Recent Trends in Energy Efficient Cloud Computing

Toni Mastelić, and Ivona Brandić
E-Commerce Group, Institute of Software Technology and Interactive Systems
Vienna University of Technology
Favoritenstrasse 9-11/188, A-1040 Vienna, Austria
{toni, ivona}@ec.tuwien.ac.at

Abstract

Cloud Computing represents the most promising Information and Communication Technology (ICT) paradigm that is directly or indirectly used by almost every online user. Its ultrascale size requires large data centers comprising several thousands of servers and other supporting equipment. The power consumption share of such infrastructures reaches 1.1% to 1.5% of the total electricity use worldwide, and is projected to rise even more. In this paper, we describe recent trends in Cloud Computing regarding the energy efficiency of its supporting infrastructure. We present the state of the art approaches found in literature and in practice covering servers, networking, Cloud management systems and appliances, which consists of a software utilized by end users. We describe benefits and trade-offs when applying energy efficiency techniques, and finally we extract existing challenges and highlight future research directions.

Index Terms

cloud computing, energy efficiency, data center, network, servers, cloud management system, appliances

I. Introduction

Today, a single data center can host up to several thousands of server units spread across over 4500 square meters. A typical 500 square meter data center can reach power consumption of 38 megawatt-hours (MWh) per day, which is more than the power consumption of over 3500 households in EU [1]. Only between year 2000 and 2007 the total power consumption of data centers worldwide went from 70 billion kWh to 330 billion kWh, and is projected to grow over 1000 billion kWh until 2020 [2]. Furthermore, in 2014 merely 8.5% data center managers estimated that their capacity will be sufficient beyond 2015, while over 75% will have to expand after 2020 [3]. Consequently, projection for data center construction size until 2020 is almost twice than it was in 2010, reaching $78 billion, which stresses the importance of dealing with the energy efficiency and the environmental impact of the Cloud Computing.

A. Cloud Computing

Cloud Computing services are powered by large data centers comprising numerous virtualized server instances, high-bandwidth networks, as well as supporting systems such as cooling and power supply. The equipment can be classified into **hardware** and **software**, which are accessed by remote **users** as shown in Figure 1. **Users** access Cloud services through *network* equipment that connects *servers* to the Internet, both part of the **hardware** equipment. User's **software**, referred to as an *appliance*, runs on top of servers and is managed by *Cloud management system* (CMS). Other supporting equipment includes power supply, cooling, as well as a data center building itself, which are out of the scope of this paper.

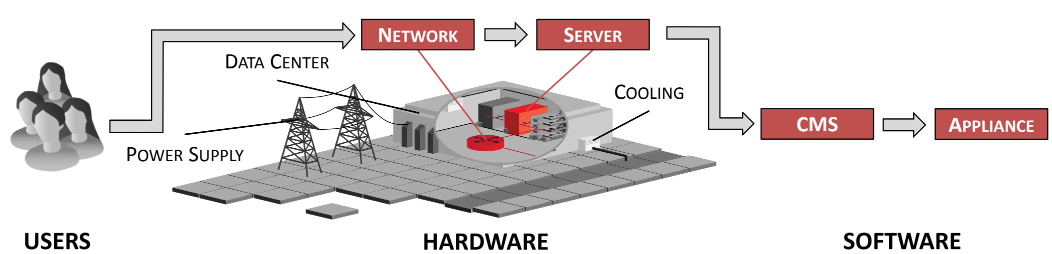

Fig. 1. Cloud Computing data center domains.

B. Energy efficiency

Energy efficiency can be defined as a reduction of energy used for a given service or level of activity [4]. However, due to scale and complexity of data center equipment it is extremely difficult to define unique service or activity that could be examined for its energy efficiency. Therefore, we identify four scenarios within a system where energy is not used in efficient way, but instead it is **lost** or **wasted**, as shown in Figure 2.

