CISTER

Research Centre in
Real-Time & Embedded
Computing Systems

Conference Paper

On Power Consumption Profiles for Data Intensive Workloads in Virtualized Hadoop Clusters

Basit Qureshi

Sultan Alwehaibi

Anis Koubâa*

*CISTER Research Centre
CISTER-TR-170304

2017/05/01

On Power Consumption Profiles for Data Intensive Workloads in Virtualized Hadoop Clusters

Basit Qureshi, Sultan Alwehaibi, Anis Koubâa*

*CISTER Research Centre

Polytechnic Institute of Porto (ISEP-IPP)

Rua Dr. António Bernardino de Almeida, 431

4200-072 Porto

Portugal

Tel.: +351.22.8340509, Fax: +351.22.8321159

E-mail: aska@isep.ipp.pt

http://www.cister.isep.ipp.pt

Abstract

Although reduction in operating costs remains to be a key motivation for migration to Cloud environments, Power consumption is a big concern for data centers and cloud service providers. Many big data applications execute on Hadoop MapReduce framework for processing large workloads. In this paper, we investigate the tradeoff between energy consumption and workload running on Hadoop clusters using multiple virtual machines. We characterize power consumption profiles for various data intensive workloads and correlate these to quality of service (QoS) metrics such as job execution time. Based on experiments, we ascertain that power consumption profiles for big data applications can be used to optimize energy efficiency in data centers. We infer that these profiles can be used by Cloud service providers and consumers to specify green metrics in Service Level Agreements (SLA).

On Power Consumption Profiles for Data Intensive Workloads in Virtualized Hadoop Clusters

B. Qureshi[*], S. Alwehaibi[*], A. Koubaa[¶*]
[*]Prince Sultan University, Riyadh, Saudi Arabia
[¶]CISTER/INESC-TEC, ISEP, Polytechnic Institute of Porto, Porto, Portugal

Abstract: **Although reduction in operating costs remains to be a key motivation for migration to Cloud environments, Power consumption is a big concern for data centers and cloud service providers. Many big data applications execute on Hadoop MapReduce framework for processing large workloads. In this paper, we investigate the tradeoff between energy consumption and workload running on Hadoop clusters using multiple virtual machines. We characterize power consumption profiles for various data intensive workloads and correlate these to quality of service (QoS) metrics such as job execution time. Based on experiments, we ascertain that power consumption profiles for big data applications can be used to optimize energy efficiency in data centers. We infer that these profiles can be used by Cloud service providers and consumers to specify green metrics in Service Level Agreements (SLA).**

Keywords: MapReduce; Energy efficiency; Virtual Hadoop clusters; Power consumption

I. INTRODUCTION

In recent times, Cloud computing technology is widely being adopted by businesses and organization. The main driver for this move is the reduction in maintenance of infrastructure, deployment and management overheads as well as overall reduced operating costs. On the other hand, environmental impact of maintaining large computational infrastructure and data centers is a big concern prompting the need for research in "greener" technologies for data centers. Cloud service providers are increasingly incorporating green metrics into service level agreements (SLA) to market their services as environmental friendly [3]. While clean energy from solar and wind power is being increasingly used for data centers by well known cloud service providers, an important challenge is to investigate how to efficiently utilize resources within data centers to optimally consume energy whereas maximizing the cost benefit to both consumers and service providers [6][14-17]. Flexibility of Cloud systems as well as variety of configuration parameters makes it difficult to understand efficient utilization of each resource in data centers.

Deploying Hadoop efficiently across a cloud environment remains an important challenge. Cloud infrastructure deployments, configuration of various parameters and virtual cluster configurations can have a major impact on energy consumption and resource utilization in a data center. Due to unavailability of any standards for efficient deployment, there is an opportunity to study the impact of various optimized deployment techniques to monitor energy consumptions and resource utilization. Apache Hadoop framework [2] is a popular platform commonly used for analysis of data intensive operations and is widely used for research in Big Data analysis where large volumes of data cannot be analyzed using traditional technologies. Hadoop's Map/Reduce [1] has become a benchmark tool for comparing performance of various architectures for compute, network, storage and IO operations

[8-9]. Recent works have provided an opportunity for further investigating efficiency of Map/Reduce workloads in a Hadoop clusters. Tiwari et.al. in [15] argue that varying MapReduce parameters have a significant impact on computation performance and energy consumption for typical MapReduce workloads. Authors in [3, 4] and [11] outline the need for understanding the potential for energy saving in MapReduce Jobs in the context of CPU-bound, IO-bound or network-bound workloads. The work presented in this paper takes motivation from the aforementioned works and follows two objectives, i) to investigate monitoring variability of power consumption for multiple executions of a data intensive application in Hadoop, ii) appreciating the correlation of number of Virtual Machines per physical server and its impact on power consumption.

