

42 Energy optimization in cloud based infrastructures

P. Karkazis

Dpt. of Computer Systems Engineering, Piraeus University of Applied Science, Athens, Greece E-mail: pkarkazis@gmail.com

A. Voulkidis²

Synelixis Solutions Ltd, Athens, Greece

E-mail: voulkidis@synelixis.com

D. Tseles³

Dpt. of Automation Engineering, Piraeus University of Applied Science, Athens, Greece E-mail: dtsele@teipir.gr

42.1 Abstract

In the last couple of years, cloud services became more popular. On the one hand, service complexity increased and on the other requirements related with the Quality of Service and Service Level Agreement negotiations moved to synergic cloud-based architectures. The number of the data centers increased and also a federated architecture of data centers were introduced. However, one of the major effects was the increase of the consumed electricity from Data Centers (DCs) which had negative impact in energy footprint and the environmental conditions. In this paper, we present an optimization procedure, focusing on minimization of the energy consumption of the data centers, based on grouping genetic algorithms and we provide experimental results from a real DC.

42.2 Introduction

The ever increasing demand for computing capacity and the resulting burgeoning of large scale Data Centers (DCs), which constitute huge energy sinks, have a direct impact on the ICT related energy consumption. This huge energy consumption poses a great challenge for the energy sector and the problem is further intensified by the volatility of the energy markets and the inability of Smart Grids to follow the electricity demand-response model, which impedes the seamless integration of large scale DCs to the energy network. Thus, Smart Grid operators need to address on the one hand the immense energy consumption of DCs and on the other hand the erratic operation of Smart Grids, caused by the inability to follow the demand-response paradigm.

To elaborate on the first of the two problems, the increase of DCs, which is accompanied by huge electricity consumption and sub-optimal energy management, directly affects their energy footprint and the environmental conditions. It is well known that the average server utilization in DCs is low, often below 30% of the maximum server load [1], [2] and only 10% in case of facilities that provide interactive services [3]. This low utilization is primarily due to two reasons: (i) the provisioning of a DC is done based on the expected peak load, rather than the average load. For interactive services, peak utilization often exceeds the average utilization by more than a factor of three [3]; and (ii) in order to provide redundancy in the

event of failure. DC operators deploy more systems than are actually needed. The over-design and over-provisioning of DCs and the increased number of low utilized servers, have significantly increased the waste of energy. In the last couple of years, the electricity consumed by DCs has doubled, representing an aggregate annual growth rate of 16.7% per year worldwide [4]. Approximately 80% of this growth is caused by the increased electricity used by servers, 10% by the growth in electricity used for DC communications and around 10% by the growth in electricity by storage equipment.

Besides electricity consumed for supporting the computational operation of the servers, a huge amount of energy is also consumed for the cooling of DC servers. To lower this waste of energy, DC containment strategies (both hot aisle and cold aisle) are widely regarded as the starting point for energy-efficiency best practices. Moreover, the so called "Green DCs" aim to use a number of green electricity sources (e.g. photovoltaic cells, geothermal power, hydroelectric energy, etc.), for normal operations and cooling purposes. The results are, in many cases, impressive, but they still represent a minority of the deployed DCs and even in those cases the intermittent nature of green electricity sources make the need for integration of green energy sources to the energy network and for stable Smart Grid operation more actual than ever.

As already highlighted above, this additional problem, that is the instability of the Smart Grids and the difficulty to follow the electricity demand-response model constitutes a major problem in the energy sector. In particular, as Europe shifts away from fossil fuels, electricity is becoming an even more important energy vector and the seamless integration of renewable energy sources to the energy network becomes imperative. More than 29 European countries have targets for a share of renewable energy in the range of 10-33% until 2020. Achieving these goals is vital for the EU internal energy market, as it will lower the dependency on importing oil and it will help towards a more sustainable growth. The implementation of more intelligent and active transmission, distribution and supply systems in the form of Smart Grids is central to the success of such a development. Thus, Smart Grids are very high on the agenda of the European energy and ICT sector. However, the problem is that Smart Grids have difficulty in following the electricity demand-response model. The introduction of Smart City technology is also being developed as a mechanism to enable intelligence in buildings, city blocks and regions. As a result, we need solutions, which can support the features of the Smart Grid, coupled with the capabilities of Smart Cities, in order to carefully manage the energy profile of DCs, especially under periods of increased demand.

