



UNIVERSITY OF LEEDS

This is a repository copy of *Towards virtual machine energy-aware cost prediction in clouds*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/124309/>

Version: Accepted Version

Proceedings Paper:

Aldossary, M, Alzamil, I and Djemame, K orcid.org/0000-0001-5811-5263 (2017) Towards virtual machine energy-aware cost prediction in clouds. In: Lecture Notes in Computer Science. GECON: International Conference on the Economics of Grids, Clouds, Systems, and Services, 19-21 Sep 2017, Biarritz, France. Springer , pp. 119-131. ISBN 9783319680651

https://doi.org/10.1007/978-3-319-68066-8_10

© Springer International Publishing AG 2017. This is an author produced version of a paper published in Lecture Notes in Computer Science. The final publication is available at Springer via http://dx.doi.org/10.1007/978-3-319-68066-8_10

Reuse

Unless indicated otherwise, fulltext items are protected by copyright with all rights reserved. The copyright exception in section 29 of the Copyright, Designs and Patents Act 1988 allows the making of a single copy solely for the purpose of non-commercial research or private study within the limits of fair dealing. The publisher or other rights-holder may allow further reproduction and re-use of this version - refer to the White Rose Research Online record for this item. Where records identify the publisher as the copyright holder, users can verify any specific terms of use on the publisher's website.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Towards Virtual Machine Energy-Aware Cost Prediction in Clouds

Mohammad Aldossary^{1,2}, Ibrahim Alzamil³, and Karim Djemame¹

¹ School of Computing, University of Leeds, Leeds, UK
{scmmald, K.Djemame}@leeds.ac.uk

² Prince Sattam Bin Abdulaziz University, KSA
MM.Aldossary@psau.edu.sa

³ Majmaah University, KSA
I.Alzamil@mu.edu.sa

Abstract. Pricing mechanisms employed by different service providers significantly influence the role of cloud computing within the IT industry. With the increasing cost of electricity, Cloud providers consider power consumption as one of the major cost factors to be maintained within their infrastructures. Consequently, modelling a new pricing mechanism that allow Cloud providers to determine the potential cost of resource usage and power consumption has attracted the attention of many researchers. Furthermore, predicting the future cost of Cloud services can help the service providers to offer the suitable services to the customers that meet their requirements. This paper introduces an Energy-Aware Cost Prediction Framework to estimate the total cost of Virtual Machines (VMs) by considering the resource usage and power consumption. The VMs' workload is firstly predicted based on an Autoregressive Integrated Moving Average (ARIMA) model. The power consumption is then predicted using regression models. The comparison between the predicted and actual results obtained in a real Cloud testbed shows that this framework is capable of predicting the workload, power consumption and total cost for different VMs with good prediction accuracy, e.g. with 0.06 absolute percentage error for the predicted total cost of the VM.

Keywords: Cloud Computing, Cost Prediction, Workload Prediction, ARIMA Model, Power Consumption, Energy Efficiency

1 Introduction

Cloud computing is an important and growing business model that has attracted the attention of many researchers. Pricing mechanisms that are employed by different service providers significantly influence the role of cloud computing within the IT industry. Billing mechanisms have become even more sophisticated, as customers are charged per month, hour or minute. Nevertheless, there are still limited as customers are charged based on a pre-defined tariff for the resource usage which include CPU, Memory, Storage and Network. This pre-defined tariff does not consider the variable cost of electricity [1]. Consequently, modelling a new pricing mechanism for services

offered that can be adjusted to the actual energy costs has become an interesting research topic.

There are limited works on cost models that measure the actual resource usage of a cloud service while taking consideration of variation in costs, power consumption, and performance together. Most cloud computing service providers charge their customers on a timely basis for the virtualised systems usage (with no performance guarantee) instead of the actual resource usage [3]. In other words, cloud service providers charge customers for the services offered on a timely basis, regardless of the actual resource usage and consideration of power consumption, which is considered one of the biggest operational cost factors by cloud infrastructure providers.

Another limitation of the cost mechanism is not only dependent on the actual resource usage and power consumption, but also on other factors that may affect the VMs total cost such as performance variation. Most of the existing studies have focused on minimising the power consumption and maximising the total resource usage, instead of improving VM performance. Further, Cloud providers (e.g. Amazon EC2) [4], have established their Service Level Agreements (SLAs) based on service availability without such an assurance of the performance. For instance, during the service operation, when the number of VMs increases on the same Physical Machine (PM)(overbooking), the resource competition may occur (e.g. once the workload exceeds the acceptable level of CPU utilisation) leading to VMs performance degradation. Thus, cloud service providers do not consider the VMs performance variation, while the VMs performance is a very important factor to satisfy cloud customers' requirements. Therefore, it is essential to consider VM performance variations in the composition of VM costs.