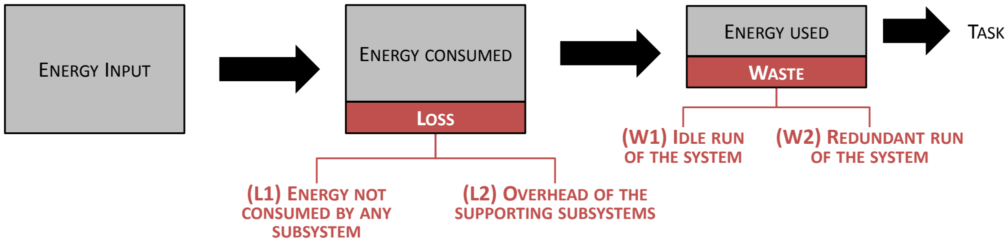

Fig. 2. Scenarios where energy is *lost* or *wasted* within a system.

Both terms define inefficient energy usage from an agnostic point of view, where **energy loss** refers to an energy brought to the system, but (L1) *not consumed by any of its subsystems*, e.g., energy lost due to transport or conversion. This also includes (L2) *energy overhead of the supporting subsystems*, such as cooling or lighting within a data center, where provisioning of Cloud services is considered as its main purpose. **Energy waste** refers to an energy used for its main purpose, however for the (W2) *idle run of the system*, e.g., processor being turned on but running idle. Additionally, (W2) *redundant run of the system* is also considered as energy waste, e.g., keeping a cooling system at maximum during night when the temperatures are lower.

II. Energy efficiency of Cloud infrastructure

In this section we put data center equipment in the context of energy efficiency by describing existing problems. Further, we describe state of the art approaches used for tackling these problems. Finally, we highlight future challenges that are yet to be solved.

A. Network

1) Context: Network serves as a bridge between end users and Cloud resources such as computation and storage. Its energy consumption is distributed over three main systems, namely the connections inside a data center, the network between data centers, and the outside fixed and wireless network that allows end users to access services via mobile devices. Within a data center, the network currently accounts for up to 10% of the operating expenses that include power consumption, and is estimated to rise to 50% due to increased Internet traffic [5]. This is due to a large number of network units, and their unproportional power usage while running idle. Moreover, poor network architectures not suited for Cloud applications can increase energy waste by unnecessary re-routing the traffic or keeping some parts of the network underutilized.

2) State of the art: First step of improving the energy efficiency is to upgrade network equipment by implementing power saving modes and adaptive transmission rates, in order to achieve proportional power usage. This can additionally reduce a head load of the equipment and hence reduce the energy consumption of the supporting cooling system, as well as improve its reliability. In case of zero utilization, unused ports, links and switches can be switched off completely. However, such approaches usually lead to performance degradation due to reduced connectivity. Tweaking the network equipment is usually insufficient, as even the most efficient equipment can waste energy due to poor network topologies. Flexible topologies that allow dynamic routing paths require fewer components, while still exhibiting high resilience to failures. Consequently, the network traffic can be spread amongst multiple paths, or sent through few energy-critical paths, while letting other parts of the network to enter low-power modes. Such a scheme is used for inter data center networks as well, by exploiting uninitialized network bandwidth for bulk transfers, which create less communication overhead compared to the separate network packages.

End user network, today typically a wireless network utilizes macrocells, which require high amount of transmit power due to relatively unfocussed transmissions, where a distance between the base station and mobile devices is usually large. On one hand, using multiple directed antennas at the base station can focus the energy to the user and hence reduce the interference. On the other hand, the distance can be tackled with small cells, which reduce path-loss by several orders of magnitude, as they can be placed closer to the end users. Both approaches exhibit lower power consumption. Applying small cells in urban areas can reduce power consumption by a factor of 46 when efficient idle modes are used [6]. Such modes provide proportional power consumption, as most of the components can be switched off while not serving active connections. Finally, unlike highly inefficient power amplifiers implemented in macrocellular networks that consume large amount of power even during low load, small cells scale far more easier due to their superiority in numbers and reduced coverage.

3) Future challenges: For wired networks, it is believed that a complete energy saving solution requires optimization strategies both on local and wide area scales, which work in conjunction between each other. On one hand, designing and producing efficient network equipment with multiple energy saving modes and adaptive transmission rate, while preserving quality of service and resilience, is still an open question. On the other hand, scalability of energy-efficient routing protocols and their implementation in real networks is yet to be solved. For wireless networks, heterogeneous networks that comprise not only different equipment, but different technologies as well, is an active area of research both for academia and industry. Such integration includes combining small cells and macrocellular networks, which requires improving idle mode control. Furthermore, due to significantly slower hardware, processing on small cells represents limiting factor. Therefore, low complexity algorithms, lightweight implementations, and low power processors represent one of the research goals for energy reductions in network systems.

