

Energy-Efficient Cloud Computing

Andreas Berl^{1*}, Erol Gelenbe², Marco di Girolamo³, Giovanni Giuliani³,
Hermann de Meer¹, Minh Quan Dang⁴ and Kostas Pentikousis⁵

¹University of Passau, Innstr. 43, 94032 Passau, Germany

²Imperial College London, South Kensington Campus, London SW7 2AZ, UK

³HP IIC (Italy Innovation Centre), Italy

⁴School of Information Technology, International University in Germany, Bruchsal, Germany

⁵VTT Technical Research Centre of Finland, Kaitoväylä 1, FI-90571 Oulu, Finland

*Corresponding author: berl@uni-passau.de, Tel.: +49 851 5093029, Fax: +49 851 5093052

Abstract. Energy efficiency is increasingly important for future information and communication technologies (ICT), because the increased usage of ICT, together with increasing energy costs and the need to reduce greenhouse gas emissions call for energy-efficient technologies that decrease the overall energy consumption of computation, storage and communications. Cloud computing has recently received considerable attention, as a promising approach for delivering ICT services by improving the utilization of data-centre resources. In principle, cloud computing can be an inherently energy-efficient technology for ICT provided that its potential for significant energy savings that have so far focused on hardware aspects, can be fully explored with respect to system operation and networking aspects. Thus this paper reviews the usage in the context of Cloud Computing, of methods and technologies currently used for energy-efficient operation of computer hardware and network infrastructure. After surveying some of the current best practice and relevant literature in this area, this paper identifies some of the remaining key research challenges that arise when such energy saving techniques are extended for use in Cloud Computing environments.

Keywords: Energy-efficient computing and networking, energy-aware data centres, cloud computing

1. Introduction

Significant savings in the energy budget of a data centre, without sacrificing Service Level Agreements, are an excellent economic incentive for data centre operators, and would also make a significant contribution greater environmental sustainability. According to Amazon.com's estimates [1], at its data centres, expenses related to the cost and operation of the servers account for 53% of the total budget (based on a 3-year amortization

schedule), while energy-related costs amount to 42% of the total, and include both direct power consumption (approx. 19%) and the cooling infrastructure (23%) amortized over a 15-year period.

Dennis Pamlin, the Global Policy Advisor of WWF-Sweden [2] highlighted different IT solutions and their beneficial impact on green house gases (GHG), which include CO₂ emissions. These opportunities include IT based solutions: e.g., Smart Buildings, Smart Transportation and Communication, Smart Commerce and Services, and Smart Industrial production. The colloquial term “smart” in this case means “with low carbon footprint”, showing that the adoption of such “smart” IT solutions will enable a potentially large GHG reduction, including ICT itself is a large power consumer (and therefore a GHG emitter), and IT solutions that have a huge potential impact in reducing GHG emissions in many sectors.




Figure 1: Energy distribution in the data centre

Based on a recent “Data Centre Energy Forecast Report” [3], it can be expected that savings of the order of 20% can be achieved in server and network energy consumption with respect to current levels, and that these savings may induce an additional 30% saving in cooling needs as detailed in a study by HP and the Uptime Institute [4]. It shows that “most of data centre power is spent on cooling ICT equipment (between 60% and 70%)”. Thus there are very significant economic and environmental gains to be obtained from a serious research thrust on energy efficiency in the general area of IT and Computer Networks. In particular, Cloud Computing is an inherently energy-efficient virtualization technique [5], in which services run remotely in a ubiquitous computing cloud that provides scalable and virtualized resources. Thus peak loads can be moved to other parts of the cloud and the aggregation of a cloud’s resources can provide higher hardware utilization.

