Characterising the Power Consumption of Hadoop Clouds
A Social Media Analysis Case Study

Javier Conejero[1], Omer Rana[2], Peter Burnap[2], Jeffrey Morgan[3], Carmen Carrión[1] and Blanca Caminero[1]

[1]*Department of Computing Systems, University of Castilla-La Mancha, Albacete, Spain*

[2]*School of Computing & Informatics, Cardiff University, Cardiff, U.K.*

[3]*School of Social Sciences, Cardiff University, Cardiff, U.K.*

Keywords: Cloud Computing, Power Consumption, Hadoop, OpenNebula, Social Media Analysis.

Abstract: Energy efficiency is often identified as one of the key reasons for migrating to Cloud environments. It is often stated that a data centre hosting the Cloud environment is likely to achieve greater energy efficiency (at a reduced cost) compared to a local deployment. With increasing energy prices, it is also estimated that a large percentage of operational costs within a Cloud environment can be attributed to energy. In this work, we investigate and measure energy consumption of a number of virtual machines running the Hadoop system, over an OpenNebula Cloud. Our workload is based on sentiment analysis undertaken over Twitter messages. Our objective is to understand the tradeoff between energy efficiency and performance for such a workload. From our results we generalise and speculate on how such an analysis could be used as a basis to establish a Service Level Agreement with a Cloud provider – especially where there is likely to be a high level of variability (both in performance and energy use) over multiple runs of the same application (at different times).

1 INTRODUCTION

Various companies (ranging in size and computing maturity) are adopting Cloud computing technology to perform their business processes, mainly driven by the fact that it reduces the cost of computing infrastructure deployment and management. At the same time, environmental concerns of many large scale computing infrastructure operators – primarily large data centres – have prompted the need for considering more energy efficient operation of computational infrastructure. This coupled with the need to consider new sources of energy, such as solar/wind energy, leads to important challenges in understanding how more energy efficient Cloud computing could be provided to end users. It is also useful to note that the business case for migrating to Cloud computing systems has often centered on the cost savings that would arise due to reduced use of energy at a client site. Currently, energy costs account for a large percentage of operational expenditure for computational infrastructure. It is often stated that due to the economies of scale, the ability to negotiate cheaper energy tariffs and the use of renewable energy sources, data centre operators are able to offer both cost and energy efficient operational systems. With increasing outsourcing of computational capability comes the need to specify Service Level Agreements (SLAs) with infrastructure providers. Such SLAs may also include support for pay-per-use scalability of backend servers, enabling a company to dynamically grow its computational usage based on demand (using an incremental charging model for the excess capacity used). Determining how such SLAs should be specified and subsequently monitored for conformance remains a challenge with many commercial Cloud providers – where repeatable performance is difficult to guarantee in many instances (due to the use of virtualisation and a variable mapping between virtual and physical resources). Increasingly, there is also the demand to include "green" metrics into an SLA, to enable a company using a data centre to display its environmentally friendly credentials to customers. Consequently, there is increasing interest in making Cloud computing environmentally sustainable (Garg and Buyya, 2012) – thereby requiring techniques to improve power efficiency at all levels of the data centre (from resource scheduling of work-

Conejero J., Rana O., Burnap P., Morgan J., Carrion C. and Caminero B..
Characterising the Power Consumption of Hadoop Clouds - A Social Media Analysis Case Study.
DOI: 10.5220/0004373502330243
In *Proceedings of the 3rd International Conference on Cloud Computing and Services Science* (CLOSER-2013), pages 233-243
ISBN: 978-989-8565-52-5

loads to the operation of Computing Room Air Conditioning (CRAC) units). Hence, understanding the behaviour of the various systems that make up a Cloud environment becomes the key in order to design a green datacenter, from the hardware deployed to the usage policies used to exploit each resource. Due to increasing flexibility of Cloud systems and the variety of configuration options now being made available, this becomes a difficult and challenging task. There is also currently significant interest in performing various types of analysis over "big–data" with Cloud-based infrastructures using Hadoop (commercial examples include Radian 6 and Palantir). Hadoop (Lam, 2010; White, 2009) is a framework for data intensive (analysis) applications on large computing clusters by the use of the Map/Reduce paradigm (Dean and Ghemawat, 2008). It has become very popular within social media data analysis projects in order to tackle the scalability of analysis required across large data volumes that could not be performed with traditional paradigms or technologies. Furthermore, MapReduce has also become a useful benchmarking tool (CloudSuite 1.0, 2012) due to its high storage, computing power and network requirements – for comparing the performance of various computing architectures.

Understanding how Hadoop could be efficiently deployed across a Cloud environment remains an important challenge, as Cloud infrastructure parameters and virtual cluster configurations can influence Hadoop performance and have an impact on resource usage. The lack of any behaviour models for achieving such resource management provides an opportunity to consider various optimization techniques for Cloud computing resource usage. The objective of this work is to address this challenge, by determining how data intensive computation could be carried out over a Cloud computing environment and how its subsequent energy footprint could be monitored and optimised. We make use of the Cardiff Online Social Media Observatory (COSMOS) platform (Cardiff Online Social Media Observatory (COSMOS), 2013), which aims to provide mechanisms to capture, analyse and visualize data harvested from online repositories and feeds, in particular interactive and openly accessible social networking sites such as Twitter. COSMOS provides a research framework for social media data analysis, in particular supporting *what-if* investigations from the social sciences community (Pang and Lee, 2008) – which are often difficult to realise in other similar commercial systems. COSMOS enables sentiment analysis (by using the SentiStrength tool (SentiStrength: The sentiment strength detection in short texts, 2012)) to be carried out, that can involve several gigabyte sized data files by using the Hadoop Map/Reduce paradigm.

