

Estimating Energy Consumption during Live Migration of Virtual Machines

Kateryna Rybina, Alexander Schill

Technical University Dresden

SFB 912 - HAEC [Chair of Computer Networks]

01062 Dresden, Germany

Email: {kateryna.rybina, alexander.schill}@tu-dresden.de

Abstract—The workload coming to the server clusters varies over the course of the day which results in disproportional resource utilisation. Server consolidation is a technique exploited by cloud service providers in order to optimally utilise available physical resources. During the time intervals of low resource usage the workload is consolidated from a big set of underutilised servers to a smaller set of optimally loaded servers. Afterwards idle servers are turned off which reduces energy consumption as well as running costs. Server consolidation can be realised via live migration of virtual machines (VMs). Though, the VMs’ migration process requires additional CPU cycles and introduces energy overhead. The quantitative knowledge of the energy consumption during migration is important in order to realise sophisticated migration decisions. This paper presents how the energy consumption of the servers during migration can be estimated based on their resource utilisation parameters using linear regression techniques. We defined the most significant parameters that influence the energy consumption of the servers during migration. These are: CPU instructions retired, last level cache line misses, and “dirty” pages observed in the source server during migration.

Index Terms—virtual machine, server consolidation, live migration, energy overhead, energy modelling, linear regression.

I. INTRODUCTION

Different techniques have been adopted by cloud service providers in order to minimise the energy consumption. First of all cloud service providers usually execute their applications within the virtual machines. Several such VMs can be running within one physical server, which allows to reduce the number of physical machines, and thus, lower the total system’s energy consumption [1]. During periods of low workloads the VMs can be consolidated from a bigger to a smaller set of servers via migrating these VMs and turning off idle servers [2], [3]. It helps to reduce the system’s energy consumption as well as running costs [4]. Though, the reconfiguration in the system via live VM’s migration has also its drawbacks. First, the service’s performance might be degraded during the VM’s migration; second, the migration process requires additional resources such as CPU cycles and network bandwidth [3]. Increased CPU activity results in increased energy consumption. Thus, migration process itself creates energy overhead

which has to be accounted for when looking for the best VM migration candidate [4]. In this paper we demonstrate how the energy cost of the VM’s migration can be estimated based on the resource utilisation parameters of the servers. The main contributions of the paper are as follows:

- 1) We demonstrate that multiple linear regression (MLR) techniques can be used to model and predict the energy consumption of the servers during the VM’s migration.
- 2) We identify the most important parameters that influence the energy consumption of the servers during migration and build our models for *memory* intensive workloads.

In previous research work [5] we made investigations into the energy consumption of the servers during live migration of the VMs and build MLR models for *CPU* intensive workloads. In this paper we extend our research towards developing energy models for *memory* intensive workloads and define important parameters that are significant for this type of workload. In [6] we previously applies MLR techniques to model migration time of virtual machines.

Strunk [7] created a model to derive energy overhead of migration applying linear regression with two parameters: RAM utilised by the VM and network bandwidth. The limitation of the model is that it was built for idle VM (running no workload), thus such models can not be applied for active VMs. The energy model of Liu et al. [8] and Akoush et al. [9] do not account for the CPU utilisation of the source server, though the power drawn by the source server depends on its CPU utilisation [10]. Liu et al. [8] assume that the energy consumption of the source and the destination servers during migration depends only on the network traffic caused by the VM’s migration process. The authors apply simple linear regression (SLR) in order to model the energy consumption of migration and simplify their model by obtaining model coefficients for two servers together. In this work we suggest to build MLR models which predict the energy consumption during the VM’s migration for each server separately, as it differs between the servers not negligibly.

The remaining part of the paper is structured as follows: In Section II we provide fundamental knowledge regarding

VM’s migration and linear regression analysis. In Section III we describe experiment set-up, benchmarks and dataset applied for training and testing energy models. In Section IV we describe our approach to modelling energy consumption and define the best MLR model, and identify the most important model’s parameters. Finally, in Section V we summarise our work.

