Energy Saving in Data Center Servers Using Optimal Scheduling to Ensure QoS

Conor McBay, Gerard Parr and Sally McClean
School of Computing and Information Engineering
Ulster University
Northern Ireland, United Kingdom
Email: mcbay-c1@email.ulster.ac.uk, {gp.parr, si.mcclean}@ulster.ac.uk

***Abstract*— With the rise in popularity of cloud computing the amount of energy consumed by the cloud computing data centres has increased dramatically. Cloud service providers are aiming to reduce their carbon footprint by reducing the energy their data centres produce, while maintain an expected Quality-of-Service adhering to set Service Level Agreements. In this paper, we present our suggested approach for using previously researched energy efficiency techniques, particularly Dynamic Voltage/Frequency Scaling and sleep states, more efficiently through the aid of an SLA-based priority scheduling algorithm, and the results we expect from our research.**

Keywords—cloud computing; energy; DVFS; sleep mode; scheduling; quality-of-service.

I. INTRODUCTION

As the use of ICT continues to become an essential aspect of modern life its increased usage has caused the high production of greenhouse gases, contributing 2-3% of global emissions, and is rising each year [1]. The recent popularisation of cloud computing is a contributing factor to this heightened usage. Cloud service providers aim to maintain an expected Quality-of-Service (QoS) and meet set Service Level Agreements (SLAs) while attempting to be cost effective. One of the biggest costs to providers is energy consumption in cloud data centers. Servers, essential for operations within the data center, consume vast amounts of energy. At times large amounts of energy is being wasted on servers that are not operating at their full capacity and may in fact be in an idle state, doing nothing.

The problem faced by cloud service providers is how to reduce the energy consumption within their data centers, while also maintaining the expected QoS and SLA requirements. There has been a lot of research into ways of dealing with the relationship between energy consumption and expected performance, but there are few incentives for providers to use the methods suggested. Reasons for this are that most have not been tested on real world applications and the providers fear that such approaches may lead to breaches in SLAs with corresponding loss of customers [2].

Our proposed solution to this problem is to create a cloud data center infrastructure taking advantage of existing energy saving techniques. Our solution will feature servers that have had these techniques applied to them, while there will be a small number of servers operating as standard as an incentive for providers, by allowing high priority jobs to be completed quickly and by acting as a buffer in case of losing operating servers. An algorithm will be created to allocate incoming tasks to either the standard or energy efficient servers dynamically based on priority scheme.

The rest of this paper is structured as follows. In Section II, we discuss the existing literature and research in this area. In Section III, we expand on our planned architecture and explain our proposed approach. In section IV, we discuss our simulation setup. In Section V, we display the results we have found thus far and what results we expect from our final version. Finally, in Section VI, we present our conclusions found at this point and outline our planned future work.

II. LITERATURE REVIEW

Various studies have been carried out into improving the energy efficiency of cloud computing [1, 3-4], some of which mentioned in Table I. Among the most frequently suggested methods are sleep modes, dynamic speed scaling, DVFS, virtualization, resource allocation, virtual machine migration, green routing, and workload optimization [3]. Our research focuses on a combination of sleep modes and DVFS, as well as introducing a scheduling algorithm to help optimize the process.

Sleep modes, wherein servers can be put into low power states, are one of the most common approaches in order to reduce energy consumption. The basis for this approach is saving energy by powering down servers when they are not needed. By turning off idle servers within a data centre, energy consumption is reduced. However, continually switching servers on and off can cause a time delay and energy penalties. Testa et al. propose a controller in addition to IEEE 802.az Energy Efficient Ethernet to manage transitions between low powered and standard powered states [5]. Using traffic forecasting to manage sleep modes has been suggest by Morosi et al. [6]. Their forecasting based algorithm, forecasting based sleep mode algorithm (FBSMA), allows for daily calculations of approximate traffic.

Another common suggested method of improving energy efficiency is DVFS. This technique allows for the frequency of the CPU to be reduced in times when the CPU load is low, meaning less voltage of power can be consumed [7]. Meisner et al. propose the PowerNap system using DVFS in conjunction with dynamic power management [8]. The Power Aware List-based Scheduling and the Power Aware Task Clustering algorithms suggested by Wang et al use DVFS to propose that non-critical jobs can be run slowly over time to allow the CPUs' frequencies to be reduced [9].