The rest of this paper is organized as follows. After reviewing the current state of energy efficiency in ICT in Section 2, we discuss the salient aspects of energy-efficient Cloud Computing in Section 3. Then in Section 4, we detail the main research challenges that lie ahead, and provide concluding remarks in Section 5

2. Current State of Energy Efficiency in ICT Infrastructures

ICT consumes an increasing amount of energy, but is also instrumental in increasing productivity and economic prosperity and in reducing energy expenditure from other sources through E-Work, E-Commerce and E-Learning. Traditional network design has sought to minimise infrastructure costs and maximise Quality of Service (QoS). However, ICT also plays a complex role in energy consumption via the “communicate more and travel less” paradigm, as well as through the use of smart devices in homes and offices to optimise energy management. Thus, ICT can reduce energy consumption and carbon emissions, but this potential reduction is partially offset by the power used by data centres and computer networks [6] which runs in the billions of dollars or euros. Thus, a fraction of energy savings in ICT and networks could lead to significant financial and carbon savings. In this section, we will review recent research in energy-efficiency for standalone hardware, and then review work that considers energy consumption as part of the cost functions to be used for scheduling in multiprocessor and grid architectures. Finally, we briefly review energy consumption in cluster servers and wired/wireless networks.

2.1 Energy-efficient hardware

One of the approaches to increase the energy efficiency is to develop more energy-efficient hardware. This effort is fostered by labels such as the US Energy Star [7] or the European TCO Certification [8] which rate IT products (mostly monitors) according to their environmental impact. Novel emerging technologies such as solid-state disks are content with much less energy than the currently used hard disk drives. Computer power can be saved by means of various well-known techniques. First, the processor can be powered down by mechanisms like SpeedStep [9], PowerNow[10], Cool’nQuiet [11] or Demand-Based Switching [12]. These measures enable slowing down CPU clock speeds (clock gating), or powering off parts of the chips (power gating), if they are idle [13],[14]. By sensing lack of user-machine interaction, different redundant hardware parts can incrementally be turned off or put in hibernating mode (display, disk, etc.).

The Advanced Configuration and Power Interface (ACPI) specification [15] defines four different power states that an ACPI-compliant computer system can be in. These states range from G0-Working to G3-Mechanical-Off. The states G1 and G2 are subdivided into further substates that describe which components are

switched off in the particular state. For devices and the CPU, separate power states (D0-D3 for devices and C0-C3 for CPUs) are defined which are similar to the global power states. Some of the mentioned techniques are usually applied to mobile devices but can be used for desktop PCs as well.

2.2 Energy-aware scheduling in multiprocessor and Grid systems

Energy-aware scheduling in multiprocessor and Grid systems is a widely discussed problem in the literature as the following overview shows. In [16] the authors present an energy-aware method to schedule multiple real-time tasks in multiprocessor systems that support dynamic voltage scaling (DVS). The key in their approach is to consider the probabilistic distribution of the tasks' execution time in order to partition the workload and reduce energy consumption. Memory energy consumption [17] can also be reduced by scheduling techniques that impact the effectiveness of frequency scaling by combining the effect of tasks running on a multicore system, including memory contention and the technical constraint of chip-wide frequency and voltage settings. The dynamic voltage scaling capability [18] addresses energy minimization for periodic pre-emptive hard real-time tasks that are scheduled on an identical multiprocessor platform. AlEnawy and Aydin [18] suggest partitioned scheduling and assume that the tasks are assigned rate-monotonic priorities. To solve this NP-hard problem, they proposed an integrated approach that consists of three different components: rate monotonic scheduling (RMS), admission control test, the partitioning heuristic and the speed assignment algorithm. All the above work proposes to control the energy consumption of hardware by adjusting voltage levels.

Liu et al. [19] consider energy-efficient scheduling for data grids supporting real-time and data-intensive applications, and use both the location of data and properties of the application properties, to design a novel Distributed Energy-Efficient Scheduler (DEES) that aims to seamlessly integrate task scheduling tasks with data placement strategies to provide energy savings. The main energy savings are obtained by reducing the amount of data replication and task transfers. In [20] an energy-constrained scheduling scheme for a grid environment is investigated both for energy minimization in mobile devices and for grid utility optimization, by formalizing energy-aware scheduling using nonlinear optimization theory under the constraints of energy budget and the job deadline. In addition, [20] proposes a distributed pricing scheme that makes trade-offs between energy and deadlines to achieve a system-wide optimization based on the preference of the grid user.

