A Robust Optimization for Proactive Energy Management in Virtualized Data Centers

Ibrahim Takouna* Wesam Dawoud* Kai Sachs± Christoph Meinel*

* Hasso Plattner Institute
University of Potsdam
Potsdam, Germany
firstname.lastname@hpi.uni-potsdam.de

±SAP AG
Walldorf, Germany
kai.sachs@sap.com

ABSTRACT

Energy management has become a significant concern in data centers to reduce operational costs and maintain systems' reliability. Using virtualization allows server consolidation, which increases server utilization and reduces energy consumption by turning off unused servers. However, server consolidation and turning off servers can cause also consequences if they are not exploited efficiently. For instance, many researchers consider a deterministic demand for capacity planning, but the demand is always subject to uncertainty. This uncertainty is an outcome of the workload prediction and the workload fluctuation. This paper presents a robust optimization for proactive capacity planning. We do not presume that the demand of VMs is deterministic. Thus, we implement a range prediction approach instead of a single point prediction. Then, we implement a robust optimization model exploiting the range-based prediction to determine the number of active servers for each capacity planning period. The results of the simulation show that our approach can mitigate undesirable changes in the power-state of the servers. Additionally, the results indicate an increase in the servers' availability for hosting new VMs and reliability against a system failure during power-state changes. As future work, we intend to apply our approach to dynamic workload such as a web application. We plan to investigate applying our approach to other resources, where we consider only the CPU demand of VMs. Finally, we compare our approach against the approaches using stochastic optimization.

Categories and Subject Descriptors

k.6 [**Management of computing and information systems**]: General; k.6.2 [**Installation Management**]: Performance and Usage Measurement

General Terms

Management, Performance, Experimentation

Keywords

Energy-aware; Robust Optimization; Prediction; Virtualization

1. INTRODUCTION

Energy efficient resource management has become a significant concern in virtualized data centers to reduce operational costs. As an idle server consumes approximately 70% of the power consumed by the server running at full capacity [1], turning off idle servers to reduce energy consumption has been widely proposed by many researchers [1][2][3]. However, these approaches considered the number of VMs and their capacity demands are deterministic. Hence, they built deterministic models that did not take into account the uncertainty of the demand. Furthermore, a typical server has multiple power states including on, off, sleep, and hibernated. Some related work has ignored the energy consumption during the change of the server's power-state, which we consider as a wasted energy. During a power-state change, a server consumes energy without performing any useful work. For instance, a normal PC takes around 25 seconds to switch from one state to off state and vice versa [4]. Furthermore, we found that a server with 1TB memory requires 5 minutes for a clean boot, which includes the hardware check stage. On the other hand, Mao et al. [5] have observed that the start-up time of a VM is proportional to its image size. The start-up time of a VM is very crucial for online applications such as web applications. Thus, implementing a proactive optimization solution can assist to avoid SLA's violations. By predicting the number of the required VMs in the next planning period, we can prepare these VMs images and the physical server in advance. Many researchers consider a deterministic demand for capacity planning, but the demand is always subject to uncertain. This uncertainty is an outcome of the workload prediction and the workload fluctuation. Ignoring the uncertainty in real world applications can make the usual optimal solution infeasible [7]. Unlike deterministic models, Dance et al. [8] have used the stochastic optimization for considering uncertainty. However, the stochastic optimization requires to know the probability distributions of the demand. In our approach, we used the robust optimization that addresses data uncertainty. The robust optimization model assumes that uncertain parameters belong to a bounded range. Importantly, the robust optimization can outperform the stochastic optimization when selecting an appropriate robustness level, and it is less computationally intensive

