

Optimization-based scheduling of data center workload in function of outside weather conditions

Roald DE MEYER, Vincent DEBUSSCHERE, Seddik BACHA

Univ. Grenoble Alpes - G2ELab

RESUME – Data centers are the fifth largest energy consumer in the world and demand for data center services, driven by cloud computing, is rising rapidly. There is also a lot of interest in using data centers for offering grid services. Here, focus is put on scheduling, or in other words, shifting workload in time. This work explores the possible gains that could be achieved if workload would be easily scheduled. An energetic model of the data-center is used, taking into account the dependency of the cooling's coefficient of performance (COP) on the outside weather conditions as well as the influence of the heat load on the power consumption of the fans and pumps. This model is used to show the possible energy savings that could be obtained by scheduling the workload in function of outside wet bulb temperatures and fan powers.

Mots-Clés – Optimization ; modeling ; data-center ; energy management.

1. INTRODUCTION

1.1. Context

Data centers represent a large part of the worldwide energy consumption. A study in 2008 estimated that worldwide expenses for data center power consumption were as high as \$30 billion in the year 2008. Furthermore, the current trend is a steep growth of data center usage, driven by an increased demand for cloud computing services [13]. Not only does the total energy consumption increase, but also the size of individual data centers. Nowadays, their capacities easily reach several MW's, consuming more than thousands of households [12]. Both the high cost of energy and current governmental regulations and incentives for energy efficiency improvements have led to a reversal in the industry. Companies like ABB and Schneider Electric are publishing white papers in which they discuss improvements of data center architecture and control [1]. The commitment of the sector also shows through formation of the Green Grid, a non-profit consortium dedicated to advancing energy efficiency in data centers. The Green Grid has over 150 members, amongst others : Google, IBM, Microsoft, Intel and HP.

Current research concerning the control can be split up into two categories. The first category is trying to find the most energy efficient way of executing the workload as given. Usually this involves the dynamic allocation of virtual machines. The main idea is to save power at partial load, concentrating all the virtual machines on a few servers and putting superfluous servers in idle mode. This has been the subject of the EnergeTIC project (FUI Minalogic 2010-2013) and similar approaches can be found in [2] and [6]. The results from the EnergeTIC project

have shown to be very promising with cuts in the electricity bill as high as 45%.

The second category contains all research that handles optimal scheduling of the workload. Scheduling implicitly assumes that requests can be delayed or cancelled. Very little information is available about workload that can be scheduled and further research on which services could offer this flexibility is necessary. Examples of possible tasks that could be flexible are for instance the generation of back-up or the installation of new software. Another situation that results in flexibility in time could be when data centers have the choice of accepting a request or not. Whether this is a viable option will depend on the quality of service (QOS) agreed upon in the service level agreement (SLA). A method that chooses to execute certain requests or not, comparing the cost of energy and the cost of not executing, can be found in [9].

This work assesses the possible efficiency gains and cost savings that could be realized if the workload is scheduled [10]. For this purpose, it introduces an energetic model of the data center's cooling system in function of the outside weather conditions and the workload intensity. The results could be used as a motivation for researchers to continue to investigate further the "schedulability" of services or to convince data center owners to increase their efforts to make their own workload able to be scheduled.

As a prospective objective, attention will be given to the question of how data centers can be integrated into the smart grid and more specifically to what would be the cost of these smart grid services [11].

Since the goal is to improve the energy efficiency, the energetic model is at the heart of the work. After that, the analysis allowing to identify the major energy consumers is conducted, quantifying the dependency of the energy consumption on workload intensity and outside air conditions. The optimization of different workload profiles for different temperature profiles and flexibilities is finally conducted for results analysis.

