
A survey into performance and energy efficiency in HPC, cloud and big data environments

Eduardo Camilo Inacio* and
Mario A.R. Dantas

Department of Informatics and Statistics (INE),
Federal University of Santa Catarina (UFSC),
Florianópolis, SC, Brazil
Email: eduardo.camilo@posgrad.ufsc.br
Email: mario.dantas@ufsc.br
*Corresponding author

Abstract: The growing demand for performance observed in many organisations can still be considered the main motivator for the evolution of high performance computing and, more recently, cloud environments. Part of this demand regards the need to deal with large and complex datasets, called *big data*. Performance improvement in such environments begins to be limited by energy consumption. Workload characterisation is a well-known approach to reproducing systems' behaviour. However, there are several methodologies, techniques and parameters that can be considered for a workload characterisation. As a result, we present a differentiated survey on workload characterisation focusing on performance and energy efficiency improvement on HPC, cloud and big data environments. After an extensive review and classification of research works, our study indicates that around 56.4% of the papers reviewed offer contributions to performance and energy efficiency improvement, and the growing interest in this subject has a rate of 7.86% per year.

Keywords: workload characterisation; performance improvement; energy efficiency; high performance computing; HPC; cloud computing; big data.

Reference to this paper should be made as follows: Inacio, E.C. and Dantas, M.A.R. (2014) 'A survey into performance and energy efficiency in HPC, cloud and big data environments', *Int. J. Networking and Virtual Organisations*, Vol. 14, No. 4, pp.299–318.

Biographical notes: Eduardo Camilo Inacio received his BSc in Information Systems from the Federal University of Santa Catarina, Florianópolis, Brazil, in 2013. He is currently an MSc student under the supervision of Professor Mario A.R. Dantas. His research interests are distributed systems and high performance computing.

Mario A.R. Dantas received his PhD in Computer Science from the University of Southampton in the UK in 1997, BSc in Electrical Engineering from Gama Filho University in 1984 and System Analyst degree in 1986 from the Catholic University, both in Rio de Janeiro, Brazil. He worked in the E&P Department at a Brazilian petroleum company for 12 years, as a System Analyst specialising in distributed scientific and engineering applications. He joined the academia in 1998 (UnB and then UFSC), where he researches and teaches, as an Associate Professor, subjects related to distributed systems and networking in undergraduate and graduate courses.

1 Introduction

As the problems attacked by several organisations become larger and more complex, so does the demand for computing resources. This has been one of the main motivations for the development of more powerful computer systems over the years. Historically, this demand for computational resources is being supplied by high performance computing (HPC) environments, such as supercomputers, computer clusters and grids. However, we have recently been observing the rise of cloud computing environments as a provider of such resource demands. Another characteristic of modern applications is the huge amount of data that needs to be processed. These data usually come from different sets of devices (e.g., satellites, telescopes or sensors) and procedures (e.g., experiments or simulations). Thus, big data techniques are being increasingly used in ordinary corporate applications.

Therefore, it is clear the need to develop computational architectures with even better performance to support current and future applications. However, this development begins to face limitations such as miniaturisation of components, thermal control, and energy consumption. As an example of this challenge, let us consider the initiative of building an *exascale* system (Kogge et al., 2008). This initiative targets to build a supercomputer capable of performing 10^{18} floating-point operations per second (FLOPS) with an energy consumption of no more than 20 MW (Stevens and White, 2009). Considering the most powerful supercomputer according to the Top500 (2013b) list of November 2013, the *Tianhe-2*, an *exascale* system would perform approximately 29.5 times more FLOPS. If we consider the ratio of FLOPS/W (floating-point operations per second per watt of energy consumption) of *Tianhe-2*, an *exascale* system would consume approximately 526 MW. This consumption is more than 26 times above the suggested target. Even considering the FLOPS/W ratio of the most energy-efficient supercomputer according to the Green500 (2013) list of November 2013, the *Tsubame-KFC*, an *exascale* system would have an energy consumption of more than 11 times above the target (approximately 222 MW).

Performance improvement can be addressed through the use of faster components, the use of more efficient algorithms, or by parallelising and distributing the computation (Dantas, 2005). Lower energy consumption of a computing environment can be achieved with the use of virtualisation and server consolidation techniques (e.g., Dhiman et al., 2010a), dynamic voltage and frequency scaling (DVFS) mechanisms (e.g., Huang and Feng, 2009), or even temporary inactivation of idle components (e.g., Meisner et al., 2011). To make better decisions about the appropriate approach in each case, it is essential a clear understanding on the behaviour of the applications and of the environments where those will be executed. Workload characterisation is a well-known approach to support that decision.

Workload characterisation is the process in which models are constructed allowing to describe and reproduce the behaviour of the workload of a computing system (Elnaffar and Martin, 2006). It is essential to understand the characteristics of applications' workloads in order to design efficient and cost-effective computing architectures (Obaidat and Boudriga, 2010). However, the development of accurate workload models involves challenges such as defining the desired level of detail, selecting a representative set of metrics, choosing the methods and tools to measure the metrics, and validating the generated model.

In this paper, we present a survey of the literature with a focus on the use of workload characterisation for performance and energy efficiency improvement for HPC, cloud and big data environments. The aim of this survey is to provide researchers and organisations with a better understanding of the state-of-the-art in order to apply workload characterisation to both improve performance and to reduce energy consumption in HPC, cloud and big data environments.

This paper is organised as follows. In Section 2, we present the main concepts related to workload characterisation, HPC, cloud and big data. We discuss related works in Section 3. Section 4 presents this research methodology. We present and discuss the results in Section 5. In Section 6, we comment on some of the contributions to HPC, cloud and big data environments observed in this study. Our conclusions and directions for future works are presented in Section 7.

2 Background

In this section, we briefly present the main conceptual aspects regarding workload characterisation, HPC, cloud and big data.

2.1 Workload characterisation

Workload characterisation is the process in which models are constructed allowing to describe and reproduce the behaviour of the workload of a computing system (Elnaffar and Martin, 2006). It is essential to understand the characteristics of applications' workloads to design efficient and cost-effective computing architectures (Obaidat and Boudriga, 2010). This fact is explained by the influence on the performance of a system through the characteristics of its hardware and software components as well as the load that must be processed. Workload analysis plays a key role in all the studies where the performance indices of a system are to be determined (Calzarossa and Serazzi, 1993).

The development of accurate workload models involves challenges such as defining the desired level of detail, selecting a representative set of metrics, choosing the methods and tools to measure the metrics, and validating the generated method. According to Ferrari et al. (1983), the main steps to produce workload models are fourfold: formulation, collection of parameters, statistical analysis of the measured data (preliminary analysis, analysis of the parameter distributions, sampling, static analysis, and dynamic analysis), and representativeness.

