Extending Energetic Potentials of Data Centers by Resource Optimization to Improve Carbon Footprint

Alexander Borgerding[1], Gunnar Schomaker[2]

Abstract

The electric power is one of the major operating expenses in data centers. Rising and varying energy costs induce the need of further solutions to use energy efficiently. The first steps to improve efficiency have already been accomplished by applying virtualization technologies.

In this paper, we address the problem of energy efficiency in data centers. Efficient and scalable power usage for data centers is needed. We present different approaches to improve efficiency and carbon footprint as background information. We propose an in-progress idea to extend the possibilities of power control in data centers and to improve efficiency. Our approach is based on virtualization technologies and live-migration to improve resource utilization by comparing different effects on virtual machine permutation on physical servers. It delivers an efficiency-aware VM Placement by assessing different virtual machine permutation.

1. Introduction

The IP traffic increases year by year worldwide. New Information and Communication Technology (ICT) services are coming up and existing services are migrating to IP technology, for example, VoIP, TV, radio, and video streaming. Following these trends, the power consumption of ICT obtains a more and more significant value. In the same way, data centers are growing in number, size, and their share of electric power consumption in order to comply with the increasing demand. The data center's power consumption has doubled in the period 2000-2006 [13]. Energy costs rise continuously and the data center operators are faced with customer questions about sustainability and carbon footprint while economical operation is an all-over goal. The electric power consumption has become one of the major expenses in data centers.

A high performance server in idle-state consumes up to 70% of its peak power [14]. To reduce the quantity of servers in idle-state, virtualization technologies are used. Virtualization technologies allow several virtual machines (VMs) to be operated on one physical server (PM). In this way the number of servers in idle-state can be reduced to save energy [9]. However, the rising energy costs lead to a rising cost pressure and further solutions are needed.

This paper is organized as follows. Section 2 motivates and defines the problem of energy efficiency and integrating renewable energy in data centers. Section 3 gives background on approaches relevant to energy efficiency, virtualization technology and improving the carbon footprint. In Section 4 we present the resource-efficient and energy-adaptive approach. Lastly, Section 5 concludes the paper with comments on our progressing work.

2. Problem definition

The increasing amount of IT services places even greater demands on data centers and energy costs are rising. These conditions induce the need to operate a maximum number of IT services with

---

[1] University of Oldenburg, 26111 Oldenburg, Germany, alexander.borgerding@uni-oldenburg.de

[2] Software Innovation Campus Paderborn, Zukunftsmeile 1, 33102 Paderborn, Germany, schomaker@sicp.de

minimal employment of resources, since the aim is an economical service operation. Therefore the effectiveness of the invested power should be at a maximum level. In this paper, we focus on the server's power consumption and define the efficiency of a server as the work done per energy unit [11].

In the related work part of this paper, we show different kinds of approaches in the context of energy consumption, energy efficiency, and integrating renewable power. In this research approach, we want to explore, which further options exist to use energy efficiently and how we can take effect on the data center's power consumption and, finally, to adapt to volatile renewable energy. The share of volatile renewable energy increases which causes a larger dynamic of the available power. To deal with variable power availability we need an approach that ensures controllable power consumption beyond general energy efficiency. We need to improve the efficiency of the data center using an intelligent and efficient VM placement in order to adapt to volatile energy availability, and improve carbon footprint while keeping the overall goal to use the invested energy as efficiently as possible.

Some approaches use geographically-distributed data centers to schedule the workload across data centers with high renewable energy availability. This methodology is only suitable in big, geographically-spread scenarios and the overall power consumption stays not affected. Hence, we do not pursue these approaches. In general, many approaches are based on strategies with focus on CPU utilization because CPU utilization correlates with the server's power consumption directly [5]. The utilization of other server components does not have such an effect on the server's power consumption. However, the application performance depends not only on CPU usage, but all required resources are needed for optimal application performance. Hence, the performance relies on other components too and we also want to focus on these other components such as NIC, RAM, and disk I/O to improve the efficiency, especially if their utilization does not have an adverse effect on the server's power consumption. Our assumption is that the optimized usage of these resources is not increasing the power consumption.

