2013 IEEE Sixth International Conference on Cloud Computing

A Lightweight Model for Estimating Energy Cost of Live Migration of Virtual Machines

Anja Strunk
Chair of Computer Networks
Faculty of Computer Science
Technical University of Dresden
01062 Dresden, Germany
Email: anja.strunk@tu-dresden.de

***Abstract*—Live migration, the process of moving a virtual machine (VM) interruption-free between physical hosts is a core concept in modern data centers. Power management strategies use live migration to consolidate services in a cluster environment and switch off underutilized machines to save power. However, most migration models do not consider the energy cost of migration. In a previous study we showed that live migration entails an energy overhead and the size of this overhead varies with the RAM size of the virtual machine and the available network bandwidth. This paper extends our previous work and proposes a lightweight mathematical model to estimate the energy cost of live migration of an idle virtual machine quantitatively. A series of experiments were conducted on KVM to profile the migration time and the power consumption during live migration. Based on these data we derived an energy cost model that predicts the energy overhead of live migration of virtual machines with an accuracy of higher than 90%.**

I. Introduction

Virtualization is a technique that enables several operating systems to run simultaneously on a single physical machine. It has become a core aspect in modern servers and data centers due to several advantages, such as flexible and efficient sharing of resources, fault tolerance, portability, and cost efficiency [1]. In a virtualized environment, virtual machines (VM) acting like real physical machines can run in parallel and in isolation from each other and yet sharing the same physical resources. A low level middleware called a hypervisor abstracts these virtual machines from the physical hardware and determines the exclusive use of resources by each VM.

One of the key features of virtualization is the live migration of virtual machines which enables an active (executing) VM to be moved from one physical machine to another in a transparent fashion [2]. This key feature has become a significant tool for a variety of scenarios. Some of which include:

- *Load balancing* [3], [4]. The aim is to adjust a virtual machine placement in order to achieve critical business goals, such as high throughput.
- *Transparent IT maintenance*. Administrators can transparently move virtual machines to free and shut down hosts for maintenance.
- *Power management* [5]–[7]. The aim is to consolidate virtual machines through live migration on an optimal number of servers and to switch off underutilized servers. The optimality criterion is the minimization of the energy consumption of the data center [8], [9].

Although live migration is widely used by the industry as well as the research community, most existing or proposed approaches disregard the cost of migration. For example, the live migration approach of Li et al. for energy-saving application placement in cloud computing environment disregards the cost of migration [10]. The same view is shared by similar approaches [11]–[13]. However in a previous work [14] we showed that the energy overhead of live migration cannot be neglected and it varies with the RAM size of the virtual machine and the network bandwidth that is available for migration.

This paper proposes a lightweight model to quantify the energy cost of live migration of virtual machines. A series of experiments were conducted on KVM [15] and we derived such a model through linear regression on recorded data. The migration cost model is able to estimate the energy overhead of live migration of an idle virtual machine within an accuracy of 90%.

The rest of the paper is organized as follows. In Section II we explain the technical aspect of live migration. In Section III, we summarize related work. In Section IV, we analyze the energy overhead due to live migration of virtual machines, and in Section V, we derive the migration cost model to estimate the energy overhead. In Section VI we discuss how to generalize the results of the paper. Finally, in Section VII, we outline some open research issues in this area and give concluding remarks.

II. Virtual Machine Live Migration

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. A common practice of current virtualization technologies (based on hypervisors) is to not use local discs to store VM images. Instead they require a network attached storage (NAS) or a storage-area-network (SAN) that can be accessible to all hosts and serves as hard drive for the virtual machines. By using a NAS/SAN, the process of live migration is limited to copying the in-memory state and the content of the virtual CPUs between the physical machines. For this task modern

978-0-7695-5028-2/13 $26.00 © 2013 IEEE
DOI 10.1109/CLOUD.2013.17

virtualization systems use a technique called pre-copy [2], consisting of the following three phases (see Figure 1):

