On the Minimization of the Energy Consumption in Federated Data Centers

Alexis I. Aravanis(✉), Panagiotis Karkazis, Artemis Voulkidis, and Theodore Zahariadis

Synelixis Solutions Ltd., Athens, Greece
{aravanis,pkarkazis,voulkidis,zahariad}@synelixis.com

Abstract. As cloud services are becoming increasingly popular, the number of operating data centers is accordingly increasing, together with the need of implementing federated data centers and clouds. In this context, we consider a framework for achieving energy efficiency in federated clouds, by means of continuous monitoring and SLA renegotiation, coupled with the operation of prediction and multi-layered optimization components. In this paper, relevant prediction and optimization components, based on Support Vector Regression and Bin-Packing solving heuristics, operating at local data center level are examined and the experimental results of their deployment in a real-life testbed are presented and discussed.

Keywords: Federated data centers · Energy minimization · Optimization · Support vector regression · Bin-packing problem

1 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

B. Mandler et al. (Eds.) IoT 360° 2015, Part I, LNICST 169, pp. 388–398, 2016.
DOI: 10.1007/978-3-319-47063-4_40

(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 [23], 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 [25], cost [24] and carbon footprint [26], the techniques employed varying among Mixed Integer Programming, Dynamic Programming, heuristics through Genetic Algorithms etc. [23]. Load balancing of Data Centers in the context of the Smart Grid are also investigated in [30], 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 [29] 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 [28], where a greedy heuristic is presented with a view towards minimizing the carbon dioxide emissions due to Data Center operation. Finally, in [27], 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 proposed framework, developed in the context of the European Union project: "Data centres Optimization for energy-efficient and enviromentalLy Friendly iNternet (DOLFIN)" [5], facilitates therefore the integration of a federated DC network to the Smart Grid within a Smart City exchange network, allowing the optimal allocation of the cumulative load, based on predictive optimization techniques, that will be presented hereafter. The efficacy of these will be further corroborated by the actual results obtained by the in-house micro DC, presented herein.

The remainder of the paper is organized as follows. Section 2 presents the DOLFIN approach, exploiting SDN to seamlessly integrate the Smart Grid and DC networks in a Smart City agglomeration. Thus, counteracting the adverse effect of erratically operating Smart Grids and allowing for the energy efficient operation of DCs. Section 3 introduces the predictive optimization techniques employed by the DOLFIN ecosystem, in order to implement the proposed approach. Section 4 presents the preliminary results obtained by our in-house micro DC, after the employment of the predictive optimization techniques described in Sect. 3. Finally, Sect. 5 concludes the paper and presents relevant perspectives.

2 DOLFIN Flexible Approach

Modern DCs are part of computing and storage clouds, offering their customers Virtual Machines (VMs) as a virtual operating environment. Exploiting this virtualization of modern DCs and capitalizing on the benefits of SDN, the present approach focuses on modelling, monitoring, and measuring the energy consumption of VMs. This real time monitoring allows for the seamless, autonomic migration of VMs between servers of the same DC or across a group of Energy-conscious, Synergetic DCs, aiming to (i) optimize the overall energy consumption by dynamically changing the percentage of active versus stand-by servers and the load per active server in a DC, and (ii) stabilize the Smart Grid energy distribution, under peak load and increased demand, by dynamically changing the energy consumption/production requirements of the local DCs.

To elaborate, the leeway provided by the independent management of VMs allows for the optimal allocation of the computing load. On the one hand the optimal allocation of VMs within the same DC could lead to the VM consolidation in favor of significant energy savings emerging from the hibernation of inactive servers. On the other hand the optimal allocation of VMs across synergetic DCs could allow for DCs to adjust their operation in accordance with the Smart Grid needs, moving VMs to DCs where energy is cheaper or abundant either due to the time difference of the respective DCs or due to the existence of renewable energy surplus, generated by renewable energy sources in the vicinity of the destination DC. The high-level architecture of such a network of synergetic DCs interacting with the energy network is depicted in Fig. 1.