2. THE ENERGETIC MODEL

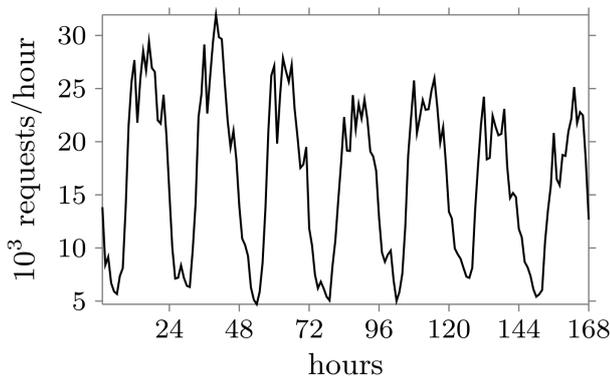
The model consists out of two main parts ; the information technology (IT) one [15] and the cooling one [14].

2.1. Information Technology (IT)

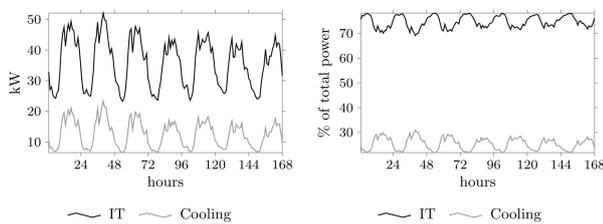
The models of the different IT components are part of the results from the EnergeTIC project. The entire IT system in this model contains three server racks, three power distribution units (PDU's) and an uninterruptible power supply (UPS) [16].

a varying workload is fed to the model.

The workload that is used was measured on actual servers of the G-scop lab. The workload for seven days is plotted in Figure 2(a). The temperature is set to the reference value of 30°C.



(a) Workload for one week



(b) IT and cooling power

Fig. 2. Workload and cooling power

Figure 2(b) shows the IT and cooling power in kW and as percentages of the total power. It can be seen that the cooling represent between 20 and 30 % of the total power consumption. At peak hours the share of the cooling increases towards 30 %.

This is because the IT power is a linear function of the data center workload, but the fans and pumps in the cooling system have a cubic relation with the workload via the volumetric flow rate. If some of the data center’s workload is flexible, ventilation energy could be saved by shifting this flexible load to off-peak periods.

3.2. The Optimizations

To benefit from the two potential energy savings, some kind of flexibility is needed. This flexibility can be offered in two different ways ; storage and flexible workload.

Because of the high power density of data centers, storage of a cooled fluid would involve a large isolated or underground storage tank (for a 1 MW installation the size would be of the order of an Olympic swimming pool) and other major changes in the cooling system and data center architecture. At this point, we should also note that even if a storage capacity is not considered in this work, it is clearly a relevant perspective.

The goal is to optimize the scheduling of the workload. This will be done for the measured workload data under the assumption that part of it is flexible. The flexibility is defined as such that a certain percentage of the hourly workload can be shifted to periods within the same day with more favorable operating conditions.

Measurements of the server workload for every hour for two months are available, but for starters only two load profiles of

two different days are selected : a day with a high workload and one with a low workload. The workload is systematically express in percentages, including in the coming results. The two profiles are shown in Figure 3.

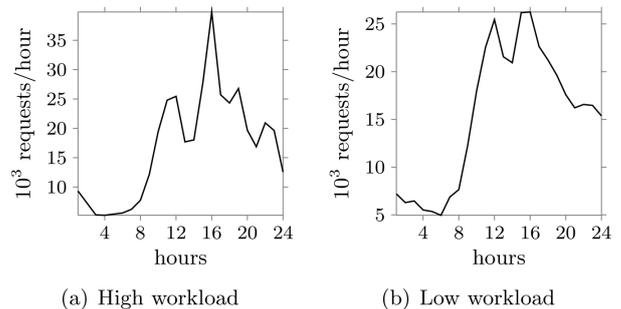


Fig. 3. Workload profiles

The workload is scheduled so that it minimizes the total data center energy consumption, taking into account the constraints. A first constraint imposes a maximum on the workload the data center can handle every hour. This maximum is set equal to 60,000requests/sec. A second constraint makes sure that the workload can only be delayed and not advanced in time. When only a part of the original load is flexible, a third constraint limits the workload that can be shifted. And finally, a fourth constraint is that at the end of the optimization window all workload has to be treated. The mathematical description of the optimization problem is given below :

$$\min \sum_{h=1}^{24} P_h(optw_h, t_h) \quad (6)$$

with :

$optw_h$ the workload at a given time (number of requests per second),

t_h the hourly step time.

