

The Effect of Cloud Workload Consolidation on Cloud Energy Consumption and Performance in Multi-Tenant Cloud Infrastructure

Kenga Mosoti Dardus
Strathmore University
Nairobi, Kenya

Vincent Oteke Omwenga
Strathmore University
Nairobi, Kenya

Patrick Job Ogao
Technical University of Kenya
Nairobi, Kenya

ABSTRACT

As energy consumption is becoming a problem in cloud data centers, cloud service providers have adopted different techniques to address this problem. One of the most attractive technique is virtual machine (VM) consolidation. Apart from reducing energy consumption in computing platforms, this technique has other advantages such as reduced infrastructure costs and ease of virtual machine management. However, VM consolidation, which does not recognize workload characteristics may, in the long run, increase energy consumption and lead energy wastage. This paper investigates the relationship between different VM workload types and server energy consumption in a multi-tenant datacenters. Experiments are conducted using well known CPU, I/O, memory and network intensive workload benchmark obtained from Phoronix Test Suite (PTS). Results obtained show that there is a noticeable difference in the amount of energy consumed when VMs run workloads, which dominate the various server physical resources. Secondly, consolidating homogeneous workloads is disastrous in terms of energy consumption and performance over heterogeneous workloads. The latter can further reduce energy consumption and achieve acceptable performance levels if an optimum workload mix is reached.

Keywords

Cloud computing, cloud workloads, data center energy consumption, cloud workload consolidation, multi-tenant cloud, IaaS

1. INTRODUCTION

In recent years, the demand for cloud computing infrastructure has increased constantly. To respond to this, cloud service providers and private organizations are putting up many servers in their datacenters to provide more processing power for their computing needs [1]. Unfortunately, data centers consume a lot of energy, which is about 3% of global electrical energy consumption [2]. The effect of this is that cloud service providers experience high operating costs [3]. Moreover, the excessive amount of energy consumed has a negative impact on the environment, which is the emission of carbon dioxide – the ICT industry contributes about 2% of global carbon dioxide emission.

The defacto technique adopted to reduce energy consumption and other operational costs in data centers is workload consolidation through virtualization [1]. Virtualization makes it possible to run many different workloads in the same physical server using VMs thus it is possible to ensure that less physical servers are needed. By the year 2013, at least 75% of enterprise workloads were executed in virtualized environments.

Cloud computing concepts encompass three major layers namely, Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS) and each of these layers play a crucial role in satisfying customers' (tenants) needs [4]. A number of multinational companies such as Google, Amazon, HP, IBM, Citrix, Rackspace, Microsoft, DigitalOcean, Linode and Vultr are already providing cloud computing services [4].

In IaaS, a customer is provided with a VM, which they use as they wish. Thus, the cloud provider has no control over the kind of applications the customers run. Because the cloud providers want to utilize the resources effectively, many VMs are co-located in the same physical machine. As the number of VMs sharing the physical resource increase, they compete for hypervisor capacity and the shared physical resources. As a result, interference is caused, which leads to violation of Service Level Agreements (SLAs) [5]. From the forgoing, it is clear that VM consolidation need to be informed by the nature of VM workloads. This paper investigates the effect of consolidation of homogeneous and heterogeneous workloads IaaS cloud. To simulate real cloud workloads, workload benchmarks available in PTS have been utilized, which have been described in Section 3 of this paper. The main contributions of this works are:

- An investigation on the difference in energy consumption by VMs running various workloads that dominate the various physical server resources.
- An investigation on the effect of consolidating heterogeneous and homogeneous cloud workloads on server power consumption and VM performance in multi-tenant cloud.
- To the best of our knowledge, this is the first academic work that has experimentally explored the use Phromatic for automating test orchestration.

The reminder of this paper is organized as follows: section 2 presents related research, which has been done in workloads types and energy consumption. Sections 3 provides a summary on PTS. Section 4 describes experimental setup including how data was collected. In section 5, results and discussion are presented. Finally, this paper is concluded in section 6.

2. RELATED WORK

To achieve better VM consolidation, [6] uses business critical workloads to study the demand and dynamism of workloads on the various computing resources. Using correlation, they also study the dependency among the various computing resources. Their results show that there is a strong correlation between high CPU and memory usage among VMs. This means that CPU-intensive workloads are likely to be memory-

intensive. This is a similar observation made by [7]. However, they discovered that it is less likely for VMs to exhibit high CPU and memory simultaneously. Correlation between CPU and other resources usage (network and disk) were discovered to be very low.