Characterisation techniques can be functionally classified as static or dynamic (Elnaffar and Martin, 2006). Static techniques explore characteristics of the workload that usually do not change over time. Descriptive statistics (e.g., average, standard deviation), single- and multi-parameter histogram, factor analysis, and clustering are examples of static techniques. Dynamic systems, on the other hand, describe not only the workload characteristics but also the way they change over time. Examples of dynamic techniques are Markov models, prediction using neural networks, moving average, exponential smoothing, regression methods, user behaviour graphs, and probabilistic attributed context free grammar.

2.2 HPC, cloud and big data

HPC environments can still be considered the major infrastructure for scientific computing. In traditional owner-centric HPC environments, the locally owned resources are managed by a single administrative domain (Mateescu et al., 2011). Clusters have been the predominant architecture for such environments since November 2003, according to the Top500 (2013a) Supercomputer site, achieving 84.6% of the system share by November 2013. In distributed HPC environments, such as grid computing, users can make use of both local and external resources, managed by virtual organisations (Foster and Kesselman, 2001).

Cloud computing denotes a distributed computing infrastructure from which businesses and users are able to access applications from anywhere in the world on demand (Buyya et al., 2009). Clouds can be private, when the resources are internally owned and available to the users of an organisation, or can be public, when the resources are externally owned and accesses are typically on a pay-as-you-go basis. Major service provisioning models in cloud computing environments are threefold: software as a service (SaaS), platform as a service (PaaS) and infrastructure as a service (IaaS). Recently, we have seen the rise of researches (e.g., Mateescu et al., 2011; Benedict, 2012) targeting the execution of HPC workloads on cloud computing environments.

According to the definition proposed by Gartner (2013), “Big data is high volume, high velocity, and/or high variety information assets that require new forms of processing to enable enhanced decision making, insight discovery and process optimization”. This type of complex and huge dataset can be observed in many scientific areas (e.g., CERN, 2013; Soares et al., 2013).

3 Related works

Calzarossa and Serazzi (1993) carried out a survey on workload characterisation focusing on the steps required for the construction of workload models and the sets of parameters to be considered in various types of systems (i.e., batch, interactive, database, network-based, parallel, and supercomputer). Calzarossa et al. (2000) extended their work in Calzarossa and Serazzi (1993), discussing methodologies and modelling techniques for workload characterisation through two use case studies: client/server environments and parallel systems.

In Elnaffar and Marin (2002), and Elnaffar and Martin (2006), some of those case studies presented in Calzarossa and Serazzi (1993) are revisited, and new ones are included (i.e., WWW and client/server), emphasising the techniques used to characterise workloads in such systems. They present a classification of the most common characterisation techniques according to their ability to extract the static and dynamic aspects of the workload. Finally, their work presents a framework for the construction of workload models based on the methodology proposed by Ferrari et al. (1983).

Our research paper presents differential contributions from the prior studies in at least three aspects. Firstly, to the best of our knowledge, this is the first time that a survey regarding the use of workload characterisation targeting energy issues is conducted. Secondly, while the works of Calzarossa and Serazzi (1993), Calzarossa et al. (2000), Elnaffar and Marin (2002), and Elnaffar and Martin (2006) focus on the methodologies and techniques for construction of workload models, and parameters associated with the

systems which are subject of the process, our survey is dedicated to providing a better understanding about the use of workload characterisation to address questions about performance and energy efficiency improvement. Finally, this study fills a gap in the state-of-the-art knowledge about workload characterisation, given that the last survey on this subject dates to 2006 (Elnaffar and Martin, 2006).

4 Research methodology

This study can be classified, according to Gil (2010), as exploratory research. These kinds of work are basically intended to provide greater familiarity with the researched subject so as to make it more explicit or facilitate hypothesis building. The development process adopted on the present survey is similar to the work in Nazário et al. (2012) and consisted of three phases: data collection, data analysis, and synthesis and representation of data.

In order to guide the process, targeting at a better understanding of the subject, we considered the following questions in our research method.

- 1 How many publications on workload characterisation were there between 2009 and 2013?
- 2 What were the main media used to communicate contributions towards workload characterisation?
- 3 Which country is leading workload characterisation research?
- 4 What was the research topics most related to workload characterisation?
- 5 How many publications addressed performance and energy issues?
- 6 Are there any trends in the contributions addressing performance and energy issues?

4.1 Data collection

The data collection phase is where the publications considered in this study were collected. This phase comprised three steps: selection of the database, definition of the query criteria, and query execution for obtaining the papers. The database selection is an important step of this kind of study because it determines the research universe. Among the available scientific literature databases, we decided on *Scopus* (Elsevier, 2013). This decision was motivated by the acceptance of the *Scopus* database in the scientific community, and by the knowledge that the main conferences and journals in the area are indexed by this database.

The next step was to define the query criteria. It aims to define a subset of the research universe that contains publications relevant to this study. To build the query criteria, we adopted an iterative approach, as suggested in Kitchenham (2004). We started by deriving search terms based on analysis of previous literature surveys (i.e., Calzarossa and Serazzi, 1993; Calzarossa et al., 2000; Elnaffar and Marin, 2002; Elnaffar and Martin, 2006). Then, we tested several combinations of search terms, reviewing the search results, and, iteratively, we built the query criteria. As a result of this step, we defined the query criteria as publications containing in their title, abstract, and keywords,

the term ‘workload characterisation’, and at least one of the following terms: ‘energy’, ‘power’, ‘performance’, ‘HPC’, ‘cluster’, ‘grid’, ‘cloud’, and ‘big data’. In order to make the analysis phase feasible, because of the amount of publications found with these criteria, we chose to restrict the analysis to a period of five years, considering the publications between 2009 and 2013. It is noteworthy that other minor restrictions, not described here, were included in the query criteria in order to exclude irrelevant registers from the search results, such as conference presentations, journal editorials, and duplicated results. Additionally, it is important to mention that research papers similar to those of Birke et al. (2013), Islam et al. (2013) and Benmoussa et al. (2013) were not part of this survey, although they regard systems’ characterisations. This is due to the absence of the term ‘workload characterisation’ in their titles, keywords and abstracts. We observed that the formulation of query criteria capable of covering the whole set of relevant works and, at the same time, filtering out the non-relevant ones, is a common problem of surveys when the subject is not so normalised.

After defining the query criteria and concluding the data collection phase, the query was executed on *Scopus* and the publications were collected for analysis. In the query performed in November 2013, we obtained 269 publications. This amount of publications can be considered significant for the purposes of this study.

