy consumption nor a similar power dispersion. Thanks to the high sampling rate and to the different measured lines, the internal wattmeters allow to better visualize some power fluctuations. However, a high sampling rate is not always necessary to understand the evolution of the power consumption during the execution of a benchmark [15].

- *Modeling the Energy Consumption.* Several techniques exist and have been used within the course of the Action; some very simple and some complicated machine learning or neural networks based. One basic observation is that the accuracy of the model has to be related to the scale of the utilization of the model. For instance, a simple model could easily give a 5–10% accuracy at almost no cost (linear combination of power and load), which is sufficient for actions on the system at the cloud level, done generally at the scale of minutes or more. At the operating system level, models have to be more precise because actions on the system (for instance, using DVFS) are in the order of milliseconds, but at the same time, this accuracy must not come at the cost of complexity and time, because the operating system has to react quickly. When it comes to model the applications and not the system, some tools (for instance, the ectop and valgreen software suite [16]) have been developed for estimating the power consumption of applications.

A very important and difficult research topic addressed by members of the action was understanding how specific applications affect power usage of servers they are executed on. This problem became particularly relevant in the advent of many-core processors, dynamic power states management techniques, and a variety of applications ranging from mobile devices, through virtualized environments, to petascale HPC codes. To solve this problem, colleagues worked on the analysis of applications using information coming from the system (performance counters, Input/Output status, temperature, etc.) and correlated these using statistical methods. Proposed methodology and models enabled estimation of power usage of servers and single applications without the use of power meters [17–21] and modeling application profiles in simulations tools [22]. Methods for application clustering and classification at runtime were also proposed [20, 23].

We also witnessed that only a few works are intended to evaluate and model the power consumption of virtual machines.

- *Taking Actions at the Communication Layer.* On the side of the networks, the potential of energy is also very huge, either in wired or in wireless networks. For instance, in wireless networks, working on energy savings between end-users and base stations has more value than working in the core network: The numbers make the savings.

Concerning wireless networks, researchers investigated in [24] the energy consumption of various wireless network technologies such as WiFi and

WiMAX, in particular considering network interfaces of end systems when receiving data from cloud services. Other researchers analyzed the trade-off between QoE and energy consumption for receiving high-quality video streams via wireless mesh networks as well as WiFi access networks [25]. On the basis of that work, a novel power saving algorithms for continuous media applications (optimized power save algorithm for continuous media applications, OPAMA) over IEEE 802.11 WiFi networks has been developed and evaluated [26].

In [27], Aleksic et al. present joint work assessing the feasibility of using a core network approach where no switching is performed at the Internet Protocol (IP) layer, but end-to-end data connections between ingress and egress nodes are realized by means of digital and optical resources at the transport layer. They introduce an analytical model for estimating the average number of required transport switch ports for different topologies, in order to assess both scalability and power consumption of three particular realizations of these static optical core networks. The results show that the concept of a static optically transparent core network promises high energy efficiency and scalability to several tens of nodes.

Another example of the front of networks is the work on developing an algorithm to solve the resource assignment problem in virtual networks (also known as virtual network embedding) with a focus on energy efficiency. In particular, the approach tries to achieve higher consolidation of resources, thereby allowing to power off part of the substrate network [28].

- *Taking Actions for Data Centers.* Beyond classical approaches for reducing energy in data centers focusing on one aspect of the problem (machines, scheduling, operating system, communication layer, etc.), some colleagues joined efforts to start research on energy efficiency of large-scale distributed systems taking into consideration the thermal aspects. To this end, they took a holistic approach including modeling of application profiles, thermal-aware resource management policies, and appropriate metrics to evaluate efficiency. This joint work was on the basis of the CoolEmAll project.¹¹ The main goal of that project is to decrease energy consumption of data centers by allowing data center designers, planners, and administrators to model and analyze energy efficiency of various configurations and solutions. The project provides models of data center building blocks and tools that apply these models to simulate, visualize, and analyze data center energy efficiency. Both building blocks and the toolkit take into account the aspects that have major impact on actual energy consumption: hardware characteristics, cooling solutions, application profiles, and workload and resource management policies. In addition to common static approaches, the proposed platform enables studies of dynamic states of data centers based on changing workloads, management policies, cooling method, and ambient temperature. To our knowledge, this is the first such holistic approach to date. Especially, combination of detailed application analysis, workload management and scheduling simulation, and heat transfer simulation are not

¹¹coolemall.eu.

available in concurrent solutions. Results have been presented in several joint publications [17, 22, 29, 30].