II. FUNDAMENTAL TECHNIQUES

Live migration of the VM is a process of physically moving the VM (copying its RAM) between two servers. The VM is executing interruption free in an unperceived to the end user manner during migration. Using *pre-copy* live migration algorithm [9], [8] the migration is realised in several consecutive rounds (iterations). During the first iteration all memory content of the VM is copied between two servers. But due to the fact that copying takes time, the applications that are executing whitening the VM manage to modify (“dirty”) some of the memory pages. Only those pages have to be resent during the consecutive iteration(s). Some stop-conditions are defined upon reaching which, the VM is shortly suspended on the source server until the remaining “dirty” pages are copied to the destination server. The stop-conditions depend on the implementation of the hypervisor but some of the common stop-conditions are: 1) The number of iterations reached a pre-defined threshold; 2) the number of remaining “dirty” pages falls below a pre-defined threshold; or 3) the number of “dirty” pages in the current iteration is higher than that one in the previous iteration [11], [9]. Afterwards the VM is resumed on the destination server and stopped on the source server. Thus, the VM migration takes some time and requires additional CPU resources, which results in an increased energy consumption of the servers during migration. Our goal is to develop models which can be used to estimate energy consumption of the servers during migration.

It is possible to estimate the dependent variable (also called response variables) from multiple independent variables (also called predictor variables) using MLR techniques. MLR assumes that the conditional expectation of a dependent variable is a linear function of multiple independent variables x_1, \dots, x_k as shown in Equation 1 [12], [13].

$$\mathbf{E}\{Y|X_1 = x_1, \dots, X_k = x_k\} = Z(x) = \beta_0 + \beta_1 x_1 + \dots + \beta_k x_k \quad (1)$$

Where β_0 is an intercept and $\beta_1 \dots \beta_k$ are slopes [12]. First, such MLR model has to be trained on the training data. Afterwards the model can be used to predict the dependent variable from new (test) data. Important is to find model’s coefficients, so that the error of the model (difference between the real data values and values that lie on regression plane) is minimised. This can be done using the method of least squares [12], [13].

We can judge how well the linear regression model fits the real data and how appropriate are the selected

independent variables by analysing R-square (R^2) or adjusted R-square (R_{Adj}^2) coefficients [12]. The closer these coefficients are to 1 the better model fits the measured data. R_{Adj}^2 is more appropriate parameter when more than one independent variable is used to model the dependent variable as it penalizes for adding insignificant parameters to the model [12], [13]. Moreover, how accurate is a model can be analysed by calculating model’s mean absolute percentage error $Perc_m$ on test data as follows¹:

$$Perc_m = \frac{\sum_{i=1}^n \left| \frac{\tilde{y}_i - y_i}{y_i} \right|}{n} \times 100\% \quad (2)$$

Where \tilde{y}_i is a predicted by the model value and y_i is true (measured) value, and n is a number of observations.

Another error, namely model’s standard error of estimate from a sample² $Pr_{st.err}$ quantifies the departure of the predicted by the trained model value of the dependent variable \tilde{y}_i from the actual value y_i :

$$Pr_{st.err} = \sqrt{\frac{\sum_{i=1}^n (y_i - \tilde{y}_i)^2}{n - k - 1}} \quad (3)$$

Where n is a number of observations, k is the number of independent variables, and minus 1 accounts for the intercept. $Pr_{st.err}$ is calculated on test data. The residual standard error of linear regression model $Res_{st.err}$ is calculated in a similar way to $Pr_{st.err}$ but on the training data. It shows the departure of the fitted by the model values (which lie on regression plane) from real data.

III. EXPERIMENT SET-UP

Our experiment set-up consists of two homogeneous servers (source and destination), a network attached storage (NAS), client machine used to trigger VM’s migration, and two Yokogawa WT210 digital power analysers. These power analysers were measuring DC and AC power consumption of the source and the destination servers at the rate of 10 Hz. Two servers between which the VM’s migration was realised have Intel 15-680 Dual Core 3.6 GHz processors, include 4 GB DDR3-1333 SDRAM, and 1 Gbps Ethernet NIC. The NAS consists of Intel Xeon E5620 Quad-Core 2.4 GHz processor, 10 GB DDR3-1333 SDRAM memory, and 1 Gbps Ethernet NIC.