1) *Pre-Copy Phase*: At this stage, the VM continues to run while its memory is iteratively copied page-wise from the source machine to the destination host. Iteratively means the algorithm works in several rounds. It starts with transferring all active memory pages. As each round takes some amount of time, some of the memory pages on the source machine may be changed (dirtied) and may no longer be in sync with the copy version on the destination host. These pages have to be re-sent to ensure memory consistency.
2) *Pre-Copy Termination Phase*: Without any stop condition, the iteratively pre-copy phase may carry on indefinitely. Stop conditions depend on the design of the hypervisor, but typically, they take one of the following thresholds into account: (1) the number of iterations exceeds a pre-defined threshold ($n > n_{th}$), (2) the total amount of memory that has already been transmitted exceeds a pre-defined threshold ($mem_{mig} > mem_{th}$), or (3) the number of pages dirtied in the previous round falls below a pre-defined threshold ($pg < pg_{th}$).
3) *Stop-and-Copy Phase*: At this stage the hypervisor suspends the VM to prevent further page writing and copies the remaining dirty pages as well as the state of the virtual CPUs to the destination host. After the migration process is completed, the hypervisor on the destination host resumes the VM.

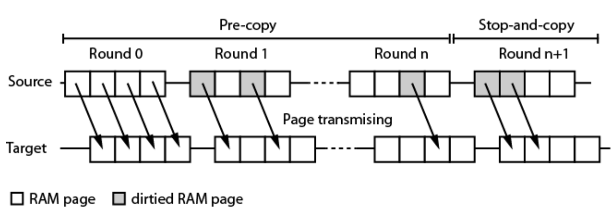

Fig. 1. Live Migration algorithm performs memory transfer page wise in several rounds [16].

III. Related Work

Live migration has been investigated in various contexts [5], [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.

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 an exception that is handled by the hypervisor. This slows down the VM's response to requests and reduces its throughput [2]. 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. investigate the processing speed of CPU-intensive and the reading speed of IO-intensive (disk) workloads. The authors learn that the performance of CPU-intensive workloads decelerates by 15% whereas the reading speed diminishes by 10% [21]. In [2], the authors demonstrate that the transmission rate of an Apache Web Server slows down by 12% to 20%. Performance loss may be problematic in systems where the response time constitutes a strict performance guarantee. For example, Voorslys et al. show 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 [24].

The additional resource utilization during live migration creates energy overhead. However, current live migration scenarios do not consider this energy cost. For example, Mistral [25] 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. This idea is shared by similar approaches which investigate service consolidation and dynamic power management in data centers [13], [26]–[28]. However, in previous work [14] we demonstrated that the energy cost of live migration of virtual machines cannot be neglected. The energy overhead rises with an increment in the RAM size of the VM and drops with an increment in the available network bandwidth.

In [20] and in [18] the authors propose two migration time models with different parameters and an accuracy of 90%. As energy is power consumed within time, these models could be used to capture energy cost of live migration. However this requires a deterministic power model, which is very difficult to specify, because it is not possible to give a complete account as to why the power consumption of a server or a component thereof behaves the way it does.

This paper addresses a lightweight energy cost model that copes without a dedicated power or time model. The model is able to quantify energy overhead of live migration with respect to the RAM size of the virtual machine as well as with respect to the servers available network bandwidth. It completes the study which focuses on performance cost of live migration of virtual machines and estimates energy cost within an accuracy of higher than 90%.

IV. Experiments

The average energy, E, of a migration is defined as the average power, P, multiplied by the migration duration τ:

$$E = P \times \tau \tag{1}$$

To derive the energy cost model, we conclude a series of experiments and apply linear regression on recorded data. We measure the power consumption of source and destination servers before and during migration and analyze the migration duration. To isolate the cost of migration from all other costs due to uncontrolled activities, we carry our migration when both servers are idle with 0% CPU utilization. As our previous work shows the energy overhead of live migration of virtual machines varies with the RAM size of the virtual machine and the available network bandwidth, we choose RAM size of virtual machine and network bandwidth as model parameter. The central research questions we would like to address can be formulated as follow:

1) How does TRAM size of virtual machine and network bandwidth available for migration, impact energy consumption during migration?
2) How does a formal model to estimate energy cost of migration with respect to network bandwidth and RAM size look like?

