Investigation into the Energy Cost of Live Migration of Virtual Machines

Kateryna Rybina, Waltenegus Dargie, Anja Strunk, and Alexander Schill
Chair of Computer Networks
Faculty of Computer Science
Technical University Dresden
01062 Dresden, Germany
Email: {kateryna.rybina, waltenegus.dargie, anja.strunk, alexander.schill}@tu-dresden.de

Abstract—One of the mechanisms to achieve energy efficiency in virtualized environments is to consolidate the workload (virtual machines) of underutilized servers and to switch-off these servers all together. Similarly, the workloads of overloaded servers can be distributed onto other servers for a load balancing reason. Central to this approach is the migration of virtual machines at runtime, which may introduce its own overhead in terms of energy consumption and service execution latency. This paper experimentally investigates the magnitude of this overhead. We use the Kernel-based Virtual Machine (KVM) hypervisor and a custom-made benchmark for our experiments. We will demonstrate that the workload of a virtual machine does not have any bearing on the power consumption of the destination server during migration but it has on the source server. Moreover, the available network bandwidth and the size of the virtual machine do indeed introduce a non-negligible energy overhead and migration latency on both the source and the destination server.

Index Terms—virtual machine, live virtual machine migration, migration time, migration cost, power consumption, energy overhead, workload types, energy-efficient computing.

I. Introduction

The energy consumption of the worldwide IT-infrastructure has been steadily increasing in the past decade [1], [2], but since the year 2007, there has been a noticeable slowdown [3]. One of the reasons for the slowdown is the introduction of resource virtualization and its adoption by the cloud computing community [3]. A resource virtualization refers to the abstraction of a hardware resource (a physical server) by a software entity (a virtual machine or, simply, a VM), so that two or more virtual machines encapsulating their own operating systems and the applications they host can run on one and the same physical machine, each virtual machine having its own execution environment which is effectively isolated from the computing environments of all the others. A low level middleware called a hypervisor makes the isolation of the virtual machines from the physical hardware as well as from each other and determines the exclusive use of hardware resources by each virtual machine.

This approach enables the dynamic consolidation of services (workload) in a data center without constraining the individual execution space of the services. The consolidation itself has several merits, for example, (1) underutilized servers can be switched-off after migrating their virtual machines elsewhere, thereby achieving energy efficiency [4]–[6]; (2) the workload of overloaded servers can be reduced by migrating some of their virtual machines to servers which have sufficient resources [7], [8]; and (3) a change of demand for computing resources can be dynamically accommodated.

In most cases, VM consolidation has to take place without violating service level agreements and transparent to the users of the virtual machines. In other words, migration should be carried out without actually "stopping" or "suspending" the virtual machines. In the literature, this process is termed as a live migration [9]. Because a live migration has to occur in parallel to the normal operation of all the virtual machines hosted by the physical machines, it introduces an additional execution overhead in terms of extra computing resources, an extra power consumption, and an execution latency. In the past this cost has been disregarded or was given little attention [10] [11]–[13]. However, there is no theoretical or experimental basis for this attitude. In this paper we aim to show our experimental observations concerning the cost of migrating virtual machines at run-time. We employed custom-made benchmarks having different CPU utilization statistics and ran the benchmarks inside virtual machines having different sizes. Moreover, we varied the available bandwidth during the migration of the virtual machines and studied how these parameters (VM size, CPU utilization statistics, and available bandwidth) influenced the cost of VM migration.

The rest of the paper is organized as follows. In Section II we explain the technical aspects of a live migration. In Section III, we introduce the benchmarks we generated to run on a VM and the experiment setting to carry our a live VM migration. In Section IV, we report the experiment results and discuss in detail the implications of the results. In Section V we give a compact summary of the related work and highlight how our contribution

complements it. Finally, in Section VI, we summarize our observations and outline some open research issues.

II. Virtual Machine Live Migration

A live migration enables a virtual machine to be physically moved from one physical host to another in a transparent fashion, while the virtual machine is still running. The current virtualization technology (based on hypervisors) does not use local discs to store VM images. Instead it requires a network attached storage (NAS) that can be accessible to all hosts and serve as hard drive for the virtual machines. By using a NAS, the process of live migration is limited to copying the in-memory state and the content of the CPU registers between the physical machines. One of the existing migration approach is the pre-copy [9] approach and it consists of the following phases (see Figure 1):

1) *Pre-Copy Phase*: This phase is executed in several rounds or iterations. The memory pages of the VM are copied page-wise to the destination host while the VM is continuously running on the source machine. In the first round all memory pages are copied to the destination machine, but as the copying process takes some amount of time, some of the memory pages on the source machine may be modified (dirtied). Thus, the dirty pages need to be re-sent in the subsequent round(s) to insure that the memory image on the destination machine is in sync with the one running on the source machine.