Operating system Fedora 15³ with Linux kernel v. 2.6.38, x86_64 was installed on both servers. Kernel-based Virtual Machine (KVM)⁴ was used as a hypervisor and libvirt⁵, as a toolkit to manage migration. As the NAS we employed FreeNAS⁶, which is a FreeBSD-based operating system (v. 8.0.1, AMD 64). The VM under test was allocated 4 GB RAM, 4 virtual CPUs and 20 GB disc

¹<http://mathworld.wolfram.com/PercentageError.html>

²<http://onlinestatbook.com/2/regression/accuracy.html>

³Fedora 15. <http://fedoraproject.org/>.

⁴http://www.linux-kvm.org/page/Main_Page.

⁵Libvirt: <http://libvirt.org/>.

⁶FreeNAS: <http://www.freenas.org/>.

space on the NAS. Operating system Ubuntu 14.04.2 LTS (Linux kernel 3.16.0-30-generic) was installed on it. The VM under test was migrated in isolation.

The experiment was organised as follows: The VM was migrated between the source and the destination servers at the data transmission rate from 70 MBps to 100 MBps in steps of 10 MBps. *Memory* intensive benchmarks from SPEC CPU2006⁷ benchmark suite were continuously executing on the VM during migration. We considered that a benchmark is *memory* intensive when it has a large number of read/write operations from/to the memory subsystem. After a coarse grained separation of benchmarks: *bzip2*, *astar*, *mcf*, *gcc*, and *perlbench* were selected as *memory intensive* benchmarks.

We realised 20 ping pong VM’s migrations from the source server to the destination server and backwards for each considered data transmission rate and each benchmark. As a result we derived a dataset of 400 entries (20 iterations \times 5 benchmarks \times 4 data transmission rates). The VM’s migrations corresponding to the data transmission rate of 70 MBps, 80 MBps, and 100 MBps (300 separate migration runs) were used as *training* data for the model. The observations corresponding to the data transmission rate of 90 MBps (100 migrations) were used as *test* data to check the accuracy of the model.

Moreover, during migration we were continuously recording resource utilisation parameters of the servers and the VM. *dstat*⁸ resource monitoring tool was executing on the VM as well as on the source and the destination servers and recording the total mean CPU utilisation and active memory (RAM) utilisation of the servers and VMs. Intel performance monitoring counters *Intel PCM* tool⁹ was used to monitor last level (L3) cache line misses $L3_{miss}$ as well as total number of instructions retired $INST$ ¹⁰. Furthermore, we recorded another statistic which indicates the total amount of “dirty” memory (measured in kilobytes) waiting to be written back to the disk *Dirty*¹¹. From it we derived the total number of 4KB “dirty” memory pages of the source server during migration $DirtyPages_{server}$. Table I presents the parameters which we monitored during migration and which will be used for predicting the energy consumption of the servers during migration. We were recording the start and the end time of each VM’s migration in order to extract afterwards the resource utilisation statistics which are related to VM migration process by analysing timestamps. The data analysis was realised in R [14] which is a language and an environment for statistical analysis.

⁷<https://www.spec.org/cpu2006/publications/>

⁸*dstat*: <http://linux.die.net/man/1/dstat/>.

⁹www.intel.com/software/pcm

¹⁰<https://software.intel.com/en-us/blogs/2014/07/18/intel-pcm-column-names-decoder-ring>

¹¹http://www.centos.org/docs/5/html/5.2/Deployment_Guide/s2-proc-meminfo.html

TABLE I
RESOURCE UTILISATION PARAMETERS (INDEPENDENT VARIABLES)
USED TO MODEL ENERGY CONSUMPTION OF THE SOURCE SERVER
DURING VM’S MIGRATION $E_{s-during}$.