The main motivation of this article is to understand the impact of high throughput computing (such as Hadoop) on Cloud computing power consumption, validated through an example of "big–data" processing using COSMOS. Our focus in this work is to highlight how the power consumption: (i) can be monitored and understood (in particular, focusing on the variability of consumption across multiple execution runs of the application) – and subsequently exposed to the user; (ii) can be related to the number of virtual machines and the associated workload generated on a physical server. We also investigate how power-usage metrics can be included within an SLA and how variability in power use can impact what can/cannot be included.

The rest of this article is organized as follows. Section 2 presents related work. In Section 3 the Cloud infrastructure, the workload base, the experiments carried out for the study and the instrumentation used to record power usage data is presented. The system behaviour is described in Section 4. Conclusions and future work are subsequently outlined in Section 6.

2 RELATED WORK

Power consumption (synonymously referred to as 'Green IT'), is being considered within a number of areas in Computer Science – for example CPU (their development is heavily conditioned by power dissipation (and consequently consumption) requirements), network and disk power management. Similarly, in the context of Cloud computing, significant efforts have already been documented to support power-consumption aware Green Clouds (Sood and Kumar, 2010).

Companies like APC (by Schneider electric) (UPS Selector Sizing Application, 2012) and VMware (Green IT Calculator, 2012) have designed static power consumption models and provide interfaces (for Uninterrupted Power Supply (UPS) sizing and Virtualization impact, for instance), in order to help users determine the power consumption of a specific computer (depending on its internal components). The approaches identified by these companies can be helpful in the general design of a datacenter's power and UPS requirements; but they are not workload aware nor take into account how such infrastructure is subsequently used. Such models often focus on developing a database of motherboards, their associated components and the recorded power consumption. Given a particular system configuration and motherboard, it is therefore possible to search through such a database to estimate the likely power consumption one is likely to see when using the same (or similar) motherboard. Such a database also contains information about idle and full workload power consumption on such hardware, but does not record any usage for particular types of workloads.

There are also various proposals that suggest the development of a Cloud architecture (Garg and Buyya, 2012; Liu et al., 2009) to provide and use power saving mechanisms while guaranteeing the performance from a user's perspective. In (Garg and Buyya, 2012) a survey of Cloud computing systems is provided, in order to support environmental sustainability and a generic Green Cloud computing architecture introducing the concept of a "Green Broker" is proposed; while (Liu et al., 2009) base their proposal on live virtual machine migration while monitoring the power consumption of resources using dedicated hardware – an expensive option to support in most cases.

Other proposals go further and try to distribute the Cloud workload amongst geographically dispersed datacenters (Ghamkhari and Mohsenian-Rad, 2012) such that they exploit different renewable energy sources depending on the time of day and the observed workloads. This enables a more effective use of green power sources depending on computing demand at particular times in the day.

Hadoop has also been extensively researched focusing on its power efficiency within clusters. In (Leverich and Kozyrakis, 2010) the authors outline the main problems and inefficiencies inherent within the Hadoop Map/Reduce paradigm, while (Goiri et al., 2012) and (Kaushik and Bhandarkar, 2010) propose developing energy saving mechanisms for the file system used in Hadoop (HDFS) and the associated job scheduling in Hadoop, known as GreenHDFS and GreenHadoop respectively. In GreenHDFS an energy-conserving, hybrid, logical multi-zoned variant of HDFS is presented, whose main objective is to reduce energy consumption costs by using low-power, high-energy-saving inactive power modes during idle periods of utilization. GreenHadoop is a framework for datacenters powered by photovoltaic solar arrays of an electrical grids, which schedules the Map/Reduce jobs depending on the solar prediction (that this framework performs) in order to maximize the green energy consumption. Both proposals try to make more effective use of computational resources and relate these to the availability of renewable energy at particular times of day.

In (Shi and Srivastava, 2010) the authors explore the impact of Hadoop based storage clusters, based on the HDFS file system, in thermal terms. It proposes a thermal and power-aware task scheduler for Hadoop that is focused on the minimisation of the total power consumption in the air conditioning (A/C) system, by balancing the load between cluster nodes in order to keep them under a definite thermal threshold. Identifying a suitable operating threshold remains a challenge in many such systems – with an ambient temperature of between 25 to 30 degrees celsius suggested by many authors.

There is also current work focusing on specifying SLA using power consumption metrics. For example (Laszewski and Wang, 2010) identifies several parameters that could be used within an SLA, such as: amount of CO_2 correlated with environmental measurements that are easier to measure and understand for a user. This work also introduces a framework where such metrics could be integrated to support decision making and resource management. Finally, there are additional efforts that aim to monitor power consumption such as PowerTop – although it is not always possible to effectively measure these (such as basic I/O operations) and associate a value with such metrics.

A key metric used in data centres to measure the effectiveness of power usage and efficiency is Power Usage Effectiveness (PUE), developed by the Green Grid Association (Green Grid Association, 2013) (a multi-industry association focusing on power efficiency of data centres). It is used as a ratio of the amount of power entering a data center divided by the power used to run the computational infrastructure within it – with an ideal value being 1.0. As such, it is much broader in scope and takes account all the various infrastructure available within a data centre (such as building, computing room air conditioning systems, etc). It is useful to note that most data centres use almost the same amount of energy to support the "non-computing" capabilities they provide (such as cooling and air conditioning) as the energy used to run their servers and networks. Our focus in this work is much more finer grained than calculating the PUE – as we attempt to determine how power consumption can be associated with a particular application workload across a server. Our attempt is therefore to characterise the impact on power usage of a particular type of VM configuration and application workload. The outcome of this work can be used to subsequently provide different PUE analysis given a particular workload.