2) *Pre-Copy Termination*: The pre-copy phase may carry on indefinitely unless some stop condition is introduced. Typically one of the following thresholds is applied in practice: (a) the number of copying rounds exceeds a pre-defined threshold ($n > n_{th}$), (b) the total amount of memory that has already been transmitted exceeds a specified threshold ($mem_{mig} > mem_{th}$), or (c) the number of pages dirtied in the previous round falls below a specified threshold ($pg < pg_{th}$). Wnen either of the pre-defined stop conditions is met, the algorithm goes into the stop-and-copy stage.

3) *Stop-and-Copy Phase*: At this last phase the hypervisor suspends the VM on the source machine in order to prevent further modifications of the memory pages. Then the remaining dirty pages as well as the state of the CPU registers are copied to the destination host. After the migration process is completed (the copy of the VM on the destination machine is in sync with the one on the source host), the VM is resumed on the destination machine and discarded on the source machine.

   The duration of the stop-and-copy phase indicates the downtime of a migration process during which the VM is inactive on both the source and the destination machine. Thus, it is important to keep this time as small as possible.
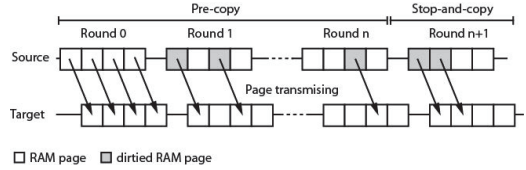

Fig. 1. Phases of the pre-copy live VM migration algorithm [14], [15].

III. Concept

The three most significant resources that influence the performance as well as the power consumption of a VM are the CPU, the memory, and the network bandwidth. Often it is to achieve the efficient utilization of these resources that VM consolidation takes place in data centres. At the time a decision for migrating a particular VM is made, the VM can be utilizing these resources at various intensities. For example, if the reason for migrating the VM is switching off an underutilized server, it is most likely that the VM does not utilize a considerable amount of resources. If, on the other hand, the reason is to relieve an overloaded server, then it is likely that the VM is utilizing a considerable amount of resources. Therefore, an investigation aiming to realistically quantify the cost of VM live migration should take these aspects into account.

A. Workload generation

The workload of most of the Internet applications hosted by virtual machines (such as video-on-demand servers) cannot be known except in a probabilistic sense [16]. In order to emulate these applications, we generated a custom-made, stochastic, CPU-bound benchmark which computes a discrete convolution of two random variables. The operation consists of integer, float point, and shift operations and executes inside an infinite loop. While the loop operation runs, it utilizes 100% of the CPU, but when the loop operation is stopped, the benchmark has 0% CPU utilization.

In order to determine the CPU utilization statistic of the benchmark, we divided time into a sequence of one-second none overlapping windows. Then we generated a set of random numbers (50,000 samples per distribution) in the interval [0, 100] (corresponding to 0 to 100% CPU utilization) using the `runif` `rnorm` `rexp` and `rpois` functions of the `R` statistical tool[1]. The random numbers correspond to a uniform, a normal, an exponential, and a Poisson distribution, respectively. For each time window, we picked out one of these random numbers and determined the portion of the one-second window the CPU is fully utilized by the convolution operation. Consequently, the benchmark generates workloads that have uniform distribution: $U(0, 100)$, normal distribution with mean utilization of 40% and a standard deviation of

[1]http://www.r-project.org/.

9%: $N(40, 9)$, an exponential distribution with an average utilization of 3%: $Exp(1/3)$, and a Poisson distribution with an average utilization of 60%: $Pois(60)$.

Moreover, for each experiment, we varied the size of the VM (i.e., its memory content, which, for our case, is a static quantity) being migrated and the available network bandwidth of both the source and the destination servers. The aim of our experiment is to investigate:

1) The statistics (distribution) of the power consumed by the source and the destination servers during a VM live migration; and,
2) The expected duration (latency) of the migration of an active VM having a varying resource utilization statistics.

As a consequence of (1) and (2), we will be able to quantify the expected energy overhead of the system under consideration due to a VM live migration, which can be expressed as:

$$E_{overhead} = (E\left[P_{mig}\right] - E\left[P_{bmig}\right]) \times E\left[t_{mig}\right] \quad (1)$$

where the $E[\cdot]$ operation refers to the expected value of a random variable, $E\left[P_{mig}\right]$ refers to the power consumption of the system during migration, $E\left[P_{bmig}\right]$ refers to the average power consumption of the system before a migration took place, and $E\left[t_{mig}\right]$ refers to the migration time (duration). Under the term 'system' we refer here to the source and destination machines.

