JITEL 2017
XIII Jornadas de Ingeniería Telemática
27-29 Septiembre 2017, Valencia (España)

Actas de las XIII Jornadas de Ingeniería Telemática (JITEL 2017), Valencia (España), 27-29 de Septiembre de 2017.

ISBN: 978-84-9048-595-8
DOI: http://dx.doi.org/10.4995/JITEL2017.2017.6504

A Study on the Energetic Viability of Single Board Computers for Cloud Computing Scenarios

Pedro Verdugo, Joaquín Salvachúa, Gabriel Huecas
Grupo de Internet de Nueva Generación,
Departamento de Ingeniería Telemática,
Escuela Técnica Superior de Ingenieros de Telecomunicación,
Universidad Politécnica de Madrid
Avenida Complutense, 30. 28040 Madrid.
pmverdugo@dit.upm.es, joaquin.salvachua@upm.es, gabriel.huecas@upm.es

***Abstract*—The following document explores the viability of the usage of consumer-grade, ARM-based single board computers as a power saving alternative to the traditional monolithic x64-full-server based approach. By taking advantage of several capabilities provided by such devices, such as low cost, low power consumption and low on-time, the authors finally propose a scalable, energy-efficient, ARM-based cloud infrastructure. To that end, we start analyzing the current offerings in terms of capabilities, net cost, processing power and power consumption, comparing them with the relevant server-oriented offerings. We subsequently explore the adequacy of several metrics to model on-budget raw data processing, considering full-system wattage under nominal usage conditions. The low initial investment and long-term affordability of this approach results in quite a relevant case of application to Edge Cloud computing scenarios.**

***Palabras Clave*—cloud computing; energy efficiency; green datacenter; microprocessor**

I. INTRODUCTION

As current Big Data loads continue to increase, datacenter processing power is constantly required to scale exponentially. Therefore, it is of the utmost importance for the current infrastructures to ensure that such increase in volume is performed in an energy-efficient way. As a result, major industry players are turning to customized hardware options, often different from the traditional monolithic x64-full-server based approach.

In the present study we will investigate the energy and cost viability of ARM architecture processors for the deployment of cloud based Big Data analysis datacenters as opposed to more traditional systems.

For starters, we will classify the current hardware offerings in terms of selected variables applicable to the investigation field in the subject matter, such as *power consumption*, *price* and *processing power*.

A higher order classification will be made possible by subdividing the detected offerings in terms of actual datacenter volume, providing three levels of performance maturity.

These first order variables will later be put into context by means of an analysis of the relevant metrics to use, where we will explore the validity of the data provided by current manufacturer documentation, extending and adapting them to our constrained field of study.

A case will be made for the selection of the main metrics employed along this case of study, which will namely consist of energy-related **Data per Joule** , time-related **Data per Second** and cost-related **Total Cost of Ownership**.

Once clarified and refined, these detected useful metrics will be used to obtain real world values with which to compare the previously selected offerings. These results will be presented in a graphic form and interpreted in terms of significance.

Finally, we will summarize the value of the original contribution as presented, as well as point out future open avenues of investigation.

II. LOW COST DEVICES VS. TRADITIONAL SYSTEMS

The last few years have seen a great deal of effort being poured into the development of embedded processors, mainly driven by the need for low power-high performance systems in the cellphone market. As a side effect, nowadays ARM (Advanced RISC Machine) architecture based devices are ubiquitous as access devices, and also as the core of embedded systems in all kind of sensor networks and home appliances.

This work is licensed under a Creative Commons 4.0 International License (CC BY-NC-ND 4.0)
EDITORIAL UNIVERSITAT POLITÈCNICA DE VALÈNCIA

Table I
PROCESSOR CHARACTERISTICS

Grade	*User*		*MicroServer*		*Datacenter*	
Processor	*Amlogic S805*$^+$	*Xeon E5404*$^+$	*Samsung Exynos5422*$^+$	*Core i7-4790K*$^+$	*Cavium ThunderX*	*Xeon E74890*
ISA	ARM Cortex-A5	x86-64	ARM Cortex-A15	x86-64	ARMv8-x64	x86-64
Number of cores	4	8	8	4	48	15
Frequency (GHz)	1.5	2.5	2.00	4.00	2.5	2.80
Dhrystone GIPS	1.57	12.10	1.78	33.435	12.53	270.73
DGIPS (Total)	9.42	96.85	14.24	133.74	601.65	4061
TDP (W)	2.3	80	4.25	88	80	155
Idle Power (W)	0.73	30	1.75	42	24	67.5
Price (USD)	37	382.84	74	699	750	5649.95

There is a case to be made, therefore, for the use of ARM machines as datacenter processing nodes[1].