The first step towards this is an Energy-Aware Cost Prediction Framework that may influence the decision making of other problems. This paper focuses on the problem of estimating the resource usage, power consumption, and the total cost of the VMs at service operation. Therefore, a framework is proposed to predict VMs workload using an Autoregressive Integrated Moving Average (ARIMA) model. The relationship between the VMs and PMs workload (CPU utilisation) is investigated using regression models in order to estimate the VMs power consumption and predict the total cost of the VMs. This paper's main contributions are summarized as follow:

- A proposed Energy-Aware Cost Modeller for Cloud system architecture to assess the actual consumption of Cloud infrastructure resources.
- Energy-Aware Cost Prediction Framework that predicts the total cost for heterogeneous VMs by considering their resource usage and power consumption.
- Evaluation of the proposed framework in an existing Cloud testbed in order to verify the capability of the prediction models.

The remainder of this paper is organised as follows: a discussion of the related work is summarised in Section 2. Section 3 presents the system architecture followed by a discussion of the Energy-Aware Cost Prediction Framework. Section 4 presents the experimental set up followed by results and discussion in Section 5. Finally, Section 6 concludes this paper and discusses the future work.

2 Related Work

This paper discusses the cost that is associated with the resource usage and power consumption of the VMs. Previous work has looked into the area of calculating the cost of running services on Cloud infrastructure. Altmann and Kashef [13] presented the service placement optimisation based on the cost model in federated clouds to guarantee the cost minimisation for Cloud customers. This approach depends on a brute-force algorithm to evaluate the cost of each possible service placement. The cost model defined in their work as the sum of the fixed costs and the variable costs. The fixed costs include the costs for hardware and the variable costs include (e.g. the electricity cost). However, the cost model proposed in their work does not consider predicting the cost in the future. Also, more factors need to be considered (e.g. performance variation) to guarantee the SLAs. Horri and Dastghaibyfarid [8] emphasised the difficulty of dealing with minimising Cloud infrastructure energy consumption while conducting the Quality of Service (QoS), especially since there is a trade-off between energy consumption and SLA. Therefore, they have proposed and implemented a cost model in CloudSim. Their approach considers the total cost including the cost of energy consumption based on (e.g. number of VMs and data size). Nonetheless, their objectives do not consider predicting the total cost or power consumption.

In terms of prediction based on historical data, estimating the resources usage and power consumption of the VMs would require understanding the characteristics of the underlying physical resources, like idle power consumption and variable power under different workload, and the projected virtual resources usage, as stated in [20]. Thus, it is essential to get the predicted VMs' workload first in order to get their predicted power. Some work has predicted future workload in a Cloud environment based on Autoregressive Integrated Moving Average (ARIMA) model; nonetheless, their objectives do not consider predicting the power consumption. For example, Calheiros et al [24] introduced a Cloud workload prediction module based on the ARIMA model to proactively and dynamically provision resources. They define their workload as the expected number of requests received by the users, which are then mapped to predict the number of VMs needed to execute customers' requests and meet the QoS. Caron et al [11] presented a resource usage prediction algorithm based on identifying similar usage patterns of the short-term workload history. The algorithm has shown a good result within 4.8% prediction error. Khan et al [16] proposed a method of characterising and predicting workload based on Hidden Markov Modeling to discover the correlations between VMs workload that can be used to predict the changes of workload patterns. Further, Wood et al [12] focused on estimating the resource requirements when deploying an application into a virtualised environment using a regression-based model to predict future CPU utilisation. While the evaluation has shown that the prediction error is less than 5%, however these approaches do not consider the prediction of costs or power consumption of the VMs.

Other work focuses on predicting power consumption based on historical data while others use performance counters, which are queried directly from the hardware or the operating system. But, relying on performance counters would not work appropriately in heterogeneous environments with different server's characteristics, as argued by

Zhang et al [17]. Therefore, they presented a best fit energy prediction model (BFPEM) that flexibly selects the best model for a given server based on a series of equations that consider only CPU utilisation [17]. Dargie [18] proposed a stochastic model to estimate the power consumption for a multi-core processor based on the CPU utilisation workload and found out that the relationship between the workload and power is best estimated using a linear function in a dual-core processor and using a quadratic function in a single-core processor. Further, Fan et al [19] have introduced a framework to estimate the power consumption of servers based on CPU utilisation only and argued in their results that the power consumption correlates well with the CPU usage. As their framework produced accurate results, they argued that it is not necessary to use more complex signals, like hardware performance counters, to model power usage.