A. Cluster Set up

The server cluster we set up for our experiment consists of two identically constructed servers (we call them Gandalf and Wuotan), a client machine, and a network attached storage (NAS). All devices are connected to each other via a 1 GBit/s switch (Figure 2). The servers run Fedora 15 [29] (Linux kernel v. 2.6.38, x86_64) in which KVM [15] is used as a hypervisor. We use the open source operating system FreeNAS [30] (v. 8.0.1, AMD_64) as a Network Attached Storage.

Each server employs an Intel I5-680 Dual Core 3.6 MHz processor, 4 GB DDR3-1333 SDRAM memory and a 1 Gbit/s Ethernet Network Interface Card (NIC). The NAS system is equipped with one Intel Xeon E5620 Quad-Core 2.4 MHz processor, 10 GB DDR3-1333 SDRAM memory and 1 Gbit/s Ethernet NIC.

Athena, the virtual machine under test, runs Fedora 15 (Linux kernel v. 2.6.38, x86_64) with one virtual processor, variable memory size and network bandwidth. The client, written in C and hosted on a third physical machine, triggers the live migration of Athena between the servers using libvirt [31], a toolkit enabling interactions with the hypervisor and the operation system.

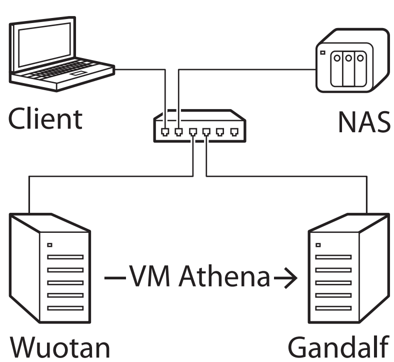

Fig. 2. The set up for migration cluster.

We run for each parameter setting 25 iterations. Each iteration starts with migrating Athena from Wuotan to Gandalf, followed by a break of 30 seconds. After that the virtual machine is moved back to Wuotan and the iteration concludes with a break of 30 seconds before the next iteration begins.

The test run of 25 iterations is controlled by a client program that uses the libvirt API to trigger migrations. The migration command returns immediately after live migration finishes allowing us to record the start and the end time as well as the duration of each migration.

We employ two Yokogawa WT210 digital power analyzers to measure the overall AC power consumption of both servers. The devices can measure DC as well as AC power consumption at a rate of 10 Hz and a DC current between 15 μA and 26 A with an accuracy of 99.9 %.

As live migration requires additional resources to perform pre-copy and stop-and-copy rounds, we installed dstat [32] to log CPU, memory and network utilization of the physical as well as the virtual machines.

To obtain the measurements belonging to the same migration, we synchronized the servers' time using the Network Time Protocol and use the timestamps of each measurement as links.

To study the influence of the RAM size of the virtual machine on the energy overhead, we set Athena's main memory stepwise from 800 MB to 1700 MB. Likewise, we limit the available network bandwidth during migration to a value between 20 MBps and 100 MBps. As KVM limits memory transfer in live migration to pages that are actually in use, we occupy Athena's entire reserved memory by running the memory allocator from [14]. In the following, we call the amount of main memory that is occupied by the virtual machine as *VM size*.

B. Preliminarily Experiments

To quantify the energy overhead of live migration, we compare the energy consumption of the source and the destination host with and without live migration. We measure the power consumption of Wuotan and Gandalf for ten hours continuously while they are idle. In this setting, Wuotan hosts Athena with 2700 MB main memory with no additional co-located virtual machines. Gandalf runs no VM or other applications.

The average overall idle state power consumption is 25.41 W for Wuotan (within a range of 23.5 W and 26 W) and 25.78 W for Gandalf (within a range of 24.5 W and 26.75 W). It is reasonable to say that both machines consume nearly the same amount of power while idle. The standard deviation for the idle power consumption is 2.32 W for both servers.

C. Power Consumption

To analyze the power consumption in more detail, we employ the probability distribution functions (CDF) of the power consumptions of the two servers during migration. The CDF approach considers the power consumption as a random variable **p**. This is a fitting consideration since it is not possible to give a complete account as to why the power consumption of a server or a component thereof behaves the way it does. The CDF, or simply the distribution function is defined as:

$$F_p(p) = P\{\mathbf{p} \leq p\} \tag{2}$$

where p is a real number. In other words, the distribution function expresses the probability that the random variable $\mathbf{p}$ (in our case the server's power consumption) has a value less than or equal to a certain real number p. Because it is a cumulative function, $F_p(p)$ is monotonic increasing, so that for any $p_2 > p_1$ the expression $F_p(p_2) \geq F_p(p_1)$ holds. Moreover, in our case, $F_p(0) = 0$ and $F_p(\infty) = 1$.