Recent literature suggests that the problem of optimizing and coordinating the energy consumption of federated Data Centers and its alignment with the Smart Grid stabilization needs, is actively researched. In [5], a survey on the existing techniques utilizing geographical load balancing for optimizing the energy consumption of Data Centers in the context of a Smart Grid is presented. The optimization may have different targets, including absolute energy consumption with respect to QoS guarantees [6], cost [7] and carbon footprint [8], the techniques employed varying among Mixed Integer Programming, Dynamic Programming, heuristics through Genetic Algorithms etc [5]. Load balancing of Data Centers in the context of the Smart Grid are also investigated in [9], where the authors present a two-stage framework for modelling the relevant interactions and formulate a cost-minimization problem based on linear programming. Similarly, the authors of [10] present a cooperation scheme between Smart Grids and Data Centers, with the aim to maximize the share of renewables in the energy mix used for Data Center operation. The problem of optimal load (VM) allocation in federated Data Centers is also tackled in [11], where a greedy heuristic is presented with a view towards minimizing the carbon dioxide emissions due to Data Center

operation. Finally, in [12], an in-depth survey of existing algorithms and techniques for orchestrated energy management and energy sustainability in federated clouds is presented. In the direction of tackling the above problems and significantly contributing toward improving the energy efficiency of DCs and stabilizing Smart Grids, the present paper introduces a holistic approach interconnecting networks of DCs and Smart Grids, addressing both problems in a complementary way. Specifically, in the context of smart city and Smart Grid integration, a network of synergetic DCs can adjust its operation shifting load to regions of renewable energy surplus, playing a key role toward Smart Grid stabilization and “Green” operation of modern DCs. Moreover, the proposed approach can be seamlessly integrated to legacy DC equipment discounting any capital expenditure employing solely the software defined networking (SDN) and software defined infrastructures (SDI) of legacy DC equipment. The paper is organized as follows Section 2 introduces the predictive optimization techniques, in order to implement the proposed approach. Section 3 presents the preliminary results obtained by an in-house micro DC, after the employment of the predictive optimization techniques described in Section 2. Finally, Section 4 concludes the paper and presents relevant perspectives.

42.3 Predictive Optimization

The employment of load prediction methods for the efficient prediction of the standalone DC load, the synergetic DC load and the user load is of paramount importance for the efficient optimization of the load allocation to the synergetic DCs. Such methods provide energy predictions based on the user habits, the behavior and the workload patterns as well as the weather forecast, allowing the devise of predictive energy patterns. Subsequently, based on these a priori devised energy patterns the relevant optimization modules can devise relevant plans optimizing the VM / load allocation to standalone and synergetic DCs, based on the Smart Grid status.

A number of techniques have been proposed for forecasting aggregated and correlated energy consumption inspired by machine learning, and have passed from linear regression and autoregressive moving average models [13] to neural networks [14] and boosting approaches [15] and finally to the Support Vector Machine For Regression (SVR) that is a state of the art forecasting method [16-18] . The SVR uses the same principles as the Support Vector Machine (SVM) for classification, but as output instead of a real number, which has infinite possibilities, it returns a margin of tolerance, to minimize error.

The SVR method is employed for energy consumption forecasting, as it combines several desirable properties compared with other existing techniques: it has a good behavior even if the ratio between the number of variables and the number of observations becomes very unfavorable, with highly correlated predictors, it makes possible to construct a nonlinear model without explicitly having to produce new descriptors (the famous "kernel trick"), while a deeply study of the characteristics of the method allows to make comparison with penalized regression such as ridge regression [19], whereas a number of pre-calibrated SVR toolkits can be found online [20], facilitating the easier fine-tuning of the SVR.