Compared with the work presented in this paper, the ARIMA model is used to predict the VMs workload, which is then mapped within the prediction framework to get the predicted VMs power consumption. Then, having predicted the VMs' workload and power consumption, the total cost of the VMs is predicted accordingly.

3 Resource Usage and Power Consumption for VMs

This section presents the proposed Energy-Aware Cost Prediction Framework to predict the resource usage, power consumption and total cost for VMs. The overall system architecture of this work will be discussed in the next subsection.

3.1 System Architecture

Cloud computing architecture consists of three standard layers, which are software as a service (SaaS), platform as a service (PaaS), and infrastructure as a service (IaaS). This paper will focus on the IaaS layer where the service operation takes place, as shown in Figure 1.

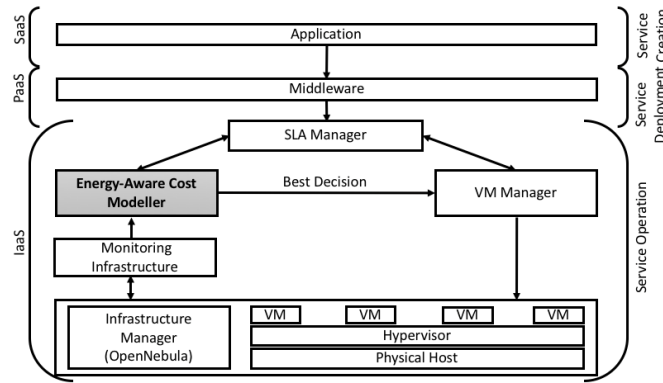


Fig. 1. System architecture

In the IaaS layer, the admission, allocation and management of VMs are performed through the interaction between the components. The highlighted component Energy-Aware Cost Modeller is the main focus of our work.

- **SLA Manager**: this component monitors and measures the SLA's agreed terms.
- **VM Manager**: considers the best decision in order to improve resource usage and reduce the power consumption cost and consequently the total cost of the VMs. For instance, if performance degrades, this component will have actuators to attempt to get the performance to the agreed level. This component interacts with the **Energy-Aware Cost Modeller** to request predictions related to the resource usage, power consumption and cost that VMs would have for a particular host.
- **Monitoring Infrastructure**: this will monitor resource usage, power consumption and performance related metrics.
- **Energy-Aware Cost Modeller**: this component supports:
 - 1) Energy-Aware Pricing Model that considers the actual resources and power consumption, as introduced in our previous work [5], and
 - 2) Energy-Aware Cost Prediction Framework that estimates the resource usage, power consumption and total cost for the VMs.

3.2 Energy-Aware Cost Prediction Framework

In our previous work [5], we introduced an **Energy-Aware Pricing Model** that considers power consumption as a key parameter with respect to performance and cost. The proposed model charges the customer based on the actual resource usage and considers the cost of power consumption of the VMs.

In this paper, we extend our work and introduce a new **Energy-Aware Cost Prediction Framework** that would predict VMs workload (CPU, RAM, Disk and Network), power consumption and total cost using the ARIMA model and regression models. This is the main focus of this paper as shown in Figure 2.

The ARIMA model is a time series prediction model that has been used widely in different domains, including finance, owing to its sophistication and accuracy; further details about the ARIMA model can be found in [14]. Unlike other prediction methods, like sample average, ARIMA takes multiple inputs as historical observations and outputs multiple future observations depicting the seasonal trend. It can be used for seasonal or non-seasonal time-series data. The type of seasonal ARIMA model is used in this work as the targeted workload patterns are reoccurring and showing seasonality in time intervals. In order to use the ARIMA model for predicting the VMs workload in this work, the historical time series workload data has to be stationary, otherwise Box and Cox transformation [15] and data differencing methods are used to make these data stationary. The model selection is based on the best fit model of ARIMA based on Akaike Information Criterion (AIC) or Bayesian Information Criterion (BIC) value.