4.2 *Data analysis*

In the data analysis phase, the main objective is to examine in detail the set of publications collected, oriented by the questions proposed in this research methodology. This phase consists of three steps: preliminary review and classification of the collected publications, refinement and consolidation of the proposed classification, and analysis of bibliometrics indices. Bibliometrics, as defined in Pritchard (1969), is “the application of mathematics and statistical methods to books and other media of communication”. It provides a way to quantitatively analyse scientific literature.

The first step taken in this analysis was a preliminary review and classification of the collected publications. We read the titles and abstracts of all 269 papers collected and attributed classes identifying the contributions presented in each one. It is important to notice that some research papers were assigned to more than one class when identified more than one contribution relevant to this analysis. Through this preliminary review, 12 publications were excluded from the analysis because they were not related to the subject of this survey, although they met the query criteria. As a result of this step, we got the remaining 257 papers grouped by 34 classes.

Due to the large number of classes and uneven distribution of publications in these classes, a second review was conducted, aiming to refine and consolidate the classification proposal. This step considered the same information used in the first step, i.e., title and abstracts of the papers. However, in this step, unlike the previous one where reading did not consider any specific ordering of publications, the papers were selected according to the class they belonged to. The goal was to verify similarities among the papers assigned to the same class, and assess the proposed classification. Another goal of this step was to merge similar classes, whose definition was not significantly relevant to this survey or little representative as a publication. As a result of this step, we obtained 14 classes which can be considered representative of the contributions of the publications collected in this study.

With the information obtained until this point, we were able to analyse several bibliometrics indices which enabled us to quantitatively answer the questions presented at the beginning of this section. The following are the bibliometrics indices considered in this survey:

- Distribution of publications per year (addressing question 1).
- Distribution of publications per source type (i.e., conference proceedings, journals, books and book series) (addressing question 2).
- Distribution of publications per country (addressing question 3).
- Keywords used in the collected publications (addressing question 4).
- Distribution of publications according to the classification proposal (addressing question 5).
- Evolution of the representativeness of publications addressing performance and energy issues per year (addressing question 6).

4.3 Synthesis and representation of data

The final phase of this research methodology deals with the synthesis and representation of data processed in the previous phases. The goal is to organise and present the data in a way that provides a better understanding of the subject and facilitates the answering of the questions proposed at the beginning of this section. To achieve this goal, we used histograms, pie charts, line charts, and tag clouds. A detailed presentation of all the results obtained in this paper is given in the next section.

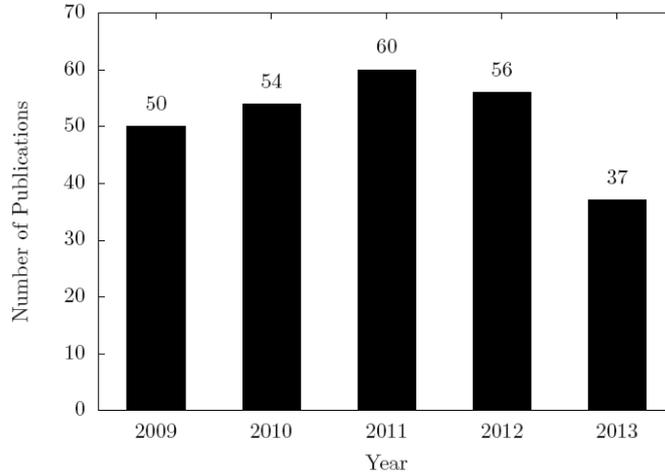
5 Results

As mentioned in the previous section, this survey was oriented by six research questions. By the end of the last section, we presented six bibliometrics indices that were chosen to answer the research questions. In this section, the results obtained and the reasoning about what was observed in these results are presented.

5.1 How many publications on workload characterisation were there between 2009 and 2013?

To answer this question, we used the distribution of publications per year along the period between 2009 and 2013. Figure 1 presents the results obtained. A statistical analysis of this index shows a small variation in the number of publications per year in the period considered (i.e., standard deviation = 8.82). The major variation is observed in 2013, when the number of publications is about 1.63 times the standard deviation below the mean (i.e., mean = 51.4). In this study, we were not able to accurately determine the cause for this reduction in 2013. However, we believe that when this study was conducted (i.e., November 2013) some 2013 publications had not been indexed yet by the *Scopus* database. If our assumption is correct, later query executions with the same criteria used in this survey should bring more publications.

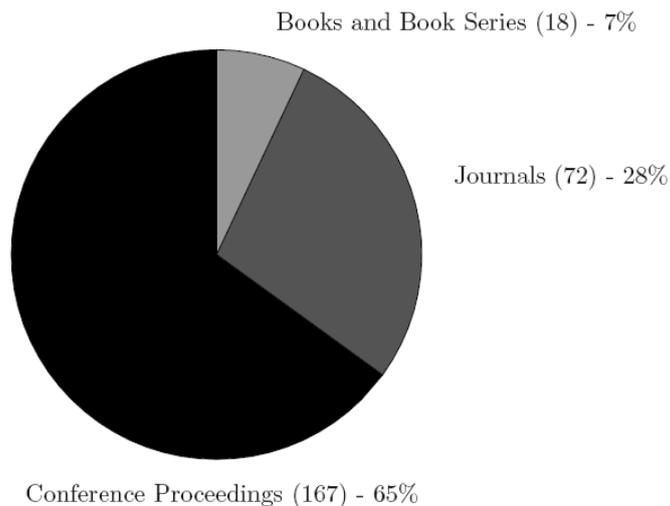
Figure 1 Distribution of publications per year along the period between 2009 and 2013



5.2 What was the main media used to communicate contributions to workload characterisation?

Figure 2 presents the distribution of publications per source type. We considered the following source types in our analysis: conference proceedings, journals, and books and book series. The behaviour observed in the publications on workload characterisation was not unexpected, according to this index. Most of the contributions were communicated to the scientific community through conference proceedings (65%), followed by journals (28%), and, to a minor extent, books and book series (7%).