Another approach that increases energy efficiency in data centres is based on server consolidation by service virtualization [21][22][23][24][25]. Virtualization partitions computational resources and allows the sharing of hardware. Many services often need only a small fraction of the available computational resources

[26] of a data-centre server. However, even when run at a low utilization, servers typically need up to 70% of their maximum power consumption [27]. Such services can be virtualized and run within a virtual machine resulting in significant increases in overall energy efficiency. Depending on their utilization, many virtual machines can run on a single hardware unit (server consolidation). Therefore, less hardware is needed overall, thus reducing energy wasted for cooling, while the deployed hardware utilization increases. This consolidation of shared hardware fosters energy efficiency, measured as workload accomplished per consumed energy [28].

Resources can be virtualized on different layers and implement different forms of virtualization: full virtualization, hosted virtualization, or operating system (OS) layer virtualization [29]. When system virtualization [30] (e.g., full virtualization) is supported, the virtualization software emulates full-featured hardware and runs on top of the local OS. In the paravirtualization approach, guest virtual machines (VMs) are modified in order to perform so called “hyper calls” instead of system calls, leading to higher performance of VMs [31], as used, for instance, in the XEN systems. In XEN 3.0 [32] guests can be virtualized without being modified by using virtualization support of X86 CPUs. OS-layer virtualization has been proposed in the Linux-VServer [33][34], a kernel based virtualization.

It is important to point out that virtualization comes at a cost which needs to be managed efficiently. When resources are virtualized, additional management of VMs is needed to create, terminate, clone, or move VMs from host to host. Migration of VMs can be done offline (the guest in the VM is powered off) or online (live-migration of a running VM to another host). The management solution Infrastructure 3 [35] of VMWare, for instance, supports live-migration.

2.3 Power minimisation in clusters of servers

Recent research has considered power minimisation in server clusters, with guaranteed throughput and response time [36]. Energy consumption depends primarily on CPU utilisation, but components such as disks, memory, network devices, also use energy so that a server that seems to remain idle may still use up to 60% of its peak power. In [37] policies are developed that use economic criteria and energy as criteria to dispatch jobs to a small set of active servers, while other servers are down to a low-power state. Similar dynamic provisioning algorithms [38] are studied for long-lived TCP connections as in instant messaging (IM) and gaming. A queuing approach to [39] dynamic provisioning technique has also been studied to obtain the minimum number of servers required to respect the required QoS and reactive provisioning can be used to compensate for sudden surges in load.

2.4 Power minimization in wireless and wired networks

According to some estimates, “the Internet” may be consuming more than 860 TWh annually [40], but such figures can only be considered as educated guesses, due to the number of assumptions one has to make. Traditionally fixed network operators have not considered energy consumption as a major cost factor. Lately, however, as sustainability is becoming a key business objective, fixed network operators are looking for ways to decrease their energy footprint. On the contrary, wireless network operators due to regulatory requirements and operational considerations regarding base station deployments have been trying to minimize energy consumption for over a decade. In fact, it is reported that [41] the radio access network (rather than the core network) is the most energy consuming part of the infrastructure, and in many cases the associated energy bills are comparable to the total costs for the personnel who work on network operations and maintenance. The ICT energy estimates in [41] report that the Vodafone Group radio access network alone consumed nearly 3 TWh in 2006.