To understand the impact of resource utilization for various loads of data intensive computation and subsequently correlating its energy footprint, we make use of TeraSort benchmark [3] that is part of Hadoop framework. TeraSort is widely used as a stress test to allow Infrastructure-as-a-service (IaaS) administrators to optimize storage and network parameters configurations for optimal Hadoop deployment using HDFS and MapReduce layers of the Hadoop cluster. To this end we study various deployment models to measure, analyze and possibly optimize power consumption behavior for data intensive applications in virtual Hadoop clusters. We utilize two cluster testbeds RIoTU and Kafala testbeds with 4 low-end servers and 8-high-end servers respectively and deploy virtual Hadoop clusters using a number of virtual machines with various configurations. Power consumption across the clusters is measured against the associated workload generated using specialized power measurement equipment. We investigate the impact of scaling the number of virtual machines per server in the virtual cluster and analyze the performance and energy consumption. Furthermore, we provide a detailed evaluation of a set of MapReduce work-loads, highlighting significant variation in both the performance and power consumption of the applications.

The contributions of this work can be categorized as follows:
- Analyze the tradeoff between scalability of virtual machines per physical server and job completion efficiency on power consumption in virtualized Hadoop clusters.
- Provide insight into significance of power consumption profiles for various cloud-based applications. We believe that these profiles can be used by Cloud service providers and consumers to specify green metrics in SLAs.

The rest of the paper is organized as follows. Section 2 presents the related work. Section 3 presents the methodology with details on Hadoop virtual cluster setup, designing the workload and configuration of the power measurement equipment. Section 4 presents analysis of results with

characterization of power consumption profiles for the virtual cluster as well as analysis of computation times for various workloads, followed by conclusions in Section 5.

II. RELATED WORK

Recently Green energy harvested from solar and wind farms is being used in data centers to lower the overall emissions and carbon footprint [5, 7, 12-13]. In [5] authors analyze cost of energy on datacenters built in cold climates. Li in [7] proposed Oasis, a datacenter expansion strategy for scaling data center infrastructure while considering power/carbon emissions constraints. Oasis allows switching between green energy power supplies for optimizing power consumption. Hadoop has been extensively researched for its power inefficiencies within clusters. GreenHadoop [12] is a framework for data centers powered by photovoltaic solar arrays. The framework describes scheduling of Map/Reduce jobs based predicting the availability of solar power to maximize the green energy consumption. GreenHDFS [13] address developing energy saving mechanisms for the Hadoop Distributed File System (HDFS).

On the other hand, many recent researches point towards optimizing workloads in order to efficiently utilize energy in existing data centers. Tiwari [15] study the impact of Hadoop replication-factor, and its interaction among block-size, Map-slots and CPU-frequency. They conclude that Hadoop power consumption optimization is dependent on many factors including CPU frequency, placement of map tasks, scheduling of jobs, HDFS block-size and workloads. Krish in [11] present oSched, a workflow scheduler that profiles the performance and the energy characteristics of applications on hardware clusters. oSched considers power utilization from server machines in determining power configurations and energy profiles for scheduling of jobs. X.Dai in [19] focus on the placement of communicating virtualized servers in the data center in an energy efficient manner and proposed two algorithms, minimum energy virtual machine scheduling algorithm (MinES) and minimum communication virtual machine scheduling algorithm (MinCS).

E.Feller et.al [3] investigated the effect of virtual machine coexistence on the disk speed and evaluate the performance and power of Hadoop with datasets obtained from Wikipedia. They conclude that both write and read throughput decreases with increased number of virtual machines. Authors in [17] present an optimization approach using dynamic placement and migration of virtual machines in green cloud computing environment. The focus of this work is to enable clients in receiving acceptable service with a limited number of active servers.

The work presented in this paper focuses on characterizing the power consumption vs. Quality of Services (QoS) metrics in data centers. We consider outlining power consumption profiles for typical Big data applications in order to optimize power consumption in Virtualized Hadoop clusters. These power consumption profiles can be used to help determine the number of virtual machines to be deployed on physical servers to achieve throughput within acceptable constraints.

III. METHODOLOGY

Power Usage Effectiveness (PUE), which was developed by the Green Grid Association is the key metric used in data centers. PUE is used as the ratio of power entering the data center divided by the power used to run the computation infrastructure. It is noticeable that the large portion of power consumption in the data center is due to (non-compute) related infrastructure such as buildings, air-conditioning systems etc). Furthermore, the utilization of physical machines of the clusters in data center is sub-optimal with nodes idling around 70% of the time [17-19]. It is important to understand the behavior of power utilization for various applications for their intensity of resource utilization. Based on these power consumption metrics, policies can be generated to optimally utilize the data center resources thus reducing the overall power consumption.

