



Alshammari, D., Singer, J. and Storer, T. (2017) Does CloudSim Accurately Model Micro Datacenters? In: IEEE International Conference on Cloud Computing (IEEE CLOUD 2017), Honolulu, HI, USA, 25-30 June 2017, pp. 705-709. (doi:[10.1109/CLOUD.2017.97](https://doi.org/10.1109/CLOUD.2017.97))

This is the author's final accepted version.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/141245/>

Deposited on: 18 May 2017

Enlighten – Research publications by members of the University of Glasgow
<http://eprints.gla.ac.uk>

Does CloudSim Accurately Model Micro Datacenters?

Dhahi Alshammari, Jeremy Singer, Timothy Storer

School of Computing Science

University of Glasgow, UK

d.alshammari.1@research.gla.ac.uk, jeremy.singer@glasgow.ac.uk, timothy.storer@glasgow.ac.uk

Abstract—Novel cloud computing algorithms and techniques are initially evaluated via testbeds, simulators and mathematical models of datacenter infrastructure. However, it can be difficult to perform cross validation of these platforms against realistic scale infrastructures due to the prohibitive costs involved. This paper describes an approach to evaluating a cloud simulator through an empirical study involving a micro datacenter of commodity Raspberry Pi devices. To demonstrate the methodology, we compare performance of real-world workloads on this physical infrastructure against corresponding models of the workloads and infrastructure on the CloudSim simulator. After modelling a Raspberry Pi micro datacenter in CloudSim, we claim that the simulator lacks sufficient accuracy for cloud infrastructure experiments.

Keywords-CloudSim; Glasgow Raspberry Pi Cloud; simulation performance; evaluation;

I. INTRODUCTION

Cloud computing is widely used in many areas of information technology, as it effectively reduces costs and maximises benefits for many contemporary businesses [1] [2] [3]. However, cloud computing features inherent challenges, such as security, cost modelling, energy management, and virtual machine migration. Researchers conduct experiments to address cloud computing issues using either physical infrastructures, simulators or mathematical models [4], [5]. Simulators such as CloudSim [6], iCanCloud [7], and Green Cloud [8] generally possess the advantage of easy experiment repetition and minimal infrastructure costs. However simulators have significant limitations in simulating advanced features of large scale datacenters. CloudSim [6] simulates and models the workload of the cloud using object-oriented Java classes. CloudSim is designed to simulate different cloud layers (software as a service, platform as a service, and infrastructure as a service). Resource allocation management is the most common aspect of cloud that has been explored using CloudSim [4]. However, according to [9], large scale systems such as cloud computing infrastructure are difficult to simulate due to advanced datacenter features such as large buffers and cache memory. This paper evaluates the CloudSim simulator to assess the accuracy of simulated results against the measured performance of a micro datacenter.

This work develops a new method to check the fidelity of any simulation models to be used for performance predicting of any realistic infrastructure. We accept that cross-

Table I: Features of the Glasgow Raspberry Pi Cloud

Nodes	1.2GHz 64bit quad core ARMv8 CPU, 4 cores: Raspberry Pi v3.
	900MHz quad core ARM CortexA7 CPU, 4 cores: Raspberry Pi v2.
	1GB RAM : Each node
	16GB SD Cards
Network	Network Topology: Tree
	Standard ethernet switches
	Bandwidth: 100 Mbits/sec

validation of simulator with large scale infrastructure is difficult; hence we perform ‘sanity check’ cross-validation of simulator with cut-down micro scale infrastructure. If there are discrepancies between the simulator and the micro infrastructure, then we know the simulator is inaccurate.

In this work, we compare the results of the actual performance by running workloads on a Raspberry Pi datacenter [10] with a simulated performance of the model sharing the same features of the workloads and infrastructure. We gathered the low level properties of the Raspberry Pi datacenter via standard Linux monitoring tools such as perf [11] and iperf [12] to be fed into CloudSim as parameters for the model of the Raspberry Pi Cloud infrastructure.

The remainder of this paper is organized as follows: we describe our the experiment and benchmarks in Sections II and III. Section IV shows the related work to clarify the difference between our work and other works in the literature. The conclusion and possible future work of this study is described in Section V.

II. WORKLOADS AND PROFILING TOOLS

We modelled the Glasgow Raspberry Pi Cloud [10] in the CloudSim framework by customizing the CloudSim Java classes that characterize the compute nodes, network connectivity and virtual machine environments of the datacenter instance to be simulated.

