ANALYSIS OF SIMULATION TOOLS FOR DETERMINING THE ENERGY CONSUMPTION OF DATA CENTERS FOR CLOUD COMPUTING

Matthias Splieth[(a)], Sascha Bosse[(b)], Klaus Turowski[(c)]

Otto von Guericke University Magdeburg
Faculty of Computer Science, Very Large Business Applications Lab
P.O. Box 4120, 39016 Magdeburg

[(a)]matthias.splieth@ovgu.de, [(b)]sascha.bosse@ovgu.de, [(c)]klaus.turowski@ovgu.de

ABSTRACT
Cloud Computing has become an important paradigm for providing IT services. An important aspect that data centers which offer cloud services need to deal with is the reduction of their energy consumption. A suitable way for reducing the energy consumption and thus the energy costs is using energy efficient load distribution algorithms. These are commonly evaluated by simulation. Up to now, there are only few simulation tools for clouds that are freely disposable and those that do exist have hardly been compared in terms of their applicability for determining the energy consumption of clouds. Therefore, this contribution presents an analysis of simulators that investigates their suitability for determining the energy consumption of cloud data centers. The findings of the analysis will be supported by a unified experiment. The results demonstrate that existing work has limitations, for example in terms of components that are not considered to be energy-consuming.

Keywords: cloud computing, energy consumption, simulation, survey

1. INTRODUCTION

In recent years, Cloud Computing has been established as a new paradigm for the dynamic provisioning of IT services. In spite of the advantages of cloud services for both customers and providers, there are also some challenges that need to be addressed. One of the most important challenges for providers is the reduction of the overall energy consumption that continuously increases due to the rapidly expanding demand for computational power (Zhang, Cheng and Boutaba 2010). Current studies have shown that the energy consumption of data centers increased by 56% from 2005 to 2010 (Koomey 2011). Taking into account that prices and demand for energy will continue to rise and therefore become the dominant factor in the total cost of ownership (Orgerie, De Assuncao and Lefevre 2014), it is desirable for operators of (cloud) data centers to even slightly decrease the energy consumption since this can have significant impacts on their profitability. A possible way to save energy is load distribution (Orgerie, De Assuncao and Lefevre 2014) which is also one of the major challenges regarding cloud computing (Zhang, Cheng and Boutaba 2010). Besides the fact that load distribution in terms of placing or migrating virtual machines (referred to as virtual machine scheduling) is an NP-hard problem (Zhang, Cheng and Boutaba 2010), additional aspects such as the free scalability of the resources according to the current workloads aggravate the task of distributing load in cloud data centers.

In recent years, numerous algorithms for load distribution in clouds have been developed, many of them with the goal to increase energy efficiency. However, their effects on energy consumption in a specific scenario cannot be predicted trivially (Nehru et al. 2013). This is because the possible improvements of an algorithm strongly depend on the parameters of the specific data center, such as the offered services, the configuration of servers, the architecture of the network, the network topology and others. Hence, applying the same algorithm in different scenarios can lead to varying results in terms of energy consumption (Liu et al. 2013). Therefore, there is a need for a solution that is able to exactly predict the energy consumption of cloud data centers and hence to also evaluate the effects of load distribution algorithms on the energy consumption of cloud data centers.

In addition to testbeds and mathematical models, simulation is often used in order to investigate clouds (Sakellari and Loukas 2013). Since tools for the simulation of grids (or other similar paradigms) are not applicable to clouds (Buyya, Ranjan and Calheiros 2009), the number of simulation tools that are available for public use is limited. The existing simulators differ in several aspects, for example in terms of components of a data center that can be modeled. As an example, cooling components (which are an important energy consumer) cannot be modeled in *CloudSim* (Calheiros et al. 2011) but in *CReST* (Cartlidge and Cliff 2013).

In order to be able to investigate the effects if load distribution algorithms on the energy consumption of data centers, a simulator is needed that is able to exactly predict the energy consumption of clouds. Therefore, a survey on and an analysis of simulation tools for cloud data centers is conducted in this paper. It aims to examine their suitability in terms of determining energy consumption in clouds. The survey provides a decision basis for choosing simulation tools for energy efficient cloud computing. Furthermore, a detailed comparison of existing simulation tools can reveal existing flaws that may be addressed in future research. Finally, a unified experiment is defined and implemented with each simulator. The respective results of each simulator will be compared in order to verify the findings of the analysis of the single simulators.

2. RELATED WORK

Cloud Computing is a hot topic and is thus subject of recent research. Since the research area is emergent, research is conducted in many directions. One of the most important areas of this research is the energy efficiency of clouds (Zhang, Cheng and Boutaba 2010). For the investigation of energy aspects, simulation is often used. But until today, an analysis and comparison of cloud computing simulators has only been carried out on a simple level focusing on general aspects of the simulators.

For example, Zhao et al. provide a brief overview of existing cloud simulators without analyzing each simulator in depth (Zhao et al. 2012). In fact, the analysis is limited to only three criteria (underlying platform, programming language and whether the simulator is either based on software or based on software and hardware), while aspects related to energy efficiency are not considered.

The analysis of Sakellari and Loukas presented in (Sakellari and Loukas 2013) is more detailed and also takes energy efficiency into account. However, the analysis in terms of energy is mainly limited to the question of whether energy aspects can be taken into account at all. Furthermore, general aspects such as the programming language are investigated. Details on energy consumption, such as the question which energy-consuming components can be modeled with the simulator are not part of the survey.

The same applies to (Kumar and Rai 2014), wherein the analysis is less detailed as in (Sakellari and Loukas 2013) since this work does not mention how energy efficiency is meant to be considered.

