Efficient workload management in geographically distributed data centers leveraging autoregressive models

Albino Altomare, Eugenio Cesario, and Carlo Mastroianni

Citation: AIP Conference Proceedings **1776**, 080004 (2016); doi: 10.1063/1.4965361
View online: http://dx.doi.org/10.1063/1.4965361
View Table of Contents: http://scitation.aip.org/content/aip/proceeding/aipcp/1776?ver=pdfcov
Published by the AIP Publishing

Efficient Workload Management in Geographically Distributed Data Centers Leveraging Autoregressive Models

Albino Altomare, Eugenio Cesario[a] and Carlo Mastroianni

Institute for High Performance Computing and Networking (ICAR-CNR).
Via P. Bucci Cubo 7-11b, 87036 Rende (CS), Italy.

[a]Corresponding author: cesario@icar.cnr.it
e-mail: {altomare,cesario,mastroianni}@icar.cnr.it

Abstract. The opportunity of using Cloud resources on a pay-as-you-go basis and the availability of powerful data centers and high bandwidth connections are speeding up the success and popularity of Cloud systems, which is making on-demand computing a common practice for enterprises and scientific communities. The reasons for this success include natural business distribution, the need for high availability and disaster tolerance, the sheer size of their computational infrastructure, and/or the desire to provide uniform access times to the infrastructure from widely distributed client sites. Nevertheless, the expansion of large data centers is resulting in a huge rise of electrical power consumed by hardware facilities and cooling systems. The geographical distribution of data centers is becoming an opportunity: the variability of electricity prices, environmental conditions and client requests, both from site to site and with time, makes it possible to intelligently and dynamically (re)distribute the computational workload and achieve as diverse business goals as: the reduction of costs, energy consumption and carbon emissions, the satisfaction of performance constraints, the adherence to Service Level Agreement established with users, etc. This paper proposes an approach that helps to achieve the business goals established by the data center administrators. The workload distribution is driven by a fitness function, evaluated for each data center, which weighs some key parameters related to business objectives, among which, the price of electricity, the carbon emission rate, the balance of load among the data centers etc. For example, the energy costs can be reduced by using a "follow the moon" approach, e.g. by migrating the workload to data centers where the price of electricity is lower at that time. Our approach uses data about historical usage of the data centers and data about environmental conditions to predict, with the help of regressive models, the values of the parameters of the fitness function, and then to appropriately tune the weights assigned to the parameters in accordance to the business goals. Preliminary experimental results, presented in this paper, show encouraging benefits.

BUSINESS GOALS IN DISTRIBUTED DATA CENTERS

The administrators of a geographically distributed data center can exploit the variability of conditions both from site to site and temporally, to achieve a number of different business goals. The variability involves the price of electricity, the weather conditions, the availability of green energy, etc. Some of the most critical business objectives are:

- *Reduction of consumed energy.* Moderns data centers are equipped with instrumentation to monitor the energy consumed in computational resources. The total energy, including that needed for cooling and power distribution, is obtained by multiplying the power used for computation by the PUE (Power Usage Efficiency) index;
- *Reduction of energy costs.* The cost of electricity is generally different from site to site and also varies with time, even on a hour-to-hour basis, therefore the overall cost may be reduced by shifting portions of the workload to more convenient sites;
- *Reduction of carbon emissions.* Companies are today strongly encouraged to reduce the amount of carbon emissions, not only to compel to laws and rules, but also to advertise their green effort and attract customers that are increasingly careful about sustainability issues.