In this work, we focus on characterizing the power consumption vs. compute performance tradeoff for virtualized Hadoop deployment over Infrastructure-as-a-service (IaaS) cloud environment. It is important to understand the relationship between power consumption and performance as QoS metric in optimizing virtual machines deployment policies. To this end, we characterize the power consumption profiles for data intensive applications. For time intervals when deployment of physical machines yields poor power consumption vs. performance tradeoffs, the optimal power consumption policies can be applied. A number of virtual machines would be deployed on the cluster to maximize the power consumption tradeoff. Consequently, if the benefit of the tradeoff between power consumption and performance outweighs the deployment with virtual machines, the user may decide not to optimize the performance. In what follows we describe the cluster environment, characterization of the workloads and power measurement process used in this study.

3.1 Hadoop Virtual Cluster Environment

The experimental investigation carried out in this paper focuses on the performance of virtualized Hadoop clusters given Data intensive workloads typically used in big data applications. We conduct a series of experiments in order to assess the impact of various parameters of virtual machine configuration applicable to workloads of varying sizes for performance and power consumption. To this end, we deploy two virtual Hadoop clusters namely RIoTU Testbed and Kafala Testbed. The RIoTU Testbed is composed of four HP ProLiant machines with single Intel Core i7 processor running at 3.67GHz connected to a Gigabit Ethernet. Each machine has 8GB of RAM with 256GB of Kingston Solid State Storage devices running windows 10 as host operating system. These machines are connected to the WattsUp .net power measurement equipment for collecting reliable power consumption data at timely intervals.

The Kafala Testbed is composed of 8 servers used in this study. Each server machine is equipped with 2 Intel Xeon E5-2667 processors running at 3.30 GHz with 48GB RAM and 2TB SCSI Storage. Each server runs Windows Server 2012R2 as Host operating system with VMware used for running virtual machines. The servers in the Kafala testbed are isolated from the rest of the datacenter for performance parameters measurement for this experimentation.

On both of these cluster testbeds we deploy virtual machines running Ubuntu 16.0 LTE and Apache Hadoop 2.6.2. Table 1 shows the various configurations of virtual machines deployment on the cluster testbeds. One of the virtual machines server as the master node running the Hadoop Namenode and

Table 1. Hadoop Virtual Machine Configurations.

RIoTU Testbed				Kafala Testbed			
No of VMs Per Server	CPU	RAM (GB)	HDD (GB)	No of VMs Per Server	CPU	RAM (GB)	HDD (GB)
1 Server (Stand-alone)	20%	4	120	1 Server (Stand-alone)	20%	48	980
1 VM	70%	4	50	1 VM	70%	4	50
2 VM	35%	2	50	2 VM	35%	2	50
4 VM	17.5%	1	50	4 VM	17.5%	1	50

YARN Resource-Manager, the rest of the virtual machines execute a single Data-node and Node-manager. In Hadoop configuration files the maximum MapReduce resource memory was set to 1GB with a replication factor of 2.

3.2 TeraSort workload

The Hadoop TeraSort benchmark suite sorts data as fast as possible to benchmark the performance of the MapReduce framework. TeraSort combines testing the HDFS and MapReduce layers of a Hadoop cluster and consists of three MapReduce programs, TeraGen, TeraSort and TeraValidate. TeraGen is typically used to generate large amounts of data blocks. This is achieved by running multiple concurrent map tasks. In our experimentation, we use TeraGen to generate large datasets to be sorted using a number of map tasks writing 100-byte rows of data to the HDFS. TeraGen divides the desired number of rows by the desired number of tasks and assigns ranges of rows to each map. Consequently, TeraGen is a write intensive I/O benchmark. The TeraSort generates set of sample keys by sampling the input data generated by TeraGen before the job is submitted and writes the list of keys into HDFS. The input and output format, which are used by all three MapReduce programs, reads and writes the text files in the correct format.

By design each TeraSort MapReduce job is executed in two steps: map and reduce. During these steps, various Computation (CPU) intensive, disk I/O intensive and Network I/O intensive subtasks with varying workloads are initiated. The workloads depend on the number of map and reduce at initiation of the job. The map tasks read input data from files generated by TeraGen and outputs intermediate data. At the completion of writing the intermediate data to the disk, the reduce step reads the indexed files from Disk to the memory referred to as shuffle buffer. The merged and sorted data is used by the reduce step to write the output to the Disk. It is important to characterize these steps into CPU intensive, IO intensive and Network Intensive operations.

i. From the launch of TeraSort job to the moment the first map task is read into memory (Disk I/O intensive).
ii. From the initiation of map input until map output is written to disk (CPU intensive)
iii. From the writing of first map until all map tasks are completed (CPU, Disk IO and Network IO intensive).
iv. From the completion of all maps until last shuffle task is done (Disk IO and Network IO intensive).
v. From the end of shuffle task until all reduce tasks are done (CPU and Disk IO intensive)
vi. From the end of reduce tasks until the job is finished (Disk IO intensive).