A recent survey on cloud simulators is presented by Ahmed and Sabyasachi in (Ahmed and Sabyasachi 2014). In the analysis of the simulators, their features as well as their respective shortcomings are discussed. The analysis captures, among other things, whether energy models are part of the simulator. However, the analysis does not cover which energy models are included for which components. In addition, the final comparison of the tool is only based on the attributes of the analysis. A detailed comparison of the capabilities of simulation tools in terms of energy aspects is not carried out.

3. ANALYSIS OF CLOUD SIMULATORS

Before starting the analysis, suitable simulators must be identified. Therefore, scientific literature is reviewed towards simulation tools for cloud computing. In order to identify relevant literature, *ACM Digital Library*, *IEEE Xplore Digital Library*, *Springer Link* and *Google Scholar* have been queried for cloud simulators using the search term "*cloud computing & simulat(e|or|ion|ing)*". Emulation environments, test beds and mathematical-analytical approaches are excluded from the survey. Mathematical-analytical approaches are not included since they only depict parts of clouds and are not suitable for modeling complex communication (Sakellari and Loukas 2013). The problem with testbeds or other real-world environments is that experiments are expensive, time-consuming and not repeatable (Zhao et al. 2012). In this paper, the focus is on tools that have been explicitly designed for the simulation of cloud computing systems. Therefore, simulation tools or frameworks that address data centers in general or areas similar to cloud computing (such as grid computing) are not included since such tools are unsuitable for simulating clouds (Buyya, Ranjan and Calheiros 2009). In order to be able to analyze the tools in depth (in terms of their source code) and in order to be able to carry out the unified experiment with each tool, simulators are only analyzed if they are available to public and open source. In order to keep the effort for the experiments low, other tools that do not provide crucial changes to the original are excluded from the analysis.

In total, the search term defined before resulted in 31 different simulators that have been mentioned in scientific literature. Most of them are proprietary implementations and not available for public. On the basis of the aforementioned aspects, four simulators have been identified that are open source and also disposable:

- *CloudSim* (Calheiros et al. 2011), version 3.0.3
- *GreenCloud* (Kliazovich, Bouvry and Khan 2012), version 2.0.2
- *iCanCloud* (Núñez et al. 2012), version 0.9
- *CReST* (Cartlidge and Cliff 2013), version 0.5

These simulators are analyzed in depth in the following.

3.1. Criteria for Analysis

The analysis of the aforementioned simulators consists of three parts: first, general aspects such as the evaluation method or the applied programming language are investigated. In the second part, the simulation tools are examined in terms of their suitability for modeling physical components of a cloud data center. In the last part of the analysis, the simulators are analyzed with respect to issues that are important for simulating load distribution and its effects on energy consumption but not directly consume energy, such as energy models that define how energy is consumed by the components of a data center.

3.1.1. General Aspects

The general criteria are loosely based on (Smit and Stroulia 2013) and are summarized and described in table 1.

Table 1: General Aspects of the Analysis

Criterion	Description
Objective	What was the simulator designed for?
Evaluation	How is the simulator evaluated?
Type of Simulation	Of which type is the simulation? (discrete-event, continuous, …?)
License	License under which the simulator was published
Language	Language the simulator is written in
QoS	Is it possible to measure effects on quality of service?
Services	Supported cloud-layer
Basis	Implementation basis of the simulator

"Objective" captures the intention of the creators, such as estimating the energy-efficiency of clouds or determining their availability. The criterion "Evaluation" describes the evaluation method that has been applied in order to verify the simulator. This can for example refer to evaluation methods such as proof-of-concept or case-study. "Type of Simulation" covers the type of the simulation, such as continuous simulation or discrete-event simulation. The license under which the simulations is released is covered by the criterion "License". This is important due to the possibility for e.g. making changes to the source code. "Language" refers to the programming language that is used for the implementation of the simulator.

"QoS" depicts whether it is possible to investigate indicators that are important for the quality of service, such as availability or response time. The criterion "Services" refers to the cloud-layer supported by the simulator as defined in (Mell and Grance 2011). "Basis" indicates the basis on which the simulator has been implemented, as far as it is based on another tool or framework.

3.1.2. Physical Components

With respect to more advanced criteria, it is necessary that all physical components of a cloud data center that consume energy can be modeled. A rough overview of the relevant components (confer for example (Brown et al. 2007; Jing et al. 2013)) is given in table 2.

For the analysis, these criteria are even more detailed as depicted in table 2. The criterion "Server" represents a physical server that has various energy-consuming components such as CPU, memory, motherboard, peripheral slots, fans and maybe a GPU (Fan, Barroso and Weber 2007; Greenberg et al. 2008).

Table 2: Physical Components of the Analysis

Criterion	Description
Server	Is it possible to model servers in detail?
Cooling	Can cooling components be modeled?
Storage Systems	Is it possible to model storage systems?
Support Systems	Is it possible to model support systems such as UPS or lightning?
Network	How is the network simulated? Can devices and topologies be modeled?

Furthermore, this criterion also includes racks as well as servers in terms of blades (Barroso and Hölzle 2009). This criterion is fulfilled if all components can be modeled; otherwise it is partly fulfilled or not fulfilled (if none of the mentioned components of a server can be modeled).

"Cooling" refers to the ability of a simulation tool to model cooling units that are common in a data center. This can for example refer to CRAC units ("computer room air conditioning"), free cooling methods or in-rack cooling as described in (Barroso and Hölzle 2009). This criterion is fulfilled if at least one cooling system can be modeled.