There are different types of applications; some applications work stand-alone, others rely on several components running on different VMs. The components of the latter communicate via network and the network utilization takes effect on such distributed applications. In our approach we want to include these communication topology topics. However, the applications' requirements are changing during operation, sometimes in large scale and in short intervals. Therefore, we need an online algorithm that acts at runtime to respond to changing values. We need to keep obstacles at a low level by acting agnostic to the applications. The capable approach should be applicable without need to change the operating applications.

3. Related work

Power consumption and energy efficiency in data centers is a topic on which a lot of work has already been done. In this section we give an overview of different approaches.

The usage of low-power components seems to offer solutions for lower energy consumption. Meisner et al. [7] handled the question whether low power consumption correlates with energy efficiency in the data center context. They discovered that the usage of low power components is not the solution. They compared low power servers with high power servers and defined the energy efficiency of a system as the work done per energy unit. They achieved better efficiency with the high power servers and found that modern servers are only maximally efficient at 100% utilization.

Another potential for improvement is to let IT requirements follow energy availability. There are some approaches [2, 4, 6] that use local energy conditions. They migrate the server workload to data center destinations with renewable power availability. These ideas are finally only suitable for distributed and widespread data centers. Data center locations at close quarters typically have the same or not significantly different energy conditions.

A different idea is mentioned by Krioukov et al. [3]. A scheduler has access to a task list, where the task with the earliest deadline is on the top. This is an earliest deadline first (EDF) schedule. If renewable energy is available, the scheduler starts tasks from the top of the task list to use the renewable energy. If less energy is available, tasks get killed. In such approaches, we have to deal with application-specific topics. To build a graded list of tasks to schedule, we need to know how long a task needs to be processed and we need a deadline for each task to be processed. Tasks get killed at less availability; this leads to application-specific issues to resolve afterwards.

Tang et al. [1] propose thermal-aware task scheduling. The ambition is to minimize the cooling requirements and to improve the data center efficiency in this way. They set up a central database with server information, especially server heat information. An EDF scheduler is placing tasks with the earliest deadline on the coldest server. Thus they avoid hot spots and the cooling requirement can be decreased to improve efficiency. The usage of a graded task list is also needed with the same disadvantages as described before. To avoid handling with application-specific topics the virtual machine is a useful container to place IT load instead of explicit application tasks. In many approaches, for example Corradi et al. [9], power consumption is reduced by concentrating VMs on a fewer number of servers, and switching off unused ones to save energy. Chen et al. [11] describe the power consumption of a server as the sum of its static power consumption and its dynamic power consumption. The static power consumption is the consumption of the server in power-on state without workload. This amount of power can be saved with this approach. The dynamic part of server’s power consumption correlates with its CPU utilization as described by Pelley et al. [5]. Thus, most methodologies are only focused on CPU utilization. Dalvanadi et al. [8] and Vu et al. [10] pointed out that network communication can also influence the overall performance of an IT service and network-aware VM placement is also an important and challenging issue. Hence, they embrace network traffic so as to minimize power consumption.

As described, many approaches use virtualization technologies to concentrate VMs on a small number of PMs. While migrating VMs onto a PM, the size of the random access memory (RAM) is a limiting factor. If the RAM-size of the PM is exhausted, further VMs cannot be migrated onto this PM. This can be an adverse effect, especially if resources such as CPU are still underutilized or unused. The memory sharing technology offers the possibility to condense the redundant memory pages on a PM to one page. Unneeded physical memory can be freed to improve the VMs memory footprint. The VMs run on top of a hypervisor, which is responsible for allocating the physical resources to individual VMs. The hypervisor identifies identical memory pages on the different VMs on a PM, it shares them among the VMs with pointers. This frees up memory for new pages. If a VM's information on that shared page changes, the hypervisor writes the memory to a new page and readdresses a pointer. The capacity of the PM can be increased to concentrate further VMs on the PM and to achieve higher server utilization. Wood et al. [12] present a memory sharing-aware placement approach for virtual machines that includes a memory fingerprinting system to determine the sharing potential among a set of VMs. In addition it makes use of live migration to optimize VM placement.

