An Analysis of Power Consumption Logs from a Monitored Grid Site

Marcos Dias de Assunção, Anne-Cécile Orgerie and Laurent Lefèvre
INRIA RESO - Université de Lyon - LIP
École Normale Supérieure
46, allée d'Italie - 69364 Lyon Cedex 07 - France
Emails: {marcos.dias.de.assuncao, annececile.orgerie, laurent.lefevre}@ens-lyon.fr

***Abstract*—This paper analyses information on the electricity consumed by a site of an experimental Grid infrastructure and its correlations with users' resource reservation requests. The power consumption of all servers of this site was monitored for a period of six months; the data was gathered to enable studies on approaches for reducing the energy footprint of Grids and Clouds.**

I. INTRODUCTION

Managing and supplying computational, storage, and network resources to user applications is one of the main challenges for the distributed computing community. Over the years, large-scale infrastructures have been built to provide the resources required by distributed applications [1], [2], [3], [4], [5]. In addition, the strict availability and computational requirements of current business and scientific services have driven the creation of data centres, and more sophisticated business models such as that recently known as "Cloud Computing" [6], [7].

Based on the economies of scale and recent developments in virtualisation and network technologies, commercial Cloud resource providers aim to offer resources to users in a pay-as-you-go manner. These providers allow users to set up and customise execution environments according to their application needs. In a similar fashion, large corporations and research organisations have adopted a Cloud-like approach to manage their infrastructure [1]. Although these centres can allow for more flexible provisioning strategies, there are increasing concerns surrounding the electrical power consumed by IT infrastructure [8].

A range of techniques can be utilised to make computing infrastructures more energy efficient, including better cooling technologies, Dynamic Voltage and Frequency Scaling (DVFS) [9], and resource virtualisation. However, to develop provisioning techniques that use these technologies efficiently, it is important to understand the resource requirements and workload of user applications, and the behaviour of the physical infrastructure in terms of power consumption. Monitoring existing infrastructure and making data available to researchers and practitioners can aid the design and evaluation of power efficient scheduling and resource provisioning mechanisms. Although monitoring data on the power consumed by computing infrastructures can help improve the design of middleware and resource allocation schemes, this type of information is not always readily available.

This paper presents an analysis of data on power consumption and reservation requests from an experimental Grid platform; the Grid'5000 [1]. The resource reservation requests resemble those of a Cloud infrastructure, since users reserve nodes on which they deploy disk images with the complete software stack required by their applications. Although the infrastructure has been created for investigating and experimenting new services and protocols, the concerns on minimising the power consumption are similar to those of a production environment. Previous work has shown that it is possible to make substantial energy savings by giving users options to change the start and finish times of their reservations [10]. We believe that greater savings can be achieved under a better understanding about how users utilise reserved resources and how much energy they consume during their reservations.

II. EXPERIMENTAL SCENARIO

This section describes the Grid infrastructure, the equipments utilised for measuring the power consumption, and the information obtained and analysed in this paper.

A. The Grid'5000 Infrastructure

The Grid'5000 platform [1] depicted in Figure 1 is an experimental testbed for research in distributed computing which offers around 5000 CPU cores geographically distributed across 9 sites in France; all sites are linked by a dedicated high-speed network. The utilisation of Grid'5000 is specific: using OAR, a user can reserve a number of nodes in advance to deploy a system image customised to her application, so that a node is entirely dedicated to the user during her reservation. Kadeploy is the tool utilised by Grid'5000 users for deploying an environment on the nodes assigned by OAR. As mentioned beforehand, although Grid'5000 is different from a production environment in terms of usage and application deployment, the concerns about power consumption are similar to those of production infrastructures.

The results presented in this paper rely on data collected from sensors (*i.e.* wattmeters) that measure the power consumption of equipments located at one of Grid'5000 sites, in Lyon. The energy sensing infrastructure deployed on this site monitors 135 servers and the network equipments. Detailed information about the energy-sensing infrastructure is

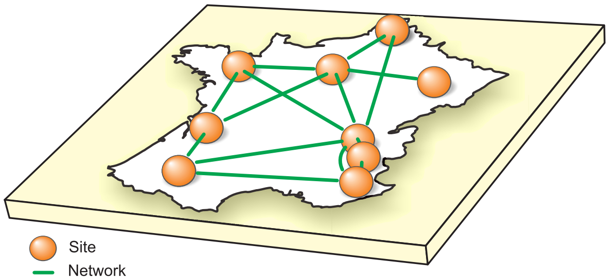

Fig. 1. The Grid'5000 infrastructure in France.

provided in previous work [11]. This paper focuses on the energy consumed by the servers. The wattmeters have been manufactured by the OMEGAWATT company. Each wattmeter is connected to a data collector via a serial port and has a precision of 0.125W. A server collects and logs the data from the wattmeters and generates graphs that help understand the behaviour of equipments and user reservations in terms of energy consumption. One measurement is carried out every second.

This site comprises two clusters namely *capricorne* and *sagittaire*, which are composed of 56 and 79 nodes respectively. Each of the capricorne machines is a IBM eServer 326 with 2 AMD Opteron CPUs, and 2GB of RAM. The sagittaire machines are Sun Fire V20z servers equipped with 2 AMD Opteron CPUs, and 16GB of RAM. All nodes have disks with SCSI interfaces.