"Storage Systems" depicts if it is possible to model central storage system since these can be important regarding the way the migration of virtual machines (VMs) is performed (Mishra et al. 2012). This criterion is fulfilled if storage systems can be modeled. The criterion "Support Systems" states whether support systems can be modeled. This refers to uninterruptible power supplies (USP) and lighting and is fulfilled if both of these can be modeled with the simulator.

The criterion "Network" addresses several aspects. The first aspect is concerned with the question how the network is simulated. This can be done using a flow model or a packet model, while the latter is more accurate (Velho et al. 2013). Further aspects concerning this criterion are dealing with the question whether network topologies can be modeled (since there are different types of topologies which have different side-effects (Barroso and Hölzle 2009)) and whether network devices as energy-consuming components are included in the simulator.

3.1.3. Non-Physical Components

In addition to the physical components, there are also non-physical aspects which are relevant for investigating load distribution and energy consumption. An overview of all the relevant non-physical aspects is given in table 3.

In order to be able to evaluate the effects of load distribution algorithms on energy consumption, certainly "Load" has to be considered in a simulator. Load distribution in clouds can either refer to virtual machine scheduling or to the distribution of the workload (Beloglazov, Abawajy and Buyya 2012).

Table 3: Non-Physical Components of the Analysis

Criterion	Description
Load	Can workload and live migration be simulated?
Virtualization / Cloud	Are virtualization and rapid elasticity part of the simulator?
Software	Can software be modeled?
Energy Models	Are energy models provided for each component?
Power Saving Techniques	Are power saving methods a component of the simulator?

VM scheduling refers to the placement or the migration of virtual machines. The distribution of workload means that tasks that are generated by user requests are assigned to applications that run in virtual machines. In order to fulfill this criterion, both aspects must be considered by the respective simulators.

Since this research focuses on the simulation of clouds, it is obvious that the relevant aspects regarding cloud computing need to be taken into account. This is covered by the criterion "Virtualization / Cloud". First, the concepts of virtualization – which is a key technology of cloud computing (Zhang, Cheng and Boutaba 2010; Mishra et al. 2012) – need to be part of the simulator. Furthermore, key aspects such as rapid elasticity and resource pooling (Mell and Grance 2011) must be considered. Resource pooling refers to the possibility to model multi-tenancy aspects with the simulator, which means that it is possible to model software that is for example capable to serve different customers independently from one another. An example for single-tenancy and multi-tenancy is given in figure 1. Rapid elasticity means that provided computing resources are scaled to the actual demand of users. If all these aspects are considered, the criterion is fulfilled.

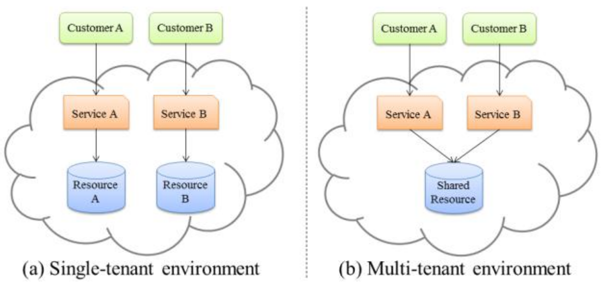

Figure 1: Example for (a) Single-Tenancy and (b) Multi-Tenancy in Clouds

An important aspect are software components in the data center since software determines the workload (Barroso and Hölzle 2009). Since applications usually do not work in isolation but frequently interact with other software components, it is also important to model such dependencies. As an example, this becomes important when a VM that runs a specific application cannot be moved to another physical host because another application would fail due to their dependence or other restrictions, such as legal issues. Therefore, software-dependencies must be taken into account. This criterion is fulfilled if both aspects – meaning the possibility to model software components as well as dependencies between software-components – are covered by the simulator.

Certainly, energy models need to be part of the simulation tool in order to determine the energy consumption of components. Thereby, the simulator should include energy models that cover all physical components. As an example, this can be archived by providing energy models for all components or by providing an energy model that combines all components in one model. This criterion is fulfilled if the simulator provides energy models for all energy-consuming components (such as in (Barroso and Hölzle 2009; Fan, Barroso and Weber 2007)), partly fulfilled if at least some energy models are provided, and not fulfilled if no energy models are provided.

Finally, power saving techniques should be considered by the simulators since they are crucial for data centers in order to save energy (Ge, Sun and Wang 2013). Such techniques can be applied at the level of servers or at the level of networking elements in a data center. Since a vast amount of techniques exists (confer (Ge, Sun and Wang 2013)), this criterion is fulfilled if at least one technique is included in the simulator.

3.2. Analysis of General Aspects

In this section, the analysis of the general criteria is presented. Table 4 summarizes the results of the analysis for each of the investigated simulation tools.

Table 4: General Aspects of the Analysis

Criterion	CloudSim	GreenCloud	CReST	iCanCloud
Objective	General Purpose	Energy-Efficiency	General Purpose	General Purpose
Evaluation	Proof-of-Concept	Proof-of-Concept	Proof-of-Concept	Proof-of-Concept
Type of Simulation	Discrete Event	Discrete Event	Discrete Event	Discrete Event
License	GPLv3	GPLv2	GPLv3	GPLv3
Language	Java	C++, OTcl	Java	C++
QoS	Yes	Yes	Yes	No
Services	IaaS	IaaS	IaaS	IaaS
Basis	-	ns-2	-	OMNet++

Except for *GreenCloud*, the tools were not implemented for a particular purpose, but can be generally used for the simulation of clouds, while *GreenCloud* was designed for determining the energy consumption in clouds. Each simulator is evaluated by a proof-of-concept. *CReST* is evaluated by conducting several experiments identified in literature. The results of the experiments are then compared to the results found in literature. *iCanCloud* is evaluated by comparing its results with results gained form different instance-types from Amazon's EC2. Kliazovich, Bouvry and Khan conduct several experiments in order to demonstrate the suitability of *GreenCloud* (Kliazovich, Bouvry and Khan 2012), but do not provide data for comparison. Also *CloudSim* is evaluated by conducting several experiments (Calheiros et al. 2011). Although Calheiros et al. mention that *CloudSim* is also used in real-world scenarios (Calheiros et al. 2011), evidence is not provided for this statement.