3 METHODOLOGY

The objective of our work is to characterise the performance-power tradeoff when deploying Hadoop over an IaaS Cloud environment. We describe the infrastructure over which our validation has been carried out – outlining the key challenges faced when attempting to measure power consumption. We also elaborate on the characteristics of the workload and the monitoring instrument we used in our experiments.

3.1 Cloud Infrastructure

The Cloud infrastructure used in this work (Figure 1) is composed of one cluster compute node, consisting of a Viglen ix4600 with 2 Xeon e5620 CPUs (4 Cores + with support for hyperthreading in each) (Intel Xeon Processor e5 Family, 2012), 24GB of main memory, and 4TB of storage). The Operating System is CentOS 6.2 Linux (CentOS: The Community ENTerprise Operating System, 2012).For the management and coordination of the Cloud environment, OpenNebula (OpenNebula: The Open Source Solution for Data Center Virtualization, 2012) software is deployed. It is a mature open source project focused on the development of an open, flexible, extensible and comprehensive management layer for building and managing Cloud infrastructures. OpenNebula was developed within the European Reservoir project and has since been extended to support a number of different application types and contexts. It provides Infrastructure as a Service (IaaS) by managing different hypervisors (such as KVM, XEN, etc.). We make use of KVM (Kernel Based Virtual Machine (KVM), 2012) as the hypervisor – which is a full virtualization open source hypervisor (with support for hardware virtualization extensions) for Linux widely used in Cloud computing environments and it is supported by RedHat. It provides support for virtualization by the use of a `/dev/kvm/` interface.

On this Cloud environment we deploy a social media data analysis application using Hadoop, described in section 3.2. The decision about using such a private and controlled Cloud environment was to enable us to more accurate measure power consumption with a variation in workload.

3.2 Social Media Analysis Workload

Social media can involve a variety of different types of content – such as video (YouTube), audio (Spotify), images (Facebook, Flickr) to text (Twitter, Facebook). The type of analysis often undertaken on such

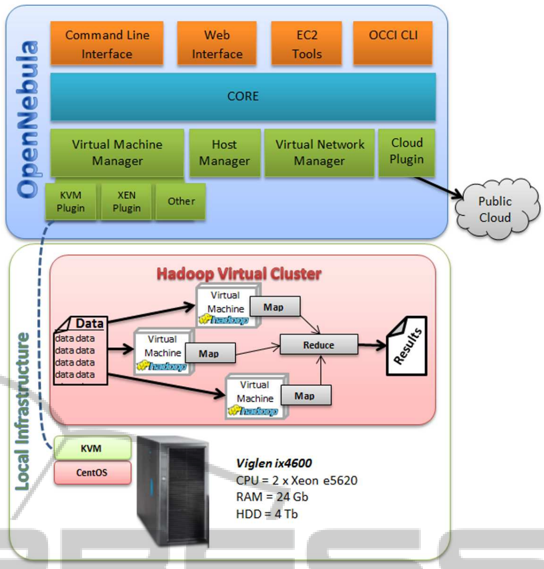

Figure 1: Cloud computing infrastructure.

content depends on the particular demands of the user community involved. In this work, we make use of text analysis using the Apache Hadoop system (Lam, 2010; White, 2009), with a user community consisting mainly of researchers in social sciences. Hadoop implements the Map/Reduce paradigm (Dean and Ghemawat, 2008), where the input data is divided into blocks and distributed across multiple computational resources during the Map stage, processed, and the results combined during the Reduce stage. Hadoop also provides data transfer transparently, limited fault tolerance mechanisms (achieved through replication) and a distributed file system to store data across the compute nodes (HDFS). Hadoop requires a cluster environment in order to perform the Map/Reduce process (Figure 1). This cluster consists of a master node and *1...n* worker nodes. In order to provide Hadoop as a service on the Cloud it is necessary to deploy multiple VMs and create a virtual cluster for Hadoop following the master/worker structure. Virtualization enables us to customize the size and characteristics of each virtual machine. Furthermore, it provides elasticity, portability and the ability to dynamically replace the underlying hardware if needed. As a drawback, a computational overhead is introduced.

The Cardiff Social Media Observatory (COSMOS) aims to support social scientists in analysing socially significant data (e.g. tweets, blogs and news stories). The volume of data produced on a daily basis requires significant computational resources to analyse. For example, COSMOS collects around 3.5 million tweets a day. To perform a longitudinal analysis of say public opinion and sentiment, around a socially significant event (e.g. a political campaign, change of legislation, world sporting event etc.) could require analysis of several weeks' worth of data. An example study may be public opinion surrounding the London 2012 Olympics, where a study of opinion for two weeks before, during and after would require the analysis of 21 million tweets. On a single desktop computer this could take approximately 20 minutes. This is perfectly acceptable as a batch-processed computational exercise, but the reality is that social media analysis may require several "tweaks" to the study parameters, and therefore requires a more interactive way of analysing data. For example, age, gender, location and topic of study within the event may change, as hypotheses are formed and tested. Therefore, the computational analysis must be able to complete much faster to give a more acceptable wait-time for the researcher. The researcher should be able to invoke the computational resources to support large-scale data analysis on-demand and resources need to be dynamically allocated depending on the size of the job. Hadoop has been used within the COSMOS system in order to scale the underlying analysis. This process undertakes analysis on a large archive of previously recorded tweets in order to determine the sentiment of each one (by using SentiStrength tool (SentiStrength: The sentiment strength detection in short texts, 2012)), following the Map/Reduce paradigm. It is a very heavy process due to the fact that it has been designed to perform this analysis through several gigabytes of data. This workload was made use of in this work in order to provide a realistic benchmark application to stress test the underlying Cloud infrastructure defined in section 3.1 by using different virtual cluster configurations.