In our experimentation, we run TeraGen and TeraSort on both clusters due to its intensive workload which is correlated to a data intensive big data application. We execute these for various runs with data size in the range of 0.1GB, 1GB, 10GB and 100GB respectively. We observe the job execution time for each run for comparison and analyze the performance on each cluster. The results and analysis of these experiments are provided in the next section.

3.3 Power measurement

Since Hadoop exploits all resources (CPU, memory, Disk and Network IO) of the compute environment it is important to analyze the power consumption of the cluster collectively. External devices such as the WattsUp Pro [1] power consumption meter are required since the collective power consumption of the entire cluster cannot be monitored from the local monitoring software. In this experimental study, we use the WattsUp Pro .net power meter that logs the power used in terms of watts at time intervals specified, into the unit's non-volatile memory. The unit allows easy download of data using the USB cable connected to an external device (such as laptop). The user can also collect data only when the power consumption exceeds a predefined threshold.

IV. EXPERIMENTAL RESULTS

In this work, we focus on attempting to find an optimal tradeoff between power consumption and data intensive MapReduce workloads using TeraSort benchmark on Hadoop virtual clusters. The power consumption of the virtual cloud environment running Hadoop can be characterized by using power consumption profiles. A power consumption profile for a cloud-based application is the characterization of its power consumption levels at different time intervals during its execution on the cloud testbed. We define and explain the various levels of power consumption obtained from both cloud testbeds used in this study to describe the power consumption profiles for TeraSort as an instance of a big data application. Furthermore, we determine the power usage of TeraSort jobs with workloads of various sizes and compare these for different configurations of virtual machine deployment in the clusters. Finally, we provide a performance comparison for these jobs in terms of CPU execution times and analyze the results.

4.1 Power usage profiles

The power usage profiles can be specified for applications executing in a data center. We observe six distinct power consumption levels for TeraSort jobs running on the cluster testbeds from Host machine running in idle mode, to initiation of TeraSort job, to completion and shutting down of the virtual cluster. The choice of the host machine operating system, virtualization software and hardware capabilities also have an impact on the overall power consumption. We therefore provide the average values for power consumption at each level to understand the behavior. Figure 1 shows the power usage profile for TeraSort with various levels of power usage at

---

[1] http://www.wattsupmeters.com

different time intervals on the RIoTU testbed. Table 2 describes power usage levels w_0 to w_6 in time intervals t_0 to t_7.

As the machines in the clusters are booted, there is a small peak in power consumption due to the use of IO operations in running the host operating system. The level W_0 is the idle mode when the machines are running with Host operating system idling without any virtual machines running ($t_0 < t_1$). An increase in the power consumption is observed for level W_1 when virtual machines are started until the guest operating system in the virtual machines is running ($t_1 < t_2$). A small but noticeable increase in power consumption is observed when Hadoop is started in each virtual machine for level W_2. This value increases as the number of virtual machines executing per node also increases. When the TeraSort job is initiated we observe a significant increase in power consumption due to the intensity of Disk IO, Network IO tasks running at the same time in all virtual machines on the cluster.

For TeraSort jobs, Hadoop stresses the system increasing the power consumption significantly for a short period of time ($t_2 < t_3$) for level W_3. As the map step begins, the map tasks start reading the data from Disk increasing the Disk I/O but reducing the overall power consumption at level W_3 to level W_4. The cluster maintains almost a constant power consumption time with a variability of ±4% in power consumption until the *Shuffling* phase is completed and the reduce jobs are started. As the reduce jobs complete, the power consumption also reduces due to the decrease in number of parallel tasks executing in the cluster. We define level W_5 to depict the completion time of TeraSort job. In our experimentation levels, W_2 and W_5 were observed to be very close. Level W_6 defines power consumption behavior when Hadoop is shutdown. Finally, the physical machine can be put to idle state when we close all the virtual machines. Table 3 shows the minimum, maximum and average power consumption (watts) for RIoTU testbed.
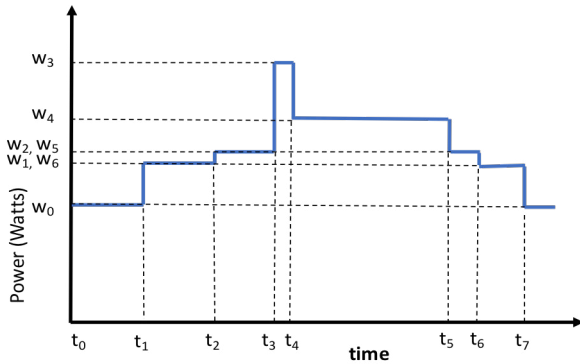
4.2 Power usage for TeraGen and TeraSort tasks