As expected due to the nature of clouds, each tool is a discrete event simulation. However, the fact that all simulators mostly generate deterministic results is surprising. Fluctuating demands for resources in the workload of a user or a random number of users are rarely taken into account. This is surprising since fluctuating and unpredictable demands are a characteristic of clouds (Cartlidge and Cliff 2013).

All tools have been released under the GPL license. The applied programming languages are C++ and Java, while *GreenCloud* uses a mix of C++ and OTcl for defining experiments since it is based on the network simulator *ns-2* which also uses C++ as well as OTcl. Except for *iCanCloud*, all simulators allow investigating impacts on quality of service, although this is limited on performance degradations. Impacts of load distribution on quality of service cannot be measured in terms of common metrics, such as defined in (Li et al. 2012). This can be a disadvantage in so far that too frequent migration of virtual machines on the one hand may significantly reduce the energy consumption, while on the other hand leading to an increasing response time because the load on the physical servers it too high. This can in turn lead to contractual penalties and thus unpick the cost savings. The only tool that considers other aspects than performance degradations is *CReST* since it is able to investigate the availability of physical components. Overall, the options for investigating QoS-aspects in clouds using simulation are limited.

Regarding the criterion "Services", every simulator addresses the infrastructure-layer (Mell and Grance 2011) of the cloud stack.

3.3. Analysis of Physical Components

Table 5 summarizes the results of the analysis regarding the physical components.

Table 5 already suggests that apparently, there are differences between the simulators in terms of which components can be modeled. For example, the model of a "Server" should consist of a CPU, main memory, hard disk drives, a motherboard, one or more peripheral slots and fans since these are the major energy consumers of a server (Greenberg et al. 2008).

Table 5: Analysis of Physical Components

Criterion		**CloudSim**	**GreenCloud**	**CReST**	**iCanCloud**
Server		Partly	Partly	Partly	Partly
Cooling		No	No	Yes	No
Storage		Yes	No	No	No
Support		No	No	No	No
Network	*Model*	Flow	Packet	Flow	Packet
	Topologies	Partly	Yes	Partly	Yes
	Devices	Yes	Yes	Yes	Yes

As mentioned before, also racks and blades are included in this criterion since these are common in data centers. Taking this fact into account with respect to simulators, it can be stated that *CloudSim* can only model CPU, memory, and disk. *GreenCloud* can model everything except for fans, motherboards, and peripheral slots. Nevertheless, main memory, disks, and network cards are not entities but attributes of a server. Furthermore, memory and disk are static values that are not varied during the runtime. The same elements as in *CloudSim* can be modeled with *CReST*. In the context of a direct comparison, only networking cards are missing in the latter. However, main memory and disks can be explicitly modeled as entities in *CReST*. Furthermore it can be stated that *iCanCloud* can also only partly model the energy consuming components of a server: CPU, main memory, and disks are included as well as servers in terms of racks and blades. Other components are not included in the simulator. Therefore, it can be concluded that none of the simulators can model servers according to their energy consuming components.

"Cooling" is only taken into account by *CReST*, whereas none of the other simulators takes this into consideration. This is surprising, since cooling accounts for a substantial part of the energy costs in data centers – confer for example (Pelley et al. 2009; Barroso and Hölzle 2009) – and is important for other aspects as well, such as the availability of data centers since these cannot operate without cooling (Barroso and Hölzle 2009). Therefore, cooling should be taken into account when investigating the energy consumption of clouds.

The only simulator that can model storage systems is *CloudSim*. Since such systems may be important for migrating virtual machines between physical servers (Mishra et al. 2012), they should also be included in a simulator; especially since the energy consumption is different when applying central storage systems in comparison to servers with integrated storage (Minas and Ellison 2009).

Regarding the network criterion, the analysis revealed that *GreenCloud* and *iCanCloud* are packet-level simulators. Packet-simulation is widely used to study network protocols and is very accurate, but also expensive. An alternative is to implement network simulation by flow-models (Velho et al. 2013) as it is done in *CloudSim* and *CReST*. Since *GreenCloud* as well as *iCanCloud* are based on famous open source network simulators (*ns-2* respectively *OMNet++*), they also provide the possibility to model different types of topologies as well as network devices, such as different types of switches. Also *CloudSim* and *CReST* provide the possibility to model topologies as well as network devices, although less detailed in comparison to *GreenCloud* and *iCanCloud*. In *CloudSim*, different types of topologies do not lead to different results, which is due to a bug in the implementation. The modeling of topologies is thus possible, but is not correctly implemented. The criterion is therefore only partially fulfilled for *CloudSim*. Support systems are not considered by any investigated simulator, so that the criterion “Support Systems” is not fulfilled.

Summing up, it can be said that none of the examined simulators covers all energy-consuming components of (cloud) data centers. Therefore, there is a high probability that the energy consumption for an entire data center estimated by one of the simulators will be inaccurate since various components are left out.