We define in the following results the amount of flexibility, noted “flex” in the text.

The objective function is nonlinear, non-smooth and non-convex and thus the optimization problem is fairly complicated. The problem was tackled using different solvers in both Matlab and AMPL. The solvers used in MATLAB are fmincon, global-search, simulated annealing and the genetic algorithm. The ones in AMPL are MINOS and COUENNE.

The best results in Matlab are obtained with fmincon. However, they are very sensitive towards the chosen initial value. The results obtained with both solvers in AMPL were superior to those in Matlab in every case and independent of the chosen initial value.

The results discussed below are always the ones obtained in AMPL with the MINOS and COUENNE solvers. MINOS is a solver suited for finding local solutions of non-linear problems and COUENNE is an open-source solver for finding global solutions of mixed integer nonlinear problems (MINLP). A positioning of COUENNE amongst other solvers and a small user manual can be found in [3, 4].

The optimized workload profiles for different levels of flexibility and for all combinations of temperature and original work-

load profiles are shown in Figure 4. It can be seen that when temperatures are low enough for free cooling, the largest part of the workload is shifted towards these periods.

It can also be seen that during these periods the maximum workload of 60,000 requests/sec is not reached and thus that free cooling is not used to its full extend. This is because increasing the peak further would mean that the increased power of the fans would be higher than the power saved with free cooling.

The reduction in total energy consumption in percentages of cooling energy and the savings in € can be found in tables 1 and 2 respectively.

Tableau 1. Energy savings in % of total cooling energy

Flexibility	High workload		Low workload	
	1st Apr	11th Jul	1st Apr	11th Jul
20 %	12.9	5.3	11.2	3.8
40 %	19.5	6.8	17.2	4.7
60 %	21.0	7.0	18.4	4.7
80 %	21.0	7.0	18.4	4.7
100 %	21.0	7.0	18.4	4.7

Because of the free-cooling in April the energy savings are much larger than the ones possible in July. It seems that for flexibilities higher than 60 % the consumption doesn't decrease further, thus the optimal scheduling can be reached when 60 % of the load is flexible.

Tableau 2. Energy savings in €/day

Flexibility	High workload		Low workload	
	1st Apr	11th Jul	1st Apr	11th Jul
20 %	9.2	4.8	7.1	3.2
40 %	13.9	6.2	10.8	3.9
60 %	14.8	6.4	11.5	3.9
80 %	14.9	6.4	11.5	3.9
100 %	14.9	6.4	11.5	3.9

For this flexibility the total energy savings vary from 3.7 to 4.5 % on the 1st of April according to the workload and from 1.2 to 1.8 % on the 11th of July. The cooling energy savings vary from 18.4 to 21.0 % on the 1st of April and from 4.7 to 7 % on the 11th of July. If a cost of electricity of 0.30 €/kWh is assumed, this results in cost savings of 11.5 to 14.9 and 3.9 to 6.4 €/day.

3.3. Shifting window

In reality the scheduling is done shifting the optimization window and solving an optimization problem for example every hour. The constraints of the optimization problem need some minor adjustments to avoid that workload keeps on being postponed endlessly.

The results of a 24 hour time window and a time step of one hour are identical to the results of the reference in the previous section, except for the particular case of 20 % flexibility, no free cooling and a high workload. Here the savings in cooling energy with the shifting window were 3.3 % less compared to the reference results.