We study the power consumption on RIoTU cluster using the TeraGen and TeraSort benchmark due to their intensive CPU and IO bound operations. To accurately measure power consumption in the cluster, a Wattsup Pro .net power meter is attached to the cluster and the power mains. The Wattsup Pro .net meter is capable of recording power consumption in terms of watts, each reading is collected every 10 seconds and is logged in the meter's onboard memory. The meter is initialized 60 seconds before each TeraGen and TeraSort job is initiated and stops reading 60 seconds after the job is completed.

In order to run TeraSort, data files need to be generated in the HDFS using TeraGen using the single, 2 and 4 virtual machine configurations. TeraGen was executed 10 times each for dataset sizes of 100MB, 1 GB and 10 GB respectively. For each of these jobs, 10 map tasks with 1 reduce tasks were provided as parameters. Figure 2 shows the power consumption (in terms of watts) against time and completion rate for the cluster setup using 2 and 4 virtual machines for 10 GB datasets. As the job initiates, we notice a spike in power usage for a short period of time for both VM configurations. We observe this behavior due to the intensive read/write Disk and Network IO operations. As the distribution of the map tasks over the clusters is completed, the map tasks start executing slightly reducing the power consumption. Since TeraGen is IO bound job, map tasks write to the HDFS and we observe steady power consumption until map tasks are completed. With the progress of map task completion, we notice a drop in power consumption due to the decrease in Disk IO and completion of the job. We compute the ratio of power usage in terms of Watts per hour for each of these configurations. For a single machine configuration, we obtain the Energy consumption E in Kilo Watts per hour (KWh) to be $16.42 * 10^{-2}$ KWh. For 2VM and 4 VM configurations we obtain $16.381 * 10^{-2}$ KWh and $12.231 * 10^{-2}$ KWh. This indicates that executing this task in

Figure 1: Power consumption profile for TeraSort

Table 2: Power usage levels for different time intervals

Time Interval	Power Usage Level	Description
$[t_0, t_1]$	w_0	Host OS idle with no VMs running
$[t_1, t_2]$	w_1	Virtual Machines started
$[t_2, t_3]$	w_2	Hadoop Started and working
$[t_3, t_4]$	w_3	TeraSort Map starting phase
$[t_4, t_5]$	w_4	TeraSort Map/Reduce in progress
$[t_5, t_6]$	w_5	TeraSort Job completed
$[t_6, t_7]$	w_6	Hadoop shut down
$[t_7, \ )$	w_0	VMs shut down, Host is idle

Table 3: Power usage for various workloads on the RIoTU testbed

Workload (MB)	No of VMs	Min power (watts)	Max power (watts)	Average (watts)	Variability
100 MB	1	85.5	91.2	88.35	±2.85
100 MB	2	105.1	114.6	109.85	±4.75
100 MB	4	238.5	246.6	242.55	±4.05
1000 MB	1	87.1	90.1	88.6	±1.50
1000 MB	2	106.4	115.3	110.85	±4.45
1000 MB	4	239.3	245.4	242.35	±3.05
10000 MB	1	86.8	91.3	89.05	±2.25
10000 MB	2	107.4	114.9	111.15	±3.75
10000 MB	4	241.6	246.5	244.05	±2.45

4VM configuration is cost efficient compared to single and 2VM configurations.

We observe similar power consumption patterns for TeraSort jobs. The TeraSort generates a set of sample keys by sampling the input data generated by TeraGen before the job is submitted, and writes the list of keys into HDFS. The input and output format, which are used by all three MapReduce programs, reads and writes the text files in the correct format. The TeraSort benchmark is CPU bound during the map phase as it reads input data and shuffles it, I/O bound during the reduce phase for writing output to HDFS. We notice a similar spike in power usage at the initiation of a TeraSort job while map tasks are written across various nodes in the cluster as can be seen in Figure 3. As the mappers continue to complete the tasks, the incoming results start processing in the reduce jobs. Before the completion of all map tasks, the reduce tasks initiate sorting and summarizing process requiring CPU as well as IO resources towards completion of the tasks. Whilst the distributed tasks complete, the power consumption drops. We notice that the trends are similar for other data sizes used in this study. As can be seen from Figure 3, the percentage of map tasks and reduce tasks completed correlates with the power consumption for both 2VM and 4 VM configurations. In particular, when the map and reduce tasks complete, the power consumption drops therefore highlighting underutilized nodes in the clusters.