Figure 3 displays the CDF of the power consumption of the two physical machines during the live migration of Athena with different configurations. A close examination of these functions reveals the following:

1) The power consumption of both servers during migration alternates between two distinct power states: 1) the idle power P_i at approximately 25 W and 2) a variable active power P_a.
2) The value of P_a grows with increasing available (utilized) network bandwidth, whereas VM size does not contribute to power consumption during migration.
3) The time both severs spend in one of the two power states varies with network bandwidth and VM size.

Figure 3 shows the distribution functions $F_P(p)$ of the power consumption of the source server (left diagrams) and the destination server (right diagrams). The upper diagrams depict the power distribution functions for migration of a 1000 MB virtual machine when the servers were using network bandwidth between 20 MBps and 100 MBps. The lower part represents the power distribution functions for migrating virtual machines of size between 800 MB and 1700 MB with available network bandwidth of 100 MBps.

The four figures clearly display the role of the network bandwidth played during migration. Regardless of the size of the virtual machine, both servers are almost in two distinct states, either they are in near-idle state, consuming approximately 25 W, or they are active, consuming a power range that grows with increasing available network bandwidth. For example, the source server requires between 31 W and 38 W using 40 MBps network bandwidth, whereas the bandwidth of 100 MBps causes an active power consumption between 41 W and 46 W. Moreover, there was no significant difference in active power during the migration of virtual machines of size between 800 MB and 1700 MB.

Table I summaries the active power states of source and destination host for different network bandwidths. The range of power as well as the average power both servers consume in their active state exhaust a linear function with respect to network bandwidth. We explain this fact with the increase in computational effort and performance a higher network bandwidth causes.

The time both servers spend either in idle or in active mode varies with network bandwidth as well as with VM size. For example, the source server using 100 MBps network bandwidth is 5% of the time near an idle state (consuming $\approx$. 25 W) if the VM size is 1500 MB, whereas it is 20% of the time in the same state if the VM size is 900 MB. In contrast to that, migrating a 1000 MB sized virtual machine at 100 MBps, the source server works 10% of the time in near-idle state. Generally, idle part-time of destination server is longer than idle part-time of source server and we assume that is, because the source server controls the migration process and has a higher computational effort.

Network BW (MBps)	Active Power Consumption Range (W)	Average (W)
20	28 - 33	30
30	29 - 35	32
40	31 - 38	34
50	32 - 41	35
60	34 - 43	37
70	36 - 45	39
80	38 - 45	40
90	40 - 46	42
100	41 - 46	43

TABLE I
SUMMARY OF ACTIVE POWER CONSUMPTION OF SOURCE AS WELL AS DESTINATION HOST DURING A LIVE MIGRATION OF A VM.

D. Migration Time

We measured the migration time when the VM size was between 800 MB and 1700 MB and when the servers available network bandwidth was between 20 and 100 MBps. As migration varies within 0.4% at a max (coefficient of variation), we decided to average over all 25 migrations of the same parameter setting. Figure 4 summarizes the arithmetic average of the migration time of Athena with different RAM size and network bandwidth configurations and makes the following observations:

1) The migration time of an idle VM decreases non-linearly with increasing available (utilized) network bandwidth.
2) The migration time of an idle VM increases linearly with increasing VM size.

