

# Experimental Evaluation of Energy Savings of Virtual Machines in the Implementation of Cloud Computing

Roberto Rojas-Cessa, Sarh Pessima, and Tingting Tian

**Abstract**—Host virtualization has become of great interest as it is a technology that can enable the implementation of cloud computing. For this, it can offer several benefits, such as high levels of sharing of computing resources, ubiquitous availability, and savings on hardware investment. Although it has been claimed that host virtualization can also offer the benefit of energy savings, the levels of savings are unknown. In this paper, we present experimental evaluations of energy expenditure of virtual machines used for computation processes. We measure the energy spent by a number of virtual machines and compare it to the energy spent by a single and multiple (real) Linux hosts. The results show that the virtual machines deliver similar performance to processes run in real machines for small loads, and achieve significant energy savings for a modest number of users. Therefore, it is then expected that there is a number of VMs for what energy savings is optimal. The values also indicate that workstations optimized for virtualization can offer significant benefits.

**Index Terms**—Cloud computing, Energy, Green Cloud, Resource Sharing, Task Scheduling.

## I. INTRODUCTION

Cloud computing has raised as a new computer paradigm that brings greater flexibility to shared and high-performance computing. Cloud computing can be translated as computing services provided by network resources accessible through the Internet and that are highly accessible and flexible. Accessibility of these services is as robust as the available access to the Internet. With such accessibility, users can process most of their computing needs everywhere, every time, and through a wide variety of network access equipment, such as smart phones or mobile computers.

Examples of services provided by the cloud are large amounts of storage, extensive computing power via grid-computing-like mechanisms, software accessibility, web services, large databases, among others. A user who makes use of these services can use low-end computer equipment to access the cloud as the more intensive computations are performed in the cloud. Moreover, users mostly require the execution of low-intensive computations and network access in general.

Examples of the major motivations for a user to be subscribed to cloud services are savings against large investments for computer infrastructure, resource provisioning flexibility, open access (a large number of user can access the same

information simultaneously), and large storage capacities. This allows users to minimally invest in computing equipment and use cloud resources for as long as needed.

Large online business have foreseen the service opportunity that the cloud can potentially provide, and cloud services are being created quickly. Examples of the available cloud services are Google Docs [1], Amazon's Elastic Compute Cloud and Simple Storage [2], Microsoft's Azure [3], IBM Smart Businesses Services [4], among others.

Cloud services are usually implemented in one or more datacenters, where a large number of servers and storage units are interconnected through a large network infrastructure. Datacenters assign resources dynamically to the subscribed users via machine virtualization. This technology provides availability of resources flexible enough to adapt themselves to both the different levels of computing users' needs or to allocate sufficient resources to the varying number of users.

A challenge for dynamic allocation of resources is to keep datacenters infrastructure in stand-by with low maintenance costs while unused. A large part of maintenance costs is associated with power consumption of the large number of servers and telecommunications systems. In a datacenter, the telecommunications infrastructure, named communication links, switching and aggregation elements, consume about one third of the total power consumption, while the remaining two thirds are consumed by the servers and storage systems [5]. Unattended power consumptions can lead to high operational costs, and thermal hot spots that can potentially decrease the performance of the datacenter and cooling systems, or else to produce costly damages to the datacenter [6]. A comprehensive evaluation of the energy expenditure of the different equipment in a datacenter has been estimated [7]. This evaluation, however, considers the maximum energy expenditure of a datacenter as the maximum power of the different subsystems are considered.

Servers and other systems with low utilization levels need to be set to an operating mode such that power consumption and startup time are minimum. This requirement is hard to achieve as a system (e.g., a server), while in sleeping mode achieves low power consumption, it may take long time to startup [8], [9]. Turning on a large number of components can decrease the startup time but at the cost of small energy savings. To reduce number of startup servers, several strategies have been studied. Incorporation of network traffic management and server workload consolidation has been investigated [10]. In this approach, most traffic is routed into the servers that are already in use. Other schemes target an indirect detection

R. Rojas-Cessa is with the Networking Research Laboratory, Dept. of Elect. and Comp. Eng., New Jersey Institute of Technology, Newark, NJ 07102. Email: rojas@njit.edu.

S. Pessima and T. Tian are with the New Jersey Institute of Technology, Newark, NJ 07102.

of power dissipation, such as temperature distribution in the datacenter [6]. These schemes work as a closed-loop system where temperature is sensed to determine the cool zones where task can be assigned.