Discussions. Both TeraGen and TeraSort exhibit different power consumption. TeraSort on both clusters has a relatively long phase of higher power consumption from initialization of map jobs until about 80% of map jobs completion indicating high CPU utilization. Afterwards, the power consumption decreases slightly fluctuating while both map and reduce jobs are executing in parallel. Finally, the power consumption steadies with minor tails and peaks in the plot towards reduce jobs completion. For TeraSort job execution on a single machine configuration, we obtain the Energy Consumption E to be 0.136 KWh. For 2VM and 4 VM configurations we obtain 0.128 KWh and 0.151 KWh. Although the runtime for the same TeraSort jobs in 4VM configuration is time efficient, however the ratio of power consumption is 17% higher. Comparing the 2 VM and single VM configurations, it is clear that 2 VM configuration is both time (27% faster) and power efficient (6% less power) than single VM configuration. Overall the results presented a tradeoff between power consumption and time efficiency for various VM configurations. In all cases, running multitenancy of VMs per server provides better power efficiency when compared to single VM or physical system configurations.

4.3 Computation Execution times

In recent studies, various Quality of Service metrics for execution of parallel jobs in a Hadoop cluster have been employed. In this study, we analyze the impact of Virtual machines configurations on CPU Execution (computation) time for executing TeraSort jobs on datasets of 100MB, 1GB and 10GB sizes. We observe the job execution time for each run for comparison and analyze the performance on both cluster testbeds. The experiments were run 10 times for each data-size on each cluster. Figure 4 shows box whisker plots for the job completion time (CPU Execution Time) for TeraGen and TeraSort for varying data payloads. Performance in terms of
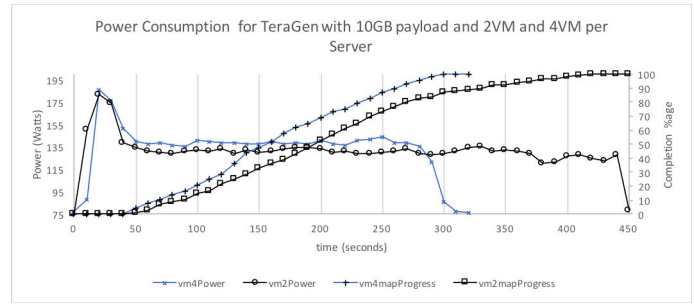

Figure 2. Comparison of Power consumption for TeraGen with 10GB workload on 2VM and 4VM configurations for RIoTU Testbed.
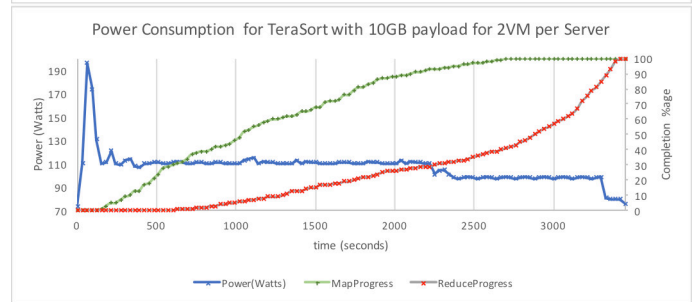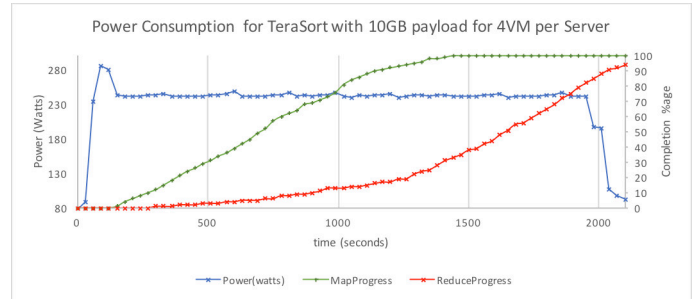

Figure 3. Power consumption and map/reduce completion rate for TeraSort with 10GB workload on 4VM and 2VM configurations for RIoTU Testbed.

job completion time is correlating in RIoTU and Kafala clusters when payloads are increased, however the completion time for these jobs is different. For TeraGen with 10GB file size and with 2 VM configuration, both clusters present similar CPU execution times, however with 4 VM configuration, the RIoTU cluster performs better. TeraSort on the other hand is CPU and IO intensive for map and reduce phase respectively. For 2 VM configuration the CPU execution time for smaller TeraSort jobs (0.1GB and 1GB) is 0.7 and 2.8 times faster for Kafala Cluster due to the increased number of servers and virtual machines.