B. Resource Reservation and Deployment Logs

OAR [12] is the Resource Management System (RMS) responsible for scheduling user reservation requests on sites of Grid'5000. We have obtained the log with information about the requests scheduled by OAR during the same period when the power consumption data was available. This information is used along with the data on power consumption described later to examine the correlations between power consumption and resource utilisation, since this work also intends to identify and study the profile of users and reservations in terms of resource requirements and the energy they consume. In addition, information from kadeploy[1] was used to identify when users deploy their computing environments once servers are reserved.

C. The Power Consumption Logs

The power consumption data spans six months, more precisely from 1 September 2009 to 27 February 2010. There are intervals over which the power consumption of the platform could not be determined. These intervals represent periods when the electricity supply was cut off, or the software and hardware responsible for collecting the power consumption data failed. Although most of these issues have been identified and addressed, they have led to small gaps in the power consumption logs during the months considered in this report. Table I summarises the periods during which the power consumption data is not available. We detail how we address the lack of measurements during these intervals when describing the results and analysing the behaviour of the platform.

[1]https://gforge.inria.fr/projects/kadeploy3

TABLE I
FAILURES WHILE COLLECTING THE POWER INFORMATION.

Start Time	Duration	Description
06-09-2009 04:25:41	01:11:34	Failure in the wattmeters.
13-09-2009 19:13:10 17-09-2009 13:30:56	90:13:40 00:02:40	Platform downtime.
19-09-2009 08:31:26 20-09-2009 04:30:25 27-09-2009 04:25:02 24-10-2009 13:33:52 25-10-2009 05:34:31 01-11-2009 05:25:16	01:11:34 01:11:34 01:12:33 00:17:51 01:11:34 01:12:05	Failure in the equipments responsible for measuring the power consumption.
06-11-2009 16:11:19	67:57:15	Platform downtime.
10-11-2009 11:47:32 15-11-2009 05:25:16 22-11-2009 05:25:05 13-12-2009 05:25:08 20-12-2009 05:25:46 10-01-2010 05:25:13 17-01-2010 05:25:13 31-01-2010 05:25:33 07-02-2010 05:25:37	01:28:38 01:11:34 01:12:56 01:13:40 01:11:34 01:13:11 01:11:34 01:13:19 01:11:34	Failure in the equipments responsible for measuring the power consumption.

III. Global Energy Consumption

The graph in Figure 2 shows the energy consumed by the servers of the Lyon site during six months; the overall consumption during this period was approximately 103.047 MWh. The intervals in green correspond to periods during which the power consumption information is not available, either due to failures in the equipments responsible for measuring the power consumption or downtimes of the Grid infrastructure. Although the energy data of some of the equipments was available during these periods, we maintain only the periods where data from all servers is on hand. The energy consumption of servers during days with failures that are shorter than two hours was obtained in the following manner. We calculate the average in KW per second of the part of the day when the data is available and use this value as the consumption during the missing seconds.

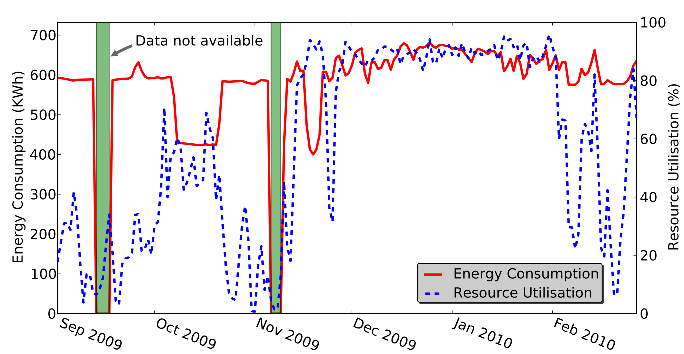

Fig. 2. Energy consumption and utilisation of nodes over six months.

Figure 2 also presents the resource utilisation according to the reservation log obtained from OAR; the utilisation indicates the percentage of reserved nodes, and does not imply that

CPUs, storage or network resources were used by reservations at the same rate, as demonstrated in later sections. Considering the intervals under low resource utilisation, one can observe that the static consumption is nearly 600 KWh; we also term the static consumption here as *idle consumption* since it corresponds to the electrical power drawn by servers when they are not actively executing user applications. Although small when compared to the idle consumption, the graph indicates that during the months of January and February the dynamic energy consumption is proportional to resource utilisation. However, there are troughs in the energy usage line when the consumption is much lower than the average. These intervals are analysed in detail in later sections.

The RMS also maintains information about which nodes are not available for serving user requests. To understand the troughs in the energy consumption graph, Figure 3 illustrates the downtime of the Grid platform according to the RMS' log. This information, however, cannot explain all the low points because a failure for the RMS may not mean that the node is switched off. For example, the low point in October does not match a failure according to the RMS. Hence, we also use the energy consumption information to identify the percentage of nodes that were switched off during the interval of six months (please see Figure 4). We consider that a node is switched off when its power consumption is below 30 Watts.

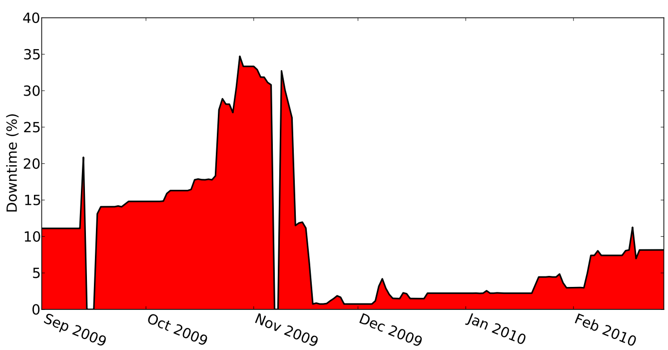

Fig. 3. Platform downtime over six months (based on the scheduler's log).