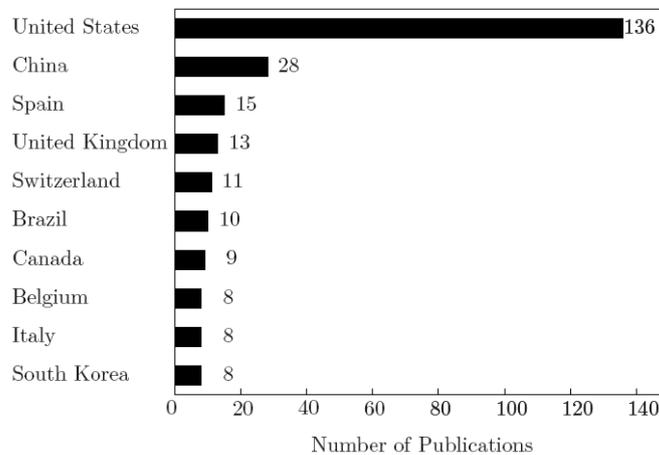
Figure 2 Distribution of publications according to their source type (i.e., conference proceedings, journals, and book and book series)



5.3 Which country is leading workload characterisation research?

The ten countries with more publications among those collected for this survey are presented in Figure 3. As can be seen in the chart, the contributions on this subject are predominantly published by US researchers. They actually correspond to 52.92% of the collected papers. It is important to notice that these numbers refer to the affiliations of the researchers, and not necessarily their country of origin. Another relevant aspect is that some publications have authors from different countries, which computes such publication to two (or more) countries. In terms of organisations, IBM Thomas J. Watson Research Center leads in number of publications (12), followed by IBM Research (9) and Oak Ridge National Laboratory (9).

Figure 3 Ten countries with more publications among those collected for this survey



5.4 What were the research topics most related to workload characterisation?

This question aims to find out which other research topics are more common in combination with workload characterisation. We used the keywords from the publications to answer this question. Keywords typically refer to other important topics in a publication. Then, evaluating the most used keywords in the set of publications, we were able to devise the research topics most related to workload characterisation.

Figure 4 presents a tag cloud with the keywords found in the collected publications. In tag clouds, the magnitude of the font of a tag represents its importance in the set. In other words, the bigger the font used for a keyword in the tag cloud, the more occurrences of this keyword were found in the collected publications. The keyword ‘workload characterisation’ was disregarded in this analysis because it does not contribute to the desired result, since we wish to observe other research topics related to it, and because it has a much higher number of occurrences than the other keywords, causing a distortion in the tag cloud. The tag cloud allows us to easily see, for example, that ‘benchmarking’ and ‘data centres’ are research topics more related to workload characterisation than ‘quality of service (QoS)’ or ‘multi core’.

5.5 *How many publications addressed performance and energy issues?*

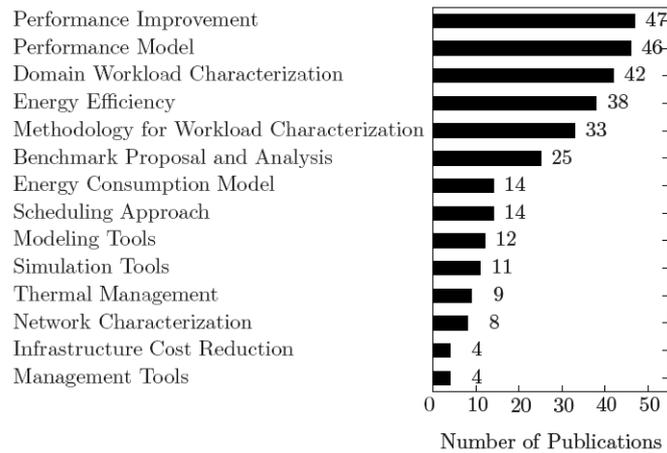
From this point on, we began to focus more on the objectives of this survey, turning our attention to the use of workload characterisation for performance and energy issues. In this subsection, our goal is to find out how many research papers were produced in this sense. To achieve this goal, we analysed the distribution of the amount of publications according to the classification proposed. But, prior to presenting the results from this analysis, it is important to describe the 14 classes proposed in this survey. They are as follows:

- *Benchmark proposal and analysis*: papers that propose or analyse benchmarks.
- *Domain workload characterisation*: it is the class that encompasses the papers that presents the workload characterisation of an application, an architecture or an specific domain.
- *Energy consumption model*: corresponds to the papers that propose models of energy consumption.
- *Energy efficiency*: papers that propose mechanisms for more efficient energy consumption in a computing environment.
- *Infrastructure cost reduction*: papers aimed at reducing infrastructure costs, mainly related to sizing and cost of equipments.
- *Management tools*: corresponds to proposals or analysis of management and performance measurement tools for computing environments.
- *Methodology for workload characterisation*: proposals of methodologies and techniques of workload characterisation.
- *Modelling tools*: papers that propose or analyse system modelling tools.
- *Network characterisation*: consists of papers that present characterisations of computer networks.
- *Performance improvement*: this class comprises papers whose objective is to increase performance in computing environments.
- *Performance Model*: papers that propose models to characterise performance of systems.
- *Scheduling approach*: corresponds to papers that propose or analyse scheduling approaches.
- *Simulation tools*: proposals or analysis of workload simulation tools.
- *Thermal management*: consists of papers that propose or analyse temperature management approaches in computing environments.

Figure 5 shows the distribution of the amount of publications according to the classification proposed. We observed in the chart a considerable amount of publications addressing performance issues, represented by the classes *performance improvement and performance model*. We observed in papers such as those of Tang et al. (2011), Saez et al. (2011), Wu and Lui (2011), and Camara et al. (2010) that significant performance

improvement was achieved by the understanding about the characteristics of the applications, environments and workloads. Works such as those of Bodík et al. (2010), Casale et al. (2010), Wu et al. (2010) and Peña-Ortiz et al. (2009) formalised this environment and workload knowledge through performance models, allowing the evaluation and prediction of performance under variable conditions.

Figure 5 Distribution of publications according to our classification proposal



Another interesting aspect to be noticed in Figure 5 is the amount of papers presenting studies on workload characterisation of specific systems and domains, represented by class domain workload characterisation. The results presented by works such as those of Iosup and Epema (2011), and Thies and Amarasinghe (2010) provide researchers with workload characterisations of specific domains, which could help them better understand their own environments and let them concentrate directly on the optimisation solutions.

One aspect of particular interest in this survey is the use of workload characterisation to address energy issues. Figure 5 shows a significant amount of publications proposing solutions for energy efficiency improvement, represented by the class *energy efficiency*. Several papers, such as those of Huang and Feng (2009), Meisner et al. (2011) and Dhiman and Rosing (2009) tackle this problem by dynamically adjusting the frequency and voltage (DVFS), or activating the low-power mode, of components such as processors, memories and disks, according to the workload characteristics. Other works such as those of Dhiman et al. (2010a) and Jiang et al. (2010) optimise the allocation of virtual machines (VM) in a cluster in order to reduce the energy consumption, which is especially relevant to cloud environments. There are even works such as those of Rodero et al. (2010) and Chetsa et al. (2013) that present resource management proposals based on workload characteristics for HPC environments. Proposals of energy consumption models were less observed in the analysed papers. These models, such as those proposed in Dhiman et al. (2010b) and Choi et al. (2010) grant significantly accurate predictions of energy consumption, allowing better optimisation decisions.