For the experimentation carried out in this work, a test tweet archive is used consisting of up to 15 million tweets to extract their sentiment. COSMOS currently harvests and archives the 'spritzer stream using the Twitter Streaming API, and makes it available to researchers for inspection and analysis. Even at 1%, the API provides COSMOS with approximately 3.5 million messages (or tweets) per day. The tweet files are archived using a specialised hierarchical filing system, which stores tweets based on the day in which the collection was made. Hadoop also replicated these files multiple times in order to improve fault tolerance – this is achieved automatically.

3.3 Instrumentation and Monitoring

We focus on the power consumption of the whole compute node, due to the fact that high throughput computing (within distributed environments) exploits almost all resources (CPU, Main memory, Storage and Network) in an aggressive way. In order to get the power consumption from a compute node, an external monitoring device is needed, since this metric cannot be obtained from local monitoring software. Although various attempts have been made to approximate this value from system configration data (such as type of CPU, disk, motherboard, operating system, etc) – generally by interpolating between the system configration and previous recorded data – such attempts remain of limited benefit and accuracy. There are several commercial products to direcly measure the energy consumption of as service – the most widely used are Kill-A-Watt[1] and WattsUp[2].

We make use of WattsUp PRO in this work to monitor and log all the information related to power consumption. This meter aims to provide an independently managed and accessed power data collecting mechanism, and must be positioned between the computing node power supply and the mains power plug (Figure 2). Therefore, a second computer is needed in order to get the logging information stored in the WattsUp PRO non-volatile memory. The reason for choosing WattsUp Pro rather than Kill-A-Watt is due to the fact that the latter does not provide the storage feature (needed for long term experimentation). The information obtained by the use of WattsUp PRO reflects the power consumption behaviour of the compute node during the monitoring interval – and enables a user to either collect data at predefined time intervals, or only record data when particular events (i.e. the power consumption exceeds a pre-defined threshold) occur. The monitoring frequency has been set to one sample per second for all experiments described in this work.

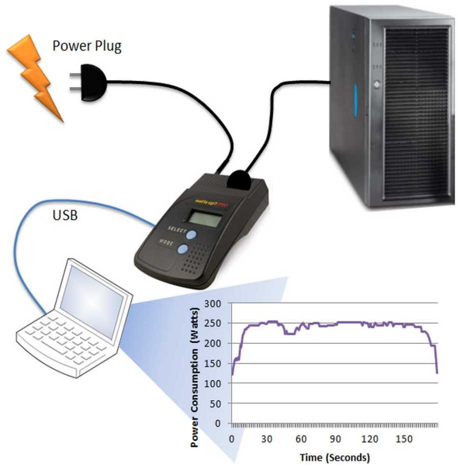

Figure 2: Power consumption monitoring.

[1]http://www.p3international.com/products/special/p4400/p4400-ce.html

[2]https://www.wattsupmeters.com/secure/index.php

4 SYSTEM BEHAVIOUR MONITORING

In order to understand the power consumption of the Cloud system, we devise a number of experiments – results of which are described in subsequent sections. Our approach consists of three main parts:

1. *Basic Power Consumption*: the power consumption of the running cluster is monitored. Monitoring in this part involves understanding the power consumption when turning on/off the system, and the power consumption of keeping the cluster up and running but without any workload (i.e. when the cluster is in the idle state).
2. *Power Consumption Range*: this part is focused on measuring the maximum and minimum power consumption, that is, the power consumption range of the cluster.
3. *Virtualization Power Consumption*: this stage involves monitoring the power consumption with different virtualized workloads executed on the cluster.

4.1 Basic Power Consumption

This stage involves analysing the power demand when turning on/off the system described in Section 3.1 without any workload. The power consumption profile, illustrated in Figure 3) behaves as expected: particularly high when the server is switched on and off and stable once it has booted (and reached stable state). There is a peak in power consumption at start up time, which stabilizes after the operating system (OS) finishes loading all services (at ∼105 Watts). Stopping the server (470-510 seconds in the figure) requires an increased use of power to stop all the OS services and to perform a controlled switch off. Furthermore, when the server is stopped, a power consumption of ∼10 Watts is observed which is due to the standby state of the server. The standby state power consumption seems to be very low for one compute node, but this fact needs to be taken into account within big infrastructures, as it increases linearly with the number of computing nodes. The excess power required to turn on/off a server needs to be balanced against the standby state power usage. Hence, if a server is likely to remain inactive over long time frames, it is possibly better to incur the extra power usage hit when it is turned on/off – alternatively, if the server is likely to be required for ad hoc, bursty workloads, keeping it in standby mode is likely to be more efficient.

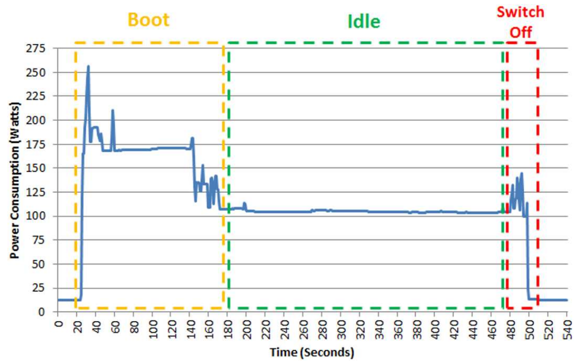

Figure 3: Simulation results.

