Minimizing Energy Costs in Federated Datacenters Under Uncertain Green Energy Availability

M. Ruiz*, A. Asensio, and L. Velasco
Optical Communications Group (GCO), Universitat Politècnica de Catalunya (UPC), Barcelona, Spain
**e-mail: mruiz@ac.upc.edu*

ABSTRACT
The cost of energy represents, by far, the largest fraction of total operational expenditures that datacenter operators ought to face. For this very reason, several studies have focused on evaluating how such energy costs can be reduced and on quantifying that reduction; using green energy sources (*e.g.* solar) that can be generated by installing infrastructures nearby datacenters is clearly an interesting option. Assuming that green energy is available, workloads consolidation in those datacenters with the highest amount of self-generated energy allows reducing remarkably the consumption of brown energy. Workload management is of paramount importance to increase green energy consumption in the context of distributed datacenters. In that scenario, a centralized and orchestrated operation leads to large energy cost savings. To this end, we firstly present a model to estimate the amount of green energy produced in each location as a function of the specific time period and the expected weather conditions. Next, the problem of minimizing energy costs by properly placing workloads in federated datacenters under uncertainty in the availability of green energy in each location is faced using stochastic programming techniques. Illustrative numerical results validate the usefulness of the proposed stochastic approach.

Keywords: federated datacenters, green energy, energy minimization, statistical models, stochastic programming.

1. INTRODUCTION

Cloud computing has transformed the IT industry, shaping the way IT hardware is designed and purchased. Datacenters (DC) are equipped with IT resources, including servers and networking devices. Because those electronic devices consume huge amount of energy, energy expenditure becomes a predominant part of total operational expenditures for their operators. Thanks to virtualization, workloads (*e.g.* web applications) can be easily consolidated and placed in the most proper server according to its performance goals. By encapsulating workloads in virtual machines (VM) a DC resource manager can migrate them from one server to another looking for optimizing some objective function, such as energy consumption, whilst ensuring the committed quality of experience (QoE).

Federating DCs can be a way for independent DC operators to, not only increase their revenue, but also reduce operational expenditures. In this regard, network providers can facilitate federated datacenters interconnection by allowing them to automate connection provisioning, thus minimizing communications costs [1], [2]. Aiming at optimizing energy costs whilst ensuring the desired QoE for users, authors in [3] described and formally stated the ELFADO problem to orchestrate federated DCs, placing workloads in the most convenient DC. Taking advantage of *green* energy coming from solar source (replacing either partially or totally energy coming from *brown* polluting sources), energy costs can be substantially reduced. To solve the ELFADO problem, two approaches were compared: the *distributed* approach based on running scheduling algorithms inside DC resource managers and the *centralized* one, where a federation orchestrator is used to compute the global optimal placement for all the VMs in the federated DCs. Results over a worldwide topology showed that energy costs savings around 50% are obtained for the centralized approach.

In this work we present a stochastic programming to solve the centralized approach of the ELFADO problem. We consider that the availability of green energy is a variable not perfectly known in advance and, therefore, its probability distribution must to be taken into account when solving ELFADO. The paper presents two contributions: first, a statistical model to predict the availability of green energy in a DC as a function of the day, hour, and weather conditions is presented. Then, a mixed integer linear programming formulation based on discrete probability scenarios obtained with the statistical model is presented and solved for real-size instances.

2. STATISTICAL MODEL FOR GREEN ENERGY AVAILABILITY

According to [3], the amount of green energy available produced in a DC can be estimated as $g = \alpha * EMD$, where α *is* the green coverage and *EMD* is the amount of energy consumed for the maximum dimensioning. The green coverage is the variable to be estimated, since it is directly related with green energy availability.

The rationale behind our statistical approach can be easily identified after observing Figure 1, where pairs $\{day, \alpha\}$ measured at 2 p.m. are depicted as a function of the day for a DC placed in Berlin (data from references in [3]). In view of the figure, five curves or levels with similar shape but different scale are identified. It is worth noting that weather conditions are the main cause of such differences and, therefore, each curve can be associated with a weather level (plotted with different markers). In light of this, our proposed statistical model for α consists in a set of K curves (one of each weather level) as a function of the day. In addition, a set of coefficients to scale predictions for every hour of the daytime are necessary. With this approach, multiple stationary effects (daily, yearly) are managed in an easier way than other similar time series models [4].