This framework is aimed towards predicting the total cost of the VMs. In order to achieve that, the VMs workload is first predicted for the next time interval using the ARIMA model based on historical workload patterns. Then, the predicted VMs CPU utilisation is correlated with the PM CPU utilisation in order to predict the power consumption of PM, from which the VMs power consumption is estimated. Finally, the total cost for the VMs is predicted based on the predicted workload and power consumption of the VMs.

As depicted in Figure 2, the framework includes five main steps in order to predict the VMs workload and power consumption, then predict the total cost of VMs. To reach this goal, the following steps are required.




Fig. 2. Energy-Aware Cost Prediction Framework

Step 1: to predict (CPU, RAM, Disk and Network) utilisations for the next time interval, ARIMA model is used to identify the best fit model. After predicting the VM workload using the ARIMA model based on historical data, the next steps take place to predict the PM workload and the PM/VM power consumption using regression models.

Before predicting the power consumption for PM/VM, understanding how the resource usage affect the power consumption is required. Therefore, we did an experimental study to investigate the effect of the resource usage (CPU, RAM, Disk and Network) on the power consumption. The findings show that the CPU utilisation correlates well with the power consumption, as this finding is supported in other work [17-19].

Step 2: to predict the PM workload which is (PM CPU utilisation), would require measuring the relationship between the number of vCPU and the PM CPU utilisation for a single PM, as shown in Figure 3. This experiment was carried out on a local Cloud Testbed (see Section 4). Linear regression model has been applied to predict the PM CPU utilisation based on the used ratio of the requested number of vCPU for the VMs with consideration of its current workload as the PM may be running other VMs already [6]. The following equation is used (1):

$$PMx_{PredUtil} = \left(\alpha \times \left(\sum_{y=1}^{VMCount} (VMy_{ReqvCPUs} \times \frac{VMy_{PredUtil}}{100}) \right) + \beta \right) + (PM_{xCurrUtil} - PM_{xIdleUtil}) \quad (1)$$

$PMx_{PredUtil}$ is the predicted PM CPU utilisation ; α is the slope and β is the intercept of the CPU utilisation. The $VMy_{ReqvCPUs}$ is the number of requested vCPU for each VM and $VMy_{PredUtil}$ is the predicted utilisation for each VMs. The $PM_{xCurrUtil}$ is the current PM utilisation and $PM_{xIdleUtil}$ is the idle PM utilisation.




Fig. 3. Number of vCPUs vs PM CPU Utilisation.




Fig. 4. The PM CPU Utilisation vs Power Consumption.

Step 3: the PM power consumption is predicted based on the relationship between the predicted PM workload (PM CPU utilisation) with PM power consumption on the same PM. Using a regression analysis, the relation is best described using polynomial model with order three for this particular PM, as shown in Figure 4. Thus, the predicted PM power consumption $PMx_{PredPwr}$ measured by Watt, can be identified using the following formula (2).

$$PMx_{PredPwr} = (\alpha(PMx_{PredUtil})^3 + \gamma(PMx_{PredUtil})^2 + \delta(PMx_{PredUtil}) + \beta) \quad (2)$$

Where α , γ and δ are all slopes, β is the intercept and $PMx_{PredUtil}$ is predicted PM CPU utilization.

Step 4: based on the requested number of vCPU and the predicted vCPU utilisation, the VM power consumption is predicted using the proposed formula in [6], as shown in equation (3).

$$VMx_{Predpwr} = PMx_{IdlePwr} \times \left(\frac{VMx_{ReqvCPUs}}{\sum_{y=1}^{VMcount} VMY_{ReqvCPUs}} \right) + (PMx_{PredPwr} - PMx_{IdlePwr}) \times \left(\frac{VMx_{(PredUtil+ReqvCPUs)}}{\sum_{y=1}^{VMcount} VMY_{(PredUtil+ReqvCPUs)}} \right) \quad (3)$$

Where $VMx_{Predpwr}$ is the predicted power consumption for one VM measured by Watt. $VMx_{ReqvCPUs}$ is the requested number of vCPU and $VMx_{PredUtil}$ is the predicted VM CPU utilisation. $\sum_{y=1}^{VMcount} VMY_{ReqvCPUs}$ is the total of vCPU for all VMs in the same PM. The $PMx_{IdlePwr}$ is idle power consumption and $PMx_{PredPwr}$ is the predicted power consumption for a single PM.