| Variable | Description |
|-----------------------|--|
| $L3_{miss}$ | Total number of last level cache line misses during migration (in millions) |
| $INST$ | Total number of instructions retired during migration (in thousands) |
| $CPU_{util_{server}}$ | Mean total CPU utilisation of the source server during migration (in percentage) |
| $MEMtoBW_{server}$ | Ratio of active memory utilised by the source server to data transmission rate |
| $DirtyPages_{server}$ | Number of “dirty” pages observed in the source server during migration |

IV. ENERGY MODEL

We apply MLR techniques to model the dependence of the energy consumption of the servers during migration on the resource utilisation parameters. In our case energy consumption of the server is a dependent variable and its resource utilisation parameters (see Table I) are independent variables. Here we build energy models for the source server as the VM is continuously executing on it during migration. The models of the energy consumption of the destination server during migration can be built using the same technique but based on its particular resource utilisation parameters.

First, in order to use linear regression analysis we have to assure that the energy consumption is linearly dependent on each of the selected independent variables. To prove it we calculate Pearson’s correlation coefficients between energy consumption of the source server during migration $E_{s-during}$ and each of the independent variables. The results can be seen in Table II. We used *cor.test()* function in R [14] which calculates Pearson’s correlation coefficients and tests the significance of these coefficients. *pvalue* in Table II shows the probability of observing the correlation coefficient as high as depicted in the table when assuming that the null-hypothesis is true. The null-hypothesis assumes that there is no correlation between the variables ($\rho = 0$). Thus, having *pvalue* $< 4.624e - 06$ in all cases we can reject the null hypothesis and state that correlation coefficients are not equal to 0 and the variables are linearly dependent. $E_{s-during}$ has the highest positive dependence on $INST$ ($\rho = 0.893$) and the least positive dependence on $MEMtoBW_{server}$ ($\rho = 0.268$) from the considered five independent variables. The closer ρ is to 1 the stronger is the positive linear dependence. Thus, these variables can be applied in linear regression models.

A. Selecting the best model

We define the best MLR model using the method called *all subsets regression* [13]. This method examines all possible parameters’ combinations and selects the best model as that one that has the highest adjusted R-square measure. We have considered five independent variables,

TABLE II
 SAMPLE PEARSON'S CORRELATION COEFFICIENTS BETWEEN ENERGY CONSUMPTION OF THE SOURCE SERVER DURING MIGRATION ($E_{s-during}$) AND RESOURCE UTILISATION PARAMETERS DEFINED IN TABLE I.

| Pearson's correlation, ρ | Coefficient | p-value |
|---|-------------|-----------|
| $cor.test(E_{s-during}, INST)$ | 0.893 | < 2.2e-16 |
| $cor.test(E_{s-during}, L3_{miss})$ | 0.763 | < 2.2e-16 |
| $cor.test(E_{s-during}, DirtyPages_{server})$ | 0.819 | < 2.2e-16 |
| $cor.test(E_{s-during}, MEMtoBW_{server})$ | 0.268 | 4.624e-06 |
| $cor.test(E_{s-during}, CPU_{util}_{server})$ | -0.321 | 2.97e-08 |

thus *all subsets regression* will analyse all combinations of one, two, three, four and five independent variables and output the best model for each subset size. The analysis were realised in R [14] using *regsubsets()* function from *leaps* package. Figure 1 depicts the best model for each subset size with respect to R^2_{Adj} as well as the mean absolute percentage error of each of these models ($Perc_m$). We can see that the best simple linear regression model is the one that uses CPU instruction retired $INST$ as independent variable. Its R^2_{Adj} is equal to 0.797. It means that 79.7% of variance in energy consumption of the source server during migration is explained by this single parameter. Though, $Perc_m$ of the model on the test data equals to 30.11%, which indicates that SLR model has low prediction accuracy. The best model with two independent variables includes $INST$ and $L3_{miss}$ independent variables (depicted as $INST_L3$ on the figure). Its R^2_{Adj} equals to 0.934 and $Perc_m$ is reduced to 16.48%. We can clearly see the benefit of MLR models compared to SLR models. The best MLR model which includes three parameters ($INST$, $L3_{miss}$, and $DirtyPages_{server}$) has R^2_{Adj} equal to 0.943 and its $Perc_m$ is the lowest and equals to 14.5%. Adding additional parameters to the model neither improve R^2_{Adj} significantly nor the prediction accuracy on the test data. On the contrary, mean absolute percentage error $Perc_m$ of the best model with five independent variables ($INST$, $L3_{miss}$, $DirtyPages_{server}$, CPU_{util}_{server} , and $MEMtoBW_{server}$) is 1.41% higher compared to the model with three parameters. The summary of all five models with respect to R^2 , R^2_{Adj} , mean average percentage error calculated on test data $Perc_m$ (in percentage), residual standard error of the model calculated on training data (in Watt per hour), and standard error of estimate from a sample calculated on test data $Pr_{st.err}$ (in Watt per hour) is given in Table III.