Investigating the energy overhead of live migration on the network devices (such as the switch) falls out of the scope of this experiment.
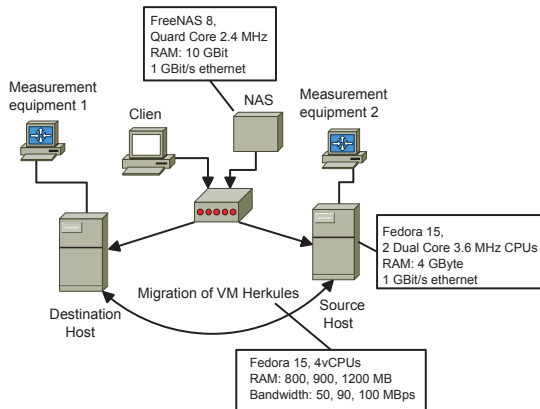

Fig. 2. Experimental setup for live migration of virtual machines.

B. Experimental settings

Our experimental setup consists of two homogeneous servers (the source and the destination servers), a client server to trigger a migration, a network attached storage (NAS), and two power analyzers as depicted in Figure 2.

Each server consists of two Intel 15-680 Dual Core 3.6 MHz processors, 4 GB DDR3-1333 SDRAM and a 1

VM size (MB)	800, 900, 1200
BW (MBps)	50, 90, 100
Workload distribution	$Pois(60), U(0, 100), N(40, 9), Exp(1/3)$

Gbit/s Ethernet NIC. The servers are connected with each other via a 1 Gbit/s switch. The NAS system employs an Intel Xeon E5620 Quad-Core 2.4 MHz processor, 10 GB DDR3-1333 SDRAM memory, and 1 Gbit/s Ethernet NIC.

On all the servers runs Fedora 15[2] (Linux kernel v. 2.6.38, x86_64). We use KVM[3] as our hypervisor and libvirt[4], as a toolkit managing the interaction between the hypervisor and the operating system during migration. Furthermore, we employed FreeNAS[5], an open source operating system (v. 8.0.1, AMD 64), as the Network Attached Storage. All the virtual machines we migrated were allocated 4 virtual CPUs and run Fedora 15 as their operating system, but we varied their size (memory content) and the available network bandwidth during the experiment. The live migration was triggered by the client server. The source and the destination hosts were managed by this same client server.

Two identical power analyzers (Yokogawa WT210 digital power analyzers) were used to measure the overall AC power consumption of the source and the destination servers under test. The devices can measure DC and AC power consumption at the rate of 10 Hz and a DC current between 15 μA and 26 A with an accuracy of 0.1%.

IV. EXPERIMENTS AND RESULTS

A. Migration settings

We ran the custom-made benchmark on each of the four processor cores of the virtual machines under different configurations and bandwidth constraints (see Table I) and migrated each virtual machine in isolation from the source server to the destination server. During each migration, we measured and recorded the power consumption and the resources utilization (CPU and memory) of both servers. For each configuration the migration was repeated eleven times to ensure that the observations we made were statistically consistent. Altogether we conducted 36 distinct experiments (four different workload types or utilization statistics, three different VM sizes, and three different available bandwidths). In all the experiments the servers were time synchronized so that we can accurately determine the beginning and the end of a VM migration.

B. Power consumption during migration

To investigate the additional power consumption introduced by a VM migration, we first measured the power

[2]Fedora 15. http://fedoraproject.org/.
[3]http://www.linux-kvm.org/page/Main_Page.
[4]Libvirt: The virtualization API. http://libvirt.org/.
[5]Freenas: Storage for open source. http://www.freenas.org/.

consumption of the source server while the VM ran on it, performing its normal operation. Similarly, we measured the idle power consumption of the destination server – throughout the experiment, the destination server was not hosting any additional virtual machine prior to a migration. Then we measured the power consumption of both servers during the VM migration and compared the cumulative distribution functions of the power consumptions.

In the beginning, we fixed the VM memory size to 1200 MB and run the benchmark with the exponentially distributed workload which required on average only 3% of the CPU utilization. Then we migrated the VM by considering two different values for the available network bandwidth: 50 and 100 MBps. The aim was to examine the influence of the available network bandwidth on the power consumption of the server during a VM migration. It should be noted that a migration process is predominantly an IO-intensive process.

Figure 3 displays the cumulative distribution functions $(F_p(p))$ of the power consumption of both servers before and during migration. In both servers, the power consumption visibly increased during migration and, as can be expected, the power consumption of both servers increased as the available bandwidth for migration increased. The migration overhead on the source machine was 9 W on average when the network bandwidth was 50 MBps and 17 W when it was 100 MBps. Similarly, the overhead of migration on the destination machine for the same set up was on average 11 W and 12 W, respectively.

In the second round of experiments, we fixed the VM size and varied the available network bandwidth as before, but we ran the benchmark with the normal distribution and the Poisson distribution workloads. The normally distributed workload had on average 40% CPU utilization while the Poisson workload had on average 60% CPU utilization. Figure 4 and 5 show the corresponding CDF of the power consumption of the source and the destination servers. As before, the power consumption of both servers increased during migration and the increment was larger when the available network bandwidth was larger.