B. Servers

1) Context: Servers include computing and storage servers placed inside enclosures/cabinets, as well as its components such as processors, memory, disks, etc., excluding the communication equipment that is part of the network. The power density of a single rack of servers started in range between 250 W to 1.5 kW in 2003. In 2014 it reached almost 10 kW and is projected to rise up to 30 kW until 2020 [7]. Moreover, most servers consume over 50% of their peek power consumption while running idle, while their average utilization is typically as low as 10-50% [8]. Consequently, an infrastructure operating at only 20% capacity may consume 80% of the energy as the same infrastructure operating at 100% capacity [1]. Such numbers become alarming when considering that only in the last quarter of 2013 over 2.5 million new servers where shipped [9]. Therefore, improving energy efficiency of the servers represents a top priority task.

2) State of the art: Similar to network equipment, servers can also benefit from power proportional components, where power and performance is automatically scaled based on a current load. Most prominent technology includes Dynamic Voltage and Frequency Scaling (DVFS) found on today's CPUs, which provide P-states (power modes while being utilized), as well as sleep C-states (power modes while being idle) [10]. The same applies for other components such as memory and storage disks, which can be put to a low power state while being idle. In order to increase idle time of disks, modern hardware utilizes onboard cache controllers, where reads and writes are performed prior from/on cache. If data cannot be found in cache, only then the disk is accessed. Furthermore, cache itself can be optimized by dynamically reconfiguring its size and associativity, or even turning off parts of cache that are not used, thus reducing static power consumption (also known as leakage power). Finally, new low-power cache technologies can be integrated along with the existing ones, thus reducing the energy consumption while maintaining the performance.

Selection of server enclosure design effects the efficiency of cooling and power supply, and consequently the server rack itself. Researchers have shown that using enclosures with localized cooling applied on the right component can provide some worthwhile improvements, as compared to room air conditioners. Example is Embedded thermoelectric cooling (eTEC) or vapour compression refrigeration system that removes hot spots generated by modern processors. Further improvements can be achieved by tweaking the components themselves, e.g., adjusting the memory throughput is used for mitigating high temperatures, and thus avoiding energy loss through heat. Limiting the energy input through power supply, also referred to as power capping, can decrease power consumption of a single component or an entire rack. Finally, energy loss can also be reduced by using a compact server configuration which simply excludes components that are not used.

3) Future challenges: Implementing energy efficient features in real world deployments may not be as straightforward as it seems. For instance, low power modes for server components are proven beneficial only for long idle periods, which may not be that common in production environments. Furthermore, regardless of server's low utilization during execution of lightweight tasks, it still requires promptness due to elasticity demands during peek loads. Therefore, self-scalability of server components must be achieved in correlation with software components. Example is false application of DVFS technique, where the execution time of an application is prolonged due to lower cpu frequency, which eventually increases overall energy consumption. Another goal is achieving a near zero power consumption of the idle components, which can perhaps be achieved by consolidating underutilized instances. Finally, combining low power components when demand is low with traditional components when high performance is required can also be an option for energy scalable infrastructure.

C. Cloud Management System (CMS)

1) Context: Due to increasing energy consumption of ICT, monitoring and management of Cloud supporting infrastructures is recognized as the main concern within a data center facilities [3]. Therefore, the main role is assigned to the Cloud management system (CMS) in order to improve efficiency, increase utilization, and thus lower the energy loss/waste within a data center. The CMS includes scheduler, monitoring system and virtualization technology, which allows running multiple virtual machines (VMs) on top of a single physical machine, and thus increases utilization. However, a trade off of the latter is an energy consumption overhead as it adds additional software layer, namely a hypervisor. Similarly, a monitoring system creates an overhead as well, through monitoring agents and probes while collecting monitoring information required for managing the infrastructure.