All the above mentioned goals are important, yet different data centers may focus on different aspects (i.e., depending on the specific operating conditions and on the priorities prescribed by the management) and it is up to the company's management to specify the objectives and their relative weights. Dealing with this issue, in [1] a

Numerical Computations: Theory and Algorithms (NUMTA–2016)
AIP Conf. Proc. 1776, 080004-1–080004-4; doi: 10.1063/1.4965361
Published by AIP Publishing. 978-0-7354-1438-9/$30.00

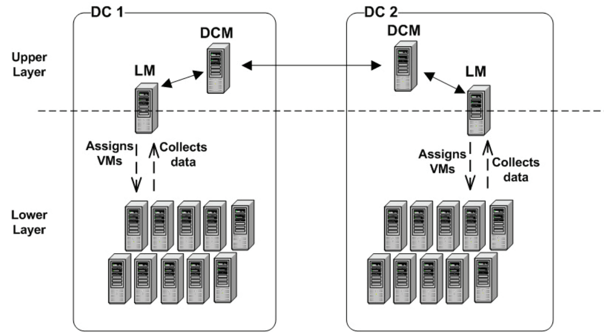

FIGURE 1. Architecture for the management of distributed data centers.

hierarchical architecture is illustrated, which aims at supporting the efficient management of the workload in a multi-site scenario. The architecture is depicted in Figure 1 for the case of two interconnected data centers. Each data center (DC) is composed of two layers: the *upper layer* and the *lower layer*. The upper layer is handled by a data center manager (DCM) that exchanges information among the different sites and takes decisions about the migration of workload from one site to another, combining the information related to single DCs. On the other side the local manager (LM), in charge of handling the functionalities of the lower layer, manages the internal workload using its own consolidation algorithm and autonomously decides whether or not to accommodate a VM (i.e., virtual machine) or trigger a VM migration within the local DC.

Since the single DCs are autonomous regarding the choice of the internal algorithms for workload management, the focus here is on the algorithms of the upper layer. Two basic algorithms are executed at each DCM: (*i*) the *assignment algorithm* that determines the appropriate target DC for each new VM; (*ii*) the *migration algorithm* that periodically evaluates whether the current load distribution is appropriate, decides whether an amount of workload should be migrated and, if so, determines from which source site to which target site.

ALGORITHMS FOR WORKLOAD ASSIGNMENT AND REDISTRIBUTION

In this section, we describe the two basics algorithms executed by the DCMs, i.e. the assignment and migration algorithms, and some strategies based on predictive data mining that help improve the effectiveness of the approach. In particular, the objective is to appropriately tune the weights of the parameters related to the different business objectives, so as to achieve the desired goals.

Assignment Algorithm. The optimal distribution of the workload among the data centers is driven by a purposely defined *assignment function*, which balances and weighs the chosen business goals. This function associates to each DC a value that represents the cost to run some workload in that DC, with low values corresponding to low overall cost of the DC. To balance the load among the data centers, the function also includes the load currently assigned to each data center. The strategy, then, is to assign each new VM to the DC with the lowest value of the function. For example, if the objectives are the reduction of consumed energy, the minimization of carbon emissions and the minimization of costs related to energy, the assignment function f^i_{assign}, for each DC i, is defined as follows:

$$f^i_{assign} = \alpha \cdot \frac{C_i}{C_{max}} + \beta \cdot \frac{PUE_i}{PUE_{max}} + \gamma \cdot \frac{E_i}{E_{max}} + k \cdot \frac{U_i}{U_{max}} \tag{1}$$

where the coefficients α, β, γ and k are positive and $\alpha + \beta + \gamma + k = 1$. The values C_i, PUE_i, E_i and U_i state for carbon emissions, achieved PUE, energy costs and relative load (fraction of the capacity of the data center that is currently loaded) at the i^{th} DC, respectively. The parameters are normalized with respect to the maximum values communicated by DCs. The mentioned goals are weighed through the values of the coefficients. After computing the values of f_{assign} for each DC, the VM is assigned to the data center having the lowest value. Once consigned to the target DC, the VM is allocated to a physical host using the local assignment algorithm, that is out of the scope of this paper.