It can be observed that a substantial part of the platform was switched off during the months of October and November. These periods correspond respectively to a failure in the air-conditioning system and a demo organised during an important conference. During the air-conditioning failure, the system administrator reserved part of the platform and switched the nodes off, whereas throughout the conference, nodes were switched off to demonstrate the impact of power-aware scheduling mechanisms.

To check whether the platform consumes different amounts of electrical power depending on the time of the day, we divide the day into three periods of eight hours each and plot the consumption of all servers during these different intervals. Figure 5 summarises the results. In this case, we do not fill the intervals lacking energy measurements with averages from other periods, hence the troughs in the 0h-8h interval. It can be observed that the consumptions at different time intervals do not differ greatly because most of the energy consumed by the platform may be idle consumption (*i.e.* the minimum

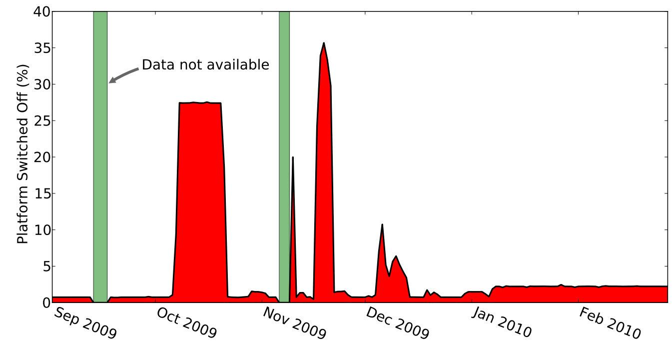

Fig. 4. Percentage of the platform switched off over six months.

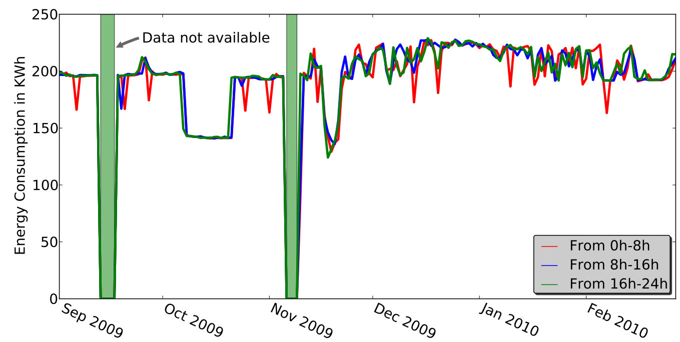

Fig. 5. Energy consumption at different periods of the day.

energy consumed by servers even when they are not executing user applications). This distance between static and dynamic energy consumptions would probably be larger if the platform employed more recent hardware with technologies such as CPU frequency scaling. In spite of the small dynamic energy consumption, a better understanding of resource usage patterns of reservation requests could lead to system optimisations for reducing the energy consumed. Hence, the next section provides details on the energy consumed by the resource reservation requests.

IV. Resource Reservations

When considering the energy consumed by individual resource reservations, the lack of measurement data over some periods may lead mainly to two problems:

- the total energy consumption of a reservation cannot be determined because measurements are not available for part of its duration; and
- the reservation falls entirely within a failure interval.

The lack of measurements when computing the energy consumed by resource reservations is addressed by:

- obtaining the average consumption in Watts of the seconds when data is available; and
- using this average as the consumption at each second when a measurement is missing.

This approach is used when measurements are available for at least 50% of the reservation's duration. The results ignore reservations whose partial power consumption data accounts for less than 50% of their duration. Table II summarises the energy consumed by different reservation categories. The second column of the table shows the overall energy consumed by the reservations under each category, whereas the third column presents the average number of watts per node. The watts per node wn_j of a reservation request j is given by $wn_j = cs_j \div m_j \div p_j$, where cs_j is energy consumed by reservation j, m_j is the number of nodes required by the reservation, p_j is the duration of the request.

TABLE II
ENERGY CONSUMED BY DIFFERENT RESERVATION CATEGORIES.

Number of nodes used	Number of reservations	Overall consumption in KWh	Average Watts per Node
1 to 2	3844	1529.77	197.23
3 to 5	681	1203.71	192.36
6 to 10	611	7384.92	200.44
11 to 30	7408	35371.17	216.49
31 to 70	205	9821.79	178.78
71 to 100	45	1918.33	185.91
101 to 135	50	5447.23	185.85

The values of average in watts demonstrate that reservations that allocate up to 30 nodes tend to consume more energy. The highest average is given by reservations that require between 11 and 30 nodes. If one considers that resource utilisation is proportional to dynamic energy usage, it can be observed that the low average watts of large reservations demonstrate: that small reservations are more resource intensive; and that users may make large reservations to ensure exclusive use of a cluster or the whole infrastructure. This is a common practice for network experiments, where concurrent usage of the platform by multiple applications could compromise the experiments and undermine the obtained results. This shows that optimisations can be made to allow unused machines to be switched off or placed on a low consumption status when they are not required.

To better understand the energy consumption of requests of different sizes (*i.e.* number of nodes) and lengths (*i.e.* duration), reservations were also grouped into four categories [13]: based on their duration – Short(S) or Long(L) – and the number of nodes requested – Narrow(N) or Wide(W). Table III provides an overview of the reservation groups and the percentages of reservations that fall in each group.