Step 5: finally, this step predicts the total cost for the VM based on the predicted VM resource usage from step 1 and the predicted VM power consumption from step 4. The energy providers usually charge by the Kilowatt per hour (kWh). Therefore, convert the power consumption to energy is required using the following equation (4):

$$VMx_{PredEnergy} = \frac{VMx_{AvgPredpwr}}{1000} \times \frac{Time_s}{3600} \quad (4)$$

To predict the total cost for the VM using the proposed model, as shown in equation (5):

$$VMx_{PredTotalCost} = \left(\left(VMx_{ReqvCPUs} \times \frac{VMx_{PredUtil}}{100} \right) \times (Cost\ per\ vCPU \times Time_s) \right) + \left(VMx_{PredRAMUsage} \times (Cost\ per\ GB \times Time_s) \right) + \left(VMx_{PredDiskUsage} \times (Cost\ per\ GB \times Time_s) \right) + \left(VMx_{PredNetUsage} \times (Cost\ per\ GB \times Time_s) \right) + \left(VMx_{PredEnergy} \times Cost\ per\ kWh \right) \quad (5)$$

Where $VMx_{PredTotalCost}$ is the predicted total cost of the VM. $VMx_{PredRAMUsage}$ is the predicted resource usage of RAM times the cost for that resource for a period of time and so on for each resource such as CPU, Disk and Network. $VMx_{PredEnergy}$ is the predicted energy consumption of the VM times the electricity price as announced by the energy providers.

4 Experimental Set Up

This section describes the environment and the details of the experiments conducted in order to evaluate the work presented in this paper.

In terms of the experimental design, the aim is to evaluate the new Energy-Aware Cost Prediction Framework presented in terms of predicting the workload, power consumption and total cost for heterogeneous VMs based on historical periodic workload. The prediction process starts by firstly predicting the VM workload using the (auto.arima) function in R package [25] and then completing the cycle of the framework and considering the correlation between the physical and virtual resources to predict power consumption of the VMs on a single PM. After that, the total cost is predicted for the VMs based on their predicted workload and power consumption.

A number of experiments have been designed and implemented on a local Cloud Testbed with the support of the Virtual Infrastructure Manager (VIM), OpenNebula [7] version 4.10, and KVM hypervisor for the Virtual Machine Manager (VMM). This Cloud Testbed includes a cluster of commodity Dell servers, and one of these servers with eight core E31230 V2 Intel Xeon CPU was used. The server includes 16GB RAM and 1000GB hard drives. Also, the server has a WattsUp meter [9] attached to directly measure the power consumption. Heterogeneous VMs are created and their monitoring is performed through Zabbix [10], which is also used for resources usage monitoring purposes. Rackspace [26] is used as a reference for the VMs configurations. Three types of VMs, small, medium and large are provided with different capacities. The VMs are allocated with 1, 2 and 4 vCPUs, 1, 2 and 4 GB RAM, 10 GB Disk and 1 GB Network, respectively. In terms of the cost of the virtual resources, ElasticHosts [27] and VMware [28] prices are followed: where 1 vCPU = £0.008/hr, 1 GB Memory = £0.016/hr, 1 GB Storage = £0.0001/hr, 1 GB Network = £0.0001/hr; and the cost of Energy = £0.14/kWh [21].

In terms of the workload patterns, Cloud applications can experience different workload patterns based on the customers' usage behaviours, and these workload patterns consume power differently based on the resources they utilise. There are several workload patterns, such as static, periodic, continuously changing, unpredicted, and once-in-a-life-time, as stated in [23]. This paper considers the periodic workload pattern as this work is driven towards solving the issue of the performance variation.

Thus, a number of direct experiments have been conducted to synthetically generate periodic workload by using Stress-ng [2] tool in order to stress all resources (CPU, RAM, Disk and Network) on different types of VMs. The generated workload of each VM type has four time intervals of 30 minutes each. The first three intervals will be used as the historical data set for prediction, and the last interval will be used as the testing data set to evaluate the predicted results.

5 Results and Discussion

This section presents the evaluation of the Energy-Aware Cost Prediction Framework. The figures below show the predicted results for three types of VMs, small, medium and large, running on a single PM based on historical periodic workload pattern. Because of space limitation, only large VMs results are shown. As mentioned earlier, the

generated VMs workload along with their power consumption and cost for the last interval are used as the testing data set.