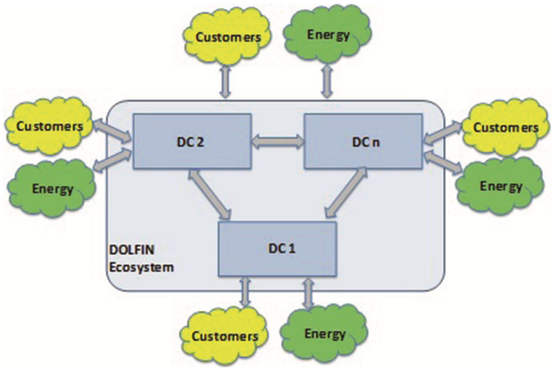

Fig. 1. Network of synergetic DCs interacting with the energy network.

Evidently, the stabilizing effect of the proposed paradigm allows the seamless integration of renewable energy sources (RES) to the energy network, counteracting the adverse effect of the intermittent green energy generation. In particular, occasional peaks and troughs of green energy generation, leading to inverse power flow or reduced system inertia can be balanced by the demand response of federated DCs, mitigating electrical grid instabilities and system outage or blackout threats.

The above energy consumption optimization approach, revolves around three main pillars, underpinning the DOLFIN ecosystem:

(a) Energy-conscious Synergetic DC-level: optimizing the energy consumption within the limits of a single DC, based on system virtualization and the optimal distribution of VMs. This is coupled with the dynamic adaptation of active and stand-by servers and the load optimization per active server. Utilizing a monitoring framework to measure the energy consumption per server module/networking component and activate low-power states on devices.
(b) Group of Energy-conscious Synergetic DCs-level: optimizing the cumulative energy consumption in a group of DCs, based on optimal distribution of VMs across

all of the servers that belong in the group of DCs using load prediction methods for the standalone DCs and the group of DCs. Measuring and predicting the energy consumption on the DC level and achieving a decreased cumulative power consumption across the whole group of DCs.

(c) Smart City-level: optimizing the energy consumption at the smart city level and providing stabilization of the local Smart Grid, based on distribution of VMs across the servers that are part of a group of DCs, following an electricity demand-response approach. In order to stabilize the Smart City Energy consumption, the percentage of active versus stand-by servers per DC is dynamically changed and the load per active server is optimized.

The above hierarchical classification gives rise to four different optimization levels that should be dealt with. In this course, a distributed spiral optimization process is assumed dealing with these four levels of optimization. Specifically, during a single optimization cycle the energy is optimized first by an internal control loop at servers rack level, next at DC segment or DC level and then at federated DC level allowing load relocation among energy conscious DCs. Thus, a network of interconnected DCs employing the spiral optimization approach, could provide energy-efficient DC operation in the context of a fully elastic cloud, while playing a key role in a Smart Grid energy network balancing the stochastic energy surplus provided by renewable energy sources, through efficient load relocation.

Moreover, even though the VM movement between servers in geographically distributed groups of DCs is not trivial, as very strict Service Level Agreements (SLAs) to the DCs' clients should be guaranteed, the real time monitoring of SLAs described above ensures that the VM movement will only exploit the leeway provided not leading to any Quality of Service (QoS) breakage. Thus, the real time monitoring of SLAs provides an additional degree of freedom to the flexible VM management allowing the exploitation of the SLA margin to its fullest.

Having outlined the optimization levels of the DOLFIN approach pertaining to the optimal VM allocation at four different levels, the relevant optimization problems need to be formally formulated, whereas the predictive optimization techniques employed for the problems in hand, need to be thoroughly described as well. In this course, the VM optimization problem and the predictive optimization techniques employed are presented hereafter.

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 [6] to neural networks [7] and boosting approaches [8] and finally to the Support Vector Machine For Regression (SVR) that is a state of the art forecasting method [9–11]. 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 by DOLFIN 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 [12], whereas a number of pre-calibrated SVR toolkits can be found online [13], 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 [14], OpenNebula [15] and Eucalyptus [16]) 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 [17] and has been formally formulated by the authors in [18]. In particular, "the problem of VM allocation can be considered as a "bin packing" problem, where given a finite set $U = \{u_1, u_2, \ldots, u_n\}$ 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 $U_1, U_2, \ldots, U_k$ 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) Si 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 512 MB can be assumed which is a typical value encountered in practice."

