

Title: Scheduling in the clouds

Advisors: Luciana Arantes (LIP6), Evripidis Bampis (LIP6) and Giorgio Lucarelli (LIG/INRIA)

Hybrid cloud platforms are composed by a set of computational resources of individual organizations (private clouds), of collaborating organizations (community clouds), and of public services from commercial providers (public clouds). Such platforms provide an easy access to an extraordinary amount of computational power at low prices. Cloud Computing is an important factor for environmental sustainable development. If, on the one hand, the increasing demand of clients leads to the creation of large data-centers, on the other hand Cloud Computing’s multitenancy” trait (through the use of virtualization, several users may actually share a single physical resource) leads to a reduced use of physical hardware, saving in that way energy. Therefore, it is important to optimize the energy consumption in such computing environments by taking into account the trade-off between energy and performance by using virtualization and migration. Another important aspect in cloud environments is the development of policies that are able to tolerate failures. We propose to study the impact of such policies, including virtualization, migration and replication, on the energy consumption and the performance of such platforms.

The objective of this “stage” will be the development and the evaluation of scheduling policies taking into account the characteristics of such environments.

Note: *The “stage” will take place at LIP6, UPMC, 4, place Jussieu 75005.*

References

1. Zoltan Adam Mann, Allocation of Virtual Machines in Cloud Data Centers—A Survey of Problem Models and Optimization Algorithms, ACM Comput. Surv. 48, 1, Article 11 (August 2015).
2. Marco E. T. Gerards, Johann L. Hurink, Philip K. F. Hölzenspies, A survey of offline algorithms for energy minimization under deadline constraints, Journal of scheduling, 2016.