3.4. Analysis of Non-Physical Components

Despite the fact that the physical components are crucial for estimating energy consumption and the effects of load distribution algorithms on energy consumption, also non-physical components are highly relevant. Regarding the determination of the energy consumption by simulation, energy models are especially important. But also other components are of importance for this purpose, as described in section 3.1. The results of the analysis concerning the non-physical components are summarized in table 6.

Table 6: Non-Physical Components of the Analysis

Criterion	**CloudSim**	**GreenCloud**	**CReST**	**iCanCloud**
Load	Yes	Yes	Yes	Partly
Virtualization / Cloud	Partly	Partly	Partly	Partly
Software	Partly	Partly	Partly	Partly
Energy Models	Partly	Partly	Partly	No
Power Saving Techniques	Yes	Yes	No	No

Except for *iCanCloud*, all simulators fulfill the criterion “Load” since they provide functionalities to model different types of workloads as well as taking virtual machine scheduling into account, even though most included workloads lead to a deterministic output. Only *iCanCloud* does not include a functionality for VM scheduling, which is why “Load” is only partially fulfilled.

Regarding the criterion “Virtualization / Cloud”, it is to say that no simulator can meet all of the aspects that are relevant to this criterion (confer section 3.1). *CloudSim* does not consider resource pooling. The same applies for *CReST* and *iCanCloud*. *GreenCloud* provides all the functionalities required for this. However, rapid elasticity is only partially fulfilled, as the demand is not subject to stochastic influences. Thus, rapid elasticity is taken into account in principle, but there are no adjustments to the demand at runtime and all information is thus known a priori, which is why the features for rapid elasticity are not used. With respect to the criterion “Software”, only *iCanCloud* allows to model both software components and dependencies between software components. The other three simulators lack the feature to model dependencies between software. *GreenCloud* is also not able to model software. It is only possible to model tasks which somehow represent software in *GreenCloud*.

Energy models are provided by *CloudSim*, *GreenCloud* and *CReST*, whereas *iCanCloud* does not include any energy models. But none of the simulators provides energy models for all energy-consuming components: *CloudSim* and *CReST* do only provide an energy model for servers and this model only bases on the CPU respectively its utilization, whereas all other components are left out. Regarding an energy model for servers, the same applies to *GreenCloud*. Additionally, *GreenCloud* provides an energy model for networking devices.

The last criterion for the non-physical components are power saving techniques. Since *iCanCloud* does not even provide energy models, also power saving techniques are not included in the simulator. Also *CReST* currently does not include such techniques. However, *CloudSim* and *GreenCloud* include such techniques. *CloudSim* provides an implementation for Dynamic Voltage and Frequency Scaling (DVFS), whereas *GreenCloud* provides implementations of DVFS and Dynamic Power Management (DPM). Therefore, the criterion is fulfilled for both *CloudSim* and *GreenCloud*.

4. EXPERIMENT

The analysis presented in section 3 revealed that there are significant differences between the simulators regarding the components and concepts they take into account. Accordingly, their results in terms of energy consumption should also differ significantly. Therefore, an experiment that aims to verify this assumption is carried out. First, a hypothetical scenario is defined. This scenario is implemented with each simulator (as far as possible). Finally, the results gained from each simulator are compared in order to verify the assumption mentioned before.

4.1. Experimental Setting

In order to keep the effort for implementing the single experiments low, a simple scenario based on a three-tier data center architecture is defined. An example for a typical three-tier architecture as illustrated in figure 2.

A three-tier architecture is common in modern data centers (Kliazovich, Bouvry and Khan 2012). It consists of different types of switches: core layer, aggregation layer and access layer. The core layer connects the data center to the Internet, the aggregation layer provides several functionalities (such as SSL or firewall) and the access layer connects the servers that are partitioned in a racks (Ge, Sun and Wang 2013).

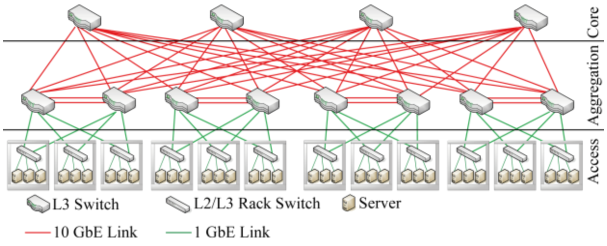

Figure 2: Typical three-tier data center architecture according to (Kliazovich, Bouvry and Khan 2012)

The scenario consists of one core layer switch, two aggregation layer switches and four access layer switches. In total, 40 physical servers are partitioned in four racks. Therefore, each rack holds ten single servers (blades). The servers are homogenous in terms of their provided resources:

- 100.000 MIPS (Million instructions per second)
- 8 GB main memory
- 250 GB hard disk drive
- 1 GbE (Gigabit Ethernet) network adapter

Additionally, a single user is part of the experiment. This user requests a unique service which is always served by a virtual machine or an application within a virtual machine. The user's requests results in 40 different virtual machines. According to the physical servers, the virtual machines have the same resource demands:

- 50.000 MIPS
- 2 GB virtual memory
- 3 GB virtual disk
- 1 GbE virtual network controller

The workload that is generated by the user requests leads to a mean utilization of 50% in terms of the CPU load of each physical server. Each VM is placed on a single server. Live migration of running virtual machines is deactivated, so that the processing of each user request will have to be finished on the server where its serving VM was initially placed. Furthermore, power saving techniques are deactivated since not every simulator provides such techniques. Additionally, the linear energy models of the respective simulators are used. Two cooling units are also part of the scenario in order to manage the temperature of the data center.