2) State of the art: The main purpose of the CMS is scheduling and load balancing the underlying infrastructure, more specifically servers, VMs and applications [10]. In the context of VM management, as the most prominent example of Cloud services, several actions can be taken for improving energy efficiency. First includes a VM self-adaption and hardware adjustment, where a VM reconfigures its resources based on the current load. Additionally, using more lightweight virtualization technologies such as Linux containers is shown beneficial in many cases. Selecting most efficient physical machines for VM placements is considered as the second optimization action. Final action includes VM migrations, which allows dynamic consolidation of physical machines by moving underutilized VMs on fewer hosts and powering of unused ones. Such rescheduling of VMs over time enables reduction in resource consumption at any given moment.

Above steps require implementation of software CPU power modes, where hardware and software is coordinated in order to provide DVFS in a virtualized environment as well. Furthermore, the physical machines themselves can be adjusted to their actual load so as to reduce their power consumption. Researchers show that even by using very simple heuristics, e.g., shutting down a server that has been idle for the last minutes, can save a significant amount of energy. Further step is going beyond a single cluster or a data center and applying management actions, such as consolidation over a federation of data centers. However, this increases the overhead of VM migrations, as they are recognized as an I/O intensive actions, which consume energy on both ends. Simulation tools such as CloudSim are used for estimating such trade-offs and benefits of CMS management actions, hence comparing the kind of constraints the actions put on the quality of service [11].

The CMS can go beyond controlling only the servers by expanding its control to the network system, or even cooling and power supply. Such approach is feasible since the CMS has a "knowledge" of which resources are required and what not, e.g., shutting down cooling equipment while servers are idle. Further optimization can be done through selection of appropriate virtualization technology with regards to the service requirements, namely full-virtualization vs. para-virtualization, or even micro-kernel architectures. Finally, taking a modular approach for building the CMS, where modules are loaded only when they are actually required, can result in lightweight systems, e.g., using a monitoring agent that loads plugins for monitoring specific components and removing them once they are no longer required.

3) Future challenges: The main challenges in CMS regarding the energy efficiency include power consumption monitoring of individual VMs. Current approaches use simplistic approach by considering only cpu utilization, while excluding resources such as network and storage that also exhibit high energy footprint. Due to different application profiles running inside the VMs that require different resources, an energy share of each VM or even individual applications must be mathematically and precisely modelled. Further challenges include investigating the interdependencies between possible leverages in the CMS. Overusing VM migrations in order to consolidate physical servers can be counter productive, as the energy overhead of optimization actions can exceed achieved savings. Similarly, deploying large monitoring systems and collecting enormous amount of monitoring data can lead to resource demands of supporting systems being larger than that of the main production environment.

D. Appliances

1) Context: Appliances represent operating systems, platforms and applications accessed and utilized by end-users. On one hand, efficiency of appliances is only considered for Software (SaaS) and Platform as a Service (PaaS), since an appliance is then under control of the provider and thus part of the Cloud Computing infrastructure. On the other hand, for lower level services such as Infrastructure as a Service, an appliance is deployed by an user, and therefore the user is responsible for its efficiency. Since applications run on top of operating systems and runtime environments such as Java Virtual Machine, latter ones usually create overhead in resource consumption and thus energy consumption. However, performance of the applications can have even greater impact on energy usage, as applications exhibiting poor performance require more instances and therefore more servers, which consume additional amount of energy.

2) State of the art: Recommendations and techniques for developing energy efficient software with focus on reducing power consumption of idle appliances are provided by Agrawal et.al [12]. Moreover, by limiting wake-up events and changing timer activities idle time can be increased, and hence power consumption reduced. For instance, using push instead of a pull mechanism allow the software to remain dormant until action is actually required. Additionally, avoiding memory leaks and removing unused features improves software efficiency in terms of ratio between resource consumption and the amount of useful work performed. Furthermore, improving compilers to include cache skipping, use of register operands, instruction clustering and re-ordering, and loop optimization can improve code execution. Finally, optimizing software with regards to hardware resources can significantly improve energy efficiency, e.g., implementing batch access to resources such as disk, hence decreasing number of unnecessary wake-ups.