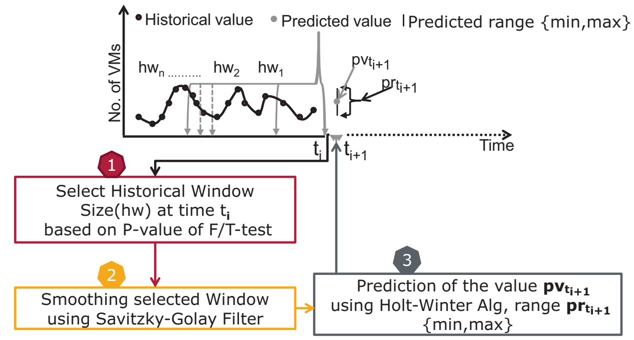

Figure 1: The stages of the range-based prediction approach

when the distribution of uncertain parameters is complicated [9]. To build the robust optimization model for proactive capacity planning, we implement an adaptive range-based workload prediction instead single point prediction for predicting the number of the requested VMs. The results of the simulation show that our approach can mitigate undesirable changes in the power-state of the servers. Additionally, the results indicate an increase in the servers' availability for hosting new VMs and reliability against a system failure during power-state changes.

2. RANGE-BASED PREDICTION

Typically, point value prediction approaches might not cover the workload fluctuation. The approaches solve the problem as a deterministic optimization, which assume the precise knowledge of the workload demand. Furthermore, optimization based on the mean-value or the max-value of the workload can produce low provision or high provision, respectively. This is costly in both cases. In this section, we present a range-based prediction approach with an adaptive window-size algorithm to predict the number of demanding VMs in a data center. Our approach consists of three stages as shown in Figure 1: (1) selecting the historical window size based on the statistical test (F/T-test); (2) smoothing the values of the selected historical window; and (3) predicting the next number of active VMs and its minimum and maximum range. Most of the related work in the context have been done for grid computing [10][11]. For example, Wu et al. [10] have proposed an adaptive prediction of grid performance with a confidence window for the historical values. They used an auto-regression to find a model for the historical interval by which predicts the future workload. However, we consider the historical data variation to enhance the prediction accuracy and bound the predicted range. As shown in Figure 2, the number of VMs shows a random behavior, but it also depicts a certain pattern. For instance, it starts low at morning then increases reaching the peak at around the midday. At evening, it starts to go down again. Our approach uses an adaptive window-size of historical values to provide a high accurate prediction range. In Figure 1, the measured workload values are shown by a series of line-dots up to time t. On the other hand, the gray dot represents the predicted workload value.

- Window selection: our interest is to predict the number of VMs for the next 5 minutes from the historical window hw. The historical window-size is determined based on the P-value of both F-test and T-test to filter out the values that are very unlikely to be in the same window. We used F-test and T-test to probe the significance of the change in variance and mean between two samples of populations, respectively. Using the P-value of F-test and T-test, we can decide whether the two samples have almost the same variance and the same mean. For example, after performing F-test, if we find out that the P-value is less than $\alpha = 0.05$, we reject the null hypothesis. This means that these values do not belong to the same historical window. Thus, the algorithm stops going back to take more historical values and moves to the next stage, which is window smoothing.

- Window smoothing: using prediction algorithms with the historical values causes errors. Thus, we used a smoothing filter to remove noise and prevent its influence on the prediction algorithm. There are many smoothing filters, but we selected Savitzky-Golay filter due its effectiveness in keeping the peak values and removing the spikes, which can be considered as noise. The filter has two significant parameters that guide the smoothing process: the frame size and polynomial degree. In our approach, the frame size is not constant, and it equals to the selected historical window-size. In contrast, Wu et al. [10] fixed the frame size.

- Range prediction: we used Holt-Winter implemented in the R-tool, because it dynamically optimizes and determines the level and the trend of the time series. Importantly, we determine the predicted range based on the single predicted point value pv and the standard deviation of the selected window σ_{hw}. The predicated range pr $\{r_l\ ,\ r_h\}$ where r_l and r_h equal $(pv - \sigma_{hw})$ and $(pv + \sigma_{hw}.)$

2.1 Planet-lab Workload Traces

The monitoring infrastructure project of Planet-lab [6] provides traces of historical data for CPU utilization, which measured every 5 minutes. These traces are for more than a thousand machines running in more than 500 locations around the world. Here, we present data for two days that have different workload fluctuations. We assume that these machines are hosted as VMs.