Having fine-tuned the SVR parameters for the load prediction at DC and user level an appropriate optimization algorithm must be selected in order to efficiently employ the SDI (through cloud managers such as OpenStack [21], OpenNebula [22] and Eucalyptus [23]) in order to minimize the reserved physical resources and the implicit operating cost. In practice, VMs reserve virtual shared CPU and shared storage, whereas the only physical resource reserved in a stringent way is the server physical memory. Thus, the problem of optimal VM

allocation across the network of synergetic DCs is reduced to that of allocating the aggregate server memory to VMs, based on their forecasted load and availability.

The above problem can be reduced to a “bin packing” problem [24] and has been formally formulated by the authors in [25]. In particular, “the problem of VM allocation can be considered as a “bin packing” problem, where given a finite set $U=\{u_1, u_2, \dots, u_n\}$ of “items”(i.e. VMs) and a rational “size” (i.e. memory) $s(u)$ for each item $u \in U$ a partition of U into disjoint subsets $U_1, U_2, \dots, U_k$ must be found such that the sum of the sizes of the items in each subset U_i is no more than a respective “bin size”(i.e. server memory) S_i and such that k is as small as possible. Thus, VMs of memory s need to be allotted to servers of memory S , while reserving the minimum number of servers, whereas a memory granularity of 512MB can be assumed which is a typical value encountered in practice.”

The above problem constitutes an NP-hard problem [24], however a number of approximation and heuristic techniques can be employed to provide solutions to the problem. The Best Fit Decreasing Algorithm (BFD) [24] constitutes one of the best approximation algorithms for the “bin packing” problem and it can be employed to achieve a consolidated VM allocation. As stressed by the authors in [25] in the direction of employing the BFD “the DC servers are indexed based on their energy-efficiency, with energy-efficient servers being assigned a lower index. Subsequently, “items” (i.e. VMs) are placed into “bins”(i.e. servers) in order of increasing index. As a result, energy-efficient servers are assigned a higher priority and for instance servers of a Green Room are reserved first, or servers of the same DC segment are reserved prior to remote DC servers in order to allow remote DC servers to hibernate, providing substantial energy savings. Next, “items” in U are sorted by size and re-indexed so that $s(u_1) \geq s(u_2) \geq \dots \geq s(u_n)$. “Items” are then placed in order of increasing index, first into the occupied “bins” of lower available capacity and then, in case they do not fit into the occupied “bins”, or in case of a tie, “items” are placed in order of increasing index into the lower indexed “bin” they fit.”

Thus, the documented success of the BFD approximation solution can be exploited to initialize the search of an appropriate heuristic approach. Specifically, the above solution is used as an initial seed to initialize the search of a Genetic Algorithm (GA) approach [26]. The GA constitutes one of the most successful heuristics [26], however, a number of factors hinder the convergence of GA when the latter is applied to grouping problems such as the “bin packing” problem in hand. In particular, grouping problems – aiming either to find a good partition of a set or better yet to group together the members of a set - challenge the cornerstone of the GA, namely the principal of minimal redundancy of each solution, as different encodings and different permutations of the groups may refer to the same solution. Also, solution clustering into groups hinders the passing of useful (i.e. standalone) information to the next generation through the crossover and mutation operators of the GAs [27-29].

In this course, the Grouping Genetic Algorithms (GGA) have been proposed [27] allowing the encoding of grouping problems like the one in hand, by using groups or in our case “bins” as the GA building blocks on which GA operators are applied. One could envisage a GGA as a simple GA where each gene of a GAs’ chromosome corresponds to a tuple of elements corresponding to the “items” of each “bin”, whereas the “bins” are the building blocks evolved by the employment of the GAs. This approach alters all GA operators significantly, however this approach outperforms the standalone GA substantially when applied to grouping problems.

The employment of the GGA, initialized by the BFD, for the optimal VM allocation, allows for the consolidated allocation of VMs at an intra-DC level as well as an inter-DC level, whenever a VM consolidation is imposed by the Smart Grid operation. Thus, the distributed application of the above optimization algorithm on DC sites, when that is deemed necessary

based on the SVR load predictions, could yield significant energy savings as well as reliable Smart Grid operation.