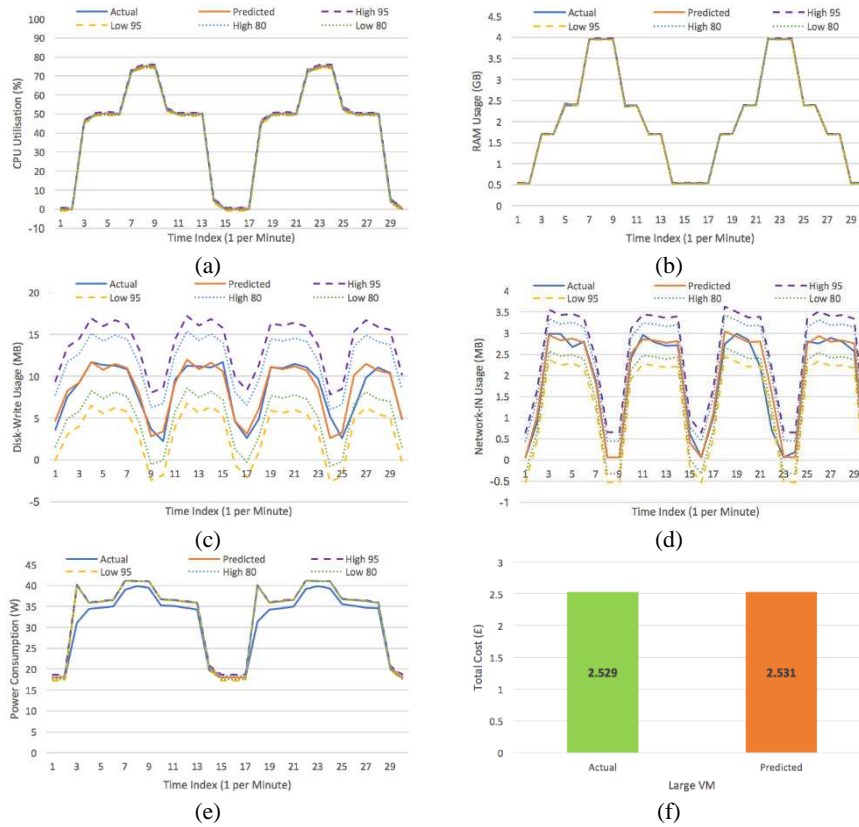


Fig. 5. The prediction Results for a Large VM.

Figure 5 (a, b, c and d) depict the results of the predicted versus the actual VMs workload, including CPU, RAM, Disk, and Network usage for the VMs. Despite the periodic utilisation peaks, the predicted VMs’ CPU and RAM workload results closely match the actual results, which reflects the capability of the ARIMA model to capture the historical seasonal trend and give a very accurate prediction accordingly. The predicted VMs’ Disk and Network workload is also matching the actual workload, but with less accuracy as compared to the CPU and RAM prediction results. This can be justified because of the high variations in the generated historical periodic workload pattern of the disk and network not closely matching in each interval, whereas the generated historical periodic workload pattern for the RAM and CPU usage are closely matched in each interval. Beside the predicted mean values, the figures also show the high and low 95% and 80% confidence intervals.

The proposed framework can predict the power consumption for a number of VMs with only a small variation as compared to the actual one as shown in Figure 5 (e). The predicted power consumption attribution for each VM is affected by the variation in the

predicted CPU utilisation of all the VMs, hence the predicted power consumption of the medium VM matches its predicted CPU utilisation as it has the highest variation than the other predicted VMs' CPU utilisation.

In terms of prediction accuracy, a number of metrics have been used to evaluate the results. These metrics include, Absolute Percentage Error (APE) which measures the absolute value of the ratio of the error to the actual observed value; Mean Error (ME) which measures the average error of the predicted values; Root Mean Squared Error (RMSE) which depicts the square root of the variance measured by the mean absolute error; Mean Absolute Error (MAE) is the average of the absolute value of the difference between predicted value and the actual value; Mean Percentage Error (MPE) is the computed average of percentage errors by which the predicted values vary from the actual values; and Mean Absolute Percent Error (MAPE) is the average of the absolute value of the difference between the predicted value and the actual value explained as a percentage of the actual value [22].

Table 1. Prediction Accuracy for a Large VM.

