How an experimental Grid is used : The Grid5000 case and its impact on energy usage

Anne-Cécile Orgerie, Laurent Lefèvre, Jean-Patrick Gelas
INRIA RESO - Université de Lyon - LIP (UMR CNRS, INRIA, ENS, UCB)
École Normale Supérieure - 46, allée d'Italie - 69364 LYON Cedex 07 - FRANCE ,
laurent.lefevre@inria.fr, {annececile.orgerie|jean-patrick.gelas}@ens-lyon.fr

1 Introduction : A day in the life of an experimental Grid

The question of energy savings is a matter of concern since a long time in the mobile distributed systems. However, for the large-scale non-mobile distributed systems, which nowadays reach impressive sizes, the energy dimension just starts to be taken into account.

Some previous work on operational Grids [IDE+06] show that grids are not utilized at their full capacity. We focus on the utilization and the energy analysis of experimental Grids by relying on the case study of Grid5000[Cea05][1], a french experimental Grid.

The Grid5000 platform is an experimental testbed for research in grid computing which owns more than 3400 processors geographically distributed on 9 sites in France. This platform can be defined as a highly reconfigurable, controllable and monitorable experimental Grid equipment.

Its utilization is very specific. Each user can indeed reserve in advance some nodes and then during its reservation time he is root on these nodes and he can deploy his own system images, collect data, reboot and so on. The node is entirely dedicated to the user during his reservation.

So Grid5000 is, by some important aspects, really different from an operational Grid. We don't need a really strong reliability and the usage is exclusive.

2 Grid usage : In search of the bursts

This poster will present our last results and explanations on experimental Grid usage by focusing on some specific sites and scenario usage. We observe that the usage greatly varies from one site to another (Figure 1, res. stands for reservation) in terms of number of reservations, number of cores, average number of cores per reservation, length of a reservation and the percentage of real work (without including the dead and absent time periods). This heterogeneous usage can be due to geographical purposes (the most distant are interesting to conduct communication experiments) and to hardware purposes:

Site	# of res.	# of core	# of core / res.	length of res.	real work
Bordeaux	45775	650	55.50	5224.59 s.	47.80%
Lille	330694	250	4.81	1446.13 s.	36.44%
Lyon	33315	322	41.64	3246.15 s.	46.38%
Nancy	63435	574	22.46	19480.49 s.	56.41%
Orsay	26448	684	47.45	4322.54 s.	18.88%
Rennes	36433	714	54.85	7973.39 s.	49.87%
Sophia	35179	568	57.93	4890.28 s.	51.43%
Toulouse	20832	434	12.89	7420.07 s.	50.57%

Figure 1: Grid5000 usage over one-year period (2007)

each site has different nodes with different architectures (storage, network capabilities...).

Figure 2 shows the example of the site of Sophia with 568 cores. The black line indicates the number of reservations per week. For each week, we have represented in red the time during which some cores are dead, that is to say they are down; in orange when they are suspected, they do not work properly; in yellow when they are absent, they do not answer and in green when they are working (a job is running).

We can see on figure 2 that during some weeks, the usage of the site is low, but the real matter of concern of such a Grid is to be able to support burst periods of work and communication specially before well-known deadlines.

3 Towards energy aware Grid

Our objective is the measurement of the power consumption of the Grid nodes in Watt in order to modelize the link between electrical cost and applications or processes.

In order to measure the real consumption of some nodes, we have used a watt-meter furnished by the SME Omegawatt[2]. Figure 3 shows our experimental architecture. It is composed of a set of power sensors (Watt meters), each connected to a node's power supply. Then, Watt meters are able to send measures to a database server through a serial link (RS232) converted in Ethernet.

Figures 4 and 5 show our first results concerning the energy consumption on the site of Lyon. We have dynamically collected the consumption in Watts of six different nodes representing the 3 hardware architectures available on this site : IBM eServer 325 (2.0GHz, 2 CPUs per node), Sun Fire v20z

[1]Some experiments of this article were performed on the Grid5000 platform, an initiative from the French Ministry of Research through the ACI GRID incentive action, INRIA, CNRS and RENATER and other contributing partners (http://www.grid5000.fr)