A. SPECIFIC HARDWARE COMPARISON

To set the testbed for the following sections, in table I we will break down the specifics of the two base processors families used for comparison, indicating their Instruction Set Architecture (from now on, **ISA**), processing power and price[1].

Single core Dhrystone MIPS have been considered as a de-facto standard for processing power calculations, avoiding the shortcomings of regular manufacturer-provided MIPS data for different system architectures. For the studied architectures, a common term of reference will be billion of instructions per second, or **GIPS**.

As illustrated with the data presented in table I, in its current state the ARM architecture offers some promising characteristics that we can relate to their most common x86 counterpart for a set number of cores:

- **Performance**: In a first approach, x86 processor performance is extremely superior to that of an ARM core.
- **Price**: The cost of a server or computer grade x86 processor and board is quite superior to an ARM system.
- **Power Consumption**: As previously mentioned, the ARM power usage is unequivocally inferior to its x86 counterpart.

III. RELEVANT METRICS

A. CLASSIC POWER METRICS

From 1982[2], classic relevant metrics for server workloads have been based on the **Thermal Design Power** (TDP), or thermal design point, defined as the maximum amount of heat generated during typical computer operation [3]. This clearly insufficient concept[2] as a measure of computer processing power has been overly relegated to a back plane, as the main cpu manufacturers tend to introduce new mainly subjective and inexact measurement to try and solve these constraints.

On the one side, in 2009 AMD proposed the **Average CPU Power** (ACP) [4], with scarce application results.

[1]All prices valid for the current date, given in USD and retrieved from *Amazon* or the manufacturer site as applicable.
indicates measured values.

In the same vein, Intel's own recently introduced[5] **Scenario Design Power** (SDP) is defined as an operating mode of certain mobile processors, revamping the TPD concept to set another metric with a lower thermal point, without any practical application or formal definition whatsoever.

B. ENERGY METRICS

As stated by Hennessy[2], processor performance metrics must necessarily be tied to energy, and not power measurements. This ensures the ability to compare in the same grounds different processor architectures as well as different families from the same one. The introduced energy metrics that will be used along the rest of this publication, are defined as follows:

- **Data Processed Per Second** (DPS): Understood as the amount of CPU processed data in a given time (in our case, one second).
- **Data Processed Per Joule** (DPJ): Defined as the amount of data the CPU is able to process with a given energy budget of 1 Joule.
- **Energy-Delay Product** (EDP): As introduced by Horowitz [6] in the transistor performance environment and subsequently expanded by Laros III[7], defines in our specific environment the time taken by the processor to output a given amount of data for a set energy budget. Because it relates to the output processing latency due to i/o artifacts, we will not consider it in our processor-only context.

C. COST CALCULATIONS

The study of cloud computing setup costs has been of wide interest throughout the literature, and there seems to be an agreement as the usage of the **Total Cost of Ownership** (TCO) as a de-facto standard. For the purposes of this paper, we will base our calculations on the TOC formulae presented in [8] as referred to **Infrastructure as a Service** (IaaS) setups, with further operational cost refinements as detailed by [9].

n	Number of nodes in cluster
t	Runtime (Hours)
C_{pi}	Provisioning Cost per node ($)
C_{ei}	Total Electricity Cost per node ($)
C_h	Electricity Cost per hour (KWh)
P_f	Full Power Usage per node (W)
P_i	Idle Power Usage per node (W)

This work is licensed under a Creative Commons 4.0 International License (CC BY-NC-ND 4.0)
EDITORIAL UNIVERSITAT POLITÈCNICA DE VALÈNCIA

U Usage Factor (%)

Equation 1 will set the main cost formula to apply to our setup:

$$TCO = \sum_{i=1}^{n}(C_{pi} + C_{ei}) \quad (1)$$

Where C_e can be furtherly detailed as follows:

$$C_e = t * C_h * (U * P_f + (1 - U) * P_i) \quad (2)$$

In our case of study, we will limit the energy aspects to a given set of cpu-intensive tasks (Mesos based MapReduce tasks), but in the interest of completeness we must remark that datacenter and input/output related costs are of the utmost importance to the correct applicability of the following calculations.