4. Resource-efficient and energy-adaptive approach

In this section the in-progress idea for resource-efficient and energy-adaptive VM placement in data centers is proposed. To optimize the server utilization, many data center operators already use server virtualization technologies and operate several virtual machines on one physical server. This technology is the base for our further optimizations. In our approach, we are at the point that the first steps of optimizations have already been done. Hence, we are running a set of VMs concentrated on a small number of potential servers. Unused servers are already switched off. As further input we get the target power consumption.

It is generally accepted that applications that can access all required server resources operate ideally. With the aim of improving the data center's efficiency, resource-competing VMs should not be operated on the same physical server together. Our approach is to create a VM allocation that concentrates VMs with suitable resource requirements on the same physical server for ideal application performance and efficiency. In this constellation, each application has access to the required server resources and operates ideally. Finally, the overall server resources are more utilized than before and the efficiency rises. These effects include the CPU utilization, so the power consumption increases, because CPU utilization is the most reliable factor regarding server power consumption as mentioned before. This situation leads to more efficiency, but also to a higher power consumption and application performance. This scenario is suitable for times of high energy availability. Following the idea of green energy usage, this technology is also capable of reducing the data center's power consumption in situations of less green power availability. Therefore the methodology can be used to explicitly reduce resource utilization by combining resource-competing applications, leading to lower power consumption but also potentially to a reduced application performance.

In data centers, applications induce specific power consumptions by their evoked server load. This required amount of power is understood as a fixed and restricted value. Our concept is to let this amount of power become a controllable value by applying a corresponding VM allocation. Hence, the power consumption is controllable; it can be increased in times of high energy availability and decreased otherwise.

Our approach is based on virtualization technology and the possibility to live-migrate VMs. The methodology is agnostic to the operating applications. This is an advantage compared to other task scheduling-based algorithms since these have to deal with task execution times and other application-specific topics. In our approach, the applications are untouched and the technology is non-invasive regarding the applications; it only takes effect on the availability of server resources. The variable availability of server resources is a usual setting that applications are confronted with. As described in the related work part of this paper, the PM's RAM can be a limiting factor while migrating further VMs to the PM. We make use of the technology to share RAM across the VMs to increase the number of VMs operated on a PM.

The following diagrams illustrate the practice, how the methodology's strategy migrates VMs between physical servers.

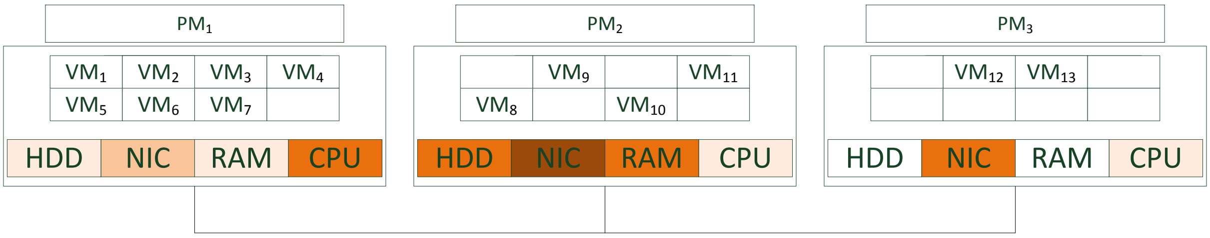

Figure 1: Schematic VM on physical server diagram: initial situation

In Figure 1 the initial, non-optimized situation is displayed showing a set of VMs operated on three physical servers. The resource utilization is highlighted (lighter colors meaning low, darker colors high utilizations). On PM_2, for example the performance is affected by high network utilization.

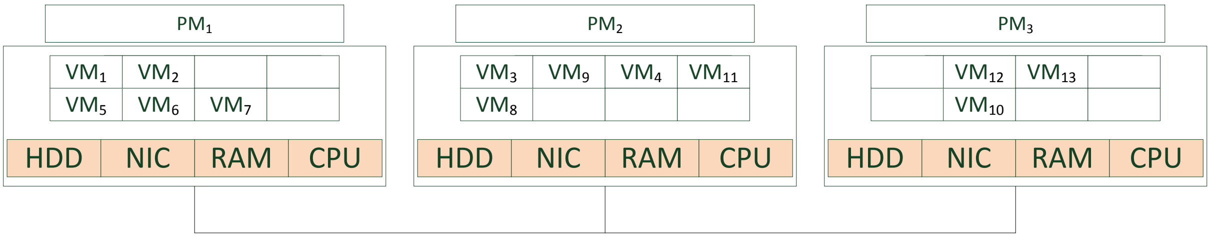

Figure 2: Schematic VM on physical server diagram: optimized situation

After the methodology is deployed, an equilibrium allocation regarding the resource utilization, as shown in Figure 2, is the result. This leads to an average utilization of all involved resources. Hence, the approach increased efficiency and power consumption by resource optimization.