TABLE III
GROUPS OF RESERVATIONS CONSIDERING THEIR SIZES AND LENGTHS.

Duration	Number of nodes	
	$\leq$ 10 nodes	$>$ 10 nodes
$\leq$ 1 hour	SN (27.18%)	SW (47.26%)
$>$ 1 hour	LN (12.81%)	LW (12.75%)

Moreover, we compute the Average Weighted Power Consumption (AWPC) for each reservation group. The AWPC measures how much energy reservations use according to their resource consumption.

$$AWPC = \frac{\sum_{j \in \tau_k} p_j \cdot m_j \cdot wn_j}{\sum_{j \in \tau_k} p_j \cdot m_j} \tag{1}$$

The AWPC is given by Equation 1, where m_j is the number of nodes required by reservation request j, p_j is the duration of the request, and wn_j is the number of watts per node of the reservation. The resource consumption $(p_j \cdot m_j)$ of each request j is used as weight. This metric aims to illustrate the impact of large reservations on the energy consumption and show whether future optimisations can be performed to reduce the energy they consume.

TABLE IV
ENERGY CONSUMED BY DIFFERENT RESERVATION GROUPS.

Reservation Group	Overall Consumption in KWh	Average Watts per Node	AWPC
SN	278.29	198.78	191.46
SW	6406.21	217.72	207.09
LN	9840.11	193.10	197.84
LW	46152.31	205.42	184.23

Table IV summarises the results. The AWEC of the different groups of reservations show a small contrast with the average watts per node. For example, when considering large requests, the AWEC of narrow reservations is larger than the average watts per node, whereas the opposite is the case for wide reservations. Once more, this behaviour can be a result of higher resource utilisation during the narrow reservations. This gives room for further optimisations to minimise the energy consumption of wide reservations, such as setting unused resources to low power-consumption modes. Hence, the metrics described in this section show that resources are not fully utilised during some reservations, specially those that require a large number of resources. In the next sections, we attempt to evaluate the cost of resource under-utilisation in terms of energy wastage.

V. THE ENERGY COST OF RESOURCE UNDER-UTILISATION

As discussed beforehand, recent computer hardware and operating systems have offered a range of techniques to manage the power drawn by servers. The challenge posed to middleware designers is how to benefit from these techniques for making distributed systems more energy efficient, hence reducing both their energy footprint and the costs associated with managing a Grid infrastructure. Tackling this challenge requires an understanding of how users utilise the infrastructure and how their applications behave in terms of resource usage and energy consumption. Previous study on the energy consumption of Grid'5000 has shown that substantial energy savings can be achieved if users accept to change the start or finish times of their reservations, so that they can be aggregated, creating free windows during which unused servers can be switched off [10]. Additional savings could be made by analysing the power consumption data and crossing it with information from other system components such as the RMS or scheduler.

The "interactive" advance reservations of Grid'5000 are used here to demonstrate how Grid middleware and policies can be improved when information on power consumption of

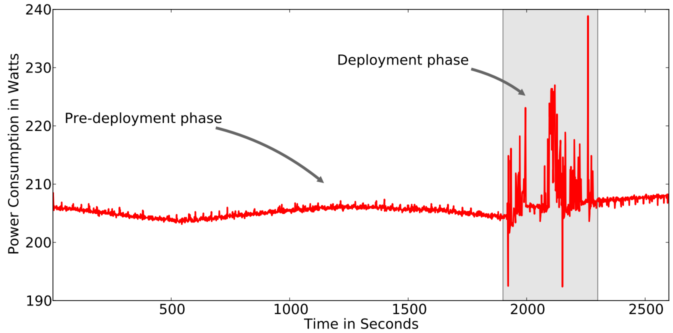

Fig. 6. A server's consumption during a reservation with late deployment.

servers and resource utilisation logs are made available to developers. Under an interactive reservation, a user reserves a group of servers on which she deploys an environment that contains the whole operating system customised to her application. The user has the option to carry out a late environment deployment, and in this case, OAR does not deploy any system image on the reserved nodes at the start of a reservation. The deployment is started at the user's request when she is ready to use the reserved nodes. Figure 6 shows the power consumed by a server of the Lyon site during part of a reservation that starts at 10am on a working day. The noisy measurements obtained after 1900 seconds are typical of an environment-deployment phase during which the disk image is copied to the node that is later rebooted with the new operating system. This figure shows that during the period that proceeds the deployment, hereafter termed as *pre-deployment phase*, the server was not utilised, hence wasting resources and energy. If mechanisms for identifying and predicting these behaviours are incorporated into middleware design, the unused servers can be switched off, thus minimising resource wastage and consequently improving the energy efficiency of the infrastructure.

Information from kadeploy is used here to compute the energy consumed during long pre-deployment phases in the Lyon site. Although kadeploy maintains a log of deployments carried out by the users, this information was not available during the whole six-month period. In the case of the Lyon site, the log contained information about the deployments carried out between 13 October 2009 and 27 February 2010. For identifying to which reservation a deployment belongs, we cross the deployment information with data from OAR's reservation log. After matching reservations and deployments, for each reservation, the pre-deployment phase is computed as the difference between the time at which the first environment deployment has been performed and the start time of the reservation. The energy consumed during the pre-deployment phase is the sum of the energy consumed by the reserved servers during this period. Reservations whose pre-deployment phase is shorter than 5 minutes are ignored because we consider it a reasonable period for obtaining the information required to start a deployment, such as which nodes have been allocated by OAR. When a reservation does not have a deployment, it is considered that the user has not utilised the reserved resources, and in this case, the pre-deployment phase is equal to the length of the reservation.

