

A Class-based Virtual Machine Placement Technique for a Greener Cloud

Claudia Canali, Riccardo Lancellotti
*Department of Engineering "Enzo Ferrari",
University of Modena and Reggio Emilia,
Via Vivarelli, 10, Modena, Italy
{riccardo.lancellotti, claudia.canali}@unimore.it*

Keywords: IaaS Cloud, VM Placement, Server Consolidation Models, Class-Based Placement

Abstract: The management of IaaS cloud systems is a challenging task, where a huge number of Virtual Machines (VMs) must be placed over a physical infrastructure with multiple nodes. Economical reasons and the need to reduce the ever-growing carbon footprint of modern data centers require an efficient VMs placement that minimizes the number of physical required nodes. As each VM is considered as a black box with independent characteristics, the placement process presents scalability issues due to the amount of involved data and to the resulting number of constraints in the underlying optimization problem. For large data centers, this excludes the possibility to reach an optimal allocation. Existing solutions typically exploit heuristics or simplified formulations to solve the allocation problem, at the price of possibly sub-optimal solutions. We introduce a novel placement technique, namely *Class-Based*, that exploits available solutions to automatically group VMs showing similar behavior. The Class-Based technique solves a placement problem that considers only some representatives for each class, and that can be replicated as a *building block* to solve the global VMs placement problem. Our experiments demonstrate that the proposed technique is a viable solution that can significantly improve the scalability of the VMs placement in IaaS Cloud systems with respect to existing alternatives.

1 INTRODUCTION

The success of cloud computing is motivated by its ability to provide computational, storage and networking resources available on-demand. However, as the success of cloud computing grows, larger and ever more powerful data centers are deployed. The management of these infrastructures is challenging, considering their impact in terms of energy consumption and carbon footprint. The EPA and NDRC reports (EPA, 2007; Whitney and Delforge, 2014) place the power consumption of data centers in the last years to 1.5% of the global demands (roughly comparable to the power consumption of countries such as Italy or Spain). Furthermore, power consumption for data centers tends to grow as information and communication technologies become pervasive in our society. To face these challenges, energy consumption has become a key performance indicator for the data centers, and the efficiency in the task of allocating Virtual Machines (VMs) over the physical nodes has become a fundamental goal for the cloud infrastructure management.

As the cloud data centers grow in size, the problem of VMs placement on the physical nodes of the infrastructure becomes challenging. The typical problem formulation is that of a multi-dimensional bin packing. The optimization goal is to minimize the number of physical nodes required to host the VMs, while the capacity requirements of each VM correspond to the expected demand for multiple resources (e.g., CPU, memory, network traffic) at future time intervals (Setzer and Bichler, 2013; Speitkamp and Bichler, 2010). Due to the *NP-hard* nature of the problem, as the number of VM increases, the number of decision variables and constraints grows. Achieving an optimal solution to problem is no longer feasible due to the huge time and memory required taken by the optimization problem solver. The state of the art solutions rely on simplifications to reduce the dimensionality of the problem and/or on simple heuristics to reduce the computational cost of the problem. A first example of dimensionality reduction is to consider only the nominal capacity of each VM (Rochwarger et al., 2011; Mills et al., 2011; Tang et al., 2007) instead of taking into account their actual re-

quirements. This approach simplifies the solution of the bin packing problem because it considers VM characteristics that do not change over time. However, it overestimates the resources that must be provided to the VMs, because the actual utilization of resources for each VM is typically below 100% (Barroso and Hölzle, 2007). A model based only on nominal capacity determines an inefficient use of the cloud data center, resulting in a higher-than-required carbon footprint of the overall infrastructure. Another approach is to reduce the dimensionality of the problem by limiting the number of resources that are considered in the bin packing problem and/or the number of time intervals that are considered for the constraints of the optimization problem (Setzer and Stage, 2010; Speitkamp and Bichler, 2010). It is worth to note that even with these approaches, the computational cost for solving the VM placement problem remains rather high, especially for large data centers. However, the time to obtain a solution for the bin packing problem must remain acceptable even at the expense of the solution quality. For these reason, simple heuristics are usually preferred to more complex and computationally expensive approaches (Wäscher et al., 2007). However, adoption of these heuristics typically reduces the placement solution quality because commonly used techniques, such as First Fit Decreasing (FFD) (Kao, 2008), can only manage few dimensions of the placement problem, thus hindering the use of multiple resources and time intervals.