Figure 4 shows that the migration time drops by nearly 50% if network bandwidth doubles. E.g. the average migration time for VM size of 1000 MB with the network bandwidth of 20 MBps is 41 seconds and decreases to 21 seconds with a network bandwidth of 40 MBps. In contrast to that, if the VM size doubles, the migration time increases by 183%, e.g. if the network bandwidth is 50 MBps from 18 seconds for a VM size of 800 MB to 33 seconds for a VM size of 1600 MB. These observations clearly demonstrate the role of VM size and network bandwidth in live migration of virtual machines with the pre-copy approach. In idle mode, if almost no page dirtying happens, no pages have to be copied twice and pre-copy is limited to only one iteration to move the VM's entire memory content from the source to the destination host. Hence, if a VM size is doubled, for example from 800 MB to 1600 MB, twice as many bits are copied, resulting in a migration time that is twice as long. In contrast, a higher network bandwidth allows

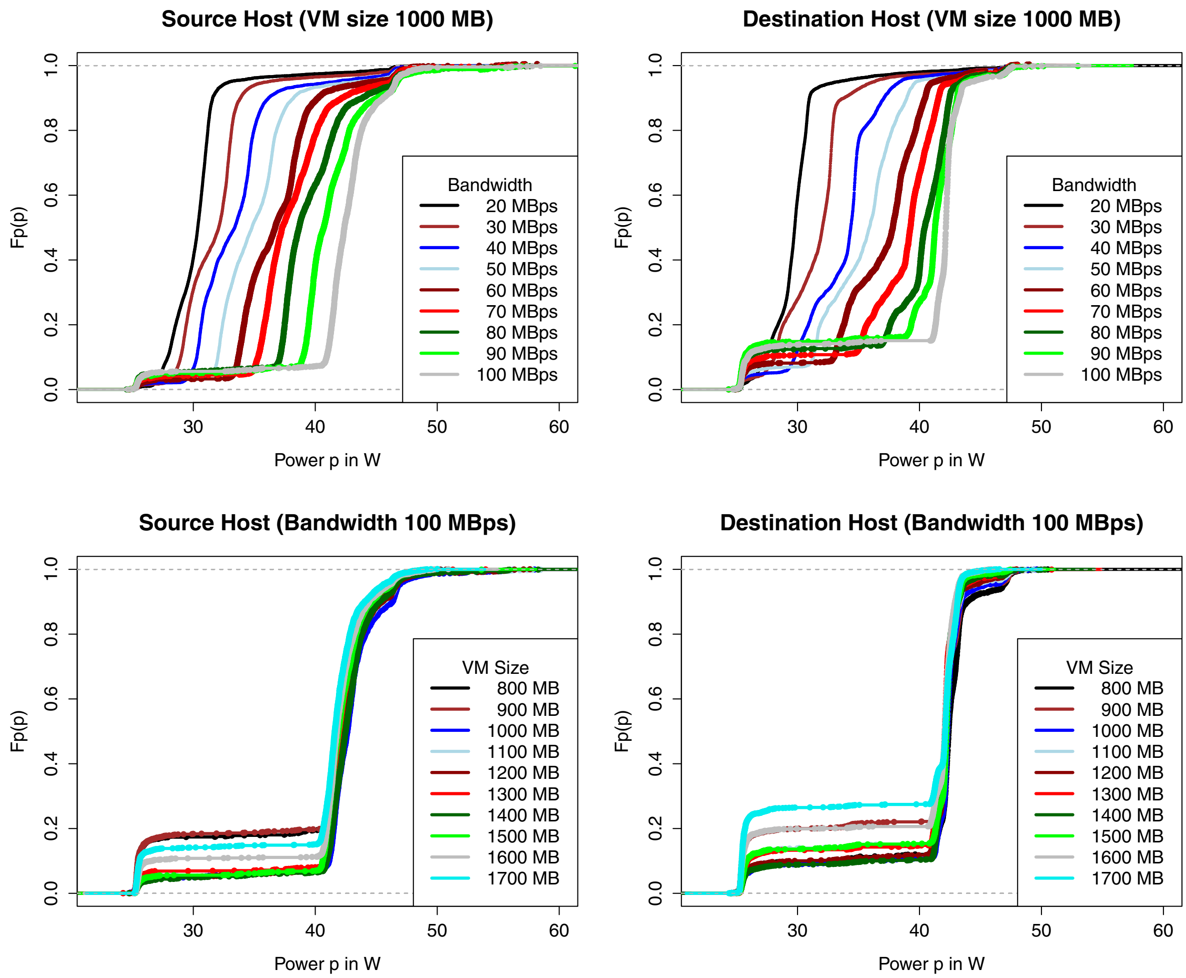

Fig. 3. The distribution function $F_P(p)$ of the power consumption of the source and the destination server during a migration. Top: When the VM size was 1000 MB. Bottom: When the network bandwidth was 100 MBps.

moving more bits within a specific time period, reducing the time required for memory transfer.