4.2. Results

The scenario described in section 4.1 has been implemented in all simulators that had been analyzed before. As already stated in the analysis conducted in section 3, not every detail of the experimental scenario can be modeled with each simulator. For example, it was defined in the previous section that a linear model is used for determining the energy consumption. For instance, such a linear model for the consumption of servers is included in *CloudSim* and in *GreenCloud*. Although both models are based on the utilization of the CPU, they are not identical (compare (Beloglazov and Buyya 2012) and (Kliazovich, Bouvry and Khan 2012)). While *CloudSim* and *GreenCloud* (as well as *CReST*) at least provide an energy model, this is completely missing in *iCanCloud*.

Another example is that, except for *CReST*, none of the simulators is able to model the cooling units (confer section 3). The implementations of the experiments are therefore not completely identical. The simulation results in terms of throughput or response times would probably only be partially comparable – for example since *CloudSim* cannot correctly simulate different topologies, which will lead to errors in the results regarding response times. However, these aspects are not critical for comparing the energy consumption determined by each simulator. Therefore, the similarity of the various implementations of the experiment is sufficient. Indeed, components that maybe cannot be modeled with a simulator, such as specific types of switches, are especially important to the outcome of the experiment. However, profound adjustments of the respective simulation tools would be necessary for other purposes than the comparison of differences in energy consumption in order to get to comparable results.

The simulation results in terms of the energy consumption determined with each simulator are depicted in figure 3. It shows the mean energy consumption determined by each simulator for a time frame of 24 hours.

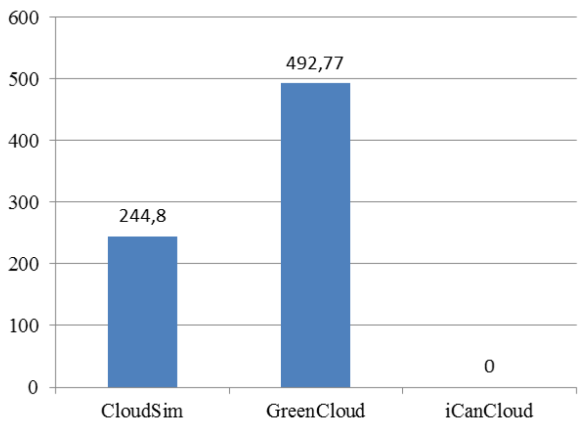

Figure 3: Mean Energy Consumption for each Simulator in kWh

Figure 3 does not contain any results regarding *CReST*. Although the experiment could be modeled in great detail with *CReST*, a program error prevents the experiment from being successfully completed. The same error also occurs in the sample scenarios provide by in *CReST*, which is why it can be foreclosed that the experiment is implemented incorrectly. Therefore, results are missing for this simulator.

As expected, *iCanCloud* is not able to determine the energy consumption of the scenario defined in the previous section. Although many components of the scenario can be modeled, energy models are not included in *iCanCloud* as already stated above. Therefore, the amount of energy consumed by the single components cannot be determined, which is why the energy consumption is 0 as shown in figure 3.

Significant consumption values could thus only be determined for *CloudSim* and *GreenCloud*. As shown in figure 3, the values differ greatly from one another. This is mainly due to the fact that *GreenCloud* simulates the energy consumption of the network devices as well. Since the consumption of both simulators regarding the servers is very similar (244,8 kWh in *CloudSim* and 256,9 kWh in *GreenCloud*), the influence of the network components – although depending on the particular network architecture and topology (Bilal et al. 2013) – is high. As an example, three-tier architectures usually consume plenty of energy. The differences in the determined consumption of the servers are the result of the functioning of the simulators. For example, there are minor differences in the workload profiles: *GreenCloud* reaches a mean utilization of 50% over all servers, but the utilization can be higher or lower for single servers. Since in all simulators, the energy consumption of servers is based on the utilization of the CPU, the determined energy consumption is not exactly the same due to differing utilization factors, although the consumption values are very similar. Since the confidence intervals of both results do not overlap, the differences in the results are significant. In fact, the results are deterministic. Thus, a significant difference exists (confer figure 3).

The assumption made at the beginning of this section that the amount components and concepts taken into account as presented in section 3 have a significant impact on the simulated energy consumption can thus be confirmed.

5. DISCUSSION

The analysis conducted in this paper has revealed some differences between the analyzed simulators.

There is a certain homogeneity between the simulators in terms of which components are included in the servers. Referring for example to the energy-consuming components defined in (Fan, Barroso and Weber 2007), the amount of considered components is not sufficient for providing an accurate prediction of the energy consumption and is hence also not sufficient for determining the effects of load distribution algorithms on energy consumption. The experiment carried out in section 4 supports this assumption, although only a simple scenario was modeled and only two simulators were able to provide results. Considering the proportioning of the energy consuming components of a data center as illustrated in figure 4, it is likely to expect that the results of a more comprehensive simulator will differ even more from the analyzed simulation tools since these do not even completely include the IT equipment (confer the analysis of physical components in section 3.3), not to speak of other consumers such as cooling.

The fact that many components are not considered by the respective simulators is thereby not surprising since appropriate energy models for the single components are often hard to find (Kansal et al. 2010). As the energy models of servers provided by the analyzed simulators are limited on the CPU, the simulators are also not applicable for all purposes.