However, the power consumption of the source server increased significantly in all the three cases (before migration as well as when the available network bandwidth was 50 and 100 MBps) when compared with the situation in which the exponential workload was executing. This is reasonable, since the CPU was busy much of the time when executing the Poisson and the normal distribution workloads. In contrast, the CDFs of the power consumption of the destination server remained similar to the previous case. This is because, during the VM migration, the VM was still executing from the source machine and there was no additional workload executing on the destination machine other than writing to memory the content of the migrated VM. This task is similar for all the workload types. To highlight this point, we put side by side in
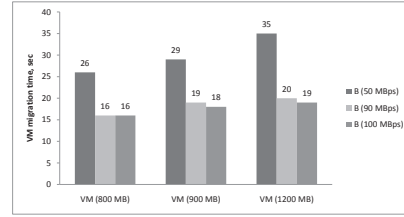

Fig. 8. The average migration time of a VM as a function of a VM's size and available network bandwidth. The type of workload executing on the virtual machines during migration was a normally distributed workload: $N(40, 9)$

Figure 6 the CDFs of the power consumption of both servers when the migrated VMs were executing the four different types of workloads. As can be seen in the figures, the power consumption of the source server during the VM migration experienced considerable variations depending on the type of workload a VM was executing, however, the CDFs of the power consumption of the destination server were remarkably similar for all the workload types. Consequently, as far as the destination server is considered, the migration cost can be taken as a fixed cost.

Unlike, the network bandwidth, the size of a VM did not influence the power consumption of neither the source nor the destination server during a VM migration. Our observation confirms an earlier observation made by Strunk and Dargie [15]. This assertion can be visualized in Figure 7 in which the workload type and the available network bandwidth were fixed but the size of the virtual machines is varied (800 MB, 900 MB, and 1200 MB). However, caution is needed to interpret this assertion, since the VM size has some bearing on the migration time which in turn affects the overall energy overhead (to be discussed shortly).

C. Migration time

The type of workload a VM executes does not affect the migration time. However, the workloads we considered in our experiments are CPU-intensive workloads and do not create contention on IO resources. Our observation may change for IO-intensive workloads, since these workloads should share the network bandwidth with the VM migrating operation, they will certainly influence the migration time.

In agreement with similar studies (for example with [15]), our investigation reveals that the VM size and the available network bandwidth influence the migration time, but of the two, the available bandwidth had a stronger influence. Migration time decreases at a faster rate with a higher network bandwidth but decreases slowly with a smaller VM size. This is displayed in Figure 8.

D. Energy overhead caused by VM migration

Independently considering the contribution of the three parameters (VM size, network bandwidth, and workload type) to the migration latency and the power consumption
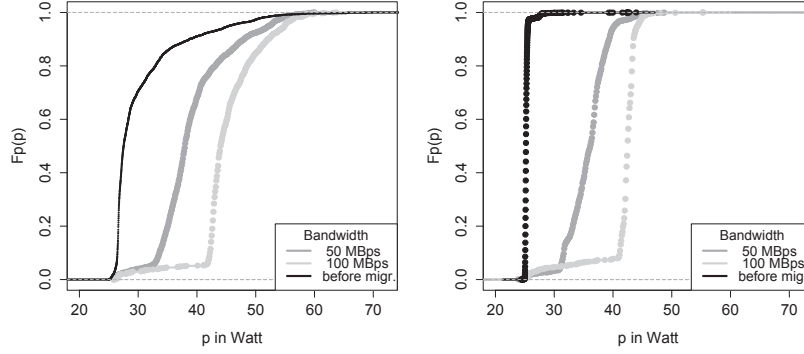
Fig. 3. The CDF ($F_P(p)$) of the power consumption of the source (left) and the destination (right) servers during the VM migration. The workload running on Hercules was an exponentially distributed workload with parameter: $Exp(1/3)$, VM size = 1200 MB.
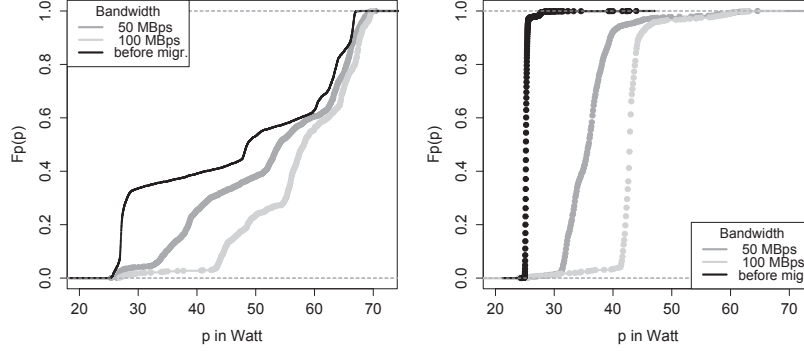