Surprisingly enough, energy savings for infrastructure networks have not received much attention until very recently, while energy-saving routing protocols in wireless sensor networks have already been studied in detail [42][43][44] because of the specific needs of battery powered networks and the related research has included the use of topology control (TC) [45][46] that modify the network graph to optimize properties such as network capacity and QoS. Since processing and transmission power in nodes are the essential consumers of energy, it is also necessary to optimize the number of hops traversed by packets. An interesting trade-off then arises between high transmission power that can reduce the number of hops, and low transmission power that can lead to more hops being necessary due to shorter ranges, and transmission interference which can be affected by power in a complex manner. Related work can be found in [42] and in [43] the idea of turning nodes on and off is also considered. In a wired node, power consumption depends and influences other factors, such as the node’s throughput; furthermore up to 60% of a node’s energy consumption can originate with peripheral devices such as link drivers. Turning wired network nodes on and off may be very difficult in a wired context because of the high volumes and rate of traffic and the stringent QoS constraints. Routing for wireless ad-hoc networks with battery-powered nodes [47] have also been considered in the selection paths so as to satisfy QoS constraints and minimise power. Generally speaking, we feel that the research community has begun to seriously consider energy consumption in infrastructure networks, and the IEEE has now focused on developing a standard for energy-efficient Ethernet (IEEE 802.3az).

In a future Internet where Cloud Computing may become a mainstay for economic growth, businesses and individuals will require energy-efficient operation that involves not only computation and storage facilities but networking as well. It is further anticipated that the majority of users will access Cloud Computing resources from mobile, battery-powered devices, which impose stringent limitations on power consumption. Clearly, one therefore needs to address not only the issues arising from individual components (such as storage and processing elements), resource utilization algorithms (job scheduling, virtualization, migration), and topology considerations, but the entire chain of services and infrastructure enablers.

3. Towards Energy-Efficient Cloud Computing

The previous discussion highlights the need to develop a comprehensive approach for energy efficiency that involves all system layers and aspects, including physical nodes, cooling of nodes, networking hardware, communication protocols, and finally the servers and services themselves. The conceptual framework of Cloud Computing may therefore be a way forward to analyze, identify and implement overall energy savings in a system to attain truly “Green” computing services.

In contrast to hardware-oriented optimization, software systems can potentially be optimized at development time by specifying their energy characteristics and by adapting the implementation. However, this requires individual adaptation of each component and it also requires understanding the interaction between individual components when they operate as a system. A major challenge is therefore to explore the relations among system components and the tradeoffs that can result in **optimal balance between performance, QoS, and energy consumption and include self-aware runtime adaptation [40][48][49][50]**. Thus in this section we will briefly discuss some areas of energy-efficiency research based on a Cloud Computing perspective.

Significant energy savings can result from using energy-aware scheduling mechanisms pervasively throughout a system. Our survey in the previous sections shows that progress has been made in this area, but that much more needs to be done in holistically examining where and how such mechanisms are needed, in applying the available mechanisms, and inventing and evaluating new ones as needed. In the case of any one service provider site, algorithms to multiplex and de-multiplex workload in order to save energy are needed, and they should incorporate the trade-off between performance and the reduction in service cost due to energy savings. In addition to scheduling and the mapping of workflows, the improvement of energy-aware cloud applications themselves can also benefit from software optimization.

In a business environment based on Cloud Computing, workflows that run over many sites will tend to be popular. Thus, developing methods that map the workflow onto resources under the constraint of energy optimization becomes a central problem of great value and novelty. Furthermore, in order to comprehensively raise the energy efficiency of a system, all of its layers have to be considered, including application layer services. Services have different needs concerning the environment they are running on or have special properties that support the energy efficiency of the underlying system (e.g. certain usage patterns). A service, for instance, might only be used weekdays, say, from 8h to 18h or have peak usage at a certain time of the day. A user also may also consider a trade-off between a more energy-efficient service and a more reliable or faster service, and compose the service in a way that fits its needs. Thus it should also be possible to develop accounting mechanisms that track and depend on the energy that has been used by a service.

