Probabilistic Consolidation of Virtual Machines in Self-Organizing Cloud Data Centers

Carlo Mastroianni[1,3], Michela Meo[2], Giuseppe Papuzzo[1,3]

[1] Institute of High Performance Computing and Networking (ICAR-CNR), Italy
[2] Department of Electronics and Communications, Politecnico di Torino
[3] Eco4Cloud srl, Italy, www.eco4cloud.com
Contacts: mastroianni@icar.cnr.it, michela.meo@polito.it, papuzzo@eco4cloud.com

Data center computational efficiency

Some of the main issues related to the management of data centers are: reduction of **energy consumption** and associated costs, minimization of **carbon emissions**, and **capacity planning**. All these issues are related to the general problem of increasing the efficiency of resources.

In the last few years important results have been achieved by improving the **efficiency of the physical infrastructure** (e.g., of supplying and cooling components), as testified by the very good values of PUE (Physical Usage Effectiveness) obtained in modern data centers.

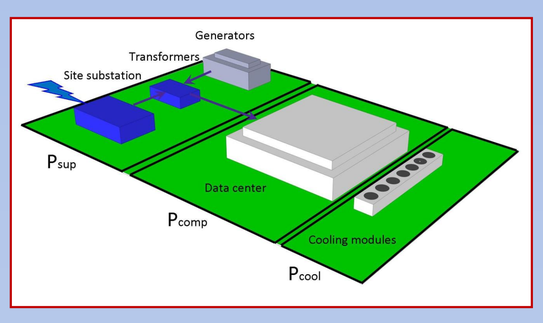

However there is still much room for improving the **computational efficiency**: very often computing resources (CPU, RAM, bandwidth) are under-utilized. They are sized to sustain load picks, but in periods of lower load (during nights or weekends) the data center should **consolidate** the VMs in order to use fewer servers and put the unused ones in a low power state, in order to save energy and, if needed, to devote the spare resources to accommodate more load.

Consolidation of VMs

The objective of consolidation is to use the minimum number of servers to sustain the VMs workload. Unfortunately, this is a very difficult problem, as it can be formulated as a **Bin Packing Problem**, known to be NP-hard.

Two more circumstances make the problem even harder:

(i) the assignment of VMs should take into account **multiple server resources** at the same time, e.g., CPU and RAM, therefore it becomes a "*multi-dimensional bin packing problem*", much more difficult than the single dimension problem;

(ii) even when a good assignment has been achieved, the VMs continuously modify their hardware requirements, potentially baffling the previous assignment decisions in a few hours. **VM migrations** are needed in this case.

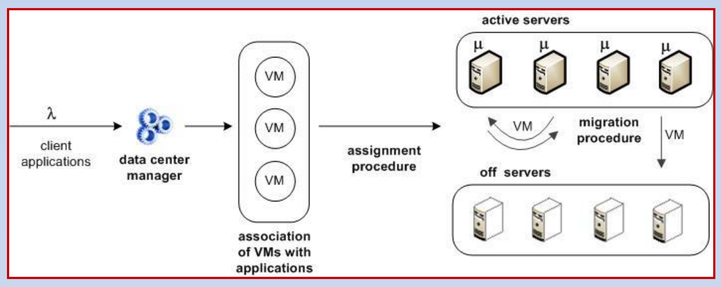

Many solutions are available for the consolidation problem, but most of them adopt centralized algorithms that may not scale well. Moreover, concurrent migrations of many Virtual Machines are generally needed, possibly deteriorating the quality of service perceived by users.

Scalable Multi-Resource Consolidation

EcoCloud is a self-organizing algorithm with the following characteristics:

- The decision whether or not to accept a VM is taken **locally** on each server, based on the utilization of single resource (CPU, RAM, etc.);
- Local decisions are **probabilistic**: e.g., a low-utilized server tends to reject a new VM, as it will try to be unloaded and switched off;

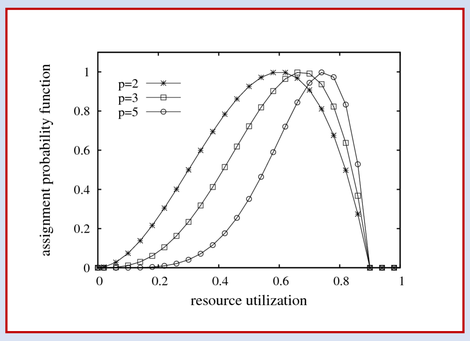

Assignment probability function

The probability of accepting a VM is high in a server with intermediate load (in order to foster consolidation), and low for high- and low-utilized servers

- The probabilistic approach efficiently solves the **multi-resource** problem, and VMs with different requirements (e.g., CPU- bound and RAM-bound) are naturally **balanced** on each server;
- Migrations are **gradual** and asynchronous;
- Despite most operations are decentralized, final decisions are taken centrally, in order to keep the complete **control** at any time;
- **Scalability** is ensured by the use of probabilities, and in larger data centers the consolidation degree is even higher.