The overall energy consumed by all servers during the period from 13 October 2009 and 27 February 2010 is approximately 81.223 MWh, whereas the consumption during the pre-deployment phase of reservations is about 11,709.18 KWh. This amount corresponds to approximately 14.41% of the overall energy consumed by the server infrastructure. This is an optimistic scenario since it does not consider deployment failures and assumes that users utilise the resources from the deployment until the end of a reservation. However, the results show that a reservation scheme without a cancelation policy for 'no shows' [14] – the case of Grid'5000 – can be costly in terms of resource under-utilisation and energy consumption.

A. *The Power Consumptions of Nodes*

As discussed beforehand, the server infrastructure consumes a substantial amount of power even when it is idle. This section attempts to analyse the static power consumed by the servers of the studied infrastructure.

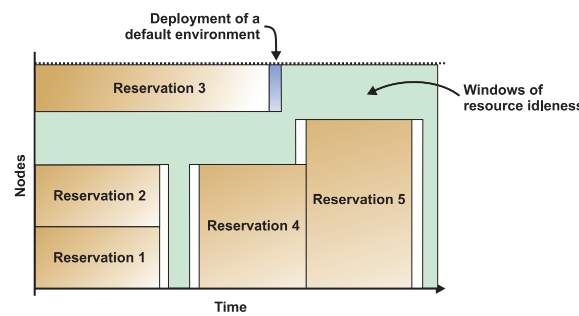

Fig. 7. Obtaining the windows of resource idleness.

In order to compute the static power consumption of servers – *i.e.* the idle power consumption – only the power consumption measurements obtained during windows of resource idleness were utilised. Figure 7 depicts the concept of windows of idleness. When computing the idle consumption of a node, all power measurements carried out during reservations for that node are discarded. We also disregard measurements within an interval of five minutes before and after reservations, and during environment deployment periods. Moreover, we disregard all power measurements carried out before 1 December 2009 in order to avoid events such as air-conditioning problems and large gaps in the power consumption data.

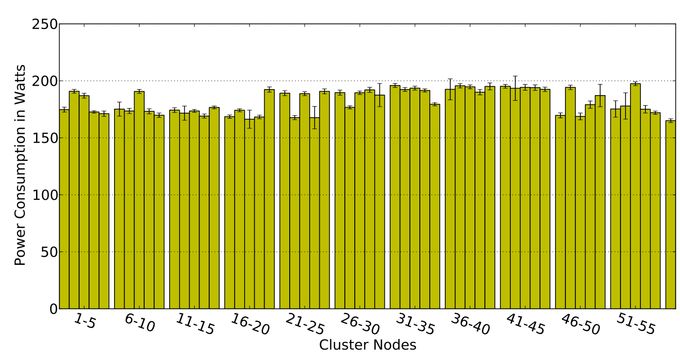

Fig. 8. Average idle power consumption of capricorne's nodes.

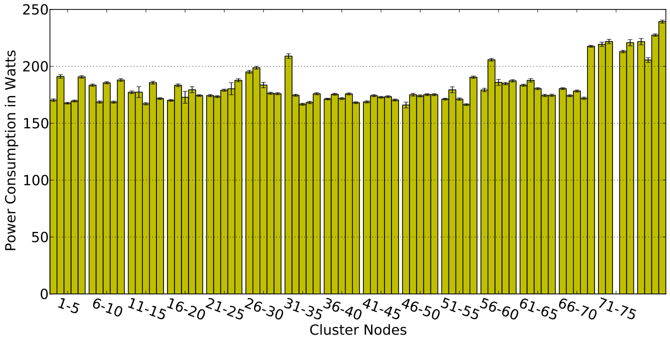

Fig. 9. Average idle power consumption of sagittaire's nodes.

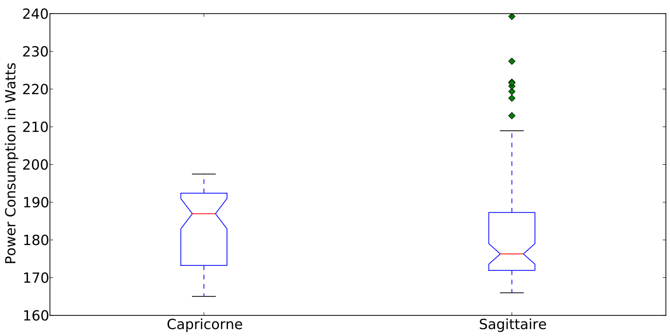

Fig. 10. Average idle power consumption of servers.

Figure 8 presents the average idle power consumption of capricorne cluster's nodes whereas Figure 9 shows the idle consumption of sagittaire's servers. The average idle consumption of sagittaire's 73th node could not be computed because it has undergone large periods of failures and, therefore, there were insufficient power consumption measurements. Figure 10 reports the same information showing how the average consumption of nodes is spread. The heterogeneous idle consumption of sagittaire nodes was expected since some nodes of this cluster (*i.e.* from 70th to 79th) have been acquired before the remaining nodes and their higher power consumption has been identified in previous work using this cluster. The different consumptions of capricorne nodes and the remaining of sagittaire's can depend on several factors. The idle power consumption of the nodes is not constant and is greatly impacted by the room temperature [15], [16]. Hence, when a node is idle, its power consumption increases if the nodes close to it are running and producing heat. The locality of a node in the rack seems to be another influencing factor because the air-conditioning pumps cold air into the room through the floor. Nodes located at the top of the rack tend to heat more and consequently consume more power. Moreover, some of the machines have been in production for a few years, specially capricorne nodes, and some of their components have been replaced (*e.g.* hard-disks, memory and network cards). Although the information about the replaced components is not on hand, it certainly leads to a hidden heterogeneity which can also influence the idle power consumption of nodes. Furthermore, we believe that mechanical parts such as fans and hard-disks can wear out and consume more energy over time. However, further investigation should be carried out to support this claim.