In order to validate the efficiency of the proposed approach and the feasibility and applicability of the proposed approach employing solely the SDI, the preliminary results of the optimized VM allocation are tested on an in-house micro DC and the obtained data corroborating the substantial benefits arising from the proposed approach are presented hereafter.

42.4 Experimental Results

In the process of developing and fine-tuning the prediction engine and optimization module, a number of attested scenarios were used as benchmark to quantify the convergence of the developed optimization algorithms and the accuracy of the developed prediction models. These optimized test scenarios to which the predictive optimization converged to, were then implemented based on the SDI of a small scale testbed, as a proof of concept, employing OpenStack for the actuation of these scenarios. The testbed consists of 4 low consumption blade servers (less than 50W of energy consumption at average load, simultaneously underclocking idle cores) which run artificial loads to emulate the operation of a commercial DC.




Figure 1: Prediction Engine performance

The performance of the implemented prediction engine is depicted in Figure 1, where the load (power) prediction is plotted against the real power demand values. The training set of the prediction engine spanned two months of data. The Root Mean Squared Error between the actual power demand values and the predicted ones is 74.3 which is considered acceptable for our value range, granted the limited training set volume; further training of the prediction engine is required in order to acquire more accurate results.




Figure 2: Test optimization scenario result

In the same context, Figure 2 presents the outcome of the load optimization on the aforementioned test setup. Specifically, under random load and granted relevant predictions from the prediction engine, the optimization component was able to reorganize the existing load in such a way that the energy consumption dropped by approximately 15%, exhibiting that through proper management, the energy consumption of DCs can be significantly lowered, to help towards assisting the operation of Smart Grids. Moreover, when considering the ability to relocate loads to geographically distant DCs when intra-DC optimization is unable to accommodate the load inside the DC boundaries, the coordination of DC loads with the Smart Grid demand response plans, is expected to contribute substantially to the achievement of Smart Grid stability.

42.5 Conclusions

In this paper, a framework for achieving energy optimization in federated DCs has been presented, employing continuous DC resources and network monitoring and scalar optimization architectures operating at local and federated levels. In the course of minimizing the energy consumption at local DC level, we employ load optimization through load re-allocation, coupled with near future load predictions, implemented with the help of support vector regression techniques. The results of the prediction and optimization processes are presented and briefly discussed, indicating significant power savings may be achieved by employing the proposed architecture.

42.6 References

- [1] C. Bash and G. Forman, "Cool job allocation: Measuring the power savings of placing jobs at cooling-efficient locations in the Data Centre," in Proc. of the 2007 USENIX Annual Technical Conference, Jan 2007.
- [2] P. Bohrer, E. Elnozahy, T. Keller, M. Kistler, C. Lefurgy, and R. Rajamony, "The case for power management in web servers," Power Aware Computing, Jan 2002.
- [3] X. Fan, W.-D. Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," in Proc. of the 34th Annual International Symposium on Computer Architecture, 007.