Improving energy efficiency of appliances also includes monitoring them, which is a non-trivial problem. Similarly to VM monitoring, mapping between resource usage and energy usage of appliance is still limited to the CPU usage. More sophisticated approaches try to detect processes responsible for hardware wake-ups, which prevent hardware components from entering the sleep mode. Further studies have shown that a substantial amount of energy is being used while systems are only slightly utilized due to hardware's unproportional energy usage. Even when appliances are idle a resource consumption overhead of a monitoring system, virtualization or operating system still keep the hardware awake. Consequently, an event based monitoring interfaces, lightweight virtualization layer and better optimization of operating system schedulers is shown to provide lower resource usage. Substantial effort was invested into research on distributed Cloud operating systems such as CloudOS [13] that exploits distributed Cloud environment. Additionally, efficient usage and deployment of shared services is achieved through energy-aware software layers that utilize micro and macro metrics for achieving optimized performance. Nevertheless, their overhead and energy-efficiency characteristics should be studied in more detail.

3) Future challenges: The main challenges concerning energy efficiency of Cloud appliances is related to its development process, namely going from lightweight design, optimized implementation and smart deployment by selecting targeted hardware. Although, main focus of developers is on high performance, some of energy efficiency goals are partially consistent with performance goals, as scalability and short execution time typically lead to low energy usage. However, another great challenge is the runtime, where developers tend to load everything in memory and keep the system awake 24/7 in order to achieve promptness of appliances and support high demand peaks. Such policies rarely comply with energy efficient goals, while their are extremely common in production environments. Nevertheless, efficiency can still be improved through close integration of an operating system and applications by removing unnecessary features. This also includes adjusting the size of virtual machines or providing sand boxes or Linux containers for applications in order to reduce the overhead. Finally, running Cloud appliances in an energy-efficient way requires communication between appliances and other domains, especially the CMS to make optimal decisions.

III. Data center system integration

Data centers that power Cloud Computing represent a tightly coupled system that can almost be viewed as a single big computer [8]. Therefore, along with covering each system separately, the entire data center has to be analysed as a whole. Here we describe correlations between parts of the infrastructure and give overview on state of the art approaches for improving energy efficiency inside of a data center.

Appliances represent the smallest manageable unit in Cloud environment that is ultimately accessed and utilized by an end user. However, its performance can produce domino effect by influencing number of required servers, which further affects the scale of network and supporting equipment such as cooling and power supply. Therefore, properly selecting underlying hardware resources for targeted applications can play significant role in overall energy consumption of a data center. For instance, porting highly parallel applications to general purpose GPUs (GPGPU), or using microservers equipped with processors such as ARMs for lightweight applications. However, this requires further investigation and detailed classification of applications in order to be able to map certain applications to the underlying resources.

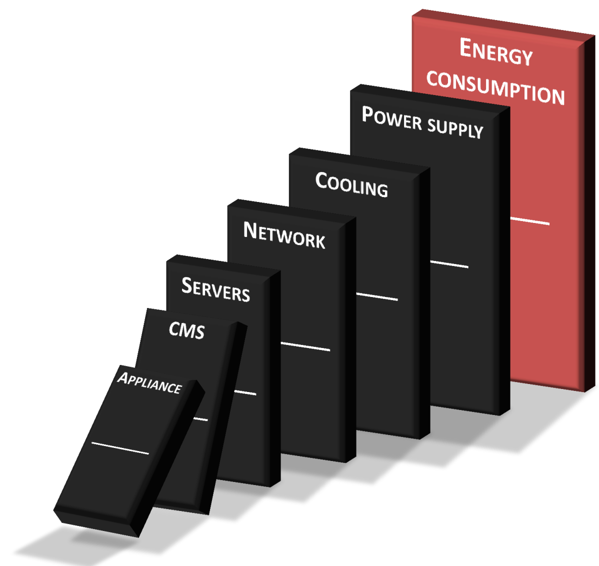

Fig. 3. Data center system integration domino effect

Management of such heterogeneous infrastructure requires flexible and comprehensive CMS with advanced monitoring capabilities. Integrating direct management of supporting systems such as cooling and power supply adds additional layers of complexity. However, even with ICT equipment management alone, CMS requires monitoring information on progress, performance, state, data size, as well as hardware metrics in order to support smart scheduling and hardware matching. Further improvements involve federation of Clouds through optimization actions over several geographically distributed data centers. Such approaches require not only interaction between software and hardware within a single data center, but an exchange of information, user load and data across the globe as well.