The features of our Raspberry Pi Cloud are given in table I. These are the features we use to customize our CloudSim model.

A. Benchmarks

In this paper we are providing two different types of application (batch and transactional processing). We applied different workloads and tested on the Raspberry Pi cluster

due to their availability to be used in experiments, possibility to be profiled on Raspberry Pi devices and the possibility to be modeled directly on Cloudsim without any modifications on CloudSim or its extensions. Therefore, we use existing benchmarks for message passing interface (MPI) applications and a web server benchmark called Apache Bench (AB). These workloads are modelled in CloudSim v3.0.3 with appropriate parameterization.

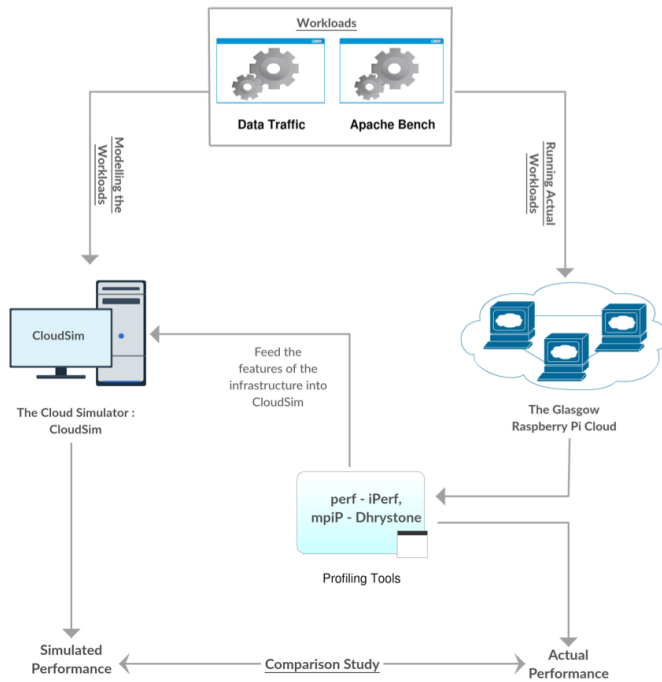


Figure 1: Cross-validate Performance Metrics for CloudSim and The Glasgow Raspberry Pi Micro Datacenter

The benchmarks run on Raspberry Pi v2 and Raspberry Pi v3 in different number of parallel nodes between effectively one rack of the micro datacenter. The experiments were repeated thirty times on each configuration.

B. Profiling Tools: Perf, iperf and mpiP

A bash script including perf [11] was implemented to generally profile any workloads run on the Raspberry Pi cloud. Each perf record is based on hardware events generated by a specific Linux process. iperf was also used [12] in order to quantify network features such as latency, throughput and bandwidth. iperf is a tool for active measurement of the maximum achievable bandwidth on IP networks. It supports the tuning of various parameters related to timing, buffers and protocols (TCP, UDP, SCTP with IPv4 and IPv6). For each test, the bandwidth as well as other specific parameters are reported.

Moreover, we have utilised mpiP [13], a profiling tool, in order to profile cluster performance by measuring the execution time of the jobs on each node in the cluster. Other

tools could be utilised to profile the cluster, for example MPIInside, however we have selected mpiP as it is open-source and free to use. mpiP is a lightweight profiling library for MPI applications that collects statistical information regarding MPI function-calls.

Figure 2 shows the profiling tools we used and a load balancing tool called HAproxy [14].

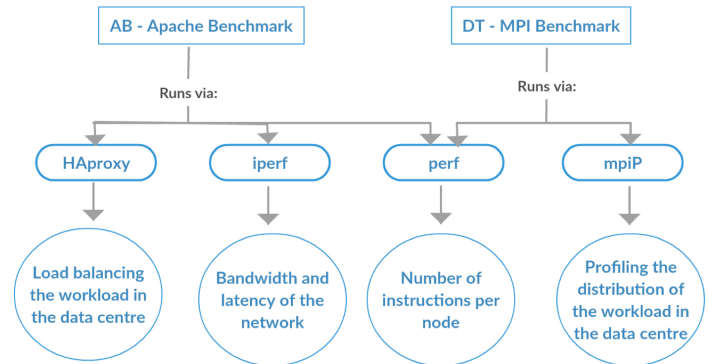


Figure 2: Cross-validate Performance Metrics for CloudSim and The Glasgow Raspberry Pi Micro Datacenter

III. EXPERIMENT AND EVALUATION

A. Evaluation of MPI Application

CloudSim is designed to model the MPI application and the communication between tasks in the datacenter as it contains the following Java classes:

- 1) Taskstage represents various stages a networkCloudlet can have during execution.
- 2) AppCloudlet class to represent an application which user submit for execution within datacenter. It consist of several networkClouds.
- 3) Workflowapp is also a class of AppCloudlet having communicating tasks [6].