Fig. 4. The CDF ($F_P(p)$) of the power consumption of the source (left) and the destination (right) servers during the VM migration. The workload running on the virtual machine was a normally distributed workload with parameters $N(40, 9)$, VM size = 1200 MB.
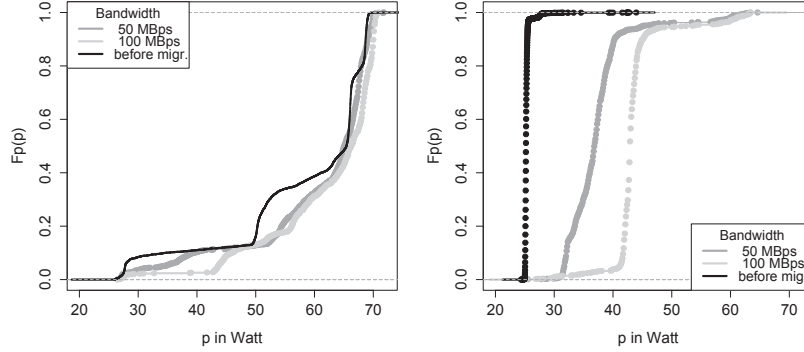

Fig. 5. The CDF ($F_P(p)$) of the power consumption of the source (left) and the destination (right) servers during the VM migration. The workload running on the virtual machine was a Poisson distributed workload with parameters $Pois(60)$, VM size = 1200 MB.

of the source and the destination servers enables to prioritize decisions pertaining to VM consolidation, but does not tell the whole story about the cost of VM migration. For example, the assertion that the VM size does not have any bearing on the power consumption of any of the servers participating in a migration process can be misleading, because a big VM requires more time to migrate than a small VM. In this case, considering the energy cost of a VM's migration is more appropriate, because as a product of the migration time and the power consumption, the energy overhead provides a more complete information.

To calculate the energy overhead, we first observed for a long time the power consumptions of both the source and the destination servers before migration to obtain their steady state power consumptions. From this observation, we constructed their cumulative distribution functions. For the source server, this refers to the statistics of the power consumption during the normal operation of the VM whereas for the destination server, this refers to the statistics of the idle power consumption of the server. Secondly, we measured the power consumptions of both servers during the VM migration and from this measurement, we established the distribution functions of the power consumptions. Thirdly, we subtracted the average
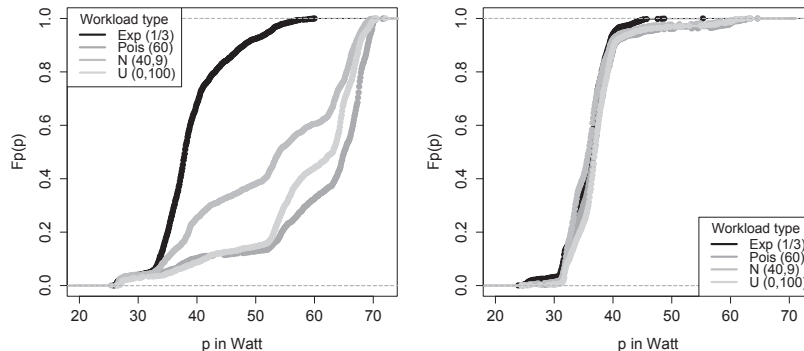
Fig. 6. The CDF $(F_P(p))$ of the power consumption of the source (left) and the destination (right) servers during a VM migration. The virtual machines were executing four different types of workloads. VM size = 1200 MB and network bandwidth = 50 MBps.
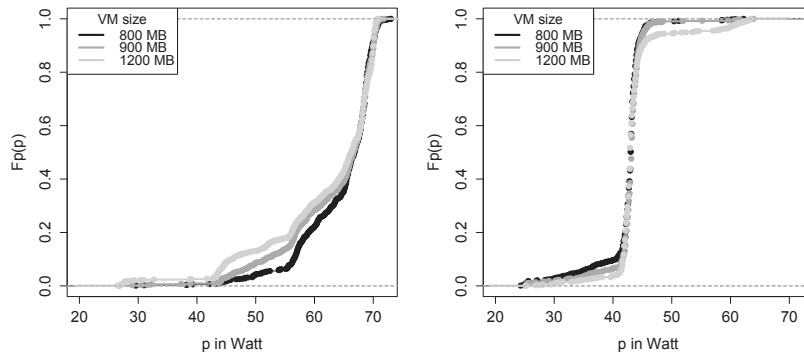

Fig. 7. The CDF $(F_P(p))$ of the power consumption of the source (left) and the destination (right) servers during the migration of four different virtual machines with the same workload type but different VM sizes. The workload running on the virtual machines was a Poisson distributed workload $(Pois(60))$ and the available bandwidth for migration was 50 MBps.

steady state power consumptions from the average power consumption of the servers during migration. Finally, using Equation 2, we computed the overall energy overhead:

$$E_{Ooverhead} = (E\,[P_{smig}] - E\,[P_{sbmig}]) \times E\,[t_{mig}] + (E\,[P_{dmig}] - E\,[P_{dbmig}]) \times E\,[t_{mig}] \quad (2)$$

where $E\,[P_{smig}]$ and $E\,[P_{dmig}]$ refer to the power consumption of the source and the destination servers, respectively, during migration and $E\,[P_{sbmig}]$ and $E\,[P_{dbmig}]$ refer to the average power consumption of the source and the destination servers, respectively, before migration took place, and $E\,[t_{mig}]$ refers to the migration time (duration).