Akoush et al. propose in [18] a migration time model for idle VMs as in Eq. 3. The model consists of two parts: (1) a constant time A for pre- and post-migration overhead and (2) the duration of transfer of VM's main memory content from source to destination host. To calculate the transfer time the VM size s has to be divided by the available network bandwidth b. The parameter B represents the protocol overhead increasing the migration volume by a constant factor.

$$\tau = A + \frac{B * s}{b} \tag{3}$$

The experimental results we got through migrating of Athena with different RAM sizes and network bandwidths look similar to the model of Akoush et al. Thus, we applied their model to our measurements using a two-part approach. At first, we utilized non-linear regression to identify the values for A and B that fit best migration time τ. By doing so, we picked the migration experiments with VM size between 800 MB and 1200 MB and network bandwidth between 20 MBps and 60 MBps. As in Eq. 4, in our cluster setup, migration overhead takes 2.07 seconds and the protocol overhead of copying main memory content from source to destination host amounts to 18%.

$$\tau = 2.07 + \frac{1.18 * s}{b} \tag{4}$$

Finally, we evaluate the fitted model. We calculated the migration time as in Eq. 4 and compared this value with the average migration time measured experimentally. However, this time, we focus on VMs of size between 1300 MB and 1700 MB and network bandwidth between 70 MBps and 100 MBps. The model estimation error was within 3% in average with a minimum of 0% and a maximum of 5%. Although the values for parameter A and B are specific for the cluster setup we used, these results demonstrate that time

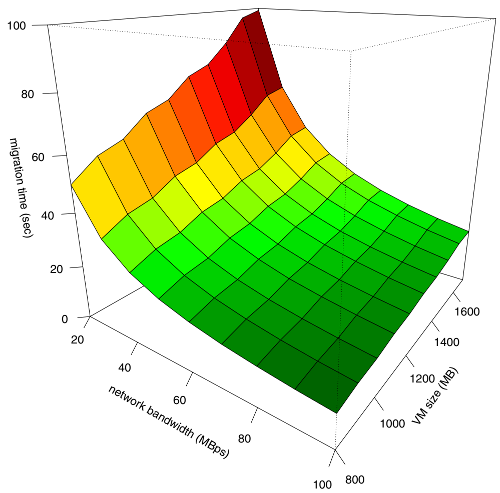

Fig. 4. Average migration time for virtual machines of size between 800 MB and 1700 MB and network bandwidth between 20 MBps and 100 MBps.

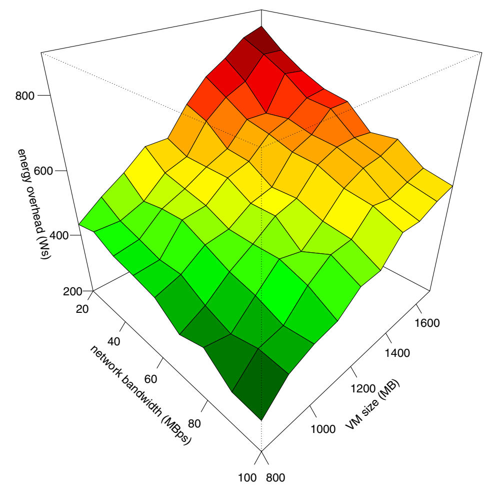

Fig. 5. Average energy overhead during migration of virtual machines of size between 800 MB and 1700 MB and network bandwidth between 20 MBps and 100 MBps.

for virtual machine live migration can be accurately estimated using a non-linear model as in Eq. 3.