We claim that the VMs placement problem can benefit from the knowledge of *classes* of VMs with similar behavior in terms of resource usage. In this paper, we sketch a novel technique, namely *Class-Based Placement*, for VMs placement over the physical nodes of the data center that exploits recent methodologies to automatically cluster into classes VMs exhibiting similar behaviors (Canali and Lancellotti, 2014b; Canali and Lancellotti, 2013b). Our proposal shifts the point of view from a single bin-packing problem, that considers the whole data center, to a much smaller problem, limited to a few representatives of each class, that can be replicated as a building block to create the solution for the global VM placement problem. The small size of the building-block problem allows us to solve to optimality problems taking into account an amount of data and constraints that would not be possible to consider in the global bin-packing problem, thus reducing the computational demand and achieving higher quality in the VM placement solution. To the best of our knowledge, no other study follows this approach to achieve in short time a high-quality solution for the VMs placement problem in cloud computing.

We apply a proof-of-concept of our technique to traces coming from a real data center to evaluate the feasibility of the proposed solution. We compare our solution with state of the art models for VMs placement (Setzer and Bichler, 2013). Preliminary results demonstrate that exploiting similarities among VMs provides a viable solution for the VM placement problem in IaaS clouds. The comparison with the alternatives demonstrates that standard techniques based on the solution of optimization problem solvers cannot reach optimal solutions for the bin-packing problem unless the number of VMs is rather small (in the order of 150-200 VMs); even worse, in the case of large problems (in the order of 1000 VMs), the solvers cannot obtain any integer feasible solution within a reasonable time frame. On the other hand, our proposal can reach optimal solutions for problems much larger (up to 700 VMs), and provides viable results even for the largest considered problems (in the order of 1000 VMs).

The remainder of this paper is organized as follows. Section 2 describes the reference scenario for our proposal, while Section 3 describes our model for solving the VM placement problem. Section 4 describes the results of the methodology evaluation. Finally, Section 5 concludes the paper with some final remarks and outlines open research problems.

2 REFERENCE SCENARIO

We now present the scenario that will be used as a reference to illustrate the characteristics of our proposal for the management of cloud data centers, focusing on the operations that decide the placement of the VMs over the physical nodes of the infrastructure. In order to apply our proposal to a data center, we make the following two assumptions.

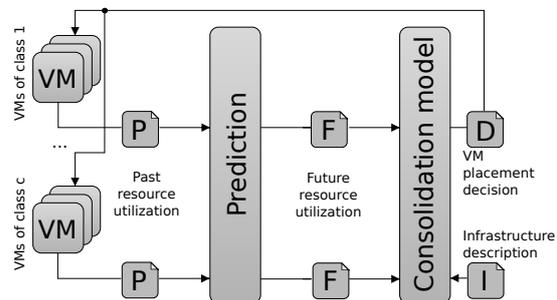


Figure 1: VMs placement in a cloud data center

First, we consider that the VMs placement is a periodic task that aims at mapping VMs over the infrastructure with the goal of minimizing the number of

required physical nodes, while ensuring that the requirements in terms of resource usage of each VM are satisfied. The details of the VM placement process are shown in Figure 1 and described in the following of this section.

Second, we assume to be able to group VMs into classes with similar behavior, where VMs of the same class will exhibit the same resource requirements. Classes containing multiple VMs occur every time an application is deployed over a distributed architecture for scalability and availability reasons: in this case, a dispatcher distributes the client requests over the VMs running an instance of the software component to balance the load, thus ensuring that every VM exhibits the same behavior in terms of resource requirements (Rabinovich and Spatscheck, 2002). Automatic methodologies to cluster VMs with similar behavior have been recently proposed in literature. Some solutions require a long time of observation to define a VM behavior model (Canali and Lancellotti, 2014b; Canali and Lancellotti, 2013a; Canali and Lancellotti, 2013b) and are more suitable for IaaS cloud characterized by long term commitment of the VMs customers (as in the case of the Amazon cloud reserved instances), while other methodologies can provide rapidly a preliminary classification (Canali and Lancellotti, 2014a) and are suitable for a more dynamical scenario. Another case where we have a knowledge of VM classes is when the cloud provider has a complete knowledge of the software running on the VMs, as in the case of private clouds or when the infrastructure supports a SaaS cloud.