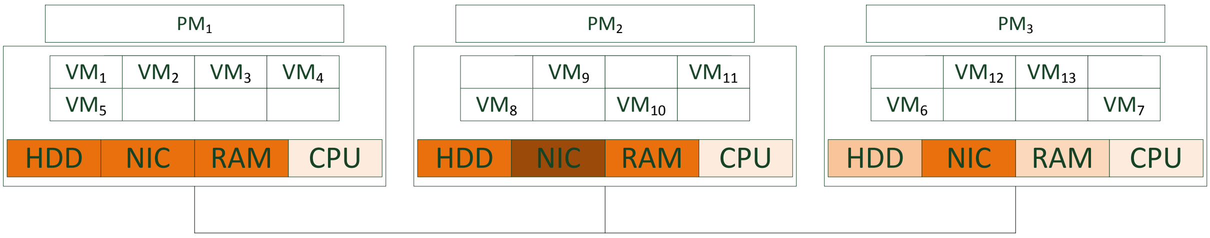

Figure 3: Schematic VM on physical server diagram: aim of reduced power consumption

The situation shown in Figure 3 is the result with reduced power consumption objectives. The CPU utilization is reduced to likewise reduce the power consumption as well while the utilization of other resources is balanced. The result is the most effective constellation at reduced power conditions.

4.1. System Model

In Figure 4 a component model of the entire system is shown. We have an application-monitoring component that delivers information about the applications and servers to the service level management (SLM). The SLM component contains all service level agreements (SLAs) and calculates new power target values for the data center to observe the SLAs. These values are propagated to all optimizers, working on every physical server. The optimizer compares the new incoming target values with its own actual value. If the difference is in range of a predefined hysteresis, the optimizer does not take action. Otherwise it starts optimization. If the target is not in the predefined range and the actual value is lower than the target, the optimizer resolves resource competing constellations and hosts additional VMs from the offer pool. In the offer pool, all the optimizers can announce VMs, for example, if they don’t fit to their actual placement strategy. The VMs in the offer pool are represented with their resource requirements that are the base for later VM placement swaps. If the actual value is higher than the target, the optimizer arranges a resource competing allocation to reduce the power consumption.