To simulate the on-demand concept of the cloud computing environment (i.e., the open system behavior), we terminate the VMs with less than 5% CPU utilization. In other words, we considered it as being destroyed and exited the system. Then, when the trace shows a VM with a CPU

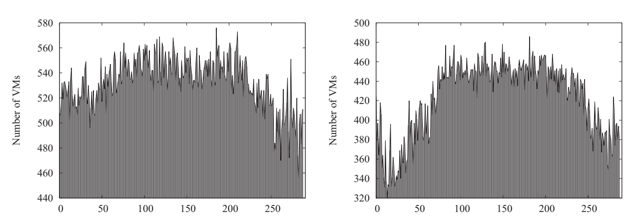

(a) Day03 (b) Day06

Figure 2: Planet-lab workload traces

utilization higher than 5%, we consider a new request for provisioning a VM.

2.2 Implementation and Evaluation

We implemented the proposed approach using Java with integration of the R-tool, which consists of many statistical functions and the required filters. Here, we present the results of our approach. Figure 3 shows the predicted range for each value of workload (i.e., number of VMs). The low predicted r_l is shown by a red dashed-line meanwhile the high predicted r_h is represented by a blue dashed-line. The purple sold-line represents the single point predicted value. The predicted range is propositional to the workload fluctuation. For instance, due to the low fluctuation of the workload around the time index 58, the predicted range is small. Contrarily, a large range is predicted around the time index 225.

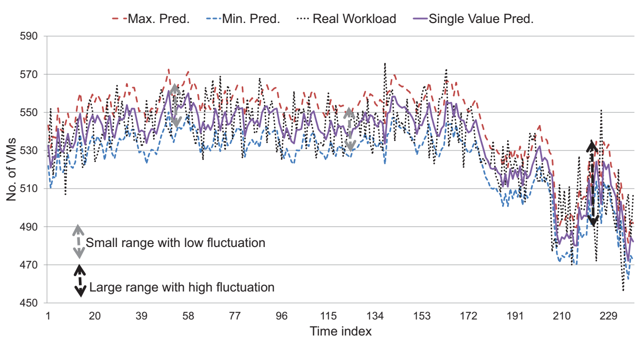

Figure 3: The results of the range-based prediction approach

3. ROBUST OPTIMIZATION

As we implemented the prediction approach based on a range not a single value, we will study the implementation of robust optimization for capacity planning. Robust optimization deals with optimization problems whereas the robustness is sought against uncertainty or deterministic variability in the value of a parameter of the problem (i.e., the workload). The principle of robust optimization considers point prediction meaningless and it is replaced by range prediction. Thus, robust optimization addresses data uncertainty by assuming that uncertain parameters belong to a bounded range. In our approach, we avoid the assumption that considers the precise knowledge of the workload demand in the planning horizon where many proposed solutions have solved the problem as a deterministic optimization [1][2][3]. We have a predicted range, centered at the nominal prediction $\bar{d}$, for the demand at each time period. The robust optimization approach replaces each deterministic demand $\bar{d}$ by an uncertain parameter $\tilde{d} = \bar{d}$ + $\hat{d}$*z, where $|z| \leq 1$. Furthermore, it guarantees that the constraints hold for a given uncertainty set. cient of its constraint matrix. to solve this problem