Taking a balance between the expressive power of the model (its R^2_{Adj}), its estimation error $Perc_m$ and complexity we select the best model in favour of the model with three independent variables. The equation of the best MLR model with three independent variables is as follows:

$$E_{s-during} = -0.12 + 0.0031 \times INST + 0.0006 \times L3_{miss} + 0.0027 \times DirtyPages_{server} \quad (4)$$

Where -0.12 is an adjustment constant (intercept) and a

coefficient 0.0031 assumes an expected increase of 0.0031 Watt per hour in energy consumption of the source server during VM's migration $E_{s-during}$ when the number of CPU instructions retired $INST$ increases for one unit (in our case one thousand) when keeping all the other parameters constant. The models' coefficients were estimated using the method of ordinary *least squares*. All model parameters are significant with p-value < 2.2e-16. Furthermore, the model satisfies main statistical assumptions of linear regression: Normality, linearity, and constant variance.




Fig. 1. The best model with respect to R^2_{Adj} for each subset size of independent variables defined with *all subsets regression* method. The prediction error on the test data of each of these models is depicted in red rectangles and is measured in percentage.

B. Defining importance of model's parameters

In order to define which of the three parameters used in the best MLR model is the most important we applied two methods. The first method calculates the relative importance of parameters by defining their *relative weights* in percentage in contributing to the model's R^2 [13]. The calculation was realised in R. The function for calculating *relative weights* was taken from [13] which is based on the method suggested by Johnson [15]. Figure 2 depicts the contribution of each independent variable to the model's R^2 . We can see that $INST$, $L3_{miss}$, and $DirtyPages_{server}$ are the first, the second, and the third most important parameters as they contribute 43.6%, 29.1%, and 27.3%, respectively to the model's R^2 of 0.943. The second method defines relative importance of indepen-




Fig. 2. Relative weights of parameters in the best MLR model with three independent variables. Relative weights were calculated based on function provided in [13].

dent variables in the model by analysing their standardised regression coefficients. According to this method the

TABLE III
THE BEST SLR AND MLR MODELS WHICH PREDICT ENERGY CONSUMPTION OF THE SOURCE SERVER DURING MIGRATION.

| The best SLR and MLR models: $lm(E_{s-during} \sim Predictors)$ | R^2 | R^2_{Adj} | $Res_{st.err}$ | $Pr_{st.err}$ | $Perc_m$ |
|---|-------|-------------|----------------|---------------|----------|
| $INST$ | 0.798 | 0.797 | 0.302 | 0.229 | 30.11 |
| $L3_{miss} + INST$ | 0.935 | 0.934 | 0.172 | 0.136 | 16.48 |
| $L3_{miss} + INST + DirtyPages_{server}$ | 0.943 | 0.943 | 0.16 | 0.12 | 14.5 |
| $L3_{miss} + INST + DirtyPages_{server} + CPU_{util_{server}}$ | 0.949 | 0.948 | 0.153 | 0.145 | 19.46 |
| $MEMtoBW_{server} + L3_{miss} + INST + CPU_{util_{server}} + DirtyPages_{server}$ | 0.95 | 0.949 | 0.152 | 0.137 | 15.91 |

regression variables are standardized to a mean of 0 and standard deviation of 1 prior to realising linear regression analysis. This can be done with the *scale()* function as suggested in [13]. The results of the second method are as follows: $INST$, $L3_{miss}$, and $DirtyPages_{server}$ are the first, the second and the third important independent variables as one standard deviation change in these variables results in 0.6, 0.4, 0.15 standard deviation change in $E_{s-during}$, respectively. Thus, the results derived by both methods are the same.