Figure 9 and 10 summarize the energy overhead of a VM migration for the different configurations. Since the energy overhead is a relative term, its significance diminishes as the power consumption of the VM during a normal operation becomes significantly large. That is why the overhead during the migration of the uniform and the Poisson workloads in Figure 9 appear smaller than the others, because these two workloads consume a large amount of power during their normal operation and in comparison, the additional power consumed during the migration of the virtual machines was small. Of all the parameters, the network bandwidth (when small, 50 MBps)

resulted in the highest energy overhead. This is because migration is an IO-intensive operation. This observation suggests that in a server that predominantly hosts virtual machines on which IO-intensive applications run (such as video streaming applications), the cost of migration can be high.
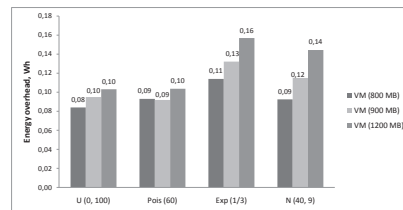

Fig. 9. The energy overhead of a VM migration as a function of the VM size and different workload types. The available network bandwidth = 100 MBps.

V. Related Work

Live migration has been investigated in various contexts [4], [17]–[22]. Most of the existing or proposed approaches focus on the performance of live migration and measure migration time and down time, under different conditions. Work that explicitly investigates the costs of migration is rare.
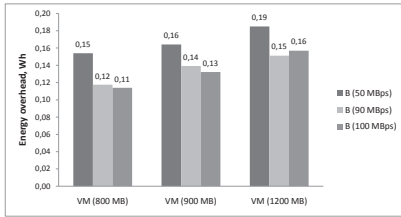
Fig. 10. The energy overhead of a VM migration as a function of the available network bandwidth and the VM size. The workload of the virtual machines for all the test cases was an exponentially distributed workload: $Exp(1/3)$.

We classify the costs of virtual machine live migration into performance loss and energy overhead. During live migration, a hypervisor labels all memory pages occupied by a VM as read-only in order to facilitate migration. All requests to overwrite some of these pages will raise exception indicating that the memory pages are faulted. This slows down the VM's response to requests and reduces its throughput [9]. Additional performance loss arises due to resource bottlenecks. The pre-copy and stop-and-copy processes require additional resources, particularly, network bandwidth and some CPU cycles. Since co-located virtual machines must not be affected by the migration, there may be a resource deficiency for the VM being migrated [23].

Kuno et al. analyse the processing speed of CPU-intensive and the reading speed of IO-intensive (disk) workloads. The authors find out that the performance of CPU-intensive workloads decreases by 15% whereas the reading speed diminishes by 10% [24]. In [9], the authors show that the transmission rate of an Apache Web Server slows down by 12 to 20%. Performance loss during live VM migration may be problematic in systems where the response time constitutes a strict performance guarantee. For example, Voorslys et al. demonstrate that 90% of the download time of home pages created with Web 2.0 technologies (PHP, Ruby on Rails, J2EE) may not be accessible during live migration [25].

The additional resource utilization cost during live migration creates an energy overhead. However, current live migration scenarios do not consider this energy cost. For example, Mistral proposes a framework to optimize the power consumption of cloud systems and uses live migration as a mechanism to consolidate virtual servers and switch off underutilized physical machines. The framework does not take the migration's additional power consumption into account [26]. This idea is shared by similar approaches which investigate service consolidation and dynamic power management in data centers [13], [27]–[29].

This paper complements studies focusing on the performance cost of migration and addresses the energy cost during the live migration of virtual machines. We showed how this cost varied as the size of the virtual machine and the available network bandwidth as well as the workload type varied.

VI. Conclusion and future work

In this paper we investigated the cost of VM migration under different operation conditions. We studied how the migration latency, the power consumption of the source and the destination server, and the overall energy overhead are affected by the workload type of a VM, the size of a VM, and the available bandwidth. We considered a discrete convolution operation as the workload of all the virtual machines we migrated. We designed the benchmark such that the CPU utilized by the convolution operation assumed a uniform, an exponential, a Poisson, and a normal distribution. We considered three VM sizes (800 MB, 900 MB, 1200 MB) and set up the available network bandwidth at 50, 90, and 100 MBps. In total, we set up 36 distinct configurations and for each configuration 11 experiments were conducted.