While all data centers have common problems to be addressed (such as thermal aspects), we also witnessed a huge difference in handling energy efficiency between cloud and HPC systems. While the first one deals with service-level agreement (SLA) and allows for its degradation, HPC system tends to be more conservative, that is, allowing less energy saving leverages.

- *Taking Actions for Cloud Computing.* Solutions for cloud computing are already on the market, some developed in startups issued from the Action members (Eco4Cloud,¹² EasyVirt¹³). Even if more precise allocation of virtual machines on the data centers can achieve better energy savings, already existing solutions can save a lot, about 40–70% of energy costs. For a specific example as results of STSMs, collaborations on energy/resource-efficient management of cloud computing infrastructures started between partners on the long term. They propose to tackle the problem in a hierarchical manner, where they structure all possible adaptation actions into the so-called escalation levels [31]. In [32], Maurer et al. devise an approach for self-adaptive and resource-efficient decision-making considering the three conflicting goals of minimizing the number of SLA violations, maximizing resource utilization, and minimizing the number of necessary time- and energy-consuming reconfiguration actions. Their approach is based on automatically detecting workload volatility in the virtual machines' demands and reaction based on rules. They introduce categorization and present cost- and volatility-based methods for self-tuning. Evaluation shows that in most cases, the self-adaptive approach outperforms the static approach. In [33], the work is extended to include different allocations strategies (vector packing, best fit, Monte Carlo, etc.).
- *Taking Actions for High-Performance Computing:* In HPC, guaranteeing high performance at all costs is a challenge. Savings are less important in terms of percentage (5%, for instance) but the size of the infrastructure is so large that the final energy savings are large as well. Here it is important to understand that a small saving at large scale has more value than large savings at small scale. Exascale will not be achievable without a significant work on energy efficiency of solutions for HPC.

Many organizations have departments and workgroups that could benefit from HPC resources to analyze, model, and visualize the growing volumes of data they need to conduct business. Up to now, most HPC systems are built on general-purpose multi-core processors that use the x86 and power instruction sets. This trend is about to change with the massive investments (several billions of dollar) voted by leading countries in 2012 to build an exascale HPC system by 2019 while staying below a power budget of around 20 MW. Such a platform requires a maximal power consumption of 0.1 W per core. In order to achieve this ambitious goal, alternative low power processor architectures are required

¹²www.eco4cloud.com.

¹³easyvirt.fr.

and two main directions are currently explored: (i) general-purpose graphics processing unit (GPGPU) accelerators or (ii) processors (ARM, Intel Atom, etc.) primarily designed for the mobile and embedded devices market. Building on preliminary results exhibited from the European Mont-Blanc project¹⁴ in an Action meeting that investigates the second approach, partners from Luxembourg and Poland cooperated to compare the performance and energy efficiency of cutting-edge high-density HPC platform enclosures featuring either very high-performing processors (such as Intel Core i7 or E7) yet having low power efficiency, or the reverse, that is, energy-efficient processors (such as Intel Atom, AMD Fusion, or ARM Cortex A9) yet with limited computing capacity [34].

- *Automating Actions.* At large scale, it is obvious that most of the actions that have to be undertaken must be automatized. Autonomic computing is key for achieving the adoption of energy savings mechanisms. Most of the approaches described earlier and in the rest of the book are relying on this concept, more or less integrated in a holistic work.

1.3.2.3 European Projects Related to the Field. During the development of the Action, a number of FP7 European projects have been initialized in the field of energy efficiency in large-scale distributed systems. Most of them have involved some members of the COST Action community. Without entering in the details of the projects themselves, the focus were on data centers (GAMES, FIT4GREEN, ALL4GREEN, COOLEMALL, primeEnergyIT), clouds (EUROCLOUD, ECO2CLOUD), HPC (Mont Blanc), and networks (EARTH, ECONET, TREND). Several results presented in this book have originated from these researches funded by the European FP7 program and developed during the course of the Action.

1.3.3 End and Future of the Action

The Action was completed in May 2013. Researchers are continuing their investigations, building on their experience gained here. Some follow-up may appear on sustainable ultrascale computing and on the smart grid problem. The main lessons learnt from the Action can be summarized as follows:

- Energy efficiency by itself is a very broad subject, encompassing hardware-related aspects and middleware, networking, and software issues. A holistic view is necessary to capture the whole picture. Unfortunately, this broad view is difficult to achieve because it requires a similar broad knowledge for the researchers as well as for the development of automation.
- Knowledge and precision of models are very important in order to describe, address, and, finally, propose solutions for better energy efficiency. While strong mathematical models help to optimize the system theoretically, it can be noted that simple solutions based on easy-to-deploy heuristics have proven to be very effective on real application domains.