V. SUMMARY AND FUTURE WORK

In this paper we developed models which can predict the energy consumption of the servers during the VM's migration based on the resource utilisation parameters of these servers. To do it we applied multiple linear regression techniques. The VMs were executing *memory* intensive benchmarks during migration. The best MLR model was identified using *all subsets regression* method based on R^2_{Adj} metric and taking into account the prediction error of the trained model on test data $Perc_m$. The best MLR model has R^2_{Adj} equal to 0.943 and its mean absolute percentage error $Perc_m$ equals to 14.5% and it satisfies the statistical assumptions of linear regression analysis. With our experiments and evaluation of results we showed that:

1) Multiple linear regression techniques can be used in order to model the energy consumption of the servers during the VM's migration.

2) The most significant parameters of the models developed for *memory* intensive workloads are: CPU instructions retired $INST$, last level cache line misses $L3_{miss}$, and "dirty" memory pages observed on the source server during migration $DirtyPages_{server}$.

The obtained MLR models can be applied in order to solve resource optimisation problem, which is our objective in the future research work. The quantitative energy cost of migration can guide us in selecting the source server from which the VMs have to be migrated and identify VMs which are the best migration candidates.

ACKNOWLEDGMENT

This work is supported (in part) by the German Research Foundation (DFG) within the Collaborative Research Center SFB 912 - HAEC.

REFERENCES

- [1] J. N. S. T. Christian Baun, Marcel Kunze, "Cloud computing: Web basierte dynamische it-services," in *Springer, ISBN 978-3-642-01593-9*, 2009.
- [2] F. Xu, F. Liu, H. Jin, and A. Vasilakos, "Managing performance overhead of virtual machines in cloud computing: A survey, state of the art, and future directions," *Proceedings of the IEEE*, vol. 102, no. 1, pp. 11–31, Jan 2014.
- [3] A. Strunk, "Costs of virtual machine live migration: A survey," in *Services (SERVICES), 2012 IEEE Eighth World Congress on*, june 2012, pp. 323–329.
- [4] A. Verma, P. Ahuja, and A. Neogi, "pmapper: Power and migration cost aware application placement in virtualized systems," in *Proceedings of the 9th ACM/IFIP/USENIX International Conference on Middleware*, ser. Middleware '08. New York, NY, USA: Springer-Verlag New York, Inc., 2008, pp. 243–264.
- [5] K. Rybina, "Modelling the live migration time of virtual machines," in *PhD thesis submitted on the 04.12.2015 at TU Dresden*, 2015.
- [6] K. Rybina, W. Dargie, S. Umashankar, and A. Schill, "Modelling the live migration time of virtual machines," in *On the Move to Meaningful Internet Systems: OTM 2015 Conferences - Confederated International Conferences: CoopIS, ODBASE, and C&TC 2015, Rhodes, Greece, October 26-30, 2015, Proceedings*, 2015, pp. 575–593.
- [7] A. Strunk, "A lightweight model for estimating energy cost of live migration of virtual machines," in *2013 IEEE Sixth International Conference on Cloud Computing, Santa Clara, CA, USA, June 28 - July 3, 2013*, 2013, pp. 510–517.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] A. Kivity, Y. Kamay, D. Laor, U. Lublin, and A. Liguori, "kvm: the linux virtual machine monitor," in *Proceedings of the Linux Symposium*, vol. 1, Ottawa, Ontario, Canada, Jun. 2007, pp. 225–230.
- [12] M. Baron, "Probability and statistics for computer scientists, second edition," 2013.
- [13] R. I. Kabacoff, "R in action: Data analysis and graphics with r," 2011.
- [14] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2015. [Online]. Available: <http://www.R-project.org/>
- [15] J. W. Johnson, "A heuristic method for estimating the relative weight of predictor variables in multiple regression," *Multivariate Behavioral Research*, vol. 35, no. 1, pp. 1–19, 2000.