E. Energy Overhead

We measured power consumption and migration time of virtual machines with size between 800 MB and 1700 MB that are migrated using network bandwidth between 20 MBps and 100 MBps. As Section IV.C showed, both, source and destination server do not consume constant power during migration. Instead power varies between 24 W and 50 W. Thus we could not apply Eq. 1 and calculate live migration's energy overhead as in Eq. 5. At this, τ_s and τ_e indicates start and end time of the migration process, respectively and $P_s(t)$ and $P_d(t)$ defines power consumption of source and destination host with respect to time t. P_i represents idle power consumption of source as well as destination host of approximately 25 W that has to be subtracted from active power to capture energy overhead.

$$E_{ov} = \int_{\tau_s}^{\tau_e} P_s(t) + P_d(t)dt - 2 * P_i * (\tau_e - \tau_s) \qquad (5)$$

Similar to migration time we use the arithmetic average to display the energy overhead of live migration of variable sized virtual machines with different network bandwidths. The coefficient of variation was 6% in average and drops with increasing VM size and network bandwidth from 12% to 4%. Figure 5 summarizes the arithmetic average of the energy overhead during live migration of Athena with different configurations. A close examination of these functions reveals the following:

1) The energy overhead of live migration of idle virtual machines drops in a minor way with increasing bandwidth and grows significantly with increasing VM size.
2) In general, from an energy point of view, live migration of virtual machines should be done with the network bandwidth as high as possible. If there are several candidates for migration, choosing the smallest VMs with respect to their occupied RAM, is the best option in case of energy efficiency.

Figure 5 displays that the average energy overhead decreases with increasing bandwidth. For example a virtual machine with a size of 1100 MB requires 525 Ws in average to be migrated using network bandwidth of 40 MBps which drops by 10% to 480 Ws if bandwidth doubles to 80 MBps. In contrast to that, if the VM's RAM size doubles, the energy overhead increases by 189%, e.g. if network bandwidth is 60 MBps, from 370 Ws for VM size of 800 MB to 720 Ws for VM size of 1600 MB. These observations clearly demonstrate the influence of migration time on energy cost of live migration of virtual machines. As energy is power consumed within time and VM size effects migration time substantially, the migration of an idle virtual machine with a size as twice as big and with almost no page dirtying takes as twice as long. As a consequence the energy overhead doubles. Moreover, migration time is more important than power consumption of source and destination host. Although power consumption increases by 6 W in average when network bandwidth doubles, energy overhead decreases, because migration time is shortened by half.

V. Energy Model

One of the major use cases of live migration is to move virtual machines between servers of data centers to consolidate load and to minimize energy consumption. Due to the fact that live migration costs energy and any reconfiguration aims to reduce energy consumption, one of the most important tasks is to select those virtual machines whose replacements save at least as much energy as their migrations cost. Optimally, a reconfiguration conserves much more energy than it costs, in order to minimize data center energy consumption. To make energy efficient decisions in terms of migration requires a migration cost model that enables to quantify the energy overhead of virtual machine live migration in advance. We derived such a model through linear regression on our experimental data of energy overhead of live migration.

Figure 5 displays the energy overhead with respect to network bandwidth and VM size. We assume a linear relationship as in Eq. 6, whereas s represents VM size and b the network bandwidth.

$$E_{ov} = A + B * s + C * b \tag{6}$$

We divide the data set we got by our experiments into two subsets. The first subset, called training subset, is used to train our model. We apply linear regression on 625 measurements, that include the raw energy overhead (not average) of migrations of virtual machines of size between 800 MB and 1200 MB using network bandwidth between 20 MBps and 60 MBps. Eq. 7 represents the fitted values for A, B and C.

$$E_{ov} = 201Ws + 0.4\frac{Ws}{MB} * s - 1.7\frac{Ws}{MBps} * b \tag{7}$$

Finally, we evaluate the estimation accuracy of the energy cost model based on a second data subset that consists of the entries for the remaining energy overhead data we measured during our experiments. This time we considered virtual machines of sizes between 1300 MB and 1700 MB and network bandwidths between 70 MBps and 100 MBps. We compared the measured energy overhead with the value estimated by the energy model in Eq. 7 and computed the estimation error. The average estimation error was below 10 %. This result evinces that energy overhead of live migration of idle virtual machines varies linear with VM size and network bandwidth and it is reasonable to capture this overhead with a lightweight linear model as in Eq. 6.