We measure the execution time for the NAS MPI Data Traffic benchmark (DT) [15], Class:A on the actual testbed. We have chosen and specified the DT benchmark as the other benchmarks from NAS require excessively large memory footprints to run on Raspberry Pi 2 and 3. We tested different sizes of high-performance computing applications on the cluster in order to gather execution times for each application on a particular node. We obtained the output of each application using mpiP [13] to generate the profiling file for each benchmark running time on the cluster in order to be compared with CloudSim.

Figures 3 and 4 show the execution time on Raspberry Pi v2 and v3 and the simulated execution time calculated by CloudSim. The points on the lines represents the mean of the execution time of running the benchmark (DT) 30 times on Raspberry Pi nodes and the simulated execution time the model of the benchmark on CloudSim in tandem with the increasing number of nodes.




Figure 3: Actual and simulated performance of DT workload on Raspberry Pi 2 and CloudSim




Figure 5: Actual and simulated performance of AB workload on Raspberry Pi 2 and CloudSim

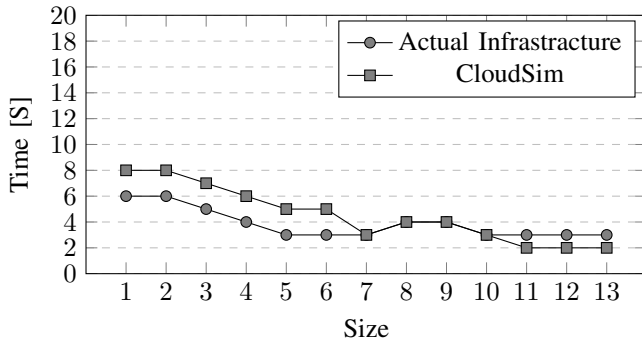


Figure 4: Actual and simulated performance of DT workload on Raspberry Pi 3 and CloudSim

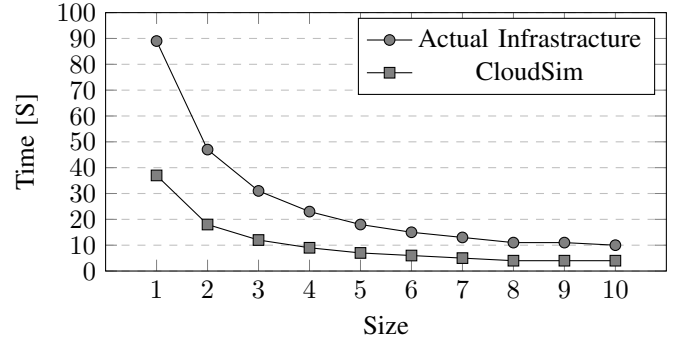


Figure 6: Actual and simulated performance of AB workload on Raspberry Pi3 and CloudSim

B. Evaluation of Webserver Application

In order to run Apache Bench (AB) [16] on our Raspberry Pi datacenter, we used an external Ubuntu machine as a client for the web servers. We also used HAProxy [14] in our datacenter as a load balancer to distribute the requests between the Raspberry Pi nodes. Using HAProxy allowed us to expand the number of nodes serving the request from the client. Figures 5 and 6 show the performance of the benchmark AB on both actual tests, which were on Glasgow Raspberry Pi Cloud and the simulated environment of CloudSim.

Table II shows the errors for the simulated environment on

Table II: Root mean square (RMS) for absolute and relative error of the simulated performance on CloudSim

DT Benchmark - MPI Workload		
	Raspberry Pi2	Raspberry Pi3
RMS for Absolute Error	1.46	1.36
RMS for Relative Error	32.03 %	17.95 %
AB Benchmark - Web Server Workload		
	Raspberry Pi2	Raspberry Pi3
RMS for Absolute Error	56.81	21.10
RMS for Relative Error	75.80 %	60.96 %

CloudSim with respect to the actual execution time on physical Raspberry Pi infrastructure. We have applied the root mean square (RMS) since the error is not consistently positive or negative. We report the mean absolute error which measures the disagreement between CloudSim performance and Raspberry Pi datacenter performance in seconds for all different cases experiments [17]. Moreover the relative error which show the percentages of the differences between the selected data sets is shown in the table.