There are various scheduling schemes suggested in research aiming to improve energy efficiency in data center servers. Dong et al. propose a scheduler that will select the most energy

TABLE I. TABLE OF ENERGY EFFICIENCY METHODS IN LITERATURE

Method	Uses			Ref
	Sleep Mode	**DVFS**	**Scheduling**	
IEEE 802.az Energy Efficient Ethernet	Yes	No	No	[5]
Forecasting Based Sleep Mode Algorithm	Yes	No	No	[6]
Dynamic Voltage/Frequency Scaling	No	Yes	No	[7]
PowerNap	Yes	Yes	No	[8]
PALS & PATC Algorithms	Yes	Yes	No	[9]
Most Energy Efficient Server First	No	No	Yes	[10]
Vacation Queueing Model	Yes	No	Yes	[11]
Separate Scheduling Algorithms	No	No	Yes	[12]
Intelligent scheduling with DVFS	Yes	Yes	Yes	[13]

efficient server first [10]. This method allocates tasks based on server energy profiles, and is shown to outperform random scheduling and least-allocated-server-first scheduling. Cheng et al. suggest a method that uses a vacation queueing model to model task schedule, then analyse task sojourn time and energy consumption of computation nodes to create algorithms [11]. Reddy and Chandan's method of using three processors to each run a different scheduling algorithm (earliest-deadline-first, earliest-deadline-late, and first-come-first-served) found energy savings when compared to existing stand-by sparing for periodic tasks, but their savings do not show in all their experiments [12]. Calheiros and Buyya propose a method that combines intelligent scheduling with DVFS for minimum frequency and power consumption, and for completion before user-defined deadline [13]. This approach focuses on urgent, CPU intensive tasks, and their results show between a 2% and 29% improvement on the baseline energy consumption [14].

III. PLANNED ARCHITECTURE

As mentioned previously, our proposed solution is to create a cloud data center infrastructure taking advantage of existing energy saving techniques. Within our infrastructure, we plan to have a combination of servers using energy efficiency techniques, along with servers that operate as normal as fall back to incentivize cloud service providers. Using the literature, we have decided to tackle the issue using a combination of sleep modes, DVFS, and energy efficient scheduling, however unlike the intelligent scheduling with DVFS method, our solution will be applied to all tasks as opposed to CPU intensive tasks only.

Our planned architecture will be made up of two sets of servers: standard servers, that is servers running as normal with no energy efficiency methods applied, and 'green' servers, that will be running using sleep modes and DVFS in order to reduce their energy consumption. A simple outline can be seen in Fig 2. Sleep modes will activate if the incoming traffic to the data center is light and a smaller number of servers can be utilized to handle the workload. DVFS will be used in much the same way, where the traffic allows for the frequency of the processers to be reduced, thereby lowering the energy consumption of the servers. In addition, the green servers will be running with lower processing speeds as we also believe that this will lead to reduced energy consumption.

Currently, we plan on having 90% of the servers within the data center operate using our green model. The final 10% will be running as normal. We believe that with most of the servers using our reduced energy method the energy savings should be sizeable. The remaining 10% will help to ensure that the QoS does not suffer. These values are just a starting point, from which we plan to experiment with different ratios to find the most beneficial setup.

In addition to the energy efficiency techniques applied to the servers, we plan to implement a scheduling algorithm that will direct tasks to the appropriate servers based on a priority scheme within the task's SLA requirements and maintain expected QoS rates for users. Outlined in Fig. 1, at a basic level we envision that tasks will have at least one of the following requirements:

- High priority requirement, meaning the task is of the highest urgency and to be completed as soon as possible with no regards to energy consumption.
- Time requirement, meaning the task can either be on a quick time limit or that it can be run slowly over time on a 'slow burn'.
- Energy requirement, meaning the task is to be as energy efficient as possible while meeting all its other SLA requirements.