Cloud computing is supported in the core by machine virtualization technology. Machine virtualization is performed by a real machine, or workstation, running a substrate software supporting virtualization software. One or multiple operative systems (or copies of the same operative system) can be then installed on executed on top of those. Each of those copies running at the same time is called a virtual machine (VM). In cloud computing, one or multiple VMs are executed on demand. In this way, the capability of adding or removing VMs allows flexible allocation and distribution of loads on different parts of the datacenter hosting the real machines.

Information about how many virtual machines are suitable for energy savings of interest. Several works on machine virtualization have been reported [11]–[16]. However, the amount of energy per virtualized machines continues to be unknown.

In this paper, we study the energy consumption of virtualized machines under small computation loads to identify the numbers of virtual machines that allow an efficient utilization of hardware while providing energy savings. We present an experimental setup where a general purpose workstation is used to host multiple virtual machines. We measure the energy and processing time of virtual machines under small computational loads and compare the energy expenditure and efficiency to a system without machine virtualization, and to that of multiple physical machines.

The remaining of this paper is organized as follows. Section II introduces the methodology used for energy expenditure and computation performance. Section III presents the obtained results from the experimental measurements. Section V present our conclusions.

## II. METHODOLOGY FOR MEASUREMENT OF POWER CONSUMPTION

The experimental setup is shown in Figure 1. The setup consists of a workstation running Windows 7 as operating system (OS). On top of Windows 7, we use Oracle VM VirtualBox as virtualization software, and Ubuntu Linux operating system as a VM. Therefore, each VM is a virtual Linux machine. The specifications of the workstation are listed in Table I.

TABLE I  
EXPERIMENT SPECIFICATIONS

	Dell Optiplex 780
Processor	Intel(R) Core 2 Duo E8500
No. of processors	2
CPU speed	3.16 GHz
RAM	4 GB
Interface speed	10/100 Mb/s
Real operating system	MS 64-bit Windows 7
Virtual operating system	Ubuntu 10.10 (kernel 2.6.22-35-generic)
Virtualization software	Oracle VM Virtualbox

The power-measuring tools are Watts Up? PRO (WUP) meter, which is a consumer-level device that measures the total energy (and power) spent by commercial and residential electrical appliances, and the Microsoft (MS) Joulemeter program. The MS Joulemeter is a program developed to work




Fig. 1. Experiment setup.




Fig. 2. Energy consumption of the Windows workstation on idle state.

in combination with the WUP meter to measure the power usage of the various workstation’s resources (CPU, disk, and monitor) in the Windows environment. The WUP meter has a data interface used by the managing software to graph energy expenditure in function of time.

Before the VM experiments are performed, we measured the energy consumption of the substrate computer, which uses Windows 7. The energy consumed by the machine is presented in Figure 2. This figure shows the energy consumed by the Windows workstation in idle mode (this mode is defined as having no other software or process running on it, except for the OS).

The energy levels on this figure also define the minimum energy expenditure of a single workstation. The WUP-Joulemeter combination identifies the total energy of the workstation in idle state as 52 W and the power consumer by the CPU as 0 W. This equipment also shows the variations of the energy consumed by the maintenance programs of the OS (i.e., Windows 7) in the workstation, as shown in Figure 3. As shown in this figure, the power consumed by the workstation remains mostly constant except for a single change, caused probably by a subroutine in the OS.

However, the power consumed by the workstation goes through a series of oscillations at startup time, as shown in Figure 4.

### A. Benchmark software.

The total energy expenditure of the workstation and VMs is tested by executing a small program that increases only




Fig. 3. Example of the power consumption of the workstation in steady idle state (after programs have started up).




Fig. 4. Energy consumption of the Windows workstation during startup.

the load of the CPU. This program is a simple counter that attempts to use CPU resources only. However, different programs for testing energy can be used, depending on the objective peripherals or workstation resources of interest, but this is out of the scope of this paper. The program is based on a simple infinite loop, written in C++. The program was executed by a real and VMs for 10 minutes and the number of operations was recorded.