For larger dataset (10 GB) the performance of Kafala cluster is slightly better. With 4VM configuration the Disk IO per physical server increases due to the larger number of virtual machines therefore affecting the read/write speeds on the local disks. This is visible in Figure 4 where the run time for TeraSort with 10GB file sizes is 0.12 times faster for RIoTU cluster. Since TeraGen and the reduce phase of TeraSort is IO intensive, the larger run time with 4VMs is due to increased Disk IO. As RIoTU servers are equipped with faster Solid State Disks, the disk speed directly correlates with TeraGen and TeraSort completion time for larger file sizes. However for smaller file sizes, the larger number of virtual machines running per server yield better run times.
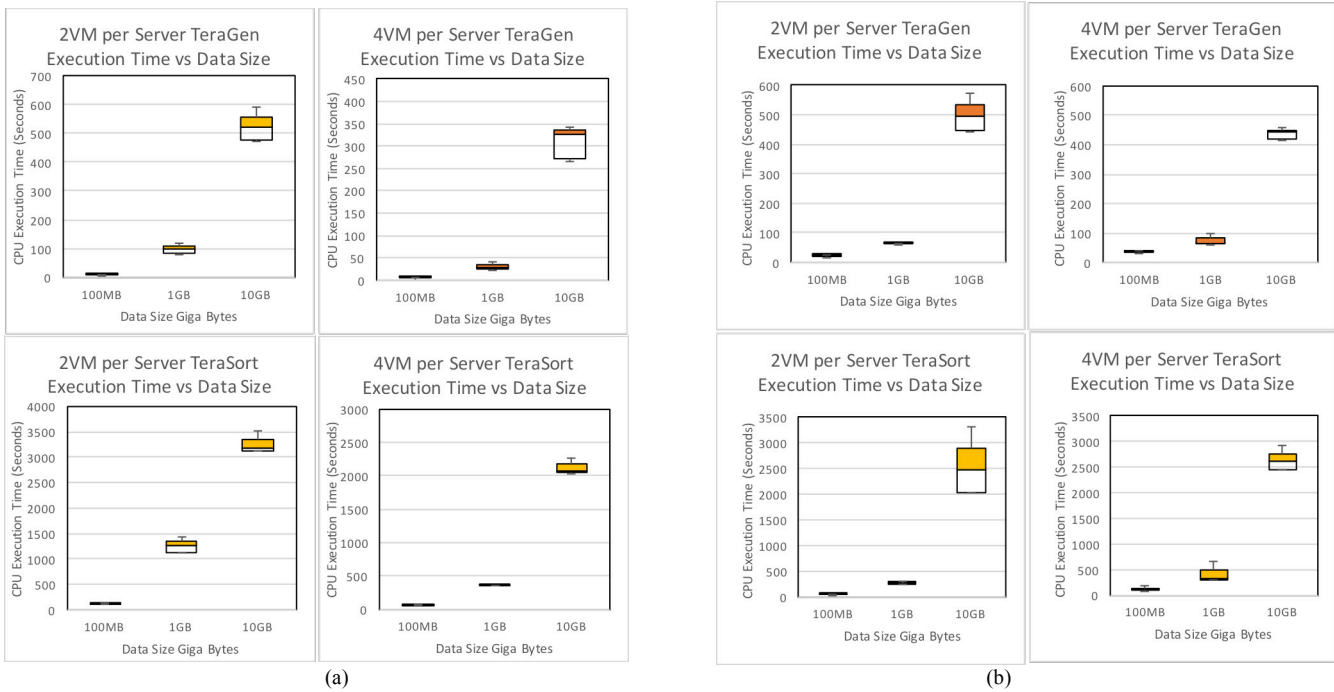
Figure 4. CPU Execution (Job completion time) for TeraGen and TeraSort for 2VM and 4VM configurations in (a) RIoTU Testbed and (b) Kafala Testbed.

V. CONCLUSIONS

Energy efficiency of data centers enabling the cloud is fast becoming a governing issue and key research direction in data center design, deployment and operation. In this paper, we investigated the issue of power consumption profiles for data intensive big data applications in determining the optimal tradeoff between power consumption and job completion time in virtualized Hadoop clusters. To this end, we deployed two virtual Hadoop cluster testbeds to analyze the power consumption behavior and time efficiency of executing TeraSort jobs with various payloads. We also observed for large file sizes the role of efficient storage media is imperative. We conclude that there is a direct correlation between the number of virtual machines and data workloads executed on these VMs compared to execution on physical machines. The work presented in this paper helps identifying how many VMs per machine can be deployed to achieve throughput at a given power consumption profile assisting decision makers in optimizing energy efficiency of the infrastructure. Although we used private cloud infrastructure as testbed for this experimental study, we believe the multitenancy in public cloud environments where workloads and number of VMs per machine greatly vary over time, can benefit from this study.