Figure 1 depicts the periodic VMs placement in a cloud data center exploiting our proposal. We start with multiple VMs grouped into classes at the left margin of the figure. We recall that VMs belonging to the same class exhibit similarity in terms of resource requirements over time. The VMs are subject to monitoring, that may take advantage from the knowledge of VM classes (Canali and Lancellotti, 2014b; Canali and Lancellotti, 2013a; Canali and Lancellotti, 2013b). We represent the output of this monitoring process as the data marked with the letter “P”. samples on past resource usage are fed into a *Prediction* task. This step can be implemented according to multiple techniques, ranging from the simplest solutions assuming that resource demands follow a periodical cycle with a length of 24 hours (Iyengar et al., 1999), to complex predictive techniques that can cope with trends, periodic behaviors and state changes (Casolari and Colajanni, 2010). The output of the prediction is an estimation of resource utilization in the future for each class of VMs (data with the letter “F”). The future demands and the description of

the infrastructure of the data center (marked with the letter “I”) are the input of the *Consolidation model*, that is the core of our proposal. The consolidation model solves the bin-packing problem and the output (marked with the letter “D” in Figure 1) is the decision on which VMs are to be placed on which physical node. The placement decision is then applied on the VMs powering on and off the servers in the cloud infrastructure.

3 PROBLEM FORMULATION

We now discuss the consolidation model that represents the core of the VMs placement technique. First, we describe the consolidation model that is typically used in literature (Setzer and Bichler, 2013; Speitkamp and Bichler, 2010); then, we discuss the possible simplifications that can be applied to improve the scalability of this task, and we describe the proposed Class-based consolidation model used in our proposal.

3.1 Multi-Dimensional Bin Packing model

The consolidation model used for VMs placement is typically based on a multi-dimensional *bin packing* problem, where one or more VM resources are considered for consolidation and the planning period is divided into a set of time intervals.

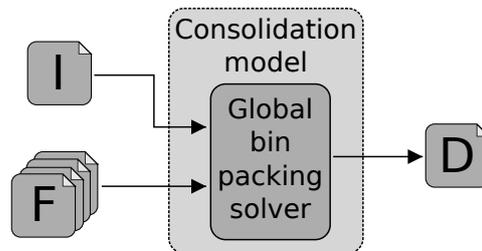


Figure 2: Consolidation model with multi-dimensional bin-packing

Figure 2 shows the simple consolidation model based on a multi-dimensional bin packing problem. The problem input is the prediction of future requirements of resources (such as CPU) for every VM in multiple time intervals (data marked with the letter “F”), and a description of the data center infrastructure, with the available physical nodes and their capacity (the data of the letter “I”). A single problem is solved for the whole data center providing the placement of VMs over the nodes of the data center (the output is represented in Figure 2 as the data with the

letter “D”). The problem formulation aims to minimize the number of used physical nodes under the following constraints:

1. Every VM is allocated exactly on one physical node;
2. On each node, the requirements for the allocated VMs must not exceed the overall capacity of the node in every considered time interval and for every resource.

When solving bin packing problems, the number of dimensions (in this case the number of considered resources and time intervals in the problem formulation) has a major impact on the time to reach a solution. To improve the scalability of VMs placement, a common approach is to reduce the dimensionality of the problem, reducing the number of considered resources and introducing a coarser grained subdivision of time (that is, we consider less time intervals of longer duration). In the extreme case, when the number of resources and time intervals is reduced to one, the multi-dimensional bin packing reverts to a one-dimensional bin packing problem. In this case, we can exploit heuristics such as the *First Fit Decreasing* algorithm to reach an approximate solution of the problem in a very short time (Kao, 2008). However, the reduction of dimensionality typically leads to sub-optimal solutions for the VM placement problem.