Before applying curve fitting, data must be normalized in order to allow estimating α for any daytime hour. To this aim, a tuple of three coefficients is needed for each hour h: $\{firstDay_h, lastDay_h, peak_h\}$. The first two coefficients store the range of days where α takes non-zero values, while the third stores the maximum α observed. In the example of *Figure 1*, these coefficients are {1, 365, 0.747}. Note that these values are visibly different for different hours: for instance, coefficients at 7 a.m. are {88, 294, 0.255}. Then, all observations are normalized from these coefficients following equations (1) and (2), thus obtaining $\{day_{norm}, \alpha_{norm}\}$ pairs. Figure 2 depicts all normalized pairs for a year (which are in the continuous range [0, 1]) for one specific k level as well as the curve that better fits the average trend, *i.e.* a polynomial of degree 4. Therefore, the statistical model shown in equation (3) consists in a polynomial with β coefficients and a Gaussian error ε (with variance σ^2).

To predict α for a given day d, hour h, and weather level k, the procedure is as follows: *i*) d is normalized using $firstDay_h$ and $lastDay_h$; *ii*) α_{norm} is estimated from equation (3); and *iii*) α is obtained from α_{norm} by reverting equation (2). The probability density function of such estimated value follows a Gaussian distribution with mean α and variance σ^2. Note that an input needed for prediction is the weather level expected for the desired day and hour. When estimations are done between consecutive hours, it is clear that a time series model for predicting next weather level from past observations is a valuable option. In this work we use a basic model consisting in a set of transition probabilities between levels in consecutive hours.

In next section, this methodology is used to generate instances for solving the stochastic approach of ELFADO.

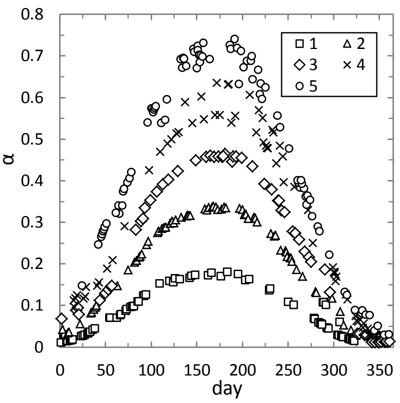

Figure 1. α vs. day for 2 p.m.

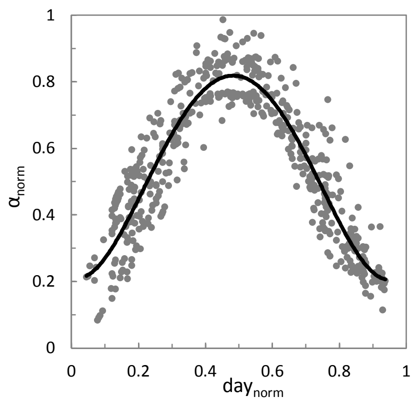

Figure 2. α_{norm} vs. day_{norm} for k = 3.

$$day_{norm} = \frac{day - firstDay_h}{lastDay_h - firstDay_h} \quad (1)$$

$$\alpha_{norm} = \alpha / peak_h \quad (2)$$

$$\alpha_{norm}(k, day_{norm}) = \sum_{i=0..4} b_{ik} \cdot (day_{norm} \wedge i) + \varepsilon(0, \sigma^2) \quad (3)$$

3. PROBLEM STATEMENT AND FORMULATION

The ELFADO problem can be formally stated as follows:

Given: a) a set of federated DC D; b) the set of optical connections E that can be established between two DCs; c) a set of VMs $V(d)$ in each DC d; d) the data volume u_v and number of cores nc_v of each VM v; e) a set of client locations L, where nl_l is the number of users in location l to be served in the next period; f) a threshold th_v for the performance required at any time for accessing the service in VM v; g) the performance q_{ld} perceived in location l when served from a VM placed in DC d; and h) the energy consumption b_d, brown energy cost c_d and green energy availability g_d in DC d for the next period.

Output: the DC where each VM will be placed the next time period;

Objective: Minimize energy and communications cost for the next period ensuring performance requirements.

We propose to solve the stochastic approach of the ELFADO problem by means of discrete probability scenarios [5]. In view of the statistical model presented in Section 2, a scenario for each weather level k can be easily obtained: the green energy available is computed from the k^{th} polynomial model and its probability comes from the transition probability from the current to the next k level. Then, the brown energy needed at each scenario is different and the cost of each scenario is weighted in the objective function by its probability.