D. ENERGY MINIMIZATION STRATEGIES

As detailed in [10], there are several layers of energy efficiency mechanisms to consider when designing a cloud computing system, from the hardware perspective to the Data Center level, including the OS and virtualization levels.

For our pretended setup, we must take into consideration the workload for which our setup will be used. For Big Data-oriented case, the following techniques have been proposed for saving energy in Hadoop MapReduce deployments[11]:

- **Covering Set** (CS): By powering off all non-essential nodes, current jobs are delegated to a given subset of nodes, sufficient to cover the task at hand. This implies an increase in the DPS and the time of task completion, but at a lower DPJ, given that less nodes are active.
- **All-In Strategy** (AIS): This strategy proposes to use all the processing power available at a given time, thus reducing the aforementioned DPS metric and increasing the DPJ to achieve the maximum available performance.
- **Berkeley Energy Efficient MapReduce** (BEEMR): Emerges as an improvement over the CS strategy, for real-time processing MapReduce systems; this time defining interactive zones, where real time processing is required and where all nodes will be active, and batch zones, that will be permitted to enter low power states depending on task load.

IV. TCO RESULTS

Most of the proposed capacity allocation algorithms and concepts can be applied without further modifications to embedded system architectures. From the aforegiven strategies, we will implement an **AIS** approach, in which tasks will be completed as soon as possible in order to maximize off-time.

In table II, we will set the values of the previously defined parameters for our calculations. The values as presented are based on mean cost calculations as illustrated in Martens[8].

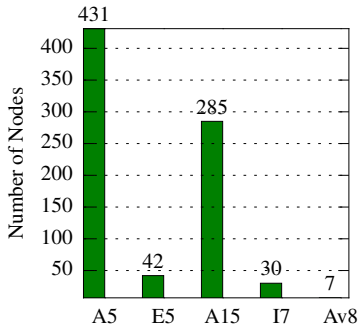

Figure 1. Nodes per DPJ budget

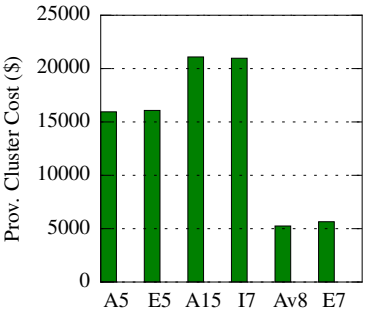

Figure 2. Initial Cost per DPJ budget

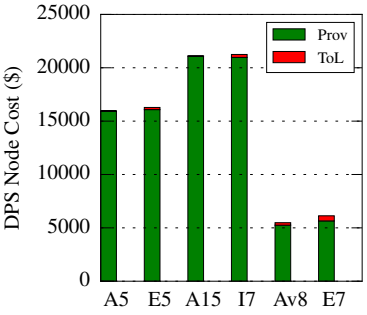

Figure 3. Node Cost per DPS

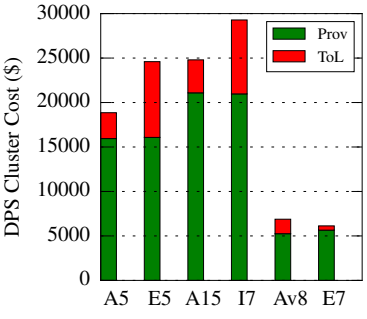

Figure 4. Cluster Cost per DPS

This work is licensed under a Creative Commons 4.0 International License (CC BY-NC-ND 4.0)
EDITORIAL UNIVERSITAT POLITÈCNICA DE VALÈNCIA