To estimate the average dynamic power consumption of nodes – *i.e.* the consumption incurred by running user applications – we use the measurements obtained during the resource reservations. All power consumption measurements obtained between reservations are discarded. In addition, only measurements obtained after 1 December 2009 are considered when computing the averages and standard deviations. We term as "busy power consumption" of a node the average electrical power that it draws during resource reservations. The "dynamic consumption" of a node is given as the difference between its busy and idle consumptions.

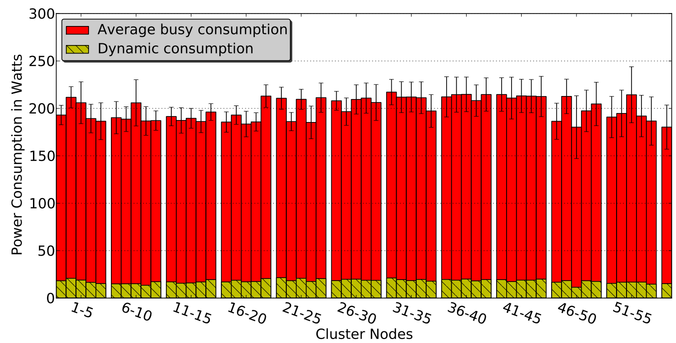

Fig. 11. Average busy power consumption of capricorne's nodes.

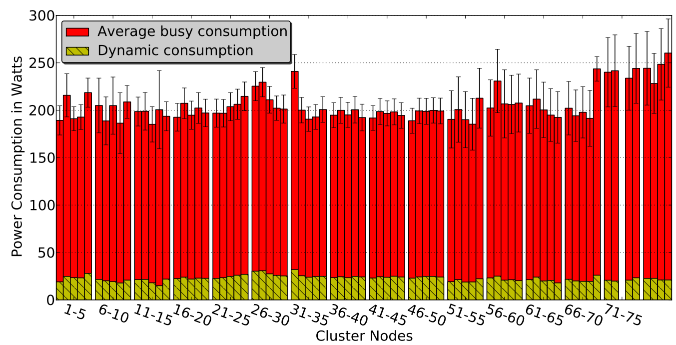

Fig. 12. Average busy power consumption of sagittaire's nodes.

Figure 11 and Figure 12 present the busy and dynamic consumptions for nodes of capricorne and sagittaire respectively. It can be observed that both the average busy consumption and the standard deviations are higher than those presented in the idle consumption graphs. Although obvious, it shows the impact of resource usage on power consumption. However, as illustrated by the dynamic consumption bars, the additional power consumed by servers during the reservations is small when compared by the overall power drawn by the servers when idle.

B. Application-Driven Energy Consumption

Having estimated the idle power consumption of servers, this section intends to calculate the amount of energy consumed by executing applications. The *application-driven energy consumption* – or dynamic consumption – is the share of energy consumed by using server resources during the reservations, hence disregarding the servers' idle consumptions.

Calculating the application-driven consumption is important for evaluating allocation policies that attempt to curb the energy consumption by using server-level power management techniques such as CPU throttling. Formally speaking, the application-driven energy consumption a_j during a resource reservation request j is given by Equation 2:

$$a_j = \int_{s_j}^{e_j} h \; dt \tag{2}$$

where s_j is the start time of reservation j, e_j is the end time of reservation j, and h is the dynamic consumption of power measurements. The dynamic power consumption h of a measurement is given by Equation 3:

$$h = \sum_{m \in M_j} w_m - i_m \tag{3}$$

where M_j is the set of nodes used by reservation j, w_m is the measured power consumption of node m and i_m is the idle power consumption of node m. When computing the application-driven consumption, we discard reservations for which partial power measurements are available.

The application-driven energy consumption for the six-month period was 1,823.47 KWh whereas the total amount of energy consumed by servers during the resource reservations was 59,688.53 KWh. Hence, the application-driven consumption accounts to approximately 3.05% of the energy consumed during the reservations. These results show that allocation policies using power management techniques such as CPU throttling are not appealing for the studied experimental infrastructure since the load posed by applications seem to be small when compared to the idle energy consumption. However, these results emphasise the need for management techniques that attempt to improve the energy efficiency of the platform by curbing the idle server consumption (*e.g.* approaches based on switching off unused resources).

VI. Related Work

Several techniques can be used to make computing infrastructures more energy efficient [17], [18], [19], [9]. The use of VMs [20] brings several benefits including environment and performance isolations; improved resource utilisation by enabling workload consolidation; and resource provisioning on demand. Nevertheless, in order to devise consolidation policies for improving the energy-efficiency of computing infrastructures [21], such power management techniques should be analysed and used carefully.