The scheduling algorithm will use these requirements to decide which server the task should be assigned to as follows. As each task arrives on the network, information on whether it is a priority task, has a time limit, or has an energy limit will be found. If the task is a priority task then it is immediately assigned to the next available server in order to meet its priority requirements. If the task has a time limit, or can be run at a lower processing rate over a longer time period, then first the target completion time will be found. For each server currently on the network, its 95 percentile completion time will be found. The first server to return a completion time that is within range of the target completion time will be assigned the task. Finally, if the task has a preferred energy limit then the energy target will be found first. Following that, the 95 percentile energy consumption per task will be found, however unlike for the time limit, only low energy servers will be used. The task will be assigned to the first low energy server to return an energy consumption per task within range of the target consumption.

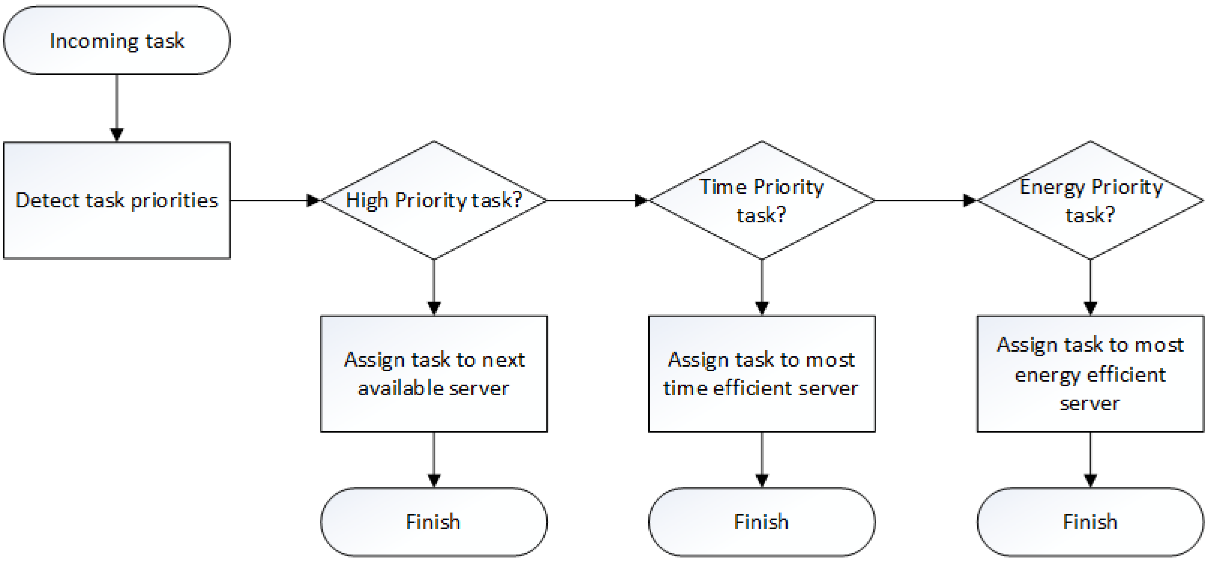

Figure 1. Flow chart of proposed algorithm functionality

IV. Simulation Setup

After testing out a range of simulation software, including CloudSim, DCSim, and iCanCloud [14], for our experiments we decided to use the Greencloud simulator [15]. Greencloud is a packet-level simulator for energy-aware cloud computing in data centers and an extension of the ns-2 simulation software.

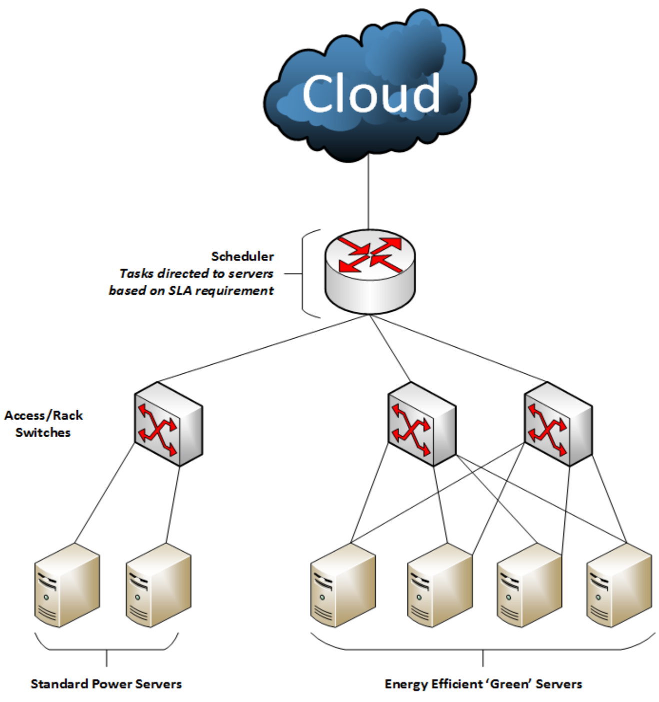

Figure 2. Outline of planned architecture.