IV. RELATED WORK

Antonescu and Braun [18] propose a validation methodology for modelling complex distribution systems using CloudSim [6]. The validation experiment is conducted with a small-scale distributed testbed. However, the comparative precision of the performance for CloudSim and the small-scale distributed test bed is not reported. Moreover, the small scale datacenter used in the experiment is not clearly described; limiting reproducibility and failing to quantify the level of accuracy. Our work compares the performance simulation workload with a replicable micro datacenter known as the Glasgow Raspberry Pi Cloud. NetworkCloudSim [19] is an extension of CloudSim that

models the concept of networking communication, it has been developed and integrated with CloudSim version 3.0.3, and it models Message Passing Interface (MPI). Our work attempts to validate CloudSim by comparing the result of running MPI benchmark DT on either a single node or multiple nodes of Raspberry Pi. Iridis-pi [20] is a small scale Beowulf cluster built using 64 Raspberry Pi Model B nodes with 700 MHz ARM processors. They give a brief description of the clustering distribution system in implementing Hadoop job on their datacenter, but our work is to profile and report the performance of the benchmarks on two different clusters of Raspberry Pi v2 and v3.

Micro datacenters have been used in different areas of the field of cloud computing for many years. For example, work in [21] focused on the realization of a micro datacenter built using Raspberry Pi devices. That work did not, however, profile the performance of the datacenter, which is the focus of this piece; providing information about the performance of the datacenter will allow further research to investigate the introduction of new techniques or algorithms to the field of cloud computing by providing a baseline for such effects to be measured from. However, the workloads that run on this infrastructure have not been fully explained and have not been profiled to demonstrate the Raspberry Pi cloud performance. Work on [22] raises the possibility of using a micro datacenter alongside a centralised datacenter to reduce energy consumption by reducing the number of hops between end users and the centralized datacenter. This work was conducted using Raspberry Pi nodes running a specialized Internet of Things (IoT) application.

V. CONCLUSION AND FUTURE WORK

This work presents an empirical approach to evaluate the software simulation tool CloudSim [6] through a comparative study with a micro datacenter known as the Glasgow Raspberry Pi Cloud [10]. Two workloads have been profiled on a per-node basis using standard Linux performance monitoring tools. We make the following two suggestions regarding CloudSim experiments:

- 1) CloudSim needs a *richer set of input features* to calculate workload execution time. Currently, CloudSim requires number of instructions for the workload, and an instruction processing rate for the CPU. Our results show that there is clear execution time dependence on the CPU architecture, based on the discrepancies between CloudSim results for datacenter nodes featuring Raspberry Pi v2 and Raspberry Pi v3.
- 2) CloudSim needs a *more complex model of inter-node communication* for distributed applications. Our experiments show high relative error between the actual and simulated performance for a non-batch-processing workload—which is more likely to be a realistic cloud application.

Finally, we note that while this paper focuses exclusively on CloudSim, the proposed cross-validation approach can be applied to any simulator.