Monitoring the power consumption of an infrastructure is the first step to understanding hardware and user behaviours and analyse the impact in energy consumption of applying different power management approaches. Moreover, monitoring the consumption can help evaluate the ratio of performance to energy spent under different application conditions [22]. Approaches for monitoring the power consumption of house and office environments have been proposed [23], [24]. Monitoring large-scale distributed systems such as Grids and Clouds still remains a challenge. Although various solutions have been proposed for monitoring the usage of resources [25], [26], [27] (*e.g.* CPU, storage and network) to improve the performance of applications, only a few systems monitor the energy usage of infrastructures [10], [28]. The present work analyses the data collected from monitoring the energy consumed by the server infrastructure of a site that is part of an experimental Grid [11].

The negative impact of reservation requests on overall resource utilisation of clusters has been demonstrated in the literature [29], [30], [31]. Solutions have been proposed for minimising this impact [32], [33], which can enable the creation of windows of resource availability during which unused servers can potentially be switched off [10]. This report demonstrates that additional energy savings can be made if we understand how users utilise resources during reservations.

Attempts to establish minimum energy efficiency requisites for data centres have also been made [22] – *e.g.* ENERGY STAR program requirements for computer servers and storage devices.[2] The definition of standard metrics and benchmarks for evaluating the efficiency of compute and storage infrastructure of data centres is desired, and monitoring the power consumption is an important aspect when comparing the effectiveness of various hardware configurations. This work intends to shed some light on the power consumption of Grids and, by understanding the application and user behaviours, show how they can benefit from the power management features offered by compute and storage hardware.

VII. Lessons Learnt

This section summarises the results obtained by analysing the power consumption data and correlating it with information from the resource management system and the toolkit responsible for deploying the execution environments.

Initially, the overall energy consumption shows a small correlation between the amount of power drawn by servers and resource utilisation. This correlation was studied in more detail later where it was shown that the dynamic energy consumption (*i.e.* the energy consumed by servers during resource reservations minus the servers' idle consumption) corresponds to 3.05% of the overall energy used by servers during the periods they are reserved. The servers that comprise the studied infrastructure consume a large amount of power when they are idle (*i.e.* not actively executing user applications). In addition, the servers, although sometimes similar in terms of hardware configuration, present heterogeneous idle power consumption as demonstrated by a study on their idle and dynamic consumption. Hence, by using CPU throttling one would save at maximum 3.05% of energy consumed during reservations because the servers do not support DVFS.

Although the idle consumption of the server infrastructure is substantial, savings can be made during phases such as the pre-deployment of execution environments. The energy used to power servers during the environment pre-deployment phase corresponds to approximately 14.41% of the overall energy consumed by the server infrastructure during the studied period. Part of this energy can be saved if servers are switched off before users deploy their execution environments or through

[2] http://www.energystar.gov/

the adoption of pro-active policies in which, for example, servers are turned off by default when the deployment does not take place within a few minutes after the start of a reservation.

An analysis of resource reservation requests showed that the average power consumption per node is higher when a small number of nodes is requested. This demonstrates that small reservations are more resource intensive and might reveal that some of the large reservations are made by users who demand exclusive use of a cluster or the entire platform. Hence, the latter reservations are not resource intensive and consequently not energy consuming. A further study considering the duration and number of nodes requested by reservations revealed that during long and wide reservations the average consumption per node is small when compared to short reservations, hence emphasising the low resource utilisation during long reservations. This analysis showed that further energy can be saved by understanding how users utilise resources during reservations and by, among other approaches, switching off unused servers in cases where users require exclusive use of the whole platform during a network experiment and not all servers are required.

VIII. Conclusion

This paper analysed information on the energy consumed by servers of a Grid'5000 site. It showed information on the energy consumed by the overall infrastructure, by different groups of reservations and by the servers. In addition, it demonstrated that optimisations can be made to save energy by switching off unused resources during environment pre-deployment phases.

In the future we will perform a more detailed statistical analysis of the power consumption of servers during reservations. In addition, we intend to use the analysis results to devise resource allocation policies for saving energy during resource reservations.

Acknowledgment

We thank Marco A. S. Netto (Pontifical Catholic University of Rio Grande do Sul) and Guilherme Koslovski (ENS de Lyon) for their feedback on a preliminary version of this work; Dominique Ponsard (ENS de Lyon) for providing the computer server to host the database and tools used for analysing the energy consumption data; and Jean-Patrick Gelas (Université de Lyon) for helping in setting up the hardware and software infrastructure required to monitor the power consumed by the computer servers. Some experiments of this paper were performed on the Grid'5000 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)

References

[1] R. Bolze et al., "Grid'5000: a large scale and highly reconfigurable experimental Grid testbed," *Intern. Journal of High Perf. Comp. App.*, vol. 20, no. 4, pp. 481–494, Nov. 2006.

[2] DAS-2, Dutch University Backbone, 2006.

[3] Enabling Grids for E-sciencE (EGEE), http://www.eu-egee.org/, 2005.

[4] C. Catlett, P. Beckman, D. Skow, and I. Foster, "Creating and operating national-scale cyberinfrastructure services," *Cyberinfrast. Technology Watch Quarterly*, vol. 2, no. 2, pp. 2–10, May 2006.

[5] T. Dunning and R. Nandkumar, "International cyberinfrastructure: Activities around the globe," *Cyberinfrast. Technology Watch Quarterly*, vol. 2, no. 1, Feb. 2006.

[6] M. Armbrust, "Above the clouds: A Berkeley view of Cloud computing," Electrical Engineering and Computer Sciences, Univ. of California at Berkeley, Berkeley, USA, Technical report UCB/EECS-2009-28, Feb. 2009.

[7] A. Weiss, "Computing in the Clouds," *netWorker*, vol. 11, no. 4, pp. 16–25, Dec. 2007.