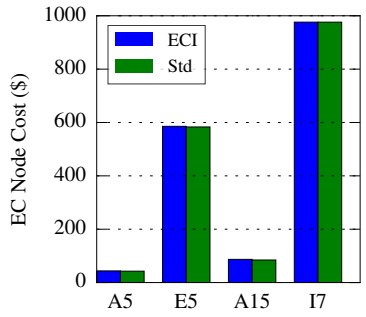

Figure 5. Node Cost with ECI

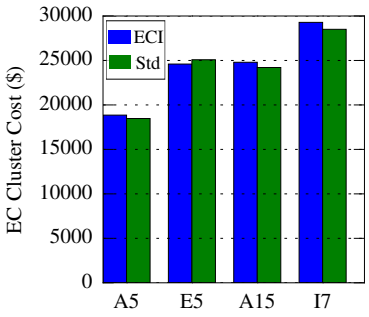

Figure 6. Cluster Cost with ECI

A. ENERGY-CONSTRAINED BUDGET

The calculations for this section will consider costs for a fixed DPJ budget set by the most powerful processor available, the Intel Xeon E7 4890. From here, we will determine how many nodes would be needed to reach said DPJ budget, as well as the provisioning costs for the obtained setup.

In figure 1 we can see how 431, 42, 285,30 and 7 nodes are respectively needed to achieve the Xeon performance level.

We also observe in figure 2 the relative initial cost advantage when provisioning user-grade infrastructure, as opposed to the high cost of deploying a microserver-grade datacenter.

B. TIME-CONSTRAINED BUDGET

In this section, we will analyze the costs associated with a given DPS budget, for a fixed time to task completion of 4 years of full use server lifetime.

Both figure 3 and 4 will represent the aforementioned systems in their X-axes, while the Y-axis will provide the results for the total cost (C_e from the previous formula 1) for the element of study in the given time period, marked in the graph as Time of Life (ToL). The base of the bar (darker colour) represents the provisioning costs (C_{pi}, marked as Prov in the graph) as opposed to the electrical cost (C_{ei} in lighter colour, upper bar).

Table II
TEST DATA

Time of Life (ToL)	4 years
Electric Power Cost (KWh)	$0.11
Mean Usage Factor	0.65

Table III
GOOGLE TCO (4 YEARS)

Grade	*User*	*MicroServer*	*DataCenter*
Google TCO	*New Startup*	*Static Enterprise App*	*Mature App*
Total Cost ($)	6258,76	153,861.6	267,632.84

Figure 3 illustrates the total cost incurred for each node for the given period of time As for figure 4, the total cost for a cluster comprised of the previously calculated number of nodes is shown.

As we can clearly see, Intel processors power usage for a cluster is quite elevated compared with their ARM counterparts (2 to 3 times).

1) ENERGY CONTROLLER INSERTION: As proposed in [12], the insertion of a power controller node with the only task of turning on or off the required nodes as needed reducing idle power consumption, renders a negligible increase in provisioning costs (figure 5) for a relatively positive increase in energy efficiency in the case of user and microserver-grade processor clusters, as seen in Figure 6 where ECI values (green, with legend ECI) are related to the previously studied case (blue, with legend Std).

C. COST-CONSTRAINED BUDGET

This last case of study will model an initial capital inversion based on the three **Google TCO Platform Calculator** (https://cloud.google.com/pricing/tco) profiles that most closely resemble our system partition, which we will detail in table III. We must note that these prices are taken as a fixed reference for processor performance comparison, not as total datacenter costs.

In the first graph from figure 7 we see how even with a limited budget we can keep a considerable number of working ARM nodes for the given timeframe. The second graph from figure 7 closely follows the behaviour introduced in the first one, and shows how a huge number of low power cores at full usage can be more efficient than a reduced number of power-hungry ones. The third graph from figure 7 is consistent with the description in [13], and shows the performance price to pay for this increase in the number of nodes, that is, a decrease in data throughtput clearly limited by the processing capabilities of each individual node.