VI. Discussion

As demonstrated in the proceeding results, VM size as well as network bandwidth affects energy overhead of live migration of virtual machines, substantially. As energy is power consumed within time, the impact results from the following observations:

1) The power consumption of both servers during migration alternates between the idle power P_i at approximately 25 W and an active power P_a, whereas the later one grows linearly with increasing network bandwidth. The VM size has no impact on power consumption during migration.
2) The migration duration varies linearly with increasing VM size and linearly with the inverse of network bandwidth. It can be estimated accurately with a migration time model as in Eq. 3.
3) By implication, the energy overhead rises linearly with an increment in the VM size and drops linearly with an increment in the available network bandwidth, whereas the impact of VM size is higher than the impact of network bandwidth.

Based on our experimental results, we deduced a lightweight, linear model to estimate the energy cost of live migration of virtual machines as in Eq. 6 and fitted the values for the model parameter through linear regression. These values are specific for the cluster setup we used in our experiments and cannot be generalized. Instead linear regression has to be repeated for all hosts in data center to make quantitative statements for the energy costs of live migration. However the energy model, i.e. the relationship between energy cost of live migration and network bandwidth as well as VM size is universal and allows deriving general migration guidelines to improve data center's energy efficiency:

1) Always use network bandwidth as high as possible during migration as energy cost reduces with increasing network bandwidth.
2) Always migrate the virtual machine with the smallest amount of occupied main memory as each additional MB in migration volume increases energy cost of live migration substantially.

Furthermore the energy model allows to make qualitative statements and to compare the virtual machines with respect to their migration costs. Moreover, as studies have shown data centers are most of time ($\geq$ 70%) in idle state [33], thus, live migration of idle virtual machines is the rule rather than the exception in nowadays data centers. Thus, despite we do not considered CPU load and memory utilization, the energy model proposed in this paper is applicable.

VII. Conclusion

In this paper we investigated the energy cost of live migration and analyzed the role of the VM size and the role of the available network bandwidth in the energy consumption of hosting servers. We migrated virtual machines of sizes between 800 MB and 1700 MB if network bandwidth is between 20 MBps and 100 MBps. We measured the migration time and the power consumption of source and destination host and compared these values with the power consumption of the hosts in an idle state to capture the energy cost of live migration. Based on the profiled data, we deduced a lightweight energy cost model through linear regression. Experimental results demonstrate that this model is able to estimate energy overhead of live migration of virtual machines with an accuracy of more than 90%.

In future work we will extend our experiments to analyze the impact of CPU and memory usage on power consumption and migration time to refine the energy model to migration of busy virtual machines. Furthermore we plan to measure energy consumption of the various subsystems of source and destination server in order to better manage when and which VM should be migrated during service consolidation.

ACKNOWLEDGMENT

This work has been partially funded by the European Regional Development Fund (ERDF), the European Social Fund (ESF), and the German Free State of Saxony under project agreement number 080949277

The work has also been partially funded by the German Research Foundation (DFG) under project agreement: SFB 912/1 2011.

REFERENCES

[1] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield, "Xen and the art of virtualization," *SIGOPS Oper. Syst. Rev.*, vol. 37, no. 5, pp. 164–177, Oct. 2003. [Online]. Available: http://doi.acm.org/10.1145/1165389.945462

[2] 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

[3] 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

[4] 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.

[5] 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

[6] 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.

[7] 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.

[8] 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.

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

[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] A. Strunk and W. Dargie, "Does live migration of virtual machines cost energy?" in *Proceedings of the 27th IEEE International Conference on Advanced Information Networking and Applications (AINA-2013)*, ser. HPDC '11, 2013.

[15] "Kvm: Kernel-based virtual machine. http://www.linux-kvm.org/page/mainpage/."

[16] 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

[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] 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.

[25] 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.

[26] 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.

[27] 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

[28] 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.

[29] "Fedora 15. http://fedoraproject.org/."

[30] "Freenas: Storage for open source. http://www.freenas.org/."

[31] "Libvirt: The virtualization api. http://fedoraproject.org/."

[32] "Dstat: Versatile resource statistics tool. http://dag.wieers.com/home-made/dstat/."

[33] F. Ahmad and T. N. Vijaykumar, "Joint optimization of idle and cooling power in data centers while maintaining response time," *SIGARCH Comput. Archit. News*, vol. 38, no. 1, pp. 243–256, Mar. 2010. [Online]. Available: http://doi.acm.org/10.1145/1735970.1736048