### III. EXPERIMENTAL MEASUREMENTS AND RESULTS

We measured the power consumed by single and multiple VMs and compared it to those of multiple workstations and multiple processes. The multiple workstations alternative recurs to using dedicated workstations while the multiple processes recurs to using shared workstations. These two alternatives are described as follows:

- **Multiple Workstations.** In this case, we measured the power consumed by a single workstation (real machine) on Linux OS, executing the benchmark program. Each benchmark program is executed in a single workstation. Multiple workstations is equivalent to as having multiple users, one per workstation, with dedicated machines.
- **Multiple Processes in a Single Linux Workstation.** In this case, a Linux workstation (real machine) is shared by multiple users, who also share the same instance of Linux.




Fig. 5. Power of a single VM processing the benchmark program.




Fig. 6. Power of two VMs processing of the benchmark program.

#### A. Power Consumption of Virtual Machines

As baseline, we created a Linux VM and measured the number of operations performed in 10 minutes. The number of operations determine when the test is completed for any number of VMs. We measured the power consumed by the single VM running the benchmark program as 71 watts, as shown in Figure 5.

The number of VMs was linearly increased in the following tests. These VMs also run the benchmark during the measurement of power. The measured power for two VMs, where each VM runs the test bench, is shown in Figure 6. The consumed power is about 81 watts. It seems that the first VM adds about 19 watts, while the second VM adds about 10 watts on top of the first VM. This may be caused by sharing the CPU (i.e., CPU re-use).

In succession, this process was repeated for three VMs and four VMs (i.e., a complete succession from 1VM to 4 VMs). Figure 7 shows the power consumption of the successive tests from 1 to 4 VMs, each one running the test bench. As the figure shows, there is a period of power oscillation during the startup of the VMs and the power consumption remains almost steady during the execution of the tesbench.

The figure shows that the addition of the third VM adds 8 watts (making a total of 88 watts), and the fourth VM adds about 3 more watts (for a total of 91 watts for all four VMs). The amount of power added by the third VM seems smaller than the other two as the load of the (real) workstation approaches to the maximum load, and therefore, the workstation uses most of the computing (CPU) resources. The reason why the fourth VM seems to consume less power




Fig. 7. Power profile for 1 VM, 2 VMs, 3 VMs, and 4 VMs, each executing the benchmark program.




Fig. 9. Number of operations for 1 to 6 virtual machines in the same test time.

than the others is because the maximum power consumption of the workstation is 91 watts. This is, no more power can be consumed by the workstation. In other words, without this cap, we may see a larger power consumption increase.

To demonstrate this, we increased the number of VMs in the system to 12. The power consumption of the successive increase of VMs is shown in Figure 8. As this figure shows, the maximum consumption power is 91watts, which is reached with four or more VMs. We then expect that 4 or more VMs would consume the maximum power but at the cost of processing performance.

To investigate the processing performance of these VMs, we evaluated the average number of operations a number of simultaneous VMs process. These average number of operations are shown in Figure 9. This figure shows that up to three simultaneous VMs can operate at the workstation performance and a larger number of VMs would have a performance penalty associated with the load of each task and the number of VMs. This results are consistent with the power consumption shown in Figure 8, which are also represented in Figure 10. We can observe that while the total consumption power is below the maximum power, the workstation can allocate resources to keep the computing performance to the peak. As power saturates, processing slows down. The slowing down of processing speed impacts the time it takes to process a

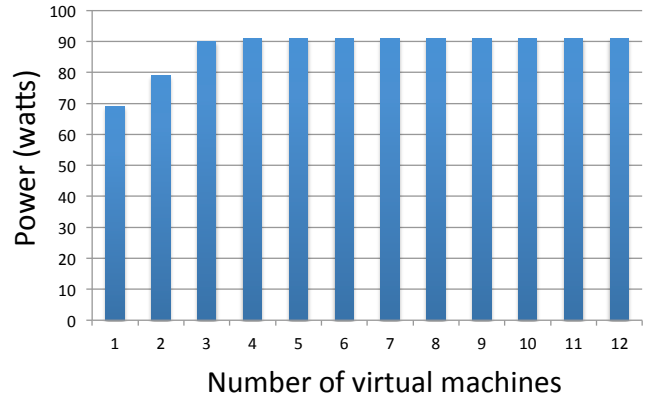


Fig. 10. Consumed power for 1 to 6 (simultaneous) virtual machines during the same test time.