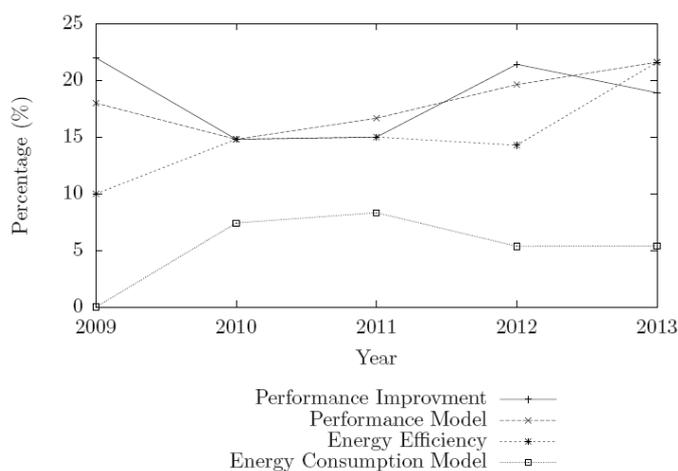
We attribute the fact that a smaller amount of the publications addressed energy issues, in comparison to the amount of publications addressing performance and other issues, to a more recently interest in this subject. This assumption is confirmed in the answer to the next research question.

5.6 Are there any trends in the contributions addressing performance and energy issues?

The answer to this question provides understanding about the behaviour of the scientific community's interest in respect to the use of workload characterisation to address performance and energy issues. It also explains why the amount of contributions to energy issues is not yet so expressive as the amount of contributions to performance issues.

We evaluated the evolution of representativeness of publications classified as *performance improvement*, *performance model*, *energy efficiency*, and *energy consumption model*, between 2009 and 2013. By representativeness we mean the proportion of the amount of publications in those specific classes compared to the total amount of publications collected. Summarised results are presented in Figure 6. As we suspected, only 10.00% of the reviewed research papers published in 2009 addressed energy issues, while this number reached 27.03% in 2013. In particular, we observed an increasing interest in energy efficiency contributions, with an average growing rate of 24.00% per year.

Figure 6 Evolution of representativeness of publications classified as *performance improvement*, *performance model*, *energy efficiency*, and *energy consumption model* between 2009 and 2013

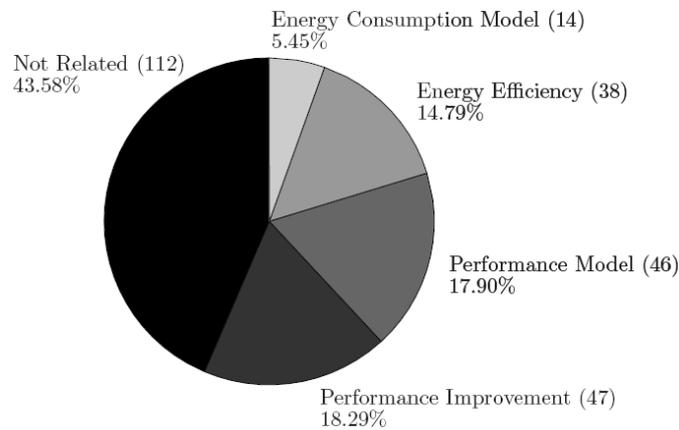


As can be seen in Figure 6, publications addressing performance issues, while having significant representativeness during the whole analysed period (i.e., with average of 36.58% and standard deviation of 5.48), do not exhibit a so expressive growing rate (i.e., in average, 2.34% per year). This can be viewed as stability in the research topic.

Closing the presentation of these bibliometrics indices, Figure 7 shows the representativeness of publications classified as *performance improvement*, *performance model*, *energy efficiency*, and *energy consumption model*, among the whole set of reviewed publications. The proportion labelled as 'not related' represents papers whose contributions were not directly related with any of the four previous classes. It can be observed that more than a half (i.e., 56.43%) of the reviewed publications presented contributions to performance and energy problems. Confirming what we observed in

Figure 6, there is a predominance of publications concerning performance issues (36.2%). From these results, it can be considered that the efforts of the scientific community are largely divided between the understanding of the characteristics and methodologies for workload characterisation and the use of this knowledge to improve performance and energy efficiency. It is important to mention that understanding about workload characteristics is the first step to proposals of performance and energy efficiency optimisations. Thus, the 43.6% of publications grouped as *not related* can be considered as important knowledge to support researchers to improve their HPC, cloud computing and big data environments.

Figure 7 Representativeness of publications classified as *performance improvement*, *performance model*, *energy efficiency*, and *energy consumption model*



Note: The proportion labelled as *not related* refers to publications whose contributions were not directly related to any of the previous classes.

6 Discussion

In this section, we briefly present some contributions to HPC, cloud and big data environments of the papers reviewed in this survey.

6.1 High performance computing

Rodero et al. (2010) studied the potential of application-centric aggressive power management of data centre's resource for HPC workloads. They proposed a *predictive and aggressive power management (PAPM)* algorithm, which transitions the subsystem to the appropriate power mode based on a-priori knowledge of the application profile. Experimental results using benchmarks with different behaviours in terms of subsystems utilisations showed average energy savings ranging from 1.67% to 10.95% per node, depending on the characteristics of resources consumption of each application.

Seelam et al. (2010) present the CPU, memory, and communication performance characteristics of three large-scale DARPA high productivity computer systems (HPCS) project benchmarks: hybrid coordinate ocean model (HYCOM), parallel ocean

programme (POP), and Lattice Boltzmann magneto-hydrodynamics code (LBMHD). The analysis used cycles per instruction (CPI) model for CPU/memory performance and MPI traces for multi-process communication performance. As a result, the research shows that: the performance of HYCOM and POP benchmarks are mostly impacted by the poorly pipelined versions of long floating point operations; load and store activities significantly contribute to CPI; and, in terms of communications, LBMHD is the most scalable between the three analysed benchmarks because its communication activities occur mostly among the nearest neighbouring tasks.

An analysis of the jobs reported by TeraGrid, a federation of 11 resource provider sites operating more than 20 HPC systems, for 2008 is presented by Hart (2011). Considering the workload from three perspectives (i.e., traditional HPC workload characteristics, grid-oriented workload characteristics, and user- and group-oriented characteristics), the study highlights the importance of distinguishing between the analyses of job patterns and work patterns. As a result, it was observed that small sets of users dominate the workload, that the impact of allocations and operational policies can be seen in the workload patterns, and that aggregate analyses across even loosely coupled grid federations reflect patterns seen in more tightly coupled grid and in single HPC systems.