Migration Algorithm. The assignment algorithm optimizes the distribution of the VMs on the basis of the chosen objectives and their respective weights. Since a data center with a lower value of f_{assign} attracts more load, and additional load will increase the value of f_{assign}, the function tends to have the same value at all the data centers. However, this balance can be altered when the conditions change, e.g., when the price of energy or the PUE varies in one or more data centers. In such cases, inter-DC VM migrations are performed to redistribute the workload. The migration algorithm is triggered when the values of the f_{assign} functions of two DCs differ by more than a predetermined threshold (e.g., 3%). In such a case, VMs are migrated from the data center having the highest value of f_{assign} to the data center with the minimum value, until the values reenter within the tolerance range.

PREDICTIVE MODELS FOR OPTIMAL WORKLOAD MANAGEMENT

The effectiveness of the f_{assign} formula described in the previous section completely depends both (*i*) on a good forecasting of the values will be assumed by C_i, PUE_i and E_i and (*ii*) on the optimal setting of the f_{assign} coefficients. The values of the coefficients α, β and γ represent de facto the weights to achieve the corresponding mentioned goals, and their best value setting is a hard task for the DCM. To deal with these two aspects, we describe an approach aimed at improving the outcome of the f_{assign} formula and driving the described hierarchical architecture towards an optimal distribution of the workload. The general idea of the approach, depicted in Figure 2, includes two steps. First, given the availability of historic environmental data (i.e., price of electricity, green energy availability, etc.) and business value trends (i.e., carbon emissions, energy costs, PUEs, consumed energy, etc.), the *Analytics Modeling* component processes such data to infer knowledge models for business value forecasts. Second, taking into account the values predicted at the first step, the *Weights Optimizer* module automatically tunes the values of the f_{assign} coefficients to achieve the business goals.

The *Analytics Modeling* component performs C_i, PUE_i and E_i forecasting through ARIMA models, that are a combination of differencing, auto-regression and moving average methods with the goal of forecasting a variable of interest using a linear combination of predictors [2]. As an example, let us consider the auto-regression of the PUE time series $\{y_t : t = 1...n\}$, where y_t is the value of the time series (the value of PUE) at the timestamp t. Then, an $ARIMA(p, d, q)$ model is written in the form

$$y_t^{(d)} = c + \phi_1 y_{t-1}^{(d)} + \ldots + \phi_p y_{t-p}^{(d)} + \theta_1 e_{t-1} + \ldots + \theta_q e_{t-q} + e_t \tag{2}$$

where:

- $y_t^{(d)}$ is the d^{th}-differenced series of y_t, that is: $y_t^{(d)} = y_t^{(d-1)} - y_{t-1}^{(d-1)}$, ... , $y_{t-p}^{(d)} = y_{t-p}^{(d-1)} - y_{t-p-1}^{(d-1)}$;
- $\phi_1, \ldots, \phi_p$ and $\theta_1, \ldots, \theta_q$ are the regression coefficients of the auto-regressive and of the moving average part, respectively;
- $e_{t-1}, \ldots, e_{t-q}$ are lagged errors, e_t is the white noise and c is a correcting factor.

The regression model above described is referred as $ARIMA(p, d, q)$, where the order of the model is stated by three parameters: p (order of the auto-regressive part), d (degree of first differencing involved) and q (order of the moving average part). The best parameter values are obtained by minimizing the *BIC* (*Bayesian Information Criterion*) and *AIC* (*Akaike Information Criterion*) measures, as described in [2].

Now, let us provide also some considerations about the optimal setting of the f_{assign} coefficients. Even though this research topic is still a work-in-progress, we outline here an approach to be further analyzed in detail. As aforementioned, the coefficients represent de facto the weights assigned to the corresponding business goals, and their optimal value settings is a crucial task for the human administrator. Therefore, it is useful to design a methodological framework (e.g., based on a genetic algorithm, or a classification algorithm, etc.) that can automatically give hints in their value setting. In particular, given a set of business goals (i.e., maintaining the overall carbon emissions under a given threshold, limiting the overall utilization in a predefined range values, etc.) and the values C_i, PUE_i and E_i forecasted

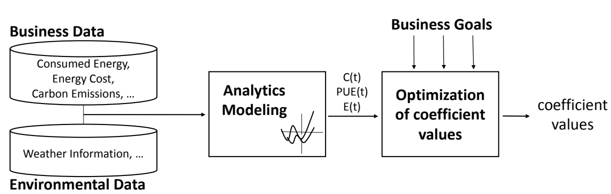

FIGURE 2. Prediction framework.