The main enabler of such distributed approaches is the network. Due to its relatively small power consumption compared to the rest of data center equipment, the network enables placement of computation and data in locations where renewable energy is available, or where less energy for cooling is required due to a cool climate. However, along with optimization of network equipment, improvements have to be done on other equipment and components as well, e.g., optimizations of appliances and CMS management actions in order to reduce network traffic. Such improvements are important in order to avoid performance trade-offs due to user load distribution, which might result in an increased latency if computations and data are placed far away from the user. Finally, overhead of CMS actions such as VM migrations can still have significant impact on energy efficiency and performance if not taken into an account.

IV. Using the Cloud in energy efficient manner

Along with improving the energy efficiency of data centers, the data centers themselves being wrapped up in the Cloud Computing concept can provide a leverage for improving the energy efficiency on a larger scale. Instead of using separate computing systems, users can consolidate their computing requirements and move their applications to the Cloud Computing environment, which is more efficient than the smaller computing systems due to its bigger scale. Energy efficiency models such as Cleer model [14] developed by the Lawrence Berkeley National Laboratory provides an open-access approach to analyse energy savings when moving to Cloud. It estimates potential of up to 95% of reduction in energy use compared to present day business software use, which includes customer relationship management (CRM), productivity and email software. Another point made here is that moving away from physical products to digital ones can also have significant energy saving impact. This includes video industry that by utilizing Cloud Computing concept can stream video over the internet, which reduces energy consumption by 15% compared to shipping it on CDs and DVDs [15].

Except using Cloud Computing concept for its main purpose, the Cloud Computing infrastructure and its flexible nature can also be utilized indirectly for energy optimization. Techniques such as DVFS and powering on/off machines can be used for frequency regulation of a power network. Using this approach, data center's dynamic load can be used for regulating electricity demand and thus production, which is important for keeping optimal frequency of the power grid. Moreover, electricity providers pay for such a service of dynamic load, enabling the data center to earn a half of million dollars annually. Additionally, instead of only consuming energy, modern data centers are becoming energy producers with on-site power generators. Not only that this approach reduces costs for the data center, it also reduces power losses due to an energy transfer and the load on an energy grid.

V. Conclusion

In this paper we analysed the energy efficiency of a data center equipment, including hardware and software that drives the Cloud Computing. We showed that state of the art approaches can cope with energy efficiency challenges. However their implementation in real world systems is typically slow and not always as straightforward as it seems. Furthermore, existing energy efficiency techniques have to be adapted in order to bring some benefits in a stratified Cloud Computing infrastructure. This includes minimizing trade-offs such as performance degradation and high costs. Finally, in order to achieve significant improvements, energy efficient solutions have to cover individual components, as well as their integration with the rest of the system.

References

[1] Emerson, "Energy logic 2.0: New strategies for cutting data center energy costs and boosting capacity," Emerson Network Power, Emerson Electric Co., Emerson Network Power, Ohio, USA, Tech. Rep., 2014.

[2] Greenpeace, "Make IT Green: Cloud computing and its contribution to climate change," Greenpeace International, Tech. Rep., March 2010.

[3] Emerson, "Data Center Users Group: Survey Results," Data Center Users Group, Emerson Electric Co., Emerson Network Power, Ohio, USA, Tech. Rep., October 2014.

[4] F. Moisan and D. Bosseboeuf, "Energy Efficiency: A Recipe for Success," World Energy Council, London, UK, Tech. Rep., 2010.

[5] D. Kliazovich, P. Bouvry, and S. Khan, "Simulation and performance analysis of data intensive and workload intensive cloud computing data centers," in *Optical Interconnects for Future Data Center Networks*, ser. Optical Networks, C. Kachris, K. Bergman, and I. Tomkos, Eds. Springer New York, 2013, pp. 47–63.

[6] R. Razavi and H. Claussen, "Urban small cell deployments: Impact on the network energy consumption," in *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, Paris, April 2012.

[7] Emerson, "Data Center 2025," Emerson Electric Co., Emerson Network Power, Ohio, USA, Tech. Rep., 2014.