[2]http://www.omegawatt.fr/gb/index.php

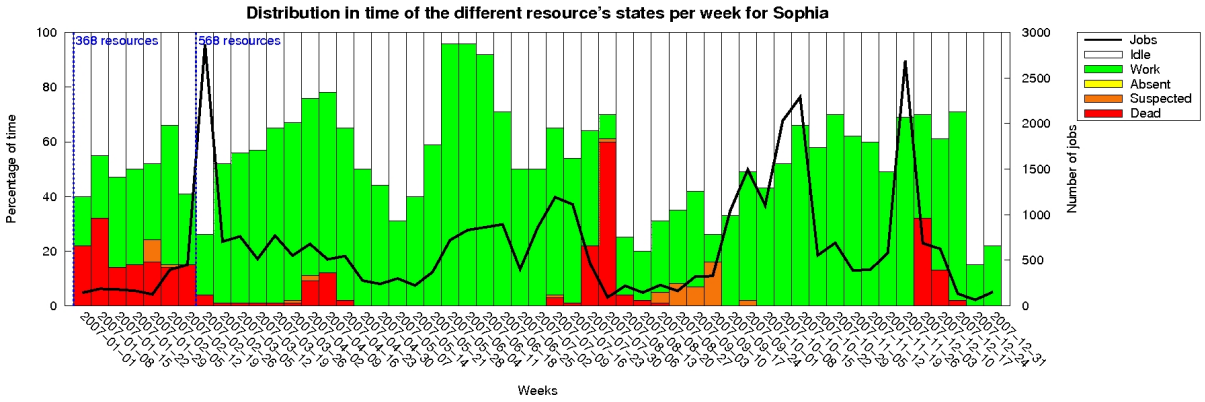

Figure 2: Weekly Sophia site usage on a one-year period

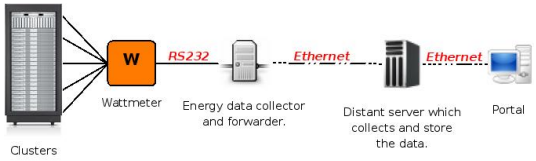

Figure 3: Experimental architecture

(2.4GHz, 2 CPUs per node) and HP Proliant 385 G2 (2.2GHz, 2 dual core CPUs per node).

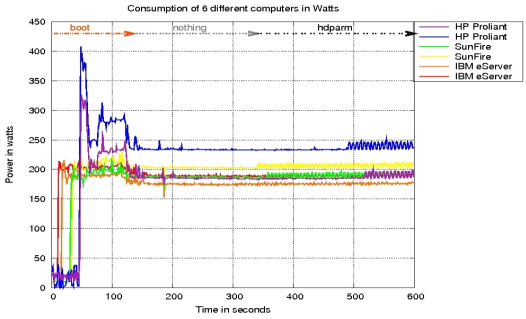

Figure 4: Node booting and disks access on six nodes

This poster will describe and analyse the energy usage of Grid nodes life during some specific scenario. These experiments represent a typical life of an experimental Grid node : nodes down but plugged in the wall socket, booting, having intensive disks access (`hdparm`), experimenting intensive high performance network communications (`Iperf` between nodes with the same architecture), or having intensive CPU usage (`cpuburn`). These first results show the impact on energy usage resulting from nodes utilization.

4 Setting objectives

This poster presents a work in progress whose goal is to better understand the usage of experimental Grids and then to propose methods and energy-aware tools to reduce the energy consumption in such Grids.

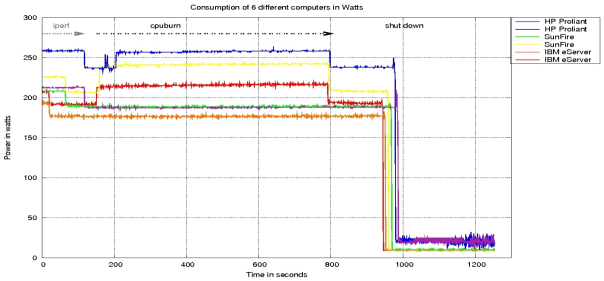

Figure 5: Intensive communication and computing on six nodes

We have seen instructive results about the utilization and the energy analysis of an experimental Grid over the example of Grid5000. We are currently working on tools proposing these results in a real-time manner to the users and grid middleware. Therefore we want to make further consumption experiments on real nodes and then we should propose a model to extend these results to all the nodes by similarity.

References

[Cea05] F. Cappello et al. Grid'5000: A large scale, reconfigurable, controlable and monitorable grid platform. In *6th IEEE/ACM International Workshop on Grid Computing, Grid'2005*, Seattle, Washington, USA, Nov. 2005.

[IDE+06] A. Iosup, C. Dumitrescu, D. Epema, Hui Li, and L. Wolters. How are real grids used? the analysis of four grid traces and its implications. In *7th IEEE/ACM International Conference on Grid Computing*, September 2006.