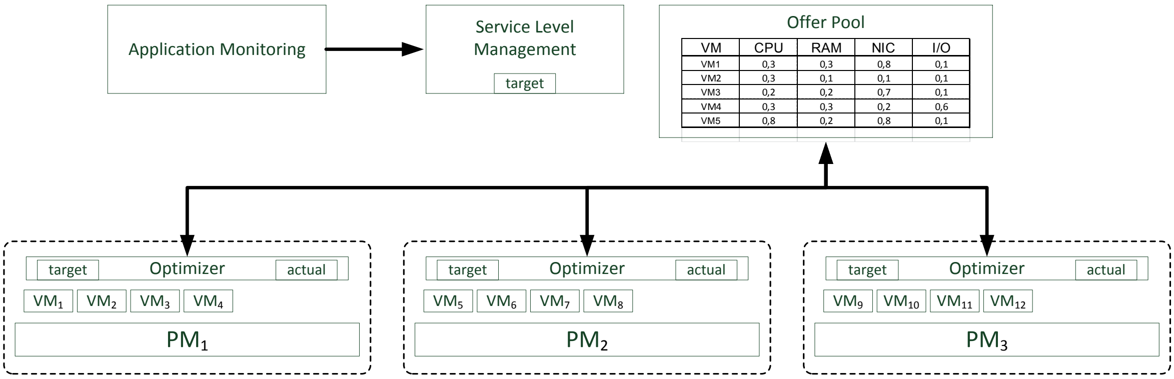

VM	CPU	RAM	NIC	I/O
VM1	0,3	0,3	0,8	0,1
VM2	0,3	0,1	0,1	0,1
VM3	0,2	0,2	0,7	0,1
VM4	0,3	0,3	0,2	0,6
VM5	0,8	0,2	0,8	0,1

Figure 4: Schematic System Model

4.2. Algorithm

CPU utilization is the most effective value regarding power consumption as mentioned before. In other words, the overall CPU utilization is the value to increase or decrease to take effect on the data center's power consumption. Our approach uses competing resource allocations to slow down applications and in series the CPU utilization. Consolidating VMs on a PM that utilize the same resources except the CPU can accomplish this. Consequently, the CPU utilization and power consumption decreases. This practice affects the application's performance and we need a feedback that is sent from the application-monitoring component to the SLM component to ensure the SLAs. With the information about the SLAs and actual application's performance, the SLM component is able to calculate power consumption target values that achieve the economical data center objectives.

The target power consumption is broadcasted to all PMs. The PM has got an optimizer component that receives the target and compares it with its actual value. If the target is similar to the actual value, the optimizer does nothing. Otherwise it starts optimizing. While doing this, the focus is kept on balanced resource utilization. Hence, the overall CPU utilization is reduced or increased but all other resources are used as efficiently as possible. Balanced resource utilization is always the goal except for CPU utilization and resources that are used to build the competing resource situation. Just the attainable CPU utilization is a variable and implicit value.

Every PM's optimizer strives to reach the target value by optimizing its own situation. We have an offer pool of VMs, which can be accessed by every PM's optimizer. The optimizer is able to read out the offered VMs from other PMs or even to offer VMs. If the target value is greater than the actual value, the optimizer removes suitable VMs from the pool to host until the target value is reached. If the target is lower than the actual value, the optimizer offers VMs to the pool to reduce the own value. Furthermore, additional VMs can be hosted from the pool to create competing resource situations to reduce the CPU utilization and to reach the target value. The process of reaching a suitable VM placement and the behavior of the optimizer is demonstrated by the following pseudo code:

Input: ***t*** target power consumption

Input: ***p*** actual PM's power consumption

Output: VM placement that reaches target power consumption

```
receive new target $t$ given by SLM component
if $t > p$ and the PM's CPU utilization is 100%, offer VMs to other PMs via offer pool
```

3. ***if t*** > ***p*** and the PM's CPU utilization is lower than 100%, and all other resources are underutilized, the PM invites VMs to shelter from other PMs with high CPU utilization
4. ***if t*** > ***p*** and the PM's CPU utilization is lower than 100%, and other resources are strong utilized, offer VMs to other PMs to solve the competing resource situation
5. ***if t*** < ***p*** and the PM's CPU utilization is lower than100%, and other resources are strong utilized, invite VMs to shelter from other PMs with high CPU utilization
6. ***if t*** < ***p*** and the PM's CPU utilization is 100%, invite VMs to shelter from other PMs to create resource competing situation
7. ***if t*** = ***p*** do nothing

4.3. Future Work and Experiments

Our primary goal is to increase the data center's power efficiency. The essential research work is to analyze the different reachable effects by combining further methodologies, for example RAM-sharing and integrating further resources (such as RAM, NIC, and HDD) into the approach as described before. Afterwards, a VM placing strategy has to be found.

The problem of determining an efficient VM placement can be formulated as an extended bin-packing problem, where VMs (objects) must be allocated to the PMs (bins). In the bin-packing problem, objects of different volumes must be fitted into a finite number of bins each of the same volume in a way that minimizes the number of bins used. The bin-packing problem has an NP-hard complexity. Hence, a global bin-packing solver will not be able to deliver a VM placement for an online acting approach. In further experiments we will point out the major effects to reduce the complexity. Finally, we have to prove the additional efficiency and to compare our methodology with the results of other existing approaches in virtualization and consolidation. We will point out the further effects of extending the algorithm's scope beyond CPU utilization.