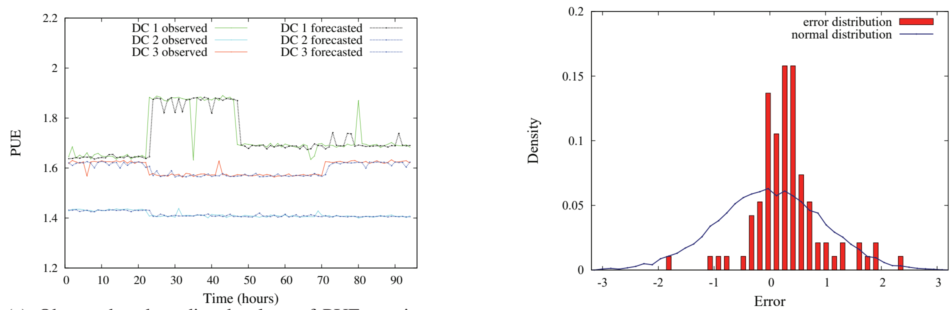

(a) Observed and predicted values of PUE vs. time, for three DCs.

(b) Distribution of the residuals (with the overlaid normal curve) on the test set.

FIGURE 3. Trends and residuals of forecasted and observed values of PUE.

by the ARIMA models, the framework offers the user some rules to tune the values of the coefficients so as to achieve the business goals.

EXPERIMENTAL EVALUATION

This section presents a preliminary evaluation of the regressive function, modeled by Equation 2, performed over synthetic data. In particular, we used an ad-hoc data generator [3] to produce PUE data series for three geographically distributed DCs, each one adhering to a predefined pattern (introducing a 5% noise). Then, the whole dataset (of each DC) has been split in training set and test set: the first one has been used to infer the regressive function, while the second one has been exploited as test bed to validate the approach. Figure 3(a) shows six curves representing, for a temporal window of the test set, the observed and forecasted data of the three DC. It is interesting to highlight that forecasted data adhere very well to the observed data and that fluctuations are well modeled and predicted by the regressive model. As forecasting error, we measured an average error of 4%, that appears to be an encouraging result. Finally, Figure 3(b) shows the distribution of the residuals (with overlaid normal curve with mean 0). The distribution of forecast errors is centered on zero and normally distributed, which is another good property of forecasting accuracy.

CONCLUSION AND FUTURE WORK

This paper proposes an approach to optimize the workload management among geographically distributed data centers, driven by a fitness function which weighs some key parameters related to business objectives. In particular, historical data about both DC resource usage and environmental conditions are exploited to predict, with the help of regressive models, the values of the parameters of the fitness function, and then to appropriately tune the weights assigned to the parameters in accordance to the business goals. Preliminary experimental results are encouraging.

ACKNOWLEDGMENTS

This research was partially supported by the MIUR project DOMUS (PON0050_2).

REFERENCES

[1] A. Forestiero, C. Mastroianni, M. Meo, G. Papuzzo, and M. Sheikhalishahi, "Hierarchical approach for green workload management in distributed data centers," in *Euro-Par 2014: Parallel Processing Workshops* (2014), pp. 323–334.

[2] R. J. Hyndman and G. Athanasopoulos, *Forecasting: principles and practice* (OTexts.com, 2014).

[3] A. Altomare and E. Cesario, "Predictive models for energy-efficient clouds: An analysis on real-life and synthetic data," in *14th IEEE International Conference on Ubiquitous Computing and Communications, IUCC 2015* (2015), pp. 1538–1543.