¹⁴www.montblanc-project.eu/.

- Developed tools for assessing the energy consumption and taking actions for reducing it exist, at middleware, software, or network layers. Their acceptance in communities is the major challenge for the coming years.
- Electrical power dissipated in data centers and networks is huge. Low hanging fruits for saving energy are still here in production sites while research has demonstrated the effectiveness of actual solutions. The lack of formal standardization and regulation in this field is a major problem [35].
- While individual hardware components become more and more energy efficient, their proliferation and the foreseen data deluge makes the problem of energy consumption more and more serious.
- As the infrastructure grows, trade-off between energy consumption and performance becomes a key issue. With a million-core machine, resilience, fault tolerance, and the expected redundancy represent opposite objectives to energy efficiency: the fight will be even more exciting.
- The link with electrical power production (in particular renewable or clean energy) is mandatory. Smart grid development will help to solve this issue. However, it must be understood that the electricity production with renewable energy is (and will be for a certain future) far from being sufficient to power a very large data center.

1.4 CHAPTERS PREVIEW

The rest of the book is organized around the major topics that have been investigated during the course of the Action.

Chapter 2 covers the technological possibilities at the hardware level and infrastructure level for energy efficiency. It discusses individual components (CPU, memory, etc.) and the global view, including large-scale infrastructure, including power supply and cooling.

Chapter 3 focuses on wired networks. After highlighting the significant energy consumption of existing wired communication networks, this chapter will examine various means of operating such networks more efficiently. The chapter examines the components that make up wired communications network and their differing characteristics between the access and core, as well as patterns of traffic behavior. Once this is done, the chapter focuses on static (network planning) and dynamic (traffic engineering) schemes that can be used to reduce the energy consumption of the network. The chapter also pays attention to a number of challenges/open research questions that need to be resolved before the implementation of such schemes. These include issues with migration and resilience. Finally, a summary draws out the key themes that have been covered.

Chapter 4 deals with wireless networks. It exhibits the metrics and the trade-offs used in these networks. It discusses the methodology for measuring or profiling the energy consumption of mobile devices and the infrastructures. It explains the different access networks Long Term Evolution(LTE), Wireless Local Area Network(WLAN) their impact in terms of energy consumption linked with the notion of quality of services. Wireless sensor networks and ad-hoc networks [vehicular ad-hoc networks

(VANET), mobile ad hoc network (MANET), opportunistic networks, delay tolerant network (DTN)] conclude this chapter.

Chapter 5 is interested in power modeling techniques and tools. It provides a broad discussion of the techniques involved in performance-based power estimation modeling, as well as the driving motivation for the use of estimation models over the alternative hardware metering devices. A discussion on the impact of providing power consumption feedback to users is given. Power modeling for single core, multi-cores, multiprocessors, and distributed systems (especially for clouds) is discussed.

Chapter 6 discusses green data centers, from their design to their operations. It describes the possibilities of energy reduction in data centers, taking into account the servers, and also the infrastructure of the data center (cooling, heating, power distribution, etc.) and the interconnect inside the data center. Solutions at architectural and middleware levels are presented.

Chapter 7 is investigating the side of energy efficiency in HPC, in particular taking into account the high utilization of HPC platforms, the hybrid architectures (including GPU), and the race for performances towards exascale.

Chapter 8 aims at providing an overview on resource management, in particular in cloud computing, in particular aspects related to the SLA enactment, economic issues, and virtualization technologies and a theoretical part on scheduling, including exact solutions, centralized heuristics, and distributed solutions.

Chapter 9 exhibits the advances on energy efficiency solutions for peer-to-peer (P2P) systems and applications with special focus on file sharing applications and epidemic protocols.

Chapter 10 concludes the book with a discussion about sustainability for large-scale computing systems: environmental, economic, and standardization aspects are outlined and open the book to less technical aspects but still mandatory to effect the transfer of researches to industry and society.

ACKNOWLEDGEMENT

This work was partially supported by the COST (European Cooperation in Science and Technology) framework, under Action IC0804 (www.cost804.org).

REFERENCES

- [1] COST European Cooperation in Science and Technology. ICT COST Action IC0804 on Energy efficiency in large scale distributed systems. Available at http://www.cost.eu/COST_Actions/ict/Actions/IC0804.
- [2] Cameron KW, Pruhs K, Irani S, Ranganathan P, Brooks D. Report of the science of power management workshop; 2009. Available at <http://scipm.cs.vt.edu/SciPM-ReportToNSF-Web.pdf>.
- [3] U.S. Environmental Protection Agency ENERGY STAR Program. Report to congress on server and data center energy efficiency; 2007. Available online www.energystar.gov/ia/partners/prod_development/downloads/epa_datacenter_report_congress_final1.pdf. Accessed 2014 Oct 16.