The above problem constitutes an NP-hard problem [17], however a number of approximation and heuristic techniques can be employed to provide solutions to the problem. The Best Fit Decreasing Algorithm (BFD) [17] 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 [18] 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 reindexed so that $s(u_1) \geq s(u_2) \geq \cdots \geq s(u_n)$. "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 [19]. The GA constitutes one of the most successful heuristics [19], 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[1], 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 [20–22].

In this course, the Grouping Genetic Algorithms (GGA) have been proposed [20] 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 DOLFIN approach employing solely the SDI, the preliminary results of the optimized VM allocation are tested on our in-house micro DC and the obtained data corroborating the substantial benefits arising from the proposed approach are presented hereafter.

[1] The principal of minimal redundancy refers to the necessary one to one relation between each encoded solution and each member of the search space.

4 Experimental Results

In the process of developing and fine-tuning the prediction engine and optimization module of the DOLFIN ecosystem 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 50 W of energy consumption at average load, simultaneously underclocking idle cores) which run artificial loads to emulate the operation of a commercial DC.

The performance of the implemented prediction engine is depicted in Fig. 2, 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.

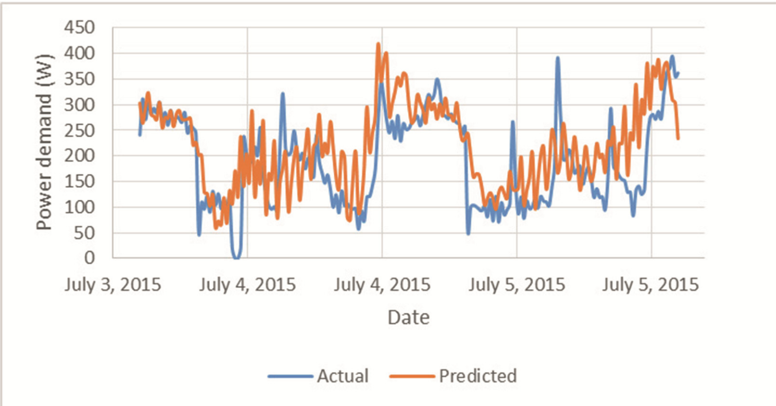

Fig. 2. Prediction engine performance

In the same context, Fig. 3 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.

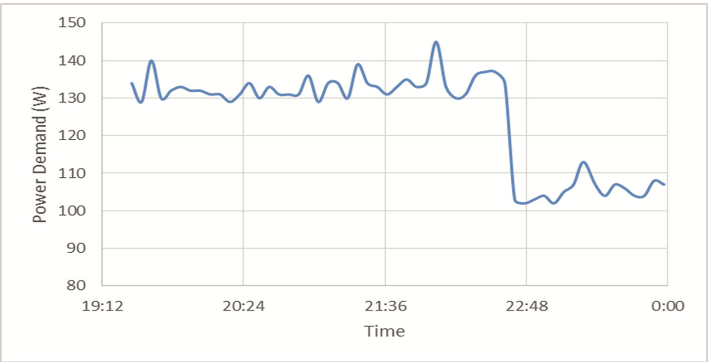

Fig. 3. Test optimization scenario result

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.

Acknowledgments. The work presented in this paper was partially supported by the EU-funded Project FP7 ICT-609140 DOLFIN.

References

1. Bash, C., Forman, G.: Cool job allocation: measuring the power savings of placing jobs at cooling-efficient locations in the Data Centre. In: Proceedings of the 2007 USENIX Annual Technical Conference, January 2007
2. Bohrer, P., Elnozahy, E., Keller, T., Kistler, M., Lefurgy, C., Rajamony, R.: The case for power management in web servers. Power Aware Comput. January 2002
3. Fan, X., Weber, W.-D., Barroso, L.A.: Power provisioning for a warehouse-sized computer. In: Proceedings of the 34th Annual International Symposium on Computer Architecture (2007)
4. Koomey, J.G.: Worldwide electricity used in Data Centres. Environ. Res. Lett. **3**, 034008 (2008). IOP Publishing
5. Dolfin.: Project Number: 609140; Strategic objective: FP7-SMARTCITIES-2013 (ICT-2013.6.2). http://www.dolfin-fp7.eu/
6. Taylor, J.W.: Triple seasonal methods for short-term electricity demand forecasting. Eur. J. Oper. Res. **204**(1), 139–152 (2010)