An allocation mechanism which considers machine affinity is proposed by Snavely et al. (2011). They propose that the CPU hours of the Department of Defense (DoD) Supercomputing Resource Centers (DSRCs) should be granted based not only on technical merit and service needs, but considering the machine affinity with the application. The approach basically characterises the high performance computing modernisation programme (HPCMP) systems and evaluates, at each request, the difference in performance between the first two machine choices. With this approach, results show that the overall throughput across all machines is boosted by 10% and that the effect is like adding a new \$10 M supercomputer to the environment.

6.2 Cloud computing

The work of Guabtni et al. (2013) presents an approach to efficient range query processing across relational databases in cloud computing environments that combines proprietary cloud-based load balancing techniques and a proposed density-based partitioning. Through a pre-processing phase, the data density of the database is calculated. Then, the range query requests are not divided into equally sized parts, but based on the density of the data. Experiments using a news database and a cloud environment with the number of database replicas (i.e., VM) ranging from 2 to 16 show that the load is better distributed among the VMs, which can lead to a better response time.

Chen et al. (2011) show their initial results in profiling the power behaviour of VM in cloud computing environments devoted to HPC workloads. A linear power model was built to represent the power behaviour of each single work node, considering their individual components (i.e., CPU, memory, HDD). Using the infrastructure of the Distributed ASCI Supercomputer 4 (DAS-4), their experiments showed that: the idle power consumption was not affected by CPU frequency scaling, but a steep rise was observed when CPU load reaches 5%; the contribution of memory power consumption accounts for only 5.5% in idle state and 3.5% on maximum load; despite the fact that

HDD power consumption is difficult to quantify, it is worthy to keep it in the proposed power model.

Di et al. (2012) present in their paper a characterisation and a comparison of work load (i.e., jobs and tasks) and host load (i.e., machine) between a real-world cloud computing environment (i.e., a Google data centre) and some Grid/HPC systems (e.g., AuverGrid, Sharcnet, Argonne National Laboratory). From the analysis of metrics such as job length, job submission frequency, and resource utilisation, it was observed a finer resource allocation, regarding CPU and memory, in Google data centre; and that the frequency of jobs submission is greater than in the Grid/HPC environments.

Ferdman et al. (2012) analyse the micro-architectural behaviour of processors in a cloud computing data centre by means of scale-out workloads. Through the use of a benchmark (i.e., CloudSuite), they observed that predominant processor micro-architecture of today's data centres is inherently inefficient for running scale-out workloads, resulting in low computational density and energy efficiency. To solve this problem, future processors micro architectures should optimise the instruction-fetch path, reduce the core aggressiveness and last-level cache capacity, and scale back the over-provisioned on- and off-chip bandwidth.

6.3 *Big data*

A set of recommendations for workload modelling for data-management in science grids is presented by Iamnitchi et al. (2009). Using real traces from a high-energy physics collaboration (i.e., DZero Experiment), they revealed the existence of a grouping of files accessed together, coined by them as *filecules*. Their experiments showed that the consideration of filecules in least recent used (LRU) caching algorithm presents significant performance improvements over more sophisticated caching techniques for the evaluated workload (i.e., from 4% to 106% of improvement in byte hit rates).

Meisner et al. (2011) examined power management for online data-intensive (OLDI) data centre workloads. Their goal was to investigate how to make OLDI systems more energy-proportional, that is, how to make power consumption more proportional to the servers load. They evaluated the applicability of active and idle low-power modes on servers primary components (i.e., processor, memory, and disk) through a characterisation of a web search workload at a cluster-wide scale. As results they found that: CPU low power modes provide the best single power-performance mechanism, but can not achieve energy-proportionality; memory system presents a great opportunity for saving energy; and coordinated full-system active low-power modes hold the greatest promise to achieve energy-proportionality.

7 **Conclusions**

This survey presents a differential contribution from prior publications in at least three aspects and successfully achieved its goal providing a better understanding of the use of workload characterisation to address issues regarding performance and energy consumption in HPC, cloud and big data environments.

After an extensive literature review and a classification proposal based on researchers' contributions, we observed that 56.4% of the papers published between 2009 and 2013 proposed solutions to performance and power consumption problems using workload characteristics. In addition, our analysis shows an average increase of 7.86% per year on the representativeness of this kind of research work. This growing interest represents the understanding by the scientific community of the importance of better knowing workload characteristics for finer optimisations. Considering the emergence of newer computer architectures and paradigms, we believe that this interest tends to grow further.

A major contribution of this paper is the knowledge about the use of workload characterisation for energy issues. Indeed, to the best of our knowledge, this is the first time that such a research work is conducted. We observed in our analysis that the use of workload characterisation for energy issues is something relatively new. Results show that representativeness of research papers proposing energy efficiency improvement and energy modelling increased from 10% in 2009 to 27% in 2013. In terms of energy efficiency improvement, the representativeness average growth rate is of 24% per year. These results can be interpreted as an understanding by the scientific community of the importance of reducing energy consumption in the HPC, cloud and big data environments, and that a comprehensive understanding of system behaviour is necessary for better results.

Performance improvement and performance modelling continues to be the main target of workload characterisation. About 36.19% of the publications reviewed presented contributions in this sense. However, the interest in this subject is more stable, presenting a slight growth of 2.34% per year, in average. In addition, we observed that no general solutions were proposed. Performance improvement proposals were characterised by well defined environments and specific conditions. Some of them even demonstrated that their proposals perform better under some conditions and poorly with little changes in the scenario.

An interesting discovery of this survey was the significant amount of contributions regarding the use of workload characterisation in cloud environments. The keyword 'cloud computing' appears in 24 publications, which corresponds to 9.34% of the reviewed papers. We believe that the recent popularity of this kind of environment is one of the reasons for this amount of papers. Another reason is that the clear understanding of the system behaviour provided by workload characterisation allows better allocation and management decisions, which is an important feature on environments with tight QoS constraints.

One of the main difficulties encountered in carrying out this survey was to define the criteria for papers' selection. The fact that the subject was not so normalised imposed an additional effort to prevent relevant papers from being discarded. However, the number of publications reviewed (i.e., 257) can be considered significant for the purposes of the present study.

For future works, we intend to further explore techniques and methodologies of workload characterisation to propose solutions for performance and energy efficiency improvement in HPC, cloud and big data environments with considerable demands of data movement.