The server needs about 140 seconds to boot and be ready to enable virtual machines to be deployment on it. But for stopping the server, the time needed is less than 15 seconds.

4.2 Power Consumption Range

In this stage we attempt to measure the power consumption during an idle period and also measure the subsequent increase when a particular workload is executed on the server. The idle power consumption value determines the minimum power required by the server to be up and running, ready for hosting any virtual machines or any equivalent workloads. Our approach therefore considers two types of workloads: (i) virtual machines deployed on the server; (ii) data analysis algorithms executed over the virtual machine. Before continuing with the analysis of the infrastructure under virtualized Hadoop workloads, we are also interested in measuring the maximum power consumption of the server. In order to achieve this, it is necessary to chose a workload that fully stresses all the physical hardware (CPU, memory and disk) available on the server. The MD5 Message-Digest Algorithm (Rivest, 1992) is a widely used cryptographic hash function that produces a 128-bit hash value. It has been utilized in a wide variety of security applications, and is also commonly used to check data integrity. MD5 is a process which exploits the CPU, main memory and disk simultaneously and can be scaled to run on large data files (in size and number). As the server has a multicore subsystem, there is a need of launching multiple MD5 threads in order to stress it completely.

The results obtained from the execution of multiple MD5 threads (Figures 4,5) show that the CPU usage increases as the number of threads also increases, reaching the top CPU usage with 16 threads and keeps this utlisation at 100% for higher number of threads (Figure 4). It is also interesting to observe the variability in CPU usage when executing less than 16

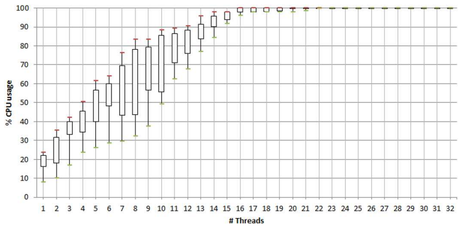

Figure 4: MD5 CPU usage per thread number.

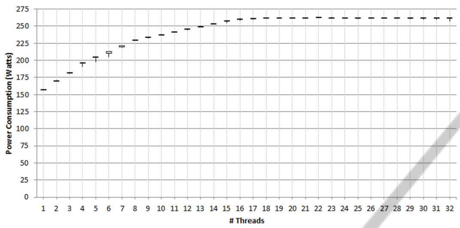

Figure 5: Power consumption under MD5 process threads.

threads, and how the CPU usage variability experiences a reduction from 16 threads and converges to 100%. The power consumption behaviour can be seen to correlated directly with CPU usage but with two major differences (Figure 5). The first one is the slope of the illustrated curve for power consumption, which reduces in value as the number of threads increases, while the second is the significantly smaller variability experienced in the power consumption (of <2%) compared with the variability in CPU usage with less than 16 threads. Using this experiment, we find that the maximum power consumption seen with the server is 268 Watts.

Once we have measured the maximum and minimum power consumption limits, the possible behaviour of a pre-determined number of threads (in terms of CPU usage and power consumption), and the cost of switching on and off the server, we now proceed to measure power consumption for virtualised workloads.

4.3 Virtualization and Power Consumption

The next stage focuses on analysing the behaviour of the system under a realistic Hadoop workload, performed across different virtual clusters. In this scenario, the number of worker nodes and their characteristics are modified, whilst keeping constant the number of resources allocated in all of the virtual clusters. The evaluation is performed with 1 Hadoop server virtual machine (VM) and 4 different VMs for Hadoop worker configurations, as described in Table 1.

These virtual cluster configurations are designed

Table 1: Hadoop virtual cluster VM configurations.

VM Conf.	# VMs	CPU	RAM (Gb)	HDD (Gb)
Server	1	20%	6	100
1 Worker	1	70%	14	200
2 Workers	2	35%	7	100
4 Workers	4	17,5%	3,5	50
8 Workers	8	8,75%	1,75	25

Figure 6: Hadoop Virtual Cluster idle power consumption.

to maximize the usage of the Cloud infrastructure and reserve 10% of the CPU power and 4GB of RAM for the Operating System and Cloud management tools within the infrastructure. Each configuration is evaluated independently from the others, and involves stopping the workers that are not going to be used and deploy the ones that will be used. Virtual clusters are composed of multiple virtual machines, each one which runs the Ubuntu 10.04 operating system. Hence, the greater the number of VMs deployed on the same physical node, the greater the size of the workload introduced, even only with the OS processes of each VM. But when running more processes, like Hadoop, the workload increases, and so does the impact on the node directly related to the number of deployed VMs.

Any deployed virtual cluster adds an overhead on the workload of the server, and as seen with the MD5 scenario, there is a relationship between workload and power consumption. With a bigger Hadoop Virtual Cluster deployed, the power consumption is increased when the Cluster is idle (Figure 6). Hence the size of the Hadoop Virtual Cluster also has an impact on power consumption, even when not performing any Hadoop application.