3.1 Energy-aware data centres

The key current technology for energy-efficient operation of servers in data centres is virtualization. Virtual machines that encapsulate virtualized services can be moved, copied, created, and deleted depending on management decisions. Consolidating hardware and reducing redundancy can achieve energy efficiency. Unused servers can be turned off (or hibernated) to save energy. Some hardware gets higher load, which reduces the number of physical servers needed. However, the degree of energy-efficient self-management in data centres is still limited today. Services should not only be virtualized and managed within a data-centre site but they should be moved to other sites if necessary. Not only the aspect of load has to be considered, also the “heat” generated by a service has to be measured and accounted for before migrating operations. Every operational physical node produces heat. When a particular node is excessively used or is near other high-loaded nodes, hotspots can appear in a given data centre. To avoid such hotspots, heat can be distributed across sites. Further, services can be moved from sites with high load or high temperature to sites with smaller loads and lower temperatures. Generally, services should be moved to those locations, where they can operate in the most energy-efficient way. This kind of energy-efficient management of resources has to be realized by an autonomous energy management that is as transparent as possible to the user of a service. Energy-related problems have to be solved according to defined policies without needing human interaction. Machine-readable descriptions of the needs and features of services, servers, networks and even whole sites have to be available to enable energy efficiency in the highly autonomous and adaptive [48] systems of the future.

3.2 Energy savings in networks and protocols

Research has shown that communications in particular is one of the largest consumers of energy, however energy optimization for communications must deal with the tradeoffs between performance, energy savings, and QoS [51].

Some hardware already offers features that create an opportunity for energy-efficient operation such as turning off network interfaces, throttling of processors. Network protocols could also be optimized, or even be redeveloped in a way that enhances the energy-efficient operation of the network elements. Network devices could be enabled to delegate services to other devices, so as to transfer services from energy inefficient to amore energy-efficient devices or to devices that need to be always on, while certain other devices are turned off. The delegating device can then become dormant and be turned off. Currently, many basic network services have to remain active to periodically confirm their availability even when no communication is taking place. These “soft states” make it impossible to turn off certain system components; therefore, new protocols need to be designed to work around such soft states so as to increase the energy efficiency of the network. Signalling can also be revisited in this context; whereas data and signalling traffic vary widely, the same technology and mechanisms are used for both (in so called in-band signalling). While signalling needs only low bandwidth but can occur anytime, data traffic occurs after signalling has taken place, and usually requires high bandwidth and traverses all network layers up to the application layer, and uses processing power through multiple layers of the network. Therefore the use of out-of-band signalling should also be evaluated to design and improve energy-aware communication protocols.

3.3 The effect of Internet applications

So far, we have considered the opportunities offered by Cloud Computing as a possible foundation for energy-efficient ICT infrastructures, but have not discussed the nature of the applications themselves. We note that one large application area for the Internet is in information dissemination. From digital cameras embedded in mobile phones to environmental sensors to Web 2.0, end users are generating and interconnecting unprecedented amounts of information and this trend is expected to continue unabated. However, the professional, expedited, and reliable distribution of content requires increasing investments in infrastructure build out and maintenance, and a matching electricity bill to run the underlying ICT. Web, peer-to-peer, and web-based video-on-demand services currently dominating Internet traffic and, taken together, consistently comprise 85% or more of the Internet traffic mix for several years. In practice, dissemination networks operate using methods and paradigms based on remote-access, replicating functionality in several parts of the protocol stack, and fail to benefit from

recent advances in wired and wireless communications, storage technologies, and Moore's Law. If Cloud Computing becomes a significant platform for producing and accessing information, the amount of data that will be transferred over the Internet will increase significantly. Content replication and dissemination algorithms will then need to consider energy as a key parameter of optimal operation, and therefore Cloud Computing calls for a thorough re-examination of the fundamentals of major computation/communication/storage and energy/performance tradeoffs.

4. Conclusions

This paper has reviewed the potential impact of energy-saving strategies for the management of integrated systems that include computer systems and networks. We have surveyed the contributions that are available in this area from recent research. We propose that Cloud Computing with Virtualization as a way forward to (i) identify the main sources of energy consumption, and the significant trade-offs between performance, QoS and energy efficiency, (ii) and offer insight into the manner in which energy savings can be achieved in large scale computer services that integrate communication needs. Based on the approaches that we have identified, we think that specific plug-ins and energy-control centres for networked large-scale hardware-software can be implemented and that they can have significant impact, including:

- Reducing the software and hardware related energy cost of single or federated data centres that execute "cloud" applications,
- Improving load balancing and hence QoS and performance of single and federated data centres,
- Reducing energy consumption due to communications,
- Saving Green House Gas and CO₂ emissions resulting from data centres and networks so as to offer computing power that is "environment protecting/conserving".