3.1 Robust Problem Formulation

In this section, we present a robust formulation of the capacity planning problem. We use the following notation: e_{idle} is the energy consumption of a server during idle state (i.e., p_{idle}*t_{idle}). The parameter e_{swt} is the energy consumption of a server to change from power-state to another, which consumes p_{swt} and takes t_{swt} seconds. So, e_{swt} the power-state change energy wastage equals p_{swt}*t_{swt}. The binary variables n and b represent currently active servers and previously active servers, respectively. The parameters $\tilde{nv}$, $\bar{nv}$, and $\hat{nv}$ represent the uncertain number of VMs, the mean of the predicted number of VMs, and the standard deviation of the predicted number of VMs, respectively. The parameters $\tilde{vc}$, $\bar{vc}$, and $\hat{vc}$ represent the uncertain utilization of a VM, the mean of the VM's utilization, and the standard deviation of the VM's utilization, respectively. Finally, the parameter $\tilde{d}$ is the total uncertain demand of the number of VMs and their utilization. The objective function in Equation 1 is to minimize the wastage energy that might result from keeping the server idle and switching the power-state of a server. Equation 2 guarantees that the constraints hold for a given uncertainty set of the demand in Equation 3. Equation 4 and Equation 5 represent uncertain demand of a number of VMs and uncertain utilization of VMs, receptively. We formulate this problem based on the following assumption. When a VM is requested, it occupies a certain portion of a server capacity. For instance, a small-instance with 1 vCPU might take 1/4 of a server has 4 logical cores. Then, after running the VM for while the real resource consumption of this VM can be revealed and will be taken into consideration for the next planning period. The constant parameters t_{swt}, t_{idle}, p_{swt} and p_{idle} are 150s, 300s, 120watts, 100watts, respectively. Observably, the switching power is slightly greater than the idle power due to the CPU utilization. The power constants were set based on SPECpower [12] results of HP ProLiant ML110 G3 server [13]. A scalar variable z models the demand uncertainty. We do not presume exact knowledge of the actual distribution of demand, but instead we assume that the distribution is characterized by the mean of the number of VMs $\bar{nv}$ and their capacity $\bar{vc}$ and standard deviation of the demand of the number of VMs $\hat{nv}$ and their capacity $\hat{vc}$. Moreover, we have an accurate estimations of the most optimistic uncertainty z_{min} and the most pessimistic uncertainty z_{max}. These parameters form lower and upper bounds of z, respectively.

Minimize:

$$\sum_{i=1}^{ns} e_{idle} + \sum_{i=1}^{ns} e_{swt} * (n * (1-b) + b(1-n)) \tag{1}$$

Subject to:

$$\tilde{d} \leq \sum_{i=1}^{ns} \tilde{sc}_i * n \tag{2}$$

$$\tilde{d} = \sum_{i=1}^{nv} \tilde{vc}_i \tag{3}$$

$$\tilde{nv} = \bar{nv} + \hat{nv} * z_{nv} \tag{4}$$

$$\tilde{vc} = \bar{vc} + \hat{vu} z_{vc} \tag{5}$$

$$|z| \leq 1, \quad z = [z_{min}, z_{max}], \quad and \quad n, b \in \{0, 1\}$$

3.2 Implementation and Evaluation

To solve a robust optimization problem, we used Robust Optimization Made Easy (ROME), which is an algebraic modeling tool implemented in the MATLAB environment [14]. ROME operates as an intermediate layer between the modeler and optimization solver engines. It helps converting the original uncertain optimization problem into its robust counterparts. Its core functionality consists of translating modeling code into an internal structure in ROME. Then, it translated into a solver-specific input format for solving by linear optimization solvers. This can be done manually, but it is tedious and error-prone [14]. We used IBM ILOG CPLEX as optimization solver.

Figure 4 depicts simulation results of a deterministic optimization and different values of uncertainty scalar z. The lift axis represents the amount of energy and the right axis represents the percentage of the energy consumption by the power-state switching and the idle state. The deterministic result means that prediction of the VMs demand is 100% accurate. So, we used the real traces to perform capacity planning using deterministic optimization. Then, we compared the results with the range-based prediction taking into account uncertainty of demand. This shows the results with different range of uncertainty scalar z. Figure 4 depicts the following observations. In deterministic, there is no idle energy consumption, because we assumed that the demand is known, and the total wastage energy results from changing a server's power-state. On the other hand, the robust optimization considering the uncertainty of the demand, we had different results by changing the uncertainty scalar z. First, the range of uncertainty is very wide, z{-1,1}. The total of the objective function and the total of idle energy are the highest, and the total of switching energy is the lowest compared to the other results. Second, when changing the uncertainty scalar z from z{-1,1} to z{-0.5,0.5}, we decreased the uncertainty range. Thus, we could save energy by reducing an unbeneficial power-state switching and keeping some server idle. This also can increase the system availability and reliability. Finally, when z was set {-0.25,0.25}, we could achieve more energy savings from both power-state switching and idle servers. However, this can cause some under provisioning of capacity and violation of allocation some VMs. The calculated mean and standard deviation of under provisioning VMs are 3VMs and 5VMs, respectively. Importantly, the execution time of our proposed approach is less than 1 second while using a computer with Pentium 2.6GHz processor.