7. Hippert, H.: Neural networks for short-term load forecasting: a review and evaluation. IEEE Trans. Power Syst. (2001)
8. Taieb, S.B., Hyndman, R.J.: A gradient boosting approach to the Kaggle load forecasting competition. Int. J. Forecast. (2013)
9. Chen, B.-J., Chang, M.-W., Lin, C.-J.: Load forecasting using support vector machines: a study on EUNITE competition 2001. IEEE Trans. Power Syst. **19**(4), 1–7 (2004)
10. Sapankevych, N., Sankar, R.: Time series prediction using support vector machines: a survey. IEEE Comput. Intell. Mag. **4**(2), 24–38 (2009)
11. Humeau, S., Wijaya, T., Vasirani, M., Aberer, K.: Electricity load forecasting for residential customers: exploiting aggregation and correlation between households. In: IEEE Sustainable Internet and ICT for Sustainability (SustainIT), Palermo, 30–31 October 2013
12. Christianni, N., Taylor, J.S.: An Introduction to Support Vector Machines and Other Kernel-Based Learning Methods. Cambridge University Press, New York (2000). Sect. 6.2.2
13. Scikit-learn. http://scikit-learn.org/stable/index.html
14. Openstack. http://www.openstack.org/
15. Opennebula. http://opennebula.org/
16. Eucalyptus. https://www.eucalyptus.com/
17. Garey, M., Johnson, D.: Computers and Intractability. A Guide to the Theory of NP-Completeness. W.H. Freeman and Company, New York (1979)
18. Aravanis, A., Velivassaki, T.H., Voulkidis, A., Zahariadis, T., Cottis, P.: 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
19. Talbi, E.G.: Metaheuristics From Design to Implementation. Wiley, Hoboken (2009)
20. Emanuel, F., Alain, D.: A genetic algorithm for bin packing and line balancing. In: Proceedings of the IEEE 1992 International Conference on Robotics and Automation (RA 1992), 10–15 May 1992, Nice, France (1992)
21. Emanuel, F.: The grouping genetic algorithms - widening the scope of the GAs in JORBEL. Belgian J. Oper. Res. Stat. Comput. Sci. **33**(1, 2), 79–102 (1993)
22. Emanuel, F.: New representation and operators for GAs applied to grouping problems. Evol. Comput. **2**(2), 123–144 (1994)
23. Rahman, A., Liu, X., Kong, F.: A survey on geographic load balancing based data center power management in the smart grid environment. IEEE Commun. Surv. Tutorials **16**(1), 214–233 (2014)
24. Rao, L., Liu, X., Ilic, M.D., Liu, J.: Distributed coordination of internet data centers under multiregional electricity markets. Proc. IEEE **100**(1), 269–282 (2012)
25. Rao, L., Liu, X., Xie, L., Liu, W.: Minimizing electricity cost: optimization of distributed internet data centers in a multi-electricity market environment. In: Proceedings of International Conference on Computer Communications, pp. 1–9. IEEE (2010)
26. Zhang, Y., Wang, Y., Wang, X.: GreenWare: greening cloud-scale data centers to maximize the use of renewable energy. In: Kon, F., Kermarrec, A.-M. (eds.) Middleware 2011. LNCS, vol. 7049, pp. 143–164. Springer, Heidelberg (2011)
27. Giacobbe, M., Celesti, A., Fazio, M., Villari, M., Puliafito, A.: Towards energy management in Cloud federation: a survey in the perspective of future sustainable and cost-saving strategies. Comput. Netw. **91**, 14 (2015)
28. 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, pp. 1–4, 14–15 April 2015

29. Niedermeier, F., et al.: Increasing data centre renewable power share via intelligent smart city power control. In: Proceedings of the 2015 ACM Sixth International Conference on Future Energy Systems. ACM (2015)
30. Wang, H., et al.: Exploring smart grid and data center interactions for electric power load balancing. ACM SIGMETRICS Perform. Eval. Rev. **41**(3), 89–94 (2014)