Such improvements can have additional impact by reducing energy utilization for transportation and work by encouraging "green" ICT based smart solutions for E-Work, E-Learning, and smart climate control for homes.

References

- [1] Hamilton, J. (2009) Cooperative Expendable Micro-Slice Servers (CEMS): Low Cost, Low Power Servers for Internet-Scale Services. *Proceedings of 4th Biennial Conference on Innovative Data Systems Research (CIDR)*, Asilomar, California, USA, January.

- [2] Pamlin, D. (2008) *The potential global CO₂ reductions from ICT use: Identifying and assessing the opportunities to reduce the first billion tonnes of CO₂*, May, WWF-Sweden.
- [3] Accenture (2008) *Data Centre Energy Forecast Report*, Final Report, Silicon Valley Leadership Group, July.
- [4] Malone, C. and Belady, C. (2006) Metrics to Characterise Data Centre & IT Equipment Energy Use, *Proceedings of Digital Power Forum*, Richardson, Texas, USA, September.
- [5] Hewitt C. (2008) ORGs for Scalable, Robust, Privacy-Friendly Client Cloud Computing, *IEEE Internet Computing*, September, pp. 96 – 99, IEEE, NJ, USA.
- [6] Fan, X., Weber, W.-D., and Barroso, L. A. (2007) Power provisioning for a warehouse-sized computer, *Proceedings of 34th Annual International Symp. on Computer Architecture*, June, pp. 13 – 23, ACM, New York.
- [7] *Energy Star*, <http://www.energystar.gov>, <http://www.eu-energystar.org>
- [8] *European TCO Certification*, <http://www.tcodevelopment.com>
- [9] Intel whitepaper 30057701 (2004) *Wireless Intel SpeedStep Power Manager: Optimizing Power Consumption for the Intel PXA27x Processor Family*, <http://sunsite.rediris.es/pub/mirror/intel/pca/applicationsprocessors/whitepapers/30057701.pdf>
- [10] AMD *PowerNow! Technology*, http://www.amd.com/us-en/Processors/ProductInformation/0,,30_118_10220_10221%5E964,00.html
- [11] AMD *Cool'n'Quiet Technology*, http://www.amd.com/us-en/Processors/ProductInformation/0,,30_118_9485_9487%5E10272,00.html
- [12] Intel Software Network (2008) *Enhanced Intel SpeedStep® Technology and Demand-Based Switching on Linux*, <http://softwarecommunity.intel.com/articles/eng/1611.htm>
- [13] Whitepaper Revision-001 (2007) *ENERGY STAR* System Implementation*, Intel with technical collaboration from the U.S. Environmental Protection Agency, http://www.energystar.gov/ia/partners/prod_development/revisions/downloads/316478-001.pdf
- [14] Windeck C. (2007) Energy Star 4.0, *C't German Magazine for Computer Techniques*, **14**, 52-53.