Our observations can be summarized as follows:

1) During a VM migration the power consumption of both the source and the destination servers is higher than the power consumption of the servers before a migration was carried out. However, the power consumption of the source server was higher than the power consumption of the destination server in all the experiments.
2) The power consumption of the source server during a migration was affected by the type of workload the VM was running. The power consumption of the destination server during a migration did not depend on the workload type, since the VM was executing on the source machine during the migration.
3) The energy overhead of a VM migration cannot be considered negligible. This is true regardless of the type of workload the VM was hosting.
4) The VM migration time was not influenced by the type of workload running on the VM and was approximately the same for all types of workloads as long as the VM size and the network bandwidth were the same.
5) The VM migration time was affected by the size of the VM and the available network bandwidth, particularly, by the network bandwidth.
6) The energy overhead of a live VM migration significantly decreases with a higher network bandwidth.

In this study we considered one type of benchmark and the VM migration was tested in isolation. In reality, most servers host heterogeneous applications and multiple virtual machines at the same time. In future we will take this fact into account to investigate the cost of VM migration.

REFERENCES

[1] J. G. Koomey, "Worldwide electricity used in data centers," *Environmental Research Letters*, vol. 3, no. 3, p. 034008, 2008. [Online]. Available: http://stacks.iop.org/1748-9326/3/i=3/a=034008

[2] W. Dargie, A. Strunk, and A. Schill, "Energy-aware service execution," in *The 36th Annual IEEE Conference on Local Computer Networks*, 2011.

[3] J. G. Koomey, "Growth in data center electricity use 2005 to 2010," Tech. Rep., July 2011. [Online]. Available: http://www.analyticspress.com/datacenters.html

[4] M. Andreolini, S. Casolari, M. Colajanni, and M. Messori, "Dynamic load management of virtual machines in cloud architectures." in *CloudComp*, ser. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, D. R. Avresky, M. Diaz, A. Bode, B. Ciciani, and E. Dekel, Eds., vol. 34. Springer, 2009, pp. 201–214. [Online]. Available: http://dblp.uni-trier.de/db/conf/cloudcomp/cloudcomp2009.html AndreoliniCCM09

[5] H. Mi, H. Wang, G. Yin, Y. Zhou, D. Shi, and L. Yuan, "Online self-reconfiguration with performance guarantee for energy-efficient large-scale cloud computing data centers," in *Services Computing (SCC), 2010 IEEE International Conference on*, july 2010, pp. 514 –521.

[6] B. Li, J. Li, J. Huai, T. Wo, Q. Li, and L. Zhong, "Enacloud: An energy-saving application live placement approach for cloud computing environments," *2012 IEEE Fifth International Conference on Cloud Computing*, vol. 0, pp. 17–24, 2009.

[7] T. Wood, P. Shenoy, A. Venkataramani, and M. Yousif, "Black-box and gray-box strategies for virtual machine migration," in *Proceedings of the 4th USENIX conference on Networked systems design anbd implementation*, ser. NSDI'07. Berkeley, CA, USA: USENIX Association, 2007, pp. 17–17. [Online]. Available: http://dl.acm.org/citation.cfm?id=1973430.1973447

[8] B. Gerofi, H. Fujita, and Y. Ishikawa, "An efficient process live migration mechanism for load balanced distributed virtual environments," in *Cluster Computing (CLUSTER), 2010 IEEE International Conference on*, sept. 2010, pp. 197 –206.

[9] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield, "Live migration of virtual machines," in *Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation - Volume 2*, ser. NSDI'05. Berkeley, CA, USA: USENIX Association, 2005, pp. 273–286. [Online]. Available: http://dl.acm.org/citation.cfm?id=1251203.1251223

[10] B. Li, J. Li, J. Huai, T. Wo, Q. Li, and L. Zhong, "Enacloud: An energy-saving application live placement approach for cloud computing environments," in *Proceedings of the 2009 IEEE International Conference on Cloud Computing*, ser. CLOUD '09. Washington, DC, USA: IEEE Computer Society, 2009, pp. 17–24. [Online]. Available: http://dx.doi.org/10.1109/CLOUD.2009.72

[11] *Efficient Resource Management for Cloud Computing Environments*. Chicago, IL: IEEE, 08/2010 2010. [Online]. Available: http://grids.ucs.indiana.edu/ptliupages/publications/10-ajy-igcc.pdf

[12] P. Padala, K. G. Shin, X. Zhu, M. Uysal, Z. Wang, S. Singhal, A. Merchant, and K. Salem, "Adaptive control of virtualized resources in utility computing environments," in *Proceedings of the 2nd ACM SIGOPS/EuroSys European Conference on Computer Systems 2007*, ser. EuroSys '07. New York, NY, USA: ACM, 2007, pp. 289–302. [Online]. Available: http://doi.acm.org/10.1145/1272996.1273026