Parameters	ME	RMSE	MAE	MPE	MAPE
CPU Utilisation	0.03765	0.299769	0.137823	0.309809	6.615192
RAM Usage	0.000004	0.008671	0.002587	-0.00675	0.107601
Disk-Write Usage	0.1838898	1.116114	0.733408	0.924781	12.64005
Network-IN Usage	0.0657477	0.225631	0.132185	-6.13982	17.56377
Power Consumption	1.648176	2.617798	1.648176	4.358135	4.358135

This framework is also capable of predicting the total cost for a number of VMs as shown in Figure 5 (f), with 0.06 of APE for predicted total cost of the large VM, 17.23 of APE for the medium VM and 14.7 of APE for the small VM as shown in Figure (6).

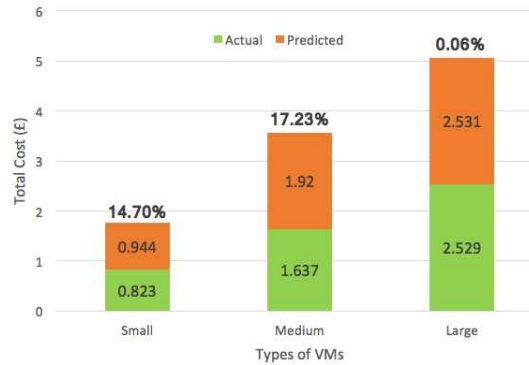


Fig. 6. The predicted versus the actual VMs total cost.

The accuracy of the predicted VMs workload (CPU, RAM, Disk, Network) and their power consumption based on periodic workload is evaluated using these accuracy metrics, as summarised in Table 1. In addition, Figure (6) shows the results of the predicted versus the actual total cost for all VMs with the absolute percent error for the predicted total cost. Despite the high variation of the workload utilisation in the periodic pattern,

the accuracy metrics indicate that the predicted VMs workload and power consumption achieve good prediction accuracy along with the predicted total cost.

6 Conclusion and Future Work

This paper has presented and evaluated a new Energy-Aware Cost Prediction Framework that predicts the total cost of VMs by considering the resource usage and power consumption of heterogeneous VMs based on their usage and size, which reflect the physical resource usage by each VM. A number of direct experiments were conducted on a local Cloud Testbed to evaluate the capability of the prediction models. Overall, the results show that the proposed Energy-Aware Cost Prediction Framework can predict the resource usage, power consumption and the total cost for the VMs with a good prediction accuracy based on periodic Cloud workload patterns.

Unlike other existing works, this approach considers the heterogeneity of VMs with respect to predicting the resource usage, power consumption and the total cost.

In future work, we intend to extend our approach and integrate it with performance prediction models to determine the costs of different scenarios. Besides, further investigation will focus on VM performance prediction models, dynamic placement of VMs, and demonstration of the trade-off between cost, power consumption and performance. Also, the scalability aspects with different prediction algorithms will be considered to further show the capability of the proposed work. Finally, as this paper has focused on predicting the VMs total cost based on periodic workload pattern, we aim to extend this by considering other workload patterns, such as static, continuously changing, unpredicted, and once-in-a-life time workload patterns.

References

1. A. Narayan and S. Rao, "Power-Aware Cloud Metering," *IEEE Transactions on Services Computing*, vol. 7, no. 3, pp. 440-451, 2014.
2. Stress-ng, [Online]. Available: <http://kernel.ubuntu.com/~cking/stress-ng/>. [Accessed: 19-Feb-2017].
3. P. Berndt and A. Maier, "Towards Sustainable IaaS Pricing," in *Economics of Grids, Clouds, Systems, and Services SE - 13*, vol. 8193, J. Altmann, K. Vanmechelen, and O. Rana, Eds. Springer International Publishing, 2013, pp. 173-184.
4. Amazon EC2 Service Level Agreement, [Online]. Available: <https://aws.amazon.com/ec2/sla/>. [Accessed: 26-Feb-2017].
5. M. Aldossary and K. Djemame, "Energy Consumption-based Pricing Model for Cloud Computing," 32nd UK Perform. Eng. Work. UKPEW'2016, Bradford, UK, 8-9 Sept., pp. 16-27, 2016.
6. I. Alzamil and K. Djemame, "Energy Prediction for Cloud Workload Patterns," *Proceedings of the 13th International Conference on Economics of Grids, Clouds, Systems and Services (GECON'2016)*, Athens, Greece, September 2016. [LNCS 10382, pp. 113-127](#)
7. "Sunstone Cloud Testbed," OpenNebula.org, [Online]. Available: <http://opennebula.org/>. [Accessed: 20-Feb-2017].
8. A. Horri and G. Dastghaibyfar, "A novel cost based model for energy consumption in cloud computing," *The Scientific World Journal*, vol. 2015, p. 724524, 2015.