We decided to use Greencloud as it is a simulator designed specifically for measuring the energy efficiency of cloud computing data centers. Greencloud uses a three-tier data center design, seen in Fig. 3, consisting of a core layer, an aggregation layer, and an access layer.

Currently our simulation setup consists of 144 servers, all of the same specification, with 1 core switch, 2 aggregation switches, and 3 access switches. The servers' specification contain commodity processors with 4 cores, 8GB of memory, and 500GB of disk storage, giving the data center a total of 576057600MIPS (Millions of Instructions Per Second) of processing power.

To examine QoS in our simulation, we are using round-trip time. We have self-imposed a limit of 2 seconds on single tasks. Tasks exceeding the 2 second limit are in breach of QoS, however we have also decided on a 5% trade-off, allowing for 5% of tasks within the experiments to exceed this limit without being in breach of SLAs.

V. Results

Using the simulator, we have tested the energy consumption of the data center when servers are run as normal and when DVFS and sleep modes have been applied. For this experiment, we used a random assignment scheduler and aimed or a data center workload of around 30%. This created 32689 tasks for the data center to process.

We found that using DVFS and sleep modes in a data center can reduce the energy consumption. As seen, the data center consumes 495.3 W*h using standard servers, with the servers consuming 332 W*h of that. In comparison, when DVFS and sleep modes are applied, the data center consumes 472.1 W*h, with the servers consuming 308.8 W*h. That is a percentage difference of 7.24% without any optimisation applied.

VI. Conclusions And Future Work

In this paper, we presented our idea to introduce a cloud data center infrastructure taking advantage of existing energy saving techniques combined with a priority scheme based scheduling algorithm in order to optimise the process and maintain QoS. We have suggested an approach wherein a major number of servers within the data center are run using

 ISBN: 978-1-61208-460-2

energy efficiency methods, such as DVFS and sleep modes, and a small minority are left to run as standard as an incentive to cloud service providers. We outlined the structure for our scheduling algorithm based upon SLA requirements and how we think it can improve the energy efficiency of cloud data centers.

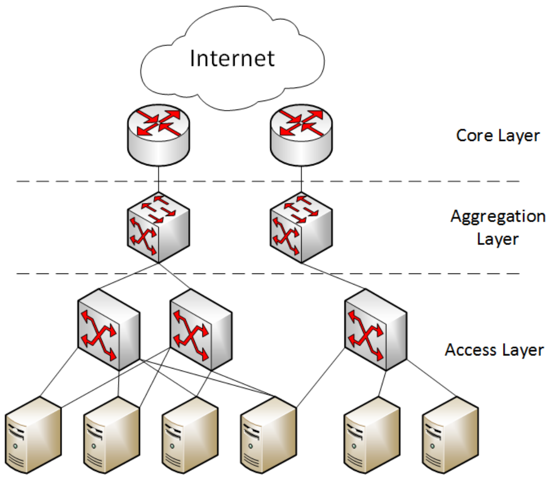

Figure 3. Example of a simple three-tier data centre architecture in Greencloud.

The results from our experiments thus far have shown that by using DVFS and sleep modes in cloud data centers the energy consumption of the servers can be reduced, and that there is still room for improvement, and we aim for our scheduling algorithm to help do this.

In the future, we aim to accomplish several key issues. Firstly, we will expand our simulation model to create our two distinct types of servers. Using our original 144 servers, approximately 10% will be left to run using the initial specifications with no energy efficiency methods applied. The remaining 90% will be run with slower processing speeds, as well as DVFS and sleep modes to allow for the servers to be placed into a lower power state then the incoming traffic is low.