In [7], the authors discovered that the amount of energy consumption in physical server depends on the number of VMs running in it. These authors have also used the transactional web benchmark TPC-W to show that the amount of energy consumption and performance depends on workload profiles such as browsing and ordering.

In [8], the authors recognize that efficient consolidation is a solution to reducing energy consumption in virtualized environments. However, blind consolidation may lead to energy wastage. In their experiments to discover, which workload profiles are good candidates for consolidation, they concluded that performance degradation caused by CPU intensive workload are ignorable. On the other hand, disk intensive workloads are not consolidation friendly. This is attributed to performance overheads caused by a layer of virtualization software. Besides in [13], computing trends depicts a growing gap between CPU and I/O performance and bandwidth. For this reason, it observed that during data intensive operations, bottleneck caused by storage and memory leads to low CPU utilization.

In [9] the authors carry out experiments to study the difference in energy consumption in the various cloud workloads on different hardware and architectures. They have used various parameters such, processor manufacturer, processor mode, processor core and threads, memory size, hard disk characteristics, to vary the hardware and architecture. They have shown that some architectures are energy efficient on CPU-intensive workloads while others are efficient on disk-intensive workloads. This effectively forms a basis for workload characterization before consolidation.

In [10] the author carried an experiment to determine the types of workloads that are beneficial in terms of energy consumption when run in VMs co-located in the same physical server. Multiple VMs running homogenous workloads (CPU-intensive), homogenous workloads (memory-intensive) and heterogeneous workloads (mixtures of CPU-intensive and memory-intensive) were run in the same physical server. The results show that co-scheduling VMs with similar profiles in terms of resource demands is not beneficial from an energy consumption point of view.

[11] is the first work that has examined difference in energy consumption by server host running various workloads that dominate the various physical server resources using PTS. Although their primary goal was to use software-bases power measurement technique as the basis for efficient consolidation of cloud workloads, they have shown that an optimal workload mix needs to be determined, which is favorable to energy consumption and performance.

Finally, the work in [1] recognizes that VM consolidation can achieve great energy savings in virtualized environments. However, when more instances of a VM resides in a single machine in multi-tenant clouds, there is competition for the shared resources leading to interference. Their solution is as a VM scheduling model used to perform a diagnosis of interference and has been implemented on top of Xen.

3. BENCHMARKING WITH PHORONIX TEST SUITE (PTS)

PTS [11, 13, 14, 15] is the most compressive open source benchmarking tool. PTS has a collection test profiles with different but known behaviors. PTS has successfully been used to test computer performance, hardware validation and as a source of cloud workload. The test suites in PTS are developed to run in a number of operating systems such as Microsoft windows, MacOSX, Linux, Solaris and BSD. Today, PTS has over 450 benchmark micro-programs which target hardware performance and stability such as disk, processor, I/O, GPU, memory, network and overall computing system [14]. When a benchmark is ran, PTS uses sensors to read different variables such CPU power consumption, system power consumption, CPU frequency, memory usage, CPU temperature and I/O speed [15]. The latest version (8.2.0) of PTS includes a server known as Phoromatic, which automates the process of test orchestration to any number of clients systems and can be used to schedule tests that require to run at particular times [15].

When PTS test suites are used as a source cloud workloads, they are tagged depending on the computing resource that is targeted. For instance, to test an aspect of a CPU such performance, a test suite known as 7zip is used [11, 16]. The following is a description of 4 selected test suite, which has been used in this paper as a source workload and are summarized in Table 1.

3.1 7zip

7zip is file compression application using 7-zip compression with an observed overheads on CPU and a bit of memory. Hence it is labeled as a CPU benchmark or CPU intensive workload. 7zip is designed to compress a large files during which the CPU demand remains constantly high. The current version of 7zip is 1.7.0.

3.2 gzip

This is a memory PTS benchmark, which measures the time it takes to compress two copies of version 4.13 of Linux kernel source tree using gzip compression. This application is observed to have a high overhead of memory and CPU, thus it is labelled as a memory-intensive workload. The size of the Linux kernel source three is about 2GB. 7zip and gzip performs almost a similar task but with an observed difference in the resources demand, which dominates the operation.