The notation needed for the MILP formulation is as follows:

	Sets
$\boldsymbol{D}$	set of federated DC, index d.
$\boldsymbol{E}$	set of optical connections, index e.
$\boldsymbol{V}$	set of VM, index v.
$\boldsymbol{V(d_1)}$	set of VMs in DC d_1.
$\boldsymbol{L}$	set of client locations, index l.
$\boldsymbol{K(d)}$	set of probability scenarios in DC d
	Users and performance
$\boldsymbol{q_{ld}}$	performance perceived in location l when accessing DC d.
$\boldsymbol{nl_l}$	number of users in location l.
$\boldsymbol{th_v}$	the threshold performance to be guaranteed for v.
	DC architecture and VMs
$\boldsymbol{M}$	number of clusters per DC.
$\boldsymbol{ns}$	number of cores per server.
$\boldsymbol{u_v}$	size in bytes of VM v.
$\boldsymbol{nc_v}$	number of cores needed by VM v.

	Connections
$\boldsymbol{U_e}$	capacity assigned in connection e.
$\boldsymbol{c_e}$	cost per Gb transmitted through connection e.
	Energy
$\boldsymbol{b_d}$	Total energy consumption in DC d
$\boldsymbol{PUE_d}$	PUE for DC d.
$\boldsymbol{c_d}$	brown energy cost per kWh in DC d.
$\boldsymbol{g_{dk}}$	amount of green energy available in DC d under probability scenario k
$\boldsymbol{p_{dk}}$	probability of scenario k in DC d
	Decision variables
$\boldsymbol{x_{vd}}$	binary, 1 if VM v is placed in DC d, 0 otherwise.
$\boldsymbol{y_{dk}}$	real positive, energy consumption in DC d in probability scenario k
$\boldsymbol{z_e}$	integer positive, bytes to transfer through optical connection e.
$\boldsymbol{\gamma_d}$	positive integer with the number of servers operating with some load in DC d.
$\boldsymbol{\rho_d}$	positive integer with the number of clusters switched on in DC d.

Finally, the formulation based on introducing discrete probability scenarios is as follows:

$$(ELFADO)\ \text{minimize} \sum_{d\in D} c_d \sum_{k\in K(d)} p_{dk}\cdot y_{dk} + \sum_{e\in E} c_e \cdot z_e \quad (4)$$

$$\frac{1}{\sum_{l\in L} nl_l}\cdot \sum_{l\in L}\sum_{d\in D} nl_l \cdot q_{ld}\cdot x_{vd} \le th_v \quad \forall v\in V \quad (5)$$

$$\sum_{d\in D} x_{vd} = 1 \quad \forall v\in V \quad (6)$$

$$\gamma_d \ge \frac{1}{ns}\cdot \sum_{v\in V} nc_v \cdot x_{vd} \quad \forall d\in D \quad (7)$$

$$\rho_d \ge \frac{4}{M^2}\cdot \gamma_d \quad \forall d\in D \quad (8)$$

$$y_{dk} \ge b_d - g_{dk} \quad \forall d\in D, k\in K(d) \quad (9)$$

$$z_{e=(d_1,d_2)} = \sum_{v\in V(d_1)} u_v \cdot x_{vd2} \quad \forall d_1, d_2 \in D, d_1 \ne d_2 \quad (10)$$

$$z_e \le U_e \quad \forall e\in E \quad (11)$$

The objective function (4) minimizes the total cost for all DC in the federation, which consists on the energy costs (weighted by the probability of each scenario) plus the communication costs for the VMs that are moved between DCs. Constraint (5) guarantees that each VM is assigned to a DC if the on-average performance perceived by the users is above the given threshold. Constraint (6) ensures that each VM is assigned to one DC. Constraint (7) computes, for each DC, the amount of servers where some VM is to be placed, whereas constraint (8) computes the number of clusters that will be switched on. Constraint (9) computes the brown energy consumption for each scenario in each DC as the difference between the effective energy consumption (assuming that DCs follow a flat-tree architecture [3]) and the amount of green energy available in the next period in each DC. Constraint (10) computes the amount of data to be transfer from each DC to some other remote DC. Finally, constraint (11) assures that the capacity of each optical connection is not exceeded.

In the next section, the performance of ELFADO is evaluated from a European-wide DC topology.

4. ILLUSTRATIVE NUMERICAL RESULTS AND CONCLUSIONS

For evaluation we consider the European topology depicted in Figure 3 consisting of 10 DC each with a number of clusters $M = 10$. To keep the DC resource utilization in a moderated value (around 25%) 5000 VMs are distributed among DCs. The values of power consumption and costs at each location have been derived from references in [3].