V. CONCLUSION

As a brief recapitulation of the exposed data as well as a comparison with the current literature, we can summarize our findings as follows:

As stated in previous works [14], ARM-based architectures are not univocally superior to traditional datacenter infrastructures neither in raw data processing nor in standard energy usage. In a first instance this seems to oppose the optimistic findings of Svandeldt-Winter[15], but we must consider the excellent improvements to x86 energy state management in the last few years, which have

This work is licensed under a Creative Commons 4.0 International License (CC BY-NC-ND 4.0)
EDITORIAL UNIVERSITAT POLITÈCNICA DE VALÈNCIA

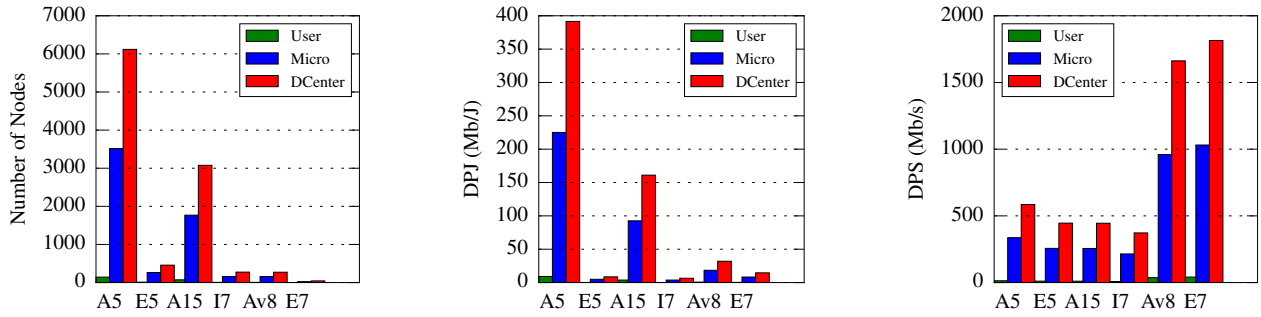

Figure 7. Cost-Driven Budget

increased the DPJ in at least 2 orders of magnitude over the previous generation[5]. However, depending on the task at hand, there seems to be a case to be made for the proposed system in whichever situation involves CPU Intensive workloads, such as the case of small, random database accesses[16]. A generalization of these ideas, based on a similar hardware platform, can also be seen in Cecowski [17] as the proposal for a modular, ARM-based datacenter. Furthermore, in the same trend of our current proposal, Big Data workloads over ARM infrastructures can already be simulated thanks to the work of Kecskemeti[18]

From a cost-based approach, we have improved on [19] by considering typical warehouse time-of-life power consumption, discovering that the accumulated electricity costs for a cluster system clearly cut on the data therein presented.

From an energy-based standpoint, and extending on the outstanding analysis of Tudor [20], we have generalized the energy studies to n-machine cluster structures. This has clearly shown the performance degradation caused by the linear scale of the underlying support hardware (ram bus speed, network, storage), which will hit on the system's performance under I/O Intensive workloads, as is the case with sequential database scans. As pointed in [21], a valid solution for these constraints is the integration of more cpu cores by board, which seems to be consistent with the current market direction.

Cluster job scheduling has also received attention as denoted in [22], as well as physical thermal design [23] and data migration considerations [24]. However insightful these technologies may be for a practical infrastructure deployment, in a first instance they escape the applicable premises for our study.

There are common shortcoming for ARM-based systems pointed at in all the researched literature related to the operating system and software layer, that we'll newly establish here:

- It's necessary to ensure the usage of a parallellization oriented OS (coreos, ranchos) to keep the underlying hardware performance degradation under check.
- The usage of a parallel task optimized management environment (mesos, nomad, hadoop) is also of paramount importance to ensure scale-growth.
- The compile-time optimization of the running code for the specific processor architecture is probably the most determinant and most often underlooked feature that can improve ARM cluster performance.

To conclude, the authors concur on the interest and feasibility of the proposed reference infrastructure, generalizing the DPS and DPJ energy studies to n-machine clusters based on ARM processors, and, given the presented results, also consider the need to consolidate the presented data with the promising current advances in the ARM64 architecture [25].