3.2 Class-based Placement model

The Class-based Placement exploits the knowledge of classes of VMs with similar behavior in terms of resource usage to improve scalability. This knowledge can be obtained even in IaaS cloud systems, where the cloud providers typically do not have any knowledge of the applications running on the VMs, by exploiting recently proposed techniques (Canali and Lancellotti, 2014b; Canali and Lancellotti, 2013b) that automatically cluster similar VMs. The basic idea is to reduce the global bin packing problem, that operates on the whole data center, to a smaller problem involving only few VMs for each class. The reduced size of the problem allows us to solve to optimality the consolidation model considering a multi-dimensional formulation with a number of resources and time intervals that would not be possible to consider for the global problem; then, the obtained solution can be replicated as a building block to determine the solution for the global VM placement problem.

Figure 3 represents our proposal. Again the input is the future resource requirements for VMs, although in this case we group VMs into classes and we assume that all the VMs of a same class present similar

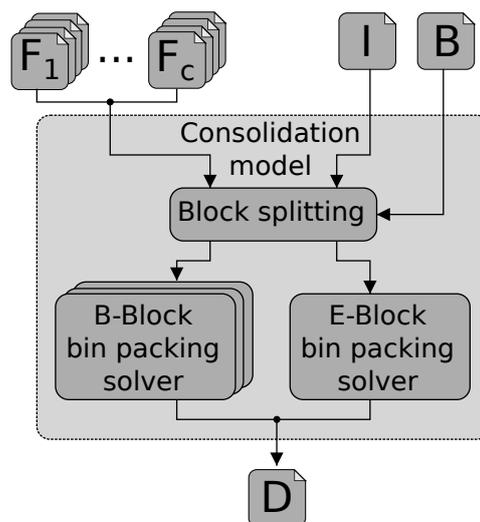


Figure 3: Consolidation model with class-based placement

resource requirements (we show different groups of data with labels from “ F_1 ” to “ F_c ” for the different classes).

The first step of our methodology divides the global set of VMs in a number B of B-blocks, all composed by the same number of VMs for each class, and one E-block containing the remaining VMs. The number of B-Blocks B is considered as an input of our consolidation model (represented with the letter “ B ” in Figure 3). For our experiments, we consider B as the cardinality of the smallest class of VMs. A more accurate analysis on the impact of B on the performance of the consolidation model is considered as an open research direction, to address as a future work.

Next, we solve the bin-packing algorithms for the B-blocks and for the E-block. Since all the VMs of a same class present similar resource requirements, the placement solution computed for a single B-block can be replicated on all the remaining B-blocks.

The reduced size of these blocks allows us to solve the corresponding placement problems considering a multi-dimensional formulation with several resources and time intervals within an amount of time that is acceptable for cloud systems management, thus achieving a scalability much higher compared to the previous approach.

4 EXPERIMENTAL EVALUATION

In this section we present the setup and the results of the experimental evaluation regarding the Class-Based placement technique.

4.1 Experimental setup

We obtain an extensive dataset from a private cloud data center. The set contains up to 1200 VMs traces for the resource usage of Web/application/database servers and ERP applications, where the VMs belongs to 44 different classes, with each class containing from 8 to 50 VMs. We use our traces as the output from the prediction step in the VM placement problem. In our experiments we consider traces with a length of 24 hours, where resource usage is measured in intervals of 5 minutes, that is a setup consistent with other experiments in literature (Addis et al., 2013). For our experiments we limit our model to a single resource, the CPU utilization, that is well-know to be the bottleneck resource for this type of applications (Andreolini et al., 2008). However, it is worth to note that an extension of our model to include multiple resource is straightforward. In the experimental evaluation we simulate data centers of different size, by changing the number of considered VMs. In particular, we consider a VMs set size ranging from 150 to 1200 VMs. For each VM the CPU utilization is in the range [0%-100%], with an average value of 54%. For each physical node, the CPU capacity is 800%, meaning that each node can host 8 VMs with CPU utilization of 100%.