- [4] Naegel B. The Data Center Journal, Energy Efficiency: The New SLA, Dec; 2008. http://datacenterjournal.com/index.php?option=com_content&task=view&id=2352&Itemid=43.
- [5] Barroso LA and Holzle U. IEEE Computer, The Case for Energy-Proportional Computing, 40; 2007. http://www.computer.org/portal/site/computer/index.jsp?pageID=computer_level1&path=computer/homepage/Dec07&file=feature.xml&xsl=article.xsl.
- [6] Bertoldi P, Atanasiu B. Electricity consumption and efficiency trends in the enlarged European Union; 2006. Available online re.jrc.ec.europa.eu/energyefficiency/pdf/eneffreport2006.pdf. Accessed 2014 Oct 16.
- [7] The Climate Group. SMART 2020: enabling the low carbon economy in the information age. Technical report; 2008. A report by The Climate Group on behalf of the Global eSustainability Initiative (GeSI), http://www.theclimategroup.org/_assets/files/Smart2020Report.pdf
- [8] Koomey J. Growth in data center electricity use 2005 to 2010; 2011.
- [9] Koomey JG, Belady C, Patterson M, Santos A. Assessing trends over time in performance, costs, and energy use for servers; 2009.
- [10] Bianco C, Cucchietti F, Griffa G. (2006) Energy consumption trends in the Next Generation - Access Network-a Telco perspective <http://www.polisave.polito.it/greenet/docs/TlcEnergyTrends.pdf> Accessed 2014 Oct 16.
- [11] IDATE. 2010, http://www.fftelecoms.org/sites/default/files/contenus_lies/007.15_idate_presentation_conference_de_presse.pdf. Accessed 2014 Oct 16.
- [12] Fan X, Weber W-D, Barroso LA. Power provisioning for a warehouse-sized computer. ISCA '07: Proceedings of the 34th Annual International Symposium on Computer Architecture. New York (NY): ACM; 2007. p 13–23.
- [13] Barroso LA, Holzle U. The case for energy-proportional computing. IEEE Comput 2007;40(12):33–37.
- [14] Herzog C, Lefevre L, Pierson J-M. Green IT for innovation and innovation for green IT: the virtuous circle. In: Hercheui MD, Whitehouse D, McIver W Jr., Phahlamohlaka J, editors, *ICT Critical Infrastructures and Society, 10th IFIP TC9 International Conference on Human Choice and Computers (HCC); 2012 Sep 27–28. Number 386 in IFIP AICT*. Amsterdam: Springer; 2012. p 79–89. Available at <http://www.springerlink.com>. Accessed 2014 Oct 16.
- [15] Diouri MEM, Dolz MF, Glück O, Lefèvre L, Alonso P, Catalán S, Mayo R, Quintana-Ortí ES. Solving some mysteries in power monitoring of servers: take care of your wattmeters! Energy Efficiency in Large Scale Distributed Systems conference (EE-LSDS); Vienne, Autriche; 2013.
- [16] Cupertino LF, Costa G, Sayah A, Pierson J-M. Energy consumption library. *Energy Efficiency in Large Scale Distributed Systems*, Vienna, Austria, April 2013. Lecture notes in computer science. Springer Berlin Heidelberg; 2013. p 51–57.
- [17] Cupertino LF, Da Costa G, Sayah A, Pierson J-M. Valgreen: an application's energy profiler. International Journal of Soft Computing and Software Engineering, SCSE'13 Conference, Volume 3(3):(on line), March 2013.
- [18] Costa GD, Hlavacs H. Methodology of measurement for energy consumption of applications. 2010 11th IEEE/ACM International Conference on Grid Computing (GRID); 2010. p 290–297.
- [19] Witkowski M, Oleksiak A, Piontek T, Węglarz J. Practical power consumption estimation for real life {HPC} applications. Future Gener Comput Syst 2013;29(1):208–217.
- [20] Jarus M, Oleksiak A, Piontek T, Węglarz J. Runtime power usage estimation of HPC servers for various classes of real-life applications, Future Generation Computer Systems Volume 36, July 2014, P 299–310, ISSN 0167–739X, <http://dx.doi.org/10.1016/j.future.2013.07.012>.