VI. ACKNOWLEDGEMENTS

The current work has been partially funded by the Ministerio de Educación, Cultura y Deporte of Spain.

REFERENCES

[1] C. Pahl, S. Helmer, L. Miori, J. Sanin, and B. Lee, "A Container-Based Edge Cloud PaaS Architecture Based on Raspberry Pi Clusters," in *2016 IEEE 4th International Conference on Future Internet of Things and Cloud Workshops (FiCloudW)*, 2016, pp. 117–124.

[2] J. L. Hennessy and D. A. Patterson, *Computer Architecture: A Quantitative Approach*. Elsevier, 2012.

[3] C. Gough, I. Steiner, and W. Saunders, "CPU Power Management," in *Energy Efficient Servers*. Apress, 2015, pp. 21–70. [Online]. Available: http://link.springer.com/chapter/10.1007/978-1-4302-6638-9_2

[4] AMD, "ACP — The Truth About Power Consumption Starts Here," AMD, Tech. Rep., 2009, 00000.

[5] Intel, *Intel Brings Core Down to 7W, Introduces a New Power Rating to Get There: Y-Series SKUs Demystified*, 2013, 00000. [Online]. Available: http://www.anandtech.com/show/6655/intel-brings-core-down-to-7w-introduces-a-new-power-rating-to-get-there-yseries-skus-demystified

[6] M. Horowitz, T. Indermaur, and R. Gonzalez, "Low-power digital design," in *Low Power Electronics, 1994. Digest of Technical Papers., IEEE Symposium*. IEEE, 1994, pp. 8–11, 00362. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=573184

[7] J. H. Laros III, K. Pedretti, S. M. Kelly, W. Shu, K. Ferreira, J. Vandyke, and C. Vaughan, "Energy delay product," in *Energy-Efficient High Performance Computing*. Springer, 2013, pp. 51–55, 00005. [Online]. Available: http://link.springer.com/chapter/10.1007/978-1-4471-4492-2_8

[8] B. Martens, M. Walterbusch, and F. Teuteberg, "Costing of cloud computing services: A total cost of ownership approach," in *System Science (HICSS), 2012 45th Hawaii International Conference on*. IEEE, 2012, pp. 1563–1572, 00050. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6149074

This work is licensed under a Creative Commons 4.0 International License (CC BY-NC-ND 4.0)
EDITORIAL UNIVERSITAT POLITÈCNICA DE VALÈNCIA

[9] L. A. Barroso, J. Clidaras, and U. Hölzle, "The datacenter as a computer: An introduction to the design of warehouse-scale machines," *Synthesis lectures on computer architecture*, vol. 8, no. 3, pp. 1–154, 2013, 01105. [Online]. Available: http://www.morganclaypool.com/doi/abs/10.2200/S00516ED2V01Y201306CAC024

[10] A. Beloglazov, "Energy-efficient management of virtual machines in data centers for cloud computing," 2013, 00028. [Online]. Available: https://minerva-access.unimelb.edu.au/handle/11343/38198

[11] A. Beloglazov and R. Buyya, "Optimal online deterministic algorithms and adaptive heuristics for energy and performance efficient dynamic consolidation of virtual machines in cloud data centers," *Concurrency and Computation: Practice and Experience*, vol. 24, no. 13, pp. 1397–1420, 2012. [Online]. Available: http://onlinelibrary.wiley.com/doi/10.1002/cpe.1867/full

[12] N. Maheshwari, R. Nanduri, and V. Varma, "Dynamic energy efficient data placement and cluster reconfiguration algorithm for MapReduce framework," *Future Generation Computer Systems*, vol. 28, no. 1, pp. 119–127, 2012, 00061. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S0167739X1100135X

[13] M. Malik and H. Homayoun, "Big data on low power cores: Are low power embedded processors a good fit for the big data workloads?" in *Computer Design (ICCD), 2015 33rd IEEE International Conference on*. IEEE, 2015, pp. 379–382, 00001. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=7357128