For each scenario, we compare different consolidation models: the proposed Class-Based Placement (CBP) is solved with 288 five-minutes time intervals and the B parameter is set to the size of the smallest class that is 8. For the Multiple Bin Packing (MBP) model we consider different number of time constraints, that are 288 (five-minutes intervals), 24 (1 hour), 2 (12 hours) and a single time interval (24 hours). All the experiments are run on 2.4 GHz Intel Xeon with 16 GB RAM, using IBM ILOG CPLEX 12.6 as the optimizer solver.

It is worth to note that for many problems, starting from a medium size (e.g 400 VMs), the resolution of the MPB consolidation models may take long times, such as hours or days, even for a limited number of time intervals. For that reason, we used a time limit of 30 minutes (1800 seconds) for each problem and considered the best integer solution found as the solution of the placement problem, as commonly done in similar research studies (Setzer and Bichler, 2013; Zhang and Ardagna, 2004).

4.2 Experimental results

In our experiments we compare the different consolidation models to evaluate if they can reach an optimal or a viable solution within the expected time limit.

Table 1 shows for which scenarios it was possible to solve the problem instances to optimality (S), reach an integer solution even if not optimal (L), or not even find any feasible integer solution (N) within the 30 minutes time limit. We evidence the cells related to unsolvable problem instances with a gray background.

Table 1: Resolvable scenarios

VMs Set Size	Consolidation Models				
	CBP 5min	MBP 1d	MBP 12h	MBP 1h	MBP 5min
150	S/S	S	S	S	S
200	S/S	S	S	S	S
250	S/S	S	L	L	L
300	S/S	S	L	L	L
400	S/S	L	L	L	N
500	S/S	L	L	L	N
600	S/S	L	L	N	N
700	S/S	L	L	N	N
800	L/S	L	L	N	N
900	L/S	L	L	N	N
1000	L/S	L	L	N	N
1100	L/S	L	N	N	N
1200	L/S	N	N	N	N

It is worth to note that the MBP consolidation model with five minute time interval (MBP-5min) represents the most complete placement formulation that exploits all the available information to find an optimal solution. However, the number of variables and constraints for this model increases rapidly with the VMs set size, producing an optimization problem instances whose computation may easily take extremely long times or may be not able to produce any feasible solution due to the huge main memory requirements that may cause the solver to abort the optimizer processing.

From the results shown in the table, we observe that only small sized problem instances (up to 200 VMs) can be solved to optimality by every consolidation model. On the other hand, starting from 250 VMs the resolution process lasts longer than the imposed time limit for every MBP model with more than one time interval. For MBP models considering short time intervals of 5 minutes and 1 hour, it is not possible to find a feasible integer solution within the time limit starting from medium sized problems of 400 and 600 VMs, respectively; for larger time of 12 hours and 1 day, the size of resolvable problems grows to 1000 and 1100, respectively. On the other hand, the breakdown in building blocks allows the CBP model to find a feasible integer solution for every VMs set size, with the possibility to solved to optimality even scenarios up to 700 VMs. From these results, it is evident that the CBP technique allows us to solve to

optimality significantly larger problem size with respect to a MBP approach, even when the MBP problem considers only one time interval of 24 hours.

5 CONCLUSIONS

In this paper, we focused on the critical problem of VMs placement in Cloud computing data centers. We pointed out the scalability challenges of this task and the impact of inefficient VM placement on the carbon footprint of Cloud systems. To cope with the scalability issues of current consolidation models, we introduced an alternative approach where VMs are not considered as black boxes each with its own resource requirements. Exploiting recent solutions that can cluster together VMs exhibiting similar behaviors, we sketched a novel VMs placement technique, namely *Class-Based*, that solves a small-size VMs placement problem and replicates it as a building block to obtain the global solution. Preliminary experiments confirmed that our proposal outperforms existing solutions, reaching optimal solutions where other solutions must relax part of the constraints to achieve a sub-optimal but viable solution. On the other hand, our proposal can easily scale to more than 1000 VMs without the need to relax any constraint.