3.3 aio-stress

This is an asynchronous I/O benchmark created by SuSE. The current profile (version 1.1.1) uses a single thread and constantly reads and writes a 2048 MB test file and a 64KB block size from and to the hard disk. For this reason, the system subcomponent that is being tested is the hard disk.

3.4 Loopback TCP network

This benchmark puts an activity on the loopback network adapter to measure its performance. It simulates network traffic and this is it labelled as a network-intensive workload.

Table 1: Summary of Selected PTS Tests

Select PTS test	Targeted resource
CPU	7zip
Memory	gzip
Network	Loopback TCP network
Hard disk	aio-stress

4. EXPERIMENT SETUP

This paper has used four selected PTS test suites shown in Table 1 as the source of workloads. The first experiment is setup in a Lenovo Personal Computer (PC) to investigate the behavior of each individual workload in terms of resource usage and consumption. The system information for the PC is shown in Table 2. Each of the test suite is executed in turn with sensors activated via PTS to collect data such as CPU usage (in percentage) over time, system power consumption (Milliwatts (MW)) over time and memory (megabytes (MB)) usage over time. The graph plots presented in the results are automatically generated by PTS.

Table 2: PC System Information

Components	Component details
CPU	Intel core i7-4600U @ 3.30GHz Core count: 2 Thread count: 4 Cache size: 4096 KB
Memory	Memory size: 8192MB
Operating system	Ubuntu 18.04 server Kernel: 4.15.0-34 (x86_64)
Disk	SSD 128GB SanDisk SD&SB3Q1

The second experiment is setup to investigate the effect of workload mixing on power consumption and performance. In this experiment, all tests are used. However, 7zip and aio-stress benchmarks are used when investigating the effects of workload mix on power consumption and performance. Power is measured in Watts (W). Performance of CPU, disk or I/O, loopback network and memory is measured Million Instructions per Second (MIPS), Megabytes per Second (MB/s), seconds (s) and seconds (s) respectively. A physical server (host) running Ubuntu server OS is set to run 8 VMs (guests). Each VM also runs Ubuntu server OS. A VM (s) either executes homogenous (similar) workloads or heterogeneous (mixed) workloads in turn during. During this period, power consumption is recorded every second over a period of time using Linux powerstat on Running Average Power Limit (RAPL) interface [17]. Workload orchestration on the VMs (clients) is made possible using PTS Phoromatic server. The characteristics of host and guests is summarized in Table 3. Tests run for the two experiments were executed multiple times and checked for consistency. It is also worth noting for every test executed, PTS repeats it 3 times before presenting the results.

Table 3: A summary of characteristics of Physical Server (host) and Virtual machine (guest)

Item	Item details
Physical server (host)	Processor: Intel Core i7-4790 @ 4.00GHz, 8 cores 8192 KB cache size. Operating system: Ubuntu 18.04, 4.15.0-36-generic (x86_64) kernel version Memory size: 12288MB Disk size: 1000GB Western Digital WD10EZEX-60M.
Virtual machine (guest)	Processor: 1 vCPU @ 3.59GHz Operating system: Ubuntu 18.04, 4.15.0-36-generic (x86_64) kernel version Memory size: 1280MB Disk size: 43GB Hypervisor/system layer: KVM
Number of	8

VMs (guests)	
--------------	--

5. 5. RESULTS

5.1 5.1. Individual Workload Examination

a) 7zip - The graphs shown in Fig. 1-3 show CPU usage, memory usage and power consumption over time when 7zip benchmark is run. 7zip is a computation intensive workload. Although the average CPU usage is at 87.3%, it is observed that CPU usage is *almost* consistently at 100% (Fig. 1). Memory usage ranges from 1272 MB to 2127 MB with an average of 1482MB and it seems to be varying a lot than percentage CPU usage. The amount of memory used by 7zip is used to show this workload is also memory intensive (Fig. 2). The power consumption over time (Fig. 3) ranges from 5875 MW to 201457 MW and has an average of 18815 MW. Power consumption can also be observed to be consistent over time with two slight dips similar to those observed with CPU usage. At this point, it can loosely be concluded that CPU is the main consumer of power in computing infrastructure.




Fig 1: 7zip CPU usage in percentage over time




Fig 2: 7zip memory usage in megabytes over time




Fig 3: 7zip power consumption in milliwatt over time

b) gzip - As observed in Fig. 3-6, gzip uses a consistent amount of memory (average 1752MB) with varying CPU (average 28%). Power consumption over times (Fig. 6) is observed to be varying slightly with an average of 18545 MW, which is lower than that of 7zip. However, a reduction




Fig 4: gzip CPU usage in percentage over time

in CPU usage does not achieve much savings in power consumption as one would expect.




