

A Multi-start Local Search Scheduler for an Energy-aware Cloud Manager

Y. Kessaci¹, N. Melab¹ and E-G. Talbi¹

Laboratoire d'Informatique Fondamentale de Lille,
INRIA Lille Nord Europe ,
Université de Lille1.

{Yacine.Kessaci,Nouredine.Melab,EI-Ghazali.Talbi}@lifl.fr

1 Introduction

The field of cloud computing uses different management techniques for data center virtualization such as OpenNebula [1]. However, computers composing the cloud infrastructure use a significant and growing portion of energy in the world specifically when dealing with virtualization for high performance computing (HPC). Therefore, energy-aware computing is crucial for large-scale systems that consume considerable amount of energy.

In this paper, we present a new work that aims to deal with the energy consumption within a realistic cloud infrastructure using OpenNebula as a software management solution. Our scheduler is based on a multi-start local search heuristic that helps to find the best scheduling by dispatching the arriving of virtual machines (VM) according to the minimum energy consumption.

The remainder of the paper is organized as follows. In Section 2 we present the related work to our approach. Our approach is presented in Section 3. The results of our experimental study are discussed in Section 4. The conclusion is drawn in Section 5.

2 Related work

The major part of the work on the literature tackle the energy topic by referring and focusing on scheduling dedicated applications. In [2, 3] for example a hardware technique (DVFS) is proposed. The drawback of this type of methods is the assumption that they make about a tight coupling between the tasks and the resources. Another way of reducing cloud computing energy footprints is proposed in [4]. This work uses the potential offered by the virtualization in order to apply a task consolidation through two heuristics in order to maximize the resource utilization.

Moreover, only few works deal with the energy aware on cloud distribution. The work in [5] is a part of them . It proposes a VM scheduling algorithm based on DVFS technique to reduce the energy consumption of a single OpenNebula virtualized cluster. The idea behind this work is to reduce the clock frequency of the cluster as low as possible to fit exactly the VMs requests.

The major difference of our work with the approach cited before is that our work deals with not only one cluster but a big number of geographically distributed machines. Indeed, our approach tends to provide a scheduler on an HPC virtualized environment. In addition, our approach is the only one which is embedded on real cloud distribution (OpenNebula) and not only using it for simulation.

3 EMLS-ONC algorithm

Before each scheduling, the scheduler waits for a fixed period of time called *scheduling cycle*. This period helps to gather a pool of VMs in order to have a larger choice and thus to optimize the future scheduling. Once this phase done the host pool is filtered out to keep only the hosts with the correct requirements. The multi-start phase launches each local search algorithm separately. The number of launched LS is equal to the minimum value between the number of hosts composing the distributed cloud and 20.

The role of the local search algorithm is to make a number of combinations from the initial solution using neighborhood relation in order to find the best scheduling according to the specified

objective. Using the multi-start adds diversity, while each local search plays the role of intensification. The local search algorithm starts by generating the initial solution. This initial solution is used to generate a neighborhood based on two neighborhood operators. Both operators are based on an exchange process. The first operator is dedicated to generate neighborhoods for small cloud configuration while the second is dedicated to big neighborhoods with huge distributed infrastructure. Each time a solution is generated, the algorithm checks the feasibility of this latter. A fitness value is also assigned to this solution. The best solution is kept to build another neighborhood with the previous operators. The algorithm stops when the number of iterations reaches the number of VMs.

In the last step, OpenNebula dispatches the VMs according to the best solution found by EMLS-ONC, updates the hosts states and a new scheduling cycle is started.

4 Experimental results

Our approach deals with VMs for HPC usage. Therefore, we conducted our experiments on only one scheduling cycle. Indeed, it gives the ability to evaluate both our approach and the default OpenNebula scheduler VMs workload capacity.

To the best of our knowledge no previous approach deals with a multi-start local search scheduler for a realistic OpenNebula managed distributed cloud. In addition, EMLS-ONC optimizes both the energy consumption and the quality of service (QoS), by satisfying the maximum number of clients (VMs requests). Therefore, the energy consumption is reduced averagely by **15,6%** compared to the OpenNebula default scheduler, while the VMs assignment is averagely **1,3%** better than this latter.

Furthermore, EMLS-ONC returns results in a small time duration, largely lower than the threshold represented by the scheduling cycle. It respects therefore, the rapidity required for schedulers.

5 Conclusion

In this paper, we investigated the energy reduction for VMs scheduling on a distributed cloud, particularly for high performance computing usage. The current cloud manager does not pay much attention to energy consumption.

We presented a new scheduler for the cloud manager OpenNebula to solve this problem. It is based on a multi-start local search algorithm to minimize the energy consumption, while it satisfies the clients' QoS by assigning the maximum VMs. The obtained results of our approach clearly improves the results obtained by the previous OpenNebula scheduler.

The energy saving of our approach exploits the disparity and the difference in the features of the hosts that compose the distributed cloud.

References

1. Fontan, J., Vazquez, T., Gonzalez, L., Montero, R.S., Llorente, I.M.: Opennebula: The open source virtual machine manager for cluster computing. In: San Francisco, CA, USA. (2008)
2. Lee, Y.C., Zomaya, A.Y.: Minimizing energy consumption for precedence-constrained applications using dynamic voltage scaling. In: CCGRID'09: Proceedings of the 2009 9th IEEE/ACM International Symposium on Cluster Computing and the Grid. (2009) 92–99
3. Rizvandi, N.B., Taheri, J., Zomaya, A.Y., Lee, Y.C.: Linear combinations of dvfs-enabled processor frequencies to modify the energy-aware scheduling algorithms. Cluster Computing and the Grid, IEEE International Symposium on **0** (2010) 388–397
4. Lee, Y., Zomaya, A.: Energy efficient utilization of resources in cloud computing systems. The Journal of Supercomputing (2010) 1–13 10.1007/s11227-010-0421-3.
5. Power-Aware Scheduling of Virtual Machines in DVFS-enabled Clusters, New Orleans, LA, IEEE Computer Society (2009)