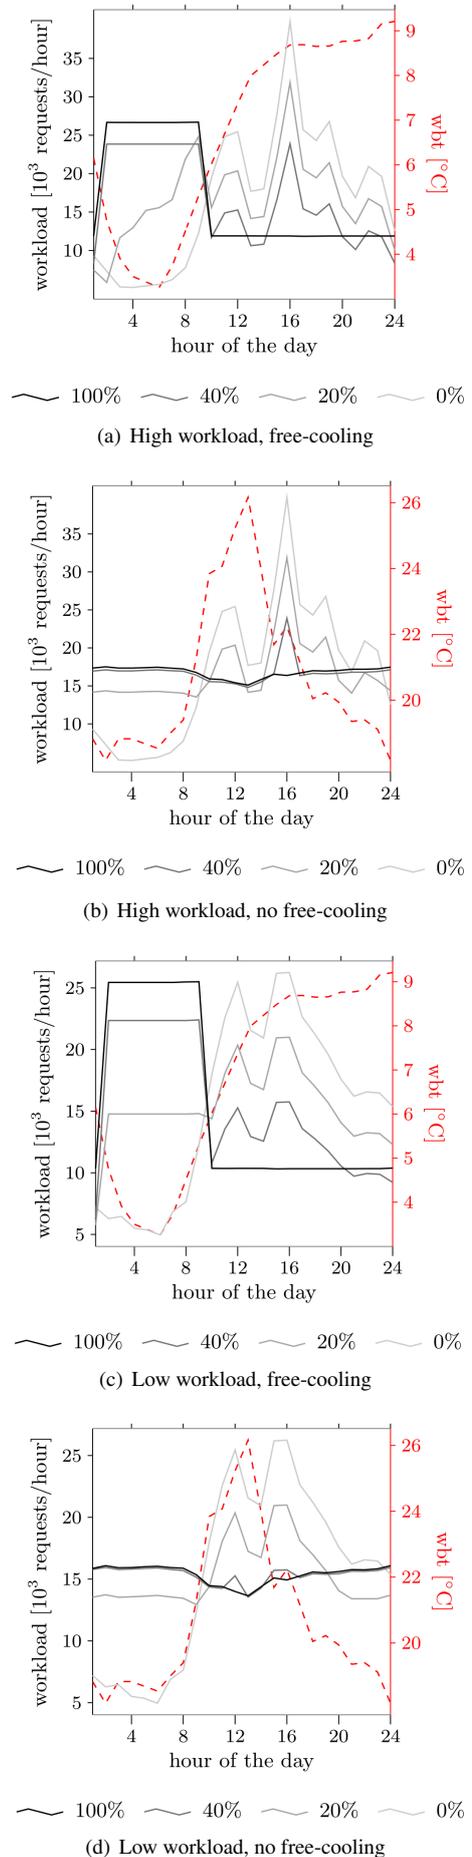


Fig. 4. Optimized workload profiles

Hence, it is possible to conclude that with perfect predictions for the next 24 hours the most optimal solution can be reached, except in some situations for low flexibilities. If the quality of the prediction endures, it might be beneficial to increase the size of the optimization window for low flexibilities.

4. CONCLUSION

In current researches there is a lot going on about improving the efficiency of operation of data centers. Most research is about trying to optimize the allocation of workload. However, very little information can be found about scheduling.

A detailed energetic model of a data center was proposed. Besides modeling the IT equipment, a lot of attention is given to the cooling system, which is new compared to the current state of the art of related research. The formulas used to describe the power consumption of the IT are derived from measurements on real data centers. The model for the cooling contains the outside air temperature dependency of the cooling efficiency and the possibility for free cooling.

All the fans are modeled as if they are equipped with variable speed drives and thus their power consumption varies with the data center workload. The strengths of the model are exactly its relation to real measurements, the dependency of the outside weather conditions and the variable fan power.