Fig. 11. Processing time for 1 to 6 virtual machines to complete the execution of the benchmark program.

processing job, which in this case is the number of operations performed by having a single VM executing the tesbench for a period of 10 minutes (73,581,597 operations). We proceeded to measure the time it takes to process this number of operation with two and more simultaneous VMs. The obtained results are shown in Figure 11. The results show that three VMs can keep up with the processing time of a single VMs (i.e., there is no significant performance loss of performance). Furthermore, the results also show that the increase of processing time is slow (linear). One of the reasons for that is that the testbench program only loads the CPU of the workstations (there are no communications or hard-drive subroutines) and all the VMs execute the same task.

### B. Comparison of VM performance and other computing alternatives

There are many parameters to consider on whether cloud computing can provide a comparable performance to dedicated and real workstations. Herein, we focus on the power and processing performance. We compare these two parameters of VMs with two other possible solutions: multiple workstations and multiple processes. We start with the measurement of power consumption, where for multiple workstations, we expect to be high and linear as the power is directly proportional to the number of workstations (and instances of the benchmark




Fig. 8. Power profile for 12 (simultaneous) VMs, each executing the benchmark program.

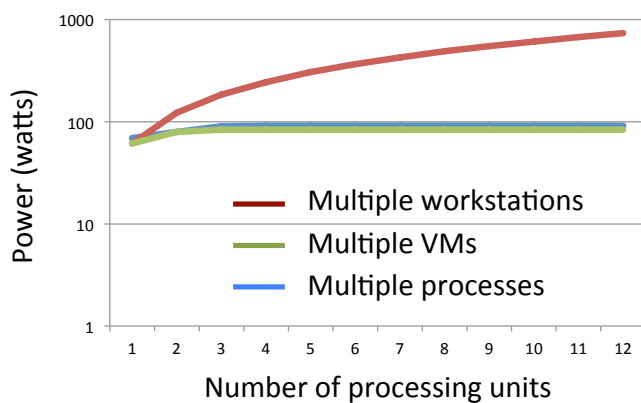


Fig. 12. Power comparison in the execution of 12 processes by a) virtual machines, b) multiple processes in a Linux machine, c) one Linux machine per process.

software). Figure 12 shows the power consumption of these three approaches. As the figure shows, the power consumed by multiple processes is very similar to that of multiple VMs. As the processing time requires running a large number of VMs longer periods of time, we measured the amount of energy if multiple VMs and compared it to that of multiple workstations and multiple processes. Figure 13 shows that the expanded time to finish the computing job used in this test is not long enough to jeopardize the energy savings VMs provide. We see that the running multiple processes requires similar amount of energy to multiple VMs as they deliver similar performance. These outcomes indicate that machine virtualization is effective to reduce power expenditure as multiple VMs can provide similar performance to running (simple) multiple processes and consume less energy than dedicated workstations. The performance of dedicated workstations may be still higher for the execution of multiple tasks, but at higher levels of power and cost.

#### IV. DISCUSSION

The virtualization environment provides an opportunity to improve the management and efficiency of utilizing computing resources. The resource sharing aspects of virtualization does,

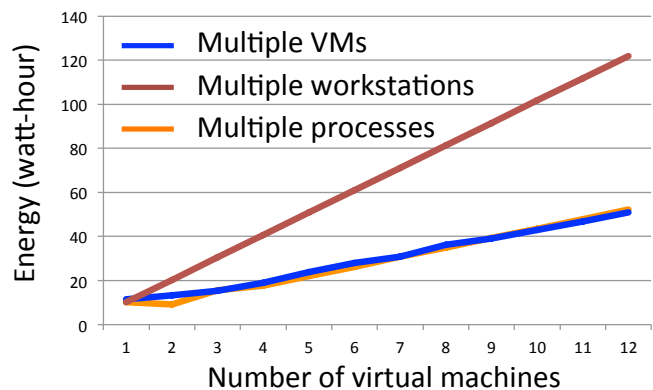


Fig. 13. Consumed energy to execute 12 processes by a) virtual machines, b) multiple processes in a Linux machine, c) one Linux machine per process.

however, create some new challenges from an application implementation and system support context. One such challenge, is the possibility that at high resource utilization levels individual programs are put in a prolonged wait state queuing for physical resources.

For example, we attempted to increase the number of VMs in the workstation under test, but the system (the combination of Windows 7 and Virtualbox) could not hold them, leading the OS into an execution crash. This indicates that there are other limits that need to be considered to find an optimum number of VMs per workstations. Nevertheless, we didn't observe a significant loss of performance for 12 simultaneous VMs, with however a benchmark program that produces a light computing load.