Fig 5: gzip memory usage in megabytes over time




Fig 6: gzip power consumption in milliwatt over time

Since there is no fluctuation in memory usage, the cause of power consumption fluctuation over time is the CPU.

c) aio-stress – Fig. 7 shows aio-stress CPU usage, which highly fluctuates but is relatively low. Memory usage as shown in Fig. 8 is fairly consistent but still low as compared to that of 7zip and gzip. Fig. 9 shows an increasing power consumption over time but is much less as compared to 7zip and gzip. The power consumption ranges from 6026 MW to 11321 MW and an average of 9071 MW. This behavior shows that the hard disk is also an important component for determining the overall power consumption. However, it needs further investigation.



Fig 7: aio-stress CPU usage in percentage over time




Fig 8: aio-stress memory usage in megabytes over time




Fig 9: aio-stress power consumption in milliwatt over time

d) Loopback network - In this test, the CPU usage (Fig. 10) fluctuates with a repeating pattern – the peaks levels before they deep again. This is close to the behavior of power consumption (Fig. 12), whose peaks are however sharp. The high power consumption as compared to aio-stress can be attributed to consistently high memory usage whose average is 1744 MB.




Fig 10: Loopback CPU usage in percentage over time



Fig 11: Loopback CPU usage in megabytes over time




Fig 12: Loopback power consumption in milliwatt over time

From the tests above, it can be concluded that, the CPU is the main consumer of power but not the only one – the other components (memory, hard disk and network connection) also consume power.

Second, similar tasks can consume different amounts of power of depending on the resource demand that dominates the operation. gzip, unlike 7zip, achieved power savings by using constant amount of memory and varying CPU.

5.2 Workload Mix

Table 4 summarizes the performance and power consumption when running homogenous and heterogeneous VM workloads. It is observed that the performance of executing a particular workload is highest when single VMs are run in the physical server than when multiple VMs, running that

particular workload (homogeneous), are executed. For example CPU performance while processing 7zip workload achieves 3737 MIPS when only one VM is excited while it achieves 3099 MIPS and 2985 MIPS when heterogeneous and homogenous workloads are independently executed in 4 VMs respectively.

As more VMs are consolidated in a single server, more VM instances compete for the physical resources as well as the hypervisor capacity. The interference caused results to tasks running longer due to low throughput and thus consume more energy – energy is a product of power and time and given in Eq 1. [17]. This interference is more pronounced with

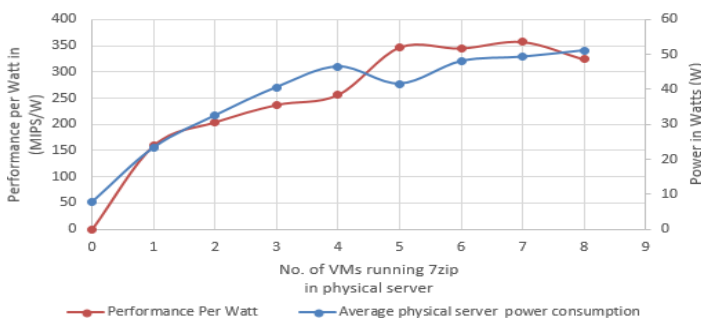
homogenous workloads because only one resource is used while the rest remain idle. This can be seen in Fig. 13 (b). A part from high power usage for homogeneous workloads due to tasks running longer, idle power for unutilized components goes to waste.

$$E = PT, \tag{1}$$

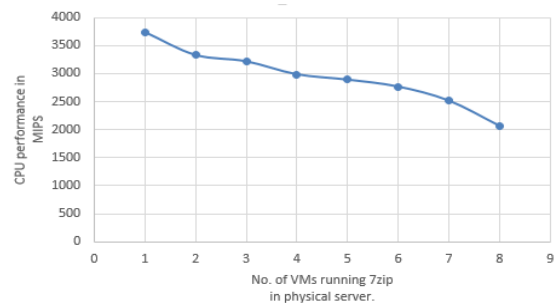
where P is average power consumption (in watts) and T is a time (in seconds) interval.

Table 4: Summary of Performance and Power Consumption for Homogenous and Heterogeneous Workloads.