The next step in our planned work is to amend our simulated tasks with our new SLA priority scheme. Tasks will be assigned at least one of the three priority requirements discussed previously. Once the priority scheme is in place we will work to implement our scheduling algorithm that will assign incoming tasks to the relevant servers based upon the priority scheme. Using this algorithm we hope to find that the energy efficiency methods enacted can be optimised to achieve the best energy reduction results possible while maintaining the QoS through measuring task completion time, failed tasks, and dropped packets.

Finally, as another incentive for cloud service providers, we are aiming to create realistic simulated workloads based on Google cluster data that was released in 2011 [16]. This data has been collected over a period of three months in one of Google's cloud data centers and offers real world information that can be adapted for our simulations. We hope that the inclusion of this data in our workload generation will show how our optimised approach would react if placed in a real world situation. This would act as another incentive for cloud service providers, along with the standard powered servers, to implement more 'green' energy applications.

REFERENCES

[1] A. Hameed et al., "A Survey and Taxonomy on Energy Efficient Resource Allocation Techniques for Cloud Computing Systems," in Computing, 2014, pp. 1–24.

[2] Y. Georgiou, D. Glesser, K. Rzadca, and D. Trystram, "A Scheduler-Level Incentive Mechanism for Energy Efficiency in HPC," in 2015 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, 2015, pp. 617–626.

[3] A. Kulseitova and A. T. Fong, "A Survey of Energy-Efficient Techniques in Cloud Data Centers," in 2013 International Conference on ICT for Smart Society (ICISS), 2013, pp. 1–5.

[4] G. L. Valentini et al., "An overview of energy efficiency techniques in cluster computing systems", Cluster Computing, vol. 16, no. 1, pp 3-15, 2013.

[5] P. Testa, A. Germoni, and M. Listanti, "QoS-aware sleep mode controller in 'Energy Efficient Ethernet,'" in 2012 IEEE Global Communications Conference (GLOBECOM), 2012, pp. 3455–3459.

[6] S. Morosi, P. Piunti, and E. Del Re, "Sleep mode management in cellular networks: a traffic based technique enabling energy saving," Trans. Emerg. Telecommun. Technol., vol. 24, no. 3, pp. 331–341, 2013.

[7] A. Hammadi and L. Mhamdi, "A survey on architectures and energy efficiency in Data Center Networks," Comput. Commun., vol. 40, no. 1, pp. 1–21, Dec. 2013.

[8] D. Meisner, B. T. Gold, and T. F. Wenisch, "The PowerNap Server Architecture," ACM Trans. Comput. Syst., vol. 29, no. 1, pp. 1–24, Feb. 2011.

[9] L. Wang et al. , "Energy-aware parallel task scheduling in a cluster," Futur. Gener. Comput. Syst., vol. 29, no. 7, pp. 1661–1670, Sep. 2013.

[10] Z. Dong, N. Liu, and R. Rojas-Cessa, "Greedy scheduling of tasks with time constraints for energy-efficient cloud-computing data centers," J. Cloud Comput., vol. 4, no. 1, pp. 1-14, 2015.

[11] C. Cheng, J. Li, and Y. Wang, "An energy-saving task scheduling strategy based on vacation queuing theory in cloud computing," Tsinghua Sci. Technol., vol. 20, no. 1, pp. 28–39, 2015.

[12] H. K. S. Reddy and S. P. Chandan, "Energy aware scheduling of real-time and non real-time tasks on cloud processors (Green Cloud Computing)," in 2014 International Conference on Information Communication and Embedded Systems (ICICES), 2014, no. 978, pp. 1–5.

[13] R. N. Calheiros and R. Buyya, "Energy-Efficient Scheduling of Urgent Bag-of-Tasks Applications in Clouds through DVFS," in 2014 IEEE 6th International Conference on Cloud Computing Technology and Science, 2014, pp. 342–349.

[14] W. Zhao, Y. Peng, F. Xie, and Z. Dai, "Modeling and simulation of cloud computing: A review," in Proceedings - 2012 IEEE Asia Pacific Cloud Computing Congress, APCloudCC 2012, 2012, pp. 20–24.

[15] D. Kliazovich, P. Bouvry, and S. U. Khan, "GreenCloud: A packet-level simulator of energy-aware cloud computing data centers," J. Supercomput., vol. 62, pp. 1263–1283, 2012.

[16] C. Reiss, J. Wilkes, and J. J. L. Hellerstein, "Google cluster-usage traces: format+ schema," 2011.