The model is used to investigate the possible gains of optimizing the scheduling of workload through the day in function of the outside wet bulb temperature. This is done in a general manner for different real workload and temperature profiles and for different levels of flexibility.

The results show that the most optimal solution can be reached for flexibilities of 60% or higher. For this flexibility up to 21% of cooling energy and 4.5% of total data center energy could be saved. The energy savings are lower than expected because the gains in cooling energy made by scheduling workload during periods with low outside wet bulb temperatures are counteracted by an increased workload peak and thus fan power.

Future works will include a grid integration of this optimized workload, considering a potential local production, through PV panels. Later on, ancillary services could be sharpened for the grids, based on the availability of installing any kind of energy storage on site. That storage will not have to be purely electrical to be used in that context, which is a great advantage.

RÉFÉRENCES

[1] ABB. White paper : data center modernization : From smart data centers to enterprise smart grids. Technical report, ABB, 2010.

[2] Anton Beloglazov, Jemal Abawajy, and Rajkumar Buyya. Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing. *Future generation computer systems*, 28(5) :755–768, 2012.

[3] Pietro Belotti. Couenne : a user’s manual. *Technical Report*, 2009.

[4] Michael R Bussieck and Stefan Vigerske. Minlp solver software. *Wiley encyclopedia of operations research and management science*, 2010.

[5] D Christy Sujatha and Sathesh Abimannan. Energy efficient free cooling system for data centers. In *Cloud Computing Technology and Science (CloudCom), 2011 IEEE Third International Conference on*, pages 646–651. IEEE, 2011.

[6] Gargi Dasgupta, Amit Sharma, Akshat Verma, Anindya Neogi, and Ravi Kothari. Workload management for power efficiency in virtualized data centers. *Communications of the ACM*, 54(7) :131–141, 2011.

[7] Tony Evans. The different types of air conditioning equipment for it envi-

ronments. *White Paper*, 59 :2004–0, 2004.

[8] Jorge Facao and Armando C Oliveira. Thermal behaviour of closed wet cooling towers for use with chilled ceilings. *Applied Thermal Engineering*, 20(13) :1225–1236, 2000.

[9] Mahdi Ghamkhari and Hamed Mohsenian-Rad. Energy and performance management of green data centers : A profit maximization approach. *IEEE transactions on Smart Grid*, 4(2), 2013.

[10] Mahdi Ghamkhari and Hamed Mohsenian-Rad. Energy and performance management of green data centers : A profit maximization approach. *Smart Grid, IEEE Transactions on*, 4(2) :1017–1025, 2013.

[11] Amir-Hamed Mohsenian-Rad and Alberto Leon-Garcia. Coordination of cloud computing and smart power grids. In *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, pages 368–372. IEEE, 2010.

[12] Asfandyar Qureshi, Rick Weber, Hari Balakrishnan, John Guttag, and Bruce Maggs. Cutting the electric bill for internet-scale systems. In *ACM SIGCOMM computer communication review*, volume 39, pages 123–134. ACM, 2009.

[13] Lei Rao, Xue Liu, Le Xie, and Wenyu Liu. Coordinated energy cost management of distributed internet data centers in smart grid. *Smart Grid, IEEE Transactions on*, 3(1) :50–58, 2012.

[14] R.A. Villamizar Saenz. Modélisation et simulation du système de refroidissement des data centres. Technical report, Technical report, 2012.

[15] Van Giang Tran, Vincent Debusschere, and Seddik Bacha. Data center energy consumption simulator from the servers to their cooling system. In *PowerTech, IEEE Grenoble*, pages 1–6, Grenoble, France, 16-20 June 2013. IEEE.

[16] Ghaith Warkozek, Elisabeth Drayer, Vincent Debusschere, and Seddik Bacha. A new approach to model energy consumption of servers in data centers. In *Industrial Technology (ICIT), IEEE International Conference on*, pages 211–216. IEEE, 19-21 March 2012.