The ELFADO problem has been solved with two different schemes: *stochastic*, where one scenario per weather level and DC are considered (total number of 50 scenarios), with expected g_{dk} and p_{dk} obtained from the

Figure 3. European network.

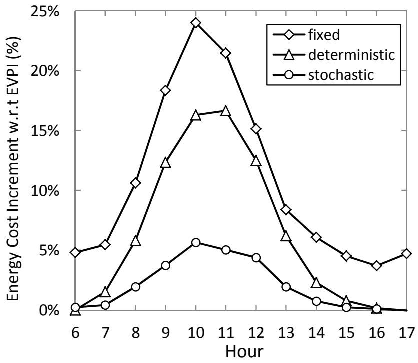

Figure 4. Energy cost w.r.t. EVPI.

TABLE 1 RELATIVE COST STATISTICS

	fixed	det.	sto.
average	12.7%	8.3%	**2.7%**
max	24.0%	16.7%	**5.7%**

TABLE 2 RELATIVE SAVINGS STATISTICS

	det.	sto.
average	43.5%	**81.8%**
max	82.2%	**94.9%**

proposed statistical model; and *deterministic*, where only one scenario per DC with the expected average green energy availability is assumed. Moreover, a *fixed* scheme, where the total workload is evenly distributed among the federated DCs (*i.e.* no migration is performed), is considered for comparison. In order to evaluate the quality of the obtained solutions, we computed the *Expected Value of Perfect Information* (*EVPI*), which assumes that the real amount of green energy is perfectly known in advance, being this the benchmark for computing energy costs and savings.

We solved the ELFADO problem (using CPLEX) at every hour of 120 consecutive days of spring and summer (total of 2880 instances). In order to estimate likely costs of each ELFADO solution, a simulation is run to obtain the green energy availability from the probability distribution obtained with the α statistical model (thus emulating a real-life behavior) and costs are eventually computed from that simulated green energy values.

Figure 4 shows the energy cost increment with respect to the cost of EVPI solutions as a function of daytime hours. Moreover, Table 1 presents some statistics extracted from Figure 4. As can be observed, the main differences among strategies occur in central hours, when α takes highest values and also fluctuations induced by weather variations are large. It is worth noting that the fixed strategy is the most expensive (up to 24% w.r.t EVPI), whereas both deterministic and stochastic reduce this increment. However, it is clear than the stochastic approach allows reaching values much closer to EVPI than the deterministic scheme. Specifically, a maximum difference around 5% is observed, which aims concluding that using few discrete probability scenarios is enough to reach very high-quality solutions.

An alternative analysis can be done in terms of energy cost savings with respect to the fixed approach (values in Table 2). We assume than by applying the fixed strategy, obviously energy savings equal to 0%. Moreover, EVPI solutions will lead to the 100% of savings. So, both deterministic and stochastic approaches will range between these two values. In light of the results, we can conclude than the stochastic approach almost doubles the average savings obtained with the deterministic one (which are in line with results in [3])

As final conclusion, we demonstrated that the presented statistical model for green coverage estimation allow easily generating discrete scenarios of green energy availability to be included in a linear programming formulation with the aim of solving the stochastic approach of the ELFADO problem. Results shown that applying the stochastic approach instead of the deterministic one leads to a significant reduction of costs and relative savings close to the values obtained if future was known in advance.

ACKNOWLEDGEMENTS

The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement n° 317999 IDEALIST project. Moreover, it was supported by the Spanish science ministry through the TEC2011-27310 ELASTIC project. The authors want to thank Mr. Jorge León for his very valuable contribution to this work.

REFERENCES

[1] L. Velasco *et al.*, "Towards a carrier SDN: An example for elastic inter-datacenter connectivity," (Invited Paper) *OSA Optics Express*, vol. 22, pp. 55-61, 2014.

[2] L. Velasco *et al.*, "Cross-stratum orchestration and flexgrid optical networks for datacenter federations," *IEEE Network Magazine*, vol. 27, pp. 23-30, 2013.

[3] L. Velasco *et al.*, "Elastic operations in federated datacenters for performance and cost optimization," *Elsevier Computer Communications*, 2014.

[4] M. Anastasopoulos, A. Tzanakaki, and D. Simeonidou, "Stochastic service provisioning in converged optical data center networks and smart grids," in *Proc. ONDM*, 2013.

[5] A. Shapiro *et al.*, "Lectures on Stochastic Programming: Modeling and Theory, " MPS-SIAM, 2009.