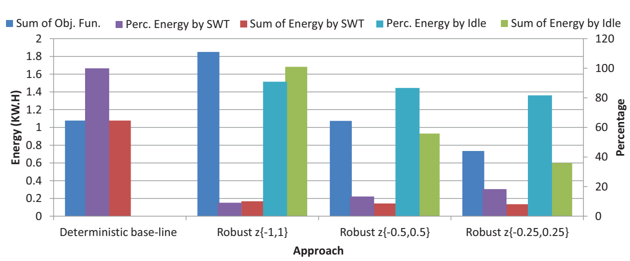

Figure 4: The result of the proposed robust optimization approach

4. CONCLUSIONS AND FUTURE WORK

In this paper, we presented a proactive robust optimization approach for capacity planning in virtualized data centers. To achieve this, we implemented a range-based prediction algorithm, which allows formulating the problem using the robust optimization. The robust optimization model takes into account the prediction uncertainty. We compared the results of deterministic and robust of capacity planning, and we found that the robust optimization more realistic to be used in data centers where VMs demand and their utilization are uncertain. Importantly, by using our approach, we could achieve energy saving and provide high availability and reliability for the system. As future work, we will consider heterogeneous servers and VMs size. Furthermore, we intend to extend our approach for a dynamic provision of web applications. Thus, we intend to compare our approach against the other approaches that use stochastic and deterministic optimization for dynamic provision of web applications. Furthermore, we intend to investigate applying our approach to other resources, where we just consider only the CPU demand of VMs.

5. REFERENCES

[1] R. Nathuji, K. Schwan, "Virtualpower: Coordinated power management in virtualized enterprise systems," *ACM SIGOPS Operating Systems Review*, 41(6) pp. 265-278, 2007

[2] A. Verma, P. Ahuja, A. Neogi, "PMAPPER: Power and Migration Cost Aware Application Placement in Virtualized Systems," *proceedings of the 9th ACM/IFIP/USENIX Middleware*,pp. 243-264,Belgium.

[3] A. Gandhi, M. Harchol-balter, R. Das, C. Lefurgy,"Optimal Power Allocation in Server Farms," *proceedings of the 11th International Joint Conference on MMCS*, ACM New York, USA, 2009; 157-168.

[4] Booting-time.[online].available:http://www.toms hardware.co.uk/ubuntu-oneiric-ocelot-benchmark-review,review-32377-15.html.[accessed: 1-jun-2012].

[5] M. Mao and M. Humphrey," A Performance Study on the VM Startup Time in the Cloud," *proc. of fifth IEEE Conference on Cloud Computing*, 2012.

[6] Planet-lab. available :http://www.planet-lab.org/

[7] A. Ben-Tal and A. Nemirovski, "Robust Solutions of Linear Programming Problems Contaminated with Uncertain Data," mathematical programming v. 88 (2000), 411-424.

[8] C. Dance and A. Gaivoronski, "Stochastic Optimization for Real Time Service Capacity Allocation under Random Service Demand," annals of operations research, vol. 193, no. 1, pp. 221-253, feb. 2011.

[9] W. Yanbo, "Empirical Comparison of Robust, Data Driven and Stochastic Optimization," MIT 2008.

[10] Y. Wu, K. Hwang, Y. Yuan , and W. Zheng, "Adaptive Workload Prediction of Grid Performance in Confidence Windows," *IEEE trans. on PDS 2009*

[11] M. J. Clement and M.J. Quinn, "Analytical Performance Prediction on Multicomputers," J. Supercomputing, pp. 886-894, 1993.

[12] K. D. Lange,"Identifying Shades of Green: The SPECpower Benchmarks," IEEE Computer 42(3), pp. 95-97, 2009.

[13] SPEC-results.[online].available. http://www.spec.org/power$_s$sj2008/results/res2011q1 /power$_s$sj2008-20110127-00342.html

[14] J. Goh and M. Sim, "Robust Optimization Made Easy with Rome," Operations Research, 2011, pp.973-985.