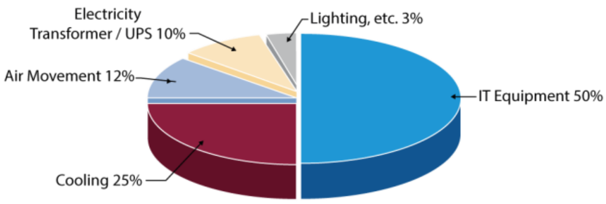

Figure 4: Sources of Energy Consumption in Data Centers (Source: (Power 2007))

For example, if storage systems are investigated, these will generate a lot more load on the disks than the on the CPU (Orgerie, Lefevre and Gelas 2010). Due to the higher load, the energy consumption of the disks would rise, but that cannot be determined by the current energy models. Another type of load would not even require a storage system; also an I/O-intensive workload would possibly increase the load of the disks and therefore their energy consumption. Especially memory is to become a dominant factor in the energy consumption of servers (Minas and Ellison 2009). Regarding an exact prediction of the energy consumption, such aspects must be taken into account.

Surprisingly, cooling is almost ignored by existing simulators, which is astonishing considering their importance for data centers and their energy consumption – cooling accounts for about 25% of the energy consumption of data centers (Orgerie, De Assuncao and Lefevre 2014; Power 2007). Only *CReST* implements a module for cooling, which however does not determine any energy consumption since only a thermal model is implemented. However, the integration of an energy model for cooling has already been prepared in the source code.

Just like cooling, also storage systems and support systems are hardly considered by the simulators. However, these are crucial both for the operation of the data center as well as for its energy consumption (confer figure 3). Especially regarding the physical components, there is a lot of potential for improvements on the side of the simulators in terms of their accuracy of predicting the energy consumption of data centers.

There is also potential for improvement in terms of the non-physical components of simulators. For example, there is a lot of backlog regarding the criterion "Virtualization / Cloud", especially in terms of multi-tenancy aspects (confer section 3.4). Furthermore, the consideration of software within the simulators needs to be improved. Especially with regard to non-functional properties such as availability, it is important to be able to model dependencies between applications. As the analysis revealed, this aspect is excluded by most simulators. Particularly aspects that are relevant for quality of service, such as the availability of a service, are ignored by most simulators.

It is also worth mentioning that none of the analyzed simulators was actually evaluated on the basis of a complete data center. Due to the fact that relevant data for this purpose is hard to get, this is not surprising. However, an evaluation that captures a complete data center is important in order to provide evidence for the validity of the simulation tool. Such an evaluation would be a significant asset with regard to the proof of the suitability of a simulator for determining the energy consumption of data centers. Furthermore, this can help when it comes to determining the effects of load distribution on other components of a data center, such as effects on response time or availability.

Regarding the simulation of clouds or cloud data centers, several issues are still unsolved and should be addressed in future work.

6. CONCLUSION

In recent years, cloud computing has become an important paradigm for providing IT services. Besides the advantages for cloud service providers, cloud computing also comes along with several challenges, among which reducing the overall energy consumption is one of the most important ones. A common method for reducing energy consumption is using load distribution in terms of VM scheduling (Zhang, Cheng and Boutaba 2010). In the past years, numerous algorithms for energy efficient load distribution in clouds have been proposed. Such algorithms are often evaluated by using simulation approaches (Sakellari and Loukas 2013). In order to investigate the applicability of existing simulators for determining the energy consumption in clouds and the effects of load distribution algorithms on energy consumption, an analysis of existing simulators was conducted in this contribution.

The four simulators analyzed in this paper have similar objectives, but diverse designs and implementations. Regarding their suitability for determining the energy consumption of cloud data centers, the simulators have several limitations when it comes to modeling the relevant energy consuming components and also non-physical concepts that are relevant in the context of cloud computing, such as resource pooling (Mell and Grance 2011) or availability. The additionally performed experiment confirms the assumption that the consideration of other components of a data center can strongly affect the outcome of the simulation in terms of energy consumption.

In order to get to more accurate results in terms of energy consumption, a more comprehensive simulator needs to be designed that addresses the shortcomings identified in the course of the analysis.

REFERENCES

Ahmed, A. and Sabyasachi, A.S., 2014. Cloud computing simulators: A detailed survey and future direction. *2014 IEEE International Advance Computing Conference (IACC)*, 866–872.

Barroso, L.A. and Hölzle, U., 2009. The datacenter as a computer: An introduction to the design of warehouse-scale machines. *Synthesis Lectures on Computer Architecture*, 4 (1), 1–108.

Beloglazov, A., Abawajy, J.H. and Buyya, R., 2012. Energy-aware resource allocation heuristics for efficient management of data centers for Cloud computing. *Future Generation Computer Systems*, 28 (5), 755–768.

Beloglazov, A. and Buyya, R., 2012. Optimal Online Deterministic Algorithms and Adaptive Heuristics for Energy and Performance Efficient Dynamic Consolidation of Virtual Machines in Cloud Data Centers. *Concurrency and Computation: Practice and Experience*, 24 (13), 1397–1420.

Bilal, K., Khan, S.U., Malik, S.U.R., Khalid, O., Hameed, A., Alvarez, E., Wijaysekara, V., Irfan, R., Shrestha, S., Dwivedy, D., Ali, M., Khan, U.S., Abbas, A. and Jalil, N., 2013. A taxonomy and survey on Green Data Center Networks. *Future Generation Computer Systems*.

Brown, R., Koomey, J., Masanet, E., Nordman, B., Tschudi, B., Shehabi, A., Stanley, J., Sartor, D., Chan, P., Loper, J., Capana, S., Hedman, B., Duff, R., Haines, E., Sass, D. and Fanara, A., 2007. *Report to Congress on Server and Data Center Energy Efficiency: Public Law 109-431*. United States Environmental Protection Agency.