	Average Performance				Average Power consumption (W)			
	CPU for 7zip (MIPS)	Memory for gzip (s)	Hard disk/IO for aio-stress (MB/s)	Network for network loopback (s)	CPU for 7zip	Memory for gzip	Hard disk/IO for aio-stress	Network for network loopback
Executing each of the 4 workloads in 1 VM running in physical server separately	3737	68	57.57	15.62	23.39	23.99	11.71	15.36
Executing heterogeneous workloads in 4 VMs	3099	160	28.99	16.36	31.74			
Executing homogenous workloads in 4 VMs	2985	501	10.35	16.98	46.48	19.00	10.36	15.39
Idle server (99.8% idle)	—	—	—	—	4.55			

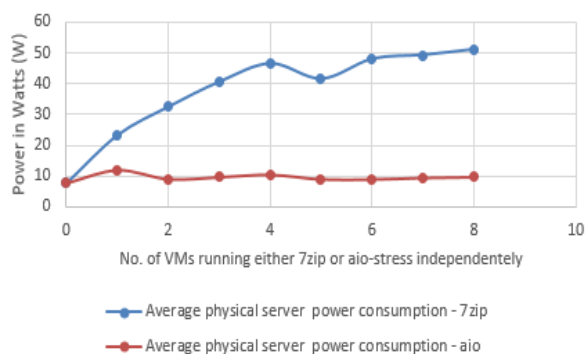


a) Energy and energy performance



(b) CPU performance

Fig 13: Effect of increasing the number of VMs on performance, power consumption and power performance.



a) Power consumption



b) Performance

Fig 14: A comparison between CPU and I/O's performance and power consumption when processing 7zip and aio-stress respectively at different levels of consolidation.

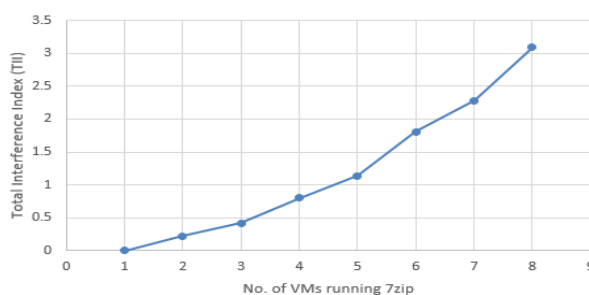


Fig 15. The Impact of Performance Interference at Different Levels of Contention on Homogenous Workloads

In Fig. 14 (a), results show that there is hardly an increase in power consumption when increasing the number of aio-stress VMs. Increase in number of VM running 7zip shows an increase in power consumption but as earlier observed, in Fig. 13 (a), CPU performance per watt increases, which is an advantage. In Fig. 14 (b), it is observed that an increase in the number of VMs causes noticeable performance degradation in I/O (worsening towards the 3rd VM) as compared to CPU performance. Thus, it can be concluded that consolidation limiting factor for CPU intensive workloads and disk intensive workloads is power consumption and performance respectively.

Further, from conclusions in [18], a value called Total Interference Index (TII) is calculated, which measures the degree of interference at different levels of resource contention on homogenous workload. TII is given as shown in Eq. 2.

$$TII = \sum_{i=1}^k \frac{P' - P}{P_i} \quad (2)$$

where k is the number of VMs co-located in the same physical server, P' is the performance of 1 VM when it runs alone in the physical server and P is the performance of the VM when it runs with other VMs in the physical server. Fig.15 shows that an increase in the number of VMs increases performance interference. It is also observed from the plot that as the number of VMs increases, the rate of growth of performance interference also increases.

In a multi-tenant virtualized environment, many different applications are host in the same server. Therefore, a mixture of workloads in not uncommon. An optimal workload mix needs to be determined that reduces energy consumption and retains high levels of performance. For this particular server, 6 VMs running 7zip and 2 VMs aio-stress achieves efficient use of resources, better performance and better performance per watt. However, this mix ratio can differ from server to server.

6. CONCLUSION

In this paper, it has been shown that, although the CPU is the main consumer of power in a server, other components such as memory, disk and network interface also consume power and needs to be considered before workload consolidation. To this end, it has been shown that if consolidation is not carefully carried out, it can lead to high power consumption in servers and even power wastage especially when homogenous workload is co-located. It has also been shown that optimal workload mix can achieve efficient power utilization and still preserve acceptable performance. In this paper, well known PTS has been as a source of workload. As future work, use of real cloud workloads can be used to propose techniques for workload characterization and scheduling instead of workload benchmarks.