[14] D. Loghin, B. M. Tudor, H. Zhang, B. C. Ooi, and Y. M. Teo, "A performance study of big data on small nodes," *Proceedings of the VLDB Endowment*, vol. 8, no. 7, pp. 762–773, 2015, 00006. [Online]. Available: http://dl.acm.org/citation.cfm?id=2752945

[15] O. Svanfeldt-Winter, S. Lafond, and J. Lilius, "Cost and energy reduction evaluation for ARM based web servers," in *Dependable, Autonomic and Secure Computing (DASC), 2011 IEEE Ninth International Conference on*. IEEE, 2011, pp. 480–487. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6118745

[16] B. Kantarci, L. Foschini, A. Corradi, and H. T. Mouftah, "Inter-and-intra data center VM-placement for energy-efficient large-scale cloud systems," in *Globecom Workshops (GC Wkshps), 2012 IEEE*. IEEE, 2012, pp. 708–713, 00019. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6477661

[17] M. Cecowski, G. Agosta, A. Oleksiak, M. Kierzynka, M. v. d. Berge, W. Christmann, S. Krupop, M. Porrmann, J. Hagemeyer, R. Griessl, M. Peykanu, L. Tigges, S. Rosinger, D. Schlitt, C. Pieper, C. Brandolese, W. Fornaciari, G. Pelosi, R. Plestenjak, J. Cinkelj, L. Cudennec, T. Goubier, J. M. Philippe, U. Janssen, and C. Adeniyi-Jones, "The M2dc Project: Modular Microserver DataCentre," in *2016 Euromicro Conference on Digital System Design (DSD)*, 2016, pp. 68–74.

[18] G. Kecskemeti, W. Hajji, and F. P. Tso, "Modelling Low Power Compute Clusters for Cloud Simulation," in *2017 25th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP)*, Mar. 2017, pp. 39–45.

[19] Z. Ou, B. Pang, Y. Deng, J. K. Nurminen, A. Yla-Jaaski, and P. Hui, "Energy-and cost-efficiency analysis of arm-based clusters," in *Cluster, Cloud and Grid Computing (CCGrid), 2012 12th IEEE/ACM International Symposium on*. IEEE, 2012, pp. 115–123. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6217412

[20] B. M. Tudor and Y. M. Teo, "On understanding the energy consumption of arm-based multicore servers," in *ACM SIGMETRICS Performance Evaluation Review*, vol. 41. ACM, 2013, pp. 267–278. [Online]. Available: http://dl.acm.org/citation.cfm?id=2465553

[21] N. Rajovic, L. Vilanova, C. Villavieja, N. Puzovic, and A. Ramirez, "The low power architecture approach towards exascale computing," *Journal of Computational Science*, vol. 4, no. 6, pp. 439–443, Nov. 2013, 00046. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S1877750313000148

[22] S. Zikos and H. D. Karatza, "Performance and energy aware cluster-level scheduling of compute-intensive jobs with unknown service times," *Simulation Modelling Practice and Theory*, vol. 19, no. 1, pp. 239–250, 2011, 00048. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S1569190X10001309

[23] R. T. Kaushik and K. Nahrstedt, "T: a data-centric cooling energy costs reduction approach for big data analytics cloud," in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*. IEEE Computer Society Press, 2012, p. 52, 00027. [Online]. Available: http://dl.acm.org/citation.cfm?id=2389067

[24] C. Ghribi, M. Hadji, and D. Zeghlache, "Energy Efficient VM Scheduling for Cloud Data Centers: Exact Allocation and Migration Algorithms." IEEE, May 2013, pp. 671–678, 00041. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6546155

[25] K. Keipert, G. Mitra, V. Sunriyal, S. S. Leang, M. Sosonkina, A. P. Rendell, and M. S. Gordon, "Energy-Efficient Computational Chemistry: Comparison of x86 and ARM Systems," *Journal of chemical theory and computation*, vol. 11, no. 11, pp. 5055–5061, 2015, 00001. [Online]. Available: http://pubs.acs.org/doi/abs/10.1021/acs.jctc.5b00713

This work is licensed under a Creative Commons 4.0 International License (CC BY-NC-ND 4.0)
EDITORIAL UNIVERSITAT POLITÈCNICA DE VALÈNCIA