9. "Watt's Up Power Meter," watts up?, [Online]. Available: <https://www.wattsupmeters.com/secure/products.php?pn=0>. [Accessed: 20-Feb-2017].
10. "Zabbix Monitoring," Zabbix, [Online]. Available: <http://www.zabbix.com/>. [Accessed: 20-Feb-2017].
11. E. Caron, F. Desprez, and A. Muresan, "Forecasting for grid and cloud computing on-demand resources based on pattern matching," Proc. - 2nd IEEE Int. Conf. Cloud Comput. Technol. Sci. CloudCom 2010, pp. 456–463, 2010.
12. T. Wood, L. Cherkasova, K. Ozonat, and P. Shenoy, "Profiling and Modeling Resource Usage of Virtualized Applications," pp. 366–387, 2008.
13. Altmann, J. and Kashef, M. M. "Cost model based service placement in federated hybrid clouds". *Futur. Gener. Comput. Syst.*, vol. 41, no. 1, pp. 79-90, 2014.
14. Box GEP, Jenkins GM, Reinsel GC. *Time series analysis: forecasting and control*. 4th ed. Hoboken, NJ: John Wiley & Sons, 2008.
15. Box, G.E.P., Cox, D.R.: An analysis of transformations. *J. R. Stat. Soc. Ser. B.* 211–252, 1964.
16. A. Khan, X. Yan, S. Tao, and N. Anerousis, "Workload Characterization and Prediction in the Cloud: A Multiple Time Series Approach", in Proc. of the IEEE Network Operations and Management Symposium (NOMS), Maui, HI, Apr. 16-20, 2012.
17. Zhang, X., Lu, J., Qin, X.: BFEPM: "Best Fit Energy Prediction Modeling Based on CPU Utilization", 2013 IEEE Eighth Int. Conf. Networking, Archit. Storage. 41–49, 2013.
18. Dargie, W.: "A stochastic model for estimating the power consumption of a processor.", *IEEE Transactions on Computers.* 1311-1322., 2015.
19. Fan, X., Weber, W.-D., Barroso, L.A.: "Power Provisioning for a Warehouse-sized Computer", In: Proceedings of the 34th Annual International Symposium on Computer Architecture. pp. 13–23. ACM, New York, NY, USA, 2007.
20. R. Kavanagh, D. Armstrong, K. Djemame, D. Sommaccampagna, L. Blasi. "Towards an Energy-Aware Cloud Architecture for Smart Grids", In Proceedings of the 12th International Conference on Economics of Grids, Clouds, Systems and Services (GECON'2015), Cluj-Napoca, Romania, September 2015.
21. Electricity Price Electricity Price per kWh Comparison of Big Six Energy Companies - CompareMySolar.co.uk Blog. [Online]. Available: <http://blog.comparemysolar.co.uk/electricity-price-per-kwh-comparison-of-big-six-energy-companies/> [Accessed: 16-Feb-2017].
22. R. J. Hyndman and G. Athanasopoulos, "Forecasting: Principles and Practice", OTexts, 2017. [Online]. Available: <https://www.otexts.org/fpp/2/5/>. [Accessed: 16-Feb-2017].
23. Fehling, C., Leymann, F., Retter, R., Schupeck, W., Arbitter, P.: *Cloud computing patterns*. Springer (2014).
24. Calheiros, R.N., Masoumi, E., Ranjan, R., Buyya, R.: Workload prediction using ARIMA model and its impact on cloud applications' QoS. *IEEE Trans. Cloud Comput.* 3, 449–458, 2015.
25. R Core Team: *R: A Language and Environment for Statistical Computing*, [Online]. Available: <https://www.r-project.org/>. [Accessed: 19-Feb-2017].
26. Rackspace, *Cloud Servers Pricing and Cloud Server Costs*, [Online]. Available: <http://www.rackspace.co.uk/cloud/servers/pricing>. [Accessed: 20-Feb-2017].
27. ElasticHosts, *Pricing - ElasticHosts Linux, Windows VPS Hosting*, [Online]. Available: <https://www.elastichosts.co.uk/pricing/>. [Accessed: 16-Feb-2017].
28. VMware - *OnDemand Pricing Calculator*, [Online]. Available: <http://vcloud.vmware.com/uk/service-offering/pricing-calculator/on-demand/>. [Accessed: 16-Feb-2017].