- [21] Landry G, Chetsa T, Lefevre L, Pierson J-M, Stolf P, Da Costa G. DNA-inspired scheme for building the energy profile of HPC systems. E2DC - 1st International Workshop on Energy-Efficient Data Centres - 2011; Madrid, Europe; 2012.
- [22] Piatek W. Dcworms - a tool for simulation of energy efficiency in data centers. In: Pierson J-M, Da Costa G, Dittmann L, editors, *Energy Efficiency in Large Scale Distributed Systems*, Vienna, Austria, April 2013. Lecture notes in computer science. Springer Berlin Heidelberg; 2013. p 118–124.
- [23] Da Costa G, Pierson J-M. Characterizing applications from power consumption: a case study for HPC benchmarks. International Conference on ICT as Key Technology for the Fight Against Global Warming (ICT-GLOW); 2011 Aug 29-Sep 02 (electronic medium). Toulouse: Springer; 2011. Available at <http://www.springerlink.com>. Accessed 2014 Oct 16.
- [24] Bernardo V, Curado M, Staub T, Braun T. Towards energy consumption measurement in a cloud computing wireless testbed. *Proceedings of the 2011 1st International Symposium on Network Cloud Computing and Applications, NCCA '11*. Washington (DC): IEEE Computer Society; 2011. p 91–98.
- [25] Hock D, Bernardo V, Zinner T, Wamser F, Hummel KA, Curado M, Pries R, Braun T, Tran-Gia P. Evaluating the trade-off between energy-efficiency and QoE in wireless mesh networks. 4th International Conference on Communications and Electronics (ICCE 12); Hue, Vietnam; 2012.
- [26] Bernardo V, Curado M, Braun T. Enhancing IEEE 802.11 energy efficiency for continuous media applications. In: Pierson J-M, Da Costa G, Dittmann L, editors. *Energy Efficiency in Large Scale Distributed Systems*, Vienna, Austria, April 2013. Lecture notes in computer science, Springer Berlin Heidelberg; 2013. p 203–217.
- [27] Aleksic S, Van Heddeghem W, Pickavet M. Scalability and power consumption of static optical core networks. Global Communications Conference (GLOBECOM), 2012 IEEE; 2012. p 3465–3471.
- [28] Botero JF, Hesselbach X, Duelli M, Schlosser D, Fischer A, De Meer H. Energy efficient virtual network embedding. *IEEE Commun Lett* 2012;16(5):756–759.
- [29] Berge M, Da Costa G, Kopecki A, Oleksiak A, Pierson J-M, Piontek T, Volk E, Wesner S. Modeling and simulation of data center energy-efficiency in coolemall. In: Huusko J, Meer H, Klingert S, Somov A, editors. *E2DC'12 Proceedings of the First international conference on Energy Efficient Data Centers*. Volume 7396 of Lecture notes in computer science. Springer Berlin Heidelberg; 2012. p 25–36.
- [30] Vor Dem Berge M, Da Costa G, Jarus M, Oleksiak A, Piontek W, Volk E. Modeling data center building blocks for energy-efficiency and thermal simulations. International Workshop on Energy-Efficient Data Centres, Co-located with E-Energy (E2DC); 2013 May 21 (electronic medium). Berkeley: Springer; 2013. Available at <http://www.springerlink.com>.
- [31] Maurer M, Brandic I, Sakellariou R. Adaptive resource configuration for cloud infrastructure management. *Future Gener Comput Syst* 2013;29(2):472–487.
- [32] Maurer M, Brandic I, Sakellariou R. Self-adaptive and resource-efficient SLA enactment for cloud computing infrastructures. 2012 IEEE 5th International Conference on Cloud Computing (CLOUD); 2012. p 368–375.
- [33] Borgetto D, Maurer M, Da Costa G, Pierson J-M, Brandic I. Energy-efficient and SLA-aware management of IAAS clouds. *Proceedings of the 3rd International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet, e-Energy '12*. New York (NY): ACM; 2012. p 25:1–25:10.

- [34] Jarus M, Varrette S, Oleksiak A, Bouvry P. Performance evaluation and energy efficiency of high-density HPC platforms based on Intel, AMD and ARM processors. In: Pierson J-M, Da Costa G, Dittmann L, editors. *Energy Efficiency in Large Scale Distributed Systems*, Vienna, Austria, April 2013. Lecture notes in computer science. Springer Berlin Heidelberg; 2013. p 182–200.
- [35] Herzog C. Standardization bodies, initiatives and their relation to green IT focused on the data centre side. *Energy Efficiency in Large Scale Distributed Systems (EE-LSDS)*; 2013 April 22-24; LNCS. Vienna: Springer; 2013. Available at <http://www.springerlink.com>. Accessed 2014 Oct 16.