REFERENCES

- [1] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, “A view of cloud computing,” *Commun. ACM*, vol. 53, no. 4, pp. 50–58, Apr. 2010. [Online]. Available: <http://doi.acm.org/10.1145/1721654.1721672>
- [2] P. Mell and T. Grance, “The nist definition of cloud computing,” Computer Security Division, Information Technology Laboratory, National Institute of Standards and Technology Gaithersburg, Tech. Rep., 2011.
- [3] S. Marston, Z. Li, S. Bandyopadhyay, J. Zhang, and A. Ghalasasi, “Cloud computing—the business perspective,” *Decision support systems*, vol. 51, no. 1, pp. 176–189, 2011.
- [4] A. Ahmed and A. Sabyasachi, “Cloud computing simulators: A detailed survey and future direction,” in *Advance Computing Conference (IACC), 2014 IEEE International*, Feb 2014, pp. 866–872.
- [5] A. Barker, B. Varghese, J. S. Ward, and I. Sommerville, “Academic cloud computing research: Five pitfalls and five opportunities,” in *Proceedings of the 6th USENIX Conference on Hot Topics in Cloud Computing*, ser. HotCloud’14. Berkeley, CA, USA: USENIX Association, 2014, pp. 2–2. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2696535.2696537>
- [6] R. N. Calheiros, R. Ranjan, A. Beloglazov, C. A. F. De Rose, and R. Buyya, “Cloudsim: A toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms,” *Softw. Pract. Exper.*, vol. 41, no. 1, pp. 23–50, Jan. 2011. [Online]. Available: <http://dx.doi.org/10.1002/spe.995>
- [7] A. Núñez, J. L. Vázquez-Poletti, A. C. Caminero, G. G. Castañé, J. Carretero, and I. M. Llorente, “icancloud: A flexible and scalable cloud infrastructure simulator,” *J. Grid Comput.*, vol. 10, no. 1, pp. 185–209, Mar. 2012. [Online]. Available: <http://dx.doi.org/10.1007/s10723-012-9208-5>
- [8] D. Kliazovich, P. Bouvry, Y. Audzevich, and S. Khan, “Greencloud: A packet-level simulator of energy-aware cloud computing data centers,” in *Global Telecommunications Conference (GLOBECOM 2010), 2010 IEEE*, Dec 2010, pp. 1–5.
- [9] J. Gustedt, E. Jeannot, and M. Quinson, “Experimental methodologies for large-scale systems: A survey,” *Parallel Processing Letters*, vol. 19, no. 03, pp. 399–418, 2009. [Online]. Available: <http://www.worldscientific.com/doi/abs/10.1142/S0129626409000304>
- [10] F. P. Tso, D. White, S. Jouet, J. Singer, and D. Pezaros, “The glasgow raspberry pi cloud: A scale model for cloud computing infrastructures,” in *Distributed Computing Systems Workshops (ICDCSW), 2013 IEEE 33rd International Conference on*, July 2013, pp. 108–112.

- [11] “PERF tutorial: Finding execution hot spots,” <http://sandsoftwaresound.net/perf/perf-tutorial-hot-spots/>, accessed: 2015-04-12.
- [12] “iPerf - The network bandwidth measurement tool ,” <http://iperf.fr/>, accessed: 2016-05-12.
- [13] “mpiP: Lightweight, Scalable MPI Profiling,” <http://mpip.sourceforge.net/>, accessed: 2016-05-22.
- [14] W. Tarreau, “Haproxy-the reliable, high-performance tcp/http load balancer,” 2012.
- [15] D. H. Bailey, E. Barszcz, J. T. Barton, D. S. Browning, R. L. Carter, L. Dagum, R. A. Fatoohi, P. O. Frederickson, T. A. Lasinski, R. S. Schreiber *et al.*, “The nas parallel benchmarks,” *International Journal of High Performance Computing Applications*, vol. 5, no. 3, pp. 63–73, 1991.
- [16] “ab - Apache HTTP server benchmarking tool ,” <https://httpd.apache.org/docs/2.4/programs/ab.html>, accessed: 2016-05-22.
- [17] C. J. Willmott and K. Matsuura, “Advantages of the mean absolute error (mae) over the root mean square error (rmse) in assessing average model performance,” *Climate research*, vol. 30, no. 1, pp. 79–82, 2005.
- [18] A.-F. Antonescu and T. Braun, “Modeling and simulation of concurrent workload processing in cloud-distributed enterprise information systems,” in *Proceedings of the 2014 ACM SIGCOMM Workshop on Distributed Cloud Computing*, ser. DCC ’14. New York, NY, USA: ACM, 2014, pp. 11–16. [Online]. Available: <http://doi.acm.org/10.1145/2627566.2627575>
- [19] S. Garg and R. Buyya, “Networkcloudsim: Modelling parallel applications in cloud simulations,” in *Utility and Cloud Computing (UCC), 2011 Fourth IEEE International Conference on*, Dec 2011, pp. 105–113.
- [20] S. J. Cox, J. T. Cox, R. P. Boardman, S. J. Johnston, M. Scott, and N. S. Obrien, “Iridis-pi: a low-cost, compact demonstration cluster,” *Cluster Computing*, vol. 17, no. 2, pp. 349–358, 2014.
- [21] P. Abrahamsson, S. Helmer, N. Phaphoom, L. Nicolodi, N. Preda, L. Miori, M. Angriman, J. Rikkila, X. Wang, K. Hamily *et al.*, “Affordable and energy-efficient cloud computing clusters: the bolzano raspberry pi cloud cluster experiment,” in *Cloud Computing Technology and Science (CloudCom), 2013 IEEE 5th International Conference on*, vol. 2. IEEE, 2013, pp. 170–175.
- [22] F. Jalali, K. Hinton, R. Ayre, T. Alpcan, and R. S. Tucker, “Fog computing may help to save energy in cloud computing,” *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 5, pp. 1728–1739, 2016.