- [15] Hewlett-Packard/Intel/Microsoft/Phoenix/Toshiba (2004) *Advanced Configuration and Power Interface Specification*, <http://www.acpi.info/DOWNLOADS/ACPIspec30.pdf>
- [16] Changjiu, X., Yung-Hsiang, L., and Zhiyuan, L. (2007) Energy-aware scheduling for real-time multiprocessor systems with uncertain task execution time, *Proceedings of the 44th annual conference on Design automation, San Diego, CA, USA, June*, pp. 664 – 669, ACM
- [17] Merkel, A. and Bellosa, F. (2008) Memory-aware Scheduling for Energy Efficiency on Multicore Processors, *Proceedings of the Workshop on Power Aware Computing and Systems (HotPower'08), San Diego, CA, USA, December*, pp. 123-130, USENIX online
- [18] AlEnawy, A.T. and Aydin, H. (2005) Energy-Aware Task Allocation for Rate Monotonic Scheduling, *11th IEEE Real Time and Embedded Technology and Applications Symposium (RTAS'05), San Francisco, CA, USA, March*, pp.213-223, IEEE
- [19] Cong, L., Xiao, Q., Kulkarni, S., Chengjun, W., Shuang, L., Manzanares, A. and Baskiyar, S. (2008) Distributed Energy-Efficient Scheduling for Data-Intensive Applications with Deadline Constraints on Data Grids, *Proceedings of the IEEE International Conference on Performance, Computing and Communications, IPCCC 2008, Austin, Texas, USA, December*, pp. 26-33, IEEE.
- [20] Chunlin, L., Layuan, L., (2008) Utility-based scheduling for grid computing under constraints of energy budget and deadline, *Journal of Computer Standards & Interfaces*.
- [21] Koomey, J. (2006) Server energy measurement protocol, Version 1.0, *Following Energy Efficiency Server Benchmark Technical Workshop*, Santa Clara, CA, <http://www.energystar.gov/ia/products/downloads/Finalserverenergyprotocol-v1.pdf>
- [22] *IBM Project Big Green*, <http://www-03.ibm.com/press/us/en/presskit/21440.wss>
- [23] Berl, A. and de Meer, H. (2009) An Energy-Efficient Distributed Office Environment. *Proceedings of European Conference on Universal Multiservice Networks (ECUMN 2009)*, Sliema, Malta, October, IEEE Computer Society Press.
- [24] Berl, A., Hlavacs, H., Weidlich, R., Schrank, M., and de Meer, H. (2009) Network Virtualization in Future Home Environments. *LNCS: Proceedings of Int. Workshop on Distributed Systems: Operations and Management (DSOM09)*, Venice, Italy, October, Springer, Berlin, Germany.

- [25] Hlavacs, H., Hummel, K. A., Weidlich, R., Houyou, A., Berl, A. and de Meer, H. (2009) Distributed Energy Efficiency in Future Home Environments. *Annals of Telecommunication: Next Generation Network and Service Management*, **63**, pp. 473-485, Springer, Paris, France.
- [26] IBM (2007) *Virtualization can help power efficiency*, <http://www-03.ibm.com/systems/virtualization/view/011607.html>
- [27] Bundesverband Informationswirtschaft (2008) *Energieeffizienz im Rechenzentrum*, 2, p. 10, Telekommunikation und neue Medien e.V.
- [28] Shah, R. and Kozyrakis, R. (2007) JouleSort: A Balanced Energy-Efficiency Benchmark. *ACM SIGMOD International Conference on Management of Data (SIGMOD)*
- [29] Siering P. (2006) Realitätsverschiebung – Virtualisierungstechniken im Vergleich, *Magazin für Computer Technik*, **13**.
- [30] Berl, A., Fischer, A., and de Meer, H. (2009) Using System Virtualization to Create Virtualized Networks. *Workshops der Wissenschaftlichen Konferenz Kommunikation in Verteilten Systemen (WowKiVS2009)*, Kassel, Germany, March, vol. 17, EASST.
- [31] Barham, P., Dragovic, B., Fraser, K., Hand, S. Harris, T., Ho, A., Neugebauer, R., Pratt, I., Warfield, A. (2003) Xen and the art of virtualization, *19th ACM Symposium on Operating Systems Principles (SOSP'03)*, New York, October, pp. 164-177, ACM Press, New York
- [32] Pratt, I. (2005) Xen 3.0 and the art of virtualization, *XEN - Computer Laboratory Architecture*, University of Cambridge, <http://www.cl.cam.ac.uk/netos/papers/2005-xen-ols.ppt>
- [33] *Linux-VServer*, <http://linux-vserver.org>
- [34] Des Ligneris, B. (2005) Virtualization of Linux based computers: the Linux-VServer project, *19th International Symposium on High Performance Computing Systems and Applications (HPCS 2005)*, Guelph, Ontario, Canada, May, pp. 340-346, IEEE.
- [35] *Brochure of VMWare Infrastructure 3*, VMWare, http://www.vmware.com/pdf/vi_brochure.pdf
- [36] Economou, D., Rivoire, S., Kozyrakis, C., and Ranganathan, P. (2006) Full-system power analysis and modeling for server environments, *Workshop on Modeling, Benchmarking, and Simulation (MoBS)*, June, Boston, USA.