[8] Greenpeace International, "Make IT green: Cloud computing and its contribution to climate change," http://www.greenpeace.org/international/press/reports/make-it-green-cloud-computing, Mar. 2010.

[9] G. Semeraro et al., "Energy-efficient processor design using multiple clock domains with dynamic voltage and frequency scaling," in *8th HPCA*, Washington, DC, USA, 2002, p. 29.

[10] A.-C. Orgerie, L. Lefèvre, and J.-P. Gelas, "Save watts in your Grid: Green strategies for energy-aware framework in large scale distributed systems," in *14th IEEE ICPADS*, Melbourne, Australia, 2008, pp. 171–178.

[11] M. D. de Assunção, J.-P. Gelas, L. Lefèvre, and A.-C. Orgerie, "The green Grid'5000: Instrumenting and using a Grid with energy sensors," in *5th INGRID*, Poznan, Poland, 2010.

[12] N. Capit et al., "A batch scheduler with high level components," in *5th IEEE CCGrid*, vol. 2, Washington, DC, USA, 2005, pp. 776–783.

[13] S. Srinivasan, R. Kettimuthu, V. Subramani, and P. Sadayappan, "Selective reservation strategies for backfill job scheduling," in *8th JSSPP*, ser. LNCS, vol. 2537, London, UK, 2002, pp. 55–71.

[14] A. Sulistio, K. H. Kim, and R. Buyya, "Managing cancellations and no-shows of reservations with overbooking to increase resource revenue," in *8th IEEE CCGrid 2008*, Lyon, France, 2008, pp. 267–276.

[15] Q. Tang, S. K. S. Gupta, and G. Varsamopoulos, "Thermal-aware task scheduling for data centers through minimizing heat recirculation," in *IEEE Cluster 2007*, Washington, DC, USA, 2007, pp. 129–138.

[16] G. Varsamopoulos, A. Banerjee, and S. K. S. Gupta, "Energy efficiency of thermal-aware job scheduling algorithms under various cooling models," in *2nd IC3*, ser. Comm. in Comp. and Inform. Science, vol. 40, Berlin/Heidelberg, 2009, pp. 568–580.

[17] J. Moore, J. Chase, P. Ranganathan, and R. Sharma, "Making scheduling "cool": Temperature-aware workload placement in data centers," in *USENIX ATEC*, Berkeley, USA, 2005, pp. 5–5.

[18] X. Fan, W.-D. Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," in *34th ISCA*, New York, USA, 2007, pp. 13–23.

[19] R. K. Sharma, C. E. Bash, C. D. Patel, R. J. Friedrich, and J. S. Chase, "Balance of power: Dynamic thermal management for internet data centers," *IEEE Internet Comp.*, vol. 9, no. 1, pp. 42–49, Jan./Feb. 2005.

[20] P. Barham et al., "Xen and the art of virtualization," in *19th ACM SOSP*, New York, USA, 2003, pp. 164–177.

[21] S. Srikantaiah, A. Kansal, and F. Zhao, "Energy aware consolidation for Cloud computing," in *HotPower*. San Diego, USA: USENIX, 2008.

[22] M. Poess et al., "Energy benchmarks: A detailed analysis," in *1st e-Energy 2010*, Passau, Germany, 2010, pp. 131–140.

[23] S. Taherian, M. Pias, G. Coulouris, and J. Crowcroft, "Profiling energy use in households and office spaces," in *1st e-Energy 2010*, Passau, Germany, 2010, pp. 21–30.

[24] F. Mattern, T. Staake, and M. Weiss, "ICT for green – how computers can help us to conserve energy," in *1st e-Energy*, Passau, Germany, 2010, pp. 1–10.

[25] S. Zanikolas and R. Sakellariou, "A taxonomy of grid monitoring systems," *Future Gen. Comp. Systems*, vol. 21, no. 1, pp. 163–188, 2005.

[26] S. Andreozzi et al., "GridICE: a monitoring service for Grid systems," *Future Gen. Comp. Systems*, vol. 21, no. 4, pp. 559–571, 2005.

[27] H.-L. Truong and T. Fahringer, "SCALEA-G: A unified monitoring and performance analysis system for the Grid," *Scientific Programming*, vol. 12, no. 4, pp. 225–237, 2004.

[28] T. Singh and P. K. Vara, "Smart metering the clouds," in *18th IEEE WETICE*, Los Alamitos, USA, 2009, pp. 66–71.

[29] M. W. Margo, K. Yoshimoto, P. Kovatch, and P. Andrews, "Impact of reservations on production job scheduling," in *JSSPP*, ser. LNCS, vol. 4942, Berlin/Heidelberg, Germany, 2007, pp. 116–131.

[30] Q. Snell, M. Clement, D. Jackson, and C. Gregory, "The performance impact of advance reservation meta-scheduling," in *JSSPP*, ser. LNCS, vol. 1911, 2000, pp. 137–153.

[31] W. Smith, I. Foster, and V. Taylor, "Scheduling with advanced reservations," in *14th IPDPS*, Cancun, Mexico, 2000, pp. 127–132.

[32] T. Röblitz, F. Schintke, and J. Wendler, "Elastic Grid reservations with user-defined optimization policies," in *AGridM*, Antibes Juan-les-Pins, France, 2004.

[33] M. A. S. Netto, K. Bubendorfer, and R. Buyya, "SLA-Based advance reservations with flexible and adaptive time QoS parameters," in *5th ICSOC*, Berlin/Heidelberg, Germany, 2007, pp. 119–131.