REFERENCES

- Addis, B., Ardagna, D., Panicucci, B., Squillante, M. S., and Zhang, L. (2013). A hierarchical approach for the resource management of very large cloud platforms. *IEEE Transactions on Dependable and Secure Computing*, 10(5):253–272.
- Andreolini, M., Casolari, S., and Colajanni, M. (2008). Models and framework for supporting runtime decisions in Web-based systems. *ACM Transactions on the Web*, 2(3):1–43.
- Barroso, L. A. and Hölzle, U. (2007). The case for energy-proportional computing. *IEEE computer*, 40(12):33–37.
- Canali, C. and Lancellotti, R. (2013a). Automatic virtual machine clustering based on Bhattacharyya distance for multi-cloud systems. In *Proc. of International Workshop on Multi-cloud Applications and Federated Clouds*, pages 45–52, Prague, Czech Republic.
- Canali, C. and Lancellotti, R. (2013b). Exploiting ensemble techniques for automatic virtual machine clustering in cloud systems. *Automated Software Engineering*, pages 1–26. Available online.
- Canali, C. and Lancellotti, R. (2014a). An Adaptive Technique to Model Virtual Machine Behavior for Scalable Cloud Monitoring. In *Proc. of IEEE Symposium on Computers and Communications (ISCC)*, Madeira, Portugal.
- Canali, C. and Lancellotti, R. (2014b). Improving Scalability of Cloud Monitoring Through PCA-Based Clustering of Virtual Machines. *Journal of Computer Science and Technology*, 29(1):38–52.
- Casolari, S. and Colajanni, M. (2010). On the selection of models for runtime prediction of system resources. In *Run-time Models for Self-managing Systems and Applications*, pages 25–44. Springer.
- EPA (2007). Report to congress on server and data center energy efficiency. Technical report, US Environmental Protection Agency.
- Iyengar, A. K., Squillante, M. S., and Zhang, L. (1999). Analysis and characterization of large-scale web server access patterns and performance. *World Wide Web*, 2(1-2):85–100.
- Kao, M. (2008). *Encyclopedia of Algorithms*. Springer.
- Mills, K., Filliben, J., and Dabrowski, C. (2011). Comparing vm-placement algorithms for on-demand clouds. In *Cloud Computing Technology and Science (Cloud-Com), 2011 IEEE Third International Conference on*, pages 91–98. IEEE.
- Rabinovich, M. and Spatscheck, O. (2002). *Web caching and replication*. Addison-Wesley Boston, USA.
- Rochwerger, B., Breitgand, D., Epstein, A., Hadas, D., Loy, I., Nagin, K., Tordsson, J., Ragusa, C., Villari, M., Clayman, S., et al. (2011). Reservoir when one cloud is not enough. *IEEE computer*, 44(3):44–51.
- Setzer, T. and Bichler, M. (2013). Using matrix approximation for high-dimensional discrete optimization problems: Server consolidation based on cyclic time-series data. *European Journal of Operational Research*, 227(1):62–75.
- Setzer, T. and Stage, A. (2010). Decision support for virtual machine reassignments in enterprise data centers. In *Proc. of Network Operations and Management Symposium*, Osaka, Japan.
- Speitkamp, B. and Bichler, M. (2010). A Mathematical Programming Approach for Server Consolidation Problems in Virtualized Data Centers. *Services Computing, IEEE Transactions on*, 3(4):266–278.
- Tang, C., Steinder, M., Spreitzer, M., and Pacifici, G. (2007). A scalable application placement controller for enterprise data centers. In *Proceedings of the 16th International Conference on World Wide Web, WWW '07*, pages 331–340, New York, NY, USA. ACM.
- Wäscher, G., Haußner, H., and Schumann, H. (2007). An improved typology of cutting and packing problems. *European Journal of Operational Research*, 183(3):1109–1130.
- Whitney, J. and Delforge, P. (2014). Data center efficiency assessment – scaling up energy efficiency across the data center industry: Evaluating key drivers and barriers. Technical report, NRDC, Anthesis. – <http://www.nrdc.org/energy/files/data-center-efficiency-assessment-IP.pdf>.
- Zhang, L. and Ardagna, D. (2004). Sla based profit optimization in autonomic computing systems. In *Proceedings of the 2Nd International Conference on Service Oriented Computing, ICSOC '04*, pages 173–182, New York, NY, USA. ACM.