- [37] Chase, J. S., Anderson, D. C., Thakar, P. N., Vahdat, A. M. and Doyle, R. P. (2001) Managing energy and server resources in hosting centres, *Proceedings of 18th ACM Symposium on Operating System Principles*, Banff, Canada, October, pp. 103 – 116, ACM press.
- [38] Chen, G., He, W., Liu, J., Nath, S., Rigas, L., Xiao, L. and Zhao, F. (2008) Energy-aware server provisioning and load dispatching for connection-intensive internet services, *NSDI'08: Proc. 5th USENIX Symposium on Networked Systems Design and Implementation*, San Francisco, CA, USA, April, pp. 337 – 350.
- [39] Urgaonkar, B., Shenoy, P., Chandra, A., Goyal, P. and Wood, T. (2008) Agile dynamic provisioning of multi-tier Internet applications, *ACM Trans. Autonomic Adaptive Systems*, **3**, no. 1.
- [40] Gelenbe, E. (2009) Steps towards self-aware networks, *Communications of the ACM*, **52**, no. 7, July 2009, pp. 66-75.
- [41] Sarokin, D. (2007) Question: Energy use of Internet, <http://uclue.com/?xq=724>.
- [42] Jia, X., Li, D. and Du, D. (2004) QoS topology control in ad hoc wireless networks, *Proc. of IEEE INFOCOM '04*, Hong Kong, China, March, IEEE.
- [43] Boukerche, A., Cheng, X. and Linus, J. (2005) A performance evaluation of a novel energy-aware data-centric routing algorithm in wireless sensor networks, *Wireless Networks*, **11**, no. 5.
- [44] Xu, Y., Heidemann, J. and Estrin, D. (2000) *Adaptive energy-conserving routing for multihop ad hoc networks - Research Report 527*, USC/Information Sciences Institute, <http://www.isi.edu/johnh/PAPERS/Xu00a.html>
- [45] Rajaraman, R. (2002) Topology control and routing in ad hoc networks: a survey, *SIGACT News*, **33**, no. 2, pp. 60 – 73.
- [46] Santi, P. (2005) Topology control in wireless ad hoc and sensor networks, *ACM Comput. Surv.*, **37**, no. 2, pp. 164 – 194.
- [47] Gelenbe, E. and Lent, R. (2004) Power-aware ad hoc cognitive packet networks, *Ad Hoc Networks*, **2**, no. 3, pp. 205-216.

- [48] Gelenbe, E., Lent, R. and Nunez, A. (2004) Self-aware networks and QoS, *Proc. of the IEEE*, **92**, no. 9, September 2004, pp. 1478 – 1489.
- [49] Dobson, S., Denazis, S. G., Fernández, A., Gäiti, D., Gelenbe, E., Massacci, F., Nixon, P., Saffre, F., Schmidt, N. and Zambonelli, F. (2006) A survey of autonomic communications, *ACM Transactions on Adaptive and Autonomous Systems*, **1**, Issue 2, pp. 223 – 259.
- [50] Gelenbe, E. (2006) Users and services in intelligent networks, *Proceedings IEE (ITS)*, **153**, no. 2, pp. 213-220.
- [51] Gelenbe, E. and Silvestri S. (2009) Reducing power consumption in wired networks, to appear in *Proc. 24th Annual International Symposium on Computer and Information Sciences (ISCIS 09)*, Cyprus, September 14-16, 2009.