References

- Benedict, S. (2012) ‘Performance issues and performance analysis tools for HPC cloud applications: a survey’, *Computing*, Vol. 95, No. 2, pp.89–108.
- Benmoussa, Y., Boukhobza, J., Senn, E. and Benazzouz, D. (2013) ‘GPP vs DSP: a performance/energy characterization and evaluation of video decoding’, in *International Symposium on Modelling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS)*, IEEE, pp.273–282.
- Birke, R., Chen, L.Y., Gribaudo, M. and Piazzolla, P. (2013) ‘Characterization analysis of resource utilization distribution’, in *International Symposium on Modelling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS)*, IEEE, pp.370–374.
- Bodík, P., Fox, A., Franklin, M., Jordan, M. and Patterson, D. (2010) ‘Characterizing, modeling, and generating workload spikes for stateful services’, in *ACM Symposium on Cloud Computing (SoCC)*, pp.241–252.
- Buyya, R., Yeo, C., Venugopal, S., Broberg, J. and Brandic, I. (2009) ‘Cloud computing and emerging it platforms: vision, hype, and reality for delivering computing as the 5th utility’, *Future Generation Computer Systems*, Vol. 25, No. 6, pp.599–616.
- Calzarossa, M. and Serazzi, G. (1993) ‘Workload characterization: a survey’, *Proceedings of the IEEE*, Vol. 81, No. 8, pp.1136–1150.
- Calzarossa, M., Massari, L. and Tessera, D. (2000) ‘Workload characterization issues and methodologies’, *Performance Evaluation: Origins and Directions (LNCS)*, Vol. 1769, pp.459–482.
- Camara, J., Moretó, M., Vallejo, E., Bevide, R., Miguel-Alonso, J., Martínez, C. and Navaridas, J. (2010) ‘Twisted torus topologies for enhanced interconnection networks’, *IEEE Transactions on Parallel and Distributed Systems*, Vol. 21, No. 12, pp.1765–1778.
- Casale, G., Zhang, E. and Smimi, E. (2010) ‘Trace data characterization and fitting for markov modeling’, *Performance Evaluation*, Vol. 67, No. 2, pp.61–79.
- CERN (2013) *The Large Hadron Collider* [online] <http://home.web.cern.ch/topics/large-hadroncollider> (accessed November 2013).
- Chen, Q., Grosso, P., Van Der Veldt, K., De Laat, C., Hofman, R. and Bal, H. (2011) ‘Profiling energy consumption of VMs for green cloud computing’, in *IEEE International Conference on Dependable, Autonomic and Secure Computing (DASC)*, pp.768–775.
- Chetsa, G., Lefevre, L. and Stolf, P. (2013) ‘A three step blind approach for improving HPC systems’ energy performance’, *Energy Efficiency in Large Scale Distributed Systems (LNCS)*, Vol. 8046, pp.168–181.
- Choi, J., Govindan, S., Jeong, J., Uргаonkar, B. and Sivasubramaniam, A. (2010) ‘Power consumption prediction and power-aware packing in consolidated environments’, *IEEE Transactions on Computers*, Vol. 59, No. 12, pp.1640–1654.
- Dantas, M. (2005) *Computação Distribuída de Alto Desempenho: Redes, Clusters e Grids Computacionais*, Axcel Books do Brasil, Rio de Janeiro, RJ, USA.
- Dhiman, G. and Rosing, T. (2009) ‘System-level power management using online learning’, *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol. 28, No. 1, pp.676–689.
- Dhiman, G., Marchetti, G. and Rosing, T. (2010a) ‘VGreen: a system for energy-efficient management of virtual machines’, *ACM Transactions on Design Automation of Electronic Systems*, Vol. 16, No. 1, pp.6:1–6:27.
- Dhiman, G., Mihic, K. and Rosing, T. (2010b) ‘A system for online power prediction in virtualized environments using Gaussian mixture models’, in *Design Automation Conference (DAC)*, pp.807–812.
- Di, S., Kondo, D. and Cirne, W. (2012) ‘Characterization and comparison of cloud versus grid workloads’, in *IEEE International Conference on Cluster Computing (CLUSTER)*, pp.230–238.

- Elnaffar, S. and Marin, P. (2002) *Characterizing Computer Systems' Workloads*, Technical report, Queen's University.
- Elnaffar, S. and Martin, P. (2006) 'Techniques and a framework for characterizing computer systems' workloads', in *Innovations in Information Technology (IIT)*, IEEE.
- Elsevier (2013) *Scopus* [online] <http://www.scopus.com> (accessed November 2013).
- Ferdman, M., Adileh, A., Kocberber, O., Volos, S., Alisafae, M., Jevdjic, D., Kaynak, C., Popescu, A., Ailamaki, A. and Falsafi, B. (2012) 'Clearing the clouds: a study of emerging scale-out workloads on modern hardware', in *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pp.37–47.
- Ferrari, D., Serazzi, G. and Zeigler, A. (1983) *Measurement and Tuning of Computer Systems*, Prentice-Hall, Englewood Cliffs, NJ, USA.
- Foster, I. and Kesselman, C. (2001) 'The anatomy of the grid: enabling scalable virtual organizations', *International Journal of High Performance Computing Applications*, Vol. 15, No. 3, pp.200–222.
- Gartner (2013) *Gartner IT Glossary – Big Data* [online] <http://www.gartner.com/it-glossary/bigdata/> (accessed November 2013).
- Gil, A.C. (2010) *Como Elaborar Projetos de Pesquisa*, Atlas, 5th ed., São Paulo, SP, Brazil.
- Green500 (2013) *Green500 List* [online] <http://www.green500.org/lists/green201311> (accessed November 2013).
- Guabtani, A., Ranjan, R. and Rabhi, F. (2013) 'A workload-driven approach to database query processing in the cloud', *Journal of Supercomputing*, Vol. 63, No. 3, pp.722–736.
- Hart, D. (2011) 'Measuring TeraGrid: workload characterization for a high-performance computing federation', *International Journal of High Performance Computing Applications*, Vol. 25, No. 4, pp.451–465.
- Huang, S. and Feng, W. (2009) 'Energy-efficient cluster computing via accurate workload characterization', in *IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGRID)*, pp.68–75.
- Iamnitchi, A., Doraimani, S. and Garzoglio, G. (2009) 'Workload characterization in a high-energy data grid and impact on resource management', *Cluster Computing*, Vol. 12, No. 2, pp.153–173.
- Iosup, A. and Epema, D. (2011) 'Grid computing workloads', *IEEE Internet Computing*, Vol. 15, No. 2, pp.19–26.
- Islam, M.A., Eager, D., Carlsson, N. and Mahanti, A. (2013) 'Revisiting popularity characterization and modeling of user-generated videos', in *International Symposium on Modelling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS)*, IEEE, pp.350–354.
- Jiang, C., Zhang, J., Wan, J., Xu, X., Yin, Y., Yu, R. and Lv, C. (2010) 'Power aware resource allocation in virtualized environments through VM behavior identification', in *IEEE/ACM International Conference on Green Computing and Communications (GreenCom)*, pp.313–318.
- Kitchenham, B. (2004) *Procedures for Performing Systematic Reviews*, Technical report, Keele University and National ICT Australia Ltd.
- Kogge, P., Borkar, S., Campbell, D., Carlson, W., Dally, W., Denneau, M., Franzon, P., Harrod, W., Hiller, J., Keckler, S., Klein, D. and Lucas, R. (2008) *ExaScale Computing Study: Technology Challenges in Achieving Exascale Systems*, Technical report, DARPA-IPTO.
- Mateescu, G., Gentsch, W. and Ribbens, C. (2011) 'Hybrid computing – where HPC meets grid and cloud computing', *Future Generation Computer Systems*, Vol. 27, No. 5, pp.440–453.
- Meisner, D., Sadlerz, C., Barrosoz, L., Weberz, W-D. and Wenischy, T. (2011) 'Power management of online data-intensive services', in *International Symposium on Computer Architecture (ISCA)*, pp.319–330.