We subsequently modify the workloads we deploy within the Hadoop virtual cluster. The results obtained from the execution of the sentiment analysis over twitter data application in the Cloud (sequential and Hadoop versions) are shown in Figure 7 – the figure shows the maximum and minimum power consumption (vertical lines) and the 90% observed val-

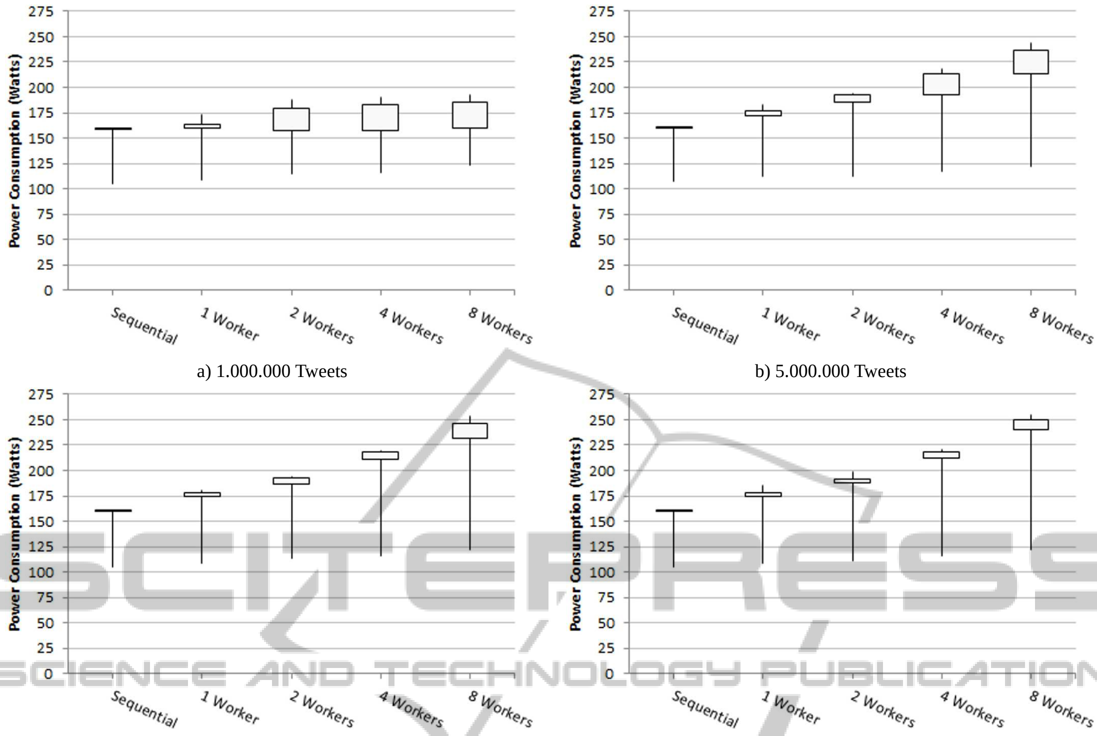

a) 1.000.000 Tweets

b) 5.000.000 Tweets

c) 10.000.000 Tweets

d) 15.000.000 Tweets

Figure 7: Sentiment Analysis for Hadoop Power Consumption.

ues (box). It can be observed that with 8 VMs, the maximum power consumption is achieved. If bigger virtual clusters are deployed, there is a corresponding increase in power consumption. It should also be noted that the highest power consumption level, as measured when applying 16 MD5 threads, can be achieved but not exceeded. The sequential version experiences less power consumption than the Hadoop version when running on any virtual cluster, but on the contrary, its performance is unacceptable (main reason to develop the Hadoop version). Hence the objective of a user deploying an analysis application over such an infrastructure should decide the particular power consumption profile they have in mind, and therefore chose the number of VMs based on this profile. In this instance, we can also map power consumption into a cost, by associating a unit cost for each KWh consumed when running the VM workload.

The size of the problem, and consequently the length of the executions, reduces the variability experienced in power consumption (evolution from Figure 7(a) to Figure 7(d)). This makes easy to forecast the power consumption expected during the execution of the Sentiment Analysis application over Hadoop.

More specifically, the variability observed (Table 2) behaves in two different ways depending on the increase in the number of tweets processed (# Tweets) or the number of virtual machines (# Workers). Repeating the experiment for the same number of tweets but increasing the number of workers, it can be seen that the variability increases in all cases. On the other hand, for the same number of workers, the increase in the number of tweets analysed produces a reduction in the variability.

Finally, it can be observed that the range of power consumption values are similar for the same configuration (i.e. 8 workers). However, when comparing across different configurations there is a clear difference; but the power consumption trend is maintained from 5M to 15M tweets (independently of the scenario and amount of tweets).

5 POWER CONSUMPTION CHARACTERISATION

The behaviour of the Cloud environment running Hadoop can be derived from the analysis undertaken in Section 4. Synthesising the results obtained, we can identify a power consumption profile as illustrated in Figure 8. Hence, six power consumption levels (clearly defined and differentiated in the figure) can be observed – starting from the physical machine

Table 2: Sentiment Analysis for Hadoop Power Consumption details (90%).

# Tweets	# Workers	Min Power (W)	Max Power (W)	Average (W)	Variability (W)
1.000.000	1	160.0	163.7	161.85	±1.85
1.000.000	2	157.5	179.8	168.65	±11.15
1.000.000	4	157.3	183.3	170.30	±13.00
1.000.000	8	159.4	185.8	172.60	±13.20
5.000.000	1	172.4	176.6	174.50	±2.10
5.000.000	2	185.2	192.3	188.75	±3.55
5.000.000	4	192.7	213.6	203.15	±10.45
5.000.000	8	212.7	236.9	224.80	±12.10
10.000.000	1	174.4	177.9	176.15	±1.75
10.000.000	2	186.4	192.2	189.30	±2.90
10.000.000	4	211.4	218.3	214.85	±3.45
10.000.000	8	231.6	246.1	238.85	±7.25
15.000.000	1	174.1	178.0	176.05	±1.95
15.000.000	2	187.8	191.3	189.55	±1.75
15.000.000	4	212.5	217.7	215.10	±2.60
15.000.000	8	239.6	250.1	244.85	±5.25

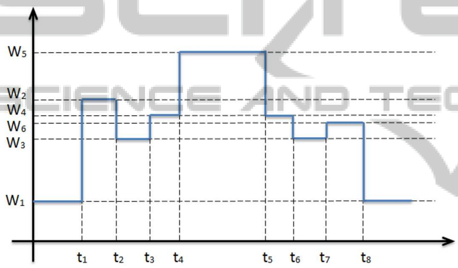

Figure 8: Energy consumption profile.

being initialised at boot up time to the point where it is eventually shut down. These difference in power consumption observed across these different stages of machine use is influenced by the choice of particular infrastructure capabilities – such as the type of hypervisor and virtualization software chosen. As there can be some variation in power consumption in each stage (as observed in Section 4.3), an average value is reported to demonstrate the representative behaviour seen. with more than one virtual machine deployed within the same physical node, as observed in Section 4.3.