[13] L. Liu, H. Wang, X. Liu, X. Jin, W. B. He, Q. B. Wang, and Y. Chen, "Greencloud: a new architecture for green data center," in *Proceedings of the 6th international conference industry session on Autonomic computing and communications industry session*, ser. ICAC-INDST '09. New York, NY, USA: ACM, 2009, pp. 29–38. [Online]. Available: http://doi.acm.org/10.1145/1555312.1555319

[14] H. Liu, C.-Z. Xu, H. Jin, J. Gong, and X. Liao, "Performance and energy modeling for live migration of virtual machines," in *Proceedings of the 20th international symposium on High performance distributed computing*, ser. HPDC '11. New York, NY, USA: ACM, 2011, pp. 171–182. [Online]. Available: http://doi.acm.org/10.1145/1996130.1996154

[15] A. Strunk and W. Dargie, "Does live migration of virtual machines cost energy?" in *The 27th IEEE International Conference on Advanced Information Networking and Applications (AINA-2013)*, 2013.

[16] W. Dargie, "Analysis of the power consumption of a multimedia server under different dvfs policies," *2012 IEEE Fifth International Conference on Cloud Computing*, vol. 0, pp. 779–785, 2012.

[17] A. Strunk, "Costs of virtual machine live migration: A survey," in *Services (SERVICES), 2012 IEEE Eighth World Congress on*, june 2012, pp. 323 –329.

[18] S. Akoush, R. Sohan, A. Rice, A. Moore, and A. Hopper, "Predicting the performance of virtual machine migration," in *Modeling, Analysis Simulation of Computer and Telecommunication Systems (MASCOTS), 2010 IEEE International Symposium on*, aug. 2010, pp. 37 –46.

[19] T. Imada, M. Sato, and H. Kimura, "Power and qos performance characteristics of virtualized servers," in *Grid Computing, 2009 10th IEEE/ACM International Conference on*, oct. 2009, pp. 232 –240.

[20] Y. Wu and M. Zhao, "Performance modeling of virtual machine live migration," in *Cloud Computing (CLOUD), 2011 IEEE International Conference on*, july 2011, pp. 492 –499.

[21] Y. Kuno, K. Nii, and S. Yamaguchi, "A study on performance of processes in migrating virtual machines," in *Autonomous Decentralized Systems (ISADS), 2011 10th International Symposium on*, march 2011, pp. 567 –572.

[22] S. Kikuchi and Y. Matsumoto, "Performance modeling of concurrent live migration operations in cloud computing systems using prism probabilistic model checker," in *Cloud Computing (CLOUD), 2011 IEEE International Conference on*, july 2011, pp. 49 –56.

[23] A. Verma, G. Kumar, and R. Koller, "The cost of reconfiguration in a cloud," in *Proceedings of the 11th International Middleware Conference Industrial track*, ser. Middleware Industrial Track '10. New York, NY, USA: ACM, 2010, pp. 11–16. [Online]. Available: http://doi.acm.org/10.1145/1891719.1891721

[24] Y. Kuno, K. Nii, and S. Yamaguchi, "A study on performance of processes in migrating virtual machines," in *Autonomous Decentralized Systems (ISADS), 2011 10th International Symposium on*, march 2011, pp. 567 –572.

[25] W. Voorsluys, J. Broberg, S. Venugopal, and R. Buyya, "Cost of virtual machine live migration in clouds: A performance evaluation," in *Proceedings of the 1st International Conference on Cloud Computing*, ser. CloudCom '09. Berlin, Heidelberg: Springer-Verlag, 2009, pp. 254–265.

[26] G. Jung, M. Hiltunen, K. Joshi, R. Schlichting, and C. Pu, "Mistral: Dynamically managing power, performance, and adaptation cost in cloud infrastructures," in *Distributed Computing Systems (ICDCS), 2010 IEEE 30th International Conference on*, june 2010, pp. 62 –73.

[27] T. V. T. Duy, Y. Sato, and Y. Inoguchi, "Performance evaluation of a green scheduling algorithm for energy savings in cloud computing," in *Parallel Distributed Processing, Workshops and Phd Forum (IPDPSW), 2010 IEEE International Symposium on*, april 2010, pp. 1 –8.

[28] G. Dhiman, G. Marchetti, and T. Rosing, "vgreen: a system for energy efficient computing in virtualized environments," in *Proceedings of the 14th ACM/IEEE international symposium on Low power electronics and design*, ser. ISLPED '09. New York, NY, USA: ACM, 2009, pp. 243–248. [Online]. Available: http://doi.acm.org/10.1145/1594233.1594292

[29] B. Li, J. Li, J. Huai, T. Wo, Q. Li, and L. Zhong, "Enacloud: An energy-saving application live placement approach for cloud computing environments," *2012 IEEE Fifth International Conference on Cloud Computing*, vol. 0, pp. 17–24, 2009.