7. REFERENCES

- [1] X. Chen, L. Rupprecht, R. Osman, P. Pietzuch, F. Franciosi and W. Knottenbelt, "CloudScope: Diagnosing and Managing Performance Interference in Multi-tenant Clouds," in 2015 IEEE 23rd International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems, 2015.
- [2] Industry Outlook, "Industry Outlook Data Center Energy Efficiency," 2014. [Online]. Available: <http://www.datacenterjournal.com/industry-outlook-data-center-energy-efficiency/>. [Accessed 10 October 2018].
- [3] M. D. Kenga, V. Omwenga and P. Ogao, "Energy Consumption in Cloud Computing Environments," in Pan African Conference on Science, Computing and Telecommunications (PACT) 2017, Nairobi, 2017.
- [4] J. Patel, V. Jindal, I.-L. Yen, F. Bastani, J. Xu and P. Garraghan, "Workload Estimation for Improving Resource

- Management Decisions in the Cloud," in 2015 IEEE Twelfth International Symposium on Autonomous Decentralized Systems, Taichung, Taiwan, 2015.
- [5] M. G. Xavier, K. J. Matteussi, F. Lorenzo and C. A. F. D. Rose, "Understanding performance interference in multi-tenant cloud databases and web applications," in 2016 IEEE International Conference on Big Data , Washington, DC, USA, 2016.
- [6] S. Shen, V. v. Beek and A. Iosup, "Statistical Characterization of Business-Critical Workloads Hosted in Cloud Datacenters," in 2015 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, Shenzhen, China, 2015.
- [7] C.-Z. Mar, S. Lavinia, A.-C. Orgerie and P. Guillaume, "An experiment-driven energy consumption model for virtual machine management systems," 2016.
- [8] S. Mohsen, S. Hadi and N. Mahsa, "Power-efficient distributed scheduling of virtual machines using workload-aware consolidation techniques," *The Journal of Supercomputing* , 2011.
- [9] M. Kurpicz, A. Sobe and P. Felber, "Using power measurements as a basis for workload placement in heterogeneous multi-cloud environments," in Proceedings of the 2nd International Workshop on CrossCloud Systems, New York, NY, USA , 2014.
- [10] F. P. Sareh, "Energy-Efficient Management of Resources in Enterprise and Container-based Clouds," The University of Melbourne , 2016.
- [11] J. Smith and I. Sommerville, "Workload Classification & Software Energy Measurement for Efficient Scheduling on Private Cloud Platforms," in Conference'10 University of St Andrews, 2011.
- [12] A. Al-Dulaimy, R. Zantout, W. Itani and A. Zekri, "Job Submission in the Cloud: Energy Aware Approaches," in Proceedings of the World Congress on Engineering and Computer Science, San Francisco, USA, 2016.
- [13] V. Vasudevan, D. Andersen, M. Kaminsky, L. Tan, J. Franklin and I. Moraru, "Energy-efficient cluster computing with FAWN: workloads and implications," in Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking, Passau, Germany, 2010.
- [14] A. Martin and V. Marangozova-Martin, "Automatic benchmark profiling through advanced workflow-based trace analysis," Wiley Online Library, 2018.
- [15] Phoronix Test Suite, "Phoronix Test Suite - Linux Testing and Benchmarking Platform, Automated Testing, Open-Source Benchmarking," Phoronix Media, 2018. [Online]. Available: <https://www.phoronix-test-suite.com/>. [Accessed 01 August 2018].
- [16] Phoronix Test Suite , "Phoronix Test Suite Suites," Phoronix Media, 2018. [Online]. Available: <https://openbenchmarking.org/suites/pts>. [Accessed 05 August 2018].
- [17] S. M. Ismael, Y. Renyu, X. Jie and W. Tianyu, "Improved Energy-Efficiency in Cloud Datacenters with Interference-Aware Virtual Machine Placement," in Autonomous Decentralized Systems (ISADS), 2013 IEEE Eleventh International Symposium, 2013.
- [18] C. Ltd, "powerstat - a tool to measure power consumption," 2018. [Online]. Available: <http://manpages.ubuntu.com/manpages/bionic/man8/powerstat.8.html>. [Accessed 05 August 2018].