Furthermore, the Cloud usage, in terms of deployment policy also has an impact on the power consumption (not only affecting variability in values) since the number of virtual machines deployed conditions the peak power consumption that the system reaches (Figure 8(W_1)).

Each of the six power consumption levels identified is related to a determined machine state (Table 3). The first of these levels is observed when the machine is connected to the electrical grid. It is known as *Standby* mode (W_1) and the machine consumes some power in this state. This consumption, although small compared to the consumption when the machine is running, is very important when the amount of physical nodes is increased and must be taken into account to achieve a particular energy usage threshold. Subsequently, when the physical machine is switched on, an important increase in power consumption is observed (W_2), but it is temporarily de-limited (this value can be further reduced by the use of solid state (hard) disks (SSD) and faster main memory technologies). Hence, there is significant power consumption due to the use of I/O operations on the machine. Once the physical machine has booted it reaches the idle state (W_3), in which the machine is ready to host virtual machines.

When virtual machines are deployed, two different power consumption levels are identified: the first as a consequence of the deployment and idle state of the virtual machines (W_4) (increasing the power consumption due to the fact that there are a higher number of services running at the same time), and the second as a consequence of Hadoop job executions (W_5). As it has been observed in the experiments performed in this work, Hadoop stresses the system and consequently the power consumption increases significantly (even reaching the maximum power consumption value for the physical machine). However, as previously discussed in section 4.3, the actual power usage value depends on the number of virtual machines currently deployed and the total execution time associated with these VMs (Table 2).

The completion of processing jobs by Hadoop makes the system go back to idle state (with idle virtual machines deployed) (W_4). The system in this state is ready to process additional Hadoop jobs. This corresponds to the state where the machine has VMs available to execute additional workload to be submitted by a master node. If the virtual machines are stopped, the system goes back to idle state (W_3). It

Table 3: Energy consumption profiles.

Time (range)	Energy Consumption	Concept
$0 \leq x < t_1$	W_1	Standby
$t_1 \leq x < t_2$	W_2	Switch On (Boot)
$t_2 \leq x < t_3$	W_3	Idle
$t_3 \leq x < t_4$	W_4	VMs running
$t_4 \leq x < t_5$	W_5	Hadoop working
$t_5 \leq x < t_6$	W_4	VMs running
$t_6 \leq x < t_7$	W_3	Idle
$t_7 \leq x < t_8$	W_6	Switch Off (Shut down)
$t_8 \leq x$	W_1	Standby

is then ready to host additional virtual machines (e.g. another virtual cluster). Finally, the physical machine can be stopped and reaches the *Standby* mode (W_1), and as a consequence due to the need to stop associated operating system services an increase in power consumption is observed (W_6) – although this surge is only temporary.

The methodology used in this article can be extrapolated to other computing systems, but the characterisation of energy profiles identified within this section can not be easily generalised for other types of infrastructure. However, we note that the general trend observed is likely to be common when deploying the OpenNebula and Hadoop environments on other physical machines.

The information that this model may be used to select and subsequently optimize a deployment policy. That is, to decide the proper virtual cluster for each user requirement taking into account the power consumption that it is going to require. Thus, it is important to decide when to stop the physical node (depending on time restrictions and user requirements) in order to estimate the additional benefit in performance vs. the corresponding increase in power. Hence, if only a small additional performance benefit can be achieved with a significant power consumption increase (and consequently the associated energy costs), then a user may decide not to optimise performance further to limit power usage. Our characterisation enables such decisions to be supported across different types of workloads.

6 CONCLUSIONS AND FUTURE WORK

The objective of this work has been to measure and characterise power consumption for high throughput workloads (using Hadoop). Such measurement can be used as the basis for developing a workload power consumption model for analysing social media data. As sentiment analysis remains one of the most widely performed operation on social media data streams, our approach can provide a useful basis for understanding how a system should be configured to achieve a particular performance-energy profile.

The main conclusion obtained from this study is that there is a non-linear relationship between the number of virtual machines, the workloads that these VMs execute and the power consumption seen on the physical machine. Identifying how many VMs are needed to achieve a particular throughput at a given power usage profile can be undertaken based on the results reported in his work. Consequently, deploying and using 8 or more VMs on the same physical machine suggests the maximum power consumption possible for the particular Cloud infrastructure we investigated in this work. The infrastructure makes use of OpenNebula, Hadoop and the KVM hypervisor – as all of these systems are widely used in the research community, we believe the outcome of this analysis is usable in a number of similar contexts. Our methodology used to analyse and compare power consumption (using the three stages outlined in section 4) could be adapted for other applications – and other Cloud environments.