It is left as future research to investigate the power performance of VMs under benchmarks programs that use the CPU intensively and under communications subroutines.

#### V. CONCLUSIONS

This paper presented an experimental evaluation of energy consumption of multiple virtual machines running a simple computing process. The experiments were designed to evaluate the level of energy savings that machine virtualization can achieve. We tested multiple virtual machines created on a general purpose computer and measured the power consumption

and processing times. The results show that the performance degradation of multiple VMs provided a small performance degradation as the workstation reached the peak power consumption. We tested up to 12 VMs, being that number the largest number of VMs the workstation could create, and showed that machine virtualization provides significant power savings. We noted that there are two power performance regions: before a workstation reaches its maximum power consumption and after that. Larger energy savings may be achieved in the former region, while the amount of savings for the latter region depends on the number of VMs and the load assigned to each.

## REFERENCES

- [1] "Google docs," <http://docs.google.com>.
- [2] "Amazon web services," <http://aws.amazon.com>.
- [3] "Azure services platform," <http://www.microsoft.com/azure>.
- [4] "Ibm smart business services," <http://www.ibm.com/ibm/cloud>.
- [5] D. Kliazovich, P. Bouvry, and S. Khan, "Dens: Data center energy-efficient network-aware scheduling," in *Green Computing and Communications (GreenCom), 2010 IEEE/ACM Int'l Conference on Int'l Conference on Cyber, Physical and Social Computing (CPSCoM)*, Dec. 2010, pp. 69–75.
- [6] Q. Tang, S. Gupta, and G. Varsamopoulos, "Energy-efficient thermal-aware task scheduling for homogeneous high-performance computing data centers: A cyber-physical approach," vol. 19, no. 11, Nov. 2008, pp. 1458–1472.
- [7] J. Baliga, R. Ayre, K. Hinton, and R. Tucker, "Green cloud computing: Balancing energy in processing, storage, and transport," vol. 99, no. 1, Jan. 2011, pp. 149–167.
- [8] A. Bohra and V. Chaudhary, "Vmeter: Power modelling for virtualized clouds," in *Parallel Distributed Processing, Workshops and Phd Forum (IPDPSW), 2010 IEEE International Symposium on*, April 2010, pp. 1–8.
- [9] I. Goiri, F. Julia? and, R. Nou, J. Berral, J. Guitart, and J. Torres, "Energy-aware scheduling in virtualized datacenters," in *Cluster Computing (CLUSTER), 2010 IEEE International Conference on*, Sept. 2010, pp. 58–67.
- [10] P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan, "Energy aware network operations," in *INFOCOM Workshops 2009, IEEE*, April 2009, pp. 1–6.
- [11] M. Marcu and D. Tudor, "Power consumption measurements of virtual machines," in *Applied Computational Intelligence and Informatics (SACI), 2011 6th IEEE International Symposium on*, may 2011, pp. 445–449.
- [12] A. Beloglazov and R. Buyya, "Energy efficient resource management in virtualized cloud data centers," in *Cluster, Cloud and Grid Computing (CCGrid), 2010 10th IEEE/ACM International Conference on*, may 2010, pp. 826–831.
- [13] J. F. Brady, "Virtualization and cpu wait times in a linux guest environment," *J. Comp. Resource Mgmt*, vol. 116, 2005.
- [14] Y. Wang and X. Wang, "Power optimization with performance assurance for multi-tier applications in virtualized data centers," in *Parallel Processing Workshops (ICPPW), 2010 39th International Conference on*, sept. 2010, pp. 512–519.
- [15] A. Kansal, J. Liu, A. Singh, R. Nathuji, and T. Abdelzaher, "Semantic-less coordination of power management and application performance," *SIGOPS Oper. Syst. Rev.*, vol. 44, pp. 66–70, March 2010. [Online]. Available: <http://doi.acm.org/10.1145/1740390.1740406>
- [16] A. Kansal, F. Zhao, J. Liu, N. Kothari, and A. A. Bhattacharya, "Virtual machine power metering and provisioning," in *Proceedings of the 1st ACM symposium on Cloud computing*, ser. SoCC '10. New York, NY, USA: ACM, 2010, pp. 39–50. [Online]. Available: <http://doi.acm.org/10.1145/1807128.1807136>