VI. ACKNOWLEDGEMENTS

VII. REFERENCES

[1]. J. Dean, S. Ghemawat, "MapReduce: simplified data processing on large clusters", Commun. ACM, 51 (1) (2008), pp. 107–113.

[2]. The Apache Hadoop Project, 2017. URL: http://www.hadoop.org.

[3]. Eugen Feller, Lavanya Ramakrishnan, Christine Morin, Performance and energy efficiency of big data applications in cloud environments: A Hadoop case study, Journal of Parallel and Distributed Computing, Volumes 79–80, Pages 80-89, May 2015.

[4]. Javier Conejero, Omer Rana, Peter Burnap, Jeffrey Morgan, Blanca Caminero, Carmen Carrión, Analyzing Hadoop power consumption and impact on application QoS, Future Generation Computer Systems, Volume 55, Pages 213-223, February 2016.

[5]. E. Sheme, N. Frashëri, S. Holmbacka, S. Lafond and D. Luzanin, "Datacenters powered by renewable energy: A case study for 60 degrees latitude north," 2016 24th International Conference on Software, Telecomm. and Computer Networks (SoftCOM), Split, Croatia, 2016.

[6]. Z. Zhou, F. Liu and Z. Li, "Bilateral Electricity Trade Between Smart Grids and Green Datacenters: Pricing Models and Performance Evaluation", IEEE J. on Selec. Areas in Com., vol. 34, no. 12, pp. 3993-4007, Dec. 2016.

[7]. C. Li; Y. Hu; J. Gu; y. jingling; T. Li, "Oasis: Scaling Out Datacenter Sustainably and Economically," in IEEE Trans. on Parallel and Distributed Systems (To appear 2017).

[8]. H. Jin, S. Ibrahim, L. Qi, H. Cao, S. Wu, X. Shi, "The MapReduce programming model and implementations", Cloud Computing: Principles and Paradigms, pp. 373–390, 2011.

[9]. Y. Chen, S. Alspaugh, D. Borthakur, R. Katz, Energy efficiency for large-scale mapreduce workloads with significant interactive analysis", in proc. 7th ACM European conf. on Comp. Systems, EuroSys, pp. 43–56, 2012.

[10]. Shadi Ibrahim, et. al., Governing energy consumption in Hadoop through CPU frequency scaling: An analysis, Future Generation Computer Systems, Volume 54, Pages 219-232, January 2016.

[11]. Krish. K. R., M. S. Iqbal, M. M. Rafique and A. R. Butt, "Towards Energy Awareness in Hadoop", 4th International Workshop on Network-Aware Data Management, New Orleans, LA, pp. 16-22, 2014.

[12]. I. Goiri, K. Le, T.D. Nguyen, J. Guitart, J. Torres, R. Bianchini, "GreenHadoop: leveraging green energy in data-processing frameworks", Proceedings of the 7th ACM European Conference on Computer Systems, EuroSys'12, pp. 57–70, 2012.

[13]. R.T. Kaushik, M. Bhandarkar, "GreenHDFS: towards an energy-conserving, storage-efficient, hybrid hadoop compute cluster", Procs. of Intl. Conf. on Power Aware Computing and Systems, HotPower'10, USENIX Association, Berkeley, CA, USA 2010.

[14]. K. Bilal, et.al., "A taxonomy and survey on green data center networks", Future Gener. Comput. Syst., 36 (0) (2014), pp. 189–208.

[15]. N. Tiwari, U. Bellur, S. Sarkar and M. Indrawan, "Identification of critical parameters for MapReduce energy efficiency using statistical Design of Experiments," 2016 IEEE Intl Parallel and Distributed Processing Symposium Workshops (IPDPSW), Chicago, IL, 2016, pp. 1170-1179.

[16]. F. Kong and X. Liu, "GreenPlanning: Optimal Energy Source Selection and Capacity Planning for Green Datacenters," 2016 ACM/IEEE 7th Interntl. Conference on Cyber-Physical Systems (ICCPS), Vienna, 2016.

[17]. M. Anan, N. Nasser, A. Ahmed and A. Alfuqaha, "Optimization of power and migration cost in virtualized data centers," 2016 IEEE Wireless Communications and Networking Conference, Doha, 2016, pp. 1-5.

[18]. Xiang Deng ; Di Wu ; Junfeng Shen ; Jian He, "Eco-Aware Online Power Management and Load Scheduling for Green Cloud Datacenters", IEEE Systems Journal, Vol. 10, Issue 1, Page(s): 78 – 87, March 2016,.

[19]. Xiangming Dai, Jason Min Wang, Brahim Bensaou. "Energy-Efficient Virtual Machines Scheduling in Multi-Tenant Data Centers" IEEE Transactions on Cloud Computing, pp. 210-221, 2016.