5. Conclusion

We pointed out the raising data center demand, the increasing energy costs, and the requirement to handle volatile energy availability respectively. In the related work part of this paper we presented different approaches related to energy efficiency, power consumption, and usage of renewable power in data centers. We defined the problem of energy efficiency and proposed a resource-optimization approach that improves overall energy efficiency and also allows controlling actual data center power consumption without application-invasive measures. We will point out the potential of this methodology in our ongoing work, especially including further resources beyond CPU and technologies such as RAM-sharing.

Our approach is an instrument to increase efficiency and to adapt to renewable power availability; both have positive effect on the carbon footprint.

6. References

[1] Tang, Q. (et al.) (2008): Thermal-Aware Task Scheduling for Data centers through Minimizing Heat Recirculation, The IMPACT Laboratory School of Computing and Informatics Arizona State University Tempe, AZ 85287.

[2] Chen, C. (et al.) (2012): Green-aware workload scheduling in geographically distributed data centers. 4th IEEE International Conference on Cloud Computing Technology and Science Proceedings, 0:82–89, 2012.

[3] Krioukov, A. (et al.) (2011): Integrating renewable energy using data analytics systems: Challenges and opportunities. IEEE Data Eng. Bull., 34(1):3–11, 2011.

[4] Liu, Z. (et al.) (2011): Greening geographical load balancing. In Proceedings of the ACM SIGMETRICS Joint International Conference on Measurement and Modeling of Computer Systems, SIGMETRICS ’11, pages 233–244, New York, NY, USA, 2011. ACM.

[5] Pelley, S. (et al.) (2009): Understanding and abstracting total data center power. Proc. of the 2009 Workshop on Energy Efficient Design (WEED), Jun. 2009, 2009.

[6] Qureshi, A. (et al.) (2009): Cutting the electric bill for internet-scale systems. SIGCOMM Comput. Commun. Rev., 39(4): 123–134, August 2009.

[7] Meisner, D. (et al.) (2011): Does low-power design imply energy efficiency for data centers? In Naehyuck Chang, Hiroshi Nakamura, Koji Inoue, Kenichi Osada, and Massimo Poncino, editors, ISLPED, pages 109–114. IEEE/ACM, 2011.

[8] Dalvandi, A., Gurusamy, M., Kee Chaing Chua (et al.) (2013): Time-Aware VM-Placement and Routing with Bandwidth Guarantees in Green Cloud Data Centers. In Proceedings of the 2013 IEEE International Conference on Cloud Computing Technology and Science - Volume 01 (CLOUDCOM '13), Vol. 1. IEEE Computer Society, Washington, DC, USA, 212-217.

[9] Corradi, A., Fanelli, M. & Foschini, L. (2014). VM consolidation: A real case based on OpenStack Cloud.. Future Generation Comp. Syst., 32, 118-127.

[10] Vu, Hieu Trong, and Soonwook Hwang. "A Traffic and Power-aware Algorithm for Virtual Machine Placement in Cloud Data Center." *International Journal of Grid & Distributed Computing* 7.1 (2014).

[11] Hao Chen; Caramanis, M.C.; Coskun, AK., "The data center as a grid load stabilizer," *Design Automation Conference (ASP-DAC), 2014 19th Asia and South Pacific* , vol., no., pp.105,112, 20-23 Jan. 2014

[12] Timothy Wood, Gabriel Tarasuk-Levin, Prashant Shenoy, Peter Desnoyers, Emmanuel Cecchet, and Mark D. Corner. 2009. Memory buddies: exploiting page sharing for smart colocation in virtualized data centers. In Proceedings of the 2009 ACM SIGPLAN/SIGOPS international conference on Virtual execution environments (VEE '09). ACM, New York, NY, USA, 31-40. DOI=10.1145/1508293.1508299 http://doi.acm.org/10.1145/1508293.1508299

[13] Smith, J. W. & Sommerville, I. (2011). Workload Classification & Software Energy Measurement for Efficient Scheduling on Private Cloud Platforms. CoRR, abs/1105.2584.

[14] Anton Beloglazov and Rajkumar Buyya. 2012. Optimal online deterministic algorithms and adaptive heuristics for energy and performance efficient dynamic consolidation of virtual machines in Cloud data centers. Concurr. Comput. : Pract. Exper. 24, 13 (September 2012), 1397-1420. DOI=10.1002/cpe.1867 http://dx.doi.org/10.1002/cpe.1867