Furthermore, we have also observed a variability in power consumption over multiple runs of the same workload. However such variation is generally small, although there are uncontrollable variations (such as, sudden drops or peaks is power usage that cannot be easily explained). This variability reduces as execution times increase. Hence, for short running jobs, using power related metrics in service level agreements can be limiting – even on private clouds. We believe such variation is likely to be significant within public Cloud environments that use a multi-tenancy approach, where workloads and number of VMs can change over time (and from a users perspective are hard to characterise). The approach we advocate in this work can also be used to include metrics such as power usage within such a Service Level Agreement – alongside more traditional performance related metrics. This is particularly important when the client requiring access to a service needs to demonstrate "green" credentials to its customers. The objective is therefore to understand how significant the change in performance is likely to be with an increase in the power consumed, to execute a particular type of workload.

Future work guidelines are: (i) perform a similar type of study over a distributed Cloud Computing infrastructure (with different VM deployment strategies) and extend the model for these environments; (ii) better understand how metrics related to key performance indicators (such as revenue, penalty, etc) can be mapped into operational metrics which include both performance and power – and subsequently how these could be used within a service level agreement; (iii) design and implement policies (by applying the power characterisation described in this work) to handle Cloud computing environments in an optimized way in terms of power saving and/or performance. Although PUE metrics already exist for a data centre, our aim is to develop similar metrics for particular applications.

ACKNOWLEDGEMENTS

This work was supported by the Spanish Government under Grant TIN2012-38341-C04-04 and through a FPI scholarship associated to TIN2009-14475-C04-03 project.

REFERENCES

Cardiff On-line Social Media Observatory (COSMOS) (Last access: January 30, 2013). Web page at http://www.cs.cf.ac.uk/cosmos/.

CentOS: The Community ENTerprise Operating System (Last access: 13th October, 2012). Web page at http://www.centos.org/.

CloudSuite 1.0 (Last access: 16th October, 2012). Web page at http://parsa.epfl.ch/cloudsuite/cloudsuite.html.

Dean, J. and Ghemawat, S. (2008). Mapreduce: simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113.

Garg, S. and Buyya, R. (2012). *Green Cloud Computing and Environmental Sustainability, Harnessing Green IT: Principles and Practices*. Wiley Press, UK.

Ghamkhari, M. and Mohsenian-Rad, H. (2012). Optimal Integration of Renewable Energy Resources in Data Centers with Behind-the-Meter Renewable Generator. In *Proc. of the IEEE International Conference in Communications (ICC'2012)*, Ottawa, Canada.

Goiri, I. n., Le, K., Nguyen, T. D., Guitart, J., Torres, J., and Bianchini, R. (2012). Greenhadoop: leveraging green energy in data-processing frameworks. In *Proceedings of the 7th ACM european conference on Computer Systems*, EuroSys '12, pages 57–70, New York, NY, USA. ACM.

Green Grid Association (Last access: January 30, 2013). Web page at http://www.thegreengrid.org/.

Green IT Calculator (Last access: 22th November, 2012). Web page at http://www.vmware.com/solutions/green/calculator.html.

Intel Xeon Processor e5 Family (Last access: 13th October, 2012). Web page at http://www.intel.com/content/www/us/en/processors/xeon/xeon-processor-5%000-sequence.html.

Kaushik, R. T. and Bhandarkar, M. (2010). Greenhdfs: towards an energy-conserving, storage-efficient, hybrid hadoop compute cluster. In *Proceedings of the 2010 international conference on Power aware computing and systems*, HotPower'10, pages 1–9, Berkeley, CA, USA. USENIX Association.

Kernel Based Virtual Machine (KVM) (Last access: October 13, 2012). Web page at http://www.linux-kvm.org/.

Lam, C. (2010). *Hadoop in Action*. Manning Publications.

Laszewski, G. and Wang, L. (2010). GreenIT Service Level Agreements. In Wieder, P., Yahyapour, R., and Ziegler, W., editors, *Grids and Service-Oriented Architectures for Service Level Agreements*, pages 77–88. Springer US.

Leverich, J. and Kozyrakis, C. (2010). On the energy (in)efficiency of hadoop clusters. *SIGOPS Oper. Syst. Rev.*, 44(1):61–65.

Liu, L., Wang, H., Liu, X., Jin, X., He, W. B., Wang, Q. B., and Chen, Y. (2009). Greencloud: a new architecture for green data center. In *Proceedings of the 6th international conference industry session on Autonomic computing and communications industry session*, ICAC-INDST '09, pages 29–38, New York, NY, USA. ACM.

OpenNebula: The Open Source Solution for Data Center Virtualization (Last access: 13th October, 2012). Web page at http://opennebula.org/.

Pang, B. and Lee, L. (2008). Opinion Mining and Sentiment Analysis. In *Foundations and Trends in Information Retrieval 2(1-2) – Available at: http://www.cs.cornell.edu/home/llee/opinion-mining-sentimen%t-analysis-survey.html*, pages 1–135.

Rivest, R. (1992). The MD5 Message-Digest Algorithm. RFC 1321 (Informational). Updated by RFC 6151.

SentiStrength: The sentiment strength detection in short texts (Last access: 10th October, 2012). Web page at http://sentistrength.wlv.ac.uk/.

Shi, B. and Srivastava, A. (2010). Thermal and power-aware task scheduling for hadoop based storage centric datacenters. In *Proceedings of the International Conference on Green Computing*, GREENCOMP '10, pages 73–83, Washington, DC, USA. IEEE Computer Society.

Sood, D. D. and Kumar, S. (2010). Cloud Computing & Green IT. Technical report.

UPS Selector Sizing Application (Last access: 22th November, 2012). Web page at http://www.apc.com/template/size/apc/.

White, T. (2009). *Hadoop: The Definitive Guide*. O'Reilly.