Buyya, R., Ranjan, R. and Calheiros, R.N., 2009. Modeling and simulation of scalable Cloud computing environments and the CloudSim toolkit: Challenges and opportunities. *International Conference on High Performance Computing & Simulation (HPCS '09)*, 1–11, IEEE.

Calheiros, R.N., Ranjan, R., Beloglazov, A., De Rose, C.A.F. and Buyya, R., 2011. CloudSim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms. *Software: Practice and Experience*, 41 (1), 23–50.

Cartlidge, J. and Cliff, D., 2013. Comparison of Cloud Middleware Protocols and Subscription Network Topologies using CReST, the Cloud Research Simulation Toolkit. *Third International Conference*

on Cloud Computing and Services Science (CLOSER '13), 58–68.
Fan, X., Barroso, L.A. and Weber, W.-D., 2007. Power Provisioning for a Warehouse-sized Computer. *ACM International Symposium on Computer Architecture*, 13–23.
Ge, C., Sun, Z. and Wang, N., 2013. A Survey of Power-Saving Techniques on Data Centers and Content Delivery Networks. *IEEE Communications Surveys & Tutorials*, 15 (3), 1334–1354.
Greenberg, A., Hamilton, J., Maltz, D.A. and Patel, P., 2008. The cost of a cloud: research problems in data center networks. *ACM SIGCOMM Computer Communication Review*, 39 (1), 68–73.
Jing, S.-Y., Ali, S., She, K. and Zhong, Y., 2013. State-of-the-art research study for green cloud computing. *The Journal of Supercomputing*, 65 (1), 445–468.
Kansal, A., Zhao, F., Liu, J., Kothari, N. and Bhattacharya, A.A., 2010. Virtual Machine Power Metering and Provisioning. *ACM Symposium on Cloud Computing*, 39–50, New York, New York, USA: ACM.
Kliazovich, D., Bouvry, P. and Khan, S.U., 2012. GreenCloud: a packet-level simulator of energy-aware cloud computing data centers. *The Journal of Supercomputing*, 62 (3), 1263–1283.
Koomey, J., 2011. *Growth in Data center electricity use 2005 to 2010*. Oakland, CA: Analytics Press.
Kumar, P. and Rai, A.K., 2014. An Overview and Survey of Various Cloud Simulation Tools. *Journal of Global Research in Computer Science*, 5 (1), 24–26.
Li, Z., O'Brien, L., Zhang, H. and Cai, R., 2012. On a Catalogue of Metrics for Evaluating Commercial Cloud Services. *ACM/IEEE 13th International Conference on Grid Computing (GRID)*, 164–173, IEEE.
Liu, H., Jin, H., Xu, C.-Z. and Liao, X., 2013. Performance and energy modeling for live migration of virtual machines. *Cluster Computing*, 16 (2), 249–264.
Mell, P. and Grance, T., 2011. *The NIST Definition of Cloud Computing*. Gaithersburg, MD, USA: National Institute of Standards and Technology (NIST).
Minas, L. and Ellison, B., 2009. *Energy Efficiency for Information Technology: How to Reduce Power Consumption in Servers and Data Centers*. Intel Press.
Mishra, M., Das, A., Kulkarni, P. and Sahoo, A., 2012. Dynamic Resource Management Using Virtual Machine Migrations. *IEEE Communications Magazine*, 50 (9), 34–40.
Nehru, E.I., Venkatalakshmi, B., Balakrishnan, R. and Nithya, R., 2013. Neural load prediction technique for power optimization in cloud management system. *2013 IEEE Conference on Information & Communication Technologies (ICT '13)*, 541–544, IEEE.
Núñez, A., Vázquez-Poletti, J.L., Caminero, A.C., Castañé, G.G., Carretero, J. and Llorente, I.M., 2012. iCanCloud: A flexible and scalable cloud infrastructure simulator. *Journal of Grid Computing*, 10 (1), 185–209.
Orgerie, A.-C., De Assuncao, M.D. and Lefevre, L., 2014. A Survey on Techniques for Improving the Energy Efficiency of Large Scale Distributed Systems. *ACM Computing Surveys*, 46 (4), 1–35.
Orgerie, A.-C., Lefevre, L. and Gelas, J.-P., 2010. Demystifying energy consumption in Grids and Clouds. *Green Computing Conference*, 335–342, IEEE.
Pelley, S., Meisner, D., Wenisch, T.F. and VanGilder, J.W., 2009. Understanding and abstracting total data center power. *Workshop on Energy-Efficient Design*,
Power, E.N., 2007. *Five Strategies for Cutting Data Center: Energy Costs Through Enhanced Cooling Efficiency*. Emerson Network Power.
Sakellari, G. and Loukas, G., 2013. A survey of mathematical models, simulation approaches and testbeds used for research in cloud computing. *Simulation Modelling Practice and Theory*, 39, 92–103.
Smit, M. and Stroulia, E., 2013. Simulating Service-Oriented Systems: A Survey and the Services-Aware Simulation Framework. *IEEE Transactions on Services Computing*, 6 (4), 443–456.
Velho, P., Schnorr, L.M., Casanova, H. and Legrand, A., 2013. On the validity of flow-level tcp network models for grid and cloud simulations. *ACM Transactions on Modeling and Computer Simulation*, 23 (4), 23.
Zhang, Q., Cheng, L. and Boutaba, R., 2010. Cloud computing: state-of-the-art and research challenges. *Journal of Internet Services and Applications*, 1 (1), 7–18.
Zhao, W., Peng, Y., Xie, F. and Dai, Z., 2012. Modeling and simulation of cloud computing: A review. *IEEE Asia Pacific Cloud Computing Congress*, 20–24, IEEE.