- Nazário, D.C., Dantas, M.A.R. and Todesco, J.L. (2012) 'Taxonomia das Publicações Sobre Qualidade de Contexto', *Sustainable Business International Journal*, Vol. 20, No. 20, pp.1–28.
- Obaidat, M. and Boudriga, N. (2010) *Fundamentals of Performance Evaluation of Computer and Telecommunication Systems*, Wiley, New York, NY, USA.
- Peña-Ortiz, R., Sahuquillo, J., Pont, A. and Gil, J. (2009) 'DWEB model: representing Web 2.0 dynamism', *Computer Communications*, Vol. 32, No. 6, pp.1118–1128.
- Pritchard, A. (1969) 'Statistical bibliography of bibliometrics?', *Journal of Documentation*, Vol. 4, No. 25, pp.348–349.
- Rodero, I., Chandra, S., Parashar, M., Muralidhar, R., Seshadri, H. and Poole, S. (2010) 'Investigating the potential of application-centric aggressive power management for HPC workloads', in *International Conference on High Performance Computing (HiPC)*.
- Saez, J., Shelepov, D., Fedorova, A. and Prieto, M. (2011) 'Leveraging workload diversity through OS scheduling to maximize performance on single-ISA heterogeneous multicore systems', *Journal of Parallel and Distributed Computing*, Vol. 71, No. 1, pp.114–131.
- Seelam, S., Chung, I-H., Cong, G., Wen, H-F. and Klepacki, D. (2010) 'Workload performance characterization of DARPA HPCS benchmarks', *Concurrency Computation Practice and Experience*, Vol. 22, No. 4, pp.441–461.
- Snavely, A., Gamst, A., Carrington, L., Tikir, M. and Laurenzano, M. (2011) 'Using machine affinity to increase science throughput (machine affinity characterization of the HPCMP workload)', in *DoD High Performance Computing Modernization Program Users Group Conference (HPCMP UGC)*, pp.472–477.
- Soares, T.S., Dantas, M., de Macedo, D.D. and Bauer, M.A. (2013) 'A data management in a private cloud storage environment utilizing high performance distributed file systems', in *Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE)*, IEEE, pp.158–163.
- Stevens, R. and White, A. (2009) 'Architectures and technology for extreme scale computing', in *ASCR Scientific Grand Challenges Workshop Series*.
- Tang, L., Mars, J., Vachharajani, N., Hundt, R. and Soffa, M. (2011) 'The impact of memory subsystem resource sharing on datacenter applications', in *International Symposium on Computer Architecture (ISCA)*, pp.283–294.
- Thies, W. and Amarasinghe, S. (2010) 'An empirical characterization of stream programs and its implications for language and compiler design', in *Parallel Architectures and Compilation Techniques (PACT)*, pp.365–376.
- Top500 (2013a) *Top500 Architecture Development over Time* [online] <http://www.top500.org/statistics/overtime/> (accessed November 2013).
- Top500 (2013b) *Top500 List* [online] <http://www.top500.org/lists/2013/11/> (accessed November 2013).
- Wu, W. and Lui, J. (2011) 'Exploring the optimal replication strategy in P2P-VoD systems: characterization and evaluation', in *International Conference on Computer Communications (INFOCOM)*, pp.1206–1214.
- Wu, Y., Hwang, K., Yuan, Y. and Zheng, W. (2010) 'Adaptive workload prediction of grid performance in confidence windows', *IEEE Transactions on Parallel and Distributed Systems*, Vol. 21, No. 7, pp.925–938.

The High Performance Computing environment is being hosted at a shared data centre for education and research, offered by digital technologies charity Jisc. The data centre has the capacity, technological capability and flexibility to future-proof and support all of eMedLab's HPC needs, with its ability to accommodate multiple and varied research projects concurrently in a highly collaborative environment. Arif Ali, Technical Director of OCF says: "The private cloud HPC environment offers a flexible solution through which virtual clusters can be deployed for specific workloads. The multi-tenancy features of the Red Hat platform enable different institutions and research groups to securely co-exist on the same hardware, and share data when appropriate." PDF | The relocation of high performance computing systems (HPC) to the cloud poses new challenges for data center architects and IT managers. These challenges are due to heterogeneity injected into data centers by cutting-edge virtualization technologies and hardware accelerators... efficiency in heterogeneous data centers on the cloud, it is best. to re-claim excess VCPUs from CUDA workloads and re-allocate them to other collocated VMs; this frees stranded resources and improves utilization. C. Micro-architectural Variation. DVFS is a prevalent power management technique in data centers to throttle server power consumption at the expense of. Scaling Up Energy Efficiency Across the Data Center Industry: Evaluating Key Drivers and Barriers. Acknowledgments. Project managed by: Pierre Delforge (NRDC) Josh Whitney (Anthesis). Others include public disclosure of data center energy and carbon performance, and the development of templates for green colocation and other multi-tenant data center contracts. | PAGE 6 Data Center Efficiency Assessment. The last segment of the data center market is high-performance computing (HPC), which comprises data centers owned and operated by universities, research facilities, and national laboratories, such as the National Renewable Energy Lab and the Lawrence Berkeley National Lab.