Experiments on a real data center

The effectiveness of **EcoCloud** was confirmed in many ways:

- through an **analytical study** (see reference 2 below)
- via **simulation** (to analyze very large data centers)
- through **experiments** on real data centers

One of the live experiments was performed on the data center of a major telco company, in the following scenario:

- **28 physical servers** virtualized with VMWare
- **447 VMs**, of which 335 are **RAM-bound** (M-type) and 112 are **CPU-bound** (C-type)

The experiments starts form a non-consolidated scenario.

Consolidation on fewer servers

11 servers, out of 28, are hibernated thanks to workload consolidation

Energy saving

EcoCloud leads to remarkable energy savings, higher than 40%!

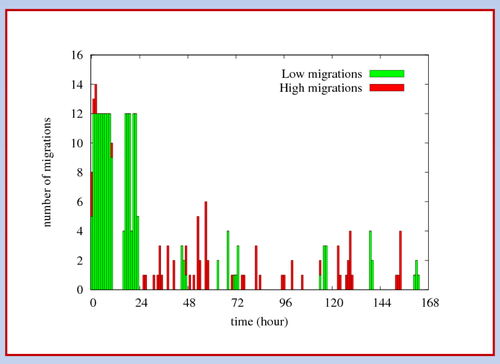

Migrations of VMs

High migrations prevent resources overload. Low migrations are made to unload under-utilized servers and switch them off.

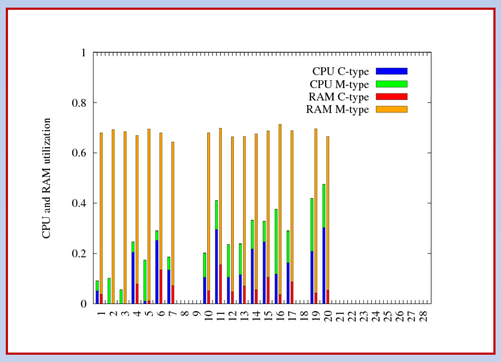

Multi-resource consolidation

EcoCloud consolidates workload w.r.t. to the more utilized resource (RAM), and balances CPU- and RAM-bound VMs on each server

Eco4Cloud start up company

This and other experiments were performed as Proofs of Concept by Eco4Cloud srl (www.eco4cloud.com), a company founded by the authors of this work.

Eco4Cloud, spinoff from Italian CNR and Univ. of Calabria, commercializes a software product that derives from this research.

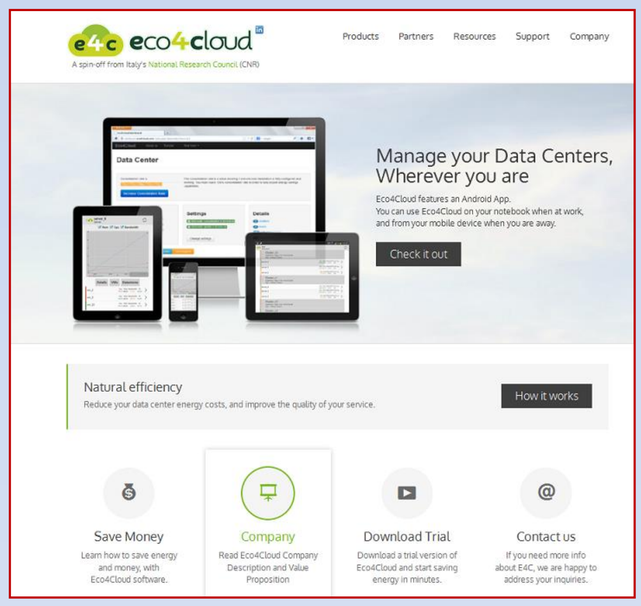

References

1) Eco4Cloud white paper:
R. Giordanelli, C. Mastroianni, M. Meo, G. Papuzzo, A. Roscetti, **Saving energy in data centers through workload consolidation**, downloadable from http://www.eco4cloud.com/download/e4c-white-paper.pdf

2) Paper on probabilistic consolidation on a single resource (CPU):
C. Mastroianni, M. Meo, G. Papuzzo, **Self-Economy in Cloud Data Centers: Statistical Assignment and Migration of Virtual Machines**, 17th International European Conference on Parallel and Distributed Computing, **Euro-Par 2011**, Bordeaux, France, vol. 6852, pp. 407-418, Springer LNCS, September 2011

3) Paper with analytical foundation for the single resource problem:
C. Mastroianni, M. Meo, G. Papuzzo, **Analysis of a Self-Organizing Algorithm for Energy Saving in Data Centers**, 9th Workshop on High-Performance, Power-Aware Computing at **IPDPS 2013**, Boston (MA), USA , May 2013

4) Paper on multiple resource consolidation, to appear:
C. Mastroianni, M. Meo, G. Papuzzo, "**Probabilistic Consolidation of Virtual Machines in Self-Organizing Cloud Data Centers**", accepted by **IEEE Transactions on Cloud Computing**