- [4] Jonathan G Koomey, "Worldwide electricity used in Data Centres", IOP Publishing Ltd, Sep. 2008.
- [5] Rahman, Aminur, Xue Liu, and Fanxin Kong. "A survey on geographic load balancing based data center power management in the smart grid environment." *Communications Surveys & Tutorials*, IEEE 16.1 (2014): 214-233.
- [6] [25] Lei Rao, Xue Liu, Le Xie, and Wenyu Liu. Minimizing electricity cost: optimization of distributed internet data centers in a multi-electricity market environment. In *Proc. International Conference on Computer Communications*, pages 1–9. IEEE, 2010
- [7] Lei Rao, Xue Liu, Marija D Ilic, and Jie Liu. Distributed coordination of internet data centers under multiregional electricity markets. *Proc. IEEE*, 100(1):269–282, 2012.
- [8] Yanwei Zhang, Yefu Wang, and Xiaorui Wang. Greenware: Greeningcloud-scale data centers to maximize the use of renewable energy. In *Middleware*, pages 143–164. Springer, 2011
- [9] Wang, Hao, et al. "Exploring smart grid and data center interactions for electric power load balancing." *ACM SIGMETRICS Performance Evaluation Review* 41.3 (2014): 89-94
- [10] Niedermeier, Florian, et al. "Increasing Data Centre Renewable Power Share via Intelligent Smart City Power Control." *Proceedings of the 2015 ACM Sixth International Conference on Future Energy Systems*. ACM, 2015.
- [11] Giacobbe, M.; Celesti, A.; Fazio, M.; Villari, M.; Puliafito, A., "An approach to reduce carbon dioxide emissions through virtual machine migrations in a sustainable cloud federation," in *Sustainable Internet and ICT for Sustainability (SustainIT)*, 2015 , vol., no., pp.1-4, 14-15 April 2015
- [12] Maurizio Giacobbe, Antonio Celesti, Maria Fazio, Massimo Villari, Antonio Puliafito, *Towards energy management in Cloud federation: A survey in the perspective of future sustainable and cost-saving strategies*, *Computer Networks*, Volume 91, 14 November 2015
- [13] J. W. Taylor, "Triple seasonal methods for short-term electricity demand forecasting," *European Journal of Operational Research*, July 2010.
- [14] H. Hippert, "Neural networks for short-term load forecasting: A review and evaluation," *IEEE Transactions on Power Systems*, February 2001.
- [15] S. B. Taieb and R. J. Hyndman, "A gradient boosting approach to the Kaggle load forecasting competition," *International Journal of Forecasting*, 2013
- [16] B.-J. Chen, M.-W. Chang, and C.-J. Lin, "Load forecasting using support vector machines: A study on EUNITE competition 2001," *IEEE Transactions on Power Systems*, vol. 19, no. 4, 2004
- [17] N. Sapankevych and R. Sankar, "Time series prediction using support vector machines: A survey," *IEEE Computational Intelligence Magazine*, May 2009
- [18] S. Humeau, T. Wijaya, M. Vasirani, K. Aberer, "Electricity Load Forecasting for Residential Customers: Exploiting Aggregation and Correlation between Households," *IEEE Sustainable Internet and ICT for Sustainability (SustainIT)*, Palermo, 30-31 Oct. 2013
- [19] N. Christianni, J. Shawe Taylor, "An introduction to Support Vector Machines and other kernel-based learning methods," Cambridge University Press, 2000, section 6.2.2
- [20] "Scikit-learn," <http://scikit-learn.org/stable/index.html>
- [21] "Openstack," <http://www.openstack.org/>
- [22] "Opennebula," <http://opennebula.org/>
- [23] "Eucalyptus," <https://www.eucalyptus.com/>
- [24] G. M. and J. D., *Computers and Intractability; A guide to the theory of NP-Completeness*. New York: W.H.Freeman and Company, 1979.
- [25] A. Aravanis, T. H. Velivassaki, A. Voulkidis, T. Zahariadis, and P. Cottis, "Federated data centers as smart city stabilizing factors," in *The 3rd IEEE ISCC 2015 International Workshop on Smart City and Ubiquitous Computing Applications*, July 2015.
- [26] E. G. Talbi, *Metaheuristics From Design to Implementation*. Hoboken, NJ, USA: Wiley, 2009.
- [27] Falkenauer Emanuel and Delchambre Alain A Genetic Algorithm for Bin Packing and Line Balancing, in "Proc. of the IEEE 1992 Int. Conference on Robotics and Automation (RA92)", May 10-15, 1992, Nice, France.
- [28] Falkenauer Emanuel *The Grouping Genetic Algorithms - Widening the Scope of the GAs in JORBEL - Belgian Journal of Operations Research, Statistics and Computer Science*, vol. 33 (1, 2), pp.79-102.
- [29] Falkenauer Emanuel *New Representation and Operators for GAs Applied to Grouping Problems*, in *Evolutionary Computation*, Vol.2, N.2, pp. 123-144.