[8] U. Hoelzle and L. A. Barroso, *The Datacenter As a Computer: An Introduction to the Design of Warehouse-Scale Machines*, 2nd ed. Morgan and Claypool Publishers, 2013.

[9] Gartner, "Servers Quarterly Statistics Worldwide: Database," Gartner Research, Stamford, USA, Tech. Rep., February 2014.

[10] G. Valentini, W. Lassonde, S. Khan, N. Min-Allah, S. Madani, J. Li, L. Zhang, L. Wang, N. Ghani, J. Kolodziej, H. Li, A. Zomaya, C.-Z. Xu, P. Balaji, A. Vishnu, F. Pinel, J. Pecero, D. Kliazovich, and P. Bouvry, "An overview of energy efficiency techniques in cluster computing systems," *Cluster Computing*, vol. 16, no. 1, pp. 3–15, 2013. [Online]. Available: http://dx.doi.org/10.1007/s10586-011-0171-x

[11] W. Long, L. Yuqing, and X. Qingxin, "Using cloudsim to model and simulate cloud computing environment," in *Computational Intelligence and Security (CIS), 2013 9th International Conference on*, Dec 2013, pp. 323–328.

[12] A. Agrawal and M. Sabharwal, "Enabling green it through energy aware software," *IT Professional*, vol. PP, no. 99, p. 1, 2012.

[13] F. Pianese, P. Bosch, A. Duminuco, N. Janssens, T. Stathopoulos, and M. Steiner, "Toward a cloud operating system," in *Network Operations and Management Symposium Workshops (NOMS Wksps), 2010 IEEE/IFIP*, 2010, pp. 335–342.

[14] Berkeley, "The energy efficiency potential of cloud-based software," http://crd.lbl.gov/assets/pubs_presos/ACS/cloud_efficiency_study.pdf, 2013.

[15] A. Shehabi, B. Walker, and E. R. Masanet, "Energy and greenhouse implications of internet video streaming in the united states," *Environmental Research Letters*, vol. 9, 05/2014 2014.

Toni Mastelic did his bachelor and masters studies in Computer Science at the University of Split, Faculty of Electrical Engineering, Mechanical Engineering and Naval Architecture (FESB), Croatia, where he received his Bachelors degree in 2009, and Masters degree in 2011. His main interests are Cloud infrastructures, specifically related to energy efficiency. He works on Cloud monitoring by developing generic and comprehensive monitoring system, as well as introducing new performance, resource, energy efficiency and security metrics. Currently, he is pursuing his PhD and also working as a research assistant at the E-Commerce Group, Institute of Software Technology and Interactive Systems, Vienna University of Technology, Vienna, Austria. For more information, please visit http://www.infosys.tuwien.ac.at/staff/toni/

Ivona Brandic is Assistant Professor at the E-Commerce Group, Institute of Software Technology and Interactive Systems, Vienna University of Technology (TU Wien). She received her PhD degree in 2007 and her venia docendi for practical computer science in 2013 - both from Vienna University of Technology. From 2003 to 2007 she participated in the special research project AURORA - Advanced Models, Applications and Software Systems for High Performance Computing and the European Union's GEMSS - Grid-Enabled Medical Simulation Services project. She was involved in the European Union's SCube project and she led the Austrian national FoSII - Foundations of Self-governing ICT Infrastructures project funded by the Vienna Science and Technology Fund (WWTF). She was the management committee member of the European Commission's COST Action on Energy Efficient Large Scale Distributed Systems. From June to August 2008 she was visiting researcher at the University of Melbourne, Australia. Dr. Brandic is on the Editorial Board of the IEEE Transactions on Cloud Computing and the Springer Journal on Scalable Computing. In 2011 she received the Distinguished Young Scientist Award from the Vienna University of Technology for her HALEY project on Holistic Energy Efficient Hybrid Clouds. Her interests comprise Service Level Agreement and Quality of Service management in large scale distributed systems, autonomic computing, workflow management for scientific applications, and energy efficient large scale distributed systems (Cloud, Grid, Cluster, etc.). I. Brandic has been invited expert evaluator of the European Commission, French National Research Organization (ANR), National Science and Engineering Research Council Canada (NSERC) and Netherlands Organization for Scientific Research (NWO).